Parameter Uncertainty in Stochastic Filtering



Andrew Allan

Keble College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

Acknowledgements

I would like to thank my supervisor, Sam Cohen, for his limitless patience, his empathetic and unending encouragement, his infinite yet unassuming wisdom, and for always being immensely generous with his time, advice and guidance.

The two other people who have contributed most to the successful completion of this thesis are my office mates, Andreas and Jono. Having the perfect combination of hilarity, wisdom and sheer monumental brain power at my disposal has been a privilege. Never has an office with any semblance of productivity been so enjoyable.

Abstract

In standard treatments of stochastic filtering one first requires the various parameters of the model. Simply running a filter with estimated parameters without considering the reliability of this estimate does not take into account this additional source of statistical uncertainty. We propose a novel approach to address this problem by making evaluations via a nonlinear expectation. We show how our approach may be reformulated as a pathwise stochastic optimal control problem, where the optimisation is performed for each fixed realisation of the observation process, and proceed to analyse the corresponding value function. We focus in particular on two finite-dimensional continuous-time filters, namely the Kalman–Bucy and Wonham filters. In each case, we present novel comparison results for the viscosity solution of the associated Hamilton–Jacobi–Bellman equation.

To enable this analysis we consider pathwise optimal control problems in generality. In particular we investigate the degeneracy induced by directly controlling the coefficient of the noise term, and we propose a simple procedure to resolve this degeneracy whilst retaining dynamic programming. As we will see, such control problems may be rephrased as the optimal control of rough differential equations, thus demonstrating an original application of rough path theory to statistics.

Contents

1 Introduction **1**

1.1 Stochastic filtering 1

1.2 Filtering under parameter uncertainty 2

1.3 Pathwise stochastic control 5

1.4 Outline of the thesis 8

2 The Kalman–Bucy Filter under Signal Uncertainty **10**

2.1 Estimation via nonlinear expectations 10

2.2 A change of variables 14

2.3 Reformulation as an optimal control problem 17

2.4 An HJB equation 21

2.5 Regularity of the value function 27

2.6 Uniqueness for the HJB equation 31

2.7 A numerical scheme 37

2.8 A numerical example 39

3 Rough Paths **43**

3.1 Motivation 43

3.2 Notation 44

3.3 Rough integration 47

3.4 Rough differential equations 48

4 Hidden Markov Models under Uncertainty **56**

4.1 The Wonham filter 56

4.2 Robust filtering with DR-expectations 58

4.3 Fixing an observation path 61

4.4 An unbounded pathwise control problem 64

4.5 Regularity of the value functions 67

4.6 A classical HJB equation 70

4.7 Comparison results 73
4.8 A rough HJB equation 77

5 Pathwise Optimal Control **80**
5.1 Degeneracy of the naive formulation 80
5.2 Avoiding degeneracy 82
5.3 Recovering dynamic programming 84
5.4 A generalised dynamic control problem 85
5.5 A smooth noise approximation 88
5.6 A rough HJB equation 89
5.7 Examples 91

6 The Kalman–Bucy Filter under General Uncertainty **93**
6.1 The Kalman–Bucy filter revisited 93
6.2 The penalty function 96
6.3 Fixing an observation path 98
6.4 Reformulation as an optimal control problem 99
6.5 A nonlinear backward control problem 102

7 A Linear-Quadratic Robust Filter **110**
7.1 A linear-quadratic pathwise control problem 110
7.2 A numerical example 112

Bibliography **115**

List of Figures

2.1 Realisation of η. 40
2.2 Solution of HJB equation at time 0. 40
2.3 Solution of HJB equation at time 1. 41
2.4 Comparison of standard and robust filter estimates. 41
2.5 Robust upper and lower expectations. 42

7.1 Solution of the robust filtering equations. 113
7.2 Estimates of S_t. 113
7.3 Estimates of S_t^2. 114
7.4 Estimates of $(S_t - 2)^+$. 114

Chapter 1

Introduction

1.1 Stochastic filtering

In many applications, from finance and biology to engineering, defence and aerospace, one is interested in the behaviour of a process evolving in time which cannot be observed directly, and must therefore rely on partial observations in the presence of noise. The problem of estimating the current state of such a hidden process from noisy observations is known as *stochastic filtering*.

Suppose for instance that we are interested in the value of a stochastic process X, referred to as the *signal*, but that we must work in the filtration generated by another process Y, referred to as the *observation process*. We also suppose that the dynamics of Y are dependant on the value of X, such as through a relationship of the form

$$\mathrm{d}Y_t = h(X_t)\,\mathrm{d}t + \mathrm{d}B_t,$$

for some observation function h and measurement noise B. It is clear from this relationship that by observing Y over a given period of time one can infer information about X. In short, the filtering problem is concerned with, at each time t, determining the best estimate (typically in the sense of best mean square) for X_t given $\mathcal{Y}_t = \sigma(Y_s : 0 \leq s \leq t)$, that is, finding the best estimate for the current value of X, given our past observations of Y.

Of course, in general one may be interested in not only the value of $\mathbb{E}[X_t \,|\, \mathcal{Y}_t]$, but in the entire conditional distribution $\pi_t \sim X_t|\mathcal{Y}_t$. The solution of the filtering problem is thus, in general, a measure-valued $\mathcal{Y}_t$-adapted process $\pi = \{\pi_t : t \geq 0\}$, satisfying

$$\int \varphi(x)\,\pi_t(\mathrm{d}x) = \mathbb{E}[\varphi(X_t) \,|\, \mathcal{Y}_t]$$

for every sufficiently regular function φ and every time $t \geq 0$.

The mathematical theory underpinning the filtering of stochastic systems is by now well understood; see Bain and Crisan [5] for a comprehensive exposition. In general the solution π satisfies a nonlinear stochastic PDE, commonly attributed both to Kushner [44, 45] and Stratonovich [59], and to Fujisaki, Kallianpur and Kunita [36]. The classical stochastic filtering problem is therefore essentially solved. We refer to Crisan and Rozovskii [22] for a survey of nonlinear stochastic filtering.

As is well known, the generic stochastic filtering problem is infinite dimensional. That is, cases in which the solution π is finite dimensional are very much the exception rather than the rule, as discussed in detail in for instance Hazewinkel, Marcus and Sussmann [40], Mitter [53] and Ocone and Pardoux [55].

Arguably the most famous finite dimensional filter is the Kalman–Bucy filter (Kalman [41], Kalman and Bucy [42]), which concerns the case where the underlying processes are Gaussian. Early applications of the Kalman–Bucy filter were to areospace engineering (see Cipra [16] or Grewal and Andrews [38]), and it has since found numerous applications in various fields from guidance and navigation to time series analysis and credit risk estimation. Another notable example is the case when the signal X is modelled as a Markov process living in a finite set; the solution to the corresponding filtering problem was first obtained by Wonham [63]. Of course, these modelling choices make strong assumptions about the properties of the processes involved, and the corresponding filters must be carefully calibrated to ensure a satisfactory performance.

1.2 Filtering under parameter uncertainty

The underlying equations in a filtering problem represent the dynamics of the (hidden) signal process and the relationship between the signal and the observation processes. These equations typically involve various parameters so, naturally, in order to run the filter one must first obtain the values of these parameters. In standard treatments of stochastic filtering one often simply runs the filter using an estimate of the parameters. However, this does not take into account the statistical uncertainty introduced by adopting this estimate. Particularly when there is limited available data, resulting in a lack of precision in the estimate, this should cast doubt as to the accuracy of the filter. This additional uncertainty also has implications for decision making based on the outcome of the filter. By evaluating, say, the posterior variance of the signal, one can take into consideration the inherent risk associated with a decision based on

the estimated values of the parameters, but this is not sufficient to account for the uncertainty associated with the error in the estimation of the parameters.

This objection is not new. Robustness in stochastic filtering has been considered in various frameworks for both linear and nonlinear systems. Note that the term 'robustness' has more than one meaning in the context of filtering. We are concerned here with robustness with respect to model uncertainty, rather than, say, continuity of the filter with respect to the observation path. Miller and Pankov [52], Siemenikhin [57], Siemenikhin, Lebedev and Platonov [58] and Verdú and Poor [61] for instance consider known linear dynamics with uncertainty in the noise covariance matrices. However, as pointed out by Martin and Mintz [51], the performance of a filter is typically more dependant on dynamic uncertainty than on uncertainties in either the signal or observation noise covariances.

A by now standard approach to robust estimation problems is to construct a minimax estimator, which generally speaking is designed to find the estimate which minimizes the maximum expected loss over a range of possible models, an idea that goes back at least as far as Wald [62]. Here one typically considers a family $\mathcal{Q}$ of probability measures, and considers an expected loss criterion of the form

$$\xi \longmapsto \operatorname*{ess\,sup}_{\mathbb{Q}\in\mathcal{Q}} \mathbb{E}^{\mathbb{Q}}[\xi \,|\, \mathcal{Y}_t]. \tag{1.1}$$

In addition to the aforementioned authors, this type of estimator has more recently been utilized by Borisov [10, 11], who considers filtering of finite state Markov processes under uncertainty of the transition intensity and observation matrices. As pointed out in [61], minimax estimators can be criticized for being too pessimistic, in the sense that, since they are dependent on the specification of an often arbitrary uncertainty class, they will tend to compensate for potentially very implausible models, and hence do not guarantee a satisfactory performance under the most statistically probable model. Moreover, the loss functions considered are invariably assumed to be quadratic which, although generally reasonable, is somewhat restrictive.

A more contemporary approach is the use of H_∞ filters, which attempt to achieve a desired degree of noise attenuation by ensuring that the maximum energy gain from the noise source to the estimation error is bounded by a prescribed level. This type of filter is suitable in situations when an accurate system model or knowledge of the noise statistics is not available, although as noted in George [37], since such estimators are designed to minimize the worst case energy gain, they generally sacrifice the average filter performance. For treatments of robust H_∞ filters for linear systems see Lewis,

Xie and Popa [46], Xie, de Souza and Fu [64], Yang and Ye [65], Zhang, Xia and Shi [67] and the references therein.

A related strategy is to combine such H_∞ optimisation with more classical least squares techniques, to form what are known as mixed H_2/H_∞ filters. Variations on this type of filter are explored in Bernstein and Haddad [9] and Khargonekar, Rotea and Baeyens [43], where in each case the goal is to minimize an L^2-type state estimation error criterion, subject to a prespecified H_∞ constraint on this error. See also Aliyu and Boukas [2] for an application of this method to nonlinear filtering, or Chen and Zhou [15] for a game approach to mixed H_2/H_∞ filter design.

Another recent technique is that of Neveux, Blanco and Thomas [54] and Zhou [68], which penalises the sensitivity of estimation errors to parameter variations, by minimizing both the estimation error as well as its gradient with respect to the uncertain parameters.

The objective of this thesis is to propose a new, substantially different approach to parameter uncertainty in stochastic filtering. Based on ideas in Cohen [17, 18], we propose to make evaluations using a convex expectation, represented in terms of a penalty function κ_t:

$$\mathcal{E}(\xi \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\mathbb{Q}\in\mathcal{Q}} \big\{ \mathbb{E}^{\mathbb{Q}}[\xi \,|\, \mathcal{Y}_t] - \kappa_t(\mathbb{Q} \,|\, \mathcal{Y}_t) \big\}.$$

In the context of model uncertainty, i.e. uncertainty of the underlying probability measure, nonlinear expectations provide an evaluation of random variables which takes into account every admissible measure. In other words, they consider every plausible view of the world, and envisage the worst case scenario. However, in contrast with sublinear expectations (cf. (1.1)), the inclusion of the penalty term κ_t means that we can penalise different measures according to how unreasonable we consider them to be, thus restricting our attention to only those measures which we consider to be realistic. Convex expectations are in this sense less pessimistic than their sublinear counterparts.

This penalty κ_t (at least after a slight reformulation) represents how our uncertainty evolves in time, and is calculated by propagating our a priori uncertainty forward through time using the filter dynamics. The main theoretical result of this body of work is to establish this penalty function as the unique solution of a Hamilton–Jacobi–Bellman (HJB) type partial differential equation (PDE). This means in particular that it can be calculated recursively through time alongside the standard filter.

The approach taken here allows us considerable flexibility in terms of calibration and filter construction. For instance, the initial calibration of the penalty function

may be based on classical statistics, or may be chosen according to the user's preference. Moreover, by prescribing the initial penalty to be very large outside a particular range of parameter values, we essentially consider perturbations of the signal dynamics which are bounded by a prescribed level, recovering a setting comparable to that of H_∞ filtering.

Our penalty function can be used to perform robust estimation by, for instance, constructing a minimax estimator with an arbitrary choice of loss function. Much of the aforementioned literature concerns estimators only of the state of the signal. On the approach taken here however, the construction of our penalty function allows for the estimation of *any* random variable which depends on the current state of the signal, including nonlinear functionals thereof. In particular, our penalty only has to be computed once, and does not have to be recalculated for each choice of variable to be estimated. As we will see, our nonlinear expectation can also be used to define an interval estimate for any such random variable, thus providing uncertainty-robust error bounds, in analogy with classical confidence intervals.

1.3 Pathwise stochastic control

Stochastic optimal control is a classical optimisation problem with numerous applications, from optimal liquidation and portfolio selection in mathematical finance to various problems in production planning, engineering and biology. Here one typically has a stochastic differential equation (SDE) of the form

$$\mathrm{d}X_s = b(X_s, \gamma_s)\,\mathrm{d}s + \sigma(X_s, \gamma_s)\,\mathrm{d}W_s, \qquad s \in [t, T], \tag{1.2}$$

with an initial condition $X_t = x$, where W is a Brownian motion and γ is an adapted control process, and the goal is to minimize (resp. maximize) a cost (resp. reward) functional of the form

$$J(t, x; \gamma) = \mathbb{E}\bigg[\int_t^T f(X_s, \gamma_s)\,\mathrm{d}s + g(X_T)\bigg] \tag{1.3}$$

over all possible choices of the control γ. The resolution of this problem is by now well understood—two primary approaches being that of the Pontryagin stochastic maximum principle and Bellman's principle of optimality (or dynamic programming principle), which allows one to characterise the value function of the control problem, defined by $v(t, x) = \inf_\gamma J(t, x; \gamma)$, as the unique solution of an HJB equation.

In their 1998 paper [48], Lions and Souganidis considered a variant of this problem, known as 'pathwise stochastic control', where the optimisation is performed pathwise.

In other words, one considers controlling the solution of an equation of the form (1.2) for each individual realisation of the Brownian motion W. Moreover, they suggest that in this case the value function should satisfy a 'stochastic HJB equation'. Indeed, at least in the case when σ does not depend on the control γ, if we pretend for the moment that the paths of W were smooth, then, at least formally, the classical theory leads one to derive a stochastic PDE of the form

$$-\mathrm{d}v - \inf_{\gamma} \left\{ b \cdot \nabla v + f \right\} \mathrm{d}t - \sigma \cdot \nabla v \, \mathrm{d}W = 0, \tag{1.4}$$

with $v(T, \cdot) = g$.

As is a common theme among robust filtering techniques, our approach will lead us to derive a particular optimisation procedure, in our case a stochastic optimal control problem. However, in our filtering context we wish to perform our optimisation based on our past observations which, having already been observed, are no longer stochastic. Thus, it will not be meaningful to consider a cost/reward which is averaged over all possible realisations of past observations, as in (1.3). Our problem will therefore turn out to be precisely one of pathwise stochastic control.

The notion of pathwise stochastic control actually goes back at least as far as the work of Davis and Burstein [23, 24], who note that pathwise control is equivalent to the classical stochastic control setting if one allows for anticipative controls, leading to the conclusion that the difference between classical and pathwise control boils down to nonanticipativity of the controls. In this view, pathwise control can be thought of as performing optimal control with the benefit of complete knowledge of both the past and future realisations of the stochastic noise. On the other hand, as shown, for instance, by Rogers [56], pathwise control can also be used to obtain duality results for classical (nonanticipative) stochastic control, thus providing an alternative approach for numerical computations.

Since pathwise control entails the optimisation of a stochastic system path by path, it is natural to fix such an (arbitrary) path and proceed to analyse the resulting deterministic problem. This invites a pathwise interpretation of the stochastic integral appearing in the controlled dynamics. The strategy followed by Buckdahn and Ma [12] circumnavigates such a technical requirement, by instead employing a Doss–Sussmann-style transformation to convert the problem into a more standard setting of 'wider-sense control problem', allowing them to establish their value function as the unique stochastic viscosity solution of the associated HJB equation. We will use a similar approach to this in Chapter 2. A more direct approach, avoiding such an

ad hoc change of variables, requires one to utilize a pathwise approach to stochastic integration.

One such deterministic approach to integration against paths of low regularity is provided by rough path theory, introduced by Lyons [50]. The basic idea here is that the notion of integration can be extended in a consistent way to paths of lower regularity such that strong stability results concerning continuity of the integration map with respect to the driving 'rough path' hold, but one requires extra information about the driving signal than is expressed in the path alone. Such a path Y must therefore be 'enhanced' by a suitable 'second order' process $\mathbb{Y}$ which captures this missing information. The addition of the process $\mathbb{Y}$ is equivalent to considering the Lévy area of the path Y, and corresponds to the addition of the iterated integral $\int_0^{\cdot} \int_0^r \mathrm{d}Y_s \otimes \mathrm{d}Y_r$, but since this integral does not exist in the classical sense, its value must be postulated, rather than being uniquely determined by the original path Y.

More recently, Diehl, Friz and Gassiat [26] gave what appears to be the first attempt to apply rough path theory to optimal control. The authors consider controlled dynamics of the form

$$\mathrm{d}X_s = b(X_s, \gamma_s)\,\mathrm{d}s + \sigma(X_s)\,\mathrm{d}\mathbf{Y}_s, \tag{1.5}$$

driven by a geometric rough path $\mathbf{Y}$. They proceed to obtain both a version of Pontryagin's maximum principle and establish their value function as the unique solution of a 'rough HJB equation', and moreover obtain a duality result for the corresponding nonanticipative stochastic control problem. Their results suggest that rough path theory is an ideal tool for the study of pathwise stochastic control—a notion that we will echo in this thesis.

Notably however, the existing literature on pathwise control invariably focuses only on the case where the control process appears in the drift term, but does not appear in the coefficient of the noise term (or rough path). That is, the controlled dynamics considered are typically of the general form in (1.5), where σ is not allowed to depend on the control γ. As observed in Diehl et al. [26], the pathwise control problem with full dynamic control (particularly control in the coefficient of the noise term), when stated in the obvious way, turns out to be degenerate, which explains the lack of results in this direction. An indication of this arises when one reruns the formal derivation that led to (1.4) in the case where σ depends on γ. In this case the resulting equation exhibits the Brownian motion W *inside* the infimum. At least heuristically, this corresponds to the ability to perfectly optimize over the path of the noise term, which, as we will see, is the source of the degeneracy.

To consider parameter uncertainty in a general filtering setting, particularly with uncertainty of the observation function, we must have control in the coefficient of the noise term. We shall therefore dedicate Chapter 5 to exploring how the degeneracy phenomenon described above can be resolved.

1.4 Outline of the thesis

In Chapter 2 we will introduce our approach to robust stochastic filtering in the context of the one-dimensional Kalman–Bucy filter with uncertainty of the dynamics of the signal process. We will see how our approach can be reformulated as an optimal control problem, and proceed to analyse this control problem and the associated (highly nonlinear) HJB equation. This chapter is based on the paper [3].

In particular, in Section 2.3 we will introduce a change of variables which in this setting allows us to deal with the pathwise nature of our formulation without appealing to rough path theory or any other such intricate machinery. However, in order to deal with more general cases it will be necessary to invoke a (deterministic) notion of calculus which is capable of handling paths of infinite total variation. Chapter 3 is therefore devoted to introducing some tools from rough path theory. In particular, in Theorem 3.4.3 we shall establish solutions to differential equations driven by p-rough paths (for $p \in [2, 3)$) with control parameters in both the drift and in the coefficient of the driving rough path.

In Chapter 4 we shall consider continuous-time hidden Markov models with uncertain transition matrices. In this setting our pathwise stochastic control problem will be formulated as the optimal control of a rough differential equation, leading to a characterisation of the associated value function as the unique solution of a rough HJB equation.

In order to deal with more general uncertainty, in particular uncertainty of the observation function, we will need to understand how to deal with the degeneracy phenomenon exhibited by general pathwise control problems described in Section 1.3 above. We will therefore dedicate Chapter 5 to the study of pathwise optimal control in a general setting, particularly with control in the coefficient of the driving noise.

We will return to considering uncertainty in the Kalman–Bucy filter in Chapter 6, but this time in a multidimensional setting with the parameters of both the signal and observation processes being considered to be uncertain. Chapters 5 and 6 are based on the paper [4].

Finally, in Chapter 7 we will consider a particular special case in which our pathwise control problem is of linear-quadratic type. We will derive the associated robust filtering equations, and conclude with a numerical example.

Chapter 2

The Kalman–Bucy Filter under Signal Uncertainty

In this chapter we shall introduce our approach in the setting of Kalman-Bucy filtering. After a brief recall of the Kalman–Bucy filter itself, we shall introduce the notion of nonlinear expectations, and see how they may be used to construct robust estimates in the context of model uncertainty. We will then see how the associated penalty may be formulated as the value function of a (deterministic) optimal control problem. This control problem turns out to be somewhat degenerate due to the highly nonlinear nature of the Hamiltonian, and the correspondingly unruly behaviour of the state trajectories. Nevertheless, we will proceed to provide a thorough and rigorous treatment of this control problem and its relationship with the associated HJB equation, via techniques which are generalisable to similarly nonlinear control problems.

We will establish important properties of the value function, as well as its counterpart obtained by restricting to uniformly bounded parameters; our main result in this direction is uniqueness of solutions to the HJB equation. We will briefly outline a numerical scheme to solve the HJB equation, and finally exhibit a numerical example to illustrate our results.

2.1 Estimation via nonlinear expectations

Let us take an underlying filtered space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0})$, and suppose that a signal process X and observation process Y satisfy the following pair of linear equations

$$\mathrm{d}X_t = \alpha_t X_t \,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B^1_t, \tag{2.1}$$

$$\mathrm{d}Y_t = c_t X_t \,\mathrm{d}t + \mathrm{d}B^2_t, \tag{2.2}$$

with the initial conditions $Y_0 = 0$ and $X_0 \sim N(\mu_0, \sigma_0^2)$, for some $\mu_0 \in \mathbb{R}$, $\sigma_0 > 0$. Here, B^1 and B^2 are standard $\mathcal{F}_t$-adapted Brownian motions which are assumed to be uncorrelated, and the parameters $\alpha\colon [0,T] \to \mathbb{R}$, $\beta\colon [0,T] \to [0,\infty)$ and $c\colon [0,T] \to \mathbb{R}$ are bounded measurable (deterministic) processes. In this chapter we shall suppose that we are uncertain of the parameters of the signal process, namely α and β. Uncertainty of c will be considered in Chapter 6.

For simplicity of presentation we shall focus on the scalar case, where X and Y are one-dimensional, however all of our results extend analogously to the multivariate case; see Chapter 6.

The standard object to study in stochastic filtering is the conditional distribution process $\{\pi_t : t \geq 0\}$, given by $\pi_t\varphi = \mathbb{E}[\varphi(X_t) \,|\, \mathcal{Y}_t]$ for any bounded measurable function φ, where $\mathcal{Y}_t$ is the (completed) σ-algebra generated by the observation process up to time t. We recall (from e.g. Bain and Crisan [5] or Liptser and Shiryaev [49]) that in this setting the conditional distribution π_t of X_t given $\mathcal{Y}_t$ is a normal distribution, and hence completely characterised by its mean $q_t = \mathbb{E}[X_t \,|\, \mathcal{Y}_t]$ and variance $R_t = \mathbb{E}\big[(X_t - q_t)^2 \,\big|\, \mathcal{Y}_t\big]$ which satisfy the equations

$$\mathrm{d}q_t = \alpha_t q_t \,\mathrm{d}t + c_t R_t \,\mathrm{d}V_t, \tag{2.3}$$

$$\mathrm{d}R_t = \big(\beta_t + 2\alpha_t R_t - c_t^2 R_t^2\big)\mathrm{d}t, \tag{2.4}$$

$$\mathrm{d}V_t = \mathrm{d}Y_t - c_t q_t \,\mathrm{d}t, \tag{2.5}$$

where V is known as the innovation process.

Naturally this distribution depends on the parameters α, β, which in our setting are unknown. Typically one finds an estimate for the parameters, uses this estimate to compute the solution of (2.3) and (2.4), and can then calculate $\pi_t\varphi = \mathbb{E}[\varphi(X_t) \,|\, \mathcal{Y}_t]$ for a given choice of φ.

Making evaluations via (standard) expectations immediately becomes problematic in the uncertain parameter setting, as uncertainty in the parameters corresponds naturally to uncertainty in the underlying probability measure. An advantage of nonlinear expectations is that they can be defined without reference to a specific measure.

Nonlinear expectations are closely linked to the theory of risk measures, which are tools used in finance to attempt to quantify the riskiness of a financial position, and are a natural object to study when dealing with model uncertainty. One typically defines a *dynamic convex expectation* as a map $\mathcal{E}(\,\cdot \,|\, \mathcal{Y}_t)\colon L^\infty(\mathcal{F}) \to L^\infty(\mathcal{Y}_t)$ which satisfies the properties of monotonicity, translation equivariance, normalization and convexity. The convexity of the expectation corresponds to the notion that, in the

context of decision making in the presence of uncertainty, diversification of a portfolio should not increase the associated risk.

A particularly nice kind of convex expectation are those which also satisfy the Fatou property. As proved separately by Föllmer and Schied [30] and by Frittelli and Rosazza Gianin [31] in the context of monetary risk measures, any dynamic convex expectation which satisfies the Fatou property admits a representation of the form

$$\mathcal{E}(\xi \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\mathbb{Q} \in \mathcal{Q}_t} \big\{ \mathbb{E}^{\mathbb{Q}}[\xi \,|\, \mathcal{Y}_t] - \kappa_t(\mathbb{Q}) \big\}, \tag{2.6}$$

where $\mathcal{Q}_t$ is a suitably chosen space of equivalent probability measures, and κ_t is a $\mathcal{Y}_t$-measurable 'penalty' function. For a more thorough discussion of the theory of nonlinear expectations and risk measures, see e.g. Cohen [17], Delbaen, Peng and Rosazza Gianin [25] or Föllmer and Schied [29, 30].

We shall now see how such nonlinear expectations can be used to estimate functions of the hidden process X given our observations of Y, taking parameter uncertainty into account. The first thing to notice is that the probability measure which governs our system depends on the choices of α, β, μ_0 and σ_0. That is, any given choice of these parameters corresponds to a probability measure $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$ which determines the behaviour of the processes X and Y. We associate to the set of all possible choices of parameters the corresponding class of all admissible probability measures, and suppose that the 'true' measure belongs to this class.

Motivated by the representation in (2.6), for a given uncertainty aversion parameter $k_1 > 0$ and exponent $k_2 \geq 1$, we consider, for $\xi \in L^\infty(\mathcal{F})$, the nonlinear expectation given by

$$\mathcal{E}(\xi \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \left\{ \mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\xi \,|\, \mathcal{Y}_t] - \left(\frac{1}{k_1} \left(\int_0^t f(s, \alpha_s, \beta_s)\, \mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \right) \right)^{k_2} \right\}, \tag{2.7}$$

where we write $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}$ for the expectation corresponding to the measure $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$. Here the essential supremum is taken over all possible parameters $\alpha, \beta, \mu_0, \sigma_0$ for the dynamics and initial distribution of the signal.

Remark 2.1.1. Due to the nature of the essential supremum, in (2.7) it is sufficient to consider only deterministic parameters α, β. We note however that a pair of parameters α, β which attains this essential supremum will naturally depend on the observation Y, and thus be $\mathcal{Y}_t$-predictable, rather than deterministic.

The first term inside the essential supremum is a version of the conditional expectation of ξ, corresponding to the measure associated with a particular choice of the

parameters. We consider the value of this conditional expectation for all choices of the parameters, and then discount this value by a penalty associated with the corresponding parameters. This penalty is composed of the nonnegative functions f and κ_0, where f may either be deterministic or $\mathcal{Y}_t$-predictable. These penalty functions represent our opinion of how plausible different values of the parameters are. That is, we penalise different choices of parameters according to how unreasonable we consider them to be. Here f and κ_0 may be chosen arbitrarily, though in practice may be based on a priori statistical analysis.

The inclusion of the parameters k_1, k_2 allows for more flexibility in terms of filter construction. Here $k_1 > 0$ measures our aversion to uncertainty, and $k_2 \geq 1$ allows one to change the shape of the penalty. In particular, if the penalty functions f, κ_0 take the minimal value zero for a given choice of reference parameters $(\alpha^*, \beta^*, \mu_0^*, \sigma_0^*)$, then, at least formally, we recover $\mathcal{E}(\xi \,|\, \mathcal{Y}_t) \to \mathbb{E}^{\alpha^*,\beta^*,\mu_0^*,\sigma_0^*}[\xi \,|\, \mathcal{Y}_t]$ as $k_1 \to 0$. See [17] or [18] for a more thorough discussion of the role of these parameters. We note furthermore that, as we will see, the inclusion of these parameters adds no additional complexity to our analysis.

If the functions f and κ_0 achieve the value zero, that is there exists a choice of reference parameters which have zero associated penalty, then it follows that $\mathcal{E}(\cdot \,|\, \mathcal{Y}_t)\colon L^\infty(\mathcal{F}) \to L^\infty(\mathcal{Y}_t)$ defines a dynamic convex expectation which satisfies the Fatou property.

Moreover, assuming that the functions f and κ_0 achieve the value zero for some choice of parameters, $(\alpha^*, \beta^*, \mu_0^*, \sigma_0^*)$ say, then we infer from (2.7) that

$$\mathcal{E}(\xi \,|\, \mathcal{Y}_t) \geq \mathbb{E}^{\alpha^*,\beta^*,\mu_0^*,\sigma_0^*}[\xi \,|\, \mathcal{Y}_t].$$

Thus, provided that these parameters are calibrated reasonably well, the nonlinear expectation $\mathcal{E}(\xi \,|\, \mathcal{Y}_t)$ will typically overestimate the 'true' value of $\mathbb{E}[\xi \,|\, \mathcal{Y}_t]$. We may therefore think of $\mathcal{E}(\xi \,|\, \mathcal{Y}_t)$ as an 'upper' expectation, the corresponding 'lower' expectation being given by $-\mathcal{E}(-\xi \,|\, \mathcal{Y}_t)$. The interval estimate $\big[-\mathcal{E}(-\xi \,|\, \mathcal{Y}_t),\, \mathcal{E}(\xi \,|\, \mathcal{Y}_t)\big]$ may then be considered as a range of plausible values that the true value of $\mathbb{E}[\xi \,|\, \mathcal{Y}_t]$ could take. This is analogous to the notion of confidence/credible intervals from classical statistics.

One can also use the nonlinear expectation to obtain a robust point estimate of ξ, by finding the $\mathcal{Y}_t$-measurable random variable $\hat{\xi}$ which minimizes $\mathcal{E}(\psi(\xi - \hat{\xi}) \,|\, \mathcal{Y}_t)$ for a given loss function ψ. This defines a type of minimax estimator. However, note that the inclusion of the penalty function ensures that parameter choices which are considered to be very unreasonable will never be considered, thus avoiding the

previously mentioned criticism of other minimax estimators of being too pessimistic. Moreover, whereas many such estimators assume that the loss function is quadratic, the function ψ in our formulation may be chosen arbitrarily.

Our focus is the application of the approach described above to the filtering context described above. Correspondingly, we shall henceforth restrict our attention to evaluations of random variables of the form $\xi = \varphi(X_t)$, where X is the (unobserved) signal process and φ is an arbitrary bounded measurable function.

2.2 A change of variables

Before proceeding, we shall define more precisely what is meant by the expression in (2.7). We adopt the following setup.

Let X and Y be $\mathcal{F}_t$-adapted processes. For each parameter choice $(\alpha, \beta, \mu_0, \sigma_0)$, we let $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$ be a probability measure such that the law of (X, Y) is equal to the law of $(\tilde{X}^{\alpha,\beta,\mu_0,\sigma_0}, \tilde{Y}^{\alpha,\beta,\mu_0,\sigma_0})$, where $\tilde{X}^{\alpha,\beta,\mu_0,\sigma_0}$, $\tilde{Y}^{\alpha,\beta,\mu_0,\sigma_0}$ satisfy the SDEs (2.1)–(2.2) driven by *some* Brownian motions B^1, B^2 on some probability space, with $\tilde{X}_0^{\alpha,\beta,\mu_0,\sigma_0} \sim N(\mu_0, \sigma_0^2)$ and $\tilde{Y}_0^{\alpha,\beta,\mu_0,\sigma_0} = 0$.

In particular, the processes X, Y as functions from $\Omega \times [0,T] \to \mathbb{R}$, and hence the (uncompleted) filtration $\sigma(Y_s : s \in [0,t])$, $t \in [0,T]$, are defined independently of the choice of parameters.

As described above, we would like to consider a nonlinear expectation of the form

$$\mathcal{E}(\varphi(X_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,\mu_0,\sigma_0} \left\{ \mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t) \,|\, \mathcal{Y}_t] - \left(\frac{1}{k_1} \left(\int_0^t f(s, \alpha_s, \beta_s)\, \mathrm{d}s + \kappa_0(\mu_0, \sigma_0^2) \right) \right)^{k_2} \right\} \tag{2.8}$$

for all bounded measurable functions φ. Note that one should take care here, as in general the measures $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$ are not equivalent on $\mathcal{F}_t$ (since $\langle X \rangle_t = \int_0^t \beta_s \,\mathrm{d}s$ is measurable), so the completed filtration $\mathcal{Y}_t = \sigma(Y_s : s \in [0,t]) \vee \mathcal{N}$ will in general depend on the choice of parameters (where $\mathcal{N}$ denotes the $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$-null sets in $\mathcal{F}$).

Recall the filtering equations (2.3)–(2.5) satisfied by the conditional mean q and variance R. The natural parameters, in the sense of exponential families, for a $N(\mu, \sigma^2)$ distribution are $\big(\mu/\sigma^2, -1/2\sigma^2\big)$. This motivates us to calculate

$$\begin{aligned} \mathrm{d}\left(\frac{q}{R}\right)_t &= -\frac{q_t}{R_t}\left(\alpha_t + \frac{\beta_t}{R_t}\right)\mathrm{d}t + c_t\,\mathrm{d}Y_t, \\ \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{1}{R_t}\right) &= -\frac{\beta_t}{R_t^2} - \frac{2\alpha_t}{R_t} + c_t^2. \end{aligned} \tag{2.9}$$

Notice that the stochastic integral in (2.9) no longer involves the variable R. Since we assume that the parameter c is known, we can (and do) fix a version of the integral

$$\eta_t := \int_0^t c_s \, \mathrm{d}Y_s$$

for $t \in [0, T]$, which, crucially, may be defined independently of the choice of parameters $(\alpha, \beta, \mu_0, \sigma_0)$ and of the functional φ.

Recall (from e.g. [5]) that the innovation process V is a $\mathcal{Y}_t$-adapted Brownian motion under the corresponding probability measure. It follows from the Burkholder–Davis–Gundy inequality and the Kolmogorov continuity criterion (see e.g. [19]), that η is almost surely Hölder continuous with any exponent strictly less than $1/2$.

We may now rewrite (2.9) as the ODE

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{q_t}{R_t} - \eta_t\right) = -\frac{q_t}{R_t}\left(\alpha_t + \frac{\beta_t}{R_t}\right).$$

It is now possible to simplify the filtering equations with a change of variables. We transform the state trajectories from (q, R) to w by the relation

$$w(t) = (w_1(t), w_2(t)) = \left(\frac{q_t}{R_t} - \eta_t, \ \frac{1}{R_t}\right). \tag{2.10}$$

By the above, w satisfies the following ODEs

$$\begin{aligned} \frac{\mathrm{d}w_1}{\mathrm{d}t}(t) &= -\big(w_1(t) + \eta_t\big)\big(\alpha_t + \beta_t w_2(t)\big), \\ \frac{\mathrm{d}w_2}{\mathrm{d}t}(t) &= -\beta_t w_2(t)^2 - 2\alpha_t w_2(t) + c_t^2, \end{aligned} \tag{2.11}$$

which may be equivalently expressed as

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = F(s, w(s), \alpha_s, \beta_s) \qquad \text{for} \quad s \in [0, t], \tag{2.12}$$

where the function F is given by

$$F(t, x, a, b) = \begin{pmatrix} -(x_1 + \eta_t)(a + bx_2) \\ -bx_2^2 - 2ax_2 + c_t^2 \end{pmatrix} \qquad \text{for} \quad x = (x_1, x_2) \in \mathbb{R} \times (0, \infty). \tag{2.13}$$

We also make the corresponding change of variables in the terminal values of q and R. That is, for a fixed time t, we make the transformation from (μ, σ^2) to x via

$$x = (x_1, x_2) = \left(\frac{\mu}{\sigma^2} - \eta_t, \ \frac{1}{\sigma^2}\right), \tag{2.14}$$

noting that the map $(\mu, \sigma) \mapsto x$ is a bijection from $\mathbb{R} \times (0, \infty)$ to itself. We also define a new initial penalty function $g \colon \mathbb{R} \times (0, \infty) \to \mathbb{R}$, defined such that

$$g(x) = \kappa_0(\mu_0, \sigma_0^2), \qquad \text{whenever} \qquad x = \left(\frac{\mu_0}{\sigma_0^2}, \frac{1}{\sigma_0^2}\right) \in \mathbb{R} \times (0, \infty).$$

Remark 2.2.1. The change of variables we have chosen is, to our knowledge, the simplest one in this setting which successfully reduces our SDE into an ODE with random coefficients, but it is by no means unique. In general this type of result can be achieved via Doss–Sussmann transformations. In particular, the transformation exhibited in Buckdahn and Ma [12] is also sufficient for our purposes, and results in an ODE with comparable complexity.

In the following we shall denote by $\Phi(\cdot\,;\mu,\sigma^2)$ the distribution function of the $N(\mu,\sigma^2)$ distribution.

Recalling the change of variables in (2.10), it follows that

$$\int_{\mathbb{R}} \varphi(s)\,\Phi\bigg(\mathrm{d}s;\frac{w_1(t)+\eta_t}{w_2(t)},\frac{1}{w_2(t)}\bigg)$$

is a version of the conditional expectation $\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(X_t)\,|\,\mathcal{Y}_t]$ for every choice of parameters and every functional φ.

Choosing this version, we rewrite the nonlinear expectation in (2.8) as

$$\begin{aligned}\mathcal{E}(\varphi(X_t)\,|\,\mathcal{Y}_t) = \operatorname*{ess\,sup}_{\alpha,\beta,w(0)} \bigg\{ &\int_{\mathbb{R}} \varphi(x)\,\Phi\bigg(\mathrm{d}x;\frac{w_1(t)+\eta_t}{w_2(t)},\frac{1}{w_2(t)}\bigg) \\ &- \bigg(\frac{1}{k_1}\bigg(\int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g(w(0))\bigg)\bigg)^{k_2}\bigg\}, \end{aligned} \tag{2.15}$$

where now the essential supremum is taken over the family of parameters α, β, and over all possible initial values $w(0)$ of w.

The reason for the change of variables performed above was to arrive at the expression (2.15), represented in terms of the variable w, the dynamics of which (2.12) are simply a random ODE, rather than a truly stochastic equation. In fact, all the stochastic elements of this expression are encoded in the process η.

In particular, by fixing a realisation of η, it is natural to consider a 'pathwise version' of the nonlinear expectation, given by

$$\begin{aligned}\mathcal{E}(\varphi(X_t)\,|\,\eta_{[0,t]}) := \sup_{\alpha,\beta,w(0)} \bigg\{ &\int_{\mathbb{R}} \varphi(x)\,\Phi\bigg(\mathrm{d}x;\frac{w_1(t)+\eta_t}{w_2(t)},\frac{1}{w_2(t)}\bigg) \\ &- \bigg(\frac{1}{k_1}\bigg(\int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g(w(0))\bigg)\bigg)^{k_2}\bigg\}, \end{aligned} \tag{2.16}$$

where here we have fixed a realisation $\eta_{[0,t]} = \eta|_{[0,t]}(\omega)$ for some $\omega\in\Omega$, and w satisfies the ODE (2.12) with this realisation of η.

Remark 2.2.2. In the above we fixed the two processes X and Y and considered their law under a family of different probability measures. An alternative setup would be

to instead fix B^1 and $Y = B^2$ as independent standard Brownian motions under a fixed measure $\tilde{\mathbb{P}}$, and let

$$\mathcal{Y}_t := \sigma(Y_s : s \in [0,t]) \vee \mathcal{N}$$

be the observation filtration, where $\mathcal{N}$ denotes the $\tilde{\mathbb{P}}$-null sets (and subsets thereof).

By the Kallianpur–Striebel formula (see [5, Chapter 3]), for a given parameter choice $(\alpha, \beta, \mu_0, \sigma_0)$, the required filter is given by

$$\text{“}\mathbb{E}^{\alpha,\beta,\mu_0,\sigma_0}[\varphi(S_t^{\alpha,\beta,\mu_0,\sigma_0}) \,|\, \mathcal{Y}_t]\text{”} = \frac{\tilde{\mathbb{E}}[\tilde{Z}_t^{\alpha,\beta,\mu_0,\sigma_0}\varphi(S_t^{\alpha,\beta,\mu_0,\sigma_0}) \,|\, \mathcal{Y}_t]}{\tilde{\mathbb{E}}[\tilde{Z}_t^{\alpha,\beta,\mu_0,\sigma_0} \,|\, \mathcal{Y}_t]},$$

where

$$\tilde{Z}_t^{\alpha,\beta,\mu_0,\sigma_0} := \exp\left(\int_0^t c_s S_s^{\alpha,\beta,\mu_0,\sigma_0}\,\mathrm{d}Y_s - \frac{1}{2}\int_0^t (c_s S_s^{\alpha,\beta,\mu_0,\sigma_0})^2\,\mathrm{d}s\right),$$

and $S^{\alpha,\beta,\mu_0,\sigma_0}$ satisfies

$$\mathrm{d}S_t^{\alpha,\beta,\mu_0,\sigma_0} = \alpha_t S_t^{\alpha,\beta,\mu_0,\sigma_0}\,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B_t^1$$

with $S_0^{\alpha,\beta,\mu_0,\sigma_0} \sim N(\mu_0, \sigma_0^2)$.

The advantage of this approach is that one only has to write down a single filtration $\mathcal{Y}_t$ and work under a single probability measure $\tilde{\mathbb{P}}$. On the other hand, as well as being somewhat more cumbersome, in later chapters we will wish to consider stochastic integrals and differential equations under our family of probability measures, and it will be more convenient to be able to work directly with the measure $\mathbb{P}^{\alpha,\beta,\mu_0,\sigma_0}$ for each choice of parameters.

2.3 Reformulation as an optimal control problem

Our next aim is to show how the variable $\mathcal{E}(\varphi(X_t) \,|\, \eta_{[0,t]})$, as given in (2.16), is related to the following ‘value function’. Here and throughout, we adopt the convention that $\inf \emptyset = \infty$. We define

$$v(t,x) = \inf\left\{ \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g(w(0)) \;\middle|\; \begin{array}{c} \alpha, \beta, w(0) \\ \text{such that} \quad w(t) = x \end{array} \right\}, \tag{2.17}$$

where the infimum is taken over all bounded measurable functions α, β and all initial values $w(0)$ such that the corresponding terminal value $w(t)$ is equal to $x \in \mathbb{R}\times(0,\infty)$. As the following lemma suggests, the function v may be viewed as the convex dual of our nonlinear expectation.

Lemma 2.3.1. *For any realisation $\eta_{[0,t]} = \eta|_{[0,t]}(\omega)$ and any functional φ, we have*

$$\mathcal{E}(\varphi(X_t) \,|\, \eta_{[0,t]}) = \sup_{x \in \mathbb{R} \times (0,\infty)} \left\{ \int_{\mathbb{R}} \varphi(s)\, \Phi\left(\mathrm{d}s; \frac{x_1 + \eta_t}{x_2}, \frac{1}{x_2} \right) - \left(\frac{1}{k_1} v(t,x) \right)^{k_2} \right\}. \quad (2.18)$$

Proof. Simply note that, for any $x \in \mathbb{R} \times (0,\infty)$,

$$\begin{aligned} &\sup \left\{ \int_{\mathbb{R}} \varphi(s)\, \Phi\left(\mathrm{d}s; \frac{w_1(t) + \eta_t}{w_2(t)}, \frac{1}{w_2(t)} \right) \right. \\ &\qquad \left. - \left(\frac{1}{k_1} \left(\int_0^t f(s, \alpha_s, \beta_s)\, \mathrm{d}s + g(w(0)) \right) \right)^{k_2} \;\middle|\; \begin{array}{c} \alpha, \beta, w(0) \\ \text{such that} \quad w(t) = x \end{array} \right\} \\ &= \int_{\mathbb{R}} \varphi(s)\, \Phi\left(\mathrm{d}s; \frac{x_1 + \eta_t}{x_2}, \frac{1}{x_2} \right) - \left(\frac{1}{k_1} v(t,x) \right)^{k_2}, \end{aligned}$$

and then take the supremum over $x \in \mathbb{R} \times (0,\infty)$. □

It also follows from the previous lemma that

$$\begin{aligned} &- \mathcal{E}(-\varphi(X_t) \,|\, \eta_{[0,t]}) \\ &= \inf_{\alpha, \beta, w(0)} \left\{ \int_{\mathbb{R}} \varphi(s)\, \Phi\left(\mathrm{d}s; \frac{w_1(t) + \eta_t}{w_2(t)}, \frac{1}{w_2(t)} \right) \right. \\ &\qquad \left. + \left(\frac{1}{k_1} \left(\int_0^t f(s, \alpha_s, \beta_s)\, \mathrm{d}s + g(w(0)) \right) \right)^{k_2} \right\}, \\ &= \inf_{x \in \mathbb{R} \times (0,\infty)} \left\{ \int_{\mathbb{R}} \varphi(s)\, \Phi\left(\mathrm{d}s; \frac{x_1 + \eta_t}{x_2}, \frac{1}{x_2} \right) + \left(\frac{1}{k_1} v(t,x) \right)^{k_2} \right\}. \end{aligned} \quad (2.19)$$

If one could obtain an explicit expression, or approximate expression, for the function $v(t,x)$, then it would be a relatively straightforward task to compute the supremum in (2.18) or the infimum in (2.19) for any sufficiently regular function φ. We have therefore reduced our problem to understanding the behaviour of the function v.

Our fundamental observation is that the expression for v given in (2.17) has the familiar form of the value function of an optimal control problem. In this view, we interpret f as the running cost, g as the initial cost, (α, β) as a control, and w as the state trajectory corresponding to a given control.

In the following we write $\mathcal{A}$ for the space of controls, that is, the set of bounded measurable functions $(\alpha, \beta) \colon [0,T] \to \mathbb{R} \times [0,\infty)$, and we denote the domain of the state trajectories by

$$U := \mathbb{R} \times (0,\infty).$$

The other key ingredient for this control problem is the dynamics of the state trajectories w, given by (2.12).

To consider the function v as a value function in a strictly classical sense, we should express it as an infimum over only the space of controls, simply noting that the initial point $w(0)$ is uniquely determined by the control (α, β) and the terminal condition $w(t) = x$. One should notice however that for a given terminal value $w(t) = x$, there may exist controls such that the corresponding state trajectory w leaves the domain U. Moreover, since the state trajectories satisfy an ODE (2.12) which has quadratic growth (note the quadratic terms in (2.13)), some choices of control may result in trajectories which 'blow up' in a finite time.

Example 2.3.2. Suppose that $c \equiv 2$, and consider the control given by $\alpha \equiv 0$ and $\beta \equiv 1$. Recall the equation (2.11) for the second component of w, which in this case becomes

$$\frac{\mathrm{d}w_2}{\mathrm{d}s}(s) = 4 - w_2(s)^2.$$

If $x_2 = w_2(t) = 1$, then the solution of this equation is given by

$$w_2(s) = \frac{6 - 2e^{4(t-s)}}{3 + e^{4(t-s)}},$$

which leaves the domain U when $s = t - \frac{1}{4}\log(3)$.

If instead $x_2 = w_2(t) = 3$, then the solution is given by

$$w_2(s) = \frac{10 + 2e^{4(t-s)}}{5 - e^{4(t-s)}},$$

which 'blows up' when $s = t - \frac{1}{4}\log(5)$.

The previous example demonstrates some undesirable behaviour of the state trajectories. However we can prevent these types of behaviour by assigning an infinite initial cost to such trajectories, as follows.

Let us extend the domain of the state trajectories w to the entirety of $\mathbb{R}^2$. We extend the function F by setting

$$F(t, x, a, b) = \begin{pmatrix} -(x_1 + \eta_t)a \\ c_t^2 \end{pmatrix} \qquad \text{for} \quad x = (x_1, x_2) \in \mathbb{R}^2 \setminus U, \tag{2.20}$$

noting that, provided c is continuous (which we shall assume), the function $F\colon [0,T] \times \mathbb{R}^2 \times \mathbb{R} \times [0,\infty) \to \mathbb{R}^2$ is continuous in each of its variables. Provided that the resulting extended function is continuous, how we choose to define the function F for $x \in \mathbb{R}^2 \setminus U$ is not important, since we will assign an infinite initial cost to any trajectory which leaves U. It will however be useful later to have F explicitly defined in the whole of $\mathbb{R}^2$.

We suppose that the trajectories w satisfy the ODE (2.12) with this extended function F. We also extend the domain of the initial cost function g to $\mathbb{R}^2$ by defining

$$g(x) = \infty \qquad \text{for} \quad x \in \mathbb{R}^2 \setminus U.$$

Although we don't actually obtain an initial point for trajectories which 'blow up' in a finite time, we may simply prescribe an infinite 'initial' cost to all such trajectories.

For each terminal value $x \in U$, we consider solutions w of (2.12) with the terminal condition

$$w(t) = x \in U. \tag{2.21}$$

We will sometimes write $w = w^{t,x,\alpha,\beta}$ to make explicit the dependence of w on t, x, α and β.

We define the cost functional J, for $t \in [0,T]$, $x \in U$, and $(\alpha,\beta) \in \mathcal{A}$, by

$$J(t,x;\alpha,\beta) = \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g(w(0)),$$

where $w(0) = w^{t,x,\alpha,\beta}(0)$ is the initial value of the solution w of (2.12) and (2.21).

We can now write v as simply

$$v(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}} J(t,x;\alpha,\beta). \tag{2.22}$$

Naturally, we have that $v(0,x) = g(x)$ for all $x \in U$. As noted above, the initial value $w(0)$ may not always lie in U, but such trajectories can be ignored when taking the infimum in (2.22) as we have assigned them an infinite cost.

To summarise, we have:

Control problem for v. *For every $t \in [0,T]$ and $x \in U$, find*

$$v(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}} J(t,x;\alpha,\beta), \tag{2.23}$$

where

$$J(t,x;\alpha,\beta) = \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g(w(0)),$$

and $w = w^{t,x,\alpha,\beta}$ satisfies $w(t) = x$ and

$$\frac{\mathrm{d}w}{\mathrm{d}s}(s) = F(s,w(s),\alpha_s,\beta_s) \qquad \textit{for} \quad s \in [0,t],$$

where the function F is given by (2.13) *and* (2.20).

We will henceforth refer to (2.23) as the definition of v.

2.4 An HJB equation

Having derived the optimal control problem described above, we now proceed to study the value function v in some detail. In particular, we aim to identify it as the unique solution of the associated Hamilton–Jacobi–Bellman (HJB) equation:

$$\frac{\partial v}{\partial t}(t,x) + H\big(t,x,\nabla v(t,x)\big) = 0, \tag{2.24}$$

$$v(0,x) = g(x), \tag{2.25}$$

where the Hamiltonian H is defined by

$$H(t,x,p) = \sup_{(a,b)\in\mathbb{R}\times[0,\infty)} \big\{F(t,x,a,b)\cdot p - f(t,a,b)\big\}. \tag{2.26}$$

Remark 2.4.1. As there is no second order derivative, equations of this type are often referred to as simply Hamilton–Jacobi (HJ), but we will use the name HJB to be consistent with, for example, [6], [7] and [66].

Recall the function F given by (2.13) (and (2.20)). Since the functions c and η are bounded, one can show that there exists a constant $L_F > 0$ such that for all $t \in [0,T]$, $x, z \in U$ and $(a,b) \in \mathbb{R}\times[0,\infty)$, one has

$$\big|F(t,x,a,b)\big| \le L_F\big(1 + |a| + |a||x| + b|x| + b|x|^2\big) \tag{2.27}$$

and

$$\big|F(t,x,a,b) - F(t,z,a,b)\big| \le L_F\big(|a| + b + b|x| + b|z|\big)|x-z|. \tag{2.28}$$

In particular we note that F has linear growth in a and b, locally uniformly in x. That is, for any compact $K \subset U$, there exists a constant $\tilde{L}_F > 0$ such that

$$\big|F(t,x,a,b)\big| \le \tilde{L}_F\big(1 + |a| + b\big) \qquad \text{for all} \quad t\in[0,T],\ x\in K,\ (a,b)\in\mathbb{R}\times[0,\infty). \tag{2.29}$$

Henceforth we shall always assume that the function c is continuous. We also make the following assumptions on the functions f and g.

Assumption 2.4.2. We suppose that the running cost function $f\colon [0,T]\times\mathbb{R}\times(0,\infty)\to\mathbb{R}$ is continuous, nonnegative, and satisfies the following coercivity condition. We suppose that

$$\inf_{t\in[0,T]} \frac{f(t,a,b)}{|a|+b} \longrightarrow \infty \qquad \text{as} \qquad |a|+b \longrightarrow \infty. \tag{2.30}$$

It does not matter, at this stage, whether f is also defined and finite for $b = 0$, or whether f approaches $+\infty$ as $b \to 0^+$.

Definition 2.4.3. We shall denote by $C^{\uparrow}(U;\mathbb{R})$ the class of functions $u\colon U \to \mathbb{R}$ which are finite-valued and continuous on U, and satisfy

$$u(x) \longrightarrow \infty \quad \text{as} \quad |x| \longrightarrow \infty,$$
$$\text{and} \quad \inf_{x_1 \in \mathbb{R}} u(x_1, x_2) \longrightarrow \infty \quad \text{as} \quad x_2 \longrightarrow 0^+.$$

Assumption 2.4.4. As in the previous section, we assume that the initial cost function $g\colon \mathbb{R}^2 \to \mathbb{R} \cup \{\infty\}$ takes the value ∞ on $\mathbb{R}^2 \setminus U$. We also assume that g belongs to the class $C^{\uparrow}(U;\mathbb{R})$, and is locally Lipschitz continuous on U.

Lemma 2.4.5. *The value function v is finite-valued and locally bounded.*

Proof. For a given point $(t,x) = (t,x_1,x_2) \in [0,T] \times U$, consider the control $(\overline{\alpha},\overline{\beta})$ given by

$$\overline{\alpha}_s = \frac{c_s^2 - x_2^2}{2x_2} \qquad \text{and} \qquad \overline{\beta}_s = 1 \qquad \text{for} \quad s \in [0,T].$$

Since c is bounded and f is continuous, we see that $\int_0^t f(s,\overline{\alpha}_s,\overline{\beta}_s)\,\mathrm{d}s < \infty$.

Recall the ODE (2.12) satisfied by the state trajectories, given in this case by

$$\frac{\mathrm{d}w_1}{\mathrm{d}s}(s) = -\big(w_1(s) + \eta_s\big)\frac{c_s^2 + x_2^2}{2x_2} \qquad \text{and} \qquad \frac{\mathrm{d}w_2}{\mathrm{d}s}(s) = 0.$$

The solution of the first equation above has at most exponential growth, and in particular has a finite speed of propagation. It is therefore clear that the initial point $w(0)$ corresponding to this control lies in U. Since g is finite-valued in U, we see that $J(t,x;\overline{\alpha},\overline{\beta}) < \infty$, and hence that $v(t,x) < \infty$.

Let K be a compact subset of U. Note that for $x \in K$, the value of x_2 is bounded away from zero, so that the control $(\overline{\alpha},\overline{\beta})$ defined above (which depends on the choice of x) is uniformly bounded on the compact time interval $[0,T]$. It is also clear that the initial point $w(0) = w^{t,x,\overline{\alpha},\overline{\beta}}(0)$ is bounded in absolute value, and bounded away from the boundary of U at $x_2 = 0$, uniformly for $x \in K$. Since g is assumed to be continuous and hence locally bounded, we deduce that $\sup_{t\in[0,T],\,x\in K} J(t,x;\overline{\alpha},\overline{\beta}) < \infty$, and hence that v is bounded on $[0,T] \times K$. □

Although we are primarily interested in the value function v, it is also worth considering the 'approximate value function', given by taking the infimum of the cost functional J over a class of uniformly bounded controls. More precisely, for a given constant $M > 0$, we define

$$\mathcal{A}^M = \big\{(\alpha,\beta) \in \mathcal{A} \,:\, |\alpha_s| + \beta_s \leq M \quad \text{for all} \quad s \in [0,T]\big\}$$

to be the set of controls which are bounded by M, and let

$$v^M(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}^M} J(t,x;\alpha,\beta) \qquad \text{for} \quad (t,x) \in [0,T]\times U. \tag{2.31}$$

Although restricting to uniformly bounded controls makes parts of our analysis easier, the approximate value function defined above is actually in some ways more awkward to deal with, as it is not finite-valued in the whole of U.

Recall from Example 2.3.2 instances of controls and terminal conditions such that the corresponding state trajectory either leaves the domain U or 'blows up' in a finite time. As stated earlier, we associate an infinite initial cost to all such unphysical trajectories. The point here is that if we adopt a uniform bound on the size of the parameters, then there will exist terminal conditions such that the corresponding trajectories exhibit these anomalous behaviours for *all* choices of control, meaning that the value function itself is infinite-valued at these points.

Recall the control $(\overline{\alpha},\overline{\beta})$ from the proof of Lemma 2.4.5. Note that $\overline{\alpha}_s \to +\infty$ as $x_2 \to 0^+$ (whenever $c_s \neq 0$). This is an indication of the fact that, as the terminal value x approaches the boundary of U at $x_2 = 0$, we require larger and larger controls to ensure that the corresponding state trajectory remains inside U. This can be seen by noticing that for small values of $w_2(t)$, the dominant term in the ODE (2.11) is c_t^2, which pushes the initial value of the solution towards (and indeed beyond) the boundary of U. Recalling that $x_2 = 1/\sigma^2$, we note that this behaviour corresponds to extremely large values of the conditional variance being considered to be very unreasonable.

Thus, for any $M > 0$, the approximate value function v^M is infinite-valued in a nontrivial (and time-dependant) subset of the domain U. In the following we denote

$$Q := (0,T]\times U$$

and write Q^M for the part of the domain where v^M is finite-valued, that is

$$\begin{aligned} Q^M &:= \big\{(t,x)\in Q \ : \ v^M(t,x) < \infty\big\} \\ &= \big\{(t,x)\in Q \ : \ \exists(\alpha,\beta)\in\mathcal{A}^M \ \text{ such that } \ w^{t,x,\alpha,\beta}(0)\in U\big\}. \end{aligned}$$

Since $g \in C^{\uparrow}(U;\mathbb{R})$, it can be seen that $v^M(t,x)$ tends to infinity as (t,x) approaches the boundary of the set $Q \setminus Q^M$.

Remark 2.4.6. It is not hard to see that Q^M is an open subset of Q. Further, by inspecting the proof of Lemma 2.4.5, we can deduce that, for any compact $K \subset U$, we will have that $(0,T)\times K \subset Q^M$ for all sufficiently large $M > 0$, so that in particular

$\bigcup_{M>0} Q^M = Q$. Since any given control $(\alpha, \beta) \in \mathcal{A}$ is bounded on the compact time interval $[0, T]$, we also have that $v(t, x) = \inf_{M>0} v^M(t, x)$ for all $(t, x) \in Q$.

When considering uniformly bounded controls, it is natural to also consider the modified HJB equation:

$$\frac{\partial u}{\partial t}(t, x) + H_M\big(t, x, \nabla u(t, x)\big) = 0, \tag{2.32}$$

$$u(0, x) = g(x), \tag{2.33}$$

where H_M is the modified Hamiltonian, given by

$$H_M(t, x, p) = \sup\left\{ F(t, x, a, b) \cdot p - f(t, a, b) \;\middle|\; \begin{array}{c} (a, b) \in \mathbb{R} \times [0, \infty) \\ \text{such that } \ |a| + b \leq M \end{array} \right\}. \tag{2.34}$$

Lemma 2.4.7. *The Hamiltonian H and modified Hamiltonian H_M are continuous.*

Proof. As the functions F and f are continuous, and the supremum in (2.34) is taken over a compact set, it is straightforward to see that H_M is continuous for any given M.

Let $R > 0$, let K be a compact subset of U, and write $\overline{B}_R(0)$ for the closed ball in $\mathbb{R}^2$, centred at 0 with radius R. It follows from the conditions (2.29) and (2.30) that

$$\begin{aligned} &\sup_{\substack{x \in K,\ 0 \leq t \leq T \\ p \in \overline{B}_R(0)}} \big\{ F(t, x, a, b) \cdot p - f(t, a, b) \big\} \\ &\qquad \leq \big(1 + |a| + b\big)\left(\tilde{L}_F R - \inf_{0 \leq t \leq T} \frac{f(t, a, b)}{1 + |a| + b} \right) \longrightarrow -\infty \end{aligned}$$

as $|a| + b \to \infty$. We therefore deduce the existence of an $M > 0$ such that $H(t, x, p) = H_M(t, x, p)$ for all $(t, x, p) \in [0, T] \times K \times \overline{B}_R(0)$. Continuity of H then clearly follows from the continuity of H_M. □

The value function v and approximate value function v^M satisfy the following dynamic programming principle.

Proposition 2.4.8. *For any t, x and $r \in [0, t]$, we have*

$$v(t, x) = \inf_{(\alpha, \beta) \in \mathcal{A}} \left\{ \int_r^t f(s, \alpha_s, \beta_s)\, \mathrm{d}s + v\big(r, w^{t, x, \alpha, \beta}(r)\big) \right\}.$$

The statement of the dynamic programming principle for the approximate value function v^M is identical to the above, except that one only considers controls $(\alpha, \beta) \in \mathcal{A}^M$. The proof of this result is standard; see e.g. Yong and Zhou [66].

Definition 2.4.9. We say that a continuous function $u\colon Q \to \mathbb{R}$ is a viscosity subsolution (resp. supersolution) of (2.24)–(2.25) if $u(0,x) \leq$ (resp. $\geq$) $g(x)$ for all $x \in U$, and, for any point $(t,x) \in Q$,

$$\frac{\partial \varphi}{\partial t}(t,x) + H\big(t,x,\nabla\varphi(t,x)\big) \leq \text{(resp. } \geq\text{) } 0 \tag{2.35}$$

for every smooth function $\varphi\colon Q \to \mathbb{R}$ such that $u-\varphi$ has a local maximum (resp. local minimum) at the point (t,x). We say that u is a viscosity solution if it is both a viscosity subsolution and a viscosity supersolution.

As noted above, the approximate value function v^M is only finite-valued in the subdomain Q^M. However, we still wish to consider v^M as being a supersolution of the modified HJB equation. We therefore consider viscosity supersolutions of (2.32)–(2.33) in the following sense.

Definition 2.4.10. We say that a function $u\colon Q \to \mathbb{R} \cup \{\infty\}$ is a viscosity supersolution of (2.32)–(2.33) if

- $u(0,x) \geq g(x)$ for all $x \in U$,
- the set $Q^u := \{(t,x) \in Q \,:\, u(t,x) < \infty\}$ is open,
- u satisfies the usual definition of viscosity supersolution (as in Definition 2.4.9 with H replaced by H_M) on Q^u,
- and moreover u is 'continuous' in the sense that, as well as being continuous on Q^u, we also have that $u(t,x)$ tends to infinity as (t,x) approaches the boundary of the set $Q \setminus Q^u$.

We can use Proposition 2.4.8 to obtain the following result, whose proof is adapted from [66, Chapter 4].

Proposition 2.4.11. *The value function v is a viscosity solution of the HJB equation* (2.24)–(2.25)*, and the approximate value function v^M is a viscosity supersolution of the modified HJB equation* (2.32)–(2.33)*.*

Proof. We first note that v is continuous on Q, and v^M is continuous on Q^M, which follows from Theorem 2.5.3 below.

Suppose that φ is a smooth function such that $v - \varphi$ has a local maximum at a point $(t,x) \in Q$. Fix $(a,b) \in \mathbb{R} \times [0,\infty)$, and define a control $(\alpha,\beta) \in \mathcal{A}$ by

$(\alpha_t, \beta_t) \equiv (a, b)$, and let $w = w^{t,x,\alpha,\beta}$ be the corresponding state trajectory. Then, by Proposition 2.4.8, for $r \in [0, t)$ with $t - r$ sufficiently small, we have

$$\begin{aligned} 0 &\leq \frac{v(t,x) - \varphi(t,x) - v(r, w(r)) + \varphi(r, w(r))}{t-r} \\ &\leq \frac{1}{t-r}\bigg(\int_r^t f(s,a,b)\,\mathrm{d}s - \varphi(t,x) + \varphi(r,w(r))\bigg) \\ &\longrightarrow f(t,a,b) - \frac{\partial \varphi}{\partial t}(t,x) - F(t,x,a,b)\cdot\nabla\varphi(t,x) \qquad \text{as} \quad r \to t. \end{aligned}$$

It follows that

$$\frac{\partial \varphi}{\partial t}(t,x) + H\big(t,x,\nabla\varphi(t,x)\big) \leq 0,$$

so v is a viscosity subsolution of (2.24).

Suppose that φ is a smooth function such that $v^M - \varphi$ has a local minimum at some point $(t,x) \in Q^M$. Let $\varepsilon > 0$. By Proposition 2.4.8, for any $r \in [0,t)$ with $t - r$ sufficiently small, we can find a control $(\alpha^r, \beta^r) \in \mathcal{A}^M$ such that, writing $w = w^{t,x,\alpha^r,\beta^r}$, we have

$$\begin{aligned} 0 &\geq v^M(t,x) - \varphi(t,x) - v^M(r,w(r)) + \varphi(r,w(r)) \\ &\geq -\varepsilon(t-r) + \int_r^t f(s,\alpha_s^r,\beta_s^r)\,\mathrm{d}s - \varphi(t,x) + \varphi(r,w(r)). \end{aligned}$$

Rearranging, it follows that

$$\begin{aligned} -\varepsilon &\leq \frac{1}{t-r}\int_r^t \bigg(\frac{\partial\varphi}{\partial t}(s,w(s)) + F(s,w(s),\alpha_s^r,\beta_s^r)\cdot\nabla\varphi(s,w(s)) - f(s,\alpha_s^r,\beta_s^r)\bigg)\mathrm{d}s \\ &\leq \frac{1}{t-r}\int_r^t \bigg(\frac{\partial\varphi}{\partial t}(s,w(s)) + H_M\big(s,w(s),\nabla\varphi(s,w(s))\big)\bigg)\mathrm{d}s \\ &\longrightarrow \frac{\partial\varphi}{\partial t}(t,x) + H_M\big(t,x,\nabla\varphi(t,x)\big) \qquad \text{as} \quad r \to t. \end{aligned}$$

One should note that the function w in the above depends on r, so care should be taken when taking the limit in the last line. The convergence is justified by the fact that the controls $\{(\alpha^r,\beta^r)\}_{0\leq r<t}$ are uniformly bounded (by M), which implies that the functions $\{w^{t,x,\alpha^r,\beta^r}\}_{0\leq r<t}$ are continuous at the point (t,x), uniformly in r.

Taking $\varepsilon \to 0$, we deduce that v^M is a viscosity supersolution of (2.32)–(2.33) in the sense of Definition 2.4.10.

Since $H_M \leq H$, we see that v^M is also a viscosity supersolution of (2.24)–(2.25) in Q^M. As noted in Remark 2.4.6, we have that $\bigcup_{M>0} Q^M = Q$, and $v(t,x) = \inf_{M>0} v^M(t,x)$ for all $(t,x) \in Q$, from which it follows that v is also a viscosity supersolution of (2.24)–(2.25). □

2.5 Regularity of the value function

Recall our assumption that the initial cost function g belongs to the class $C^{\uparrow}(U;\mathbb{R})$ (defined in Definition 2.4.3). We now show that the value function v also satisfies a similar type of 'weak coercivity', which one may also think of as an 'explosive boundary condition'.

Definition 2.5.1. We shall denote by $C^{\uparrow}(Q;\mathbb{R})$ the class of functions $u\colon Q \to \mathbb{R}\cup\{\infty\}$ which are continuous, and satisfy

$$\inf_{t\in[0,T]} u(t,x) \longrightarrow \infty \qquad \text{as} \quad |x| \longrightarrow \infty, \tag{2.36}$$

$$\inf_{t\in[0,T],\ x_1\in\mathbb{R}} u(t,x_1,x_2) \longrightarrow \infty \qquad \text{as} \quad x_2 \longrightarrow 0^+. \tag{2.37}$$

Lemma 2.5.2. $v \in C^{\uparrow}(Q;\mathbb{R})$.

Proof. Using the fact that the functions c and η are bounded, we have

$$\begin{aligned}\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}s}|w(s)|^2 &= w(s)\cdot F(s,w(s),\alpha_s,\beta_s)\\ &= -w_1(s)\big(w_1(s)+\eta_s\big)\big(\alpha_s+\beta_s w_2(s)\big) - \beta_s w_2(s)^3 - 2\alpha_s w_2(s)^2 + c_s^2 w_2(s)\\ &\le C\big(1+|\alpha_s|+\beta_s\big)\big(1+|w(s)|^2\big),\end{aligned}$$

where crucially we notice that, since $w_2(s) > 0$ and $\beta_s \ge 0$, we are able to drop the highest order terms in w. Rearranging and integrating this expression, we obtain

$$\log\big(1+|x|^2\big) \le \log\big(1+|w(0)|^2\big) + C\int_0^t \big(1+|\alpha_s|+\beta_s\big)\,\mathrm{d}s, \tag{2.38}$$

which holds for any choice of control (α,β). Using (2.30), we then obtain

$$\log\big(1+|x|^2\big) \le \log\big(1+|w(0)|^2\big) + C\int_0^t \big(1+f(s,\alpha_s,\beta_s)\big)\,\mathrm{d}s, \tag{2.39}$$

with a possibly different constant C. Recalling the definition of the value function,

$$v(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}}\left\{\int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g\big(w^{t,x,\alpha,\beta}(0)\big)\right\},$$

and the fact that $g \in C^{\uparrow}(U;\mathbb{R})$, we deduce that (2.36) holds for v.

The second component of a given state trajectory satisfies

$$\begin{aligned}\frac{\mathrm{d}w_2}{\mathrm{d}s}(s) &= -\beta_s w_2(s)^2 - 2\alpha_s w_2(s) + c_s^2\\ &\ge -2\big(|\alpha_s|+\beta_s\big)w_2(s)\big(1+w_2(s)\big),\end{aligned}$$

from which, using (2.30) again, we obtain

$$\log\left(\frac{w_2(0)}{1+w_2(0)}\right) - C\int_0^t \big(1+f(s,\alpha_s,\beta_s)\big)\,\mathrm{d}s \le \log\left(\frac{x_2}{1+x_2}\right) \tag{2.40}$$

for some constant $C > 0$. From the definition of the value function and fact that $g \in C^{\uparrow}(U;\mathbb{R})$, we similarly deduce that (2.37) holds for v. □

Recall Assumptions 2.4.2 and 2.4.4, in particular the assumption that the initial cost function g is locally Lipschitz continuous on U. Under these assumptions we have the following.

Theorem 2.5.3. *The value function v is locally Lipschitz continuous.*

Proof. The following argument takes inspiration from Section 2 of Bardi and Da Lio [7].

Step 1. Let K be a compact subset of U.

For notational simplicity, throughout the proof, whenever we use the symbol C we shall mean a positive constant which can be chosen independently of $t \in [0,T]$ and $x \in K$, but we will allow this constant to be different in each occurrence.

Let us recall the definition of the value function:

$$v(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}} \left\{ \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g\big(w^{t,x,\alpha,\beta}(0)\big) \right\}.$$

By Lemma 2.4.5, v is bounded on $[0,T]\times K$. Since f and g are bounded below, it follows that there exists a constant $\lambda > 0$ such that, when considering the value of $v(t,x)$ for $t \in [0,T]$ and $x \in K$, it is sufficient to consider only those controls (α,β) which satisfy

$$\int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s < \lambda \qquad \text{and} \qquad g\big(w^{t,x,\alpha,\beta}(0)\big) < \lambda. \tag{2.41}$$

We then deduce from (2.30) the following bound which holds for all such controls:

$$\int_0^t |\alpha_s| + \beta_s\,\mathrm{d}s \le C. \tag{2.42}$$

It follows from the second bound in (2.41) and the fact that $g \in C^{\uparrow}(U;\mathbb{R})$ (recall Definition 2.4.3) that the initial point $w^{t,x,\alpha,\beta}(0)$ is bounded in absolute value, and also bounded away from the boundary of U at $x_2 = 0$, uniformly for $x \in K$.

Recalling the estimate (2.39),

$$\log\big(1+|x|^2\big) \le \log\big(1+|w(0)|^2\big) + C\int_0^t \big(1+f(s,\alpha_s,\beta_s)\big)\,\mathrm{d}s,$$

we deduce a (partial) finite speed of propagation property. That is, although the state trajectories w may have an infinite speed of propagation 'toward infinity', they can only return 'toward zero' with a finite speed.

Since the terminal points $w^{t,x,\alpha,\beta}(t) = x$ lie in the compact set K and, as noted above, the initial points $w^{t,x,\alpha,\beta}(0)$ are bounded in absolute value, it follows from this finite speed of propagation property that the entire trajectory $\{w(s) : 0 \le s \le t\}$ is uniformly bounded in absolute value, that is

$$\big|w^{t,x,\alpha,\beta}(s)\big| \le C \tag{2.43}$$

for all $t \in [0,T]$, $x \in K$ and $0 \le s \le t$.

Similarly, we can deduce from the estimate (2.40) that the trajectory $\{w(s) : 0 \le s \le t\}$ is bounded away from the boundary of U at $x_2 = 0$, uniformly in $t \in [0,T]$ and $x \in K$.

Step 2. Let $x, z \in K$. Writing $w^x = w^{t,x,\alpha,\beta}$ and $w^z = w^{t,z,\alpha,\beta}$, we have, by (2.28),

$$\begin{aligned}
&-\frac{\mathrm{d}}{\mathrm{d}s}|w^x(s) - w^z(s)|^2 \\
&\quad = -2\big(w^x(s) - w^z(s)\big) \cdot \big(F(s, w^x(s), \alpha_s, \beta_s) - F(s, w^z(s), \alpha_s, \beta_s)\big) \\
&\quad \le 2\big|w^x(s) - w^z(s)\big|\big|F(s, w^x(s), \alpha_s, \beta_s) - F(s, w^z(s), \alpha_s, \beta_s)\big| \\
&\quad \le 2L_F\big(|\alpha_s| + \beta_s + \beta_s|w^x(s)| + \beta_s|w^z(s)|\big)|w^x(s) - w^z(s)|^2,
\end{aligned}$$

and so by (2.43) we obtain

$$\frac{\mathrm{d}}{\mathrm{d}s}|w^x(s) - w^z(s)|^2 \ge -C(|\alpha_s| + \beta_s)|w^x(s) - w^z(s)|^2.$$

Integrating this expression and using (2.42), it follows that

$$\big|w^{t,x,\alpha,\beta}(0) - w^{t,z,\alpha,\beta}(0)\big| \le C|x - z|, \tag{2.44}$$

and that this bound holds uniformly for $t \in [0,T]$ and $x, z \in K$.

We have

$$\begin{aligned}
\big|v(t,x) - v(t,z)\big| &\le \sup_{\alpha,\beta} \big|J(t,x;\alpha,\beta) - J(t,z;\alpha,\beta)\big| \\
&= \sup_{\alpha,\beta} \big|g\big(w^{t,x,\alpha,\beta}(0)\big) - g\big(w^{t,z,\alpha,\beta}(0)\big)\big| \\
&\le \sup_{\alpha,\beta} C\big|w^{t,x,\alpha,\beta}(0) - w^{t,z,\alpha,\beta}(0)\big| \\
&\le C|x - z|,
\end{aligned} \tag{2.45}$$

which holds uniformly for $t \in [0,T]$ and $x, z \in K$, where we have used (2.44), and the assumption that g is locally Lipschitz. Thus, $v = v(t,x)$ is locally Lipschitz in x, uniformly in t.

Step 3. Fix some control $(\overline{\alpha}, \overline{\beta})$ such that (2.41) holds. By Proposition 2.4.8, for any $r \in [0,t]$, we may restrict ourselves to controls (α, β) which satisfy

$$\int_r^t f(s, \alpha_s, \beta_s)\,\mathrm{d}s \leq \int_r^t f(s, \overline{\alpha}_s, \overline{\beta}_s)\,\mathrm{d}s + \big|v\big(w^{t,x,\alpha,\beta}(r)\big) - v\big(w^{t,x,\overline{\alpha},\overline{\beta}}(r)\big)\big|. \tag{2.46}$$

Our next aim is to obtain estimates for the terms on the right-hand side of (2.46).

We showed in Step 1 that the trajectories are bounded in absolute value, and bounded away from the boundary at $x_2 = 0$. Since we also showed in Step 2 that v is locally Lipschitz in x, we deduce that

$$\big|v\big(r, w^{t,x,\alpha,\beta}(r)\big) - v\big(r, w^{t,x,\overline{\alpha},\overline{\beta}}(r)\big)\big| \leq C\big|w^{t,x,\alpha,\beta}(r) - w^{t,x,\overline{\alpha},\overline{\beta}}(r)\big|. \tag{2.47}$$

We have

$$\begin{aligned}\big|w^{t,x,\alpha,\beta}(r) - x\big| &= \left|\int_r^t \frac{\mathrm{d}w}{\mathrm{d}s}(s)\,\mathrm{d}s\right| \leq \int_r^t \big|F(s, w(s), \alpha_s, \beta_s)\big|\,\mathrm{d}s \\ &\leq C\int_r^t \big(1 + |\alpha_s| + \beta_s\big)\,\mathrm{d}s,\end{aligned} \tag{2.48}$$

and similarly, since the control $(\overline{\alpha}, \overline{\beta})$ is fixed (and bounded),

$$\big|w^{t,x,\overline{\alpha},\overline{\beta}}(r) - x\big| \leq C\int_r^t \big(1 + |\overline{\alpha}_s| + \overline{\beta}_s\big)\,\mathrm{d}s \leq C(t-r), \tag{2.49}$$

and

$$\int_r^t f(s, \overline{\alpha}_s, \overline{\beta}_s)\,\mathrm{d}s \leq C(t-r). \tag{2.50}$$

Combining (2.48) and (2.49) with (2.47), and then substituting this and (2.50) into (2.46), we obtain

$$\int_r^t f(s, \alpha_s, \beta_s)\,\mathrm{d}s \leq C\left((t-r) + \int_r^t |\alpha_s| + \beta_s\,\mathrm{d}s\right).$$

For any $\varepsilon > 0$, by (2.30), there exists some new constant C_ε depending on ε, such that

$$\int_r^t f(s, \alpha_s, \beta_s)\,\mathrm{d}s \leq C_\varepsilon(t-r) + \varepsilon\int_r^t f(s, \alpha_s, \beta_s)\,\mathrm{d}s$$

Choosing $\varepsilon < 1$ and rearranging, we obtain

$$\int_r^t f(s, \alpha_s, \beta_s)\,\mathrm{d}s \leq C(t-r), \tag{2.51}$$

from which it also easily follows, using (2.30) again, that

$$\int_r^t |\alpha_s| + \beta_s \, \mathrm{d}s \leq C(t-r). \tag{2.52}$$

Step 4. By Proposition 2.4.8, there exists a sequence of controls $(\alpha^k, \beta^k)_{k\geq 1}$ such that

$$v(t,x) = \lim_{k\to\infty} \int_r^t f(s, \alpha_s^k, \beta_s^k) \, \mathrm{d}s + v\big(r, w^{t,x,\alpha^k,\beta^k}(r)\big).$$

We have, using the fact that v is locally Lipschitz in x and the bounds (2.48), (2.51) and (2.52),

$$\begin{aligned}
\Big| \int_r^t f(s, \alpha_s^k, \beta_s^k) \, \mathrm{d}s + v\big(r, w^{t,x,\alpha^k,\beta^k}(r)\big) - v(r,x) \Big| \\
&\leq \int_r^t f(s, \alpha_s^k, \beta_s^k) \, \mathrm{d}s + C\big|w^{t,x,\alpha^k,\beta^k}(r) - x\big| \\
&\leq \int_r^t f(s, \alpha_s^k, \beta_s^k) \, \mathrm{d}s + C \int_r^t \big(1 + |\alpha_s^k| + \beta_s^k\big) \, \mathrm{d}s \\
&\leq C(t-r),
\end{aligned}$$

and taking the limit as $k \to \infty$ we conclude that

$$\big|v(t,x) - v(r,x)\big| \leq C(t-r),$$

which holds uniformly for $0 \leq r \leq t \leq T$ and $x \in K$. Thus, $v = v(t,x)$ is Lipschitz in t, locally uniformly in x. □

One can also show that the approximate value function v^M is locally Lipschitz continuous on Q^M, but we will not make use of this result.

An immediate consequence of Theorem 2.5.3 is that the value function v belongs to the Sobolev space $W^{1,\infty}_{\mathrm{loc}}(Q)$. In particular, v has a weak derivative, has a classical derivative almost everywhere, and these two notions of derivative coincide almost everywhere (see Evans [28]). Moreover, v satisfies the HJB equation (2.24) in the classical sense at each point where its classical derivative exists.

2.6 Uniqueness for the HJB equation

In the previous section we saw that the value function v is a viscosity solution of the HJB equation (2.24)–(2.25). We now turn our attention to the question of whether this solution is unique.

There is a substantial amount of existing literature on uniqueness results for solutions of first-order HJB equations. The standard approach to obtaining such results is via a 'doubling the variables' type argument, which provides what is essentially a maximum principle for viscosity solutions of certain PDEs (see Crandall, Ishii and Lions [20]). However, typical treatments of such equations on an unbounded domain (e.g. in Barles [1] or Evans [28]) assume that the solutions are bounded and uniformly continuous, whereas our solution v exhibits explosive behaviour near the boundary of U. In Yong and Zhou [66] solutions are only assumed to be continuous, but the Hamiltonian is assumed to grow at most linearly in the spatial variable x. In our case the state trajectories obey an ODE (2.12) which is quadratic in x, which means that our Hamiltonian also has quadratic growth in x. Moreover, although the function η is locally Hölder continuous, it is not locally Lipschitz, as would be required in order to apply the sub/super-optimality principle (see Lions and Souganidis [47]).

On the other hand, when working on a bounded domain where the comparison is assumed to hold on the parabolic boundary of the domain, standard comparison results are less restrictive in their assumptions on the Hamiltonian. Our strategy here is therefore to exploit properties of the value function in order to be able to restrict attention to a bounded subdomain, and then apply a standard comparison theorem from Barles [1].

The rest of this section is dedicated to the proof of Theorem 2.6.2 below, which we shall prove under the following additional mild assumption.

Assumption 2.6.1. We assume that $f(t,0,0) < \infty$ for all $t \in [0,T]$.

To the best of our knowledge our technique is new, and is applicable to more general nonlinear optimal control problems. The key is the fact that the value function v depends only on local values of the initial cost function g. We can therefore manipulate the asymptotic behaviour of g, and hence that of v, without changing the local behaviour of v. This is the motivation for Propositions 2.6.3 and 2.6.5 below. We will then be able to deduce the following result.

Theorem 2.6.2. *The value function v is the unique viscosity solution of the HJB equation* (2.24)–(2.25) *in the class $C^{\uparrow}(Q;\mathbb{R})$.*

Note that, rather than any sort of integrability or regularity condition in the usual sense, the additional condition we impose to establish uniqueness for our equation is actually an 'explosive boundary condition'. Further, notice that although we insist

that the relevant solutions converge to infinity as we approach the boundary of U, we make no restriction whatsoever on the *speed* of this convergence.

In the following, by a viscosity (sub/super)solution we will always mean a continuous viscosity (sub/super)solution. As in the previous section, we consider supersolutions in the sense of Definition 2.4.10. In particular, we allow viscosity supersolutions to be infinite-valued, provided that they are still 'continuous' in the sense that they tend to infinity as they approach the part of the domain where they are infinite-valued.

It will be helpful to first restrict ourselves to the case of uniformly bounded controls, and consider solutions of the modified HJB equation (2.32)–(2.33). From the bound (2.28) on the function F, we easily deduce that there exists a constant $L_M > 0$ such that

$$\big|H_M(t,x,p) - H_M(t,z,p)\big| \leq L_M\big(1 + |x| + |z|\big)|x - z||p| \tag{2.53}$$

for all $t \in [0,T]$, $x, z \in U$ and $p \in \mathbb{R}^2$.

Proposition 2.6.3. *Fix an $M > 0$, and let u be a viscosity subsolution of the modified HJB equation* (2.32)–(2.33)*. Then $u \leq v^M$.*

Proof. Fix an arbitrary compact set $K \subset U$, and choose $T_0 \in (0,T]$ sufficiently small such that $(0,T_0) \times K \subset Q^M$. Let us recall the definition of the approximate value function:

$$v^M(t,x) = \inf_{(\alpha,\beta)\in\mathcal{A}^M} \left\{ \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + g\big(w^{t,x,\alpha,\beta}(0)\big) \right\}.$$

As v^M is continuous on Q^M, it is bounded on $(0,T_0) \times K$. Since f and g are bounded below, it follows that there exists a constant $\lambda > 0$ such that, when considering the value of $v^M(t,x)$ for $t \in (0,T_0)$ and $x \in K$, it is sufficient to consider only those controls (α,β) which satisfy

$$g\big(w^{t,x,\alpha,\beta}(0)\big) < \lambda. \tag{2.54}$$

For $\delta > 0$, define the set $U_\delta := \{x = (x_1,x_2) \in U \,:\, |x| < 1/\delta,\ x_2 > \delta\}$. Since $g \in C^{\uparrow}(U;\mathbb{R})$ (recall Definition 2.4.3), there exists a $\delta > 0$ sufficiently small such that

- $K \subset U_\delta$, and
- $g > \lambda$ on $U \setminus U_\delta$.

For this δ, let G be a continuous function belonging to the class $C^{\uparrow}(U;\mathbb{R})$, such that

- $G \geq g$,
- $G = g$ on U_δ,

- $G > \lambda$ on $U \setminus U_\delta$,
- $G = \infty$ on $\mathbb{R}^2 \setminus U$.

We now define

$$V(t,x) := \inf_{(\alpha,\beta)\in\mathcal{A}^M} \left\{ \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + G\big(w^{t,x,\alpha,\beta}(0)\big) \right\} \tag{2.55}$$

where, just as in Section 2.3 when we considered the value function v, we associate an infinite 'initial' cost to state trajectories which 'blow up' after a finite time.

The function V is defined in the same way as the approximate value function v^M, except with a different initial cost function, namely G instead of g. We note that V is also finite-valued precisely in the subdomain Q^M. Further, the result of Proposition 2.4.11 applies equally well to the function V, which is hence a viscosity supersolution of the modified HJB equation (2.32)–(2.33).

Since $g > \lambda$ and $G > \lambda$ in $\mathbb{R}^2 \setminus U_\delta$, it follows from (2.54) that when considering the values of $v^M(t,x)$ and $V(t,x)$ for $t \in (0,T_0)$ and $x \in K$, we may ignore all controls (α,β) such that the corresponding initial point $w^{t,x,\alpha,\beta}(0)$ lies in $\mathbb{R}^2 \setminus U_\delta$. Since, $G = g$ on U_δ, we deduce that

$$v^M(t,x) = V(t,x) \qquad \text{for all} \quad (t,x) \in (0,T_0) \times K. \tag{2.56}$$

Recall the estimates (2.39) and (2.40) from the proof of Lemma 2.5.2, which in particular hold for all controls $(\alpha,\beta) \in \mathcal{A}^M$. Since $G \in C^\uparrow(U;\mathbb{R})$, we can argue exactly as in Lemma 2.5.2 that $V \in C^\uparrow(Q;\mathbb{R})$. Moreover, since u is continuous and hence locally bounded, and since we are only considering controls which are uniformly bounded[1] (by M), we see from (2.39) and (2.40) that by choosing G to grow sufficiently quickly as $|x| \to \infty$ and as $x_2 \to 0$, we can ensure that there exists an open subset E, compactly contained in U, such that

- $K \subset E$, and
- $V \geq u$ on $[0,T_0] \times \partial E$ (where ∂E denotes the boundary of E in U).

Since $V(0,x) = G(x) \geq g(x) \geq u(0,x)$ for all $x \in U$, we see that $V \geq u$ on $([0,T_0] \times \partial E) \cup (\{0\} \times \overline{E})$, i.e. on the parabolic boundary of $(0,T_0) \times E$. Since (2.53) holds,

[1]The assumption of uniformly bounded controls is essential here. This argument could not therefore be applied directly to the value function v.

we may apply Theorem 5.1 in Barles [1] on the subdomain[2] $(0, T_0) \times E$ with the subsolution u and the supersolution V, to deduce that $u \leq V$ on $(0, T_0) \times E$.

By (2.56), we then have that

$$u \leq V = v^M \quad \text{on} \quad (0, T_0) \times K.$$

Since any point $(t, x) \in Q^M$ belongs to a set of the form $(0, T_0) \times K$ for some compact $K \subset U$ and some $T_0 \in (0, T]$, we have the result. □

Corollary 2.6.4. *Let u be a viscosity subsolution of the HJB equation* (2.24)–(2.25). *Then $u \leq v$.*

Proof. Fix an arbitrary point $(t, x) \in Q$, and let $\varepsilon > 0$. Then there exists a control $(\alpha, \beta) \in \mathcal{A}$ such that $J(t, x; \alpha, \beta) < v(t, x) + \varepsilon$.

As we assume that each individual control is bounded, there exists an $M > 0$ such that $(\alpha, \beta) \in \mathcal{A}^M$.

Since u is a viscosity subsolution of the HJB equation (2.24)–(2.25), and $H_M \leq H$, it follows that u is also a viscosity subsolution of the modified HJB equation (2.32)–(2.33). By Proposition 2.6.3, we have that $u \leq v^M$. In particular, since $(\alpha, \beta) \in \mathcal{A}^M$, we have that

$$u(t, x) \leq v^M(t, x) \leq J(t, x; \alpha, \beta) < v(t, x) + \varepsilon.$$

Letting $\varepsilon \to 0$, we deduce the result. □

Proposition 2.6.5. *Suppose that $\hat{u}$ is a viscosity supersolution of the HJB equation* (2.24)–(2.25)*, and assume that $\hat{u} \in C^{\uparrow}(Q; \mathbb{R})$. Then $\hat{u} \geq v$.*

Proof. Fix an arbitrary compact set $K \subset U$. We can argue, exactly as in the proof of Proposition 2.6.3, that there exists a constant $\lambda > 0$ such that, when considering the value of $v(t, x)$ for $x \in K$, it is sufficient to consider only those controls (α, β) which satisfy

$$g\big(w^{t,x,\alpha,\beta}(0)\big) < \lambda. \tag{2.57}$$

Since $g \in C^{\uparrow}(U; \mathbb{R})$, there exists $\delta > 0$ sufficiently small such that

- $K \subset U_\delta$, and
- $g > \lambda$ on $U \setminus U_\delta$,

[2]Strictly speaking the supersolution V could be infinite-valued in part of this domain, but since V is continuous in the sense of Definition 2.4.10, this requires only a trivial extension of the theorem, with essentially no change to its proof.

where $U_\delta := \{x = (x_1, x_2) \in U : |x| < 1/\delta,\ x_2 > \delta\}$. For this δ, let G be a finite-valued, locally Lipschitz function on $\mathbb{R}^2$ such that

- $G \leq g$,
- $G = g$ on U_δ,
- $G > \lambda$ on $\mathbb{R}^2 \setminus U_\delta$.

We then define

$$V(t,x) := \inf_{(\alpha,\beta)\in\mathcal{A}} \left\{ \int_0^t f(s,\alpha_s,\beta_s)\,\mathrm{d}s + G\big(w^{t,x,\alpha,\beta}(0)\big) \right\}. \tag{2.58}$$

As usual, we associate an infinite 'initial' cost to state trajectories which 'blow up' after a finite time. The function V is defined similarly to the value function v, except with initial cost function G instead of g. Note however that, since G is finite-valued in the entirety of $\mathbb{R}^2$, the state trajectories are allowed to leave the domain U. It is for this reason that we explicitly defined the function F for values of $x \in \mathbb{R}^2 \setminus U$.

We can argue, exactly as in the first part of the proof of Proposition 2.4.11, that V is a viscosity subsolution of the HJB equation (2.24)–(2.25).

Since $g > \lambda$ and $G > \lambda$ in $\mathbb{R}^2 \setminus U_\delta$, it follows from (2.57) that when considering the values of $v(t,x)$ and $V(t,x)$ for $x \in K$, we may ignore all controls (α,β) such that the corresponding initial point $w^{t,x,\alpha,\beta}(0)$ lies in $\mathbb{R}^2 \setminus U_\delta$. Since, $G = g$ on U_δ, we deduce that

$$v(t,x) = V(t,x) \qquad \text{for all} \quad (t,x) \in (0,T) \times K. \tag{2.59}$$

Heuristically, if G grows slowly as $|x| \to \infty$, then V should also grow slowly as $|x| \to \infty$. In particular, since $\hat{u} \in C^\uparrow(Q;\mathbb{R})$, if G grows sufficiently slowly, then we should be able to ensure that V grows slower than $\hat{u}$, so that $\hat{u}$ dominates V asymptotically. We will now make this more precise.

Let us write (α^0,β^0) for the control satisfying $(\alpha_s,\beta_s) = (0,0)$ for all $s \in [0,T]$, and $w^0 = w^{t,x,\alpha^0,\beta^0}$ for the corresponding state trajectory. Recall the function F which defines the ODE satisfied by the state trajectories, given by (2.13) and (2.20). In this case we have

$$\frac{\mathrm{d}w^0}{\mathrm{d}s}(s) = F(s,w^0(s),0,0) = \begin{pmatrix} 0 \\ c_s^2 \end{pmatrix}$$

so that

$$w^0(0) = x - \begin{pmatrix} 0 \\ \int_0^t c_s^2\,\mathrm{d}s \end{pmatrix}.$$

In other words, for any point $(t,x) \in Q$, there exists a control such that the corresponding state trajectory w 'doesn't move very far'.

Recalling Assumption 2.6.1, since

$$V(t,x) \leq G(w^0(0)) + T \sup_{0 \leq s \leq T} f(s,0,0),$$

and $\hat{u} \in C^{\uparrow}(Q;\mathbb{R})$, we deduce that the function G may be chosen so that $G(x) \to \infty$ as $|x| \to \infty$, but such that this convergence to infinity is sufficiently slow to ensure that there exists an open subset E, compactly contained in U, such that

- $K \subset E$, and
- $V \leq \hat{u}$ on $[0,T] \times \partial E$.

Since G is locally Lipschitz continuous and $G(x) \to \infty$ as $|x| \to \infty$, the result of Theorem 2.5.3 also applies to V, and implies that it is also locally Lipschitz continuous. The only addition to the proof of Theorem 2.5.3 in this case is to check that the estimate (2.38) still holds when the trajectories are in $\mathbb{R}^2 \setminus U$; however this is easy, as F is simply given by (2.20) for $x \in \mathbb{R}^2 \setminus U$.

As $V(0,x) = G(x) \leq g(x) \leq \hat{u}(0,x)$ for all $x \in U$, we see that $V \leq \hat{u}$ on the parabolic boundary of $(0,T) \times E$. Since V is Lipschitz continuous on $[0,T] \times \overline{E}$, we may apply Theorem 5.1 in Barles [1] on the subdomain $(0,T) \times E$ with the subsolution V and the supersolution $\hat{u}$, to deduce that $V \leq \hat{u}$ in $(0,T) \times E$.

By (2.59) we have that

$$v = V \leq \hat{u} \quad \text{on} \quad (0,T) \times K.$$

Since K was arbitrary, we have the result. □

Theorem 2.6.2 now follows by combining the results of Corollary 2.6.4 and Proposition 2.6.5.

2.7 A numerical scheme

We briefly outline a numerical scheme to compute the solution of the HJB equation (2.24).

As the 'reference' parameters $\alpha^*, \beta^*, \mu_0^*, \sigma_0^*$ are a priori considered to be the best estimates for the true parameters, we set their penalty to zero. That is, our penalty functions f and g will naturally be equal to zero for this choice of parameters. For an example of the functions f, g, see (2.61)–(2.62).

Thus, recalling the change of variables in (2.14), and the definition of the value function v in (2.23), we note that the transformed reference dynamics, given by

$$w^*(t) = \left(\frac{q_t^*}{R_t^*} - \eta_t,\ \frac{1}{R_t^*} \right),$$

where q^* and R^* are the conditional mean and variance calculated using the reference parameters, traces out the minimum point of the value function. That is, we have that $v(t, w^*(t)) = 0$ for every $t \geq 0$. Since in practice we are only interested in the local behaviour of v around its minimum, we only need to solve the PDE in a finite subdomain centred at this (moving) minimum point.

Defining $\bar{v}(t, \zeta) = v(t, w^*(t) + \zeta)$, the HJB equation (2.24) becomes

$$\frac{\partial \bar{v}}{\partial t}(t, \zeta) + \sup_{(a,b) \in \mathbb{R} \times [0,\infty)} \left\{ \bar{F}(t, \zeta, a, b) \cdot \nabla \bar{v}(t, \zeta) - f(t, a, b) \right\} = 0,$$

where

$$\bar{F}(t, \zeta, a, b) = F(t, w^*(t) + \zeta, a, b) - F(t, w^*(t), \alpha_t^*, \beta_t^*).$$

Since the value function v is unbounded (indeed, it 'blows up' at the boundary in the sense of Definition 2.5.1) we instead solve for the function $\lambda = -1/(1 + \bar{v})$, which takes values in the bounded interval $[-1, 0)$, and satisfies the equation

$$\frac{\partial \lambda}{\partial t}(t, \zeta) + \sup_{(a,b) \in \mathbb{R} \times [0,\infty)} \left\{ \bar{F}(t, \zeta, a, b) \cdot \nabla \lambda(t, \zeta) - \lambda(t, \zeta)^2 f(t, a, b) \right\} = 0, \tag{2.60}$$

with the boundary condition that $\lambda(t, \zeta) \to 0$ as $|\zeta| \to \infty$ and as $\zeta_2 \searrow -w_2^*(t)$. Since the transformation from $\bar{v}$ to λ is monotone increasing, the corresponding viscosity theory carries over.

We approximate the solution of (2.60) with an explicit finite difference scheme on a finite rectangular domain, centred at $\zeta = 0$, on the boundary of which we impose a zero boundary condition.

Remark 2.7.1. Although the function λ is not actually equal to zero on the boundary of a finite subdomain, this procedure does lead to a convergent scheme as we increase the size of the domain. This can be seen by recalling, from the proofs of Propositions 2.6.3 and 2.6.5, the fact that v (and hence λ) only depends on the local values of its initial condition, and thus, on any given inner region, is unaffected by perturbations in any sufficiently large outer region, provided that such perturbations only increase the value of v (or λ).

By the comparison results obtained in Section 2.6, since λ is bounded, we recover the setting of Barles and Souganidis [8]. To establish monotonicity we use an upwind scheme, truncating large values of the drift $\bar{F}$ in order to ensure a suitable Courant–Friedrichs–Lewy condition holds. By Theorem 2.1 in [8], the scheme is uniformly convergent, and by the previous remark, the entire solution λ may be obtained by increasing the size of the domain. The value function v may then be recovered via $v(t,x) = -1 - 1/\lambda(t, x - w^*(t))$.

Remark 2.7.2. We could have made the above transformation earlier and instead shown well-posedness of the PDE (2.60). This leads to a different proof of uniqueness, but one of comparable complexity, and one which does not straightforwardly generalise to the settings we will encounter in later chapters.

2.8 A numerical example

We conclude with a numerical example. Let us recall our original problem of filtering the state of the unobserved process X from observations of Y, where X and Y satisfy the equations

$$\begin{aligned} \mathrm{d}X_t &= \alpha_t X_t \,\mathrm{d}t + \sqrt{\beta_t}\,\mathrm{d}B_t^1, \\ \mathrm{d}Y_t &= c_t X_t \,\mathrm{d}t + \mathrm{d}B_t^2. \end{aligned}$$

We simulate a realisation of the signal process X with 'true' parameters as in the table (where for simplicity α and β are taken to be constant in time), and the observation process Y with $c = 1$. We then consider the problem of estimating X using the 'estimated' parameters as shown.

true parameters	$\alpha = 0.5$	$\beta = 1.5$	$\mu_0 = 1$	$\sigma_0 = 0.2$
estimated parameters	$\alpha^* = 0$	$\beta^* = 1$	$\mu_0^* = 0$	$\sigma_0^* = 1$

We use the penalty functions given by

$$f(t,a,b) = 5(a-\alpha^*)^2 + 10(b-\beta^*)^2, \tag{2.61}$$

$$g(x_1,x_2) = 15(x_1 - x_1^*)^2 + 15(x_2 - x_2^*)^2, \tag{2.62}$$

where $x_1^* = \frac{\mu_0^*}{(\sigma_0^*)^2}$ and $x_2^* = \frac{1}{(\sigma_0^*)^2}$.

In Figure 2.1 we plot the simulated realisation of the process $\eta = \int_0^{\cdot} c_s \,\mathrm{d}Y_s$.

We solve the HJB equation (2.24) for the value function v using the numerical scheme exhibited in Section 2.7. A plot of v at time 0 and at time 1 are shown in Figures 2.2 and 2.3 respectively.

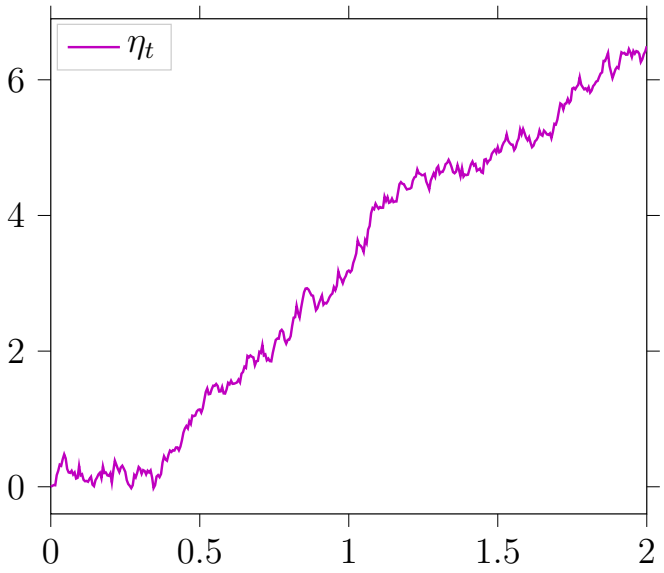


Figure 2.1: Realisation of η.

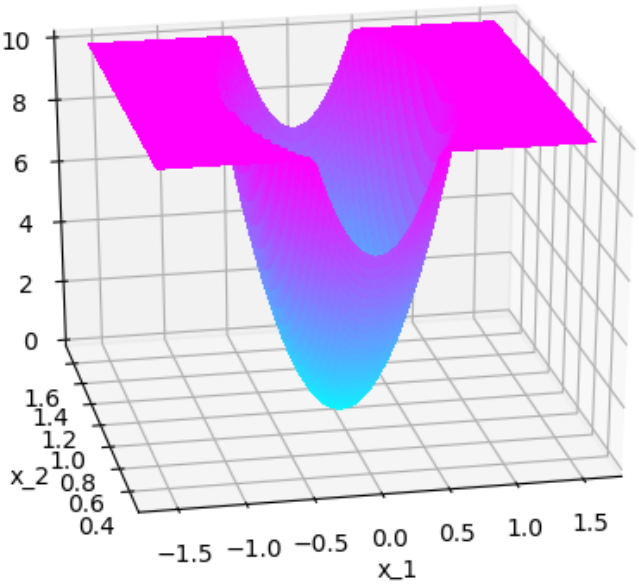


Figure 2.2: Solution of HJB equation at time 0.

Having solved for v, we reverse the change of variables performed in Section 2.3 to obtain the penalty function κ. We are then in a position to calculate the nonlinear expectation via (2.18). In this example we choose $k_1 = 10$ and $k_2 = 5$.

In Figure 2.4a we compare a robust minimax estimate of the signal, calculated as $\arg\min_\xi \mathcal{E}\big((X_t - \xi)^2 \,\big|\, \mathcal{Y}_t\big)$, with the standard Kalman–Bucy filter, itself calculated with both the estimated parameters ($\mathbb{E}^{\text{est}}$) and with the true parameters ($\mathbb{E}^{\text{true}}$). The latter is of course only computable with knowledge of the true parameters, and represents

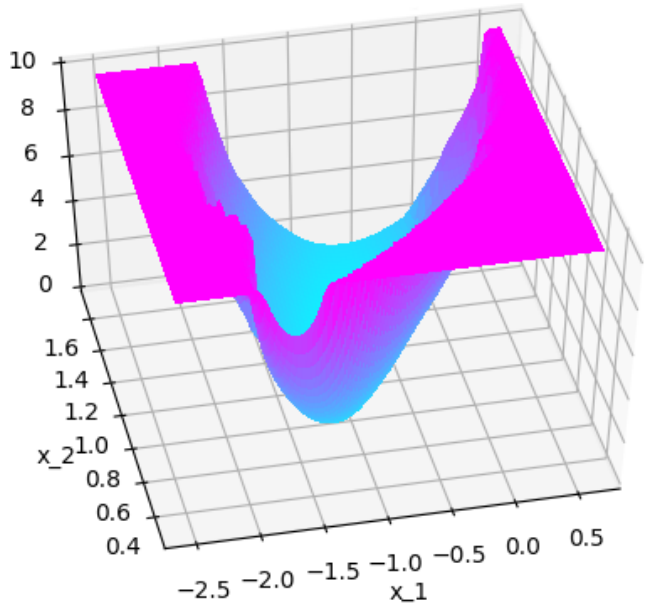
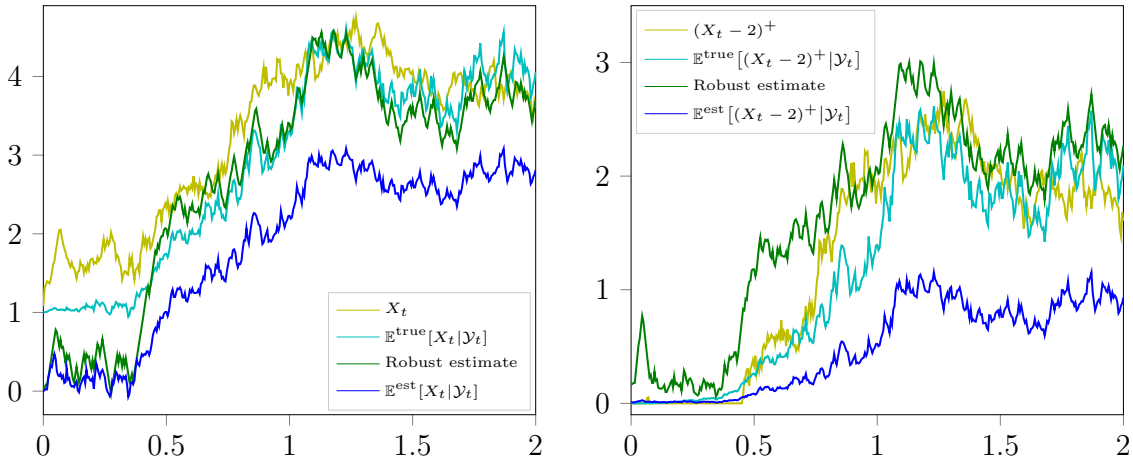


Figure 2.3: Solution of HJB equation at time 1.

the theoretical 'best' estimate one could hope to achieve.



(a) Estimates of X_t.

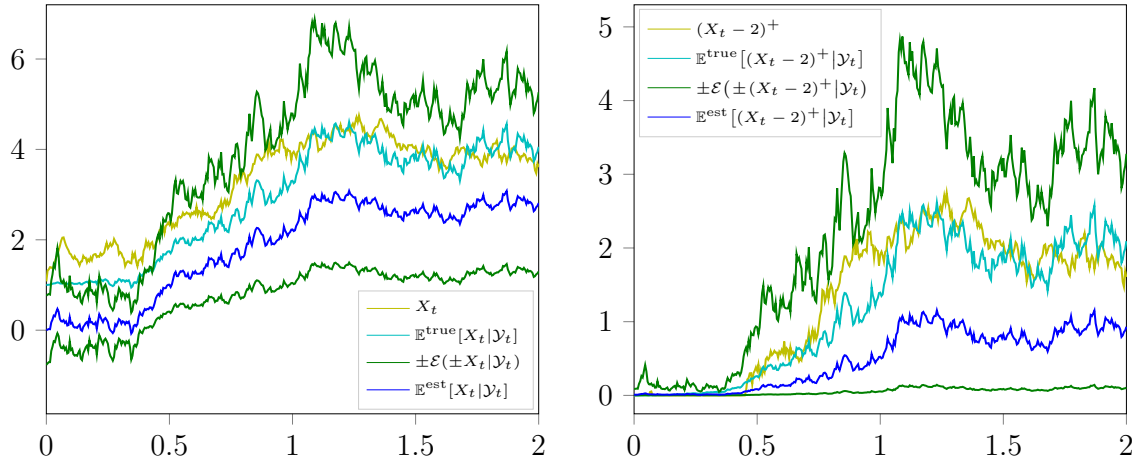
(b) Estimates of $(X_t-2)^+$.

Figure 2.4: Comparison of standard and robust filter estimates.

As mentioned in the introduction, the results of the majority of the existing literature on robust filtering allow for estimation only of the state of the signal, or linear functions thereof. A significant strength of the current approach is that it also enables one to compute estimates of nonlinear functions of the signal. Moreover, this can be done using the same penalty function—that is, we do not have to solve the PDE again in order to perform each subsequent estimate. To illustrate this, in Figure 2.4b we compute a robust estimate of the variable $(X_t-2)^+$, i.e. $\max(X_t-2,0)$,

as $\arg\min_\xi \mathcal{E}\big(\big((X_t-2)^+-\xi\big)^2 \,\big|\, \mathcal{Y}_t\big)$, and compare it with standard filter estimates.

We also calculate the 'upper' and 'lower' expectations of the signal, i.e. $\mathcal{E}(X_t \,|\, \mathcal{Y}_t)$ and $-\mathcal{E}(-X_t \,|\, \mathcal{Y}_t)$ respectively, which can be thought of as providing error bounds for the Kalman–Bucy filter, and are shown in Figure 2.5a. Finally, Figure 2.5b shows the upper and lower expectations of $(X_t-2)^+$. We note in particular the asymmetry between the upper and lower expectations around the Kalman–Bucy filter, in contrast to classical equal-sided confidence intervals.



(a) Expectations of X_t.

(b) Expectations of $(X_t-2)^+$.

Figure 2.5: Robust upper and lower expectations.

Chapter 3

Rough Paths

In Chapter 2 we were led to an optimal control problem in which the optimisation was performed for each fixed realisation of the observation process. In other words, our control problem had to be solved pathwise. In Section 2.3 we introduced a change of variables (similar to the one in Buckdahn and Ma [12]) to essentially eliminate the stochastic integral from our dynamics, which reduced the problem to one of classical (deterministic) optimal control. In more general settings however, such a change of variables will not be available, and we will need to deal directly with the pathwise nature of our problem. We will therefore devote the current chapter to equip ourselves with some tools from rough path theory, which we will then make thorough use of during the remainder of the thesis.

3.1 Motivation

We would like to consider an $\mathbb{R}^m$-valued process X which, for each choice of parameters $\alpha\colon [0,T] \to \mathbb{R}^l$, $\gamma\colon [0,T] \to \mathbb{R}^k$, satisfies an equation of the form

$$\mathrm{d}X_t = \big(a(X_t)\alpha_t + b(X_t, \gamma_t)\big)\,\mathrm{d}t + \phi(X_t, \gamma_t)\,\mathrm{d}Y_t, \qquad t \in [0,T], \tag{3.1}$$

where Y is a continuous (deterministic) $\mathbb{R}^d$-valued path of infinite variation.

Suppose that a, b and ϕ are Lipschitz continuous. In the case when γ is of finite variation and ϕ does not depend on the solution X, the integral against Y then exists in the classical Riemann–Stieltjes sense by integration by parts (see e.g. Theorem 1.2.3 in Stroock [60]). The equation (3.1) then has a unique solution, and moreover the solution map from the driver Y to the corresponding solution X^Y is continuous with respect to the supremum norm.

On the other hand, in the general case when ϕ depends on X, the integral against Y in (3.1) does not even exist in the Riemann–Stieltjes sense. Moreover, even if Y

were smooth, the solution map $Y \mapsto X^Y$ is known to lack continuity, which later would be fatal to the derivation of our HJB equation.

A deterministic approach to integration against very general classes of signals is provided by rough path theory, which moreover allows the continuity property mentioned above to be recovered. As mentioned in the introduction, the key here is, rather than to simply integrate against the path Y, to first enhance Y by a suitable 'second order' process $\mathbb{Y}$, which contains the missing information required to construct the so-called 'rough integral' against the enhanced path $\mathbf{Y} := (Y, \mathbb{Y})$. There are by now a number of monographs on this subject, such as Friz and Hairer [32] and Friz and Victoir [34].

3.2 Notation

We consider a finite time interval $[0, T]$, and write $\Delta_{[0,T]} := \{(s, t) : 0 \leq s \leq t \leq T\}$ for the standard 2-simplex. For any path X on $[0, T]$, we write the increment of X over the interval $[s, t]$ as $X_{s,t} := X_t - X_s$, and write $\|X\|_\infty := \sup_{s\in[0,T]} |X_s|$ for the supremum norm. We also define the following function spaces. We write

- $\mathcal{L}(\mathbb{R}^d; \mathbb{R}^m)$ for the space of linear maps from $\mathbb{R}^d$ to $\mathbb{R}^m$,
- Lip_b for the space of bounded Lipschitz functions, either from $\mathbb{R}^m \to \mathcal{L}(\mathbb{R}^l; \mathbb{R}^m)$ or from $\mathbb{R}^m \times \mathbb{R}^k \to \mathbb{R}^m$, which will always be clear from the context,
- C_b^n ($n \in \mathbb{N}$) for the space of n times continuously differentiable (in the Fréchet sense) functions ϕ, either from $\mathbb{R}^m \times \mathbb{R}^k \to \mathcal{L}(\mathbb{R}^d; \mathbb{R}^m)$ or from $\mathbb{R}^m \times \mathbb{R}^k \to \mathcal{L}(\mathbb{R}^d; \mathbb{R})$ (which will always be clear from the context) such that ϕ and all its derivatives up to order n are uniformly bounded,
- $\mathcal{C}^{p\text{-var}} = \mathcal{C}^{p\text{-var}}([0, T]; \mathbb{R}^d)$ for the space of continuous paths $Y \colon [0, T] \to \mathbb{R}^d$ with finite p-variation:
$$\|Y\|_p := \left(\sup_{\mathcal{P}} \sum_{[s,t]\in\mathcal{P}} |Y_{s,t}|^p \right)^{\frac{1}{p}} < \infty,$$
where the supremum is taken over all finite partitions $\mathcal{P}$ of the interval $[0, T]$.
- $\mathcal{C}^{0,p\text{-var}} = \mathcal{C}^{0,p\text{-var}}([0, T]; \mathbb{R}^k)$ for the closure of smooth paths from $[0, T] \to \mathbb{R}^k$ with respect to the p-variation seminorm.

For $p \in [2,3)$ we write $\mathscr{C}^p = \mathscr{C}^p([0,T];\mathbb{R}^d)$ for the space of continuous p-rough paths, that is, pairs $\mathbf{Y} = (Y, \mathbb{Y})$ such that $Y \in \mathcal{C}^{p\text{-var}}([0,T];\mathbb{R}^d)$ is a continuous path of finite p-variation, its 'enhancement' $\mathbb{Y}\colon \Delta_{[0,T]} \to \mathbb{R}^{d\times d}$ is continuous and satisfies

$$\|\mathbb{Y}\|_{\frac{p}{2}} := \left(\sup_{\mathcal{P}} \sum_{[s,t]\in\mathcal{P}} |\mathbb{Y}_{s,t}|^{\frac{p}{2}} \right)^{\frac{2}{p}} < \infty,$$

and such that Chen's relation,

$$\mathbb{Y}_{s,t} = \mathbb{Y}_{s,r} + \mathbb{Y}_{r,t} + Y_{s,r} \otimes Y_{r,t}, \tag{3.2}$$

holds for all times $s \le r \le t$, where $\otimes$ denotes the standard tensor product from $\mathbb{R}^d \times \mathbb{R}^d$ to $\mathbb{R}^d \otimes \mathbb{R}^d \subset \mathbb{R}^{d\times d}$.

Remark 3.2.1. As can be readily checked, any *smooth* path $Y\colon [0,T] \to \mathbb{R}^d$ can be 'lifted' in a canonical way to a rough path $\mathbf{Y} = (Y, \mathbb{Y})$ by enhancing it with the integral

$$\mathbb{Y}_{s,t} = \int_s^t Y_{s,r} \otimes \mathrm{d}Y_r. \tag{3.3}$$

On the other hand, for a general path Y of finite p-variation, the integral in (3.3) does not exist in the classical sense. The point here then is that the value of this integral is postulated by the enhancement $\mathbb{Y}$, which in practice is often constructed using stochastic integration.

For example, given a continuous semimartingale Y, for $p \in (2,3)$ one can construct an enhancement via Stratonovich integration:

$$\mathbb{Y}_{s,t} = \int_s^t Y_{s,r} \otimes \circ\, \mathrm{d}Y_r, \tag{3.4}$$

that is, $\mathbb{Y}^{ij}_{s,t} = \int_s^t Y^i_{s,r} \circ \mathrm{d}Y^j_r$ for $i,j = 1,\dots,d$, and the resulting lift $\mathbf{Y} = (Y,\mathbb{Y})$ then defines a random rough path, so that $\mathbf{Y}(\omega) \in \mathscr{C}^p$ for almost every $\omega \in \Omega$. Of course, when $i = j$ we have simply $\mathbb{Y}^{ii}_{s,t} = \int_s^t Y^i_{s,r} \circ \mathrm{d}Y^i_r = \frac{1}{2}(Y^i_{s,t})^2$. More generally,

$$\mathbb{Y}^{ij}_{s,t} + \mathbb{Y}^{ji}_{s,t} = \int_s^t Y^i_{s,r} \circ \mathrm{d}Y^j_r + \int_s^t Y^j_{s,r} \circ \mathrm{d}Y^i_r = Y^i_{s,t} Y^j_{s,t}.$$

Thus, the additional information encoded by this lift is contained in the antisymmetric part of $\mathbb{Y}$, which corresponds to the Lévy area of the process Y.

For rough paths $\mathbf{Y} = (Y,\mathbb{Y})$ and $\tilde{\mathbf{Y}} = (\tilde{Y},\tilde{\mathbb{Y}})$, we write

$$|||\mathbf{Y}|||_p := \|Y\|_p + \|\mathbb{Y}\|_{\frac{p}{2}}$$

for the (inhomogeneous) rough path norm[1], and

$$\|\mathbf{Y};\tilde{\mathbf{Y}}\|_p := \|Y-\tilde{Y}\|_p + \|\mathbb{Y}-\tilde{\mathbb{Y}}\|_{\frac{p}{2}} \tag{3.5}$$

for the (inhomogeneous) rough path distance.

We will also consider the space of geometric rough paths $\mathscr{C}_g^{0,p} \subset \mathscr{C}^p$, defined as the closure of canonical lifts of smooth paths with respect to the pseudometric in (3.5). For example, when Y is a semimartingale and we lift using Stratonovich integration, as in (3.4), the resulting lift turns out to be a (random) geometric rough path. This property of being well approximated by smooth paths allows one to make sense of solutions to a wide class of rough ODEs and PDEs—we will see an example of this in Definition 4.8.1 below.

We will sometimes write e.g. $\|Y\|_{p;[s,t]}$ for the p-variation of Y over the subinterval $[s,t]$.

Lemma 3.2.2. *For some $n \geq 1$, let $0 = t_0 < t_1 < \ldots < t_{n-1} < t_n = T$, be a partition of the interval $[0,T]$. Then, for any path X, one has that*

$$\|X\|_{p;[0,T]} \leq n^{\frac{p-1}{p}} \left(\sum_{i=1}^{n} \|X\|^p_{p;[t_{i-1},t_i]} \right)^{\frac{1}{p}}.$$

Proof. Let $0 = s_0 < s_1 < \ldots < s_{N-1} < s_N = T$ be another partition of the interval $[0,T]$. We can label the union of these two partitions in two different ways as follows. We can either write

$$\begin{aligned} & s_{j-1} = t^j_0 < t^j_1 < \ldots < t^j_{n_j} = s_j \quad \text{for each} \quad j = 1,\ldots,N, \\ \text{or} \quad & t_{i-1} = s^i_0 < s^i_1 < \ldots < s^i_{N_i} = t_i \quad \text{for each} \quad i = 1,\ldots,n, \end{aligned}$$

where, crucially, $n_j \leq n$ for every j. For every $j = 1,\ldots,N$, we have, by Hölder's inequality,

$$\sum_{i=1}^{n_j} \left| X_{t^j_i} - X_{t^j_{i-1}} \right| \leq n_j^{\frac{p-1}{p}} \left(\sum_{i=1}^{n_j} \left| X_{t^j_i} - X_{t^j_{i-1}} \right|^p \right)^{\frac{1}{p}} \leq n^{\frac{p-1}{p}} \left(\sum_{i=1}^{n_j} \left| X_{t^j_i} - X_{t^j_{i-1}} \right|^p \right)^{\frac{1}{p}}.$$

Then

$$\begin{aligned} \sum_{j=1}^{N} \left| X_{s_j} - X_{s_{j-1}} \right|^p &\leq \sum_{j=1}^{N} \left(\sum_{i=1}^{n_j} \left| X_{t^j_i} - X_{t^j_{i-1}} \right| \right)^p \leq n^{p-1} \sum_{j=1}^{N} \sum_{i=1}^{n_j} \left| X_{t^j_i} - X_{t^j_{i-1}} \right|^p \\ &= n^{p-1} \sum_{i=1}^{n} \sum_{j=1}^{N_i} \left| X_{s^i_j} - X_{s^i_{j-1}} \right|^p \leq n^{p-1} \sum_{i=1}^{n} \|X\|^p_{p;[t_{i-1},t_i]}. \end{aligned}$$

The result then follows from taking the supremum over all possible partitions $s_0 < s_1 < \ldots < s_N$ of the interval $[0,T]$. □

[1]Of course, $\mathscr{C}^p$ is not actually a vector space due to the nonlinear structure of (3.2).

3.3 Rough integration

As noted above, the enhancement $\mathbb{Y}$ may be seen as postulating the value of the integrals $\mathbb{Y}^{ij}_{s,t} = \int_s^t Y^i_{s,r}\,\mathrm{d}Y^j_r$ for $i,j = 1,\dots,d$. More generally, $\mathbf{Y}^j = (Y^j, (\mathbb{Y}^{ij})_{i=1}^d)$ contains all the information required to define integrals against Y^j, for any integrand in the space of paths 'controlled' by $(Y^1,\dots,Y^d)$ in the following sense.

Let $\mathbf{Y} = (Y,\mathbb{Y})$ be a rough path for some $p \in [2,3)$. We say that a path $X \in \mathcal{C}^{p\text{-var}}([0,T];\mathbb{R}^m)$ is a controlled rough path (in the sense of Gubinelli [39]), if there exists a path $X' \in \mathcal{C}^{p\text{-var}}([0,T];\mathcal{L}(\mathbb{R}^d;\mathbb{R}^m))$, known as the Gubinelli derivative of X with respect to Y, such that the remainder term $R^X\colon \Delta_{[0,T]} \to \mathbb{R}^m$, defined implicitly by

$$X_{s,t} = X'_s Y_{s,t} + R^X_{s,t},$$

satisfies $\|R^X\|_{\frac{p}{2}} < \infty$. We write $\mathscr{D}^p_Y = \mathscr{D}^p_Y([0,T];\mathbb{R}^m)$ for the space of controlled rough paths, which becomes a Banach space when equipped with the norm $(X,X') \mapsto |X_0| + |X'_0| + \|X'\|_p + \|R^X\|_{\frac{p}{2}}$.

Given any controlled path $(X,X') \in \mathscr{D}^p_Y$, parameter $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}$ and function $\phi \in C^2_b$, the composition $\phi(X,\gamma)$ is then itself a controlled path with Gubinelli derivative $\phi(X,\gamma)' = D_x\phi(X,\gamma)X'$, where $D_x\phi$ denotes the Fréchet derivative of ϕ in its first argument.

Proposition 3.3.1 (Proposition 2.6 in [35])**.** *Let $\mathbf{Y} = (Y,\mathbb{Y}) \in \mathscr{C}^p([0,T];\mathbb{R}^d)$, and let $(X,X') \in \mathscr{D}^p_Y([0,T];\mathcal{L}(\mathbb{R}^d;\mathbb{R}^m))$ be a controlled rough path. Then the limit*

$$\int_0^T X_r\,\mathrm{d}\mathbf{Y}_r := \lim_{|\mathcal{P}|\to 0} \sum_{[s,t]\in\mathcal{P}} X_s Y_{s,t} + X'_s \mathbb{Y}_{s,t}$$

exists[2]*, where the limit is taken over any sequence of partitions $\mathcal{P}$ of the interval $[0,T]$ such that the mesh size $|\mathcal{P}| \to 0$. This limit (which does not depend on the choice of sequence of partitions) is called the rough integral of X against $\mathbf{Y}$.*

Moreover, for any $0 \le s < t \le T$, we have the estimate

$$\left| \int_s^t X_r\,\mathrm{d}\mathbf{Y}_r - X_s Y_{s,t} - X'_s \mathbb{Y}_{s,t} \right| \le C_p\Big(\|R^X\|_{\frac{p}{2};[s,t]} \|Y\|_{p;[s,t]} + \|X'\|_{p;[s,t]} \|\mathbb{Y}\|_{\frac{p}{2};[s,t]} \Big) \quad (3.6)$$

where the constant C_p depends only on p.

[2]Strictly speaking, in making precise sense of the product $X'_s\mathbb{Y}_{s,t}$, we use the natural identification of $\mathcal{L}(\mathbb{R}^d;\mathcal{L}(\mathbb{R}^d;\mathbb{R}^m))$ with $\mathcal{L}(\mathbb{R}^d \otimes \mathbb{R}^d;\mathbb{R}^m)$.

3.4 Rough differential equations

We consider the rough differential equation (RDE) given by

$$\mathrm{d}X_t = \big(a(X_t)\alpha_t + b(X_t, \gamma_t)\big)\,\mathrm{d}t + \phi(X_t, \gamma_t)\,\mathrm{d}\mathbf{Y}_t \tag{3.7}$$

driven by a rough path $\mathbf{Y}$, where $\alpha \in L^1([0,T];\mathbb{R}^l)$ and $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}([0,T];\mathbb{R}^k)$. In particular, we consider a drift which is linear in the parameter α, on which we assume no regularity other than integrability.

Here we exhibit a direct approach to obtaining solutions in the controlled path setting of Gubinelli, based on the fixed point argument of Friz and Zhang [35].

Lemma 3.4.1. *Let $\phi \in C_b^3$, $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}$, and let $\mathbf{Y} = (Y, \mathbb{Y}) \in \mathscr{C}^p$ with $|||\mathbf{Y}|||_p \le L$ for some $L > 0$. For any controlled path $(X, X') \in \mathscr{D}_Y^p$, we have that*

$$\left(\int_0^{\cdot} \phi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r,\ \phi(X,\gamma)\right) \in \mathscr{D}_Y^p,$$

and the estimates

$$\begin{aligned}
\|\phi(X,\gamma)\|_p &\le C\Big(\big(|X_0'| + \|X'\|_p\big)\|Y\|_p + \|R^X\|_{\frac{p}{2}} + \|\gamma\|_{\frac{p}{2}}\Big),\\
\big\|R^{\int_0^{\cdot}\phi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r}\big\|_{\frac{p}{2}} &\le C\Big(1 + |X_0'| + \|X'\|_p + \|R^X\|_{\frac{p}{2}} + \|\gamma\|_{\frac{p}{2}}\Big)^2 |||\mathbf{Y}|||_p,
\end{aligned}$$

where the constant C depends only on ϕ, p and L.

Proof. From Lemma 3.6 in [35], we immediately have

$$\begin{aligned}
\|\phi(X,\gamma)\|_p &\le C\Big(\big(|X_0'| + |\gamma_0'| + \|X'\|_p + \|\gamma'\|_p\big)\|Y\|_p + \|R^X\|_{\frac{p}{2}} + \|R^\gamma\|_{\frac{p}{2}}\Big),\\
\big\|R^{\int_0^{\cdot}\phi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r}\big\|_{\frac{p}{2}} &\le C\Big(1 + |X_0'| + |\gamma_0'| + \|X'\|_p + \|\gamma'\|_p + \|R^X\|_{\frac{p}{2}} + \|R^\gamma\|_{\frac{p}{2}}\Big)^2 |||\mathbf{Y}|||_p,
\end{aligned}$$

for some constant C depending only on ϕ, p and $\|Y\|_p$. Since $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}$, it is immediately controlled by Y, with $\gamma_s' = 0$ and $R^\gamma_{s,t} = \gamma_{s,t}$. Substituting into the above, we obtain the desired estimates. □

The following lemma follows similarly from [35, Lemma 3.7].

Lemma 3.4.2. *Let $\phi \in C_b^3$, $\gamma, \tilde\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}$, and let $\mathbf{Y} = (Y, \mathbb{Y})$ and $\tilde{\mathbf{Y}} = (\tilde Y, \tilde{\mathbb{Y}})$ be rough paths with $|||\mathbf{Y}|||_p, |||\tilde{\mathbf{Y}}|||_p \le L$ for some $L > 0$. Suppose that $(X, X') \in \mathscr{D}_Y^p$ and $(\tilde X, \tilde X') \in \mathscr{D}_{\tilde Y}^p$ satisfy $|X_0'| + \|X'\|_p + \|R^X\|_p + \|\gamma\|_{\frac{p}{2}} \le M$ and $|\tilde X_0'| + \|\tilde X'\|_p + \|R^{\tilde X}\|_p + \|\tilde\gamma\|_{\frac{p}{2}} \le M$ for some $M > 0$. Then we have the estimates*

$$\begin{aligned}
\big\|\phi(X,\gamma) - \phi(\tilde X, \tilde\gamma)\big\|_p \le C\Big(&|X_0 - \tilde X_0| + |X_0' - \tilde X_0'| + \|X' - \tilde X'\|_p\|Y\|_p\\
&+ \|R^X - R^{\tilde X}\|_{\frac{p}{2}} + \|\gamma - \tilde\gamma\|_\infty + \|\gamma - \tilde\gamma\|_{\frac{p}{2}} + \|Y - \tilde Y\|_p\Big),
\end{aligned}$$

$$\begin{aligned}
&\big\|R^{\int_0^\cdot \phi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r} - R^{\int_0^\cdot \phi(\tilde{X}_r,\tilde{\gamma}_r)\,\mathrm{d}\tilde{\mathbf{Y}}_r}\big\|_{\frac{p}{2}} \\
&\quad\le C\Big(\|\mathbf{Y};\tilde{\mathbf{Y}}\|_p + \Big(|X_0-\tilde{X}_0| + |X_0'-\tilde{X}_0'| + \|X'-\tilde{X}'\|_p \\
&\qquad\qquad + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|\gamma-\tilde{\gamma}\|_\infty + \|\gamma-\tilde{\gamma}\|_{\frac{p}{2}}\Big)\vertiii{\mathbf{Y}}_p\Big),
\end{aligned}$$

where the constant C depends on ϕ, p, L and M.

Theorem 3.4.3. *Let $a, b \in \mathrm{Lip}_b$ and $\phi, \psi \in C_b^3$. Let $p \in [2,3)$ and $T > 0$.*

(i) For any rough path $\mathbf{Y} = (Y, \mathbb{Y}) \in \mathscr{C}^p$, parameters $\alpha \in L^1$, $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}$ and any $x \in \mathbb{R}^m$, there exists a unique $X \in \mathcal{C}^{p\text{-var}}$ such that the controlled path $(X, X') = (X, \phi(X,\gamma)) \in \mathscr{D}_Y^p$ solves the RDE (3.7) driven by $\mathbf{Y}$ with parameters α, γ and initial condition $X_0 = x$.

(ii) Suppose that $\vertiii{\mathbf{Y}}_{p;[0,T]} \le L$ for some $L > 0$. We have the bound

$$\left\|\int_0^\cdot \psi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r\right\|_{p;[0,T]} \le C\left(1 + \|\gamma\|_{\frac{p}{2};[0,T]}^{\frac{p}{2}} + \int_0^T \big(1+|\alpha_r|\big)\,\mathrm{d}r\right)^{\frac{p-1}{p}} \vertiii{\mathbf{Y}}_{p;[0,T]}, \tag{3.8}$$

where the constant C depends on a, b, ϕ, ψ, p and L.

(iii) Suppose that $\tilde{X}$ is the solution of (3.7) driven by $\tilde{\mathbf{Y}}$ with parameters $\tilde{\alpha}, \tilde{\gamma}$, and suppose that $\vertiii{\mathbf{Y}}_{p;[0,T]}, \vertiii{\tilde{\mathbf{Y}}}_{p;[0,T]} \le L$. Assume also that

$$\|\gamma\|_{\frac{p}{2}}, \quad \|\tilde{\gamma}\|_{\frac{p}{2}}, \quad \int_0^T (1+|\alpha_r|)\,\mathrm{d}r \quad \text{and} \quad \int_0^T (1+|\tilde{\alpha}_r|)\,\mathrm{d}r$$

are all bounded by M, for some $M > 0$. Then we have the estimates

$$\begin{aligned}
\|X - \tilde{X}\|_{p;[0,T]} \le C'\Big(&|X_0 - \tilde{X}_0| + \|\mathbf{Y};\tilde{\mathbf{Y}}\|_{p;[0,T]} \\
&+ \|\gamma-\tilde{\gamma}\|_\infty + \|\gamma-\tilde{\gamma}\|_{\frac{p}{2};[0,T]} + \int_0^T |\alpha_r - \tilde{\alpha}_r|\,\mathrm{d}r\Big),
\end{aligned} \tag{3.9}$$

$$\begin{aligned}
&\left\|\int_0^\cdot \psi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r - \int_0^\cdot \psi(\tilde{X}_r,\tilde{\gamma}_r)\,\mathrm{d}\tilde{\mathbf{Y}}_r\right\|_{p;[0,T]} \\
&\quad\le C''\Big(|X_0 - \tilde{X}_0| + \|\mathbf{Y};\tilde{\mathbf{Y}}\|_{p;[0,T]} \\
&\qquad + \|\gamma-\tilde{\gamma}\|_\infty + \|\gamma-\tilde{\gamma}\|_{\frac{p}{2};[0,T]} + \int_0^T |\alpha_r - \tilde{\alpha}_r|\,\mathrm{d}r\Big),
\end{aligned} \tag{3.10}$$

where the constants C', C'' depend on a, b, ϕ, p, L and M, and C'' also depends on ψ.

(iv) *Let $p \in (2,3)$. Suppose that Y is a continuous semimartingale and $\mathbf{Y} = (Y, \mathbb{Y})$ is its Stratonovich lift, as in* (3.4)*, so that $\mathbf{Y}(\omega) = (Y(\omega), \mathbb{Y}(\omega)) \in \mathscr{C}_g^{0,p}$ for almost every $\omega \in \Omega$. Then the solution of the random RDE* (3.7) *is indistinguishable from the solution of the Stratonovich SDE*

$$\mathrm{d}X_t = \big(a(X_t)\alpha_t + b(X_t, \gamma_t)\big)\,\mathrm{d}t + \phi(X_t, \gamma_t) \circ \mathrm{d}Y_t,$$

and moreover the rough and Stratonovich integrals coincide almost surely, that is,

$$\int_0^t \psi(X_r(\omega), \gamma_r)\,\mathrm{d}\mathbf{Y}_r(\omega) = \bigg(\int_0^t \psi(X_r, \gamma_r) \circ \mathrm{d}Y_r\bigg)(\omega)$$

for almost every $\omega \in \Omega$.

Remark 3.4.4. By setting $\hat{Y}_t = Y_{T-t}$ and $\hat{\mathbb{Y}}_{s,t} = -\mathbb{Y}_{T-t,T-s} + Y_{T-t,T-s} \otimes Y_{T-t,T-s}$, we obtain a rough path $\hat{\mathbf{Y}} = (\hat{Y}, \hat{\mathbb{Y}}) \in \mathscr{C}^p$ which is in fact the time reversal of the original path $\mathbf{Y}$. Letting $\hat{\alpha}_t = \alpha_{T-t}$ and $\hat{\gamma}_t = \gamma_{T-t}$, the solution $\hat{X}$ of the RDE (3.7) (with $\mathrm{d}t$ replaced by $-\mathrm{d}t$) driven by $\hat{\mathbf{Y}}$ with parameters $\hat{\alpha}$, $\hat{\gamma}$ and initial condition $\hat{X}_0 = X_T$, is then simply the time reversal of the solution X of the original RDE. By applying part (iii) of Theorem 3.4.3 to the solution $\hat{X}$, it follows that we may rewrite the estimates (3.9)–(3.10) in terms of the difference of the terminal values $|X_T - \tilde{X}_T|$, instead of the initial values $|X_0 - \tilde{X}_0|$.

Proof. **Step 1.** We define the map $\mathcal{M}_t \colon \mathscr{D}_Y^p([0,t];\mathbb{R}^m) \to \mathscr{D}_Y^p([0,t];\mathbb{R}^m)$ by

$$\mathcal{M}_t(X, X') = \bigg(x + \int_0^{\cdot} \big(a(X_r)\alpha_r + b(X_r, \gamma_r)\big)\,\mathrm{d}r + \int_0^{\cdot} \phi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r,\ \phi(X, \gamma)\bigg),$$

and, for $\delta \geq 1$, introduce the subset of controlled paths

$$\mathcal{B}_t^{(\delta)} := \Big\{(X, X') \in \mathscr{D}_Y^p([0,t];\mathbb{R}^m) : (X_0, X_0') = (x, \phi(x, \gamma_0)),\ \|X, X'\|_{Y,p}^{(\delta)} \leq 1\Big\},$$

where

$$\|X, X'\|_{Y,p}^{(\delta)} := \|X'\|_p + \delta \|R^X\|_{\frac{p}{2}}.$$

Let $L > 0$ such that $\|\!|\!|\mathbf{Y}|\!|\!|_p \leq L$. Provided that

$$\|\gamma\|_{\frac{p}{2};[0,t]} \leq 1, \tag{3.11}$$

by Lemma 3.4.1, for any $(X, X') \in \mathcal{B}_t^{(\delta)}$, we have

$$\begin{aligned}
&\|\mathcal{M}_t(X, X')\|_{Y,p}^{(\delta)} \\
&\quad \leq \|\phi(X, \gamma)\|_p + \delta \big\| R^{\int_0^{\cdot} \phi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r} \big\|_{\frac{p}{2}} + \delta \bigg\| \int_0^{\cdot} \big(a(X_r)\alpha_r + b(X_r, \gamma_r)\big)\,\mathrm{d}r \bigg\|_{\frac{p}{2}} \\
&\quad \leq C_1 \bigg(\frac{1}{\delta} + \delta \bigg(|\!|\!|\mathbf{Y}|\!|\!|_p + \|\gamma\|_{\frac{p}{2}} + \int_0^t \big(1 + |\alpha_r|\big)\,\mathrm{d}r\bigg)\bigg)
\end{aligned}$$

for some constant $C_1 \geq \frac{1}{2}$ depending only on a, b, ϕ, p and L. Let $\delta = \delta_1 := 2C_1 \geq 1$, so that

$$\|\mathcal{M}_t(X, X')\|_{Y,p}^{(\delta_1)} \leq \frac{1}{2} + 2C_1^2 \bigg(|||\mathbf{Y}|||_{p;[0,t]} + \|\gamma\|_{\frac{p}{2};[0,t]} + \int_0^t \big(1 + |\alpha_r|\big)\, \mathrm{d}r \bigg).$$

By taking $t = t_1$ sufficiently small such that

$$|||\mathbf{Y}|||_{p;[0,t_1]} + \|\gamma\|_{\frac{p}{2};[0,t_1]} + \int_0^{t_1} \big(1 + |\alpha_r|\big)\, \mathrm{d}r \leq \frac{1}{4C_1^2}, \tag{3.12}$$

which in particular ensures that (3.11) holds, we obtain $\|\mathcal{M}_{t_1}(X, X')\|_{Y,p}^{(\delta_1)} \leq 1$. That is, $\mathcal{B}_{t_1}^{(\delta_1)}$ is invariant under $\mathcal{M}_{t_1}$.

Step 2. Let $(X, X'), (\tilde{X}, \tilde{X}') \in \mathcal{B}_t^{(\delta_1)}$ with δ_1 as above and some $t \leq t_1$. Note that $|X_0'| + \|X'\|_p + \|R^X\|_{\frac{p}{2}} + \|\gamma\|_{\frac{p}{2}} \leq |\phi(x, \gamma_0)| + 2$ and similarly for $\tilde{X}$, so that the hypotheses of Lemma 3.4.2 are satisfied. For any (new) $\delta \geq 1$, we then have

$$\begin{aligned}
&\big\|\mathcal{M}_t(X, X') - \mathcal{M}_t(\tilde{X}, \tilde{X}')\big\|_{Y,p}^{(\delta)} \\
&\quad \leq \big\|\phi(X, \gamma) - \phi(\tilde{X}, \gamma)\big\|_p + \delta \big\| R^{\int_0^\cdot \phi(X_r, \gamma_r)\, \mathrm{d}\mathbf{Y}_r} - R^{\int_0^\cdot \phi(\tilde{X}_r, \gamma_r)\, \mathrm{d}\mathbf{Y}_r} \big\|_{\frac{p}{2}} \\
&\qquad + \delta \bigg\| \int_0^\cdot \big(a(X_r)\alpha_r + b(X_r, \gamma_r)\big)\, \mathrm{d}r - \int_0^\cdot \big(a(\tilde{X}_r)\alpha_r + b(\tilde{X}_r, \gamma_r)\big)\, \mathrm{d}r \bigg\|_{\frac{p}{2}} \\
&\quad \leq C_2 \bigg(\|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\
&\qquad\qquad + \delta \Big(\|X' - \tilde{X}'\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \Big) \bigg(|||\mathbf{Y}|||_p + \int_0^t \big(1 + |\alpha_r|\big)\, \mathrm{d}r \bigg) \bigg),
\end{aligned}$$

for some constant $C_2 > \frac{1}{2}$ depending only on a, b, ϕ, p and L, where in bounding the drift terms we used the fact[3] that $\|X - \tilde{X}\|_\infty \leq \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}}$. Let $\delta = \delta_2 := 2C_2 > 1$, so that

$$\begin{aligned}
\big\|\mathcal{M}_t(X, X') - \mathcal{M}_t(\tilde{X}, \tilde{X}')\big\|_{Y,p}^{(\delta_2)} &\leq \frac{\delta_2}{2} \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\
&\quad + 2C_2^2 \Big(\|X' - \tilde{X}'\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \Big) \bigg(|||\mathbf{Y}|||_{p;[0,t]} + \int_0^t \big(1 + |\alpha_r|\big)\, \mathrm{d}r \bigg).
\end{aligned}$$

Taking $t = t_2 \leq t_1$ sufficiently small such that

$$|||\mathbf{Y}|||_{p;[0,t_2]} + \int_0^{t_2} \big(1 + |\alpha_r|\big)\, \mathrm{d}r \leq \frac{1}{4C_2^2}, \tag{3.13}$$

[3]This follows by taking the difference of the equalities $X_{0,t} = X_0' Y_{0,t} + R_{0,t}^X$ and $\tilde{X}_{0,t} = \tilde{X}_0' Y_{0,t} + R_{0,t}^{\tilde{X}}$.

we obtain

$$\begin{aligned}\big\|\mathcal{M}_{t_2}(X,X') - \mathcal{M}_{t_2}(\tilde{X},\tilde{X}')\big\|_{Y,p}^{(\delta_2)} &\leq \frac{1}{2}\|X' - \tilde{X}'\|_p + \frac{\delta_2+1}{2}\|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\ &\leq \frac{\delta_2+1}{2\delta_2}\big\|(X,X') - (\tilde{X},\tilde{X}')\big\|_{Y,p}^{(\delta_2)},\end{aligned}$$

and we have thus established that $\mathcal{M}_{t_2}$ is a contraction on $\mathcal{B}_{t_2}^{(\delta_1)}$. The fixed point of this map is then the unique solution of the RDE (3.7) with initial condition $X_0 = x$ over the time interval $[0, t_2]$ satisfying $X' = \phi(X, \gamma)$.

Step 3. Let $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{\mathbb{Y}}) \in \mathscr{C}^p$ be another rough path with $|||\tilde{\mathbf{Y}}|||_p \leq L$, let $\tilde{\alpha} \in L^1$, $\tilde{\gamma} \in \mathcal{C}^{\frac{p}{2}\text{-var}}$ and $\tilde{x} \in \mathbb{R}^m$. Let $(X, X') = (X, \phi(X,\gamma)) \in \mathscr{D}_Y^p$ (resp. $(\tilde{X}, \tilde{X}') = (\tilde{X}, \phi(\tilde{X}, \tilde{\gamma})) \in \mathscr{D}_{\tilde{Y}}^p$) be the unique solution of the RDE (3.7) driven by $\mathbf{Y}$ (resp. $\tilde{\mathbf{Y}}$) with parameters α (resp. $\tilde{\alpha}$) and γ (resp. $\tilde{\gamma}$) and initial value x (resp. $\tilde{x}$), which exists over the time interval $[0, t_2]$ (resp. $[0, \tilde{t}_2]$). In the following we denote by $\lesssim$ inequality up to a multiplicative constant which may depend on a, b, ϕ, p and L.

From the equality $X_{s,t} - \tilde{X}_{s,t} = X'_s Y_{s,t} - \tilde{X}'_s \tilde{Y}_{s,t} + R^X_{s,t} - R^{\tilde{X}}_{s,t}$, we deduce

$$\begin{aligned}\|X - \tilde{X}\|_p &\leq \|X' - \tilde{X}'\|_\infty \|Y\|_p + \|\phi(\tilde{X},\tilde{\gamma})\|_\infty \|Y - \tilde{Y}\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\ &\lesssim |X_0 - \tilde{X}_0| + \|X' - \tilde{X}'\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|Y - \tilde{Y}\|_p + \|\gamma - \tilde{\gamma}\|_\infty.\end{aligned} \tag{3.14}$$

By Lemma 3.4.2, for any (new) $\delta \geq 1$, we have

$$\begin{aligned}&\|X' - \tilde{X}'\|_p + \delta\|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\ &\leq \big\|\phi(X,\gamma) - \phi(\tilde{X},\tilde{\gamma})\big\|_p + \delta\big\|R^{\int_0^\cdot \phi(X_r,\gamma_r)\,\mathrm{d}\mathbf{Y}_r} - R^{\int_0^\cdot \phi(\tilde{X}_r,\tilde{\gamma}_r)\,\mathrm{d}\tilde{\mathbf{Y}}_r}\big\|_{\frac{p}{2}} \\ &\quad + \delta\bigg\|\int_0^\cdot \big(a(X_r)\alpha_r + b(X_r,\gamma_r)\big)\,\mathrm{d}r - \int_0^\cdot \big(a(\tilde{X}_r)\tilde{\alpha}_r + b(\tilde{X}_r,\tilde{\gamma}_r)\big)\,\mathrm{d}r\bigg\|_{\frac{p}{2}} \\ &\lesssim |X_0 - \tilde{X}_0| + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2}} \\ &\quad + \delta\bigg(\Big(|X_0 - \tilde{X}_0| + \|X' - \tilde{X}'\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2}}\Big)|||\mathbf{Y}|||_p \\ &\qquad + \|\mathbf{Y};\tilde{\mathbf{Y}}\|_p + \big(\|X - \tilde{X}\|_\infty + \|\gamma - \tilde{\gamma}\|_\infty\big)\int_0^t \big(1 + |\alpha_r|\big)\,\mathrm{d}r + \int_0^t |\alpha_r - \tilde{\alpha}_r|\,\mathrm{d}r\bigg).\end{aligned}$$

Writing $\|X - \tilde{X}\|_\infty \le |X_0 - \tilde{X}_0| + \|X - \tilde{X}\|_p$ and using (3.14), we then have

$$\begin{aligned}
&\|X' - \tilde{X}'\|_p + \delta \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} \\
&\quad \le C_3 \Bigg(|X_0 - \tilde{X}_0| + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2}} \\
&\qquad + \delta \Bigg(\|\mathbf{Y}; \tilde{\mathbf{Y}}\|_p + \int_0^t |\alpha_r - \tilde{\alpha}_r| \, \mathrm{d}r \Bigg) \\
&\qquad + \delta \Bigg(|X_0 - \tilde{X}_0| + \|X' - \tilde{X}'\|_p + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2}} + \|Y - \tilde{Y}\|_p \\
&\qquad\quad + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2}} \Bigg) \Bigg(|||\mathbf{Y}|||_{p;[0,t]} + \int_0^t \big(1 + |\alpha_r|\big) \, \mathrm{d}r \Bigg) \Bigg)
\end{aligned}$$

for some constant $C_3 > \frac{1}{2}$ depending only on a, b, ϕ, p and L. Let $\delta = \delta_3 := 2C_3 > 1$. Taking $t = t_3 \le t_2 \wedge \tilde{t}_2$ sufficiently small such that

$$|||\mathbf{Y}|||_{p;[0,t_3]} + \int_0^{t_3} \big(1 + |\alpha_r|\big) \, \mathrm{d}r \le \frac{1}{4C_3^2} \tag{3.15}$$

and rearranging, we obtain

$$\begin{aligned}
&\|X' - \tilde{X}'\|_{p;[0,t_3]} + (\delta_3 - 1) \|R^X - R^{\tilde{X}}\|_{\frac{p}{2};[0,t_3]} \\
&\quad \le (\delta_3 + 1) \Big(|X_0 - \tilde{X}_0| + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2};[0,t_3]} \Big) \\
&\qquad + (\delta_3^2 + 1) \Bigg(\|\mathbf{Y}; \tilde{\mathbf{Y}}\|_{p;[0,t_3]} + \int_0^{t_3} |\alpha_r - \tilde{\alpha}_r| \, \mathrm{d}r \Bigg).
\end{aligned} \tag{3.16}$$

Combining this with (3.14), we obtain

$$\|X - \tilde{X}\|_{p;[0,t_3]} \lesssim |X_0 - \tilde{X}_0| + \|\mathbf{Y}; \tilde{\mathbf{Y}}\|_{p;[0,t_3]} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2};[0,t_3]} + \int_0^{t_3} |\alpha_r - \tilde{\alpha}_r| \, \mathrm{d}r. \tag{3.17}$$

Step 4. From now on we allow the multiplicative constant indicated by $\lesssim$ to also depend on ψ. We have

$$\begin{aligned}
&\Bigg\| \int_0^{\cdot} \psi(X_r, \gamma_r) \, \mathrm{d}\mathbf{Y}_r - \int_0^{\cdot} \psi(\tilde{X}_r, \tilde{\gamma}_r) \, \mathrm{d}\tilde{\mathbf{Y}}_r \Bigg\|_{p;[0,t_3]} \\
&\lesssim \big(\|X - \tilde{X}\|_\infty + \|\gamma - \tilde{\gamma}\|_\infty \big) \|Y\|_{p;[0,t_3]} + \|Y - \tilde{Y}\|_{p;[0,t_3]} \\
&\qquad + \big\| R^{\int_0^{\cdot} \psi(X_r) \, \mathrm{d}\mathbf{Y}_r} - R^{\int_0^{\cdot} \psi(\tilde{X}_r) \, \mathrm{d}\tilde{\mathbf{Y}}_r} \big\|_{\frac{p}{2};[0,t_3]} \\
&\lesssim |X_0 - \tilde{X}_0| + \|X - \tilde{X}\|_{p;[0,t_3]} + \|X' - \tilde{X}'\|_{p;[0,t_3]} + \|R^X - R^{\tilde{X}}\|_{\frac{p}{2};[0,t_3]} \\
&\qquad + \|\mathbf{Y}; \tilde{\mathbf{Y}}\|_{p;[0,t_3]} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2};[0,t_3]} \\
&\lesssim |X_0 - \tilde{X}_0| + \|\mathbf{Y}; \tilde{\mathbf{Y}}\|_{p;[0,t_3]} + \|\gamma - \tilde{\gamma}\|_\infty + \|\gamma - \tilde{\gamma}\|_{\frac{p}{2};[0,t_3]} + \int_0^{t_3} |\alpha_r - \tilde{\alpha}_r| \, \mathrm{d}r,
\end{aligned} \tag{3.18}$$

where we have applied Lemma 3.4.2, and then used (3.16) and (3.17) to obtain the last line.

Step 5. Recalling (3.12) and (3.13), we have shown that on any time interval $[s,t]$ sufficiently small that

$$|||\mathbf{Y}|||_{p;[s,t]} + \|\gamma\|_{\frac{p}{2};[s,t]} + \int_s^t \big(1+|\alpha_r|\big)\,\mathrm{d}r \leq h := \min\left\{\frac{1}{4C_1^2}, \frac{1}{4C_2^2}\right\}, \tag{3.19}$$

given an initial value X_s, there exists a unique solution X to the RDE (3.7) over the interval $[s,t]$ corresponding to the driving rough path $\mathbf{Y}$ and the parameters α, γ. Moreover, if (3.12), (3.13) and additionally (3.15) hold for both $(\mathbf{Y}, \alpha, \gamma)$ and $(\tilde{\mathbf{Y}}, \tilde{\alpha}, \tilde{\gamma})$, then the estimates (3.17) and (3.18) also hold over the interval $[s,t]$.

We can divide the (arbitrary) time interval $[0,T]$ into a partition $0 = s_0 < s_1 < \ldots < s_N = T$ such that (3.19) holds on each subinterval $[s_{i-1}, s_i]$, $i = 1, \ldots, N$. (That this partition is finite follows from the compactness of the interval $[0,T]$).

Note that, crucially, all of our estimates hold independently of the initial values $X_0 = x$ and γ_0. Thus, by pasting solutions on each of these subintervals together, we obtain a unique global solution X, which holds over the entire interval $[0,T]$.

By possibly increasing the length of each subinterval, we may assume without loss of generality that the inequality in (3.19) is actually an equality on every subinterval $[s_{i-1}, s_i]$ in our partition (except possibly on one such subinterval). It then follows that, on each subinterval $[s_{i-1}, s_i]$, at least one of $\|Y\|_{p;[s_{i-1},s_i]} \geq h/4$ or $\|\mathbb{Y}\|_{\frac{p}{2};[s_{i-1},s_i]} \geq h/4$ or $\|\gamma\|_{\frac{p}{2};[s_{i-1},s_i]} \geq h/4$ or $\int_{s_{i-1}}^{s_i} \big(1+|\alpha_r|\big)\,\mathrm{d}r \geq h/4$ holds. Since

$$\sum_{i=1}^k \|Y\|_{p;[s_{i-1},s_i]}^p \leq \|Y\|_{p;[0,T]}^p, \qquad \sum_{i=1}^k \|\mathbb{Y}\|_{\frac{p}{2};[s_{i-1},s_i]}^{\frac{p}{2}} \leq \|\mathbb{Y}\|_{\frac{p}{2};[0,T]}^{\frac{p}{2}},$$

$$\sum_{i=1}^k \|\gamma\|_{\frac{p}{2};[s_{i-1},s_i]}^{\frac{p}{2}} \leq \|\gamma\|_{\frac{p}{2};[0,T]}^{\frac{p}{2}}, \qquad \text{and} \qquad \sum_{i=1}^k \int_{s_{i-1}}^{s_i} \big(1+|\alpha_r|\big)\,\mathrm{d}r \leq \int_0^T \big(1+|\alpha_r|\big)\,\mathrm{d}r$$

for every $k \geq 1$, we infer an upper bound on the total number of subintervals:

$$N \lesssim 1 + \|\gamma\|_{\frac{p}{2};[0,T]}^{\frac{p}{2}} + \int_0^T \big(1+|\alpha_r|\big)\,\mathrm{d}r. \tag{3.20}$$

It follows from Lemma 3.4.1 that

$$\left\| \int_0^\cdot \psi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r \right\|_{p;[s_{i-1},s_i]} \lesssim |||\mathbf{Y}|||_{p;[s_{i-1},s_i]}.$$

Hence, by Lemma 3.2.2,

$$\begin{aligned}\left\| \int_0^\cdot \psi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r \right\|_{p;[0,T]} &\leq N^{\frac{p-1}{p}} \left(\sum_{i=1}^N \left\| \int_0^\cdot \psi(X_r, \gamma_r)\,\mathrm{d}\mathbf{Y}_r \right\|_{p;[s_{i-1},s_i]}^p \right)^{\frac{1}{p}} \\ &\lesssim N^{\frac{p-1}{p}} |||\mathbf{Y}|||_{p;[0,T]}\end{aligned}$$

and, combining this with (3.20), we obtain (3.8).

We now choose the partition $0 = s_0 < s_1 < \ldots < s_N = T$ so that (3.12), (3.13) and (3.15) all hold on each subinterval $[s_{i-1}, s_i]$ for both $(\mathbf{Y}, \alpha, \gamma)$ and $(\tilde{\mathbf{Y}}, \tilde{\alpha}, \tilde{\gamma})$. Since

$$\|X - \tilde{X}\|_{p;[0,T]} \leq N^{\frac{p-1}{p}} \left(\sum_{i=1}^{N} \|X - \tilde{X}\|^p_{p;[s_{i-1},s_i]} \right)^{\frac{1}{p}}$$

(again by Lemma 3.2.2), we can paste the local estimates (3.17)–(3.18) over each subinterval to obtain the global estimates (3.9)–(3.10).

The statements on consistency of rough and stochastic integrals and differential equations follow as in Corollary 5.2 and Theorem 9.1 in [32]. □

Chapter 4

Hidden Markov Models under Uncertainty

In this chapter we turn our attention to parameter uncertainty in the filtering of another class of stochastic processes, namely continuous-time hidden Markov models. Moreover, rather than just being based on our a priori beliefs, as in Chapter 2, here the penalty function of our nonlinear expectation will be based on a posterior density, allowing the penalty to 'learn' from new observations.

4.1 The Wonham filter

Let X be a continuous-time Markov chain taking values in the standard basis $\mathcal{X} = \{e_1, \dots, e_m\}$ of $\mathbb{R}^m$. Writing $A_t = [a_{ij}(t)]_{m\times m}$ for the rate matrix[1] of X, the Doob–Meyer decomposition of X is given by

$$\mathrm{d}X_t = A_t X_t \,\mathrm{d}t + \mathrm{d}M_t,$$

where M is a càdlàg martingale. We assume that $\mathbb{E}[X_0] = \pi_0$, for some element π_0 of the open probability simplex

$$S^m := \left\{ x = (x_1, \dots, x_m)^\top \in (0,1)^m \;\middle|\; \sum_{j=1}^m x_j = 1 \right\}.$$

We consider the problem of estimating the current state of the chain X from observations of the $\mathbb{R}^d$-valued process $Y = (Y^1, \dots, Y^d)$, with $Y_0 = 0$ and dynamics

$$\mathrm{d}Y_t^i = (h^i)^\top X_t \,\mathrm{d}t + \mathrm{d}B_t^i, \qquad i = 1, \dots, d, \tag{4.1}$$

[1] This is defined as the transpose of the Q-matrix of the chain, so that $\mathbb{P}(X_{t+\varepsilon} = e_i \,|\, X_t = e_j) = \delta_{ij} + a_{ij}(t)\varepsilon + o(\varepsilon)$ as $\varepsilon \to 0$, where δ_{ij} is the Kronecker delta.

where $h^i \in \mathbb{R}^m$, $i = 1, \dots, d$, and $B = (B^1, \dots, B^d)$ is a standard d-dimensional Brownian motion, independent of X. In the following it will be convenient to write

$$H^i = \operatorname{diag}(h^i),$$

i.e. the diagonal matrix with diagonal elements given by the vector h^i.

Remark 4.1.1. For clarity we assume the parameters h^i to be constant and known. They may also be taken to be time-dependent, provided that they are taken to be continuous and with finite $\frac{p}{2}$-variation; cf. Theorem 3.4.3.

We will denote by $(\mathcal{Y}_t)_{t\geq 0}$ the (completed) natural filtration generated by the observation process Y. The goal of the associated filtering problem is to determine, at each time t, the posterior distribution $\pi_t = \mathbb{E}[X_t \,|\, \mathcal{Y}_t]$. In the present setting, the filtering problem was resolved by Wonham [63]; see also Bain and Crisan [5, Chapter 3]. The optimal filter π is the unique continuous $\mathcal{Y}_t$-adapted solution of the stochastic differential equation

$$\mathrm{d}\pi_s = A_s\pi_s\,\mathrm{d}s + \sum_{i=1}^{d} \big(H^i - (h^i)^\top \pi_s I\big)\pi_s\big(\mathrm{d}Y_s^i - (h^i)^\top \pi_s\,\mathrm{d}s\big) \tag{4.2}$$

where I denotes the $m \times m$-identity matrix.

In component form, omitting the time parameter for clarity, this is

$$\mathrm{d}\pi_k = \sum_{j=1}^{m} a_{kj}\pi_j\,\mathrm{d}s + \sum_{i=1}^{d}\left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right)\pi_k\left(\mathrm{d}Y^i - \sum_{j=1}^{m} h_j^i\pi_j\,\mathrm{d}s\right), \qquad k = 1, \dots, m. \tag{4.3}$$

It will be convenient later to interpret the stochastic integrals appearing in (4.2) and (4.3) in the sense of Stratonovich, rather than that of Itô. The difference between these two interpretations is given by the quadratic covariation:

$$\begin{aligned}
&\mathrm{d}\left\langle \left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right)\pi_k, Y^i\right\rangle \\
&\quad = \left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right)\mathrm{d}\langle \pi_k, Y^i\rangle + \pi_k\,\mathrm{d}\left\langle \left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right), Y^i\right\rangle \\
&\quad = \left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right)^2\pi_k\,\mathrm{d}s - \pi_k\sum_{j=1}^{m} h_j^i\left(h_j^i - \sum_{l=1}^{m} h_l^i\pi_l\right)\pi_j\,\mathrm{d}s \\
&\quad = \left(h_k^i - \sum_{j=1}^{m} h_j^i\pi_j\right)^2\pi_k\,\mathrm{d}s - \pi_k\left(\sum_{j=1}^{m}(h_j^i)^2\pi_j - \left(\sum_{j=1}^{m} h_j^i\pi_j\right)^2\right)\mathrm{d}s.
\end{aligned}$$

We thus have that

$$\begin{aligned}\mathrm{d}\pi_s = A_s\pi_s\,\mathrm{d}s + \sum_{i=1}^{d}\big(H^i - (h^i)^\top\pi_s I\big)\pi_s\big(\circ\,\mathrm{d}Y_s^i - (h^i)^\top\pi_s\,\mathrm{d}s\big)\\ -\frac{1}{2}\sum_{i=1}^{d}\Big(\big(H^i - (h^i)^\top\pi_s I\big)^2 - \big((h^i)^\top H^i\pi_s\big)I + \big((h^i)^\top\pi_s\big)^2 I\Big)\pi_s\,\mathrm{d}s,\end{aligned}$$

which may be simplified to

$$\mathrm{d}\pi_s = A_s\pi_s\,\mathrm{d}s + \frac{1}{2}\sum_{i=1}^{d}\Big(\big((h^i)^\top H^i\pi_s\big)I - (H^i)^2\Big)\pi_s\,\mathrm{d}s + \sum_{i=1}^{d}\big(H^i - (h^i)^\top\pi_s I\big)\pi_s\circ\mathrm{d}Y_s^i. \tag{4.4}$$

For notational simplicity, we rewrite this equation in the form

$$\mathrm{d}\pi_s = \big(A_s\pi_s + b(\pi_s)\big)\,\mathrm{d}s + \phi(\pi_s)\circ\mathrm{d}Y_s, \tag{4.5}$$

where the coefficients $b\colon S^m \to \mathbb{R}^m$ and $\phi\colon S^m \to \mathcal{L}(\mathbb{R}^d;\mathbb{R}^m)$ are chosen such as to make this equation consistent with (4.4). Note that, since π lives in the bounded domain S^m, we have that $b \in \mathrm{Lip}_b$ and $\phi \in C_b^3$.

4.2 Robust filtering with DR-expectations

We shall consider uncertainty of the rate matrix A of the Markov chain X.

Notation 4.2.1. We write $\mathcal{R}$ for the set of admissible rate matrices, which we assume to be a known (possibly proper) subset of the space of $m \times m$ rate matrices, that is, matrices $A = [a_{ij}]_{m\times m}$ such that

$$\begin{aligned} a_{ij} \geq 0 \quad &\text{for all} \quad i \neq j, \quad \text{and}\\ \sum_{i=1}^{m} a_{ij} = 0 \quad &\text{for all} \quad j = 1,\dots,m.\end{aligned}$$

We also write $\mathcal{A}$ for the space of bounded measurable paths taking values in $\mathcal{R}$.

Example 4.2.2. As a simple example, one might have $m = 2$ and

$$\mathcal{R} = \left\{ \begin{pmatrix} -\lambda & \lambda \\ \lambda & -\lambda \end{pmatrix} \,\middle|\, \lambda \in [0,1] \right\}.$$

Assumption 4.2.3. We assume that $0 \in \mathcal{R}$, i.e. the zero matrix is admissible.

Our setup is the following. Let X, Y be $\mathcal{F}_t$-adapted processes. For each admissible rate matrix $A \in \mathcal{A}$ and initial distribution $\pi_0 \in S^m$, we let $\mathbb{P}^{A,\pi_0}$ be a probability measure such that the law of (X, Y) is equal to the law of $(\tilde{X}^{A,\pi_0}, \tilde{Y}^{A,\pi_0})$, where, on some probability space, $\tilde{X}^{A,\pi_0}$ is a Markov chain with transition matrix A and initial distribution $\mathbb{E}[\tilde{X}_0^{A,\pi_0}] = \pi_0$, and $\tilde{Y}^{A,\pi_0}$ satisfies the SDE (4.1) driven by some Brownian motion B with $\tilde{Y}_0^{A,\pi_0} = 0$. In particular, the processes X, Y as functions on $\Omega \times [0,T]$, and hence the (uncompleted) filtration $\sigma(Y_s : s \in [0,t])$, $t \in [0,T]$, are defined independently of the choice of parameters.

Similarly to Chapter 2, we introduce the following nonlinear expectation:

$$\mathcal{E}(\varphi(X_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{A \in \mathcal{A},\, \pi_0 \in S^m} \left\{ \mathbb{E}^{A,\pi_0}[\varphi(X_t) \,|\, \mathcal{Y}_t] - \left(\frac{1}{k_1} \beta_t(A, \pi_0 \,|\, \mathcal{Y}_t) \right)^{k_2} \right\}. \tag{4.6}$$

In Chapter 2 the penalty β was assumed to be fixed a priori, i.e. it only took our prior beliefs into account. Although the parameters of the underlying system are unknown, as we make new observations we may wish to use these observations to update our opinion of how reasonable different parameter choices are. We shall therefore suppose that this penalty takes the form of a negative log-posterior density. That is, we take

$$\beta_t(A, \pi_0 \,|\, \mathcal{Y}_t) = -\log\Big(\vartheta_t(A, \pi_0) L_t(A, \pi_0 \,|\, \mathcal{Y}_t) \Big), \tag{4.7}$$

where ϑ and $L(\cdot \,|\, \mathcal{Y}_t)$ denote the prior and likelihood respectively.

Remark 4.2.4. Since the posterior is only proportional to the product of prior and likelihood, (4.7) is correct up to an additive constant. For simplicity we will omit this constant from our analysis, conceding that our penalty function is correct up to an additive constant. This constant may be reintroduced upon numerical computation, chosen to ensure that the penalty function always takes the value zero at its minimum.

The penalty function in (4.7) is built from the log-likelihood function, a familiar object from classical statistics. Penalties based on log-likelihoods form the basis of the data-driven robust (DR) expectation of [18], which allows the level of penalisation of different parameter choices to be recursively updated through time as we collect new observations. We refer to [18] for further discussion. Here we add to this an additional penalty based on our prior beliefs, which may be calibrated accordingly.

We shall assume that the prior takes the form

$$-\log \vartheta_t(A, \pi_0) = \int_0^t \mathfrak{f}(\pi_s, A_s) \, \mathrm{d}s + g(\pi_0), \tag{4.8}$$

where the functions $\mathfrak{f}\colon S^m \times \mathcal{R} \to \mathbb{R}$ and $g\colon S^m \to \mathbb{R}$ may be calibrated to represent our prior beliefs about the plausibility of different parameter choices. Here $\pi = (\pi_s)_{s\in[0,t]}$ is the posterior distribution corresponding to the parameters A and π_0.

Note that the measures $\mathbb{P}^{A,\pi_0}$ for different choices of A and π_0 are all equivalent on $\mathcal{Y}_t$. A natural choice for $L_t(\cdot \,|\, \mathcal{Y}_t)$ is thus the Radon–Nikodym derivative

$$L_t(A, \pi_0 \,|\, \mathcal{Y}_t) = \left(\frac{\mathrm{d}\mathbb{P}^{A,\pi_0}}{\mathrm{d}\mathbb{P}^{\bar{A},\bar{\pi}_0}} \right)_{\mathcal{Y}_t},$$

that is, the likelihood ratio of the (arbitrary) parameter choice A, π_0, with respect to a (fixed) choice of reference parameters $\bar{A}, \bar{\pi}_0$. We will now derive an explicit expression for this likelihood.

Recall (from e.g. Bain and Crisan [5, Chapter 2]) that for a given choice of parameters A, π_0, the *innovation* process $V = (V^1, \ldots, V^d)$, given in this setting by

$$\mathrm{d}V_s^i = \mathrm{d}Y_s^i - (h^i)^\top \pi_s \,\mathrm{d}s, \qquad i = 1, \ldots, d,$$

is a $\mathcal{Y}_t$-adapted Brownian motion under $\mathbb{P}^{A,\pi_0}$. Writing $\bar{\pi}$ (resp. $\bar{V}$) for the posterior distribution (resp. innovation process) under the reference measure $\mathbb{P}^{\bar{A},\bar{\pi}_0}$, we have

$$\mathrm{d}V_s^i = \mathrm{d}\bar{V}_s^i - (h^i)^\top (\pi_s - \bar{\pi}_s) \,\mathrm{d}s, \qquad i = 1, \ldots, d.$$

Thus, by Girsanov's theorem[2] (see e.g. [19, Chapter 15]), we can represent the likelihood as a stochastic exponential, namely

$$L_t(A, \pi_0 \,|\, \mathcal{Y}_t) = \exp\left(\sum_{i=1}^d \left(\int_0^t (h^i)^\top (\pi_s - \bar{\pi}_s) \,\mathrm{d}\bar{V}_s^i - \frac{1}{2} \int_0^t \left| (h^i)^\top (\pi_s - \bar{\pi}_s) \right|^2 \mathrm{d}s \right) \right).$$

Substituting $\mathrm{d}\bar{V}_s^i = \mathrm{d}Y_s^i - (h^i)^\top \bar{\pi}_s \,\mathrm{d}s$, a short calculation yields

$$-\log L_t(A, \pi_0 \,|\, \mathcal{Y}_t) = \sum_{i=1}^d \left(-\int_0^t (h^i)^\top (\pi_s - \bar{\pi}_s) \,\mathrm{d}Y_s^i + \frac{1}{2} \int_0^t \left(\left| (h^i)^\top \pi_s \right|^2 - \left| (h^i)^\top \bar{\pi}_s \right|^2 \right) \mathrm{d}s \right).$$

Since the reference parameters are taken to be fixed, they simply amount to an additive constant in the above expression. That is,

$$-\log L_t(A, \pi_0 \,|\, \mathcal{Y}_t) = \sum_{i=1}^d \left(-\int_0^t (h^i)^\top \pi_s \,\mathrm{d}Y_s^i + \frac{1}{2} \int_0^t \left| (h^i)^\top \pi_s \right|^2 \mathrm{d}s \right) + \text{const.} \tag{4.9}$$

As in Remark 4.2.4, we shall henceforth omit this constant.

[2] Although in general the innovation process does not generate the observation filtration, it does have the predictable representation property; see Theorem 3.1 in [36].

For later convenience, we make the transformation

$$
\begin{aligned}
-\int_0^t (h^i)^\top \pi_s \,\mathrm{d}Y_s^i &= -\int_0^t (h^i)^\top \pi_s \circ \mathrm{d}Y_s^i + \frac{1}{2}\big\langle (h^i)^\top \pi, Y^i \big\rangle_t \\
&= -\int_0^t (h^i)^\top \pi_s \circ \mathrm{d}Y_s^i + \frac{1}{2}\int_0^t (h^i)^\top \big(H^i - (h^i)^\top \pi_s I\big)\pi_s \,\mathrm{d}s.
\end{aligned}
$$

Substituting back into (4.9), we obtain

$$
\begin{aligned}
&-\log L_t(A, \pi_0 \,|\, \mathcal{Y}_t) \\
&\quad= \sum_{i=1}^d \bigg(-\int_0^t (h^i)^\top \pi_s \circ \mathrm{d}Y_s^i + \frac{1}{2}\int_0^t \big|(h^i)^\top \pi_s\big|^2 + (h^i)^\top \big(H^i - (h^i)^\top \pi_s I\big)\pi_s \,\mathrm{d}s \bigg) \\
&\quad= \sum_{i=1}^d \bigg(-\int_0^t (h^i)^\top \pi_s \circ \mathrm{d}Y_s^i + \frac{1}{2}\int_0^t (h^i)^\top H^i \pi_s \,\mathrm{d}s \bigg). \qquad (4.10)
\end{aligned}
$$

Combining (4.8) and (4.10) with (4.7), and then substituting into (4.6), we obtain

$$
\begin{aligned}
\mathcal{E}(\varphi(X_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{A \in \mathcal{A},\, \pi_0 \in S^m} \Bigg\{ &\mathbb{E}^{A,\pi_0}[\varphi(X_t) \,|\, \mathcal{Y}_t] \qquad (4.11) \\
&- \bigg(\frac{1}{k_1} \bigg(\int_0^t f(\pi_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s) \circ \mathrm{d}Y_s + g(\pi_0) \bigg) \bigg)^{k_2} \Bigg\},
\end{aligned}
$$

where, for notational simplicity, we have introduced the functions $f\colon S^m \times \mathcal{R} \to \mathbb{R}$ and $\psi\colon S^m \to \mathcal{L}(\mathbb{R}^d; \mathbb{R})$, defined by

$$
f(\pi, A) = \mathfrak{f}(\pi, A) + \frac{1}{2}\sum_{i=1}^d (h^i)^\top H^i \pi, \qquad \text{and} \qquad \psi^i(\pi) = -(h^i)^\top \pi, \quad i = 1, \dots, d.
$$

4.3 Fixing an observation path

We must take care, as in general the measures $\mathbb{P}^{A,\pi_0}$ are not equivalent on $\mathcal{F}_t$ (since we do not assume that the entries of the rate matrix A are strictly positive), so the completed filtration $\mathcal{Y}_t = \sigma(Y_s : s \in [0,t]) \vee \mathcal{N}$ will in general depend on the choice of parameters.

Fix a reference measure $\mathbb{P}^{\bar{A},\bar{\pi}_0}$. We then enhance the observation process Y using Stratonovich integration:

$$
\mathbb{Y}_{s,t} := \int_s^t Y_{s,r} \otimes \circ\, \mathrm{d}Y_r \qquad \text{for all} \quad (s,t) \in \Delta_{[0,T]}, \qquad (4.12)
$$

defined under the measure $\mathbb{P}^{\bar{A},\bar{\pi}_0}$, which we recall defines a random geometric rough path $\mathbf{Y} = (Y, \mathbb{Y}) \in \mathscr{C}_g^{0,p}$.

Note that, since the integrand is $\sigma(Y_s : s \in [0,t])$-adapted, $\mathbb{Y}$ coincides almost surely with the same integral defined under any other choice of measure $\mathbb{P}^{A,\pi_0}$ (even though the corresponding completed filtrations $\mathcal{Y}_t$ may not agree). Thus, defining π as the solution of the rough filtering equation

$$\mathrm{d}\pi_s = \big(A_s\pi_s + b(\pi_s)\big)\,\mathrm{d}s + \phi(\pi_s)\,\mathrm{d}\mathbf{Y}_s \tag{4.13}$$

(which exists by part (i) of Theorem 3.4.3), then, for each parameter choice (A, π_0), π is indistinguishable from the solution of the Stratonovich equation

$$\mathrm{d}\pi_s = \big(A_s\pi_s + b(\pi_s)\big)\,\mathrm{d}s + \phi(\pi_s) \circ \mathrm{d}Y_s$$

(cf. (4.5)), and moreover the rough integral

$$\int_0^t \psi(\pi_s)\,\mathrm{d}\mathbf{Y}_s$$

coincides almost surely with the stochastic integral

$$\int_0^t \psi(\pi_s) \circ \mathrm{d}Y_s.$$

In particular, it follows that

$$\sum_{j=1}^m \pi_{j,t}\varphi(e_j)$$

is a version of the conditional expectation $\mathbb{E}^{A,\pi_0}[\varphi(X_t) \,|\, \mathcal{Y}_t]$ for every choice of parameters and every functional φ.

Choosing this version, we then have that

$$\begin{aligned}&\mathcal{E}(\varphi(X_t) \,|\, \mathcal{Y}_t)\\ &= \operatorname*{ess\,sup}_{A\in\mathcal{A},\,\pi_0\in S^m} \left\{ \sum_{j=1}^m \pi_{j,t}\varphi(e_j) - \left(\frac{1}{k_1}\left(\int_0^t f(\pi_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s)\,\mathrm{d}\mathbf{Y}_s + g(\pi_0)\right)\right)^{k_2}\right\}.\end{aligned}$$

It is then clear that the 'pathwise version' of the nonlinear expectation should be given by

$$\begin{aligned}&\mathcal{E}(\varphi(X_t) \,|\, \mathbf{Y}_{[0,t]})\\ &:= \sup_{A\in\mathcal{A},\,\pi_0\in S^m} \left\{ \sum_{j=1}^m \pi_{j,t}\varphi(e_j) - \left(\frac{1}{k_1}\left(\int_0^t f(\pi_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s)\,\mathrm{d}\mathbf{Y}_s + g(\pi_0)\right)\right)^{k_2}\right\},\end{aligned}$$

where here we fix a realisation $\mathbf{Y}_{[0,t]} = (Y|_{[0,t]}(\omega), \mathbb{Y}|_{\Delta_{[0,t]}}(\omega))$ of the enhanced observation path, and π satisfies the rough filtering equation (4.13) driven by this realisation of $\mathbf{Y}$.

Remark 4.3.1. Writing $\pi = \pi^{\mathbf{Y}}$ for the unique solution of (4.13), the stability estimate (3.9) implies that the map $\mathbf{Y} \mapsto \pi^{\mathbf{Y}}$ is continuous with respect to the p-variation rough path topology. In particular, we immediately deduce 'robustness' of the Wonham filter in the sense of Crisan, Diehl, Friz and Oberhauser [21], that is, continuity of the filter with respect to the (enhanced) observation path. Incidentally, one can actually choose a version of the Wonham filter which is robust with respect to the (unenhanced) observation path Y alone; see [5, Chapter 5].

For a given $\mathbf{Y} \in \mathscr{C}_g^{0,p}$, we define the functional $\kappa(\cdot \,|\, \mathbf{Y})\colon [0,T] \times S^m \to \mathbb{R}$ by

$$\kappa(t, x \,|\, \mathbf{Y}) := \inf \left\{ \int_0^t f(\pi_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s)\,\mathrm{d}\mathbf{Y}_s + g(\pi_0) \;\middle|\; \begin{array}{c} A, \pi_0 \text{ such that} \\ \pi_t = x \end{array} \right\}, \quad (4.14)$$

where the infimum is taken over all parameters $A \in \mathcal{A}$ and $\pi_0 \in S^m$ such that the solution $\pi = \pi^{\mathbf{Y}}$ of (4.13) takes the value x at time t.

Lemma 4.3.2. *For any $\mathbf{Y} \in \mathscr{C}_g^{0,p}$, $t \in [0,T]$ and functional $\varphi\colon \mathcal{X} \to \mathbb{R}$, we have that*

$$\mathcal{E}\big(\varphi(X_t) \,\big|\, \mathbf{Y}_{[0,t]}\big) = \sup_{x \in S^m} \left\{ \sum_{j=1}^m x_j \varphi(e_j) - \left(\frac{1}{k_1} \kappa(t, x \,|\, \mathbf{Y}) \right)^{k_2} \right\}.$$

Proof. For any $x \in S^m$, we have that

$$\begin{aligned} \sup \Bigg\{ & \sum_{j=1}^m \pi_{j,t} \varphi(e_j) \\ & - \left(\frac{1}{k_1} \left(\int_0^t f(\pi_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s)\,\mathrm{d}\mathbf{Y}_s + g(\pi_0) \right) \right)^{k_2} \;\middle|\; \begin{array}{c} A, \pi_0 \text{ such that} \\ \pi_t = x \end{array} \Bigg\} \\ = & \sum_{j=1}^m x_j \varphi(e_j) - \left(\frac{1}{k_1} \kappa(t, x \,|\, \mathbf{Y}) \right)^{k_2}. \end{aligned}$$

Then simply take the supremum over $x \in S^m$. □

As is clear from the result of Lemma 4.3.2, given κ, one can evaluate the DR-expectation for any (and all) functionals φ. We shall therefore turn our attention to understanding the function κ, and in particular how it may be recursively computed as we make new observations.

The key insight of our approach is that the function κ, as defined in (4.14), has the form of the value function of an optimal control problem. In this view, we can interpret the parameter A as the control and π as the state variable, which satisfies the controlled dynamics given by (4.13).

In order to make this interpretation exact, we should think of taking the infimum in (4.14) over only the parameter A, with the initial value π_0 then being uniquely determined by the path $\mathbf{Y}$, the control A, and the terminal condition $\pi_t = x$. That is, we rewrite (4.14) as

$$\kappa(t, x \,|\, \mathbf{Y}) = \inf_{A \in \mathcal{A}} \left\{ \int_0^t f(\pi_s, A_s)\, \mathrm{d}s + \int_0^t \psi(\pi_s)\, \mathrm{d}\mathbf{Y}_s + g(\pi_0) \right\}. \tag{4.15}$$

However, we need to take care here to exclude unphysical trajectories. That is, given a path $\mathbf{Y}$ and terminal condition $\pi_t = x$, there exist choices of control A for which the solution $\pi = (\pi_s)_{s \in [0,t]}$ of (4.13) leaves the domain S^m, and such controls should therefore be discarded.

Example 4.3.3. Suppose that, at some time $s \in (0, t]$, we have $\pi_s = (1, 0, 0, \ldots, 0)^\top$. Then (4.13) reduces to $\mathrm{d}\pi_s = A_s \pi_s \, \mathrm{d}s$, so the dynamics of π at time s are simply given by the first column of A_s. Since A_s is a rate matrix, we see that, unless its first column is zero, $\pi_r \notin S^m$ for $r < s$.

Heuristically, the boundary of the domain S^m corresponds to absolute certainty with regard to the signal process (not) being in a particular state, which is not a reasonable inference. Although we don't obtain an initial value $\pi_0 \in S^m$ for trajectories which leave the domain, we can simply assign an infinite initial cost to all such trajectories, so that the corresponding controls are never considered when taking the infimum in (4.15).

We also note that, thanks to Assumption 4.2.3, given any path $\mathbf{Y}$ and terminal condition $\pi_t = x$, there always exists a choice of control (for example $A = 0$) such that $\pi_0 \in S^m$. Thus, the function $\kappa(\cdot \,|\, \mathbf{Y})$ is finite-valued.

4.4 An unbounded pathwise control problem

We shall henceforth consider the (enhanced) observation path $\mathbf{Y}$ to be fixed, and will therefore suppress the dependency of κ on $\mathbf{Y}$ in our notation.

We have derived the following optimal control problem. We consider the value function

$$\kappa(t, x) = \inf_{A \in \mathcal{A}} \left\{ \int_0^t f(\pi_s^{t,x,A}, A_s)\, \mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\, \mathrm{d}\mathbf{Y}_s + g(\pi_0^{t,x,A}) \right\}, \tag{4.16}$$

where $\pi^{t,x,A}$ satisfies

$$\mathrm{d}\pi_s^{t,x,A} = \left(A_s \pi_s^{t,x,A} + b(\pi_s^{t,x,A})\right) \mathrm{d}s + \phi(\pi_s^{t,x,A})\, \mathrm{d}\mathbf{Y}_s \tag{4.17}$$

with the terminal condition $\pi_t^{t,x,A} = x$. As discussed in the previous section, in (4.16) we assign an infinite initial cost to all trajectories which leave the domain S^m.

We note that this is a 'backward' control problem, in the sense that we prescribe a terminal condition $\pi_t = x$ for the state trajectories and, for each choice of control A, solve the controlled dynamics backwards in time to obtain the corresponding initial value π_0.

More significantly, here we wish to perform the optimisation for every fixed (enhanced) realisation $\mathbf{Y}_{[0,T]}$ of the stochastic process Y. This type of problem is known as 'pathwise stochastic control'. In fact, we have formulated our problem in terms of the optimal control of a rough differential equation, which we wish to perform for an arbitrary geometric rough path $\mathbf{Y} \in \mathscr{C}_g^{0,p}$. Control problems of this type were first studied in Diehl et al. [26].

In the following we shall denote by $\|A\|$ the norm of a given matrix A. The choice of norm here is not important.[3]

Assumption 4.4.1. We assume that the running cost function $f = f(\pi, A)$ is continuous, bounded below, locally Lipschitz continuous in π, uniformly in A, and satisfies

$$\inf_{\pi \in S^m} \frac{f(\pi, A)}{\|A\|} \longrightarrow \infty \qquad \text{as} \quad \|A\| \longrightarrow \infty.$$

Notation 4.4.2. We denote by $C^{\uparrow}(S^m; \mathbb{R})$ the space of continuous functions $\tilde{g} \colon S^m \to \mathbb{R}$ which explode at the boundary of the domain S^m, i.e. such that $\tilde{g}(x) \to \infty$ as $\min_{1 \le j \le m} x_j \to 0$.

Assumption 4.4.3. We assume that the initial cost function $g = g(x)$ is locally Lipschitz continuous, and that $g \in C^{\uparrow}(S^m; \mathbb{R})$.

Notation 4.4.4. Let $L > 0$ such that $|||\mathbf{Y}|||_{p;[0,T]} \le L$. In the following we shall use $\lesssim$ to denote inequality up to a multiplicative constant which may depend on any of the dimensions m, d, the functions f, g, the measure of regularity p, the bound L, the terminal time T, and the parameters h^i, $i = 1, \dots, d$.

Lemma 4.4.5. *The value function κ is bounded below, and locally bounded above, uniformly in paths $\mathbf{Y}$ satisfying $|||\mathbf{Y}|||_{p;[0,T]} \le L$.*

Proof. By part (ii) of Theorem 3.4.3, we have

$$\left| \int_0^t \psi(\pi_s^{t,x,A}) \, \mathrm{d}\mathbf{Y}_s \right| \le \left\| \int_0^{\cdot} \psi(\pi_s^{t,x,A}) \, \mathrm{d}\mathbf{Y}_s \right\|_{p;[0,t]} \lesssim 1 + \int_0^t \|A_s\|^{\frac{p-1}{p}} \, \mathrm{d}s.$$

[3]Indeed, as this space is finite-dimensional, all norms are equivalent.

By Assumption 4.4.1, for any $\varepsilon > 0$ we infer the existence of a constant C_ε such that

$$\|A\|^{\frac{p-1}{p}} \leq C_\varepsilon + \varepsilon f(\pi, A) \tag{4.18}$$

for all $\pi \in S^m$ and all $A \in \mathcal{R}$. We then have that

$$\left| \int_0^t \psi(\pi_s^{t,x,A}) \, \mathrm{d}\mathbf{Y}_s \right| \leq C + \frac{1}{2} \int_0^t f(\pi_s^{t,x,A}, A_s) \, \mathrm{d}s \tag{4.19}$$

for some constant C. Thus, recalling (4.16), we have

$$\kappa(t,x) \geq \inf_{A \in \mathcal{A}} \left\{ \frac{1}{2} \int_0^t f(\pi_s^{t,x,A}, A_s) \, \mathrm{d}s - C + g(\pi_0^{t,x,A}) \right\} \tag{4.20}$$

and, since f and g are bounded below, the same is true of κ.

Now let K be a compact subset of S^m. It is straightforward to see that there exists another compact subset $Q \subset S^m$, such that $\pi_s^{t,x,0} \in Q$ for all $(t,x) \in [0,T] \times K$ and $s \in [0,t]$ (where here we have taken the zero control $A = 0$, which is valid by Assumption 4.2.3). Since $f(\cdot, 0)$ and g are assumed to be continuous, they are bounded on Q, so that $\sup_{0 \leq s \leq t \leq T,\, x \in K} f(\pi_s^{t,x,0}, 0) < \infty$ and $\sup_{0 \leq t \leq T,\, x \in K} g(\pi_0^{t,x,0}) < \infty$. Using (4.19) again, we have

$$\sup_{(t,x) \in [0,T] \times K} \kappa(t,x) \leq \sup_{(t,x) \in [0,T] \times K} \left\{ C + \frac{3}{2} \int_0^t f(\pi_s^{t,x,0}, 0) \, \mathrm{d}s + g(\pi_0^{t,x,0}) \right\} < \infty,$$

which implies that κ is locally bounded above. □

Corollary 4.4.6. *Let K be a compact subset of S^m. There exists a bound $M > 0$ such that, when taking the infimum in* (4.16) *for $(t,x) \in [0,T] \times K$, one may restrict to controls $A \in \mathcal{A}_M$, where*

$$\mathcal{A}_M := \left\{ A \in \mathcal{A} \;\middle|\; \int_0^T \big(1 + \|A_s\|\big) \, \mathrm{d}s \leq M \right\}, \tag{4.21}$$

without changing the value of $\kappa(t,x)$.

Proof. Recalling (4.19) and the fact that g is bounded below, we have

$$\int_0^t f(\pi_s^{t,x,A}, A_s) \, \mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A}) \, \mathrm{d}\mathbf{Y}_s + g(\pi_0^{t,x,A}) \geq \frac{1}{2} \int_0^t f(\pi_s^{t,x,A}, A_s) \, \mathrm{d}s - C$$

for some new constant C. Since κ is bounded above on $[0,T] \times K$, we infer an upper bound on $\int_0^t f(\pi_s^{t,x,A}, A_s) \, \mathrm{d}s$. Recalling (4.18), we see that this then translates into an upper bound on $\int_0^t \|A_s\| \, \mathrm{d}s$, and the result follows. □

Notation 4.4.7. Recall Notation 4.4.4. For the rest of this chapter, whenever we have identified a compact subset $K \subset S^m$, we shall allow the multiplicative constant indicated by the symbol $\lesssim$ to also depend on the constant M in (4.21).

Our main aim is to establish κ as the solution of a rough HJB equation (see (4.45) below). However, it does not seem tractable, nor necessary, to show directly that κ solves such an equation. Instead, as in [26], we first approximate the rough path $\mathbf{Y}$ by a smooth path η. Having solved the associated classical control problem, as is a standard strategy for rough ODEs and PDEs, we can define solutions to our HJB equation for genuinely rough driving paths by taking the closure of smooth paths in rough path topology.

Accordingly, given a smooth path $\eta\colon [0,T] \to \mathbb{R}^d$, we define the approximate value function:

$$\kappa^\eta(t,x) = \inf_{A\in\mathcal{A}} \left\{ \int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + g(\pi_0^{t,x,A}) \right\}, \tag{4.22}$$

where here $\pi^{t,x,A}$ satisfies

$$\mathrm{d}\pi_s^{t,x,A} = \big(A_s\pi_s^{t,x,A} + b(\pi_s^{t,x,A})\big)\,\mathrm{d}s + \phi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s \tag{4.23}$$

with the terminal condition $\pi_t^{t,x,A} = x$.

In the following we will write $\pi^{t,x,A}$ both for the solution of (4.17) and for the solution of (4.23), but the intended solution should always be clear from the context. We will also keep the assumption that $|||\boldsymbol{\eta}|||_{p;[0,T]} \le L$, where $\boldsymbol{\eta}$ is defined as the canonical lift of η, i.e. $\boldsymbol{\eta} = (\eta, \eta^{(2)}) \in \mathscr{C}_g^{0,p}$ with

$$\eta_{s,t}^{(2)} := \int_s^t \eta_{s,r} \otimes \mathrm{d}\eta_r \qquad \text{for all} \quad (s,t) \in \Delta_{[0,T]}. \tag{4.24}$$

4.5 Regularity of the value functions

We have the following dynamic programming principle.

Lemma 4.5.1. *For any t, x and $r \in [0,t]$, the value function κ (as defined in* (4.16) *above), satisfies*

$$\kappa(t,x) = \inf_{A\in\mathcal{A}} \left\{ \int_r^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_r^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s + \kappa(r, \pi_r^{t,x,A}) \right\}. \tag{4.25}$$

Similarly, the approximate value function κ^η satisfies (4.25) *with $\mathbf{Y}$ replaced by η.*

The result of Lemma 4.5.1 follows the same proof as that of Theorem 2.1 in [66, Chapter 4]. In particular, the rough integrals appearing in the controlled dynamics and value function do not cause any additional difficulty.

Proposition 4.5.2. *The value function $\kappa = \kappa(t,x)$ and approximate value function $\kappa^\eta = \kappa^\eta(t,x)$ are locally Lipschitz continuous in x, uniformly in $t \in [0,T]$.*

Proof. Let K be a compact subset of S^m. Let $t \in [0,T]$ and $x, \tilde{x} \in K$. Fix a control $A \in \mathcal{A}$. By Corollary 4.4.6, we may assume that $A \in \mathcal{A}_M$, i.e. $\int_0^T (1 + \|A_s\|)\,\mathrm{d}s \leq M$ for some $M > 0$. By part (iii) of Theorem 3.4.3, in view of Remark 3.4.4, we have

$$\|\pi^{t,x,A} - \pi^{t,\tilde{x},A}\|_{p;[0,t]} \lesssim |x - \tilde{x}|, \tag{4.26}$$

$$\left\| \int_0^\cdot \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s - \int_0^\cdot \psi(\pi_s^{t,\tilde{x},A})\,\mathrm{d}\mathbf{Y}_s \right\|_{p;[0,t]} \lesssim |x - \tilde{x}|. \tag{4.27}$$

Since $g \in C^\uparrow(S^m;\mathbb{R})$ and $\int_0^T (1+\|A_s\|)\,\mathrm{d}s \leq M$ for every control A, it follows that, for $x \in K$, the trajectories $\pi^{t,x,A}$ are uniformly bounded away from the boundary of S^m. That is, controls A which cause $\pi^{t,x,A}$ to get too close to the boundary of S^m are too expensive, and are thus excluded. We may therefore take f and g to be Lipschitz in π. Using (4.26) and (4.27), we then have

$$\begin{aligned}
&\left|\kappa(t,x) - \kappa(t,\tilde{x})\right| \\
&\quad = \sup_{A \in \mathcal{A}_M} \left| \int_0^t \left(f(\pi_s^{t,x,A}, A_s) - f(\pi_s^{t,\tilde{x},A}, A_s)\right)\mathrm{d}s \right. \\
&\qquad\qquad \left. + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s - \int_0^t \psi(\pi_s^{t,\tilde{x},A})\,\mathrm{d}\mathbf{Y}_s + g(\pi_0^{t,x,A}) - g(\pi_0^{t,\tilde{x},A}) \right| \\
&\quad \lesssim \sup_{A \in \mathcal{A}_M} \left(\int_0^t |\pi_s^{t,x,A} - \pi_s^{t,\tilde{x},A}|\,\mathrm{d}s + |x - \tilde{x}| + |\pi_0^{t,x,A} - \pi_0^{t,\tilde{x},A}| \right) \\
&\quad \lesssim |x - \tilde{x}|.
\end{aligned}$$

An identical argument applies to κ^η, by simply replacing $\mathbf{Y}$ by $\boldsymbol{\eta}$ in the above. □

Proposition 4.5.3. *The value function $\kappa = \kappa(t,x)$ and approximate value function $\kappa^\eta = \kappa^\eta(t,x)$ are continuous in t. Moreover, κ^η is Lipschitz continuous in t, locally uniformly in x.*

Proof. Let K be a compact subset of S^m. Let $0 \leq r \leq t \leq T$ and $x \in K$. By Corollary 4.4.6, we know that we may restrict to controls $A \in \mathcal{A}_M$, so that $\int_0^T (1 +$

$\|A_s\|)\,\mathrm{d}s \leq M$. Similarly to the proof of Theorem 2.2 in Bardi and Da Lio [7], by Lemma 4.5.1 we can further restrict to controls $A \in \mathcal{A}_M$ such that

$$\int_r^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s \leq \left| \int_r^t f(\pi_s^{t,x,0}, 0)\,\mathrm{d}s + \int_r^t \psi(\pi_s^{t,x,0})\,\mathrm{d}\mathbf{Y}_s - \int_r^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s \right| + \left|\kappa(r, \pi_r^{t,x,0}) - \kappa(r, \pi_r^{t,x,A})\right|.$$

By part (ii) of Theorem 3.4.3, we have

$$\left| \int_r^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s \right| \leq \left\| \int_0^{\cdot} \psi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s \right\|_{p;[r,t]} \lesssim |||\mathbf{Y}|||_{p;[r,t]}.$$

As in the proof of Proposition 4.5.2 above, for $x \in K$, the trajectories $\pi^{t,x,A}$ are bounded away from the boundary of S^m. By Proposition 4.5.2, we may therefore take κ to be Lipschitz in x. We then have that

$$\left|\kappa(r, \pi_r^{t,x,0}) - \kappa(r, \pi_r^{t,x,A})\right| \lesssim |\pi_r^{t,x,0} - \pi_r^{t,x,A}|,$$

and, by part (iii) of Theorem 3.4.3,

$$|\pi_r^{t,x,0} - \pi_r^{t,x,A}| \leq \|\pi^{t,x,0} - \pi^{t,x,A}\|_{p;[r,t]} \lesssim \int_r^t \|A_s\|\,\mathrm{d}s.$$

By Assumption 4.4.1, for any $\varepsilon > 0$, there exists a constant C_ε such that

$$\|A\| \leq C_\varepsilon + \varepsilon f(\pi, A) \tag{4.28}$$

for all π and A. Putting all of this together, we have

$$\int_r^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s \lesssim (1 + C_\varepsilon)|t - r| + |||\mathbf{Y}|||_{p;[r,t]} + \varepsilon \int_r^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s,$$

and choosing ε sufficiently small, we obtain

$$\int_r^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s \lesssim |t - r| + |||\mathbf{Y}|||_{p;[r,t]}. \tag{4.29}$$

Using (4.28) again, we deduce that

$$\int_r^t \|A_s\|\,\mathrm{d}s \lesssim |t - r| + |||\mathbf{Y}|||_{p;[r,t]}. \tag{4.30}$$

It follows that

$$\begin{aligned} |\pi_r^{t,x,A} - x| &= \left| \int_r^t \left(A_s \pi_s^{t,x,A} + b(\pi_s^{t,x,A})\right)\mathrm{d}s + \int_r^t \phi(\pi_s^{t,x,A})\,\mathrm{d}\mathbf{Y}_s \right| \\ &\lesssim |t - r| + |||\mathbf{Y}|||_{p;[r,t]}. \end{aligned} \tag{4.31}$$

By Lemma 4.5.1, we can take a sequence of controls $\{A^k\}_{k\geq 1}$ such that

$$\kappa(t,x) = \lim_{k\to\infty} \bigg(\int_r^t f(\pi_s^{t,x,A^k}, A_s^k)\,\mathrm{d}s + \int_r^t \psi(\pi_s^{t,x,A^k})\,\mathrm{d}\mathbf{Y}_s + \kappa(r, \pi_r^{t,x,A^k}) \bigg).$$

Using (4.29) and the fact that κ is locally Lipschitz in x, we have

$$\begin{aligned} &\bigg| \int_r^t f(\pi_s^{t,x,A^k}, A_s^k)\,\mathrm{d}s + \int_r^t \psi(\pi_s^{t,x,A^k})\,\mathrm{d}\mathbf{Y}_s + \kappa(r, \pi_r^{t,x,A^k}) - \kappa(r,x) \bigg| \\ &\lesssim |t-r| + |||\mathbf{Y}|||_{p;[r,t]} + |\pi_r^{t,x,A^k} - x|. \end{aligned}$$

Using (4.31), and taking the limit as $k \to \infty$, we obtain

$$\big|\kappa(t,x) - \kappa(r,x)\big| \lesssim |t-r| + |||\mathbf{Y}|||_{p;[r,t]}.$$

An identical argument applies to κ^η, by replacing $\mathbf{Y}$ by $\boldsymbol{\eta}$ in the above. Moreover, since η is smooth, a short calculation shows that $|||\boldsymbol{\eta}|||_{p;[r,t]} = \|\eta\|_{p;[r,t]} + \|\eta^{(2)}\|_{\frac{p}{2};[r,t]} \lesssim (1 + \|\eta\|_\infty)\|\dot\eta\|_\infty |t-r|$, so that κ^η is Lipschitz in t. □

4.6 A classical HJB equation

We will return to the (rough) value function κ in the next section. For now we will restrict our attention to the smoothed version κ^η, and introduce the associated HJB equation:

$$\frac{\partial \kappa^\eta}{\partial t} + b\cdot\nabla\kappa^\eta + \sup_{A\in\mathcal{R}} \big\{ Ax\cdot\nabla\kappa^\eta - f \big\} + \big(\phi\cdot\nabla\kappa^\eta - \psi\big)\dot\eta = 0, \tag{4.32}$$

$$\kappa^\eta(0,\cdot) = g, \tag{4.33}$$

for a smooth approximation η of $\mathbf{Y}$, where as usual $\dot\eta$ denotes the derivative of η.

Remark 4.6.1. Since the spatial variable x is confined to the simplex $S^m \subset \mathbb{R}^m$, it may not seem meaningful to consider taking the gradient ∇. However, since the coefficients b, Ax and ϕ always remain directed within the simplex, the directional derivatives $b\cdot\nabla$, $Ax\cdot\nabla$ and $\phi\cdot\nabla$ always exist.

For convenience we recall here the notion of viscosity solution. The unfamiliar reader is referred to Barles [1] or Crandall, Ishii and Lions [20] for a detailed explanation.

Definition 4.6.2. We say that a continuous function $v\colon [0,T]\times S^m \to \mathbb{R}$ is a viscosity subsolution (resp. supersolution) of (4.32)–(4.33) if $v(0,x) \leq$ (resp. $\geq$) $g(x)$ for all $x \in S^m$, and, for any point $(t,x) \in (0,T] \times S^m$,

$$\frac{\partial \varphi}{\partial t}(t,x) + b(x)\cdot\nabla\varphi(t,x) + \sup_{A\in\mathcal{R}}\big\{Ax\cdot\nabla\varphi(t,x) - f(x,A)\big\} \tag{4.34}$$
$$+\big(\phi(x)\cdot\nabla\varphi(t,x) - \psi(x)\big)\dot\eta_t \leq (\text{resp. } \geq)\ 0$$

for every smooth function $\varphi\colon (0,T]\times S^m \to \mathbb{R}$ such that $v-\varphi$ has a local maximum (resp. local minimum) at the point (t,x). We say that v is a viscosity solution if it is both a viscosity subsolution and a viscosity supersolution.

In the following it will sometimes be useful to restrict ourselves to controls bounded by a fixed constant. For $\theta > 0$, we define $\mathcal{R}_\theta = \{A \in \mathcal{R} : \|A\| \leq \theta\}$ and

$$\mathcal{A}_{\infty,\theta} = \big\{A \in \mathcal{A} \;\big|\; A_s \in \mathcal{R}_\theta \text{ for all } s \in [0,T]\big\}.$$

We then define the modified value function:

$$\kappa^\eta_\theta(t,x) = \inf_{A\in\mathcal{A}_{\infty,\theta}}\left\{\int_0^t f(\pi^{t,x,A}_s, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi^{t,x,A}_s)\,\mathrm{d}\eta_s + g(\pi^{t,x,A}_0)\right\}, \tag{4.35}$$

where $\pi^{t,x,A}$ satisfies (4.23) with $\pi^{t,x,A}_t = x$, and as well as the modified HJB equation:

$$\frac{\partial\kappa^\eta_\theta}{\partial t} + b\cdot\nabla\kappa^\eta_\theta + \sup_{A\in\mathcal{R}_\theta}\big\{Ax\cdot\nabla\kappa^\eta_\theta - f\big\} + \big(\phi\cdot\nabla\kappa^\eta_\theta - \psi\big)\dot\eta = 0, \tag{4.36}$$

$$\kappa^\eta_\theta(0,\cdot) = g. \tag{4.37}$$

Proposition 4.6.3. *The approximate value function κ^η is a viscosity solution of* (4.32)–(4.33)*, and the modified value function κ^η_θ is a viscosity solution of* (4.36)–(4.37)*.*

Proof. We first note that κ^η and κ^η_θ are continuous by Propositions 4.5.2 and 4.5.3. It is also clear that $\kappa^\eta(0,x) = g(x) = \kappa^\eta_\theta(0,x)$ for all $x \in S^m$.

The following argument is adapted from [66, Chapter 4]. Let φ be a smooth function such that $\kappa^\eta - \varphi$ has a local maximum at a point $(t,x) \in (0,T]\times S^m$. Let A be a constant valued control. By Lemma 4.5.1, for any $r \in [0,t)$ with $t-r$ sufficiently small, we have

$$\begin{aligned}
0 &\leq \frac{\kappa^\eta(t,x) - \varphi(t,x) - \kappa^\eta(r,\pi^{t,x,A}_r) + \varphi(r,\pi^{t,x,A}_r)}{t-r}\\
&\leq \frac{1}{t-r}\left(\int_r^t \big(f(\pi^{t,x,A}_s, A) + \psi(\pi^{t,x,A}_s)\dot\eta_s\big)\,\mathrm{d}s - \varphi(t,x) + \varphi(r,\pi^{t,x,A}_r)\right)\\
&\longrightarrow f(x,A) + \psi(x)\dot\eta_t - \frac{\partial\varphi}{\partial t}(t,x) - \big(Ax + b(x) + \phi(x)\dot\eta_t\big)\cdot\nabla\varphi(t,x)
\end{aligned}$$

as $r \to t$. Rearranging, we obtain (4.34), so that κ^η is a viscosity subsolution of (4.32)–(4.33). An identical argument shows that κ^η_θ is a viscosity subsolution of (4.36)–(4.37).

Now let φ be a smooth function such that $\kappa^\eta_\theta - \varphi$ has a local minimum at some point $(t,x) \in (0,T] \times S^m$. Let $\varepsilon > 0$. By Lemma 4.5.1, for any $r \in [0,t)$ with $t-r$ sufficiently small, we can find a control $A^{\varepsilon,r} \in \mathcal{A}_{\infty,\theta}$ such that

$$\begin{aligned}
0 &\geq \kappa^\eta_\theta(t,x) - \varphi(t,x) - \kappa^\eta_\theta(r, \pi_r^{t,x,A^{\varepsilon,r}}) + \varphi(r, \pi_r^{t,x,A^{\varepsilon,r}}) \\
&\geq -\varepsilon(t-r) + \int_r^t \big(f(\pi_s^{t,x,A^{\varepsilon,r}}, A_s^{\varepsilon,r}) + \psi(\pi_s^{t,x,A^{\varepsilon,r}})\dot{\eta}_s \big)\,\mathrm{d}s - \varphi(t,x) + \varphi(r, \pi_r^{t,x,A^{\varepsilon,r}}).
\end{aligned}$$

Rearranging, it follows that

$$\begin{aligned}
-\varepsilon &\leq \frac{1}{t-r}\int_r^t \Bigg(- f(\pi_s^{t,x,A^{\varepsilon,r}}, A_s^{\varepsilon,r}) - \psi(\pi_s^{t,x,A^{\varepsilon,r}})\dot{\eta}_s + \frac{\partial \varphi}{\partial t}(s, \pi_s^{t,x,A^{\varepsilon,r}}) \\
&\qquad + \Big(A_s^{\varepsilon,r}\pi_s^{t,x,A^{\varepsilon,r}} + b(\pi_s^{t,x,A^{\varepsilon,r}}) + \phi(\pi_s^{t,x,A^{\varepsilon,r}})\dot{\eta}_s \Big) \cdot \nabla\varphi(s, \pi_s^{t,x,A^{\varepsilon,r}}) \Bigg)\mathrm{d}s \\
&\leq \frac{1}{t-r}\int_r^t \Bigg(\frac{\partial \varphi}{\partial t}(s, \pi_s^{t,x,A^{\varepsilon,r}}) + b(\pi_s^{t,x,A^{\varepsilon,r}}) \cdot \nabla\varphi(s, \pi_s^{t,x,A^{\varepsilon,r}}) \\
&\qquad + \sup_{A \in \mathcal{R}_\theta} \Big\{ A\pi_s^{t,x,A^{\varepsilon,r}} \cdot \nabla\varphi(s, \pi_s^{t,x,A^{\varepsilon,r}}) - f(\pi_s^{t,x,A^{\varepsilon,r}}, A) \Big\} \\
&\qquad + \Big(\phi(\pi_s^{t,x,A^{\varepsilon,r}}) \cdot \nabla\varphi(s, \pi_s^{t,x,A^{\varepsilon,r}}) - \psi(\pi_s^{t,x,A^{\varepsilon,r}}) \Big)\dot{\eta}_s \Bigg)\mathrm{d}s \\
&\longrightarrow \frac{\partial \varphi}{\partial t}(t,x) + b(x) \cdot \nabla\varphi(t,x) + \sup_{A \in \mathcal{R}_\theta} \big\{ Ax \cdot \nabla\varphi(t,x) - f(x,A) \big\} \\
&\qquad + \big(\phi(x) \cdot \nabla\varphi(t,x) - \psi(x) \big)\dot{\eta}_t
\end{aligned}$$

as $r \to t$. Since the controls $A^{\varepsilon,r}$ depend on r, here we required the continuity of the map $s \mapsto \pi_s^{t,x,A^{\varepsilon,r}}$ at time $s = t$ to be uniform in r, which holds due to the fact that the controls are uniformly bounded (by θ). It was for this reason that we restricted to controls $A \in \mathcal{A}_{\infty,\theta}$.

Taking $\varepsilon \to 0$ we obtain (4.34), so that κ^η_θ is a viscosity supersolution of (4.36)–(4.37).

Since $\mathcal{R}_\theta \subseteq \mathcal{R}$, we immediately have that κ^η_θ is also a viscosity supersolution of (4.32)–(4.33). Since every control $A \in \mathcal{A}$ is bounded (by some $\theta > 0$), it is straightforward to see that $\kappa^\eta(t,x) = \inf_{\theta > 0} \kappa^\eta_\theta(t,x)$ for all $(t,x) \in [0,T] \times S^m$, and we conclude by recalling that the infimum of a family of viscosity supersolutions is itself a viscosity supersolution. □

Notation 4.6.4. We denote by $C^{\uparrow}([0,T]\times S^m;\mathbb{R})$ the space of continuous functions $u\colon [0,T]\times S^m \to \mathbb{R}$ which explode near the boundary of the domain S^m, i.e. such that $\inf_{t\in[0,T]} u(t,x) \to \infty$ as $\min_{1\le j\le m} x_j \to 0$.

Lemma 4.6.5. *Write $\pi_s^{t,x,A} = (\pi_{1,s}^{t,x,A},\dots,\pi_{m,s}^{t,x,A})$ for the solution of the smoothed dynamics* (4.23)*, so that $\pi_{k,t}^{t,x,A} = x_k$ for each $k = 1,\dots,m$. We have the inequalities:*

$$x_k \ge \pi_{k,0}^{t,x,A} \exp\Big(- C\int_0^t \big(1+\|A_s\|\big)\,\mathrm{d}s\Big), \tag{4.38}$$

$$x_k \ge \pi_{k,0}^{t,x,A} \exp\Big(- C\int_0^t \big(1+f(\pi_s^{t,x,A},A_s)\big)\,\mathrm{d}s\Big), \tag{4.39}$$

where the constant C depends on $\|\dot\eta\|_\infty$, f, and on the parameters h^i, $i=1,\dots,m$.

Proof. Recalling the filtering equation (4.13) (in this instance with $\mathbf{Y}$ replaced by η), we have that

$$\frac{\mathrm{d}\pi_{k,s}^{t,x,A}}{\mathrm{d}s} \ge -C\big(1+\|A_s\|\big)\pi_{k,s}^{t,x,A}$$

for some constant C depending on $\|\dot\eta\|_\infty$ and h^i, $i=1,\dots,m$. By Grönwall's inequality, we obtain (4.38). By Assumption 4.4.1 (in particular (4.28)), we then deduce (4.39). □

Corollary 4.6.6. $\kappa^\eta \in C^{\uparrow}([0,T]\times S^m;\mathbb{R})$.

Proof. It follows from (4.39) that, if $x_k \to 0$ for some k, then either $\pi_{k,0}^{t,x,A} \to 0$ or $\int_0^t f(\pi_s^{t,x,A},A_s)\,\mathrm{d}s \to \infty$. By Assumption 4.4.3, in the case that $\pi_{k,0}^{t,x,A}\to 0$ we then also have that $g(\pi_0^{t,x,A}) \to \infty$. Recalling (4.20), and the fact that f and g are both bounded below, in either case it follows that $\inf_{t\in[0,T]}\kappa^\eta(t,x)\to\infty$ as $x_k\to 0$. □

4.7 Comparison results

The previous result shows that the approximate value function explodes near the boundary of the domain S^m. Similarly to Section 2.6, this 'explosive boundary condition' is precisely the extra condition needed to obtain the following comparison result for the corresponding HJB equation.

Proposition 4.7.1. *The approximate value function κ^η is the minimal supersolution of* (4.32)–(4.33) *in the class $C^{\uparrow}([0,T]\times S^m;\mathbb{R})$.*

Proof. Let $u \in C^{\uparrow}([0,T] \times S^m; \mathbb{R})$ be another viscosity supersolution of (4.32)–(4.33), and let K be a compact subset of S^m.

By Lemma 4.4.5, κ^η is locally bounded above, so that $\sup_{(t,x)\in[0,T]\times K} \kappa^\eta(t,x) < \infty$. By (4.19) (with $\mathbf{Y}$ replaced by η) and the assumption that f is bounded below, it follows that

$$\int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + g(\pi_0^{t,x,A}) \geq g(\pi_0^{t,x,A}) - C$$

for some constant C. Recalling (4.22), it follows that there exists a $\lambda > 0$ such that, for $(t,x) \in [0,T] \times K$, we can restrict to controls $A \in \mathcal{A}$ such that

$$g(\pi_0^{t,x,A}) < \lambda. \tag{4.40}$$

For $\delta > 0$, we define the subdomain

$$S_\delta^m := \Big\{ x \in S^m \;\Big|\; \min_{1\leq j\leq m} x_j \geq \delta \Big\}. \tag{4.41}$$

As $g \in C^{\uparrow}(S^m; \mathbb{R})$ (by Assumption 4.4.3), there exists a $\delta > 0$ sufficiently small such that

- $K \subset S_\delta^m$, and
- $g > \lambda$ on $S^m \setminus S_\delta^m$.

In particular, by (4.40), we have that $\pi_0^{t,x,A} \in S_\delta^m$ whenever $x \in K$.

Let $G \in C^{\uparrow}(S^m; \mathbb{R})$ be a locally Lipschitz function such that

- $G \leq g$,
- $G = g$ on S_δ^m, and
- $G > \lambda$ on $S^m \setminus S_\delta^m$.

We define a new value function ν, with initial cost function G, viz.

$$\nu(t,x) = \inf_{A\in\mathcal{A}} \Big\{ \int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + G(\pi_0^{t,x,A}) \Big\}.$$

In particular, ν is a viscosity subsolution of (4.32)–(4.33) (which can be proved exactly as it was for κ^η). Moreover, by Proposition 4.5.2, ν is locally Lipschitz continuous in x, uniformly in $t \in [0,T]$.

As noted above, we have that $\pi_0^{t,x,A} \in S_\delta^m$ for all $(t,x) \in [0,T] \times K$. Since $G = g$ on S_δ^m, it follows that

$$\nu = \kappa^\eta \quad \text{on} \quad [0,T] \times K. \tag{4.42}$$

The asymptotic growth rate of $\nu(t,x)$ as x approaches the boundary of S^m depends on the growth rate of the initial cost G. By taking G to grow sufficiently slowly, we can ensure that ν is dominated asymptotically by u. More precisely, since $u \in C^{\uparrow}([0,T] \times S^m; \mathbb{R})$, we can choose G so that there exists an open set E, compactly contained in S^m, such that

- $K \subset E$, and
- $\nu \leq u$ on $[0,T] \times \partial E$ (where ∂E denotes the boundary of E in S^m).

As $\nu(0,x) = G(x) \leq g(x) \leq u(0,x)$ for all $x \in S^m$, we have that $\nu \leq u$ on $([0,T] \times \partial E) \cup (\{0\} \times \overline{E})$, i.e. on the parabolic boundary of the subdomain $(0,T) \times E$. By Theorem 5.1 in Barles [1], applied with the subsolution ν and the supersolution u, it follows that $\nu \leq u$ on $[0,T] \times E$.

By (4.42), we then have

$$\kappa^\eta = \nu \leq u \quad \text{on} \quad [0,T] \times K.$$

Since K was arbitrary, we conclude that $\kappa^\eta \leq u$ on $[0,T] \times S^m$. □

Proposition 4.7.2. *The modified value function κ^η_θ, as defined in* (4.35)*, is the maximal subsolution of* (4.36)–(4.37)*.*

Proof. Let u be another viscosity subsolution of (4.36)–(4.37), and let K be a compact subset of S^m.

By the same argument as in the proof of Proposition 4.7.1 above, that there exists a $\lambda > 0$ such that, for $(t,x) \in [0,T] \times K$, we can restrict to controls $A \in \mathcal{A}_{\infty,\theta}$ such that

$$g(\pi_0^{t,x,A}) < \lambda. \tag{4.43}$$

For $\delta > 0$, we define the subdomain S^m_δ as in (4.41). As $g \in C^{\uparrow}(S^m; \mathbb{R})$ (by Assumption 4.4.3), there exists a $\delta > 0$ sufficiently small such that

- $K \subset S^m_\delta$, and
- $g > \lambda$ on $S^m \setminus S^m_\delta$.

In particular, by (4.43), we have that $\pi_0^{t,x,A} \in S^m_\delta$ whenever $x \in K$.

Let $G \in C^{\uparrow}(S^m; \mathbb{R})$ be a locally Lipschitz function such that

- $G \geq g$,

- $G = g$ on S^m_δ, and
- $G > \lambda$ on $S^m \setminus S^m_\delta$.

We define a new value function ν_θ, with initial cost function G, viz.

$$\nu_\theta(t,x) = \inf_{A \in \mathcal{A}_{\infty,\theta}} \left\{ \int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + G(\pi_0^{t,x,A}) \right\}.$$

In particular, ν_θ is a viscosity supersolution of (4.36)–(4.37). Moreover, by Proposition 4.5.2, ν_θ is locally Lipschitz continuous in x, uniformly in $t \in [0,T]$.

Since $\pi_0^{t,x,A} \in S^m_\delta$ for all $(t,x) \in [0,T] \times K$, and $G = g$ on S^m_δ, it follows that

$$\nu_\theta = \kappa^\eta_\theta \quad \text{on} \quad [0,T] \times K. \tag{4.44}$$

The asymptotic growth rate of $\nu_\theta(t,x)$ as x approaches the boundary of S^m depends on the growth rate of the initial cost G. By taking G to grow sufficiently quickly, we can ensure that ν_θ asymptotically dominates u.

More precisely, by the inequality (4.38) in Lemma 4.6.5, combined with the fact that the controls $A \in \mathcal{A}_{\infty,\theta}$ are bounded by θ, it follows that $x_k \geq C_\theta \pi_{k,0}^{t,x,A}$ for some constant C_θ. We can therefore choose G so that there exists an open set E, compactly contained in S^m, such that

- $K \subset E$, and
- $\nu_\theta \geq u$ on $[0,T] \times \partial E$.

As $\nu_\theta(0,x) = G(x) \geq g(x) \geq u(0,x)$ for all $x \in S^m$, we have that $\nu_\theta \geq u$ on $([0,T] \times \partial E) \cup (\{0\} \times \overline{E})$, i.e. on the parabolic boundary of $(0,T) \times E$. By Theorem 5.1 in Barles [1], applied with the subsolution u and the supersolution ν_θ, it follows that $\nu_\theta \geq u$ on $[0,T] \times E$.

By (4.44), we then have

$$\kappa^\eta_\theta = \nu_\theta \geq u \quad \text{on} \quad [0,T] \times K.$$

Since K was arbitrary, we conclude that $\kappa^\eta_\theta \geq u$ on $[0,T] \times S^m$. □

Corollary 4.7.3. *The approximate value function κ^η is the maximal subsolution of* (4.32)–(4.33).

Proof. Let $u \in C^{\uparrow}([0,T] \times S^m; \mathbb{R})$ be another viscosity subsolution of (4.32)–(4.33).

Fix an arbitrary point $(t,x) \in [0,T] \times S^m$, and let $\varepsilon > 0$. Let $A \in \mathcal{A}$ be an ε-optimal control for $\kappa^\eta(t,x)$, so that

$$\int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + g(\pi_0^{t,x,A}) < \kappa^\eta(t,x) + \varepsilon.$$

As A is bounded (by the definition of $\mathcal{A}$), there exists a $\theta > 0$ such that $A \in \mathcal{A}_{\infty,\theta}$.

Since u is a viscosity subsolution of (4.32)–(4.33), it is clear from an inspection of Definition 4.6.2 that u is also a viscosity subsolution of (4.36)–(4.37). It therefore follows from Proposition 4.7.2 that $u \le \kappa_\theta^\eta$. In particular, since $A \in \mathcal{A}_{\infty,\theta}$, we have that

$$\begin{aligned} u(t,x) &\le \kappa_\theta^\eta(t,x) \\ &\le \int_0^t f(\pi_s^{t,x,A}, A_s)\,\mathrm{d}s + \int_0^t \psi(\pi_s^{t,x,A})\,\mathrm{d}\eta_s + g(\pi_0^{t,x,A}) < \kappa^\eta(t,x) + \varepsilon. \end{aligned}$$

Letting $\varepsilon \to 0$, we deduce the result. □

Combining the results of Proposition 4.7.1 and Corollary 4.7.3, have the following uniqueness result.

Theorem 4.7.4. *The approximate value function κ^η is the unique viscosity solution of (4.32)–(4.33) in the class $C^{\uparrow}([0,T] \times S^m; \mathbb{R})$.*

4.8 A rough HJB equation

Replacing the smooth path η by the rough path $\mathbf{Y}$ in (4.32)–(4.33), we obtain the rough HJB equation:

$$\mathrm{d}\kappa + b \cdot \nabla\kappa\,\mathrm{d}t + \sup_{A \in \mathcal{R}} \big\{ Ax \cdot \nabla\kappa - f \big\}\mathrm{d}t + \big(\phi \cdot \nabla\kappa - \psi\big)\mathrm{d}\mathbf{Y} = 0, \tag{4.45}$$

$$\kappa(0,\cdot) = g. \tag{4.46}$$

We understand a solution to this equation in the sense of the following definition, sometimes known as a rough viscosity solution, used by Diehl et al. [26], as well as for instance by Caruana, Friz and Oberhauser [13, 14, 33] (see also Chapter 12 in Friz and Hairer [32]).

Definition 4.8.1. Given a smooth path η, we write $\boldsymbol{\eta} = (\eta, \eta^{(2)}) \in \mathscr{C}_g^{0,p}$ for its canonical lift, with $\eta^{(2)}$ defined as in (4.24). We write κ^η for the unique viscosity solution of (4.32)–(4.33) in the class $C^{\uparrow}([0,T] \times S^m; \mathbb{R})$, which by Theorem 4.7.4

is precisely the approximate value function, as defined in (4.22). We say that a continuous function v solves the rough HJB equation (4.45)–(4.46) if

$$\kappa^{\eta^n} \longrightarrow v \qquad \text{as} \qquad n \longrightarrow \infty$$

locally uniformly on $[0,T] \times S^m$, whenever $(\eta^n)_{n \geq 1}$ is a sequence of smooth paths such that $\boldsymbol{\eta}^n \to \mathbf{Y}$ with respect to the p-variation rough path distance, i.e. $\|\boldsymbol{\eta}^n; \mathbf{Y}\|_p \to 0$ as $n \to \infty$.

We note that if such a solution of (4.45)–(4.46) exists, then it is unique.

Theorem 4.8.2. *The value function κ, as defined in* (4.16)*, solves the rough HJB equation* (4.45)–(4.46) *in the sense of Definition 4.8.1. Moreover, writing $\kappa = \kappa^{\mathbf{Y}}$, the map from $\mathscr{C}_g^{0,p} \to \mathbb{R}$ given by $\mathbf{Y} \mapsto \kappa^{\mathbf{Y}}(t,x)$ is locally uniformly continuous with respect to the p-variation rough path distance, locally uniformly in (t,x).*

Proof. We first note that κ is continuous by Propositions 4.5.2 and 4.5.3.

Let K be a compact subset of S^m. Let $\mathbf{Z} \in \mathscr{C}^p$ be another rough path such that $\|\mathbf{Y}; \mathbf{Z}\|_p \leq 1$. By possibly replacing L by $L+1$, we may assume that $|||\mathbf{Z}|||_p \leq L$. Let us write $\pi^{t,x,A,\mathbf{Y}}$ (resp. $\pi^{t,x,A,\mathbf{Z}}$) for the solution of the RDE (4.17) driven by $\mathbf{Y}$ (resp. $\mathbf{Z}$), and write $\kappa^{\mathbf{Y}}$ (resp. $\kappa^{\mathbf{Z}}$) for the corresponding value function, as defined in (4.16).

By Corollary 4.4.6, we can restrict to controls $A \in \mathcal{A}_M$ for some $M > 0$, so that $\int_0^T (1 + \|A_s\|)\, \mathrm{d}s \leq M$. We then deduce from part (iii) of Theorem 3.4.3, in view of Remark 3.4.4, that

$$\begin{aligned} \|\pi^{t,x,A,\mathbf{Y}} - \pi^{t,x,A,\mathbf{Z}}\|_{p;[0,t]} &\lesssim \|\mathbf{Y}; \mathbf{Z}\|_{p;[0,t]}, \\ \left\| \int_0^{\cdot} \psi(\pi_s^{t,x,A,\mathbf{Y}})\, \mathrm{d}\mathbf{Y}_s - \int_0^{\cdot} \psi(\pi_s^{t,x,A,\mathbf{Z}})\, \mathrm{d}\mathbf{Z}_s \right\|_{p;[0,t]} &\lesssim \|\mathbf{Y}; \mathbf{Z}\|_{p;[0,t]}. \end{aligned}$$

As in the proof of Proposition 4.5.2, for $x \in K$, the trajectories $\pi^{t,x,A,\mathbf{Y}}$, $\pi^{t,x,A,\mathbf{Z}}$ are bounded away from the boundary of S^m. We can therefore take f and g to be Lipschitz in π. Then, for any $(t,x) \in [0,T] \times K$,

$$\begin{aligned} &\left|\kappa^{\mathbf{Y}}(t,x) - \kappa^{\mathbf{Z}}(t,x)\right| \\ &\leq \sup_{A \in \mathcal{A}_M} \left| \int_0^t \left(f(\pi_s^{t,x,A,\mathbf{Y}}, A_s) - f(\pi_s^{t,x,A,\mathbf{Z}}, A_s)\right) \mathrm{d}s \right. \\ &\qquad \left. + \int_0^t \psi(\pi_s^{t,x,A,\mathbf{Y}})\, \mathrm{d}\mathbf{Y}_s - \int_0^t \psi(\pi_s^{t,x,A,\mathbf{Z}})\, \mathrm{d}\mathbf{Z}_s + g(\pi_0^{t,x,A,\mathbf{Y}}) - g(\pi_0^{t,x,A,\mathbf{Z}}) \right| \\ &\lesssim \sup_{A \in \mathcal{A}_M} \left\{ \int_0^t |\pi_s^{t,x,A,\mathbf{Y}} - \pi_s^{t,x,A,\mathbf{Z}}|\, \mathrm{d}s + \|\mathbf{Y}; \mathbf{Z}\|_{p;[0,t]} + |\pi_0^{t,x,A,\mathbf{Y}} - \pi_0^{t,x,A,\mathbf{Z}}| \right\} \\ &\lesssim \|\mathbf{Y}; \mathbf{Z}\|_{p;[0,t]}. \end{aligned}$$

Taking a sequence of smooth paths $(\eta^n)_{n\geq 1}$ such that $\|\boldsymbol{\eta}^n; \mathbf{Y}\|_p \to 0$, the required convergence follows by taking $\mathbf{Z} = \boldsymbol{\eta}^n$ in the above. The stated continuity of the value function with respect to the driving rough path is also immediate from the above. $\square$

Having established the value function κ as the unique solution of the rough PDE (4.45)–(4.46), one can look to compute κ with a numerical scheme similar to the one given in Section 2.7. The derivation and application of the corresponding numerical scheme is left to future work.

Chapter 5

Pathwise Optimal Control

As we have seen, our approach to robust stochastic filtering leads to a pathwise stochastic control problem, which in the previous chapter was formulated as the optimal control of a rough differential equation. In this chapter we shall examine this type of problem in a general setting. The ideas developed in this chapter will then be applied to our robust filtering approach in Chapter 6.

5.1 Degeneracy of the naive formulation

Example 5.1.1 (Insider trading)**.** Let us suppose that an agent is trading a stock with price composed of a diffusive term with volatility σ, representing usual market uncertainty, and a deterministic path ζ which represents some additional information known only to the agent. We suppose that $\zeta\colon [0,T] \to \mathbb{R}$ is a continuous path which has infinite variation on any interval. Denoting the size of the agent's investment by γ and their wealth process by X, we have the controlled dynamics[1]

$$\mathrm{d}X_s^{t,x,\gamma} = \gamma_s(\sigma\,\mathrm{d}W_s + \mathrm{d}\zeta_s), \qquad s \in [t,T], \tag{5.1}$$

with $X_t^{t,x,\gamma} = x$, where W is a standard Brownian motion and x is the agent's initial wealth. We assume that at each time the agent can only hold a finite amount of stock. More precisely, we impose that γ takes values in the finite interval $[-\varepsilon,\varepsilon]$, for some $\varepsilon > 0$. The agent's expected terminal wealth is given by the value function

$$v(t,x) = \sup_{\gamma} \mathbb{E}\big[X_T^{t,x,\gamma}\big], \tag{5.2}$$

[1]Let us suppose for the moment that we have employed a suitable notion of integration such that the integral against ζ is well-posed.

where the supremum is taken over the collection of progressively measurable $[-\varepsilon, \varepsilon]$-valued processes. In order to study the nature of this problem, let us approximate ζ by a smooth function η. The dynamics (5.1) are then approximated by

$$\mathrm{d}X_s^{t,x,\gamma,\eta} = \gamma_s(\sigma\,\mathrm{d}W_s + \dot{\eta}_s\,\mathrm{d}s),$$

where $\dot{\eta}$ denotes the derivative of η. The resulting control problem is classical. The associated HJB equation is given by

$$\frac{\partial v^\eta}{\partial t} + \sup_{\gamma\in[-\varepsilon,\varepsilon]}\left\{\frac{1}{2}\gamma^2\sigma^2\frac{\partial^2 v^\eta}{\partial x^2} + \gamma\dot{\eta}_t\frac{\partial v^\eta}{\partial x}\right\} = 0,$$

with terminal condition $v^\eta(T,x) = x$. The solution v^η is seen to be

$$v^\eta(t,x) = x + \varepsilon\int_t^T |\dot{\eta}_s|\,\mathrm{d}s, \tag{5.3}$$

and we infer that the optimal control γ^* is given by

$$\gamma_t^* = \varepsilon\,\mathrm{sgn}(\dot{\eta}_t). \tag{5.4}$$

Suppose now that we were to repeatedly refine the approximation η so that it better captures the fast fluctuations of ζ. Even without giving a precise definition of what we might mean by the limit as $\eta \to \zeta$, it is clear that the solution (5.3) should diverge to infinity in this limit. In other words, the original value function, as defined in (5.2), should be simply given by $v(t,x) = \infty$ whenever $t < T$. Thus, we infer that our original control problem, with the infinite variation signal ζ, is degenerate.

The phenomenon exhibited here is typical for such control problems, where we attempt to control the coefficient of the infinite variation term in the controlled dynamics. The problem in the previous example is that, in contrast to a classical stochastic setting, since the controller can 'see' the path ζ in advance, they can choose controls γ with very small, but extremely quick, fluctuations, which allow the solution X to take full advantage of the infinite variation of ζ. Indeed, notice that the sign of the optimal control in (5.4) changes at the same rate as that of $\dot{\eta}$, which varies 'infinitely quickly' in the limit as $\eta \to \zeta$.

To paraphrase Diehl et al. [26], if the coefficient of the driving signal has enough dependence on the control, and this signal has unbounded variation on any interval, then the controller can drive the solution to reach any point instantly whilst incurring an arbitrarily low cost.

In this chapter we investigate this degeneracy phenomenon in more detail, and see how it may be resolved by introducing an artificial cost to penalise the variation of the controls. We will see how this cost may be chosen to ensure that a dynamic programming principle is retained, thus allowing one to recover a setting comparable to that of [26]. As an extension of that paper, we proceed to consider pathwise control with unbounded cost functions and, by obtaining locally uniform bounds on the controls, establish the value function as the unique solution of a rough HJB equation.

5.2 Avoiding degeneracy

Our set-up is the following. We let $b \in \mathrm{Lip}_b$, $\phi \in C_b^3$ and $T > 0$. We fix a geometric rough path $\boldsymbol{\zeta} = (\zeta, \zeta^{(2)}) \in \mathscr{C}_g^{0,p}([0,T];\mathbb{R}^d)$ such that $|||\boldsymbol{\zeta}|||_p \leq L$ for some $p \in [2,3)$ and $L > 0$. We consider, for each $\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}([0,T];\mathbb{R}^k)$, the controlled dynamics

$$\mathrm{d}X_s^{t,x,\gamma} = b(X_s^{t,x,\gamma}, \gamma_s)\,\mathrm{d}s + \phi(X_s^{t,x,\gamma}, \gamma_s)\,\mathrm{d}\boldsymbol{\zeta}_s, \qquad X_t^{t,x,\gamma} = x, \tag{5.5}$$

driven by $\boldsymbol{\zeta}$, which has a unique solution $(X^{t,x,\gamma}, \phi(X^{t,x,\gamma}, \gamma)) \in \mathscr{D}_\zeta^p$ by part (i) of Theorem 3.4.3. We then consider the control problem with value function given by

$$v(t,x) := \inf_{\gamma \in \mathcal{C}^{\frac{p}{2}\text{-var}}} J(t,x;\gamma) \tag{5.6}$$

for $(t,x) \in [0,T] \times \mathbb{R}^m$, where the cost functional J is defined as

$$J(t,x;\gamma) := \int_t^T f(X_s^{t,x,\gamma}, \gamma_s)\,\mathrm{d}s + \int_t^T \psi(X_s^{t,x,\gamma}, \gamma_s)\,\mathrm{d}\boldsymbol{\zeta}_s + g(X_T^{t,x,\gamma}). \tag{5.7}$$

Here $f\colon \mathbb{R}^m \times \mathbb{R}^k \to \mathbb{R}$, $\psi\colon \mathbb{R}^m \times \mathbb{R}^k \to \mathcal{L}(\mathbb{R}^d;\mathbb{R})$ and $g\colon \mathbb{R}^m \to \mathbb{R}$.

Let us assume that $\psi \in C_b^3$. By part (ii) of Theorem 3.4.3, for any t, x and γ, we have that

$$\left| \int_t^T \psi(X_s^{t,x,\gamma}, \gamma_s)\,\mathrm{d}\boldsymbol{\zeta}_s \right| \leq C\Big(1 + \|\gamma\|_{\frac{p}{2};[t,T]}^{\frac{p-1}{2}}\Big), \tag{5.8}$$

where the constant C depends only on b, ϕ, ψ, p, T and L.

Remark 5.2.1. The choice to measure the regularity of the controls using q-variation for $q = \frac{p}{2}$ ensures that $\frac{1}{p} + \frac{1}{q} > 1$, so that the corresponding Young integral of γ against ζ always exists. It may be tempting to wonder whether a bound of the form in (5.8) could still hold based on the q-variation of controls for a larger value of q. However, this is not true in general. Indeed, if $\frac{1}{p} + \frac{1}{q} < 1$, then one can construct a sequence $\{(\zeta^n, \gamma^n)\}_{n \geq 1}$ of pairs of bounded variation paths such that $\|\zeta^n\|_p = 1 = \|\gamma^n\|_q$ for all $n \geq 1$, but such that $\int_0^T \gamma_s^n\,\mathrm{d}\zeta_s^n \to \infty$ as $n \to \infty$, which would contradict (5.8).

Preventing degeneracy of this control problem can essentially be thought of as preventing the size of the rough integral above from becoming arbitrarily large. The estimate (5.8) shows that one can control the size of this integral by the $\frac{p}{2}$-variation of the controls. However, since controls can exhibit arbitrarily large $\frac{p}{2}$-variation whilst remaining uniformly bounded, the cost functional in (5.7) is not able to adequately penalise this variation. In view of Example 5.1.1, for a typical choice of ψ, one should expect the value function in (5.6) to be simply given by

$$v(t,x) = -\infty \qquad \text{for all} \quad (t,x) \in [0,T) \times \mathbb{R}^m.$$

We also point out that merely restricting the class of controls γ to, say, smooth functions does nothing to resolve this problem.

The estimate (5.8) implies that one could prevent degeneracy by imposing a uniform bound on the $\frac{p}{2}$-variation of the controls but, as appreciated in Diehl et al. [26], this would not be a very natural condition. Instead, we first propose to introduce an artificial cost in order to penalise this variation.

Definition 5.2.2. Let $\mathcal{S} \subseteq \mathcal{C}^{\frac{p}{2}\text{-var}}$ be a Banach space of functions from $[0,T] \to \mathbb{R}^k$ (with a possibly stronger topology). We shall call a function $\beta\colon \Delta_{[0,T]} \times \mathcal{C}^{\frac{p}{2}\text{-var}} \to \mathbb{R} \cup \{+\infty\}$ a *regularising cost on* $\mathcal{S}$, if it is bounded below, takes the value $+\infty$ on $\Delta_{[0,T]} \times (\mathcal{C}^{\frac{p}{2}\text{-var}} \setminus \mathcal{S})$, and, for every $0 \le r < t \le T$, the map $\beta_{r,t}\colon \mathcal{S} \to \mathbb{R}$ is continuous, and satisfies

$$\frac{\beta_{r,t}(\gamma)}{\|\gamma\|^{\frac{p-1}{2}}_{\frac{p}{2};[r,t]}} \longrightarrow \infty \qquad \text{as} \qquad \|\gamma\|_{\frac{p}{2};[r,t]} \longrightarrow \infty. \tag{5.9}$$

An example of such a cost on $\mathcal{C}^{\frac{p}{2}\text{-var}}$ is given by

$$\beta_{r,t}(\gamma) = \varepsilon \|\gamma\|^q_{\frac{p}{2};[r,t]} \tag{5.10}$$

for any $\varepsilon > 0$ and $q > \frac{p-1}{2}$.

Instead of the naive value function in (5.6), we consider the modified function given by[2]

$$V(t,x) := \inf_{\gamma \in \mathcal{C}^{0,\frac{p}{2}\text{-var}}} \big\{ J(t,x;\gamma) + \beta_{t,T}(\gamma) \big\}, \tag{5.11}$$

for some regularising cost β. In practice, the justification of the introduction of this 'artificial cost' depends on the application one has in mind; we will see examples of this later in Sections 5.7 and 6.4.

The following proposition demonstrates the nondegeneracy of this modified control problem.

[2]The restriction to controls $\gamma \in \mathcal{C}^{0,\frac{p}{2}\text{-var}} \subset \mathcal{C}^{\frac{p}{2}\text{-var}}$ is negligible. Indeed, we recall that $\mathcal{C}^{\frac{q}{2}\text{-var}} \subset \mathcal{C}^{0,\frac{p}{2}\text{-var}}$ for any $q \in [2,p)$.

Proposition 5.2.3. *Under the natural assumption that f and g are bounded below, the same is true of the value function V.*

Proof. It follows from (5.8) and (5.9) that

$$\left| \int_t^T \psi(X_s^{t,x,\gamma}, \gamma_s) \, \mathrm{d}\boldsymbol{\zeta}_s \right| \leq C + \frac{\beta_{t,T}(\gamma)}{2}$$

for some new constant C, and hence that

$$J(t,x;\gamma) + \beta_{t,T}(\gamma) \geq \int_t^T f(X_s^{t,x,\gamma}, \gamma_s) \, \mathrm{d}s + g(X_T^{t,x,\gamma}) + \frac{\beta_{t,T}(\gamma)}{2} - C.$$

Since the cost functions f, g and β are all bounded below, the result follows. □

5.3 Recovering dynamic programming

We have seen that one can resolve the degeneracy of the optimal control problem by introducing an artificial cost to penalise the variation of the controls. In Definition 5.2.2 we introduced a rather general class of cost functions which provide a sufficient penalisation. The problem with such cost functions, such as the one in (5.10), is that typically they are not additive, in the sense that $\beta_{r,s} + \beta_{s,t} \neq \beta_{r,t}$. A consequence of this is that the corresponding control problem is no longer dynamic. That is, the value function in (5.11) is not generally amenable to dynamic programming, and thus one cannot necessarily write down a PDE associated with the control problem. Our next aim will be to demonstrate the existence of an additive regularising cost on a more regular space of controls, which allows dynamic programming to be recovered.

Lemma 5.3.1. *Let β be a regularising cost on $\mathcal{C}^{0,\frac{p}{2}\text{-var}}$. Let $\mathcal{S} \subseteq \mathcal{C}^{0,\frac{p}{2}\text{-var}}$ be a subset which contains all smooth functions from $[0,T] \to \mathbb{R}^k$. Then the value function defined in* (5.11) *satisfies*

$$V(t,x) = \inf_{\gamma \in \mathcal{S}} \big\{ J(t,x;\gamma) + \beta_{t,T}(\gamma) \big\}. \tag{5.12}$$

Proof. By definition, for any $\gamma \in \mathcal{C}^{0,\frac{p}{2}\text{-var}}$, there exists a sequence of smooth controls $\{\gamma^n\}_{n \geq 1}$ such that $\|\gamma^n - \gamma\|_\infty + \|\gamma^n - \gamma\|_{\frac{p}{2}} \to 0$ as $n \to \infty$. The result then follows from the continuity of $\beta_{t,T}$ and the stability estimates in part (iii) of Theorem 3.4.3. □

In particular, the equality (5.12) holds with $\mathcal{S} = W^{1,q}$ for any $q \geq 1$, where $W^{1,q} = W^{1,q}([0,T];\mathbb{R}^k)$ denotes the usual Sobolev space. We recall the continuous

embeddings $W^{1,q} \hookrightarrow \mathcal{C}^{1\text{-var}} \hookrightarrow \mathcal{C}^{0,\frac{p}{2}\text{-var}}$, exhibited by the inequalities

$$T^{\frac{q-1}{q}} \left(\int_r^t |\dot{\gamma}_s|^q \, \mathrm{d}s \right)^{\frac{1}{q}} \geq \|\gamma\|_{1;[r,t]} \geq \|\gamma\|_{\frac{p}{2};[r,t]},$$

where we write $\dot{\gamma}$ for the unique element $\dot{\gamma} \in L^q([0,T];\mathbb{R}^k)$ such that $\mathrm{d}\gamma_s = \dot{\gamma}_s \, \mathrm{d}s$. It follows that, for any $\varepsilon > 0$ and $q > \frac{p-1}{2}$, the choice

$$\beta_{r,t}(\gamma) = \varepsilon \int_r^t |\dot{\gamma}_s|^q \, \mathrm{d}s$$

for $\gamma \in W^{1,q}$ (and $\beta(\gamma) \equiv \infty$ otherwise), defines a regularising cost on $W^{1,q}$. Moreover, β is additive, in the sense that $\beta_{r,s} + \beta_{s,t} = \beta_{r,t}$ for all $r \leq s \leq t$; in other words, for each $\gamma \in W^{1,q}$, the two-parameter functional $\beta(\gamma) \colon \Delta_{[0,T]} \to \mathbb{R}$ is uniquely characterised by the path $t \mapsto \beta_{0,t}(\gamma)$. With this choice of β, we can now write

$$V(t,x) = \inf_{a \in \mathbb{R}^k} v(t,x,a)$$

where, for $(t,x,a) \in [0,T] \times \mathbb{R}^m \times \mathbb{R}^k$,

$$v(t,x,a) := \inf_{u \in L^q} \left\{ J(t,x;\gamma^{t,a,u}) + \varepsilon \int_t^T |u_s|^q \, \mathrm{d}s \right\}, \tag{5.13}$$

with $\gamma_r^{t,a,u} := a + \int_t^r u_s \, \mathrm{d}s$ for $r \in [t,T]$. The function v is both nondegenerate, and satisfies the following dynamic programming principle:

Proposition 5.3.2. *Let us write $X^{t,x,a,u} := X^{t,x,\gamma^{t,a,u}}$. Then, for any t, x, a and $r \in [t,T]$, with v as in* (5.13)*, we have*

$$v(t,x,a) = \inf_{u \in L^q} \left\{ v(r, X_r^{t,x,a,u}, \gamma_r^{t,a,u}) + \int_t^r f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}) \, \mathrm{d}s \right.$$
$$\left. + \int_t^r \psi(X_s^{t,x,a,u}, \gamma_s^{t,a,u}) \, \mathrm{d}\boldsymbol{\zeta}_s + \varepsilon \int_t^r |u_s|^q \, \mathrm{d}s \right\}.$$

This result follows the same proof as that of Theorem 2.1 in [66, Chapter 4]. In particular, the rough integrals appearing in the controlled dynamics and value function do not cause any additional difficulty.

5.4 A generalised dynamic control problem

To summarise the previous sections, we propose to reformulate the naive control problem, given originally by (5.5)–(5.7), to resolve the degeneracy problem whilst retaining enough dynamic structure to retain dynamic programming, by restricting

to a sufficiently regular space of controls, and introducing an additive artificial cost function, written in terms of the derivative of the controls. Rather than merely (5.5), by including γ as part of the state trajectory, we instead consider the controlled dynamics

$$\mathrm{d}X_s^{t,x,a,u} = b(X_s^{t,x,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}s + \phi(X_s^{t,x,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}\boldsymbol{\zeta}_s, \qquad X_t^{t,x,a,u} = x, \tag{5.14}$$

$$\mathrm{d}\gamma_s^{t,a,u} = h(\gamma_s^{t,a,u}, u_s)\,\mathrm{d}s, \qquad \gamma_t^{t,a,u} = a. \tag{5.15}$$

For generality, we have introduced the function $h\colon \mathbb{R}^k \times U \to \mathbb{R}^k$, where here $(U, \|\cdot\|_U)$ is a finite dimensional Banach space, and the control u belongs to the space $\mathcal{U}$ of bounded measurable functions $u\colon [0,T] \to U$.

We shall henceforth consider the cost functional

$$J(t,x,a;u) := \int_t^T f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s + \int_t^T \psi(X_s^{t,x,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}\boldsymbol{\zeta}_s + g(X_T^{t,x,a,u}, \gamma_T^{t,a,u})$$

and the value function

$$v(t,x,a) := \inf_{u\in\mathcal{U}} J(t,x,a;u), \tag{5.16}$$

where we have absorbed a regularising cost into the function $f\colon \mathbb{R}^m \times \mathbb{R}^k \times U \to \mathbb{R}$, which crucially is now also allowed to depend on u. There is also no harm in allowing the terminal cost $g\colon \mathbb{R}^m \times \mathbb{R}^k \to \mathbb{R}$ to depend on the terminal value of γ.

Remark 5.4.1. If the function $h = h(a,u)$ were bounded in u, then setting $\tilde{X} = (X, \gamma)$ would now put us into a comparable setting to Diehl et al. [26]. However, it is more natural here to allow h to be unbounded in u, meaning that [26, Theorem 5] does not directly apply[3]. Moreover, in [26] the cost functions f and g are assumed to be bounded, but here we will relax this assumption.

The inclusion of the integral $\int \psi(X,\gamma)\,\mathrm{d}\boldsymbol{\zeta}$ in the value function also takes us outside the setting of [26]. This term could be included in the terminal cost by setting $\tilde{X} = (X, \gamma, Z)$ and $\tilde{g}(x,a,z) = g(x,a) + z$ with $Z_r = z + \int_t^r \psi(X,\gamma)\,\mathrm{d}\boldsymbol{\zeta}$, albeit with the additional complication that the terminal cost $\tilde{g}$ would then be neither bounded from above nor below.

Assumption 5.4.2. We assume that

- $b \in \mathrm{Lip}_b$ and $\phi, \psi \in C_b^3$,

[3]This boundedness condition is not stated explicitly in [26], but is necessary for the application of [6, Corollary III.3.6] in the proof of [26, Theorem 5]; see Assumption (A_1) in [6, Chapter III].

- $f = f(x,a,u)$ and $g = g(x,a)$ are continuous, bounded below, and Lipschitz continuous in (x,a), uniformly in u,
- $h = h(a,u)$ is continuous, Lipschitz in a, uniformly in u, and is bounded in a, locally uniformly in u, and moreover, for some $\delta \geq 1$, satisfies

$$\sup_{a\in\mathbb{R}^k} \frac{\big|h(a,u)\big|}{\|u\|_U^\delta} \longrightarrow 0 \qquad \text{as} \quad \|u\|_U \longrightarrow \infty, \tag{5.17}$$

- with the same δ as in (5.17), the running cost f satisfies

$$\inf_{x\in\mathbb{R}^m,\, a\in\mathbb{R}^k} \frac{f(x,a,u)}{\|u\|_U^{\frac{\delta(p-1)}{2}}} \longrightarrow \infty \qquad \text{as} \quad \|u\|_U \longrightarrow \infty. \tag{5.18}$$

The following lemma demonstrates the nondegeneracy of our newly formulated control problem.

Lemma 5.4.3. *For any t, x, a and u, we have that*

$$\left| \int_t^T \psi(X_s^{t,x,a,u}, \gamma_s^{t,a,u})\, \mathrm{d}\boldsymbol{\zeta}_s \right| \leq C + \frac{1}{2}\int_t^T f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}, u_s)\, \mathrm{d}s, \tag{5.19}$$

where the constant C depends only on b, ϕ, ψ, h, p, T and L.

Proof. By (5.8), (5.15) and Hölder's inequality, we have that

$$\begin{aligned}
&\left| \int_t^T \psi(X_s^{t,x,a,u}, \gamma_s^{t,a,u})\, \mathrm{d}\boldsymbol{\zeta}_s \right| \leq C\Big(1 + \|\gamma^{t,a,u}\|_{\frac{p}{2};[t,T]}^{\frac{p-1}{2}}\Big) \leq C\Big(1 + \|\gamma^{t,a,u}\|_{1;[t,T]}^{\frac{p-1}{2}}\Big) \\
&= C\bigg(1 + \bigg(\int_t^T \big|h(\gamma_s^{t,a,u}, u_s)\big|\, \mathrm{d}s\bigg)^{\frac{p-1}{2}}\bigg) \leq C\bigg(1 + T^{\frac{p-3}{2}} \int_t^T \big|h(\gamma_s^{t,a,u}, u_s)\big|^{\frac{p-1}{2}}\, \mathrm{d}s\bigg).
\end{aligned}$$

Then, by (5.17) and (5.18) (noting that, since U is finite dimensional, h is uniformly bounded on bounded subsets of U), we can ensure that (5.19) holds for a new constant C. □

Corollary 5.4.4. *Let K be a compact subset of $\mathbb{R}^m \times \mathbb{R}^k$. There exists an $M > 0$ such that, when taking the infimum over $u \in \mathcal{U}$ in (5.16) for $(t,x,a) \in [0,T] \times K$, one may restrict to controls u satisfying $\|\gamma^{t,a,u}\|_{\frac{p}{2}} \leq M$, without changing the value of $v(t,x,a)$.*

Proof. By Lemma 5.4.3 and the assumption that g is bounded below, we have that

$$J(t,x,a;u) \geq \frac{1}{2}\int_t^T f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}, u_s)\, \mathrm{d}s - C$$

for some possibly new constant C. Let $u^* \in \mathcal{U}$ be an arbitrary control. By the above, we may ignore all controls u such that

$$\frac{1}{2}\int_t^T f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s - C > \sup_{(\hat{t},\hat{x},\hat{a})\in[0,T]\times K} J(\hat{t},\hat{x},\hat{a};u^*).$$

This gives an upper bound on $\int_t^T f(X_s^{t,x,a,u}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s$, which we observe, by the proof of Lemma 5.4.3, also implies an upper bound on $\|\gamma^{t,a,u}\|_{\frac{p}{2}}$. □

5.5 A smooth noise approximation

Although a dynamic programming principle of the form in Proposition 5.3.2 holds for the value function v in (5.16), the appearance of the rough integrals makes it less straightforward to derive a PDE directly from this result. As in Example 5.1.1, we will therefore proceed by first approximating ζ by a smooth function η. We then define the corresponding approximate control problem, with dynamics

$$\mathrm{d}X_s^{t,x,a,u,\eta} = b(X_s^{t,x,a,u,\eta}, \gamma_s^{t,a,u})\,\mathrm{d}s + \phi(X_s^{t,x,a,u,\eta}, \gamma_s^{t,a,u})\,\mathrm{d}\eta_s, \qquad X_t^{t,x,a,u,\eta} = x, \quad (5.20)$$

where $\gamma^{t,a,u}$ satisfies (5.15). Naturally equation (5.20) has a unique C^1 solution. However, in the following it will be useful to also embed this solution in rough path space. As η is smooth, we can simply enhance it with its iterated integrals in the classical Lebesgue–Stieltjes sense,

$$\eta_{s,t}^{(2)} := \int_s^t \eta_{s,r} \otimes \mathrm{d}\eta_r, \quad (5.21)$$

so that $\boldsymbol{\eta} = (\eta, \eta^{(2)})$ is itself a rough path. Any continuous path with finite p-variation would make a valid candidate for the Gubinelli derivative of X^η (with respect to η), but to be consistent with the genuinely rough case above we insist on the choice $(X^\eta)' = \phi(X^\eta, \gamma)$. We can then consider $(X^\eta, \phi(X^\eta, \gamma))$ as the solution of (5.20) in the sense of Theorem 3.4.3.

We also define the corresponding approximate value function v^η as

$$v^\eta(t,x,a) := \inf_{u\in\mathcal{U}} \Bigg\{ \int_t^T f(X_s^{t,x,a,u,\eta}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s + \int_t^T \psi(X_s^{t,x,a,u,\eta}, \gamma_s^{t,a,u})\,\mathrm{d}\eta_s + g(X_T^{t,x,a,u,\eta}, \gamma_T^{t,a,u}) \Bigg\}.$$

Writing $\dot{\eta}$ for the derivative of η, under Assumption 5.4.2, we can apply Theorem 3.2 in Bardi and Da Lio [7], to obtain that v^η is the unique viscosity solution of the HJB equation

$$-\frac{\partial v^\eta}{\partial t}(t,x,a) - b(x,a)\cdot\nabla_x v^\eta(t,x,a) - \inf_{u\in U}\big\{h(a,u)\cdot\nabla_a v^\eta(t,x,a) + f(x,a,u)\big\} \\ - \big(\phi(x,a)\cdot\nabla_x v^\eta(t,x,a) + \psi(x,a)\big)\dot{\eta}_t = 0, \tag{5.22}$$

with the terminal condition

$$v^\eta(T,x,a) = g(x,a). \tag{5.23}$$

Moreover, by Theorem 2.2 in [7], this solution is locally Lipschitz continuous.

5.6 A rough HJB equation

Replacing η in (5.22) with $\boldsymbol{\zeta}$, we formally derive the rough PDE given by

$$-\mathrm{d}v - b\cdot\nabla_x v\,\mathrm{d}t - \inf_{u\in U}\big\{h\cdot\nabla_a v + f\big\}\mathrm{d}t - \big(\phi\cdot\nabla_x v + \psi\big)\mathrm{d}\boldsymbol{\zeta} = 0, \tag{5.24}$$

with

$$v(T,x,a) = g(x,a). \tag{5.25}$$

We point out that as written equation (5.24) is only formal, and is given a precise meaning in Definition 5.6.1 below; cf. Definition 4.8.1.

Definition 5.6.1. For any smooth function $\eta\colon [0,T]\to\mathbb{R}^d$, write v^η for the unique viscosity solution of (5.22) and (5.23). Moreover, write $\boldsymbol{\eta}$ for the rough path obtained by enhancing η with its iterated integrals in the Lebesgue–Stieltjes sense, as in (5.21). We say that a continuous function v solves (5.24) and (5.25) if

$$v^{\eta^n} \longrightarrow v \qquad \text{as} \qquad n\longrightarrow\infty$$

locally uniformly on $[0,T]\times\mathbb{R}^m\times\mathbb{R}^k$, whenever $(\eta^n)_{n\geq 1}$ is a sequence of smooth paths such that $\boldsymbol{\eta}^n\to\boldsymbol{\zeta}$ with respect to the p-variation rough path distance, i.e. $\|\boldsymbol{\eta}^n;\boldsymbol{\zeta}\|_p\to 0$ as $n\to\infty$.

Note that uniqueness of such a solution is built into the definition. Moreover, note that since we assumed that $\boldsymbol{\zeta}$ is a geometric rough path, there certainly exists such a sequence of smooth paths $(\eta^n)_{n\geq 1}$.

Theorem 5.6.2. *Under Assumption 5.4.2, the value function v defined in* (5.16) *solves* (5.24) *and* (5.25) *in the sense of Definition 5.6.1. Moreover, writing $v = v^\zeta$, the map from $\mathscr{C}_g^{0,p}([0,T];\mathbb{R}^d) \to \mathbb{R}$ given by $\boldsymbol{\zeta} \mapsto v^\zeta(t,x,a)$ is locally uniformly continuous with respect to the p-variation rough path distance, locally uniformly in (t,x,a).*

Proof. Let K be a compact subset of $\mathbb{R}^m \times \mathbb{R}^k$ and let $\boldsymbol{\eta} \in \mathscr{C}^p$ be another rough path such that $\|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p \leq 1$. By possibly replacing L by $L+1$, we may assume that $|||\boldsymbol{\eta}|||_p \leq L$. Let us write $X^\eta = X^{t,x,a,u,\eta}$ (resp. $X^\zeta = X^{t,x,a,u,\zeta}$) for the solution of the RDE (5.14) driven by $\boldsymbol{\eta}$ (resp. $\boldsymbol{\zeta}$), and write v^η (resp. v^ζ) for the corresponding value function, as defined in (5.16).

By Corollary 5.4.4, there exists an $M > 0$ such that, for $(t,x,a) \in [0,T] \times K$, we may restrict to controls $u \in \mathcal{U}^M \subseteq \mathcal{U}$ satisfying $\|\gamma^{t,a,u}\|_{\frac{p}{2}} \leq M$, so that in particular the hypotheses of part (iii) of Theorem 3.4.3 are satisfied.

In the following we shall use $\lesssim$ to denote inequality up to a multiplicative constant which may depend on $b, \phi, \psi, f, g, h, p, T, L$ and M. It follows from part (iii) of Theorem 3.4.3 that

$$\|X^\eta - X^\zeta\|_\infty \lesssim \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p,$$

and

$$\left\| \int_t^\cdot \psi(X_s^\eta, \gamma_s)\, \mathrm{d}\boldsymbol{\eta}_s - \int_t^\cdot \psi(X_s^\zeta, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s \right\|_\infty \lesssim \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p.$$

By the Lipschitz assumptions on f and g, for any $(t,x,a) \in [0,T] \times K$, we have

$$\begin{aligned}
&\left|v^\eta(t,x,a) - v^\zeta(t,x,a)\right| \\
&\quad\leq \sup_{u\in\mathcal{U}^M} \Bigg| \int_t^T \big(f(X_s^{u,\eta}, \gamma_s^u, u_s) - f(X_s^{u,\zeta}, \gamma_s^u, u_s)\big)\, \mathrm{d}s \\
&\qquad\qquad + \int_t^T \psi(X_s^{u,\eta}, \gamma_s^u)\, \mathrm{d}\boldsymbol{\eta}_s - \int_t^T \psi(X_s^{u,\zeta}, \gamma_s^u)\, \mathrm{d}\boldsymbol{\zeta}_s + g(X_T^{u,\eta}, \gamma_T^u) - g(X_T^{u,\zeta}, \gamma_T^u) \Bigg| \\
&\quad\lesssim \sup_{u\in\mathcal{U}^M} \left(\int_t^T |X_s^{u,\eta} - X_s^{u,\zeta}|\, \mathrm{d}s + \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p + |X_T^{u,\eta} - X_T^{u,\zeta}| \right) \\
&\quad\lesssim \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p.
\end{aligned}$$

Taking a sequence of smooth paths $(\eta^n)_{n\geq 1}$ such that $\|\boldsymbol{\eta}^n;\boldsymbol{\zeta}\|_p \to 0$, the required convergence follows by taking $\boldsymbol{\eta} = \boldsymbol{\eta}^n$ in the above. Since the approximate value functions v^{η^n} are continuous, continuity of the function v with respect to (t,x,a) also follows from this convergence. The stated continuity of the value function with respect to the driving rough path is also immediate from the above. □

Remark 5.6.3. One could also introduce another Brownian motion W and consider as controlled dynamics the hybrid Itô-rough differential equation $\mathrm{d}X = b(X,\gamma)\,\mathrm{d}s + \sigma(X,\gamma)\,\mathrm{d}W + \phi(X,\gamma)\,\mathrm{d}\boldsymbol{\zeta}$. Just as in the classical case, the value function is then defined as the infimum (or supremum) over adapted controls of an expected cost function, and the associated HJB equation is then of second order; see Example 5.7.2 below.

5.7 Examples

Example 5.7.1. When a Brownian motion W is enhanced with its iterated integrals in the sense of Stratonovich integration, i.e.

$$W^{(2)}_{s,t} := \int_s^t W_{s,r} \otimes \circ\,\mathrm{d}W_r, \tag{5.26}$$

then, almost surely, $\mathbf{W} = (W, W^{(2)})$ defines a geometric p-rough path for any $p \in (2,3)$. The choice $\boldsymbol{\zeta} = \mathbf{W}$ thus leads to the stochastic PDE

$$-\mathrm{d}v - b\cdot\nabla_x v\,\mathrm{d}t - \inf_{u\in U}\big\{h\cdot\nabla_a v + f\big\}\mathrm{d}t - \big(\phi\cdot\nabla_x v + \psi\big)\circ\mathrm{d}\mathbf{W} = 0.$$

Example 5.7.2 (Insider trading revisited)**.** Let us return to the setting of Example 5.1.1, where we recall that an agent is trading a stock with the benefit of some extra information not available to the rest of the market. We denote the agent's initial investment by a, and the rate at which they purchase new stock by u. The dynamics of the agent's wealth process X and investment γ are given by[4]

$$\begin{aligned} \mathrm{d}X^{t,x,a,u}_s &= \gamma^{t,a,u}_s(\sigma\,\mathrm{d}W_s + \mathrm{d}\zeta_s), & X^{t,x,a,u}_t &= x,\\ \mathrm{d}\gamma^{t,a,u}_s &= u_s\,\mathrm{d}s, & \gamma^{t,a,u}_t &= a, \end{aligned}$$

where W is a Brownian motion and ζ is an arbitrary continuous path. Note that, since ζ is continuous and $\gamma^{t,a,u}$ is of finite variation, the integral $\int_t^{\cdot}\gamma^{t,a,u}_s\,\mathrm{d}\zeta_s$ exists in the Riemann–Stieltjes sense, so there is no need here to lift ζ into rough path space.

Let us suppose that the agent must pay a transaction cost of εu^2. The agent's expected terminal wealth is then given by the value function

$$v(t,x,a) = \sup_{u\in\mathcal{U}}\mathbb{E}\bigg[X^{t,x,a,u}_T - \int_t^T \varepsilon u_s^2\,\mathrm{d}s\bigg],$$

[4]When $\sigma \neq 0$ the inclusion of the Brownian motion takes us outside the class of problems considered above, but there is no conceptual change and we expect all of the analysis to follow with appropriate technical adjustments.

where $\mathcal{U}$ is the space of progressively measurable $\mathbb{R}$-valued processes. In this case the HJB equation (5.24) takes the form

$$-\mathrm{d}v - \frac{1}{2}a^2\sigma^2\frac{\partial^2 v}{\partial x^2}\,\mathrm{d}t - \frac{1}{4\varepsilon}\left(\frac{\partial v}{\partial a}\right)^2\mathrm{d}t - a\frac{\partial v}{\partial x}\,\mathrm{d}\zeta = 0 \tag{5.27}$$

with

$$v(T, x, a) = x. \tag{5.28}$$

Approximating ζ by a smooth function η, we obtain the classical HJB equation

$$-\frac{\partial v^\eta}{\partial t} - \frac{1}{2}a^2\sigma^2\frac{\partial^2 v^\eta}{\partial x^2} - \frac{1}{4\varepsilon}\left(\frac{\partial v^\eta}{\partial a}\right)^2 - a\dot{\eta}\frac{\partial v^\eta}{\partial x} = 0.$$

The solution of this equation along with the terminal condition (5.28) is given by

$$v^\eta(t, x, a) = x + (\eta_T - \eta_t)a + \frac{1}{4\varepsilon}\int_t^T (\eta_T - \eta_s)^2\,\mathrm{d}s.$$

Recalling Definition 5.6.1, we obtain the solution of (5.27) and (5.28) as

$$v(t, x, a) = x + (\zeta_T - \zeta_t)a + \frac{1}{4\varepsilon}\int_t^T (\zeta_T - \zeta_s)^2\,\mathrm{d}s.$$

Note that this quantity remains finite even when ζ is of infinite variation. Thus, an agent, even with perfect knowledge of the future stock price, subject to sufficient transaction costs, can only make a finite profit.

Chapter 6

The Kalman–Bucy Filter under General Uncertainty

In this chapter we turn our attention back to the problem of Kalman–Bucy filtering under model uncertainty. In Chapter 2 we constructed a robust filter, the calculation of which involves the derivation of a pathwise stochastic control problem. In that setting the control terms did not appear (in any crucial way) in the coefficient of the driving noise in the controlled dynamics. Similarly to the approach in Buckdahn and Ma [12], a change of variables could therefore be used to 'hide' the rough noise term in the drift coefficient, thus recovering a more classical optimal control setting. In this chapter we aim to significantly extend the theoretical results of Chapter 2, which will require us to be able to handle full control of the dynamics. We will therefore make use of the central ideas of the previous chapter in order to handle our control problem.

6.1 The Kalman–Bucy filter revisited

Let us take an underlying filtered space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0})$. We suppose that an $\mathbb{R}^m$-valued signal process S and an $\mathbb{R}^d$-valued observation process Y satisfy the following pair of linear equations

$$\mathrm{d}S_t = \big(\theta_t + \alpha_t S_t\big)\,\mathrm{d}t + \sigma_t\,\mathrm{d}B^1_t, \tag{6.1}$$

$$\mathrm{d}Y_t = c_t S_t\,\mathrm{d}t + \mathrm{d}B^2_t, \tag{6.2}$$

with the initial conditions $Y_0 = 0$ and $S_0 \sim N(\mu_0, \Sigma_0)$ for some $\mu_0 \in \mathbb{R}^m$ and $\Sigma_0 \in \mathcal{S}^m_+$, where $\mathcal{S}^m_+$ denotes the set of symmetric, positive definite $m \times m$-matrices. Here B^1 (resp. B^2) is a standard $\mathbb{R}^l$(resp. $\mathbb{R}^d$)-valued Brownian motion, and $\theta\colon [0,T] \to \mathbb{R}^m$, $\alpha\colon [0,T] \to \mathbb{R}^{m\times m}$, $\sigma\colon [0,T] \to \mathbb{R}^{m\times l}$ and $c\colon [0,T] \to \mathbb{R}^{d\times m}$ are parameters. Here

we include the case when the signal noise and observation noise are correlated; we suppose that their quadratic covariation is given by

$$\mathrm{d}\langle B^1, B^2\rangle_t = \rho_t \,\mathrm{d}t, \tag{6.3}$$

for some correlation matrix $\rho\colon [0,T] \to \mathbb{R}^{l\times d}$. In the scalar case, the correlation should naturally satisfy $\rho^2 \leq 1$. The analogous assumption here is that the matrix $I - \rho\rho^\top$ be positive semi-definite, where I denotes the $l \times l$ identity matrix.

As usual, we shall denote by $(\mathcal{Y}_t)_{t\geq 0}$ the (completed) natural filtration generated by Y. In this setting, the conditional mean $q_t = \mathbb{E}[S_t \,|\, \mathcal{Y}_t]$ satisfies the SDE

$$\mathrm{d}q_t = \big(\theta_t + \alpha_t q_t\big)\,\mathrm{d}t + (R_t c_t^\top + \sigma_t \rho_t)(\mathrm{d}Y_t - c_t q_t \,\mathrm{d}t) \tag{6.4}$$

with $q_0 = \mu_0$, and the conditional variance $R_t = \mathbb{E}[(S_t - q_t)(S_t - q_t)^\top \,|\, \mathcal{Y}_t]$ satisfies the deterministic matrix Riccati equation

$$\frac{\mathrm{d}R_t}{\mathrm{d}t} = \sigma_t\sigma_t^\top + \alpha_t R_t + R_t\alpha_t^\top - (R_t c_t^\top + \sigma_t\rho_t)(c_t R_t + \rho_t^\top \sigma_t^\top) \tag{6.5}$$

with $R_0 = \Sigma_0$.

In Chapter 2 we considered a similar setting to the one described above, but there the parameter c was assumed to be known, and the signal and observation noises were assumed to be uncorrelated (so that $\rho = 0$). Now we shall relax these assumptions, and also allow a more general penalty, analogous to the one introduced in Chapter 4, which in particular takes into account the statistical likelihood for different parameter choices. As we will see, this approach will lead to the derivation of a pathwise stochastic control problem with control in the coefficient of the (rough) noise term, and thus require the central ideas of the previous chapter in order to proceed.

In our current setting, the class of admissible measures simply corresponds to the family of possible parameters $\theta, \alpha, \sigma, c, \rho, \mu_0$ and Σ_0 of the dynamics of the signal and observation processes. For notational brevity, we shall denote[1]

$$\gamma := (\theta, \alpha, \sigma, c, \rho),$$

and write

$$\Gamma := \mathbb{R}^m \times \mathbb{R}^{m\times m} \times \mathbb{R}^{m\times l} \times \mathbb{R}^{d\times m} \times \Upsilon$$

[1]One is not obliged to consider all of these parameters as being uncertain, but we will focus on this, the most general case.

for the space in which γ takes values, where Υ denotes the space of valid correlation matrices:

$$\Upsilon := \big\{\rho \in \mathbb{R}^{l\times d} : I - \rho\rho^\top \text{ is positive definite}\big\} = \big\{\rho \in \mathbb{R}^{l\times d} : \lambda_{\max}(\rho\rho^\top) < 1\big\}, \quad (6.6)$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue.

Let S, Y be $\mathcal{F}_t$-adapted processes. For each parameter choice $(\gamma, \mu_0, \Sigma_0)$, we let $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$ be a probability measure such that the law of (S, Y) is equal to the law of $(\tilde{S}^{\gamma,\mu_0,\Sigma_0}, \tilde{Y}^{\gamma,\mu_0,\Sigma_0})$, where $\tilde{S}^{\gamma,\mu_0,\Sigma_0}$, $\tilde{Y}^{\gamma,\mu_0,\Sigma_0}$ satisfy the SDEs (6.1)–(6.3) driven by some Brownian motions B^1, B^2 on some probability space, with $\tilde{S}_0^{\gamma,\mu_0,\Sigma_0} \sim N(\mu_0, \Sigma_0)$ and $\tilde{Y}_0^{\gamma,\mu_0,\Sigma_0} = 0$. In particular, the processes S, Y as functions on $\Omega \times [0, T]$, and hence the (uncompleted) filtration $\sigma(Y_s : s \in [0, t])$, $t \in [0, T]$, are defined independently of the choice of parameters.

For a given uncertainty aversion parameter $k_1 > 0$ and exponent $k_2 \geq 1$, we define, for any real-valued bounded measurable function φ, the convex expectation with the representation

$$\mathcal{E}(\varphi(S_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\gamma,\mu_0,\Sigma_0} \left\{ \mathbb{E}^{\gamma,\mu_0,\Sigma_0}[\varphi(S_t) \,|\, \mathcal{Y}_t] - \left(\frac{1}{k_1}\beta_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t)\right)^{k_2} \right\}. \quad (6.7)$$

Here the essential supremum is taken over all possible parameters $(\mu_0, \Sigma_0) \in \mathbb{R}^m \times \mathcal{S}_+^m$ for the initial distribution of the signal, and over all choices of parameters γ governing the dynamics of S and Y.

One should take care here, as in general the measures $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$ are not equivalent on $\mathcal{F}_t$ (since $\langle S\rangle_t = \int_0^t \sigma_s\sigma_s^\top \,\mathrm{d}s$ is measurable), so the completed filtration $\mathcal{Y}_t = \sigma(Y_s : s \in [0, t]) \vee \mathcal{N}$ will in general depend on the choice of parameters.

In view of the insights of the previous chapter, we anticipate the eventual need to restrict to a sufficiently regular space of parameters γ (which we will later refer to as controls). We consequently make the following assumption.

Assumption 6.1.1. We shall take the space of possible parameters γ to be the family of all absolutely continuous functions $\gamma\colon [0, T] \to \Gamma$ with bounded derivative.

The penalty function β represents our opinion of how unreasonable different values of the parameters are. We shall discuss this term further in the next section. As usual, the uncertainty aversion parameters k_1, k_2 are included for generality, but will play no significant role in our analysis.

6.2 The penalty function

Similarly to in Chapter 4, we shall suppose that this penalty takes the form of a negative log-posterior density. That is, we take

$$\beta_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = -\log\Big(\vartheta_t(\gamma, \mu_0, \Sigma_0) L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t)\Big), \tag{6.8}$$

where ϑ and $L(\cdot \,|\, \mathcal{Y}_t)$ denote the prior and likelihood respectively.

We shall assume that the prior takes the form

$$-\log \vartheta_t(\gamma, \mu_0, \Sigma_0) = \int_0^t \mathfrak{f}(q_s, R_s, \gamma_s)\,\mathrm{d}s + g(\mu_0, \Sigma_0), \tag{6.9}$$

where the functions $\mathfrak{f}$ and g may be calibrated to represent our prior beliefs about the plausibility of different parameter choices. Here q and R are the conditional mean and variance corresponding to the parameters γ, μ_0 and Σ_0, given by the solutions of (6.4) and (6.5).

Note that the measures $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$ for different choices of γ, μ_0 and Σ_0 are all equivalent on $\mathcal{Y}_t$. A natural choice for $L_t(\cdot \,|\, \mathcal{Y}_t)$ is thus the Radon–Nikodym derivative

$$L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = \left(\frac{\mathrm{d}\mathbb{P}^{\gamma,\mu_0,\Sigma_0}}{\mathrm{d}\mathbb{P}^{\gamma^*,\mu_0^*,\Sigma_0^*}}\right)_{\mathcal{Y}_t}$$

which is precisely the likelihood ratio of the (arbitrary) parameter choice γ, μ_0, Σ_0, with respect to a (fixed) choice of reference parameters $\gamma^*, \mu_0^*, \Sigma_0^*$. We will now derive an explicit expression for this likelihood.

Recall (from e.g. Bain and Crisan [5, Chapter 2]) that for a given choice of parameters γ, μ_0, Σ_0, the innovation process V, given in this setting by

$$\mathrm{d}V_s = \mathrm{d}Y_s - c_s q_s\,\mathrm{d}s,$$

is a $\mathcal{Y}_t$-adapted Brownian motion under $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$. Writing q^* (resp. V^*) for the conditional mean (resp. innovation process) under the reference measure $\mathbb{P}^{\gamma^*,\mu_0^*,\Sigma_0^*}$, we have

$$\mathrm{d}V_s = \mathrm{d}V_s^* - (c_s q_s - c_s^* q_s^*)\,\mathrm{d}s.$$

Thus, by Girsanov's theorem[2] (see e.g. [19, Chapter 15]), we can represent the likelihood as a stochastic exponential, namely

$$L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = \exp\left(\int_0^t (c_s q_s - c_s^* q_s^*) \cdot \mathrm{d}V_s^* - \frac{1}{2}\int_0^t \left|c_s q_s - c_s^* q_s^*\right|^2 \mathrm{d}s\right).$$

[2]We recall that V and V^* have the predictable representation property under their respective measures.

Substituting $\mathrm{d}V_s^* = \mathrm{d}Y_s - c_s^* q_s^* \,\mathrm{d}s$, a short calculation yields

$$-\log L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = -\int_0^t (c_s q_s - c_s^* q_s^*) \cdot \mathrm{d}Y_s + \frac{1}{2}\int_0^t \Big(|c_s q_s|^2 - |c_s^* q_s^*|^2\Big)\,\mathrm{d}s.$$

Since the reference parameters are taken to be fixed, they simply amount to an additive constant in the above expression. That is,

$$-\log L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = -\int_0^t c_s q_s \cdot \mathrm{d}Y_s + \frac{1}{2}\int_0^t |c_s q_s|^2\,\mathrm{d}s + \text{const.} \tag{6.10}$$

As in Chapter 4, for simplicity we will henceforth omit this constant from our analysis, conceding that our penalty function is correct up to an additive constant.

As in previous chapters, it will be useful later to interpret the stochastic integral in (6.10) in the sense of Stratonovich, rather than that of Itô. We therefore make the transformation

$$-\int_0^t c_s q_s \cdot \mathrm{d}Y_s = -\int_0^t c_s q_s \circ \mathrm{d}Y_s + \frac{1}{2}\big\langle cq, Y\big\rangle_t.$$

Recalling (6.4), and using the fact that c is absolutely continuous and in particular of bounded variation, after some calculation we deduce that the quadratic covariation term is given by

$$\big\langle cq, Y\big\rangle_t = \int_0^t \mathrm{tr}\big(c_s(R_s c_s^\top + \sigma_s \rho_s)\big)\,\mathrm{d}s,$$

where $\mathrm{tr}(\,\cdot\,)$ denotes the trace. Note that $c_s R_s c_s^\top$ is positive semi-definite and therefore has nonnegative trace. Substituting back into (6.10), we obtain

$$-\log L_t(\gamma, \mu_0, \Sigma_0 \,|\, \mathcal{Y}_t) = -\int_0^t c_s q_s \circ \,\mathrm{d}Y_s + \frac{1}{2}\int_0^t \Big(|c_s q_s|^2 + \mathrm{tr}\big(c_s(R_s c_s^\top + \sigma_s \rho_s)\big)\Big)\mathrm{d}s. \tag{6.11}$$

For notational consistency with the previous chapters, we introduce the functions f and ψ, given by

$$f(q, R, \gamma) := \mathfrak{f}(q, R, \gamma) + \frac{1}{2}\Big(|cq|^2 + \mathrm{tr}\big(c(Rc^\top + \sigma\rho)\big)\Big) \qquad \text{and} \qquad \psi(q, \gamma) := -cq,$$

where we recall $\gamma = (\theta, \alpha, \sigma, c, \rho)$. Combining (6.8), (6.9) and (6.11), and substituting into (6.7), we then obtain the following representation:

$$\begin{aligned}\mathcal{E}(\varphi(S_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\gamma, \mu_0, \Sigma_0} \Bigg\{ & \mathbb{E}^{\gamma, \mu_0, \Sigma_0}[\varphi(S_t) \,|\, \mathcal{Y}_t] \\ & - \bigg(\frac{1}{k_1}\bigg(\int_0^t f(q_s, R_s, \gamma_s)\,\mathrm{d}s + \int_0^t \psi(q_s, \gamma_s) \circ \mathrm{d}Y_s + g(\mu_0, \Sigma_0)\bigg)\bigg)^{k_2}\Bigg\}.\end{aligned} \tag{6.12}$$

6.3 Fixing an observation path

Since the parameters $\theta, \alpha, \sigma, c$ and ρ are assumed to be absolutely continuous, R is then the C^1 solution of (6.5), and we can then let q be the solution of the filtering equation (6.4), where, by integration by parts (see e.g. Theorem 1.2.3 in [60]), the stochastic integral is interpreted as a pathwise Riemann–Stieltjes integral against Y.

Let us denote by $\Phi(\cdot; \mu, \Sigma)$ the distribution function of a $N(\mu, \Sigma)$ distribution. Then,

$$\int_{\mathbb{R}^m} \varphi(x)\, \Phi(\mathrm{d}x; q_t, R_t)$$

is a version of the conditional expectation $\mathbb{E}^{\gamma,\mu_0,\Sigma_0}[\varphi(S_t) \,|\, \mathcal{Y}_t]$ for every choice of parameters and for every functional φ.

Recall that the representation in (6.12) for the nonlinear expectation involves the stochastic integral of $\psi(q, \gamma)$ against Y. Unlike the stochastic integral in (6.4), since the paths of q and Y both have Brownian-type regularity, in general this integral does not exist in the pathwise Riemann–Stieltjes sense. As in the previous chapters, we instead aim to interpret it as a rough integral.

We therefore fix a reference measure $\mathbb{P}^{\gamma^*,\mu_0^*,\Sigma_0^*}$, and let

$$\mathbb{Y}_{s,t} := \int_s^t Y_{s,r} \otimes \circ\, \mathrm{d}Y_r,$$

defined under the measure $\mathbb{P}^{\gamma^*,\mu_0^*,\Sigma_0^*}$, so that $\mathbf{Y} = (Y, \mathbb{Y}) \in \mathscr{C}_g^{0,p}$ is a geometric rough path. Note that, since the integrand is $\sigma(Y_s : s \in [0,t])$-adapted, $\mathbb{Y}$ coincides almost surely with the same integral defined under any other choice of measure $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$ (even though the corresponding completed filtrations $\mathcal{Y}_t$ may not agree).

It remains to establish $\psi(q, \gamma)$ as being controlled (in the sense of Gubinelli) by Y. The Gubinelli derivative of q with respect to Y can be inferred by simply inspecting (6.4). Indeed, we have that

$$q_{s,t} = \int_s^t (R_r c_r^\top + \sigma_r \rho_r)\, \mathrm{d}Y_r + \mathrm{O}\big(|t-s|\big) = (R_s c_s^\top + \sigma_s \rho_s) Y_{s,t} + \mathrm{O}\big(|t-s|\big).$$

Since c is of bounded variation, it is trivially controlled by Y with derivative zero, and we conclude (from e.g. Corollary 7.4 in [32]) that $\psi(q, \gamma) = -cq$ is indeed controlled by Y with Gubinelli derivative $\psi(q, \gamma)' = -c(Rc^\top + \sigma\rho)$.

Thus, for each parameter choice $(\gamma, \mu_0, \Sigma_0)$, the rough integral

$$\int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\mathbf{Y}_s$$

exists, and is a version of the stochastic integral

$$\int_0^t \psi(q_s, \gamma_s) \circ \mathrm{d}Y_s$$

defined under the measure $\mathbb{P}^{\gamma,\mu_0,\Sigma_0}$.

Choosing this version, we then have

$$\mathcal{E}(\varphi(S_t) \,|\, \mathcal{Y}_t) = \operatorname*{ess\,sup}_{\gamma,q_0,R_0} \Bigg\{ \int_{\mathbb{R}^m} \varphi(x)\, \Phi(\mathrm{d}x; q_t, R_t) - \left(\frac{1}{k_1} \left(\int_0^t f(q_s, R_s, \gamma_s)\, \mathrm{d}s + \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\mathbf{Y}_s + g(q_0, R_0) \right) \right)^{k_2} \Bigg\}.$$

As in previous chapters, we will now restrict our attention to a particular realisation of the enhanced observation path $\mathbf{Y}$, which here we denote by $\boldsymbol{\zeta}$. That is, we define

$$\boldsymbol{\zeta} = (\zeta, \zeta^{(2)}) := (Y(\omega), \mathbb{Y}(\omega)) = \mathbf{Y}(\omega) \in \mathscr{C}_g^{0,p}$$

for some fixed $\omega \in \Omega$. By the previous paragraph, we can then consider the filter dynamics (6.4)–(6.5) with Y replaced by ζ, namely

$$\mathrm{d}q_s = \big(\theta_s + \alpha_s q_s\big)\, \mathrm{d}s + (R_s c_s^\top + \sigma_s \rho_s)(\mathrm{d}\zeta_s - c_s q_s\, \mathrm{d}s), \tag{6.13}$$

$$\frac{\mathrm{d}R_s}{\mathrm{d}s} = \sigma_s \sigma_s^\top + \alpha_s R_s + R_s \alpha_s^\top - (R_s c_s^\top + \sigma_s \rho_s)(c_s R_s + \rho_s^\top \sigma_s^\top). \tag{6.14}$$

It is then clear that the 'pathwise version' of the nonlinear expectation should be given by

$$\mathcal{E}(\varphi(S_t) \,|\, \boldsymbol{\zeta}) := \sup_{\gamma,q_0,R_0} \Bigg\{ \int_{\mathbb{R}^m} \varphi(x)\, \Phi(\mathrm{d}x; q_t, R_t) - \left(\frac{1}{k_1} \left(\int_0^t f(q_s, R_s, \gamma_s)\, \mathrm{d}s + \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s + g(q_0, R_0) \right) \right)^{k_2} \Bigg\},$$

where here q, R satisfy the equations (6.13)–(6.14) driven by ζ.

6.4 Reformulation as an optimal control problem

Writing $\gamma = (\theta, \alpha, \sigma, c, \rho)$ as usual, consider the functional $\kappa_t \colon \mathbb{R}^m \times \mathcal{S}_+^m \to \mathbb{R}$ defined by

$$\kappa_t(\mu, \Sigma) = \inf \left\{ \int_0^t f(q_s, R_s, \gamma_s)\, \mathrm{d}s + \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s + g(q_0, R_0) \;\middle|\; \begin{array}{c} \gamma, q_0, R_0 \text{ such that} \\ (q_t, R_t) = (\mu, \Sigma) \end{array} \right\}, \tag{6.15}$$

where q and R satisfy (6.13)–(6.14) with the *terminal* condition $(q_t, R_t) = (\mu, \Sigma)$. The function κ_t is related to the nonlinear expectation (6.12) by the following lemma.

Lemma 6.4.1. *For any rough path $\boldsymbol{\zeta} = (\zeta, \zeta^{(2)})$, and any bounded measurable function φ, we have the equality*

$$\mathcal{E}(\varphi(S_t) \,|\, \boldsymbol{\zeta}) = \sup_{(\mu,\Sigma)\in\mathbb{R}^m\times\mathcal{S}_+^m} \left\{ \int_{\mathbb{R}^m} \varphi(x)\, \mathrm{d}\Phi(x;\mu,\Sigma) - \left(\frac{1}{k_1}\kappa_t(\mu,\Sigma)\right)^{k_2} \right\}. \tag{6.16}$$

Proof. We have, for any $(\mu, \Sigma) \in \mathbb{R}^m \times \mathcal{S}_+^m$,

$$\begin{aligned}
&\sup \left\{ \int_{\mathbb{R}^m} \varphi(x)\, \Phi(\mathrm{d}x; q_t, R_t) - \left(\frac{1}{k_1}\left(\int_0^t f(q_s, R_s, \gamma_s)\, \mathrm{d}s \right.\right.\right. \\
&\qquad\qquad \left.\left.\left. + \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s + g(q_0, R_0) \right)\right)^{k_2} \;\middle|\; \begin{array}{c} \gamma, q_0, R_0 \;\text{ such that} \\ (q_t, R_t) = (\mu, \Sigma) \end{array} \right\}, \\
&= \int_{\mathbb{R}^m} \varphi(x)\, \mathrm{d}\Phi(x;\mu,\Sigma) - \left(\frac{1}{k_1}\kappa_t(\mu,\Sigma)\right)^{k_2},
\end{aligned}$$

and we conclude by taking the supremum over $(\mu, \Sigma) \in \mathbb{R}^m \times \mathcal{S}_+^m$. □

The expression for κ in (6.15) looks very much like that of the value function of an optimal control problem with state trajectories governed by (6.13)–(6.14). To make this exact we should write κ as an infimum over the 'control' γ alone. This is easy, but one should note that, for certain choices of control γ and terminal condition $(\mu, \Sigma) \in \mathbb{R}^m \times \mathcal{S}_+^m$, there will not actually exist a corresponding initial value $(q_0, R_0) \in \mathbb{R}^m \times \mathcal{S}_+^m$. This can happen for one of two reasons. First, due to the term $\sigma_s \sigma_s^\top$, the solution to (6.14) may no longer be positive semi-definite, so that R_0 does not correspond to a covariance matrix. Second, the solution to (6.14) may 'blow up' in finite time, due to the quadratic term (in R) in the final term on the right-hand side of (6.14). An example of such behaviour was exhibited in Example 2.3.2.

Heuristically, the Kalman–Bucy filter is well behaved when run forwards in time from an initial condition, but here we instead fix a terminal condition and run the filtering equations backwards in time, which introduces the abnormalities described above. To prevent this unphysical behaviour we simply prescribe the value $g(\mu_0, \Sigma_0) = \infty$ for any initial value $(\mu_0, \Sigma_0) \notin \mathbb{R}^m \times \mathcal{S}_+^m$ and, although we don't actually obtain a physical initial value for solutions which 'blow up' in a finite time, we assign an infinite 'initial' cost to all such trajectories.

We can now write

$$\kappa_t(\mu,\Sigma) = \inf_\gamma \left\{ \int_0^t f(q_s^{t,\mu,\Sigma,\gamma}, R_s^{t,\Sigma,\gamma}, \gamma_s)\, \mathrm{d}s + \int_0^t \psi(q_s^{t,\mu,\Sigma,\gamma}, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s + g\big(q_0^{t,\mu,\Sigma,\gamma}, R_0^{t,\Sigma,\gamma}\big) \right\} \tag{6.17}$$

where $q^{t,\mu,\Sigma,\gamma}, R^{t,\Sigma,\gamma}$ satisfy (6.13)–(6.14) with the terminal condition

$$\left(q_t^{t,\mu,\Sigma,\gamma}, R_t^{t,\Sigma,\gamma}\right) = (\mu, \Sigma),$$

noting that trajectories with the undesired behaviour described above will never be considered when taking the infimum in (6.17).

We have derived an optimal control problem, with the controlled dynamics (6.13)–(6.14), and the value function defined in (6.17). Moreover, the appearance of the 'Brownian-like' path ζ in (6.13), and indeed the rough path $\boldsymbol{\zeta}$ in (6.17), puts us back into the setting of pathwise stochastic control. In the case where the parameter c is known, the signal and observation noises are uncorrelated (so that $\rho = 0$), and if we omit the likelihood term in the penalty of our nonlinear expectation, then we are not directly controlling the coefficient of the rough term ζ. This was the case in the setting of Chapter 2, where a change of variables was then used to completely isolate the observation path from the controlled terms.

In the current setting however we cannot escape the need to control the coefficient of ζ. As described in Chapter 5, if the variation of the controls γ is not sufficiently penalised then the control problem degenerates. The physical interpretation here is the following: even if we suppose that the parameters $\theta, \alpha, \sigma, c$ and ρ are able to fluctuate at the same rate as the observation path ζ, it is not reasonable to suppose that we should be able to calibrate these parameters over time scales that are so small that our observations are dominated by measurement noise.

Accordingly, we employ the strategy introduced in the previous chapter of introducing a regularising cost, and rewriting the problem in terms of an abstract control process u. We consider the dynamics (6.13)–(6.14) along with

$$\mathrm{d}\gamma_s^{t,a,u} = h(\gamma_s^{t,a,u}, u_s)\,\mathrm{d}s,$$

for some function $h\colon \Gamma \times U \to U$, where u belongs to the class $\mathcal{U}$ of bounded measurable functions $u\colon [0,T] \to U := \mathbb{R}^m \times \mathbb{R}^{m\times m} \times \mathbb{R}^{m\times l} \times \mathbb{R}^{d\times m} \times \mathbb{R}^{l\times d}$. The terminal condition is now given by

$$\left(q_t^{t,\mu,\Sigma,a,u}, R_t^{t,\Sigma,a,u}, \gamma_t^{t,a,u}\right) = (\mu, \Sigma, a) \in \mathbb{R}^m \times \mathcal{S}_+^m \times \Gamma.$$

Allowing f to depend on u, and g to depend on γ_0 (which makes no difference to the proof of Lemma 6.4.1), and writing $\tilde{\kappa}$ for the regularised version of κ, we can write

$$\tilde{\kappa}_t(\mu, \Sigma) = \inf_{a\in\Gamma} v(t, \mu, \Sigma, a), \tag{6.18}$$

where

$$v(t,\mu,\Sigma,a) := \inf_{u\in\mathcal{U}} \Big\{ \int_0^t f(q_s^{t,\mu,\Sigma,a,u}, R_s^{t,\Sigma,a,u}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s \\ + \int_0^t \psi(q_s^{t,\mu,\Sigma,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}\boldsymbol{\zeta}_s + g\big(q_0^{t,\mu,\Sigma,a,u}, R_0^{t,\Sigma,a,u}, \gamma_0^{t,a,u}\big)\Big\} \tag{6.19}$$

is the value function of our new control problem. As before, to avoid unphysical trajectories, we assign an infinite cost to any controls u such that $(\mu_0,\Sigma_0)\notin\mathbb{R}^m\times\mathcal{S}_+^m$ or such that the solution to (6.14) 'blows up' in a finite time. Moreover, we assign an infinite cost to those controls which lead to ρ leaving the space of valid correlation matrices Υ (as defined in (6.6)).

Our uncertainty is thus represented by the function v. Once the value of this function has been determined, one can use (6.18) and then (6.16) to evaluate arbitrary functions of the signal process S under the nonlinear expectation $\mathcal{E}(\cdot\,|\mathcal{Y}_t)$.

6.5 A nonlinear backward control problem

It remains to characterise the value function v defined in (6.19) as the unique solution of a rough HJB equation. For convenience, we rewrite the controlled dynamics in full as

$$\begin{aligned}
\mathrm{d}q_s^{t,\mu,\Sigma,a,u} &= b_\mu(q_s^{t,\mu,\Sigma,a,u}, R_s^{t,\Sigma,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}s + \lambda(R_s^{t,\Sigma,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}\zeta_s, & q_t^{t,\mu,\Sigma,a,u} &= \mu,\\
\mathrm{d}R_s^{t,\Sigma,a,u} &= b_\Sigma(R_s^{t,\Sigma,a,u}, \gamma_s^{t,a,u})\,\mathrm{d}s, & R_t^{t,\Sigma,a,u} &= \Sigma,\\
\mathrm{d}\gamma_s^{t,a,u} &= h(\gamma_s^{t,a,u}, u_s)\,\mathrm{d}s, & \gamma_t^{t,a,u} &= a,
\end{aligned}$$

where $\gamma=(\alpha,\sigma,c,\rho)$, and we define

$$\begin{aligned}
b_\mu(q,R,\gamma) &= \theta + \alpha q - (Rc^\top + \sigma\rho)cq,\\
b_\Sigma(R,\gamma) &= \sigma\sigma^\top + \alpha R + R\alpha^\top - (Rc^\top+\sigma\rho)(cR+\rho^\top\sigma^\top),\\
\lambda(R,\gamma) &= Rc^\top + \sigma\rho.
\end{aligned}$$

As usual this is a 'backward' control problem in the sense that, in contrast to the classical setting of optimal control, here we prescribe a terminal condition for the state trajectories, and consider a cost associated with their initial value. More significantly, we note that $b_\mu, b_\Sigma \notin \mathrm{Lip}_b$ and $\lambda,\psi\notin C_b^3$, so we cannot immediately apply the results of the previous chapter. Nevertheless, as we will see, the desired results can be recovered with some modifications.

Notation 6.5.1. In the following we write $|\cdot|$ for the usual Euclidean norm, and $\|A\|$ for the Frobenius norm of a given matrix A, i.e. $\|A\|^2 = \mathrm{tr}(A^\top A)$. Given an element $\gamma = (\theta, \alpha, \sigma, c, \rho)$ of $U = \mathbb{R}^m \times \mathbb{R}^{m\times m} \times \mathbb{R}^{m\times l} \times \mathbb{R}^{d\times m} \times \mathbb{R}^{l\times d}$, we write $\|\gamma\| = \max\big\{|\theta|, \|\alpha\|, \|\sigma\|, \|c\|, \|\rho\|\big\}$. We point out however that, since the space Υ of correlation matrices is uniformly bounded[3], the dependence on $\|\rho\|$ is not particularly crucial.

If $A \in \mathcal{S}_+^m$, so that in particular A is symmetric and positive definite, we write $\lambda_{\min}(A)$ (resp. $\lambda_{\max}(A)$) for the smallest (resp. largest) eigenvalue of A.

Where there is no risk of ambiguity, we will omit the superscripts from the state variables q, R and γ. Finally, we will use the symbol $\lesssim$ to denote inequality up to a multiplicative constant which may depend on any of the dimensions d, l, m, the functions f, g, h, the measure of regularity p, the terminal time T, and the bound L, where as usual $L > 0$ is chosen such that $\|\!|\!|\boldsymbol{\zeta}|\!|\!|_p \leq L$.

Assumption 6.5.2. We assume that

- $f = f(q, R, \gamma, u)$ and $g = g(q, R, \gamma)$ are continuous, bounded below, and locally Lipschitz in (q, R, γ), uniformly in u,
- $h = h(\gamma, u)$ is continuous, surjective in u, i.e. $\{h(\gamma, u) : u \in U\} = U$ for every $\gamma \in \Gamma$, Lipschitz in γ, uniformly in u, and bounded in γ, locally uniformly in u, and moreover, for some $\delta \geq 1$, satisfies
$$\sup_{\gamma\in\Gamma} \frac{\big\|h(\gamma, u)\big\|}{\|u\|^\delta} \longrightarrow 0 \qquad \text{as} \quad \|u\| \longrightarrow \infty, \tag{6.20}$$
- the running cost f satisfies the asymptotic growth condition:
$$\frac{f(q, R, \gamma, u)}{\big(1 + |q| + \|R\|^2 + \|\gamma\|^2\big)\|u\|^\delta + \big(1 + |q|^2 + \|R\|^2\big)\big(1 + \|\gamma\|^4\big)} \longrightarrow \infty \tag{6.21}$$
as $|q| + \|R\| + \|\gamma\| + \|u\| \to \infty$, with the same δ as above,
- and the initial cost g satisfies:
$$\frac{g(q, R, \gamma)}{|q|^2 + \big(1 + \|R\|\big)\big(1 + \|\gamma\|^2\big)} \longrightarrow \infty \qquad \text{as} \qquad |q| + \|R\| + \|\gamma\| \longrightarrow \infty, \tag{6.22}$$
$$\inf_{(q,\gamma)\in\mathbb{R}^m\times\Gamma} g(q, R, \gamma) \longrightarrow \infty \qquad \text{as} \qquad \lambda_{\min}(R) \longrightarrow 0. \tag{6.23}$$

[3]Indeed, one can show that $\|\rho\| \leq \sqrt{l}$ for every $\rho \in \Upsilon$.

If the correlation ρ is known, then one can simply take, for example, $h(\gamma, u) = u$. If ρ is uncertain then, for mostly technical reasons, one must take a little extra care to ensure that correlations close to the boundary of Υ (where $\lambda_{\max}(\rho\rho^\top) = 1$) are sufficiently penalised. In this case we assume in addition that

- $$\inf_{q,R,\alpha,\sigma,c} g(q, R, \gamma) \longrightarrow \infty \qquad \text{as} \quad \lambda_{\max}(\rho\rho^\top) \longrightarrow 1, \tag{6.24}$$
$$\big\|h(\gamma, u)\big\| \leq (1 - \lambda_{\max}(\rho\rho^\top))\|u\| \qquad \text{for all} \quad (\gamma, u) \in \Gamma \times U. \tag{6.25}$$

Remark 6.5.3. The surjectivity of h in u is assumed to ensure that, no matter the choice of terminal condition (t, μ, Σ, a), there always exists a choice of control u such that the state trajectories remain inside their respective domains, so that in particular $R_0 \in \mathcal{S}_+^m$ and $\rho_0 \in \Upsilon$. This guarantees that the value function v is finite-valued.

The result of Lemma 5.4.3 can be recovered in the current setting as follows.

Lemma 6.5.4. *Under Assumption 6.5.2, for any terminal condition (t, μ, Σ, a) and control u, we have that*

$$\left| \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s \right| \leq C + \frac{1}{2}\bigg(\int_0^t f(q_s, R_s, \gamma_s, u_s)\, \mathrm{d}s + g(q_0, R_0, \gamma_0) \bigg),$$

where the constant C depends on d, l, m, f, g, h, p, T and L.

Proof. From the estimate (3.6) in Proposition 3.3.1, we have that

$$\left| \int_0^t \psi(q_s, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s \right| \lesssim \big|\psi(q_0, \gamma_0)\big| + \big\|\psi(q, \gamma)_0'\big\| + \big\|R^{\psi(q,\gamma)}\big\|_{\frac{p}{2};[0,t]} + \big\|\psi(q, \gamma)'\big\|_{p;[0,t]}. \tag{6.26}$$

We aim to bound each of the terms on the right-hand side. Recalling that $\psi(q, \gamma) = -cq$ and $\psi(q, \gamma)' = -c(Rc^\top + \sigma\rho)$, we have that

$$\begin{aligned} \big|\psi(q_0, \gamma_0)\big| + \big\|\psi(q, \gamma)_0'\big\| &\lesssim \|c_0\||q_0| + \|c_0\|\big\|R_0 c_0^\top + \sigma_0\rho_0\big\| \\ &\lesssim |q_0|^2 + \big(1 + \|R_0\|\big)\|\gamma_0\|^2. \end{aligned} \tag{6.27}$$

Writing $(\dot\theta, \dot\alpha, \dot\sigma, \dot c, \dot\rho) = \dot\gamma = h(\gamma, u)$, we have

$$\begin{aligned} \big\|\psi(q, \gamma)'\big\|_{p;[0,t]} &\leq \big\|\psi(q, \gamma)'\big\|_{1;[0,t]} = \big\|c(Rc^\top + \sigma\rho)\big\|_{1;[0,t]} \\ &= \int_0^t \Big| \dot c_s\big(R_s c_s^\top + \sigma_s\rho_s\big) + c_s\big(b_\Sigma(R_s, \gamma_s)c_s^\top + R_s\dot c_s^\top + \dot\sigma_s\rho_s + \sigma_s\dot\rho_s\big) \Big|\, \mathrm{d}s \\ &\lesssim \int_0^t \big(1 + \|R_s\|^2 + \|\gamma_s\|^2\big)\big\|h(\gamma_s, u_s)\big\| + \big(1 + \|R_s\|^2\big)\big(1 + \|\gamma_s\|^4\big)\, \mathrm{d}s. \end{aligned} \tag{6.28}$$

By the Young–Lóeve inequality (see e.g. Theorem 6.8 in [34]), we have

$$\left| \int_h^r (R_s c_s^\top + \sigma_s \rho_s) \, \mathrm{d}\zeta_s - (R_h c_h^\top + \sigma_h \rho_h) \zeta_{h,r} \right| \leq \frac{1}{1 - 2^{-\frac{1}{p}}} \left\| Rc^\top + \sigma\rho \right\|_{1;[h,r]} \|\zeta\|_{p;[h,r]} \quad (6.29)$$

for any interval $[h, r] \subset [0, t]$. We calculate

$$\begin{aligned} -R_{h,r}^{\psi(q,\gamma)} &= -\psi(q_r, \gamma_r) + \psi(q_h, \gamma_h) + \psi(q, \gamma)'_h \zeta_{h,r} \\ &= c_h q_{h,r} + c_{h,r} q_r + \psi(q, \gamma)'_h \zeta_{h,r} \\ &= c_h \left(\int_h^r \mathrm{d}q_s - (R_h c_h^\top + \sigma_h \rho_h) \zeta_{h,r} \right) + \int_h^r \dot{c}_s q_r \, \mathrm{d}s. \end{aligned}$$

Recalling (6.13) and using (6.29), we have

$$\begin{aligned} \left| R_{h,r}^{\psi(q,\gamma)} \right| &\lesssim \|c_h\| \left\| Rc^\top + \sigma\rho \right\|_{1;[h,r]} + \left| \int_h^r c_h \big(\alpha_s q_s - (R_s c_s^\top + \sigma_s \rho_s) c_s q_s \big) + \dot{c}_s q_r \, \mathrm{d}s \right| \\ &\lesssim \int_h^r \Big(\|c_h\| \left| b_\Sigma(R_s, \gamma_s) c_s^\top + R_s \dot{c}_s^\top + \dot{\sigma}_s \rho_s + \sigma_s \dot{\rho}_s \right| \\ &\qquad + \|c_h\| \left| \alpha_s q_s - (R_s c_s^\top + \sigma_s \rho_s) c_s q_s \right| + \|\dot{c}_s\| |q_r| \Big) \, \mathrm{d}s. \end{aligned}$$

We then obtain

$$\begin{aligned} &\left\| R^{\psi(q,\gamma)} \right\|_{\frac{p}{2};[0,t]} \leq \left\| R^{\psi(q,\gamma)} \right\|_{1;[0,t]} = \lim_{|\mathcal{P}| \to 0} \sum_{[h,r] \in \mathcal{P}} \left| R_{h,r}^{\psi(q,\gamma)} \right| \\ &\lesssim \int_0^t \Big(\|c_s\| \left| b_\Sigma(R_s, \gamma_s) c_s^\top + R_s \dot{c}_s^\top + \dot{\sigma}_s \rho_s + \sigma_s \dot{\rho}_s \right| \\ &\qquad + \|c_s\| \left| \alpha_s q_s - (R_s c_s^\top + \sigma_s \rho_s) c_s q_s \right| + \|\dot{c}_s\| |q_s| \Big) \, \mathrm{d}s \\ &\lesssim \int_0^t \big(1 + |q_s| + \|R_s\|^2 + \|\gamma_s\|^2 \big) \left\| h(\gamma_s, u_s) \right\| + \big(1 + |q_s|^2 + \|R_s\|^2 \big) \big(1 + \|\gamma_s\|^4 \big) \, \mathrm{d}s, \end{aligned} \quad (6.30)$$

where the limit in the above is taken over any sequence of partitions of the interval $[0, t]$ with mesh size tending to zero. Substituting (6.27), (6.28) and (6.30) into (6.26), and using the growth conditions (6.20)–(6.22) in Assumption 6.5.2, we deduce the result. □

Corollary 6.5.5. *Let K be a compact subset of $\mathbb{R}^m \times \mathcal{S}_+^m \times \Gamma$. There exists an $M > 0$ such that, when taking the infimum over $u \in \mathcal{U}$ in (6.19) for $(t, \mu, \Sigma, a) \in [0, T] \times K$, one may restrict to controls u such that the norms*

$$\|q\|_\infty, \ \|R\|_\infty, \ \|\gamma\|_\infty, \ \|R\|_{1;[0,t]}, \ \|\gamma\|_{1;[0,t]}$$

are all bounded by M, without changing the value of $v(t, \mu, \Sigma, a)$.

Proof. One can obtain a bound for $\|\gamma\|_{1;[0,t]}$ by a similar argument to that in the proof of Corollary 5.4.4. Since $\gamma_t = a$ lives in a compact set, we immediately also have a bound for $\|\gamma\|_\infty$.

We infer from (6.22) that both the terminal value $R_t = \Sigma$ and initial value R_0 of R must lie in some bounded set, and by inspecting the ODE (6.14) satisfied by R, we deduce that the entire path of R must also live in a bounded set, giving a bound for $\|R\|_\infty$.

Given the bounds for $\|\gamma\|_\infty$ and $\|R\|_\infty$, a bound for $\|R\|_{1;[0,t]}$ follows easily from (6.14). Finally, inspecting the equation (6.13) satisfied by q, in view of (6.29), we deduce a bound for $\|q\|_\infty$. □

As in the previous chapters, let us approximate the rough path $\boldsymbol{\zeta}$ by a smooth path η. We then obtain the approximate value function

$$v^\eta(t,\mu,\Sigma,a) := \inf_{u\in\mathcal{U}} \Big\{ \int_0^t f(q_s^{t,\mu,\Sigma,a,u,\eta}, R_s^{t,\Sigma,a,u}, \gamma_s^{t,a,u}, u_s)\,\mathrm{d}s \\ + \int_0^t \psi(q_s^{t,\mu,\Sigma,a,u,\eta}, \gamma_s^{t,a,u})\,\mathrm{d}\eta_s + g\big(q_0^{t,\mu,\Sigma,a,u,\eta}, R_0^{t,\Sigma,a,u}, \gamma_0^{t,a,u}\big)\Big\}. \tag{6.31}$$

Notation 6.5.6. In the following, we will write ∇_μ for the usual gradient with respect to μ, and write ∇_Σ and ∇_a for the generalised gradients with respect to each of the components of Σ and a respectively. We will also write $A:B$ for the inner product of two elements A, B from the same vector space. In particular, when A, B are matrices, $A:B = \mathrm{tr}(A^\top B)$ denotes the Frobenius inner product of A and B.

We shall denote by $C^\uparrow([0,T]\times\mathbb{R}^m\times\mathcal{S}_+^m\times\Gamma;\mathbb{R})$ the class of functions $\tilde{v}\colon [0,T]\times\mathbb{R}^m\times\mathcal{S}_+^m\times\Gamma\to\mathbb{R}$ which explode asymptotically; that is, those functions $\tilde{v}$ such that

$$\tilde{v}(t,\mu,\Sigma,a) \longrightarrow \infty$$

as $|\mu| + \|\Sigma\| + \|a\| \to \infty$, and as $\lambda_{\min}(\Sigma)\to 0$, and, in the case when ρ is uncertain, as $\lambda_{\max}(\rho\rho^\top)\to 1$.

Proposition 6.5.7. *Under Assumption 6.5.2, the approximate value function v^η, as defined in* (6.31)*, is locally Lipschitz continuous with respect to (t,μ,Σ,a), and is the unique viscosity solution of the HJB equation*

$$\frac{\partial v^\eta}{\partial t} + b_\mu\cdot\nabla_\mu v^\eta + b_\Sigma : \nabla_\Sigma v^\eta + \sup_{u\in U}\big\{h : \nabla_a v^\eta - f\big\} + (\lambda\cdot\nabla_\mu v^\eta - \psi)\dot{\eta}_t = 0 \tag{6.32}$$

in the class $C^\uparrow([0,T]\times\mathbb{R}^m\times\mathcal{S}_+^m\times\Gamma;\mathbb{R})$ with the initial condition

$$v^\eta(0,\mu,\Sigma,a) = g(\mu,\Sigma,a).$$

Proof. As the path η is smooth, the associated PDE (6.32) is classical, except for the nonlinearities inherited from the filtering equations. As the proof of this result is lengthy, and has already been presented in detail in similar contexts in Chapters 2 and 4, we will only give a sketch of the proof in this setting.

That v^η is a viscosity solution of (6.32) is a standard application of the dynamic programming principle; we refer to the proof of Proposition 2.4.11 for precise details.

Heuristically, as a result of (6.21) and (6.25), for terminal conditions (μ, Σ, a) which take extreme or close to degenerate values, i.e. when either $|\mu| + \|\Sigma\| + \|a\| \gg 1$ or $\lambda_{\min}(\Sigma) \approx 0$ or $\lambda_{\max}(\rho\rho^\top) \approx 1$, it takes very expensive controls to allow the state trajectories (q, R, γ) to escape these parts of their domain. It then follows from the growth conditions (6.22)–(6.24) that the value function itself must explode as one approaches these extreme and degenerate values; that is, $v^\eta \in C^\uparrow([0,T] \times \mathbb{R}^m \times \mathcal{S}_+^m \times \Gamma; \mathbb{R})$.

One can prove that v^η is locally Lipschitz in all of its arguments by adapting the proof of Theorem 2.5.3.

The controlled dynamics do not satisfy the standard Lipschitz condition which would be required to be able to apply a standard uniqueness result for Hamilton–Jacobi equations on unbounded domains, as in e.g. Yong and Zhou [66, Chapter 4]. Nevertheless, uniqueness for an equation of the same form as (6.32) was established in Section 2.6, and an analogous argument may be used here. The main insight of this result is that the extra condition one should impose to obtain uniqueness is that solutions belong to the space $C^\uparrow([0,T] \times \mathbb{R}^m \times \mathcal{S}_+^m \times \Gamma; \mathbb{R})$; that is, one should restrict to solutions which explode as they approach the boundary. □

The main result of this chapter is given by the following theorem.

Theorem 6.5.8. *Under Assumption 6.5.2, the value function v, as defined in* (6.19)*, solves the rough HJB equation*

$$\mathrm{d}v + \big(b_\mu \cdot \nabla_\mu v + b_\Sigma : \nabla_\Sigma v\big)\,\mathrm{d}t + \sup_{u \in U} \big\{h : \nabla_a v - f\big\}\mathrm{d}t + \lambda \cdot \nabla_\mu v \,\mathrm{d}\zeta - \psi \,\mathrm{d}\boldsymbol{\zeta} = 0$$

with

$$v(0, \mu, \Sigma, a) = g(\mu, \Sigma, a)$$

in the sense of Definition 5.6.1.

Moreover, writing $v = v^\zeta$, the map from $\mathscr{C}_g^{0,p}([0,T]; \mathbb{R}^d) \to \mathbb{R}$ given by $\boldsymbol{\zeta} \mapsto v^\zeta(t, \mu, \Sigma, a)$ is locally uniformly continuous with respect to the p-variation rough path distance, locally uniformly in (t, μ, Σ, a).

Proof. Let K be a compact subset of $\mathbb{R}^m \times \mathcal{S}_+^m \times \Gamma$ and let $\boldsymbol{\eta} \in \mathscr{C}^p$ be another rough path such that $\|\boldsymbol{\eta}; \boldsymbol{\zeta}\|_p \leq 1$. By possibly replacing L by $L+1$, we may assume that $|||\boldsymbol{\eta}|||_p \leq L$. Let us write q^η (resp. q^ζ) for the solution of (6.13) driven by η (resp. ζ), and write v^η (resp. v^ζ) for the corresponding value function, as defined in (6.19).

By Corollary 6.5.5, there exists a constant $M > 0$ such that, for terminal conditions $(t, \mu, \Sigma, a) \in [0,T] \times K$, we may restrict to controls $u \in \mathcal{U}^M \subseteq \mathcal{U}$ such that

$$\|q^\eta\|_\infty,\ \|q^\zeta\|_\infty,\ \|R\|_\infty,\ \|\gamma\|_\infty,\ \|R\|_{1;[0,t]},\ \|\gamma\|_{1;[0,t]}$$

are all bounded by M. In the following we will allow the multiplicative constant indicated by the symbol $\lesssim$ to also depend on M.

By the Young–Lóeve inequality (see e.g. Theorem 6.8 in [34]), we have

$$\begin{aligned}\left| \int_s^t (R_r c_r^\top + \sigma_r \rho_r)\, \mathrm{d}(\eta - \zeta)_r \right| &\lesssim \left\| R_s c_s^\top + \sigma_s \rho_s \right\| |(\eta_t - \zeta_t) - (\eta_s - \zeta_s)| \\ &\quad + \left\| R c^\top + \sigma \rho \right\|_{1;[s,t]} \|\eta - \zeta\|_{p;[s,t]} \\ &\lesssim \|\eta - \zeta\|_{p;[s,t]},\end{aligned}$$

from which we deduce that

$$|q_s^\eta - q_s^\zeta| \lesssim \int_s^t |q_r^\eta - q_r^\zeta|\, \mathrm{d}r + \|\eta - \zeta\|_{p;[s,t]}$$

for all $s \in [0,t]$, and thus, by Grönwall's inequality, that

$$\|q^\eta - q^\zeta\|_\infty \lesssim \|\eta - \zeta\|_{p;[0,t]}. \tag{6.33}$$

Since the state variables q^η, q^ζ, R and γ are uniformly bounded, we are free to modify the coefficients b_μ, λ and ψ outside of some large ball containing the domain of the state variables in its interior, without affecting the solutions q^η, q^ζ. We may therefore pretend that actually $b_\mu \in \mathrm{Lip}_b$ and $\lambda, \psi \in C_b^3$, so that in particular the hypotheses of part (iii) of Theorem 3.4.3 are satisfied. By the same argument, we may also suppose that f and g are Lipschitz in q.

By part (iii) of Theorem 3.4.3 combined with (6.33), we obtain

$$\left\| \int_0^\cdot \psi(q_s^\eta, \gamma_s)\, \mathrm{d}\boldsymbol{\eta}_s - \int_0^\cdot \psi(q_s^\zeta, \gamma_s)\, \mathrm{d}\boldsymbol{\zeta}_s \right\|_{p;[0,t]} \lesssim \|\boldsymbol{\eta}; \boldsymbol{\zeta}\|_p. \tag{6.34}$$

Using (6.33) and (6.34), we have, for any terminal condition $(t, \mu, \Sigma, a) \in [0,T] \times K$,

that

$$\begin{aligned}
&\left|v^{\eta}(t,\mu,\Sigma,a) - v^{\zeta}(t,\mu,\Sigma,a)\right| \\
&\quad\leq \sup_{u\in\mathcal{U}^M}\left|\int_0^t \left(f(q_s^{\eta},R_s,\gamma_s,u_s) - f(q_s^{\zeta},R_s,\gamma_s,u_s)\right)\mathrm{d}s\right. \\
&\qquad\qquad \left. + \int_0^t \psi(q_s^{\eta},\gamma_s)\,\mathrm{d}\boldsymbol{\eta}_s - \int_0^t \psi(q_s^{\zeta},\gamma_s)\,\mathrm{d}\boldsymbol{\zeta}_s + g(q_0^{\eta},R_0,\gamma_0) - g(q_0^{\zeta},R_0,\gamma_0)\right| \\
&\quad\lesssim \sup_{u\in\mathcal{U}^M}\left(\int_0^t |q_s^{\eta} - q_s^{\zeta}|\,\mathrm{d}s + \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p + |q_0^{\eta} - q_0^{\zeta}|\right) \\
&\quad\lesssim \|\boldsymbol{\eta};\boldsymbol{\zeta}\|_p,
\end{aligned}$$

and we conclude as we did in the proof of Theorem 5.6.2. □

Remark 6.5.9. As we have seen, in order to prevent degeneracy of the control problem it is necessary to control the derivative of the parameters, rather than controlling them directly. This allows us to calibrate, not only beliefs about reasonable values the parameters could take, but also at what rate they should able to vary. For example, if one believes that the true parameters should remain fairly constant then one can put a large penalty on the magnitude of this derivative. In fact, by taking the penalty to be infinite for all non-zero controls (derivatives), we obtain a setting with unknown parameters which are constant in time. The discrete-time results of [18] suggest that we should then expect the resulting filter to converge to the true parameter. (Although our observations are not independent and identically distributed as in [18], under reasonable conditions they are ergodic, and this leads to consistency properties in the likelihood function; see Douc et al. [27]. In this case, we expect that this would lead to the nonlinear expectation asymptotically converging to the 'true' expectation, and the analysis of [18] further suggests an interpretation of the nonlinear expectation in terms of confidence intervals.) Establishing precise convergence results could be the subject of future research.

Chapter 7

A Linear-Quadratic Robust Filter

In this short chapter we consider a special case of uncertainty in the Kalman–Bucy filter, in which our control problem turns out to be particularly amenable to efficient numerical computations.

7.1 A linear-quadratic pathwise control problem

Let us suppose that the signal S and observation Y are one-dimensional and, as in Chapter 6, satisfy the equations

$$\begin{aligned} \mathrm{d}S_t &= (\theta_t + \alpha_t S_t)\,\mathrm{d}t + \sigma_t\,\mathrm{d}B^1_t, \\ \mathrm{d}Y_t &= c_t S_t\,\mathrm{d}t + \mathrm{d}B^2_t, \\ \mathrm{d}\langle B^1, B^2\rangle_t &= \rho_t\,\mathrm{d}t, \end{aligned}$$

with $S_0 \sim N(\mu_0, \Sigma_0)$ for some $\mu_0 \in \mathbb{R}$ and $\Sigma_0 > 0$.

We recall again the standard Kalman–Bucy filtering equations:

$$\mathrm{d}q_s = (\theta_s + \alpha_s q_s)\,\mathrm{d}s + (R_s c_s + \sigma_s \rho_s)(\mathrm{d}Y_s - c_s q_s\,\mathrm{d}s), \tag{7.1}$$

$$\frac{\mathrm{d}R_s}{\mathrm{d}s} = \sigma_s^2 + 2\alpha_s R_s - (R_s c_s + \sigma_s \rho_s)^2. \tag{7.2}$$

In this section we shall suppose that the parameters θ and μ_0 are uncertain, but that all other parameters are known smooth functions. In this special case there is no uncertainty in the equation (7.2), so the trajectory of the posterior variance R is known. Accordingly, we shall take $g = g(q_0)$ to be a function of q_0 only.

As usual we have the value function κ, which here takes the form:

$$\kappa(t, \mu) = \inf_\theta \left\{ \int_0^t f(s, \theta_s)\,\mathrm{d}s - \int_0^t c_s q_s\,\mathrm{d}\mathbf{Y}_s + \frac{1}{2}\int_0^t (c_s q_s)^2\,\mathrm{d}s + g(q_0) \right\}.$$

We shall consider the quadratic penalty functions

$$f(s,\theta) = \frac{\nu}{2}(\theta - \hat{\theta}_s)^2 \qquad \text{and} \qquad g(q_0) = \delta(q_0 - \hat{\mu}_0)^2$$

for some $\nu, \delta > 0$, $\hat{\mu}_0 \in \mathbb{R}$ and some continuous function $\hat{\theta}\colon [0,T] \to \mathbb{R}$.

We thus have the linear controlled equation

$$\mathrm{d}q_s = (\theta_s + \alpha_s q_s)\,\mathrm{d}s + (R_s c_s + \sigma_s \rho_s)(\mathrm{d}Y_s - c_s q_s\,\mathrm{d}s)$$

with the terminal condition

$$q_t = \mu,$$

and the value function

$$\kappa(t,\mu) = \inf_\theta \left\{ \frac{1}{2}\int_0^t \left(\nu(\theta_s - \hat{\theta}_s)^2 + c_s^2 q_s^2\right)\mathrm{d}s - \int_0^t c_s q_s\,\mathrm{d}\mathbf{Y}_s + \delta(q_0 - \hat{\mu}_0)^2 \right\}.$$

We note in particular that this is a linear-quadratic pathwise control problem.

Approximating $\mathbf{Y}$ by a smooth path η, and writing κ^η for the approximate value function, we recall that κ^η satisfies the HJB equation:

$$\begin{aligned}
0 &= \frac{\partial \kappa^\eta}{\partial t} \\
&\quad + \sup_\theta \left\{ \Big(\theta + \alpha_t\mu + (R_t c_t + \sigma_t\rho_t)(\dot{\eta}_t - c_t\mu)\Big)\frac{\partial \kappa^\eta}{\partial \mu} - \left(\frac{\nu}{2}(\theta - \hat{\theta}_t)^2 + \frac{c_t^2\mu^2}{2} - c_t\mu\dot{\eta}_t\right)\right\} \\
&= \frac{\partial \kappa^\eta}{\partial t} + \Big(\hat{\theta}_t + \alpha_t\mu + (R_t c_t + \sigma_t\rho_t)(\dot{\eta}_t - c_t\mu)\Big)\frac{\partial \kappa^\eta}{\partial \mu} + \frac{1}{2\nu}\left(\frac{\partial \kappa^\eta}{\partial \mu}\right)^2 - \frac{c_t^2\mu^2}{2} + c_t\mu\dot{\eta}_t,
\end{aligned}$$

with $\kappa^\eta(0,\mu) = \delta(\mu - \hat{\mu}_0)^2$.

Adopting the ansatz

$$\kappa^\eta(t,\mu) = \xi_t(\mu - \beta_t)^2 + \tau_t$$

and substituting into the above, we obtain

$$\begin{aligned}
0 = {} & \Big(\dot{\xi}_t\big(\mu^2 - 2\mu\beta_t + \beta_t^2\big) - 2\xi_t\dot{\beta}_t(\mu - \beta_t) + \dot{\tau}_t\Big) \\
& + \Big(\hat{\theta}_t + \alpha_t\mu + (R_t c_t + \sigma_t\rho_t)(\dot{\eta}_t - c_t\mu)\Big)\big(2\xi_t(\mu - \beta_t)\big) \\
& + \frac{2\xi_t^2}{\nu}\big(\mu^2 - 2\mu\beta_t + \beta_t^2\big) - \frac{c_t^2\mu^2}{2} + c_t\mu\dot{\eta}_t.
\end{aligned}$$

Equating coefficients of μ^2 and of μ, we have

$$\begin{aligned}
0 &= \dot{\xi}_t + 2\xi_t\big(\alpha_t - c_t(R_t c_t + \sigma_t\rho_t)\big) + \frac{2\xi_t^2}{\nu} - \frac{c_t^2}{2}, \\
0 &= -2\dot{\xi}_t\beta_t - 2\xi_t\dot{\beta}_t - 2\xi_t\beta_t\big(\alpha_t - c_t(R_t c_t + \sigma_t\rho_t)\big) \\
&\qquad + 2\xi_t\big(\hat{\theta}_t + (R_t c_t + \sigma_t\rho_t)\dot{\eta}_t\big) - \frac{4\xi_t^2\beta_t}{\nu} + c_t\dot{\eta}_t.
\end{aligned}$$

Combining these equations, we obtain

$$0 = -2\xi_t\dot{\beta}_t + 2\xi_t\beta_t\big(\alpha_t - c_t(R_tc_t + \sigma_t\rho_t)\big) + 2\xi_t\big(\hat{\theta}_t + (R_tc_t + \sigma_t\rho_t)\dot{\eta}_t\big) + c_t\dot{\eta}_t - c_t^2\beta_t.$$

Rearranging then gives

$$\begin{aligned}\dot{\beta}_t &= \beta_t\big(\alpha_t - c_t(R_tc_t + \sigma_t\rho_t)\big) + \big(\hat{\theta}_t + (R_tc_t + \sigma_t\rho_t)\dot{\eta}_t\big) + \frac{c_t\dot{\eta}_t}{2\xi_t} - \frac{c_t^2\beta_t}{2\xi_t} \\ &= \hat{\theta}_t + \alpha_t\beta_t + \left((R_tc_t + \sigma_t\rho_t) + \frac{c_t}{2\xi_t}\right)\big(\dot{\eta}_t - c_t\beta_t\big).\end{aligned}$$

Replacing η with[1] Y, we have thus derived the robust filtering equations:

$$\begin{aligned}\frac{\mathrm{d}R_t}{\mathrm{d}t} &= \sigma_t^2 + 2\alpha_tR_t - (R_tc_t + \sigma_t\rho_t)^2, && R_0 = \Sigma_0, \\ \frac{\mathrm{d}\xi_t}{\mathrm{d}t} &= -2\xi_t\big(\alpha_t - c_t(R_tc_t + \sigma_t\rho_t)\big) - \frac{2\xi_t^2}{\nu} + \frac{c_t^2}{2}, && \xi_0 = \delta, \qquad (7.3) \\ \mathrm{d}\beta_t &= \big(\hat{\theta}_t + \alpha_t\beta_t\big)\,\mathrm{d}t + \left(R_tc_t + \sigma_t\rho_t + \frac{c_t}{2\xi_t}\right)\big(\mathrm{d}Y_t - c_t\beta_t\,\mathrm{d}t\big), && \beta_0 = \hat{\mu}_0. \qquad (7.4)\end{aligned}$$

The value function is then given by $\kappa(t,\mu) = \xi_t(\mu - \beta_t)^2$, up to an additive constant, depending only on time t, which we may set to zero.

Remark 7.1.1. We see that in this case both the 'risk' R and 'uncertainty' ξ satisfy Riccati equations. Moreover, comparing the standard Kalman–Bucy filtering equation (7.1) with the robust version (7.4), we note that the predictor term remains unchanged, but that the Kalman gain has increased by the amount $c_t/2\xi_t$, where ξ_t is calculated via (7.3).

7.2 A numerical example

As an example, we solve the equations derived above with $\alpha = -0.2$, $\sigma = 0.3$, $c = 1.5$, $\rho = 0$, $\Sigma_0 = 1$, $\nu = 0.5$ and $\delta = 1$. We shall suppose that the 'true' parameters are given by

$$\theta_t = 2\sin\left(\frac{t}{10}\right) \qquad \text{and} \qquad \mu_0 = 0,$$

but we shall take as 'estimated' parameters

$$\hat{\theta}_t = 0 \qquad \text{and} \qquad \hat{\mu}_0 = 0.$$

In particular, the estimate $\hat{\theta}$ is simply the long-time average of θ.

In Figure 7.1 we plot the value function κ as a function of t and μ.

[1]Since the integrand of the integral against Y has finite variation, the second order process $\mathbb{Y}$ plays no further role.

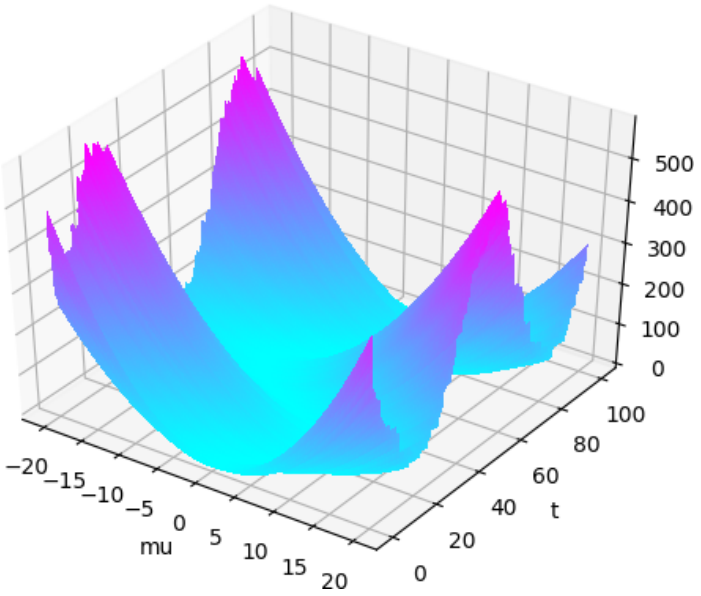


Figure 7.1: Solution of the robust filtering equations.

We can then use this solution to compute robust estimates of the signal. In Figure 7.2 we compare the standard Kalman–Bucy filter, computed using both the 'true' and 'estimated' parameters, with the minimum point of the value function, which we recall represents the posterior distribution which we consider to be the most reasonable given our observations.

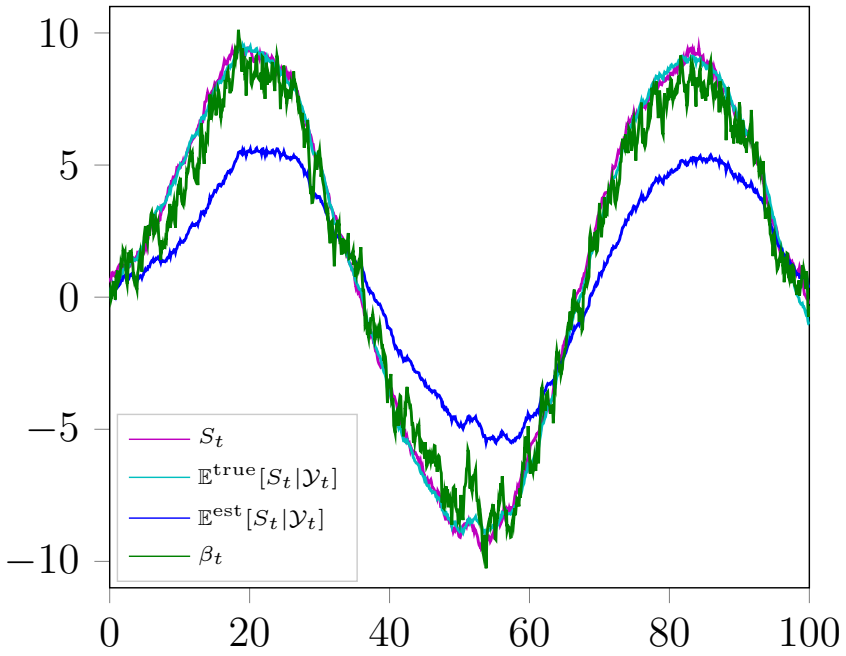


Figure 7.2: Estimates of S_t.

In Figures 7.3 and 7.4 we estimate two nonlinear functionals of the signal, namely S^2 and $(S-2)^+$ respectively.

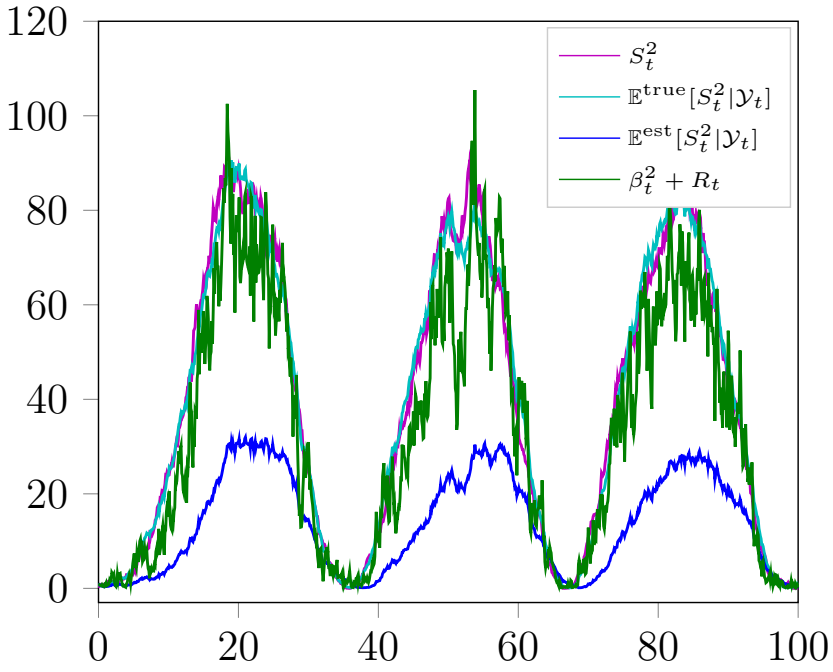


Figure 7.3: Estimates of S_t^2.

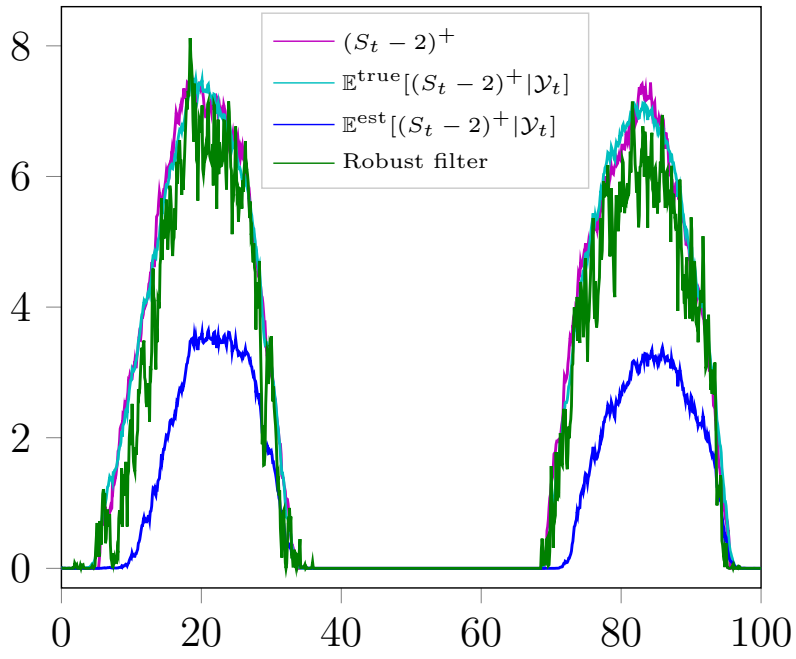


Figure 7.4: Estimates of $(S_t-2)^+$.

Bibliography

[1] Y. Achdou, G. Barles, H. Ishii, and G. L. Litvinov. *Hamilton–Jacobi Equations: Approximations, Numerical Analysis and Applications.* Springer-Verlag, Berlin, Heidelberg, 2013.

[2] M. Aliyu and E. Boukas. Mixed H_2/H_∞ nonlinear filtering. *Int. J. Robust Nonlinear Control*, 19:394–417, 2009.

[3] A. L. Allan and S. N. Cohen. Parameter uncertainty in the Kalman–Bucy filter. *SIAM J. Control Optim.*, 57:1646–1671, 2019.

[4] A. L. Allan and S. N. Cohen. Pathwise stochastic control with applications to robust filtering. arXiv:1902.05434, 2019.

[5] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering.* Springer, New York, 2009.

[6] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations.* Birkhäuser, Boston, 2008.

[7] M. Bardi and F. Da Lio. On the Bellman equation for some unbounded control problems. *Nonlinear Differ. Equ. Appl.*, 4:491–510, 1997.

[8] G. Barles and P. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4:271–283, 1991.

[9] D. S. Bernstein and W. M. Haddad. Steady-state Kalman filtering with an H_∞ error bound. *Systems & Control Letters*, 12:9–16, 1989.

[10] A. V. Borisov. Minimax a posteriori estimation of the Markov processes with finite state spaces. *Autom. Remote Control*, 69:233–246, 2008.

[11] A. V. Borisov. The Wonham filter under uncertainty: a game-theoretic approach. *Automatica*, 47:1015–1019, 2011.

[12] R. Buckdahn and J. Ma. Pathwise stochastic control problems and stochastic HJB equations. *SIAM J. Control Optim.*, 45:2224–2256, 2007.

[13] M. Caruana and P. K. Friz. Partial differential equations driven by rough paths. *J. Differential Equations*, 247:140–173, 2009.

[14] M. Caruana, P. K. Friz, and H. Oberhauser. A (rough) pathwise approach to a class of non-linear stochastic partial differential equations. *Ann. Inst. H. Poincaré Anal. Non Linéaire*, 28:27–46, 2011.

[15] X. Chen and K. Zhou. H_∞ Gaussian filter on infinite time horizon. *IEEE Trans. Circuits Syst. I: Fundam. Theory Appl.*, 49:674–679, 2002.

[16] B. Cipra. Engineers look to Kalman filtering for guidance. *SIAM News*, 26, 1993.

[17] S. N. Cohen. Uncertainty and filtering of hidden Markov models in discrete time. arXiv:1606.00229v4, 2016.

[18] S. N. Cohen. Data-driven nonlinear expectations for statistical uncertainty in decisions. *Electron. J. Stat.*, 11:1858–1889, 2017.

[19] S. N. Cohen and R. J. Elliott. *Stochastic Calculus and Applications*. Springer, New York, 2nd edition, 2015.

[20] M. G. Crandall, H. Ishii, and P.-L. Lions. User's guide to viscosity solutions of second order partial differential equations. *Bull. Amer. Math. Soc.*, 27:1–67, 1992.

[21] D. Crisan, J. Diehl, P. K. Friz, and H. Oberhauser. Robust filtering: Correlated noise and multidimensional observation. *Annals of Applied Probability*, 23:2139–2160, 2013.

[22] D. Crisan and B. Rozovskii. *The Oxford handbook of nonlinear filtering*. Oxford University Press, 2011.

[23] M. H. A. Davis and G. Burstein. Anticipative stochastic control. *Proceedings of 30^{th} IEEE Conf. Decision Control*, pages 1830–1835, 1991.

[24] M. H. A. Davis and G. Burstein. A deterministic approach to stochastic optimal control with application to anticipative control. *Stoch. Stoch. Rep.*, 40:203–256, 1992.

[25] F. Delbaen, S. Peng, and E. Rosazza Gianin. Representation of the penalty term of dynamic concave utilities. *Finance and Stochastics*, 14:449–472, 2010.

[26] J. Diehl, P. K. Friz, and P. Gassiat. Stochastic control with rough paths. *Appl. Math. Optim.*, 75:285–315, 2017.

[27] R. Douc, E. Moulines, J. Olsson, and R. van Handel. Consistency of the maximum likelihood estimator for general hidden Markov models. *Ann. Stat.*, 39:474–513, 2011.

[28] L. C. Evans. *Partial Differential Equations.* American Mathematical Society, 2nd edition, 2010.

[29] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6:429–447, 2002.

[30] H. Föllmer and A. Schied. *Stochastic Finance, An Introduction in Discrete Time.* De Gruyter, Berlin, 4th edition, 2016.

[31] M. Frittelli and E. Rosazza Gianin. Putting order in risk measures. *Journal of Banking and Finance*, 26:1473–1486, 2002.

[32] P. K. Friz and M. Hairer. *A Course on Rough Paths, With an Introduction to Regularity Structures.* Springer, Switzerland, 2014.

[33] P. K. Friz and H. Oberhauser. Rough path stability of (semi-)linear SPDEs. *Probab. Theory Relat. Fields*, 158:401–434, 2014.

[34] P. K. Friz and N. B. Victoir. *Multidimensional Stochastic Processes as Rough Paths.* Cambridge University Press, 2010.

[35] P. K. Friz and H. Zhang. Differential equations driven by rough paths with jumps. *J. Differential Equations*, 264:6226–6301, 2018.

[36] M. Fujisaki, G. Kallianpur, and H. Kunita. Stochastic differential equations for the non linear filtering problem. *Osaka J. Math.*, 9(1):19–40, 1972.

[37] J. George. Robust Kalman–Bucy filter. *IEEE Trans. Automat. Contr.*, 58:174–180, 2013.

[38] M. S. Grewal and A. P. Andrews. Applications of Kalman filtering in aerospace 1960 to the present. *IEEE Control Systems Magazine*, 30:69–78, 2010.

[39] M. Gubinelli. Controlling rough paths. *J. Funct. Anal.*, 216:86–140, 2004.

[40] M. Hazewinkel, S. I. Marcus, and H. J. Sussmann. Nonexistence of finite-dimensional filters for conditional statistics of the cubic sensor problem. *Syst. Contr. Lett.*, 3:331–340, 1983.

[41] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Basic Eng.*, 82:35–45, 1960.

[42] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *J. Basic Eng.*, 83:95–108, 1961.

[43] P. P. Khargonekar, M. A. Rotea, and E. Baeyens. Mixed H_2/H_∞ filtering. *Int. J. Robust Nonlinear Control*, 6:313–330, 1996.

[44] H. J. Kushner. On the differential equations satisfied by conditional probability densities of Markov processes, with applications. *SIAM J. Control*, 2(1):106–119, 1964.

[45] H. J. Kushner. Dynamical equations for optimal nonlinear filtering. *J. Differential Equations*, 3(2):179–190, 1967.

[46] F. L. Lewis, L. Xie, and D. Popa. *Optimal and Robust Estimation.* CRC Press, Taylor & Francis Group, 2nd edition, 2008.

[47] P.-L. Lions and P. Souganidis. Differential games, optimal control and directional derivatives of viscosity solutions of Bellman's and Isaacs' equations. *SIAM J. Control Optim.*, 23:566–583, 1985.

[48] P.-L. Lions and P. E. Souganidis. Fully nonlinear stochastic partial differential equations: non-smooth equations and applications. *C. R. Acad. Sci. Paris Sér. I*, 327:735–741, 1998.

[49] R. Liptser and A. Shiryaev. *Statistics of Random Processes I-II.* Springer-Verlag, Berlin, 2nd edition, 2001.

[50] T. J. Lyons. Differential equations driven by rough signals. *Rev. Mat. Iberoamericana*, 14:215–310, 1998.

[51] C. J. Martin and M. Mintz. Robust filtering and prediction for linear systems with uncertain dynamics: a game-theoretic approach. *IEEE Trans. Automat. Contr.*, 28:888–896, 1983.

[52] G. Miller and A. Pankov. Filtration of a random process in a statistically uncertain linear stochastic differential system. *Autom. Remote Control*, 66:53–64, 2005.

[53] S. K. Mitter. Existence and nonexistence of finite-dimensional filters. *Rend. Sem. Mat. Univ. Politec. Torino*, Special Issue:173–188, 1982.

[54] P. Neveux, E. Blanco, and G. Thomas. Robust filtering for linear time-invariant continuous systems. *IEEE Trans. Sig. Process.*, 55:4752–4757, 2007.

[55] D. Ocone and E. Pardoux. *A Lie algebraic criterion for nonexistence of finite-dimensionally computable filters, in Stochastic Partial Differential Equations and Applications II*, pages 197–204. Springer, Berlin, 1989.

[56] L. C. G. Rogers. Pathwise stochastic optimal control. *SIAM J. Control Optim.*, 46:1116–1132, 2007.

[57] K. V. Siemenikhin. Minimax linear filtering of a random sequence with uncertain covariance function. *Autom. Remote Control*, 77:226–241, 2016.

[58] K. V. Siemenikhin, M. V. Lebedev, and E. P. Platonov. Kalman filtering by minimax criterion with uncertain noise intensity functions. *Proceedings of the 44^{th} IEEE CDC-ECC*, pages 1929–1934, 2005.

[59] R. L. Stratonovich. Conditional Markov processes. *Theory Probab. Appl.*, 5(2):156–178, 1960.

[60] D. W. Stroock. *Essentials of Integration Theory for Analysis.* Springer, New York, 2011.

[61] S. Verdú and H. V. Poor. Minimax linear observers and regulators for stochastic systems with uncertain second-order statistics. *IEEE Trans. Automat. Contr.*, 29:499–511, 1984.

[62] A. Wald. Statistical decision functions which minimize the maximum risk. *Ann. of Math.*, 46:265–280, 1945.

[63] W. M. Wonham. Some applications of stochastic differential equations to optimal nonlinear filtering. *J. SIAM Control Ser. A*, 2:347–369, 1965.

[64] L. Xie, C. E. de Souza, and M. Fu. H_∞ estimation for discrete-time linear uncertain systems. *Int. J. Robust Nonlin.*, 1:111–123, 1991.

[65] G.-H. Yang and D. Ye. Robust H_∞ filter design for linear systems with time-varying uncertainty. *Internat. J. Control*, 82:517–524, 2009.

[66] J. Yong and X. Y. Zhou. *Stochastic Controls, Hamiltonian Systems and HJB Equations*. Springer-Verlag, New York, 1999.

[67] J. Zhang, Y. Xia, and P. Shi. Parameter-dependent robust H_∞ filtering for uncertain discrete-time systems. *Automatica*, 45:560–565, 2009.

[68] T. Zhou. Sensitivity penalization based robust state estimation for uncertain linear systems. *IEEE Trans. Automat. Contr.*, 55:1018–1024, 2010.