Lone Amalie Storaker Eriksen
Tonje Josefine Olsen

A Multi-Layered Learning Approach for Sequential Decision Problems with Multiple Uncertainties

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel
Co-supervisor: Reidar B. Bratvold
June 2021

Master's thesis

NTNU
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

Lone Amalie Storaker Eriksen
Tonje Josefine Olsen

A Multi-Layered Learning Approach for Sequential Decision Problems with Multiple Uncertainties

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel
Co-supervisor: Reidar B. Bratvold
June 2021

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

Preface

This thesis serves as the concluding part of our Master of Science at the Norwegian University of Science and Technology (NTNU). The degree is a specialisation in Financial Engineering at the Department of Industrial Economics and Technology Management. It was written during the spring of 2021.

We would like to express our deepest gratitude towards our supervisors Verena Hagspiel of NTNU and Reidar B. Bratvold of UiS. Thank you for your excellent guidance and challenging questions, making us more reflective and observant of our own choices. In other words, thank you for teaching us how to make good (better) decisions. We are very grateful to you both for your time and interest in our thesis, and for the many interesting and fun discussions. You have both inspired and encouraged us, and you have made us see things from new perspectives. Working with you has been a pleasure.

Trondheim, 11th of June, 2021

Lone Amalie Storaker Eriksen

Tonje Josefine Olsen

Abstract

In this paper, we contribute to the existing real options valuation literature by adding Bayesian updating to the Least Squares Monte Carlo (LSM) solution approach. In addition to the learning that occurs over time from the stochastic processes used in the LSM approach, the decision maker will periodically receive signals providing information about the hyperparameters of the stochastic processes. Bayesian inference is used to update the decision maker's prior beliefs about the hyperparameters, which are used to set the distributions in the stochastic processes. The need for parameter updates arises from regime shifts as indicated by the signals available to the decision maker. We provide further insights into the decision context by examining how different beliefs and parameter choices affect the optimal decision policy for an illustrative example using sensitivity analysis.

We find that receiving signals and updating the beliefs can notably impact the investment value and decision policy for the investment problem. If the signals differ sufficiently from the prior beliefs, a different decision policy and investment value are often reached. The signals will have a more prominent effect when the decision maker's uncertainty is high. The methodology is flexible and versatile and is applicable to a broad set of problems.

Sammendrag

I denne masteroppgaven bidrar vi til realopsjonslitteraturen ved å kombinere Bayesiansk læring og Least Squares Monte Carlo (LSM)-metoden. I tillegg til den læringen som oppstår over tid med stokastiske prosesser brukt i den tradisjonelle LSM-metoden, så vil beslutningstakeren periodisk motta signaler med informasjon om prosessenes hyperparametere. Bayesiansk inferens er brukt til å oppdatere beslutningstakerens tro på hyperparameterne, som videre er brukt til å angi distribusjonene i de stokastiske prosessene. Behovet for å oppdatere parameterne kommer fra regimeskifter, indikert av signalene beslutningstakeren mottar. Vi gir videre innsikt i beslutningskonteksten ved å undersøke hvordan ulik tro og parametervalg påvirker den optimale beslutningsstrategien i et illustrativt eksempel ved bruk av sensitivitetsanalyse.

Vi observerer at å motta signaler og oppdatere troen kan betydelig påvirke investeringsmulighetens verdi og beslutningsstrategi. Hvis signalene avviker betydelig fra den opprinnelige troen, så oppnår beslutningstakeren ofte en annen investeringsverdi og beslutningsstrategi. Signalene vil ha en større effekt når beslutningstakerens usikkerhet er høy. Metodologien er allsidig og fleksibel, og den er relevant for et stort utvalg applikasjoner.

Contents

List of Figures

List of Tables

1 Introduction

In times of fast-paced technological advancements, decision makers are constantly facing the decision of whether to invest in new technology now or wait and learn about its future development. Let's consider a firm that is evaluating adoption of a new technology necessary to launch an innovative product. The firm must consider whether and when to invest in the new technology. These types of decision problems include sources of significant uncertainty. The decision maker has beliefs about these uncertainties based on past data and experience. As decisions are about the future, uncertainty in decision making is inevitable and will be present until after the decision is made (Bratvold and Begg, 2010). However, there may be value in reducing this uncertainty. This can be done by collecting additional information about the potential success of the new technology. For example, the firm considering launching the innovative product could delay the adoption decision and perform market surveys to map the market's interest. It is also possible that Governments impose new regulations or a competitor introduce a disruptive innovation to the market, affecting the profitability of the new technology. These events cause regime shifts, resulting in new information the decision maker can observe. With new information, probabilities used to quantify the decision maker's uncertainty should be updated. This updating can be accomplished through Bayesian Inference. Hence, it may be valuable for the decision maker to wait and gain further knowledge that can impact the decision.

The decision maker uses the arrival of information over time to learn about her uncertainties. We introduce two levels of learning. The first level arises from the Markov property embedded in the traditional real option theory. This type of learning happens through observing revelations of the state of, e.g., the market price as new price information becomes available. Note that the Markov property is implicit in uncertainties modelled by stochastic processes. At this level, the parameters of the stochastic process are unchanged over the investment horizon. The first level of learning will also be referred to as Markov learning. We define the second level of learning to be the processing of new information and explicit updating of the decision maker's subjective beliefs in the same way as Martzoukos (2003) and Dalby et al. (2018). They defined this form of learning as "active learning", whereas the first level was considered “passive” learning. In the second level, the parameters of the stochastic processes are updated using signals with new information. This means that the decision maker can use signals with new information to update their subjective beliefs over the considered investment horizon. We will refer to this second level of learning as Bayesian learning or updating throughout this paper. A combination of both levels of learning will be referred to as a multi-layered learning approach.

In this paper, we present a new methodology to study sequential technology adoption problems that accounts for multiple uncertainties as well as Bayesian updating of hyperparameters. Our solution method extends the classical Least-Squares Monte Carlo (LSM) approach to update the model parameters of the stochastic processes underlying the decision problem using Bayesian inference. The methodology allows us to derive the value of the investment opportunity and inform the decision maker whether the optimal decision at a given point in time is to adopt or reject a new technology or to delay the decision. We apply the methodology to an illustrative example and perform sensitivity analysis to understand how and when the results differ from those obtained when not applying a multi-layered approach. Furthermore, we investigate in what contexts accounting for signals is of significance.

We find that the investment value and optimal decision often differ when accounting for signals and Bayesian updating. If the decision maker does not account for signals from regime changes, there is a risk that she will underestimate or overestimate the investment value of the new technology. Consequently, the decision maker may invest or reject at a suboptimal time. Whether our methodology yields different results compared to when we exclude Bayesian learning depends on how much the received information differs from the decision maker’s initial beliefs and how uncertain she is about them. If the decision maker has strong beliefs, she is less susceptible to signals and is less likely to change her beliefs. On the other hand, signals will have a big impact if she is very unsure of her beliefs. However, if her prior beliefs of the uncertainty equal the mean of the signal distribution, accounting for both levels of learning leads to similar results compared to when only Markov learning is considered.

The main contribution of this paper to the literature is the new methodology extending the classical sequential decision problem where the beliefs are updated in a Bayesian manner. The solution approach is based on the LSM methodology and includes an additional layer of Bayesian updating. By applying Bayesian updating of the hyperparameters of the stochastic processes in LSM, we allow for an extra level of uncertainty and learning. We thereby recognise that estimated parameters are likely to change over time. To the best of our knowledge, this has not been done before. The solution approach easily allows for multiple sources of uncertainty, which is a limitation of many papers on technology adoption. It can also incorporate learning for multiple uncertainties.

The remainder of this paper is organised as follows. A literature review is provided in Section 2. In Section 3, the model setting and the new solution approach are presented. Then, in Section 4, we present an illustrative application of the model and discuss results from sensitivity analysis. Lastly, Section 5 summarises and concludes.

2 Literature review

With this work, we aim to contribute to several strands of literature. First, our sequential decision problem is relevant in both the traditional fields of real options theory as well as decision analysis. These two research communities, as well as the problems they study, are overlapping. However, both tend to have different preferences related to modelling choices and terminology. As a result, participants from the distinct fields might not be aware of the similarities as highlighted by Smith and Nau (1995). Problems in both fields could be categorised as a real option when applying Dixit and Pindyck's (1994) definition of a real option. Dixit and Pindyck (1994) describe how real options can be valued using either contingent claims analysis or dynamic programming. They point out that both techniques are "closely related to each other, and lead to identical results in many applications(...)[, but] make different assumptions about financial markets, and the discount rates that firms use to value future cash flows." (Dixit and Pindyck, 1994, p. 93). In Smith and Nau (1995) and Smith and McCardle (1998), contingent claims analysis is described as the solution method used in financial markets, and dynamic programming the solution method in decision analysis. These methods yield consistent results under certain conditions (Insley and Wirjanto, 2010). We can therefore describe many real options problems as decision analysis problems and vice versa. Common problems studied in both fields are adoption decisions of new technologies with uncertain key factors influencing future profitability. Such decisions are often modelled as optimal stopping problems where a firm must decide whether the investment should be made now or delayed. By delaying the investment decision, the decision maker can observe the evolution of the key factors and gain new knowledge. Therefore, one can say that the decision maker learns with time.

There are multiple papers addressing the problem of technology adoption under uncertainty. Farzin et al. (1998) study the optimal timing of technology adoption for a firm that faces a stochastic innovation process. They consider uncertainty about both the speed of arrival and the degree of improvement of the new technologies. They find that a slower speed of arrival for the new technologies will lead to quicker adoption. Hagspiel et al. (2015) investigate a firm considering investing in a new technology from an old one while facing uncertain timing of future technology improvements. They relax the assumption of constant arrival from Farzin et al. (1998). They find that a firm's investment strategy changes significantly when the arrival rate of new technology is changing instead of constant. Doraszelski (2004) investigates technology adoption under the distinction between innovations and improvements. Improvements are defined as engineering refinements that follow new technologies and enhance the basic innovation's efficiency. Multiple such improvements could add up to significant efficiency gains and make it optimal to delay the adoption of the new technology. He finds that the possibility of further improvements gives an incentive for the firm to delay the adoption of the new technology until the new technology is sufficiently advanced.

All of the papers from the previous paragraph apply what we defined as the first level of learning in the introduction. That is, they account for technological uncertainty but do not update subjective beliefs over the investment horizon. Uncertainty is constant over time in this stream of literature (Dalby et al., 2018). An interesting extension of the traditional real options literature on technology adoption is the introduction of Bayesian learning. In this literature stream, the decision maker receives signals by which she updates her beliefs and reduces her uncertainty. One of the pioneers to apply Bayesian inference when analysing optimal technology adoption decisions is Jensen (1982). He considers the profitability of a new technology to be uncertain. The decision maker can gather information in the form of signals regarding the true profitability to update her beliefs. At each time step, the decision maker can either adopt the new technology or delay the decision to learn more. This model was further developed by McCardle (1985), who accounts for a cost to delay the decision and receive more signals. Since the decision to delay is costly, the option to reject is introduced as the waiting cost can be larger than the potential gain. A limitation to both of these models is that they are restricted to conjugate pairs. Conjugate pairs provide ease of computation but can make the model less applicable to certain situations. Processes in the real world cannot always be sufficiently described by conjugate pairs but can be better approximated by other distributions. Ulu and Smith (2009) address this further, generalising and extending McCardle (1985) by allowing general probability distributions and signal processes. All of these papers only apply for a single uncertainty. In contrast, our methodology allows us to account for multiple uncertainties without suffering from the curse of dimensionality. However, McCardle partly inspires our problem setting through the implementation of costly signals and Bayesian inference.

Ryan and Lippman (2003) and Kwon and Lippman (2011) both analyse decision making under Bayesian updating by studying investment in a project with uncertain profit streams. Ryan and Lippman (2003) seek to find the optimal exiting strategy. The cumulative profit stream follows a Brownian motion with unknown drift, interpreted as the profit rate. The firm can receive signals over time to update its beliefs of the project's profit rate. Kwon and Lippman (2011) study optimal investment time in a project where the firm can learn about the project's true profitability through investment in small scale pilot projects. Both papers assume the signals are subject to noise and arrive at a cost. Harrison and Sunar (2015) formulate an investment problem where the project's value is uncertain. They use a continuous-time Bayesian framework to model the project's uncertainty. The uncertainty can be reduced through the use of one or several learning modes. Signals about the project's value arrive continuously following a Brownian motion. Thijssen et al. (2004) consider a similar investment opportunity with uncertain profitability, but signals can be gathered costlessly and arrive according to a Poisson process. Common to all these papers is that the profitability of the project can be in either of two states; high (profitable) or low (unprofitable). In this case, it is not optimal to wait unless it is possible for additional signals to change the adoption decision. There is no reason to pay for information if it cannot alter the decision outcome (McCardle, 1985). These approaches differ from ours as we assume that the uncertain parameter the decision maker updates her beliefs of can take any value from a given distribution and is not limited to two states. Similarly, we calculate the investment value instead of quantifying it as either of two known values. Further, we model the signals to arrive at regular intervals, and they are neither assumed continuous nor to follow a Poisson process. However, this is a modelling choice and not a limitation of our model.

Most papers in the field of technology adoption with Bayesian learning account for at most one uncertainty. An exception is Diendorfer (2019) who accounts for multiple uncertainties while implementing Bayesian inference. He presents a decision analysis model that allows for four uncertainties, where a firm's beliefs of the outcome of one uncertainty can be updated. The model is applied to analyse blockchain adoption in the energy sector. The decision problem he studies is the most similar to this paper as far as we know. However, the uncertainties are not modelled as stochastic processes, and his solution approach is based on a decision tree with continuous and discrete chance nodes. Furthermore, he only applies the second level of learning and can at most learn about one of the uncertainties.

We also add to the strand of literature concerning simulation based approaches to value American options. Our solution approach of the classical sequential decision problem with Bayesian updating contributes as a numerical approach for valuing American options that allow for multiple uncertainties. Tilley (1993) is the first to use Monte Carlo simulations to value American options. A drawback of his methodology is that it is not apparent how the method can be applied to multiple variables. Barraquand and Martineau (1995) aim to overcome Tilley's drawbacks. Their presented algorithm is able to value American options depending on multiple underlying assets. However, since the algorithm does not follow an optimal exercise policy, it underestimates the option value. Broadie and Glasserman (1997) develop a simulation algorithm for estimating the prices of American-style securities by estimating two values of the option price; one that is biased high and one that is biased low. Their methodology is also suitable for dealing with multiple underlying assets.

Longstaff and Schwartz (2001) introduce a new method based on Tilley's idea for valuing American options called Least-Squares Monte Carlo (LSM). This method can easily be applied to multiple variables without suffering from the curse of dimensionality and therefore overcomes one of the issues with Tilley's approach. To price the American option and find the optimal exercise timing, they estimate the conditional expected payoff to the option holder from continuation using least squares. More specifically, a conditional expectation function is regressed using the Monte Carlo simulations. The optimal decision is then made by comparing the value of the conditional expected payoff and immediate exercise. When the optimal decision is found for each Monte Carlo simulation, the cash flows are discounted back to time zero, and the option value equals the average cash flow.

The LSM approach of Longstaff and Schwartz (2001) has been applied to a large variety of real options problems (e.g., Rodrigues and Armada (2006) value portfolios of real options, Willigers and Bratvold (2009) value oil and gas options, and Blanco et al. (2011) values FACTS investments). Abdel Sabour and Poulin (2006) and Cortazar et al. (2008) contribute to the literature by presenting an application of the LSM method to value real capital investments and illustrate its efficiency for higher dimensions. The approaches of Abdel Sabour and Poulin (2006) and Cortazar et al. (2008) differ in what state variables they use to regress the continuation function. Both papers conclude that the LSM method is suitable for valuing real options when comparing the results with those obtained using traditional methods. We contribute to this strand of literature by extending the LSM to allow for Bayesian updating of the hyperparameters of the stochastic processes. To the best of our knowledge, we are the first ones to develop an appropriate methodology for this. Furthermore, we contribute by solving a sequential technology adoption problem using a methodology based on LSM, allowing for more than one uncertainty. To our knowledge, LSM has not been applied to technology adoption problems before.

3 Model

In this section, we first introduce the considered decision problem and the general model setup. We then suggest and explain a solution approach for the model.

3.1 Model setup

We consider a decision maker, hereafter referred to as DM, who has the option to invest, at a cost K, in a new and innovative technology. The DM is uncertain about the future profitability of the new technology. It is assumed that the investment decision can be made at the latest at time T. At any earlier time t, the DM can either adopt or reject the new technology or postpone the decision at a waiting cost, denoted by W. If the DM decides to wait, she can receive signals yielding information regarding the profitability of the new technology. The information is caused by a regime shift in the environment the DM is operating in, making her previous beliefs inappropriate. She can use the information to update her beliefs about the profitability and make a more informed decision. The decision to adopt or reject the new technology is assumed to be at least partially irreversible. When the new technology is either adopted or rejected, the investment opportunity no longer exists. This sequential decision problem is illustrated in Figure 3.1.

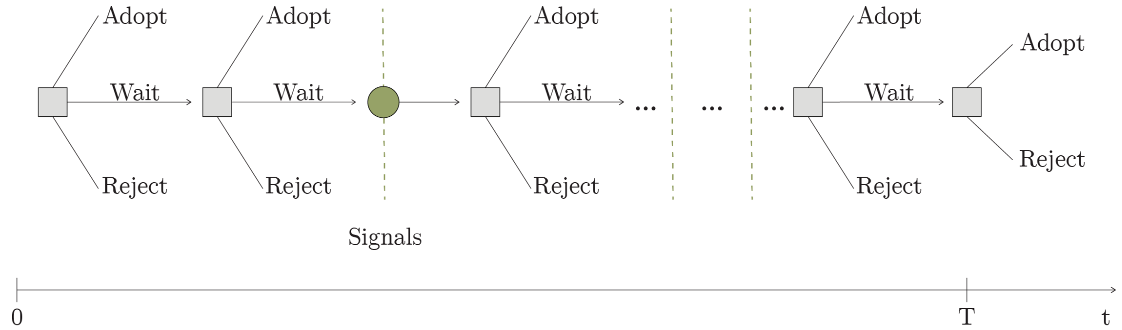

Figure 3.1: Illustration of the sequential decision tree when receiving signals.

The trade-off for the DM in terms of timing is the following. If the DM decides to wait, the technology level might rise, which could result in a higher payoff of investment later. Another possibility is that the new technology remains relevant for a longer period of time, with little change to its value. In this case, the DM should invest as early as possible to reap most of the technology's benefits. A third option is that the technology soon becomes outdated. Then the profitability is overestimated, and the DM should not invest. Therefore, the DM risks both investing at a suboptimal time and making the wrong decision.

The general model setting serves numerous practical applications and is of high relevance. The technology of interest could be a new production process, a new Covid-19 vaccine or the purchase of an electric vehicle. Common to these is that the value of investing in the technology depends on uncertain variables given exogenously to the firm. The DM evaluates the investment decision until she either adopts or rejects it. At time T, there is no longer an option to delay the decision, and the DM must either adopt or reject the technology. Adoption or rejection are mutually exclusive. The decision problem considered here represents a typical real options problem (Dixit and Pindyck, 1994).

The value resulting from immediate adoption of the technology is given by a payoff function, $P_t(., ..., .)$, dependent on M stochastic variables. The values of the M variables, denoted by $(S_{1,t}...., S_{M,t})$, are known at time zero, $t = 0$, but their future values are uncertain and are modelled by stochastic processes. Therefore, the payoff from immediate adoption is known at time zero, but its future value is uncertain. For simplicity, we assume that the payoff is a one-time payment received at the time of adoption. However, the model can easily be extended to a setting where the payoff is received over time.

Our model is applicable for a broad range of investment problems subject to uncertainty. The payoff function is problem specific and must be modified to fit the problem at hand. It can easily be extended for many uncertainties and different functional dependencies. For illustrative purposes, we formulate a simple, general payoff function in order to introduce the methodology. To cover a broad class of problems, we introduce a payoff function that includes uncertainty related to both revenue and cost. Specifically, we assume that the payoff function is given by

$$P_t(S_{1,t}, S_{2,t}) = Q \times (S_{1,t} - S_{2,t}) - K$$

for a production problem, where Q is the number of units produced, $S_{1,t}$ denotes the uncertain unit price and $S_{2,t}$ the uncertain unit cost. Note that the payoff can be both positive and negative. Similarly, one could define a function for the rejection value. Rejection could also result in the payment of a sunk cost or a salvage value that is received. In our model, the value of rejection is set to zero to keep the model as simple as possible to make the methodology clearer. Lastly, one must define the value of waiting, which is the expected value of the investment if the decision is delayed.

Considering the presented example-payoff function, we have two variables. The first variable, $S_{1,t}$ represents the unit sales price at time t, and the second variable, $S_{2,t}$ represents the unit cost at time t. The unit sales price is assumed to follow a standard geometric Brownian motion (GBM), and the unit cost follows a geometric Ornstein-Uhlenbeck (OU) process. We have chosen these stochastic one-factor processes as they are relatively simple yet rich enough to illustrate the important aspects of our model. The unit price process then follows the SDE given by

$$dS_{1,t} = \mu_{S_1} S_{1,t} dt + \sigma_{S_1} dZ_t \tag{3.1}$$

where μ_{S_1} is the drift and σ_{S_1} is the standard deviation. For the GBM, the drift represents the expected growth the price will have over time. The unit cost process then follows the SDE given by

$$dS_{2,t} = \eta(log(\mu_{S_2}) - log(S_{2,t}))dt + \sigma_{S_2} dZ_t \tag{3.2}$$

where μ_{S_2} and σ_{S_2} are the drift and standard deviation, respectively, and η is the mean-reverting constant. For the OU, the drift represents what cost the variable will revert to over time. The methodology is not limited to these stochastic processes. It could easily be extended for multiple uncertainties modelled by any stochastic processes, including multi-factor processes.

Standard real options models commonly assume constant drift and variance parameters throughout the investment horizon. Instead of assuming these variables are constant over the entire investment period, we model the arrival of signals with information. The signals are denoted by $\mathbb{X}$ and can be collected over time if the decision is delayed. These signals can be viewed as results of regime shifts. When a regime shift occurs, the DM observes new information that she uses to update her beliefs and reduce her uncertainty about the profitability. For example, the signals could result from a regime shift such as the Norwegian government's introducing subsidies for electric vehicles due to their focus on facilitating greener lifestyles, thereby making electric vehicles more attractive. For the new production process, the signals could be tests of its produced units or signs of other competing systems likely to enter the market. If a competitor launched a similar product, the market environment for the original firm's product would change. It is likely that the firm would have to change its sales price and marketing strategy, and it can therefore be seen as a regime shift. We assume these regime shifts happen at regular intervals. However, the model can be extended to account for shifts at unknown times. These regime shifts split the entire investment horizon into smaller intervals of different regimes. Different regimes may have different process parameters from each other, but each regime is stationary.

We implement learning through signals with Bayesian inference due to the richness and flexibility of the Bayesian framework. By applying Bayes' rule to update a prior belief into a posterior belief, the DM's updated beliefs are a combination of her old beliefs and her new knowledge. In order to keep the model as simple as possible, we exploit conjugate pairs to enhance the understanding of the methodology. However, the model could easily be extended to account for more complex distributions. For example, it is possible to apply Markov Chain Monte Carlo (MCMC) to this problem setting. MCMC is a powerful algorithm for complex models and can approximate any distributions where the target distribution is unknown (Kruschke, 2014).

To exemplify the Bayesian learning layer in the methodology, the underlying drift of the price process is unknown. This adds an uncertain parameter in addition to the uncertain future variable values, making the model more complex. The underlying drift is also called the true[1] drift, denoted by μ_{true}. The DM has beliefs or knowledge of this parameter, denoted $\mu_{S_{1,t}}$ in Equation 3.1. The beliefs are represented by a normal distribution, $N(\mu, \tau^2)$. μ is what the DM believes the drift to be, and τ is the beliefs' variance describing how uncertain she is that μ is a good estimate of μ_{true}. The DM has initial beliefs denoted by $N(\mu_0, \tau_0^2)$ before any signals are received. The signals, $\mathbb{X}$, provide information on the unknown drift parameter. $\mathbb{X}$ are assumed to be normally distributed with the true drift and are subjected to noise, which is known and represented by σ_{signal}. That is, $\mathbb{X} \sim N(\mu_{true}, \sigma_{signal}^2)$. As the prior and the signals are normally distributed, they form a conjugate pair, and the resulting posterior is then normally distributed. After observing the signals, $\mathbb{X}$, we have $N(\mu_{post}, \tau_{post}^2)$, where μ_{post} is the posterior beliefs after updating and the DM's new beliefs of μ_{true}. τ_{post} tells how strongly the DM believes that μ_{post} is equal to μ_{true}. The more signals received, the smaller τ_{post} is. If the DM receives another signal, then the previous posterior becomes the new prior for this signal, and μ_{post} and τ_{post} then represent the new prior. This is illustrated in Figure 3.2.

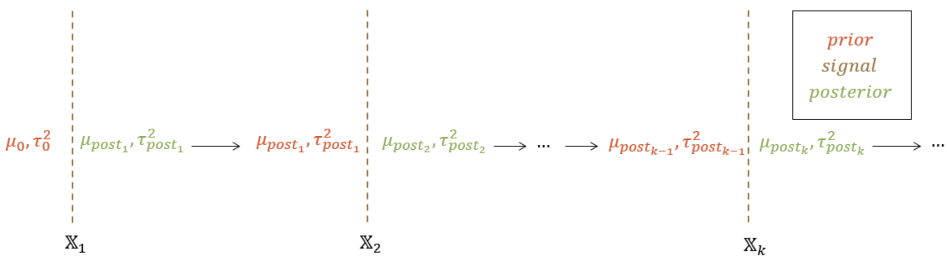

Figure 3.2: Sequentially updating the DM's beliefs with received signals and explanation of how the posterior becomes the next prior.

[1]Note that we are referring to this value as the true drift. This is a common phrase in engineering and economics for the uncertain parameter when applying Bayes' rule. However, the statistics and decision analysis field generally do not approve of this reference. Their argumentation is that the uncertainty is with regards to the future, and signals are not able to predict the future. Signals either reflect the past or the current situation, and the DM uses these to make inference about the future. We choose to use the reference conventional for the engineering and economics discipline and will refer to it as "the true drift"

A low variance of the beliefs, τ, indicates that the DM is very confident that μ is a good estimate of μ_{true}, and a high variance indicates the opposite. If the DM has little or poor prior knowledge to estimate the true value, τ is large. A prior and signal, and the resulting posterior are illustrated in Figure 3.3. Implementing learning on one of the drift parameters is not a limitation of the model, as it can be extended to account for learning on multiple uncertainties. Additionally, the DM can receive signals with information on other parameters than the drift, e.g. the volatility. It is also possible to obtain the benefits of conjugate pairs with other distributions, such as Gamma-Poisson or Beta-Bernoulli.

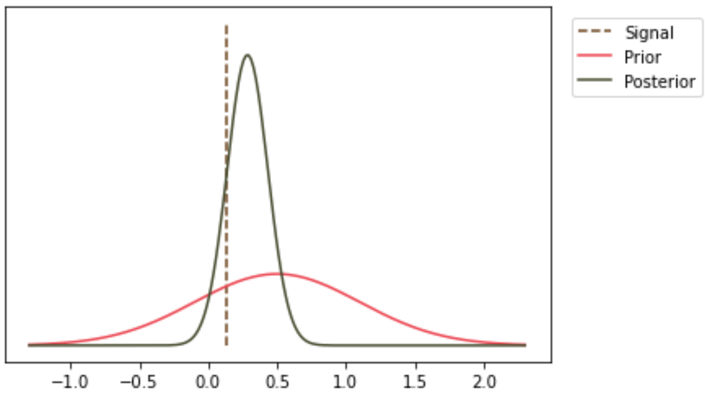

Figure 3.3: Updating the DM's beliefs using one received signal. Notice that the posterior is narrower than the prior, because the DM is more confident in her beliefs after receiving a signal.

We denote the value of the investment opportunity at a given time t by V_t. To find the optimal action for the DM, we derive the value of rejection, adoption and waiting at time t. The value V_t can be described by the maximum value of these three actions. This can be expressed in the following objective function at time t

$$V_t = max[0, P_t, -W + \delta V_{t+1}]$$

where δ is a discount factor to compare values at different time steps. We define $\delta = e^{-rdt}$, where r is the annual risk-free rate, and dt is the time since last evaluation. To summarise, we want to obtain the value of the investment opportunity and the optimal decision at the current time.

3.2 Solution approach

In this section, we present a novel solution method for the model described above in Section 3.1. Our solution approach is based on the classic LSM method. We apply this method because this simulation-based approach is a well-known and recognised solution approach for investment valuation problems where the investment decision depends on multiple sources of uncertainty and must be made in each time step. Notably, LSM does not suffer from the curse of dimensionality. By not having to limit the number of uncertainties due to computational complexity, we can obtain a more general and practice-relevant model. However, different from classical applications to real options problems, we need to alter the method to account for both rejection and learning.

3.2.1 Monte Carlo simulations

The first step in the solution approach is to use Monte Carlo simulations to generate the future uncertain values of the M variables that are modelled by stochastic processes. To this end, we must discretise the stochastic processes. The investment horizon from $t = 0$ to $t = T$ is split into Ω time periods. The decision of whether to invest, wait or reject is revisited at the beginning of each time period. We simulate N paths of each of the M variables over the investment horizon. This yields multiple realisations for the uncertain variables, now denoted by $S^n_{m,t}$ for $n \in [1, N]$. For our illustrative example introduced in Section 3.1, we have a price and a cost, $S^n_{1,t}$ and $S^n_{2,t}$, for each time step t and path n.

It is possible for several regime shifts to occur over the considered investment horizon, but they do not necessarily happen at every time step. When a regime shift arises, the DM receives a signal she can update her beliefs with. Depending on the investment environment, the DM can receive signals regarding multiple uncertain parameters for which she has beliefs about. For our illustrative example, the signals contain information about, and are used to update the DM's beliefs of, the drift parameter in the price process, μ_{S_1}, as introduced in Section 3.1. Until an update is triggered by the arrival of a signal, the simulations are based on the DM's previous beliefs about the drift. Upon the aforementioned update, the evolution of the uncertain price value is simulated with the updated belief of the drift. The simulated variable value for the next time step is based on the simulated value of the preceding step and the updated drift parameter. We simulate the values this way because we find it likely that the DM will change her perception of the parameters over the investment horizon when accounting for regime shifts. The regime shifts and corresponding evolution for four example simulation paths are illustrated in Figure 3.4.

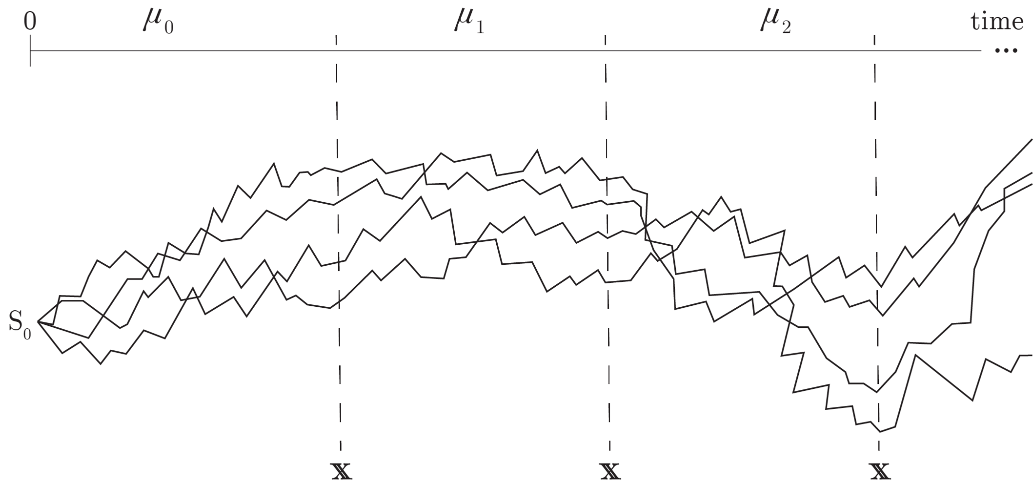

Figure 3.4: Illustration of the effects of signals on simulation paths of the price process.

Given that we update the prior belief using Bayesian inference for a normal-normal conjugate pair, the mean and variance of the posterior will be given by Equations 3.3 and 3.4 when B signals are received at once in a time step. For simplicity, the DM receive one signal every j^{th} time step, i.e. $B = 1$.

$$\mu_{post} = \left(\frac{B}{\sigma_s^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{B}{\sigma_s^2}\left(\frac{1}{B}\sum_{i=1}^{B} x_i + \frac{1}{\tau_0^2}\mu_0\right)\right] \tag{3.3}$$

$$\tau_{post}^2 = \left(\frac{B}{\sigma_s^2}\right)^{-1} \tag{3.4}$$

For illustration purposes of the methodology, the signals are implemented through random samples from a known distribution. This is done to infer how the possibility of receiving signals affects the decision problem. Alternatively, one could use observed signals with an estimated variance. We use the first approach since this paper is an analysis of the new methodology in a hypothetical setting, where we do not have any observed data. By knowing the true signal distribution, we can also analyse whether the results derived with the solution approach behave as expected. With regards to the model implementation, there is no difference between these two signals. Implementing a signal distribution with known parameters only facilitates sensitivity analysis we will perform later in Section 4.

Since the signals are random draws from a normal distribution, there are many possible signal outcomes. H random signal samples are drawn and perceived as a single signal to decrease the variability. The more signals are drawn from the distribution, the closer to the true drift the received signal will be. There are two possible ways to update the prior with this perceived single signal. The first alternative is to calculate the average signal, $\bar{x} = \frac{1}{H}\sum_{i=1}^{H} x_i$, and use this one signal to update the prior and receive one posterior. The second solution method is to update the prior once for each signal sample draw, resulting in H different posteriors from the same prior. The posterior to use for further Monte Carlo simulations would be the average of these H posteriors. The two methods yield identical posteriors, and as such we state

Lemma 3.1. *When you have H random signal draws, updating the prior into H posteriors and taking their average is the same as having one average signal and updating the prior into one posterior.*

See Appendix A for derivations.

Next, we use the simulated variable values for the entire investment horizon to calculate the investment value and derive the optimal decision strategy.

3.2.2 Investment value and optimal decision strategy

We solve the optimal decision strategy for the investment problem using dynamic programming in a backwards fashion. First, we find the optimal decision strategy at the end of the investment horizon, when $t = T$. Postponing the investment is not an option, and the optimal exercise strategy is either to adopt or reject the investment opportunity. The value of adoption is equal to the payoff, P_T^n, calculated with the simulated variable values, and the value of rejection is zero. The value of the investment opportunity in the last time step is equal to

$$V_T^n = max[0, P_T^n].$$

Then the intermediate decision points are evaluated by iterating backwards over the investment horizon, from $t = T - 1$ to $t = 0$. At each decision point, we find the optimal decision between adoption, rejection and waiting. To find the value of waiting, we approximate the value of the investment opportunity in the next time step based on the current variable values. To approximate this continuation value, denoted CV_t^n, we use a conditional expectation function, denoted F_t. We regress the conditional expectation function of a specified order, d, for some given basis functions. We use polynomials[2] which have the advantage of being able to approximate a broad range of functions. The general equation for a multivariate polynomial function of order d and M variables, denoted by F_t, is given by[3]

$$F_t = \sum_{i=0}^{d} \sum_{j=1}^{\binom{i+M-1}{i}} \alpha_{ij} \cdot \prod_{m=1}^{M} S_{m,t}^{\,{}^{i}L_j^m}, \qquad \text{where} \sum_{m=1}^{M} {}^{i}L_j^m = i$$

[2] Examples of other possible basis functions are the Hermite, Legendre, Chebyshev, Gegenbauer, and Jacobi polynomials

[3] Thanks to Dr. Rouholah Ahmed for providing this formula

where $S_{m,t}$ is the value of the uncertain variable m at time t, and the a_{ij}'s are coefficients. ${}^{i}L_{j}$ is the set of all possible m-tuple combinations of non-negative integers from 0 to i, which is constrained by their summations needed to be equal to i. ${}^{i}L$ can be defined iteratively through the following equations

$$
{}^{i}L_{j}^{m} = \begin{cases} m_1 & = 0 : i \\ m_2, ..., m_l, ..., m_{M-1} & = 0 : \left(i - \sum_{q=1}^{l-1} m_q\right) \\ m_M & = i - \sum_{q=1}^{M-1} m_q \end{cases}
$$

We start with explaining the procedure for $T-1$. First, we regress the discounted investment value in time step T, δV_T^n, on the simulated variable values in $T-1$, $(S_{1,T-1}, ..., S_{M,T-1})$. $(S_{1,T-1}, ..., S_{M,T-1})$ are used as the independent variables. The discounted investment value, δV_T^n, is the dependent variable. We regress on all paths, resulting in N vectors, $\mathbb{S}_{T-1}^n$, and N corresponding dependent variables. Here we deviate from the procedure introduced by Longstaff and Schwartz (2001) who only consider the in-the-money paths (i.e. the payoff function is positive) for the regression. We regress on all paths because we introduce rejection as a third alternative. The DM will only reject the investment opportunity and receive nothing when both the payoff and the expected value of waiting are less than or equal to zero. The waiting cost must be paid each period the DM postpones the decision, so the option to wait is no longer for free as in Longstaff and Schwartz' model. The coefficients of the regression function are estimated using least squares. This gives the regressed conditional expectation function $\hat{F}_{T-1}(\mathbb{S}_{T-1}^n)$.

Using $\hat{F}_{T-1}$ and the corresponding asset values, $\mathbb{S}_{T-1}^n$, we calculate the continuation value, CV_{T-1}^n. If the DM decides to wait, she must pay the waiting cost. This results in the following expected value of waiting,

$$CV_{T-1}^n - W.$$

The optimal decision at time $T-1$ is the decision that corresponds to

$$max[0, P_{T-1}^n, CV_{T-1}^n - W].$$

Depending on the action taken in step $T-1$, the optimal value V_{T-1}^n equals either 0, P_{T-1}^n, or δV_T^n. These correspond to the optimal decision being to reject, adopt or wait respectively. Note that when the optimal decision is to wait, the corresponding optimal value, V_{T-1}^n, will be the investment value from the next period and not the regressed continuation value. The regressed continuation values are used to find the optimal decision only. Since V_{T-1}^n will be used to regress the continuation value in $T-2$,

V^n_{T-1} needs to equal the obtained payoff and not the regressed estimate. The presented procedure for $t = T - 1$ is repeated for the remaining intermediate decision points until $t = 1$.

In order to derive the value of the investment opportunity at time zero, we discount the identified values from the optimal decisions at time $t = 1$. This results in a discounted potential future cash flow (PFCF) for each path, denoted CF^n. We then find the potential future cash flow by taking the average for all the paths,

$$\text{PFCF} = \frac{1}{N} \sum_{n=1}^{N} CF^n.$$

The value of the investment opportunity, V_0, is then equal to

$$V_0 = max[0, P_0, PFCF]. \tag{3.5}$$

4 Results and discussion

In this section, we contextualise the model with an illustrative example and perform sensitivity analysis. First, we present the decision problem and the interpretation of the parameters in a specific decision context. The parameters are then assigned values for which the context is solved and analysed, and the results presented. We then perform sensitivity analysis where we compare the results to a single-layered learning approach where we only account for Markov learning and do not update the DM's beliefs. The aim of this section is to gain further insight into our model and its possible applications.

4.1 Illustrative example

We present an illustrative example to demonstrate a possible application of the model. The example also serves as a foundation for further considerations of the methodology. A firm has the option to buy an innovative production technology for an investment cost, K. The firm intends to use the new technology to introduce new innovative products to the market. However, they are unsure of the market potential for the new products. With the new technology, the firm will have the capacity to produce Q units. Here, adoption results in the firm investing in and using the production technology to sell the new product. It is assumed that all units will be sold, and the potential revenue from selling these units can be seen as a one-off payment if the firm carries out the investment. Alternatively, the firm can reject the new technology and will then not be able to produce the product. The firm must decide whether they want to invest in the production technology by time T. They consider the investment decision at every time step t. If the firm decides to delay the decision, it costs them W to compensate for the flexibility of the investment option. It is assumed that if the firm decides to invest in the new product, the production and sales will be executed instantaneously, and the firm will receive an immediate payoff.

The firm knows the price they would sell the product for today, but they are unsure about what price they can expect to receive in the future. The price is an uncertain variable, denoted S_1. Based on their beliefs, they use a GBM to represent the price the market is willing to pay. The GBM is modelled with a drift μ_{S_1} and standard deviation σ_{S_1}, and its SDE is given by Equation 3.1. The drift μ_{S_1} is uncertain, of which the firm has an initial belief, denoted μ_0. With time, the firm can update their initial beliefs of the price drift with signals they receive. This could be product statements from competitors. The signals are random draws from a normal distribution with a mean, μ_{true}, and a standard deviation, σ_{signal}. To implement this problem, we need to know the true price drift, μ_{true}. The signals are set to arrive every j^{th} time step.

The production process has a variable cost per product, which is another uncertain variable and denoted by S_2. For simplicity, the product consists primarily of a single raw material (e.g. copper), and the variable cost depends on the price of this material. The variable cost is known today but will change over time as the price of the raw material changes. We model the cost uncertainty using a mean-reverting process, which follows the SDE given by Equation 3.2. It is expected that the cost will revert to the price μ_{S_2} over time with a standard deviation σ_{S_2}. The payoff function is then given by

$$P_t = Q \times \left(S_{1,t} - S_{2,t}\right) - K.$$

4.1.1 Numerical results

In the following, we present numerical results of the described decision problem for parameters with assigned expository values, as shown in Table 4.1.

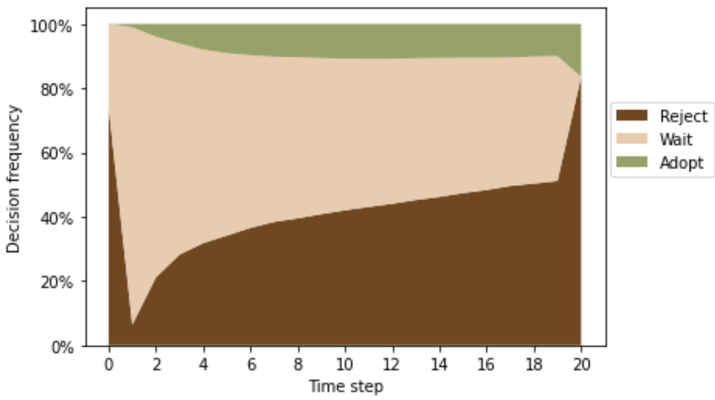

Figure 4.1: The decision frequency diagram resulting from solving the model using the parameters in Table 4.1.

Figure 4.1 illustrates the resulting decision frequency diagram. The optimal decision is to wait at an investment value of NOK 263,402, obtained by Equation 3.5. Although nearly 80% of the paths end in rejection at time zero, it is optimal to wait. This is due to the fact that the average value only needs to be higher than the waiting cost to make it more attractive to wait than to reject. In this illustrative example, the waiting cost is relatively low compared to potential future payoffs if the price increases.

Parameter	Interpretation	Value
Q	Number of units	10,000 units
T	Time horizon	5 years
W	Waiting cost	NOK 10,000
K	Investment cost	NOK 450,000
Δt	How often the investment decision is evaluated	0.25 years (quarterly)
j	Time steps between each signal received	4 (Receive signal annually)
r	Risk-free rate	1.05%
$S_{1,0}$	Initial value for the unit price variable	NOK 200
μ_0	Initial belief of price drift	0
τ_0	Initial variance of the belief, describing how uncertain the DM is in her initial prior	10
σ_{S_1}	Standard deviation of the price process	0,2
$S_{2,0}$	Initial value for the unit cost variable	NOK 200
μ_{S_2}	Expected mean-reverting unit cost	NOK 150
σ_{S_2}	Standard deviation of the cost process	0.32
η	Mean-reverting constant	0.22
μ_{true}	The signal drift, also the true price drift	-0.1
σ_{signal}	Standard deviation of the signal distribution	0.6
N	The number of Monte Carlo simulations	100,000
B	The number of signals received at a time	1
H	Signal sample size, the number of random draws that make up each signal	5
Ω	Number of intervals the entire investment horizon is divided into	20

Table 4.1: Summary of the model's parameters and their expository values.

Furthermore, the expected cost is lower than the current price. Therefore, the DM expects the cost to decrease in the future, which will contribute to obtaining a positive payoff. However, if the DM receives information signalling that the price will decrease in the future, the profit margin is reduced and will not cover the costs. Consequently, it is optimal to reject the new technology in most cases, even if the cost is decreasing. The signals will change the DM's beliefs in a negative direction over time. This can be seen in the decision frequency diagram as the rejection region increases in later time steps.

4.2 Sensitivity analysis

According to Howard and Abbas (2015), sensitivity analysis is an important feature of professional decision analysis. It enables us to investigate how the decision will change if we change certain numbers in the decision basis. This can help determine whether additional effort should be expended in increasing the parameter precision. As such, we perform a sensitivity analysis to obtain a broader understanding of the model's applicability and the effects of different parameter values on the investment decision. The sensitivity analysis is based on the presented illustrative example and its results. First, we analyse the impact of the parameters of the signal distribution. Second, we examine the impact of the parameters of the stochastic processes. Lastly, we present a situation where receiving signals results in a different decision compared to the single-layered learning approach without Bayesian updating. The parameter values are as given in Table 4.1 if nothing else is specified.

4.2.1 Impact of the signal distribution

The primary contribution of our model in terms of versatility is the opportunity to update the DM's beliefs about the parameters of the underlying stochastic processes. The updating is accomplished through Bayesian learning from signals given by a signal distribution. Therefore, it is of interest to examine the effects of changes to the signal sample size and to the initial mean of the signal distribution, and to infer how they affect the DM's beliefs.

First, we see how the signal sample size, H, impacts the results. The signal sample size is set to five to reduce extreme impacts of outliers, which are more likely when there is a single draw from the distribution. Since the signals are randomly sampled from the distribution, they may result in very different decision frequency diagrams. It is therefore of interest to map how changes to the signal sample size affect the investment decision. This is illustrated in Figure 4.2, where the normalised decision frequencies are shown for different signal sample sizes.

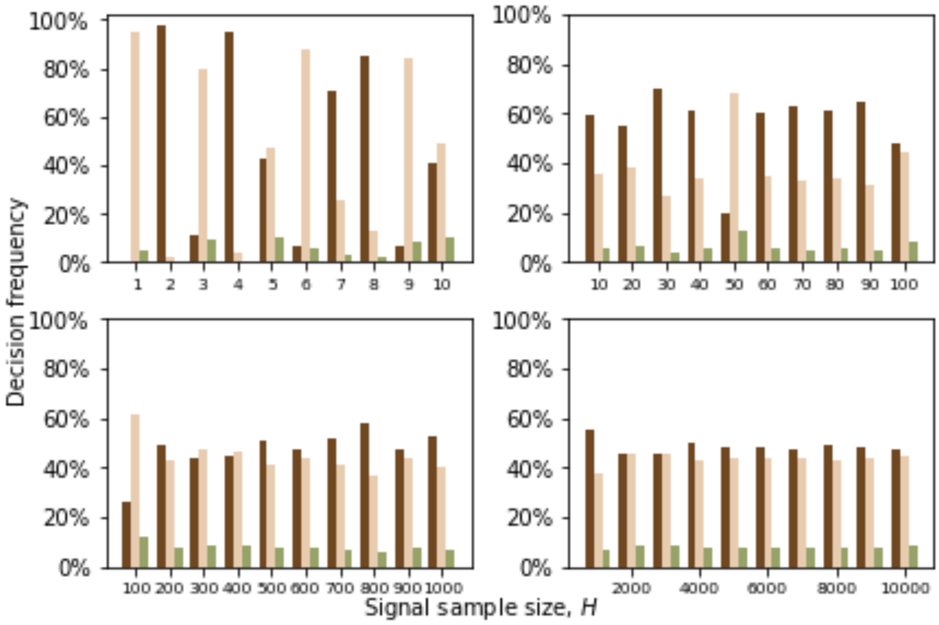

Figure 4.2: Decision frequency for all simulations over the investment horizon for different signal sample sizes depicted on the x-axis. The total percentage is of the different decisions for different signal sample sizes.

When the average signal is based on a smaller signal sample size, the variability in decision frequencies is higher, as illustrated by the upper left plot in Figure 4.2. When the signal sample size is increased, the decision frequencies stabilise. However, there is still apparent variability in the decision frequency for a signal sample size in the range of 100 to $1,000$. For signal sample sizes larger than $1,000$, the decision frequencies for the three alternatives are relatively stable. In later analyses, we will use $H = 1,000$ when we want to reduce the difference between the plots that arise due to variability in signal draws.

Next, we examine the impact of the beliefs' initial variance and signal variance on convergence. The beliefs' initial variance is the parameter that describes the DM's uncertainty in her prior beliefs, τ_0. τ_0 affects how fast the DM's beliefs converge. By converging, we mean when the variance of the DM's beliefs approaches zero, $\tau \longrightarrow 0$. The DM can become more or less confident in her beliefs by observing signals, and the uncertainty is the lowest when the beliefs' variance is 0. The lower τ_{post} gets, the less impact another signal has on the beliefs as the DM becomes increasingly confident in her estimate. Figure 4.3 illustrates how different start values of τ_0 converges when updated with the same signal volatility, σ_{signal}.

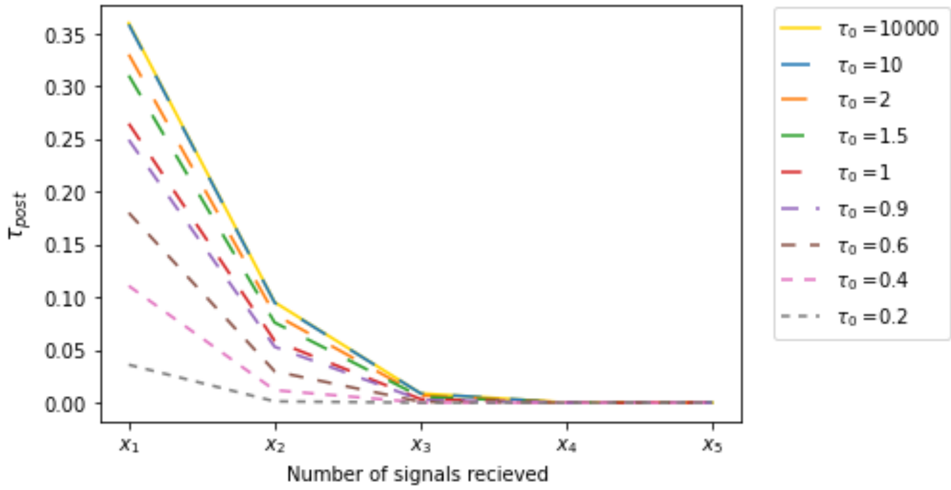

Figure 4.3: The evolution of τ_{post} with time for different values of τ_0 when signals with the same volatility, σ_{signal}, are observed over the investment horizon. Both $\tau_0 = 10,000$ and $\tau_0 = 10$ have the same development after the first signal is received. Note the large change from the initial uncertainty to τ_{post} after the first signal, x_1, is received, especially for large values of initial uncertainty, τ_0. All paths have values lower than 0.4 after updating on the first signal.

From Figure 4.3, we can see that a τ_0 larger than 10 converges to the same value after only one signal. Therefore, setting the initial τ_0 to an uninformative prior of 10 denotes that the DM has no knowledge to help her assess the true drift. When the DM has received the third signal, she will become confident enough in her beliefs that another signal makes little difference. Additional signals would then yield insignificant changes in the DM's uncertainty of her beliefs. Further, we investigate if and how the investment value, optimal decision and decision frequency diagrams change for different values of τ_0. However, we increased the signal sample size from $H = 5$ to $H = 1,000$ to reduce the difference between decision frequency distributions arising due to the variability of signal draws. We solved the problem for $\tau_0 \in [0.2, 0.6, 1, 1.5, 5, 10]$. The results are shown in Figure 4.4 below. We observe that the shape of the decision frequency diagram changes little when $\tau_0 \geq 1$, though the values change. The investment value changes for a τ_0 less than 5, while there are less drastic differences for τ_0 higher than 5.

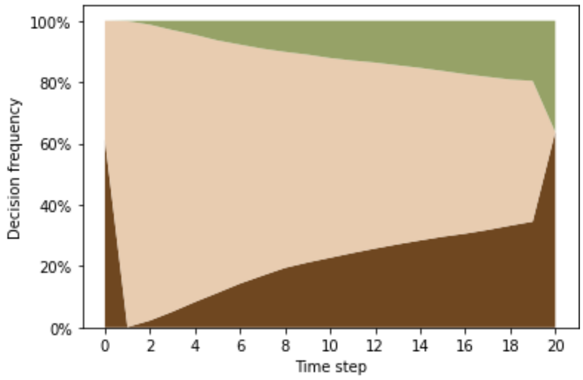

(a) $\tau_0 = 0.2$

Decision: Wait

Invetment value: NOK 201,655

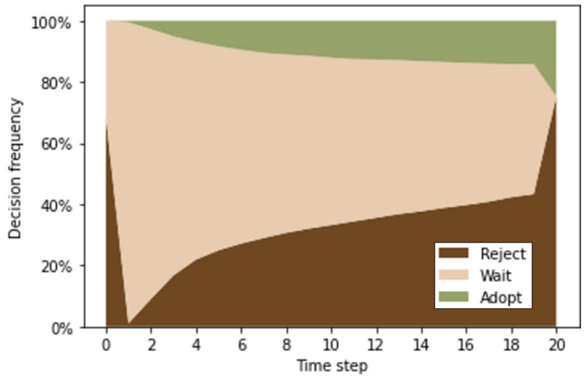

(b) $\tau_0 = 0.6$

Decision: Wait

Invetment value: NOK 110,684

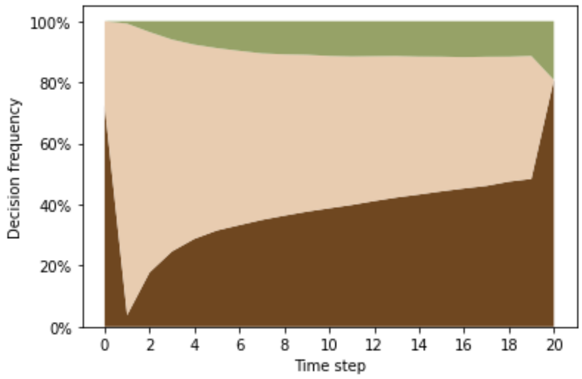

(c) $\tau_0 = 1$

Decision: Wait

Invetment value: NOK 79,845

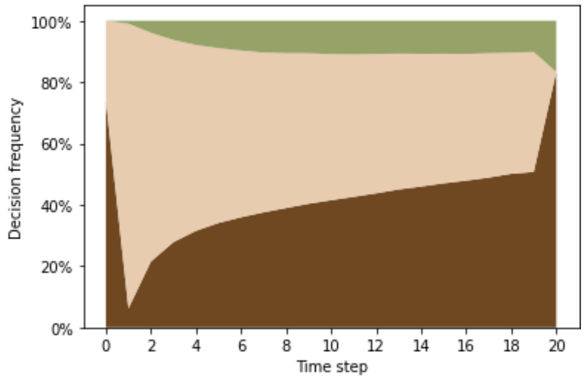

(d) $\tau_0 = 1.5$

Decision: Wait

Invetment value: NOK 68,701

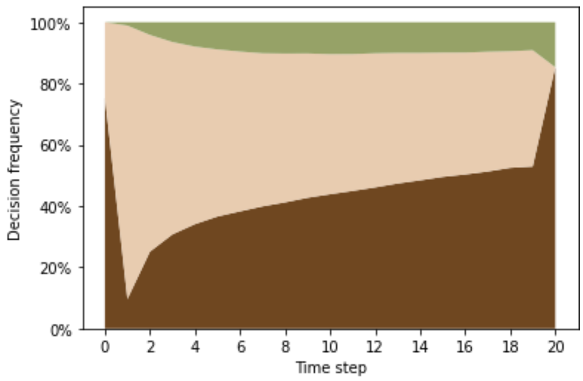

(e) $\tau_0 = 5$

Decision: Wait

Invetment value: NOK 59,778

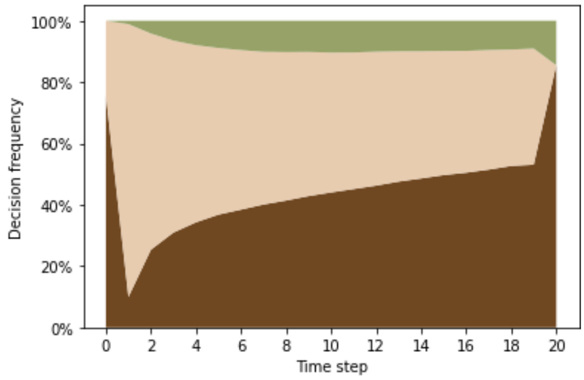

(f) $\tau_0 = 10$

Decision: Wait

Invetment value: NOK 59,034

Figure 4.4: Decision frequency diagrams with the corresponding optimal decision and investment value for different values of the DM's initial uncertainty, τ_0.

Continuing with this line of investigation, it is of interest to examine how the signal volatility affects the convergence speed. The signal volatility, σ_{signal}, is a measure of how far each sampled signal can be from the distribution's centre, μ_{true}. Since the signals vary more for a higher σ_{signal}, the DM must receive more signals for the variance of her beliefs' to converge. Figure 4.5 illustrates convergence speed for τ given different values of σ_{signal}.

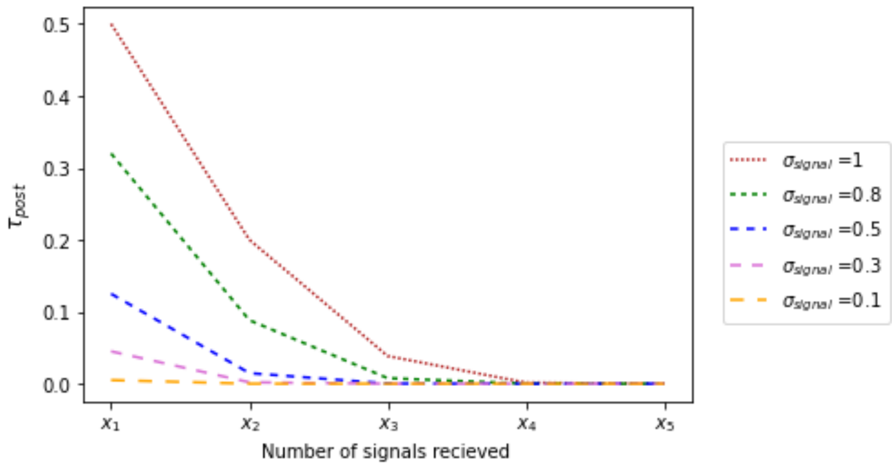

(a)

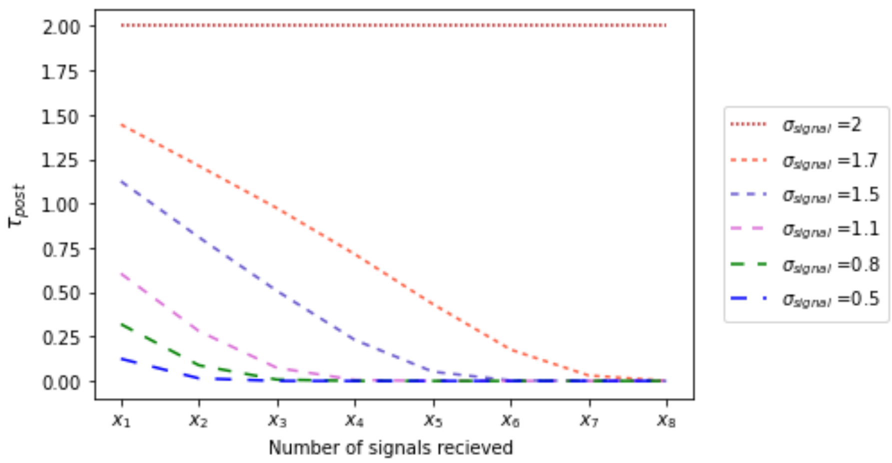

(b)

Figure 4.5: Change in the beliefs' variance, τ, with time for different values of signal volatility, σ_{signal}, and the same initial variance, $\tau_0 = 10$, as stated in Table 4.1 over the investment horizon. The top plot shows the beliefs' variance for $\sigma_{signal} \in [0.1, 1]$. The bottom plot shows the beliefs' variance for $\sigma_{signal} \in [0.5, 2]$.

As shown in Figure 4.5, the volatility of the DM's beliefs has converged by the fourth signal for all $\sigma_{signal} \in [0.1, 1]$. However, when σ_{signal} is above 1, convergence is slower and requires the observation of more signals. The top plot, where $\sigma_{signal} \in [0.1, 1]$, shows that all paths converge after receiving four signals. The bottom plot, where $\sigma_{signal} \in [0.5, 2]$, shows that seven signals are needed to obtain approximately the same value of τ for all $\sigma_{signal} < 2$. A combination of $\sigma_{signal} \geq 2$ and $\tau_0 \geq 2$ will never converge towards zero, as shown in Figure 4.6. In this case, the DM will never approach complete confidence in her beliefs of the drift, regardless of how many signals she collects. We can conclude that the combinations of σ_{signal}, τ_0 and H affects the decision frequencies. A combination of higher values of σ_{signal} and τ_0 will give more variable decision frequency distributions with lower H values, while a higher H reduces this behaviour.

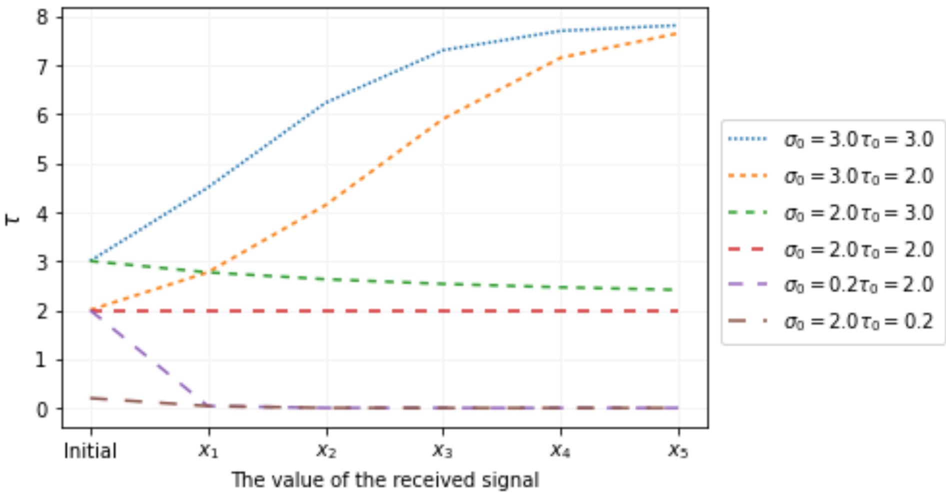

Figure 4.6: The change in the beliefs' variance for different combinations of signal volatility, σ_{signal}, and initial variance of her beliefs, τ_0.

Third, we examine the impact on the investment decision of changing the combination of the signal distribution's mean and variance, μ_{true} and σ_{signal}. We vary μ_{true} from -0.5 to 0.5 and σ_{signal} from 0.1 to 1.7 to assess these effects. As inferred previously in this section, the signal sample size is of great significance for the results. It often impacts the decisions in later time periods, although it may not change the DM's immediate decision at time step zero. Figure 4.7 shows how the decision frequency diagram changes for different values of μ_{true} and σ_{signal} for multiple values of H.

A lower-valued σ_{signal} have a greater impact on situations where μ_{true} is closer to zero. This can be explained by the fact that the signals will more easily change the believed direction of price evolution. When the absolute value of μ_{true} is large, the σ_{signal} must also be large to receive signals that will change the DM's beliefs of the direction of the evolution in unit price. Regarding the impact on the decision frequency, σ_{signal} has less impact for higher values of H. For higher values of H, the true drift of the price process is the signal parameter with the most considerable impact. However, σ_{signal} still influences the convergence speed for the DM's beliefs, as discussed above.

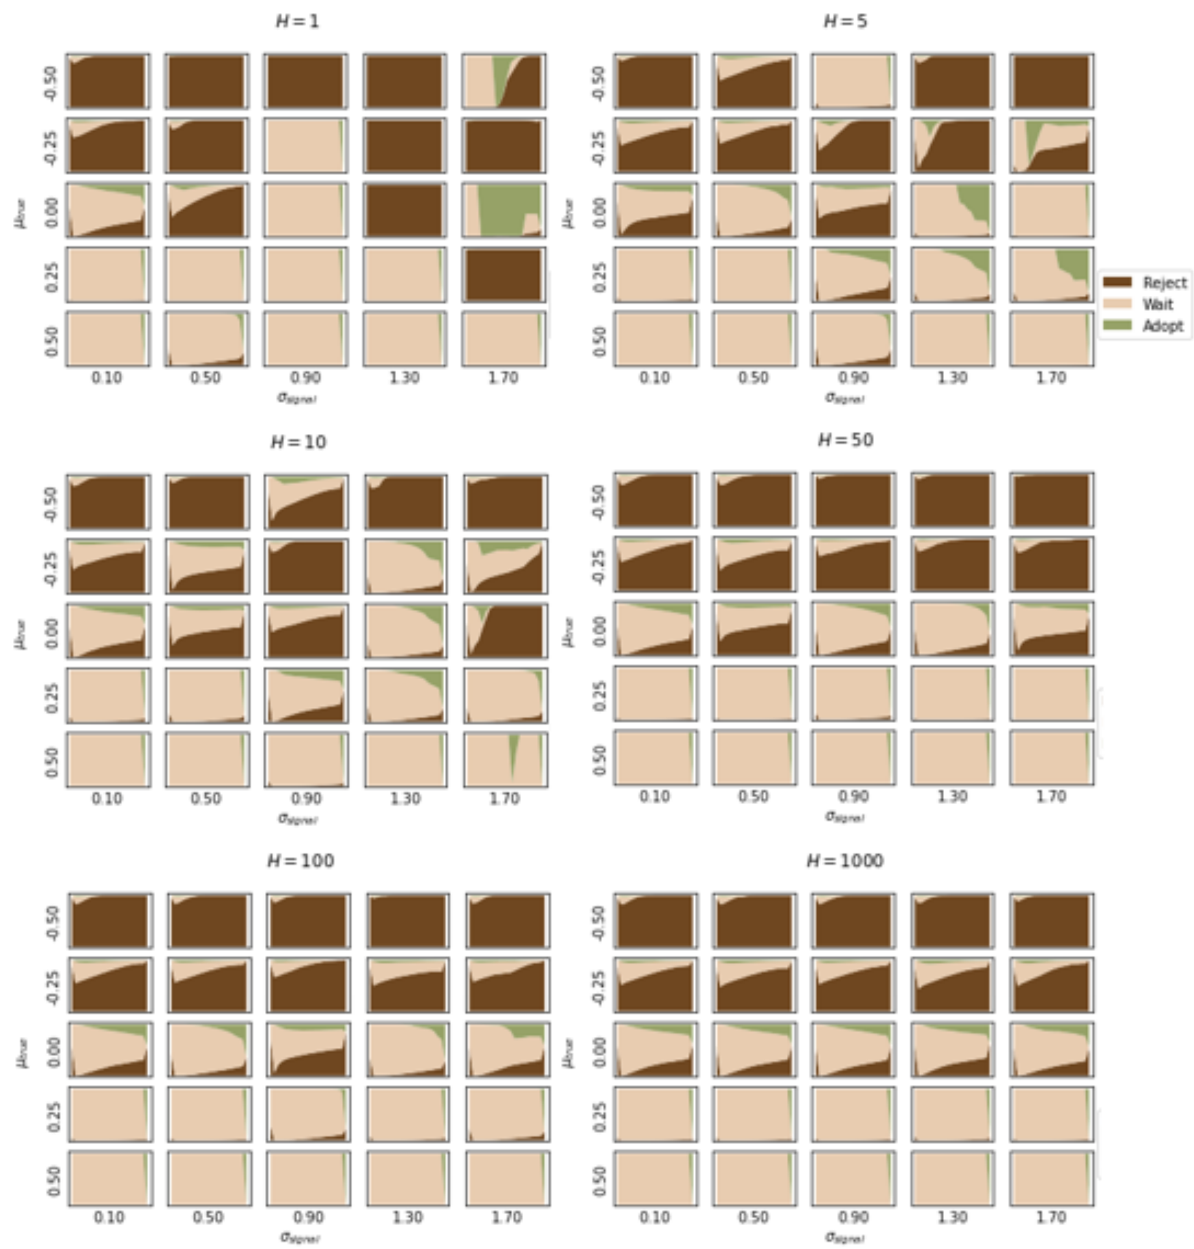

Figure 4.7: Decision frequency diagrams for different combinations of signal sample size, H, true drift, μ_{true}, and signal variance, σ_{signal}.

Lastly, we examine the impact of the signals given the initial variance of the DM's beliefs regarding the drift of the stochastic price process, τ_0. We start by examining when the DM is very confident in her beliefs, setting τ_0 to 0.3. Figure 4.8 shows how the beliefs change for five different consecutive signals received. The different lines represent different values of the initial prior, μ_0. The numbers on the horizontal axis are the signal's value and are shown in the order they are received. We can observe that when the DM is relatively confident in her prior, the impact of the first signal is minimal, even if the difference in values between the prior and the signal is relatively big. The resulting posteriors after five signals are close to their original priors and do not converge to a unified posterior. In other words, the prior has a significant impact when the DM is confident in her beliefs. Previously we found that the τ parameter converges rapidly for a σ_{signal} of 0.6 and that the DM will be very confident after only two signals. This is confirmed in Figure 4.8, where it can be seen that the posteriors of the different lines change insignificantly after the first signal of -0.20. In addition, later signals have little impact when the DM is confident. To summarise, we can say that when the DM is very confident in her initial beliefs, the signals have little effect on the resulting posterior.

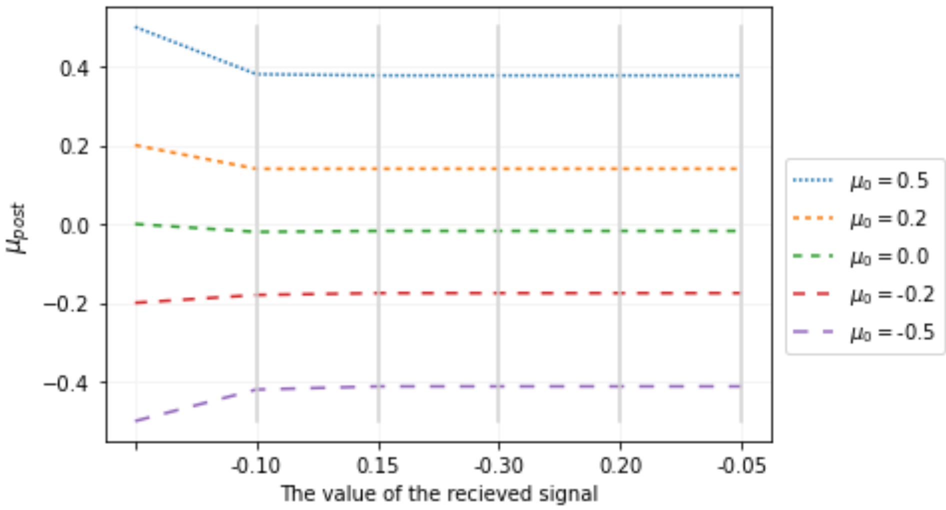

Figure 4.8: The resulting posterior for different priors, μ_0, given an initial uncertainty $\tau_0 = 0.3$ and the signals received in the stated order over the investment horizon.

Now we examine the situation where the DM is very uncertain and has little confidence in her beliefs, setting τ_0 to 10. The resulting graph is given in Figure 4.9 for the same signals received and the same initial priors as in Figure 4.8. Independent of the initial prior, we see that all five paths end in the same posterior value. In the light of this, we can similarly conclude that when the DM has little confidence in her initial beliefs, the prior has little impact on the resulting posterior. In addition, we can make the same observation as from Figure 4.8 that the posterior does not change much after the second signal.

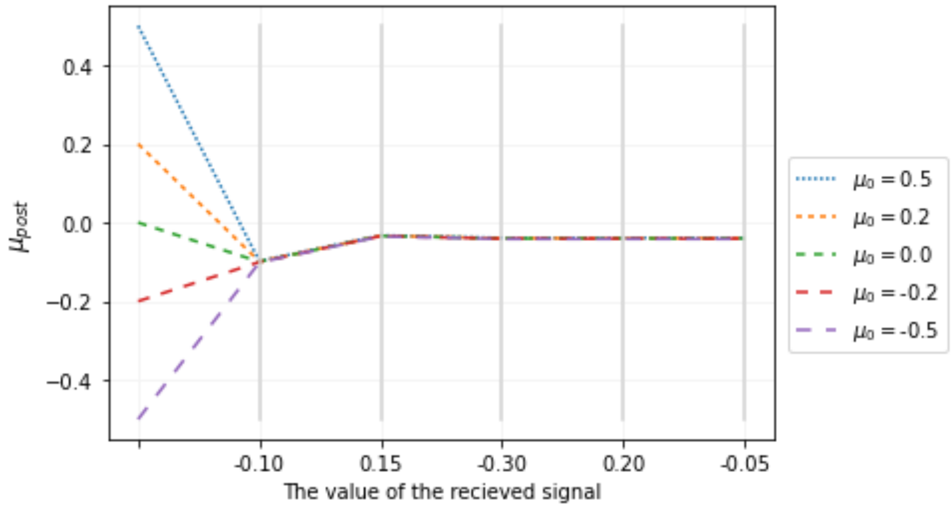

Figure 4.9: The resulting posterior given an initial uncertainty $\tau_0 = 10$ and the signals received in the stated order over the investment horizon for different priors, μ_0.

Further, in Figure 4.10, we can see how the posterior changes and evolves for different values of the initial uncertainty of the DM's beliefs, τ_0, and a prior belief, μ_0, of zero. It can be seen that the less confident the DM is about her initial prior, the more the posterior changes when signals are received. Meaning that, for the larger values of τ_0, later signals have more impact than for lower values of τ_0. However, later signals have less impact in general because of convergence, as discussed previously.

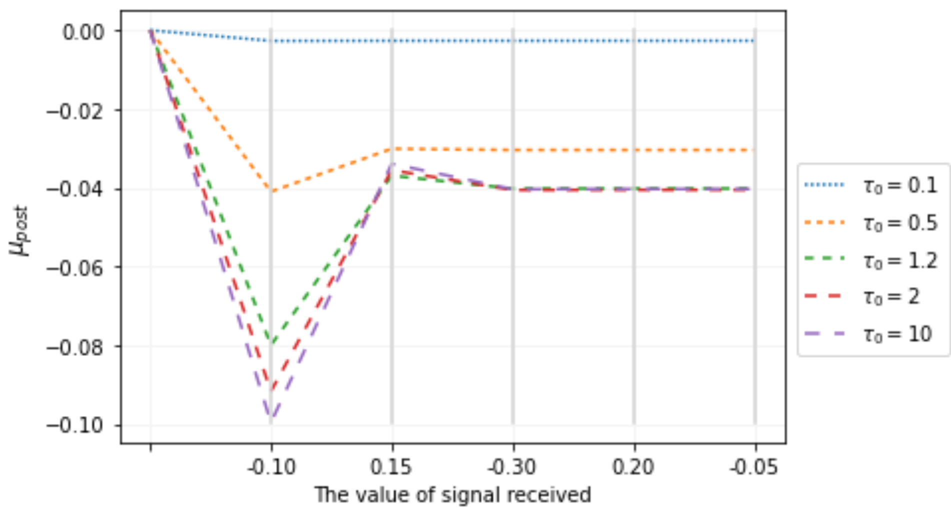

Figure 4.10: The resulting posterior given an initial prior of zero and σ_{signal} of 0.6 for different initial uncertainties and the signals received in the stated order over the investment horizon.

4.2.2 Impact of the stochastic process parameters

In this section, we examine the effects of changing some of the parameter values in the stochastic processes of the uncertain variables. The signal sample size is increased from $H = 5$ to $H = 1,000$ to decrease the variability between the plots. We compare the effects to the same changes applied in a traditional option valuation model where the parameter updating feature is removed. That is, Bayesian learning is not performed, and the DM does not update her beliefs over time.

We start by looking at sensitivity to the initial unit price, $S_{1,0}$. When increasing the current price from NOK 0 to 1,000, the results in Figure 4.11a indicate that the investment value shown by the dotted pink line increases linearly once the price is sufficiently higher than the unit cost to cover the fixed investment cost. Similarly, in Figure 4.11b, we see that the optimal decision is to reject for prices up to approximately NOK 200. This is reasonable as the immediate payoff is negative. In addition, the DM will receive signals that the price will decrease further in the future. Then there is a small interval where the optimal decision is to wait before it changes to adoption for even higher prices.

If we omit the signals and do not account for future learning, then the corresponding plots are given as the brown dotted lines in Figures 4.11a and 4.11b. We can see that the investment value increases in a similar fashion as before and that the optimal decision changes from rejection to waiting at approximately the same spot. However, the optimal decision is to wait for higher prices compared to the models with signals where the decision changes to adopt. This is expected as the DM now believes that the price will stay unchanged while the cost will fall in the future. Even if the potential gain is positive today, it is expected to be higher in the future, and the DM will therefore prefer to wait.

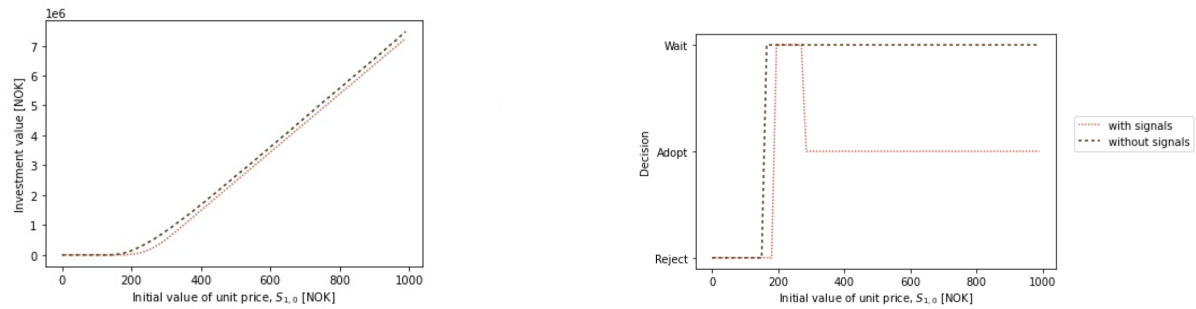

(a) (b)

Figure 4.11: Changes to the investment value and optimal decision when the current unit price varies from NOK 0 to 1,000. The pink dotted line is the situation where the DM receives signals. The brown dotted line is for the situation where signals are ignored.

Further, we examine the effect of changing the initial beliefs of the true drift. If we vary μ_0 from -1 to 1, we see from the pink dotted lines in Figures 4.12a and 4.12b that neither the investment value nor decision changes much. Changing this variable is equal to changing the prior beliefs, and the signals can therefore explain the unaffected investment value and decision. If the DM is uncertain in her prior beliefs, this has less impact as the signals will change her beliefs. When the signals indicate a drift of -0.1, the investment value will not increase even for high values of μ_0 because the DM is very uncertain in her beliefs. The impact of the prior beliefs with a high uncertainty was discussed in Section 4.2.1.

On the other hand, if we omit the signals, the graphs change to those given as the brown dotted lines in the same figures, and the results differ significantly. The investment value is increasing exponentially, and the optimal decision changes from reject to wait right before μ_0 reaches zero. This is reasonable because when the DM's initial beliefs are that the true drift is negative, she believes the price will continue to decrease further in the future. There is a small region where the drift is negative, and the DM believes the unit cost may decrease faster than the unit price. As such, she will wait to learn what the realisations will be. Consequently, it will never be optimal to produce the new product as the current payoff is negative. When the prior beliefs are higher than 0.5, the DM believes the price will increase substantially in the future. This results in a large potential gain in the future, therefore the high investment value and that the optimal decision is to wait.

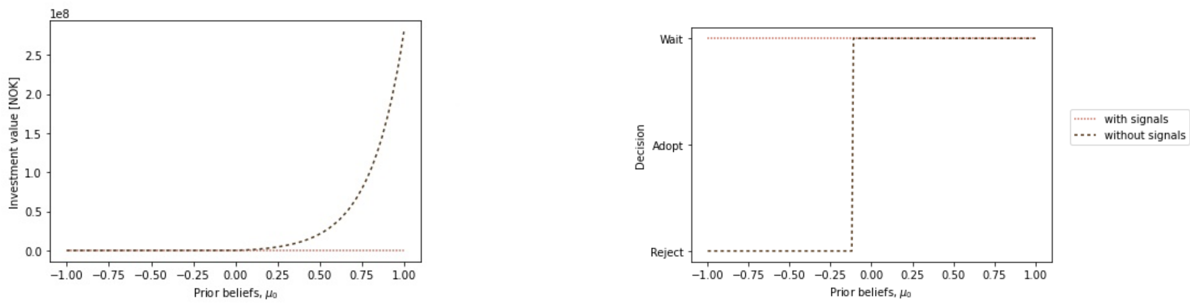

(a) **(b)**

Figure 4.12: Changes to the investment value and optimal decision when the prior beliefs vary from -1 to 1. The pink dotted line is the situation where the DM receives signals. The brown dotted line is for the situation where signals are ignored.

Next, we examine the effects of changing the parameters of the mean-reverting stochastic process used to represent the cost. Figures 4.13a and 4.13b show the impact of varying the current unit cost, $S_{2,0}$, from NOK 0 to 500 presented as dotted pink lines. We see that the investment value decreases as the cost increases, which is reasonable as a higher cost reduces the payoff and makes the investment opportunity less attractive. Furthermore, we see in Figure 4.13b that the optimal decision is to adopt for low costs. This can be explained that a low cost yields a positive immediate payoff and that the DM expects the cost to increase in the future. Because the DM will receive signals over time indicating that the price will decrease, it is optimal to produce the product now rather than delaying the decision. Thereafter, it is optimal to wait for a small interval around the current sales price, which can be explained by the uncertainty of an expected decrease in both unit cost and unit price. Depending on which decreases fastest, the DM could change her decision whether to invest or reject. Therefore, it is optimal to wait. If the current unit cost is sufficiently higher than the current unit price, it is optimal to reject.

When we omit the signals, the results are given as dotted brown lines in Figures 4.13a and 4.13b. The investment value decreases in a similar fashion, but the waiting interval is longer in the single-layered model. This is due to the fact that, without the signals, the DM believes the price will stay relatively unchanged over the investment horizon while the cost will decrease, as explained before. This increases the interval of unit cost where the DM is willing to wait.

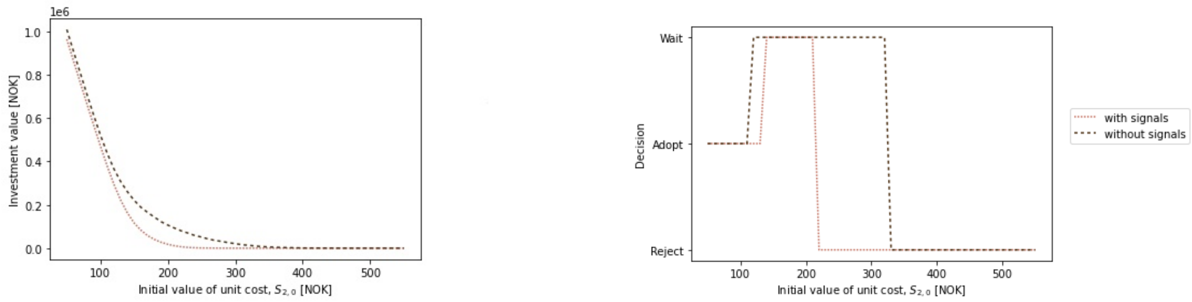

(a) **(b)**

Figure 4.13: Changes to the investment value and optimal decision when the current unit cost varies from NOK 0 to 500. The pink dotted line is the situation where the DM receives signals. The brown dotted line is for the situation where signals are ignored.

Lastly, we vary the expected mean-reverting cost, μ_{S_2}, from NOK 50 to 500. From Figure 4.14a, we see that the investment value decreases when μ_{S_2} increases, similar to Figure 4.13a above. From Figure 4.14b, we see that the optimal decision changes from wait to reject as the expected cost increases, occurring right before NOK 200. This is because when the expected cost is lower than the current cost, the DM expects the cost to decrease in the future. Even if she receives signals that the price will decrease too, the DM should wait as she expects the new technology to be profitable soon if the cost is expected to decrease more. The immediate payoff is negative, and it will therefore never be optimal to adopt at the current time with the current price, independent of the expected cost. When the expected cost is high, the DM expects the cost to increase, and she receives signals indicating that the price will decrease. This makes the investment less attractive with time, and it should therefore be rejected immediately.

When the signals are omitted, the values and decisions change to the dotted brown lines in Figures 4.14a and 4.14b for the specified values of μ_{S_2}. The investment value is of a similar form, but not value, and the optimal decision policy changes. In this situation, the DM will be willing to wait for a higher expected cost. Intuitively, we would expect the DM to reject the investment opportunity for these higher costs since the cost is expected to increase over time and the price to stay unchanged. However, the DM needs only a few realisations of a higher unit price to receive a bigger payoff than either of the immediate exercise actions. Since the uncertain variables are modelled as stochastic processes, and the expected values are calculated using Monte Carlo simulations, the payoff may be positive in some of the paths. Given a relatively low waiting cost, the DM is willing to wait because the potential future payoff is higher than the current one.

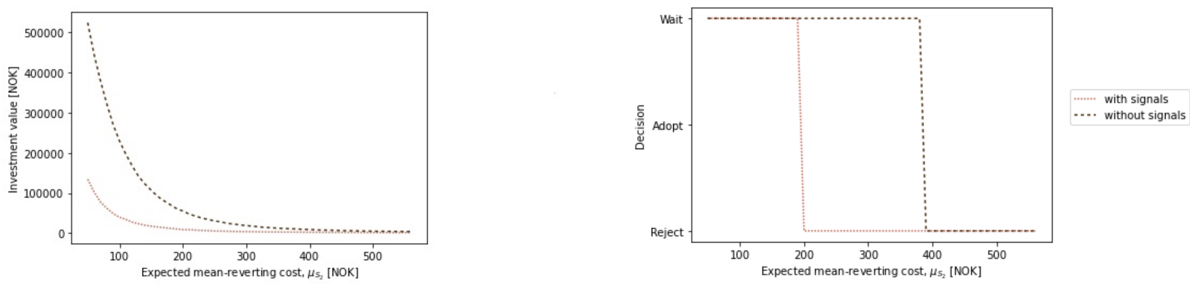

(a) **(b)**

Figure 4.14: Changes to the investment value and optimal decision when the expected mean-reverting cost varies from NOK 0 to 500. The pink dotted line is the situation where the DM receives signals. The brown dotted line is for the situation where signals are ignored.

How the decision changes for different combinations of initial unit cost and initial unit price when signals are received are illustrated in Figure 4.15a. In accordance with expectations, the optimal decision is to adopt the technology for higher values of unit prices in combination with low unit costs. Similarly, for lower unit prices in combination with higher unit costs, it is optimal to reject the technology. Both are reasonable as the combinations respectively increase and decrease the profit margin, and thereby the profitability. Figure 4.11b and 4.13b are excerpts of these figures. Similarly, Figure 4.15b illustrates the decision changes when signals are not received. The area where it is optimal to reject is quite similar in both figures. The waiting and adoption area, on the other hand, are distinctly different. It is much more often optimal to wait at time zero when signals are not received compared to Figure 4.15a. Again, this can be explained by the DM believing the investment will be more valuable in the future. She would therefore prefer to delay the investment than investing now. However, we can see that she will make several suboptimal decisions when she does not receive signals.

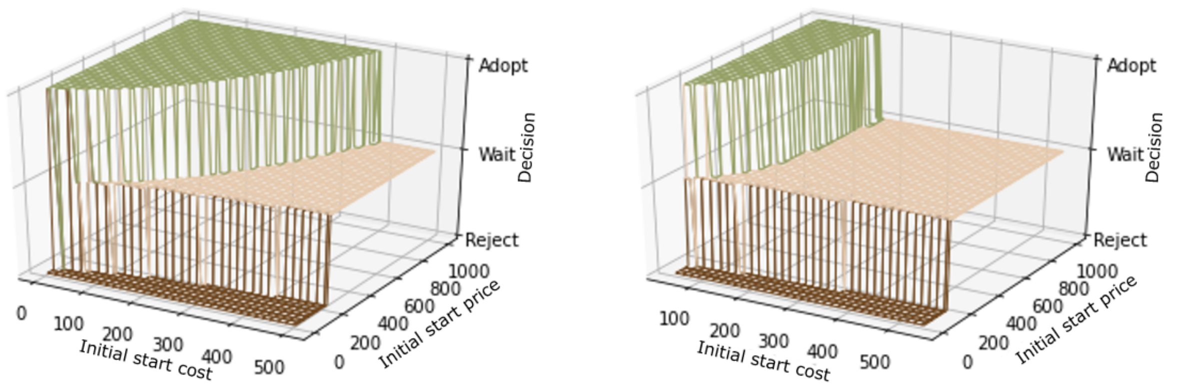

(a) **(b)**

Figure 4.15: Illustrates how the optimal decision changes a) when accounting for signals and b) without signals for different combinations of initial unit cost and initial unit price.

4.2.3 Different decision compared to a single-layered learning model

This section examines a situation that yields different optimal decisions when accounting for signals and not. For this purpose, the initial sales price, $S_{1,0}$, is set to NOK 400, where the immediate payoff is positive. The decision frequency diagram is presented in Figure 4.16a and has a signal sample size of five. Here the optimal decision was to adopt the process technology and produce the new product today with an investment value of NOK 1,550,000. The DM receives information that the price will decrease with time. Since the immediate payoff is positive, it is expected that adoption is the most frequent decision in the earlier time steps and rejection in the later time steps. Note that this diagram is affected by the signal sample size. If the signal sample size instead is set to 1,000, the average signals are closer to the true drift, and the decision frequency diagram would look like the one in Figure 4.16b. The optimal decision is to adopt, and therefore the investment value equals the immediate payoff and is the same for both signal sample sizes. However, there are some apparent differences between the two graphs. The relative frequency of waiting is larger at time zero for the diagram with the largest signal sample size. Furthermore, the rejection region is significantly smaller in this diagram. The findings in Section 4.2.1 and 4.2.2 can explain these differences.

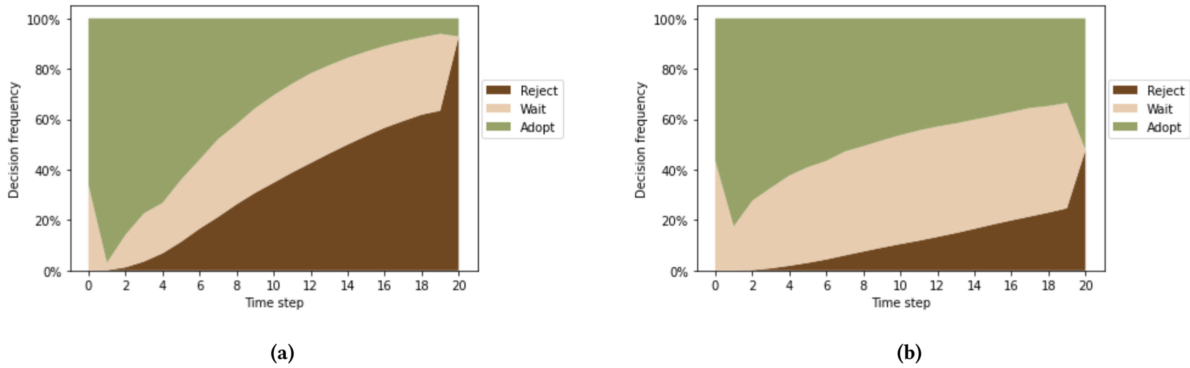

Figure 4.16: Decision frequency diagrams where signals are received for two different signal sample sizes, respectively $H = 5$ and $H = 1000$.

When not receiving signals, the optimal decision is to wait at an investment value of NOK 1,772,366. The decision frequency diagram is given in Figure 4.17. Now the DM believes that the unit sales price will stay relatively unchanged over time. However, she believes that the unit cost will fall in the future because the current cost of NOK 200 is above the expected cost of NOK 150. In this specific case, where the waiting cost is low relative to the potential future payoff, the optimal decision is to wait. The DM will no longer receive information that the drift is negative, which yields a different result than when applying the multi-layered learning approach. The investment value when signals are not received is higher than when signals are received. Intuitively, one might think that this is odd as the value of real options usually increases compared to the net present value as one introduces more flexibility. The values resulting from solving our model with and without signals cannot be compared as such. Instead, the positive difference tells us that the DM overestimates the investment value when signals are not accounted for. The DM's initial beliefs of the price drift are higher than its true value, and she finds it optimal to delay the decision instead of adopting now. However, if she delays the decision, she will observe that the price drift falls, and she will receive less than if she adopted at time zero.

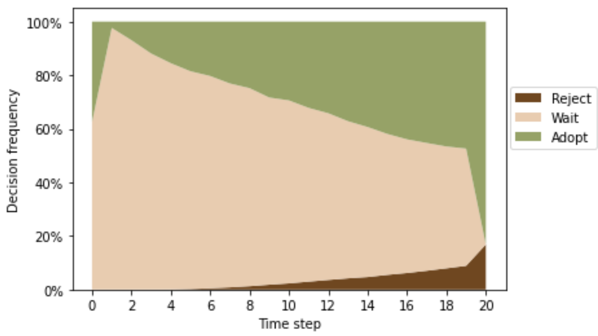

Figure 4.17: Decision frequency diagram where signals are not received. The signal sample size is therefore irrelevant.

In the situation where the DM's prior beliefs are equal to the true drift, she would reach the same decision using both the single-layered and the multi-layered learning model. This is shown in Figure 4.18, where the decision makers' prior beliefs are equal to the true drift, i.e. $\mu_0 = \mu_{true} = -0.1$. In this case, both models result in the optimal decision to adopt immediately at an investment value of NOK 1,550,000. The signals will indicate to the DM that the uncertain drift is the same as her beliefs and as a result she will not change her mind. Given that the signals' variance is small enough such that the DM's beliefs will converge (as inferred and discussed in Section 4.2.1), the optimal decision and the corresponding investment value will be the same for both models.

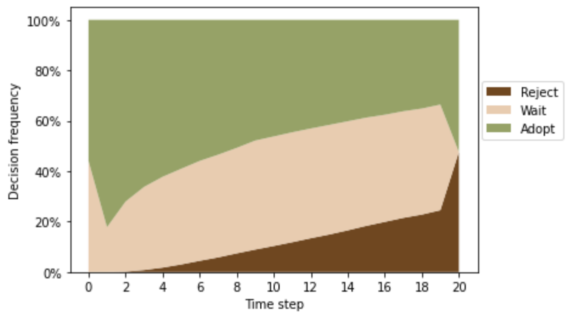

(a) Single-layered model result

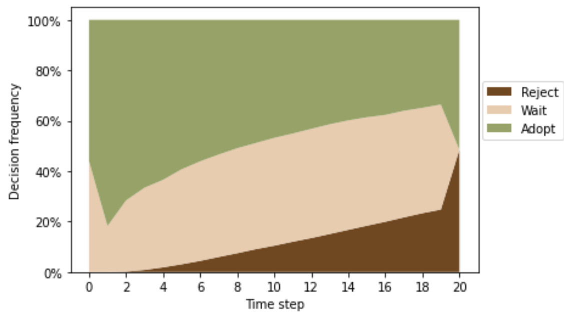

(b) Multi-layered model result

Figure 4.18: Decision frequency diagram when the DM's initial prior equals the true drift of the price process.

5 Conclusion

This paper presents a new methodology for solving a sequential decision making problem under multiple uncertainties where a second level of learning is enabled through receiving signals. The first level of learning occurs over time from the Markov property of the stochastic processes used in the Least Squares Monte Carlo approach. With the second level, the decision maker will periodically receive signals providing information about the hyperparameters of the stochastic processes. Using the signals, the decision maker can update her prior beliefs for multiple hyperparameters in a Bayesian manner.

Our methodology provides insight for a decision maker into the value of information by identifying the optimal decision policy for a new technology investment and its corresponding value. The methodology is flexible and can be applied to a broad range of decision problems with different payoff structures. It allows for multiple uncertain variables without increasing the complexity, in contrast to other models that often suffer from the curse of dimensionality. The added flexibility of allowing the decision maker to update her beliefs about one or more of the underlying uncertain variables is providing additional value and insight into the decision context. For simplicity we illustrate our methodology based on an example where the uncertainties are modelled as one-factor stochastic processes. However, our approach can easily be extended to account for uncertainties modelled as multi-factor processes.

We find that receiving signals and updating the beliefs can significantly impact the investment value and optimal decision policy of the investment problem. If the signals differ sufficiently from the prior beliefs, a different decision policy and investment value are often reached compared to the single-layered learning approach commonly used in traditional real options valuation. This happens because the decision maker can overestimate or underestimate the investment value when Bayesian learning is not applied. It is not easy to draw general conclusions of the model because it depends greatly on the context and chosen parameter values, its payoff function, and the modelling choice of the signals. The main aim of this paper is to present the methodology. However, some general insights have been established. First, when the prior beliefs equal the true signal distribution, the model yields the same results as the model without the signal feature. Furthermore, if the decision maker's prior beliefs imply a significant lack of knowledge, the prior has limited impact on the decision policy. Likewise, if the decision maker is very confident in her prior beliefs, the signals are of less significance and have limited impact.

The paper lays the groundwork for various extensions and future research. First, the model framework allows for other ways of Bayesian updating than through conjugate pairs. If the signal distribution is better represented by a non-conjugate prior, Markov Chain Monte Carlo methods can be used. This would extend the versatility and applicability of the methodology. Furthermore, it is of interest to analyse how updating the beliefs of multiple, possibly dependent, uncertain parameters impacts the decision policy. Another possible extension is to apply our novel solution approach to technology adoption problems where the uncertainties are modelled by multi-factor processes to obtain further methodology validation. Lastly, it could be of interest to extend the model to accommodate for random arrival of signals, e.g., by using a Poisson process. It would then be possible to have a varying number of signals arriving over the investment horizon. This would also allow varying lengths of the regimes.

References

Barraquand, J. and Martineau, D. (1995). Numerical valuation of high dimensional multivariate american securities. *The Journal of Financial and Quantitative Analysis*, *30*(3), 383–405. https://doi.org/10.2307/2331347

Blanco, G., Olsina, F., Garces, F. and Rehtanz, C. (2011). Real option valuation of facts investments based on the least square monte carlo method. *IEEE Transactions on Power Systems*, *26*(3), 1389–1398. https://doi.org/10.1109/TPWRS.2010.2094211

Bratvold, R. B. and Begg, S. H. (2010). *Making good decisions.* Society of Petroleum Engineers.

Broadie, M. and Glasserman, P. (1997). Pricing american-style securities using simulation. *Journal of Economic Dynamics and Control*, *21*(8-9), 1323–1352. https://doi.org/10.1016/S0165-1889(97)00029-8

Cortazar, G., Gravet, M. and Urzua, J. (2008). The valuation of multidimensional american real options using the lsm simulation method. *Computers Operations Research*, *35*(1), 113–129. https://doi.org/https://doi.org/10.1016/j.cor.2006.02.016

Dalby, P. A., Gillerhaugen, G. R., Hagspiel, V., Leth-Olsen, T. and Thijssen, J. J. (2018). Green investment under policy uncertainty and bayesian learning. *Energy*, *161*, 1262–1281. https://doi.org/10.1016/j.energy.2018.07.137

Diendorfer, C. (2019). Investment decision making for blockchain applications in the energy sector using real options valuation and bayesian learning [Master Thesis]. *Vienna University.*

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

Doraszelski, U. (2004). Innovations, improvements, and the optimal adoption of new technologies. *Journal of Economic Dynamics and Control*, *28*(7), 1461–1480. https://doi.org/10.1016/S0165-1889(03)00112-X

Farzin, Y. H., Huisman, K. J. M. and Kort, P. M. (1998). Optimal timing of technology adoption. *Journal of Economic Dynamics and Control*, *22*(5), 779–799. https://doi.org/10.1016/S0165-1889(97)00097-3

Hagspiel, V., Huisman, K. J. M. and Nunes, C. (2015). Optimal technology adoption when the arrival rate of new technology changes. *European Journal of Operational Research*, *243*(3), 897–911. https://doi.org/10.1016/j.ejor.2014.12.024

Harrison, J. M. and Sunar, N. (2015). Investment timing with incomplete information and multiple means of learning. *Operations Research*, *63*(2), 442–457. https://doi.org/10.1287/opre.2015.1344

Howard, R. A. and Abbas, A. E. (2015). *Foundations of decision analysis*. Pearson.

Insley, M. C. and Wirjanto, T. S. (2010). Contrasting two approaches in real options valuation: Contingent claims versus dynamic programming. *Journal of Forest Economics*, *16*(2), 157–176. https://doi.org/10.1016/j.jfe.2009.11.002

Jensen, R. A. (1982). Adoption and diffusion of an innovation of uncertain profitability. *Journal of Economic Theory*, *27*(1), 182–193. https://doi.org/10.1016/0022-0531(82)90021-7

Kruschke, J. (2014). *Doing bayesian data analysis*. Academic Press Inc.

Kwon, H. D. and Lippman, S. A. (2011). Acquisition of project-specific assets with bayesian updating. *Operations Research*, *59*(5), 1119–1130. https://doi.org/10.1287/opre.1110.0949

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

Martzoukos, S. H. (2003). Real rd options with endogenous and exogenous learning. Butterworth-Heinemann. https://doi.org/10.1016/B978-075065332-9.50007-0

McCardle, K. F. (1985). Information acquisition and the adoption of new technology. *Management Science*, *31*(11), 1335–1473. https://doi.org/10.1287/mnsc.31.11.1372

Rodrigues, A. and Armada, M. J. R. (2006). The valuation of real options with the least squares monte carlo simulation method. *Risk Management eJournal*. https://doi.org/10.2139/ssrn.887953

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

Sabour, S. A. A. and Poulin, R. (2006). Valuing real capital investments using the least-squares monte carlo method. *The Engineering Economist*, *51*(2), 141–160. https://doi.org/10.1080/00137910600705210

Smith, J. E. and McCardle, K. F. (1998). Valuing oil properties: Integrating option pricing and decision analysis approaches. *Operations Reasearch*, *46*(2), 198–217. https://doi.org/10.1287/opre.46.2.198

Smith, J. E. and Nau, R. F. (1995). Valuing risky projects: Option pricing theory and decision analysis. *Management Science*, *41*(5), 795–816. https://doi.org/10.1287/mnsc.41.5.795

Thijssen, J. J. J., Huisman, K. J. M. and Kort, P. M. (2004). The effect of information streams on capital budgeting decisions. *European Journal of Operational Research*, *157*(3), 759–774. https://doi.org/10.1016/S0377-2217(03)00257-1

Tilley, J. (1993). Valuing american options in a path simulation model. *Transactions of the Society of Actuaries*, *45*, 83–104.

Ulu, C. and Smith, J. E. (2009). Uncertainty, information acquisition, and technology adoption. *Operations Research*, *57*(3), 740–752. https://doi.org/10.1287/opre.1080.0611

Willigers, B. J. A. and Bratvold, R. B. (2009). Valuing oil and gas options by least-squares monte carlo simulation. *SPE Projects, Facilities Construction*, *4*(4), 146–155. https://doi.org/10.2118/116026-PA

A Derivation of the posterior

Derivation of the mean and variance of the posterior distribution when an average signal is received.

$$\begin{aligned}
\mu_1 &= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\left(\frac{1}{1}\sum_{i=1}^{1}\overline{x}\right) + \frac{1}{\tau_0^2}\mu_0\right] \\
&= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\overline{x} + \frac{1}{\tau_0^2}\mu_0\right] \\
&= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\left(\frac{1}{H}\sum_{i=1}^{H}x_i\right) + \frac{1}{\tau_0^2}\mu_0\right] \\
\tau_1^2 &= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1}
\end{aligned} \tag{A.1}$$

Derivation of the mean and variance of the posterior distribution when an average posterior is calculated from multiple posteriors.

$$\begin{aligned}
\mu_1 &= \frac{1}{H}\sum_{i=1}^{H}\left(\left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\left(\frac{1}{1}\sum_{i=1}^{1}x_i\right) + \frac{1}{\tau_0^2}\mu_0\right]\right) \\
&= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \frac{1}{H}\sum_{i=1}^{H}\left[\frac{1}{\sigma^2}x_i + \frac{1}{\tau_0^2}\mu_0\right] \\
&= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\left(\frac{1}{H}\sum_{i=1}^{H}x_i\right) + \frac{1}{\tau_0^2}\mu_0\frac{1}{H}\sum_{i=1}^{H}\right] \\
&= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1} \left[\frac{1}{\sigma^2}\left(\frac{1}{H}\sum_{i=1}^{H}x_i\right) + \frac{1}{\tau_0^2}\mu_0\right] \\
\tau_1^2 &= \left(\frac{1}{\sigma^2} + \frac{1}{\tau_0^2}\right)^{-1}
\end{aligned} \tag{A.2}$$

NTNU
Norwegian University of
Science and Technology