Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

Sebastian Sund

Real options with active learning

An approach based on the Kalman filter

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel, Lars H. Sendstad, Jacco Thijssen and Roel Nagy
July 2020

NTNU
Norwegian University of
Science and Technology

Sebastian Sund

Real options with active learning

An approach based on the Kalman filter

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel, Lars H. Sendstad, Jacco Thijssen and Roel Nagy
July 2020

Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

NTNU
Norwegian University of
Science and Technology

Abstract

Although traditional real options models are able to appropriately account for an uncertain investment environment, they do not allow the decision maker to actively seek information in order to learn about future market conditions. This opportunity to learn introduces considerations that have received increased attention in recent literature. With active learning, factors of the project under consideration that have traditionally been assumed to stay constant are allowed to change as a result of observations. This thesis extends the literature by discussing how an option holder may learn about a stochastic process that evolves with time. The Kalman filter is applied to derive a time-varying estimate of the process, and the option is valued as dependent on this estimation. We focus our attention on linear stochastic processes with normally distributed noise. The resulting estimation has a risk that is assumed to be undiversifiable. A utility model is therefore applied to model the option holder's relationship to risk. The main purpose of the thesis is to illustrate the applicability of this approach and to discuss properties of the resulting optimal strategy. With a numerical example, we find that the marginal benefit of learning decreases rapidly over time, and that the majority of exercises occur early in the option holding period, after the holder has realized the main benefits of learning. We also illustrate how the distribution of exercise times varies with changes in the initial belief.

Sammendrag

Selv om tradisjonelle realopsjonsmodeller tar hensyn til usikkerheter i markedet, åpner de ikke for en beslutningstaker som aktivt søker ny informasjon for å lære om fremtidige markedsforhold. Muligheten til å lære fra omgivelsene introduserer nye betraktninger som har fått økt oppmerksomhet i nyere litteratur. Med aktiv læring tillates faktorer som tradisjonelt sett har vært antatt å forbli konstante å endre seg med nye observasjoner. Denne oppgaven utvider literaturen om realopsjoner ved å diskutere hvordan en opsjonsholder kan lære om stokastiske prosesser som utvikler seg over tid. Vi anvender et Kalmanfilter til å formulere et tidsavhengig estimat av en observert stokastisk prosess. Vi begrenser oss til lineære stokastiske prosesser med normalfordelt støy. Den resulterende estimeringsprosessen har en usikkerhet som antas å ikke kunne diversifiseres. Vi formulerer derfor opsjonsholderens forhold til risiko ved hjelp av en nyttefunksjon. Hovedformålet med denne oppgaven er å demonstrere hvordan denne metoden kan anvendes, og å diskutere egenskapene til den optimale beslutningsstrategien. Ved hjelp av et numerisk eksempel ser vi at marginalverdien av læring synker fort over tid, og at majoriteten av investeringsbeslutninger forekommer tidlig i levetiden til opsjonen, etter opsjonsholderen har realisert de største læringsfordelene. Vi viser også hvordan fordelingen av investeringsbeslutninger varierer med endringer i det opprinnelige estimatet.

Preface

This thesis concludes my studies towards the degree of Master of Science in Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU). The work is based in part on research topics raised while working on a specialization project in the fall of 2019. I would like to thank my supervisors Verena Hagspiel, Lars H. Sendstad, Jacco Thijssen and Roel Nagy for bringing these topics to my attention, and for excellent guidance and extensive support throughout all stages of the thesis. I would also like to thank my family and Mona for their continuous support.

Contents

List of figures

1 Introduction

Investment decisions, such as the launch of nascent technologies, or whether to enter an emerging market, often involve considerable uncertainty. If the investment opportunity allows for it, waiting to invest may play a crucial role in reducing this uncertainty (Dixit & Pindyck, 1994), especially if more information about the project presents itself over time (Dixit, 1992). By observing and incorporating this information, the decision maker may form better grounds for decision making, especially when the information regards factors of high sensitivity to the project value. Often, the most pertinent underlying factor to project valuation is stochastic over time and may consequently be modeled by a stochastic process. A decision maker with the ability to learn from new information may form an estimation of the current value of such processes, effectively forming an estimation that also evolves as a stochastic process. Individual observations may not give an accurate picture of reality, and so the decision maker should account for some measurement error. The main purpose of this thesis is to derive and analyze the optimal strategy of a decision maker who holds a continuous time estimation of certain kinds of stochastic processes that are not perfectly observable, but may be estimated by incorporating noisy observations.

We examine the optimal investment behavior of a firm that has the option to invest in a project against a fixed sunk cost, and consider a situation in which the underlying variable of the project value is assumed to be either a constant or an Ornstein-Uhlenbeck process, with parameter values known to the firm at the beginning of the option holding period. The initial process value is uncertain, and since the process is not perfectly observable, its value remains uncertain throughout the holding period. The firm may engage in noisy observations of the process over time in order to obtain an estimate of its value, and consequently, the project value. This noise represents a firm-specific uncertainty, and the firm is consequently assumed to face an incomplete market. We therefore resort to maximizing expected utility, assuming a firm with a known constant relative risk aversion and rate of time preference.

We begin our analysis by formulating a continuous time subjective belief of a stochastic process with a certain structure based on historic observations. We then apply this model to an Ornstein-Uhlenbeck process. In order to compare the resulting estimation, another estimation is designed with the assumption of an underlying process that is constant at the mean value of the Ornstein-Uhlenbeck process, unless explicitly stated otherwise. We then derive the expected net present utility of investing in the project, and proceed to discuss properties of the option value. The optimal investment policy is derived with a simulation-based method. Finally, we present a case study in order to demonstrate the dynamics of the model.

In contrast to investment strategies based on the net present value (NPV), traditional real options models incorporate the flexibility of waiting to invest, and are consequently often valued at a premium over their myopic counterparts. Seminal contributions to this field include Dixit and Pindyck (1994) and McDonald and Siegel (1985, 1986). Recently, the real options literature has been extended by the incorporation of learning. This new methodology considers information acquisition as a conscious activity by the firm, and not a passive consequence of waiting to exercise the option. In contrast to traditional models in which parameters that establish the project value are assumed to be known at the beginning of the option holding period, a learning firm changes its estimations throughout this period, either discretely or continuously, based on received information. If this information is in fact an accurate representation of reality, a firm applying "active learning" will have better grounds for decision making throughout the holding period than a firm that remains passive, and consequently an improved investment strategy. This additional benefit of holding the option therefore introduces another quantifiable factor to consider when deciding whether to exercise.

There are several ways of modeling active learning. One widely applied method is by Bayesian updating of a parameter's prior probability distribution when new observations arrive. Examples of this approach include Singh, Ghosh, and Adhikari (2018), who apply Bayesian updating at discrete observations to estimate the drift and diffusion parameters of an Ornstein-Uhlenbeck process, and Blanke and Bosq (2012) who study a similar problem in both discrete and continuous time. Although the majority of papers that model learning in a real options setting do so through Bayesian updating, another equally feasible estimation method is the Kalman filter algorithm, as outlined in Øksendal (2013, p. 2). Instead of sequentially updating the entire probability distribution of the estimated parameter, the Kalman filter algorithm is generally applied to only update the conditional expectation, circumventing some of the analytical difficulties with Bayesian updating. When the observations are considered to be independent of the observed parameter, and observation noise is considered to be normally distributed, the two approaches result in identical estimations, as demonstrated by Barker, Brown, and Martin (1995) and Soyer (2018).

A familiar result in real options analysis is the nondecreasing relationship between uncertainty and option value. Incorporating learning in a real options framework often reduces this uncertainty, and a natural question is whether learning is valuable in this context. As Martzoukos and Trigeorgis (2001) conclude, the value lost from lower uncertainty is compensated by the value gained from better information. This may also be the case if learning is costly, as supported by the results of Harrison and Sunar (2015) and Bellalah (2001), among others, supporting the

notion that incorporating learning adds additional value to the traditional real options framework. Nevertheless, the work on learning in a real options context is arguably scarce. Among the existing literature, Kwon and Lippman (2011), Ryan and Lippman (2003), Kwon, Xu, Agrawal, and Muthulingam (2016) and Thijssen, Huisman, and Kort (2004) take a theoretical approach and illustrate how the optimal investment strategy is influenced by different aspects of learning about some aspect of future project profitability when the investor is assumed to be able to enter, exit or expand the project, or a combination of the three. Herath and Herath (2008), Kwon (2014) and Dalby, Gillerhaugen, Hagspiel, Leth-Olsen, and Thijssen (2018) examine how learning with real options can be applied to specific situations in industry. More specifically, Herath and Herath (2008) consider how the framework may help in valuing certain types of information security systems, and conclude that the incorporation of learning leads to a reduction of upward bias in estimates, as well as a specific implications for security system management. Kwon (2014) models the optimal decision policy of a firm that has the option to invest in order to protect a project against disruption and may continuously learn about the probability of this disruption from trends in the market. His model illustrates the sensitivity of optimal decision to the probability of disruption. Dalby et al. (2018) consider a firm which may invest in a project in the renewable energy industry that is subject to an expected adjustment of the support scheme it is currently backed by. The firm is assumed to be able to learn about the arrival rate of the adjustment from a continuous information stream. The authors illustrate how the optimal investment threshold varies with the desired learning rate and the corresponding effect on option value, and, notably, how the relative time to optimal investment decreases with learning rate. A key contribution of this thesis in the context of the aforementioned papers is that we estimate a stochastic process, and not a parameter that is assumed to be constant in time.

A factor that is often of interest in relation to active learning is the "rate of learning", or, "learning rate", represented by Kwon et al. (2016) as a parameter that reflects the magnitude of the difference between the prior and posterior probability distributions when applying Bayesian updating. With a Kalman filter, the equivalent measure would be the magnitude of the difference between the prior and posterior conditional expectations. In both cases, the learning rate may be defined as a function of the volatility of the observation process. The volatility describes the uncertainty of the estimate, and a higher rate of learning is intended to translate to a faster decrease in estimator uncertainty. When learning is considered to be costly, the cost function may be expressed in terms of this learning rate, as demonstrated in Moscarini and Smith (2001), and the investor is consequently faced with selecting the optimal learning rate. We consider a fixed learning rate in this thesis, but acknowledge the importance of discussing the optimal learning

rate when considering practical applications, especially when learning is costly, as discussed in Hagspiel, Nagy, Sund, and Thijssen (2019). We assume costless learning to simplify the analysis and to illustrate the dynamics of the optimal strategy more clearly. For further discussion on costly learning within a real options framework, consider e.g. Bellalah (2001); Bergemann and Välimäki (2008); Harrison and Sunar (2015); Keller and Rady (1999); Moscarini and Smith (2001); Pertile, Forster, and Torre (2014) and Thijssen and Bregantini (2017).

We focus our discussion on a certain class of stochastic processes known as Ornstein-Uhlenbeck processes. These processes are mean-reverting, and have been applied to model a wide range of scientific phenomena. Within finance, evidence for mean-reversion is abundant (Wong & Lo, 2009), and the Ornstein-Uhlenbeck process has been used to model commodity prices, as in Schwartz (1997) and Lucia and Schwartz (2002), exchange rates, as in Jorion and Sweeney (1996), and interest rates, as in Vasicek (1977). In a real options framework, Ekström, Lindberg, and Tysk (2011) formulate the problem of when to liquidate a position in a pairs trade by modeling a mean-reverting price spread with an Ornstein-Uhlenbeck process. Their model is extended by Leung and Li (2015) with the incorporation of a strategy for optimal entry into the position. In an industrial context, Näsäkkälä and Fleten (2005) analyse a real options problem of investment in a power plant when the spread between the electricity price and cost of gas is assumed to follow the sum of an arithmetic Brownian motion and an Ornstein-Uhlenbeck process, similar to the method applied in Lucia and Schwartz (2002). Overall, if an observable underlying process of a project is an Ornstein-Uhlenbeck, and the investor has derived an expression for the expected value of the project, our model may be applied to devise an optimal investment strategy. As an example of current relevance, Gray, Greenhalgh, Hu, Mao, and Pan (2011) demonstrate that the disease transmission coefficient in an epidemiological "susceptible-infected-susceptible (SIS)"-model may be expressed by an Ornstein-Uhlenbeck process. If an investor formulates a project value in terms of this coefficient, our model may be applied to value the opportunity of investing in it. We introduce a case study that assumes an industry with mean reverting prices following an Ornstein-Uhlenbeck model. It should be noted, however, that our main concern with this thesis is to derive a model of general applicability. The case study is intended solely to demonstrate the properties and dynamics of the model.

Although the aforementioned literature offers meaningful insights on optimal investment decisions and learning, it is developed under the assumption of risk neutrality, which relies on the assumption that the underlying asset may be spanned or replicated by assets in the market. Hugonnier and Morellec (2007) point out the following: "While assumptions of risk neutrality or market completeness are convenient to characterize investment decisions under uncertainty, they are not

particularly relevant to most real-world applications. In particular, corporate executives and entrepreneurs typically have to make investment decisions in situations where the cash flows from the project are not spanned by those of existing assets or under other constraints which make them face incomplete markets. In such environments, we can expect their risk aversion to affect firms' investment decisions." In this thesis, the stochastic process underlying the option value consists of a volatility component that changes with time. If the firm were to attempt to create a replicating portfolio, it would have to continuously update the portfolio composition in order to accurately reproduce the dynamics of the process. As noted by Leland (1985) and, more recently, Kolm and Ritter (2019), the presence of transaction costs makes a continuously updated portfolio infinitely costly in theory. In practice, a dynamically replicating portfolio would be updated discretely, which limits total transaction costs, at the expense of a lower replication accuracy. Although there are ways of optimizing this trade-off, as both Leland (1985) and Kolm and Ritter (2019) show, we have decided on a different modelling approach that avoids these difficulties altogether. Similar to Henderson and Hobson (2002), we assume a firm with known, constant relative risk aversion (CRRA), as well as a constant rate of time preference. Following Hugonnier and Morellec (2007), we consider the firm's net present utility of investing in the project rather than its expected net present value of cash flows as the relevant condition for investment decisions. With this approach, the risk originating from the volatility of the estimated process is incorporated in the valuation of the investment opportunity. It should be noted that the utility function may easily be converted to its risk-neutral equivalent by letting the firm's risk aversion be equal to zero, in case the estimation is in fact spanned by existing assets. Our model therefore has wider applicability than one that expresses its exercise condition in terms of expected NPV.

We consider our approach to model learning about a stochastic process in a utility-based real options framework to be the main contribution of this thesis. To the best of our knowledge, this has not yet been explored in literature. Our results show how the main benefits from learning occur early in the option lifetime, and that the distribution of exercise times has a positive relationship with both the mean and variance of the distribution of the initial estimate.

The remainder of this thesis is organized as follows. We outline our model in section 2, and our solution approach in section 3, introducing different relevant concepts along the way. A case study is then introduced and discussed in section 4, followed by a conclusion in section 5.

2 Model

We consider a firm with an option to invest in a project for a sunk, fixed cost. The option is considered to have a known lifetime. Upon investment, the firm is assumed to receive the entire value of the project, although the model may be extended by simple measures to allow for a project with profit streams that arrive over time. We assume that the value of the project depends on a single stochastic process that the firm may acquire information about through noisy observations. Furthermore, in line with Harrison and Sunar (2015), we assume that information is received by the firm frequently enough to be modelled to arrive continuously. The firm accumulates this information in order to form an estimate of the process. We will refer to this estimate as the firm's "belief process". At the beginning of the planning horizon, the firm holds an initial belief about the true process value, which is assumed to be normally distributed. This functions as the starting point of the belief process. We assume that the firm makes investment decisions based on its subjective utility function, and that it has a constant relative risk aversion, as well as a constant rate of time preference.

2.1 The belief process

Since the firm continuously receives information about the underlying process of the project value, the parameters of the belief process change in time, and consequently the firm's expected value of the project. Because the option value depends on this expectation, we derive a stochastic differential equation (SDE) that describes the evolution of the belief process through time. Our method in this section is largely based on the procedure presented in Øksendal (2013, Chapter 6). We begin by introducing the underlying process and how observations are made, and proceed to derive a general SDE for the belief process. We then apply our model to two examples of underlying processes.

2.1.1 State and observations

Consider the stochastic process X_t, representing the underlying process of the project value. In the context of Kalman filters, X_t is sometimes referred to as the "state" of the observed system. We assume that $X_t \in \mathbb{R}$ evolves according to the stochastic differential equation

$$dX_t = \mu(X_t, t)dt + \sigma(X_t, t)dU_t, \tag{2.1}$$

with initial value X_0, where $t \geq 0$ and U_t is a Brownian motion. X_t is assumed to be continuously observable by the firm, with individual observations of the form

$$H_t = \beta(X_t, t) + \gamma(X_t, t)W_t, \tag{2.2}$$

where W_t represents white noise, independent of U_t. We then define

$$Z_t = \int_0^t H_s ds \tag{2.3}$$

as the "observation process", representing the incorporation of all observations since the process was initiated at time $t = 0$. The advantage of considering Z_t as the observation process rather than H_t is that the transformation makes the mathematical modelling more tractable, while in fact no information is lost or gained by doing so (Øksendal, 2013, p. 86). We may now represent the observation process in the following differential form

$$dZ_t = \beta(X_t, t)dt + \gamma(X_t, t)dV_t, \tag{2.4}$$

where $dV_t = W_t dt$ such that V_t is a Brownian motion independent of U_t.

2.1.2 The filtering problem

Let $(\Omega, \mathcal{F}, P)$ be the probability space of the Brownian motions U_t and V_t. We now consider the "filtering problem" represented in Øksendal (2013, p. 86): "Given observations Z_s satisfying equation (2.4) for $0 \leq s \leq t$, what is the best estimate $\widehat{X}_t$ of the parameter X_t based on these observations? In other words, what is the most accurate belief $\widehat{X}_t$ about the value of X_t?" Formally,

- $\widehat{X}_t$ is $\mathcal{G}_t$-measurable, where $\mathcal{G}_t$ is the σ-algebra generated by observations $\{Z_s\}$,
- $\mathbb{E}\left[(X_t - \widehat{X}_t)^2\right] = \inf\left\{\mathbb{E}\left[(X_t - \widehat{X}_t)^2\right]; Y \in \mathcal{K}\right\}$ where $\mathcal{K} \equiv \mathcal{K}_t \equiv \{Y : \Omega \to \mathbb{R}; Y \text{ is } \mathcal{G}_t\text{-measurable}\}$, or, $\widehat{X}_t$ is an efficient estimator, having the smallest mean squared error among estimators based on observations $\{Z_s\}$.

The belief may now be expressed as $\widehat{X}_t = \mathbb{E}\left[X_t \middle| \mathcal{G}_t\right]$ (Øksendal, 2013, Theorem 6.1.2), and we may turn to the problem of finding a suitable expression for it. We apply the Kalman filter in order to obtain this expression. The Kalman filter is a procedure for estimating the state of a parameter or system based on a series of noisy observations (Øksendal, 2013, p. 2) and has been applied to a wide range of estimation problems (Grewal, 2011). Generally, if observations of a

certain parameter are subject to measurement inaccuracies that may be represented by normally distributed "noise", the Kalman filter allows one to identify the estimator with the smallest mean squared error among candidate estimators. In the situation presented in subsection 2.1.1, at measurement times $s \in [0,\ t]$, the noise that arises when measuring $\beta(s, X_s)$ is expressed by the term $\gamma(s, X_s)W_s$ from equation (2.2). Note that, depending on its structure, a measurement of $\beta(s, X_s)$ may be transformed to a measurement of X_s.

The Kalman filter represented in Øksendal (2013) only allows for observations of a linear dynamical system, in which the aforementioned processes take the form

$$dZ_t = \beta(X_t, t)dt + \gamma(X_t, t)dV_t = G(t)X_t dt + D(t)dV_t, \quad \text{and} \tag{2.5}$$

$$dX_t = \mu(X_t, t)dt + \sigma(X_t, t)dU_t = F(t)X_t dt + C(t)dU_t, \tag{2.6}$$

where $F(t), C(t), G(t) \in \mathbb{R}$ and $D(t) \in \mathbb{R} \setminus \{0\}$. We focus our discussion on processes with this structure. As noted by Soyer (2018), a Kalman filter algorithm applied to linear dynamical systems with Gaussian noise results in a Gaussian distribution that is identical to the filtering distribution obtained by application of sequential Bayesian updating. This distribution may consequently be used to obtain expectations of functions of the observed process at a given t, either analytically or numerically.

If the linear dynamical system takes the form of equations (2.5)-(2.6), Øksendal (2013, Theorem 6.2.8) shows that the application of a Kalman filter results in a stochastic differential equation for $\widehat{X}_t$ of the form

$$d\widehat{X}_t = \left(F(t) - \frac{G^2(t)S(t)}{D^2(t)}\right)\widehat{X}_t dt + \frac{G^2(t)S(t)}{D^2(t)}dZ_t, \tag{2.7}$$

where $\widehat{X}_0 = \mathbb{E}[X_0]$, and

$$S(t) = \mathbb{E}\left[(X_t - \widehat{X}_t)^2\right] \tag{2.8}$$

satisfies the differential equation

$$\frac{dS(t)}{dt} = 2F(t)S(t) - \frac{G^2(t)}{D^2(t)}S^2(t) + C^2(t). \tag{2.9}$$

Note that equation (2.9) is deterministic. If functions $F(t), C(t), G(t)$ and $D(t)$ are known and equations (2.7), and (2.9) are solvable analytically, we may derive $\widehat{X}_t$ explicitly. However, the stochastic differential equation (2.7) itself is sufficient to derive an option value, and we will therefore focus our attention on this equation. We simplify equation (2.7) by introducing

coefficient functions $L_1(t)$ and $L_2(t)$, given as

$$L_1(t) = F(t) - \frac{G^2(t)S(t)}{D^2(t)} \quad \text{and} \tag{2.10}$$

$$L_2(t) = \frac{G^2(t)S(t)}{D^2(t)}, \tag{2.11}$$

such that

$$d\widehat{X}_t = L_1(t)\widehat{X}_t dt + L_2(t)dZ_t. \tag{2.12}$$

We apply our model to two different situations. In the first, the observed process is assumed to be constant. In the other, it is assumed to follow an Ornstein-Uhlenbeck process. For clarity, we denote the coefficient functions of the first application as $L_{1,c}(t)$ and $L_{2,c}(t)$ and those of the second application as $L_{1,o}(t)$ and $L_{2,o}(t)$. Any general discussions of the coefficients will drop the subscripts.

2.1.3 Application I: A constant process

Consider a filtering problem in which X_t is assumed to be constant, meaning $dX_t = 0$ and $X_t = X_0$. We let $\mathbb{V}[X_0] = a^2$. Hence, at $t = 0$, the observer holds an estimate of X_0 with a mean of $\widehat{X}_0 = \mathbb{E}[X_0]$ and variance $\mathbb{V}[X_0] = a^2$. Observations are assumed to be of the form

$$H_t = X_t + mW_t \tag{2.13}$$

where $m \in \mathbb{R} \setminus \{0\}$, resulting in

$$dZ_t = X_t dt + m dV_t. \tag{2.14}$$

Following the steps laid out in subsection 2.1.2, we obtain the following stochastic differential equation

$$\begin{aligned} d\widehat{X}_t &= L_{1,c}(t)\widehat{X}_t dt + L_{2,c}(t)dZ_t \\ &= -\frac{a^2}{m^2 + a^2 t}\widehat{X}_t dt + \frac{a^2}{m^2 + a^2 t}dZ_t. \end{aligned} \tag{2.15}$$

Derivations are given in Appendix A. Note that by expanding dZ_t, the process is expressed in terms of the Brownian motion differential, such that

$$d\widehat{X}_t = mL_{2,c}(t)dV_t. \tag{2.16}$$

2.1.4 Application II: An Ornstein-Uhlenbeck process

Now consider a filtering problem in which X_t is assumed to follow an Ornstein-Uhlenbeck process, with $dX_t = -pX_t dt + qdU_t$, where $p > q > 0$. As before, the observer holds an estimate of X_0 with a mean of $\widehat{X}_0 = \mathbb{E}[X_0]$ and variance $\mathbb{V}[X_0] = a^2$. Observations are again of the form

$$H_t = X_t + mW_t \tag{2.17}$$

where $m \in \mathbb{R} \setminus \{0\}$, resulting in

$$dZ_t = X_t dt + mdV_t. \tag{2.18}$$

Following the steps laid out in subsection 2.1.2, we obtain the following stochastic differential equation

$$\begin{aligned} d\widehat{X}_t &= L_{1,o}(t)\widehat{X}_t dt + L_{2,o}(t)dZ_t \\ &= -\frac{\left(p + \frac{a^2}{m^2}\right)\sqrt{p^2 + \frac{q^2}{m^2}} + \left(p^2 + \frac{q^2}{m^2}\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} + p\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}\widehat{X}_t dt \\ &\quad + \frac{a^2}{m^2}\frac{\sqrt{p^2 + \frac{q^2}{m^2}} - \left(p - \frac{q^2}{a^2}\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} + p\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}dZ_t, \end{aligned} \tag{2.19}$$

for the belief process $\widehat{X}_t$. Derivations are given in Appendix B. Again, note that by expanding dZ_t, the process is expressed in terms of the Brownian motion differential, such that

$$d\widehat{X}_t = -p\widehat{X}_t dt + mL_{2,o}(t)dV_t. \tag{2.20}$$

This shows that the process $d\widehat{X}_t$ takes the same form as the Ornstein-Uhlenbeck process it is intended to estimate, dX_t.

We have so far assumed that the process reverts to zero. However, certain applications require the process to revert to a specific constant μ. Following Hull (2015, Chapter 31.7), without loss of generality, we may shift the process by μ and analyze $\widehat{X}_t + \mu$, while still modeling $\widehat{X}_t$ as an Ornstein-Uhlenbeck process reverting to zero. Note that by shifting the Ornstein-Uhlenbeck in such a way, it effectively becomes structurally equivalent to the model in Vasicek (1977), which allows for mean reversion to a nonzero constant. Furthermore, it is worth noting that by allowing μ to be time dependent such that $\mu = \mu(t)$, we may model observations of processes

that are assumed to have a time-dependent long-run mean as the sum of $\mu(t)$ and a non-shifted Ornstein-Uhlenbeck model. This may for example be applicable to situations in which the process is influenced by seasonal effects.

2.1.5 Comparisons

As evident from equations (2.16) and (2.20), $L_2(t)$ plays a crucial role in how the estimates evolve with time. The component has similar characteristics for both processes. Specifically,

i. $L_2(0) = \frac{a^2}{m^2}$,

ii. $0 \leq \lim_{t\to\infty} L_2(t) < \infty$,

iii. $L_2(t) > 0$ and $L_2'(t) < 0$ when $t > 0$,

iv. a has a positive relationship with $L_2(t)$, and

v. m has a negative relationship with $L_2(t)$.

In short, this means that the uncertainty in the initial estimate is equally large for both processes, and decreases strictly towards zero. The negative gradient illustrates how learning affects the estimate, by allowing for greater certainty as time passes. Since $L_2(t)$ is strictly positive as well as strictly decreasing, the gradient must decrease in absolute magnitude with increasing t, which may be interpreted as a decreasing marginal benefit of additional observations.

Furthermore, it may be shown that $L_{2,c}(t) \geq L_{2,o}(t)$ for $0 \leq t \leq t'$, where t' represents the time of intercept between the curves. The intercept exists due a nonzero limiting value of $L_{2,o}(t)$. Specifically,

$$\lim_{t\to\infty} L_{2,c}(t) = 0, \quad \text{and} \tag{2.21}$$

$$\lim_{t\to\infty} L_{2,o}(t) = \frac{a^2}{m^2} \frac{\sqrt{p^2 + \frac{q^2}{m^2}} - \left(p - \frac{q^2}{a^2}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} + p\right)} > 0. \tag{2.22}$$

This is to be expected, as the uncertainty in the observation process can never be completely eliminated if the observed process is stochastic. Note that if $q = 0$, then $L_{2,o}(t)$ does in fact limit to zero. However, we only consider $p > q > 0$ in this thesis, and the observation process uncertainty is therefore higher for the constant process than the Ornstein-Uhlenbeck process initially, and lower after t'. Whether $t' < T$ depends on parameters.

Considering the fact that m represents the volatility of single observations, as represented in equations (2.13) and (2.17), its negative relationship with $L_2(t)$ may be surprising. However, when the observation process is expressed in terms of the Brownian motion differential, the volatility is given by $mL_2(t)$, and it becomes apparent that m has an increasing relationship with observation process volatility.

2.2 The option to invest

We define the value of the firm's option to invest in the project as $f = f(\widehat{X}_t, t)$, and the project value function to be known, differentiable and strictly increasing in X_t. The firm may decide to invest in, or, "exercise" the option at any time t for a sunk, fixed cost k. The option has a given lifetime T. If the firm does not invest at time $t < T$, the option value continues to evolve. At each point in time, therefore, the firm must choose whether to exercise. At exercise, the firm obtains a net present value of $V(X_t, t) - k$, where V is the present value of the project. As noted in section 1, we consider the firm's utility as given by the CRRA model to be the relevant metric, and instead consider the expected net present utility $\Psi = \Psi(\widehat{X}_t, t)$ at exercise, equal to the expectation of the utility of paying k to obtain V. Consequently, and similar to Hugonnier and Morellec (2007), we seek to identify an investment policy which maximizes the expected net present utility obtained from exercising the option. The "option value" is therefore a quantity that reflects the firm's utility holding the option.[1] Since the value of the option depends on the stochastic belief process $\widehat{X}_t$, and the CRRA model for utility is differentiable and strictly increasing, the optimal investment policy will be characterized by a deterministic threshold $\widehat{X}_t^*$.

2.2.1 Risk aversion and the expected utility function

The CRRA utility function u for a cash flow w is given by

$$u(w) = \begin{cases} \frac{w^{1-\gamma}}{1-\gamma} & \text{when } 0 \leq \gamma < 1, \text{ and} \\ \ln(w) & \text{when } \gamma = 1, \end{cases} \tag{2.23}$$

for $w \in \mathbb{R}$. The utility function is therefore concave in w, as opposed to a risk-neutral situation in which $\gamma = 0$ and the function is linear in w. Following Hugonnier and Morellec (2007) and Conejo, Morales, Kazempour, and Siddiqui (2016, Chapter 7.7), the utility of investing in the

[1] Hereafter, if not explicitly stated otherwise, we consider "value" to refer to the firm's utility of the corresponding cash-based value.

project with present value $V(X_t, t)$, at a cost of k, can therefore be expressed as

$$u\left[V(X_t,t)\right] - u(k) = \frac{V^{1-\gamma}(X_t,t) - k^{1-\gamma}}{1-\gamma}. \tag{2.24}$$

The firm would never accept an investment with a negative expected net present utility. We may therefore formulate the expected net present utility of exercising the option as

$$\Psi(\widehat{X}_t, t) = \max\left\{\frac{\mathbb{E}\left[V^{1-\gamma}(X_t,t)\right] - k^{1-\gamma}}{1-\gamma},\ 0\right\}. \tag{2.25}$$

For simplicity, we will refer to Ψ as the "exercise value" hereafter. The expectation in Ψ requires knowledge about the distribution of X_t. We may use a simulation procedure to estimate this distribution, based on realizations of the observation process. Further details are given in Appendix C.

2.2.2 The Bellman equation

The Bellman equation is a necessary and sufficient condition for optimality in dynamic programming. At each point in time, the firm maximizes the current expected net present utility of either exercising or holding the option. Since the utility function is strictly increasing in $\widehat{X}_t$, the firm will exercise the option if the belief process $\widehat{X}_t$ is greater than or equal to the threshold belief $\widehat{X}_t^*$, and wait otherwise. We define these regions of the belief process as the "exercise region" and "continuation region", respectively. If the firm exercises the option, the option value is equal to the exercise value. If the firm holds the option, however, the option value is equal to the discounted expected future option value, which may be expressed in terms of a differential equation in $f(\widehat{X}_t, t)$. The firm is therefore interested in identifying the threshold belief $\widehat{X}_t^*$ such that it may know whether exercising the option is optimal at a given point in time. Hence, the firm faces a decision problem with a Bellman equation of the form

$$f(\widehat{X}_t, t) = \max_{\widehat{X}_t}\left\{\Psi(\widehat{X}_t, t),\ \frac{f(\widehat{X}_t, t) + \mathbb{E}[df(\widehat{X}_t, t)]}{1+\rho dt}\right\}. \tag{2.26}$$

Equation (2.26) is a sufficient condition to characterize the threshold belief $\widehat{X}_t^*$. Rewriting the expression for $f(\widehat{X}_t, t)$ in the continuation region, we obtain

$$\rho f(\widehat{X}_t, t)dt = \mathbb{E}[df(\widehat{X}_t, t)] \quad \text{when } \widehat{X}_t < \widehat{X}_t^*. \tag{2.27}$$

Note that since $f(\widehat{X}_t, t)$ is defined by $\mathbb{E}[df(\widehat{X}_t, t)]$, $f(\widehat{X}_t, t)$ must be deterministic.[2] In order to determine the value of the option to invest, we consider a general belief process of the form

$$d\widehat{X}_t = L_1(t)\widehat{X}_t dt + L_2(t)dZ_t, \tag{2.28}$$

with an observation process of the form

$$dZ_t = X_t dt + m dV_t, \tag{2.29}$$

we apply Itô's lemma and obtain

$$\frac{\partial f}{\partial t} + [L_1(t) + L_2(t)]\widehat{X}_t \frac{\partial f}{\partial \widehat{X}_t} + \frac{1}{2}m^2 L_2^2(t)\frac{\partial^2 f}{\partial \widehat{X}_t^2} - \rho f = 0 \quad \text{when } \widehat{X}_t < \widehat{X}_t^*. \tag{2.30}$$

Derivations are given in Appendix D.

2.2.3 Valuing the option to invest

The option value $f(\widehat{X}_t, t)$ is constrained by the boundary conditions

$$\lim_{\widehat{X}_t \to -\infty} f(\widehat{X}_t, t) = 0, \tag{2.31}$$

$$f(\widehat{X}_T, T) = \Psi(\widehat{X}_T,\ T), \tag{2.32}$$

$$f(\widehat{X}_t^*, t) = \Psi(\widehat{X}_t^*, t) \quad \text{and} \tag{2.33}$$

$$\frac{\partial f(\widehat{X}_t^*, t)}{\partial \widehat{X}_t} = \frac{\partial \Psi(\widehat{X}_t^*, t)}{\partial \widehat{X}_t}. \tag{2.34}$$

Condition (2.31) states that the option is worthless if the belief process becomes infinitely negative. The belief process being an unbiased estimator with a finite variance, the limit implies that the observed process also approaches negative infinity in expectation. Due to the strictly increasing nature of the exercise value, the option value consequently approaches its minimum value, zero. Condition (2.32) states that the firm makes a now-or-never investment decision at the expiration time T of the option. Equations (2.33)-(2.34) are value-matching and smooth pasting conditions that are necessary to ensure an optimal free boundary, as argued in Dixit and Pindyck (1994, Chapter 3, Appendix C).

[2] As a sidenote, for a given terminal condition $\Psi(\widehat{X}_T, T)$, equation (2.27) may be interpreted as the differential of the Feynman-Kac formula for parabolic partial differential equations, here represented by equation (2.30), as pointed out by Dixit and Pindyck (1994, Chapter 4, Section 3).

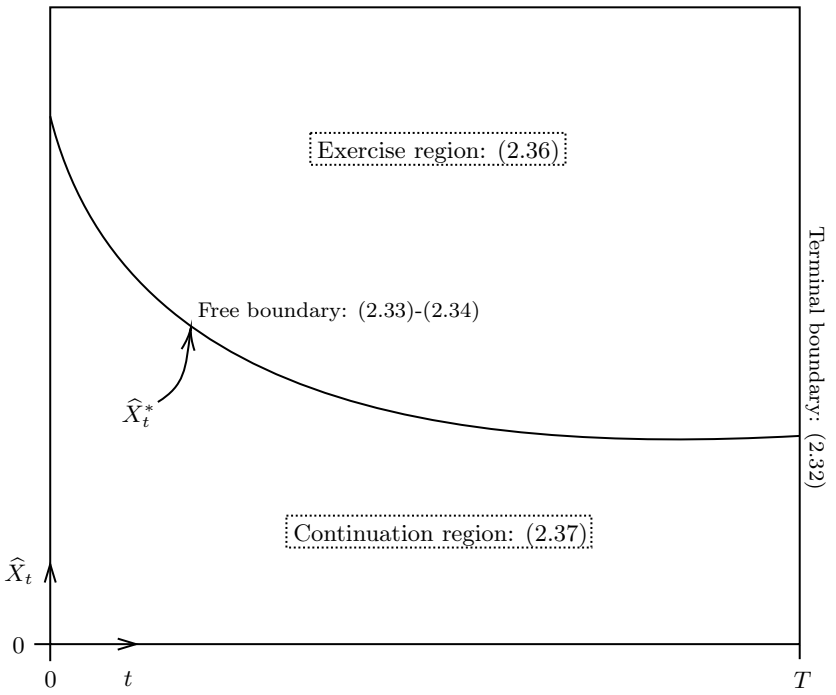

Figure 2.1: Domain of the option value function $f(\widehat{X}_t, t)$.

Figure 2.1 summarizes the boundary value problem. Since the free boundary $\widehat{X}_t^*$ is initially unknown, it must be identified together with the option value function $f(\widehat{X}_t, t)$ in the solution procedure. Rewriting equation (2.26), we obtain

$$\max_{\widehat{X}_t} \left\{ \Psi(\widehat{X}_t, t) - f(\widehat{X}_t, t), \frac{\mathbb{E}[df(\widehat{X}_t, t)]}{dt} - \rho f(\widehat{X}_t, t) \right\} = 0, \tag{2.35}$$

which applies across the entire domain of f. Hence, we seek to identify a free boundary $\widehat{X}_t^*$ that satisfies

$$\Psi(\widehat{X}_t, t) - f(\widehat{X}_t, t) = 0 \quad \text{when } \widehat{X}_t \geq \widehat{X}_t^*, \text{ and} \tag{2.36}$$

$$\frac{\mathbb{E}[df(\widehat{X}_t, t)]}{dt} - \rho f(\widehat{X}_t, t) = 0 \quad \text{when } \widehat{X}_t < \widehat{X}_t^*, \tag{2.37}$$

for all $t \in [0,\ T]$. These constraints, together with boundary conditions (2.31)-(2.34), are sufficient to identify the free boundary and the option value function.

3 Solution approach

As Longstaff and Schwartz (2001) point out, "[...] the valuation and optimal exercise of American options remains one of the most challenging problems in derivatives finance, particularly when more than one factor affects the value of the option." Although the option considered in this thesis only contains one factor, the nature of the belief process introduces a time-dependence in the option value that arguably makes the problem more mathematically complex than a multi-factor model, as the PDE governing the option value is time-inhomogeneous.

The boundary value problem introduced in subsection 2.2 is such a problem. As the PDE (2.30) introduced in subsection 2.2.2 shows, the equation governing the option value contains a derivative in time, and a first and second derivative in space with coefficient terms both in time and space. Due to the coefficients being bivariate, there are no clear variable substitutions to simplify the problem. There exist a range of numerical techniques to handle situations in which an analytical solution is intractable. Finite difference methods, for example, segment the option domain into discrete elements and express the derivatives of the PDE as linear approximations. The transformed equation is then expressed as a matrix equation which is solved recursively in time, starting at the terminal boundary. Implicit difference schemes provide the greatest numerical stability, examples of which include the implicit Euler scheme and the Crank-Nicolson scheme. Following Thijssen (2020), in order to identify the free boundary, an initial boundary would be postulated, and the total absolute deviations from conditions (2.36) and (2.37) would be checked against a specified threshold. If this error term is too large, a new boundary would be specified such that the conditions hold, and the procedure would repeat. Upon implementation, the implicit Euler and Crank-Nicolson schemes proved to give non-converging results. This is thought to be caused by the aforementioned time-dependence of the PDE coefficients. We have included our modeling approach of the Crank-Nicolson scheme in Appendix E for reference. In formulating the temporal matrix equations, one will note that the coefficients of the discretized option function change in time. As illustrated in Kushner and Dupois (2001, Chapter 5.1), for continuous stochastic control problems with varying coefficients, an adaptive grid should be applied such that the size of the time step also depends on time. Their example does not allow for a time-varying volatility component in the underlying process, however, and it falls outside the area of competence of this author to apply stochastic control theory in order to derive a solution for such a case. Another approach is to apply a method of lines, which only discretizes the PDE in time and solves the remaining ODE recursively, using the previous solution as a boundary condition. The procedure would apply conditions (2.33)-(2.34) at each time step in order to

identify the free boundary, and would therefore not require the repeating procedure described above. This method lead to promising option surfaces for small time intervals. However, with increasing iterations, the boundary condition becomes increasingly complex, and the resulting ODE too time consuming to solve, either analytically or numerically.

We have consequently decided to use a simulation-based approach by applying the method introduced in Longstaff and Schwartz (2001), in which various American options are valued by first simulating a series of trials of the underlying stochastic process, then obtaining stopping rules for each trial, and finally averaging over the discounted exercise values. The model begins by first calculating the exercise values at the terminal boundary T, and thereafter stepping backwards in time, using previously discounted exercise values as a regressand, and current belief process values as a regressor, in order to obtain least-squares coefficients of a conditional expectation function for the continuation value of the option at the given point in time, which is then compared to the value of immediate exercise. As a useful comparison, this method may be viewed as analogous to the decision problem presented by the Bellman equation in subsection 2.2.2, in which the expected continuation value is compared to the exercise value in continuous time. As noted by Longstaff and Schwartz (2001), the functions used in the regression and the resulting conditional expectation function need to form an orthonormal basis. By increasing the number of orthonormal basis functions, the accuracy of the option value improves. If we denote Δt and $N = \lfloor T/\Delta t \rfloor$ as the magnitude and number of time-steps, respectively, the regression equation at time-steps $j \in \{1,\ \ldots,\ N-1\}$ may be represented as

$$Y(\omega_i, j\Delta t) = \sum_{k=0}^{B} a_{n,j} \mathcal{L}_n(\widehat{X}_j), \tag{3.1}$$

where ω_i indicates the ith trial, Y indicates the regressand, B denotes the number of basis functions, $\{a_{n,j}\}$ the regression coefficients and $\mathcal{L}_n$ the nth order basis function. Further details are provided in Appendix F.

3.1 Simulation of the belief process by the Euler method

We simulate M trials of the belief process by the Euler method, discretizing the expanded SDE (2.28), such that

$$\widehat{X}_j = \widehat{X}_{j-1} + [L_1(t-\Delta t) + L_2(t-\Delta t)]\widehat{X}_{j-1}\Delta t + mL_2(t-\Delta t)\sqrt{\Delta t}\ \zeta_j, \tag{3.2}$$

where $j = \{1, \ldots, N\}$ and $\zeta_j \sim \mathcal{N}(0,1)$. All trial paths begin at the initial belief $\widehat{X}_0$. In order to reduce the time complexity of the algorithm, we apply antithetic variates when sampling the standard normal distribution. For each trial i with a given vector of realized standard normal variables $\{\zeta_{i,j}\}_j$ we design a second trial i' with antithetical realizations $\{-\zeta_{i',j}\}_j$. This results in a negative covariance between path values of trials i and i' at any given j, which, when applied to all trials, reduces the variance across all path values at j, resulting in a lower required number of trials for a desired level of accuracy.

3.2 The Longstaff-Schwartz algorithm

Following Longstaff and Schwartz (2001), we first order the simulated trials of the belief process in a matrix $\mathbb{Q}$ with dimensions $M \times N$, such that each row of the matrix represents the values taken by a given trial for all $j = \{1, \ldots, N\}$. We also initiate $\mathbb{C}$ as a matrix of zeros of the same size to store the exercise values of the option at a given belief and time, using the path values in $\mathbb{Q}$. At $j = N$, $\mathbb{C}$ is set equal to the exercise value. The algorithm then proceeds accordingly:

At $1 < j < N$:

i. Trials that have a nonzero exercise value at j are selected from $\mathbb{Q}$.

ii. For each selected trial, the discounted value of the first nonzero element in $\mathbb{C}$ at times $s > j$ is stored in a vector.

iii. This vector is then used as a regressand in equation (3.1), with the corresponding vector of belief process values at j as a regressor.

iv. The regression results in a vector of estimated coefficients, which are substituted into equation (3.1) in order to obtain the conditional expectation of the continuation value at j.

v. The continuation values obtained for each trial at j are then compared to the corresponding exercise values at j. If the exercise value exceeds the continuation value, the exercise value is stored in $\mathbb{C}$, and any nonzero values at times $s > j$ are set to 0.

At $j = 1$:

$\mathbb{C}$ now contains stopping times and corresponding exercise values for each trial if the option has in fact been exercised in that trial. The option value f_0 at $t = 0$ is estimated as the expected sum of discounted exercise values over all paths.

At each time-step, the free boundary $\widehat{X}_t^*$ is estimated by solving for the root of the regression equation (3.1) after coefficients $\{\boldsymbol{a}_j\}$ have been replaced with their respective estimations $\{\hat{\boldsymbol{a}}_j\}$, following Longstaff and Schwartz (2001). This is done numerically, as the resulting function may contain polynomials of a high order. A smoothing technique is then applied to obtain a graph from the estimated free boundary points. Further details are given in Appendix F, together with our full implementation of the algorithm.

4 Results

This section illustrates the applicability of the approach in section 3 through a case study. We consider a situation in which the firm is planning the launch of a product. The prospective product is assumed to be categorized as part of an industry with product prices that are mean reverting and may be modeled by a shifted Ornstein-Uhlenbeck model around a mean μ known to the firm. We include a model that assumes a constant price as a comparison. At the beginning of the planning horizon, the firm holds a prior normal distribution of the price of its product based on prices of similar products in the industry. The firm then proceeds to observe product prices continuously throughout the option holding period in order to obtain a better estimate at the impending time of launch. We assume that the firm obtains a one-time payment that depends linearly on the product price at the time of launch, and that the cost of launch is fixed throughout the holding period and paid in a single transaction.

With the given assumptions, we may express the project value function V as

$$V\left(X_t\right) = c\left(X_t + \mu\right), \tag{4.1}$$

where $c > 0$ is the quantity (Q) sold at launch, X_t is expressed in terms of unit price, and μ is the deterministic long run mean of the price process. The expected net present utility function is consequently given as

$$\Psi(\widehat{X}_t, t) = \max\left\{\frac{1}{1-\gamma}\left[\mathbb{E}_t\left(\left[c\left(X_t+\mu\right)\right]^{1-\gamma}\right) - k^{1-\gamma}\right],\ 0\right\}. \tag{4.2}$$

We apply the parameter values in Table 4.1 as a base, and proceed to look at how different properties change with variations of specific parameters.[3]

We separate our discussion into a risk neutral and risk averse case, respectively. The risk neutral case removes a layer of complexity by allowing for a simpler exercise value function, and is included to better illustrate the properties of the model. The risk averse case is then discussed in terms of deviations from the risk neutral case. We focus our attention on the free boundary, the value of waiting to learn, and expected exercise times.

[3]For comparability, we have used the same random number generator seed in all sections, specifically, `rng(1108)` in MATLAB v. R2020a.

Table 4.1: Base parameter values under risk neutrality and risk aversion, respectively.

Parameter	Value	Unit	Description
T	5	Years	Option lifetime
k	5 000	MNOK	Investment cost
c	100	MQ	Proportionality of project value function
γ	0 (0.1)	-	Relative risk aversion
ρ	0.05	-	Rate of time preference (discount rate)
$\widehat{X}_0$	40	NOK/Q	Initial belief
a	60	NOK/Q	Volatility of initial belief
m	20	NOK/Q	Volatility of observations
p	0.05	Years^{-1}	Mean reversion of observed process
q	0.02	NOK/Q	Volatility of observed process
μ	40	NOK/Q	Mean of observed process
$\bar{\epsilon}$	0.1%	-	Error tolerance
Δt	0.05	Years	Magnitude of time discretizations
Δx	(0.2)	NOK/Q	Magnitude of integral eval. discretizations
M	500 000	-	Number of trials
B	5 (6)	-	Number of Laguerre basis functions
η	16 (10)	-	Order of polynomial smoothing function

4.1 Under risk neutrality

As discussed in subsection 2.2.1, the CRRA model under risk neutrality is modeled with $\gamma = 0$. This simplifies the structure of the exercise value function and clarifies certain characteristics of the option. Specifically,

$$\Psi(\widehat{X}_t) = \max\left\{c\left(\widehat{X}_t + \mu\right),\ 0\right\}, \tag{4.3}$$

which is equivalent to the expected net present value of investing.

We begin with selecting the number of Laguerre basis functions B to regress on by analyzing the relative changes in option values for increasing B. Selecting $B = 5$ ensures that the simulated option value changes by less than $\bar{\epsilon} = 0.1\%$ when $B = 6$, which we have deemed a high enough accuracy for the purpose of this case study. An illustrating figure is provided in Appendix G. It should be noted that computational limitations such as the number of trials M and the magnitude of discrete time steps Δt naturally also affect the accuracy of the results. We have consequently selected M iteratively, such that for $\Delta t = 0.05$ and $\Delta x = 0.2$, $M = 500\ 000$ trials also gives an option value that changes by less than $\bar{\epsilon} = 0.1\%$ when using $M + 1$ trials.

Figure 4.1 illustrates the free boundary for both the shifted Ornstein-Uhlenbeck process and the constant process, with initial values $\widehat{X}_0$ equal to the long-term average μ of the processes. The discrete points have been smoothed by a polynomial function.[4] The boundaries are similar in shape for both processes. The decreasingly negative gradient may be attributed to a diminishing value of waiting to learn, since the observer obtains better estimates of the underlying process as time goes by, reflected by the decreasing marginal benefit of new observations as outlined in subsection 2.1.5. Since the value of holding the option decreases, the observer requires lower thresholds to invest.

With a finite lifetime, the option is terminated before the observer has decreased the estimator uncertainty to its limiting value, or, before the observer has realized the full value of waiting to learn. We would therefore expect a discontinuity in the boundary at T such that $\lim_{t \to T^-} \widehat{X}^*_t \neq \widehat{X}^*_T$, since the terminal boundary condition is a now-or-never decision. This discontinuity is thought to be the source of some numerical issues that occur close to the terminal boundary. We have therefore plotted the boundaries for $t \leq 4.5$, and stress that the characteristics and dynamics of the boundary are obtainable from examining this region alone. Note that for an infinite lifetime $T \to \infty$, following the discussion in subsection 2.1.5, we would expect the free boundary to converge smoothly towards the terminal boundary for the constant process, since the marginal benefit of new observations approaches zero. For the Ornstein-Uhlenbeck, however, we would expect to see this discontinuity regardless of option lifetime, since the underlying process has a stochastic component. Intuitively, an investor who engages in noisy observations of a stochastic process always has more to learn, due to a process uncertainty that never disappears.

The shifted Ornstein-Uhlenbeck process has a lower free boundary than the constant process, which may be counterintuitive at first sight. We point out that $t' > T$ for all parameter values, which implies that the marginal benefit of learning is greater for the constant process than the Ornstein-Uhlenbeck process throughout the option lifetime.[5] Hence, the difference between the boundaries is thought to be due to a lower uncertainty in its estimate, as the process reverts to a known value with time, while the constant process has no such convergence.

A quantity of interest to the observer is the value of waiting to learn, as opposed to investing immediately. We may quantify this as the difference between the option value f_0 and the exercise value at $t = 0$. Figure 4.2 shows the expected value of waiting to learn for an increasing difference between the estimated initial values $\widehat{X}_0$ and μ, with the latter held constant. The relationship

[4] We have selected a polynomial order of $\eta = 16$ for the smoothing function in the risk neutral case. The discrete points will be included in the initial free boundary plot to illustrate the fit, but will be omitted for convenience in later sections.

[5] This also applies in the risk averse case of subsection 4.2.

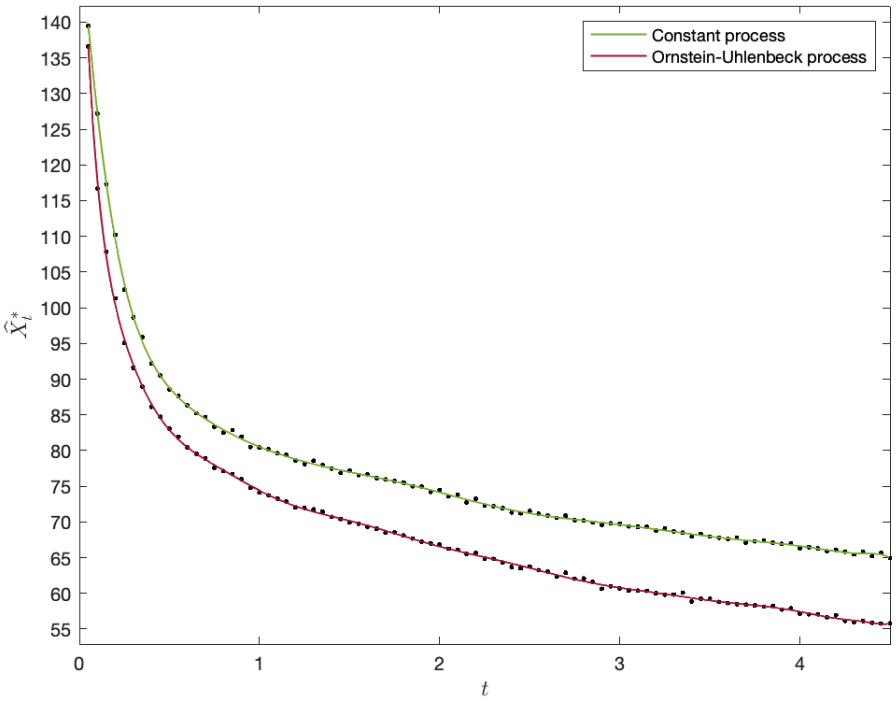

Figure 4.1: Exercise thresholds.

is seemingly linear, implying that the value of waiting to learn decreases with a higher initial estimate of the underlying process. The option value alone has an increasing relationship with $\widehat{X}_0$, but as the figure shows, this trend is offset by an increase in the value of immediate exercise. This becomes more intuitive when viewed in the context of our case study. If the firm has an initial estimate that is high, its expected value of launching early will also be high, and it will seem more lucrative to invest immediately than to hold off the launch.

The vertical gap between the processes is thought to arise because of the difference in magnitude between the volatilities of the processes, as discussed in subsection 2.1.5. The constant process has a higher volatility than the Ornstein-Uhlenbeck process throughout the holding period, and the value of waiting to learn is consequently higher.

Figure 4.3 shows the expected value of waiting to learn against the standard deviation – or, uncertainty – a of the initial estimate. The relationship is linear here as well, but increasing, which implies that the value of waiting to learn increases with uncertainty at the beginning of the holding period. This result is comparable to those of standard real options models, in which option values generally increase with underlying process uncertainty.

Figure 4.4 shows the distribution of exercise times conditioned on the trial having been exercised before option termination. With base parameter values, approximately 40.7% and 44.0% of the

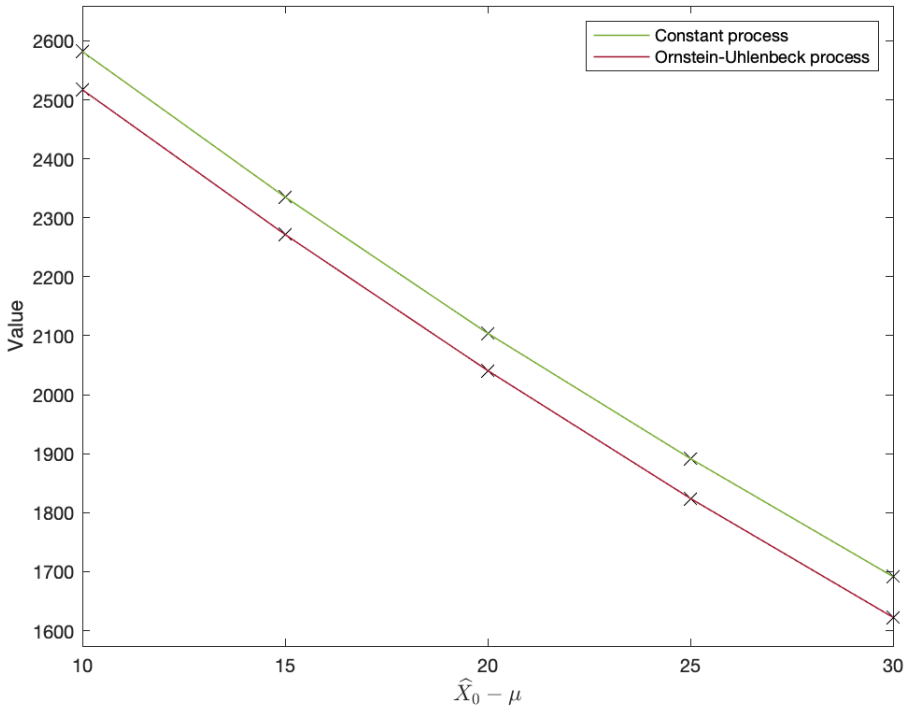

Figure 4.2: Values of waiting to learn against $\widehat{X}_0 - \mu$.

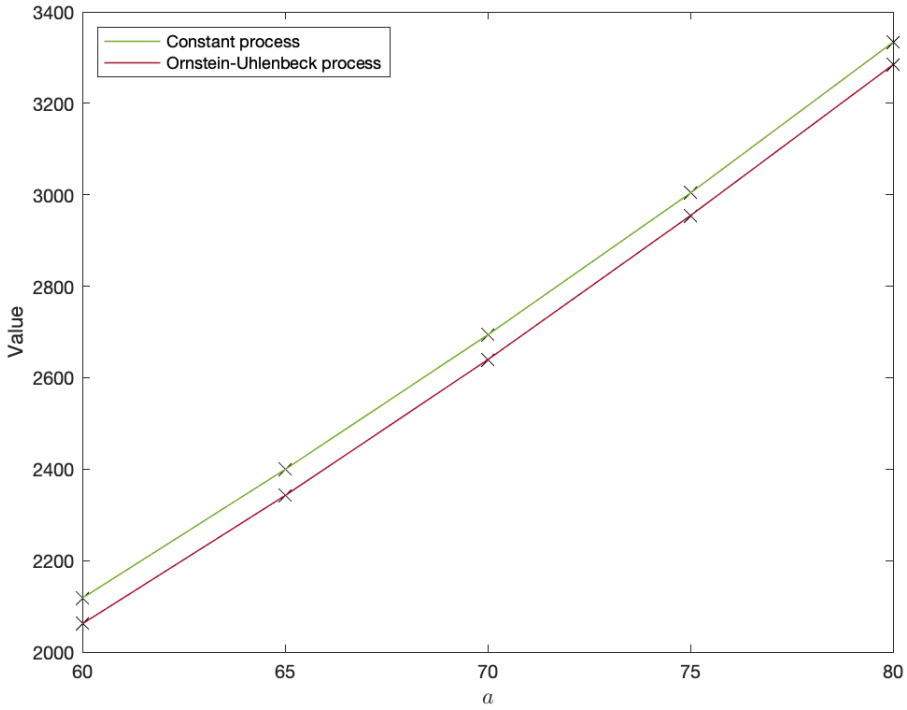

Figure 4.3: Values of waiting to learn against a.

trials were exercised for the constant process and Ornstein-Uhlenbeck process, respectively. The histogram shows a clear trend of early exercise, with medians 0.25 and 0.30. This result is more intuitive when viewed in tandem with the discussion in subsection 2.1.5, which examines the properties of the volatilities of the observation processes. Due to the structure of the volatilities, we can conclude that most of the learning happens early in the holding period. The remaining benefit of waiting, as opposed to exercising, is therefore generally at its highest early in the holding period. Notice that the histograms do not start out at their maximum values, but rather increase after some time has passed since $t = 0$. This illustrates how the investor is generally too uncertain at the outset of the holding period to forego the option to wait and learn for the base parameter values. With a high initial marginal benefit of learning, however, the estimations improve quickly, and the remaining benefit of waiting decreases, which is thought to explain the skewness.

With such a heavy skewness there is a significant distance between the means and medians. This is arguably not highly relevant for the firm contemplating a product launch, but it can be an important observation to an institution that makes policy decisions based on market trends, for example, to not only consider the mean time to investment.

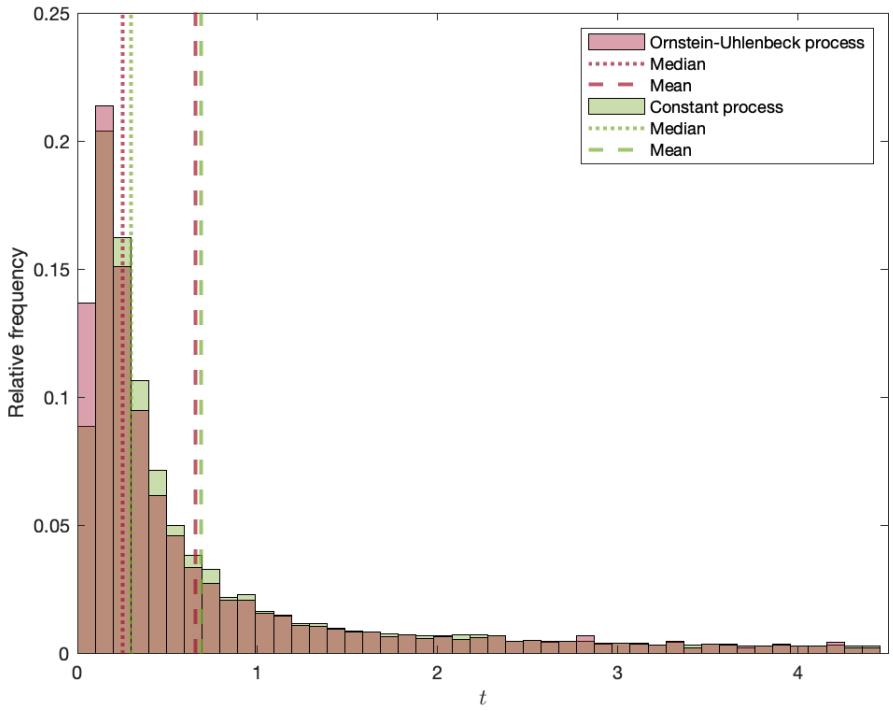

Figure 4.4: Distribution of exercise times conditioned on exercise before T, with corresponding medians and means.

We would expect a mean and median further to the right for lower $\widehat{X}_0$ since the expected value of immediate exercise is lower, as discussed in relation to figure 4.2. This trend can be seen in figure 4.5, with corresponding means and medians in figure 4.6.

Figure 4.7 shows how the distributions change with increasing uncertainty in the initial estimate, with corresponding means and medians in figure 4.8. Interestingly, even though the value of waiting to learn increases with higher a as discussed in relation to figure 4.3, the means and medians show the opposite trend, indicating that the skewness in fact shifts leftwards with higher a. This is thought to be a consequence of the positive relationship between the initial uncertainty and the initial marginal benefit of learning. With a higher initial uncertainty, the firm is more susceptible to change its estimate and hence has a steeper gradient in the volatility of observations, as early observations influences the firm to a greater extent than they would with a lower a. We therefore expect the remaining value of waiting to learn to decrease faster over time with higher a, and consequently result in earlier exercise times.

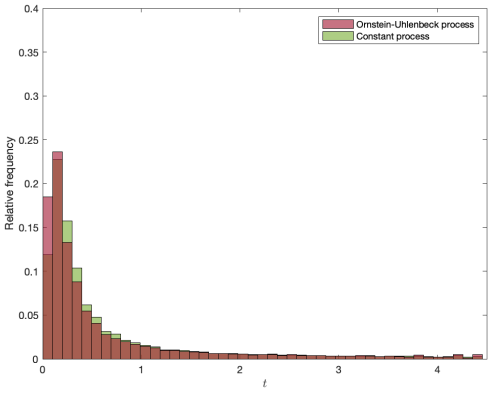

(a) $\widehat{X}_0 - \mu = 10$

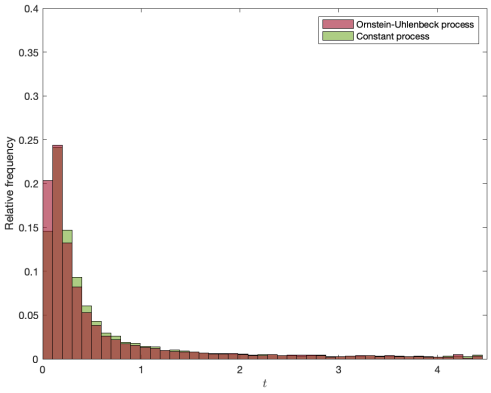

(b) $\widehat{X}_0 - \mu = 15$

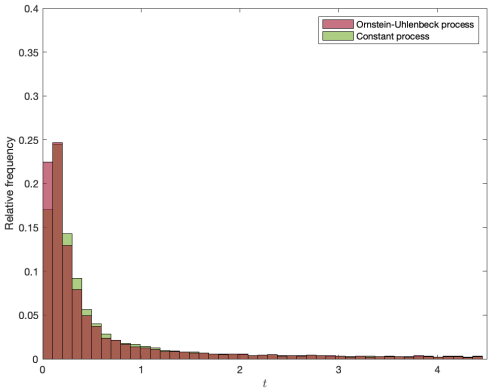

(c) $\widehat{X}_0 - \mu = 20$

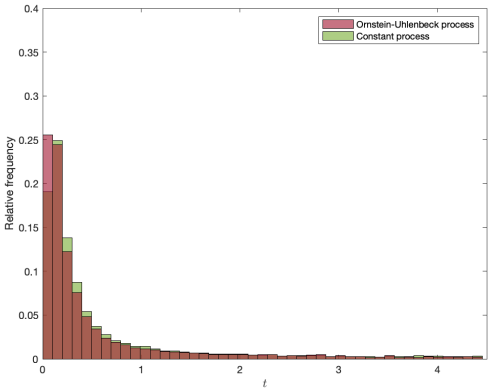

(d) $\widehat{X}_0 - \mu = 25$

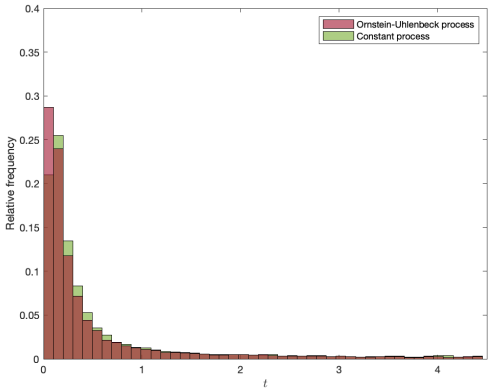

(e) $\widehat{X}_0 - \mu = 30$

Figure 4.5: Changes in the distribution of exercise times under increasing $\widehat{X}_0 - \mu$.

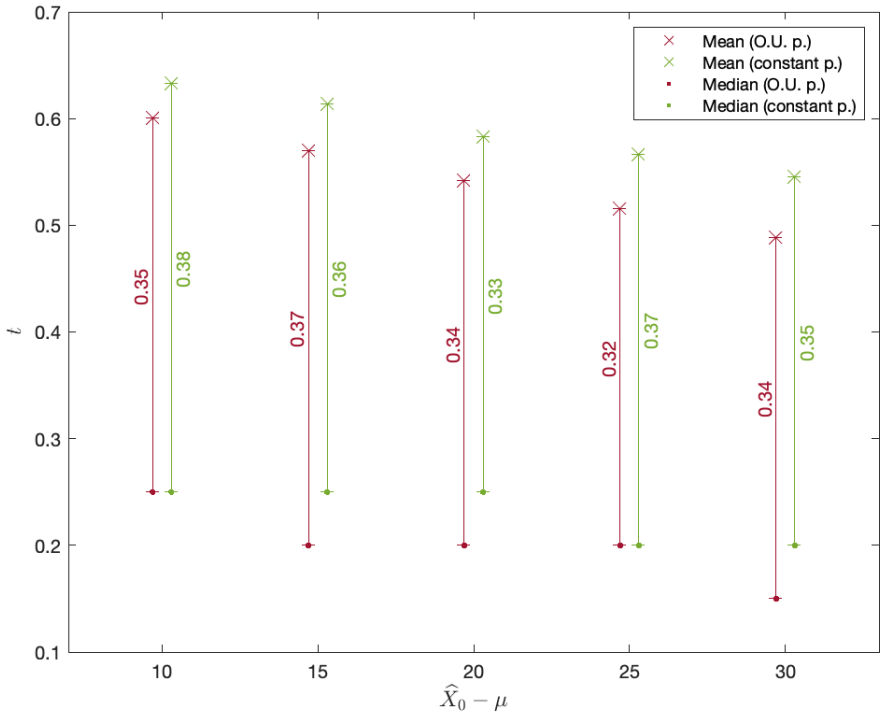

Figure 4.6: Means and medians of exercise times conditioned on exercise before T, against $\widehat{X}_0 - \mu$.

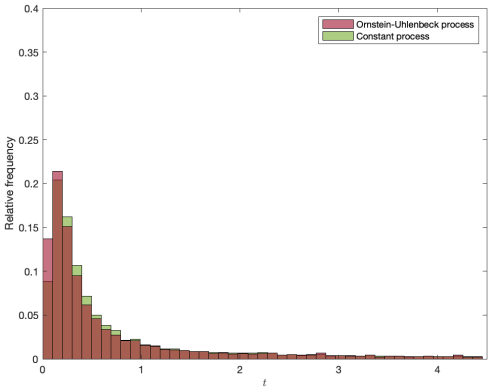

(a) $a = 60$

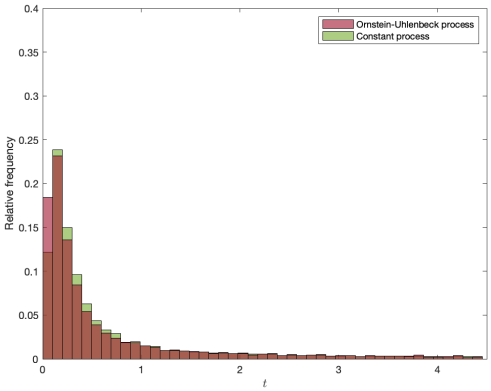

(b) $a = 65$

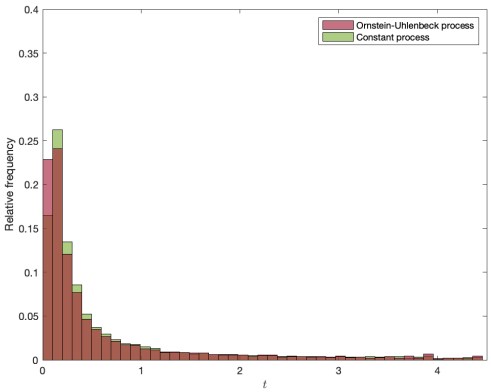

(c) $a = 70$

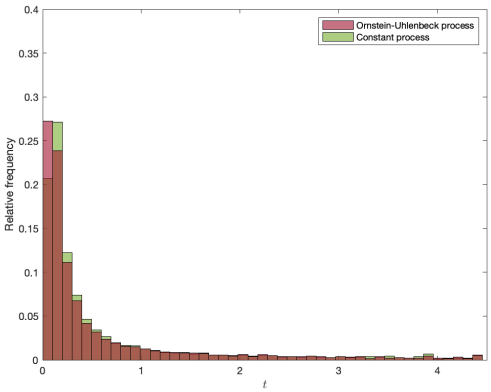

(d) $a = 75$

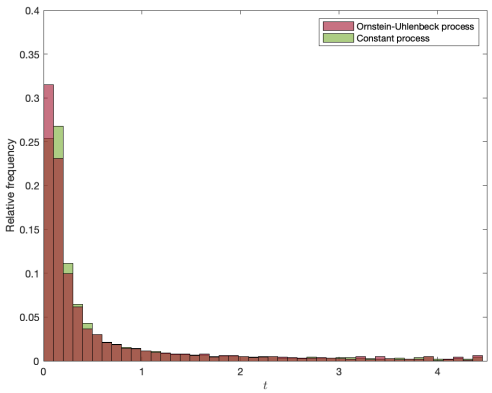

(e) $a = 80$

Figure 4.7: Changes in the distribution of exercise times under increasing a.

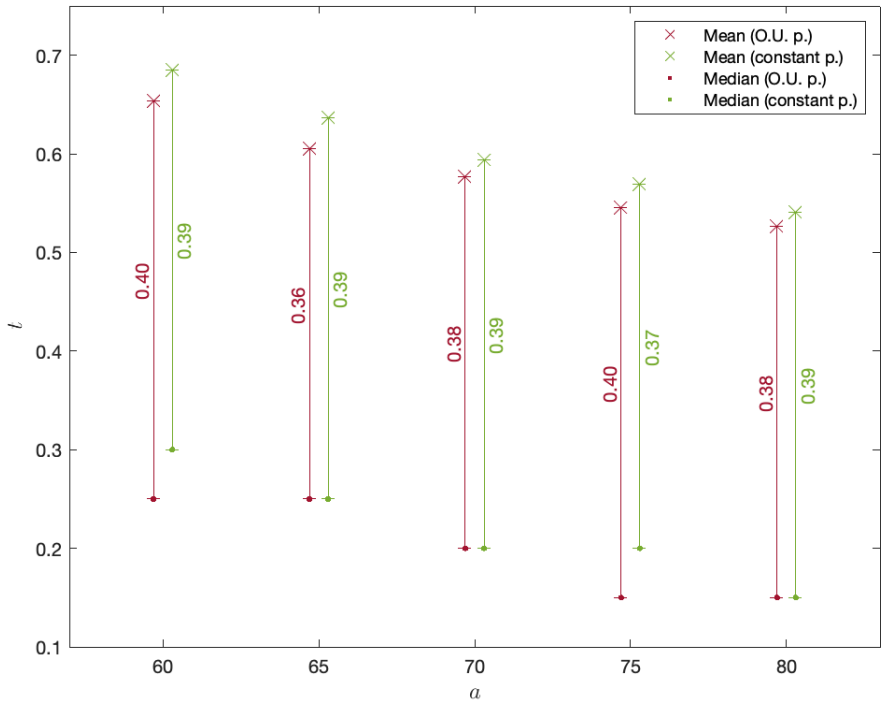

Figure 4.8: Means and medians of exercise times conditioned on exercise before T, against a.

4.2 Under risk aversion

When the investor is risk averse, $\gamma > 0$ and the exercise value function is derived by total expectation, expressed as

$$\Psi(\widehat{X}_t, t) = \max\left\{ \frac{1}{1-\gamma} \left[\int_{\mathbb{R}} \left[c\,(x+\mu) \right]^{1-\gamma} \widehat{\mathbb{P}}(X_t = x \mid \mathcal{G}_t) dx - k^{1-\gamma} \right],\ 0 \right\}, \tag{4.4}$$

where $\widehat{\mathbb{P}}$ is the estimated Gaussian distribution of X_t given previous observations, with mean $\widehat{X}_t$ and variance estimated by the simulated trials, as detailed in Appendix C. Note that Ψ now has a direct time dependence, as opposed to the risk neutral case, due to the time dependence of the variance in the probability distribution. We again analyze the number of necessary basis functions, and find that the risk averse case requires $B = 6$ to satisfy $\bar{\epsilon}$, as illustrated in Appendix G.

The free boundary in the risk averse case shows some fluctuations for $t \gtrsim 1$.[6] This is thought to be caused by the inaccuracy arising from the approximation of the exercise value function. The boundaries display similar characteristics as they do under risk neutrality, with the exception of having lower values for early t. Interestingly, this implies that risk aversion decreases the early investment threshold, compared to the risk neutral case. Hugonnier and Morellec (2007) show the opposite trend for a case in which the investor is risk averse, but active learning is not incorporated. Hence, this result is believed to illustrate that the value of learning decreases under risk aversion, effectively outweighing the trend of an increasing boundary.

This conclusion is supported in part by figures 4.10 and 4.11, which show the initial value of waiting to learn against $\widehat{X}_0 - \mu$ and in a, respectively. The curves are greater in magnitude under risk neutrality, implying that the initial value of waiting to learn is in fact less valuable under risk aversion than under risk neutrality in our case study. Additionally, the curves are more sensitive towards the parameters due to greater slopes. The initial distribution is therefore generally less critical to the value of waiting to learn under risk aversion than under risk neutrality with equivalent parameters.

Figure 4.12 displays the histogram of exercise times for base parameters, which is largely similar to the distribution under risk neutrality, with 41.3% and 44.0% exercised before T for the constant process and Ornstein-Uhlenbeck process, respectively. However, early exercises are more evenly distributed in the risk averse case. We see the same trends as under risk neutrality with regards to sensitivity in $\widehat{X}_0 - \mu$ and a, as displayed in figures 4.13 to 4.16.

[6] We have applied a polynomial order of $\eta = 10$ for the smoothing function under risk aversion.

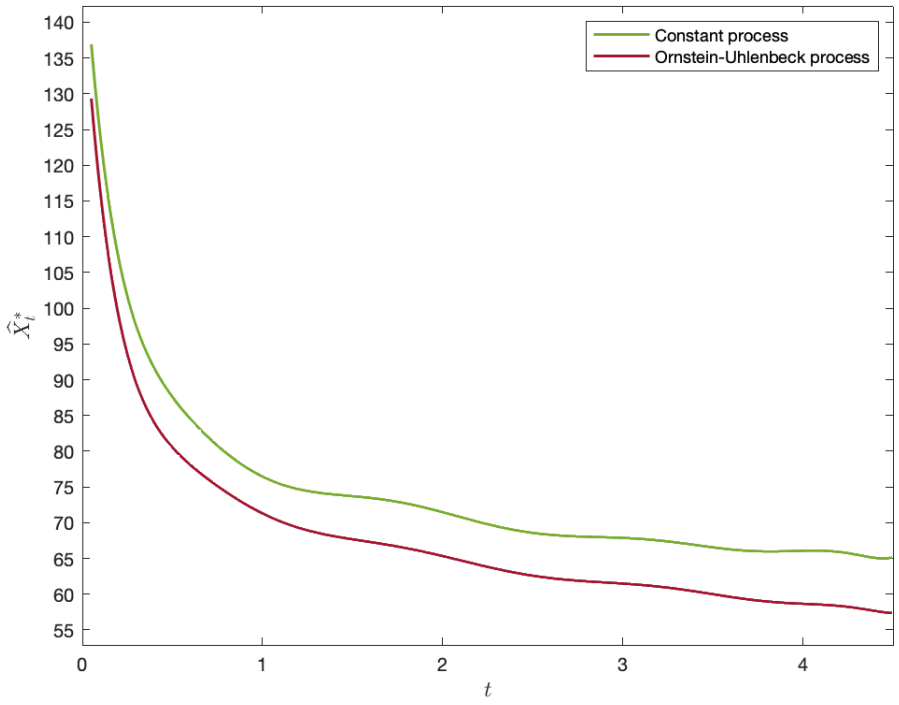

Figure 4.9: Exercise thresholds.

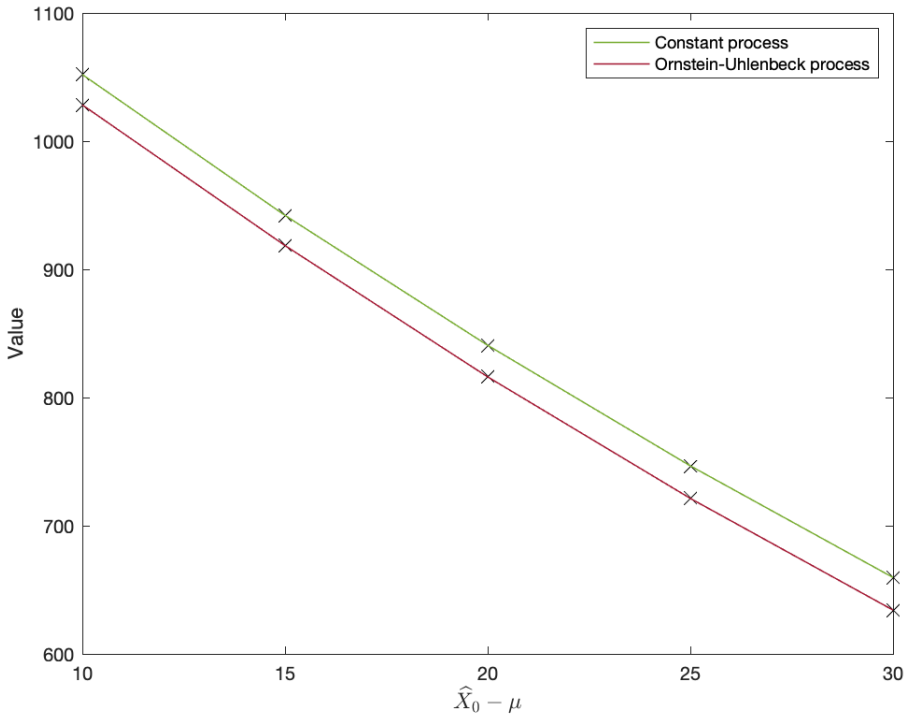

Figure 4.10: Values of waiting to learn against $\widehat{X}_0 - \mu$.

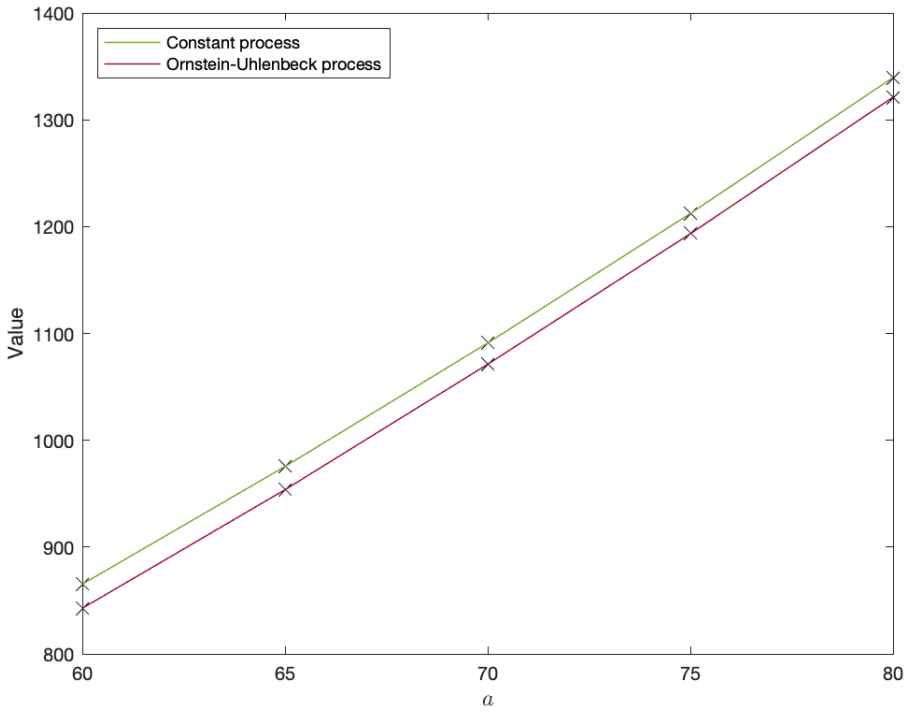

Figure 4.11: Values of waiting to learn against a.

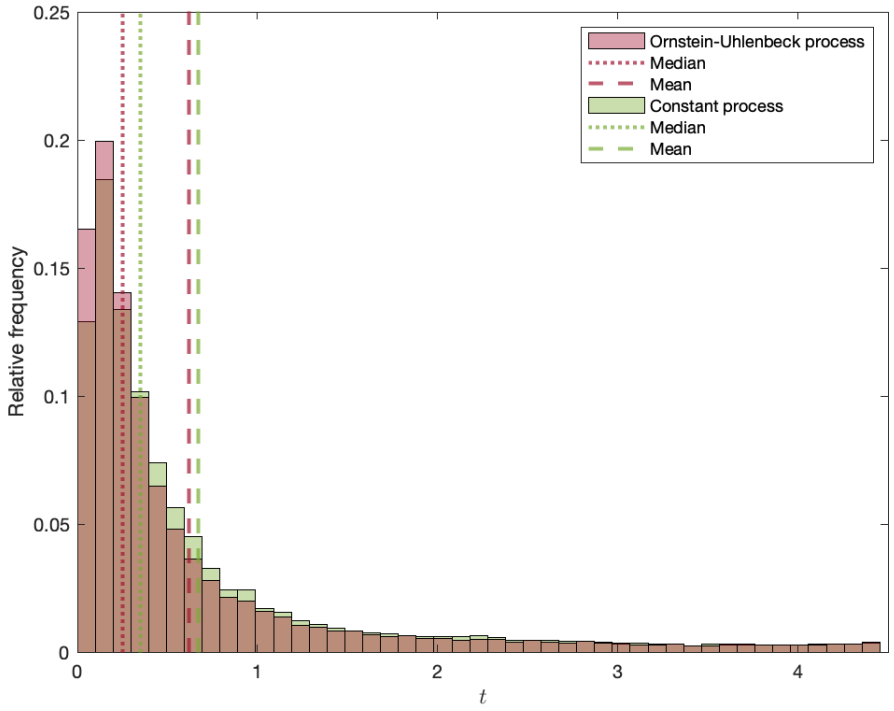

Figure 4.12: Distribution of exercise times conditioned on exercise before T, with corresponding medians and means.

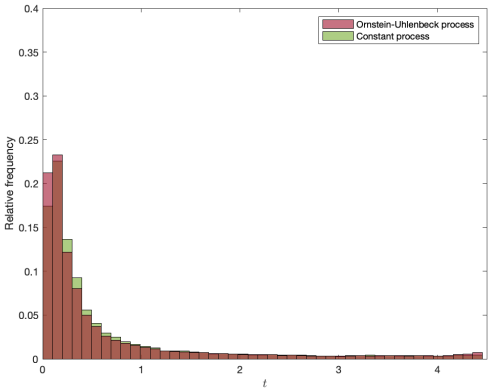

(a) $\widehat{X}_0 - \mu = 10$

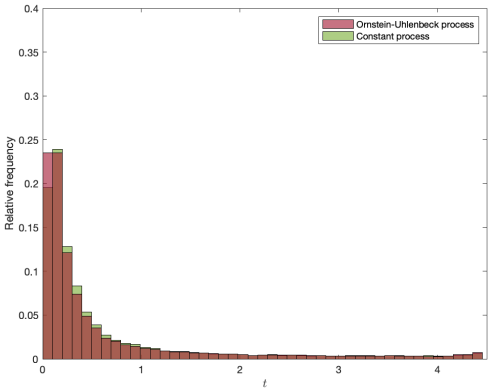

(b) $\widehat{X}_0 - \mu = 15$

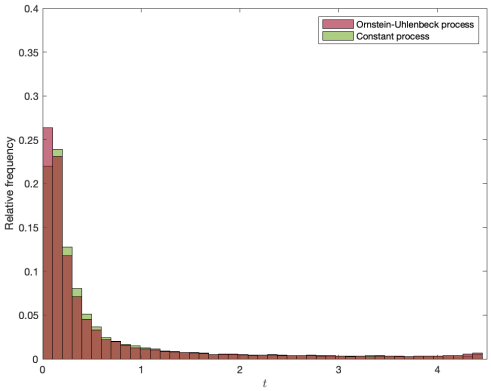

(c) $\widehat{X}_0 - \mu = 20$

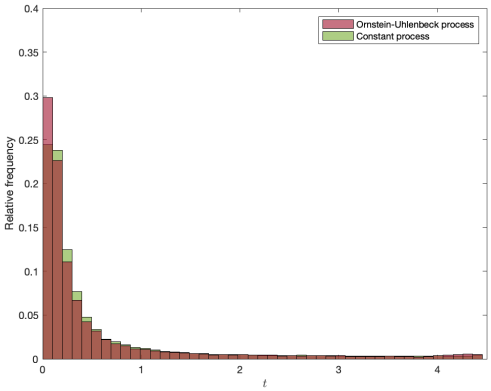

(d) $\widehat{X}_0 - \mu = 25$

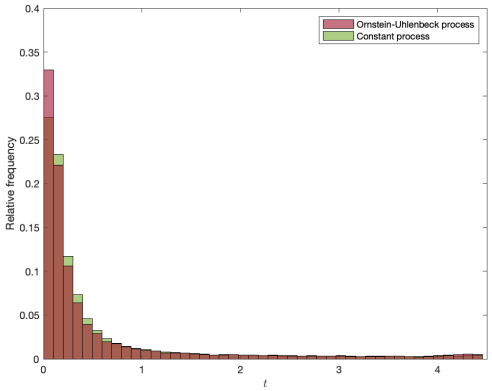

(e) $\widehat{X}_0 - \mu = 30$

Figure 4.13: Changes in the distribution of exercise times under increasing $\widehat{X}_0 - \mu$.

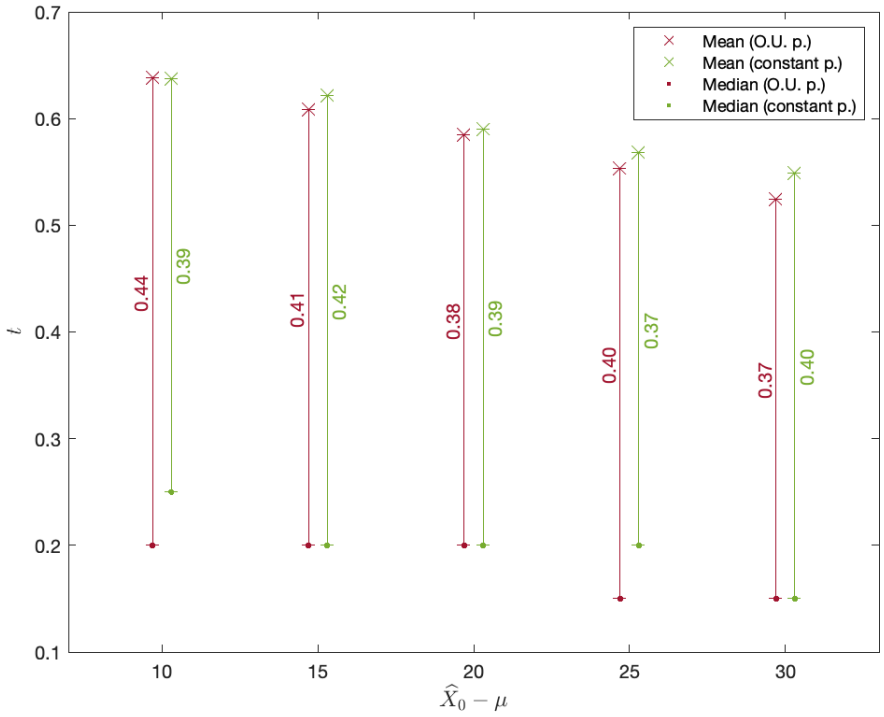

Figure 4.14: Means and medians of exercise times conditioned on exercise before T, against $\widehat{X}_0 - \mu$.

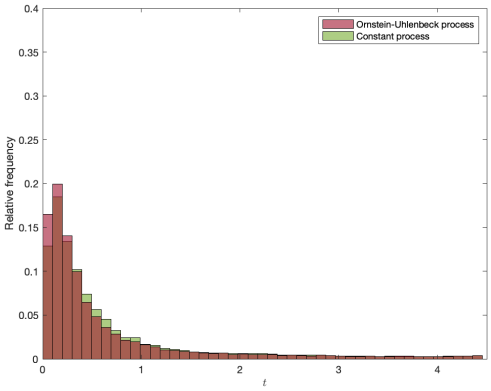

(a) $a = 60$

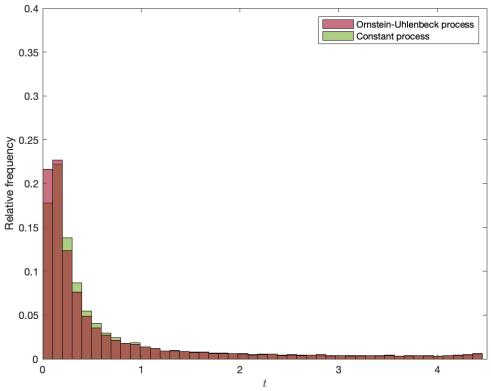

(b) $a = 65$

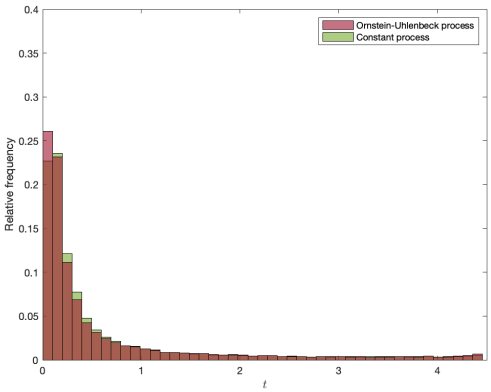

(c) $a = 70$

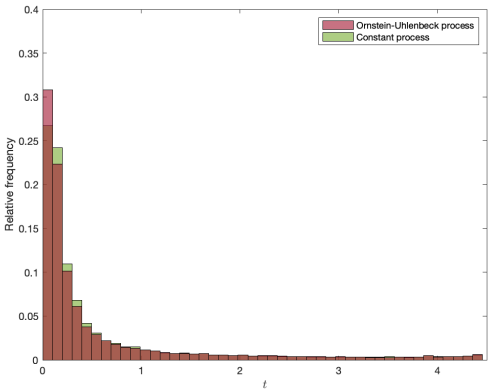

(d) $a = 75$

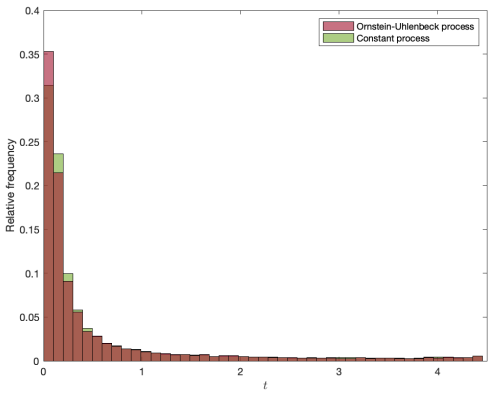

(e) $a = 80$

Figure 4.15: Changes in the distribution of exercise times under increasing a.

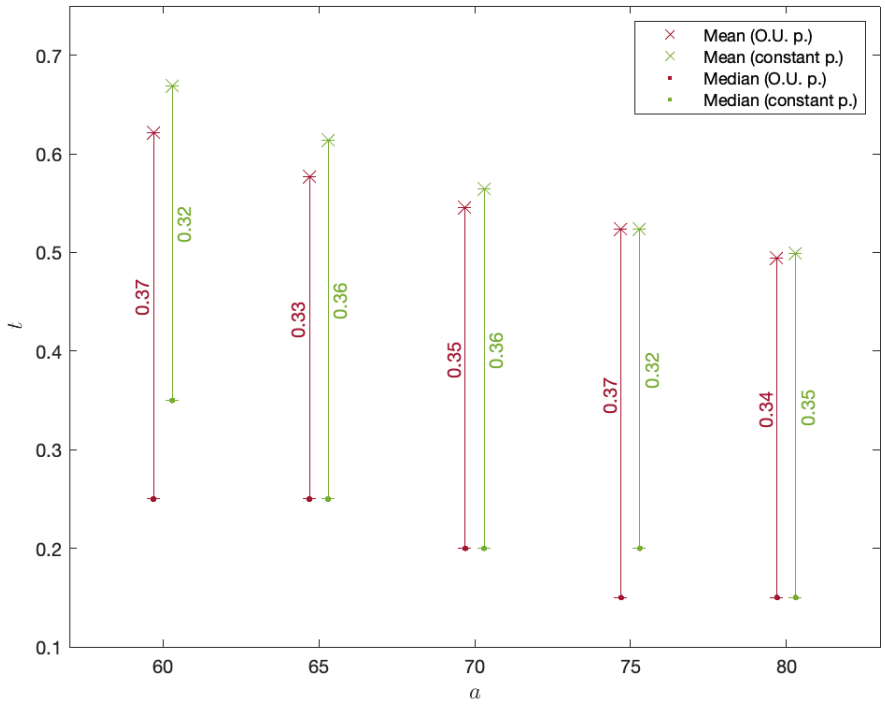

Figure 4.16: Means and medians of exercise times conditioned on exercise before T, against a.

5 Conclusion

This thesis extends real options literature by introducing the flexibility to learn about a stochastic process underlying the project value with a risk averse decision maker. We demonstrate how the exercise threshold changes with the marginal value of learning, and how the marginal value of learning is consistently higher for a constant process than an Ornstein-Uhlenbeck process, within an initial time interval. Numerical results indicate that the value of waiting to learn has a negative relationship with the initial process estimate, and a positive relationship with the uncertainty of its distribution. We show how the skewness of the distribution of exercise times shifts leftwards when increasing the initial estimate and the uncertainty of the initial distribution, respectively, even if the initial value of waiting to learn increases in the latter case over the same interval. We also show how the value of learning seemingly decreases when the investor is assumed to be risk averse as opposed to risk neutral.

We identify three potential areas of further research. The first is to relax the continuity assumption of observations, allowing the decision maker to incorporate observations at discrete times. This allows the model to be applied to settings in which observations happen less frequently. The second is to let the project value depend on multiple stochastic processes, all subject to noisy observations. This opens a range of possible extensions such as, for example, a stochastic investment cost that is not directly observable. Our final suggestion is to apply the model to a logarithmic transform of a geometric Brownian motion. This allows for a greater range of applications, as geometric Brownian motion is applied extensively within finance, notably to model the behaviour of stock prices.

Bibliography

Barker, A. L., Brown, D. E., & Martin, W. N. (1995). Bayesian estimation and the Kalman filter. *Computers & Mathematics with Applications*, *30*, 55-77.

Bellalah, M. (2001). Irreversibility, sunk costs and investment under incomplete information. *R&D Management*, *31*(1), 115-126.

Bergemann, D., & Välimäki, J. (2008). Bandit problems. In S. N. Durlauf & L. E. Blume (Eds.), *The new palgrave dictionary of economics* (2nd ed., Vol. 1, p. 336.340). Macmillan Press.

Blanke, D., & Bosq, D. (2012). Bayesian prediction for stochastic processes: Theory and applications. *arXiv e-prints*, 1211.2300.

Conejo, A. J., Morales, L. B., Kazempour, S. J., & Siddiqui, A. S. (2016). *Investment in electricity generation and transmission - decision making under uncertainty* (1st ed.). Switzerland: Springer International Publishing.

Dalby, P. A. O., Gillerhaugen, G. R., Hagspiel, V., Leth-Olsen, T., & Thijssen, J. J. J. (2018). Green investment under policy uncertainty and Bayesian learning. *Energy*, *161*(1), 1262-1281.

Dixit, A. K. (1992). Real options with constant relative risk aversion. *Journal of Economic Perspectives*, *6*, 107-132.

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty.* Princeton, New Jersey: Princeton University Press.

Ekström, E., Lindberg, C., & Tysk, J. (2011). Optimal liquidation of a pairs trade. In G. Di Nunno & B. Øksendal (Eds.), (p. 247-255). Berlin, Heidelberg: Springer Berlin Heidelberg.

Gray, A., Greenhalgh, D., Hu, L., Mao, X., & Pan, J. (2011). A stochastic differential equation SIS epidemic model. *SIAM Journal on Applied Mathematics*, *71*(3), 876-902.

Grewal, M. S. (2011). Kalman filtering. In *International encyclopedia of statistical science* (p. 705-708). Berlin, Heidelberg: Springer Berlin Heidelberg.

Hagspiel, V., Nagy, R. L. G., Sund, S., & Thijssen, J. J. J. (2019). *Investment under uncertainty with costly Bayesian learning: The optimal choice of learning rate.* (Working paper)

Harrison, J. M., & Sunar, N. (2015). Investment timing with incomplete information and multiple means of learning. *Operations Research*, *63*(2), 442-457.

Henderson, V., & Hobson, D. G. (2002). Real options with constant relative risk aversion. *Journal of Economic Dynamics & Control*, *27*, 329-355.

Herath, H. S. B., & Herath, T. C. (2008). Investments in information security: A real options perspective with Bayesian postaudit. *Journal of Management Information Systems*, *25*(3), 337-375.

Hugonnier, J., & Morellec, E. (2007). Real options and risk aversion. *Swiss Finance Institute Research Paper*.

Hull, J. C. (2015). *Options, futures and other derivatives* (9th ed.). Boston, Massachusetts: Pearson.

Jorion, P., & Sweeney, R. J. (1996). Mean reversion in real exchange rates: Evidence and implications for forecasting. *Journal of International Money and Finance*, *15*(4), 535-550.

Keller, R., & Rady, S. (1999). Optimal experimentation in a changing environment. *Review of Economic Studies*, *66*(3), 475-507.

Kolm, P. N., & Ritter, G. (2019). Dynamic replication and hedging: A reinforcement learning approach. *The Journal of Financial Data Science*, *1*(1), 159-171.

Kushner, H. J., & Dupois, P. (2001). *Numerical methods for stochastic control problems in continuous time* (2nd ed.). New York, New York: Springer-Verlag.

Kwon, H. D. (2014). Prevention of catastrophic failures with weak forewarning signals. *Probability in the Engineering and Informational Sciences*, *28*(1), 121-144.

Kwon, H. D., & Lippman, S. A. (2011). Acquisition of project-specific assets with Bayesian updating. *Operations Research*, *59*(5), 1119-1130.

Kwon, H. D., Xu, W., Agrawal, A., & Muthulingam, S. (2016). Impact of Bayesian learning and externalities on strategic investment. *Management Science*, *62*(2), 550-570.

Leland, H. E. (1985). Option pricing and replication with transactions costs. *The Journal of Finance*, *40*(5), 1283-1301.

Leung, T., & Li, X. (2015). Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance*, *18*(3).

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies*, *14*(1), 113-147.

Lucia, J. J., & Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research*, *5*, 5-50.

Martzoukos, S., & Trigeorgis, L. (2001). *Resolving a real options paradox with incomplete information: After all, why learn?* (Working paper)

McDonald, R., & Siegel, D. (1985). Investment and the valuation of firms when there is an option to shut down. *International Economic Review*, *26*, 331-349.

McDonald, R., & Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, *101*, 707-728.

Mjaavatten, A. (2020). polyfix (v. 1.3.1.2). Retrieved 27/6/2020, from `https://www.mathworks.com/matlabcentral/fileexchange/54207-polyfix-x-y-n-xfix-yfix-xder-dydx`

Moscarini, G., & Smith, L. (2001). The optimal level of experimentation. *Econometrica*, *69*(6), 1629-1644.

Näsäkkälä, E., & Fleten, S.-E. (2005). Flexibility and technology choice in gas fired power plant investments. *Review of Financial Economics*, *14*(3), 371-393.

Øksendal, B. (2013). *Stochastic differential equations: An introduction with applications* (6th ed.). Dordrecht, London: Springer Heidelberg New York.

Pertile, P., Forster, M., & Torre, D. L. (2014). Optimal Bayesian sequential sampling rules for the economic evaluation of health technologies. *Journal of the Royal Statistical Society*, *177*(2), 419-438.

Ryan, R., & Lippman, S. A. (2003). Optimal exit from a project with noisy returns. *Probability in the Engineering and Informational Sciences*, *17*(4), 435-458.

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, *52*(3), 923-973.

Singh, R., Ghosh, D., & Adhikari, R. (2018). Fast Bayesian inference of the multivariate Ornstein-Uhlenbeck process. *Physical Review E*, *98*, 1-9.

Soyer, R. (2018). Kalman filtering and sequential Bayesian analysis. *WIREs Computational Statistics*, *10*, e1438.

Thijssen, J. J. J. (2020). *Value-for-money appraisal of large-scale infrastructure investment in a real options framework.* (mimeo, University of York)

Thijssen, J. J. J., & Bregantini, D. (2017). Costly sequential experimentation and project valuation with an application to health technology assessment. *Journal of Economic Dynamics and Control*, *77*(C), 202-229.

Thijssen, J. J. J., Huisman, K. J. M., & Kort, P. M. (2004). The effect of information streams on capital budgeting decisions. *European Journal of Operational Research*, *157*(3), 759-774.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, *5*(2), 177-188.

Wong, H. Y., & Lo, Y. W. (2009). Option pricing with mean reversion and stochastic volatility. *European Journal of Operational Research*, *197*(1), 179-187.

Appendices

Appendix A Deriving the belief process for a constant parameter

We first solve the ordinary differential equation (ODE) of $S(t)$ given in equation (2.9) with substitutions $F(t) = C(t) = 0$, $G(t) = 1$ and $D(t) = m$, which simplifies to

$$dS(t) = -\frac{1}{m^2}S^2(t)dt, \quad \text{where } S(0) = a^2, \tag{A.1}$$

resulting in

$$S(t) = \frac{m^2a^2}{m^2 + a^2t}. \tag{A.2}$$

Building on equation (2.7), we obtain

$$L_{1,c}(t) = -\frac{a^2}{m^2 + a^2t} \quad \text{and} \tag{A.3}$$

$$L_{2,c}(t) = \frac{a^2}{m^2 + a^2t}. \tag{A.4}$$

Appendix B Deriving the belief process for a parameter following an Ornstein-Uhlenbeck process

We first solve differential equation (2.9) with substitutions $F(t) = -p$, $C(t) = q$, $G(t) = 1$ and $D(t) = m$. This results in

$$\frac{dS(t)}{dt} = -2pS(t) - \frac{1}{m^2}S^2(t) + q^2, \quad \text{where } S(0) = a^2, \tag{B.1}$$

which gives

$$S(t) = a^2\frac{\sqrt{p^2 + \frac{q^2}{m^2}} - \left(p - \frac{q^2}{a^2}\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} + p\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}. \tag{B.2}$$

Building on equation (2.7), we obtain

$$L_{1,o}(t) = -\frac{\left(p + \frac{a^2}{m^2}\right)\sqrt{p^2 + \frac{q^2}{m^2}} + \left(p^2 + \frac{q^2}{m^2}\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} - p\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)} \quad \text{and} \tag{B.3}$$

$$L_{2,o}(t) = \frac{a^2}{m^2}\frac{\sqrt{p^2 + \frac{q^2}{m^2}} - \left(p - \frac{q^2}{a^2}\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}{\sqrt{p^2 + \frac{q^2}{m^2}} + \left(\frac{a^2}{m^2} + p\right)\tanh\left(t\sqrt{p^2 + \frac{q^2}{m^2}}\right)}. \tag{B.4}$$

Appendix C The expected value function

As demonstrated in subsection 2.2.1, the expected net present utility of investment includes an expectation of the project value function raised to a power, due to the firm's risk-aversion. We may derive this expectation by application of the law of total probability.

As noted in subsection 2.1.2, the Kalman filter applied to a system of linear state and observation processes with Gaussian sources of uncertainty gives a Gaussian distribution that is equivalent to the distribution of the observed process conditional on historic observations. Hence, we seek the expectation and variance of this distribution at time t, given previous observations, for each of the model applications.

The expected value of the observed process is already established as $\mathbb{E}[X_t|\mathcal{G}_t] = \widehat{X}_t$. For both observation processes, regardless of the initial $\widehat{X}_0$, the variance equals the integrated volatility of the process since observations began. This is equivalent to the variance of the observed process conditioned on previous observations. For the constant process, we first integrate (2.16) to obtain

$$\widehat{X}_t = \int_0^t mL_{2,c}(\tau)dU_\tau + \widehat{X}_0. \tag{C.1}$$

The variance may then be found by

$$\begin{aligned}
\mathbb{V}[X_t|\mathcal{G}_t] &= \mathbb{E}\left[\left(\int_0^t mL_{2,c}(\tau)dU_\tau + \widehat{X}_0\right)^2 \;\middle|\; \mathcal{G}_t\right] - \widehat{X}_0^2 \\
&= \left[\int_0^t mL_{2,c}(\tau)dU_\tau\right]^2 \\
&= m^2\int_0^t L_{2,c}^2(\tau)d\tau,
\end{aligned} \tag{C.2}$$

where the third equality is derived by application of Itô isometry. Similarly, for the Ornstein-

Uhlenbeck process, using equation (2.20),

$$\widehat{X}_t = \widehat{X}_0 e^{-pt} + \int_0^t m L_{2,c} e^{-p(t-\tau)} dU_\tau, \tag{C.3}$$

using a variable substitution $\widehat{Y}_t = \widehat{X}_t e^{pt}$. Variance can then be found by

$$\begin{aligned} \mathbb{V}[X_t|\mathcal{G}_t] &= \mathbb{E}\left[\left(\widehat{X}_0 e^{-pt} + \int_0^t m L_{2,c} e^{-p(t-\tau)} dU_\tau\right)^2 \;\middle|\; \mathcal{G}_t\right] - \left(\widehat{X}_0 e^{-pt}\right)^2 \\ &= \left[\int_0^t m L_{2,o}(\tau) e^{-p(t-\tau)} dU_\tau\right]^2 \\ &= m^2 \int_0^t L_{2,o}^2(\tau) e^{-2p(t-\tau)} d\tau \end{aligned} \tag{C.4}$$

again by Itô isometry. Since both observation processes are Gaussian for all t, we now have all the elements necessary to describe the distributions at a given point in time.

The integral in (C.2) is easily solvable, but this is not the case with the integral in equation (C.4), due to the complexity of $L_{2,o}(t)$. We have therefore used simulated trials of the observation process to estimate the probability distribution variances. Specifically, if $\mathbb{P}$ denotes the Gaussian probability distribution of X_t given previous observations at t, we may denote the equivalent distribution using a simulated variance as $\widehat{\mathbb{P}}$. By total probability,

$$\mathbb{E}\left[V^{1-\gamma}(X_t)\right] \approx \int_{\mathbb{R}} V^{1-\gamma}(x)\, \widehat{\mathbb{P}}(X_t = x \mid \mathcal{G}_t)\, dx, \tag{C.5}$$

where

$$\widehat{\mathbb{P}} \sim \mathcal{N}\left(\widehat{X}_t, \widehat{\mathbb{V}}\right), \tag{C.6}$$

and $\widehat{\mathbb{V}}$ is the sample variance of the trial data $\{\widehat{X}_{t,i}\}_i$ for $i = 1, \ldots, M$ given previous observations at t, where M is the number of trials such that $\lim_{M\to\infty} \widehat{\mathbb{P}} = \mathbb{P}$ by the central limit theorem.[7]

The expectation of equation (C.5) assumes that the entire project value is received at the time of investment. As noted in the introductory paragraph of section 2, our model may be extended to allow for profit streams that arrive over time. The utility function in equation (2.23) would then be applied over the profit stream function, and the expectation would be taken over the integral of discounted utilities of profits. This would simplify to an integral over expectations of utilities of profits. However, the expectation is taken at time t for profits that arrive in the future, so

[7] We calculate the integral of equation (C.5) numerically, using the function `integral` in MATLAB v. R2020a, and evaluate the resulting expression at discrete points of $\widehat{X}_t$, with discretization Δx. Exercise values across paths are then approximated by using the value among discrete integral evaluations that is closest to the given path value.

we would have to derive the future distribution conditioned on a starting point at the current process estimate in order to apply total probability once again. This is not an issue, however, since for our applications, the time dependent distribution of the processes may be derived if an initial value is specified.

Appendix D The Bellman equation

Considering the belief and observation processes (2.28)-(2.29), and building on equation (2.27), we obtain

$$\begin{aligned}
\rho f dt &= \mathbb{E}[df] \\
&= \mathbb{E}\left[\frac{\partial f}{\partial t}dt + \frac{\partial f}{\partial \widehat{X}_t}d\widehat{X}_t + \frac{1}{2}\frac{\partial^2 f}{\partial \widehat{X}_t^2}(d\widehat{X}_t)^2\right] \\
&= \frac{\partial f}{\partial t}dt + \mathbb{E}\left[L_1(t)\widehat{X}_t dt + L_2(t)dZ_t\right]\frac{\partial f}{\partial \widehat{X}_t} + \frac{1}{2}\mathbb{E}\left[L_2^2(t)(dZ_t)^2\right]\frac{\partial^2 f}{\partial \widehat{X}_t^2} \qquad \text{(D.1)} \\
&= \frac{\partial f}{\partial t}dt + [L_1(t) + L_2(t)]\widehat{X}_t\frac{\partial f}{\partial \widehat{X}_t}dt + \frac{1}{2}m^2 L_2^2(t)\frac{\partial^2 f}{\partial \widehat{X}_t^2}dt \\
&\text{when } \widehat{X}_t < \widehat{X}_t^*,
\end{aligned}$$

by applying Itô's lemma on $df = df(t, \widehat{X}_t)$. This simplifies to a partial differential equation (PDE) that characterizes the continuation region of f,

$$\frac{\partial f}{\partial t} + [L_1(t) + L_2(t)]\widehat{X}_t\frac{\partial f}{\partial \widehat{X}_t} + \frac{1}{2}m^2 L_2^2(t)\frac{\partial^2 f}{\partial \widehat{X}_t^2} - \rho f = 0 \quad \text{when } \widehat{X}_t < \widehat{X}_t^*. \qquad \text{(D.2)}$$

Appendix E A finite element method

We have added this appendix to illustrate one of the finite element methods applied in attempting to obtain a surface for the option value. We chose this method as an example as we consider it to be the most illustrative of the challenges that arise when attempting a numerical solution. The notation introduced here is only intended to illustrate the method, and should not be confused with the notation used otherwise throughout the thesis.

For convenience, we let $x = \widehat{X}_t$. We will refer to x as the "spatial" dimension and t as the "temporal" dimension. We let $x \in [0, \bar{x}]$, where $\bar{x}$ is some large value relative to the expected value of x, and $t \in [0, T]$. The grid is divided into M' and N' equidistant points in the spatial and temporal directions, with distances $\Delta x = \bar{x}/M'$ and $\Delta t = T/N'$, respectively. Let i be the

spatial and j the temporal indices, such that $x = i\Delta x$ and $t = j\Delta t$. We discretize f and its domain by applying a Crank-Nicolson scheme using the points displayed in Figure E.1.

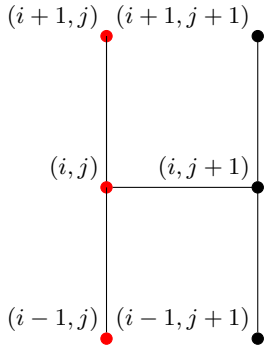

Figure E.1: The Crank-Nicolson stencil.

Let $u_{i,j}$ be the discrete approximation of $f(x,t)$, and $\mathbb{U} = [u_{i,j}]$. The Crank-Nicolson scheme is a linear combination of the forward and backward Euler scheme. With a forward Euler scheme, the spatial and temporal differences are centered around (i,j) and $(i,j+1/2)$, respectively, giving

$$\frac{\partial f}{\partial t} \approx \frac{u_{i,j+1} - u_{i,j}}{\Delta t}, \tag{E.1}$$

$$\frac{\partial f}{\partial x} \approx \frac{u_{i+1,j} - u_{i-1,j}}{2\Delta x} \quad \text{and} \tag{E.2}$$

$$\frac{\partial^2 f}{\partial x^2} \approx \frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{(\Delta x)^2}. \tag{E.3}$$

Substituting these approximations into the PDE (2.30), we obtain

$$\begin{aligned} &\frac{u_{i,j+1} - u_{i,j}}{\Delta t} + [L_1(j\Delta t) + L_2(j\Delta t)]i\Delta x \frac{u_{i+1,j} - u_{i-1,j}}{2\Delta x} \\ &\quad + \frac{1}{2}m^2 L_2^2(j\Delta t)\frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{(\Delta x)^2} - \rho u_{i,j} = 0 \quad \text{when } i(j) < i_j^*, \end{aligned} \tag{E.4}$$

which serves as the discrete approximation of the PDE of the option value. Gathering terms,

$$\begin{aligned} u_{i,j+1} = (1 + \rho\Delta t)u_{i,j} \; + \\ \left[-\frac{\Delta t}{2}[L_1(j\Delta t) + L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2 L_2^2(j\Delta t)\right] u_{i+1,j} \; + \\ \left[\frac{\Delta t}{(\Delta x)^2}m^2 L_2^2(j\Delta t)\right] u_{i,j} \; + \\ \left[\frac{\Delta t}{2}[L_1(j\Delta t) + L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2 L_2^2(j\Delta t)\right] u_{i-1,j}. \end{aligned} \tag{E.5}$$

Similarly, with a backward Euler scheme,

$$
\begin{aligned}
(1-\rho\Delta t)u_{i,j+1} = u_{i,j} \; + \\
&\left[-\frac{\Delta t}{2}[L_1(j\Delta t)+L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2L_2^2(j\Delta t)\right]u_{i+1,j+1} \; + \\
&\left[\frac{\Delta t}{(\Delta x)^2}m^2L_2^2(j\Delta t)\right]u_{i,j+1} \; + \\
&\left[\frac{\Delta t}{2}[L_1(j\Delta t)+L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2L_2^2(j\Delta t)\right]u_{i-1,j+1}.
\end{aligned}
\tag{E.6}
$$

We now let $p^u_{i+1,j}$, $p^s_{i,j}$ and $p^d_{i-1,j}$ represent the coefficients in the last three terms of equation (E.5), such that

$$
p^u_{i,j} = \left[-\frac{\Delta t}{2}[L_1(j\Delta t)+L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2L_2^2(j\Delta t)\right], \tag{E.7}
$$

$$
p^s_{i,j} = \left[(1+\rho\Delta t) + \frac{\Delta t}{(\Delta x)^2}m^2L_2^2(j\Delta t)\right] \quad \text{and} \tag{E.8}
$$

$$
p^d_{i,j} = \left[\frac{\Delta t}{2}[L_1(j\Delta t)+L_2(j\Delta t)]i - \frac{\Delta t}{2(\Delta x)^2}m^2L_2^2(j\Delta t)\right]. \tag{E.9}
$$

Note that the coefficients change in time due to the dependence of L_1 and L_2 on j. Combining equations (E.5) and (E.6), from the forward and backward Euler schemes, respectively, we obtain

$$
\begin{aligned}
& p^d_{i,j}u_{i-1,j} \; + (2+\rho\Delta t + p^s_{i,j})u_{i,j} \; + p^u_{i,j}u_{i+1,j} \\
= \quad & - p^d_{i,j+1}u_{i-1,j+1} + (2-\rho\Delta t - p^s_{i,j+1})u_{i,j+1} \; - p^u_{i,j+1}u_{i+1,j+1}.
\end{aligned}
\tag{E.10}
$$

The free boundary varies across time. We represent the spatial index of the free boundary by i^*_j. At the free boundary, value-matching, from equation (2.33), is obtained by letting $u_{i+1,j} = \Psi_{i^*_j,j}$, where $\Psi_{i^*_j,j}$ is the discretized exercise value function at the free boundary i^*_j, when the stencil is centered one spatial point below i^*_j, meaning $i = i^*_j - 1$. Smooth pasting, from equation (2.34), is implemented by letting the derivative approximation of f be equal to $\partial\Psi_{i^*_j,j}/\partial x$ with a similarly centered stencil. We may in fact improve on the derivative approximation in equation (E.2) at the boundary by allowing for a second order approximation, obtaining

$$
\frac{\partial f}{\partial x} \approx \frac{3u_{i+1,j} - 4u_{i,j} + u_{i-1,j}}{2\Delta x} = \frac{\partial \Psi_{i^*_j,j}}{\partial x}. \tag{E.11}
$$

Hence, equation (E.11), and not (E.10), applies at points centered one spatial point below the boundary. Ordering the spatial elements of the grid in vectors and matrices, we let $\mathbf{u}_j = [u_{i,j}]$

for $i = 1 \dots i_j^* - 1$ and $\mathbb{A}_j$ be a diagonal matrix with elements

$$a_{l,l-1} = p_{l,j}^d \quad \text{when } 1 < l < i_j^* - 1, \tag{E.12}$$

$$a_{l,l} = 2 + \rho\Delta t + p_{l,j}^s \quad \text{when } l < i_j^* - 1 \quad \text{and} \tag{E.13}$$

$$a_{l,l+1} = p_{l,j}^u \quad \text{when } l < i_j^* - 1, \tag{E.14}$$

along the subdiagonal, main diagonal and superdiagonal, respectively. At the spatial point below the free boundary, due to equation (E.11), we have

$$a_{l,l-2} = 0 \quad \text{and} \tag{E.15}$$

$$a_{l,l} = 4p_{i_j^*-1,j}^d + 2 + \rho\Delta t + p_{i_j^*-1,j}^s, \tag{E.16}$$

when $l = i_j^* - 1$. We now define $\mathbf{b}_{j+1} = [b_{i,j+1} + c_{i,j}]$ for $i = 1 \dots (i_j^* - 1)$ such that $b_{i,j+1}$ represents the right-hand side of equation (E.10), and

$$c_{i,j} = 0 \quad \text{when } i < i_j^* - 1 \quad \text{and} \tag{E.17}$$

$$c_{i,j} = -p_{i,j}^d \left(\frac{\partial \Psi_{i_j^*,j}}{\partial x} \cdot 2\Delta x - 3u_{i+1,j} \right) - q_{i,j}^u u_{i+1,j} \quad \text{when } i = i_j^* - 1. \tag{E.18}$$

We may now express the option values for all $i = 1 \dots (i_j^* - 1)$ as a solution to the matrix equation

$$\mathbb{A}_j \mathbf{u}_j = \mathbf{b}_j \quad \text{when } 1 \leq j < N' - 1. \tag{E.19}$$

At the terminal boundary, where $j = N'$, the condition from (2.32) applies such that $u_{i,N'} = \Psi_{i,N'}$ for $i = 1 \dots M'$. Equation (E.19) may therefore be solved recursively for $j = (N' - 1) \dots 1$.

In order to identify a suitable free boundary and corresponding option value surface, we apply the above solution procedure for an initial matrix $\mathbb{U}^0 = [\Psi_{i,j}]$ and boundary $i_j^* = M'$ for $j = 1 \dots N'$. $\mathbb{U}^1$ is found by equation (E.19), after $\mathbb{A}_j$ and $\mathbf{b}_j$ have been designed for the given free boundary. The free boundary is selected such that continuation is optimal for all points $i < i_j^*$. Accordingly, following subsection 2.2.3, we seek a free boundary that satisfies

$$\frac{\mathbb{E}[df(x,t)]}{dt} - \rho f(x,t) > \Psi(x,t) - f(x,t) \quad \text{when } x < x^* \tag{E.20}$$

for all points in the continuation region, or, by approximation,

$$\mathcal{A}^{\Delta} f > \Psi_{i,j} - u_{i,j} \quad \text{when } i < i_j^* \text{ and } j < N' - 1, \tag{E.21}$$

where $\mathcal{A}^{\Delta} f$ denotes equation (2.30) approximated by applying equations (E.1)-(E.3). For all points in the continuation region, we calculate the maximum of the absolute deviations that equations (2.36) and (2.37) take from zero. If the deviation is found to be larger than a specified value, the process is repeated until a satisfying free boundary is found, with corresponding option surface $\mathbb{U}^n$, after n iterations.

Appendix F The Longstaff-Schwartz method

F.1 Convergence and the choice of basis functions

Longstaff and Schwartz (2001) show that the simulated option value resulting from the algorithm outlined in subsection 3.2 is bounded by the true option value from above when the number if trials M approach infinity, or,

$$f_0 \geq \lim_{M \to \infty} \frac{1}{M} \sum_{i=1}^{M} LSM(\omega_i), \tag{F.1}$$

where ω_i indicates the ith trial, and $LSM(\omega_i)$ is the discounted exercise value of ω_i when following the algorithm investment rules. Note that any set of orthonormal basis functions may be used in the regressions so long as their linear combination span the range of the continuation value. Longstaff and Schwartz (2001) use Laguerre polynomials, but point to other possibilities such as Hermite, Legindre, Chebyshev and Jacobi polynomials. They then proceed to show that the option value resulting from following the algorithm converges to the true option value when the number of basis functions increases. Hence, for a high enough number of trials, the simulated option value will approach the true option value from below when the number of basis functions increases. This result is useful in that it allows the user to iteratively increase the number of basis functions until the estimated option value f_0 increases from the previous estimation by an amount $\bar{\epsilon}$ that is deemed acceptably small.

Following Longstaff and Schwartz (2001), we have chosen the set of Laguerre polynomials as basis functions. Using similar notation, we denote the nth order polynomial as $\mathcal{L}_n(\widehat{X}_j)$, expressed in closed form by

$$\mathcal{L}_n(\widehat{X}_j) = \sum_{k=0}^{n} \binom{n}{k} \frac{(-1)^k}{k!} \widehat{X}_j^k. \tag{F.2}$$

The regression equation may then be written as follows, with $\widehat{X}_j$ as the regressor and $Y(\omega_i, j\Delta t)$ the regressand, at $j = \{1, \ldots, N-1\}$,

$$Y(\omega_i, j\Delta t) = \sum_{n=0}^{B} a_{n,j}\mathcal{L}_n(\widehat{X}_j), \tag{F.3}$$

where B denotes the number of basis functions and $\{\boldsymbol{a}_j\}$ the corresponding coefficients.

F.2 The exercise boundary

Following Longstaff and Schwartz (2001) again, we identify the exercise boundary by solving for the belief process values that equate the regression equation – or, the conditional expectation of the continuation value – with the exercise value, at each discrete time point $j = \{1, \ldots, N-1\}$, such that

$$\widehat{X}_j^* = \inf\left\{\widehat{X}_j \in \mathbb{R} : \sum_{n=0}^{B} \hat{a}_{n,j}\mathcal{L}_n(\widehat{X}_j) = 0\right\}, \tag{F.4}$$

where $\{\hat{\boldsymbol{a}}_j\}$ are the estimated coefficients resulting from the regression. The free boundary at $j = N$ is set equal to the indifference value of exercising the option. Since the regression functions consist of weighted polynomials of order potentially higher than 4, and the Abel-Ruffini theorem states that such polynomials do not have an algebraic solution when equated to zero, we find the boundary points by a numerical solution procedure. After the boundary points have been identified, we apply a polynomial smoothing technique that identifies the least squares coefficients of a polynomial of user-defined order η, with fixed endpoints.[8]

F.3 Implementation of the algorithm

We have written our implementation of the Longstaff-Schwartz algorithm from subsection 3.2 in MATLAB v. R2020a with parameter values as given in Table 4.1.

```
%% PARAMETERS
    ornstein = false;       % indicator of process to simulate
    rho = 0.05;     % rate of time–preference (discount rate)
    p = 0.05;       % rate of mean–reversion of OU process
    q = 0.02;       % volatility of OU process
    m = 20;         % observation volatility
    a = 60;         % initial value uncertainty
    life = 0;
```

[8] We have applied the functions `fsolve` and `polyfix` (Mjaavatten, 2020) in MATLAB v. R2020a for the numerical solution procedure and polynomial smoothing technique, respectively.

```
    profit = 100;
    k = 5000;       % investment cost
    gamma = 0;      % relative risk-aversion
    T = 5;          % option lifetime
    mu = 40;
    x_0 = 0 + mu; % initial state
    trials = 500000;
    dt = 0.05;
    dx = 0.2;
    periods = floor(T/dt);
    basis_functions = 5;        % includes constant
%% SIMULATIONS OF SAMPLE PATHS
    if ornstein == true
        % Ornstein-Uhlenbeck
        L1 = @(t) -((p + (a/m)^2)*sqrt(p^2 + (q/m)^2) + (p^2 + (q/m)^2)*tanh(t*
            sqrt(p^2+(q/m)^2)))/(sqrt(p^2 + (q/m)^2) + ((a/m)^2 + p)*tanh(t*sqrt(
            p^2+(q/m)^2)));
        L2 = @(t) ((a/m)^2)*(sqrt(p^2 + (q/m)^2) - (p - (q/a)^2)*tanh(t*sqrt(p
            ^2+(q/m)^2)))/(sqrt(p^2 + (q/m)^2) + ((a/m)^2 + p)*tanh(t*sqrt(p^2+(q
            /m)^2)));

        F = @(t, x) p*(mu - x);     % this drift (combined with x_0) shifts the
            process
        G = @(t, x) m*L2(t);
    else
        % Constant process
        L1 = @(t) -(a^2)/(m^2 + (a^2)*t);
        L2 = @(t) (a^2)/(m^2 + (a^2)*t);

        F = @(t, x) 0;
        G = @(t, x) m*L2(t);
    end
    process = sde(F, G, 'StartTime', 0, 'StartState', x_0);
    if ~hold_simulations
```

```
        default_rng = rng('default');       % save default random number
            generator
        rng(1108);      % set own generator
        [paths, times, ~] = simByEuler(process, periods, 'DeltaTime', dt, '
            nTrials', trials, 'Antithetic', true);      % simulate
        rng(default_rng);       % reset generator
        paths = squeeze(paths);     % remove unary dimensions
    end
%% MATRICES
    PATHS = paths';
    psi = @(x) (((profit.*x).^(1 - gamma))./(1 - gamma)) - ((k^(1 - gamma))/(1 -
        gamma));
    if gamma > 0    % if risk averse, derive expected value by total expectation
        for jj = 1:length(times)
            estimated_std = std(PATHS(:, jj));
            psi_averse = @(x) integral(@(chi) ((((profit.*x).^(1 - gamma))./(1 -
                gamma)) - ((k^(1 - gamma))/(1 - gamma))).*(1/(sqrt(2*pi)*
                estimated_std)).*exp(-0.5*((chi - x)./estimated_std).^2), -Inf,
                Inf, 'AbsTol', 0, 'RelTol', 10^(-3));
            matlab_fun_psi = @(x) max(psi_averse(x), 0);

            x_disc = 0:dx:max(PATHS, [], 'all');
            disc_matlab_fun_psi = zeros(1, length(x_disc));
            for ii = 1:length(x_disc)       % calculate psi at discrete points
                disc_matlab_fun_psi(ii) = matlab_fun_psi(x_disc(ii));
            end

            disc_matrix = [x_disc; disc_matlab_fun_psi];
            [~, closestIndex] = min(abs(x_disc' - PATHS(:, jj).'));     % find
                closest values to PATHS in x_disc
            EXERCISE(:, jj) = disc_matrix(2, closestIndex);     % use
                corresponding discrete psi as exercise values
        end
    else    % if risk neutral, take expectation directly
```

```
64          matlab_fun_psi = @(x) max(psi(x), 0);
65          EXERCISE = matlab_fun_psi(PATHS);
66      end
67      EXERCISE(imag(EXERCISE) ~= 0) = zeros(size(EXERCISE(imag(EXERCISE) ~= 0)));
68      EXERCISE = max(EXERCISE, 0);
69      CONTINUATION = NaN(size(PATHS));
70      CONTINUATION(:, end) = EXERCISE(:, end);
71      CONTINUATION(CONTINUATION == 0) = NaN(size(CONTINUATION(CONTINUATION == 0)));
                % convert all 0-elements to NaN
72      CASHFLOW = zeros(size(PATHS));
73      CASHFLOW(:, end) = EXERCISE(:, end);
74      laguerre = cell(1, basis_functions);
75      syms x n;
76      for i = 1:basis_functions
77          laguerre{i} = char(symsum(nchoosek(i - 1, n)*(((-1)^n)/factorial(n))*x^n,
                n, 0, i - 1));
78      end
79      clear x n;
80  %% ITERATIONS
81      free = NaN(1, length(times));
82      regressand_prediction_int = NaN(2, length(times));
83      free_prediction_int = NaN(2, length(times));
84      if gamma > 0
85          matlab_fun_psi = @(x) psi_averse(x);
86          free(end) = fsolve(matlab_fun_psi, x_0, optimoptions('fsolve', 'Display',
                'off'));
87      else
88          matlab_fun_psi = @(x) psi(x);
89          free(end) = fsolve(matlab_fun_psi, x_0, optimoptions('fsolve', 'Display',
                'off'));
90      end
91      for j = (length(times) - 1):(-1):2
92          % Update discounted cash flows at given j
93              DISC_CASHFLOW = zeros(trials, 1);
```

```
94              for s = length(times):(-1):(j + 1)
95                  DISC_CASHFLOW = DISC_CASHFLOW + CASHFLOW(:, s).*exp(-rho*(times(s
                        ) - times(j)));
96              end
97          % Find new regressands and regressors
98              REGRESSAND = DISC_CASHFLOW((EXERCISE(:, j) > 0));       % only
                    selects paths that are nonzero at j
99              REGRESSOR = PATHS(EXERCISE(:, j) > 0, j);       % only selects paths
                    at j with exercise values that are nonzero at j
100         % Perform regression using Laguerre and update continuation
101             regression_model = fit(REGRESSOR, REGRESSAND, fittype(laguerre), '
                    Normalize', 'off');
102             CONTINUATION(EXERCISE(:, j) > 0, j) = regression_model(REGRESSOR);
103         % Update cash flows
104             CASHFLOW(EXERCISE(:, j) >= CONTINUATION(:, j), j) = EXERCISE(EXERCISE
                    (:, j) >= CONTINUATION(:, j), j);
105             CASHFLOW(EXERCISE(:, j) >= CONTINUATION(:, j), (j + 1):end) = zeros(
                    size(CASHFLOW(EXERCISE(:, j) >= CONTINUATION(:, j), (j + 1):end))
                    );      % set all future cash flows to zero (option is exercised)
106         % Solve for free boundary
107             % 1. Create specific regression function
108             syms x;
109             symbolic_regression_model = subs(str2sym(char(regression_model)),
                    coeffnames(regression_model), coeffvalues(regression_model)');
110             matlab_fun_regression_model = matlabFunction(
                    symbolic_regression_model);
111             clear x;
112             % 2. Update psi (without a maximum function, following the
                    formulation in L&S)
113             if gamma > 0
114                 estimated_std = std(PATHS(:, j));
115                 psi_averse = @(x) integral(@(chi) ((((profit.*x).^(1 - gamma))
                        ./(1 - gamma)) - ((k^(1 - gamma))/(1 - gamma))).*(1/(sqrt(2*
                        pi)*estimated_std)).*exp(-0.5*((chi - x)./estimated_std).^2),
```

```
                    -Inf, Inf, 'AbsTol', 0, 'RelTol', 10^(-3));
116             matlab_fun_psi = @(x) psi_averse(x);
117         else
118             matlab_fun_psi = @(x) psi(x);
119         end
120         % 3. Solve the  boundary equation
121         boundary_equation = @(x) matlab_fun_regression_model(x) -
                matlab_fun_psi(x);
122         free(j) = lsqnonlin(boundary_equation, free(j + 1), -1000000,
                1000000, optimoptions('lsqnonlin', 'Display', 'off'));
123     end
124     % Update discounted cash flows
125     j = 1;
126     DISC_CASHFLOW = zeros(trials, 1);
127     for s = length(times):(-1):(j + 1)
128         DISC_CASHFLOW = DISC_CASHFLOW + CASHFLOW(:, s).*exp(-rho*(times(s) -
                times(j)));
129     end
130     % Find option value at 0
131     option_value = mean(DISC_CASHFLOW);
132     % Find exercise value at free boundary
133     if gamma > 0
134         matlab_fun_psi = @(x) max(psi_averse(x), 0);
135         exercise_free = zeros(1, length(free));
136         for ii = 1:length(free)
137             exercise_free(ii) = matlab_fun_psi(free(ii));
138         end
139     else
140         matlab_fun_psi = @(x) max(psi(x), 0);
141         exercise_free = matlab_fun_psi(free')';
142     end
143 %% OUTPUT
144     fprintf('Option value = %.2f\n', option_value);
```

Appendix G Relative error against number of basis functions

Figure G.1 shows the relative error of the option values plotted against the number of basis functions used in simulations of the base case. Note that the plots use different horizontal axis ranges.

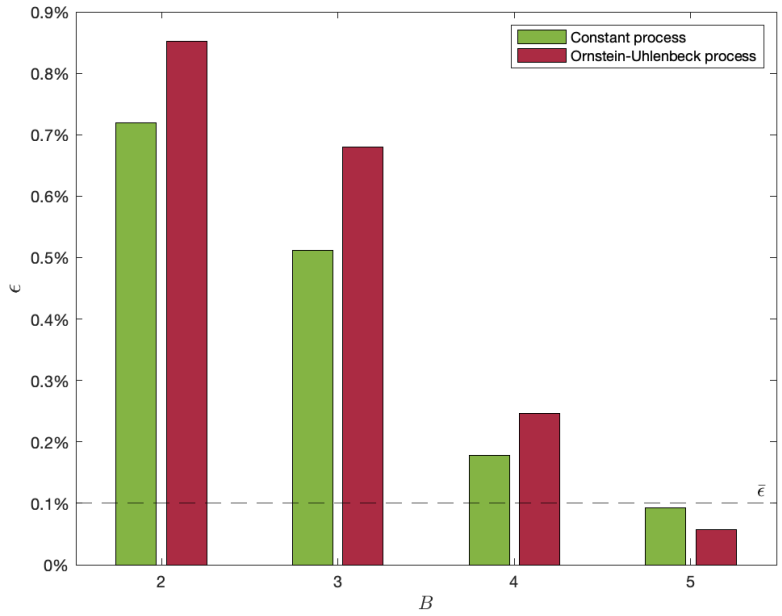

(a)

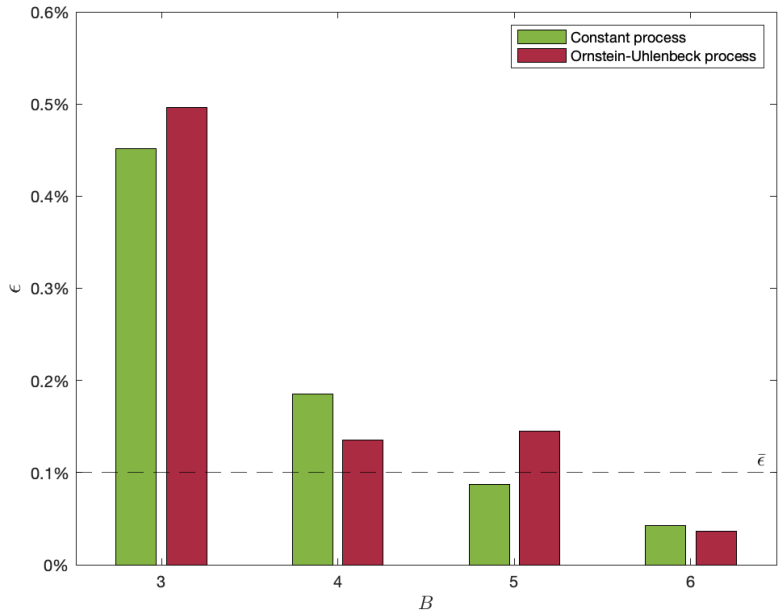

(b)

Figure G.1: Relative errors of option values ϵ against the number of basis functions B under (a) risk neutrality, and (b) risk aversion.

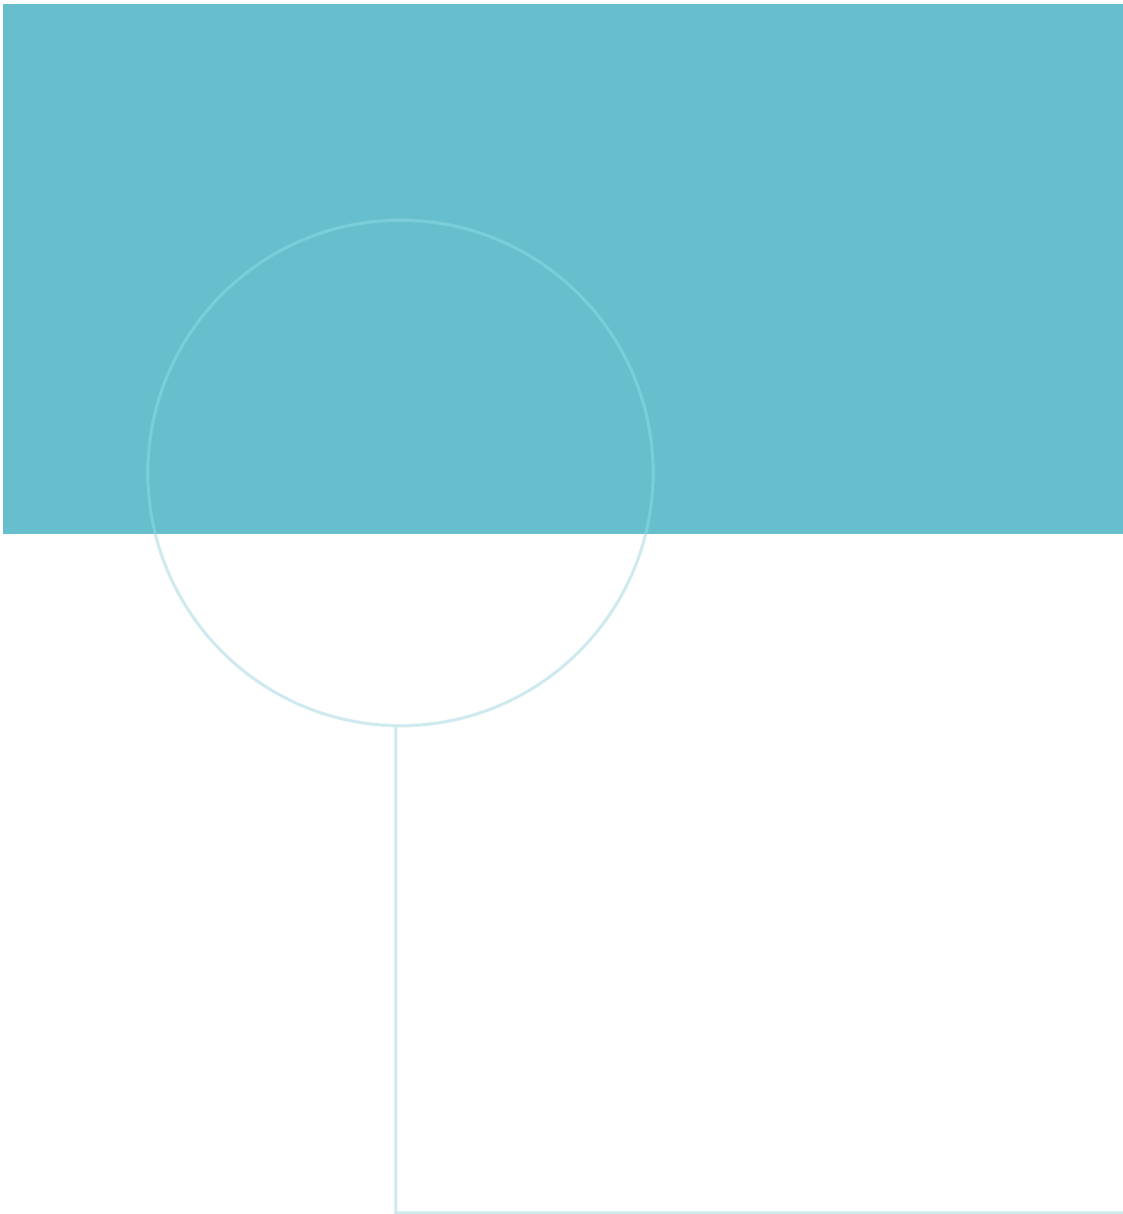

NTNU

Norwegian University of
Science and Technology