

Filtering, Stability, and Robustness

Thesis by

Ramon van Handel

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2007

(Defended December 12, 2006)

© 2007

Ramon van Handel

All Rights Reserved

Acknowledgments

First and foremost I would like to thank my advisor, Hideo Mabuchi, who has not only given me the complete freedom to pursue my own research interests, but has always encouraged me to do so. It takes an exceptional advisor to provide trust and support for a student whose main interests lie outside those of his research group, and I am very grateful for the opportunity he has given me to be a part of that group throughout these exciting years. Hideo's inspiring scientific vision and helpful advice have always been very valuable to me, and have contributed to a unique graduate school experience.

I am indebted to all of my collaborators over the years—Luc Bouten, Pavel Chigansky, John Gough, Matt James, Mazyar Mirrahimi, and John Stockton—I have learned much from all of you, and I look forward with anticipation to our future interactions. I would like to thank particularly my two closest colleagues, Luc Bouten and Pavel Chigansky, for our fruitful collaborations and for their friendship.

At Caltech, I have benefited from the exceptional group of students and post-docs, past and present, that have been part of MabuchiLab. I have had many interesting interactions with Mike Armen, John Au, Nicole Czakon, Nathan Hodas, Asa Hopkins, Joe Kerckhoff, Ben Lev, Tim McGarvey, Kevin McHale, Tony Miller, Gopal Sarma, and John Stockton, at least on those occasions when they were not squirreled away in the lab. On the postdoc side, I am grateful to André Conjusteau, JM Geremia, Jen Sokol, Jon Williams, Naoki Yamamoto, and particularly Andrew Doherty, who provided much inspiration when I was getting started in my first year at Caltech. Special thanks goes to Sheri Stoll, without whose care and unfailing competence we would likely all be quite helpless; but I am particularly grateful for the many interesting coffee-break-time discussions, which have often provided an enjoyable and much needed diversion from long days filled with science.

During my graduate career I have interacted with a number of faculty outside our research group whose comments and insight have been particularly helpful to me; in this context, I would like to thank in particular Profs. P. S. Krishnaprasad, Jerry Marsden, Richard Murray, Houman Owhadi, and Anders Rantzer. I am indebted to Prof. Krishnaprasad, who arranged for me to visit the Institute for Systems Research at the University of Maryland in 2005. This visit was extremely interesting and fruitful for me, and many of the ideas that have led to this thesis germinated at that time. I highly appreciate the interest he has shown in my research, and am honored that he has agreed to serve on my thesis committee.

To my Dutch friends: Michiel, Pim, Wim, Daan, Astrid, Eva, en alle andere bewoners van Villa Buitenrust—het is vaak moeilijk tussen alle drukte goed contact te houden, en ik ben niet trots op de gelegenheden wanneer ik door de drukte niet meer heb (terug) geschreven. Jullie weten allemaal, echter, dat jullie altijd een

speciale plek in mijn hart innemen. Ik stel jullie vriendschap ontzettend op prijs en hoop jullie snel weer te zien, waar op de wereld we dan ook terecht komen.

Most of all, my thanks go to my family at home, consisting of three bears, a kitten, a seal, a lion, and, most importantly, Paige, the love of my life—without whom I could not do anything, including this thesis.

2000 *Mathematics Subject Classification*. Primary 93E11; secondary 93E20, 93E15, 81S25, 81P15, 62F15, 60G35, 60H07, 60J27, 34D40

ABSTRACT. The theory of nonlinear filtering concerns the optimal estimation of a Markov signal in noisy observations. Such estimates necessarily depend on the model that is chosen for the signal and observations processes. This thesis studies the sensitivity of the filter to the choice of underlying model over long periods of time, in the framework of continuous time filtering with white noise type observations. A good understanding of filter sensitivity is important for the practical application of nonlinear filters: a physical system is never exactly described by an idealized Markov model, so that high sensitivity to the underlying model would be fatal to the reliable application of nonlinear filters on all but very short time scales. Sensitivity is also directly related to the approximation error in approximate nonlinear filtering.

The first topic of this thesis is the asymptotic stability of the filter—i.e., the long-time sensitivity of the filter to the initial measure. This is a much-studied problem in nonlinear filtering. Here I study filter stability using the theory of conditional diffusions. This leads to some improvements on pathwise stability bounds, and to new insight into existing stability results in a fully probabilistic setting. As a matter of independent interest, I develop in detail the theory of conditional diffusions for finite-state Markov signals and clarify the duality between estimation and stochastic control in this context.

The second topic of this thesis is the sensitivity of the nonlinear filter to the model parameters of the signal and observations processes. This section concentrates on the finite state case, where the corresponding model parameters are the jump rates of the signal, the observation function, and the initial measure. The main result is that the expected difference between the filters with the true and modified model parameters is bounded uniformly on the infinite time interval, and converges to zero uniformly in time as the approximate model converges to the true model, provided that the signal process satisfies a mixing property. Such filters are consequently extremely robust to variations in the model parameters. The proof of this result uses properties of the stochastic flow generated by the filter on the simplex, as well as the Malliavin calculus and the associated anticipative stochastic calculus.

The third and final topic of this thesis is the asymptotic stability of non-commutative (quantum) filters. I begin by developing quantum filtering theory using reference probability methods. The stability of the resulting filters is not easily studied using the preceding methods, as smoothing violates the non-demolition requirement. Fortunately, progress can be made by randomizing the initial state of the filter. Using this technique, I prove that the filtered estimate of the measurement observable is stable regardless of the underlying model, provided that the initial states are absolutely continuous in a suitable sense.

Contents

Acknowledgments	iii
Abstract	v
List of Notations	xi
Introduction	1
0.1. Nonlinear filtering: Stability, robustness, and applications	1
0.1.1. Optimal filtering	1
0.1.2. Approximations	3
0.1.3. Filter stability	3
0.1.4. Conditional diffusions and filter stability	5
0.1.5. Filter robustness	7
0.1.6. Quantum filtering and filter stability	9
0.1.7. Applications of nonlinear filtering	10
0.2. A suggestive numerical experiment	11
0.3. Outline of this thesis, main results, and outlook	16
0.4. Other work	18
Chapter 1. Fundamentals of Nonlinear Filtering	19
1.1. Nonlinear filtering of a signal in white noise	20
1.1.1. The basic model	20
1.1.2. An explicit construction	20
1.1.3. The Kallianpur-Striebel formula	21
1.1.4. The Zakai and Kushner-Stratonovich equations	23
1.1.5. The innovations process	25
1.2. Finite-state Markov signals and the Wonham filter	26
1.3. Nonlinear filtering for diffusions	28
1.4. Pathwise filtering	29
1.4.1. The pathwise Kallianpur-Striebel formula	30
1.4.2. The pathwise Wonham filter	31
1.4.3. Pathwise diffusion filtering	33
Chapter 2. Filter Stability and Conditional Signals: Finite State Space	35
2.1. Conditional finite-state signals: A direct approach	35
2.1.1. Change of measure for finite-state signals	36
2.1.2. Conditional signal—forward case	38
2.1.3. Conditional signal—time-reversed case	41
2.1.4. Smoothing and path estimation	42
2.1.5. A note on general Markov signals	43
2.2. On the duality between estimation and stochastic control	43

2.2.1.	The variational Kallianpur-Striebel formula	44
2.2.2.	Dynamic programming—forward case	45
2.2.3.	Dynamic programming—time-reversed case	49
2.2.4.	Some more variations on the same theme	51
2.3.	Exponential stability of the Wonham filter	53
2.3.1.	Filter stability	53
2.3.2.	Exponential stability: A coupling proof	55
Chapter 3.	Model Robustness of the Nonlinear Filter: Finite State Space	59
3.1.	Introduction	59
3.1.1.	A little deterministic intuition	60
3.1.2.	Model robustness of the Wonham filter	61
3.1.3.	Notation	63
3.2.	Stochastic semiflow of the Wonham filter	64
3.3.	Exponential estimates for the derivative of the filter	67
3.4.	Proof of the main result	72
3.5.	Some technical lemmas	79
Chapter 4.	Filter Stability and Conditional Signals: Continuous State Space	83
4.1.	The filtering model	83
4.2.	Preliminaries: Some conditional signal theory	84
4.2.1.	Conditional signal—time-reversed case	84
4.2.2.	Conditional signal—gauge transformation	86
4.3.	Filter stability for potential diffusions	89
4.3.1.	Filter stability	89
4.3.2.	Uniform convexity of the value function	91
4.3.3.	Proof of the main result	94
4.3.4.	A probabilistic proof of a Brascamp-Lieb inequality	96
4.4.	A strong stability result for stable signals	99
Chapter 5.	Quantum Filtering and Filter Stability	103
5.1.	Conditional expectations and the Bayes formula	103
5.1.1.	Conditional expectations in quantum probability	103
5.1.2.	The Bayes formula	105
5.2.	Quantum filtering: A reference probability approach	108
5.2.1.	The basic model	108
5.2.2.	The Kallianpur-Striebel formula	110
5.2.3.	The quantum filtering equations	114
5.2.4.	Imperfect detection	116
5.3.	A filter stability result	118
5.3.1.	Incorrect initialization and randomization of the initial state	118
5.3.2.	A generic stability result	121
Appendix A.	Elements of the Malliavin Calculus	125
A.1.	The Malliavin derivative: Definition and elementary properties	125
A.2.	The Malliavin derivative of a stochastic flow	126
A.3.	The Clark-Haussmann-Ocone formula	127
A.4.	The Skorokhod integral: Definition and elementary properties	127
A.5.	Anticipative stochastic calculus	128

Appendix B. Elements of Quantum Probability	129
B.1. Quantum probability	129
B.1.1. Quantum probability spaces	129
B.1.2. The spectral theorem	132
B.1.3. Unbounded observables	134
B.2. Quantum noise	136
B.2.1. Fock space and the fundamental processes	136
B.2.2. Quantum stochastic calculus on the exponential domain	139
B.2.3. Quantum stochastic differential equations	144
Bibliography	147

List of Notations

$(\mathcal{A}, \mathbb{P})$	Quantum probability space	130
$A_t, A_t^\dagger, \Lambda_t$	Fundamental noises	138
B_t	Observation noise	20
$C(\mathbb{R}_+; \mathbb{R}^p)$	Space of continuous \mathbb{R}^p -valued paths	20
$D(\mathbb{R}_+; \mathbb{S})$	Skorokhod space of \mathbb{S} -valued càdlàg paths	20
$D(\mu \ \tilde{\mu})$	Relative entropy	44
$D\pi_{s,t}(\mu)$	Directional derivative of Wonham filter	65
$E_X(\cdot)$	Spectral measure of the self-adjoint operator X	134
H	Observation matrix for finite-state signal	27
$H^y(t, x)$	Pathwise filtering density for diffusion signal	34
$H_T(x, y)$	Energy function $-\log \tilde{Z}_T(x, y)$	44
$I(H)$	Total information	44
Q_t, P_t, Λ_t	Fundamental noises (self-adjoint form)	138
$T\mathcal{S}^{d-1}$	Tangent space to \mathcal{S}^{d-1}	64
$U_{s,t}$	Stochastic flow of unnormalized filter $\sigma_{s,t}(\mu) = U_{s,t}\mu$	64
$X \eta \mathcal{C}$	X is affiliated to \mathcal{C}	135
X^*	Adjoint of X	134
X^\dagger	Adjoint of an allowable operator on \mathbb{D}	141
X_t	Signal process	20
Y_t	Observation process	20
$Z_T(x, y)$	Girsanov transformation in Kallianpur-Striebel formula	22
Z_t, Y_t	Quantum observation noise and process	108
Δ^{d-1}	The d -simplex	64
Λ, λ_{ij}	Transition intensities of finite-state Markov signal	26
Φ	Vacuum vector	138
$\Pi_T(X, y)$	Pathwise conditional measure	31
$\Pi_T^\mu(X, y)$	Conditional measure for signal with initial measure μ	55
$\Sigma(x)$	Normalization map $\Sigma(x) = x/ x $	65
$\ v\ $	ℓ_2 -norm of vector v	63
$\ v\ _p$	ℓ_p -norm of vector v	63
$\prec\prec$	Absolutely continuous states	119
\overline{W}_t	Innovations process	25
\overline{X}_t	Time-reversed finite-state signal	41
\overline{z}_t	Classical innovations process	114
$\overline{\lambda}_{ij}^{T,y}(t)$	Conditional transition intensities, time-reversed case	41
$\overline{\pi}_t$	Quantum filter	114
$\overline{\sigma}_t$	Quantum unnormalized filter	114
Γ	Standard symmetric Fock space	137
$\Gamma(\mathbb{H}_1)$	Symmetric Fock space over \mathbb{H}_1	136

$\Gamma_{[s]}, \Gamma_{[s,t]}, \Gamma_{[t]}$	Natural splitting of the Fock space	137
$\hat{\cdot}, \hat{+}$	Closed multiplication and addition	136
\mathbf{H}	Hilbert space	129
ι	*-isomorphism from spectral theorem	133
$\lambda_{ij}^{T,y}(t)$	Conditional transition intensities	39
\rightsquigarrow	Equivalent states	119
\mathbf{D}	Malliavin derivative	125
\mathbb{D}^∞	Random variables with arbitrary Malliavin derivatives	126
$\mathbb{D}^{k,p}$	Domains of iterated Malliavin derivatives	125
$\mathbb{D}_{\text{loc}}^{k,p}$	Localized domains of iterated Malliavin derivatives	126
$\mathbb{L}^{k,p}$	Regular Skorokhod integrable processes	128
$\mathbb{L}_{\text{loc}}^{k,p}$	Locally regular Skorokhod integrable processes	128
\mathbb{Q}^t	Quantum reference measure	110
\mathbb{R}_{++}^d	Positive orthant	64
\mathbb{S}	Signal state space	20
\mathbf{P}	Standard probability measure	20
\mathbf{P}^μ	Probability measure under which $X_0 \sim \mu$	54
\mathbf{Q}	Reference measure	21
$\tilde{\mathbf{P}}_{t,x}$	Probability measure under which $\tilde{X}_t = x$	85
\mathcal{F}_t	Standard filtration	20
\mathcal{F}_t^Y	Observation filtration	20
$\mathcal{L}(X)$	Lindblad generator	115
$\mathcal{L}^*(\varrho)$	Adjoint Lindblad generator	115
\mathcal{S}^{d-1}	Interior of the d -simplex	64
V_t, V_t	State transformation in quantum Kallianpur-Striebel formula	112
\mathcal{A}_t	Standard quantum filtration	108
\mathcal{B}	Initial system algebra	140
$\mathcal{B}(\mathbf{H})$	Von Neumann algebra of bounded operators on \mathbf{H}	129
\mathcal{C}'	The commutant of \mathcal{C}	104
\mathcal{L}	Generator of Markov signal process	23
$\mathcal{M}_d(\cdot)$	Algebra of $d \times d$ matrices	106
$\mathcal{N}(\mathcal{C})$	Algebra of normal operators affiliated to \mathcal{C}	136
$\mathcal{N}(\tilde{\mathcal{C}})$	Algebra of unbounded operators affiliated to $\tilde{\mathcal{C}}$	107
\mathcal{W}	Field algebra	138
\mathcal{Y}_t	Observation process (Y_t) filtration	109
\mathcal{Z}_t	Observation noise (Z_t) filtration	109
\mathbf{D}	Restricted exponential domain	140
$\mathbf{D}(X)$	Domain of X	134
\mathbf{E}	Exponential domain	137
\mathbf{h}	Initial system Hilbert space	140
$\mathbf{vN}(\mathcal{S})$	Von Neumann algebra generated by \mathcal{S}	131
μ_w	Wiener measure	21
$\frac{\mu_x}{\bar{X}}$	Law of the signal process	21
\bar{X}	Closure of X	136
$\pi_t(\mu)$	Wonham filter with initial condition μ	54
$\pi_t(f)$	Filtered estimate $\mathbf{E}(f(X_t) \mathcal{F}_t^Y)$	20
π_t^i	Conditional probability of finite-state signal	27
$\pi_{s,t}(\mu)$	Wonham filter with initial condition μ at time s	63

ρ	Initial state	140
$\sigma_t(f)$	Unnormalized filtered estimate	23
σ_t^i	Unnormalized conditional probability of finite-state signal	27
δ	Skorokhod integral	127
\tilde{H}	Misspecified observation matrix	63
$\tilde{\Lambda}, \tilde{\lambda}_{ij}$	Misspecified transition intensities	63
$\tau_t(f, y)$	Pathwise filtered estimate	31
$\tau_t^i(y)$	Pathwise conditional probability of finite-state signal	31
\tilde{W}_t	Time-reversed Wiener process	84
\tilde{X}_t	Time-reversed diffusion signal	84
\tilde{X}_t^*	Time-reversed conditional diffusion signal	85
\tilde{X}_t^u	Controlled time-reversed diffusion signal	84
$\tilde{Z}_T(x, y)$	Pathwise version of $Z(x, y)$	31
$\tilde{v}_{t,T}^i(y)$	Pathwise dual filter for finite-state signal	38
$\tilde{\sigma}_t$	Misspecified unnormalized filter	63
$\tilde{\pi}_t$	Misspecified Wonham filter	63
\otimes	Algebraic tensor product	140
φ	Vacuum state	138
ϱ_t	Conditional density matrix	115
$\varrho_t^y(x)$	(Pathwise) unnormalized filtering density for diffusion signal	85
ζ_t	Unnormalized conditional density matrix	115
$ v $	ℓ_1 -norm of vector v	27
$e(f)$	Exponential vector with amplitude function f	137
$f_t $	Restriction of f to $[0, t]$	137
$h(\cdot)$	Observation function	20
$j_t(X)$	Quantum time evolution	108
$p_t(x)$	Unconditional density of diffusion signal	84
v^*, M^*	Adjoint of vector v or matrix M	27
$v_{t,T}^i(y)$	Dual filter for finite-state signal	38
y_t	Classical observation process	114
z_t	Classical observation noise	112

Introduction

“Just what kind of a doctor are you?” he asked suspiciously.
“Well, you might say I’m a specialist,” said the doctor. “I specialize in noise—all kinds—from the loudest to the softest, and from the slightly annoying to the terribly unpleasant.”

—Norton Juster, *The Phantom Tollbooth*

0.1. Nonlinear filtering: Stability, robustness, and applications

Noise is ubiquitous both in nature and in engineering. Many systems are prone to some form of noise in their dynamics, be it due to the coupling of the system to a complex and unpredictable environment, due to our fundamental lack of knowledge about the system dynamics at the fast time scales, or due to the inherent quantum mechanical uncertainty that is unavoidably present at the smallest energy scales. Similarly, any observations we make of a system are likely to be corrupted by some amount of noise. Depending on the particular application, the presence of noise may be terribly unpleasant or only slightly annoying; but in either case one needs methodology to separate the useful signal from the noise. As such these methods play an important role across a broad spectrum of science and technology.

This thesis is about one such method—the method of *optimal filtering*, whose goal is optimal estimation, in the L^2 sense, of a Markov signal in noisy observations. To fix some ideas, let us briefly discuss this method in a fairly general context.

0.1.1. Optimal filtering. As is usual in probability theory, we fix some underlying probability space $(\Omega, \mathcal{F}, \mathbf{P})$. We presume that there is a Markov process X_t that lives on this probability space and takes values in some state space \mathbb{S} . In this thesis we will consider the case of continuous time $t \in \mathbb{R}_+$, though other time sets are common as well (discrete time \mathbb{N} , two-sided discrete or continuous time \mathbb{Z} or \mathbb{R} , or even higher-dimensional “times”, e.g., \mathbb{Z}^d , for image processing or the estimation of random fields). The process X_t represents the signal of interest—e.g., the dynamics of a physical system that we wish to observe.

Also on this probability space lives the observation process Y_t , which takes values in some space $Y_t \in \mathbb{O}$. Y_t need not be a Markov process itself, but the joint process (X_t, Y_t) is Markov. Y_t is the process we are allowed to observe, and should be correlated in some way with the signal process X_t . In this thesis, we will study *white noise type* observations: i.e., we will take (with $\mathbb{O} = \mathbb{R}$ or \mathbb{R}^p)

$$\dot{Y}_t = h(X_t) + \xi(t),$$

where h is called the *observation function* and ξ is white noise. In practice, as is usual in stochastic analysis, we will work with the integrated version

$$Y_t = \int_0^t h(X_s) ds + B_t,$$

where B_t is a Wiener process or Brownian motion. This circumvents questions of mathematical well-posedness of white noise.

Now suppose we have been observing the observation process Y_t from the initial time $t = 0$ up to some time $t = T$. The information we have collected through these observations is characterized by the σ -algebra $\mathcal{F}_T^Y = \sigma\{Y_t : 0 \leq t \leq T\}$. The goal of optimal filtering theory is to find a best L^2 -estimate $\pi_T(f)$ of some function of the signal process $f(X_T)$ based on the observations: i.e., we wish to find an \mathcal{F}_T^Y -measurable random variable $\pi_T(f)$ that minimizes the L^2 -error $\|\pi - f(X_T)\|_2$ over all \mathcal{F}_T^Y -measurable random variables π . It is well known that this estimate is uniquely determined by the conditional expectation

$$\pi_T(f) = \mathbf{E}(f(X_T)|\mathcal{F}_T^Y),$$

up to a.s. equivalence. The goal of filtering theory then becomes to provide an explicit expression for $\pi_T(f)$ in terms of the observation history $\{Y_t : 0 \leq t \leq T\}$.

If we wish to estimate $f(X_t)$ for some $t < T$, or if we are interested in estimating some functional $f(X_{[0,T]})$ of the entire history of the signal, we could form the estimates $\mathbf{E}(f(X_t)|\mathcal{F}_T^Y)$ and $\mathbf{E}(f(X_{[0,T]})|\mathcal{F}_T^Y)$, respectively. This is known as *smoothing* (as opposed to filtering), and the latter case is also called *path estimation*. Even when we are interested in properties of the optimal filter, the consideration of these smoothers will play an important role in our proofs.

One of the main results of optimal filtering theory with white noise observations can be stated as follows. Under suitable technical conditions, The filtered estimate $\pi_t(f)$ satisfies the *Kushner-Stratonovich equation* (in Itô form)

$$\pi_T(f) = \nu(f) + \int_0^T \pi_t(\mathcal{L}f) dt + \int_0^T \{\pi_t(hf) - \pi_t(h)\pi_t(f)\} (dY_t - \pi_t(h) dt),$$

where \mathcal{L} is the infinitesimal generator of the Markov process X_t and ν is the law of X_0 . Note that the expression for $\pi_T(f)$ does not close, i.e., it is not expressed only in terms of $\pi_t(f)$ for $t \leq T$; rather, we have to know $\pi_t(\mathcal{L}f)$ as well, etc. To obtain a closed form expression, define formally

$$\pi_t(f) = \int_{\mathbb{S}} f(x) p_t(x) dx, \quad \nu(f) = \int f d\nu = \int_{\mathbb{S}} f(x) p_0(x) dx,$$

i.e., $p_t(x)$ is the density of the conditional law of X_t and $p_0(x)$ is the density of ν . Formally integrating by parts gives

$$p_T(x) = p_0(x) + \int_0^T \mathcal{L}^* p_t(x) dt + \int_0^T \{h(x) - \pi_t(h)\} p_t(x) (dY_t - \pi_t(h) dt),$$

which is a stochastic integro-differential equation (\mathcal{L}^* is the formal adjoint of \mathcal{L}). Note that everything here is formal, as we have not imposed any technical conditions, etc., but these manipulations can be given a precise meaning in many cases of practical interest. Further details and references can be found in chapter 1.

0.1.2. Approximations. The discussion above directly raises two points.

- (1) To implement the filtering equation one would have to propagate the entire conditional density $p_t(x)$, which is usually an infinite-dimensional object. Hence filtering often suffers from the *curse of dimensionality*.
- (2) The filter depends explicitly on the model chosen for the signal process X_t (through ν and \mathcal{L}) and for the observations (through h).

There are two exceptions to the curse of dimensionality. First, it can be the case that there is a finite-dimensional family of densities $\{p^\theta(x) : \theta \in \Theta \subset \mathbb{R}^q\}$ that is invariant for the Kushner-Stratonovich equation, i.e., $p_t(x) = p^{\theta_t}(x)$. The optimal filter can then be reduced to a finite-dimensional equation for θ_t . With the exception of a few cases of marginal practical interest (such as the Beneš filter [Ben81]), this is only the case if (X_t, Y_t) is a Gauss-Markov process (i.e., (X_t, Y_t) is the solution of a linear stochastic differential equation). The family p^θ then consists of the Gaussian distributions, and the corresponding filter is the celebrated *Kalman-Bucy filter* which is used extensively in many engineering applications. Much is known about linear filtering (an excellent reference is [KSH00]), and the special structure of this problem makes it amenable to simple and elegant methods of analysis that usually do not generalize to the much more poorly understood nonlinear setting. It is the latter case that we are chiefly concerned with in this thesis.

The second class of filters that escape the curse of dimensionality are those where \mathbb{S} is a finite set. In this case, the density $p_t(x)$ takes values in a finite-dimensional simplex and the Kushner-Stratonovich equation reduces to a finite-dimensional stochastic differential equation called the *Wonham equation*. Approximately half of this thesis is devoted to the finite state case. There are two good reasons for this: First, being the simplest type of nonlinear filtering problem, we can obtain much valuable intuition about more general nonlinear situations by studying the technically less demanding finite state case. Second, being one of the few nonlinear filters that admit direct implementation, the Wonham filter is of significant practical value in engineering applications. In contrast, a “truly” nonlinear filter with continuous state space \mathbb{S} is infinite dimensional [CM84] and hence can not be implemented without some form of (suboptimal) approximation.

Let us now turn to the second issue, i.e., to the dependence of the filter on the model parameters ν , \mathcal{L} , and h . Any mathematical model is necessarily an idealization of a true physical system; but even if we are willing to accept the validity of the model, as we will do, one rarely has access to the exact model parameters: these have to be obtained through some system identification procedure. Hence even if the curse of dimensionality can be overcome, one still deals in practice with suboptimal filters where the model parameters match only approximately those of the underlying system. The main theme of this thesis is the study of the sensitivity of various filters with respect to variations in the model parameters ν , \mathcal{L} , and h . This is a matter of significant practical importance, as almost all applications of nonlinear filtering necessarily operate with approximate model parameters. We also note at this point that the approximation of the model parameters is not unrelated to the type of approximation needed to circumvent the curse of dimensionality; this is mostly a topic for further research, but see section 0.2 and chapter 3.

0.1.3. Filter stability. A much studied problem in nonlinear filtering theory is the sensitivity of the filter to the initial measure ν . In particular, if the signal and

observations are sufficiently nice, the filter will “forget” the initial measure ν after some period of time: i.e., the sensitivity of the filter $\pi_t^\mu(f)$ (where $\pi_0^\mu(f) = \mu(f)$) to changes in the initial measure μ decays to zero as $t \rightarrow \infty$. In this case, the filter is said to be *stable*. Clearly this is a desirable state of affairs: we do not want a small error made at $t = 0$ to haunt us forever.

Filter stability gains extra significance if we interpret the filter as a statistical procedure. Nonlinear filtering can be thought of in the context of Bayesian inference, where μ plays the role of the prior distribution in our setup. Unlike in the probabilistic context, however, where $\mu = \nu$ is the *correct* initial measure and any other choice of μ leads to suboptimal results (i.e., $\pi_t^\mu(f)$ has a larger mean-square error than $\pi_t^\nu(f)$), the choice of prior is much more subjective in Bayesian inference. As a Bayesian estimator, the filter can only be used reliably if the information gained from the observations over a long period of time completely supersedes the prior, i.e., if the filter is stable. (In the remainder of this thesis we will always consider filtering in the probabilistic context, not in a statistical sense.)

The main questions are now:

- (1) Under what conditions is the optimal filter guaranteed to be stable?
- (2) If the filter is stable, how fast is the initial measure forgotten?

In the Kalman-Bucy case, there are some simple and powerful results in this direction, see, e.g., [OP96]. Roughly speaking, if the linear system (X_t, Y_t) is controllable and observable, then the filter is stable and the initial condition is forgotten at an explicitly computable exponential rate.

In the nonlinear setting, matters are not as clear cut. An excellent survey of the results to date can be found in [Chi06], to which we refer for a detailed discussion and an exhaustive list of references. Let us briefly highlight, however, some of the major results that are currently available (concentrating on the continuous time case). Many of the results are variations on a statement of the following form, first considered in [DZ91]. These results state that under certain conditions

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \|\pi_t^\mu - \pi_t^\nu\|_{\text{TV}} \leq \kappa < 0,$$

where $\|\cdot\|_{\text{TV}}$ denotes the total variation norm of a signed measure. It follows immediately that $\|\pi_t^\mu - \pi_t^\nu\|_{\text{TV}} \rightarrow 0$ at an exponential rate as $t \rightarrow \infty$. The analysis is significantly simplified when a particularly convenient metric, the Hilbert projective metric, is used in the proofs. This fact is exploited in [AZ97b, AZ97a] and in much of the subsequent literature. In the finite state case, stability is guaranteed in the following main cases: (i) if the signal process obeys a *mixing condition*, i.e., if all its transition rates are strictly positive [AZ97b] (see [BCL04] for slightly weaker conditions); (ii) in the low signal-to-noise limit [DZ91]; (iii) in the high signal-to-noise limit, with a nondegeneracy condition on the observation function [AZ97b]. In all these cases, estimates on the rate κ are available. The case that X_t is a strictly elliptic diffusion on a compact manifold also implies stability [AZ97a], but no useful estimate on κ is given. In a noncompact state space these methods are much less useful, as the Hilbert metric is not well suited to this situation; some progress can be made, however (see [Chi06] for discussion and references).

A much stronger form of filter stability is implied by the following statement:

$$\|\pi_t^\mu - \pi_t^\nu\|_{\text{TV}} \leq C(\mu, \nu) e^{-\kappa t},$$

for some rate $\kappa > 0$ and deterministic positive constant $C(\mu, \nu)$. Such a bound was first obtained in the finite state case in [BCL04], where the condition for stability is the mixing condition mentioned above. The technique used to prove this bound relies on the fact that the filter $\pi_t^\mu(f)$ with initial measure μ can be related to a *smoothing* problem under the initial measure ν , at least in the case that the two initial measures are equivalent $\mu \sim \nu$ (see also [COC99]). One can thus express $\|\pi_t^\mu - \pi_t^\nu\|_{\text{TV}}$ in terms of conditional expectations with respect to a single initial measure ν , and the exponential bound follows from analysis of the corresponding smoothing problem. In the diffusion case (state space \mathbb{R}^d), bounds of this form are obtained in [Sta04, Sta05, Sta06]. The method used in these papers is analytic in flavor, relying on transformations of PDEs and properties of certain Feynman-Kac integrals. This is the only method to date that can accommodate nonergodic signal processes X_t on a noncompact state space, and that provides explicit stability bounds for diffusions. On the other hand, the method only works for a very restrictive class of diffusions, observations and initial measures.

Finally, we mention some general results of a more qualitative type. In [OP96], it is argued that (roughly) if the signal process X_t is ergodic, then

$$\lim_{t \rightarrow \infty} \mathbf{E}(\pi_t^\mu(f) - \pi_t^\nu(f))^2 = 0,$$

for any bounded, continuous f . This is intuitively quite plausible, as ergodicity implies that the unconditional law of X_t becomes insensitive to the initial measure μ at long times; the statement is then, essentially, that this property is inherited by the filter. Unfortunately, there is a serious gap in the proof of this result, see [BCL04], so that the extent to which this statement holds remains unclear. A much weaker result still is proved in [COC99]: these authors show that

$$\int_0^\infty (\pi_t^\mu(h) - \pi_t^\nu(h))^2 dt < \infty,$$

where h is the observation function, provided that $\nu \ll \mu$. Hence in a very weak sense, at least the estimate of the observation function is always stable. The beauty of this result is its generality: other than the absolute continuity condition on the initial measures, the result holds for any Markov signal observed in white noise, without any further assumption on the structure of the signal or observations.

It can be concluded from the discussion above that the stability problem in nonlinear filtering theory is still far from being completely understood. Conditions that guarantee filter stability are only known in a restricted set of cases, and it is often difficult to obtain explicit estimates. As such, the development of new approaches to studying these problems is still of significant interest.

0.1.4. Conditional diffusions and filter stability. In this thesis, I propose to study filter stability using the method of conditional diffusions. The application of this method to the filter stability problem appears to be new, though we will see that it is closely related to the seemingly quite different methods of P. Chigansky et al. [BCL04] and of W. Stannat [Sta04, Sta05, Sta06].

To introduce the notion of a conditional diffusion, let us consider the path estimation problem for a signal process that is an Itô diffusion in \mathbb{R}^d

$$dX_t = b(X_t) dt + \sigma(X_t) dW_t, \quad X_0 = x,$$

where W_t is a k -dimensional Wiener process, x is some nonrandom vector in \mathbb{R}^d , and $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times k}$. We consider the usual white noise type observations Y_t . In the path estimation problem we are interested in calculating conditional expectations of the form $\mathbf{E}(f(X_{[0,T]})|\mathcal{F}_T^Y)$, where $f : C \rightarrow \mathbb{R}$ is a measurable functional on the space $C = C([0, T]; \mathbb{R}^d)$ of signal sample paths. Equivalently, we are interested in conditional probabilities of the form $\mathbf{P}(X_{[0,T]}^{-1}(A)|\mathcal{F}_T^Y)$, where A is an event in the Borel σ -algebra \mathcal{C} of C (under the uniform topology). As in the rest of this introductory chapter, we forgo any form of technical precision here and below.

Now assume that we have chosen a regular version of the conditional probability $\mathbf{P}(\cdot|\mathcal{F}_T^Y)$ (a technicality that we do not worry about at this point). Then for any event $A \in \mathcal{C}$, the quantity $\mathbf{P}(X_{[0,T]}^{-1}(A)|\mathcal{F}_T^Y)$ can be expressed as a measurable functional $\Pi_T(A, \cdot)$ of the sample paths of the observation, i.e.

$$\mathbf{P}(X_{[0,T]}^{-1}(A)|\mathcal{F}_T^Y) = \Pi_T(A, Y_{[0,T]}) \quad \text{a.s.} \quad \forall A \in \mathcal{C},$$

such that for every fixed observation sample path $y \in C([0, T])$ the map $\Pi_T(\cdot, y)$ is a probability measure on the space of signal sample paths (C, \mathcal{C}) . The path estimation problem can now be stated more precisely: for any fixed $y \in C([0, T])$, we would like to find an explicit way to calculate $\Pi_T(A, y)$ for any event $A \in \mathcal{C}$.

An elegant solution to this problem follows from the remarkable fact that for fixed $y \in C([0, T])$, the measure $\Pi_T(\cdot, y)$ can be characterized as the measure induced on C by a diffusion process $X_t^{T,y}$ which is a simple modification of the original diffusion X_t . In fact, $X_t^{T,y}$ may be obtained as the solution of

$$dX_t^{T,y} = b(X_t^{T,y}) dt + \sigma(X_t^{T,y})(d\tilde{W}_t + u^{T,y}(t, X_t^{T,y}) dt), \quad X_0^{T,y} = x,$$

where \tilde{W}_t is a Wiener process that is independent of Y_t and $u^{T,y}(t, x)$ is a time-dependent drift which depends on the observation sample path $y \in C([0, T])$ (this function can be obtained, e.g., by solving backwards in time a PDE that is driven by y ; we postpone the details for later). The key point is that for every $y \in C([0, T])$, *the law of the conditional diffusion $X_t^{T,y}$ on $[0, T]$ is precisely the conditional law $\Pi_T(\cdot, y)$ of the signal process X_t given the observation path y .* Explicitly,

$$\mathbf{E}(f(X_{[0,T]})|Y_{[0,T]} = y) = \int_C f(\xi) \Pi_T(d\xi, y) = \tilde{\mathbf{E}}(f(X_{[0,T]}^{T,y})),$$

for any measurable functional $f : C \rightarrow \mathbb{R}$ and any $y \in C([0, T])$, where $\tilde{\mathbf{E}}$ denotes the expectation with respect to the Wiener process \tilde{W}_t .

How are these ideas related to filter stability? Let us outline a potential approach. In the discussion above $X_0 = x$ was deterministic, so that the initial measure is the Dirac measure $\nu = \delta_x$. In particular, we can now write

$$\pi_T(g) = \pi_T^{\delta_x}(g) = \mathbf{E}(g(X_T)|\mathcal{F}_T^Y) = \tilde{\mathbf{E}}(g(X_T^{T,Y_{[0,T]}})).$$

As it turns out, however, $u^{T,y}(t, x)$ is a functional of the observation sample paths only, i.e., it does not depend on the initial measure ν for fixed y . Hence if we choose $X_0^{T,y} = x'$ instead of $X_0^{T,y} = x$, it is evident that $X_T^{T,y}$ must be distributed according to the filter with misspecified initial measure $\delta_{x'}$:

$$\pi_T^{\delta_{x'}}(g) = \tilde{\mathbf{E}}(g(X_T^{T,Y_{[0,T]}}(x'))),$$

where $X_T^{T,y}(x')$ denotes the solution of the equation for $X_t^{T,y}$ with initial condition $X_0^{T,y} = x'$. We could now investigate a form of filter stability by studying the

properties of the conditional diffusion $X_t^{T,y}$. For example, let g be a Lipschitz continuous function. If we could show that the stochastic flow $X_t^{T,y}(\cdot)$ is contracting at an exponential rate independent of y and T , i.e., that

$$\|X_t^{T,y}(x) - X_t^{T,y}(x')\| \leq C \|x - x'\| e^{-\kappa t},$$

then we would immediately obtain the stability result

$$|\pi_T^{\delta_x}(g) - \pi_T^{\delta_{x'}}(g)| \leq C \|g\|_{\text{lip}} \|x - x'\| e^{-\kappa T}.$$

The significant simplification that is gained by using conditional diffusions is that all the conditioning has been hidden in the drift $u^{T,y}(t, x)$, and we are left with proving properties of an ordinary diffusion process under the ordinary Wiener measure $\tilde{\mathbf{E}}$, for which a large number of methods are already available. The problem does not trivialize, of course, as the properties of the conditional diffusion depend on the properties of the drift function $u^{T,y}(t, x)$. Nonetheless, this problem can be much more tractable than the filter stability problem in its original form.

The study of conditional Markov processes was initiated by R. L. Stratonovich [Str60, Str68], at least on a formal level, but the theory of conditional diffusions in the spirit of the above discussion appears to have its origins in the paper [BM82] of J.-M. Bismut and D. Michel. These authors show that the Radon-Nikodym derivative of $\Pi_T(\cdot, y)$ with respect to the measure of the unconditional diffusion $X_{[0,T]}$ is of Girsanov type. It is evident that $X_t^{T,y}$ should then have the above form. To find the drift term explicitly, they apply the Clark-Haussmann-Ocone formula (see, e.g., [Nua95]) and the theory of stochastic flows.

Recently, the theory of conditional diffusions was developed from a rather different perspective in a remarkable paper by S. K. Mitter and N. J. Newton [MN03]. These authors show that nonlinear estimation can be expressed in a variational form, where the abstract Bayes formula (which is at the heart of filtering theory) obtains a natural information-theoretic interpretation. In the case of path estimation, this variational problem can be expressed as a stochastic control problem for which the function $u^{T,y}(t, x)$ is precisely the optimal control strategy. Beside the aesthetic appeal of such a formulation, we will find that the stochastic control perspective gives us an additional technical tool that can be used in the analysis of the filter stability problem. In different contexts, the use of stochastic control as an analytic tool is not unusual [FM83, She91, DE97, Bor00, FS06].

In the sequel, we will apply the theory of conditional diffusions to study filter stability for finite-state signals (chapter 2) and for diffusions (chapter 4). In the former case, we first need to develop conditional diffusions for finite-state signals; only the diffusion theory is currently available in the literature. Filter stability can then be studied essentially as in the example above. In the diffusion case, it can be more convenient to use a time-reversed version of the theory described above. In particular, one can introduce a process $\tilde{X}_t^{T,y}$ on $[0, T]$, $\tilde{X}_0^{T,y} = x$, such that

$$\mathbf{E}(f(X_{[0,T]}) | Y_{[0,T]} = y, X_T = x) = \tilde{\mathbf{E}}(f(\tilde{X}_{T-[0,T]}^{T,y})),$$

for any measurable functional f . As we will see, this form of the theory is closely related to the methods of P. Chigansky et al. and of W. Stannat.

0.1.5. Filter robustness. Let us now turn to the issue of robustness, i.e., the sensitivity of the filter simultaneously to the model parameters ν , \mathcal{L} , and h .

Continuity with respect to the model parameters of nonlinear filtering estimates on a fixed *finite* time interval is well established, e.g., [BKK95, BKK99, BJP02, GY06]; generally speaking, it is known that the error incurred in a finite time interval due to the choice of incorrect model parameters can be made arbitrarily small if the model parameters are chosen sufficiently close to those of the true model. As the corresponding error bounds grow rapidly with the length of the time interval, however, such estimates are of little use if we are interested in robustness of the filter over a long period of time. One would like to show that the approximation errors do not accumulate, so that the error remains bounded uniformly over an *infinite* time interval. This requires a more subtle analysis.

The model robustness of nonlinear filters on the infinite time horizon has been investigated in discrete time in [BK98, LO03, LO04, OR05]. The main idea of the approach is simple and intuitive: if the filter is exponentially stable, i.e., if the initial measure is forgotten at a geometric rate, then the filter is also robust to the remaining model parameters. Let us give a rough indication as to why this is the case. Consider a discrete filter of the form $p_{n+1} = \Pi(y_n, p_n)$, where p_n is the conditional density at time step n , y_n is the observation obtained in time step n , and $\Pi(y, p)$ is the one time step map of the filter. Similarly, $\bar{p}_{n+1} = \bar{\Pi}(y_n, \bar{p}_n)$ is the filter with misspecified model parameters. Furthermore, let $d(p_1, p_2)$ be a distance metric under which the filter is contracting: $d(\Pi(y, p_1), \Pi(y, p_2)) \leq \kappa d(p_1, p_2)$ for some $\kappa < 1$, i.e., the filter is exponentially stable. Using the triangle inequality,

$$\begin{aligned} d(p_{n+1}, \bar{p}_{n+1}) &\leq d(p_{n+1}, \Pi(y_n, \bar{p}_n)) + d(\Pi(y_n, \bar{p}_n), \bar{p}_{n+1}) \\ &\leq \kappa d(p_n, \bar{p}_n) + d(\Pi(y_n, \bar{p}_n), \bar{\Pi}(y_n, \bar{p}_n)). \end{aligned}$$

The second term on the right-hand side measures the local error due to the misspecification of the model parameters, i.e., the distance after one time step between the correct and misspecified filters that are started at the same point. Let us call this local error $\delta(y_n, \bar{p}_n)$. By iterating the bound above, we obtain

$$d(p_{n+1}, \bar{p}_{n+1}) \leq \sum_{\ell=0}^n \kappa^{n-\ell} \delta(y_\ell, \bar{p}_\ell).$$

Now suppose that we can show that the local error is bounded in some sense, e.g., suppose that $\sup_{k \geq 0} \mathbf{E} \delta(y_k, \bar{p}_k) \leq C$ for some constant $C < \infty$. Then

$$\sup_{n \geq 0} \mathbf{E} d(p_n, \bar{p}_n) \leq C \sum_{\ell=0}^{\infty} \kappa^\ell = \frac{C}{1 - \kappa}.$$

In particular, if $C \rightarrow 0$ as we let the misspecified model parameters get closer and closer to the true model parameters, then the error between the true and approximate filter vanishes uniformly over the infinite time interval. In this case, we will call the filter *robust* to misspecification of the model parameters. The existing results proceed roughly along these lines, and are thus restricted to discrete time filtering; with some additional work, the same technique can be made to work in the case of point process observations in continuous time [BK98] (this is essentially like the discrete time case, but the times are now random).

In chapter 3, we will demonstrate how to extend these ideas to the context of continuous time filtering with white noise observations. We restrict ourselves to the case of a finite-state signal process that obeys the mixing condition, where strong exponential stability results can be obtained. The analysis in the continuous time

case is significantly more subtle than the discrete time case outlined above. The analysis begins by considering the stochastic flow that is generated by the filter on the simplex. Rather than requiring the filter to forget its initial measure at an exponential rate, we now need to require that the *derivative* of the filter flow with respect to the initial measure admits an exponential bound. In the case of a mixing signal process, we show that this is indeed the case. An additional complication in proving robustness is that the corresponding error bounds can only be expressed in terms of *anticipative* stochastic integrals. Hence the usual Itô calculus breaks down, and we resort to using tools from the Malliavin calculus [Nua95] and the associated anticipative calculus for Skorokhod integrals [NP88, Nua95].

We finish this section by mentioning a related result. In [BK99], the object of interest is the estimation error of the filter over a long time interval; in particular, these authors consider the pathwise time average of the squared difference between the filtered estimate of the signal and the actual signal. It is shown that if the filter is approximated in some way—e.g., by misspecifying the model parameters, though much more general approximations are covered—then under suitable conditions, the pathwise time average estimation error of the approximate filter converges to that of the exact filter as the approximate filter converges to the exact filter. This result is somewhat different in spirit than our previous discussion, however. In particular, it does not show that the approximate filter is close to the exact filter at any particular time; rather, it is shown that as an estimator, the time average performance of the approximate filter is close to that of the exact filter.

0.1.6. Quantum filtering and filter stability. In the models which we considered above, the signal process was modelled as a classical Markov process. A typical example is a stochastic differential equation, which could model a physical system that is somehow driven by an auxiliary, independent white noise input. In applications, such noisy driving is often introduced to model the effects of unavoidable wide bandwidth disturbances, caused by the coupling of the system to an unknown environment or by other factors that introduce uncertainty in the model. Some applications of nonlinear filtering are listed at the end of this section.

The increasing miniaturization of technology, however, suggests a different set of models of interest: those consisting of small numbers of atoms or photons. Such models do not only suffer from unknown auxiliary disturbances, but also from the inherent quantum mechanical uncertainty that is present at these scales. There is yet a long way to go before such systems can be said to be useful as a practical technology. Even at the time of writing, however, concrete technological applications to precision sensing and detection are emerging and are being explored in a laboratory setting [AAS⁺02, GSMD03, GSM05]. It is perhaps not surprising that filtering plays an important role in these applications.

In quantum mechanics, physical quantities are described by random variables (called observables) that may not commute with each other. In every realization, one must choose a commuting set of observables (as determined by the method of observation) and at this level, the model reduces to a classical probability model. This idea is the starting point for the theory of *noncommutative* or *quantum probability*. Within this framework, one can develop Markov models; these are quite ubiquitous in the physics literature, as is easily seen from the large number of physical systems that are routinely modeled by Lindblad-type master equations (the direct noncommutative analog of the Kolmogorov forward equation). These Markov

models can be obtained as solutions to quantum stochastic differential equations, which model, e.g., an atomic system in interaction with the electromagnetic field. The field serves a dual purpose: it acts as a source of noise (the inherent quantum vacuum fluctuations), but it also carries off energy and hence information from the atoms. If we place a photodetector in the field, we can subsequently estimate (in least-mean-square) the atomic observables based on the observations history; this is the domain of *quantum filtering theory*.

Quantum filtering has its origins in [Bel80, Bel88, Bel92], and has found its way into the physics literature from a different perspective [Dav76, Car93]. The development of the theory strongly resembles that of classical filtering theory: in particular, the derivation of the filtering equations using martingale methods [Bel92, BGM04, BVJ06b] is almost identical to the classical martingale method [Kal80, LS01, Kri05]. In the first half of chapter 5 we will develop the theory using a reference probability method, as outlined in [BV06]. This approach parallels the approach we take to classical filtering theory in chapter 1.

There are some significant differences with the classical theory that make quantum filtering theory interesting in its own right. First, unlike *filtering*, there is no natural noncommutative counterpart to *smoothing* (let alone path estimation); the reason for this will become evident when we develop the theory. Second, nothing is known about filter stability for quantum nonlinear filtering. Essentially all of the classical methods for proving filter stability fail in the quantum setting: methods based on some form of smoothing suffer from the lack of a quantum smoother, whereas the lack of a satisfactory noncommutative analog of the Hilbert projective metric limits the applicability of much of the remaining literature.

In the second half of chapter 5, we obtain a first filter stability result for quantum filters. This result is in the same spirit as the classical result of J.M.C. Clark, D.L. Ocone and C. Coumarbatch cited previously [COC99]: we will show that the filtered estimate of a particular system observable, called the *measurement observable*, is always stable regardless of the details of the underlying model. The only requirement is a certain absolute continuity condition on the correct and misspecified initial states of the filter. The key insight that allows us to obtain this result is the realization that some of the “smoothing” theory can be recovered by suitably randomizing the initial state of the system. The stability result then follows by using change of measure techniques and some elementary analysis.

0.1.7. Applications of nonlinear filtering. To complete this introductory section, we list below a set of selected applications of nonlinear filtering. The list includes applications of nonlinear filters with discrete and continuous time, state and observation spaces, and is far from exhaustive: it serves mainly to indicate the breadth of potential applications. Though this thesis is concerned with the case of continuous time and white noise type observations, it is not unusual for lessons learned in one form of nonlinear filtering to be illuminating also in other contexts. We have excluded the Kalman-Bucy filter and its relatives (including the extended Kalman filter), however, for which a seemingly inexhaustible number of applications can be found throughout the engineering and scientific literature.

Navigation and target tracking: Guidance is one of the classic applications of filtering theory [BJ87]. Nonlinear filtering can be used to navigate a vehicle through an unknown terrain using noisy sensor data or to track a moving object using noisy data; see, e.g., [GGB⁺02]. One particular

recent application is the use of filtering in GPS navigation in the case when access to the satellites varies with time [CDMS97, ASK05].

Changepoint detection: When a device or an industrial process breaks down, one would like to detect this as quickly as possible so that the fault can be repaired. When only noisy data are available, however, this can be a challenging task. Nonlinear filtering provides a good method to detect a fault with noisy observations [Shi73, HC99, VC03, ASK04].

Stochastic control: If we wish to use feedback to control a stochastic system, but only noisy observations are available, the controller design is usually split into a filtering step and a control step. In optimal control theory, this in fact turns out to be the *optimal* controller design: the filter is an *information state* for such a control problem [Ben92, EAM95]. Filtering is also the basis for certain adaptive control strategies which may be easier to implement than the optimal control [CZ95].

Finance: In models where the volatility is taken to be a random process, filtering can be used to estimate the volatility from stock market data [FR01, CLR06]. Filtering is also used for portfolio optimization in a market with randomly varying instantaneous returns [SH04].

Audio and image enhancement: Nonlinear filters can be used for noise removal and signal enhancement of audio signals and images; see, e.g., [FGDW02, Eph92, Bes86, EAM95].

Biology: Nonlinear filtering is applied in patch-clamp experiments to estimate neuronal spike trains from noisy patch-clamp data [CMX⁺90, CKM91, FR92]. Other applications areas in biology include electrocardiography (ECG) [CCB90] and DNA sequencing [LB98].

Quantum optics: Quantum filtering theory has found a large number of applications in quantum optics, both as an optimal estimator [AAS⁺02, GSDM03, GSM05] and for the modelling and simulation of optical photocurrents [Car93, GZ04]. As in the classical theory, quantum filtering plays a central role in quantum feedback control [BV06, BVJ06a] as it provides a suitable information state for quantum control problems.

Various: Other applications include speech recognition [Rab89, EM02], communication theory [EM02, BP03], and data assimilation for weather prediction and ocean current modelling [AA99, AHSV06].

0.2. A suggestive numerical experiment

The theme of this thesis was originally inspired by an attempt to find accurate approximate filters for an interesting quantum optical system [VM05]. Numerical experiments showed a much better performance of the approximate filter than was originally expected; particularly surprising was the observation that the performance does not degrade at all with time, in contrast to commonly used approximation methods where nonlinear filters are approximated by Kalman filters through local linearization. The existence of uniform approximations is crucial for real-time applications, as it allows the curse of dimensionality to be circumvented using an approximate filter of fixed complexity (requiring limited computational resources) regardless of the length of the time interval over which the filter is used.

In this section we will briefly describe some of the numerical results obtained in this system. This serves both as an illustration of an application of filtering in the




FIGURE 0.1. Schematic of an optical phase bistability experiment. A single two-level atom is placed in a high-Q optical cavity and is strongly driven using an external laser. The radiation emitted from the cavity in the forward direction is detected using a homodyne detection setup.

context of quantum optics, and as a motivation for the development of new methods to study filter stability and robustness. The results obtained in the sequel do not apply directly to this model; nonetheless, it is the stability property of this filter that allows one to obtain approximations with uniform in time performance. Further development of the technical machinery for studying filter stability and robustness could allow an important application, the design of approximations with guaranteed uniform error bounds for (a class of) infinite-dimensional nonlinear filters, to be realized. We are not close to achieving this goal, but the analysis of chapter 3 provides a modest step in this direction (in a finite-dimensional context).

The model is illustrated in figure 0.1. A strongly coupled two-level atom in a resonant, single-mode optical cavity is strongly driven by a resonant driving laser. One of the cavity mirrors is leaky, and the radiation emitted from that mirror is detected using a homodyne detection setup. Spontaneous emission is also taken into account. In this operating regime, the atom-cavity system exhibits bistable behavior: when the atom spontaneously emits a photon of a certain frequency (this photon goes off in a random direction and is not detected), the phase of the intracavity field switches. Our goal is to detect these switches using the noisy photocurrent signal obtained from the homodyne detector.

The use of optical bistability in technological applications was suggested a long time ago [AS82]. In particular, such systems could be used to build optical transistors, optical memory elements or similar devices for applications in high-speed optical signal processing. The actual engineering of devices of this type has only recently become feasible [YFSJ03] through advances in fabrication of optical technology. A different application of bistability has been suggested for low-noise detection [SVP⁺04]. Even experiments involving a single atom are well within the reach of current technology [MYK99], and an experiment implementing the setup of figure 0.1 is now under way [Mab06]. Progress in this direction, using either real or artificial atoms (quantum dots), may enable the development of ultrafast (picosecond) and extremely low energy (attojoule) optical switches, a regime that is effectively out of reach for more conventional technology [AM06].

The physical modelling of the setup of figure 0.1 was done in detail in [VM06], and we will not repeat it here. In fact, as we have not yet introduced quantum filtering theory, we will side-step quantum filtering in the context of this example. This is impossible for most quantum models, but in this particular case it turns out that the filter can be expressed, by a suitable change of variables, as a stochastic

PDE that coincides with a *classical* nonlinear filter.¹ The state space for this filter is $\mathbb{S} = \mathbb{R} \times \{-1, +1\} \ni (x, j)$, and the corresponding PDE is given by

$$dP_t(x, j) = \mathcal{L}^* P_t(x, j) dt + \sqrt{2\kappa\eta} (x - \langle x \rangle_t) P_t(x, j) (dY_t - \sqrt{2\kappa\eta} \langle x \rangle_t dt),$$

where the adjoint generator \mathcal{L}^* is defined as

$$\mathcal{L}^* P(x, j) = \frac{\partial}{\partial x} [(jg + \kappa x)P(x, j)] + \frac{\gamma}{2} (P(x, -j) - P(x, j)),$$

and we have written

$$\langle x \rangle_t = \sum_{j=\pm 1} \int_{-\infty}^{\infty} x P_t(x, j) dx.$$

Here $\kappa > 0$ is the cavity decay rate, $g > 0$ is the coupling strength between the atom and the cavity mode, $\gamma > 0$ is the spontaneous emission rate, $\eta \in [0, 1]$ is the detection efficiency of the homodyne detector, and Y_t is the (semimartingale) photocurrent obtained from the homodyne detector.

Now note that \mathcal{L}^* is the adjoint generator of the following Markov process:

$$\begin{cases} J(t) \text{ is a Markov jump process that switches between } \pm 1 \text{ at rate } \gamma/2, \\ \frac{d}{dt} X(t) = -\kappa X(t) - g J(t). \end{cases}$$

These equations have a natural physical interpretation. $J(t)$ represents the state of the atom at time t —recall that we are dealing with a two-level atom—which switches repeatedly due to the combined effect of the strong drive and the spontaneous emission (the drive has been eliminated in the description, so that $J(t)$ represents a “dressed state” of the atom). $X(t)$ represents a quadrature of the intracavity field; it is damped at the cavity decay rate κ , and is coupled to the atomic state with strength g . Every time the atomic state switches, $X(t)$ decays to a fixed point with opposite sign (phase bistability). It now appears that $P_t(x, j)$ is the density of the conditional expectation $\pi_t(f) = \mathbf{E}(f(X(t), J(t)) | \mathcal{F}_t^Y)$ if we define

$$dY_t = \sqrt{2\kappa\eta} X(t) dt + dB_t,$$

as then the equation for $P_t(x, j)$ above is precisely the corresponding Kushner-Stratonovich equation. In particular, $\langle x \rangle_t = \pi_t(x)$ is then the conditional expectation of the phase quadrature of the intracavity field. This is in fact the correct interpretation [VM05], but one should keep in mind that the classical system above was reverse engineered from a fully quantum mechanical model.

The problem that we are facing is that the filter for our system suffers from the curse of dimensionality. After all, $P_t(x, j)$ is an infinite-dimensional object, and the filtering equation does not leave invariant some finite-dimensional family of densities. If we wish to build a device that relies on real-time processing of the photocurrent coming from the cavity, the filter in the above form would be practically useless. Hence we need to find a suitable approximation. To this end, [VM05] uses the approximation method introduced by D. Brigo, B. Hanzon and F. Le Gland [BHL99], adapted to the current context. The idea is very simple: if the filter does not leave invariant some family $\{P^\theta(x) : \theta \in \Theta\}$, we can always fix

¹ This is done by using a special (P -function) representation of the conditional density operator [MW98, VM05]. This representation will not be used in the sequel. There are also other special cases where quantum filters can be expressed as classical filtering equations, most notably the linear case [Bel80, DJ99] and certain types of switching filters [VM06, Van07].

such a family and then constrain the filter always to remain in this family. We will choose the following finite-dimensional family of densities:

$$P^\theta(x, j) = \nu^j \delta(x - \mu^j), \quad \theta \in \Theta = \{(\mu^\pm, \nu^\pm) : \mu^\pm \in \mathbb{R}, \nu^\pm > 0, \nu^+ + \nu^- = 1\}.$$

We now proceed as follows. First, as our densities are singular, we need to smooth the problem out a little. To this end, we transfer our attention to the function

$$Q_t(x, j) = \frac{1}{\sqrt{4\pi}} \int P_t(x', j) e^{-(x-x')^2/4} dx'.$$

This is another normalized probability density, but is smoother than $P_t(x, j)$.² This new density once again satisfies a PDE, which is easily found. We would like to constrain this PDE to leave the following (smooth) manifold invariant:

$$S = \left\{ Q^\theta(x, j) = \frac{\nu^j}{\sqrt{4\pi}} e^{-(x-\mu^j)^2/4} : \theta \in \Theta \right\}.$$

To this end, let us write the PDE for $Q_t(x, j)$ suggestively as

$$dQ_t = A[Q_t] dt + B[Q_t] \circ dY_t,$$

where \circ denotes the Stratonovich differential. We use the Stratonovich form here as it allows us to interpret the equation geometrically (see, e.g., [Bis81]). If Q_t leaves S invariant, then the “vector fields” $A[Q]$ and $B[Q]$ should lie in the tangent space $T_Q S$ for every point $Q \in S$. As this is not the case, we will constrain the equation to leave S invariant by projecting A and B onto $T_Q S$. To project, however, we need a suitable inner product, and $T_Q S \subset L^1(\mathbb{R} \times \{-1, +1\})$ does not have a natural inner product. Hence we perform a final transformation:

$$d\sqrt{Q_t} = \frac{A[Q_t]}{2\sqrt{Q_t}} dt + \frac{B[Q_t]}{2\sqrt{Q_t}} \circ dY_t.$$

The tangent space to $S^{1/2} = \{\sqrt{Q^\theta} : \theta \in \Theta\}$ at $\theta \in \Theta$ is given by

$$T_\theta S^{1/2} = \text{span} \left\{ \frac{\partial \sqrt{Q^\theta}}{\partial \mu^+}, \frac{\partial \sqrt{Q^\theta}}{\partial \mu^-}, \frac{\partial \sqrt{Q^\theta}}{\partial \nu^+} \right\} \subset L^2(\mathbb{R} \times \{-1, +1\}).$$

(Recall that $\nu^+ + \nu^- = 1$, so we do not need to include ν^-). The usual L^2 inner product gives the (Fisher) metric $g(\theta)$ with matrix elements

$$\left\langle \frac{\partial \sqrt{Q^\theta}}{\partial \theta^k}, \frac{\partial \sqrt{Q^\theta}}{\partial \theta^l} \right\rangle_{L^2} = g_{kl}(\theta), \quad (\theta^1, \theta^2, \theta^3) = (\mu^+, \mu^-, \nu^+).$$

Using the corresponding orthogonal projection

$$\Pi_\theta X = \sum_{k,l=1}^3 (g^{-1}(\theta))_{kl} \left\langle \frac{\partial \sqrt{Q^\theta}}{\partial \theta^l}, X \right\rangle_{L^2} \frac{\partial \sqrt{Q^\theta}}{\partial \theta^k}, \quad X \in L^2(\mathbb{R} \times \{-1, +1\}),$$

we define the approximate filter

$$d\sqrt{Q^{\theta_t}} = \Pi_{\theta_t} \frac{A[Q^{\theta_t}]}{2\sqrt{Q^{\theta_t}}} dt + \Pi_{\theta_t} \frac{B[Q^{\theta_t}]}{2\sqrt{Q^{\theta_t}}} \circ dY_t.$$

² This transformation is rather odd from the classical perspective, but it is quite natural in the quantum context. In fact, it is yet another representation, called the Q -function representation, of the conditional density operator, and is equivalent to the P -function representation. See, e.g., [MW95, WM94] for details on the definitions and properties of these representations.




FIGURE 0.2. Simulation of the filters for the phase bistability model. The dashed lines denote $I_{J(t)=+1}$ (top plot) and $X(t)$ (bottom plot). The solid lines are the filtered estimates, and the dotted lines are obtained from the approximate filter. Here $\eta = 1$, $g = 120$, $\kappa = 40$, and $\gamma = 20$.

This expression leaves $S^{1/2}$ invariant by construction. Explicitly working out the various steps, we arrive at the approximate filter (in Itô form)

$$\begin{aligned} d\nu_t^+ &= -\gamma(\nu_t^+ - \tfrac{1}{2}) dt + \sqrt{2\kappa\eta} \nu_t^+ (1 - \nu_t^+) (\mu_t^+ - \mu_t^-) \\ &\quad \times \{dY_t - \sqrt{2\kappa\eta} (\mu_t^+ \nu_t^+ + \mu_t^- (1 - \nu_t^+)) dt\}, \\ \frac{d}{dt} \mu_t^j &= -\kappa \mu_t^j - g j + \frac{\gamma}{2} \left(\frac{1 - \nu_t^+}{\nu_t^+} \right)^j (\mu_t^{-j} - \mu_t^j). \end{aligned}$$

We refer to [BHL99] for further details on the *projection filter* approximation method, and to [VM05] for details on the calculations above.

Despite that we have followed a relatively clean procedure to approximate our full infinite-dimensional filter, there is no particular reason to expect at this point that the approximation will be a good one. We have chosen a rather arbitrary finite-dimensional family to constrain to, based only on intuition, computational convenience, and some inspiration from numerical experiments on the exact filter. Moreover, there is not a unique way to constrain an equation to remain in some low-dimensional space. Most of all, perhaps, there is no particular reason to expect that constraining the exact filter (essentially a geometric procedure) should give a reasonable approximation for the conditional expectations (a probabilistic notion).

In contrast to these concerns, numerical simulations demonstrate remarkable performance for our approximate filter. A glance at a typical simulation using even a not-too-flattering set of parameters, see figure 0.2, should convince the reader that we must be doing something right. Numerical experiments in a higher signal-to-noise regime, see figure 0.3, show essentially indistinguishable performance between the plotted estimates obtained from the optimal and approximate filters. Not only is the approximation error small, but it is also remarkably stable over time. Small excursions between the approximate and exact filters are visible in figure 0.2, but these errors tend to correct themselves rapidly rather than accumulating over time. Evidently the exact filter is extremely robust to approximation. The analysis of chapter 3, albeit in a finite-dimensional context, shows how this suppression of the approximation error can be related to the exponential stability of the optimal filter.

The example which we have discussed in this section is not the only example in which behavior of this kind was observed. In [VC03] the projection filter approach was used to obtain a finite-dimensional filter for the purpose of changepoint detection, an application of significant practical importance. Also here, almost indistinguishable performance was observed between the optimal and approximate




FIGURE 0.3. Simulation with $\eta = 1$, $g = 600$, $\kappa = 200$, and $\gamma = 20$.

filters. A better understanding of this phenomenon could both help to design better filter approximations, as well as to obtain error bounds for existing approximations such as the filter of this section or the changepoint detection filter of [VC03]. The example of this section and that of [VC03] are similar in other ways as well. In both cases, the underlying system is of jump-linear type (see also [HC99]), and the approximate filters can be interpreted as *adaptive* versions of the Wonham filter. Considering its relative simplicity and practical importance, this class of systems in particular merits further investigation.

0.3. Outline of this thesis, main results, and outlook

This thesis consists of five chapters and two appendices.

Chapter 1 contains an introduction to nonlinear filtering theory using the reference probability method. The goal of this chapter is to provide most of the necessary background material for the remaining chapters. As such I have provided a fairly detailed development, together with a large number of references to the literature where further details can be found.

Chapter 2 develops in detail the theory of conditional signals in the case of a finite state space. This material is new, and complements the existing diffusion theory. The theory is developed from two different perspectives: a direct approach is based on a Girsanov-type theorem for finite-state Markov processes, while a second approach uses stochastic control methods. The duality between estimation and stochastic control is fundamental in this context, and is developed in detail. The last part of the chapter applies the conditional signal theory to obtain a new nonasymptotic bound on the stability of the Wonham filter. This bound has the same decay rate as previous bounds, but, unlike previous bounds, the prefactor is constant and does not diverge for initial measures on the boundary of the simplex.

Chapter 3 is devoted to studying the robustness of the Wonham filter, on the infinite time interval, to changes in the underlying model parameters: the initial measure, the transition intensities of the signal process, and the observation function. Beside the filter stability theory developed in chapter 2, this chapter uses technical tools from the Malliavin calculus and the associated stochastic calculus of Skorokhod integrals. Additional background on this topic is given in appendix A. This chapter is based on the paper [CV06] (joint work with P. Chigansky).

Chapter 4 is devoted to the study of filter stability in a class of diffusion models (of gradient form), using conditional signal and stochastic control methods. Some of the results resemble (and are inspired by) those obtained by W. Stannat in [Sta04, Sta05, Sta06], but the methods used here are new and provide significant insight into the structure of the problem in a completely probabilistic setting. I also show that much stronger results can be obtained, using similar techniques,

if we are willing to impose more stringent requirements on the signal process. Of independent interest is a new probabilistic proof of a Brascamp-Lieb type inequality using stochastic differential equations and the Malliavin calculus.

Finally, chapter 5 is devoted to quantum filtering theory. The first part of the chapter develops the theory of quantum filtering using reference probability methods, which parallel those used in chapter 1. This method is due to joint work with L. Bouten [BV06], but many of the technical details appear here for the first time. The necessary background on quantum probability theory is provided in appendix B. The second part of the chapter develops a simple filter stability result for quantum filters, which is the first such result available in the literature. Key to the proof of this result is the development of a suitable notion of absolute continuity for quantum states, which allows changes of measure to be implemented in the quantum case through randomization of the initial state.

The methods and results developed in this thesis suggest several interesting directions for further research. Some of these are listed below.

- An important topic for further research is the improvement of the bounds obtained in chapter 3. For the reasons explained there, the estimates used in this thesis do not give quantitative bounds on the approximation error of the filter. Improved estimates could open the door to the development of near-optimal filter approximations, the feasibility of which is demonstrated numerically in section 0.2, with guaranteed error bounds.
- A particularly interesting class of models for further investigation are jump-linear systems, which consist of linear observations and a linear signal whose coefficients are switched according to a finite-state Markov process. Such models are widely used in applications, and appear to admit excellent filter approximations, but neither filter stability nor robustness have been studied for this class of models. Beside the practical importance of such results, these models are attractive as a testbed for the study of filter stability and robustness, as they are intermediate between the simpler linear and finite state space cases and the difficult general nonlinear case. Initial progress in this direction can be made by combining and extending the results of chapters 2 and 4 in this thesis.
- Though we have described several ways in which the conditional signal theory can be used to study filter stability, the deeper connections between filter stability and the stochastic control method merit further investigation. The fact that filter stability is intimately related with controllability and observability in the linear case [OP96], makes it tempting to look for such connections in the nonlinear case. Another interesting question is whether the stochastic control method can be extended to the infinite time horizon through, e.g., an ergodic cost criterion, so that the asymptotic behavior of the conditional signal can be studied.³ Finally, the well-known connection between stochastic control and large deviations [DE97] could be interesting in the filtering context (see also [JB88]).

³ A connection between “filter stability” and asymptotic properties of the conditional signal can be found, in a rather different context, in the old paper [Bež71]. The problem studied there is quite different from what we are interested in, and this paper predates much of the modern filtering theory; but some of the objects that appear in the paper look surprisingly familiar.

- The conditional signal theory developed in chapter 2 appears to be quite universal. The generalization of these techniques to a class of general Markov signals could be of fundamental interest.
- Though we present a first result on quantum filter stability in this thesis, this topic is still essentially uncharted territory. We will demonstrate in chapter 5 how some of the problems can be overcome through randomization; it could be fruitful to build on this method, or to apply other methods such as Lyapunov exponent techniques.

0.4. Other work

This thesis grew out of the desire to present as a coherent story a set of ideas which have collected on my desk, in the form of illegible notes on backs of envelopes and other varieties of scratch paper, over the course of about a year. Many of the details were worked out in the months leading up to the defense of this thesis. On the other hand, I have omitted the majority of the work which I have done throughout my graduate career, mostly on the topic of quantum filtering and feedback control, as this work can now be found in various publications and preprints. An exception to this omission policy is the paper [CV06], of which chapter 3 is an edited version. Below I provide some brief comments on and references to the remaining work which is not presented in this thesis.

Early work concentrated on the design of feedback controls for deterministic state preparation in atomic systems using stochastic stability methods, a topic motivated by ongoing laboratory experiments. The basic ideas were presented in [VSM05a], inspired by earlier numerical work [SVM04]. A full development of the theory appears in [MV06], while a unified overview of the physical modelling aspects of the theory appear in [VSM05b].

A new type of stochastic stability theory, the notion of almost global stability for stochastic differential equations and the corresponding Lyapunov-type methods, was developed in the paper [Van06].

Later work has concentrated on the foundations of quantum filtering and feedback control. Both the reference probability method, and the fundamental notion of controlled quantum flows, were developed in [BV06]. An overview of quantum filtering theory appears in [BVJ06b], while [BVJ06a] gives a unified treatment of quantum filtering and feedback control in the context of a discrete model, similar to the Cox-Ross-Rubinstein models developed in mathematical finance.

The topic of quantum filter approximation using the projection filter method was studied in [VM05] in the context of an optical bistability model. These results were discussed briefly in section 0.2. Further investigations include the use of quantum filtering for the estimation of external disturbances to a quantum memory [VM06], and the development of singular perturbation results for quantum stochastic differential equations [GV06].

The following articles based on graduate work are currently in preparation. The quantum filtering and control framework was developed in [BV06] with emphasis on the conceptual framework; technical details in the uncontrolled case can be found in chapter 5. A full technical account of the controlled case will appear in [BV07]. Finally, the fundamental role of quantum stopping times in quantum filtering and control, and in particular the application to optimal stopping and impulse control problems in the quantum context, is developed in [Van07].

Fundamentals of Nonlinear Filtering

This chapter is intended as an introduction to nonlinear filtering theory. It is impossible to do justice here to the full breadth and depth of the theory; such an exposition would (and indeed does) fill several textbooks. Our brief introduction serves to set the stage for the following chapters, and to make this thesis somewhat self-contained. We will not attempt to provide an introduction to the stochastic calculus of semimartingales, which is fundamental and will be used throughout. Detailed expositions can be found in, e.g., the excellent textbooks by P. Protter [Pro04], L. C. G. Rogers and D. Williams [RW00], R. J. Elliott [Eli82], or the classic tome by C. Dellacherie and P.-A. Meyer [DM82].

The study of nonlinear filtering has its origins in the work of R. L. Stratonovich [Str60, Str68]. Early work on nonlinear filtering in continuous time was largely heuristic. The Stratonovich theory suffered particularly from the lack of a stochastic integral (Stratonovich appears to have been unaware of Itô's theory, and later introduced a new stochastic integral [Str66] to justify his earlier work). H. J. Kushner [Kus64] developed the theory from the Itô perspective, and obtained a corrected version of the filtering equation (still using largely heuristic methods). Other important early contributions to nonlinear filtering include the work of W. M. Wonham [Won65], A. N. Shiryaev [Shi63], and R. Bucy [Buc65].

The modern theory of nonlinear filtering was developed in the late 1960s and early 1970s. There are two main approaches to the theory. The *martingale method* is based on the innovations process (see below) and on the systematic use of martingale representation theorems. The central role of the innovations process in filtering problems was developed by T. Kailath and P. Frost [FK71], but the definitive treatment of this approach (resolving a delicate representation issue) was given by M. Fujisaki, G. Kallianpur and H. Kunita [FKK72]. The martingale method is the focus of several textbooks on nonlinear filtering [LS01, Kal80, Kri05].

In this thesis we will use a different approach to nonlinear filtering theory—the *reference probability method*. This method relies fundamentally on the transformation of the filtering problem under absolutely continuous changes of the underlying probability measure, a technique that will be of central importance in the sequel. The reference probability approach was developed in fundamental work by G. Kallianpur and C. Striebel [KS68, KS69], H. J. Kushner [Kus67], T. E. Duncan [Dun68], R. E. Mortensen [Mor66], and M. Zakai [Zak69]. Textbook accounts can be found in [Par91, Ben92, EAM95].

1.1. Nonlinear filtering of a signal in white noise

1.1.1. The basic model. We will work on a standard filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbf{P})$ satisfying the usual hypotheses (see, e.g., [Pro04]). On this space are defined the following \mathcal{F}_t -adapted processes:

- (1) A càdlàg process X_t in a Polish state space \mathbb{S} .
- (2) A p -dimensional Wiener process B_t independent of X_t .

X_t is called the *signal process*, and B_t is called the *observation noise*. We assume for sake of technical concreteness that X_t is càdlàg (right-continuous with left limits) with values in a Polish (complete separable metric) space; this ensures, for example, that regular conditional probabilities always exist. In practice, any reasonable model encountered in applications will be of this form. For the time being we impose no further structure on the signal process.

We now introduce the p -dimensional *observation process*, defined as

$$(1.1) \quad Y_t = \int_0^t h(X_s) ds + B_t,$$

where $h : \mathbb{S} \rightarrow \mathbb{R}^p$ is a measurable map called the *observation function*. We denote by $\mathcal{F}_t^Y = \sigma\{Y_s : 0 \leq s \leq t\}$ the filtration generated by Y_t . The goal of filtering theory is to compute the filtered estimates

$$(1.2) \quad \pi_t(f) = \mathbf{E}(f(X_t) | \mathcal{F}_t^Y),$$

for a class of sufficiently nice measurable functions $f : \mathbb{S} \rightarrow \mathbb{R}$. Throughout this thesis, \mathbf{E} denotes the expectation $\mathbf{E}_{\mathbf{P}}$ with respect to the measure \mathbf{P} .

1.1.2. An explicit construction. Though it is not strictly necessary at this point, it is convenient to work with an explicit construction of the probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbf{P})$ and the processes X_t , Y_t and B_t . There is no loss of generality in doing so, and some of the arguments in the following become particularly transparent in this context. Let us therefore introduce this setup now, and adopt it consistently throughout the rest of this thesis.

In most applications of filtering theory, the real quantities of interest are the sample paths of the observation process and the sample paths of the signal process. It is thus convenient to define Ω as the set of all possible signal/observation sample paths. To this end, we will set $\Omega = D(\mathbb{R}_+; \mathbb{S}) \times C(\mathbb{R}_+; \mathbb{R}^p)$. Here $D(\mathbb{R}_+; \mathbb{S})$ is the space of càdlàg paths with values in \mathbb{S} , endowed with the Skorokhod topology which turns it into a Polish space, and $C(\mathbb{R}_+; \mathbb{R}^p)$ is the space of continuous paths with values in \mathbb{R}^p , endowed with the topology of uniform convergence on bounded intervals. For a detailed study of these spaces, see [Bil99, EK86]. We now define \mathcal{F} as the Borel σ -algebra on Ω (with respect to the above topology).

Note that any $\omega \in \Omega$ can be represented as a pair of sample paths $\omega = (x, y)$, where $x \in D(\mathbb{R}_+; \mathbb{S})$ and $y \in C(\mathbb{R}_+; \mathbb{R}^p)$. Define the coordinate processes

$$X_t : \Omega \rightarrow \mathbb{S}, \quad X_t(x, y) = x_t, \quad Y_t : \Omega \rightarrow \mathbb{R}^p, \quad Y_t(x, y) = y_t \quad \forall t \in \mathbb{R}_+.$$

Hence we can really think of a point $(x, y) = \omega \in \Omega$ as representing a joint signal/observation sample path. We can now define $\mathcal{F}_t^{XY} = \sigma\{(X_s, Y_s) : 0 \leq s \leq t\}$ and $\mathcal{F}_t^Y = \sigma\{Y_s : 0 \leq s \leq t\}$. Next, we introduce the measure \mathbf{P} on (Ω, \mathcal{F}) such that (X_t, Y_t) has the correct law; such a measure could be obtained, e.g., as the

measure induced on the path space $D(\mathbb{R}_+; \mathbb{S}) \times C(\mathbb{R}_+; \mathbb{R}^p)$ by a more general model as in the previous subsection. In particular, if we define

$$B_t = Y_t - \int_0^t h(X_s) ds,$$

then B_t is an $(\mathcal{F}_t^{XY}, \mathbf{P})$ -Wiener process independent of X_t . Finally, we define \mathcal{F}_t as the \mathbf{P} -augmentation of \mathcal{F}_t^{XY} ; this guarantees that this filtration satisfies the usual hypotheses (see [KS91, sec. 2.7] for extensive discussion). We have now obtained an explicit construction for the probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbf{P})$.

1.1.3. The Kallianpur-Striebel formula. The very simplest filtering scenario occurs when $h(x) = 0$. In this case, we can write $\mathbf{P} = \mu_x \times \mu_w$, where μ_x is the marginal of \mathbf{P} on $D(\mathbb{R}_+; \mathbb{S})$ (i.e., μ_x is the law of the signal process) and μ_w is the Wiener measure on $C(\mathbb{R}_+; \mathbb{R}^p)$. But then for any bounded measurable f ,

$$(1.3) \quad \mathbf{E}(f(X_{[0,T]}, Y_{[0,T]} | \mathcal{F}_T^Y) = \int f(x, Y_{[0,T]}) \mu_x(dx)$$

is clearly a version of the conditional expectation. This scenario is of course extraordinarily uninteresting, as the observations are completely uninformative.

Nonetheless, this is not at all a bad idea; it would be extremely convenient to express conditional expectations with respect to \mathcal{F}_T^Y by simply integrating out the signal process against a suitable measure. The problem is: how do we obtain the appropriate measure? The key point, as we will shortly see, is that the measure \mathbf{P} and the *reference measure* $\mathbf{Q} = \mu_x \times \mu_w$ are always equivalent—at least when restricted to a finite observation interval. Hence we can express conditional expectations under \mathbf{P} as conditional expectations under \mathbf{Q} , and the latter are precisely of the form (1.3). This is the main idea of the reference probability method.

Let us begin by recalling how conditional expectations transform under an equivalent change of measure.

LEMMA 1.1.1. *Let $\mu \sim \nu$ be two equivalent probability measures on a probability space (Ω, \mathcal{F}) , let $X \in L^1(\nu)$, and let $\mathcal{G} \subset \mathcal{F}$ be a sub- σ -algebra. Then*

$$\mathbf{E}_\nu(X | \mathcal{G}) = \frac{\mathbf{E}_\mu\left(\frac{d\nu}{d\mu} X | \mathcal{G}\right)}{\mathbf{E}_\mu\left(\frac{d\nu}{d\mu} | \mathcal{G}\right)} \quad \nu\text{-a.s., } \mu\text{-a.s.}$$

PROOF. Choose any $A \in \mathcal{G}$. Then

$$\begin{aligned} \mathbf{E}_\mu(\mathbf{E}_\mu\left(\frac{d\nu}{d\mu} X | \mathcal{G}\right) I_A) &= \mathbf{E}_\mu\left(\frac{d\nu}{d\mu} X I_A\right) = \mathbf{E}_\nu(X I_A) \\ &= \mathbf{E}_\nu(\mathbf{E}_\nu(X | \mathcal{G}) I_A) = \mathbf{E}_\mu\left(\frac{d\nu}{d\mu} \mathbf{E}_\nu(X | \mathcal{G}) I_A\right) = \mathbf{E}_\mu(\mathbf{E}_\mu\left(\frac{d\nu}{d\mu} | \mathcal{G}\right) \mathbf{E}_\nu(X | \mathcal{G}) I_A). \end{aligned}$$

As $A \in \mathcal{G}$ is arbitrary, we must have

$$\mathbf{E}_\mu\left(\frac{d\nu}{d\mu} X | \mathcal{G}\right) = \mathbf{E}_\mu\left(\frac{d\nu}{d\mu} | \mathcal{G}\right) \mathbf{E}_\nu(X | \mathcal{G}) \quad \mu\text{-a.s.}$$

(and ν -a.s. as $\mu \sim \nu$). It remains to note that $\mu \sim \nu$ implies that $\frac{d\nu}{d\mu} > 0$ a.s. \square

COROLLARY 1.1.2. *Let μ_1 be a probability measure on $(\Omega_1, \mathcal{F}_1)$, μ_2 be a probability measure on $(\Omega_2, \mathcal{F}_2)$, and let μ be a probability measure on $(\Omega_1 \times \Omega_2, \mathcal{F}_1 \times \mathcal{F}_2)$ such that $\mu \sim \mu_1 \times \mu_2$. Then for $\mathcal{G} = \{\emptyset, \Omega_1\} \times \mathcal{F}_2$ and any $X \in L^1(\mu)$,*

$$\mathbf{E}_\mu(X | \mathcal{G})(\omega_1, \omega_2) = \frac{\int \frac{d\mu}{d[\mu_1 \times \mu_2]}(\omega, \omega_2) X(\omega, \omega_2) \mu_1(d\omega)}{\int \frac{d\mu}{d[\mu_1 \times \mu_2]}(\omega, \omega_2) \mu_1(d\omega)} \quad \text{a.s.}$$

This is one particular form of the well-known *Bayes formula* (for a more general discussion of the abstract Bayes formula, see, e.g., [LS01, sec. 7.9]).

To apply these ideas in the filtering context, we could attempt to show that $\mathbf{P} \sim \mathbf{Q}$, where $\mathbf{Q} = \mu_x \times \mu_w$ is the reference measure defined above. We have to be a bit careful, however: the following standard example [KS91, pp. 192–193] shows that this is unlikely to be the case.

EXAMPLE 1.1.3. Consider another trivial filtering scenario: $h(x) = 1$. Then $Y_t = t + B_t$ where B_t is a Brownian motion under \mathbf{P} . On the other hand, Y_t is itself a Brownian motion under \mathbf{Q} . Hence we obtain

$$\lim_{t \rightarrow \infty} \frac{Y_t}{t} = 0 \quad \mathbf{Q}\text{-a.s.}, \quad \lim_{t \rightarrow \infty} \frac{Y_t}{t} = 1 \quad \mathbf{P}\text{-a.s.}$$

by the law of large numbers. Clearly $\mathbf{P} \sim \mathbf{Q}$ is impossible!

Fortunately, this example does not present any problems in practice, provided that we eliminate these pathological events. To this end, let us define the filtration $\mathcal{F}_T^\infty = \mathcal{D} \times \sigma\{y_t : 0 \leq t \leq T\}$, where \mathcal{D} is the Borel σ -algebra on the Skorokhod space $D(\mathbb{R}_+; \mathbb{S})$. We denote by \mathbf{P}_T^∞ and \mathbf{Q}_T^∞ the restriction to \mathcal{F}_T^∞ of the measures \mathbf{P} and \mathbf{Q} , respectively. Under a mild condition, we will show that a.s.

$$(1.4) \quad \frac{d\mathbf{P}_T^\infty}{d\mathbf{Q}_T^\infty} = Z_T(X_{[0,T]}, Y_{[0,T]}) = \exp\left(\int_0^T h(X_s) \cdot dY_s - \frac{1}{2} \int_0^T \|h(X_s)\|^2 ds\right)$$

($h(X_s) \cdot dY_s$ denotes the Itô differential $\sum_i h^i(X_s) dY_s^i$). This result is not surprising: $Z_T(X_{[0,T]}, Y_{[0,T]})$ is precisely the Girsanov transformation that makes $Y_{[0,T]}$ a Wiener process under \mathbf{Q} . The assumption we need for this to work is the basic assumption of the Girsanov theorem, i.e., that $Z_t(x, Y_{[0,t]})$ is an $(\mathcal{F}_t^Y, \mathbf{Q})$ -martingale for μ_x -a.e. $x \in D(\mathbb{R}_+; \mathbb{S})$. To guarantee that this is the case, let us impose the following integrability condition on the signal.

ASSUMPTION 1.1.4. μ_x satisfies the following integrability assumption:

$$\int_0^T \|h(x_s)\|^2 ds < \infty \quad \forall T < \infty, \quad \text{for } \mu_x\text{-a.e. } x \in D(\mathbb{R}_+; \mathbb{S}).$$

Note in particular that this assumption holds if h is continuous, as any path $x \in D([0, T]; \mathbb{S})$ is necessarily bounded [Bil99]. The assumption is trivial if h is bounded. In the sequel, we will always presume that Assumption 1.1.4 holds.

We are now in the position to prove (1.4).

LEMMA 1.1.5. Equation (1.4) holds, provided Assumption 1.1.4 is satisfied.

PROOF. Let $f : D(\mathbb{R}_+; \mathbb{S}) \rightarrow \mathbb{R}$ and $g : C([0, T]; \mathbb{R}^p) \rightarrow \mathbb{R}$ be nonnegative, bounded and measurable. Then $f(X_{[0,\infty)})$ and $g(B_{[0,T]})$ are \mathcal{F}_T^∞ -measurable, and indeed \mathcal{F}_T^∞ is generated by all random variables of this type. To prove the Lemma, it thus suffices to show that

$$\mathbf{E}_{\mathbf{P}}(f(X_{[0,\infty)})g(B_{[0,T]})) = \mathbf{E}_{\mathbf{Q}}(f(X_{[0,\infty)})g(B_{[0,T]})Z_T(X_{[0,T]}, Y_{[0,T]}).$$

To this end, recall that under \mathbf{P} , B_t is a Wiener process which is independent of X_t . Hence the left-hand side reduces to

$$\mathbf{E}_{\mathbf{P}}(f(X_{[0,\infty)})g(B_{[0,T]})) = \int_{D(\mathbb{R}_+; \mathbb{S})} f(x) \mu_x(dx) \times \int_{C([0,T]; \mathbb{R}^p)} g(b) \mu_w(db).$$

The right-hand side, on the other hand, reduces to

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}}(f(X_{[0,\infty)}) g(B_{[0,T]}) Z_T(X_{[0,T]}, Y_{[0,T]})) \\ = \int_{D(\mathbb{R}_+; \mathbb{S})} f(x) \left[\int_{C([0,T]; \mathbb{R}^p)} g(B(x, y)) Z_T(x, y) \mu_w(dy) \right] \mu_x(dx), \end{aligned}$$

where we have used Tonelli's theorem. Here we have written $B(x, y)$ to denote the dependence of the process $B_{[0,T]} \in C([0, T]; \mathbb{R}^p)$ on $X_{[0,\infty)}$ and $Y_{[0,T]}$. But by Assumption 1.1.4, Novikov's criterion and Girsanov's theorem, we obtain

$$\int_{C([0,T]; \mathbb{R}^p)} g(B(x, y)) Z_T(x, y) \mu_w(dy) = \int_{C([0,T]; \mathbb{R}^p)} g(b) \mu_w(db)$$

for μ_x -a.e. $x \in D(\mathbb{R}_+; \mathbb{S})$. Hence the proof is complete. \square

REMARK 1.1.6. Lemma 1.1.5 implies directly that $Z_t = Z_t(X_{[0,t]}, Y_{[0,t]})$ is a martingale under \mathbf{Q} . Indeed, Z_t is a positive local martingale, and hence a positive supermartingale by Fatou's Lemma [Pro04, p. 138]. To show that Z_t is a martingale, it suffices to note that $\mathbf{E}_{\mathbf{Q}}(Z_t) = \mathbf{E}_{\mathbf{P}}(1) = 1$ for all $t < \infty$. After all, suppose that Z_t is not a martingale; then we can find $s < t$ such that $\mathbf{E}_{\mathbf{Q}}(Z_t | \mathcal{F}_s) < Z_s$ with positive probability. But then $\mathbf{E}_{\mathbf{Q}}(Z_t) < \mathbf{E}_{\mathbf{Q}}(Z_s)$, which is a contradiction.

Using these results, we immediately obtain an explicit representation of the conditional expectation. We always presume that Assumption 1.1.4 holds.

COROLLARY 1.1.7 (Kallianpur-Striebel). *If $X \in L^1(\mathbf{P})$ is \mathcal{F}_T^∞ -measurable, then*

$$(1.5) \quad \mathbf{E}(X | \mathcal{F}_T^Y)(x, y) = \frac{\int X(x', y) Z_T(x', y) \mu_x(dx')}{\int Z_T(x', y) \mu_x(dx')} \quad \text{a.s.}$$

1.1.4. The Zakai and Kushner-Stratonovich equations. Let us return to the filtering problem. We begin by defining for any measurable $f : \mathbb{S} \rightarrow \mathbb{R}$

$$\sigma_t(f)(x, y) = \int f(x'_t) Z_t(x', y) \mu_x(dx').$$

By the Kallianpur-Striebel formula, it is evidently the case that a.s.

$$\pi_t(f) = \mathbf{E}(f(X_t) | \mathcal{F}_t^Y) = \frac{\sigma_t(f)}{\sigma_t(1)}.$$

It is often convenient to work with an equation for $\sigma_t(f)$ rather than with $\pi_t(f)$ directly. The former is called the *Zakai equation* (or Duncan-Mortensen-Zakai equation), while we have already encountered the latter as the Kushner-Stratonovich equation. We will actually derive the latter by applying the Kallianpur-Striebel formula and Itô's rule to the Zakai equation.

Before we can proceed along this path, we need to impose some more structure on our signal process—up to the present point we have required almost nothing of the signal. The theory can be set up for general semimartingale signals (as in [FKK72]), but this does not generally give rise to *recursive* filters (i.e., filters that can be updated at each time step using the new observations only). To obtain the latter, we restrict our attention to Markov signals.

ASSUMPTION 1.1.8. X_t is a Feller-Markov process with generator \mathcal{L} .

What does this do for us? Of particular interest at this point is the following elementary result, known as Dynkin's formula [Kal97, Lemma 17.21]: for any $f \in D(\mathcal{L})$ ($D(\mathcal{L})$ is the domain of the generator \mathcal{L}), the process

$$M_t^f = f(X_t) - f(X_0) - \int_0^t \mathcal{L}f(X_s) ds$$

is a martingale. Hence in particular $f(X_t)$ is a semimartingale, and we can apply the Itô calculus. The basic approach will now consist of applying Itô's rule to the integrand in the definition of $\sigma_t(f)$; after taking the μ_x -expectation, this gives the Zakai equation. Let us work out the details.

REMARK 1.1.9. In the following, we will impose significantly more restrictive conditions on the signal and observations than is strictly necessary; however, these will suffice for our purposes, and the proofs are highly simplified. For a reference probability proof that does not suffer from these restrictions, see [Ben92]. The martingale method can also be used under weaker conditions, see, e.g., [LS01].

PROPOSITION 1.1.10 (Zakai equation). *Let $f \in D(\mathcal{L})$. Let $\mathbf{E}((M_t^f)^2) < \infty$,*

$$\mathbf{E}_{\mathbf{Q}} \int_0^t Z_s^2 ds < \infty, \quad \mathbf{E}_{\mathbf{Q}} \int_0^t (Z_s f(X_s) h(X_s))^2 ds < \infty, \quad \mathbf{E} \int_0^t |\mathcal{L}f(X_s)| ds < \infty.$$

Then the Zakai equation holds (here $\nu(f) = \mathbf{E}(f(X_0))$):

$$(1.6) \quad \sigma_t(f) = \nu(f) + \int_0^t \sigma_s(\mathcal{L}f) ds + \int_0^t \sigma_s(hf) \cdot dY_s.$$

PROOF. Using Itô's rule and Dynkin's formula, we obtain

$$f(X_t) Z_t = f(X_0) + \int_0^t Z_s f(X_s) h(X_s) \cdot dY_s + \int_0^t Z_s \mathcal{L}f(X_s) ds + \int_0^t Z_s dM_s^f,$$

where $Z_t = Z_t(X_{[0,t]}, Y_{[0,t]})$. The usual quadratic variation term $[M^f, Y]$ does not appear in this expression, as $M_t^f Y_t$ is a martingale under \mathbf{Q} and hence $[M^f, Y]_t = 0$ [Pro04, p. 73]. We now take the μ_x -expectation on both sides. The last term on the right vanishes, as by our first two assumptions this is a martingale under μ_x for μ_w -a.e. observation sample path (see, e.g., [Ell82, Ch. 11]). It remains to show that the order of integration in the remaining terms on the right may be exchanged. That μ_x and the time integral may be exchanged follows from our last assumption and Fubini's theorem. The exchange of integration order in the remaining stochastic integral follows from our assumption that the integrand is square-integrable, so that we can apply the Fubini theorem for stochastic integrals [LS01, Thm. 5.15]. \square

Note that the (rather restrictive) integrability assumptions in this result are needed to have the various integrals behave nicely under the measure μ_x . If all the integrands were bounded, this would not be an issue; but Z_t will not be bounded regardless of what we choose for h ! A trick to dispose of the integrability conditions is to approximate Z_t by the bounded quantity $Z_t^\varepsilon = Z_t/(1 + \varepsilon Z_t)$, go through the procedure above, and then take the limit $\varepsilon \rightarrow 0$ at the end of the day. This establishes the Zakai equation under weak conditions—enough to make the limits converge—but requires a lot of work and is not very insightful. We refer to [Ben92] for the details. A similar but less tedious approach to weaken the conditions of Prop. 1.1.10 proceeds through localization, see [Par91]. Rather than taking this route,

we will content ourselves by obtaining some sufficient conditions that ensure that the requirements of Prop. 1.1.10 are satisfied.

LEMMA 1.1.11. *For any measurable $g : \mathbb{S} \rightarrow \mathbb{R}$, we have*

$$\mathbf{E}_{\mathbf{Q}} \left(\int_0^t Z_s^2 g(X_s)^2 ds \right) = \mathbf{E} \int_0^t \exp \left(\int_0^s \|h(X_u)\|^2 du \right) g(X_s)^2 ds.$$

PROOF. As the integrand is nonnegative, we may freely exchange the order of integration by Tonelli's theorem. Let us write

$$\mathbf{E}_{\mathbf{Q}} \left(\int_0^t Z_s^2 g(X_s)^2 ds \right) = \int_0^t \left[\int \left[\int Z_s(x, y)^2 \mu_w(dy) \right] g(x_s)^2 \mu_x(dx) \right] ds,$$

But using Assumption 1.1.4, the inner integral evaluates to

$$\int Z_s(x, y)^2 \mu_w(dy) = \exp \left(\int_0^s \|h(x_u)\|^2 du \right).$$

This completes the proof. \square

COROLLARY 1.1.12. *The conditions of Proposition 1.1.10 are satisfied if h , f and $\mathcal{L}f$ are bounded, or alternatively if h is bounded and*

$$\mathbf{E} \int_0^t f(X_s)^2 ds < \infty, \quad \mathbf{E} \int_0^t |\mathcal{L}f(X_s)| ds < \infty, \quad \mathbf{E}((M_t^f)^2) < \infty.$$

Let us finally show how to obtain the Kushner-Stratonovich equation.

COROLLARY 1.1.13 (Kushner-Stratonovich equation). *Under the conditions of Proposition 1.1.10 or Corollary 1.1.12, the filtered estimate $\pi_t(f)$ satisfies*

$$(1.7) \quad \pi_t(f) = \nu(f) + \int_0^t \pi_s(\mathcal{L}f) ds + \int_0^t \{ \pi_s(hf) - \pi_s(h) \pi_s(f) \} (dY_s - \pi_s(h) ds).$$

PROOF. We simply apply Itô's rule to the Kallianpur-Striebel formula $\pi_t(f) = \sigma_t(f)/\sigma_t(1)$, using the Zakai equation for $\sigma_t(\cdot)$. Note that $\sigma_t(1) > 0$ a.s. as $Z_t > 0$ a.s., so that we can directly apply the function x^{-1} , being C^2 on $\{x \in \mathbb{R} : x > 0\}$, using Itô's rule. The result now follows by straightforward calculation. \square

1.1.5. The innovations process. The *innovations process* $d\bar{W}_t = dY_t - \pi_t(h) dt$ that appears in the Kushner-Stratonovich equation has a special significance: \bar{W}_t is in fact an $(\mathcal{F}_t^Y, \mathbf{P})$ -Wiener process. Let us prove this.

PROPOSITION 1.1.14 (Innovations process). *Suppose that*

$$\mathbf{E} \int_0^T \|h(X_t)\| dt < \infty \quad \forall T < \infty.$$

Then the innovations process \bar{W}_t , defined by

$$\bar{W}_t = Y_t - \int_0^t \pi_s(h) ds,$$

is an $(\mathcal{F}_t^Y, \mathbf{P})$ -Wiener process.

PROOF. First, we show that \bar{W}_t is a martingale. Note that

$$\mathbf{E}(\bar{W}_t - \bar{W}_s | \mathcal{F}_s) = \mathbf{E}(B_t - B_s | \mathcal{F}_s) + \mathbf{E} \left(\int_s^t (h(X_u) - \pi_u(h)) du \middle| \mathcal{F}_s \right).$$

The first term vanishes by the martingale property of the \mathcal{F}_t -Wiener process B_t . Hence we obtain using $\mathcal{F}_t^Y \subset \mathcal{F}_t$

$$\mathbf{E}(\overline{W}_t - \overline{W}_s | \mathcal{F}_s^Y) = \mathbf{E} \left(\int_s^t (h(X_u) - \mathbf{E}(h(X_u) | \mathcal{F}_u^Y)) du \middle| \mathcal{F}_s^Y \right).$$

Using the integrability condition to apply Fubini's theorem, we can conclude that the right-hand side vanishes. Hence \overline{W}_t is a \mathcal{F}_t^Y -martingale. But \overline{W}_t has continuous paths and evidently $[\overline{W}^i, \overline{W}^j]_t = t\delta_{ij}$, so the claim follows from Lévy's theorem. \square

We will need the following important corollary in the sequel.

COROLLARY 1.1.15. *Under the condition of Proposition 1.1.14,*

$$\frac{d\mathbf{P}|_{\mathcal{F}_T^Y}}{d\mathbf{Q}|_{\mathcal{F}_T^Y}} = \mathbf{Q}(Z_T | \mathcal{F}_T^Y) = \sigma_T(1) = \exp \left(\int_0^T \pi_s(h) \cdot dY_s - \frac{1}{2} \int_0^T \|\pi_s(h)\|^2 ds \right).$$

PROOF. The first two equalities are immediate from (1.4) and Lemma 1.1.5. The last equality follows from the fact that by Proposition 1.1.14, the exponential expression is a Girsanov transformation that makes Y_t a Wiener process under \mathbf{Q} . But this expression is \mathcal{F}_T^Y -measurable, so that the claim follows from the uniqueness of the Radon-Nikodym derivative. \square

1.2. Finite-state Markov signals and the Wonham filter

In this section we will specialize the general filtering theory of the previous section to the case that \mathbb{S} is a finite set. This particularly simple case is widely used in applications, and forms the basis for chapters 2 and 3.

Let $\mathbb{S} = \{a_1, \dots, a_d\}$, $d < \infty$. We choose the law μ_x of the signal process such that X_t is a Feller-Markov process with the following generator:

$$\mathcal{L}f(a_i) = \sum_{j=1}^d \lambda_{ij} f(a_j),$$

where the matrix $\Lambda = (\lambda_{ij})$ is such that $\lambda_{ij} \geq 0$ for $i \neq j$, and $\lambda_{ii} = -\sum_{j \neq i} \lambda_{ij}$. Λ is called the *transition intensities matrix*, as the Markov property implies

$$\mathbf{P}(X_{t+\Delta} = a_j | X_t = a_i) = \lambda_{ij} \Delta + o(\Delta), \quad i \neq j.$$

We will denote the initial distribution of the process by $\nu^i = \mathbf{P}(X_0 = a_i)$. Note that in the finite state case, we can naturally associate to any function $g : \mathbb{S} \rightarrow \mathbb{R}$ a corresponding vector $g^i = g(a_i)$. We will implicitly identify these two notations whenever no confusion can arise. Note also that any such function is necessarily bounded; a finite set always attains its extrema!

We now introduce, as usual, the white noise observations

$$Y_t = \int_0^t h(X_s) ds + B_t.$$

For convenience and notational simplicity we will work with one-dimensional observations $h : \mathbb{S} \rightarrow \mathbb{R}$, but everything we will do extends without trouble to observations of higher dimension.

The filtering theory of the previous section applies immediately; in particular, as any function on \mathbb{S} is bounded, all the technical assumptions (e.g., those of Corollary 1.1.12) are automatically satisfied. These filters, however, are not yet

in recursive form: e.g., the Zakai equation for $\sigma_t(f)$ depends on $\sigma_t(\mathcal{L}f)$, etc. To obtain a recursive expression, we introduce the quantities

$$\pi_t^i = \mathbf{P}(X_t = a_i | \mathcal{F}_t^Y), \quad \sigma_t^i = \mathbf{E}_{\mathbf{Q}}(I_{X_t=a_i} Z_t | \mathcal{F}_t^Y),$$

so that we can write

$$\pi_t(f) = \sum_{i=1}^d \pi_t^i f(a_i) = f^* \pi_t, \quad \sigma_t(f) = \sum_{i=1}^d \sigma_t^i f(a_i) = f^* \sigma_t, \quad \sigma_t(1) = \sum_{i=1}^d \sigma_t^i = |\sigma_t|.$$

Here v^* (M^*) is the adjoint of the vector v (matrix M), and $|v|$ is the ℓ_1 -norm of v . From Proposition 1.1.10 and Corollary 1.1.13, we now obtain the following result.

COROLLARY 1.2.1 (Wonham filter). *Let $h : \mathbb{S} \rightarrow \mathbb{R}$. Then the conditional density π_t obeys the Wonham equation*

$$(1.8) \quad \pi_t = \nu + \int_0^t \Lambda^* \pi_s ds + \int_0^t (H - h^* \pi_s) \pi_s (dY_s - h^* \pi_s ds),$$

and the unnormalized conditional density σ_t obeys the Wonham-Zakai equation

$$(1.9) \quad \sigma_t = \nu + \int_0^t \Lambda^* \sigma_s ds + \int_0^t H \sigma_s dY_s.$$

Here $H = \text{diag}(h)$ is the diagonal matrix with $(H)_{ii} = h(a_i)$.

This is very convenient: the filter is just a finite-dimensional stochastic differential equation, driven by the semimartingale Y_t . If we can establish uniqueness of the solution, then the utility of these equations is clear: a digital signal processing device updates the current value of π_t recursively at every time step; subsequently any estimate of the signal at that time, i.e., the best estimate of $f(X_t)$, can be obtained through the simple inner product $\pi_t(f) = f^* \pi_t$.

EXAMPLE 1.2.2 (Binary symmetric signal). Consider $\mathbb{S} = \{a_1, a_2\}$ with equal transition rates $\lambda_{12} = \lambda_{21} = \lambda$. Then (1.8) becomes

$$d\pi_t^1 = \lambda (\pi_t^2 - \pi_t^1) dt + (h^1 - h^1 \pi_t^1 - h^2 \pi_t^2) \pi_t^1 \{dY_t - (h^1 \pi_t^1 + h^2 \pi_t^2) dt\}.$$

But $\pi_t^1 + \pi_t^2 = 1$, so we obtain

$$d\pi_t^1 = -2\lambda (\pi_t^1 - \frac{1}{2}) dt + (h^1 - h^2) \pi_t^1 (1 - \pi_t^1) \{dY_t - (h^1 \pi_t^1 + h^2 (1 - \pi_t^1)) dt\}.$$

Compare this expression to the discussion of section 0.2.

We conclude this section with a statement of uniqueness. This is important, because we have only established that the conditional density π_t must satisfy the Wonham equation. Uniqueness, however, guarantees that we can always obtain the conditional density by solving the Wonham equation.

LEMMA 1.2.3. *Equations (1.8) and (1.9) have unique solutions.*

PROOF. Equations (1.8) and (1.9) have locally and globally Lipschitz coefficients, respectively. Hence existence and uniqueness of the solution to (1.9) is immediate from standard results, whereas (1.8) must have a unique solution up to an accessible explosion time ζ (see, e.g., [Pro04, Ch. V]). But ζ must be infinite a.s., as $\pi_t^i = \mathbf{P}(X_t = a_i | \mathcal{F}_t^Y)$ for $t < \zeta$ by uniqueness and $\mathbf{P}(X_t = a_i | \mathcal{F}_t^Y) \leq 1$. \square

1.3. Nonlinear filtering for diffusions

Let us now turn to the diffusion case. We are interested in the state space $\mathbb{S} = \mathbb{R}^d$. The signal process is taken to be an Itô diffusion

$$(1.10) \quad dX_t = b(X_t) dt + a(X_t) dW_t, \quad X_0 \sim \nu,$$

where $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $a : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times q}$ are assumed to be of class C^2 with bounded first and second derivatives, W_t is a q -dimensional Wiener process, and the initial measure ν has finite moments. This defines a Markov diffusion, whose generator \mathcal{L} is given by

$$\mathcal{L}f(x) = \sum_{i=1}^d b^i(x) \frac{\partial f(x)}{\partial x^i} + \frac{1}{2} \sum_{i,j=1}^d \sum_{k=1}^q a^{ik}(x) a^{jk}(x) \frac{\partial^2 f(x)}{\partial x^i \partial x^j}$$

at least for bounded f with bounded first and second derivatives.

REMARK 1.3.1. We have not taken into account the possible existence of W_t in our construction of the probability space Ω , but this is not an essential issue. We could either extend Ω to be slightly larger; or we could construct the strong solution to (1.10) on a separate Wiener space, then choose μ_x to be the measure induced by this diffusion on the path space of X_t . Perhaps the most elegant solution is to characterize the weak solution to (1.10) as the solution of a Stroock-Varadhan martingale problem: μ_x is defined as the (unique) measure that makes

$$M_t^f = f(X_t) - f(X_0) - \int_0^t \mathcal{L}f(X_s) ds$$

a martingale for any $f \in C^2$ with compact support, and no auxiliary spaces are necessary (see, e.g., [RW00, Sec. V.19]). Ultimately these methods are equivalent (at least for the conditions we have placed on b and a), so we will not worry too much about the specific construction.

We now introduce the observation function $h : \mathbb{R}^d \rightarrow \mathbb{R}^p$ and the observations

$$Y_t = \int_0^t h(X_s) ds + B_t,$$

as usual. If h is bounded, then we can apply Corollary 1.1.12. Let us state here without proof that under the conditions we have imposed, the boundedness requirement of h can be weakened to a linear growth condition: that is, we will assume that $\|h(x)\| \leq K(1 + \|x\|)$ for some finite constant K and all $x \in \mathbb{R}^d$. The proof of this fact, which proceeds as we have outlined in section 1.1.12 (approximate Z_t by $Z_t^\varepsilon = Z_t/(1 + \varepsilon Z_t)$), can be found in [Ben92, Ch. 4]. We then obtain:

PROPOSITION 1.3.2. *Under the above assumptions, the Zakai and Kushner-Stratonovich equations (1.6) and (1.7) hold for any bounded $f \in C^2$ with bounded first and second derivatives.*

In the finite state case, we converted the filtering equation to recursive form by introducing the vectors σ_t and π_t such that $\sigma_t(f) = f^* \sigma_t$ and $\pi_t(f) = f^* \pi_t$. We would like to do something similar here. Proceeding formally, we could introduce the densities $\varrho_t(x)$ and $p_t(x)$ by setting

$$\sigma_t(f) = \int_{\mathbb{R}^d} f(x) \varrho_t(x) dx, \quad \pi_t(f) = \int_{\mathbb{R}^d} f(x) p_t(x) dx,$$

(recall that $\nu(f) = \pi_0(f)$, so $p_0(x)$ is the density of ν). Substituting these expressions into the Zakai and Kushner-Stratonovich equations, and formally integrating by parts, we obtain

$$\varrho_t(x) = p_0(x) + \int_0^t \mathcal{L}^* \varrho_s(x) ds + \int_0^t h(x) \varrho_s(x) \cdot dY_s,$$

$$\begin{aligned} p_t(x) &= p_0(x) + \int_0^t \mathcal{L}^* p_s(x) ds \\ &+ \int_0^t \left\{ h(x) - \int_{\mathbb{R}^d} h(x) p_s(x) dx \right\} p_s(x) \cdot \left(dY_s - \int_{\mathbb{R}^d} h(x) p_s(x) dx dt \right). \end{aligned}$$

Here \mathcal{L}^* denotes the formal adjoint of \mathcal{L} , i.e.

$$\mathcal{L}^* f(x) = - \sum_{i=1}^d \frac{\partial}{\partial x^i} (b^i(x) f(x)) + \frac{1}{2} \sum_{i,j=1}^d \sum_{k=1}^q \frac{\partial^2}{\partial x^i \partial x^j} (a^{ik}(x) a^{jk}(x) f(x)).$$

The rigorous justification of this procedure is not at all straightforward, however. It is far from clear that $\pi_t(\cdot)$ and $\sigma_t(\cdot)$ even possess a density with respect to the Lebesgue measure, let alone that it is sufficiently smooth to apply \mathcal{L}^* to. Moreover, the issue of uniqueness of the solutions to these equations is delicate. The well-posedness of the equations for $\varrho_t(x)$ and $p_t(x)$ has been studied extensively in the literature in various settings and using a range of methods; see H. Kunita [Kun90], É. Pardoux [Par82, Par91], B. L. Rozovskiĭ [Roz90], A. Bensoussan [Ben92], and S. J. Sheu [She83], for extensive discussion and further references.

1.4. Pathwise filtering

We conclude this introductory chapter with a discussion of the *pathwise* formulation of nonlinear filtering problems. The goal of this theory is to give unambiguous meaning to the statement: given an observation sample path $y \in C([0, T]; \mathbb{R}^p)$, what is the corresponding filtered estimate $\pi_T(f)$? We have not answered this question above; in fact, we have only defined the filter $\pi_T(f)$ on an unspecified set in the Borel σ -algebra of $C(\mathbb{R}_+; \mathbb{R}^p)$ of measure one. This is all we can hope for, of course, if we define the filter purely as a conditional expectation; in practice, on the other hand, it could be important to demand a little more.

Consider the practical implementation of a nonlinear filter. When we observe a signal in an engineering or laboratory environment, we just obtain one sample path $y \in C(\mathbb{R}_+; \mathbb{R}^p)$ at a time. Moreover, such a sample path will never look exactly like our white noise observation model: most real-world signals are either smooth or sampled at a high frequency (the latter can be included in our setup by joining the sampled points with straight line segments). In particular, real-world sample paths are of finite variation, and the set of finite variation paths has measure zero under the Wiener measure. Hence it is possible, within the framework of the previous sections, to have a perfectly respectable version of the filter which takes arbitrary values on any sample path that could possibly be obtained in the real world.

Clearly this is not a very desirable situation, but such a construction would also be highly artificial. On the other hand, the fact that we are forced in real life to work with approximate sample paths suggests that we would like to define our nonlinear filter as a *continuous* functional on $C(\mathbb{R}_+; \mathbb{R}^p)$; this way, the filter will be close to optimal even if we only have smoothed or sampled paths at our disposal. This

is the motivation behind the notion of pathwise filtering as introduced by J. M. C. Clark [Cla78], and subsequently studied by many authors including M. H. A. Davis [Dav80, Dav82], H. J. Kushner [Kus79], and É. Pardoux [Par82]. It is not at all obvious a priori that such a continuous version of the filter can be obtained, but it turns out that this is in fact possible under quite general conditions (see [CC05] for a recent exposition with careful attention to the technical subtleties).

The usefulness of the pathwise filtering theory significantly exceeds its original intention. Certainly pathwise filters are important for the study of filtering with smoothed observation sample paths, as well as for the design of good discretization schemes (so-called robust discretizations, see, e.g., [Kus79]). The continuity property of the filter is important for these applications. Another important application, however, is the study of nonlinear filtering equations for diffusions. As we have seen in the previous section, the question of well-posedness of the density forms of the Zakai and Kushner-Stratonovich equations is quite delicate; one of the problems here is the relative difficulty inherent in the study of stochastic partial differential equations (SPDEs). The pathwise filtering method, on the other hand, allows us to study filtering problems with a fixed observation path, and hence the SPDEs reduce to ordinary PDEs with random coefficients. This provides an important technical tool for the study of nonlinear filtering for diffusions (see, e.g., [Par82]); it is not the continuity property that is important here, but the pathwise nature of these filters. In a similar spirit, we will find it useful on occasion to study properties of the nonlinear filter for some “frozen” observation path. The pathwise filtering method provides an invaluable tool for this purpose, as it allows us to define the filters unambiguously for an arbitrary observation $y \in C(\mathbb{R}_+; \mathbb{R}^p)$.

1.4.1. The pathwise Kallianpur-Striebel formula. Let us return for the moment to the general filtering setup of section 1.1 and in particular to Corollary 1.1.7, the Kallianpur-Striebel formula. We would like to formulate this central result in an unambiguous way for a fixed observation path $y \in C(\mathbb{R}_+, \mathbb{R}^p)$. An examination of the Kallianpur-Striebel formula quickly uncovers the basic problem: the quantity $Z_T(x, y)$ is defined in terms of a stochastic integral with the observations as the integrator. Such a stochastic integral can be defined as a limit in probability at best, and certainly cannot be given an unambiguous meaning for every observation sample path individually (see [Pro04, Ch. I] for a nice discussion).

We are in luck, however: in this particular case the integrand is independent of the observations, i.e., the stochastic integral can be interpreted pathwise as a Wiener integral for every signal sample path. We could reverse the role of the observation process and the signal process through a simple integration by parts, provided that $h(X_t)$ is a semimartingale: after all, by Itô’s rule

$$h(X_T) \cdot Y_T = \int_0^T h(X_s) \cdot dY_s + \int_0^T Y_s \cdot dh(X_s) \quad \text{a.s.}$$

The left-hand side and the last term on the right-hand side can both be interpreted pathwise with respect to the observations. This is precisely what we will do.

For the rest of this section, we make the following basic assumption:

ASSUMPTION 1.4.1. $h(X_s)$ is a càdlàg semimartingale.

PROPOSITION 1.4.2 (Pathwise Kallianpur-Striebel formula). *Define*

$$\tilde{Z}_T(X_{[0,T]}, y) = \exp \left(h(X_T) \cdot y_T - \int_0^T y_s \cdot dh(X_s) - \frac{1}{2} \int_0^T \|h(X_s)\|^2 ds \right)$$

for any $y \in C_0(\mathbb{R}_+, \mathbb{R}^p)$ (the space of continuous paths such that $y_0 = 0$, with the topology of uniform convergence on bounded intervals). Let X be any μ_x -integrable functional of the signal process. Then

$$(1.11) \quad \Pi_T(X, y) = \frac{\int X(x) \tilde{Z}_T(x, y) \mu_x(dx)}{\int \tilde{Z}_T(x, y) \mu_x(dx)}$$

defines a unique regular version of the conditional expectation $\mathbf{E}(\cdot | \mathcal{F}_T^Y)$ on $D(\mathbb{R}_+; \mathbb{S})$.

PROOF. This follows from Itô's rule and the Kallianpur-Striebel formula. \square

For the purposes of filtering, let us introduce the pathwise filtering functional

$$(1.12) \quad \tau_t(f, y) = \mathbf{E} \left(f(X_t) \exp \left(- \int_0^t y_s \cdot dh(X_s) - \frac{1}{2} \int_0^t \|h(X_s)\|^2 ds \right) \right)$$

for any $y \in C_0(\mathbb{R}_+, \mathbb{R}^p)$. Then by the Kallianpur-Striebel formula,

$$\pi_t(f)(x, y) = \frac{\sigma_t(f)(x, y)}{\sigma_t(1)(x, y)} = \frac{\tau_t(f e^{h \cdot y_t}, y)}{\tau_t(e^{h \cdot y_t}, y)} \quad \text{a.e. } (x, y) \in \Omega.$$

Under sufficient regularity conditions, the pathwise filtering functional is continuously differentiable in time: applying the Itô rule to $\tau_t(f, y)$, we see that the latter is given by the expectation over time integrals and local martingales. If the local martingales are martingales, and Fubini's theorem can be applied, then we obtain a pathwise form of the Zakai equation where the stochastic integral term has disappeared. Rather than working out general conditions under which this can be done, let us briefly investigate the two cases which will be of interest to us in the sequel.

REMARK 1.4.3. In the theory of stochastic differential equations, H. Doss [Doss77] and H. J. Sussmann [Sus78] have suggested a general method to define pathwise solutions. Their methods demonstrate that for a wide class of stochastic differential equations, a change of variables can be constructed which transforms such equations into ordinary differential equations. When applied to the unnormalized Wonham equation (1.9), for example, this method indeed gives rise to the same pathwise filtering equation as in Proposition 1.4.5 below. In the filtering context, however, the approach through integration by parts in the Kallianpur-Striebel formula is much more fundamental. It is also somewhat simpler, as it does not require us to prove well-posedness of the solutions of filtering equations, which is particularly advantageous in the infinite-dimensional setting.

1.4.2. The pathwise Wonham filter. Consider case where the signal process has a finite state space $\mathbb{S} = \{a_1, \dots, a_d\}$, as in section 1.2. Let us define

$$\tau_t^i(y) = \mathbf{E} \left(I_{X_t=a_i} \exp \left(- \int_0^t y_s dh(X_s) - \frac{1}{2} \int_0^t h(X_s)^2 ds \right) \right).$$

Then by the pathwise Kallianpur-Striebel formula, we evidently have

$$\sigma_t^i = e^{h(a_i)Y_t} \tau_t^i(Y_{[0,t]}), \quad \pi_t^i = \frac{\sigma_t^i}{|\sigma_t|} = \frac{e^{h(a_i)Y_t} \tau_t^i(Y_{[0,t]})}{\sum_i e^{h(a_i)Y_t} \tau_t^i(Y_{[0,t]})}, \quad \text{a.s.}$$

Let us begin by showing that the functional $\tau_t^i(y)$ is continuous.

PROPOSITION 1.4.4. $\tau_t^i(y)$ is locally Lipschitz continuous in y .

PROOF. We can readily estimate

$$|\tau_t^i(y) - \tau_t^i(y')| \leq \mathbf{E} \left| \exp \left(- \int_0^t y_s dh(X_s) \right) - \exp \left(- \int_0^t y'_s dh(X_s) \right) \right|.$$

Hence we obtain, using $|e^a - e^b| \leq e^{a \vee b} |a - b|$,

$$|\tau_t^i(y) - \tau_t^i(y')| \leq \mathbf{E} \left(e^{|\int_0^t y_s dh(X_s)| \vee |\int_0^t y'_s dh(X_s)|} \left| \int_0^t (y_s - y'_s) dh(X_s) \right| \right).$$

Next, we estimate as follows: for any $y \in C_0([0, t]; \mathbb{R})$,

$$\left| \int_0^t y_s dh(X_s) \right| \leq K J(t) \sup_{s \in [0, t]} |y_s|,$$

where $K = \max\{h(a_i) - h(a_j) : i \neq j\}$ and $J(t)$ is the number of jumps made by X_s in the interval $s \in [0, t]$. Hence we obtain

$$|\tau_t^i(y) - \tau_t^i(y')| \leq \mathbf{E} \left(K J(t) e^{K R J(t)} \sup_{s \in [0, t]} |y_s - y'_s| \right),$$

where $R = \max\{\sup_{s \in [0, t]} |y_s|, \sup_{s \in [0, t]} |y'_s|\}$. But it follows from standard results (see, e.g., [RW00, sec. IV.21]) that the expectation is finite. Hence we have shown that for every $R < \infty$, there is some constant K_R such that

$$|\tau_t^i(y) - \tau_t^i(y')| \leq K_R \sup_{s \in [0, t]} |y_s - y'_s| \quad \forall y_s, y'_s \text{ s.t. } \sup_{s \in [0, t]} |y_s|, \sup_{s \in [0, t]} |y'_s| \leq R.$$

Hence the functional $\tau_t^i(y)$ is locally Lipschitz continuous. \square

We can now find a differential equation for $\tau_t^i(y)$.

PROPOSITION 1.4.5 (Pathwise Wonham filter). $\tau_t(y)$ satisfies

$$\frac{d\tau_t^i(y)}{dt} = \sum_{j=1}^d \tau_t^j(y) \lambda_{ji} e^{(h(a_j) - h(a_i))y_t} - \frac{1}{2} h(a_i)^2 \tau_t^i(y), \quad \tau_0^i = \nu^i.$$

PROOF. Let us first establish the result for $y \in C^1$. In this case, we can integrate by parts on a pathwise basis:

$$\begin{aligned} \tau_t^i(y) &= \mathbf{E} \left(I_{X_t = a_i} \exp \left(-h(X_t) y_t + \int_0^t h(X_s) \dot{y}_s ds - \frac{1}{2} \int_0^t h(X_s)^2 ds \right) \right) \\ &= e^{-h(a_i)y_t} \mathbf{E} \left(I_{X_t = a_i} \exp \left(\int_0^t h(X_s) \dot{y}_s ds - \frac{1}{2} \int_0^t h(X_s)^2 ds \right) \right). \end{aligned}$$

Using Dynkin's formula, we have

$$I_{X_t = a_i} = I_{X_0 = a_i} + \sum_{j=1}^d \lambda_{ji} \int_0^t I_{X_s = a_j} ds + N_t,$$

where N_t is a bounded martingale on any bounded time interval. Define

$$T_t^i(y) = \mathbf{E} \left(I_{X_t = a_i} \exp \left(\int_0^t h(X_s) \dot{y}_s ds - \frac{1}{2} \int_0^t h(X_s)^2 ds \right) \right).$$

Using Itô's rule and the fact that the exponential term is bounded, we easily obtain

$$T_t^i(y) = \nu^i + \sum_{j=1}^d \lambda_{ji} \int_0^t T_s^j(y) ds - \frac{1}{2} \int_0^t h(a_i)^2 T_s^i(y) ds + \int_0^t h(a_i) T_s^i(y) \dot{y}_s ds.$$

Calculating the time derivative of $\tau_t^i(y) = e^{-h(a_i)y_t} T_t^i(y)$, the desired result follows for $y \in C^1$. To generalize the result to arbitrary $y \in C_0(\mathbb{R}_+, \mathbb{R})$, recall that continuously differentiable paths are dense in the set of all continuous paths. Hence we can approximate y by a sequence $y^n \in C^1$, then apply Proposition 1.4.4. \square

1.4.3. Pathwise diffusion filtering. We finally consider the diffusion case, with the signal and observations as defined in section 1.3. To develop the pathwise theory, let us investigate the pathwise filtering functional (1.12).

PROPOSITION 1.4.6. *Assume that $f \in C^2$ has compact support, $h \in C^2$ is Lipschitz continuous, and a is bounded. Define $h^y(t, x) = y_t^* h(x)$, $g^y(t, x) = \nabla h^y(t, x)$,*

$$V^y(t, x) = \frac{1}{2} \|a(x)^* g^y(t, x)\|^2 - \mathcal{L}h^y(t, x) - \frac{1}{2} \|h(x)\|^2,$$

and define the time-dependent generator

$$\mathcal{L}_t^y f(x) = \mathcal{L}f(x) - g^y(x, t)^* a(x) a(x)^* \nabla f(x).$$

Then the pathwise filtering functional $\tau_t(f, y)$ satisfies for any $y \in C_0(\mathbb{R}_+, \mathbb{R}^p)$

$$\frac{d}{dt} \tau_t(f, y) = \tau_t((\mathcal{L}_t^y + V^y(t, \cdot))f, y), \quad \tau_0(f, y) = \nu(f).$$

PROOF. Applying Itô's rule to $h(X_t)$ and rearranging, we can write

$$\begin{aligned} \tau_t(f, y) &= \mathbf{E} \left(f(X_t) \exp \left(\int_0^t V^y(s, X_s) ds \right) \right) \\ &\quad \times \exp \left(- \int_0^t g^y(s, X_s)^* a(X_s) dW_s - \frac{1}{2} \int_0^t \|a(X_s)^* g^y(s, X_s)\|^2 ds \right). \end{aligned}$$

Introduce a new measure $\tilde{\mathbf{P}}$ such that

$$\frac{d\tilde{\mathbf{P}}}{d\mathbf{P}} = \exp \left(- \int_0^t g^y(s, X_s)^* a(X_s) dW_s - \frac{1}{2} \int_0^t \|a(X_s)^* g^y(s, X_s)\|^2 ds \right).$$

By our assumptions, g^y and a are bounded. Hence we find that

$$\tau_t(f, y) = \tilde{\mathbf{E}} \left(f(X_t) \exp \left(\int_0^t V^y(s, X_s) ds \right) \right) = \mathbf{E} \left(f(\tilde{X}_t) \exp \left(\int_0^t V^y(s, \tilde{X}_s) ds \right) \right)$$

by Girsanov's theorem and Novikov's condition, where

$$d\tilde{X}_t = b(\tilde{X}_t) dt - a(\tilde{X}_t) a(\tilde{X}_t)^* g^y(t, \tilde{X}_t) dt + a(\tilde{X}_t) dW_t, \quad \tilde{X}_0 = X_0.$$

Using Itô's rule, we find that

$$\begin{aligned} f(\tilde{X}_t) \exp \left(\int_0^t V^y(s, \tilde{X}_s) ds \right) &= f(X_0) + \int_0^t e^{\int_0^s V^y(u, \tilde{X}_u) du} (\nabla f(\tilde{X}_s))^* a(\tilde{X}_s) dW_s \\ &\quad + \int_0^t (\mathcal{L}_s^y + V^y(s, \tilde{X}_s)) f(\tilde{X}_s) \exp \left(\int_0^s V^y(u, \tilde{X}_u) du \right) ds. \end{aligned}$$

As f has compact support, all the integrands are bounded. Hence we can take the expectation and apply Fubini's theorem to the time integral. \square

Now suppose that we can establish that the measure $\tau_t(f, y)$ has a $C^{1,2}$ density $H^y(t, x)$ with respect to the Lebesgue measure, i.e., that

$$\tau_t(f, y) = \int f(x) H^y(t, x) dx, \quad H^y(t, x) \in C^{1,2} \forall y \in C_0(\mathbb{R}_+, \mathbb{R}^p).$$

Then it follows from the pathwise filtering equation that $H^y(t, x)$ must satisfy the parabolic equation

$$\frac{\partial}{\partial t} H^y(t, x) = (\mathcal{L}_t^y)^* H^y(t, x) + V^y(t, x) H^y(t, x).$$

The question of uniqueness of the solutions for this type of equation is a standard one; see, e.g., [Fri75, sec. 6.4]. If the equation admits a unique solution, then we have obtained a filtering equation for diffusions in recursive form (though in practice, of course, this equation suffers from the curse of dimensionality). We can then compute any filtered estimate as

$$\pi_t(f) = \frac{\int f(x) e^{h(x) \cdot Y_t} H^{Y_{[0,t]}}(t, x) dx}{\int e^{h(x) \cdot Y_t} H^{Y_{[0,t]}}(t, x) dx} \quad \text{a.s.}$$

This highlights the importance of the pathwise filtering theory in the treatment of filtering problems for diffusions, as the filtering problem then reduces to a problem in the relatively tractable theory of parabolic equations. For explicit conditions under which this PDE approach is successful, see, e.g., [Par82, Kun90].

Finally, we remark that continuity of the pathwise filter with respect to the observation paths can be established also in the diffusion case, under suitable conditions. For example, we can repeat the proof of local Lipschitz continuity (see Proposition 1.4.4) if b , a , h , and their derivatives are assumed to be bounded. For a study of continuity of the pathwise filters in a general setup, see [CC05].

Filter Stability and Conditional Signals: Finite State Space

2.1. Conditional finite-state signals: A direct approach

In chapter 1, we showed how to calculate the conditional expectation of any functional of the signal process through the Kallianpur-Striebel formula. In particular, we defined uniquely a regular version $\Pi_T(X, y)$ of the conditional expectation, equation (1.11): for every $y \in C_0([0, T]; \mathbb{R}^p)$

$$\Pi_T(X, y) = \frac{\int X(x) \tilde{Z}_T(x, y) \mu_x(dx)}{\int \tilde{Z}_T(x, y) \mu_x(dx)},$$

so that $\mathbf{E}(X | \mathcal{F}_T^Y) = \Pi_T(X, Y_{[0, T]})$ a.s. For fixed y and T , we can thus consider $\Pi_T(X, y)$ as a measure on the space $(D(\mathbb{R}_+; \mathbb{S}), \mathcal{D})$ of signal sample paths (with its Borel σ -algebra \mathcal{D}) by setting $\Pi_T(A, y) = \Pi_T(I_A, y)$ for every $A \in \mathcal{D}$. Evidently

$$\frac{d\Pi_T(\cdot, y)}{d\mu_x} = \frac{\tilde{Z}_T(x, y)}{\int \tilde{Z}_T(x, y) \mu_x(dx)}.$$

The idea which we will explore in this chapter is that this change of measure can be interpreted as a Girsanov-type transformation: in particular, if the signal process X_t is a Feller-Markov process, then X_t is still an (albeit time-nonhomogeneous) Markov process under $\Pi_T(\cdot, y)$ for fixed y and T , and we can explicitly find its *conditional generator* $\mathcal{L}_t^{T, y}$. We can now obtain the conditional expectation of any functional of the signal process X_t (with generator \mathcal{L}), given the observation path $y \in C_0([0, T]; \mathbb{R}^p)$, by calculating the (unconditional) expectation of a new Markov process $X_t^{T, y}$ with the modified generator $\mathcal{L}_t^{T, y}$.

R. L. Stratonovich initiated the investigation of conditional Markov processes [Str60, Str68], at least on a formal level. The study of conditional diffusions in the spirit of this chapter has its origins in the paper of J.-M. Bismut and D. Michel [BM82], where the diffusion case is treated using stochastic flow methods. A different perspective appears in the paper of S. K. Mitter and N. J. Newton [MN03], also in the context of diffusions, where $X_t^{T, y}$ is characterized as the solution of an optimal control problem. Both these treatments are restricted to diffusion processes and use the properties of stochastic flows. In this chapter, we develop a parallel theory for finite-state Markov signals. In fact, it appears that the theory can be extended to general Markov signals, and we will briefly comment on this case as well. In this section we pursue a direct approach to the theory; section 2.2 explores

the stochastic control perspective. Finally, section 2.3 applies the theory to obtain results on filter stability for the Wonham filter.

2.1.1. Change of measure for finite-state signals. In the following we will be interested in the Radon-Nikodym derivative between the laws of two Markov processes on a finite state space. There is a simple analog of the Girsanov theory for this case, see [RW00, sec. IV.22]. Let us briefly develop and extend some results along these lines that are needed in the sequel.

We work in the following context. The signal state space is $\mathbb{S} = \{a_1, \dots, a_d\}$, $d < \infty$. We let X_t be the coordinate process on $D([0, T]; \mathbb{S})$ for some $T < \infty$, and let μ_x be the measure on $D([0, T]; \mathbb{S})$ that makes X_t a Markov process with transition intensities matrix $\Lambda = (\lambda_{ij})$ and initial density ν . Let μ'_x be the law under which X_t is a time-nonhomogeneous Markov process with intensities matrix $\Lambda'(t) = (\lambda'_{ij}(t))$, which we assume to be bounded, and with the same initial distribution ν . We also assume $\lambda'_{ij}(t) > 0$ iff $\lambda_{ij} > 0$; if this is not the case, then $\mu'_x \not\sim \mu_x$! We now have the following result (cf. [RW00, sec. IV.22] for a more general setup).

PROPOSITION 2.1.1. *Assume $\lambda'_{ij}(t) > 0$ iff $\lambda_{ij} > 0$. Define*

$$\gamma_{ij}(t) = \begin{cases} \lambda'_{ij}(t)/\lambda_{ij} & i \neq j, \lambda_{ij} > 0, \\ 1 & \text{otherwise.} \end{cases}$$

Then $d\mu'_x/d\mu_x = \Upsilon_T$, where Υ_t is the martingale (under μ_x) defined by

$$\Upsilon_t = \exp\left(-\sum_{i,j=1}^d \int_0^t \lambda_{ij} \gamma_{ij}(s) I_{X_{s-}=a_i} ds\right) \prod_{0 < s \leq t} \sum_{i \neq j} \gamma_{ij}(s^-) I_{X_{s-}=a_i} I_{X_s=a_j}.$$

PROOF. Let us begin by showing that Υ_t is martingale. Clearly $\Upsilon_t > 0$ a.s., so we can define a process N_t by $dN_t = \Upsilon_{t-}^{-1} d\Upsilon_t$, $N_0 = 0$. Using Itô's rule, we find

$$\begin{aligned} N_t &= 1 + \sum_{0 < s \leq t} \sum_{i \neq j} (\gamma_{ij}(s^-) - 1) I_{X_{s-}=a_i} I_{X_s=a_j} - \int_0^t \sum_{i,j=1}^d \lambda_{ij} \gamma_{ij}(s) I_{X_{s-}=a_i} ds \\ &= 1 + \sum_{i \neq j} \int_{0+}^t (\gamma_{ij}(s^-) - 1) I_{X_{s-}=a_i} (dI_{X_s=a_j} - \lambda_{ij} ds) \\ &= 1 + \sum_{i \neq j} \int_{0+}^t (\gamma_{ij}(s^-) - 1) I_{X_{s-}=a_i} \left(dI_{X_s=a_j} - \sum_{k=1}^d \lambda_{kj} I_{X_{s-}=a_k} ds \right). \end{aligned}$$

But the integrator in the last integral is a bounded martingale (by Dynkin's formula) and the integrand is bounded. Hence N_t is a martingale, and consequently Υ_t (being the Doléans exponential of N_t) is a martingale as well. It is now evident that

$$\mu'_x(X_0 = a_i) = \mathbf{E}_{\mu_x}(\Upsilon_T I_{X_0=a_i}) = \mathbf{E}_{\mu_x}(\mathbf{E}_{\mu_x}(\Upsilon_T | \sigma\{X_0\}) I_{X_0=a_i}) = \mu_x(X_0 = a_i).$$

To show that X_t is a Markov process with generator $\Lambda'(t)$ under μ'_x , it suffices to show that Dynkin's formula holds: i.e., we will show that

$$M'_t(i) = I_{X_t=a_i} - I_{X_0=a_i} - \sum_{j=1}^d \int_0^t \lambda'_{ji}(s) I_{X_{s-}=a_j} ds$$

is a martingale under μ'_x for any i , i.e., we need to show that $M'_t(i)\Upsilon_t$ is a martingale under μ_x (by the Bayes formula, Lemma 1.1.1). Using the Itô rule, we have

$$M'_t(i)\Upsilon_t = \int_{0+}^t M'_{s-}(i) d\Upsilon_s + \int_{0+}^t \Upsilon_{s-} dM'_{s-}(i) + [M'(i), \Upsilon]_t.$$

As Υ_t and $M'_t(i)$ are finite variation processes,

$$\begin{aligned} [M'(i), \Upsilon]_t &= \sum_{0 < s \leq t} (M'_s(i) - M'_{s-}(i))(\Upsilon_s - \Upsilon_{s-}) \\ &= \sum_{0 < s \leq t} \sum_{j \neq k} \Upsilon_{s-} (\gamma_{kj}(s^-) - 1) I_{X_{s-}=a_k} I_{X_s=a_j} (I_{X_s=a_i} - I_{X_{s-}=a_i}) \\ &= \sum_{0 < s \leq t} \sum_{j \neq k} \Upsilon_{s-} (\gamma_{kj}(s^-) - 1) (\delta_{ji} - \delta_{ki}) I_{X_{s-}=a_k} I_{X_s=a_j} \\ &= \sum_{j \neq k} \int_{0+}^t \Upsilon_{s-} (\gamma_{kj}(s^-) - 1) (\delta_{ji} - \delta_{ki}) I_{X_{s-}=a_k} dI_{X_s=a_j}. \end{aligned}$$

But as in the proof that N_t is a martingale, we find that

$$N'_t = \sum_{j \neq k} \int_{0+}^t \Upsilon_{s-} (\gamma_{kj}(s^-) - 1) (\delta_{ji} - \delta_{ki}) I_{X_{s-}=a_k} (dI_{X_s=a_j} - \lambda_{kj} ds)$$

is a martingale under μ_x . We can now write

$$\begin{aligned} [M'(i), \Upsilon]_t &= N'_t + \sum_{j \neq k} \int_{0+}^t \Upsilon_s (\lambda'_{kj}(s) - \lambda_{kj}) (\delta_{ji} - \delta_{ki}) I_{X_{s-}=a_k} ds \\ &= N'_t + \sum_{j=1}^d \int_{0+}^t \Upsilon_s (\lambda'_{ji}(s) - \lambda_{ji}) I_{X_{s-}=a_j} ds. \end{aligned}$$

But then we obtain

$$M'_t(i)\Upsilon_t = \int_{0+}^t M'_{s-}(i) d\Upsilon_s + N'_t + \int_{0+}^t \Upsilon_{s-} \left(dI_{X_s=a_i} - \sum_{j=1}^d \lambda_{ji} I_{X_{s-}=a_j} ds \right).$$

The first two terms are evidently martingales under μ_x , whereas the last term is a martingale by Dynkin's formula. Hence the proof is complete. \square

COROLLARY 2.1.2. *Suppose that there is a time-dependent vector $v(t)$ such that*

$$\lambda'_{ij}(t) = \lambda_{ij} \frac{v^j(t)}{v^i(t)}, \quad i \neq j.$$

Then we can write

$$\frac{d\mu'_x}{d\mu_x} = \exp \left(\sum_{i=1}^d \left[\int_{0+}^T \log v^i(s^-) dI_{X_s=a_i} - \int_0^T \frac{(\Lambda v)^i(s)}{v^i(s)} I_{X_s=a_i} ds \right] \right).$$

PROOF. Apply the following manipulations:

$$\begin{aligned} \prod_{0 < s \leq t} \sum_{i \neq j} \gamma_{ij}(s^-) I_{X_{s^-} = a_i} I_{X_s = a_j} &= \exp \left(\sum_{0 < s \leq T} \sum_{i, j=1}^d \log \gamma_{ij}(s^-) I_{X_{s^-} = a_i} I_{X_s = a_j} \right) \\ &= \exp \left(\sum_{0 < s \leq T} \sum_{i=1}^d \log v^i(s^-) \Delta I_{X_s = a_i} \right) = \exp \left(\sum_{i=1}^d \int_{0+}^T \log v^i(s^-) dI_{X_s = a_i} \right). \end{aligned}$$

The statement now follows immediately. \square

Let us now introduce yet another measure $\tilde{\mu}_x$, under which X_t has the same transition intensities matrix Λ as under μ_x , but a different initial distribution $\tilde{\nu}$. Then we have the following result.

PROPOSITION 2.1.3. *Suppose that $\tilde{\nu} \ll \nu$. Then we have*

$$\frac{d\tilde{\mu}_x}{d\mu_x} = \frac{d\tilde{\nu}}{d\nu}(X_0) = \sum_{i=1}^d \frac{\tilde{\nu}^i}{\nu^i} I_{X_0 = a_i}.$$

PROOF. We first check that the initial law is correct:

$$\tilde{\mu}_x(X_0 = a_i) = E_{\mu_x} \left(\frac{\tilde{\nu}^i}{\nu^i} I_{X_0 = a_i} \right) = \tilde{\nu}^i.$$

Hence it only remains to verify that X_t has the same transition intensities, i.e., we should check Dynkin's formula under $\tilde{\mu}_x$: i.e.,

$$M_t^i = I_{X_t = a_i} - I_{X_0 = a_i} - \sum_{j=1}^d \lambda_{ji} \int_0^t I_{X_s = a_j} ds$$

should be a martingale under $\tilde{\mu}_x$ for any i . But

$$\mathbf{E}_{\tilde{\mu}_x}(M_t^i | \mathcal{F}_s) = \frac{\mathbf{E}_{\mu_x} \left(\frac{d\tilde{\nu}}{d\nu}(X_0) M_t^i | \mathcal{F}_s \right)}{\mathbf{E}_{\mu_x} \left(\frac{d\tilde{\nu}}{d\nu}(X_0) | \mathcal{F}_s \right)} = \mathbf{E}_{\mu_x}(M_t^i | \mathcal{F}_s) = M_s^i,$$

where we have used the Bayes formula. The proof is complete. \square

2.1.2. Conditional signal—forward case. We begin by introducing a useful device in the treatment of smoothing problems—the dual equation [Par82].

DEFINITION 2.1.4. The (pathwise) dual filter $v_{t,T}^i(y)$ is defined by

$$\begin{aligned} v_{t,T}^i(y) &= \mathbf{E} \left(\frac{\tilde{Z}_T(X_{[0,T]}, y)}{\tilde{Z}_t(X_{[0,t]}, y)} \middle| X_t = a_i \right) \\ &= e^{-h(a_i)y_t} \mathbf{E} \left[\exp \left(h(X_T) y_T - \int_{t+}^T y_s dh(X_s) - \frac{1}{2} \int_t^T h(X_s)^2 ds \right) \middle| X_t = a_i \right]. \end{aligned}$$

LEMMA 2.1.5. $\tilde{v}_{t,T}^i(y) = e^{h(a_i)y_t} v_{t,T}^i(y)$ solves the pathwise dual Wonham filter:

$$\frac{d\tilde{v}_{t,T}^i(y)}{dt} = - \sum_{j=1}^d \lambda_{ij} \tilde{v}_{t,T}^j(y) e^{(h(a_i) - h(a_j))y_t} + \frac{1}{2} h(a_i)^2 \tilde{v}_{t,T}^i(y), \quad \tilde{v}_{T,T}^i(y) = e^{h(a_i)y_T}.$$

PROOF. We proceed as in the proof of Proposition 1.4.5. For $y \in C^1$,

$$\tilde{v}_{t,T}^i(y) = e^{h(a_i)y_t} \mathbf{E} \left[\exp \left(\int_t^T h(X_s) \dot{y}_s ds - \frac{1}{2} \int_t^T h(X_s)^2 ds \right) \middle| X_t = a_i \right].$$

We need a time-reversed form of Dynkin's formula [Eil86]:

$$I_{X_{t-}=a_i} = I_{X_{T-}=a_i} + \sum_{j=1}^d \int_t^T \left[\frac{p_{X_s=a_i}}{p_{X_s=a_j}} \lambda_{ij} I_{X_s=a_j} - \frac{p_{X_s=a_j}}{p_{X_s=a_i}} \lambda_{ji} I_{X_s=a_i} \right] ds + \overleftarrow{M}_t,$$

where $p_{X_t=a_i} = \mathbf{P}(X_t = a_i)$ and \overleftarrow{M}_t is an $\overleftarrow{\mathcal{F}}_t^X = \sigma\{X_{s-} : t \leq s \leq T\}$ -adapted backward càglàd martingale. We may presume that $p_{X_0=a_i} > 0$ for all i , as we are ultimately going to condition on X_0 ; hence $p_{X_t=a_i} > 0$ for all t and i , and the equation above is well defined. We can now calculate the quantity

$$R_t^i = I_{X_{t-}=a_i} \exp \left(\int_t^T h(X_s) \dot{y}_s ds - \frac{1}{2} \int_t^T h(X_s)^2 ds \right)$$

using the Itô calculus backwards in time. There is no problem in doing this, as we can just work within the usual framework with the time-reversed filtration $\overleftarrow{\mathcal{F}}_t^X$, provided we take care to make the integrators càglàd and the integrands càdlàg (i.e., precisely reversed compared to the time-forward theory). This gives

$$\begin{aligned} R_t^i &= I_{X_{T-}=a_i} + \int_t^T h(a_i) R_s^i \dot{y}_s ds - \frac{1}{2} \int_t^T h(a_i)^2 R_s^i ds \\ &\quad + \sum_{j=1}^d \int_t^T \left[\frac{p_{X_s=a_i}}{p_{X_s=a_j}} \lambda_{ij} R_s^j - \frac{p_{X_s=a_j}}{p_{X_s=a_i}} \lambda_{ji} R_s^i \right] ds + \overleftarrow{M}'_t, \end{aligned}$$

where \overleftarrow{M}'_t is another backward martingale. Now note that

$$v_{t,T}^i(y) = \frac{\mathbf{E}(R_t^i)}{\mathbf{P}(X_t = a_i)}, \quad \tilde{v}_{t,T}^i(y) = e^{h(a_i)y_t} v_{t,T}^i(y).$$

Using the ordinary chain rule, we easily find

$$\frac{d}{dt} v_{t,T}^i(y) = - \sum_{j=1}^d \lambda_{ij} v_{t,T}^j(y) + \frac{1}{2} h(a_i)^2 v_{t,T}^i(y) - h(a_i) v_{t,T}^i(y) \dot{y}_t.$$

It remains to use the chain rule once more, and the result is established for $y \in C^1$. To generalize to arbitrary $y \in C_0([0, T], \mathbb{R})$, we can proceed exactly as in Proposition 1.4.5 by taking limits of smooth approximations. The required continuity property follows in an identical manner to the proof of Proposition 1.4.4. \square

We now arrive at our first characterization of the conditional signal.

THEOREM 2.1.6. *Under $\Pi_T(\cdot, y)$, $X_{[0,T]}$ is a time-nonhomogeneous Markov process with conditional intensities matrix and initial distribution given by*

$$\lambda_{ij}^{T,y}(t) = \lambda_{ij} \frac{v_{t,T}^j(y)}{v_{t,T}^i(y)} \quad (i \neq j), \quad \tilde{v}^i = \Pi_T(X_0 = a_i, y).$$

PROOF. Clearly $\log v_{t,T}^i(y) = \log \tilde{v}_{t,T}^i(y) - h(a_i) y_t$. By Lemma 2.1.5,

$$\frac{d}{dt} \log \tilde{v}_{t,T}^i(y) = -\frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} + \frac{1}{2} h(a_i)^2.$$

Hence we can calculate

$$\begin{aligned} & \int_{0+}^T \log \tilde{v}_{t,T}^i(y) dI_{X_t=a_i} - h(a_i) I_{X_T=a_i} y_T + I_{X_0=a_i} \log v_{0,T}^i(y) = \\ & - \int_0^T I_{X_t=a_i} d \log \tilde{v}_{t,T}^i(y) = \int_0^T \frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} I_{X_t=a_i} dt - \frac{1}{2} \int_0^T h(a_i)^2 I_{X_t=a_i} dt. \end{aligned}$$

Summing over i and rearranging, we obtain

$$\begin{aligned} & \sum_{i=1}^d \left[\int_{0+}^T \log \tilde{v}_{t,T}^i(y) dI_{X_t=a_i} - \int_0^T \frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} I_{X_t=a_i} dt \right] \\ & = h(X_T) y_T - \frac{1}{2} \int_0^T h(X_t)^2 dt - \sum_{i=1}^d I_{X_0=a_i} \log v_{0,T}^i(y). \end{aligned}$$

But note also that

$$\sum_{i=1}^d \int_{0+}^T h(a_i) y_t dI_{X_t=a_i} = \int_{0+}^T y_t dh(X_t).$$

Hence we obtain

$$\begin{aligned} & \sum_{i=1}^d \left[\int_{0+}^T \log v_{t,T}^i(y) dI_{X_t=a_i} - \int_0^T \frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} I_{X_t=a_i} dt \right] \\ & = h(X_T) y_T - \int_{0+}^T y_t dh(X_t) - \frac{1}{2} \int_0^T h(X_t)^2 dt - \sum_{i=1}^d I_{X_0=a_i} \log v_{0,T}^i(y). \end{aligned}$$

Exponentiating and substituting our standard notation, we obtain

$$e^{\sum_{i=1}^d [\int_{0+}^T \log v_{t,T}^i(y) dI_{X_t=a_i} - \int_0^T \frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} I_{X_t=a_i} dt]} = \frac{\tilde{Z}_T(X_{[0,T]}, y)}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) | \sigma\{X_0\})}.$$

Next, consider the following manipulations:

$$\begin{aligned} & \frac{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) | \sigma\{X_0\})}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))} = \sum_{i=1}^d I_{X_0=a_i} \frac{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) | X_0 = a_i)}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))} \\ & = \sum_{i=1}^d I_{X_0=a_i} \frac{\mathbf{E}(I_{X_0=a_i} \tilde{Z}_T(X_{[0,T]}, y))}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y)) \mathbf{P}(X_0 = a_i)} = \sum_{i=1}^d I_{X_0=a_i} \frac{\Pi_T(X_0 = a_i, y)}{\mathbf{P}(X_0 = a_i)}. \end{aligned}$$

Putting together all of the above, and using the pathwise Kallianpur-Striebel formula, we find that

$$\begin{aligned} & \frac{d\Pi_T(\cdot, y)}{d\mu_x} = \frac{\tilde{Z}_T(x, y)}{\int \tilde{Z}_T(x, y) \mu_x(dx)} = \\ & e^{\sum_{i=1}^d [\int_{0+}^T \log v_{t,T}^i(y) dI_{X_t=a_i} - \int_0^T \frac{(\Lambda v_{t,T}(y))^i}{v_{t,T}^i(y)} I_{X_t=a_i} dt]} \sum_{i=1}^d I_{X_0=a_i} \frac{\Pi_T(X_0 = a_i, y)}{\mathbf{P}(X_0 = a_i)}. \end{aligned}$$

The statement now follows from Proposition 2.1.3 and Corollary 2.1.2. \square

REMARK 2.1.7. One is not restricted to the interval $[0, T]$: it is easy to show that Theorem 2.1.6 extends to an arbitrary finite time interval by setting $\lambda_{ij}^{T,y}(t) = \lambda_{ij}$ for $t > T$. Indeed, all we have to verify is that

$$M_t^i = I_{X_t=a_i} - I_{X_0=a_i} - \sum_{j=1}^d \int_0^t \lambda_{ji}^{T,y}(s) I_{X_s=a_j} ds$$

is a martingale under $\Pi_T(\cdot, y)$ even for $t > T$, which follows immediately from the Bayes formula and the fact that $d\Pi_T(\cdot, y)/d\mu_x$ is \mathcal{F}_T -measurable.

2.1.3. Conditional signal—time-reversed case. It is sometimes convenient to work with Theorem 2.1.6 in reverse time. This gives another characterization of the conditional signal.

THEOREM 2.1.8. *Define the càdlàg reverse time signal process \bar{X}_t such that $\bar{X}_{t-} = X_{T-t}$. Under $\Pi_T(\cdot, y)$, $\bar{X}_{[0,T]}$ is a time-nonhomogeneous Markov process with conditional intensities matrix and initial distribution given by*

$$\bar{\lambda}_{ij}^{T,y}(t) = \lambda_{ji} \frac{\Pi_{T-t}(X_{T-t} = a_j, y)}{\Pi_{T-t}(X_{T-t} = a_i, y)} \quad (i \neq j), \quad \bar{\nu}^i = \Pi_T(X_T = a_i, y).$$

PROOF. By standard time reversal arguments, see, e.g., [Eil86], \bar{X}_t is a Markov process with the transition intensities

$$\bar{\lambda}_{ij}^{T,y}(t) = \lambda_{ji}^{T,y}(T-t) \frac{\Pi_T(X_{T-t} = a_j, y)}{\Pi_T(X_{T-t} = a_i, y)} \quad (i \neq j).$$

Hence the result follows if we can establish that for any j

$$\frac{\Pi_T(X_{T-t} = a_j, y)}{v_{T-t,T}^j(y)} = K \Pi_{T-t}(X_{T-t} = a_j, y),$$

where K does not depend on j . This show this, write

$$\begin{aligned} \mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) I_{X_{T-t}=a_j}) &= \mathbf{E}(\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) | \mathcal{F}_{T-t}) I_{X_{T-t}=a_j}) \\ &= \mathbf{E} \left[\mathbf{E} \left(\frac{\tilde{Z}_T(X_{[0,T]}, y)}{\tilde{Z}_{T-t}(X_{[0,T-t]}, y)} \Bigg| \mathcal{F}_{T-t} \right) \tilde{Z}_{T-t}(X_{[0,T-t]}, y) I_{X_{T-t}=a_j} \right] \\ &= \mathbf{E} \left[\mathbf{E} \left(\frac{\tilde{Z}_T(X_{[0,T]}, y)}{\tilde{Z}_{T-t}(X_{[0,T-t]}, y)} \Bigg| X_{T-t} = a_j \right) \tilde{Z}_{T-t}(X_{[0,T-t]}, y) I_{X_{T-t}=a_j} \right] \\ &= v_{T-t,T}^j(y) \mathbf{E}(\tilde{Z}_{T-t}(X_{[0,T-t]}, y) I_{X_{T-t}=a_j}). \end{aligned}$$

Note that this is precisely the characterization of [Par82] of the unnormalized smoothing density as a product of the unnormalized filter and the dual filter. Dividing both sides by $\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))$, we obtain

$$\frac{\Pi_T(X_{T-t} = a_j, y)}{v_{T-t,T}^j(y)} = \frac{\mathbf{E}(\tilde{Z}_{T-t}(X_{[0,T-t]}, y))}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))} \Pi_{T-t}(X_{T-t} = a_j, y),$$

which is the desired result. The proof is now complete. \square

2.1.4. Smoothing and path estimation. The characterization of the conditional law of the signal process in terms of the conditional signals (in forward or reverse time) allows us to sample signal sample paths according to the conditional measure using a simple Monte Carlo approach. For example, one could proceed as follows. Given an observed sample path y , one first computes $\Pi_t(X_t = a_i, y)$ on the interval $[0, T]$ using, e.g., the pathwise Wonham filter, Proposition 1.4.5. It is then straightforward to sample paths of X_t which are distributed according to $\Pi_T(\cdot, y)$ using Theorem 2.1.8. If X_0 is deterministic, one could instead calculate $\tilde{v}_{t,T}^j(y)$ on the interval $[0, T]$ using the pathwise dual Wonham filter (Lemma 2.1.5), then apply Theorem 2.1.6. If X_0 is not deterministic, Theorem 2.1.6 is not as convenient as we would have to somehow calculate $\Pi_T(X_0 = a_i, y)$. For a different approach to sampling the conditional measure, see [SVW04, HSV06].

The path estimation problem becomes a smoothing problem when we restrict to estimates at a fixed time $t < T$. In this context, the smoothing equations are precisely the Kolmogorov forward (or backward) equations associated to the conditional signals of Theorems 2.1.6 and 2.1.8. In the latter case, we obtain

$$\frac{d}{dt} \Pi_T(\bar{X}_t = a_i, y) = \sum_{j=1}^d \left[\lambda_{ij} \frac{\pi_{T-t}^i}{\pi_{T-t}^j} \Pi_T(\bar{X}_t = a_j, y) - \lambda_{ji} \frac{\pi_{T-t}^j}{\pi_{T-t}^i} \Pi_T(\bar{X}_t = a_i, y) \right],$$

where we have written $\pi_t^i = \Pi_t(X_t = a_i, y)$. Reversing time gives

$$-\frac{d}{dt} \Pi_T(X_t = a_i, y) = \sum_{j=1}^d \left[\lambda_{ij} \frac{\pi_t^i}{\pi_t^j} \Pi_T(X_t = a_j, y) - \lambda_{ji} \frac{\pi_t^j}{\pi_t^i} \Pi_T(X_t = a_i, y) \right].$$

This is a well-known smoothing equation: see [LS01, Thm. 9.5]. The conditional signal of Theorem 2.1.6, on the other hand, gives rise to the smoothing equation

$$\frac{d}{dt} \Pi_T(X_t = a_i, y) = \sum_{j=1}^d \left[\lambda_{ji} \frac{v_{t,T}^i(y)}{v_{t,T}^j(y)} \Pi_T(X_t = a_j, y) - \lambda_{ij} \frac{v_{t,T}^j(y)}{v_{t,T}^i(y)} \Pi_T(X_t = a_i, y) \right]$$

which does not appear to have been considered previously (to my knowledge).

REMARK 2.1.9. If we are only interested in smoothing, both smoothing equations can be obtained in a more direct manner. Recall that

$$\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) I_{X_t=a_i}) = v_{t,T}^i(y) \mathbf{E}(\tilde{Z}_t(X_{[0,t]}, y) I_{X_t=a_i}),$$

as established in the proof of Theorem 2.1.8. In terms of the pathwise filter $\tau_t^i(y)$ (Proposition 1.4.5) and dual filter $\tilde{v}_{t,T}^i(y)$ (Lemma 2.1.5), this can be written as

$$\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y) I_{X_t=a_i}) = \tilde{v}_{t,T}^i(y) \tau_t^i(y),$$

and in particular we can write

$$\Pi_T(X_t = a_i, y) = \frac{\tilde{v}_{t,T}^i(y) \tau_t^i(y)}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))}.$$

Using Proposition 1.4.5, Lemma 2.1.5, and the ordinary chain rule, we easily find

$$\frac{d}{dt} \Pi_T(X_t = a_i, y) = \sum_{j=1}^d \left[\frac{\lambda_{ji} v_{t,T}^i(y) \sigma_t^j(y) - \lambda_{ij} v_{t,T}^j(y) \sigma_t^i(y)}{\mathbf{E}(\tilde{Z}_T(X_{[0,T]}, y))} \right],$$

where we have written $\sigma_t^i(y) = e^{h(a_i)y_t} \tau_t^i(y)$. But recall that $\pi_t^i(y) = \sigma_t^i(y)/|\sigma_t(y)|$, so that both smoothing equations above follow directly. Theorems 2.1.6 and 2.1.8

provide a much stronger statement, however, as they characterize the conditional statistics of entire sample paths of the signal. They also elucidate the reason behind the conspicuous fact that the smoothing equations have precisely the form of Kolmogorov forward equations. It is not difficult, on the other hand, to establish directly that the conditional signal is itself a Markov process (see, e.g., Lemma 2.2.3 below), so that the transition intensities of the conditional signal are in some sense implicit when the smoothing equations are written in the above form.

2.1.5. A note on general Markov signals. The theory developed in this chapter for finite-state signals complements the existing theory for diffusion signals [BM82, MN03]. One could wonder whether these ideas are universal, in the sense that the theory can be developed for an arbitrary Markov signal process under some mild regularity conditions (thus developing and significantly extending in scope the program initiated by R. L. Stratonovich [Str60, Str68]). In fact, all that is needed is in essence the Kallianpur-Striebel formula (which holds under exceedingly general circumstances) and a change of measure result in the spirit of Corollary 2.1.2.

It turns out that the latter change of measure technique is quite natural within the context of the general theory of Markov processes; the time-homogeneous case ($v(t)$ is independent of time) is investigated in a recent paper by Z. Palmowski and T. Rolski [PR02], and the extension to the time-nonhomogeneous case does not appear to be problematic. It would thus appear that the generalization of the theory in this section to general Markov signals is mainly of a technical nature (e.g., one has to show that a generalized analog of $v_{t,T}(y)$ can be defined in such a way that it is in the domain of the generator \mathcal{L} of the signal, etc.) Similarly, the stochastic control approach explored in the next section should extend naturally to the general Markov case (in this context, see also [She85]). The details required for the completion of this program is an interesting topic for further investigation.

2.2. On the duality between estimation and stochastic control

In the previous section we characterized the conditional law of the signal process directly using only the Kallianpur-Striebel formula, a Girsanov-type result and some stochastic calculus. In this section we will take a rather different point of view, following S. K. Mitter and N. J. Newton [MN03]. The starting point is a variational formulation of the Kallianpur-Striebel formula, which characterizes the conditional law as the minimizer of a certain information-theoretic cost function. This minimization is subsequently expressed as a stochastic optimal control problem, so that the law of the optimally controlled signal process coincides with the conditional law of the signal process. The optimally controlled signal is then precisely the conditional signal process which we encountered in the previous section.

The striking similarities between linear filtering and control were noticed already by Kalman [Kal60] (see [KSH00, Ch. 15] for a modern point of view). In the nonlinear case, it was noticed by S. K. Mitter [Mit79, Mit81] in the context of diffusion signals that the PDE for the logarithm of the nonlinear filtering density coincides with the Hamilton-Jacobi-Bellman equation for the value function of a particular stochastic control problem. This observation was developed further by W. H. Fleming and S. K. Mitter in [FM83, Mit82]. The underlying reason for this phenomenon was not elucidated, however, until the appearance of [MN03]. We will see that similar things happen in the finite state case. The corresponding Bellman equations are somewhat difficult to recognize using the naked eye, but

nonetheless we will find that they emerge naturally. As before, the extension to general Markov signals is an interesting problem for further investigation.

We begin by recalling the variational formulation of the Kallianpur-Striebel formula from [MN03]; there is nothing new in this section, but the proof is short and illuminating. Specializing to a finite-state signal, we then develop the conditional signal theory in the forward case by converting the variational estimation problem into a dynamic programming problem. Finally, we develop also the reverse case. Here we depart significantly from the corresponding discussion in [MN03] for diffusions: we will show that the reverse case follows from a time-reversed version of the dynamic programming method used in the forward case. This seems more natural than the treatment of [MN03] using duality with a deterministic control problem, and emphasizes that the two conditional signals are simply time-reversed versions of each other. (It would be interesting to compare these approaches to the various dualities which are known in the linear case [KSH00, Ch. 15]).

2.2.1. The variational Kallianpur-Striebel formula. We follow [MN03], which treats both a much more general model and an additional variational formula. We will not need the latter, however, and we restrict to the finite-state signal and white noise observation model for notational convenience.

Let us begin by introducing the “energy function”

$$H_T(x, y) = -\log \tilde{Z}_T(x, y).$$

Then the Kallianpur-Striebel formula can be expressed as

$$\frac{d\Pi_T(\cdot, y)}{d\mu_x} = \frac{\exp(-H_T(x, y))}{\int \exp(-H_T(x, y)) \mu_x(dx)}.$$

In the language of statistical mechanics, we have written the conditional measure in the form of a Gibbs-type distribution. Let us now introduce the following notation. \mathcal{P}_x denotes the set of all probability measures on the space of signal sample paths $D(\mathbb{R}_+; \mathbb{S})$. For two measures $\mu, \tilde{\mu} \in \mathcal{P}_x$ we define the relative entropy

$$D(\mu || \tilde{\mu}) = \int \log \left(\frac{d\mu}{d\tilde{\mu}} \right) d\mu \quad \text{if } \mu \ll \tilde{\mu}, \quad +\infty \quad \text{otherwise.}$$

For any measurable function \tilde{H} on $D(\mathbb{R}_+; \mathbb{S})$, we define the total information

$$I(\tilde{H}) = -\log \left(\int \exp(-\tilde{H}) d\mu_x \right).$$

Finally, we define the expectation

$$\langle \tilde{H}, \mu \rangle = \int \tilde{H} d\mu \quad \text{if the integral is finite,} \quad +\infty \quad \text{otherwise.}$$

We now obtain the following fundamental result.

LEMMA 2.2.1 (Variational Kallianpur-Striebel formula). *The conditional law $\Pi_T(\cdot, y)$ is the unique minimizer in the variational expression*

$$I(H_T(\cdot, y)) = \min_{\tilde{\mu} \in \mathcal{P}_x} \{D(\tilde{\mu} || \mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle\}.$$

PROOF. It is not difficult to establish, along the lines of the proof of Proposition 1.4.4, that $I(H_T(\cdot, y))$ and $\langle H_T(\cdot, y), \Pi_T(\cdot, y) \rangle$ are finite quantities for any observation sample path $y \in C_0(\mathbb{R}_+, \mathbb{R})$. Hence we obtain

$$D(\Pi_T(\cdot, y) || \mu_x) = -\langle H_T(\cdot, y), \Pi_T(\cdot, y) \rangle + I(H_T(\cdot, y)).$$

It remains to show that

$$I(H_T(\cdot, y)) < D(\tilde{\mu}|\mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle$$

for any $\tilde{\mu} \in \mathcal{P}_x$ such that $\tilde{\mu} \neq \Pi_T(\cdot, y)$. If either of the terms on the right-hand side are infinite, then the statement is clearly true. Let us thus suppose that the right-hand side is finite. Note that we can split the relative entropy as

$$D(\tilde{\mu}|\mu_x) = \int \left[\log \left(\frac{d\tilde{\mu}}{d\Pi_T(\cdot, y)} \right) + \log \left(\frac{d\Pi_T(\cdot, y)}{d\mu_x} \right) \right] d\tilde{\mu}.$$

Using the Kallianpur-Striebel formula, we thus obtain

$$I(H_T(\cdot, y)) + D(\tilde{\mu}|\Pi_T(\cdot, y)) = D(\tilde{\mu}|\mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle.$$

But by Jensen's inequality $D(\tilde{\mu}|\Pi_T(\cdot, y)) \geq 0$ and $D(\tilde{\mu}|\Pi_T(\cdot, y))$ is strictly convex in $\tilde{\mu}$ (as $x \log x$ is a strictly convex function). Hence we have strict inequality $D(\tilde{\mu}|\Pi_T(\cdot, y)) > 0$ for $\tilde{\mu} \neq \Pi_T(\cdot, y)$, and the result follows. \square

REMARK 2.2.2. Mitter and Newton attach an information-theoretic interpretation to this result. They interpret $I(H_T(\cdot, y))$ as the total information available to the estimator through the sample path y . On the other hand, they call the quantity $F(\tilde{\mu}, y) = D(\tilde{\mu}|\mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle$ the ‘‘apparent information’’ of the estimator. In this sense, a suboptimal estimator appears to have access to more information than is actually available. Regardless of the information-theoretic interpretation, it is evident from the statement of the Lemma that the optimal estimator must find a balance between being close to the prior law μ_x (the relative entropy term) and being matched to the observed data (the expected log-likelihood term). We will find a similar structure in the associated stochastic control problem.

2.2.2. Dynamic programming—forward case. We are going to express the minimization of Lemma 2.2.1 as an optimal control problem. It is convenient to restrict ourselves to minimizing over measures $\tilde{\mu}$ under which X_t is a Markov process, as we can then use Proposition 2.1.1 to characterize any such measure that is equivalent to μ_x (and clearly $\Pi_T(\cdot, y)$ must be equivalent to μ_x). Let us thus begin by showing that X_t is Markov under $\Pi_T(\cdot, y)$.

LEMMA 2.2.3. X_t is a Markov process under $\Pi_T(\cdot, y)$.

PROOF. Let $t \leq T$, and let $f(X_{[t, \infty)})$ be nonnegative and bounded. Then

$$\Pi_T(f(X_{[t, \infty)})|\mathcal{F}_t, y) = \frac{\mathbf{E}(f(X_{[t, \infty)}) e^{-H_T(X_{[0, T], y})}|\mathcal{F}_t)}{\mathbf{E}(e^{-H_T(X_{[0, T], y})}|\mathcal{F}_t)}.$$

Using that fact that $H_t(X_{[0, t], y})$ is \mathcal{F}_t -measurable, we can write

$$\Pi_T(f(X_{[t, \infty)})|\mathcal{F}_t, y) = \frac{\mathbf{E}(f(X_{[t, \infty)}) e^{-H_T(X_{[0, T], y}) + H_t(X_{[0, t], y})}|\mathcal{F}_t)}{\mathbf{E}(e^{-H_T(X_{[0, T], y}) + H_t(X_{[0, t], y})}|\mathcal{F}_t)}.$$

But note that $H_T(X_{[0, T], y}) - H_t(X_{[0, t], y})$ is a functional of $X_{[t, T]}$ only. Hence by the Markov property of X_t under μ_x , the expression above is $\sigma\{X_t\}$ -measurable. If $t > T$, then the Bayes formula gives trivially $\Pi_T(f(X_{[t, \infty)})|\mathcal{F}_t, y) = \mathbf{E}(f(X_{[t, \infty)})|\mathcal{F}_t)$, which is $\sigma\{X_t\}$ -measurable. Hence the Markov property is established. \square

REMARK 2.2.4. It is trivial to establish that X_t must have the same transition intensities under μ_x and $\Pi_T(\cdot, y)$ for $t > T$: this is an immediate consequence of the fact that then $\Pi_T(f(X_{[t,\infty)}))|\mathcal{F}_t, y = \mathbf{E}(f(X_{[t,\infty)}))|\mathcal{F}_t$. We can thus restrict the discussion below to the time interval $[0, T]$.

Let $\tilde{\mu}$ be the measure on $D([0, T]; \mathbb{S})$ under which X_t has transition intensities $\tilde{\lambda}_{ij}(t)$ and initial measure $X_0 \sim \tilde{\nu}$. By Propositions 2.1.1 and 2.1.3, we find that

$$\frac{d\tilde{\mu}}{d\mu_x} = \frac{d\tilde{\nu}}{d\nu}(X_0) e^{\sum_{i,j=1}^d [\int_0^T \log \gamma_{ij}(s^-) I_{X_{s^-}=a_i} dI_{X_s=a_j} - \int_0^T \lambda_{ij} \gamma_{ij}(s) I_{X_s=a_i} ds]}.$$

Hence we can calculate using Dynkin's formula

$$D(\tilde{\mu}|\mu_x) = D(\tilde{\nu}|\nu) + \sum_{i,j=1}^d \int_0^T \lambda_{ij} \gamma_{ij}(s) (\log \gamma_{ij}(s) - 1) \tilde{\mu}(X_s = a_i) ds.$$

For notational convenience, define the function

$$C(a_k; \{\gamma_{ij}\}) = \sum_{j=1}^d \lambda_{kj} \gamma_{kj} (\log \gamma_{kj} - 1) = \sum_{j \neq k} \lambda_{kj} \gamma_{kj} (\log \gamma_{kj} - 1) - \lambda_{kk}.$$

Then we can write

$$D(\tilde{\mu}|\mu_x) = D(\tilde{\nu}|\nu) + \mathbf{E}_{\tilde{\mu}} \left[\int_0^T C(X_s; \{\gamma_{ij}(s)\}) ds \right].$$

Similarly, we can write explicitly

$$\langle H_T(\cdot, y), \tilde{\mu} \rangle = \mathbf{E}_{\tilde{\mu}} \left[\frac{1}{2} \int_0^T h(X_s)^2 ds + \int_{0+}^T y_s dh(X_s) - h(X_T) y_T \right].$$

Now suppose that $y \in C^1$. Then we obtain, after integrating by parts,

$$\begin{aligned} D(\tilde{\mu}|\mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle &= \mathbf{E}_{\tilde{\mu}} \left[\log \frac{d\tilde{\nu}}{d\nu}(X_0) + \int_0^T C(X_s; \{\gamma_{ij}(s)\}) ds \right. \\ &\quad \left. + \frac{1}{2} \int_0^T (\dot{y}_s - h(X_s))^2 ds - \frac{1}{2} \int_0^T (\dot{y}_s)^2 ds \right]. \end{aligned}$$

By Lemma 2.2.1, the law of X_t under $\Pi_T(\cdot, y)$ can be found by minimizing this expression with respect to $\tilde{\nu}$ and $\gamma_{ij}(s)$ at least for $y \in C^1$ (we can restrict to this case, and take limits at the end of the day). But this is a stochastic optimal control problem for a finite-state Markov chain!

REMARK 2.2.5. Recall that in our observation model $\dot{y}_s = h(X_s) + \text{noise}$, where X_s is the uncontrolled signal. Let us call the controlled signal \tilde{X}_s . If $(\dot{y}_s - h(\tilde{X}_s))^2$ is too large then the estimator is not doing a good job. On the other hand, if $(\dot{y}_s - h(\tilde{X}_s))^2$ is too small then we are likely tracking the fluctuations of the noise rather than the signal itself. The cost above vividly reflects this intuition. The estimator tries to minimize $(\dot{y}_s - h(\tilde{X}_s))^2$, but penalizes a large control effort through the $C(\tilde{X}_s; \{\gamma_{ij}(s)\})$ term in order to avoid tracking the noise. The first term of the cost only controls the initial distribution, and we will get rid of it presently.

To solve the control problem, we can proceed using dynamic programming. To this end, we define the cost-to-go

$$J^{\{\gamma_{ij}\}}(t, a_i) = \mathbf{E}_{\tilde{\mu}} \left(\int_t^T [C(X_s; \{\gamma_{ij}(s)\}) + h(X_s) (\frac{1}{2}h(X_s) - \dot{y}_s)] ds \middle| X_t = a_i \right).$$

Note that we can optimize separately over $\gamma_{ij}(s)$ and $\tilde{\nu}$: after all,

$$D(\tilde{\mu} || \mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle = D(\tilde{\nu} || \nu) + \sum_{i=1}^d \tilde{\nu}^i J^{\{\gamma_{ij}\}}(0, a_i).$$

Hence the optimal $\gamma_{ij}(s)$ can be found by minimizing $J^{\{\gamma_{ij}\}}(0, a_i)$. Now define the value function $V(t, a_i) = \min_{\{\gamma_{ij}\}} J^{\{\gamma_{ij}\}}(t, a_i)$. Standard dynamic programming arguments suggest that $V(t, a_i)$ should satisfy the Bellman equation

$$-\frac{d}{dt} V(t, a_i) = \min_{\{\gamma_{ij}\}} \left[\sum_{j \neq i} \lambda_{ij} \gamma_{ij} (V(t, a_j) - V(t, a_i)) + C(a_i; \{\gamma_{ij}\}) \right] + h(a_i) (\frac{1}{2}h(a_i) - \dot{y}_t),$$

with the terminal condition $V(T, a_i) = 0$; moreover, we expect that the optimal control $\gamma_{ij}^*(t)$ will be given by

$$\{\gamma_{ij}^*(t)\} = \operatorname{argmin}_{\{\gamma_{ij}\}} \left[\sum_{j \neq i} \lambda_{ij} \gamma_{ij} (V(t, a_j) - V(t, a_i)) + C(a_i; \{\gamma_{ij}\}) \right].$$

Following a standard device of stochastic control theory, we will justify these statements by proving an appropriate verification lemma.

LEMMA 2.2.6. *Suppose there exists $\{\gamma_{ij}^*(t)\}$ and $V : [0, T] \times \mathbb{S} \rightarrow \mathbb{R}$ such that*

(1) *For all $t \in [0, T]$ and $i = 1, \dots, d$, the function V satisfies*

$$\begin{aligned} \frac{d}{dt} V(t, a_i) + \sum_{j \neq i} \lambda_{ij} \gamma_{ij}^*(t) (V(t, a_j) - V(t, a_i)) \\ + C(a_i; \{\gamma_{ij}^*(t)\}) + h(a_i) (\frac{1}{2}h(a_i) - \dot{y}_t) = 0. \end{aligned}$$

(2) *For all $t \in [0, T]$ and $i = 1, \dots, d$, and any $\{\gamma_{ij}(t)\}$, V satisfies*

$$\begin{aligned} \frac{d}{dt} V(t, a_i) + \sum_{j \neq i} \lambda_{ij} \gamma_{ij}(t) (V(t, a_j) - V(t, a_i)) \\ + C(a_i; \{\gamma_{ij}(t)\}) + h(a_i) (\frac{1}{2}h(a_i) - \dot{y}_t) \geq 0. \end{aligned}$$

(3) *For all $i = 1, \dots, d$, V satisfies the terminal condition $V(T, a_i) = 0$.*

Then $V(t, a_i) = \min_{\{\gamma_{ij}\}} J^{\{\gamma_{ij}\}}(t, a_i)$ and $\{\gamma_{ij}^\} \in \operatorname{argmin}_{\{\gamma_{ij}\}} J^{\{\gamma_{ij}\}}(t, a_i)$.*

PROOF. The basic idea is to evaluate the quantity $V(t, X_t)$. Using Dynkin's formula and the chain rule, we find that

$$V(T, X_T) = V(t, X_t) + \int_t^T \left[\frac{\partial}{\partial s} V(s, X_s) + \mathcal{L}^{\{\gamma_{ij}(s)\}} V(s, X_s) \right] ds + M_T - M_t,$$

where M_T is a martingale under $\tilde{\mu}$ and the generator $\mathcal{L}^{\{\gamma_{ij}\}}$ is defined as

$$\mathcal{L}^{\{\gamma_{ij}\}} f(a_i) = \sum_{j \neq i} \lambda_{ij} \gamma_{ij} (f(a_j) - f(a_i)).$$

Rearranging, taking the expectation, and using part (3), we find that

$$V(t, a_i) = \mathbf{E}_{\tilde{\mu}} \left[\int_t^T \left\{ -\frac{\partial}{\partial s} V(s, X_s) - \mathcal{L}^{\{\gamma_{ij}(s)\}} V(s, X_s) \right\} ds \middle| X_t = a_i \right].$$

First, suppose that $\gamma_{ij}(t) = \gamma_{ij}^*(t)$. Then we obtain, from part (1) of the statement of the Lemma, that $V(t, a_i) = J^{\{\gamma_{ij}^*\}}(t, a_i)$. For any $\gamma_{ij}(t)$, on the other hand, we find from part (2) that $V(t, a_i) \leq J^{\{\gamma_{ij}\}}(t, a_i)$. Hence we have established that

$$V(t, a_i) = J^{\{\gamma_{ij}^*\}}(t, a_i) \leq J^{\{\gamma_{ij}\}}(t, a_i)$$

for any control strategy $\{\gamma_{ij}\}$. The result follows immediately. \square

Let us now proceed to evaluate these expressions explicitly. The minimum in the Bellman equation is easily calculated; we find that

$$\gamma_{ij}^*(t) = \exp(V(t, a_i) - V(t, a_j)).$$

Substituting into the Bellman equation, we find the equation

$$\frac{d}{dt} V(t, a_i) = \sum_{j=1}^d \lambda_{ij} e^{V(t, a_i) - V(t, a_j)} - \frac{1}{2} h(a_i)^2 + h(a_i) \dot{y}_t.$$

In particular, note that we have

$$\frac{d}{dt} e^{-V(t, a_i)} = - \sum_{j=1}^d \lambda_{ij} e^{-V(t, a_j)} + \frac{1}{2} h(a_i)^2 e^{-V(t, a_i)} - h(a_i) e^{-V(t, a_i)} \dot{y}_t.$$

Comparing with the identical expression in the proof of Lemma 2.1.5, we can conclude that evidently the value function for our control problem is given by the expression $V(t, a_i) = -\log v_{t,T}^i(y)$.

REMARK 2.2.7. This should not come as a surprise, as it can be concluded directly from Lemma 2.2.1: indeed, it is easily verified that

$$V(t, a_i) = \min_{\tilde{\mu} \in \mathcal{P}_x^t} \left\{ D(\tilde{\mu} \| \mu_x^{t,i}) + \langle H_T(\cdot, y) - H_t(\cdot, y), \tilde{\mu} \rangle \right\},$$

where \mathcal{P}_x^t denotes the set of probability measures on $D([t, T]; \mathbb{S})$ and $\mu_x^{t,i} \in \mathcal{P}_x^t$ denotes the measure under which X_s is a Markov process with transition intensities λ_{ij} and initial state $X_t = a_i$ a.s. Hence by Lemma 2.2.1, $V(t, a_i) = -\log v_{t,T}^i(y)$.

We are now essentially done: after all, we have found that

$$\tilde{\lambda}_{ij}(t) = \lambda_{ij} \gamma_{ij}^*(t) = \lambda_{ij} \frac{v_{t,T}^j(y)}{v_{t,T}^i(y)}$$

are the transition intensities of the Markov process X_t under $\Pi_T(\cdot, y)$, which is precisely the statement of Theorem 2.1.6 (the corresponding initial distribution $\Pi_T(X_0 = a_i, y)$ is a tautology). We see that this result emerges naturally in the variational/stochastic control perspective adopted in this section from the minimization of the cost function $J^{\{\gamma_{ij}\}}(0, \cdot)$.

We are still missing one technical detail: we have only established the result for $y \in C^1$. This is not an issue, however. A simple limiting argument, of the type we

have frequently used in the previous sections, can be used to extend the result to all observation paths. The details are straightforward and we will omit them here.

2.2.3. Dynamic programming—time-reversed case. Having obtained a forward characterization of the conditional signal, we could proceed to reverse time exactly as in section 2.1.3. Here, however, we are interested in the variational characterization of the conditional signals. Let us thus show how the time-reversed conditional signal can be obtained directly from the solution of a reverse time stochastic control problem. In essence, we proceed exactly as in the forward case, only using the reverse time signal \bar{X}_t defined in Theorem 2.1.8 with its reversed filtration $\bar{\mathcal{F}}_t = \sigma\{\bar{X}_s : 0 \leq s \leq t\}$ instead of the usual signal process (X_t, \mathcal{F}_t) .

LEMMA 2.2.8. $(\bar{X}_t, \bar{\mathcal{F}}_t)$ is a Markov process under $\Pi_T(\cdot, y)$.

PROOF. The proof is identical to that of Lemma 2.2.3. \square

Now recall that under μ_x , \bar{X}_t has transition intensities and initial measure

$$\bar{\lambda}_{ij}(t) = \lambda_{ji} \frac{\mu_x(\bar{X}_{T-t} = a_j)}{\mu_x(\bar{X}_{T-t} = a_i)} = \lambda_{ji} \frac{p_{\bar{X}_t = a_j}}{p_{\bar{X}_t = a_i}}, \quad \bar{\nu}^i = \mu_x(X_T = a_i) = p_{\bar{X}_0 = a_i},$$

see, e.g., [Eli86]. Now let $\tilde{\mu}$ be the measure on $D([0, T]; \mathbb{S})$ under which \bar{X}_t has transition intensities $\tilde{\lambda}_{ij}(t) = \bar{\lambda}_{ij}(t) \gamma_{ij}(t)$ ($i \neq j$) and initial measure $\tilde{\nu}$. Then

$$\frac{d\tilde{\mu}}{d\tilde{\nu}} = \frac{d\tilde{\nu}}{d\tilde{\nu}}(\bar{X}_0) e^{\sum_{i,j=1}^d \int_0^T \log \gamma_{ij}(s^-) I_{\bar{X}_{s-} = a_i} dI_{\bar{X}_s = a_j} - \int_0^T \bar{\lambda}_{ij}(s) \gamma_{ij}(s) I_{\bar{X}_s = a_i} ds}.$$

This follows, as usual, from Propositions 2.1.1 and 2.1.3 (nothing changes in the proofs of these results if we reverse time or have time-dependent intensities under the prior measure μ_x). Using Dynkin's formula, we calculate explicitly

$$D(\tilde{\mu} || \mu_x) = D(\tilde{\nu} || \tilde{\nu}) + \sum_{i,j=1}^d \int_0^T \bar{\lambda}_{ij}(s) \gamma_{ij}(s) (\log \gamma_{ij}(s) - 1) \tilde{\mu}(\bar{X}_s = a_i) ds.$$

For notational convenience, define the function

$$\bar{C}_t(a_k; \{\gamma_{ij}\}) = \sum_{j=1}^d \bar{\lambda}_{kj}(t) \gamma_{kj} (\log \gamma_{kj} - 1).$$

Now suppose that $y \in C^1$. Then we obtain, after integrating by parts,

$$\begin{aligned} D(\tilde{\mu} || \mu_x) + \langle H_T(\cdot, y), \tilde{\mu} \rangle &= \mathbf{E}_{\tilde{\mu}} \left[\log \frac{d\tilde{\nu}}{d\tilde{\nu}}(\bar{X}_0) + \int_0^T \bar{C}_s(\bar{X}_s; \{\gamma_{ij}(s)\}) ds \right. \\ &\quad \left. + \frac{1}{2} \int_0^T (\dot{y}_{T-s} - h(\bar{X}_s))^2 ds - \frac{1}{2} \int_0^T (\dot{y}_{T-s})^2 ds \right]. \end{aligned}$$

By Lemma 2.2.1, the law of \bar{X}_t under $\Pi_T(\cdot, y)$ can be found by minimizing this expression with respect to $\tilde{\nu}$ and $\gamma_{ij}(s)$ at least for $y \in C^1$. Hence everything is exactly as in the previous section. Introduce the cost-to-go

$$\bar{J}^{\{\gamma_{ij}\}}(t, a_i) = \mathbf{E}_{\tilde{\mu}} \left[\int_t^T [\bar{C}_s(\bar{X}_s; \{\gamma_{ij}(s)\}) + h(\bar{X}_s) (\frac{1}{2} h(\bar{X}_s) - \dot{y}_{T-s})] ds \middle| \bar{X}_t = a_i \right]$$

and the associated value function $\bar{V}(t, a_i) = \min_{\{\gamma_{ij}\}} \bar{J}^{\{\gamma_{ij}\}}(t, a_i)$. Proceeding as in the previous section, we find that

$$\{\gamma_{ij}^*\} = \operatorname{argmin}_{\{\gamma_{ij}\}} \bar{J}^{\{\gamma_{ij}\}}(t, a_i), \quad \gamma_{ij}^*(t) = \exp(\bar{V}(t, a_i) - \bar{V}(t, a_j)),$$

and that the value function satisfies the ODE

$$\frac{d}{dt} \bar{V}(t, a_i) = \sum_{j=1}^d \bar{\lambda}_{ij}(t) e^{\bar{V}(t, a_i) - \bar{V}(t, a_j)} - \frac{1}{2} h(a_i)^2 + h(a_i) \dot{y}_{T-t}.$$

Can we make sense of this quantity? Let us look at Lemma 2.2.1 for some guidance. It is not difficult to see that $\bar{V}(t, a_i)$ is characterized by

$$\bar{V}(t, a_i) = \min_{\tilde{\mu} \in \mathcal{P}_x^t} \{D(\tilde{\mu} \parallel \bar{\mu}_x^{t, i}) + \langle H_{T-t}(\cdot, y), \tilde{\mu} \rangle\},$$

where \mathcal{P}_x^t denotes the set of probability measures on $D([0, T-t]; \mathbb{S})$ and $\bar{\mu}_x^{t, i} \in \mathcal{P}_x^t$ denotes the measure under which \bar{X}_s is a Markov process with the same transition intensities as under μ_x , but with initial state $\bar{X}_t = a_i$ a.s. Hence by Lemma 2.2.1

$$\bar{V}(t, a_i) = -\log \mathbf{E}_{\mu_x} [e^{-H_{T-t}(X_{[0, T-t]}, y)} \mid X_{T-t} = a_i] = -\log \left[\frac{\sigma_{T-t}^i(y)}{\mu_x(X_{T-t} = a_i)} \right].$$

Let us verify this explicitly. Using the ODE for $\bar{V}(t, a_i)$ above, we find

$$\begin{aligned} \frac{d}{dt} e^{-\bar{V}(T-t, a_i)} &= \sum_{j=1}^d \frac{p_{X_t=a_j}}{p_{X_t=a_i}} \lambda_{ji} (e^{-\bar{V}(T-t, a_j)} - e^{-\bar{V}(T-t, a_i)}) \\ &\quad - \frac{1}{2} h(a_i)^2 e^{-\bar{V}(T-t, a_i)} + h(a_i) e^{-\bar{V}(T-t, a_i)} \dot{y}_t. \end{aligned}$$

In anticipation of the result, define $\sigma_t^i(y) = p_{X_t=a_i} e^{-\bar{V}(T-t, a_i)}$. Using the chain rule, we find the expression

$$\frac{d}{dt} \sigma_t^i(y) = \sum_{j=1}^d \lambda_{ji} \sigma_t^j(y) - \frac{1}{2} h(a_i)^2 \sigma_t^i(y) + h(a_i) \sigma_t^i(y) \dot{y}_t.$$

This is precisely the correct answer, compare with the proof of Proposition 1.4.5. What have we achieved? Using our newfound expression for $\bar{V}(t, a_i)$, we can write

$$\gamma_{ij}^*(t) = \exp(\bar{V}(t, a_i) - \bar{V}(t, a_j)) = \frac{\sigma_{T-t}^j(y)}{\sigma_{T-t}^i(y)} \frac{p_{X_{T-t}=a_i}}{p_{X_{T-t}=a_j}} = \frac{\pi_{T-t}^j(y)}{\pi_{T-t}^i(y)} \frac{p_{X_{T-t}=a_i}}{p_{X_{T-t}=a_j}}$$

(the latter equality follows from $\pi_t^i(y) = \sigma_t^i(y)/|\sigma_t(y)|$). Hence under $\Pi_T(\cdot, y)$, the process \bar{X}_t is a Markov process with transition intensities matrix

$$\tilde{\lambda}_{ij}(t) = \bar{\lambda}_{ij}(t) \gamma_{ij}^*(t) = \lambda_{ji} \frac{\pi_{T-t}^j(y)}{\pi_{T-t}^i(y)}.$$

We have thus reproduced also the statement of Theorem 2.1.8 from the stochastic control perspective (at least for $y \in C^1$, which is generalized in a straightforward manner). The corresponding cost was given by $\bar{J}^{\{\gamma_{ij}^*\}}(0, \cdot)$.

2.2.4. Some more variations on the same theme. In the papers [Mit81, FM83, Mit82], the connection between nonlinear filtering and stochastic control was made by observing that in the case of diffusion signals, the logarithm of the unnormalized (pathwise) filtering density satisfies a Hamilton-Jacobi-Bellman PDE. Such equations have a characteristic form and it is not difficult to reconstruct an underlying control problem. In the finite state space case the Bellman equations are not as easily recognizable using the naked eye (but see [She85]). Nonetheless, the logarithm of the unnormalized filter can be expressed as a Bellman equation, just like in the diffusion case. Let us take a moment to deduce these equations, and hence the corresponding control problems for which the logarithm of the filter is the value function, from the results obtained in the previous sections.

In the previous section, we established that

$$\hat{V}(t, a_i) = -\log \sigma_{T-t}^i(y) = \bar{V}(t, a_i) - \log p_{\bar{X}_t=a_i},$$

where $\sigma_{T-t}^i(y)$ is the unnormalized filter (it could be obtained from the Zakai equation if we were not interested in the pathwise form). Let us define

$$\hat{\gamma}_{ij}(t) = \gamma_{ij}(t) \frac{p_{\bar{X}_t=a_j}}{p_{\bar{X}_t=a_i}}, \quad \hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i) = \bar{J}^{\{\gamma_{ij}\}}(t, a_i) - \log p_{\bar{X}_t=a_i}.$$

Then it is evidently the case that

$$\tilde{\lambda}_{ij}(t) = \lambda_{ji} \hat{\gamma}_{ij}(t) \quad (i \neq j), \quad \hat{V}(t, a_i) = \min_{\{\hat{\gamma}_{ij}\}} \hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i).$$

Call $Q(t, a_i) = -\log p_{\bar{X}_t=a_i}$. Using Dynkin's formula, we can write

$$Q(T, \bar{X}_T) = Q(t, \bar{X}_t) + \int_t^T G(s, \bar{X}_s) ds + M_T - M_t,$$

where M_t is an $\bar{\mathcal{F}}_t$ -martingale under $\tilde{\mu}$ and

$$G(t, a_i) = \sum_{j \neq i} \tilde{\lambda}_{ij}(t) \gamma_{ij}(t) \log \left(\frac{p_{\bar{X}_t=a_i}}{p_{\bar{X}_t=a_j}} \right) - \tilde{\lambda}_{ii}(t) + \lambda_{ii}.$$

Rearranging and taking the conditional expectation, we find that

$$Q(t, a_i) = \mathbf{E}_{\tilde{\mu}} \left[Q(T, \bar{X}_T) - \int_t^T G(s, \bar{X}_s) ds \middle| \bar{X}_t = a_i \right].$$

Using $\hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i) = \bar{J}^{\{\gamma_{ij}\}}(t, a_i) + Q(t, a_i)$, we obtain

$$\begin{aligned} \hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i) = \\ \mathbf{E}_{\tilde{\mu}} \left[\int_t^T \left[\hat{C}(\bar{X}_s; \{\hat{\gamma}_{ij}(s)\}) + h(\bar{X}_s) \left(\frac{1}{2} h(\bar{X}_s) - \dot{y}_{T-s} \right) \right] ds + Q(T, \bar{X}_T) \middle| \bar{X}_t = a_i \right] \end{aligned}$$

where we have written

$$\hat{C}(a_k; \{\hat{\gamma}_{ij}(t)\}) = \bar{C}_s(a_k; \{\gamma_{ij}(t)\}) - G(t, a_k) = \sum_{j=1}^d \lambda_{jk} \hat{\gamma}_{kj}(t) (\log \hat{\gamma}_{kj}(t) - 1).$$

We see that minus the logarithm of the unnormalized filter is the value function of the stochastic control problem of minimizing $\hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i)$. Unlike the control

problems of the previous sections, this control problem has a terminal cost $Q(T, \bar{X}_T)$ as well as a running cost. The associated Bellman equation is given by

$$-\frac{d}{dt} \hat{V}(t, a_i) = \min_{\{\hat{\gamma}_{ij}\}} \left[\sum_{j \neq i} \lambda_{ji} \hat{\gamma}_{ij} (\hat{V}(t, a_j) - \hat{V}(t, a_i)) + \hat{C}(a_i; \{\hat{\gamma}_{ij}\}) \right] + h(a_i) \left(\frac{1}{2} h(a_i) - \dot{y}_{T-t} \right),$$

with the terminal condition $\hat{V}(T, a_i) = Q(T, a_i)$.

To check the result, let us verify explicitly that this Bellman equation reduces to the equation for minus the logarithm of the unnormalized filter. The minimum in the Bellman equation is easily seen to be attained at

$$\hat{\gamma}_{ij}^*(t) = \exp(\hat{V}(t, a_i) - \hat{V}(t, a_j)).$$

Hence we obtain

$$\frac{d}{dt} \hat{V}(t, a_i) = \sum_{j=1}^d \lambda_{ji} e^{\hat{V}(t, a_i) - \hat{V}(t, a_j)} - \frac{1}{2} h(a_i)^2 + h(a_i) \dot{y}_{T-t}.$$

Using the chain rule, we now calculate

$$\frac{d}{dt} e^{-\hat{V}(T-t, a_i)} = \sum_{j=1}^d \lambda_{ji} e^{-\hat{V}(T-t, a_j)} - \frac{1}{2} h(a_i)^2 e^{-\hat{V}(T-t, a_i)} + h(a_i) e^{-\hat{V}(T-t, a_i)} \dot{y}_t,$$

which is precisely the equation for $e^{-\hat{V}(T-t, a_i)} = \sigma_t^i(y)$. Hence the equation for $-\log \sigma_{T-t}^i(y)$ is indeed a Bellman equation just like in the diffusion case, and we have constructed the corresponding control problem above.

REMARK 2.2.9. Note that the controlled signal process for this control problem is still the time-reversed conditional signal; i.e., its law is given by $\Pi_T(\cdot, y)$. After all, we have done nothing to change the optimum of the control problem; we have only performed a change of variables, shifting the cost \bar{J} and rescaling the control parameters γ_{ij} by constants. We are thus free to choose either this formulation, or the one of the previous section, if we are interested in the conditional signal. The formulation of the previous section follows directly from the variational Kallianpur-Striebel formula, while this section makes the connection to the earlier literature (where the interpretation of the controlled signal as the conditional signal is absent).

Finally, we consider the control problem for which the logarithm of the pathwise Wonham filter $\tau_t^i(y)$ is the value function (compare the following to the expressions in **[Mit81, FM83]** for diffusions). Recall that $\tau_t^i(y) = e^{-h(a_i)y_t} \sigma_t^i(y)$, so that

$$\tilde{V}(t, a_i) = -\log \tau_{T-t}^i(y) = \hat{V}(t, a_i) + h(a_i) y_{T-t}.$$

By this point, the way forward should not come as a surprise. We calculate

$$0 = h(\bar{X}_T) y_0 = h(\bar{X}_t) y_{T-t} + \int_t^T \tilde{G}(s, \bar{X}_s) ds + \tilde{M}_T - \tilde{M}_t,$$

where \tilde{M}_t is an $\bar{\mathcal{F}}_t$ -martingale under $\tilde{\mu}$ and

$$\tilde{G}(t, a_i) = -h(a_i) \dot{y}_{T-t} + \sum_{j \neq i} \lambda_{ji} \hat{\gamma}_{ij}(t) (h(a_j) - h(a_i)) y_{T-t}.$$

Taking the conditional expectation, we obtain

$$h(a_i) y_{T-t} = \mathbf{E}_{\bar{\mu}} \left[- \int_t^T \tilde{G}(s, \bar{X}_s) ds \middle| \bar{X}_t = a_i \right].$$

Hence we can write

$$\begin{aligned} \tilde{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i) &= \hat{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i) + h(a_i) y_{T-t} = \\ &\mathbf{E}_{\bar{\mu}} \left[\int_t^T \left[\tilde{C}_s(\bar{X}_s; \{\hat{\gamma}_{ij}(s)\}) + \frac{1}{2} h(\bar{X}_s)^2 \right] ds + Q(T, \bar{X}_T) \middle| \bar{X}_t = a_i \right], \end{aligned}$$

where we have written

$$\tilde{C}_t(a_k; \{\hat{\gamma}_{ij}\}) = \sum_{j=1}^d \lambda_{jk} \hat{\gamma}_{kj} [\log \hat{\gamma}_{kj} + (h(a_k) - h(a_j)) y_{T-t} - 1].$$

Evidently $\tilde{V}(t, a_i) = \min_{\{\hat{\gamma}_{ij}\}} \tilde{J}^{\{\hat{\gamma}_{ij}\}}(t, a_i)$ is the value function for the stochastic control problem of minimizing the cost $\tilde{J}^{\{\hat{\gamma}_{ij}\}}(0, \cdot)$.

REMARK 2.2.10. Note that unlike in the previous control problems, this cost is not restricted to $y \in C^1$ as the derivative term $\propto \dot{y}_t$ has been transformed away. This is entirely in the spirit of the pathwise filtering theory, and could be a possible advantage of working with this form of the control problem.

It remains to find the corresponding Bellman equation. This is easily done:

$$-\frac{d}{dt} \tilde{V}(t, a_i) = \min_{\{\hat{\gamma}_{ij}\}} \left[\sum_{j \neq i} \lambda_{ji} \hat{\gamma}_{ij} (\tilde{V}(t, a_j) - \tilde{V}(t, a_i)) + \tilde{C}_t(a_i; \{\hat{\gamma}_{ij}\}) \right] + \frac{1}{2} h(a_i)^2,$$

with the terminal condition $\tilde{V}(T, a_i) = Q(T, a_i)$. For completeness, let us check once more that this is indeed the correct equation. The minimum in the Bellman equation is attained at

$$\hat{\gamma}_{ij}^*(t) = \exp(\tilde{V}(t, a_i) - \tilde{V}(t, a_j) + (h(a_j) - h(a_i)) y_{T-t}).$$

Hence we obtain

$$\frac{d}{dt} \tilde{V}(t, a_i) = \sum_{j=1}^d \lambda_{ji} e^{\tilde{V}(t, a_i) - \tilde{V}(t, a_j)} e^{(h(a_j) - h(a_i)) y_{T-t}} - \frac{1}{2} h(a_i)^2,$$

and in particular we find that

$$\frac{d}{dt} e^{-\tilde{V}(T-t, a_i)} = \sum_{j=1}^d e^{-\tilde{V}(T-t, a_j)} \lambda_{ji} e^{(h(a_j) - h(a_i)) y_t} - \frac{1}{2} h(a_i)^2 e^{-\tilde{V}(T-t, a_i)}.$$

This is precisely the pathwise Wonham filter, see Proposition 1.4.5.

2.3. Exponential stability of the Wonham filter

2.3.1. Filter stability. In this section we are going to apply the results of the previous sections to the filter stability problem. Let us begin by describing what we are trying to achieve.

Consider, for example, the Wonham filter:

$$d\pi_t = \Lambda^* \pi_t dt + (H - h^* \pi_t) \pi_t (dY_t - h^* \pi_t dt), \quad \pi_0 = \nu.$$

The solution π_t of this equation satisfies $\pi_t^i = \mathbf{P}(X_t = a_i | \mathcal{F}_t^Y)$ a.s. if the signal process X_t is a finite-state Markov process with initial measure $\mathbf{P}(X_0 = a_i) = \nu^i$, and the observations Y_t are given by the usual expression. In practice, however, we may not know exactly what the initial measure ν is. Hence we could erroneously process the observations Y_t using a filter with the wrong initial condition:

$$d\pi_t(\mu) = \Lambda^* \pi_t(\mu) dt + (H - h^* \pi_t(\mu)) \pi_t(\mu) (dY_t - h^* \pi_t(\mu) dt), \quad \pi_0 = \mu.$$

$\pi_t(\mu)$ is no longer the least-mean-square estimate of the signal process X_t with initial measure ν . However, it is intuitively plausible that the initial measure should not contribute much to the estimate of the signal at time t for large t : the information obtained from the observations should supersede the prior information of the signal. Hence one could hope that the wrongly initialized filter $\pi_t(\mu)$ will converge to the optimal filter π_t as $t \rightarrow \infty$.

To study the wrongly initialized filter, it is useful to understand the meaning of this quantity. Throughout this thesis, we will use three different representations of the wrongly initialized filter. Let us list them here.

- (1) $\pi_t(\mu)$ is the solution of the Wonham equation with the initial condition $\pi_0(\mu) = \mu$; i.e., the wrongly initialized filter is characterized as the solution of a certain stochastic differential equation.
- (2) Note that the Wonham filter with initial condition μ would in fact be the optimal filter, if only the signal process X_t had initial measure μ . Let us call \mathbf{P}^μ the measure under which this is the case. If $\mu \ll \nu$, then we obtain using Proposition 2.1.3 and the Bayes formula

$$\pi_t^i(\mu) = \mathbf{P}^\mu(X_t = a_i | \mathcal{F}_t^Y) = \frac{\mathbf{E}(\frac{d\mu}{d\nu}(X_0) I_{X_t=a_i} | \mathcal{F}_t^Y)}{\mathbf{E}(\frac{d\mu}{d\nu}(X_0) | \mathcal{F}_t^Y)}.$$

- (3) Similarly, we can write using the Kallianpur-Striebel formula

$$\pi_t^i(\mu) = \frac{\mathbf{E}^\mu(Z_t(X_{[0,t]}, y) I_{X_t=a_i})}{\mathbf{E}^\mu(Z_t(X_{[0,t]}, y))},$$

or using $\tilde{Z}_t(x, y)$ instead of $Z_t(x, y)$ if a pathwise version is desired.

All three representations have their advantages and disadvantages. The first representation is practically motivated: the filter is implemented using a differential equation, so that operationally the wrongly initialized filter is the solution of this equation with the wrong initial condition. A disadvantage of this representation is that it is not probabilistic, i.e., it is divorced from the underlying estimation problem, which makes the direct analysis of this equation more challenging. The second representation is probabilistic in nature, and directly demonstrates the connection between the wrongly initialized filter and a smoothing problem. A disadvantage of this representation is that it only makes sense for absolutely continuous initial distributions $\mu \ll \nu$. The third representation is also probabilistic in nature, but does not suffer from this problem. The drawback compared to the second representation is that the dependence on the initial condition is not as explicit.

Note that the first and third representations are equivalent for μ_w -a.e. observation sample path (and hence \mathbf{P}^μ - and \mathbf{P}^ν -a.s., as the Kallianpur-Striebel formula demonstrates that these measures are equivalent to the Wiener measure when restricted to the σ -algebra generated by the observation sample paths). All three

representations are equivalent if $\mu \ll \nu$. With this in mind, we may use any combination of these representations to study the filter stability problem.

2.3.2. Exponential stability: A coupling proof. In this section we will reformulate the filter stability problem, using the third representation above, in terms of the conditional signal theory studied in the first part of the chapter. We will see that this provides an intuitive way to study filter stability, and will allow us to obtain an explicit exponential bound on the filtering error.

Let us denote by $\Pi_t^\mu(\cdot, y)$ the pathwise conditional measure on $D(\mathbb{R}_+; \mathbb{S})$ for the case where the unconditional signal process X_t has initial measure μ and transition intensities matrix Λ . To be precise,

$$\Pi_t^\mu(A, y) = \frac{\mathbf{E}^\mu(I_A \tilde{Z}_t(X_{[0,t]}, y))}{\mathbf{E}^\mu(\tilde{Z}_t(X_{[0,t]}, y))}.$$

We have seen above (Theorem 2.1.6) that under $\Pi_t^\mu(\cdot, y)$, the signal process X_t is again a Markov process with transition intensities and initial measure (for $t \leq T$)

$$\lambda_{ij}^{t,y}(s) = \lambda_{ij} \frac{v_{s,t}^j(y)}{v_{s,t}^i(y)} \quad (i \neq j), \quad \tilde{\mu}^i = \Pi_t^\mu(X_0 = a_i, y).$$

Now note the following key point:

The dual filter $v_{s,t}^i(y)$ is independent of the initial measure μ of the signal process.

This is true by construction, as the dual filter is obtained by conditioning on the initial point of the signal—see Definition 2.1.4. As a consequence, the signal process X_t has the same transition intensities under $\Pi_t^\mu(\cdot, y)$ and $\Pi_t^\nu(\cdot, y)$, and only the respective initial measures differ. Hence in order to prove exponential stability of the filter, it would be sufficient to prove that the Markov process with transition intensities $\lambda_{ij}^{t,y}(s)$ is geometrically ergodic. This is precisely what we will do, using a coupling approach similar to that of D. Griffeath [Gri75].

PROPOSITION 2.3.1. *The following holds for any $y \in C_0(\mathbb{R}_+, \mathbb{R})$ and $t < \infty$:*

$$\sum_{k=1}^d |\Pi_t^\mu(X_t = a_k, y) - \Pi_t^\nu(X_t = a_k, y)| \leq 2 \exp\left(-\int_0^t \min_{i \neq j} \{\lambda_{ij}^{t,y}(s) + \lambda_{ji}^{t,y}(s)\} ds\right).$$

PROOF. Fix $y \in C_0(\mathbb{R}_+; \mathbb{R})$. We are interested in studying the relative behavior of two Markov processes with equal transition intensities $\lambda_{ij}^{t,y}(s)$ ($s \leq t$) but different initial measures $\Pi_t^\mu(X_0 = a_i, y)$ and $\Pi_t^\nu(X_0 = a_i, y)$. To this end, let us construct two such processes on the same probability space. That is, we introduce a space $(\Omega, \mathcal{F}, \mathbf{P}_c)$, on which are defined two stochastic processes X_s^1 and X_s^2 , such that both X_s^1 and X_s^2 have transition intensities $\lambda_{ij}^{t,y}(s)$ for $s \leq t$ and 0 for $s > t$, $\mathbf{P}_c(X_0^1 = a_i) = \Pi_t^\mu(X_0 = a_i, y)$, and $\mathbf{P}_c(X_0^2 = a_i) = \Pi_t^\nu(X_0 = a_i, y)$.

We would like to bound $\sum_i |\Pi_t^\mu(X_t = a_i, y) - \Pi_t^\nu(X_t = a_i, y)|$. Note that

$$\begin{aligned} & \sum_{i=1}^d |\Pi_t^\mu(X_t = a_i, y) - \Pi_t^\nu(X_t = a_i, y)| \\ &= 2 \max_{A \subset \mathbb{S}} |\Pi_t^\mu(X_t \in A, y) - \Pi_t^\nu(X_t \in A, y)| \\ &= 2 \max_{A \subset \mathbb{S}} |\mathbf{P}_c(X_t^1 \in A) - \mathbf{P}_c(X_t^2 \in A)| \end{aligned}$$

$$\begin{aligned}
&= 2 \max_{A \subset \mathbb{S}} |\mathbf{P}_c(X_t^1 \in A, X_t^1 = X_t^2) + \mathbf{P}_c(X_t^1 \in A, X_t^1 \neq X_t^2) \\
&\quad - \mathbf{P}_c(X_t^2 \in A, X_t^1 = X_t^2) - \mathbf{P}_c(X_t^2 \in A, X_t^1 \neq X_t^2)| \\
&= 2 \max_{A \subset \mathbb{S}} |\mathbf{P}_c(X_t^1 \in A, X_t^1 \neq X_t^2) - \mathbf{P}_c(X_t^2 \in A, X_t^1 \neq X_t^2)| \\
&\leq 2 \max_{A \subset \mathbb{S}} \mathbf{E}_c(|I_{X_t^1 \in A} - I_{X_t^2 \in A}| I_{X_t^1 \neq X_t^2}) \\
&\leq 2 \mathbf{P}_c(X_t^1 \neq X_t^2),
\end{aligned}$$

where we have used the usual identity of the ℓ_1 -norm and the total variation norm. The basic idea behind the coupling method [Lin02] is now as follows. Suppose that \mathbf{P}_c is such that there is some random time ζ , called the coupling time, such that $X_t^1 = X_t^2$ a.s. for all $t \geq \zeta$. Then $\mathbf{P}_c(X_t^1 = X_t^2) \geq \mathbf{P}_c(t \geq \zeta)$, and hence

$$\sum_{i=1}^d |\Pi_t^\mu(X_t = a_i, y) - \Pi_t^\nu(X_t = a_i, y)| \leq 2 \mathbf{P}_c(X_t^1 \neq X_t^2) \leq 2 \mathbf{P}_c(t < \zeta).$$

This is the well-known *coupling inequality*. Our goal is to choose a convenient measure \mathbf{P}_c under which we can obtain an explicit bound on $\mathbf{P}_c(t < \zeta)$.

We will now construct a convenient measure \mathbf{P}_c . Let (X_t^1, X_t^2, Ξ_t) be a Markov process on $(\Omega, \mathcal{F}, \mathbf{P}_c)$ with the state space $S = \mathbb{S} \times \mathbb{S} \times \{0, 1\}$, initial measure $\Pi_t^\mu(X_0 \in \cdot, y) \times \Pi_t^\nu(X_0 \in \cdot, y) \times \delta_{\{0\}}$, and the following nonzero transition intensities:

$u \in S$	$v \in S$	intensity $\xi_{uv}(s)$
	$(j, j, 0)$	$\lambda_{ij}^{t,y}(s)$
$(i, i, 0)$	$(i, i, 1)$	$\kappa(s)$
	$(i, i, 0)$	$\lambda_{ii}^{t,y}(s) - \kappa(s)$
	$(i, i, 1)$	$\lambda_{ji}^{t,y}(s) \kappa(s) / \kappa_{ij}(s)$
	$(i, i, 0)$	$\lambda_{ji}^{t,y}(s) (1 - \kappa(s) / \kappa_{ij}(s))$
	$(j, j, 1)$	$\lambda_{ij}^{t,y}(s) \kappa(s) / \kappa_{ij}(s)$
$(i, j, 0)$	$(j, j, 0)$	$\lambda_{ij}^{t,y}(s) (1 - \kappa(s) / \kappa_{ij}(s))$
	$(i, k, 0)$	$\lambda_{jk}^{t,y}(s)$
	$(k, j, 0)$	$\lambda_{ik}^{t,y}(s)$
	$(i, j, 0)$	$\lambda_{ii}^{t,y}(s) + \lambda_{jj}^{t,y}(s)$
$(i, i, 1)$	$(j, j, 1)$	$\lambda_{ij}^{t,y}(s)$
	$(i, i, 1)$	$\lambda_{ii}^{t,y}(s)$

Here we have used $i \neq j \neq k$, $\kappa_{ij}(s) = \lambda_{ij}^{t,y}(s) + \lambda_{ji}^{t,y}(s)$, and $\kappa(s) = \min_{i \neq j} \kappa_{ij}(s)$. It is easily verified that the marginal processes X_t^1 and X_t^2 are themselves Markov processes with the correct statistics, whereas Ξ_t is a Markov process on $\{0, 1\}$ with transition intensities $\chi_{01}(s) = \kappa(s)$ and $\chi_{10}(s) = 0$. Ξ_t serves as a coupling indicator: Ξ_t only switches once, from zero to one, and after that time $X_t^1 = X_t^2$ a.s. Hence evidently $\zeta = \inf\{t > 0 : \Xi_t = 1\}$ may serve as the coupling time. But Ξ_t is an exceedingly simple Markov process, and we easily evaluate explicitly

$$\mathbf{P}_c(t < \zeta) = \mathbf{P}_c(\Xi_t = 0) = \exp\left(-\int_0^t \kappa(s) ds\right).$$

The statement now follows immediately from the coupling inequality. \square

COROLLARY 2.3.2. *The following bound always holds a.s.:*

$$|\pi_t(\mu) - \pi_t(\nu)| \leq 2 \exp \left(-2t \min_{i \neq j} \sqrt{\lambda_{ij} \lambda_{ji}} \right).$$

PROOF. This follows from the following simple identity:

$$\lambda_{ij}^{t,y}(s) + \lambda_{ji}^{t,y}(s) = \lambda_{ij} \frac{v_{s,t}^j(y)}{v_{s,t}^i(y)} + \lambda_{ji} \frac{v_{s,t}^i(y)}{v_{s,t}^j(y)} \geq \inf_{x>0} \left\{ \lambda_{ij} x + \lambda_{ji} \frac{1}{x} \right\} = 2\sqrt{\lambda_{ij} \lambda_{ji}}.$$

The proof is complete. \square

This result should be compared with [BCL04, Thm. 4.3], which has the same exponential rate but a prefactor that blows up as μ or ν approach the boundary of the simplex. The constant prefactor 2 is a significant improvement, and shows that nothing bad happens even when the two initial measures are mutually singular (as we will see in chapter 3, the *derivative* of the filter with respect to its initial condition does not share this nice property).

REMARK 2.3.3. The bound can be improved a little more; see Remark 3.3.5.

In all fairness, it should be noted that the conditional signal theory is not key to this improvement: the important difference between our result and that of [BCL04] is that we use the forward rather than the time reverse characterization of the conditional signal. As noted in Remark 2.1.9, the forward smoothing equation (i.e., the Kolmogorov forward equation of the forward conditional signal) can be obtained by more elementary means, after which the bound of Corollary 2.3.2 could be obtained through direct analysis of this ODE similar to [BCL04]. Nonetheless, I think that the conditional signal point of view is valuable. First, it gives a completely probabilistic proof of the stability result, with a very intuitive interpretation. Second, it is very close in spirit to the method which we will employ to obtain stability bounds for diffusions (in that case a natural “coupling” is provided by the theory of stochastic flows), so that a parallel development of the finite state case is quite insightful. Finally, a parallel development of the finite state and diffusion cases suggests that we might be able to combine these results to obtain filter stability bounds for hybrid signals consisting of a diffusion and a switching component. Such signals are important in applications and require further investigation.

Finally, we note that the exponential rate in Corollary 2.3.2 depends only on the transition intensities of the signal process, and not on the observation structure. This does not capture the intuitive idea that the stability of the filter should improve if the signal-to-noise ratio increases (this need not always be the case, however: see the examples in [DZ91, VM06]). The observation function only enters the statistics of the conditional signal through the dual filter, however, so that in order to include the effect of the observations we would have to quantify the dependence of the dual filter on the observation structure. This adds significant difficulty to the problem. It is even unclear whether a bound of the form of Corollary 2.3.2 can be significantly improved; even if a high signal-to-noise ratio leads to improved stability for a typical observation sample path, this does not necessarily imply that such a bound can be obtained for any $y \in C_0(\mathbb{R}_+; \mathbb{R})$.

In the diffusion case (chapter 4) we will nonetheless obtain pathwise upper bounds on filter stability. As emphasized above, this requires an analysis of the dependence of the generator of the conditional signal on the observation structure. It will turn out that it is easier to quantify the dependence of the filter on the

observations than to quantify the dependence of the dual filter on the observations. As such it will pay off, unlike in the case of Corollary 2.3.2, to use the time-reversed conditional signal. There we will also put the stochastic control formulation to good use, as it will play a central role in the analysis of the conditional generator.

Model Robustness of the Nonlinear Filter: Finite State Space

3.1. Introduction

In the previous chapter, we showed that the Wonham filter started from the wrong initial condition converges at an exponential rate $2 \min_{i \neq j} (\lambda_{ij} \lambda_{ji})^{1/2}$ to the Wonham filter started from the correct initial condition. This bodes well for the applicability of the Wonham filter: even if we do not know the initial measure exactly this would not matter on the long run, provided of course that the rate above is positive. The latter requires the *mixing condition* $\lambda_{ij} > 0$ for all $i \neq j$, which we will assume throughout this chapter.

The initial measure is not the only model parameter, however, that is needed as input for the filter. In practice, any of the parameters that determine our model are likely to be at least slightly misspecified: the transition intensities λ_{ij} for a particular signal source and the observation function h for a particular measurement device are rarely known precisely, and we can only do our best to characterize them. Unlike the misspecification of the initial measure, which introduces an error at time $t = 0$ only, it is clear that any errors in Λ or h are continuously perpetrated by the misspecified filter. Hence such a filter can never become optimal as $t \rightarrow \infty$. This need not be a major problem, as long as the error can be made sufficiently small when Λ and h are sufficiently well characterized. The danger, however, is the errors could accumulate over time, making the filter essentially useless after a short and potentially uninteresting transient period. This sort of behavior is very common in numerical approximations—for example, approximations of differential equations typically have errors that grow rapidly in time—but would be quite unacceptable in most signal processing applications.

The goal of this chapter is to show that this accumulation of errors can not occur—i.e., even if we misspecify Λ , h and ν , the error remains bounded on the infinite time interval and can be made arbitrarily small if the error in the model parameters is sufficiently small. In section 0.1.5 we gave some intuition as to why this should be the case in a discrete time setting. When the filter is exponentially stable, it has a mechanism to suppress the error made in every time step, and when this suppression happens at an exponential rate the total error is summable as a geometric series. The method which we will use in the continuous time case is based on similar ideas; the actual implementation of these ideas is much more

This chapter is based on the paper “Model robustness of finite state nonlinear filtering over the infinite time horizon” by Pavel Chigansky and the author [CV06].

subtle, however, and requires a different set of tools. To get some intuition for the problems we will be faced with, let us begin by running through the argument in a toy deterministic example.

3.1.1. A little deterministic intuition. Let us forget about filtering for the moment and consider the simpler problem of approximating a deterministic, time-nonhomogeneous differential equation. Let us define

$$\frac{dx_t}{dt} = f(t, x_t), \quad x_t \in \mathbb{R}^n.$$

We denote by $\varphi_{s,t}(x)$ the solution x_t of this equation when it is started at the initial condition $x_s = x$ ($s \leq t$). We will suppose that f is sufficiently regular so that the flow $\varphi_{s,t}(x)$ exists, is unique, and is a diffeomorphism for all $0 \leq s \leq t < \infty$. Now consider another differential equation

$$\frac{d\tilde{x}_t}{dt} = \tilde{f}(t, \tilde{x}_t), \quad \tilde{x}_t \in \mathbb{R}^n,$$

where again we assume that \tilde{f} is sufficiently regular so that this equation generates a nice flow $\tilde{\varphi}_{s,t}(x)$ as above. We are now interested in obtaining a bound on the error $\|x_t - \tilde{x}_t\| = \|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\|$ in terms of the distance between f and \tilde{f} . This means we need to somehow relate the distance between the flows to the time derivatives of the flows. The following Lemma shows how this can be done.

LEMMA 3.1.1. *The distance between the flows can be bounded as follows:*

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| \leq \int_0^t \left\| \frac{d}{ds} \varphi_{s,t}(\tilde{\varphi}_{0,s}(x)) \right\| ds.$$

PROOF. Note that

$$\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x) = - \int_0^t \frac{d}{ds} \varphi_{s,t}(\tilde{\varphi}_{0,s}(x)) ds,$$

so we get, using the triangle inequality,

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| = \left\| \int_0^t \frac{d}{ds} \varphi_{s,t}(\tilde{\varphi}_{0,s}(x)) ds \right\| \leq \int_0^t \left\| \frac{d}{ds} \varphi_{s,t}(\tilde{\varphi}_{0,s}(x)) \right\| ds$$

which is the desired result. \square

Next, we need to relate the time derivative of the flows to f and \tilde{f} . This is simple for the derivative of $\tilde{\varphi}_{0,s}$ with respect to s , and the remaining derivative follows from the following elementary result.

LEMMA 3.1.2. *The backward time derivative of the flow satisfies*

$$\frac{\partial}{\partial s} \varphi_{s,t}(x) = -D\varphi_{s,t}(x) \cdot f(s, x),$$

where $D\varphi_{s,t}(x) \cdot v$ is the directional derivative of $\varphi_{s,t}(x)$ at x in the direction v .

PROOF. As $\varphi_{s,t}(\varphi_{0,s}(x)) = \varphi_{0,t}(x)$, we know that $d\varphi_{s,t}(\varphi_{0,s}(x))/ds = 0$. Hence

$$\frac{d}{ds} \varphi_{s,t}(\varphi_{0,s}(x')) = \frac{\partial \varphi_{s,t}}{\partial s}(\varphi_{0,s}(x')) + D\varphi_{s,t}(\varphi_{0,s}(x')) \cdot f(s, \varphi_{0,s}(x')) = 0.$$

Substituting $x' = \varphi_{0,s}^{-1}(x)$, we obtain the desired result. \square

Putting together these two results, we obtain our bound:

PROPOSITION 3.1.3. *The following error bound holds:*

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| \leq \int_0^t \|D\varphi_{s,t}(\tilde{\varphi}_{0,s}(x))\| \|f(s, \tilde{\varphi}_{0,s}(x)) - \tilde{f}(s, \tilde{\varphi}_{0,s}(x))\| ds.$$

PROOF. From Lemmas 3.1.1 and 3.1.2 we obtain

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| \leq \int_0^t \|D\varphi_{s,t}(\tilde{\varphi}_{0,s}(x)) \cdot (f(s, \tilde{\varphi}_{0,s}(x)) - \tilde{f}(s, \tilde{\varphi}_{0,s}(x)))\| ds.$$

The result follows from the definition $\|D\varphi_{s,t}(x)\| = \sup_{\|v\|=1} \|D\varphi_{s,t}(x) \cdot v\|$. \square

Note that the bound depends on two separate quantities: the local error $\|f(s, \tilde{\varphi}_{0,s}(x)) - \tilde{f}(s, \tilde{\varphi}_{0,s}(x))\|$ (compare with the quantity $\delta(y_n, \bar{p}_n)$ in section 0.1.5) and the term $\|D\varphi_{s,t}(\tilde{\varphi}_{0,s}(x))\|$, which bounds the sensitivity of the flow $\varphi_{s,t}(x)$ to an infinitesimal perturbation of its initial condition.

Now suppose that the flow $\varphi_{s,t}(x)$ is exponentially stable in the sense that its derivative is exponentially bounded:

$$\|D\varphi_{s,t}(x)\| \leq K e^{-\kappa(t-s)} \quad \text{for some } \kappa, K > 0.$$

Then we obtain the error bound

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| \leq K \int_0^t e^{-\kappa(t-s)} \|f(s, \tilde{\varphi}_{0,s}(x)) - \tilde{f}(s, \tilde{\varphi}_{0,s}(x))\| ds.$$

Hence we see that if the flow $\varphi_{s,t}(x)$ is stable in the above sense, then the local error is suppressed at an exponential rate. In particular, if the local error is bounded $\sup_{t,x} \|f(t,x) - \tilde{f}(t,x)\| \leq M$, then

$$\|\varphi_{0,t}(x) - \tilde{\varphi}_{0,t}(x)\| \leq KM \int_0^t e^{-\kappa(t-s)} ds = \frac{KM}{\kappa},$$

which is precisely what we want. Note that the conceptual similarities are immediately evident when these expressions are compared to those of section 0.1.5.

REMARK 3.1.4. Let us make the connection to the projection filter method. Suppose that we would like to approximate the flow $\varphi_{s,t}(x)$ by a flow $\tilde{\varphi}_{s,t}(x)$ that leaves some low-dimensional manifold $S \subset \mathbb{R}^n$ invariant; i.e., we are interested in model reduction. Given that we have fixed the manifold S , how should we choose $\tilde{f}(t,x)$ to make this a good approximation? Note that the requirement that $\tilde{\varphi}_{s,t}(x)$ leaves S invariant is equivalent to requiring that $\tilde{f}(t,x) \in T_x S$, the tangent space of S at x , for every $x \in S$. Hence to minimize the approximation error, the error bound above suggests that we should choose $\tilde{f}(t,x)$ to be the element of $T_x S$ that minimizes the local error $\|f(t,x) - \tilde{f}(t,x)\|$. But then $\tilde{f}(t,x)$ must be precisely the orthogonal projection of $f(t,x)$ onto $T_x S$ (note that there is no need to define $\tilde{f}(t,x)$ outside S , as the error bound only requires us to evaluate \tilde{f} at $\tilde{\varphi}_{0,s}(x) \in S$). Hence evidently our error bound is naturally related to the form of approximation on which the projection filter technique is based.

3.1.2. Model robustness of the Wonham filter. Having investigated the deterministic case, what problems can we expect to run into when we apply these ideas to the Wonham filter? There are two main issues that need to be overcome.

- (1) An exponential bound on the *derivative* of the filter with respect to its initial condition is a significantly stronger result than a bound on the distance between the optimal and wrongly initialized filters, as obtained

in the previous chapter. We will find, in particular, that this quantity behaves rather unpleasantly near the boundary of the simplex, making a uniform bound of the form $\sup_x \|D\varphi_{s,t}(x)\| \leq K e^{-\kappa(t-s)}$ impossible.

- (2) As the Wonham filter is a stochastic differential equation, we would expect the procedure above to give rise to a time integral as well as a stochastic integral. In principle this need not be a problem, as we could just bound the expectation of the error rather than the error itself. Things are not so straightforward, however, due to the following troublesome fact: the quantity $\|D\varphi_{s,t}(x)\|$ depends on the entire observation history in the interval $[s, t]$. Hence when we try to repeat the above procedure in the stochastic case, this gives rise to *anticipating* (nonadapted) stochastic integrals.

We will begin by tackling the first issue. Before we can begin, we need to establish that the Wonham filter actually generates a flow which is differentiable etc. This is rather straightforward, as the Wonham equation can be obtained by normalizing a linear SDE (the Zakai equation), and linear SDE generate linear flows (which are as regular as it gets). To bound the derivative of the flow, we employ the second representation of the wrongly initialized filter as given in section 2.3.1. This gives an exponential bound on the derivative of the filter when combined with Corollary 2.3.2. Unfortunately, the prefactor in this bound blows up on the boundary of the simplex, so that we need to do some extra work to bound the expectation of this prefactor. This is a little tedious but essentially straightforward.

To deal with the anticipativity problem, we distinguish between two cases. If the observation function of the exact and approximate filters are equal, then the stochastic integral term in the error bound cancels. Hence most of the unpleasantness is circumvented, and we obtain a simple error bound directly from elementary manipulations of the flow of the filter. In the general case, however, we have no recourse but to dump the Itô theory (which is firmly rooted in the nonanticipativity requirement) and to seek an anticipative replacement. The theory of Skorokhod integrals, and in particular the associated stochastic calculus developed by D. Nualart and É. Pardoux [NP88, Nua95] using Malliavin calculus techniques, provides a suitable replacement which allows us to proceed to prove our main result. As this theory is not as widely known as the traditional Itô theory, we have provided an overview in appendix A of those results that will be needed in the proofs.

Let us briefly state the main result of this chapter. We consider the usual Wonham filter setup, i.e., the signal process X_t is Markov process on the state space $\mathbb{S} = \{a_1, \dots, a_d\}$ with transition intensities $\Lambda = (\lambda_{ij})$ and initial distribution $\nu^i = \mathbf{P}(X_0 = a_i)$. The observation process Y_t is given by the usual expression

$$Y_t = \int_0^t h(X_s) ds + B_t,$$

where $h : \mathbb{S} \rightarrow \mathbb{R}$ is the observation function (we will also write $h^i = h(a_i)$) and B is a Wiener process that is independent of X . The conditional probabilities $\pi_t^i = \mathbf{P}(X_t = a_i | \mathcal{F}_t^Y)$ satisfy the Wonham equation

$$(3.1) \quad d\pi_t = \Lambda^* \pi_t dt + (H - h^* \pi_t) \pi_t (dY_t - h^* \pi_t dt), \quad \pi_0 = \nu,$$

where $H = \text{diag } h$. We will denote by $\pi_t(\mu)$ the solution of the Wonham equation at time t with an arbitrary initial distribution $\pi_0 = \mu$, and by $\pi_{s,t}(\mu)$ the solution of the Wonham equation at time $t \geq s$ with the initial condition $\pi_s = \mu$. Now

consider the Wonham filter with incorrect model parameters:

$$(3.2) \quad d\tilde{\pi}_t = \tilde{\Lambda}^* \tilde{\pi}_t dt + (\tilde{H} - \tilde{h}^* \tilde{\pi}_t) \tilde{\pi}_t (dY_t - \tilde{h}^* \tilde{\pi}_t dt), \quad \tilde{\pi}_0 = \nu,$$

where $\tilde{\Lambda}$ and \tilde{h} denote a transition intensities matrix and observation function that do not match the underlying signal-observation model (X, Y) , $\tilde{H} = \text{diag } \tilde{h}$, and we denote by $\tilde{\pi}_t(\mu)$ the solution of this equation with initial condition $\tilde{\pi}_0 = \mu$ and by $\tilde{\pi}_{s,t}(\mu)$ the solution with $\tilde{\pi}_s = \mu$. The goal of this chapter is to prove the following.

THEOREM 3.1.5. *Suppose $\nu^i, \mu^i > 0 \forall i$ and $\lambda_{ij}, \tilde{\lambda}_{ij} > 0 \forall i \neq j$. Then*

$$\sup_{t \geq 0} \mathbf{E} \|\tilde{\pi}_t(\mu) - \pi_t(\nu)\|^2 \leq C_1 |\mu - \nu| + C_2 |\tilde{h} - h| + C_3 |\tilde{\Lambda}^* - \Lambda^*|,$$

where $|\tilde{\Lambda}^* - \Lambda^*| = \sup\{ |(\tilde{\Lambda}^* - \Lambda^*)\tau| : \tau^i > 0 \forall i, |\tau| = 1 \}$ and the quantities C_1, C_2, C_3 are bounded on any compact subset of parameters $\{(\nu, \Lambda, h, \mu, \tilde{\Lambda}, \tilde{h}) : \nu^i, \mu^i > 0 \forall i, |\nu| = |\mu| = 1, \lambda_{ij}, \tilde{\lambda}_{ij} > 0 \forall i \neq j, \sum_j \lambda_{ij} = \sum_j \tilde{\lambda}_{ij} = 0 \forall i\}$. Additionally we have the asymptotic estimate

$$\limsup_{t \rightarrow \infty} \mathbf{E} \|\tilde{\pi}_t(\mu) - \pi_t(\nu)\|^2 \leq C_2 |\tilde{h} - h| + C_3 |\tilde{\Lambda}^* - \Lambda^*|.$$

In particular, this implies that if $\nu^i > 0 \forall i, \lambda_{ij} > 0 \forall i \neq j$, then

$$\lim_{\tilde{h} \rightarrow h} \lim_{\tilde{\Lambda} \rightarrow \Lambda} \limsup_{\mu \rightarrow \nu} \sup_{t \geq 0} \mathbf{E} \|\tilde{\pi}_t(\mu) - \pi_t(\nu)\| = \lim_{\tilde{h} \rightarrow h} \lim_{\tilde{\Lambda} \rightarrow \Lambda} \limsup_{t \rightarrow \infty} \mathbf{E} \|\tilde{\pi}_t(\mu) - \pi_t(\nu)\| = 0.$$

It should be mentioned that there is nothing in principle about our method that prohibits us from obtaining quantitative bounds on filter robustness; in fact, the constants $C_{1,2,3}$ have explicit expressions which can be found in the proofs. The real limiting factor in the sequel is our bound on the decay of the derivative of the filter with respect to its initial condition, which is highly suboptimal. We will discuss this point further in Remark 3.3.8. The unfortunate consequence is that Theorem 3.1.5, though of significant interest to the filter robustness problem, does not provide useful quantitative estimates on the filtering error. For this reason we have not bothered to optimize the constants $C_{1,2,3}$, nor have we considered other approximations in the spirit of the projection filter for which a qualitative result is not of particular interest. Nonetheless, there is no fundamental underlying limitation, and numerical evidence suggests that improved stability bounds could lead directly to quantitative estimates on filter robustness and approximation errors.

3.1.3. Notation. Let us fix some notation that will be used throughout the chapter. We already defined the solutions $\pi_{s,t}(\mu)$ and $\tilde{\pi}_{s,t}(\mu)$ of the Wonham equations. We will also use the Zakai equation

$$d\sigma_t = \Lambda^* \sigma_t dt + H \sigma_t dY_t, \quad \sigma_0 = \nu,$$

where as before $\sigma_t(\mu)$ and $\sigma_{s,t}(\mu)$ ($t \geq s$) denote the solutions at time t with the initial conditions $\sigma_0 = \mu$ and $\sigma_s = \mu$, respectively. Note that $\pi_{s,t}(\mu) = \sigma_{s,t}(\mu)/|\sigma_{s,t}(\mu)|$. The misspecified Zakai equation

$$d\tilde{\sigma}_t = \tilde{\Lambda}^* \tilde{\sigma}_t dt + \tilde{H} \tilde{\sigma}_t dY_t, \quad \tilde{\sigma}_0 = \nu,$$

and the associated solutions $\tilde{\sigma}_t(\mu)$ and $\tilde{\sigma}_{s,t}(\mu)$, are defined analogously. Finally, let us introduce some vector notation. For $x \in \mathbb{R}^d$, we denote by $|x|$ the ℓ_1 -norm, by $\|x\|$ the ℓ_2 -norm, and by $\|x\|_p$ the ℓ_p -norm. We write $x \succ y$ (resp. \prec, \succeq, \preceq) if $x_i > y_i$ ($<, \geq, \leq$) $\forall i$. We will repeatedly use the following spaces. Probability distributions on \mathbb{S} are elements of the simplex $\Delta^{d-1} = \{x \in \mathbb{R}^d : x \succeq 0, |x| = 1\}$. Usually, we

will be interested in the interior of the simplex $\mathcal{S}^{d-1} = \{x \in \mathbb{R}^d : x \succ 0, |x| = 1\}$. The space of vectors tangent to \mathcal{S}^{d-1} is denoted by $T\mathcal{S}^{d-1} = \{x \in \mathbb{R}^d : \sum_i x_i = 0\}$. The positive orthant will be denoted by $\mathbb{R}_{++}^d = \{x \in \mathbb{R}^d : x \succ 0\}$.

3.2. Stochastic semiflow of the Wonham filter

The main goal of this section is to establish some regularity properties of the solutions of the Wonham and Zakai equations. In particular, as we will want to calculate the derivative of the filter with respect to its initial condition, we have to establish that $\pi_{s,t}(\mu)$ is in fact differentiable. We will avoid problems at the boundary of the simplex by disposing of it altogether: we begin by proving that if $\mu \in \mathcal{S}^{d-1}$, then a.s. $\pi_{s,t}(\mu) \in \mathcal{S}^{d-1}$ for all times $t > s$.

LEMMA 3.2.1. $\mathbf{P}(\sigma_{s,t}(\mu) \in \mathbb{R}_{++}^d \text{ for all } \mu \in \mathbb{R}_{++}^d, 0 \leq s \leq t < \infty) = 1$.

PROOF. The following transformation (which is related to the pathwise filtering method) reduces the Zakai equation to a random differential equation. First, we write $\Lambda^* = S + T$ where S is the diagonal matrix with $S_{ii} = \lambda_{ii}$. Note that the matrix T has only nonnegative entries. We now perform the transformation $f_{s,t}(\mu) = L_{s,t}\sigma_{s,t}(\mu)$ where we have written

$$L_{s,t} = \exp\left(\left(\frac{1}{2}H^2 - S\right)(t-s) - H(Y_t - Y_s)\right).$$

Then $f_{s,t}(\mu)$ satisfies

$$(3.3) \quad \frac{df_{s,t}}{dt} = L_{s,t}TL_{s,t}^{-1}f_{s,t}, \quad f_{s,s} = \mu.$$

Now note that $t \mapsto B_t(\omega)$ is continuous for every $\omega \in \Omega$. Hence $t \mapsto L_{s,t}$, $t \mapsto L_{s,t}^{-1}$ are continuous in t and have strictly positive diagonal elements for every $\omega \in \Omega$. By standard arguments, there exists for every $\omega \in \Omega$, $\mu \in \mathbb{R}^d$ and $s \geq 0$ a unique solution $f_{s,t}(\mu)$ to (3.3) where $t \mapsto f_{s,t}(\mu)$ is a C^1 -curve. Moreover, note that $L_{s,t}TL_{s,t}^{-1}$ has nonnegative matrix elements for every $\omega \in \Omega$, $s \leq t < \infty$. Hence if $\mu \in \mathbb{R}_{++}^d$ then clearly $f_{s,t}(\mu)$ must be nondecreasing, i.e., $f_{s,t} \succeq f_{s,r}$ for every $t \geq r \geq s$ and $\omega \in \Omega$. But then \mathbb{R}_{++}^d must be forward invariant under (3.3), and as $L_{s,t}$ has strictly positive diagonal elements the result follows. \square

COROLLARY 3.2.2. $\mathbf{P}(\pi_{s,t}(\mu) \in \mathcal{S}^{d-1} \text{ for all } \mu \in \mathcal{S}^{d-1}, 0 \leq s \leq t < \infty) = 1$.

Let us now investigate the map $\sigma_{s,t}(\mu)$. As this map is linear in μ , we can write $\sigma_{s,t}(\mu) = U_{s,t}\mu$ a.s. where the $d \times d$ matrix $U_{s,t}$ is the solution of

$$(3.4) \quad dU_{s,t} = \Lambda^*U_{s,t} dt + HU_{s,t} dY_t, \quad U_{s,s} = I.$$

The following lemma establishes that $U_{s,t}$ defines a linear stochastic flow in \mathbb{R}^d .

LEMMA 3.2.3. For a.e. $\omega \in \Omega$ (i) $\sigma_{s,t}(\mu) = U_{s,t}\mu$ for all $s \leq t$; (ii) $U_{s,t}$ is continuous in (s, t) ; (iii) $U_{s,t}$ is invertible for all $s \leq t$, where $U_{s,t}^{-1}$ is given by

$$(3.5) \quad dU_{s,t}^{-1} = -U_{s,t}^{-1}\Lambda^* dt + U_{s,t}^{-1}H^2 dt - U_{s,t}^{-1}H dY_t, \quad U_{s,s}^{-1} = I;$$

(iv) $U_{r,t}U_{s,r} = U_{s,t}$ (and hence $U_{s,t}U_{s,r}^{-1} = U_{r,t}$) for all $s \leq r \leq t$.

PROOF. Continuity of $U_{s,t}$ (and $U_{s,t}^{-1}$) is a standard property of solution of Lipschitz stochastic differential equations. Invertibility of $U_{0,t}$ for all $0 \leq t < \infty$ is established in [Pro04, p. 326], and it is evident that $U_{s,t} = U_{0,t}U_{0,s}^{-1}$ satisfies (3.4). The remaining statements follow, where we can use continuity to remove the time dependence of the exceptional set as in the proof of [Pro04, p. 326]. \square

We now turn to the properties of the map $\pi_{s,t}(\mu)$.

LEMMA 3.2.4. *The Wonham filter generates a smooth stochastic semiflow in \mathcal{S}^{d-1} , i.e., the solutions $\pi_{s,t}(\mu)$ satisfy the following conditions:*

- (1) For a.e. $\omega \in \Omega$, $\pi_{s,t}(\mu) = \pi_{r,t}(\pi_{s,r}(\mu))$ for all $s \leq r \leq t$ and μ .
- (2) For a.e. $\omega \in \Omega$, $\pi_{s,t}(\mu)$ is continuous in (s, t, μ) .
- (3) For a.e. $\omega \in \Omega$, the injective map $\pi_{s,t}(\cdot) : \mathcal{S}^{d-1} \rightarrow \mathcal{S}^{d-1}$ is $C^\infty \forall s \leq t$.

PROOF. For $x \in \mathbb{R}_{++}^d$ define the function $\Sigma(x) = x/|x|$, so that we can write $\pi_{s,t}(\mu) = \Sigma(\sigma_{s,t}(\mu))$ ($\mu \in \mathcal{S}^{d-1}$). Note that Σ is smooth on \mathbb{R}_{++}^d . Hence continuity in (s, t, μ) and smoothness with respect to μ follow directly from the corresponding properties of $\sigma_{s,t}(\mu)$. The semiflow property $\pi_{s,t}(\mu) = \pi_{r,t}(\pi_{s,r}(\mu))$ follows directly from Lemma 3.2.3. It remains to prove injectivity.

Suppose that $\pi_{s,t}(\mu) = \pi_{s,t}(\nu)$ for some $\mu, \nu \in \mathcal{S}^{d-1}$. Then $U_{s,t}\mu/|U_{s,t}\mu| = U_{s,t}\nu/|U_{s,t}\nu|$, and as $U_{s,t}$ is invertible we have $\mu = (|U_{s,t}\mu|/|U_{s,t}\nu|)\nu$. But as μ and ν must lie in \mathcal{S}^{d-1} , it follows that $\mu = \nu$. Hence $\pi_t(\cdot)$ is injective. \square

REMARK 3.2.5. These results hold identically if we replace Λ by $\tilde{\Lambda}$, h by \tilde{h} . We will use the obvious notation $\tilde{\pi}_{s,t}(\mu)$, $\tilde{\sigma}_{s,t}(\mu)$, $\tilde{U}_{s,t}$, etc.

We finish this section by obtaining an expression for the approximation error in the case $\tilde{h} = h$; in fact, we will demonstrate the bound for this simple case in a more general setting than is considered in the following. Rather than considering the approximate Wonham filter with modified Λ , we will allow the approximate filter to have an arbitrary finite variation term, provided \mathcal{S}^{d-1} is left invariant.

PROPOSITION 3.2.6. *Let $\check{\pi}_t$ be a process with continuous paths in \mathcal{S}^{d-1} where*

$$(3.6) \quad d\check{\pi}_t = f(\check{\pi}_t) dt + (H - h^* \check{\pi}_t) \check{\pi}_t (dY_t - h^* \check{\pi}_t dt), \quad \check{\pi}_0 = \mu \in \mathcal{S}^{d-1}.$$

Then the difference between $\check{\pi}_t$ and the Wonham filter started at μ is a.s. given by

$$\check{\pi}_t - \pi_t(\mu) = \int_0^t D\pi_{s,t}(\check{\pi}_s) \cdot (f(\check{\pi}_s) - \Lambda^* \check{\pi}_s) ds,$$

where $D\pi_{s,t}(\mu) \cdot v$ is the derivative of $\pi_{s,t}(\mu)$ in the direction $v \in T\mathcal{S}^{d-1}$.

PROOF. Define the (scalar) process Γ_t by

$$\Gamma_t = \exp \left(\int_0^t h^* \check{\pi}_s dY_s - \frac{1}{2} \int_0^t (h^* \check{\pi}_s)^2 ds \right).$$

Using Itô's rule, we evaluate

$$(3.7) \quad \frac{d}{ds} (\Gamma_s U_{0,s}^{-1} \check{\pi}_s) = \Gamma_s U_{0,s}^{-1} (f(\check{\pi}_s) - \Lambda^* \check{\pi}_s).$$

Multiplying both sides by $U_{0,t}$, we obtain

$$\frac{d}{ds} (\Gamma_s U_{s,t} \check{\pi}_s) = \Gamma_s U_{s,t} (f(\check{\pi}_s) - \Lambda^* \check{\pi}_s).$$

Now introduce as before the map $\Sigma : \mathbb{R}_{++}^d \rightarrow \mathcal{S}^{d-1}$, $\Sigma(x) = x/|x|$, which is smooth on \mathbb{R}_{++}^d . Define the matrix $D\Sigma(x)$ with elements

$$[D\Sigma(x)]^{ij} = \frac{\partial \Sigma^i(x)}{\partial x^j} = \frac{1}{|x|} [\delta_{ij} - \Sigma^i(x)].$$

Note that $\Sigma(\alpha x) = \Sigma(x)$ for any $\alpha > 0$. Hence

$$\frac{d}{ds}\Sigma(U_{s,t}\check{\pi}_s) = \frac{d}{ds}\Sigma(\Gamma_s U_{s,t}\check{\pi}_s) = D\Sigma(\Gamma_s U_{s,t}\check{\pi}_s)\frac{d}{ds}(\Gamma_s U_{s,t}\check{\pi}_s).$$

But then we have, using $D\Sigma(\alpha x) = \alpha^{-1}D\Sigma(x)$ ($\alpha > 0$),

$$\begin{aligned} \frac{d}{ds}\Sigma(U_{s,t}\check{\pi}_s) &= D\Sigma(\Gamma_s U_{s,t}\check{\pi}_s)\Gamma_s U_{s,t}(f(\check{\pi}_s) - \Lambda^*\check{\pi}_s) \\ &= D\Sigma(U_{s,t}\check{\pi}_s)U_{s,t}(f(\check{\pi}_s) - \Lambda^*\check{\pi}_s). \end{aligned}$$

On the other hand, we obtain from the representation $\pi_{s,t}(\mu) = \Sigma(U_{s,t}\mu)$

$$D\pi_{s,t}(\mu) \cdot v = D\Sigma(U_{s,t}\mu)U_{s,t}v, \quad \mu \in \mathcal{S}^{d-1}, v \in T\mathcal{S}^{d-1}.$$

Note that $f(\check{\pi}_s) - \Lambda^*\check{\pi}_s$ is necessarily in $T\mathcal{S}^{d-1}$ as $\check{\pi}_t$ evolves in \mathcal{S}^{d-1} , so that $D\Sigma(U_{s,t}\check{\pi}_s)U_{s,t}(f(\check{\pi}_s) - \Lambda^*\check{\pi}_s) = D\pi_{s,t}(\check{\pi}_s) \cdot (f(\check{\pi}_s) - \Lambda^*\check{\pi}_s)$. Finally, note that

$$\int_0^t \frac{d}{ds}\Sigma(U_{s,t}\check{\pi}_s) ds = \Sigma(\check{\pi}_t) - \Sigma(U_{0,t}\check{\pi}_0) = \check{\pi}_t - \pi_t(\mu),$$

and the proof is complete. \square

COROLLARY 3.2.7. *Using the triangle inequality we obtain*

$$|\check{\pi}_t - \pi_t(\mu)| \leq \int_0^t |D\pi_{s,t}(\check{\pi}_s)| |f(\check{\pi}_s) - \Lambda^*\check{\pi}_s| ds,$$

where $|D\pi_{s,t}(\mu)| = \sup\{|D\pi_{s,t}(\mu) \cdot v| : v \in T\mathcal{S}^{d-1}, |v| = 1\}$. Moreover

$$|\check{\pi}_t - \pi_t(\nu)| \leq |\pi_t(\mu) - \pi_t(\nu)| + \int_0^t |D\pi_{s,t}(\check{\pi}_s)| |f(\check{\pi}_s) - \Lambda^*\check{\pi}_s| ds.$$

REMARK 3.2.8. The result of Corollary 3.2.7 is essentially identical to the deterministic example which we studied in section 3.1. The equality of h and \tilde{h} is key to this result: the stochastic integral term cancels in the proof, leaving only a finite variation part (note that we circumvented anticipativity in the intermediate steps of the proof by working initially with the inverse flow $U_{0,s}^{-1}$). In principle the result is applicable to any approximation of the Wonham filter that leaves untouched its stochastic integral part. This could be achieved, for example, using a finite-dimensional form of the projection filter method. In order to obtain useful error bounds for such approximations, however, one would need to have a fairly tight estimate on the derivative of the filter with respect to its initial condition, and as mentioned before we do not currently have such an estimate at our disposal. In the remainder of this chapter we will restrict ourselves to studying the robustness problem, but we will drop the requirement $h = \tilde{h}$ in section 3.4.

In the following, it will be convenient to turn around the role of the exact and approximate filters in Corollary 3.2.7, i.e., we will use the estimate

$$(3.8) \quad |\pi_t - \tilde{\pi}_t(\mu)| \leq |\tilde{\pi}_t(\nu) - \tilde{\pi}_t(\mu)| + \int_0^t |D\tilde{\pi}_{s,t}(\pi_s)| |(\Lambda^* - \tilde{\Lambda}^*)\pi_s| ds,$$

which holds provided $\tilde{h} = h$. The proof is identical to the one given above.

3.3. Exponential estimates for the derivative of the filter

In order for the bound (3.8) to be useful, we must have an exponential estimate for $|D\tilde{\pi}_{s,t}(\cdot)|$. The goal of this section is to obtain such an estimate. We proceed in two steps. First, we use the methods introduced in section 2.3 to obtain a pathwise exponential estimate for $|D\pi_{0,t}(\nu)|$. As the laws of the observation process for different initial measures, jump rates and observation functions are equivalent, we can extend this a.s. bound to $|D\tilde{\pi}_{s,t}(\mu)|$. We find, however, that the proportionality constant in the exponential estimate depends on μ and diverges as μ approaches the boundary of the simplex. This makes a pathwise bound on $|D\tilde{\pi}_{s,t}(\pi_s)|$ difficult to obtain, as π_s can get arbitrarily close to the boundary of the simplex on the infinite time interval. Instead, we proceed to find a uniform bound on $\mathbf{E}|D\tilde{\pi}_{s,t}(\pi_s)|$ by bounding the expectation of the prefactor.

We will use the following Lemma, which is similar to [BCL04, Lemma 5.7]. The coupling proof given here is new, however.

LEMMA 3.3.1. *For any $0 \leq s \leq t$, we have the following bound a.s.:*

$$\sum_{k=1}^d |\mathbf{P}(X_{t-s} = a_k | \mathcal{F}_t^Y, X_t = a_i) - \mathbf{P}(X_{t-s} = a_k | \mathcal{F}_t^Y)| \leq 2 \exp\left(-2s \min_{p \neq q} \sqrt{\lambda_{pq} \lambda_{qp}}\right).$$

PROOF. We proceed exactly as in the proof of Proposition 2.3.1, only using the time-reversed signal \bar{X}_t rather than the signal itself. Evidently $\mathbf{P}(X_{t-s} = a_k | \mathcal{F}_t^Y)$ is the expectation of \bar{X}_s under the conditional measure $\Pi_t(\cdot, Y_{[0,t]})$, whereas on the other hand $\mathbf{P}(X_{t-s} = a_k | \mathcal{F}_t^Y, X_t = a_i)$ is the conditional expectation of \bar{X}_s under $\Pi_t(\cdot, Y_{[0,t]})$ with respect to the event $\bar{X}_0 = a_i$:

$$\mathbf{P}(X_{t-s} = a_k | \mathcal{F}_t^Y, X_t = a_i) = \frac{\mathbf{P}(X_{t-s} = a_k, X_t = a_i | \mathcal{F}_t^Y)}{\mathbf{P}(X_t = a_i | \mathcal{F}_t^Y)}.$$

But as \bar{X} is a Markov process under $\Pi_t(\cdot, y)$, the latter is equal to the expectation of a Markov process with the same transition intensities as the time-reversed conditional signal, but with initial measure $\bar{X}_0 \sim \delta_{\{a_i\}}$. The result now follows identically as in the proofs of Proposition 2.3.1 and Corollary 2.3.2. \square

We can now obtain a bound on the derivative of the filter.

PROPOSITION 3.3.2. *Let $\lambda_{ij} > 0 \forall i \neq j$ and $\nu \in \mathcal{S}^{d-1}$, $v \in T\mathcal{S}^{d-1}$. Then a.s.*

$$|D\pi_t(\nu) \cdot v| \leq \sum_k \frac{|v^k|}{\nu^k} \exp\left(-2t \min_{p \neq q} \sqrt{\lambda_{pq} \lambda_{qp}}\right).$$

PROOF. Recall from section 2.3.1 that we can write

$$(3.9) \quad \pi_t^i(\mu) = \frac{\mathbf{E}\left(\frac{d\mu}{d\nu}(X_0) I_{X_t=a_i} | \mathcal{F}_t^Y\right)}{\mathbf{E}\left(\frac{d\mu}{d\nu}(X_0) | \mathcal{F}_t^Y\right)} = \frac{\sum_{j=1}^d (\mu^j / \nu^j) \mathbf{P}(X_0 = a_j, X_t = a_i | \mathcal{F}_t^Y)}{\sum_{j=1}^d (\mu^j / \nu^j) \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y)}.$$

It is straightforward to calculate the directional derivative of this expression:

$$(D\pi_t(\mu) \cdot v)^i = \frac{\sum_j \frac{v^j}{\nu^j} (\mathbf{P}(X_0 = a_j, X_t = a_i | \mathcal{F}_t^Y) - \pi_t^i(\mu) \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y))}{\sum_j \frac{\mu^j}{\nu^j} \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y)}.$$

Setting $\mu = \nu$, we obtain after some simple manipulations

$$(D\pi_t(\nu) \cdot v)^i = \pi_t^i(\nu) \sum_j \frac{v^j}{\nu^j} \{ \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y, X_t = a_i) - \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y) \}.$$

By Lemma 3.3.1 and the fact that $|\mu_1 - \mu_2| = 2 \max_{A \subset S} |\mu_1(A) - \mu_2(A)|$, we obtain

$$|(D\pi_t(\nu) \cdot v)^i| \leq \pi_t^i(\nu) \sum_j \frac{|v^j|}{\nu^j} \exp \left(-2t \min_{p \neq q} \sqrt{\lambda_{pq} \lambda_{qp}} \right).$$

The result now follows immediately. \square

To obtain this bound we had to use the true initial distribution ν , jump rates λ_{ij} and observation function h . However, the almost sure nature of the result allows us to drop these requirements.

COROLLARY 3.3.3. *Let $\tilde{\lambda}_{ij} > 0 \forall i \neq j$ and $\mu \in \mathcal{S}^{d-1}$, $v \in T\mathcal{S}^{d-1}$. Then a.s.*

$$(3.10) \quad |D\tilde{\pi}_{s,t}(\mu) \cdot v| \leq \sum_k \frac{|v^k|}{\mu^k} \exp \left(-2(t-s) \min_{p \neq q} \sqrt{\tilde{\lambda}_{pq} \tilde{\lambda}_{qp}} \right).$$

Moreover, the result still holds if μ, v are \mathcal{F}_s^Y -measurable random variables with values a.s. in \mathcal{S}^{d-1} and $T\mathcal{S}^{d-1}$, respectively.

PROOF. Let $\tilde{\mathbf{P}}$ be the measure under which X_t is a Markov process with transition intensities $\tilde{\lambda}_{ij}$ and initial measure μ , and such that $d\tilde{B}_t = dY_t - \tilde{h}(X_t) dt$ defines a Wiener process \tilde{B}_t independent from X_t . Such a measure is easily constructed using Girsanov's theorem and Propositions 2.1.1 and 2.1.3, and is certainly equivalent to \mathbf{P} . Then $\tilde{\pi}_t^i(\mu)$ is precisely the optimal filter $\tilde{\mathbf{P}}(X_t = a_i | \mathcal{F}_t^Y)$, and we can invoke Proposition 3.3.2 under the new measure. But this bound holds with unit probability, and $\tilde{\mathbf{P}} \sim \mathbf{P}$. Hence we have established the result for $s = 0$. The result for $s > 0$ follows directly as the Wonham equation is time homogeneous.

To show that the result still holds when μ, v are random, note that $\tilde{\pi}_{s,t}$ only depends on the observation increments in the interval $[s, t]$, i.e., $D\tilde{\pi}_{s,t}(\mu) \cdot v$ is $\mathcal{F}_{[s,t]}^Y$ -measurable where $\mathcal{F}_{[s,t]}^Y = \sigma\{Y_r - Y_s : s \leq r \leq t\}$. Under the reference measure \mathbf{Q} , Y is a Wiener process and hence $\mathcal{F}_{[s,t]}^Y$ and \mathcal{F}_s^Y are independent. It follows from the bound with constant μ, v that

$$\mathbf{E}_{\mathbf{Q}}(I_{|D\tilde{\pi}_{s,t}(\mu) \cdot v| \leq (*)} | \sigma\{\mu, v\}) = 1 \quad \mathbf{Q}\text{-a.s.},$$

where $(*)$ is the right-hand side of (3.10). Hence $\mathbf{E}_{\mathbf{Q}}(I_{|D\tilde{\pi}_{s,t}(\mu) \cdot v| \leq (*)}) = 1$, and the statement follows from $\mathbf{P} \sim \mathbf{Q}$. \square

Next, let us obtain a filter stability bound. Unlike the bound of Corollary 2.3.2 which has a constant prefactor, this bound vanishes as $\mu_1 \rightarrow \mu_2$.

PROPOSITION 3.3.4. *Let $\tilde{\lambda}_{ij} > 0 \forall i \neq j$ and $\mu_1, \mu_2 \in \mathcal{S}^{d-1}$. Then a.s.*

$$|\tilde{\pi}_{s,t}(\mu_2) - \tilde{\pi}_{s,t}(\mu_1)| \leq C |\mu_2 - \mu_1| \exp \left(-2(t-s) \min_{p \neq q} \sqrt{\tilde{\lambda}_{pq} \tilde{\lambda}_{qp}} \right),$$

where $C = \max\{1/\mu_1^k, 1/\mu_2^k : k = 1, \dots, d\}$.

PROOF. Define $\gamma(u) = \tilde{\pi}_{s,t}(\mu_1 + u(\mu_2 - \mu_1))$, $u \in [0, 1]$. Then

$$\tilde{\pi}_{s,t}(\mu_2) - \tilde{\pi}_{s,t}(\mu_1) = \int_0^1 \frac{d\gamma}{du} du = \int_0^1 D\tilde{\pi}_{s,t}(\mu_1 + u(\mu_2 - \mu_1)) \cdot (\mu_2 - \mu_1) du.$$

Using the triangle inequality, we obtain

$$|\tilde{\pi}_{s,t}(\mu_2) - \tilde{\pi}_{s,t}(\mu_1)| \leq \sup_{u \in [0,1]} |D\tilde{\pi}_{s,t}(\mu_1 + u(\mu_2 - \mu_1)) \cdot (\mu_2 - \mu_1)|.$$

The result now follows from Corollary 3.3.3. \square

REMARK 3.3.5. Corollary 2.3.2 and Proposition 3.3.4 can be combined:

$$|\pi_t(\mu_1) - \pi_t(\mu_2)| \leq \left[\left(\max_k \left\{ \frac{1}{\mu_1^k}, \frac{1}{\mu_2^k} \right\} |\mu_2 - \mu_1| \right) \wedge 2 \right] e^{-2t \min_{p \neq q} \sqrt{\lambda_{pq} \lambda_{qp}}}.$$

To my knowledge, this is to date the best nonasymptotic filter stability bound that is available for the Wonham filter.

Corollary 3.3.3 and Proposition 3.3.4 are exactly what we need to establish boundedness of (3.8). Note, however, that the right-hand side of (3.10) is proportional to $1/\mu^i$, and we must estimate $|D\tilde{\pi}_{s,t}(\pi_s)|$. Though we established in Section 3.2 that π_s cannot hit the boundary of the simplex in finite time, it can get arbitrarily close to the boundary during the infinite time interval, thus rendering the right-hand side of (3.10) arbitrarily large. If we can establish that $\sup_{s \geq 0} \mathbf{E}(1/\min_k \pi_s^k) < \infty$, however, then we can control $\mathbf{E}|D\tilde{\pi}_{s,t}(\pi_s)|$ to obtain a useful bound. We begin with an auxiliary integrability property of π_t :

LEMMA 3.3.6. *Let $\nu \in \mathcal{S}^{d-1}$ and $T < \infty$. Then*

$$\mathbf{E} \int_0^T (\pi_s^i)^{-k} ds < \infty, \quad \forall i = 1, \dots, d, k \geq 1.$$

PROOF. Applying Itô's rule to the Wonham equation gives

$$d \log \pi_t^i = \left(\lambda_{ii} - \frac{1}{2}(h^i - h^* \pi_t)^2 \right) dt + \sum_{j \neq i} \lambda_{ji} \frac{\pi_t^j}{\pi_t^i} dt + (h^i - h^* \pi_t) d\overline{W}_t,$$

where the innovation $d\overline{W}_t = dY_t - h^* \pi_t dt$ is a Wiener process (recall Proposition 1.1.14). The application of Itô's rule is justified by a standard localization argument, as π_t is in \mathcal{S}^{d-1} for all $t \geq 0$ a.s. and $\log x$ is smooth in $(0, 1)$. As $\lambda_{ij} \geq 0$ for $j \neq i$,

$$-k \log \pi_t^i \leq -k \log \nu^i - k \lambda_{ii} t + \frac{k}{2} \max_j (h^i - h^j)^2 t - k \int_0^t (h^i - h^* \pi_s) d\overline{W}_s.$$

But as $h^i - h^* \pi_t$ is bounded, Novikov's condition is satisfied and hence

$$\mathbf{E} \exp \left(-k \int_0^t (h^i - h^* \pi_s) d\overline{W}_s - \frac{k^2}{2} \int_0^t (h^i - h^* \pi_s)^2 ds \right) = 1.$$

Estimating the time integral, we obtain

$$\mathbf{E}(\pi_t^i)^{-k} \leq (\nu^i)^{-k} \exp \left(-k \lambda_{ii} t + \frac{1}{2} k(k+1) \max_j (h^i - h^j)^2 t \right).$$

The Lemma now follows by the Tonelli theorem, as $(\pi_s^i)^{-k} \geq 0$ a.s. \square

We are now in a position to bound $\sup_{t \geq 0} \mathbf{E}(1/\min_i \pi_t^i)$.

PROPOSITION 3.3.7. *Let $\nu \in \mathcal{S}^{d-1}$ and suppose that $\lambda_{ij} > 0 \forall i \neq j$. Then*

$$\sup_{t \geq 0} \mathbf{E} \left(\frac{1}{\min_i \pi_t^i} \right) < \infty.$$

PROOF. By Itô's rule and using the standard localization argument, we obtain

$$\begin{aligned} (\pi_t^i)^{-1} &= (\nu^i)^{-1} - \int_0^t \lambda_{ii} (\pi_s^i)^{-1} ds - \int_0^t (\pi_s^i)^{-2} \sum_{j \neq i} \lambda_{ji} \pi_s^j ds \\ &\quad - \int_0^t (\pi_s^i)^{-1} (h^i - h^* \pi_s) d\bar{W}_s + \int_0^t (\pi_s^i)^{-1} (h^i - h^* \pi_s)^2 ds. \end{aligned}$$

Using Lemma 3.3.6 we find

$$\mathbf{E} \int_0^t (\pi_s^i)^{-2} (h^i - h^* \pi_s)^2 ds \leq \max_j (h^i - h^j)^2 \mathbf{E} \int_0^t (\pi_s^i)^{-2} ds < \infty,$$

so the expectation of the stochastic integral term vanishes. Using the Tonelli theorem, we can thus write

$$\begin{aligned} \mathbf{E}((\pi_t^i)^{-1}) &= (\nu^i)^{-1} - \int_0^t \lambda_{ii} \mathbf{E}((\pi_s^i)^{-1}) ds \\ &\quad - \int_0^t \mathbf{E} \left((\pi_s^i)^{-2} \sum_{j \neq i} \lambda_{ji} \pi_s^j \right) ds + \int_0^t \mathbf{E}((\pi_s^i)^{-1} (h^i - h^* \pi_s)^2) ds. \end{aligned}$$

Taking the derivative and estimating each of the terms, we obtain

$$\frac{dM_t^i}{dt} \leq -\min_{j \neq i} \lambda_{ji} (M_t^i)^2 + \left(|\lambda_{ii}| + \min_{j \neq i} \lambda_{ji} + \max_j (h^i - h^j)^2 \right) M_t^i,$$

where we have written $M_t^i = \mathbf{E}((\pi_t^i)^{-1})$ and we have used $(M_t^i)^2 \leq \mathbf{E}(\pi_t^i)^{-2}$ by Jensen's inequality. Using the estimate

$$-K_1^i (M_t^i)^2 + K_2^i M_t^i \leq -K_2^i M_t^i + \frac{(K_2^i)^2}{K_1^i} \quad \text{for } K_1^i > 0,$$

we now obtain

$$\frac{dM_t^i}{dt} \leq K_2^i \left(\frac{K_2^i}{K_1^i} - M_t^i \right), \quad K_2^i = |\lambda_{ii}| + \min_{j \neq i} \lambda_{ji} + \max_j (h^i - h^j)^2,$$

where $K_1^i = \min_{j \neq i} \lambda_{ji} > 0$. Consequently we obtain

$$M_t^i \leq e^{-K_2^i t} (\nu^i)^{-1} + \frac{(K_2^i)^2}{K_1^i} e^{-K_2^i t} \int_0^t e^{K_2^i s} ds = e^{-K_2^i t} (\nu^i)^{-1} + \frac{K_2^i}{K_1^i} (1 - e^{-K_2^i t}).$$

We can now estimate

$$\sup_{t \geq 0} \mathbf{E} \left(\frac{1}{\min_i \pi_t^i} \right) \leq \sum_{i=1}^d \sup_{t \geq 0} \mathbf{E} \left(\frac{1}{\pi_t^i} \right) \leq \sum_{i=1}^d \left(\frac{1}{\nu^i} \vee \frac{K_2^i}{K_1^i} \right) < \infty,$$

which is what we set out to prove. \square

We can now prove Theorem 3.1.5 for the special case $\tilde{h} = h$. Using (3.8), Corollary 3.3.3, Proposition 3.3.4, and Proposition 3.3.7, we obtain

$$\begin{aligned} \mathbf{E}|\pi_t - \tilde{\pi}_t(\mu)| &\leq |\mu - \nu| \max_k \left\{ \frac{1}{\mu^k} \vee \frac{1}{\nu^k} \right\} \exp \left(-2t \min_{p \neq q} \sqrt{\tilde{\lambda}_{pq} \tilde{\lambda}_{qp}} \right) \\ &\quad + |\Lambda^* - \tilde{\Lambda}^*| \sup_{s \geq 0} \mathbf{E}(1 / \min_k \pi_s^k) \int_0^t \exp \left(-2(t-s) \min_{p \neq q} \sqrt{\tilde{\lambda}_{pq} \tilde{\lambda}_{qp}} \right) ds, \end{aligned}$$

where $|\Lambda^* - \tilde{\Lambda}^*| = \sup\{ |(\Lambda^* - \tilde{\Lambda}^*)\mu| : \mu \in \mathcal{S}^{d-1} \}$. Thus

$$\mathbf{E}|\pi_t - \tilde{\pi}_t(\mu)| \leq |\mu - \nu| \max_k \left\{ \frac{1}{\mu^k} \vee \frac{1}{\nu^k} \right\} e^{-\beta t} + |\Lambda^* - \tilde{\Lambda}^*| \frac{\sup_{s \geq 0} \mathbf{E}(1 / \min_k \pi_s^k)}{\beta},$$

where we have written $\beta = 2 \min_{p \neq q} (\tilde{\lambda}_{pq} \tilde{\lambda}_{qp})^{1/2}$. The result follows directly using $\|\pi_t - \tilde{\pi}_t(\mu)\|^2 \leq |\pi_t - \tilde{\pi}_t(\mu)|$ (as $|\pi_t^i - \tilde{\pi}_t(\mu)^i| \leq 1$).

REMARK 3.3.8. As the constants in the bound above are easily computable, it is interesting compare the bound to simulations. Unfortunately, it appears that this bound is essentially useless as a quantitative bound, except in the case of extremely low signal-to-noise. In the high signal-to-noise case, the error between the approximate and exact filters appears to be much smaller than the bound above for virtually every sample path. Numerical simulations suggest that a major reason for this fact is the rather disappointing divergence of the bound of Proposition 3.3.2 at the boundary of the simplex. By taking the expectation of $1/\pi_t^i$, we have managed to keep things finite. On the other hand, when the signal-to-noise ratio is reasonably high, the filter will spend most of its time near the boundary of the simplex so that the expectation of $1/\pi_t^i$ and hence our error bound is huge.

Is this a real effect or an artefact of our method? After all, previous filter stability bounds suggested that the prefactor in the exponential decay of $|\pi_t(\mu) - \pi_t(\nu)|$ should blow up near the boundary of the simplex, see [BCL04] or Proposition 3.3.4, but Corollary 2.3.2 corrects this problem. Similarly one might think that the divergence of the bound of Proposition 3.3.2 could be an artefact of our method of proof. Unfortunately, a simple simulation suggests that this is not the case. In figure 3.1 we show 250 sample paths of the derivative of the filter for a symmetric binary signal with unit transition intensities and observation function $h(a_1) = -5$, $h(a_2) = 5$. Evidently the bound of Proposition 3.3.2 is quite decent at the center of the simplex. On the boundary, however, rare but extremely large excursions are observed, which is not conclusive but suggestive of an unbounded random variable.

On the other hand, the expectation of $|D\pi_t(\mu)|$ appears to behave very nicely, even in the worst case scenario where μ and ν are mutually singular. Our approach, which is based on an a.s. bound on $|D\pi_t(\mu)|$, could never capture this nice behavior: after all, an a.s. bound has to guarantee that *every* sample path lies below it, which is an extremely harsh requirement. Hence it appears that a different method is called for, which bounds directly the expectation of $|D\pi_t(\mu)|$. This brings up the question: with respect to what measure should the expectation be taken? A simple result is easy to obtain: following the proof of Proposition 3.3.2, note that

$$(D\pi_t(\nu) \cdot v)^i = \sum_j \frac{v^j}{\nu^j} \{ \mathbf{P}(X_t = a_i | \mathcal{F}_t^Y, X_0 = a_j) - \pi_t^i(\nu) \} \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y).$$




FIGURE 3.1. Simulation of $|D\pi_t(\mu)|$ for the binary signal model in Remark 3.3.8. The top figure is for $\mu^1 = \mu^2 = 0.5$, the bottom figure is for $\mu^1 = 1, \mu^2 = 0$. The actual initial distribution is $\nu^1 = 0, \nu^2 = 1$. In each case 250 sample paths of $|D\tilde{\pi}_{s,t}(\mu)|$ are shown in green, and their average is shown in red. The blue line is the bound of Proposition 3.3.2.

Hence we can obtain, using the triangle inequality and Corollary 2.3.2,

$$\mathbf{E}|D\pi_t(\nu) \cdot v| \leq 2|v| \exp\left(-2t \min_{p \neq q} \sqrt{\lambda_{pq}\lambda_{qp}}\right).$$

A bound on $\mathbf{E}|D\pi_t(\mu) \cdot v|$ is much more difficult to obtain, however. If such a bound could be obtained, on the other hand, then this could open the door to quantitative bounds on the approximation error in nonlinear filtering.

3.4. Proof of the main result

We are now ready to proceed to the general case where the initial density, the transition intensities matrix and the observation function can all be misspecified. The simplicity of the special case $\tilde{h} = h$ that we have treated up to this point is due to the fact that in the calculation of (3.7), the stochastic integral term drops out and we can proceed with the calculation using only ordinary calculus. In the general case we can not get rid of the stochastic integral, and hence we run into anticipativity problems in the next step of the calculation.

We solve this problem by using anticipative stochastic integrals in the sense of Skorokhod, rather than the usual Itô integral (which is a special case of the Skorokhod integral defined for adapted processes only). Though the Skorokhod integral is more general than the Itô integral in the sense that it allows some anticipating integrands, it is less general in that we have to integrate against a Wiener process (rather than against an arbitrary semimartingale), and that the integrands should be functionals of the driving Wiener process.

In our setup, the most convenient way to apply this theory is to operate exclusively under the measure \mathbf{Q} which is defined as follows:

$$(3.11) \quad \frac{d\mathbf{P}}{d\mathbf{Q}} = |\sigma_T(\nu)| = |\sigma_T|$$

for some fixed $T < \infty$. Recall that this \mathbf{Q} is precisely the reference measure restricted to \mathcal{F}_T^Y , see Corollary 1.1.15. We will not feel bad about denoting the restricted reference measure by the same symbol, as we will only work with \mathcal{F}_T^Y -measurable random variables from this point onwards: note that both the approximate and exact filters are functionals of the observations only.

Note in particular that $Y_{[0,T]}$ is a Wiener process under \mathbf{Q} . Hence we can interpret the Wonham and Zakai equations under \mathbf{Q} as Itô stochastic differential equations which are driven by the Wiener process Y_t . The usual Skorokhod integral coincides with the Itô integral for adapted processes, but is only defined when the integrator is a Wiener process—hence the importance of working under \mathbf{Q} (this is in contrast to the Itô integral, for which any semimartingale can serve as an integrator). We can thus reinterpret the Wonham and Zakai equations as Skorokhod stochastic differential equations. What we have gained by this is that we are now allowed to perform anticipative transformations. When we wish to calculate expectations with respect to \mathbf{P} at the end of the day, we can do this by using explicitly the above expression for $d\mathbf{P}/d\mathbf{Q}$.

Our setup is further detailed in appendix A, together with a review of the relevant results from the Malliavin calculus and anticipative stochastic calculus. Below we will use the notations and results from this appendix without further comment. We will also have the need for several estimates of a technical nature, which are not very insightful to the structure of the proofs. These results have been relegated to section 3.5 of this chapter. This organization will allow us to proceed through the proofs below with minimal interruptions, but the reader should probably take at least a brief glance at appendix A at this point.

We begin by obtaining an anticipative version of Proposition 3.2.6.

PROPOSITION 3.4.1. *The difference between π_t and $\tilde{\pi}_t$ satisfies*

$$\begin{aligned} \pi_t - \tilde{\pi}_t &= \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_\Lambda \pi_r \, dr + \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) \, dY_r \\ &\quad - \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \left[h^* \pi_r (H - h^* \pi_r) \pi_r - \tilde{h}^* \pi_r (\tilde{H} - \tilde{h}^* \pi_r) \pi_r \right] \, dr \\ &\quad + \frac{1}{2} \int_0^t \left[D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (H - h^* \pi_r) \pi_r - D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (\tilde{H} - \tilde{h}^* \pi_r) \pi_r \right] \, dr, \end{aligned}$$

where the stochastic integral is a Skorokhod integral and we have written $\Delta_\Lambda = \Lambda^* - \tilde{\Lambda}^*$, $\Delta_H(\pi) = (H - h^* \pi) \pi - (\tilde{H} - \tilde{h}^* \pi) \pi$, and $D^2 \tilde{\pi}_{r,t}(\mu) \cdot v$ is the directional derivative of $D\tilde{\pi}_{r,t}(\mu) \cdot v$ with respect to $\mu \in \mathcal{S}^{d-1}$ in the direction $v \in T\mathcal{S}^{d-1}$.

Note that this result is precisely of the form one would expect. The first two lines follow the formula for the distance between two flows as one would guess, e.g., from the discussion in section 3.1.1; the last line is an “Itô correction term” which contains second derivatives of the filter with respect to its initial condition.

PROOF. Fix $T > t$. We begin by evaluating, using Itô’s rule and (3.5),

$$\begin{aligned} \tilde{U}_{0,s}^{-1} U_{0,s} \nu &= \nu + \int_0^s \tilde{U}_{0,r}^{-1} (\Lambda^* - \tilde{\Lambda}^*) U_{0,r} \nu \, dr \\ &\quad - \int_0^s \tilde{U}_{0,r}^{-1} \tilde{H} (H - \tilde{H}) U_{0,r} \nu \, dr + \int_0^s \tilde{U}_{0,r}^{-1} (H - \tilde{H}) U_{0,r} \nu \, dY_r. \end{aligned}$$

Now multiply from the left by $\tilde{U}_{0,t}$; we wish to use Lemma A.4.2 to bring $\tilde{U}_{0,t}$ into the Skorokhod integral term, i.e., we claim that

$$\begin{aligned} \tilde{U}_{s,t}U_{0,s}\nu &= \tilde{U}_{0,t}\nu + \int_0^s \tilde{U}_{r,t}(\Lambda^* - \tilde{\Lambda}^*)U_{0,r}\nu dr - \int_0^s \tilde{U}_{r,t}\tilde{H}(H - \tilde{H})U_{0,r}\nu dr \\ &\quad + \int_0^s \tilde{U}_{r,t}(H - \tilde{H})U_{0,r}\nu dY_r + \int_0^s (\mathbf{D}_r\tilde{U}_{0,t})\tilde{U}_{0,r}^{-1}(H - \tilde{H})U_{0,r}\nu dr. \end{aligned}$$

To justify this expression we need to verify the integrability conditions of Lemma A.4.2. Note that all matrix elements of $\tilde{U}_{s,t}$ are in $\mathbb{D}^\infty \forall 0 \leq s \leq t < T$, and that

$$\mathbf{D}_r\tilde{U}_{s,t} = \begin{cases} 0 & \text{a.e. } r \notin [s, t], \\ \tilde{U}_{r,t}\tilde{H}\tilde{U}_{s,r} & \text{a.e. } r \in [s, t]. \end{cases}$$

This follows directly from Proposition A.2.1 and Lemma 3.2.3 (note that the same result holds for $U_{s,t}$ if we replace \tilde{H} by H and \tilde{U} by U). Once we plug this result into the expression above, the corresponding integrability conditions can be verified explicitly, see Lemma 3.5.1, and hence we have verified that

$$\tilde{U}_{s,t}U_{0,s}\nu = \tilde{U}_{0,t}\nu + \int_0^s \tilde{U}_{r,t}(\Lambda^* - \tilde{\Lambda}^*)U_{0,r}\nu dr + \int_0^s \tilde{U}_{r,t}(H - \tilde{H})U_{0,r}\nu dY_r.$$

Next we would like to apply the anticipating Itô rule, Proposition A.5.1, with the function $\Sigma : \mathbb{R}_{++}^d \rightarrow \mathcal{S}^{d-1}$, $\Sigma(x) = x/|x|$. To this end we have to verify a set of technical conditions, see Lemma 3.5.2. We obtain

$$\begin{aligned} \Sigma(\tilde{U}_{s,t}U_{0,s}\nu) &= \Sigma(\tilde{U}_{0,t}\nu) + \int_0^s D\Sigma(\tilde{U}_{r,t}U_{0,r}\nu)\tilde{U}_{r,t}(\Lambda^* - \tilde{\Lambda}^*)U_{0,r}\nu dr \\ &\quad + \frac{1}{2} \sum_{k,\ell} \int_0^s \frac{\partial^2 \Sigma}{\partial x^k \partial x^\ell}(\tilde{U}_{r,t}U_{0,r}\nu)(\nabla_r \tilde{U}_{r,t}U_{0,r}\nu)^k (\tilde{U}_{r,t}(H - \tilde{H})U_{0,r}\nu)^\ell dr \\ &\quad + \int_0^s D\Sigma(\tilde{U}_{r,t}U_{0,r}\nu)\tilde{U}_{r,t}(H - \tilde{H})U_{0,r}\nu dY_r. \end{aligned}$$

We need to evaluate $\nabla_r \tilde{U}_{r,t}U_{0,r}\nu$. Using Prop. A.1.2 and Lemma A.1.3, we calculate

$$\lim_{\varepsilon \searrow 0} \mathbf{D}_r \tilde{U}_{r+\varepsilon,t}U_{0,r+\varepsilon}\nu = \lim_{\varepsilon \searrow 0} \tilde{U}_{r+\varepsilon,t}U_{r,r+\varepsilon}HU_{0,r}\nu = \tilde{U}_{r,t}HU_{0,r}\nu,$$

and similarly

$$\lim_{\varepsilon \searrow 0} \mathbf{D}_r \tilde{U}_{r-\varepsilon,t}U_{0,r-\varepsilon}\nu = \lim_{\varepsilon \searrow 0} \tilde{U}_{r,t}\tilde{H}\tilde{U}_{r-\varepsilon,r}U_{0,r-\varepsilon}\nu = \tilde{U}_{r,t}\tilde{H}U_{0,r}\nu.$$

After some rearranging, we obtain

$$\begin{aligned} \Sigma(\tilde{U}_{s,t}U_{0,s}\nu) &= \Sigma(\tilde{U}_{0,t}\nu) + \int_0^s D\Sigma(\tilde{U}_{r,t}U_{0,r}\nu)\tilde{U}_{r,t}(\Lambda^* - \tilde{\Lambda}^*)U_{0,r}\nu dr \\ &\quad + \frac{1}{2} \sum_{k,\ell} \int_0^s \frac{\partial^2 \Sigma}{\partial x^k \partial x^\ell}(\tilde{U}_{r,t}U_{0,r}\nu)(\tilde{U}_{r,t}HU_{0,r}\nu)^k (\tilde{U}_{r,t}HU_{0,r}\nu)^\ell dr \\ &\quad - \frac{1}{2} \sum_{k,\ell} \int_0^s \frac{\partial^2 \Sigma}{\partial x^k \partial x^\ell}(\tilde{U}_{r,t}U_{0,r}\nu)(\tilde{U}_{r,t}\tilde{H}U_{0,r}\nu)^k (\tilde{U}_{r,t}\tilde{H}U_{0,r}\nu)^\ell dr \\ &\quad + \int_0^s D\Sigma(\tilde{U}_{r,t}U_{0,r}\nu)\tilde{U}_{r,t}(H - \tilde{H})U_{0,r}\nu dY_r. \end{aligned}$$

From this point onwards we will set $s = t$. We will need (on \mathbb{R}_{++}^d)

$$D^2\Sigma^{ik\ell}(x) = \frac{\partial^2\Sigma^i(x)}{\partial x^k\partial x^\ell} = -\frac{1}{|x|}(D\Sigma^{ik}(x) + D\Sigma^{i\ell}(x)).$$

Recall that $D\Sigma(\alpha x) = \alpha^{-1}D\Sigma(x)$; it follows that also $D^2\Sigma(\alpha x) = \alpha^{-2}D^2\Sigma(x)$ for $\alpha > 0$. Using these expressions with $\alpha = |U_{0,r\nu}|$, we get

$$\begin{aligned} \pi_t - \tilde{\pi}_t &= \int_0^t D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}\Delta_\Lambda\pi_r dr + \int_0^t D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}(H - \tilde{H})\pi_r dY_r \\ &\quad + \frac{1}{2} \sum_{k,\ell} \int_0^t \frac{\partial^2\Sigma}{\partial x^k\partial x^\ell}(\tilde{U}_{r,t}\pi_r)(\tilde{U}_{r,t}H\pi_r)^k(\tilde{U}_{r,t}H\pi_r)^\ell dr \\ &\quad - \frac{1}{2} \sum_{k,\ell} \int_0^t \frac{\partial^2\Sigma}{\partial x^k\partial x^\ell}(\tilde{U}_{r,t}\pi_r)(\tilde{U}_{r,t}\tilde{H}\pi_r)^k(\tilde{U}_{r,t}\tilde{H}\pi_r)^\ell dr. \end{aligned}$$

Next we want to express the integrands in terms of $D\tilde{\pi}_{r,t}(\pi_r) \cdot v$, etc., rather than in terms of $D\Sigma(x)$. Recall that $D\tilde{\pi}_{r,t}(\pi_r) \cdot v = D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}v$ when $v \in T\mathcal{S}^{d-1}$. Similar terms appear in the expression above, but, e.g., $\tilde{H}\pi_r \notin T\mathcal{S}^{d-1}$. To rewrite the expression in the desired form, we use that $D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}\pi_r = 0$. Hence

$$D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}\tilde{H}\pi_r = D\Sigma(\tilde{U}_{r,t}\pi_r)\tilde{U}_{r,t}(\tilde{H} - \tilde{h}^*\pi_r)\pi_r = D\tilde{\pi}_{r,t}(\pi_r) \cdot (\tilde{H} - \tilde{h}^*\pi_r)\pi_r$$

and similarly for the other terms. Note also that

$$\sum_k D^2\Sigma^{ik\ell}(\tilde{U}_{r,t}\pi_r)(\tilde{U}_{r,t}\pi_r)^k = -D\Sigma^{i\ell}(\tilde{U}_{r,t}\pi_r).$$

Substituting this into the expression for $\pi_t - \tilde{\pi}_t$ and rearranging, we obtain

$$\begin{aligned} \pi_t - \tilde{\pi}_t &= \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_\Lambda\pi_r dr + \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \\ &\quad - \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \left[h^*\pi_r(H - h^*\pi_r)\pi_r - \tilde{h}^*\pi_r(\tilde{H} - \tilde{h}^*\pi_r)\pi_r \right] dr \\ &\quad + \frac{1}{2} \sum_{k,\ell} \int_0^t \frac{\partial^2\Sigma}{\partial x^k\partial x^\ell}(\tilde{U}_{r,t}\pi_r)(\tilde{U}_{r,t}(H - h^*\pi_r)\pi_r)^k(\tilde{U}_{r,t}(H - h^*\pi_r)\pi_r)^\ell dr \\ &\quad - \frac{1}{2} \sum_{k,\ell} \int_0^t \frac{\partial^2\Sigma}{\partial x^k\partial x^\ell}(\tilde{U}_{r,t}\pi_r)(\tilde{U}_{r,t}(\tilde{H} - \tilde{h}^*\pi_r)\pi_r)^k(\tilde{U}_{r,t}(\tilde{H} - \tilde{h}^*\pi_r)\pi_r)^\ell dr. \end{aligned}$$

It remains to note that we can write

$$(D^2\tilde{\pi}_{s,t}(\mu) \cdot v)^i = \sum_{k,\ell} D^2\Sigma^{ik\ell}(\tilde{U}_{s,t}\mu)(\tilde{U}_{s,t}v)^k(\tilde{U}_{s,t}v)^\ell.$$

The result follows immediately. \square

Let $e_t = \pi_t - \tilde{\pi}_t$. We wish to estimate the norm of e_t . Unfortunately, we can no longer use the triangle inequality as in Section 3.2 due to the presence of the stochastic integral; instead, we choose to calculate $\|e_t\|^2$, which is readily estimated.

LEMMA 3.4.2. *The filtering error can be estimated by*

$$\begin{aligned} \mathbf{E}_{\mathbf{P}} \|e_t\|^2 &\leq \int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_{\Lambda} \pi_r| dr + K \int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r)| dr \\ &\quad + \int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot (h^* \pi_r (H - h^* \pi_r) \pi_r - \tilde{h}^* \pi_r (\tilde{H} - \tilde{h}^* \pi_r) \pi_r)| dr \\ &\quad + \frac{1}{2} \int_0^t \mathbf{E}_{\mathbf{P}} |D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (H - h^* \pi_r) \pi_r - D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (\tilde{H} - \tilde{h}^* \pi_r) \pi_r| dr, \end{aligned}$$

where $K = 2 \max_k |h^k| + \max_k |\tilde{h}^k|$.

PROOF. We wish to calculate $\mathbf{E}_{\mathbf{P}} \|e_t\|^2 = \mathbf{E}_{\mathbf{P}} e_t^* e_t$. Using Prop. 3.4.1, we obtain

$$\begin{aligned} \mathbf{E}_{\mathbf{P}} \|e_t\|^2 &= \int_0^t \mathbf{E}_{\mathbf{P}} e_t^* D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_{\Lambda} \pi_r dr \\ &\quad + \mathbf{E}_{\mathbf{P}} \left[e_t^* \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \right] \\ &\quad - \int_0^t \mathbf{E}_{\mathbf{P}} e_t^* D\tilde{\pi}_{r,t}(\pi_r) \cdot \left[h^* \pi_r (H - h^* \pi_r) \pi_r - \tilde{h}^* \pi_r (\tilde{H} - \tilde{h}^* \pi_r) \pi_r \right] dr \\ &\quad + \frac{1}{2} \int_0^t \mathbf{E}_{\mathbf{P}} e_t^* \left[D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (H - h^* \pi_r) \pi_r - D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (\tilde{H} - \tilde{h}^* \pi_r) \pi_r \right] dr. \end{aligned}$$

The chief difficulty is the stochastic integral term. Using (3.11), we can write

$$\mathbf{E}_{\mathbf{P}} \left[e_t^* \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \right] = \mathbf{E}_{\mathbf{Q}} \left[|U_{0,t}\nu| e_t^* \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \right].$$

We would like to apply (A.1) to evaluate this expression. First, we must establish that the integrand is in $\text{Dom } \delta$; this does not follow directly from Proposition 3.4.1, as the anticipative Itô rule which was used to obtain that result can yield integrands which are only in $\mathbb{L}_{\text{loc}}^{1,2}$. We can verify directly, however, that the integrand in this case is indeed in $\text{Dom } \delta$, see Lemma 3.5.3. Next, we must establish that $|U_{0,t}\nu| e_t^i$ is in $\mathbb{D}^{1,2}$ for every i . Note that $|U_{0,t}\nu| = \sum_i (U_{0,t}\nu)^i$, so $|U_{0,t}\nu|$ is in \mathbb{D}^{∞} . Moreover, we establish in Lemma 3.5.4 that $e_t \in \mathbb{D}^{1,2}$ and that $\mathbf{D}_r e_t$ is a bounded random variable for every t . Hence it follows from Proposition A.1.1 that $|U_{0,t}\nu| e_t^i \in \mathbb{D}^{1,2}$. Consequently we can apply (A.1), and we obtain

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}} \left[|U_{0,t}\nu| e_t^* \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \right] &= \int_0^t \mathbf{E}_{\mathbf{Q}} \left[(|U_{0,t}\nu| \mathbf{D}_r e_t^* + \mathbf{D}_r |U_{0,t}\nu| e_t^*) D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) \right] dr \\ &= \int_0^t \mathbf{E}_{\mathbf{Q}} \left[|U_{0,t}\nu| (\mathbf{D}_r \pi_t - \mathbf{D}_r \tilde{\pi}_t)^* D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) \right] dr \\ &\quad + \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\sum_i (U_{r,t} H U_{0,r}\nu)^i e_t^* D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) \right] dr. \end{aligned}$$

Now note that $|e_t^i| \leq 1$, and that by Lemma 3.5.4

$$|(\mathbf{D}_r \pi_t - \mathbf{D}_r \tilde{\pi}_t)^i| \leq |(\mathbf{D}_r \pi_t)^i| + |(\mathbf{D}_r \tilde{\pi}_t)^i| \leq \max_k |h^k| + \max_k |\tilde{h}^k|.$$

Furthermore we can estimate

$$\left| \frac{\sum_i (U_{r,t} H U_{0,r} \nu)^i}{|U_{0,t} \nu|} \right| \leq \frac{1}{|U_{0,t} \nu|} \sum_{i,j,k} U_{r,t}^{ij} |h^j| U_{0,r}^{jk} \nu^k \leq \max_k |h^k|,$$

where we have used a.s. nonnegativity of the matrix elements of $U_{0,r}$ and $U_{r,t}$ (this must be the case, as, e.g., $U_{r,t} \mu$ has nonnegative entries for any vector μ with nonnegative entries). Hence we obtain, using the triangle inequality,

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}} \left[|U_{0,t} \nu| e_t^* \int_0^t D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r) dY_r \right] \\ \leq (2 \max_k |h^k| + \max_k |\tilde{h}^k|) \int_0^t \mathbf{E}_{\mathbf{Q}} |U_{0,t} \nu| |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r)| dr. \end{aligned}$$

The result follows after straightforward manipulations. \square

Unlike in the case $\tilde{h} = h$, we now have to deal also with second derivatives of the filter with respect to its initial condition. These can be estimated much in the same way as we dealt with the first derivatives.

LEMMA 3.4.3. *Let $\tilde{\lambda}_{ij} > 0 \forall i \neq j$ and $\mu \in \mathcal{S}^{d-1}$, $v, w \in T\mathcal{S}^{d-1}$. Then a.s.*

$$\begin{aligned} |D^2 \tilde{\pi}_{s,t}(\mu) \cdot v - D^2 \tilde{\pi}_{s,t}(\mu) \cdot w| \\ \leq 2 \sum_k \frac{|v^k + w^k|}{\mu^k} \sum_j \frac{|v^j - w^j|}{\mu^j} \exp \left(-2(t-s) \min_{p \neq q} \sqrt{\tilde{\lambda}_{pq} \tilde{\lambda}_{qp}} \right). \end{aligned}$$

Moreover, the result still holds if μ, v, w are \mathcal{F}_s^Y -measurable random variables with values a.s. in \mathcal{S}^{d-1} and $T\mathcal{S}^{d-1}$, respectively.

PROOF. Proceeding as in the proof of Proposition 3.3.2, we can calculate directly the second derivative of (3.9):

$$(D^2 \pi_t(\mu) \cdot v)^i = -2 (D\pi_t(\mu) \cdot v)^i \frac{\sum_j (v^j / \nu^j) \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y)}{\sum_j (\mu^j / \nu^j) \mathbf{P}(X_0 = a_j | \mathcal{F}_t^Y)}.$$

Setting $\mu = \nu$ and using the triangle inequality, we obtain

$$|D^2 \pi_t(\nu) \cdot v - D^2 \pi_t(\nu) \cdot w| \leq 2 \sum_{i,j} \frac{|v^j (D\pi_t(\nu) \cdot v)^i - w^j (D\pi_t(\nu) \cdot w)^i|}{\nu^j}.$$

Another application of the triangle inequality and using Proposition 3.3.2 gives

$$\begin{aligned} |D^2 \pi_t(\nu) \cdot v - D^2 \pi_t(\nu) \cdot w| \\ \leq \sum_k \frac{|v^k + w^k|}{\nu^k} |D\pi_t(\nu) \cdot (v - w)| + \sum_k \frac{|v^k - w^k|}{\nu^k} |D\pi_t(\nu) \cdot (v + w)| \\ \leq 2 \sum_k \frac{|v^k + w^k|}{\nu^k} \sum_j \frac{|v^j - w^j|}{\nu^j} \exp \left(-2t \min_{p \neq q} \sqrt{\lambda_{pq} \lambda_{qp}} \right). \end{aligned}$$

We can now repeat the arguments of Corollary 3.3.3 to establish that the result still holds if we replace $\pi_{0,t}$ by $\tilde{\pi}_{s,t}$, λ_{pq} by $\tilde{\lambda}_{pq}$, and ν, v, w by \mathcal{F}_s^Y -measurable random variables μ, v, w . This completes the proof. \square

We are now ready to complete the proof of Theorem 3.1.5.

PROOF OF THEOREM 3.1.5. Set $\beta = 2 \min_{p,q \neq p} (\tilde{\lambda}_{pq} \tilde{\lambda}_{qp})^{1/2}$. Let us collect all the necessary estimates. First, we have

$$\int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_{\Lambda} \pi_r| dr \leq \beta^{-1} \sup_{s \geq 0} \mathbf{E}_{\mathbf{P}} (1 / \min_k \pi_s^k) |\Lambda^* - \tilde{\Lambda}^*|,$$

as we showed in Section 3.3. Next, we obtain

$$\int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r)| dr \leq \beta^{-1} \sup_{\pi \in \mathcal{S}^{d-1}} \sum_k |h^k - \tilde{h}^k + \tilde{h}^* \pi - h^* \pi|$$

using Corollary 3.3.3. Using the triangle inequality, we can estimate this by

$$\int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot \Delta_H(\pi_r)| dr \leq (d+1)\beta^{-1} |h - \tilde{h}|.$$

Next, we estimate using Corollary 3.3.3

$$\begin{aligned} \int_0^t \mathbf{E}_{\mathbf{P}} |D\tilde{\pi}_{r,t}(\pi_r) \cdot (h^* \pi_r (H - h^* \pi_r) \pi_r - \tilde{h}^* \pi_r (\tilde{H} - \tilde{h}^* \pi_r) \pi_r)| dr \\ \leq \beta^{-1} \sup_{\pi \in \mathcal{S}^{d-1}} \sum_k |h^* \pi (h^k - h^* \pi) - \tilde{h}^* \pi (\tilde{h}^k - \tilde{h}^* \pi)| \\ \leq \beta^{-1} \left((d+1) \max_k |h^k| + d \max_{k,\ell} |\tilde{h}^k - \tilde{h}^\ell| \right) |h - \tilde{h}|, \end{aligned}$$

where we have used the estimate

$$\begin{aligned} \sum_k |h^* \pi (h^k - h^* \pi) - \tilde{h}^* \pi (\tilde{h}^k - \tilde{h}^* \pi)| \\ \leq |h^* \pi| \sum_k |h^k - \tilde{h}^k + \tilde{h}^* \pi - h^* \pi| + |h^* \pi - \tilde{h}^* \pi| \sum_k |\tilde{h}^k - \tilde{h}^* \pi| \\ \leq (d+1) \max_k |h^k| |h - \tilde{h}| + |h - \tilde{h}| \sum_k |\tilde{h}^k - \tilde{h}^* \pi| \\ \leq \left((d+1) \max_k |h^k| + d \max_{k,\ell} |\tilde{h}^k - \tilde{h}^\ell| \right) |h - \tilde{h}|. \end{aligned}$$

Next we estimate using Lemma 3.4.3

$$\begin{aligned} \frac{1}{2} \int_0^t \mathbf{E}_{\mathbf{P}} |D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (H - h^* \pi_r) \pi_r - D^2 \tilde{\pi}_{r,t}(\pi_r) \cdot (\tilde{H} - \tilde{h}^* \pi_r) \pi_r| dr \\ \leq \beta^{-1} \sup_{\pi \in \mathcal{S}^{d-1}} \sum_k |h^k - h^* \pi + \tilde{h}^k - \tilde{h}^* \pi| \sum_j |h^j - \tilde{h}^j + \tilde{h}^* \pi - h^* \pi| \\ \leq d(d+1)\beta^{-1} \left(\max_{k,\ell} |h^k - h^\ell| + \max_{k,\ell} |\tilde{h}^k - \tilde{h}^\ell| \right) |h - \tilde{h}|. \end{aligned}$$

We have now estimated all the terms in Lemma 3.4.2, and hence we have bounded $\mathbf{E}_{\mathbf{P}} \|e_t\|^2 = \mathbf{E}_{\mathbf{P}} \|\pi_t(\nu) - \tilde{\pi}_t(\nu)\|^2$. It remains to allow for misspecified initial conditions. To this end, we estimate

$$\|\pi_t(\nu) - \tilde{\pi}_t(\mu)\|^2 \leq \|e_t\|^2 + \|\tilde{\pi}_t(\nu) - \tilde{\pi}_t(\mu)\| (\|\tilde{\pi}_t(\nu) - \tilde{\pi}_t(\mu)\| + 2\|\pi_t(\nu) - \tilde{\pi}_t(\nu)\|).$$

Hence we obtain using the equivalence of finite-dimensional norms $\|x\| \leq K_{21} |x|$

$$\|\pi_t(\nu) - \tilde{\pi}_t(\mu)\|^2 \leq \|e_t\|^2 + 6K_{21} \|\tilde{\pi}_t(\nu) - \tilde{\pi}_t(\mu)\|$$

where we have used that the simplex is contained in the $(d-1)$ -dimensional unit sphere, so $\|\mu_1 - \mu_2\| \leq 2 \forall \mu_1, \mu_2 \in \Delta^{d-1}$. The statement of the Theorem now follows directly from Lemma 3.4.2, Proposition 3.3.4, and the estimates above. \square

3.5. Some technical lemmas

LEMMA 3.5.1. *The following equality holds:*

$$\begin{aligned} \tilde{U}_{0,t} \int_0^s \tilde{U}_{0,r}^{-1} (H - \tilde{H}) U_{0,r} \nu dY_r = \\ \int_0^s \tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu dY_r + \int_0^s \tilde{U}_{r,t} \tilde{H} (H - \tilde{H}) U_{0,r} \nu dr. \end{aligned}$$

The integral on the left is an Itô integral, on the right a Skorokhod integral.

PROOF. We have already established in the proof of Proposition 3.4.1 that the matrix elements of $\tilde{U}_{0,t}$ are in $\mathbb{D}^\infty \subset \mathbb{D}^{1,2}$. Moreover,

$$\begin{aligned} \mathbf{E}_{\mathbf{Q}} \|\tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu\|^2 &\leq \|H - \tilde{H}\|^2 \mathbf{E}_{\mathbf{Q}} (\|\tilde{U}_{r,t}\|^2 \|U_{0,r}\|^2) \\ &\leq \|H - \tilde{H}\|^2 \sqrt{\mathbf{E}_{\mathbf{Q}} \|\tilde{U}_{r,t}\|^4 \mathbf{E}_{\mathbf{Q}} \|U_{0,r}\|^4} \\ &\leq C_4^4 \|H - \tilde{H}\|^2 \sqrt{\mathbf{E}_{\mathbf{Q}} \|\tilde{U}_{r,t}\|_4^4 \mathbf{E}_{\mathbf{Q}} \|U_{0,r}\|_4^4}, \end{aligned}$$

where we have used the Cauchy-Schwarz inequality and $\|\nu\| \leq 1$ for $\nu \in \mathcal{S}^{d-1}$. Here $\|U\|_p = (\sum_{ij} U_{ij}^p)^{1/p}$ is the elementwise p -norm of U , $\|U\|$ is the usual matrix 2-norm, and C_p matches the norms $\|U\| \leq C_p \|U\|_p$ (recall that all norms on a finite-dimensional space are equivalent). As $U_{0,r}, \tilde{U}_{r,t}$ are solutions of linear stochastic differential equations, standard estimates give for any integer $p \geq 2$

$$\mathbf{E}_{\mathbf{Q}} \left(\sup_{0 \leq r \leq t} \|\tilde{U}_{r,t}\|_p^p \right) \leq D_1(p) < \infty, \quad \mathbf{E}_{\mathbf{Q}} \left(\sup_{0 \leq r \leq t} \|U_{0,r}\|_p^p \right) \leq D_2(p) < \infty,$$

and we obtain

$$\int_0^s \mathbf{E}_{\mathbf{Q}} \|\tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu\|^2 dr \leq s \sup_{0 \leq r \leq s} \mathbf{E}_{\mathbf{Q}} \|\tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu\|^2 < \infty.$$

Hence we can apply Lemma A.4.2 to obtain the result. By a similar calculation we can establish that the right-hand side of the expression in Lemma A.4.2 for our case is square integrable, so that the Skorokhod integral is well defined. \square

LEMMA 3.5.2. *The anticipating Itô rule with $\Sigma(x) = x/|x|$ can be applied to*

$$\tilde{U}_{s,t} U_{0,s} \nu = \tilde{U}_{0,t} \nu + \int_0^s \tilde{U}_{r,t} (\Lambda^* - \tilde{\Lambda}^*) U_{0,r} \nu dr + \int_0^s \tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu dY_r.$$

PROOF. Clearly the Skorokhod integral term has a.s. continuous sample paths, as both $\tilde{U}_{s,t} U_{0,s} \nu$ and the time integrals do; moreover, $\tilde{U}_{0,t} \nu \in (\mathbb{D}^\infty)^d$. In order to be able to apply Proposition A.5.1, it remains to check the technical conditions $v_r = \tilde{U}_{r,t} (\Lambda^* - \tilde{\Lambda}^*) U_{0,r} \nu \in (\mathbb{L}^{1,4})^d$, $u_r = \tilde{U}_{r,t} (H - \tilde{H}) U_{0,r} \nu \in (\mathbb{L}^{2,4})^d$.

As \mathbb{D}^∞ is an algebra, u_t and v_t take values in \mathbb{D}^∞ . Moreover, we can establish exactly as in the proof of Lemma 3.5.1 that u and v are in $L^4(\Omega \times [0, t])$. To

complete the proof we must establish that

$$\sum_i \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\int_0^t (\mathbf{D}_s u_r^i)^2 ds \right]^2 dr < \infty, \quad \sum_i \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\int_0^t (\mathbf{D}_s v_r^i)^2 ds \right]^2 dr < \infty,$$

thus ensuring that $u, v \in (\mathbb{L}^{1,4})^d$, and

$$\sum_i \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\int_0^t \int_0^t (\mathbf{D}_\sigma \mathbf{D}_s u_r^i)^2 ds d\sigma \right]^2 dr < \infty$$

which ensures that $u \in (\mathbb{L}^{2,4})^d$. Using the Cauchy-Schwarz inequality we have

$$\begin{aligned} \sum_i \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\int_0^t (\mathbf{D}_s u_r^i)^2 ds \right]^2 dr \\ \leq t \int_0^t \int_0^t \mathbf{E}_{\mathbf{Q}} \|\mathbf{D}_s u_r\|_4^4 ds dr \leq t^3 \sup_{0 \leq r, s \leq t} \mathbf{E}_{\mathbf{Q}} \|\mathbf{D}_s u_r\|_4^4, \end{aligned}$$

and similarly for v . Moreover, we obtain

$$\sum_i \int_0^t \mathbf{E}_{\mathbf{Q}} \left[\int_0^t \int_0^t (\mathbf{D}_\sigma \mathbf{D}_s u_r^i)^2 ds d\sigma \right]^2 dr \leq t^5 \sup_{0 \leq r, s, \sigma \leq t} \mathbf{E}_{\mathbf{Q}} \|\mathbf{D}_\sigma \mathbf{D}_s u_r\|_4^4.$$

But using the chain rule Proposition A.1.2 we can easily establish that

$$\mathbf{D}_s u_r = \begin{cases} \tilde{U}_{r,t}(H - \tilde{H})U_{s,r}HU_{0,s}\nu & \text{a.e. } 0 < s < r < t, \\ \tilde{U}_{s,t}\tilde{H}\tilde{U}_{r,s}(H - \tilde{H})U_{0,r}\nu & \text{a.e. } 0 < r < s < t, \end{cases}$$

and similarly

$$\mathbf{D}_\sigma \mathbf{D}_s u_r = \begin{cases} \tilde{U}_{r,t}(H - \tilde{H})U_{s,r}HU_{\sigma,s}HU_{0,\sigma}\nu & \text{a.e. } 0 < \sigma < s < r < t, \\ \tilde{U}_{r,t}(H - \tilde{H})U_{\sigma,r}HU_{s,\sigma}HU_{0,s}\nu & \text{a.e. } 0 < s < \sigma < r < t, \\ \tilde{U}_{\sigma,t}\tilde{H}\tilde{U}_{r,\sigma}(H - \tilde{H})U_{s,r}HU_{0,s}\nu & \text{a.e. } 0 < s < r < \sigma < t, \\ \tilde{U}_{s,t}\tilde{H}\tilde{U}_{r,s}(H - \tilde{H})U_{\sigma,r}HU_{0,\sigma}\nu & \text{a.e. } 0 < \sigma < r < s < t, \\ \tilde{U}_{s,t}\tilde{H}\tilde{U}_{\sigma,s}\tilde{H}\tilde{U}_{r,\sigma}(H - \tilde{H})U_{0,r}\nu & \text{a.e. } 0 < r < \sigma < s < t, \\ \tilde{U}_{\sigma,t}\tilde{H}\tilde{U}_{s,\sigma}\tilde{H}\tilde{U}_{r,s}(H - \tilde{H})U_{0,r}\nu & \text{a.e. } 0 < r < s < \sigma < t. \end{cases}$$

The desired estimates now follow as in the proof of Lemma 3.5.1. \square

LEMMA 3.5.3. *The Skorokhod integrand obtained by applying the anticipative Itô formula as in Lemma 3.5.2 is in $\text{Dom } \delta$.*

PROOF. We use the notation $\sigma_r = U_{0,r}\nu$. The integral in question is

$$\int_0^s D\Sigma(\tilde{U}_{r,t}\sigma_r)\tilde{U}_{r,t}(H - \tilde{H})\sigma_r dY_r = \int_0^s f_r dY_r.$$

To establish $f \in \text{Dom } \delta$, it suffices to show that $f \in \mathbb{L}^{1,2}$. We begin by showing

$$\begin{aligned} |D\Sigma(\tilde{U}_{r,t}\sigma_r)\tilde{U}_{r,t}(H - \tilde{H})\sigma_r| &= \sum_i \left| \sum_{j,k} \frac{\delta^{ij} - \Sigma^i(\tilde{U}_{r,t}\sigma_r)}{|\tilde{U}_{r,t}\sigma_r|} \tilde{U}_{r,t}^{jk} (h^k - \tilde{h}^k) \sigma_r^k \right| \leq \\ & \frac{1}{|\tilde{U}_{r,t}\sigma_r|} \sum_{i,j,k} \tilde{U}_{r,t}^{jk} |h^k - \tilde{h}^k| \sigma_r^k \leq \frac{\max_k |h^k - \tilde{h}^k|}{|\tilde{U}_{r,t}\sigma_r|} \sum_{i,j,k} \tilde{U}_{r,t}^{jk} \sigma_r^k = d \max_k |h^k - \tilde{h}^k|, \end{aligned}$$

where we have used the triangle inequality, $|\delta^{ij} - \Sigma^i(x)| \leq 1$ for any $x \in \mathbb{R}_{++}^d$, and the fact that $U_{r,t}$ and σ_r have nonnegative entries a.s. Hence f_r is a bounded process. Similarly, we will show that $\mathbf{D}_s f_r$ is a bounded process. Note that f_r is a smooth function on \mathbb{R}_{++}^d of positive random variables in \mathbb{D}^∞ ; hence we can apply the chain rule Proposition A.1.1. This gives

$$(\mathbf{D}_s f_r)^i = \begin{cases} \sum_{jk} D^2 \Sigma^{ijk}(\tilde{U}_{r,t} \sigma_r) (\tilde{U}_{r,t} (H - \tilde{H}) \sigma_r)^j (\tilde{U}_{r,t} U_{s,r} H \sigma_s)^k \\ \quad + \sum_j D \Sigma^{ij}(\tilde{U}_{r,t} \sigma_r) (\tilde{U}_{r,t} (H - \tilde{H}) U_{s,r} H \sigma_s)^j & \text{a.e. } s < r, \\ \sum_{jk} D^2 \Sigma^{ijk}(\tilde{U}_{r,t} \sigma_r) (\tilde{U}_{r,t} (H - \tilde{H}) \sigma_r)^j (\tilde{U}_{s,t} \tilde{H} \tilde{U}_{r,s} \sigma_r)^k \\ \quad + \sum_j D \Sigma^{ij}(\tilde{U}_{r,t} \sigma_r) (\tilde{U}_{s,t} \tilde{H} \tilde{U}_{r,s} (H - \tilde{H}) \sigma_r)^j & \text{a.e. } s > r. \end{cases}$$

Proceeding exactly as before, we find that $\mathbf{D}f \in L^\infty(\Omega \times [0, t]^2)$. But then by Proposition A.1.1 we can conclude that $\mathbf{D}_s f_r \in \mathbb{D}^{1,2}$ for a.e. $(s, t) \in [0, t]^2$, and in particular $f \in \mathbb{L}^{1,2}$. Hence the proof is complete. \square

LEMMA 3.5.4. $\mathbf{D}_r \pi_s = D\pi_{r,s}(\pi_r) \cdot (H - h^* \pi_r) \pi_r$ a.e. $r < s$, $\mathbf{D}_r \pi_s = 0$ a.e. $r > s$. Moreover $|(\mathbf{D}_r \pi_s)^i| \leq \max_k |h^k|$ for every i . The equivalent results hold for $\mathbf{D}_r \tilde{\pi}_s$. In particular, this implies that π_s and $\tilde{\pi}_s$ are in $\mathbb{D}^{1,2}$.

PROOF. The case $r > s$ is immediate from adaptedness of π_s . For $r < s$, apply the chain rule to $\pi_s = \Sigma(U_{0,s} \nu) \in \mathbb{D}_{\text{loc}}^{1,2}$. Boundedness of the resulting expression follows, e.g., as in the proof of Lemma 3.5.3, and hence it follows that $\pi_s \in \mathbb{D}^{1,2}$. \square

Filter Stability and Conditional Signals: Continuous State Space

The goal of this chapter is to obtain some filter stability results for a class of filtering problems with a diffusive signal, similar to the filtering model considered by W. Stannat [Sta04, Sta05, Sta06]. We both extend some results of Stannat using new probabilistic proofs, and are also able to obtain a much stronger stability result (under much more stringent conditions). Our method is completely probabilistic and, in my opinion, effectively demystifies the clever but not so intuitive PDE methods used in [Sta04, Sta05, Sta06] and gives a deeper insight into the structure and limitations of this type of result. We will be a little less careful with technicalities in this chapter in order to be able to focus on the main line of the argument. It is fairly clear, however, how one could proceed to make certain arguments rigorous, which I have tried to indicate in the text.

4.1. The filtering model

Consider the following signal-observation pair:

$$(4.1) \quad dX_t = BX_t dt + \nabla V(X_t) dt + dW_t,$$

$$(4.2) \quad dY_t = HX_t dt + dB_t,$$

where $X_t, W_t \in \mathbb{R}^d$, $Y_t, B_t \in \mathbb{R}^p$. Here W_t and B_t are independent Wiener processes, B is a $d \times d$ matrix, H is a $p \times d$ matrix and $V(x)$ is a $C^{4,\alpha}$ potential function. We denote by $b(x) = Bx + \nabla V(x)$ the drift of X_t , and assume that it is of linear growth and has bounded derivatives up to third order. Finally, let the initial law $X_0 \sim \nu \ll dx$ have a C^3 density $p_0(x)$ with respect to the Lebesgue measure, and we assume that $p_0(x)$ is bounded and positive.

The technical conditions above could probably be weakened, though they are not particularly restrictive. In particular, these conditions guarantee that (i) the stochastic differential equation for the signal has a unique (strong) solution; and (ii) that the pathwise filtering equation, see section 1.4.3, has a unique solution with C^2 density with respect to the Lebesgue measure (see [Kun90, Thm. 6.2.2]). On the other hand, we have chosen a rather restrictive set of signal/observation models: the signal drift must be of gradient plus linear type, the signal diffusion coefficient must be constant and nondegenerate, and the observation function is linear. The reason for these choices will become clear in due course.

REMARK 4.1.1. The model above extends trivially to any signal of the form

$$dX_t' = AA^*(B'X_t' + \nabla V'(X_t')) dt + A dW_t,$$

where A is an invertible matrix. Indeed, setting $X'_t = AX_t$, we find that

$$dX'_t = ABA^{-1}X'_t dt + A\nabla V(A^{-1}X'_t) dt + A dW_t,$$

so that we can apply the results of this chapter with $V(x) = V'(Ax)$, $B = A^*B'A$ (and a suitable rescaling of the final result). Extension to time-dependent $B(t)$, $V(t, x)$ is not difficult, but we restrict to the model above for notational convenience.

4.2. Preliminaries: Some conditional signal theory

Before we turn our attention to the filter stability problem, let us take a moment to develop the time-reversed conditional signal theory that will be needed in the following. The stochastic control problems which we are about to investigate already appear in the papers [FM83, Mit81, Par81], but the interpretation of the controlled process as the time-reversed conditional signal still appears to be lacking even in [MN03]. It is not our goal to develop the theory here in its full generality, but we will sketch how to proceed in the particular case under investigation.

4.2.1. Conditional signal—time-reversed case. First, let us recall a standard result on time reversal for diffusions [HP86]. Fix a terminal time $T > 0$; we are interested in the time-reversed diffusion process $\tilde{X}_t = X_{T-t}$. Under the technical conditions which we have already imposed, there exists a Wiener process \tilde{W}_t , $t \in [0, T]$ such that \tilde{X}_t is again a diffusion with

$$d\tilde{X}_t = -b(\tilde{X}_t) dt + \nabla \log p_{T-t}(\tilde{X}_t) dt + d\tilde{W}_t.$$

Here $p_t(x)$ is the (unconditional) density of X_t with respect to the Lebesgue measure, which is C^2 and positive¹ by our assumptions (for an investigation of the growth properties of $\nabla \log p_t(x)$, see [Hau85]). The Wiener process \tilde{W}_t can be identified explicitly, see [Par86], but we will forgo this issue here.

In the above, \tilde{W}_t is a Wiener process under the usual measure \mathbf{P} . To develop the conditional diffusion theory, we would like to describe the law of \tilde{X}_t under the conditional measure $\Pi_T(\cdot, y)$. We now basically proceed exactly as in section 2.2.3. First, note that $\Pi_T(\cdot, y)$ is equivalent to μ_x , the unconditional law of the signal process. The Girsanov theory [Pro04, sec. III.8] tells us that any equivalent change of measure on the space of signal sample paths is equivalently described by the addition of a drift term, i.e., by the process

$$d\tilde{X}_t^u = -b(\tilde{X}_t^u) dt + \nabla \log p_{T-t}(\tilde{X}_t^u) dt + u_t dt + d\tilde{W}_t$$

where $u_t \in \mathbb{R}^d$ is (backwards) progressively measurable. It is our goal to find the control u^* under which the law of $\tilde{X}_t^{u^*}$ under \mathbf{P} is precisely the law of \tilde{X}_t under

¹ The smoothness of $p_t(x)$ follows from standard results, while the fact that $p_t(x)$ is positive can be seen, for example, as follows. As $b(x)$ was assumed to be sufficiently regular, the solution $\xi_t(x)$ of the signal process X_t with initial condition $X_0 = x$ is a random diffeomorphism (by the theory of stochastic flows [Kun84]). Hence the conditional law $\mathbf{P}(X_t \in \cdot | \sigma\{W_s, s \in [0, t]\})$ admits a density $p_{t|W}(x)$ which is simply the initial density transformed by the flow:

$$p_{t|W}(x) = p_0(\xi_t^{-1}(x)) \det|\nabla \xi_t^{-1}(x)|.$$

But $p_0(x) > 0$ by assumption, while the Jacobian determinant must a.s. be positive by the a.s. continuity and invertibility of the flow. Hence $p_t(x) = \mathbf{E}(p_{t|W}(x))$ is positive.

$\Pi_T(\cdot, y)$. Let us pose this problem as an optimal control problem, i.e., as usual we will find u^* from the variational Kallianpur-Striebel formula

$$I(H_T(\cdot, y)) = \min_{\mu' \in \mathcal{P}_x} \{D(\mu' || \mu_x) + \langle H_T(\cdot, y), \mu' \rangle\}.$$

(We refer to [MN03] for a thorough discussion of the technical details in the application of this expression in the diffusion case.) The Girsanov theorem gives

$$D(\mu' || \mu_x) = D(\nu'_T || \nu_T) + \mathbf{E} \left[\frac{1}{2} \int_0^T \|u_s\|^2 ds \right],$$

where ν'_T is the measure of \tilde{X}_0 under μ' , and ν_T is the measure with density $p_T(x)$. Similarly, we find that for $y \in C^1$

$$\langle H_T(\cdot, y), \mu' \rangle = \mathbf{E} \left[\frac{1}{2} \int_0^T \|\dot{y}_{T-s} - H\tilde{X}_s^u\|^2 ds - \frac{1}{2} \int_0^T \|\dot{y}_{T-s}\|^2 ds \right],$$

and the expression for any $y \in C_0([0, T]; \mathbb{R}^p)$ is easily found by integration by parts.

To convert the problem to a dynamic programming problem we proceed, still following section 2.2.3, by defining the cost-to-go

$$\tilde{J}^u(t, x) = \tilde{\mathbf{E}}_{t,x} \left[\frac{1}{2} \int_t^T \left(\|u_s\|^2 + \|\dot{y}_{T-s} - H\tilde{X}_s^u\|^2 - \|\dot{y}_{T-s}\|^2 \right) ds \right],$$

where $\tilde{\mathbf{P}}_{t,x}$ is the measure under which $\tilde{X}_t^u = x$. An explicit expression for the corresponding value function² $\tilde{V}(t, x) = \min_u \tilde{J}^u(t, x)$ can be obtained directly from the variational Kallianpur-Striebel formula:

$$\tilde{V}(t, x) = -\log \tilde{\mathbf{E}}_{t,x} [\exp(-H_{T-t}(X_{[0, T-t]}, y))] = -\log \varrho_{T-t}^y(x) + \log p_{T-t}(x),$$

where $\varrho_t^y(x)$ is the (pathwise) unnormalized filtering density, which is defined as $\varrho_t^y(x) = e^{y_t^* H x} H^y(t, x)$ where $H^y(t, x)$ is the solution of the pathwise filtering equation (see section 1.4.3). Moreover, it is easily established using standard optimal control methods [FR75] (we will elaborate below for a different control problem) that the optimal control is given by $u_t^* = -\nabla \tilde{V}(t, \tilde{X}_t^{u^*})$, so that the diffusion

$$(4.3) \quad d\tilde{X}_t^* = -b(\tilde{X}_t^*) dt + \nabla \log \varrho_{T-t}^y(\tilde{X}_t^*) dt + d\tilde{W}_t$$

has the same law under \mathbf{P} as does \tilde{X}_t under $\Pi_T(\cdot, y)$. We note that the Kolmogorov backward equation for this diffusion is a well-known smoothing equation, see, e.g., [LSBS96] and the references therein.

REMARK 4.2.1. The conditional diffusion admits an interesting physical picture, which is illustrated in cartoon form in figure 4.1. Consider the reversible case where $B = 0$, $V(x) = -U(x)$, and $p_0(x)$ is taken to be the stationary density $p_0(x) \propto e^{-2U(x)}$. Then \tilde{X}_t has the same transition intensities as X_t , and describes a (noninertial) particle diffusing, in reverse time, on the potential surface $U(x)$. When we change to the measure $\Pi_T(\cdot, y)$, this simply corresponds to a change in the potential surface by the addition of the negative log-likelihood ratio $-\log(\varrho_{T-t}^y(x)/p_{T-t}(x))$. In essence we must introduce, beside the physical forces acting on the particle, additional forces that localize the particle in the regions of enhanced likelihood. In this way, the “effective potential” $U(x) + \tilde{V}(t, x)$ takes into

² Here and in the following, the minimum of the cost-to-go $\tilde{J}^u(t, x)$ at time t is understood to be taken over all $\tilde{\mathcal{F}}_s^t$ -adapted control strategies, $\tilde{\mathcal{F}}_s^t = \sigma\{\tilde{W}_r - \tilde{W}_t : r \in [t, s]\}$.




FIGURE 4.1. Cartoon of a conditional diffusion. Left figure: under \mathbf{P} , a particle diffuses on the potential $U(x)$ in reverse time. The observation path y is used to calculate the likelihood ratio $\varrho_{T-t}(x)/p_{T-t}(x)$. Right figure: under $\Pi_T(\cdot, y)$, the particle sees an “effective potential” $U(x) + \tilde{V}(t, x)$, where $\tilde{V}(t, x) = -\log(\varrho_{T-t}(x)/p_{T-t}(x))$ is the value function. The lightly shaded curves depict the individual terms $U(x)$ and $\tilde{V}(t, x)$.

account the information encoded in the observation sample path y . The forward time case (see [MN03]) is very similar, only the likelihood ratio is replaced by its time-reversed counterpart (the dual filter in the sense of [Par82]).

4.2.2. Conditional signal—gauge transformation. We can now proceed exactly as in section 2.2.4 to find the control problems for which the logarithm of the unnormalized filtering density or of the pathwise filtering density are the respective value functions. For our purposes, however, it will be convenient to go one step further by absorbing the gradient part of the drift $b(x)$ into the value function as well. This essentially corresponds to the “gauge transformation” (a term borrowed from quantum mechanics) introduced in [Mit79], or a “parabolic ground state transform” in the operator-theoretic language of [Sta04]. Though it is not yet obvious at this point, it will be quite crucial in the following that this method can be implemented in some form or another. We content ourselves for the moment by showing how this is done, and we will elaborate further on the significance of this procedure in section 4.3.2.

We would like to be in the situation where (i) the value function serves as the “effective potential” described in the previous section, rather than just the log-likelihood portion of it; and (ii) where in the absence of control, the signal process is a linear diffusion. To this end, we should consider replacing the control input by³ $u_t \mapsto u_t + \nabla \log p_{T-t}(\tilde{X}_t^u) - \nabla V(\tilde{X}_t^u) - \frac{1}{2}(B + B^*)\tilde{X}_t^u$, while we should replace the value function by $\tilde{V}(t, x) \mapsto \tilde{V}(t, x) - \log p_{T-t}(x) + V(x) + \frac{1}{2}x^* Bx$. For additional convenience, let us also add the term $y_{T-t}^* Hx$, so that we can integrate by parts and define the control problem directly for any $y \in C_0([0, T]; \mathbb{R}^p)$ rather than having to worry about taking limits of $y \in C^1$ at the end of the day. We thus define

$$\tilde{u}_t = u_t + \nabla \log p_{T-t}(\tilde{X}_t^u) - \nabla V(\tilde{X}_t^u) - \frac{1}{2}(B + B^*)\tilde{X}_t^u - H^* y_{T-t},$$

³ We incorporate the term $\frac{1}{2}(B + B^*)x$ as this is the symmetric part of Bx , i.e., we have $Bx = \frac{1}{2}(B + B^*)x + \frac{1}{2}(B - B^*)x = \nabla(\frac{1}{2}x^* Bx) + \frac{1}{2}(B - B^*)x$. This separates the linear drift into a gradient portion, which is absorbed into the value function, and a purely antisymmetric part $\frac{1}{2}(B - B^*)x$ which, as we will see, does not contribute directly to the stability of the filter.

so that we can write the controlled diffusion as

$$d\tilde{X}_t^{\tilde{u}} = H^* y_{T-t} dt - \frac{1}{2}(B - B^*)\tilde{X}_t^{\tilde{u}} dt + \tilde{u}_t dt + d\tilde{W}_t.$$

We must now redefine the cost-to-go, and hence the value function, accordingly:

$$\tilde{J}^{\tilde{u}}(t, x) = \tilde{J}^u(t, x) - \log p_{T-t}(x) + V(x) + \frac{1}{2}x^* Bx + y_{T-t}^* Hx,$$

where we define the new value function by $\tilde{V}(t, x) = \min_{\tilde{u}} \tilde{J}^{\tilde{u}}(t, x)$. As none of the added terms depend on u , the cost still attains its minimum at the control strategy \tilde{u}^* which turns $\tilde{X}_t^{\tilde{u}^*}$ into the time-reversed conditional signal, but \tilde{u} and \tilde{J} have been redefined in such a way that $\tilde{u}_t^* = -\nabla\tilde{V}(t, \tilde{X}_t^{\tilde{u}^*})$ (we will see this in a little more detail below). Moreover the uncontrolled process \tilde{X}_t^0 is now a linear diffusion, and we have pushed all the nonlinearity into the “effective potential” $\tilde{V}(t, x)$. This is precisely what we were aiming for.

Our first order of business is to bring the cost $\tilde{J}^{\tilde{u}}(t, x)$ back to standard form. We have already seen how to do this in section 2.2.4; this is nothing more than an exercise in the application of Dynkin’s formula. First, we calculate

$$\begin{aligned} \tilde{\mathbf{E}}_{t,x}(\log p_0(\tilde{X}_T^u)) &= \log p_{T-t}(x) + \tilde{\mathbf{E}}_{t,x} \int_t^T \frac{1}{2} \|\nabla \log p_{T-s}(\tilde{X}_s^u)\|^2 ds \\ &\quad + \tilde{\mathbf{E}}_{t,x} \int_t^T (\nabla \cdot b)(\tilde{X}_s^u) ds + \tilde{\mathbf{E}}_{t,x} \int_t^T (\nabla \log p_{T-s}(\tilde{X}_s^u))^* u_s ds. \end{aligned}$$

Similarly, it is not difficult to calculate that

$$\begin{aligned} \tilde{\mathbf{E}}_{t,x}(V(\tilde{X}_T^u)) &= V(x) + \tilde{\mathbf{E}}_{t,x} \int_t^T \left(\frac{1}{2}(\nabla \cdot b)(\tilde{X}_s^u) - \frac{1}{2}\text{Tr}(B) \right) ds \\ &\quad + \tilde{\mathbf{E}}_{t,x} \int_t^T (\nabla V(\tilde{X}_s^u))^* (u_s + \nabla \log p_{T-s}(\tilde{X}_s^u) - \nabla V(\tilde{X}_s^u) - B\tilde{X}_s^u) ds. \end{aligned}$$

We obtain in the same way

$$\begin{aligned} 0 = \tilde{\mathbf{E}}_{t,x}(y_0^* H \tilde{X}_T^u) &= y_{T-t}^* Hx - \tilde{\mathbf{E}}_{t,x} \int_t^T \left(\dot{y}_{T-s}^* H \tilde{X}_s^u \right) ds \\ &\quad + \tilde{\mathbf{E}}_{t,x} \int_t^T y_{T-s}^* H (u_s + \nabla \log p_{T-s}(\tilde{X}_s^u) - \nabla V(\tilde{X}_s^u) - B\tilde{X}_s^u) ds. \end{aligned}$$

Finally, note that we can write

$$\begin{aligned} \tilde{\mathbf{E}}_{t,x}((\tilde{X}_T^u)^* B \tilde{X}_T^u) &= x^* Bx + \int_t^T \text{Tr}(B) ds \\ &\quad + \tilde{\mathbf{E}}_{t,x} \int_t^T (\tilde{X}_s^u)^* (B + B^*) (u_s + \nabla \log p_{T-s}(\tilde{X}_s^u) - \nabla V(\tilde{X}_s^u) - B\tilde{X}_s^u) ds. \end{aligned}$$

Adding up these expressions, we can now obtain an expression for the cost-to-go in the usual form (running cost integral plus terminal cost):

$$\begin{aligned} \tilde{J}^{\tilde{u}}(t, x) &= \tilde{J}^u(t, x) - \log p_{T-t}(x) + V(x) + \frac{1}{2}x^* Bx + y_{T-t}^* Hx \\ &= \tilde{\mathbf{E}}_{t,x} \left[\frac{1}{2} \int_t^T \{ \|\tilde{u}_s\|^2 + W(\tilde{X}_s^{\tilde{u}}) + G_s^y(\tilde{X}_s^{\tilde{u}}) \} ds + R(\tilde{X}_T^{\tilde{u}}) \right]. \end{aligned}$$

Here we have defined the running cost terms $W(x)$ and $G_t^y(x)$ by

$$(4.4) \quad W(x) = \|Hx\|^2 + \|b(x)\|^2 + \nabla \cdot b(x) - \frac{1}{4}\|(B^* - B)x\|^2,$$

$$(4.5) \quad G_t^y(x) = y_{T-t}^* H(B - B^*)x - \|H^* y_{T-t}\|^2,$$

and we have introduced the terminal cost

$$(4.6) \quad R(x) = -\log p_0(x) + V(x) + \frac{1}{2}x^* Bx.$$

Note in particular that the derivative terms \dot{y}_{T-t} have disappeared in these expressions, so that they make sense for any observation path $y \in C_0([0, T]; \mathbb{R}^p)$.

REMARK 4.2.2. We have not been completely rigorous here; in particular, it is not entirely clear at the outset whether Dynkin's formula can be applied to $-\log p_{T-t}(x)$, $V(x)$ or $x^* Bx$ (as these functions need not have bounded derivatives), or whether the controlled diffusion, equation (4.3), even has a well-defined (weak) solution. Hence this and the previous section should be considered to be of a formal nature. However, our manipulations can be made rigorous using approximation techniques similar to the method used in [MN03].

Let us briefly indicate how the usual verification argument works (e.g., [FR75]). The Bellman equation corresponding to the control problem above is given by

$$\begin{aligned} -\frac{\partial \check{V}(t, x)}{\partial t} &= \frac{1}{2} \nabla^2 \check{V}(t, x) + (H^* y_{T-t} - \frac{1}{2}(B - B^*)x) \cdot \nabla \check{V}(t, x) \\ &\quad + \frac{1}{2} W(x) + \frac{1}{2} G_t^y(x) + \min_{u \in \mathbb{R}^d} \left\{ u \cdot \nabla \check{V}(t, x) + \frac{1}{2} \|u\|^2 \right\}, \end{aligned}$$

with the terminal condition $\check{V}(T, x) = R(x)$. Note that the minimum in the Bellman equation is attained at $u = -\nabla \check{V}(t, x)$; substituting into the expression above, we obtain the Hamilton-Jacobi-Bellman equation

$$\begin{aligned} -\frac{\partial \check{V}(t, x)}{\partial t} &= \frac{1}{2} \nabla^2 \check{V}(t, x) + (H^* y_{T-t} - \frac{1}{2}(B - B^*)x) \cdot \nabla \check{V}(t, x) \\ &\quad - \frac{1}{2} \|\nabla \check{V}(t, x)\|^2 + \frac{1}{2} W(x) + \frac{1}{2} G_t^y(x). \end{aligned}$$

This is easily verified to coincide with what we know the value function should be,

$$\check{V}(t, x) = -\log H^y(T - t, x) + V(x) + \frac{1}{2}x^* Bx,$$

using the pathwise filtering formula of section 1.4.3. To show that the control strategy $\check{u}_t^* = -\nabla \check{V}(t, \check{X}_t^{\check{u}^*})$ is indeed optimal, we proceed formally as follows. For a control strategy \check{u} , Dynkin's formula gives

$$\begin{aligned} \check{V}(0, x) &= \tilde{\mathbf{E}}_{0,x} \left[\check{V}(T, \check{X}_T^{\check{u}}) - \int_0^T \left\{ \frac{\partial}{\partial t} \check{V}(t, \check{X}_t^{\check{u}}) + \frac{1}{2} \nabla^2 \check{V}(t, \check{X}_t^{\check{u}}) \right\} dt \right. \\ &\quad \left. - \int_0^T (H^* y_{T-t} - \frac{1}{2}(B - B^*) \check{X}_t^{\check{u}} + \check{u}_t) \cdot \nabla \check{V}(t, \check{X}_t^{\check{u}}) dt \right]. \end{aligned}$$

Substituting $\check{u}_t^* = -\nabla \check{V}(t, \check{X}_t^{\check{u}^*})$ and using the Bellman equation, we find that $\check{V}(0, x) = \check{J}^{\check{u}^*}(0, x)$, whereas for any \check{u} it follows from the Bellman equation that $\check{V}(0, x) \leq \check{J}^{\check{u}}(0, x)$. To make the argument rigorous we should introduce a suitable class of admissible controls, prove that \check{u}^* is admissible, and justify the application of Dynkin's formula. This is most easily done as in [MN03] by solving the problem

for a particularly simple case in which these technicalities can be dealt with in a straightforward manner, then taking limits to establish the general result.

4.3. Filter stability for potential diffusions

4.3.1. Filter stability. Having worked our way through the preliminaries, we now turn our attention to the problem of filter stability. We are interested in bounding, in total variation norm, the distance between the filters started from two initial measures. To this end, we begin by finding an explicit expression for this quantity. The following result can be found in [COC99, Lemma 2.1].

LEMMA 4.3.1. *Let $\beta \ll \nu$ be measures on \mathbb{R}^d . Then*

$$\frac{d\mathbf{P}^\beta}{d\mathbf{P}^\nu} = \frac{d\beta}{d\nu}(X_0), \quad \frac{d\Pi_t^\beta(X_t \in \cdot, y)}{d\Pi_t^\nu(X_t \in \cdot, y)} = \frac{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0)|X_t, y\right)}{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0), y\right)},$$

where \mathbf{P}^β is the measure under which the Markov process X has initial measure β , and $\Pi_t^\beta(\cdot, y)$ is the corresponding pathwise conditional measure.

PROOF. The first statement follows as in the proof of Proposition 2.1.3. To prove the second statement, note that by the Kallianpur-Striebel formula

$$\Pi_t^\beta(A, y) = \frac{\mathbf{E}^\beta(I_A \tilde{Z}_t(X_{[0,t]}, y))}{\mathbf{E}^\beta(\tilde{Z}_t(X_{[0,t]}, y))} = \frac{\mathbf{E}^\nu\left(\frac{d\beta}{d\nu}(X_0) I_A \tilde{Z}_t(X_{[0,t]}, y)\right)}{\mathbf{E}^\nu\left(\frac{d\beta}{d\nu}(X_0) \tilde{Z}_t(X_{[0,t]}, y)\right)} = \frac{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0) I_A, y\right)}{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0), y\right)}.$$

Hence we can write

$$\frac{d\Pi_t^\beta(\cdot, y)}{d\Pi_t^\nu(\cdot, y)} = \frac{\frac{d\beta}{d\nu}(X_0)}{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0), y\right)},$$

and the result follows as $\Pi_t(X_t \in \cdot, y)$ is the restriction of $\Pi_t(\cdot, y)$ to $\sigma\{X_t\}$. \square

COROLLARY 4.3.2. *For $\beta \ll \nu$, the total variation distance is given by*

$$\|\Pi_t^\beta(X_t \in \cdot, y) - \Pi_t^\nu(X_t \in \cdot, y)\|_{\text{TV}} = \frac{\Pi_t^\nu\left(\left|\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0)|X_t, y\right) - \Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0), y\right)\right|, y\right)}{\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0), y\right)}.$$

In particular, if $\Pi_t^\nu(X_t \in \cdot, y) \ll dx$ with $d\Pi_t^\nu(X_t \in \cdot, y) = \pi_t^\nu(x) dx$, and if $\Lambda_t^{\beta\nu}(x)$ is any version of $\Pi_t^\nu\left(\frac{d\beta}{d\nu}(X_0)|X_t = x, y\right)$, then

$$\|\Pi_t^\beta(X_t \in \cdot, y) - \Pi_t^\nu(X_t \in \cdot, y)\|_{\text{TV}} = \frac{\int_{\mathbb{R}^d} \left| \Lambda_t^{\beta\nu}(x) - \int_{\mathbb{R}^d} \Lambda_t^{\beta\nu}(y) \pi_t^\nu(y) dy \right| \pi_t^\nu(x) dx}{\int_{\mathbb{R}^d} \Lambda_t^{\beta\nu}(y) \pi_t^\nu(y) dy}.$$

PROOF. This follows immediately from the definition of the TV distance. \square

We can now proceed as follows. Note that

$$\|\Pi_t^\beta(X_t \in \cdot, y) - \Pi_t^\nu(X_t \in \cdot, y)\|_{\text{TV}} \leq \frac{\int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y) \right| \pi_t^\nu(x) \pi_t^\nu(y) dx dy}{\int_{\mathbb{R}^d} \Lambda_t^{\beta\nu}(y) \pi_t^\nu(y) dy}.$$

Consider the quantity $|\Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y)|$. If the filter is to forget its initial condition, one could expect that the smoothed estimate of the signal process at the initial time X_0 should become independent of X_t for large t ; in particular, this would imply that $|\Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y)| \rightarrow 0$ as $t \rightarrow \infty$. If we can prove this, we are well on our way to proving filter stability: assuming that $d\beta/d\nu$ is bounded from below we can easily estimate the denominator in the above expression, while we could claim that the

numerator converges to zero provided that we have some control on the integral with respect to $\pi_t^\nu(x) dx \times \pi_t^\nu(y) dy$. We thus initially concentrate on bounding $|\Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y)|$, and work out the remaining details at the end.

REMARK 4.3.3. We have already encountered the finite-state counterpart of $|\Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y)|$: see, e.g., Lemma 3.3.1. In particular, our logic in the above parallels the method used in [BCL04] to prove filter stability for signals on a finite state space. The hard part, of course, is getting the requisite bound on $|\Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y)|$, which we will approach below using the conditional signal theory.

Recall that the the time-reversed conditional signal,

$$d\tilde{X}_t^* = (H^* y_{T-t} - \frac{1}{2}(B - B^*)\tilde{X}_t^* - \nabla\check{V}(t, \tilde{X}_t^*)) dt + d\tilde{W}_t,$$

has the same law under \mathbf{P} as does X_{T-t} under $\Pi_T(\cdot, y)$. In particular \tilde{X}_t^* is a diffusion, so that we can condition it on its initial point as follows⁴:

$$d\tilde{X}_t^*(x) = (H^* y_{T-t} - \frac{1}{2}(B - B^*)\tilde{X}_t^*(x) - \nabla\check{V}(t, \tilde{X}_t^*(x))) dt + d\tilde{W}_t, \quad \tilde{X}_0^*(x) = x.$$

$\tilde{X}_t^*(x)$ is now defined on the same probability space for any $x \in \mathbb{R}^d$, $t \in [0, T]$, and

$$\Lambda_T^{\beta\nu}(x) = \mathbf{E} \left(\frac{d\beta}{d\nu}(\tilde{X}_T^*(x)) \right)$$

defines a version of the conditional expectation $\Pi_T^\nu(\frac{d\beta}{d\nu}(X_0)|X_T = x, y)$. Suppose that $d\beta/d\nu$ is Lipschitz continuous. Then evidently

$$|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)| \leq \mathbf{E} \left(\left| \frac{d\beta}{d\nu}(\tilde{X}_T^*(x)) - \frac{d\beta}{d\nu}(\tilde{X}_T^*(y)) \right| \right) \leq \left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} \mathbf{E} \|\tilde{X}_T^*(x) - \tilde{X}_T^*(y)\|$$

where $\|f\|_{\text{Lip}}$ is the Lipschitz constant of f . Hence evidently $|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)| \rightarrow 0$ if the flow $\tilde{X}_T^*(x)$ of the time-reversed conditional signal is contracting, at least on average. We would like to give conditions under which this can be proved.

The filtering model which we study in this chapter has a nice property that allows us to obtain such a contraction in a straightforward manner. Note that

$$\begin{aligned} \frac{d}{dt} (\tilde{X}_t^*(x) - \tilde{X}_t^*(y)) &= -\frac{1}{2}(B - B^*)(\tilde{X}_t^*(x) - \tilde{X}_t^*(y)) \\ &\quad - (\nabla\check{V}(t, \tilde{X}_t^*(x)) - \nabla\check{V}(t, \tilde{X}_t^*(y))). \end{aligned}$$

Hence we obtain using the ordinary chain rule

$$\frac{d}{dt} \|\tilde{X}_t^*(x) - \tilde{X}_t^*(y)\|^2 = -2(\tilde{X}_t^*(x) - \tilde{X}_t^*(y))^*(\nabla\check{V}(t, \tilde{X}_t^*(x)) - \nabla\check{V}(t, \tilde{X}_t^*(y))),$$

where we have used the fact that $x^*(B - B^*)x = 0$ for any vector x . But this means that the distance $\|\tilde{X}_t^*(x) - \tilde{X}_t^*(y)\|$ is strictly contracting, provided that we require $(x - y)^*(\nabla\check{V}(t, x) - \nabla\check{V}(t, y)) > 0$ for all t, x, y —which is identical to requiring that the value function $\check{V}(t, x)$ be strictly convex for every t by the well-known

⁴ What we have done here is essentially to construct a stochastic flow [Kun84] associated to \tilde{X}_t^* , though we will not need any of its interesting properties. What we do need for this to work, however, is that the equation for \tilde{X}_t^* has a strong solution. But recall that by our assumptions $\check{V}(t, x)$ is twice continuously differentiable in x , so that $\nabla\check{V}(t, x)$ is locally Lipschitz continuous and hence there is a unique strong solution up to an explosion time ζ . Moreover, the laws of \tilde{X}_t^* and X_{T-t} are equivalent, so we see that X_t^* is a.s. finite on $[0, T]$. Hence $\zeta = \infty$ a.s., and this is sufficient to construct $X_t^*(x)$ on the same probability space as a function of x and t .

first order condition for convexity [HUL01]. This is precisely what we are going to show, using the characterization of $\check{V}(t, x)$ as the infimum of the cost-to-go.

Let us briefly summarize our strategy up to this point. Our next step will be to prove that the value function $\check{V}(t, x)$ is uniformly convex. When this is the case, the time-reversed conditional signal \check{X}_t^* is a diffusion in a uniformly convex potential, and hence its flow is contracting at an exponential rate. This in turn implies that the quantity $|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)|$ decays exponentially. Finally we will substitute this bound into the expression for the total variation distance between the differently initialized filters, and it remains to control the double integral in the expression for the total variation distance (which is not entirely trivial, as we will see).

4.3.2. Uniform convexity of the value function. In this section, we will show that under certain conditions the value function $\check{V}(t, x)$ is uniformly convex. The control-theoretic method lends itself particularly well to this task, as convexity of the value function is essentially inherited from the cost-to-go. A similar method was used by C. Borell [Bor00] to obtain certain geometric inequalities for diffusions. We begin by recalling a basic definition.

DEFINITION 4.3.4. $f(x)$ is called κ -uniformly convex if $f(x) - \frac{1}{2}\kappa\|x\|^2$ is convex.

This concept is important to us for the following elementary reason.

LEMMA 4.3.5. *The following statements are equivalent:*

- (1) $f(x)$ is κ -uniformly convex and differentiable;
- (2) $(x - y)^*(\nabla f(x) - \nabla f(y)) \geq \kappa\|x - y\|^2$ for all x, y .

PROOF. This follows immediately from the well-known first order conditions for convexity [HUL01], applied to the convex function $f(x) - \frac{1}{2}\kappa\|x\|^2$. \square

To explain the basic method, let us begin by proving when the value function is simply convex. Recall that the cost-to-go is given by

$$\check{J}^{\check{u}}(t, x) = \tilde{\mathbf{E}}_{t,x} \left[\frac{1}{2} \int_t^T \{ \|\check{u}_s\|^2 + W(\check{X}_s^{\check{u}}) + G_s^y(\check{X}_s^{\check{u}}) \} ds + R(\check{X}_T^{\check{u}}) \right].$$

We can write this differently, as follows: define

$$d\check{X}_s^{\check{u},t,x} = H^* y_{T-s} ds - \frac{1}{2}(B - B^*)\check{X}_s^{\check{u},t,x} ds + \check{u}_s ds + d\check{W}_s, \quad \check{X}_t^{\check{u},t,x} = x,$$

where \check{u}_s is any $\sigma\{\check{W}_r - \check{W}_t : r \in [t, s]\}$ -adapted control strategy. Then

$$\check{J}^{\check{u}}(t, x) = \mathbf{E} \left[\frac{1}{2} \int_t^T \{ \|\check{u}_s\|^2 + W(\check{X}_s^{\check{u},t,x}) + G_s^y(\check{X}_s^{\check{u},t,x}) \} ds + R(\check{X}_T^{\check{u},t,x}) \right],$$

and $\check{V}(t, x) = \min_{\check{u}} \check{J}^{\check{u}}(t, x)$. Now note that as the equation for $\check{X}_s^{\check{u},t,x}$ is linear in x and \check{u} , it follows immediately that

$$\check{X}_s^{\lambda\check{u} + (1-\lambda)\check{v}, t, \lambda x + (1-\lambda)y} = \lambda\check{X}_s^{\check{u},t,x} + (1-\lambda)\check{X}_s^{\check{v},t,y}$$

for any pair of controls \check{u}, \check{v} , initial points x, y , and $\lambda \in [0, 1]$. We can exploit this fact to prove that the cost-to-go is convex, provided that $W(x)$ and $R(x)$ are convex. Indeed, $\|\check{u}\|^2$ and $G_t^y(x)$ are obviously already convex, so we obtain

$$\check{J}^{\lambda\check{u} + (1-\lambda)\check{v}}(t, \lambda x + (1-\lambda)y) \leq \lambda\check{J}^{\check{u}}(t, x) + (1-\lambda)\check{J}^{\check{v}}(t, y).$$

Now take the infimum over \check{u} and \check{v} . This gives

$$\check{V}(t, \lambda x + (1-\lambda)y) \leq \lambda\check{V}(t, x) + (1-\lambda)\check{V}(t, y).$$

Hence evidently the value function $\check{V}(t, x)$ is convex, provided that $W(x)$ and $R(x)$ are chosen to be convex. We now extend this idea to the uniformly convex case.

PROPOSITION 4.3.6. *Suppose that $R(x)$ is κ -uniformly convex and that $W(x)$ is κ' -uniformly convex, $\kappa' \geq 2\kappa^2$. Then $\check{V}(t, x)$ is κ -uniformly convex for any t .*

PROOF. We would like to show that $\check{J}^{\check{u}}(t, x) - \frac{1}{2}\kappa\|x\|^2$ is convex in (\check{u}, x) . To this end, we adopt our seasoned approach, i.e., apply Dynkin's formula to $\|x\|^2$:

$$\mathbf{E}(\|\check{X}_T^{\check{u}, t, x}\|^2) = \|x\|^2 + (T-t)d + \mathbf{E} \left[\int_t^T 2\check{X}_s^{\check{u}, t, x} \cdot (H^* y_{T-s} + \check{u}_s) ds \right].$$

Hence we can write, rearranging the terms suggestively,

$$\begin{aligned} \check{J}^{\check{u}}(t, x) - \frac{1}{2}\kappa\|x\|^2 &= \frac{1}{2}\kappa(T-t)d + \mathbf{E} \left[\frac{1}{2} \int_t^T \{W(\check{X}_s^{\check{u}, t, x}) - \frac{1}{2}\kappa\|\check{X}_s^{\check{u}, t, x}\|^2\} ds \right] \\ &+ \mathbf{E} \left[\frac{1}{2} \int_t^T \{G_s^y(\check{X}_s^{\check{u}, t, x}) + 2\kappa y_{T-s}^* H \check{X}_s^{\check{u}, t, x}\} ds + R(\check{X}_T^{\check{u}, t, x}) - \frac{1}{2}\kappa\|\check{X}_T^{\check{u}, t, x}\|^2 \right] \\ &+ \mathbf{E} \left[\frac{1}{2} \int_t^T \{\|\check{u}_s\|^2 + 2\kappa \check{X}_s^{\check{u}, t, x} \cdot \check{u}_s + \frac{1}{2}\kappa'\|\check{X}_s^{\check{u}, t, x}\|^2\} ds \right]. \end{aligned}$$

By our assumptions all the terms in this expression are convex, except possibly the last line. We are thus interested in making $\|u\|^2 + 2\kappa x \cdot u + \frac{1}{2}\kappa'\|x\|^2$ a convex function in (x, u) . But it is easily verified, by requiring the Hessian of this expression to be positive semidefinite, that this is the case iff $\kappa' \geq 2\kappa^2$. Proceeding as in the proof of simple convexity above, the statement follows. \square

We will resume the proof of the main result in section 4.3.3. The rest of this section is devoted to a discussion of the role of the ‘‘gauge transformation’’ in the model under investigation; in particular, we will aim to clarify why these results do not easily extend to more general signal models.

Let us consider the simplest possible generalization: the signal-observation pair

$$\begin{aligned} dX_t &= b(X_t) dt + dW_t, \\ dY_t &= h(X_t) dt + dB_t, \end{aligned}$$

where $b(x)$ and $h(x)$ are not necessarily of linear-gradient form and of linear form, respectively. We forgo all technicalities and assume that everything is regular. To find the time-reversed conditional signal, we once again consider the control system

$$d\check{X}_t^u = -b(\check{X}_t^u) dt + \nabla \log p_{T-t}(\check{X}_t^u) dt + u_t dt + d\check{W}_t,$$

with the corresponding cost-to-go

$$\tilde{J}^u(t, x) = \tilde{\mathbf{E}}_{t,x} \left[\frac{1}{2} \int_t^T \left(\|u_s\|^2 + \|\dot{y}_{T-s} - h(\check{X}_s^u)\|^2 - \|\dot{y}_{T-s}\|^2 \right) ds \right].$$

It is convenient to absorb the term $-\log p_{T-t}(x)$ into the value function. It is easily seen that the corresponding control system can be written as

$$d\check{X}_t^{\check{u}} = \check{u}_t dt + d\check{W}_t,$$

with the modified cost-to-go

$$\begin{aligned} \check{J}^{\check{u}}(t, x) = & \tilde{\mathbf{E}}_{t,x} \left[\int_t^T \left(\frac{1}{2} \|\check{u}_s + b(\check{X}_s^{\check{u}})\|^2 + (\nabla \cdot b)(\check{X}_s^{\check{u}}) \right) ds \right] \\ & + \tilde{\mathbf{E}}_{t,x} \left[\frac{1}{2} \int_t^T (\|\dot{y}_{T-s} - h(\check{X}_s^{\check{u}})\|^2 - \|\dot{y}_{T-s}\|^2) ds - \log p_0(\check{X}_T^{\check{u}}) \right]. \end{aligned}$$

The corresponding value function is $\check{V}(t, x) = \min_{\check{u}} \check{J}^{\check{u}}(t, x) = -\log \varrho_{T-t}^y(x)$, and the optimal control is given by $\check{u}_t^* = -b(\check{X}_t^{\check{u}^*}) - \nabla \check{V}(t, \check{X}_t^{\check{u}^*})$. It would appear that we can now proceed exactly as before: note that

$$\begin{aligned} \frac{d}{dt} \|\check{X}_t^{\check{u}^*}(x) - \check{X}_t^{\check{u}^*}(y)\|^2 = & -2(\check{X}_t^{\check{u}^*}(x) - \check{X}_t^{\check{u}^*}(y))^*(b(\check{X}_t^{\check{u}^*}(x)) - b(\check{X}_t^{\check{u}^*}(y))) \\ & - 2(\check{X}_t^{\check{u}^*}(x) - \check{X}_t^{\check{u}^*}(y))^*(\nabla \check{V}(t, \check{X}_t^{\check{u}^*}(x)) - \nabla \check{V}(t, \check{X}_t^{\check{u}^*}(y))). \end{aligned}$$

Hence if we impose the mild condition that $(x - y)^*(b(x) - b(y)) \geq -\alpha\|x - y\|^2$ for some $\alpha \in \mathbb{R}$ —essentially the requirement that the rate of expansion of the uncontrolled diffusion is bounded—then contractivity of the time-reversed conditional diffusion would follow if $\check{V}(t, x)$ could be made κ -uniformly convex with $\kappa > \alpha$.

The problems quickly become evident, however, when we try to show that the cost-to-go is uniformly convex. Consider even the simple convex case, and let us assume for simplicity that $d = p = 1$ (one-dimensional signal and observations). In order for the observation term in the cost to be convex, we must have

$$\frac{d^2}{dx^2} (y - h(x))^2 = (h(x) - y) \frac{d^2 h}{dx^2} + \left(\frac{dh}{dx} \right)^2 \geq 0.$$

But this can clearly never be the case uniformly in y , except if $d^2 h/dx^2 = 0$ which implies $h(x) = Hx$. Hence evidently linear observations are essential in the proof of convexity of $\check{V}(t, x)$. Similarly, we need to show that $\frac{1}{2}(u + b(x))^2 + db(x)/dx$ is convex in (x, u) , but this is easily verified to be impossible unless $b(x)$ is linear. In this case, however, we have another trick at our disposal. Clearly the difficult term is $b(\check{X}_s^{\check{u}}) \check{u}_s$, which is obtained by expanding the square. But note that

$$\tilde{\mathbf{E}}_{t,x} \int_t^T b(\check{X}_s^{\check{u}}) \check{u}_s ds = \tilde{\mathbf{E}}_{t,x} \int_t^T b(\check{X}_s^{\check{u}}) d\check{X}_s^{\check{u}}.$$

Hence if $b(x) = dV(x)/dx$ for some function $V(x)$, we can transform away the problematic term:

$$\tilde{\mathbf{E}}_{t,x} \int_t^T b(\check{X}_s^{\check{u}}) d\check{X}_s^{\check{u}} = \tilde{\mathbf{E}}_{t,x} \left[V(\check{X}_T^{\check{u}}) - V(x) - \frac{1}{2} \int_t^T \frac{db}{dx}(\check{X}_s^{\check{u}}) ds \right].$$

But this corresponds precisely to the (essential portion of) the gauge transformation which we performed previously, and requires the signal drift to be of gradient type.

In conclusion, there is no a priori necessity to restrict to our linear-gradient model, nor to perform a gauge transformation at the beginning of the argument as we have done. However, once we start working our way through the proofs, it quickly becomes evident that the gauge transformation, and hence the linear-gradient model, are essential tools in proving that the value function is (uniformly) convex. Of course, using convexity in order to prove contractivity of the conditional signal is a rather primitive method; in particular, we do not really need the almost

sure contractivity that is guaranteed by this method, as contractivity in the mean would certainly suffice. Proving the latter, however, requires more sophisticated tools, and these still remain to be developed.

4.3.3. Proof of the main result. We have seen that $|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)|$ decays at an exponential rate, provided that the value function $\check{V}(t, x)$ is uniformly convex. Moreover, we showed that the latter holds if $W(x)$ and $R(x)$ are uniformly convex. The requirement that $W(x)$ is uniformly convex will be a basic requirement of our main result. The requirement that $R(x)$ be uniformly convex might seem rather strange, however, as this places a significant restriction on the initial measure ν .

It is important to realize that this is not in fact particularly restrictive. As we will obtain an a.s. filter stability bound, it is essentially irrelevant what the *true* initial measure is. In particular, we can bound the filter discrepancy given any pair of initial measures β, β' using the triangle inequality:

$$\|\pi_t^\beta - \pi_t^{\beta'}\|_{\text{TV}} \leq \|\pi_t^\beta - \pi_t^\nu\|_{\text{TV}} + \|\pi_t^{\beta'} - \pi_t^\nu\|_{\text{TV}}.$$

Hence as long as $\beta, \beta' \ll \nu$, the choice of ν need not at all reflect the true initial measure of the system. Instead, ν plays the role of a reference measure, and will be chosen specifically to make $R(x)$ uniformly convex.

Let us begin by introducing the basic assumptions that are needed for the main result. We will add additional assumptions later when necessary.

ASSUMPTION 4.3.7. The function $W(x)$, defined in (4.4), is $2\kappa^2$ -uniformly convex for some positive constant $\kappa > 0$.

ASSUMPTION 4.3.8. The density $p_0(x)$ of ν , satisfying the conditions of section 4.1, is chosen such that the function $R(x)$ of (4.6) is κ -uniformly convex.

ASSUMPTION 4.3.9. $\beta, \beta' \ll \nu$ are such that $d\beta/d\nu, d\beta'/d\nu$ are globally Lipschitz continuous and bounded from below by some constant $\varepsilon > 0$.

The result of the previous sections can now be written as follows.

LEMMA 4.3.10. *Under Assumptions 4.3.7–4.3.9, we have*

$$|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)| \leq \left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} e^{-\kappa T} \|x - y\|.$$

The equivalent result holds when β is replaced by β' .

PROOF. Recall that we can write

$$|\Lambda_T^{\beta\nu}(x) - \Lambda_T^{\beta\nu}(y)| \leq \left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} \mathbf{E} \|\tilde{X}_T^*(x) - \tilde{X}_T^*(y)\|,$$

and that $\|\tilde{X}_T^*(x) - \tilde{X}_T^*(y)\|$ has the following time derivative:

$$\frac{d}{dt} \|\tilde{X}_t^*(x) - \tilde{X}_t^*(y)\|^2 = -2(\tilde{X}_t^*(x) - \tilde{X}_t^*(y))^* (\nabla \check{V}(t, \tilde{X}_t^*(x)) - \nabla \check{V}(t, \tilde{X}_t^*(y))).$$

By Proposition 4.3.6, $\check{V}(t, x)$ is κ -uniformly convex for every $t \in [0, T]$. Hence

$$\frac{d}{dt} \|\tilde{X}_t^*(x) - \tilde{X}_t^*(y)\|^2 \leq -2\kappa \|\tilde{X}_t^*(x) - \tilde{X}_t^*(y)\|^2.$$

The result follows immediately. \square

To proceed, recall that we obtained the estimate

$$\|\pi_t^\beta - \pi_t^\nu\|_{\text{TV}} \leq \frac{\int_{\mathbb{R}^d \times \mathbb{R}^d} \left| \Lambda_t^{\beta\nu}(x) - \Lambda_t^{\beta\nu}(y) \right| \pi_t^\nu(x) \pi_t^\nu(y) dx dy}{\int_{\mathbb{R}^d} \Lambda_t^{\beta\nu}(y) \pi_t^\nu(y) dy}.$$

Using the previous Lemma and Assumption 4.3.9, we obtain

$$\|\pi_t^\beta - \pi_t^\nu\|_{\text{TV}} \leq \frac{1}{\varepsilon} \left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} e^{-\kappa t} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\| \pi_t^\nu(x) \pi_t^\nu(y) dx dy.$$

It remains to estimate the double integral. Note that

$$\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\| \pi_t^\nu(x) \pi_t^\nu(y) dx dy \leq \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 \pi_t^\nu(x) \pi_t^\nu(y) dx dy \right]^{1/2}$$

by Jensen's inequality. But clearly

$$\frac{1}{2} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 \pi_t^\nu(x) \pi_t^\nu(y) dx dy = \int \|x\|^2 \pi_t^\nu(x) dx - \left\| \int x \pi_t^\nu(x) dx \right\|^2.$$

Hence the double integral term is a measure of the conditional variance of the signal. If the filter is stable, then it is not implausible that this quantity is uniformly bounded; i.e., we have a guaranteed bound on the estimation quality. Proving this, however, is not necessarily a straightforward exercise.

To obtain this final estimate, we use the idea of W. Stannat [Sta06] of using an inequality of H. J. Brascamp and E. H. Lieb [BL76]. In this often cited paper, those authors prove a Poincaré-type inequality for log-concave measures which can be stated as follows. Let μ be a measure which has density $p(x)$ with respect to the Lebesgue measure on \mathbb{R}^n . We suppose that $-\log p(x)$ is C^2 and κ -uniformly convex. Moreover, let $h(x) \in C^1$ be such that $\text{var}_\mu(h) = \mu(h^2) - \mu(h)^2 < \infty$. Then

$$\text{var}_\mu(h) \leq \frac{\mu(\|\nabla h\|^2)}{\kappa}.$$

(In fact the statement of Brascamp and Lieb is a little more general, but we will only need this immediate corollary.) In section 4.3.4 we will give a new proof of this result which is completely probabilistic. For the time being, however, let us complete the proof of the main result by applying the Brascamp-Lieb inequality.

If we could guarantee that $-\log \pi_t^\nu(x)$ is uniformly convex, then we could directly apply the Brascamp-Lieb inequality. This is easily done provided that the signal process is not excessively unstable.

DEFINITION 4.3.11. $f(x)$ is called α -semiconcave if $f(x) - \frac{1}{2}\alpha\|x\|^2$ is concave.

THEOREM 4.3.12. *Suppose that Assumptions 4.3.7–4.3.9 hold, and assume furthermore that $V(x) + \frac{1}{2}x^*Bx$ is α -semiconcave with $\alpha < \kappa$. Then*

$$\|\pi_t^\beta - \pi_t^{\beta'}\|_{\text{TV}} \leq \frac{1}{\varepsilon} \sqrt{\frac{2d}{\kappa - \alpha}} \left(\left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} + \left\| \frac{d\beta'}{d\nu} \right\|_{\text{Lip}} \right) e^{-\kappa t}.$$

REMARK 4.3.13. Note that α is a measure of how unstable the signal process is, while κ measures the stability of the reverse time conditional signal. The current bound is finite when the reverse time conditional signal is more stable than the signal is unstable, i.e., when the information gain from the observations makes up for the inherent information loss in the signal. This is a very intuitive idea.

PROOF. It is easily seen that we can write

$$\left\| \int \|x\|^2 \pi_T^\nu(x) dx - \left\| \int x \pi_T^\nu(x) dx \right\|^2 \right\| = \sum_{i=1}^d \text{var}_{\pi_T^\nu}(x^i).$$

Now recall that $\tilde{V}(0, x) = -\log \varrho_T^y(x) + y_T^* Hx + V(x) + \frac{1}{2} x^* Bx$ is κ -uniformly convex. Hence $-\log \varrho_T^y(x)$ is $(\kappa - \alpha)$ -uniformly convex, and this holds for $-\log \pi_T^\nu(x)$ as well (as $\pi_T^\nu(x)$ and $\varrho_T^y(x)$ coincide up to normalization). But then we obtain

$$\left\| \int \|x\|^2 \pi_T^\nu(x) dx - \left\| \int x \pi_T^\nu(x) dx \right\|^2 \right\| \leq \frac{d}{\kappa - \alpha}$$

using the Brascamp-Lieb inequality. The result follows immediately. \square

We note that Stannat uses the Brascamp-Lieb inequality in a somewhat different way. Suppose there exists a concave function $V'(x)$ and some $M > 0$ s.t.

$$V'(x) - M \leq V(x) + \frac{1}{2} x^* Bx \leq V'(x) + M.$$

Then we can estimate the conditional density as follows:

$$\pi_T^\nu(x) = \frac{e^{y_T^* Hx + V(x) + \frac{1}{2} x^* Bx - \tilde{V}(0, x)}}{\int e^{y_T^* Hx + V(x) + \frac{1}{2} x^* Bx - \tilde{V}(0, x)} dx} \leq e^{2M} \frac{e^{y_T^* Hx + V'(x) - \tilde{V}(0, x)}}{\int e^{y_T^* Hx + V'(x) - \tilde{V}(0, x)} dx}.$$

The right-hand side is, up to the factor e^{2M} , another probability density $\bar{\pi}(x)$ such that $-\log \bar{\pi}(x)$ is κ -uniformly convex. Writing

$$\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\| \pi_T^\nu(x) \pi_T^\nu(y) dx dy \leq e^{4M} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\| \bar{\pi}(x) \bar{\pi}(y) dx dy,$$

and proceeding as before, we obtain the bound

$$\|\pi_t^\beta - \pi_t^{\beta'}\|_{\text{TV}} \leq \frac{e^{4M}}{\varepsilon} \sqrt{\frac{2d}{\kappa}} \left(\left\| \frac{d\beta}{d\nu} \right\|_{\text{Lip}} + \left\| \frac{d\beta'}{d\nu} \right\|_{\text{Lip}} \right) e^{-\kappa t}.$$

The condition of Theorem 4.3.12 gives a much more vivid picture, however, of what is actually going on in terms of “information flows”: in essence the contraction rate of the time-reversed conditional signal must exceed the rate of expansion of the signal process itself. In particular, in that case the conditional density $\pi_T^\nu(x)$ is itself uniformly log-concave, and there is no need for any additional estimates.

4.3.4. A probabilistic proof of a Brascamp-Lieb inequality. In this section we are going to give a probabilistic proof of the Brascamp-Lieb type inequality used in the previous section. This section is independent from the rest of this chapter, though we will use surprisingly familiar techniques in the proofs. As the method is of independent interest, I will give a detailed proof here. The idea behind this section is inspired by a recent note of A. N. Shiryaev [Shi06], who applies similar techniques to provide a probabilistic proof of the Poincaré-Chernoff inequality for a Gaussian random variable (related ideas can be found in [CHL97, HS87]).

The main idea behind the proof is the following. Consider the equation

$$(4.7) \quad dx_t = -\frac{1}{2} \nabla U(x_t) dt + dW_t, \quad x_0 = x,$$

where $U : \mathbb{R}^d \rightarrow \mathbb{R}$ is a C^2 potential and W_t is a d -dimensional Wiener process. It is easily verified that this equation has an invariant measure $\mu \ll dx$ with density

$$(4.8) \quad d\mu = \frac{\exp(-U(x)) dx}{\int \exp(-U(x)) dx},$$

provided that the denominator is finite. We will prove that the variance of $h(x_t)$ satisfies a Brascamp-Lieb type inequality, provided that U is uniformly convex. The result then follows by taking the limit $t \rightarrow \infty$.

Let us first prove the result in a case with plenty of regularity. We will remove the restrictions below.

PROPOSITION 4.3.14. *Let $U \in C^2$ be κ -uniformly convex and have bounded second derivatives. Let $h \in C^1$, and let $h(x)$ and $\nabla h(x)$ be bounded. Then*

$$\text{var}_\mu(h) = \mu(h^2) - \mu(h)^2 \leq \frac{\mu(\|\nabla h\|^2)}{\kappa}.$$

PROOF. First, note that $U(x)$ being uniformly convex implies that $\exp(-U(x))$ decays exponentially in all directions. Hence the denominator in (4.8) is finite, and μ defines a log-concave probability measure. Moreover, the fact that $U(x)$ has bounded second derivative means that the drift in (4.7) is globally Lipschitz continuous, and we can apply standard stochastic machinery. In particular, this equation generates a C^1 stochastic flow $\xi_t(x)$.

Consider the quantity $h(x_t)$. We are going to obtain an explicit expression for the variance of this random variable by applying the Clark-Ocone formula, Prop. A.3.1. Note that by Prop. A.2.2, x_t is in $\mathbb{D}^{1,2}$ for any t . Moreover, by Prop. A.1.1 and by the boundedness of h and its derivatives, $h(x_t)$ is also in $\mathbb{D}^{1,2}$. Applying Propositions A.3.1, A.1.1 and A.2.2, we obtain

$$h(x_t) - \mathbf{E}(h(x_t)) = \int_0^t \mathbf{E}(\nabla h(x_t)^* D\xi_t(x) D\xi_s(x)^{-1} | \mathcal{F}_s) dW_s.$$

Hence we obtain by the Itô isometry (we write $\text{var}(h(x_t)) = \mathbf{E}(h(x_t)^2) - [\mathbf{E}(h(x_t))]^2$)

$$\text{var}(h(x_t)) = \int_0^t \mathbf{E} \|\mathbf{E}(\nabla h(x_t)^* D\xi_t(x) D\xi_s(x)^{-1} | \mathcal{F}_s)\|^2 ds.$$

We now obtain the straightforward estimate

$$\text{var}(h(x_t)) \leq \int_0^t \mathbf{E}(\|\nabla h(x_t)\|^2 \|D\xi_t(x) D\xi_s(x)^{-1}\|^2) ds.$$

We need to bound $\|D\xi_t(x) D\xi_s(x)^{-1}\|^2$. It follows from standard results that

$$\frac{d}{dt} D\xi_t(x) D\xi_s(x)^{-1} = -\frac{1}{2} U_{xx}(x_t) D\xi_t(x) D\xi_s(x)^{-1},$$

where U_{xx} denotes the Hessian of U . We easily find that

$$\frac{d}{dt} \|D\xi_t(x) D\xi_s(x)^{-1} v\|^2 \leq -\kappa \|D\xi_t(x) D\xi_s(x)^{-1} v\|^2,$$

where $v \in \mathbb{R}^d$ is any vector and we have used the uniform convexity of U . Hence

$$\|D\xi_t(x) D\xi_s(x)^{-1}\|^2 \leq e^{-\kappa(t-s)}.$$

Substituting into the bound above, we obtain

$$\text{var}(h(x_t)) \leq \mathbf{E}(\|\nabla h(x_t)\|^2) \int_0^t e^{-\kappa(t-s)} ds = \frac{\mathbf{E}(\|\nabla h(x_t)\|^2)}{\kappa} (1 - e^{-\kappa t}).$$

The result is now almost obvious: we only need to justify the limit as $t \rightarrow \infty$. To this end, let \tilde{x}_t be the solution of

$$d\tilde{x}_t = -\frac{1}{2} \nabla U(\tilde{x}_t) dt + dW_t, \quad \tilde{x}_0 \sim \mu.$$

As μ is an invariant measure for this equation, $\mathbf{E}(f(\tilde{x}_t)) = \mu(f)$ for any f and t . On the other hand, it is easily verified (as we have done repeatedly in this chapter) that $\|x_t - \tilde{x}_t\| \rightarrow 0$ a.s. as $t \rightarrow \infty$. Hence $\mathbf{E}(f(x_t)) \rightarrow \mu(f)$ as $t \rightarrow \infty$ for any bounded continuous f by dominated convergence. The result follows. \square

We now remove the technical restrictions in this result.

COROLLARY 4.3.15. *Let $U \in C^2$ be κ -uniformly convex, and let $h \in C^1$ be such that $\text{var}_\mu(h) < \infty$. Then $\text{var}_\mu(h) \leq \mu(\|\nabla h\|^2)/\kappa$.*

PROOF. This is essentially an exercise in approximation; we give here one possible approach. Let us begin by eliminating the requirement that U have bounded second derivatives. To this end, we construct a sequence of C^2 functions U_ℓ such that U_ℓ has bounded second derivatives and is κ_ℓ -uniformly convex for any ℓ , and $\kappa_\ell \rightarrow \kappa$, $U_\ell \nearrow U$ as $\ell \rightarrow \infty$. Such a sequence of approximations exists and can be constructed through the Moreau-Yosida regularization [**HUL93**, **LS97**]:

$$U_\ell(x) = \min\{U(y) + \frac{1}{2}\ell\|x - y\|^2 : y \in \mathbb{R}^d\}.$$

See [**LS97**] for the C^2 property and uniform convexity of U_ℓ , and for the remaining properties see [**HUL93**, pp. 317–321]. Denote by μ_ℓ the measure corresponding to the potential function U_ℓ . By Proposition 4.3.14, we have

$$\text{var}_{\mu_\ell}(h) \leq \frac{\mu_\ell(\|\nabla h\|^2)}{\kappa_\ell}, \quad h, \nabla h \text{ bounded.}$$

Letting $\ell \rightarrow \infty$ and using dominated convergence, we establish the claim of the corollary for $h \in C^1$ such that $h(x)$ and $\nabla h(x)$ are bounded.

It remains to drop the restriction on h . To this end, introduce a smooth function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $0 \leq \phi \leq 1$, $\phi(x) = 1$ for $\|x\| \leq 1$ and ϕ has compact support. Let $\phi_n(x) = \phi(x/n)$; then $h(x)\phi_n(x)$ is bounded and has bounded derivatives, so

$$\text{var}_\mu(h\phi_n) \leq \frac{\mu(\|\nabla(h\phi_n)\|^2)}{\kappa} \quad \text{for all } n.$$

Evidently $h\phi_n \rightarrow h$, $\nabla(h\phi_n) \rightarrow \nabla h$ as $n \rightarrow \infty$. We will justify the exchange of limit and μ -expectation by dominated convergence. Note that

$$\|\nabla(h\phi_n)\|^2 = \|\phi_n \nabla h + h \nabla \phi_n\|^2 \leq 2\phi_n^2 \|\nabla h\|^2 + 2h^2 \|\nabla \phi_n\|^2 \leq 2\|\nabla h\|^2 + 2Ch^2,$$

where $C = \sup_{x \in \mathbb{R}^d} \|\nabla \phi(x)\|^2 < \infty$. But recall that we require $\text{var}_\mu(h) < \infty$, and moreover we may presume that $\mu(\|\nabla h\|^2) < \infty$ (otherwise the statement is trivial). Hence we have established that $\|\nabla(h\phi_n)\|^2$ is dominated by the integrable function $2\|\nabla h\|^2 + 2Ch^2$, and the statement follows as $n \rightarrow \infty$. \square

REMARK 4.3.16. The condition $\text{var}_\mu(h) < \infty$ is automatically satisfied if $h(x)$ grows at most polynomially as $\|x\| \rightarrow \infty$; this follows from the fact that e^{-U} decays exponentially in all directions by uniform convexity of U .

4.4. A strong stability result for stable signals

The filter stability bound obtained in the previous section has several drawbacks. First, the densities of β, β' are required to be Lipschitz continuous. This is not a significant problem, however; the Lipschitz requirement can be relaxed by using a suitably modified procedure to bound the numerator of the Bayes formula (see, e.g., [Sta04, Sta05]). A more severe problem is that the bound blows up as $\varepsilon \rightarrow 0$, i.e., as β, β' approach the boundary of the space of measures on \mathbb{R}^d . Hence we cannot say anything, for example, about initial measures that are compactly supported or about deterministic initial conditions.

The ε^{-1} factor can be traced to the fact that we used the Bayes formula to express the incorrectly initialized filter in terms of the time-reversed conditional signal. The denominator of the Bayes formula had to be bounded, and this prompted us to impose a lower bound on the densities of β, β' . This problem would not occur, however, if we used the forward conditional signal rather than the time-reversed conditional signal. The finite state space analysis of chapter 2 suggests that this is not only feasible, but is a much preferable approach.

In this section we will pursue this approach in the diffusion signal case. We will see that we can get a bound that does not blow up at the boundary, but we pay a heavy price: we have to require that the signal process is itself stable, in which case the forgetting rate of the filter is inherited from the signal. This is in contrast to the approach taken in the previous sections, where a sufficiently good observation could make up for the instability of the signal. The reason for this difference is that we lose a valuable tool in the forward conditional signal case—the use of a conveniently chosen ν to force uniform convexity of the value function.

We begin by quoting the necessary conditional signal results from [MN03]. Consider the controlled signal process

$$dX_t^u = b(X_t^u) dt + u_t dt + dW_t,$$

and introduce the following cost-to-go (for $y \in C^1$):

$$J^u(t, x) = \mathbf{E}_{t,x} \left[\frac{1}{2} \int_t^T \{ \|u_t\|^2 + \|\dot{y}_t - HX_t^u\|^2 - \|\dot{y}_t\|^2 \} dt \right].$$

The value function $V(t, x) = \min_u J^u(t, x)$ is minus the logarithm of the dual filter (see [Par82]), and the optimal control strategy is given by $u_t^* = -\nabla V(t, X_t^{u^*})$. Moreover, the conditional signal

$$dX_t^* = b(X_t^*) dt - \nabla V(t, X_t^*) dt + dW_t, \quad X_0^* \sim \Pi_T^\beta(X_0 \in \cdot, y),$$

has the same law under \mathbf{P} as does X_t under $\Pi_T^\beta(\cdot, y)$. Note in particular that $V(t, x)$ does not depend on β . As before, we proceed by absorbing the gradient part of $b(x)$ into the cost-to-go. That is, we define

$$\check{u}_t = u_t + \nabla V(X_t^u) + \frac{1}{2}(B + B^*)X_t^u + H^*y_t,$$

so that we can write the controlled diffusion as

$$d\check{X}_t^{\check{u}} = -H^*y_t dt + \frac{1}{2}(B - B^*)\check{X}_t^{\check{u}} dt + \check{u}_t dt + dW_t.$$

Similarly, we redefine the cost-to-go as

$$\check{J}^{\check{u}}(t, x) = J^u(t, x) - V(x) - \frac{1}{2}x^* Bx - y_t^* Hx,$$

and we define the new value function by $\check{V}(t, x) = \min_{\check{u}} \check{J}^{\check{u}}(t, x)$. The optimal control can now be written as $\check{u}_t^* = -\nabla \check{V}(t, \check{X}_t^{\check{u}})$. Applying Dynkin's formula, it is a straightforward exercise to write the cost-to-go $\check{J}^{\check{u}}(t, x)$ in the standard form:

$$\check{J}^{\check{u}}(t, x) = \mathbf{E}_{t,x} \left[\frac{1}{2} \int_t^T \left\{ \|\check{u}_t\|^2 + W(\check{X}_t^{\check{u}}) + G_{T-t}^y(\check{X}_t^{\check{u}}) \right\} dt + \check{R}(\check{X}_T^{\check{u}}) \right],$$

where $W(x)$ and $G_t^y(x)$ are as before, and the terminal cost is given by

$$\check{R}(x) = -V(x) - \frac{1}{2}x^*Bx - y_T^*Hx.$$

We can now obtain the following result in the same way as Proposition 4.3.6.

PROPOSITION 4.4.1. *Suppose that $V(x) + \frac{1}{2}x^*Bx$ is γ -uniformly concave, and $W(x)$ is κ' -uniformly convex with $\kappa' \geq 2\gamma^2$. Then $\check{V}(t, x)$ is γ -uniformly convex.*

Let us concentrate, for sake of demonstration, on the discrepancy between the filters with deterministic initial measures δ_x and $\delta_{x'}$ (with a little more work more general bounds are obtainable; this particularly simple case is in some sense a worst case scenario, where the initial measures are extremal and mutually singular). Denoting by $X_t^*(x)$ the conditional signal process started at $X_0^* = x$, evidently

$$\Pi_T^{\delta_x}(f(X_T), y) = \mathbf{E}(f(X_T^*(x))), \quad \Pi_T^{\delta_{x'}}(f(X_T), y) = \mathbf{E}(f(X_T^*(x'))).$$

It is easily verified that by Proposition 4.4.1 (under the appropriate conditions)

$$\|X_T^*(x) - X_T^*(x')\| \leq e^{-\gamma T} \|x - x'\|,$$

so that we obtain (assuming that f is Lipschitz continuous)

$$|\Pi_T^{\delta_x}(f(X_T), y) - \Pi_T^{\delta_{x'}}(f(X_T), y)| \leq \|f\|_{\text{Lip}} \|x - x'\| e^{-\gamma T}.$$

In particular, we immediately obtain stability in terms of the Wasserstein 1-distance

$$W_1(\Pi_t^{\delta_x}(X_t \in \cdot, y), \Pi_t^{\delta_{x'}}(X_t \in \cdot, y)) \leq \|x - x'\| e^{-\gamma t}.$$

These arguments are trivially extended at least to the case where two initial measures β, β' are supported in a ball of radius $R < \infty$.

In conclusion, we have seen two stability bounds for diffusions. Both bounds required that $W(x)$ be uniformly convex. If, in addition, $V(x) + \frac{1}{2}x^*Bx$ is semiconcave with a not-too-large constant (i.e., the signal process is not too unstable), then stability is guaranteed for initial densities that are bounded from below. If, on the other hand, $V(x) + \frac{1}{2}x^*Bx$ is uniformly concave (i.e., the signal process is stable), then we have a bound even for extremal/mutually singular initial measures, using the forward conditional signal rather than its time-reversed counterpart.

The restrictions on the initial measures may be artefacts of our crude bounding procedure. Indeed, there is no reason to believe that something bad should happen to filter stability when the support of the initial measures is not all of \mathbb{R}^d , even when the signal itself is not stable. It is not so easy to obtain a uniform bound, however. It is interesting, in this context, to note the relative merits of the forward and time-reversed conditional signals. The forward conditional signal allows us to bound stability without using the Bayes formula, which makes it easier to apply to a wider class of initial conditions. Unfortunately, it is impossible for the forward value function $\check{V}(t, x)$ to be uniformly convex for all t unless $V(x) + \frac{1}{2}x^*Bx$ is uniformly concave: after all, $\check{V}(T, x) = \check{R}(x)$ would not be uniformly convex otherwise.

This is a real effect. Consider the observations on the interval $[0, T]$. The transition probabilities of the conditional signal at time t can clearly depend only on the observations in the interval $[t, T]$. But when t is close to the terminal time T , there is very little information in the remaining observations, so that the transition probabilities should then be close to those of the unconditional signal process. Consequently, if the unconditional signal is unstable, then the forward conditional signal should always be unstable close to the terminal time T —even if the signal is strictly contracting during most of the interval $[0, T]$.

Evidently it is very well possible for the forward conditional signal to be contracting overall, but it is not possible for this to be the case uniformly on the entire interval $[0, T]$. This highlights the difficulty of using the forward conditional signal to bound filter stability; a much more sophisticated analysis would be required to quantify the time-dependent contraction rate of the forward conditional signal. The time-reversed conditional signal does not suffer from these problems; indeed, by choosing a convenient reference measure ν , we can make the time-reversed conditional signal contracting uniformly on the interval $[0, T]$. The drawback is now that the initial measures are restricted by the necessity of using the Bayes formula.

Quantum Filtering and Filter Stability

The goal of this chapter is twofold. First, we will develop in detail the theory of quantum filtering using the reference probability method. Quantum filtering theory has its origins in the work of V. P. Belavkin [Bel88, Bel92] (using various methods), but the reference probability method, which parallels the classical theory described in chapter 1 and makes systematic use of change of measure techniques, is due to L. Bouten and the author [BV06] (see also [BVJ06b, BVJ06a]). Here we follow [BV06], but with increased attention to technical details. The second part of the chapter develops a first filter stability result for quantum filters.

Throughout this chapter, we concentrate on the case where the initial system lives on a finite-dimensional Hilbert space. This case, which is the quantum counterpart of the Wonham filter, is very common in the literature and is the most useful in applications. Our results can be generalized to more general systems, however, with no conceptual and only relatively straightforward technical improvements; a full account of the general case will appear in [BV07].

Before starting this chapter, the reader is strongly encouraged to read the background material in appendix B. The reader who is unfamiliar with quantum probability theory, which will be used throughout this chapter, will find there a crash course on the necessary topics. The appendix also serves to set the notation which we will use throughout this chapter, so that even the quantum probability veteran would be advised at least to skim the appendix before proceeding.

5.1. Conditional expectations and the Bayes formula

5.1.1. Conditional expectations in quantum probability. Before we can develop filtering theory, we need a suitable notion of conditional expectations in quantum probability. Many such notions have been introduced over the past few decades, but most of these are motivated more by formal similarity to the classical case than by truly probabilistic considerations. The one exception is the notion of conditional expectation that is introduced in quantum filtering theory. This notion is very restrictive—actually, it is just the classical conditional expectation, “pulled back” using the spectral theorem—but it is nonetheless the most general notion of conditional expectation that has a well-defined operational meaning.

There are two considerations that play a role in defining a probabilistic conditional expectation. First, we can only condition on something that can actually be observed; if we can not observe the events which are being conditioned on, what does conditioning even mean? In the noncommutative case, this imposes the following restriction: *we can only condition on commutative algebras*. Next, we should

consider what can be conditioned. Suppose that we would like to find the conditional probability of the occurrence of a certain event, given a particular set of (commuting) observations. Such a concept is meaningless if the event in question does not commute with the observations, as in that case the joint statistics of the event and the observations is an undefined concept (i.e., if the probability of an event is meaningless in a realization where we have made our observations, then certainly the conditional probability of that event is a meaningless concept). This imposes the second restriction: *we can only condition observables that commute with the observations on which we are conditioning.*

With these restrictions in place, there is little noncommutativity left. Let $(\mathcal{A}, \mathbb{P})$ be a quantum probability space, $X \in \mathcal{A}$ an observable and $\mathcal{C} \subset \mathcal{A}$ a Von Neumann subalgebra. By our discussion above, the conditional expectation $\mathbb{P}(X|\mathcal{C})$ can only be defined if \mathcal{C} is commutative and if X commutes with every element of \mathcal{C} . But then $\mathcal{C}_X = \text{vN}\{X, \mathcal{C}\}$ is again a commutative algebra, and the conditional expectation can be defined on a completely classical basis: i.e., if ι_X is the *-isomorphism obtained by applying the spectral theorem to \mathcal{C}_X and \mathbf{P}_X is the corresponding measure, then we can simply define $\iota(\mathbb{P}(X|\mathcal{C})) = \mathbf{P}_X(\iota(X)|\sigma\{\iota(\mathcal{C})\})$. In this sense, the quantum conditional expectation is completely classical; but this is how it should be, as probability theory lives precisely on the commutative subalgebras of \mathcal{A} . Beware, however, of the following: if X and Y both commute with every element of \mathcal{C} , this need not imply that X and Y commute with each other. In this sense, we should see $\mathbb{P}(\cdot|\mathcal{C})$ as a truly noncommutative conditional state.

Let us first define the conditional expectation for bounded operators.

DEFINITION 5.1.1. Let $(\mathcal{A}, \mathbb{P})$ be a quantum probab. space, $\mathcal{C} \subset \mathcal{A}$ be a commutative subalgebra, and $\mathcal{C}' = \{X \in \mathcal{A} : XC = CX \ \forall C \in \mathcal{C}\}$ (the commutant of \mathcal{C} in \mathcal{A}). Let $X \in \mathcal{C}'$. Then any $\mathbb{P}(X|\mathcal{C}) \in \mathcal{C}$ such that $\mathbb{P}(\mathbb{P}(X|\mathcal{C})C) = \mathbb{P}(XC)$ for all $C \in \mathcal{C}$ is called (a version of) the conditional expectation of X given \mathcal{C} .

Let us show that the conditional expectation is well defined.

LEMMA 5.1.2. For any commutative \mathcal{C} and $X \in \mathcal{C}'$, $\mathbb{P}(X|\mathcal{C})$ exists and is unique \mathbb{P} -a.s., i.e., $\mathbb{P}((A - B)^2) = 0$ for any two versions A, B of $\mathbb{P}(X|\mathcal{C})$.

PROOF. For self-adjoint X , existence follows directly from the classical conditional expectation and the spectral theorem. For $X \neq X^*$, define $\mathbb{P}(X|\mathcal{C}) = \mathbb{P}(\frac{1}{2}(X + X^*)|\mathcal{C}) + i\mathbb{P}(\frac{i}{2}(X^* - X)|\mathcal{C})$; it is easily verified that this satisfies the definition of the conditional expectation. To verify uniqueness, let A and B be two versions of $\mathbb{P}(X|\mathcal{C})$. Then $\mathbb{P}((A - B)C) = 0$ for all $C \in \mathcal{C}$ by the definition of the conditional expectation. But $A, B \in \mathcal{C}$ by construction, so $\mathbb{P}((A - B)^2) = 0$. \square

The conditional expectation only has a truly probabilistic meaning for self-adjoint X ; however, the extension to all of \mathcal{C}' allows us to think of the conditional expectation as a sort of conditional state. All of the standard properties of the classical conditional expectation carry over to its quantum counterpart, including a.s. linearity and positivity, the tower property $\mathbb{P}(\mathbb{P}(X|\mathcal{C})|\mathcal{D}) = \mathbb{P}(X|\mathcal{D})$ for $\mathcal{D} \subset \mathcal{C}$, the module property $\mathbb{P}(AC|\mathcal{C}) = \mathbb{P}(A|\mathcal{C})C$ for $C \in \mathcal{C}$, etc. This follows immediately by copying the classical proofs of these properties; see, e.g., [Wil91].

A simple property of the conditional expectation which is worth emphasizing is its L^2 property. This provides an attractive interpretation of the conditional expectation as the optimal L^2 estimator in the quantum setting.

LEMMA 5.1.3. $\mathbb{P}(X|\mathcal{C})$ is the least-mean-square estimate of X given \mathcal{C} , i.e., for any $C \in \mathcal{C}$ we have $\|X - \mathbb{P}(X|\mathcal{C})\|_2 \leq \|X - C\|_2$, where $\|X\|_2 = (\mathbb{P}(X^*X))^{1/2}$.

PROOF. Note that $\|X - C\|_2^2 = \|X - \mathbb{P}(X|\mathcal{C}) + \mathbb{P}(X|\mathcal{C}) - C\|_2^2$. But by the definition of the conditional expectation $\mathbb{P}(C^*(X - \mathbb{P}(X|\mathcal{C}))) = 0$ for any $C \in \mathcal{C}$; so $\|X - C\|_2^2 = \|X - \mathbb{P}(X|\mathcal{C})\|_2^2 + \|\mathbb{P}(X|\mathcal{C}) - C\|_2^2 \geq \|X - \mathbb{P}(X|\mathcal{C})\|_2^2$. \square

Let us now extend the conditional expectation to unbounded observables.

DEFINITION 5.1.4. Let $(\mathcal{A}, \mathbb{P})$ be a quantum probability space and $\mathcal{C} \subset \mathcal{A}$ be a commutative subalgebra. Let $X \eta \mathcal{C}'$ be self-adjoint and suppose that $\mathbb{P}(|X|) < \infty$. Then any observable $\mathbb{P}(X|\mathcal{C}) \eta \mathcal{C}$ which satisfies $\mathbb{P}(\mathbb{P}(X|\mathcal{C}) \hat{\cdot} C) = \mathbb{P}(X \hat{\cdot} C)$ for all $C \in \mathcal{C}$ is called (a version of) the conditional expectation of X given \mathcal{C} .

The definition is well posed as \mathcal{C}_X is commutative, so that the class of self-adjoint operators $\mathcal{S}(\mathcal{C}_X)$ forms an algebra under $\hat{\cdot}$ and $\hat{+}$. Moreover, the conditional expectation exists and is unique up to \mathbb{P} -a.s. equivalence; this follows directly from the classical conditional expectation and the spectral theorem.

The quantum conditional expectation has some elementary properties which are meaningless in the classical case. The following simple result, which shows how $\mathbb{P}(\cdot|\mathcal{C})$ transforms under unitary rotations, will be put to good use in the sequel.

LEMMA 5.1.5. Let $(\mathcal{A}, \mathbb{P})$ be a quantum probability space, and let $\mathcal{C} \subset \mathcal{A}$ be a commutative subalgebra. Let U be a unitary operator such that $U^*\mathcal{A}U = \mathcal{A}$, and define the rotated state $\mathbb{Q}(X) = \mathbb{P}(U^*XU)$ on \mathcal{A} . Moreover, let $X \eta \mathcal{C}'$ be self-adjoint with $\mathbb{Q}(|X|) < \infty$. Then $\mathbb{P}(U^*XU|U^*\mathcal{C}U) = U^*\mathbb{Q}(X|\mathcal{C})U$ a.s.

PROOF. It suffices to show that $U^*\mathbb{Q}(X|\mathcal{C})U$ satisfies the definition of the conditional expectation. Note that for any $C \in \mathcal{C}$

$$\mathbb{P}(U^*\mathbb{Q}(X|\mathcal{C})U \hat{\cdot} U^*CU) = \mathbb{Q}(\mathbb{Q}(X|\mathcal{C}) \hat{\cdot} C) = \mathbb{Q}(X \hat{\cdot} C) = \mathbb{P}(U^*XU \hat{\cdot} U^*CU).$$

The claim follows from a.s. uniqueness of the conditional expectation. \square

5.1.2. The Bayes formula. As in chapter 1, the Bayes formula, which relates the conditional expectation with respect to different measures, will play an important role in the sequel. A simple analog of Lemma 1.1.1 is easy to establish in the quantum case; the main insight which is needed for this result is that it is crucial that the ‘‘Radon-Nikodym operator’’ be in the commutant of the algebra on which we are conditioning. The reader should convince himself that the statement of the result does not even make sense otherwise!

LEMMA 5.1.6. Let $(\mathcal{A}, \mathbb{P})$ be a quantum probability space, and let $\mathcal{C} \subset \mathcal{A}$ be a commutative subalgebra. Let $V \in \mathcal{C}'$ be such that $V^*V > 0$ and $\mathbb{P}(V^*V) = 1$. Then we can define a new normal state by $\mathbb{Q}(X) = \mathbb{P}(V^*XV)$, and

$$\mathbb{Q}(X|\mathcal{C}) = \frac{\mathbb{P}(V^*XV|\mathcal{C})}{\mathbb{P}(V^*V|\mathcal{C})} \quad \text{for any } X \in \mathcal{C}'.$$

PROOF. The fact that \mathbb{Q} is a normal state is easily verified. Now let K be an element of \mathcal{C} . Then for any $X \in \mathcal{C}'$, we can write

$$\begin{aligned} \mathbb{P}(\mathbb{P}(V^*XV|\mathcal{C})K) &= \mathbb{P}(V^*XVK) = \mathbb{Q}(XK) = \mathbb{Q}(\mathbb{Q}(X|\mathcal{C})K) = \\ &= \mathbb{P}(V^*V\mathbb{Q}(X|\mathcal{C})K) = \mathbb{P}(\mathbb{P}(V^*V\mathbb{Q}(X|\mathcal{C})K|\mathcal{C})) = \mathbb{P}(\mathbb{P}(V^*V|\mathcal{C})\mathbb{Q}(X|\mathcal{C})K). \end{aligned}$$

Here we have used the definition of the conditional expectation in the first step, the fact that V is in the commutant of \mathcal{C} in the second and fourth steps, and the tower

and module properties in the third, fifth, and sixth steps. The result now follows from the definition of the conditional expectation and the positivity of V^*V . \square

Unfortunately, this simple and elegant result will not be sufficient for our purposes. The problem is that all the operators involved must be bounded—the operator X which is being conditioned but also, more importantly, the change-of-state operator V . When we develop a quantum analog of the Kallianpur-Striebel formula, however, we will necessarily encounter an unbounded change-of-state operator V . It is not clear at this point that for unbounded V , the statement of the Lemma even makes sense; after all, how can we guarantee that V^*XV is well defined even for bounded X ? As is so often the case, domain problems rear their ugly head and we have to find a suitable way to deal with them.

The solution which I have chosen here is based on two ideas.

- (1) Rather than developing a full counterpart of Lemma 1.1.1, we concentrate on the special situation of Corollary 1.1.2. As this is precisely the situation encountered in filtering theory, this approach is particularly convenient.
- (2) In this context, we will extend the result of Lemma B.1.17 to a class of unbounded operators affiliated to the commutant \mathcal{C}' . When V and X are required to be of this type, this result absorbs the domain problems.

Throughout this chapter we will concentrate on the case where the initial system is finite dimensional, which admits a particularly transparent treatment. We now proceed to develop the necessary machinery. The generalization to the infinite-dimensional case, which is purely technical in nature, will appear in [BV07].

Let us presume that we are in the following situation.

- The Hilbert space underlying our problem is of the form $\mathbf{H} = \mathbf{h}_1 \otimes \mathbf{h}_2$, where \mathbf{h}_1 is finite dimensional, $\dim \mathbf{h}_1 = d$, and \mathbf{h}_2 is separable.
- We consider the quantum probability space $(\mathcal{A}, \mathbb{P})$ with the Von Neumann algebra $\mathcal{A} = \mathcal{B}(\mathbf{h}_1) \otimes \mathcal{B}(\mathbf{h}_2)$ and normal state $\mathbb{P} = \rho_1 \otimes \rho_2$.
- We consider a commutative subalgebra of the form $\mathcal{C} = I \otimes \mathcal{C}_2$.
- We wish to condition observables affiliated to $\tilde{\mathcal{C}} = \mathcal{B}(\mathbf{h}_1) \otimes \mathcal{C}_2 \subset \mathcal{C}'$.

What we are going to do is the following. As we are chiefly interested in the (noncommutative) Von Neumann algebra $\tilde{\mathcal{C}}$, we are going to build a nice $*$ -algebra $\mathcal{N}(\tilde{\mathcal{C}})$ of unbounded operators that are affiliated to $\tilde{\mathcal{C}}$ in a certain sense. We will then choose some $V \in \mathcal{N}(\tilde{\mathcal{C}})$ and define a new state $\mathbb{Q}(X) = \mathbb{P}(V^* \cdot X \cdot V)$ on the algebra $\tilde{\mathcal{C}}$ only. Finally, we will prove the Bayes formula for this state.

REMARK 5.1.7. Note that this situation is the direct quantum analog of the classical case considered in Corollary 1.1.2 (with the exception of the simplifying assumption that \mathbf{h}_1 be finite dimensional, which would correspond to Ω_1 being a finite set; this assumption will be dropped in the next section). Not surprisingly, the resulting Bayes formula will also have a form very similar to Corollary 1.1.2.

Let us get down to business. As $\dim \mathbf{h}_1 = d < \infty$, there is a natural isomorphism between $\mathcal{B}(\mathbf{h}_1)$ and the algebra $\mathcal{M}_d(\mathbb{C})$ of $d \times d$ complex matrices, obtained through an arbitrary choice of orthonormal basis in \mathbf{h}_1 , and the state ρ_1 can be represented in this basis by $\rho_1(X) = \text{Tr}(\varrho X)$ where ϱ is a $d \times d$ density matrix. Hence there is also a natural isomorphism between \mathcal{A} and the algebra $\mathcal{M}_d(\mathcal{B}(\mathbf{h}_2))$ of $d \times d$ matrices with $\mathcal{B}(\mathbf{h}_2)$ -valued entries, with the obvious addition, multiplication and

adjoint operations, and the state \mathbb{P} can be expressed as

$$\mathbb{P}(X) = \text{Tr} \left[\varrho \begin{pmatrix} \rho_2(X_{11}) & \cdots & \rho_2(X_{1d}) \\ \vdots & \ddots & \vdots \\ \rho_2(X_{d1}) & \cdots & \rho_2(X_{dd}) \end{pmatrix} \right], \quad X = \begin{pmatrix} X_{11} & \cdots & X_{1d} \\ \vdots & \ddots & \vdots \\ X_{d1} & \cdots & X_{dd} \end{pmatrix} \in \mathcal{A}.$$

From this point onwards we will make this identification without further comment, taking care to identify \mathbb{H} with the corresponding Hilbert space of d -dimensional \mathfrak{h}_2 -valued vectors. We will also denote by $\rho_2(X) \in \mathcal{M}_d(\mathbb{C})$ and $\rho_1(X) \in \mathcal{B}(\mathfrak{h}_2)$ the obvious “marginals” of the state \mathbb{P} , so that $\mathbb{P}(X) = \rho_1(\rho_2(X)) = \rho_2(\rho_1(X))$.

Let us now consider the algebra $\tilde{\mathcal{E}}$, which we have explicitly identified with $\mathcal{M}_d(\mathcal{E}_2)$. Recall from Lemma B.1.17 and the subsequent discussion that $\mathcal{N}(\mathcal{E}_2)$, the set of normal operators affiliated to \mathcal{E}_2 , forms a $*$ -algebra under the operations $\hat{+}$ and $\hat{\cdot}$. We now introduce the following definition.

DEFINITION 5.1.8. Define the $*$ -algebra $\mathcal{N}(\tilde{\mathcal{E}}) = \mathcal{M}_d(\mathcal{N}(\mathcal{E}_2))$, endowed with the obvious extensions of the addition $\hat{+}$, multiplication $\hat{\cdot}$ and adjoint.

It is now straightforward to establish the following: every $X \in \mathcal{N}(\tilde{\mathcal{E}})$ defines a closed, densely defined operator on \mathbb{H} , $X = X^*$ implies that X is self-adjoint, and the operations $\hat{+}$ and $\hat{\cdot}$ precisely coincide with the closures of the ordinary operator addition and multiplication in $\mathcal{N}(\tilde{\mathcal{E}})$. Hence $\mathcal{N}(\tilde{\mathcal{E}})$ is a true extension of $\mathcal{N}(\mathcal{E}_2)$ to a class of noncommuting unbounded operators. In essence we have done something fairly intuitive: the spectral theorem maps $\mathcal{N}(\tilde{\mathcal{E}})$ to a $*$ -algebra of operator-valued functions by diagonalizing only the commutative part of $\tilde{\mathcal{E}}$.

Our next order of business is to define a new state \mathbb{Q} .

LEMMA 5.1.9. *Let $V \in \mathcal{N}(\tilde{\mathcal{E}})$ be such that $\mathbb{P}(V^*V) = 1$. Define the functional $\mathbb{Q}(X) = \rho_2(\rho_1(V^* \hat{\cdot} X \hat{\cdot} V))$ for any $X \in \tilde{\mathcal{E}}$. Then \mathbb{Q} is a normal state on $\tilde{\mathcal{E}}$.*

PROOF. Linearity is easily verified. Now recall that V^*V is self-adjoint and nonnegative as V is closed [Rud73, Thm. 13.13], hence $V^*V = V^* \hat{\cdot} V$. Thus normalization follows immediately. Similarly, $V^* \hat{\cdot} X^* X \hat{\cdot} V = (X \hat{\cdot} V)^*(X \hat{\cdot} V)$ is self-adjoint and nonnegative. Hence positivity $\mathbb{Q}(X^*X) \geq 0$ follows. It remains to prove that \mathbb{Q} defines a normal state.

To this end, let us use the spectral theorem to find a measure space $(\Omega, \mathcal{G}, \mu)$, a σ -algebra $\mathcal{F} \subset \mathcal{G}$, a probability measure \mathbf{P} and a $*$ -isomorphism ι that maps (\mathcal{E}_2, ρ_2) to $(L^\infty(\Omega, \mathcal{F}, \mu), \mathbf{P})$. Let $x(\omega)$ be the $d \times d$ random matrix $[x]_{ij} = \iota(X_{ij})$, and let $f(\omega)$ be the $d \times d$ random matrix $[f]_{ij} = \iota(V_{ij})$. Then we can write

$$\rho_2(\rho_1(X)) = \int_{\Omega} \rho_1(f(\omega)^* x(\omega) f(\omega)) \mathbf{P}(d\omega).$$

The statement that a state ρ on an algebra \mathcal{A} is normal is equivalent to the statement that for every orthogonal family $\{P_\alpha\}$ of projections in \mathcal{A} we have $\rho(\sum_\alpha P_\alpha) = \sum_\alpha \rho(P_\alpha)$; see [KR97b, Thm. 7.1.12]. It suffices to restrict to countable families as \mathbb{H} is separable [BR87, sec. 2.5.1]. Hence let $\{X_n\}$ be a countable family of orthogonal projections in $\tilde{\mathcal{E}}$; then for μ -a.e. ω , the corresponding functions $x_n(\omega)$ must form a family of orthogonal projections in $\mathcal{B}(\mathfrak{h}_1)$ for all n . By normality of ρ_1 , we have $\rho_1(f(\omega)^* \sum_n x_n(\omega) f(\omega)) = \sum_n \rho_1(f(\omega)^* x_n(\omega) f(\omega))$ for μ -a.e. ω . The result now follows by monotone convergence. \square

We can now extend the Bayes formula to this unbounded case.

LEMMA 5.1.10. *Let \mathbb{Q} be as in Lemma 5.1.9, and assume that $V^*V > 0$. Then*

$$\mathbb{Q}(X|\mathcal{C}) = I \otimes \frac{\rho_1(V^* \hat{\cdot} X \hat{\cdot} V)}{\rho_1(V^*V)} \quad \text{for any } X \in \tilde{\mathcal{C}}.$$

PROOF. Let $K_2 \in \mathcal{C}_2$, $K = I \otimes K_2 \in \mathcal{C}$, and let $X \in \tilde{\mathcal{C}}$. Then

$$\begin{aligned} \rho_2(\rho_1(V^* \hat{\cdot} X \hat{\cdot} V) \hat{\cdot} K_2) &= \rho_2(\rho_1(V^* \hat{\cdot} X \hat{\cdot} K \hat{\cdot} V)) = \mathbb{Q}(XK) = \mathbb{Q}(\mathbb{Q}(X|\mathcal{C})K) = \\ &= \rho_2(\rho_1(V^* \hat{\cdot} \mathbb{Q}(X|\mathcal{C}) \hat{\cdot} K \hat{\cdot} V)) = \rho_2(\rho_1(V^* \hat{\cdot} V) \hat{\cdot} \mathbb{Q}(X|\mathcal{C})|_{\mathcal{C}_2} \hat{\cdot} K_2). \end{aligned}$$

(Recall that $\mathbb{Q}(X|\mathcal{C}) \in \mathcal{C}$, so $\mathbb{Q}(X|\mathcal{C})|_{\mathcal{C}_2}$ makes sense.) Note that this has to hold for any $K_2 \in \mathcal{C}_2$; hence it follows from a classical argument, e.g., [Wil91, sec. 9.5], that it must be the case that $\mathbb{Q}(X|\mathcal{C})|_{\mathcal{C}_2} = \rho_1(V^* \hat{\cdot} X \hat{\cdot} V)/\rho_1(V^* \hat{\cdot} V)$. \square

REMARK 5.1.11. It is evident that this result is a direct quantum counterpart of Corollary 1.1.2 in the classical case. Note that $\mathbb{P}(X|\mathcal{C}) = \rho_1(X)$ defines a version of the conditional expectation, as is easily verified.

5.2. Quantum filtering: A reference probability approach

5.2.1. The basic model. In this section, we are going to introduce the basic model that will be used throughout this chapter. Though this is by no means the most general model, this particular model is widely used and describes a wide variety of quantum systems. Our methods extend easily to more complicated situations; hence it is worth choosing a simple model here in order to keep the notation as simple and transparent as possible.

We work in the following setup. The basic Hilbert space is $\mathbf{H} = \mathfrak{h} \otimes \Gamma$, where \mathfrak{h} is assumed to be finite dimensional with $\dim \mathfrak{h} = d$, and Γ is the usual Fock space (see appendix B for details). Our basic quantum probability space is $(\mathcal{A}, \mathbb{P})$, where $\mathcal{A} = \mathcal{B} \otimes \mathcal{W}$ ($\mathcal{B} = \mathcal{B}(\mathfrak{h})$, $\mathcal{W} = \mathcal{B}(\Gamma)$) and $\mathbb{P} = \rho \otimes \varphi$. Here φ is the usual vacuum state, and ρ is an arbitrary state on \mathcal{B} (the initial state). Let us also introduce the fundamental filtration $\mathcal{A}_t = \mathcal{B} \otimes \mathcal{W}_t$.

To model time evolution in our system, we introduce a QSDE of Hudson-Parthasarathy type: i.e., we fix some $W, L, H \in \mathcal{A}_0$ where W is unitary and H is self-adjoint, and define the unitary process $\{U_t\}_{t \in \mathbb{R}_+}$ as the solution of

$$dU_t = \left\{ (W - I) d\Lambda_t + L dA_t^\dagger - L^*W dA_t - \frac{1}{2}L^*L dt - iH dt \right\} U_t, \quad U_0 = I.$$

Though U_t is formally defined only on the restricted exponential domain $\mathfrak{h} \otimes \mathbf{D}$, we will conveniently identify U_t with its closure (which is a unitary operator defined on all of \mathbf{H}) whenever this makes sense.

We can now introduce the basic “signal” and observations processes.

- (1) Let $X \in \mathcal{A}_0$ describe some physical property of the initial system at time $t = 0$. Then $j_t(X) = U_t^* X U_t$ is the corresponding observable at time t .
- (2) The observation process is given by $Y_t = U_t^* Z_t U_t$, with $Z_t = A_t + A_t^\dagger$.

Does this make sense? In particular, the process Y_t is a little dubious at this point—it is not obvious that it is commutative (i.e., that the spectral measures of Y_t and Y_s commute for all t and s). If this were not the case, then we could never observe this process in the laboratory. But fortunately we are in luck.

LEMMA 5.2.1. Define $\mathcal{Y}_t = \text{vN}\{Y_s : s \leq t\}$ and $\mathcal{Z}_t = \text{vN}\{Z_s : s \leq t\}$. Then $\mathcal{Y}_s = U_t^* \mathcal{Z}_s U_t$ for any $t \geq s$. In particular, this implies that \mathcal{Y}_t is commutative for any t , as $\{Z_t\}$ is a commutative process and hence \mathcal{Z}_t is commutative.

PROOF. Let E be an arbitrary projection operator in the range of the spectral measure of Z_s . We will show that $U_t^* E U_t = U_s^* E U_s$ for any $t \geq s$. As U_t and E are bounded, we can use the quantum Itô rules for this purpose without worrying about domains. Using the quantum Itô formula (Corollary B.2.17), we obtain

$$\begin{aligned} j_t(E) &= j_s(E) + \int_s^t j_\sigma(i[H, E] + L^* E L - \frac{1}{2}\{L^* L E + E L^* L\}) d\sigma \\ &\quad + \int_s^t j_\sigma(W^* [E, L]) dA_\sigma^\dagger + \int_s^t j_\sigma([L^*, E] W) dA_\sigma + \int_s^t j_\sigma(W^* E W - E) d\Lambda_\sigma. \end{aligned}$$

But by construction E commutes with H , L and W , hence we obtain $j_t(E) = j_s(E)$. It remains to note that as $Y_s = U_s^* Z_s U_s$, any spectral projection of Y_s can be written in the form $U_s^* E U_s$ where E is a spectral projection of Z_s . \square

Evidently $\{Y_t\}$ is a commutative process, and in particular the spectral theorem allows us to define a corresponding classical stochastic process $y_t = \iota(Y_t)$ on some probability space. The process y_t describes the statistics of what is measured in a laboratory experiment. Now suppose we have observed $\{Y_s : s \leq t\}$. Based on our observations, we would like to estimate some physical property of the initial system at time t (its position, dipole moment, etc.) With this motivation in mind, we introduce our basic filtering problem.

Define $\pi_t(X) = \mathbb{P}(j_t(X \otimes I) | \mathcal{Y}_t)$ for any $X \in \mathcal{B}$. The goal of the quantum filtering problem is to obtain an explicit expression for $\pi_t(X)$ in terms of the observations $\{Y_s : s \leq t\}$.

At this point we should start worrying again. Recall that the conditional expectation is only well defined if $j_t(X \otimes I) \in \mathcal{Y}_t'$; after all, if the observable of the system in which we are interested does not commute with what we have already observed, we could never design an experiment to test our predictions in the first place! Fortunately this is never a problem in our model: every system observable commutes with observations performed in the past.

LEMMA 5.2.2. $\text{vN}\{j_t(X \otimes I) : X \in \mathcal{B}\} \subset \mathcal{Y}_t'$.

PROOF. Note that $\text{vN}\{j_t(X \otimes I) : X \in \mathcal{B}\} = U_t^* \mathcal{A}_0 U_t$. The result follows immediately from Lemma 5.2.1, as \mathcal{A}_0 clearly commutes with \mathcal{Z}_t . \square

Lemmas 5.2.1 and 5.2.2 are usually called the self-nondemolition and the non-demolition properties, respectively. Evidently these are crucial for the physical interpretation of our model, and for the well-posedness of the filtering problem.

REMARK 5.2.3. It is important to note that it is almost always untrue that $U_s^* \mathcal{A}_0 U_s \subset \mathcal{Y}_t'$ for $t \geq s$. Hence *filtering* makes perfect sense in the quantum case, but *smoothing* does not. In some sense this enforces physical causality: once we have measured the observations up to time t , we can not subsequently travel back in time to observe some property of the system at an earlier time.

Before we move on to the solution of the filtering problem, let us say a word or two about the physical significance of our model. A common area of application for models of this type is in the field of quantum optics, which aims to describe the




FIGURE 5.1. Cartoon of the quantum filtering setup in quantum optics. An optical field, described by the field operators A_t , A_t^\dagger , interacts with a system, e.g., a cloud of atoms. The atom-field interaction is described by the unitary U_t . The field is subsequently detected, giving rise to the observation Y_t . Finally, the quantum filter (implemented, e.g., on a DSP) estimates atomic observables based on the field observations.

interactions of atomic-optical systems with the electromagnetic vacuum. A typical experimental setup is depicted in cartoon form in figure 5.1. A cloud of atoms—usually placed in an atom trap inside a vacuum chamber—is interrogated by probing it with a laser beam. The scattered light is detected through an optical detection setup; the observations $Y_t = U_t^* Z_t U_t$ which we have introduced here correspond to homodyne detection, while the commutative process $U_t^* \Lambda_t U_t$ describes what one would observe using a photon counter. Having detected the scattered light, we might want to estimate some property of the atomic cloud—its collective angular momentum, say. In order to find the optimal estimate (in the L^2 sense), this means we have to solve the filtering problem. The filter is also an important ingredient in the case where we would like to control some physical property of the atom through feedback. For an introduction to the physical modelling within our framework see [VSM05b]. The controlled case is developed in detail in [BV06, BV07].

5.2.2. The Kallianpur-Striebel formula. To solve the filtering problem, we would like to apply Lemma 5.1.10. In the classical case, this was particularly convenient as we had defined our probability space with the explicit product structure $D(\mathbb{R}_+; \mathbb{S}) \times C(\mathbb{R}_+; \mathbb{R}^p)$, so that the signal lives on the first component and the observations live on the second component. In the quantum case we also have a convenient product structure $\mathcal{B} \otimes \mathcal{W}$, but the product structure is not preserved by the unitary U_t . Of course, we could define a “rotated” tensor product: there is a natural isomorphism $U_t^*(\mathcal{B} \otimes \mathcal{W})U_t \simeq U_t^* \mathcal{B} U_t \otimes U_t^* \mathcal{W} U_t$. We have a much more convenient trick up our sleeve, however: we will apply Lemma 5.1.5 to rotate the conditional expectation back onto our existing product structure, and then we will apply Lemma 5.1.10 in this rotated basis. In essence, the idea is that even though our model is not initially defined with a convenient product structure as in the classical case, we are free to work in a different basis in which the model does have this form. Let us work out the details of this procedure.

Recall that we are interested in conditional expectations of the form $\mathbb{P}(j_t(X) | \mathcal{Z}_t)$ with $X \in \mathcal{A}_0$. Let us fix some time t , and define a new state by $\mathbb{Q}^t(X) = \mathbb{P}(U_t^* X U_t)$. Then we have, using Lemmas 5.2.1 and 5.1.5,

$$\mathbb{P}(j_t(X) | \mathcal{Z}_t) = \mathbb{P}(U_t^* X U_t | U_t^* \mathcal{Z}_t U_t) = U_t^* \mathbb{Q}^t(X | \mathcal{Z}_t) U_t.$$

If only it were the case that $U_t \in \mathcal{B} \otimes \mathcal{L}_t$, we would now immediately be able to apply the Bayes formula of Lemma 5.1.10 to $\mathbb{Q}^t(X|\mathcal{L}_t)$. This is most certainly not the case, though, in any nontrivial situation: after all, this would imply that $U_t^* X U_t = X$ for any $X \in \mathcal{A}_0$, which means that the initial system does not interact with the field. If we want to apply the Bayes formula, we need to find a $V_t \in \mathcal{N}(\mathcal{B} \otimes \mathcal{L}_t)$ such that $\mathbb{P}(U_t^* X U_t) = \mathbb{P}(V_t^* \hat{X} V_t)$ for every $X \in \mathcal{B} \otimes \mathcal{L}_t$. This is precisely what we are going to do. Note that in the classical case, this step is taken care of by Girsanov's theorem. Unfortunately, we do not have such a general result at our disposal in the noncommutative case.

As a first step in this direction, let us prove the following result. To my knowledge, this idea first appears in a paper by A. S. Holevo [Hol91]. Note that we are here in the realm of the Hudson-Parthasarathy theory, so the objects in this result live on the restricted domain $\mathfrak{h} \otimes \mathbb{D}$. In particular, we should not draw premature conclusions about what the closures of these operators look like.

LEMMA 5.2.4. *Let $C, D, F, G, \tilde{C}, \tilde{F} \in \mathcal{A}_0$, and let V_t and \tilde{V}_t be the solutions of*

$$\begin{aligned} dV_t &= \{C d\Lambda_t + D dA_t^\dagger + F dA_t + G dt\} V_t, \\ d\tilde{V}_t &= \{\tilde{C} d\Lambda_t + D dA_t^\dagger + \tilde{F} dA_t + G dt\} \tilde{V}_t, \quad V_0 = \tilde{V}_0. \end{aligned}$$

Then $V_t v \otimes \Phi = \tilde{V}_t v \otimes \Phi$ for any $v \in \mathfrak{h}$, where $\Phi \in \Gamma$ is the vacuum vector.

PROOF. By Theorem B.2.19, V_t and \tilde{V}_t are both uniquely defined admissible adapted processes. Consider the quantity

$$\|(V_t - \tilde{V}_t) v \otimes \Phi\|^2 = \langle (V_t - \tilde{V}_t) v \otimes \Phi, (V_t - \tilde{V}_t) v \otimes \Phi \rangle.$$

Using the quantum Itô rule we obtain

$$\begin{aligned} \|(V_t - \tilde{V}_t) v \otimes \Phi\|^2 &= \int_0^t \langle (V_s - \tilde{V}_s) v \otimes \Phi, (G + G^*)(V_s - \tilde{V}_s) v \otimes \Phi \rangle ds \\ &\quad + \int_0^t \langle D(V_s - \tilde{V}_s) v \otimes \Phi, D(V_s - \tilde{V}_s) v \otimes \Phi \rangle ds. \end{aligned}$$

Note that the last integrand is trivially estimated by

$$\|D(V_s - \tilde{V}_s) v \otimes \Phi\|^2 \leq \|D\|^2 \|(V_s - \tilde{V}_s) v \otimes \Phi\|^2.$$

Similarly, the first integrand can be estimated by

$$\langle (V_s - \tilde{V}_s) v \otimes \Phi, (G + G^*)(V_s - \tilde{V}_s) v \otimes \Phi \rangle \leq \|G + G^*\| \|(V_s - \tilde{V}_s) v \otimes \Phi\|^2.$$

Setting $C = \|D\|^2 + \|G + G^*\|$, we obtain

$$\|(V_t - \tilde{V}_t) v \otimes \Phi\|^2 \leq C \int_0^t \|(V_s - \tilde{V}_s) v \otimes \Phi\|^2 ds,$$

But then by Gronwall's lemma $\|(V_t - \tilde{V}_t) v \otimes \Phi\| = 0$, and the Lemma is proved. \square

Let us take a moment to show where we are going with this. Without loss of generality, we can restrict to the case where the initial state is defined as $\rho(X) = \langle v, X v \rangle$ for some $v \in \mathfrak{h}$. Any state ρ on \mathcal{B} can be written as a (finite) convex combination of such vector states, so that once we have established our results for all vector states we can trivially extend to arbitrary ρ . Now consider the QSDE

$$dV_t = \left\{ L(dA_t^\dagger + dA_t) - \frac{1}{2} L^* L dt - iH dt \right\} V_t, \quad V_0 = I.$$

By the previous Lemma, we see that $V_t v \otimes \Phi = U_t v \otimes \Phi$. Hence we can write

$$\mathbb{Q}^t(X) = \mathbb{P}(U_t^* X U_t) = \langle U_t v \otimes \Phi, X U_t v \otimes \Phi \rangle = \langle V_t v \otimes \Phi, X V_t v \otimes \Phi \rangle.$$

It is now tempting to bring V_t to the other side of the inner product, i.e., to write “ $\mathbb{Q}^t(X) = \mathbb{P}(V_t^* X V_t)$ ”; moreover, the QSDE for V_t is only driven by the commutative noise Z_t and has coefficients in \mathcal{A}_0 , so that it seems almost obvious that V_t should be affiliated to $\mathcal{B} \otimes \mathcal{Z}_t$ in some sense. These are, of course, precisely the things we need in order to apply the Bayes formula. We are not quite there yet, however, as these naive manipulations are not, in fact, well justified. Our goal is to find a suitable $V_t \in \mathcal{N}(\mathcal{B} \otimes \mathcal{Z}_t)$ to replace the restricted operator V_t defined above. If we can find such an operator, our formal manipulations can be made rigorous.

Because the QSDE for V_t is essentially commutative (it is driven by a single commutative noise Z_t), there is an obvious candidate for its replacement. By applying the spectral theorem to \mathcal{Z}_T , we can represent Z_t ($t \in [0, T]$) as a classical stochastic process $z_t = \iota(Z_t)$ on some probability space, with the corresponding measure \mathbf{P} induced by φ under which z_t is a Wiener process. We can now construct classical Itô integrals with respect to z_t , and pull these back to self-adjoint operators affiliated to \mathcal{Z}_T . The question is whether such integrals are extensions of the operators obtained using the Hudson-Parthasarathy integral when we restrict the domain of the integrand and integrator to $\mathfrak{h} \otimes \mathfrak{D}$, and similarly whether the solutions of the QSDE and the corresponding classical Itô SDE have this property. Note that for our purposes, it is sufficient to establish that these operators act in the same way on some domain that contains the vacuum, as we do not particularly care about the rest of $\mathfrak{h} \otimes \mathfrak{D}$. The following Lemma gives the desired result.

LEMMA 5.2.5. *Let $(\Omega, \mathcal{G}, \mu)$ be the probability space obtained from the spectral theorem applied to \mathcal{Z}_T , and let $z_t = \iota(Z_t)$. Let $L, H \in \mathcal{B}$ be such that $H = H \otimes I$, $L = L \otimes I$, and define V_t as the solution of the matrix-valued Itô SDE*

$$dV_t = \left\{ L dz_t - \frac{1}{2} L^* L dt - iH dt \right\} V_t, \quad V_0 = I.$$

Then V_t coincides with $\iota^{-1}(V_t) \in \mathcal{N}(\mathcal{B} \otimes \mathcal{Z}_t)$ at least on a dense subdomain $\mathfrak{h} \otimes \mathfrak{D}' \subset \mathfrak{h} \otimes \mathfrak{D}$ which contains all vectors of the form $v \otimes \Phi$.

PROOF. First, let us recall that there is a unique solution V_t (up to a.s. equivalence), which is adapted and square integrable. This follows from standard results on Lipschitz stochastic differential equations. Next, recall that we have made a natural identification of the $*$ -algebra $\mathcal{N}(\mathcal{B} \otimes \mathcal{Z}_t)$ with the $*$ -algebra of $\mathcal{N}(\mathcal{Z}_t)$ -valued matrices. Similarly, we can naturally identify all operators on some domain $\mathfrak{h} \otimes \mathfrak{D}'$ with the set of matrices whose matrix elements take values in the set of operators on \mathfrak{D}' . In [HP84, sec. 5], it is shown¹ that there exists a dense domain $\mathfrak{D}' \subset \mathfrak{D}$ (with $\Phi \in \mathfrak{D}'$) on which any Itô integral of a square integrable adapted process, when considered as a multiplication operator, coincides with the corresponding Hudson-Parthasarathy integral. Applying this result to every matrix element of

¹ Hudson and Parthasarathy do not use the general spectral theorem, but explicitly diagonalize their operators by using the natural isomorphism between the Fock space and the Wiener space. This particular representation is evidently equivalent to any other representation constructed using the spectral theorem. Alternatively, their proofs are easily repeated without using the Wiener-Fock isomorphism. The corresponding domain \mathfrak{D}' is then constructed by generating this domain from the vacuum using the convenient set of diagonal Weyl operators.

V_t , this implies that $\iota^{-1}(V_t)|_{\mathfrak{h} \otimes \mathcal{D}'}$ satisfies the same QSDE as does V_t . The result now follows from the uniqueness of the solution of linear QSDE. \square

As our previously defined restricted operator V_t has now outlived its usefulness, we feel justified in replacing it by its extension $\iota^{-1}(V_t)$. So without further apology:

From this point onward we identify V_t with $\iota^{-1}(V_t)$.

This is little more than a domain extension. Note that it is not entirely obvious whether $\iota^{-1}(V_t)$ coincides with the closure of our old V_t ; we also do not care, as we are only looking to reproduce its action on the vacuum.

Armed with our promising operator V_t , we can now start making progress.

LEMMA 5.2.6. $\mathbb{Q}^t(X) = \mathbb{P}(V_t^* \hat{\cdot} X \hat{\cdot} V_t)$ for all $X \in \mathcal{B} \otimes \mathcal{Z}_t$.

PROOF. Let us first prove that $\mathbb{P}(V_t^* V_t) = 1$. By the classical Itô rule,

$$d(V_t^* V_t) = V_t^*(L + L^*)V_t dz_t, \quad V_0^* V_0 = I,$$

so that $\mathbb{P}(V_t^* V_t) = \rho(\mathbf{E}_{\mathbf{P}}(V_t^* V_t)) = \rho(I) = 1$. Hence by Lemma 5.1.9, the functional $\mathbb{P}(V_t^* \hat{\cdot} X \hat{\cdot} V_t)$ defines a normal state on $\mathcal{B} \otimes \mathcal{Z}_t$. We would like to show that this state coincides with \mathbb{Q}^t ; we can assume without loss of generality that $\rho(X) = \langle v, Xv \rangle$ is a vector state. To this end, denote by $\psi \in L^2(\Omega, \mathcal{G}, \mu; \mathfrak{h})$ the vector obtained by applying the spectral theorem to $v \otimes \Phi$, by $\theta \in L^2(\Omega, \mathcal{G}, \mu; \mathfrak{h})$ the vector obtained by applying the spectral theorem to $U_t v \otimes \Phi$, and write $X = \iota(X)$. Then

$$\begin{aligned} \mathbb{P}(V_t^* \hat{\cdot} X \hat{\cdot} V_t) &= \int_{\Omega} \langle \psi(\omega), V_t(\omega)^* X(\omega) V_t(\omega) \psi(\omega) \rangle \mu(d\omega) \\ &= \int_{\Omega} \langle V_t(\omega) \psi(\omega), X(\omega) V_t(\omega) \psi(\omega) \rangle \mu(d\omega) \\ &= \int_{\Omega} \langle \theta(\omega), X(\omega) \theta(\omega) \rangle \mu(d\omega) = \langle U_t v \otimes \Phi, X U_t v \otimes \Phi \rangle = \mathbb{Q}^t(X). \end{aligned}$$

This completes the proof. \square

LEMMA 5.2.7. $V^* V > 0$.

PROOF. It is sufficient to establish that V_t is invertible. But this is standard (see, e.g., Lemma 3.2.3), so the result follows directly. \square

We can now finally apply the Bayes formula. This results in the following, which is a noncommutative counterpart of the Kallianpur-Striebel formula.

THEOREM 5.2.8. $\pi_t(X) = U_t^* \mathbb{Q}^t(X \otimes I | \mathcal{Z}_t) U_t$, where

$$\mathbb{Q}^t(X \otimes I | \mathcal{Z}_t) = I \otimes \frac{\rho(V_t^* \hat{\cdot} (X \otimes I) \hat{\cdot} V_t)}{\rho(V_t^* V_t)} \quad \text{for all } X \in \mathcal{B}.$$

PROOF. This follows from the preceding results and Lemma 5.1.10. \square

COROLLARY 5.2.9. Let $L^\infty(\Omega, \mathcal{F}, \mu)$ be the $*$ -algebra of random variables obtained by applying the spectral theorem to \mathcal{Y}_T , let \mathbf{P} be the measure on \mathcal{F} induced by \mathbb{P} , and define the classical observations process $y_t = \iota(Y_t)$ for $t \in [0, T]$. Furthermore, let \bar{V}_t be the solution of the matrix-valued Itô SDE

$$d\bar{V}_t = \left\{ L dy_t - \frac{1}{2} L^* L dt - iH dt \right\} \bar{V}_t, \quad \bar{V}_0 = I.$$

Define for any $X \in \mathcal{B}$ the following quantities:

$$\bar{\sigma}_t(X) = \rho(\bar{V}_t^* X \bar{V}_t), \quad \bar{\pi}_t(X) = \frac{\bar{\sigma}_t(X)}{\bar{\sigma}_t(I)}.$$

Then $\bar{\pi}_t(X) = \iota(\pi_t(X))$. Moreover, under the probability measure \mathbf{Q} defined by $d\mathbf{P} = \bar{\sigma}_T(I) d\mathbf{Q}$, the observations process $\{y_t\}_{t \in [0, T]}$ is a Wiener process.

PROOF. We begin by noting that the quantum probability spaces $(\mathcal{Y}_T, \mathbb{P})$ and $(\mathcal{Z}_T, \mathbb{Q}^T)$ are equivalent. To see this, define $\mathbf{i} : \mathcal{Z}_T \rightarrow \mathcal{Y}_T$, $\mathbf{i}(X) = U_T^* X U_T$. Clearly \mathbf{i} defines a *-isomorphism between \mathcal{Z}_T and \mathcal{Y}_T ; hence in order to extend this to an equivalence between quantum probability spaces, we only need to pull back the state \mathbb{P} on \mathcal{Y}_T to \mathcal{Z}_T . But the corresponding state is precisely $\mathbb{Q}^T(X) = \mathbb{P}(\mathbf{i}(X))$.

Let $\iota_Z : \mathcal{Z}_T \rightarrow L^\infty(\Omega_Z, \mathcal{F}_Z, \mu_Z)$ and $\iota_Y : \mathcal{Y}_T \rightarrow L^\infty(\Omega_Y, \mathcal{F}_Y, \mu_Y)$ be the *-isomorphisms obtained by appropriate application of the spectral theorem. We will define $\iota : \mathcal{Y}_T \rightarrow L^\infty(\Omega_Z, \mathcal{F}_Z, \mu_Z)$ by $\iota(X) = \iota_Z(\mathbf{i}^{-1}(X))$. Similarly, let us define the following probability measures on \mathcal{F}_Z : $\mathbf{P}(\iota(X)) = \mathbb{P}(X)$, $\mathbf{Q}(\iota_Z(X)) = \mathbb{P}(X)$. We are clearly free to use this ι rather than ι_Y as described in the statement of the Corollary, as the map \mathbf{i} makes everything completely isomorphic.

By construction $y_t = \iota(Y_t)$ has the law of the physical observations process under \mathbf{P} . But note that $y_t = \iota(Y_t) = \iota_Z(Z_t) = z_t$, so that evidently y_t is a Wiener process under \mathbf{Q} . Moreover, $\iota(\pi_t(X)) = \iota_Z(\mathbb{Q}^t(X \otimes I | \mathcal{Z}_t))$, as is easily established using Lemma 5.2.1, so that we obtain $\iota(\pi_t(X)) = \bar{\pi}_t(X)$ from Theorem 5.2.8. It remains to note that for any functional X of $\{y_t\}$

$$\mathbf{E}_{\mathbf{P}}(X) = \mathbb{Q}^T(\iota_Z^{-1}(X)) = \mathbb{P}(V_T^* \hat{\iota}_Z^{-1}(X) \hat{V}_T) = \mathbf{E}_{\mathbf{Q}}(\rho(\bar{V}_T^* \bar{V}_T) X).$$

Hence $d\mathbf{P}/d\mathbf{Q} = \rho(\bar{V}_T^* \bar{V}_T) = \bar{\sigma}_T(I)$, and the proof is complete. \square

PROPOSITION 5.2.10. *The innovations process \bar{z}_t , defined by*

$$\bar{z}_t = y_t - \int_0^t \bar{\pi}_s(L + L^*) ds,$$

is a Wiener process under the probability measure \mathbf{P} .

PROOF. Using the Itô rules, it is straightforward to establish that

$$d\bar{\sigma}_t(I) = \bar{\sigma}_t(L + L^*) dy_t = \bar{\pi}_t(L + L^*) \bar{\sigma}_t(I) dy_t.$$

Hence $\sigma_T(I)$ defines a Girsanov transformation, and the result follows from the Girsanov theorem and the fact that $\{y_t\}$ is a Wiener process under \mathbf{Q} . \square

5.2.3. The quantum filtering equations. Beside being the key result in the reference probability method, the quantum Kallianpur-Striebel formula has a convenient side effect: it reduces the quantum filtering problem to the realm of classical stochastic processes. From this point onward, we will work exclusively within the framework of Corollary 5.2.9. This means that we will use only the classical Itô calculus, and the resulting filtering equations will be classical Itô stochastic differential equations which are driven by the observations process y_t . This is entirely within the spirit of filtering theory; indeed, the filter should, by construction, be a functional of the observations only! Hence these classical equations provide precisely what should be implemented as a signal processing device in the applications of the theory, e.g., in a laboratory setting (cf. figure 5.1).

Let us begin by obtaining the quantum filtering equations, both in unnormalized form (the quantum counterpart of the Zakai equation) and in normalized form

(the quantum counterpart of the Kushner-Stratonovich equation). Given the results of the previous section, this is a straightforward exercise.

PROPOSITION 5.2.11. *The unnormalized filter $\bar{\sigma}_t(X)$ satisfies*

$$d\bar{\sigma}_t(X) = \bar{\sigma}_t(\mathcal{L}(X)) dt + \bar{\sigma}_t(L^*X + XL) dy_t, \quad \bar{\sigma}_0(X) = \rho(X),$$

while the normalized filter $\bar{\pi}_t(X)$ satisfies the equation

$$d\bar{\pi}_t(X) = \bar{\pi}_t(\mathcal{L}(X)) dt + (\bar{\pi}_t(L^*X + XL) - \bar{\pi}_t(L + L^*)\bar{\pi}_t(X)) d\bar{z}_t, \quad \bar{\pi}_0(X) = \rho(X).$$

Here we have used the Lindblad generator $\mathcal{L}(X)$, which is defined as

$$\mathcal{L}(X) = i[H, X] + L^*XL - \frac{1}{2}L^*LX - \frac{1}{2}XL^*L.$$

PROOF. Let $X \in \mathcal{B}$. Using the Itô rules, we find that

$$\bar{V}_t^*X\bar{V}_t = X + \int_0^t \bar{V}_s^*\mathcal{L}(X)\bar{V}_s ds + \int_0^t \bar{V}_s^*(L^*X + XL)\bar{V}_s dy_s.$$

As $\rho(X) = \text{Tr}[\varrho X]$ is just a finite linear combination of matrix elements, we can trivially pull ρ into the integrals. This gives the unnormalized filter. The normalized filter is easily found by applying the Itô rule to $\bar{\pi}_t(X) = \bar{\sigma}_t(X)/\bar{\sigma}_t(I)$. \square

As in the classical case, the density form of these equations is particularly convenient, as this gives rise to closed-form recursive equations.

COROLLARY 5.2.12. *Let ϱ_t be the (unique) random density matrix that satisfies $\bar{\pi}_t(X) = \text{Tr}[\varrho_t X]$ for all $X \in \mathcal{B}$. Then ϱ_t satisfies the equation*

$$d\varrho_t = \mathcal{L}^*(\varrho_t) dt + \{L\varrho_t + \varrho_t L^* - \text{Tr}[(L + L^*)\varrho_t]\varrho_t\} (dy_t - \text{Tr}[(L + L^*)\varrho_t] dt),$$

with the initial condition $\varrho_0 = \varrho$. Similarly, define the random nonnegative self-adjoint matrix ς_t by $\bar{\sigma}_t(X) = \text{Tr}[\varsigma_t X]$. Then ς_t satisfies the equation

$$d\varsigma_t = \mathcal{L}^*(\varsigma_t) dt + \{L\varsigma_t + \varsigma_t L^*\} dy_t, \quad \varsigma_0 = \varrho.$$

Here we have used the adjoint Lindblad generator $\mathcal{L}^*(X)$, which is defined as

$$\mathcal{L}^*(\varrho) = -i[H, \varrho] + L\varrho L^* - \frac{1}{2}L^*L\varrho - \frac{1}{2}\varrho L^*L.$$

To be useful, these filtering equations should have unique solutions.

LEMMA 5.2.13. *The equations for ς_t and ϱ_t have unique strong solutions.*

PROOF. The fact that the equation for ς_t has a unique solution is standard, as this is a Lipschitz stochastic differential equation. Moreover, the equation for ϱ_t is locally Lipschitz and hence has a unique solution up to some accessible explosion time ζ . But it is easily verified using the Itô rules that $\varsigma_t/\text{Tr}[\varsigma_t]$ satisfies the equation for ϱ_t at every time $t \in [0, T]$, and moreover $\varsigma_t/\text{Tr}[\varsigma_t]$ evolves in a compact set. Hence there can be no accessible explosion, and the claim is established. \square

Finally, let us consider the equation for ϱ_t as being driven by the innovations \bar{z}_t rather than the observations y_t . This gives some additional insight.

PROPOSITION 5.2.14. *The following Itô stochastic differential equation has a unique strong solution which coincides with ϱ_t as defined above:*

$$d\varrho_t = \mathcal{L}^*(\varrho_t) dt + \{L\varrho_t + \varrho_t L^* - \text{Tr}[(L + L^*)\varrho_t]\varrho_t\} d\bar{z}_t, \quad \varrho_0 = \varrho.$$

In particular, this implies that ϱ_t is a Feller-Markov process under \mathbf{P} and that the innovations conjecture $\sigma\{y_t : t \in [0, T]\} = \sigma\{\bar{z}_t : t \in [0, T]\}$ holds.

PROOF. As the coefficients of the equation are locally Lipschitz continuous, there is a unique solution up to some explosion time ζ . But as ϱ_t defined previously satisfies this equation, ζ must be infinite and the solutions must coincide. As \bar{z}_t is a Wiener process under \mathbf{P} , the Feller-Markov property follows from standard results. To prove the innovations conjecture, note that it is trivially the case that $\sigma\{y_t : t \in [0, T]\} \supset \sigma\{\bar{z}_t : t \in [0, T]\}$. But $\sigma\{y_t : t \in [0, T]\} \subset \sigma\{\bar{z}_t : t \in [0, T]\}$ follows from the current result, as we can reconstruct y_t from \bar{z}_t in a measurable fashion by solving the above SDE (see, e.g., [LS01, p. 276]). \square

5.2.4. Imperfect detection. The reference probability method developed above can be adapted to a large number of quantum filtering scenarios. As mentioned before, we restrict ourselves in this chapter to the model introduced above for sake of notational simplicity and of transparent treatment. There is a minor extension of this model, however, that is readily introduced and that is widely used. Let us briefly sketch how to incorporate this addition.

In the model which we have been investigating there is no external corrupting noise—all the noise in the observations originates from the intrinsic quantum fluctuations in the (electromagnetic) vacuum. This can not be circumvented; in particular, there is no quantum analog of the “perfect detection case” in classical systems theory. This highlights the reason that filtering theory is particularly important in quantum systems theory. On the other hand, there is no particular reason why there should not be any external corrupting noise in addition to the intrinsic quantum noise in the system. In fact, any realistic experiment or device is subject to technical noise, of environmental or of electronic origin. Hence it is useful to reconsider the above filtering problem, using the same underlying model but adding some independent corrupting noise to the observations. The goal of this section is to obtain the corresponding filtering equation.

To model the corrupting noise we simply extend our original Hilbert space by tensoring on another Fock space: i.e., we consider the Hilbert space $\mathfrak{h} \otimes \Gamma \otimes \Gamma$, where we denote the fundamental noises on the first copy of the Fock space by A_t , A_t^\dagger , and Λ_t , and we denote the fundamental noises on the second copy of the Fock space by B_t , B_t^\dagger , and Ξ_t (we will not use the latter). We endow the second copy of the Fock space with the vacuum, i.e., our state is $\mathbb{P} = \rho \otimes \varphi \otimes \varphi$. The interaction unitary U_t is still defined in the same way as above; in particular, it is of the form $U_t \otimes I$ on $(\mathfrak{h} \otimes \Gamma) \otimes \Gamma$, and the “signal process” is given by $j_t(X)$ for $X \in \mathscr{A}_0$. The only difference with the previous model is the observations process: we now define $Y_t = U_t^* Z_t U_t + \varepsilon(B_t + B_t^\dagger)$. Recall that under \mathbb{P} , the process $B_t + B_t^\dagger$ is a Wiener process independent of $U_t^* Z_t U_t$. Hence our new observations implement precisely the idea described above, where ε is the strength of the external corrupting noise.

To obtain the filtering equations for this case, there is no need to start from scratch. Let us define $\tilde{Y}_t = U_t^* Z_t U_t$, i.e., these are the observations which we have been considering previously, and define the corresponding filtration $\tilde{\mathscr{Y}}_t$. We can now simply invoke the tower property of the conditional expectation: $\mathbb{P}(j_t(X)|\mathscr{Y}_t) = \mathbb{P}(\mathbb{P}(j_t(X)|\mathfrak{vN}\{\mathscr{Y}_t, \tilde{\mathscr{Y}}_t\})|\mathscr{Y}_t)$. But by the independence of $j_t(X)$, $\tilde{\mathscr{Y}}_t$ and $B_{s \leq t}$, this reduces to $\mathbb{P}(j_t(X)|\mathscr{Y}_t) = \mathbb{P}(\mathbb{P}(j_t(X)|\tilde{\mathscr{Y}}_t)|\mathscr{Y}_t)$. The quantity $\mathbb{P}(j_t(X)|\tilde{\mathscr{Y}}_t)$ obeys the filtering equation which we have obtained previously; so in order to obtain our new filtering equation, it suffices to perform some additional classical conditioning.

Define $\tilde{\pi}_t(X) = \iota(\mathbb{P}(j_t(X)|\tilde{\mathcal{Y}}_t))$. We have obtained the equation for $\tilde{\pi}_t(X)$ in the previous section; unfortunately, this equation is nonlinear, so it is difficult to apply the tower property. Instead, we will apply the classical Bayes formula:

$$\tilde{\pi}_t(X) = \mathbf{E}_{\mathbf{P}}(\tilde{\pi}_t(X)|\mathcal{Y}_t) = \frac{\mathbf{E}_{\mathbf{Q}}(\tilde{\pi}_t(X) \tilde{\sigma}_t(I)|\mathcal{Y}_t)}{\mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(I)|\mathcal{Y}_t)} = \frac{\mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(X)|\mathcal{Y}_t)}{\mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(I)|\mathcal{Y}_t)}.$$

Here $\mathcal{Y}_t = \sigma\{y_s : s \leq t\}$, and \mathbf{Q} is the measure defined previously under which $\tilde{y}_t = \iota(\tilde{Y}_t)$ is a Wiener process. Now recall from the previous section that

$$d\tilde{\sigma}_t(X) = \tilde{\sigma}_t(\mathcal{L}(X)) dt + \tilde{\sigma}_t(L^*X + XL) d\tilde{y}_t, \quad \tilde{\sigma}_0(X) = \rho(X).$$

Denote by $b_t = \iota(B_t + B_t^\dagger)$ the corrupting Wiener process, which is independent of \tilde{y}_t and where $y_t = \tilde{y}_t + \varepsilon b_t$. To complete the argument, notice that if we define

$$O = \frac{1}{\sqrt{1+\varepsilon^2}} \begin{pmatrix} 1 & \varepsilon \\ \varepsilon & -1 \end{pmatrix}, \quad O \begin{pmatrix} \tilde{y}_t \\ b_t \end{pmatrix} = \frac{1}{\sqrt{1+\varepsilon^2}} \begin{pmatrix} y_t \\ y_t^\perp \end{pmatrix},$$

then O is an orthogonal matrix and hence y_t and y_t^\perp are independent Wiener processes under \mathbf{Q} (with quadratic variation $(1+\varepsilon^2)t$). Let us write suggestively

$$d\tilde{\sigma}_t(X) = \tilde{\sigma}_t(\mathcal{L}(X)) dt + \frac{1}{1+\varepsilon^2} \tilde{\sigma}_t(L^*X + XL) dy_t + \frac{\varepsilon}{1+\varepsilon^2} \tilde{\sigma}_t(L^*X + XL) dy_t^\perp.$$

We can now take the conditional expectation of this expression, using standard results to exchange the order of integration and conditioning [LS01, Ch. 5] (note that the integrands have moments of all orders; this is easily established, as \tilde{V}_t is the solution of a linear SDE). We thus obtain

$$d\mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(X)|\mathcal{Y}_t) = \mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(\mathcal{L}(X))|\mathcal{Y}_t) dt + \frac{1}{1+\varepsilon^2} \mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(L^*X + XL)|\mathcal{Y}_t) dy_t.$$

It is not difficult to see, in fact, that $\mathbf{E}_{\mathbf{Q}}(\tilde{\sigma}_t(X)|\mathcal{Y}_t) = \bar{\sigma}_t(X)$, so we have obtained the linear filtering equation for the model considered in this section. Normalization is straightforward, and we can summarize with the following statement.

PROPOSITION 5.2.15. *Let $\eta = (1+\varepsilon^2)^{-1}$, and define the normalized observations $\bar{y}_t = \sqrt{\eta} y_t$. Then the linear filtering equation for the model considered in this section is given by the Itô stochastic differential equation*

$$d\bar{\sigma}_t(X) = \bar{\sigma}_t(\mathcal{L}(X)) dt + \sqrt{\eta} \bar{\sigma}_t(L^*X + XL) d\bar{y}_t, \quad \bar{\sigma}_0(X) = \rho(X),$$

while the normalized filter $\bar{\pi}_t(X)$ satisfies the equation

$$d\bar{\pi}_t(X) = \bar{\pi}_t(\mathcal{L}(X)) dt + \sqrt{\eta} (\bar{\pi}_t(L^*X + XL) - \bar{\pi}_t(L+L^*) \bar{\pi}_t(X)) d\bar{z}_t, \quad \bar{\pi}_0(X) = \rho(X).$$

Here the innovations process \bar{z}_t is given by

$$\bar{z}_t = \bar{y}_t - \sqrt{\eta} \int_0^t \bar{\pi}_s(L+L^*) ds,$$

and \bar{z}_t defines a standard Wiener process under the measure \mathbf{P} .

The density form of these equations is trivially established, and the various existence and uniqueness results follow as in the previous section.

5.3. A filter stability result

We now turn to the second topic of this chapter: the issue of filter stability for quantum filters. The problem here is the same as in the classical case. The optimal filter, which propagates the conditional expectation of system observables and hence is an optimal estimator in the L^2 sense, requires us to initialize the filtering equation with the true initial state ρ . We would like to show that if a different initial state ρ' was used instead to initialize the filter (so that the filtered estimate is suboptimal), this misspecification is forgotten and the filter becomes optimal as $t \rightarrow \infty$. Nothing is known about this question in the literature to date beside numerical evidence. Our goal here is to obtain a first result in this direction.

5.3.1. Incorrect initialization and randomization of the initial state.

Let us reflect for a moment on the methods which we have introduced to study classical filter stability. All of these methods relied rather heavily on a study of a related smoothing problem. Even if we do not explicitly engage in smoothing, we would still need to condition on the signal process at an earlier time in order to even make sense of, e.g., the second representation described in section 2.3.1. In the quantum setting this is meaningless, as such an expression would violate the nondemolition condition. If we wish to make use of some of the logic developed in the classical case, we thus have to find a way to circumvent this problem.²

Let us consider in a little more detail the problem we are facing. We would like to be able to relate the filters initialized with the initial states ρ and ρ' under the same measure. To do this, we need two things: first, we need to be able to obtain a change-of-measure operator V such that $\rho'(X) = \rho(V^*XV)$; and second, we must make sure that V is in the commutant of the observation algebra. The latter will never be the case within the model which we have introduced.

We will use a remarkably simple trick to avoid these problems. Let us suppose that there is a third state $\tilde{\rho}$, so that we can write ρ' as a convex combination $\rho' = \lambda\rho + (1 - \lambda)\tilde{\rho}$ with some $\lambda > 0$. Then we could generate ρ' as follows: in every realization, we flip a coin that has probability λ of coming up heads; if the coin comes up heads, we produce a system with state ρ , whereas if the coin comes up tails, we produce a system with state $\tilde{\rho}$. Clearly any observable of the system (i.e., an observable that has no access to our coin) will have the same expectation under this *randomized* state as under ρ' . But we can now produce a whole family of states by changing the probability of our coin coming up heads, including the state ρ by letting the probability of heads be one. These classical absolutely continuous changes of measure commute with all the system observables, and hence provide us with precisely what we need. We will detail the procedure further below.

Before we apply this idea to the filtering problem, let us take a moment to investigate the requirement that one of the initial states can be obtained as a

² In this thesis we have not discussed techniques which are used to obtain asymptotic bounds on filter stability (e.g., Lyapunov exponents of the filtering equation), which are not based on smoothing. The majority of these techniques rely on application of the Hilbert projective metric. In principle this metric can be defined in the setting of the positive cone in any vector lattice [Liv95], e.g., in the positive cone of the predual of a Von Neumann algebra. However, the explicit bounds on the Birkhoff contraction coefficient used in the filter stability literature do not appear to be easily established in the noncommutative case. Nonetheless, such an approach could prove to be fruitful in the future, using either the Hilbert metric or some other suitable metric.




FIGURE 5.2. Cartoon illustration of the convex notion of absolute continuity. Shown are five simplices, each with two highlighted states (thick dots). The grey line segments contain all convex combinations of the two states, while the red line segments show the possible extensions. The five cases correspond to: a. and b. $\mathbb{P} \leftrightarrow \mathbb{Q}$; c. $\mathbb{Q} \ll \mathbb{P}$; d. $\mathbb{P} \ll \mathbb{Q}$; and e. satisfies neither $\mathbb{Q} \ll \mathbb{P}$ nor $\mathbb{P} \ll \mathbb{Q}$.

convex combination. This natural idea provides a surprisingly appealing definition of absolute continuity for states on a Von Neumann algebra.

DEFINITION 5.3.1. Let \mathbb{P}, \mathbb{Q} be normal states on a Von Neumann algebra \mathcal{A} . We say that \mathbb{Q} is absolutely continuous with respect to \mathbb{P} , denoted as $\mathbb{Q} \ll \mathbb{P}$, if there exists a normal state \mathbb{Q}' such that $\mathbb{P} = \lambda\mathbb{Q} + (1 - \lambda)\mathbb{Q}'$ for some $\lambda \in (0, 1]$. We say that \mathbb{P} and \mathbb{Q} are equivalent, denoted as $\mathbb{P} \leftrightarrow \mathbb{Q}$, if $\mathbb{Q} \ll \mathbb{P}$ and $\mathbb{P} \ll \mathbb{Q}$.

The definition is illustrated graphically in figure 5.2. Shown are five convex sets, which represent the state space of a Von Neumann algebra. Evidently $\mathbb{P} \ll \mathbb{Q}$ implies that the line of convex combinations of these two states can be extended outwards at \mathbb{Q} ; the extensions are precisely the possible choices for \mathbb{Q}' .

The following characterization reduces our definition to a notion which has been studied in the theory of operator algebras [Sak98, sec. 1.24]. Hence the Sakai-Radon-Nikodym theorems apply, but we will not use them.

LEMMA 5.3.2. $\mathbb{Q} \ll \mathbb{P}$ iff there is an $\varepsilon > 0$ s.t. $\mathbb{P}(X) \geq \varepsilon\mathbb{Q}(X)$ for all $X \geq 0$.

PROOF. Suppose that $\mathbb{P} = \lambda\mathbb{Q} + (1 - \lambda)\mathbb{Q}'$ for some \mathbb{Q}' and $\lambda \in (0, 1]$. Then clearly $\mathbb{P} \geq \lambda\mathbb{Q}$, so we can set $\varepsilon = \lambda$. Conversely, suppose that there is an $\varepsilon \in (0, 1]$ s.t. $\mathbb{P} \geq \varepsilon\mathbb{Q}$; then any functional $\mathbb{Q}' = (\mathbb{P} - \varepsilon\mathbb{Q})/(\varepsilon)$ with $\lambda \in (0, \varepsilon)$ defines a normal state such that $\mathbb{P} = \lambda\mathbb{Q} + (1 - \lambda)\mathbb{Q}'$. Note that $\varepsilon > 1$ is impossible, as then $1 = \mathbb{P}(I) \geq \varepsilon\mathbb{Q}(I) = \varepsilon$ provides a contradiction. Hence the Lemma is proved. \square

The simplices drawn above could represent the space of measures on the probability space $\Omega = \{1, 2, 3\}$; from the pictures, it should be evident that the convex notion of absolute continuity coincides with the classical notion. In fact, this is generally the case, provided that the Radon-Nikodym derivative is bounded.

LEMMA 5.3.3. Let \mathbb{P}, \mathbb{Q} be two probability measures on some measurable space (Ω, \mathcal{F}) . Then $\mathbb{Q} \ll \mathbb{P}$ if and only if $\mathbb{Q} \ll \mathbb{P}$ with $\|d\mathbb{Q}/d\mathbb{P}\|_\infty < \infty$.

PROOF. We show that $\mathbb{Q} \ll \mathbb{P}$ with $\|d\mathbb{Q}/d\mathbb{P}\|_\infty < \infty$ is necessary and sufficient for $\mathbb{Q} \ll \mathbb{P}$ to hold. To show sufficiency, note that for any $X \geq 0$ we obtain $\mathbf{E}_\mathbb{Q}(X) = \mathbf{E}_\mathbb{P}(X d\mathbb{Q}/d\mathbb{P}) \leq \mathbf{E}_\mathbb{P}(X) \|d\mathbb{Q}/d\mathbb{P}\|_\infty$. Hence $\mathbb{Q} \ll \mathbb{P}$ follows from Lemma 5.3.2. To show necessity, consider first the case where $\mathbb{Q} \not\ll \mathbb{P}$. Then clearly $\mathbb{Q} \ll \mathbb{P}$ can not be true, as there exists some $X \geq 0$ such that $\mathbf{E}_\mathbb{P}(X) = 0$ but $\mathbf{E}_\mathbb{Q}(X) > 0$. Hence it remains to show that if $d\mathbb{Q}/d\mathbb{P}$ is unbounded, then $\mathbb{Q} \ll \mathbb{P}$ can not hold. To this end, consider the sets $S_n = \{\omega \in \Omega : d\mathbb{Q}/d\mathbb{P} \geq n\}$.

Then $\mathbf{Q}(S_n) = \mathbf{E}_{\mathbf{P}}(I_{S_n} d\mathbf{Q}/d\mathbf{P}) \geq n \mathbf{P}(S_n)$, so by Lemma 5.3.2 $\mathbf{Q} \prec\prec \mathbf{P}$ can not hold. The result follows immediately. \square

In the noncommutative case, let us characterize absolute continuity for states on $\mathcal{B}(\mathfrak{h})$ when \mathfrak{h} is a finite dimensional Hilbert space. Recall that any such state can be written uniquely as $\rho(X) = \text{Tr}[\varrho X]$ for some density matrix ϱ .

LEMMA 5.3.4. *Let ρ, ρ' be states on $\mathcal{B}(\mathfrak{h})$, $\dim \mathfrak{h} < \infty$, which are defined by the density matrices ϱ, ϱ' . Then $\rho \prec\prec \rho'$ iff $\ker \varrho \supset \ker \varrho'$.*

PROOF. We show that $\ker \varrho \supset \ker \varrho'$ is necessary and sufficient for $\rho \prec\prec \rho'$ to hold. Recall from Lemma 5.3.2 that $\rho \prec\prec \rho'$ is equivalent to $\rho' \geq \varepsilon \rho$, or alternatively that $\rho' \geq \varepsilon \varrho$, for some $\varepsilon > 0$. To prove necessity, let $v \in \ker \varrho'$ but $v \notin \ker \varrho$. Then $\langle v, \varrho'v \rangle = 0$, while $\langle v, \varrho v \rangle = \|\varrho^{1/2}v\|^2 > 0$. But then $\langle v, \varrho'v \rangle < \varepsilon \langle v, \varrho v \rangle$ for any $\varepsilon > 0$, and this contradicts $\rho \prec\prec \rho'$. It thus remains to show sufficiency. To this end, let us restrict ϱ, ϱ' to the subspace $\mathfrak{h}' = (\ker \varrho')^\perp \subset \mathfrak{h}$. It suffices to show that $\varrho'|_{\mathfrak{h}'} \geq \varepsilon \varrho|_{\mathfrak{h}'}$ for some $\varepsilon > 0$. But note that $\varrho'|_{\mathfrak{h}'}$ has full rank and hence is positive definite, so there is some $\varepsilon > 0$ such that $\langle v, \varrho'v \rangle \geq \varepsilon \|v\|^2$ for all $v \in \mathfrak{h}'$. But the eigenvalues of ϱ must be contained in $[0, 1]$, so that $\langle v, \varrho v \rangle \leq \|v\|^2$ for any $v \in \mathfrak{h}'$. Thus we find that $\langle v, \varrho'v \rangle \geq \varepsilon \langle v, \varrho v \rangle$ for all $v \in \mathfrak{h}$, and the Lemma follows. \square

Note in particular the two extremes: any state is absolutely continuous with respect to a state with density matrix of full rank; whereas no state is absolutely continuous with respect to a vector state, except that state itself.

We now return to the reason for introducing the notion of absolute continuity. Here $D_n(\mathbb{C})$ denotes the Von Neumann algebra of $n \times n$ diagonal matrices.

LEMMA 5.3.5. *Let \mathcal{A} be a Von Neumann algebra and $\mathbb{Q} \prec\prec \mathbb{P}$ be normal states on \mathcal{A} . Then there exists a normal state \mathbb{P}^\sim on $\mathcal{A}^\sim = D_2(\mathbb{C}) \otimes \mathcal{A}$, and a nonnegative $\Lambda \in D_2(\mathbb{C})$, such that $\mathbb{P}^\sim(I \otimes X) = \mathbb{P}(X)$ and $\mathbb{P}^\sim(\Lambda \otimes X) = \mathbb{Q}(X)$ for all $X \in \mathcal{A}$. Moreover, if $\mathbb{P} \rightsquigarrow \mathbb{Q}$, then Λ can be chosen to be strictly positive.*

PROOF. Let \mathbb{Q}' and $\lambda \in (0, 1]$ be such that $\mathbb{P} = \lambda \mathbb{Q} + (1 - \lambda) \mathbb{Q}'$, let $P_1 = \text{diag}\{1, 0\}$, $P_2 = \text{diag}\{0, 1\}$, and let ρ_1, ρ_2 be the states on $D_2(\mathbb{C})$ for which P_1 and P_2 are the corresponding density matrices. We define $\mathbb{P}^\sim = \lambda \rho_1 \otimes \mathbb{Q} + (1 - \lambda) \rho_2 \otimes \mathbb{Q}'$, and set $\Lambda = P_1/\lambda$. The statement of the Lemma is now easily verified. In the case that $\mathbb{P} \rightsquigarrow \mathbb{Q}$, we can choose another state \mathbb{Q}'' such that $\mathbb{P} = \lambda \mathbb{Q}' + (1 - \lambda) \mathbb{Q}''$, $\mathbb{Q} = \lambda' \mathbb{Q}' + (1 - \lambda') \mathbb{Q}''$ with $\lambda, \lambda' > 0$, and repeat the above construction. \square

We can now easily obtain the following Bayes-type formula.

LEMMA 5.3.6. *Let \mathcal{A} be a Von Neumann algebra, let $\mathbb{Q} \rightsquigarrow \mathbb{P}$ be normal states on \mathcal{A} , and let $\mathcal{C} \subset \mathcal{A}$ be a commutative subalgebra. Then for any $X \in \mathcal{C}$*

$$I \otimes \mathbb{P}(X|\mathcal{C}) = \mathbb{P}^\sim(I \otimes X|I \otimes \mathcal{C}), \quad I \otimes \mathbb{Q}(X|\mathcal{C}) = \frac{\mathbb{P}^\sim(\Lambda \otimes X|I \otimes \mathcal{C})}{\mathbb{P}^\sim(\Lambda \otimes I|I \otimes \mathcal{C})}.$$

PROOF. The first statement follows from the definition of the conditional expectation, the second statement follows from Lemma 5.1.6 with $V = (\Lambda)^{1/2} \otimes I$. \square

Using this result, we can now proceed to obtain a noncommutative counterpart of the second representation described in section 2.3.1; i.e., we can obtain expressions for differently initialized filters in terms of a single underlying state \mathbb{P}^\sim , provided the initial states are equivalent in the convex sense. This could provide a

starting point for further investigation of the filter stability problem. In the next section, however, we go down a slightly different path.

5.3.2. A generic stability result. The goal of this section is to prove a simple stability result: a quantum counterpart of [COC99, Thm. 3.1] (see also [CL06] for a different perspective). We will use a different proof and we even go a little further than the type of result described in [COC99], but we restrict ourselves for simplicity to finite-dimensional initial systems, as in the rest of this chapter.

REMARK 5.3.7. It should be emphasized that the proof given in [COC99] for the classical case can be extended also to the quantum case using the machinery developed above. Rather than repeat this proof here, we give a different proof which is more direct. However, the original proof of [COC99] is more convenient in the infinite-dimensional case, as very little regularity is required for their argument.

To obtain some insight on what we are going to prove, recall that the innovations process $d\bar{z}_t = dy_t - \bar{\pi}_t(L + L^*) dt$ is a Wiener process under \mathbf{P} . Hence the observation process y_t has the law of a Wiener process with drift $\bar{\pi}_t(L + L^*)$; the observable $L + L^*$ is called the *measurement observable*. It thus seems plausible that at least the filtered estimate $\bar{\pi}_t(L + L^*)$ should be fairly easy to obtain, even if the entire filter is not; intuitively, the observations provide some “direct” information on the measurement observable, while even naive estimation of other observables requires us to use our knowledge of the underlying model. (Note that the measurement observable plays a similar role in the quantum case as does the observation function in the classical case.)

What we will show is that the filtered estimate of the measurement observable is always stable, regardless of the underlying model, provided that the initial states are absolutely continuous. This does not guarantee that all filtered estimates are stable, nor is there a bound on the rate of stability. The resolution of such issues would necessarily have to take into account the details of the underlying model. On the other hand, the beauty of this result is its generic applicability—we have to know absolutely nothing about the underlying model in order for this to hold!

Let us get down to the details. As before, $\bar{\pi}_t(X)$ denotes the filtered estimate of the system observable X when the initial state is ρ . The measure \mathbf{P} is the measure on the space of observation sample paths corresponding to the initial state ρ . We also introduce $\bar{\pi}'_t(X)$, which is obtained by solving the filtering equation with the misspecified initial state ρ' ; in other words, $\bar{\pi}_t(X) = \iota(\mathbb{P}(j_t(X \otimes I) | \mathcal{Y}_t))$, while $\bar{\pi}'_t(X) = \iota(\mathbb{P}'(j_t(X \otimes I) | \mathcal{Y}_t))$ where $\mathbb{P}' = \rho' \otimes \varphi$.

The following result is the key step to establish our claim.

PROPOSITION 5.3.8. *Suppose that $\rho \rightsquigarrow \rho'$. Then there is some finite constant $C < \infty$, depending only on ρ and ρ' , such that for any $t < \infty$*

$$\mathbf{E}_{\mathbf{P}} \left[\int_0^t |\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2 ds \right] \leq C.$$

PROOF. As $\rho \rightsquigarrow \rho'$, there is by Lemma 5.3.5 a state ρ^\sim on $\mathcal{B}^\sim = D_2(\mathbb{C}) \otimes \mathcal{B}$ and a strictly positive $\Lambda \in D_2(\mathbb{C})$ s.t. $\rho^\sim(I \otimes X) = \rho(X)$ and $\rho^\sim(\Lambda \otimes X) = \rho'(X)$. This implies also that $\mathbb{P} \rightsquigarrow \mathbb{P}'$, so that the state $\mathbb{P}^\sim = \rho^\sim \otimes \varphi$ on $\mathcal{A}^\sim = \mathcal{B}^\sim \otimes \mathcal{W}$ satisfies $\mathbb{P}^\sim(I \otimes X) = \mathbb{P}(X)$ and $\mathbb{P}^\sim(\Lambda \otimes X) = \mathbb{P}'(X)$ for all $I \otimes X \in D_2(\mathbb{C}) \otimes \mathcal{A}$. We extend $U_t^\sim = I \otimes U_t$ in $D_2(\mathbb{C}) \otimes \mathcal{A}$, and define $j_t^\sim(X) = U_t^{\sim*}(X \otimes I)U_t^\sim$ for $X \otimes I \in \mathcal{B}^\sim \otimes \mathcal{W}$. We also extend $\mathcal{Y}_t^\sim = I \otimes \mathcal{Y}_t$ in $D_2(\mathbb{C}) \otimes \mathcal{A}$.

To prove the claim, we consider the extended filtering problem of calculating $\bar{\pi}_t^\sim(X) = \iota(\mathbb{P}^\sim(j_t^\sim(X)|\mathcal{Y}_t^\sim))$. This filtering problem can be interpreted as exactly our previous filtering problem, only with an enlarged initial system $D_2(\mathbb{C}) \otimes \mathcal{B}$. In particular, we easily obtain the filtering equation for $\bar{\pi}_t^\sim(X)$:

$$d\bar{\pi}_t^\sim(X) = \bar{\pi}_t^\sim(\mathcal{L}^\sim(X)) dt + (\bar{\pi}_t^\sim(\tilde{L}^* X + X \tilde{L}) - \bar{\pi}_t^\sim(\tilde{L} + \tilde{L}^*) \bar{\pi}_t^\sim(X)) d\bar{z}_t,$$

where $\bar{\pi}_0^\sim(X) = \rho^\sim(X)$, $\tilde{L} = I \otimes L$, $\tilde{H} = I \otimes H$, and

$$\mathcal{L}^\sim(X) = i[\tilde{H}, X] + \tilde{L}^* X \tilde{L} - \frac{1}{2} \tilde{L}^* \tilde{L} X - \frac{1}{2} X \tilde{L}^* \tilde{L}.$$

Then we obtain the following equation:

$$\begin{aligned} d\bar{\pi}_t^\sim(\Lambda \otimes I) &= (\bar{\pi}_t^\sim(\Lambda \otimes (L + L^*)) - \bar{\pi}_t^\sim(I \otimes (L + L^*)) \bar{\pi}_t^\sim(\Lambda \otimes I)) d\bar{z}_t \\ &= \left\{ \frac{\bar{\pi}_t^\sim(\Lambda \otimes (L + L^*))}{\bar{\pi}_t^\sim(\Lambda \otimes I)} - \bar{\pi}_t^\sim(I \otimes (L + L^*)) \right\} \bar{\pi}_t^\sim(\Lambda \otimes I) d\bar{z}_t. \end{aligned}$$

But by Lemma 5.3.6 we can simplify this to

$$d\bar{\pi}_t^\sim(\Lambda \otimes I) = \{\bar{\pi}'_t(L + L^*) - \bar{\pi}_t(L + L^*)\} \bar{\pi}_t^\sim(\Lambda \otimes I) d\bar{z}_t,$$

so that evidently it is the case that

$$\bar{\pi}_t^\sim(\Lambda \otimes I) = \exp \left[\int_0^t \Delta_s d\bar{z}_s - \frac{1}{2} \int_0^t |\Delta_s|^2 ds \right],$$

with $\Delta_t = \bar{\pi}'_t(L + L^*) - \bar{\pi}_t(L + L^*)$. We obtain immediately

$$-2 \mathbf{E}_{\mathbf{P}}(\log(\bar{\pi}_t^\sim(\Lambda \otimes I))) = \mathbf{E}_{\mathbf{P}} \left[\int_0^t |\bar{\pi}'_s(L + L^*) - \bar{\pi}_s(L + L^*)|^2 ds \right].$$

But as $0 < \Lambda < \infty$, the left-hand side is finite and is bounded by some constant that only depends on Λ . This establishes the claim. \square

We immediately obtain the following weak stability result.

COROLLARY 5.3.9. *Suppose that $\rho \rightsquigarrow \rho'$. Then*

$$\int_0^\infty \mathbf{E}_{\mathbf{P}} |\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2 ds < \infty.$$

REMARK 5.3.10. Note that we have previously defined the random processes $\bar{\pi}_t(X)$, $\bar{\pi}'_t(X)$ on some finite time interval $[0, T]$; this is not crucial, but technically very convenient (see, e.g., Example 1.1.3 for the sort of issues we could run into otherwise). Fortunately there is no reason to extend these random processes to the infinite time interval, as we are only interested in expectations of the quantity $|\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2$ for fixed finite times s . Hence it suffices, for any fixed s , to choose $T \geq s$ in order to calculate $\mathbf{E}_{\mathbf{P}} |\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2$. The integral of these expectations can be taken subsequently, and Corollary 5.3.9 applies.

The remainder of this section is devoted to strengthening the above result to the statement that $\mathbf{E}_{\mathbf{P}} |\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2 \rightarrow 0$ as $t \rightarrow \infty$. This seems almost obvious from the above result; certainly this quantity can only spend a finite amount of time outside any neighborhood of zero, but this is not sufficient to establish the claim. Ostensibly, it could be that $\mathbf{E}_{\mathbf{P}} |\bar{\pi}_s(L + L^*) - \bar{\pi}'_s(L + L^*)|^2$ has a positive limit superior, while still being integrable over the infinite time horizon. This type of problem is often encountered in the study of stochastic stability theory [Kus72, DKW01, Van06], where the strong Markov property of the underlying

dynamics can be used to eliminate this possibility. A simpler technique suffices in our case, however, as we can establish the following simple result.

LEMMA 5.3.11. $\mathbf{E}_{\mathbf{P}}|\bar{\pi}_t(\mathbf{L} + \mathbf{L}^*) - \bar{\pi}'_t(\mathbf{L} + \mathbf{L}^*)|^2$ is Lipschitz continuous in t .

PROOF. This follows immediately by applying the Itô rule to the quantity $|\bar{\pi}_t(\mathbf{L} + \mathbf{L}^*) - \bar{\pi}'_t(\mathbf{L} + \mathbf{L}^*)|^2$, taking the expectation (so that the integral with respect to the innovations process vanishes), and noting that the right-hand side is bounded, i.e., $\mathbf{E}_{\mathbf{P}}|\bar{\pi}_t(\mathbf{L} + \mathbf{L}^*) - \bar{\pi}'_t(\mathbf{L} + \mathbf{L}^*)|^2$ has bounded time derivative. \square

We can now establish the main result of this section.

THEOREM 5.3.12. Suppose that $\rho \rightsquigarrow \rho'$. Then

$$\mathbf{E}_{\mathbf{P}}|\bar{\pi}_t(\mathbf{L} + \mathbf{L}^*) - \bar{\pi}'_t(\mathbf{L} + \mathbf{L}^*)|^2 \longrightarrow 0 \quad \text{as } t \longrightarrow \infty.$$

PROOF. Define $R_t = \mathbf{E}_{\mathbf{P}}|\bar{\pi}_t(\mathbf{L} + \mathbf{L}^*) - \bar{\pi}'_t(\mathbf{L} + \mathbf{L}^*)|^2$. From Corollary 5.3.9, it is immediately clear that $\liminf_{t \rightarrow \infty} R_t = 0$. We will prove $\limsup_{t \rightarrow \infty} R_t = 0$ by contradiction. To this end, suppose that $\limsup_{t \rightarrow \infty} R_t = k > 0$. Then there exist $0 < k' < k'' < k$ such that R_t crosses k' and k'' infinitely often. Let t'' be a time such that $R_{t''} = k''$, and let t' be the latest time previous to t'' such that $R_{t'} = k'$. By Lipschitz continuity of R_t , we find that $|R_{t''} - R_{t'}| \leq K|t'' - t'|$, where the Lipschitz constant $K > 0$ is independent of t', t'' . But by construction, this implies that $|t'' - t'| \geq K^{-1}|k'' - k'|$. As this happens infinitely often, this means that R_t must infinitely often spend a time in excess of $K^{-1}|k'' - k'|$ being larger than k' . But this would surely contradict Corollary 5.3.9, and the result is established. \square

REMARK 5.3.13. The results in this section extend trivially to the imperfect detection case of section 5.2.4, and extensions to the case of multiple channels, etc., are similarly straightforward. In the case of an infinite-dimensional initial system the result is also readily established, though the method of proof used in [COC99] may be more convenient in this case. In fact, as in the classical case, the quantum case appears to enjoy a general filter stability principle in the context of diffusive (homodyne) detection: the filtered estimate of the measurement observable (or observables in the case of higher multiplicity) is always stable.

Elements of the Malliavin Calculus

The goal of this appendix is to recall briefly the main results of the Malliavin calculus, Skorokhod integrals and anticipative stochastic calculus that are needed in chapter 3 (and briefly in section 4.3.4). In our application of the theory we wish to deal with functionals of the observation process $(Y_t)_{t \in [0, T]}$, where T is some finite time (usually we will calculate integrals from 0 to t , so we can choose any $T > t$). Recall that Y is an \mathcal{F}_t^Y -Wiener process under the measure \mathbf{Q} ; it will thus be convenient to work always under \mathbf{Q} , as this puts us directly in the framework used, e.g., in [Nua95]. As the theory described below is defined \mathbf{Q} -a.s. and as $\mathbf{P} \sim \mathbf{Q}$, the corresponding properties under \mathbf{P} are unambiguously obtained by using (3.11). We will presume this setup whenever the theory described here is applied.

A.1. The Malliavin derivative: Definition and elementary properties

A smooth random variable F is one of the form $f(Y(h_1), \dots, Y(h_n))$, where $Y(h)$ denotes the Wiener integral of the deterministic function $h \in L^2([0, T])$ with respect to Y and f is a smooth function which is of polynomial growth together with all its derivatives. For smooth F the Malliavin derivative $\mathbf{D}F$ is defined by

$$\mathbf{D}_t F = \sum_{i=1}^n \frac{\partial f}{\partial x^i}(Y(h_1), \dots, Y(h_n)) h_i(t).$$

The Malliavin derivative \mathbf{D} can be shown [Nua95, p. 26] to be closable as an operator from $L^p(\Omega, \mathcal{F}_T^Y, \mathbf{Q})$ to $L^p(\Omega, \mathcal{F}_T^Y, \mathbf{Q}; L^2([0, T]))$ for any $p \geq 1$, and we denote the domain of \mathbf{D} in $L^p(\Omega)$ by $\mathbb{D}^{1,p}$ (for notational convenience we will drop the measure \mathbf{Q} and σ -algebra \mathcal{F}_T^Y throughout this section, where it is understood that $L^p(\Omega)$ denotes $L^p(\Omega, \mathcal{F}_T^Y, \mathbf{Q})$, etc.). In fact, $\mathbb{D}^{1,p}$ is simply the closure of the set of smooth random variables in $L^p(\Omega)$ with respect to the norm

$$\|F\|_{1,p} = \left[\mathbf{E}_{\mathbf{Q}}|F|^p + \mathbf{E}_{\mathbf{Q}}\|\mathbf{D}F\|_{L^2([0, T])}^p \right]^{1/p}.$$

More generally, we consider iterated derivatives $\mathbf{D}^k F \in L^p(\Omega; L^2([0, T]^k))$ defined by $\mathbf{D}_{t_1, \dots, t_k}^k F = \mathbf{D}_{t_1} \cdots \mathbf{D}_{t_k} F$. The domain of \mathbf{D}^k in $L^p(\Omega)$ is denoted by $\mathbb{D}^{k,p}$, and coincides with the closure in $L^p(\Omega)$ of the smooth random variables with respect to the norm

$$\|F\|_{k,p} = \left[\mathbf{E}_{\mathbf{Q}}|F|^p + \sum_{j=1}^k \mathbf{E}_{\mathbf{Q}}\|\mathbf{D}^j F\|_{L^2([0, T]^j)}^p \right]^{1/p}.$$

The local property of the Malliavin derivative allows us to localize these domains [Nua95, pp. 44–45]. For $F \in L^2(\Omega)$, suppose there exists a sequence $(\Omega_n, F_n)_{n \geq 1}$ with $\Omega_n \in \mathcal{F}_T^Y$ and $F_n \in \mathbb{D}^{k,p}$, such that $\Omega_n \nearrow \Omega$ a.s. and $F = F_n$ a.s. on Ω_n . Then $(\Omega_n, F_n)_{n \geq 1}$ localizes F in $\mathbb{D}^{k,p}$, and we define $\mathbf{D}F = \mathbf{D}F_n$ on Ω_n . The space of random variables that can be localized in $\mathbb{D}^{k,p}$ is denoted by $\mathbb{D}_{\text{loc}}^{k,p}$.

The first result we will need is a chain rule for the Malliavin derivative.

PROPOSITION A.1.1. *Let $\varphi : \mathbb{R}^m \rightarrow \mathbb{R}$ be C^1 and $F = (F^1, \dots, F^m)$ be a random vector with components in $\mathbb{D}^{1,2}$. Then $\varphi(F) \in \mathbb{D}_{\text{loc}}^{1,2}$ and*

$$\mathbf{D}\varphi(F) = \sum_{i=1}^m \frac{\partial \varphi}{\partial x^i}(F) \mathbf{D}F^i.$$

If $\varphi(F) \in L^2(\Omega)$ and $\mathbf{D}\varphi(F) \in L^2(\Omega \times [0, T])$, then $\varphi(F) \in \mathbb{D}^{1,2}$. These results still hold if F a.s. takes values in an open domain $V \subset \mathbb{R}^m$ and φ is $C^1(V)$.

The first (local) statement can be found in [NP88, Prop. 2.9]; the second statement can be proved in the same way as [OK91, Lemma A.1], and the proofs are easily adapted to the case where F a.s. takes values in some domain.

A useful class of random variables is $\mathbb{D}^\infty = \bigcap_{p \geq 1} \bigcap_{k \geq 1} \mathbb{D}^{k,p}$. Then $\mathbf{D}_t F \in \mathbb{D}^\infty$ for any $F \in \mathbb{D}^\infty$, and the chain rule extends as follows [Nua95, p. 62].

PROPOSITION A.1.2. *Let $\varphi : \mathbb{R}^m \rightarrow \mathbb{R}$ be a smooth function which is of polynomial growth together with all its derivatives, and let $F = (F^1, \dots, F^m)$ be a random vector with components in \mathbb{D}^∞ . Then $\varphi(F) \in \mathbb{D}^\infty$ and the usual chain rule holds. This implies that \mathbb{D}^∞ is an algebra, i.e., $FG \in \mathbb{D}^\infty$ for $F, G \in \mathbb{D}^\infty$.*

We will also need the following property [Nua95, p. 32].

LEMMA A.1.3. *For a Borel set $A \subset [0, T]$, denote by \mathcal{F}_A^Y the σ -algebra generated by the random variables $\{Y(I_B) : B \subset A \text{ Borel}\}$. Let $F \in \mathbb{D}^{1,2}$ be \mathcal{F}_A^Y -measurable. Then $\mathbf{D}_t F = 0$ a.e. in $\Omega \times ([0, T] \setminus A)$.*

A.2. The Malliavin derivative of a stochastic flow

It is useful to be able to calculate explicitly the Malliavin derivative of the solution of a stochastic differential equation. Consider

$$dx_t = f(x_t) dt + \sigma(x_t) dY_t, \quad x_0 \in \mathbb{R}^m,$$

where $f(x)$ and $\sigma(x)$ are smooth functions of x with bounded derivatives of all orders. It is well known that such equations generate a smooth stochastic flow of diffeomorphisms $x_t = \xi_t(x)$ [Kun84]. We now have the following result.

PROPOSITION A.2.1. *All components of x_t belong to \mathbb{D}^∞ for every $t \in [0, T]$. We have $\mathbf{D}_r x_t = D\xi_t(x_0)D\xi_r(x_0)^{-1}\sigma(x_r)$ a.e. $r < t$, where $(D\xi_t(x))^{ij} = \partial \xi_t^i(x) / \partial x^j$ is the Jacobian matrix of the flow, and $\mathbf{D}_r x_t = 0$ a.e. $r > t$.*

The first statement is given in [Nua95, Theorem 2.2.2, p. 105], the second on [Nua95, eq. (2.38), p. 109]. $\mathbf{D}_r x_t = 0$ a.e. $r > t$ follows from adaptedness.

If f, σ are only C^1 with bounded derivative, then above result still holds with appropriately reduced regularity, see [Nua95, Theorem 2.2.1, p. 102].

PROPOSITION A.2.2. *If $f, \sigma \in C^1$ with bounded derivative, then $x_t \in \mathbb{D}^{1,p}$ for all $p \geq 1$ and $t \in [0, T]$. We have $\mathbf{D}_r x_t = D\xi_t(x_0)D\xi_r(x_0)^{-1}\sigma(x_r)$ a.e. $r < t$, where $D\xi_t(x)$ is the Jacobian of the flow, and $\mathbf{D}_r x_t = 0$ a.e. $r > t$.*

A.3. The Clark-Haussmann-Ocone formula

Let $F \in \mathbb{D}^{1,2}$. Then by the usual Itô representation theorem, it is known that we can write F as the sum of a constant and an Itô integral. The Malliavin calculus gives an explicit expression for the integrand, and this result is known as the Clark-Ocone formula, see [Nua95, Prop. 1.3.5, p. 42] or [OK91].

PROPOSITION A.3.1. *Let $F \in \mathbb{D}^{1,2}$. Then we can write*

$$F = \mathbf{E}_{\mathbf{Q}}(F) + \int_0^T \mathbf{E}_{\mathbf{Q}}(\mathbf{D}_t F | \mathcal{F}_t^Y) dY_t.$$

In the particular case that $F = f(\xi_t(x))$, where $\xi_t(x)$ is the stochastic flow generated by a stochastic differential equation, this representation formula becomes even more explicit by substituting the explicit expression for the Malliavin derivative of a stochastic flow. This is called the Clark-Haussmann-Ocone formula.

A.4. The Skorokhod integral: Definition and elementary properties

We now consider the Malliavin derivative as a closed operator from $L^2(\Omega)$ to $L^2(\Omega \times [0, T])$ with domain $\mathbb{D}^{1,2}$. Its Hilbert space adjoint $\delta = \mathbf{D}^*$ is well defined in the usual sense as a closed operator from $L^2(\Omega \times [0, T])$ to $L^2(\Omega)$, and we denote its domain by $\text{Dom } \delta$. The operator δ is called the Skorokhod integral, and coincides with the Itô integral on the subspace $L_a^2(\Omega \times [0, T]) \subset L^2(\Omega \times [0, T])$ of adapted square integrable processes [Nua95, Prop. 1.3.4, p. 41].

LEMMA A.4.1. *$L_a^2(\Omega \times [0, T]) \subset \text{Dom } \delta$, and $\delta|_{L_a^2}$ coincides with the Itô integral*

$$\delta(u) = \int_0^T u_t dY_t, \quad \forall u \in L_a^2(\Omega \times [0, T]).$$

The Skorokhod integral is thus an extension of the Itô integral to a class of possibly anticipative integrands. To emphasize this point we will use the same notation for Skorokhod integrals as for Itô integrals, i.e., we will write

$$\delta(uI_{[s,t]}) = \int_s^t u_r dY_r, \quad uI_{[s,t]} \in \text{Dom } \delta.$$

The Skorokhod integral has the following properties. First, its expectation vanishes $\mathbf{E}_{\mathbf{Q}}\delta(u) = 0$ if $u \in \text{Dom } \delta$. Second, by its definition as the adjoint of \mathbf{D} we have

$$(A.1) \quad \mathbf{E}_{\mathbf{Q}}(F\delta(u)) = \mathbf{E}_{\mathbf{Q}} \left[\int_0^T (\mathbf{D}_t F) u_t dt \right]$$

if $u \in \text{Dom } \delta$, $F \in \mathbb{D}^{1,2}$. We will also use the following result, the proof of which proceeds in the same way as its one-dimensional counterpart [Nua95, p. 40].

LEMMA A.4.2. *If u is an n -vector of processes in $\text{Dom } \delta$ and F is an $m \times n$ matrix of random variables in $\mathbb{D}^{1,2}$ such that $\mathbf{E}_{\mathbf{Q}} \int_0^T \|Fu_t\|^2 dt < \infty$, then*

$$\int_0^T Fu_t dY_t = F \int_0^T u_t dY_t - \int_0^T (\mathbf{D}_t F) u_t dt$$

in the sense that $Fu \in \text{Dom } \delta$ iff the right-hand side of this expression is in $L^2(\Omega)$.

A.5. Anticipative stochastic calculus

As it is difficult to obtain general statements for integrands in $\text{Dom } \delta$, it is useful to single out restricted classes of integrands that are easier to deal with. To this end, define the space $\mathbb{L}^{1,2} = L^2([0, T]; \mathbb{D}^{1,2})$, i.e., the space of processes u such that $u_t \in \mathbb{D}^{1,2}$ and such that the norm

$$\|u\|_{1,2} = \left[\|u\|_{L^2(\Omega \times [0, T])}^2 + \|\mathbf{D}u\|_{L^2(\Omega \times [0, T]^2)}^2 \right]^{1/2}$$

is finite. Similarly, we define $\mathbb{L}^{k,p} = L^p([0, T]; \mathbb{D}^{k,p})$ for $k \geq 1$, $p \geq 2$. Then $\mathbb{L}^{k,p} \subset \mathbb{L}^{1,2} \subset \text{Dom } \delta$ [Nua95, p. 38]. Moreover, the Skorokhod integral satisfies the local property on $\mathbb{L}^{1,2}$, so that the domains $\mathbb{L}^{k,p}$ can be localized to $\mathbb{L}_{\text{loc}}^{k,p}$ in the same way as we localized $\mathbb{D}^{k,p}$ to $\mathbb{D}_{\text{loc}}^{k,p}$ [Nua95, pp. 43–45].

We are now in the position to state the Itô change of variables formula for Skorokhod integral processes. Various versions of the formula can be found in [NP88, OP89, Nua95]. The extension to processes that a.s. take values in some domain is straightforward through localization.

PROPOSITION A.5.1. *Consider an m -dimensional process of the form*

$$x_t = x_0 + \int_0^t v_s ds + \int_0^t u_s dY_s,$$

where we assume that x_t has a continuous version and $x_0 \in (\mathbb{D}_{\text{loc}}^{1,4})^m$, $v \in (\mathbb{L}_{\text{loc}}^{1,4})^m$, and $u \in (\mathbb{L}_{\text{loc}}^{2,4})^m$. Let $\varphi : \mathbb{R}^m \rightarrow \mathbb{R}$ be a C^2 function. Then

$$\varphi(x_t) = \varphi(x_0) + \int_0^t D\varphi(x_s)v_s ds + \int_0^t D\varphi(x_s)u_s dY_s + \frac{1}{2} \int_0^t (D^2\varphi(x_s)\nabla_s x_s, u_s) ds,$$

where we write $\nabla_s x_s = \lim_{\varepsilon \searrow 0} \mathbf{D}_s(x_{s+\varepsilon} + x_{s-\varepsilon})$, $D\varphi(x_s)u_s = \sum_i (\partial\varphi/\partial x^i)(x_s)u_s^i$, $(D^2\varphi(x_s)\nabla_s x_s, u_s) = \sum_{ij} (\partial^2\varphi/\partial x^i \partial x^j)(x_s)u_s^i \nabla_s x_s^j$. The result still holds if x_s a.s. takes values in an open domain $V \subset \mathbb{R}^m \forall s \in [0, t]$ and φ is $C^2(V)$.

Elements of Quantum Probability

The goal of this appendix is to provide a brief introduction to the concepts from quantum probability that are needed in chapter 5: quantum probability spaces and the spectral theorem, Fock space quantum noises and the Hudson-Parthasarathy quantum stochastic calculus. The focus here is on recalling the necessary technical machinery. Introductions to quantum probability can be found in the papers [BVJ06b, BVJ06a], or in the lecture notes by H. Maassen [Maa03]. An excellent introduction to quantum stochastic calculus appears in the lecture notes by R. L. Hudson [Hud03], and extensive developments appear in the books by K. R. Parthasarathy [Par92], P.-A. Meyer [Mey93], P. Biane [Bia95], and A. M. Chebotarev [Che00]. The necessary background on functional analysis can be found in the classic textbook by M. Reed and B. Simon [RS72] or in J. B. Conway [Con85], and an extensive study of operator algebras appears in the textbooks of R. V. Kadison and J. R. Ringrose [KR97a, KR97b], O. Bratteli and D. W. Robinson [BR87], S. Sakai [Sak98], G. K. Pedersen [Ped79], or M. Takesaki [Tak02].

B.1. Quantum probability

B.1.1. Quantum probability spaces. In the setting of elementary quantum mechanics, the basic constructions are usually introduced as follows. We begin with a Hilbert space H . A bounded, self-adjoint operator on H is called an observable, and plays the role of a bounded, real-valued random variable (we will deal with unbounded observables later on). A unit trace nonnegative operator ρ (the density operator) is used to represent the state of the system, and plays the role of a probability density. In particular, $\text{Tr}[\rho X]$ is the expectation of the observable X , and $\text{Tr}[\rho I_A(X)]$ is the probability of the event $X \in A$ (A is a Borel set on \mathbb{R}).

In order to engage in probability theory, we need a little more substance. We will need the counterparts of σ -algebras and filtrations in the classical theory, and we will need machinery to transform between the operator picture of observables and a more traditional probabilistic description of random variables on a certain measure space. In the next few sections we will introduce the necessary concepts.

Let us first introduce some important notions.

DEFINITION B.1.1. A $*$ -algebra is a set \mathcal{A} of operators that contains the identity ($I \in \mathcal{A}$), is closed under linear combinations ($X, Y \in \mathcal{A}$, $\alpha, \beta \in \mathbb{C} \Rightarrow \alpha X + \beta Y \in \mathcal{A}$), multiplication ($X, Y \in \mathcal{A} \Rightarrow XY \in \mathcal{A}$), and involution ($X \in \mathcal{A} \Rightarrow X^* \in \mathcal{A}$).

DEFINITION B.1.2. Let H be a complex Hilbert space. We denote by $\mathcal{B}(H)$ the $*$ -algebra of all bounded linear operators on H .

DEFINITION B.1.3. A $*$ -subalgebra \mathcal{A} of $\mathcal{B}(\mathbf{H})$ is said to be monotone if for any upper bounded increasing net $\{A_\alpha\} \subset \mathcal{A}$ with nonnegative elements $A_\alpha \geq 0$, we have $\sup_\alpha A_\alpha \in \mathcal{A}$. A positive linear functional $\mu : \mathcal{A} \rightarrow \mathbb{C}$ is said to be normal if we have $\mu(\sup_\alpha A_\alpha) = \sup_\alpha \mu(A_\alpha)$ for any such net $\{A_\alpha\}$.

DEFINITION B.1.4. A Von Neumann algebra (on \mathbf{H}) is a $*$ -subalgebra \mathcal{A} of $\mathcal{B}(\mathbf{H})$ which is monotone. A normal state on \mathcal{A} is a normal positive linear functional $\mathbb{P} : \mathcal{A} \rightarrow \mathbb{C}$ such that $\mathbb{P}(I) = 1$.

Von Neumann algebras and states play a fundamental role in quantum probability. A Von Neumann algebra can be thought of as a σ -algebra, except that it holds observables rather than only events. To get some intuition for this idea, recall the monotone class theorem in classical real analysis [Pro04, page 7]:

Let \mathcal{M} be a $*$ -algebra of bounded functions on some set Ω , and suppose that for any bounded increasing sequence $\{f_n\} \subset \mathcal{M}$, $f_n \geq 0$, we have $\lim_{n \rightarrow \infty} f_n \in \mathcal{M}$. Then \mathcal{M} is precisely the set of all bounded $\sigma\{\mathcal{M}\}$ -measurable functions.

Hence in the classical case, we see that there is a one-to-one correspondence between σ -algebras and monotone $*$ -algebras of bounded random variables. A Von Neumann algebra is the direct noncommutative analog of this idea; the only difference is that a Von Neumann algebra can contain elements that do not commute with each other. In fact, we will see in the next section that if a Von Neumann algebra happens to be commutative, then it is entirely equivalent to the algebra of bounded functions on some measure space. This provides the key link between quantum and classical probability theory and will be of central importance.

The state \mathbb{P} plays the role of a probability measure; indeed, if some element $X \in \mathcal{A}$ is self-adjoint, then it can be interpreted as an observable and $\mathbb{P}(X)$ is its expectation value. Once again, the classical analogy is clear: any probability measure \mathbf{P} induces a positive linear map on the monotone $*$ -algebra \mathcal{M} through the expectation map, and $\mathbf{E}(1) = 1$. Vice versa, any positive linear map on \mathcal{M} such that $\mathbf{E}(1) = 1$ can be used to define a finitely additive measure on $\sigma\{\mathcal{M}\}$ through $\mathbf{P}(A) = \mathbf{E}(I_A)$, $A \in \sigma\{\mathcal{M}\}$. In order to make sure that \mathbf{P} is countably additive (and hence a true probability measure), we have to impose an additional requirement: necessary and sufficient is the requirement that \mathbf{E} satisfies the monotone convergence theorem. But this corresponds precisely to a normal state on a Von Neumann algebra.

The following definition should not come as a great surprise.

DEFINITION B.1.5. A quantum probability space is a pair $(\mathcal{A}, \mathbb{P})$, where \mathcal{A} is a Von Neumann algebra (on a Hilbert space \mathbf{H}) and \mathbb{P} is a normal state on \mathcal{A} .

Most traditional quantum mechanics takes place on the space $(\mathcal{B}(\mathbf{H}), \text{Tr}[\rho \cdot])$. An important reason to introduce the more abstract notion of a quantum probability space is that we are then in the position to keep track of information in terms of subalgebras, and in particular through filtrations of Von Neumann algebras.

REMARK B.1.6. The Hilbert space \mathbf{H} is usually dropped in the notation when we are dealing with algebras. An underlying Hilbert space is always implied, and such a space is generally fixed at the outset.

Let $\mathcal{S} \subset \mathcal{B}(\mathbf{H})$ be any set of operators. Then we denote by $\text{vN}(\mathcal{S})$ the smallest Von Neumann algebra that contains \mathcal{S} , which is called the Von Neumann algebra generated by \mathcal{S} . There is a well-known algebraic characterization of this notion, see, e.g., [KR97a, Thm. 5.3.1]: a fundamental theorem of Von Neumann states

that $\text{vN}(\mathcal{S}) = (\mathcal{S} \cup \mathcal{S}^*)''$. Here we have written $\mathcal{S}^* = \{X \in \mathcal{B}(\mathbf{H}) : X^* \in \mathcal{S}\}$, and $\mathcal{S}' = \{X \in \mathcal{B}(\mathbf{H}) : XS = SX \ \forall S \in \mathcal{S}\}$. \mathcal{S}' is called the commutant of \mathcal{S} .

The reader is invited to skip the remainder of this section, moving on to section B.1.2 where we introduce the spectral theorem. The goal of the following discussion is to justify the definition of a Von Neumann algebra used above. The reader who is familiar with operator algebras will not have failed to notice that my definition is a little unusual; indeed, I have not been able to find it in the literature. Nonetheless, it is quite easily established to be equivalent to the usual approach, where a Von Neumann algebra is defined to be a $*$ -algebra that is closed in the strong operator topology (easily, that is, once we use a highly nontrivial result of R. V. Kadison). I find the definition above much more appealing from a probabilistic point of view, particularly considering the deep connection with the monotone class theorem.

Let us show that our definition reduces to the usual one. We need the following.

LEMMA B.1.7 (R. V. Kadison [Kad56]). *Let \mathcal{A} be a $*$ -subalgebra of $\mathcal{B}(\mathbf{H})$ that is closed under the uniform topology, and assume moreover that any upper bounded increasing net $\{A_\alpha\}$ of positive elements in \mathcal{A} converges strongly to an element $A \in \mathcal{A}$. Then \mathcal{A} is closed in the strong operator topology.*

The following result can be found, e.g., in [Mey93, page 247].

LEMMA B.1.8. *Let $\{A_\alpha\}$ be an upper bounded increasing net of positive elements in $\mathcal{B}(\mathbf{H})$. Then the strong limit of A_α and the least upper bound $\sup_\alpha A_\alpha$ coincide.*

The following argument follows closely the equivalent classical result, which can be used as a step in the proof of the monotone class theorem (see, e.g., [Fit05]).

LEMMA B.1.9. *Let $\mathcal{A} \subset \mathcal{B}(\mathbf{H})$ be a monotone $*$ -algebra in the sense defined above. Then \mathcal{A} is closed in the uniform topology.*

PROOF. As the uniform topology on $\mathcal{B}(\mathbf{H})$ is induced by the norm $\|\cdot\|$ (i.e., it is a metric topology), it is sufficient to consider sequences rather than nets—in particular, \mathcal{A} is closed iff every uniformly convergent sequence $\{X_n\} \subset \mathcal{A}$ has its limit in \mathcal{A} . Note that it suffices to prove this claim for sequences of self-adjoint operators Y_n . After all, suppose there exists a uniformly convergent sequence $\{X_n\} \subset \mathcal{A}$ whose limit point $X \notin \mathcal{A}$. Setting $Y_n = X_n + X_n^* \in \mathcal{A}$, we see that

$$\|X + X^* - Y_n\| \leq \|X - X_n\| + \|X^* - X_n^*\| = 2\|X - X_n\| \rightarrow 0,$$

so $Y_n \rightarrow X + X^*$. Similarly, $Z_n = i(X_n - X_n^*) \in \mathcal{A}$ converges to $i(X - X^*)$. But \mathcal{A} is a $*$ -algebra, so $X \notin \mathcal{A}$ implies that either $X + X^*$ or $i(X - X^*)$ cannot lie in \mathcal{A} . But then either $\{Y_n\}$ or $\{Z_n\}$ is a uniformly convergent sequence of self-adjoint elements of \mathcal{A} whose limit is not in \mathcal{A} . Hence if we require that this can not be the case, this guarantees that \mathcal{A} is closed in the uniform topology.

Let $\{Y_n\} \subset \mathcal{A}$ be a uniformly convergent sequence of self-adjoint operators, and denote its limit point by Y . We will prove that necessarily $Y \in \mathcal{A}$ by constructing a nondecreasing sequence $\{Y'_n\} \subset \mathcal{A}$ of nonnegative operators that converges to $Y + cI$, where $c \in \mathbb{R}$ is a constant. The monotonicity assumption then guarantees that $Y \in \mathcal{A}$, and the statement of the Lemma follows. To construct Y'_n , let us assume that $\|Y_{n+1} - Y_n\| \leq 2^{-n}$ for every n . If this is not the case, we can always choose a subsequence of $\{Y_n\}$ that does have this property and proceed with this subsequence. Now define $Y'_n = Y_n + (1 - 2^{-n+1} + \max_k \|Y_k\|)I$. Evidently $Y'_n \in \mathcal{A}$ and $Y'_n \geq 0$. Moreover, we find that for any vector $v \in \mathbf{H}$

$$\langle v, (Y'_{n+1} - Y'_n)v \rangle = \langle v, (Y_{n+1} - Y_n)v \rangle + 2^{-n}\|v\|^2 \geq (-\|Y_{n+1} - Y_n\| + 2^{-n})\|v\|^2 \geq 0,$$

so $Y'_{n+1} \geq Y'_n$. Finally, $Y'_n \rightarrow Y + (1 + \max_k \|Y_k\|)I$, and the result follows. \square

We can now conclude what we are trying to show.

LEMMA B.1.10. *Let $\mathcal{A} \subset \mathcal{B}(\mathbf{H})$ be a $*$ -algebra. The following are equivalent:*

- (1) \mathcal{A} is closed under the strong operator topology;
- (2) \mathcal{A} is monotone.

PROOF. Both directions follow immediately from the preceding Lemmas. \square

Finally, note that the definition which we have used for a normal state is also equivalent to the usual one: this is well known, see [Mey93, pp. 247–248].

REMARK B.1.11. If the underlying Hilbert space \mathbf{H} is separable, then the definitions of a monotone $*$ -algebra and of a normal state simplify significantly: in this case, it is sufficient to consider monotone increasing sequences rather than nets. See [Ped79, sec. 4.5.5] or [Thi03] for further details.

B.1.2. The spectral theorem. Classical probability spaces can be introduced as a special case of quantum probability spaces. To see this, consider the probability space $(\Omega, \mathcal{F}, \mathbf{P})$. The Banach space $\mathcal{A} = L^\infty(\Omega, \mathcal{F}, \mathbf{P})$ is clearly a monotone $*$ -algebra. To interpret it as a Von Neumann algebra, however, we need to introduce an underlying Hilbert space on which these functions act as operators. To this end, let us set $\mathbf{H} = L^2(\Omega, \mathcal{F}, \mathbf{P})$, and we define the action of \mathcal{A} on \mathbf{H} by pointwise multiplication: i.e., we define $(X\psi)(\omega) = X(\omega)\psi(\omega)$ for all $X \in \mathcal{A}$, $\psi \in \mathbf{H}$. Then \mathcal{A} is a Von Neumann algebra on the Hilbert space \mathbf{H} . Similarly, \mathbf{P} defines a normal state on \mathcal{A} through the expectation map \mathbf{E} , and we have constructed a quantum probability space $(\mathcal{A}, \mathbf{E})$ which is completely equivalent to a classical probability model (up to null sets).

REMARK B.1.12. We will always denote by L^∞ the space of bounded measurable *complex* functions (up to a.s. equivalence), and similarly we will take L^2 to be a complex Hilbert space. This means that in the classical case we will trivially deal with complex random variables. Note that *real* random variables correspond precisely to self-adjoint elements in the algebra (i.e., these are the observables).

To get some intuition, consider the case where Ω is an n -point set, $n < \infty$. Then $\mathbf{H} \simeq \mathbb{C}^n$, and \mathcal{A} above is represented as the set of diagonal $n \times n$ matrices. The defining characteristic of \mathcal{A} is that it is commutative; indeed, the fact that the corresponding matrices are diagonal is only visible in one particular choice of basis (the canonical basis of \mathbb{C}^n). It is an elementary fact of linear algebra that for any commutative Von Neumann algebra of $n \times n$ matrices, we can choose a basis in \mathbb{C}^n such that all these matrices are simultaneously diagonalized in that basis. Hence any commutative Von Neumann algebra looks like a classical probability space, provided that we look at it in the appropriate basis. Our goal is to extend this idea to the infinite-dimensional case.

We state the following Theorem for separable Hilbert spaces; we will always use only separable spaces, so this is not a restriction. The spectral theory can also be developed for nonseparable spaces, see, e.g., [Sak98, Prop. 1.18.1].

THEOREM B.1.13. *Let \mathbf{H} be separable, and let \mathcal{A} be a commutative Von Neumann algebra on \mathbf{H} . Then there exists a finite measure space $(\Omega, \mathcal{G}, \mu)$, a unitary operator $U : \mathbf{H} \rightarrow L^2(\Omega, \mathcal{G}, \mu)$, and a σ -algebra $\mathcal{F} \subset \mathcal{G}$, such that $U\mathcal{A}U^* = L^\infty(\Omega, \mathcal{F}, \mu)$, where the latter acts on $L^2(\Omega, \mathcal{G}, \mu)$ by pointwise multiplication.*

OUTLINE OF PROOF. We will not prove this deep result here; however, let us give an indication how it can be obtained from standard functional analytic results.

It is possible to establish the following: if \mathbf{H} is separable and \mathcal{A} is commutative, then there exists a single self-adjoint element $A \in \mathcal{A}$ such that $\mathcal{A} = \text{vN}(A)$. This was already established by J. Von Neumann [Von29, Thm. 10], see also [Tak02, Thm. III.1.21]. By a well-known form of the spectral theorem of functional analysis [RS72, page 227], there exists a finite measure space $(\Omega, \mathcal{G}, \mu)$, a bounded measurable function a on Ω , and a unitary map $U : \mathbf{H} \rightarrow L^2(\Omega, \mathcal{G}, \mu)$, such that $(UAU^*\psi)(\omega) = a(\omega)\psi(\omega)$ for any $\psi \in L^2(\Omega, \mathcal{G}, \mu)$.

Now define $\mathcal{C} = \{f(A) : f \text{ Borel}\}$, and note that \mathcal{C} is clearly a Von Neumann algebra and \mathcal{C} contains A . Hence $\mathcal{A} \subset \mathcal{C}$. Furthermore, if we introduce $\mathcal{F} = \sigma\{a\}$, then we can define a bijection $\iota : \mathcal{C} \rightarrow L^\infty(\Omega, \mathcal{F}, \mu)$ by $\iota : X \mapsto UXU^*$; moreover ι is a *-isomorphism (i.e., it preserves the *-algebraic structure) and is order preserving. We now argue that $\mathcal{C} = \mathcal{A}$. If this were not the case, then \mathcal{A} would be a Von Neumann algebra that is strictly contained in \mathcal{C} , so that $\mathcal{M} = \iota(\mathcal{A})$ would be a monotone class that is strictly contained in $L^\infty(\Omega, \mathcal{F}, \mu)$. But $a \in \mathcal{M}$ so $\sigma\{\mathcal{M}\} = \mathcal{F}$, and the monotone class theorem gives the desired contradiction. \square

The objects defined in this Theorem are not entirely obvious: the σ -algebra \mathcal{G} and the measure μ will not play a physical role at the end of the day (they carry no probabilistic content). The reason we need these objects is analytic in nature. It is necessary to use \mathcal{G} rather than \mathcal{F} so that the Hilbert space $L^2(\Omega, \mathcal{G}, \mu)$ will be “large enough” to construct the unitary $U : \mathbf{H} \rightarrow L^2(\Omega, \mathcal{G}, \mu)$.

The measure μ also has a different task. Recall that L^∞ only defines functions up to a.s. equivalence. The role of μ is to define which elements of \mathcal{G} are null sets. Note that our state \mathbb{P} (which did not play a role in the above Theorem) could well be defined in such a way that $\mathbb{P}(P) = 0$ for some projection $P \in \mathcal{A}$; this would mean that $\iota(P)$ is an indicator function of a set of physical probability zero. If we have used the measure induced by \mathbb{P} to define the L^∞ space, then the function $\iota(P)$ would belong to the same equivalence class of L^∞ as the zero function. This would prohibit ι from being an isomorphism. This cannot happen in the current setting ($\iota(P)$ must have finite μ -expectation for any projector $P \in \mathcal{A}$), but nothing prohibits us from defining a new measure $\mathbf{P} \ll \mu$ on (Ω, \mathcal{F}) which coincides with the physical state \mathbb{P} . This is precisely what we will do.

COROLLARY B.1.14. *Let $(\mathcal{A}, \mathbb{P})$ be a commutative quantum probability space (on a separable Hilbert space \mathbf{H}), i.e., \mathcal{A} is a commutative Von Neumann algebra. Then there exists a measure space $(\Omega, \mathcal{F}, \mu)$, a probability measure $\mathbf{P} \ll \mu$, and a *-isomorphism $\iota : \mathcal{A} \rightarrow L^\infty(\Omega, \mathcal{F}, \mu)$, such that $\mathbb{P}(X) = \mathbf{E}_{\mathbf{P}}(\iota(X))$ for all $X \in \mathcal{A}$.*

The proof of this statement, which we will generally refer to as the *spectral theorem*, is evident from the discussions above. This is a key theorem in quantum probability, which establishes the equivalence between commutative quantum probability spaces and classical probability spaces. The spectral theorem is also central to the interpretation of quantum mechanics. Typically, quantum mechanical models will admit many observables that do not commute (the standard example being the position and momentum of a free particle). In any single realization of an experiment we can only choose to measure a commuting set \mathcal{S} of observables; these generate a commutative Von Neumann algebra $\text{vN}(\mathcal{S})$, and the spectral theorem furnishes a full-blown probabilistic interpretation.

In a different realization we may choose to measure a different set of commuting observables (by using a different measurement apparatus), and we have to reapply the spectral theorem in order to make predictions within the new setup. We can never measure two noncommuting observables in the same realization, so that there is never any need to represent them simultaneously as random variables on some classical probability space (which the spectral theorem can not do). Noncommuting observables are called incompatible, and their joint statistics are undefined.

REMARK B.1.15. When using the spectral theorem in practice, there can be a lot of switching back and forth between commutative subalgebras of some larger noncommutative quantum probability space and the corresponding classical probability models. We will often be a little sloppy in applying the theorem, and simply use the magic symbol ι without being careful to specify which commutative algebra the spectral theorem is being applied to. In most cases this is immediately clear from the context; in particular, if we apply ι to several commuting observables or sets of observables, and then proceed to manipulate the corresponding classical objects simultaneously, it is understood that the spectral theorem is being applied to a larger commutative algebra that contains all the relevant objects. The significant improvement in readability justifies this little flexibility of notation.

B.1.3. Unbounded observables. Up to this point we have only worked with bounded operators. It would be extremely nice if we could keep it this way—unbounded operators are a pain in the neck, and bring with them a host of unpleasant domain problems. The extension to unbounded operators is necessary, unfortunately, as we will routinely encounter unbounded observables in physical models. Let us thus briefly introduce some of the necessary concepts.

Recall that an unbounded operator X can not be defined on the entire Hilbert space; at best, it can be defined on some dense linear set $D(X) \subset H$, i.e., $X : D(X) \rightarrow H$. $D(X)$ is called the domain of X . When we define an unbounded operator we should specify its domain; the same operator may have very different properties if its domain is changed. An operator X is called closed if the set $\{(\psi, X\psi) : \psi \in D(X)\} \subset H \times H$ is closed (in the norm topology). X is called closable if it can be made closed by enlarging its domain, and the closure of a closable operator is obtained by choosing the smallest such extension of its domain. The adjoint X^* of a densely defined operator X is defined by $\langle X\psi, \phi \rangle = \langle \psi, X^*\phi \rangle \forall \psi \in D(X)$ for those $\phi \in H$ for which this definition makes sense, and the set of all such ϕ is taken to be the domain $D(X^*)$. X is called self-adjoint if $X = X^*$ (by which we mean $D(X) = D(X^*)$, and $X\psi = X^*\psi$ for all $\psi \in D(X)$). X is called essentially self-adjoint if it is closable and its closure is self-adjoint. All these definitions are standard, see [RS72, Ch. VIII] for this and much more.

The notion of self-adjointness is extremely important: it is only to self-adjoint operators, in the very strict sense described above, that the spectral theorem can be applied. A particular representation of the spectral theorem will be useful to us. This states that for any self-adjoint operator X , there exists a spectral measure E_X on \mathbb{R} (i.e., for any Borel set A , $E_X(A)$ is a projection operator; $E_X(\emptyset) = 0$, $E_X(\mathbb{R}) = I$; for a countable sequence of disjoint sets A_n with union A , $E_X(A) = s\text{-}\lim_{k \rightarrow \infty} \sum_{n=1}^k E_X(A_k)$; and $E_X(A_1)E_X(A_2) = E_X(A_1 \cap A_2)$), such that

$$X = \int_{\mathbb{R}} \lambda E_X(d\lambda), \quad \text{meaning that } \langle \psi, X\psi \rangle = \int_{\mathbb{R}} \lambda \langle \psi, E_X(d\lambda)\psi \rangle \quad \forall \psi \in D(X).$$

The latter integral is meant in the Lebesgue sense, and $D(X)$ actually coincides with the set of $\psi \in \mathbf{H}$ for which the integral of λ^2 with respect to $\langle \psi, E_X(\cdot)\psi \rangle$ is finite. The spectral measure E_X has an important probabilistic interpretation: $\iota(E_X(A)) = \{\iota(X) \in A\}$, i.e., the projector $E_X(A)$ corresponds to the event $X \in A$, and $\mathbb{P}(E_X(A))$ is the probability that measurement of X will return a value in A . Note that we should extend the state \mathbb{P} to self-adjoint operators as follows:

$$\mathbb{P}(X) = \int_{\mathbb{R}} \lambda \mathbb{P}(E_X(d\lambda)) \quad \text{in the Lebesgue sense;}$$

after all, $\mathbb{P}(E_X(\cdot))$ is evidently a probability measure on \mathbb{R} (the law of X).

Let us now consider unbounded observables in the setting of quantum probability spaces. Our Von Neumann algebra \mathcal{A} only holds bounded elements; we would like to introduce a notion that parallels the idea of a random variable being measurable with respect to a σ -algebra. The following definition is completely natural in this context, and indeed reduces to the notion of measurability if $\mathcal{A} = L^\infty(\Omega, \mathcal{F}, \mathbf{P})$.

DEFINITION B.1.16. A (not necessarily bounded) self-adjoint operator X is said to be affiliated to a Von Neumann algebra \mathcal{A} , denoted by $X \eta \mathcal{A}$, if its spectral measure E_X satisfies $E_X(A) \in \mathcal{A}$ for any Borel set A of \mathbb{R} .

That is, X is affiliated to \mathcal{A} if the corresponding events are in \mathcal{A} . We mention that this probabilistic definition is equivalent to a more algebraic definition that is usually preferred by operator algebraists; see [Mey93, page 245].

We can now define the Von Neumann algebra generated by a self-adjoint element X : we set $\text{vN}(X) = \text{vN}\{E_X(A) : A \text{ Borel set}\}$. This definition coincides with the usual definition for bounded X , and it is trivially the case that $X \eta \text{vN}(X)$.

Suppose that $X \eta \mathcal{C}$, where \mathcal{C} is a commutative Von Neumann algebra. Then we can apply the spectral theorem to \mathcal{C} to get a $*$ -isomorphism ι with some classical probability space. We would like to extend ι also to affiliated observables; that is, we would like to define $\iota(X)$ as an unbounded random variable on the aforementioned probability space. This is not difficult to do, as it can be shown that the unitary transformation U of Theorem B.1.3 also diagonalizes affiliated X . Indeed, this follows from standard functional analysis arguments [RS72, Ch. VIII], once we note that $X \eta \mathcal{C}$ implies that $X = f(C)$, where f is an unbounded function and $\mathcal{C} = \text{vN}(C)$ (the existence of such a C was discussed in the proof of Theorem B.1.3).

The real unpleasantness of unbounded operators emerges when we try to sum or multiply them. Consider two self-adjoint operators X and Y . It is not at all clear that $X + Y$ is self-adjoint or even well defined, let alone that XY or YX are well defined. The domains of these operators may not even be dense; at the very least, it is unlikely that these operators would still be closed. In fact, it is very difficult in general to manipulate unbounded operators, and intricate domain problems crop up in unexpected places. This is perhaps understandable when we try to manipulate noncommuting operators—when X and Y do not commute (i.e., their spectral measures do not commute), the sum or products of X and Y do not necessarily have a useful physical interpretation. On the other hand, when X and Y commute we can represent them simultaneously as classical random variables $\iota(X)$, $\iota(Y)$ on some probability space, where we can meaningfully add and multiply at will. Hence, somehow, the domain problems should “work out” in this case.

Things do indeed work out, but we have to take some care; see [KR97a, pp. 351–356]. Let \mathcal{A} be a commutative Von Neumann algebra, and let X and Y be

self-adjoint operators affiliated to \mathcal{A} . It turns out that $X + Y$ is neither closed nor self-adjoint, but it is essentially self-adjoint. In particular, if we introduce the operation $\hat{+}$ to mean $X \hat{+} Y = \overline{X + Y}$ (\overline{X} is the closure of X), then $X \hat{+} Y$ is self-adjoint for any pair of self-adjoint operators $X, Y \eta \mathcal{A}$. Similarly XY is essentially self-adjoint, and if we define $X \hat{\cdot} Y = \overline{XY}$, then $X \hat{\cdot} Y$ defines a self-adjoint operator. Let us summarize these statements as follows.

LEMMA B.1.17. *Let \mathcal{A} be a commutative Von Neumann algebra. Then the set $\mathcal{S}(\mathcal{A})$ of self-adjoint operators $X \eta \mathcal{A}$ forms a real algebra under the addition $\hat{+}$ and the multiplication $\hat{\cdot}$. Moreover, the *-isomorphism $\iota : \mathcal{A} \rightarrow L^\infty(\Omega, \mathcal{F}, \mu)$, obtained by applying the spectral theorem to \mathcal{A} , extends to an isomorphism between $\mathcal{S}(\mathcal{A})$ and the set of μ -a.s. finite \mathcal{F} -measurable random variables on Ω .*

This result can be extended to a slightly larger class of unbounded operators. A closed, but not necessarily self-adjoint, operator X is said to be normal if $X \hat{+} X^*$ and $i(X^* \hat{-} X)$ are self-adjoint and commute with each other. A normal operator X is said to be affiliated to \mathcal{A} if $X \hat{+} X^*$ and $i(X^* \hat{-} X)$ are affiliated to \mathcal{A} , and the set of all normal operators affiliated to a commutative algebra \mathcal{A} is denoted by $\mathcal{N}(\mathcal{A})$. The previous Lemma now extends as follows: if \mathcal{A} is commutative, then $\iota : \mathcal{A} \rightarrow L^\infty(\Omega, \mathcal{F}, \mu)$ extends to a *-isomorphism between $\mathcal{N}(\mathcal{A})$ and the set of μ -a.s. finite \mathcal{F} -measurable complex random variables on Ω (up to μ -a.s. equivalence).

To recap: we can add and multiply any bounded operators at will, and we can essentially add and multiply commuting unbounded operators at will. In any other case, utmost care should be exercised.

B.2. Quantum noise

B.2.1. Fock space and the fundamental processes. In the following sections we are going to introduce a quantum probability model that is widely used in physics, in particular in quantum optics, for modelling the interaction of quantum systems (e.g., atoms, semiconductor quantum dots, optical cavities) with quantum fields (e.g., the electromagnetic field). The model is Markov in a certain sense (which we will not emphasize), and admits a reasonable stochastic calculus which is comparable to the Itô calculus (though much less powerful, as we will see). We begin in this section by introducing the field part of the picture, and the corresponding quantum noises.

A basic building block in the theory is the (symmetric) Fock space over \mathbf{H}_1

$$\Gamma(\mathbf{H}_1) = \mathbb{C} \oplus \bigoplus_{n=1}^{\infty} \mathbf{H}_1^{\odot n}.$$

Here \mathbf{H}_1 is called the single-particle Hilbert space, and \odot denotes the symmetrized tensor product. Note that $\Gamma(\mathbf{H}_1)$ is a separable Hilbert space if \mathbf{H}_1 is separable; we will always use only separable single-particle spaces.

The operation $\Gamma(\cdot)$ behaves much like an exponential map for Hilbert spaces; in particular, there is the following natural isomorphism:¹

$$\Gamma(\mathbf{H}_1 \oplus \mathbf{H}'_1) \simeq \Gamma(\mathbf{H}_1) \otimes \Gamma(\mathbf{H}'_1).$$

¹ This (unitary) isomorphism is constructed in such a way that the exponential vectors defined below have the same property; see, e.g., [Par92, Prop. 19.6] for details.

In our application of the theory we will mainly use the spaces $\Gamma_{[s,t]} = \Gamma(L^2([s,t]))$, where we write $\Gamma_t = \Gamma_{[0,t]}$, $\Gamma_{[t]} = \Gamma_{[t,\infty)}$, and $\Gamma = \Gamma_{[0]}$. Then we have

$$\Gamma \simeq \Gamma_s \otimes \Gamma_{[s,t]} \otimes \Gamma_t.$$

This property is called the continuous tensor product structure of the Fock space. We note that there is a natural isomorphism between Γ and $L^2(W, \mu)$, where (W, μ) is the canonical Wiener space, through Wiener chaos expansion. There is also a natural isomorphism between Γ and $L^2(P, \mu)$, where (P, μ) is the probability space of a canonical Poisson process. We will not use these isomorphisms, but we will shortly see that Wiener and Poisson processes emerge naturally within this model.

A particularly useful set of vectors in $\Gamma(\mathbf{H}_1)$ is the set of exponential vectors. For any $f \in \mathbf{H}_1$, we define the exponential vector $e(f) \in \Gamma(\mathbf{H}_1)$ by

$$e(f) = 1 \oplus \bigoplus_{n=1}^{\infty} \frac{f^{\otimes n}}{\sqrt{n!}}.$$

Evidently, we have the following relation:

$$\langle e(g), e(f) \rangle_{\Gamma(\mathbf{H}_1)} = \exp(\langle g, f \rangle_{\mathbf{H}_1}).$$

(In the future, we will refrain from labeling the inner product when no confusion can occur.) Denote by $\mathbf{E}(\mathbf{H}_1) = \text{span}\{e(f) : f \in \mathbf{H}_1\}$ the linear span of all exponential vectors (and set $\mathbf{E} = \mathbf{E}(L^2(\mathbb{R}_+))$, $\mathbf{E}_{[s,t]} = \mathbf{E}(L^2([s,t]))$, etc.) Then $\mathbf{E}(\mathbf{H}_1)$ forms a dense linear manifold in $\Gamma(\mathbf{H}_1)$, and is called the exponential domain.

Exponential vectors have an important property: they respect the continuous tensor product structure of the Fock space. In particular,

$$e(f \oplus g) = e(f) \otimes e(g) \quad \text{in} \quad \Gamma(\mathbf{H}_1 \oplus \mathbf{H}'_1) \simeq \Gamma(\mathbf{H}_1) \otimes \Gamma(\mathbf{H}'_1).$$

For $f \in L^2(\mathbb{R}_+)$, we will denote its restriction to $L^2([s,t])$ by $f_{[s,t]}$, and f_s, f_t are defined similarly. Then $e(f) \in \Gamma$ satisfies $e(f) = e(f_s) \otimes e(f_{[s,t]}) \otimes e(f_t)$.

We are now going to define some observables. To do this, we need to have a way of generating self-adjoint operators. This is usually not a trivial task, considering the domain problems involved. The technique which we will use is the definition of self-adjoint operators through Stone's theorem [RS72, Thm. VIII.8]: if U_t is a strongly continuous one-parameter unitary group on a Hilbert space \mathbf{H} , then there is a unique self-adjoint operator A on \mathbf{H} , called the infinitesimal generator of U_t , such that $U_t = \exp(itA)$. Unitary operators are bounded and are thus much easier to deal with; subsequently Stone's theorem takes care of the immediate domain issues in the definition of the self-adjoint operator A .

First, consider the "translation group" on \mathbf{H}_1 . This group can be lifted to the level of $\Gamma(\mathbf{H}_1)$ by introducing the Weyl operator

$$W(g)e(f) = \exp\left(-\langle g, f \rangle - \frac{1}{2}\|g\|^2\right)e(f+g), \quad f, g \in \mathbf{H}_1.$$

The prefactor ensures that $W(g)$ is an isometry. We have only specified the action of $W(g)$ on the exponential domain $\mathbf{E}(\mathbf{H}_1)$, but as $W(g)$ is an isometry it can be uniquely extended to a unitary operator on the entire Fock space $\Gamma(\mathbf{H}_1)$. The notation $W(g)$ denotes this unitary operator. Note that we have the Weyl relation

$$W(f)W(g) = \exp(-i \text{Im}\langle f, g \rangle)W(f+g), \quad f, g \in \mathbf{H}_1.$$

Hence evidently $W(tf)$, $t \in \mathbb{R}$ forms a one-parameter unitary group. Strong continuity follows from the fact that $\langle \psi, W(tf)\phi \rangle$ is a measurable function of t for any $\psi, \phi \in \Gamma(\mathbf{H}_1)$ (this is easily verified for $\psi, \phi \in \mathbf{E}(\mathbf{H}_1)$, and follows for general ψ, ϕ by taking limits), together with [RS72, Thm. VIII.9]. Hence by Stone's theorem, there exists for any $f \in \mathbf{H}_1$ a self-adjoint operator $B(f)$ such that $W(tf) = \exp(itB(f))$. Note that clearly $B(f)$ and $B(g)$ commute (i.e., their spectral measures commute) if and only if $\langle f, g \rangle$ is real, as then $W(f)$ and $W(g)$ commute.

A different way to generate a one-parameter unitary group on $\Gamma(\mathbf{H}_1)$ is by lifting a one-parameter unitary group on \mathbf{H}_1 . Define

$$\Gamma(U) e(f) = e(Uf), \quad f \in \mathbf{H}_1, U \in \mathcal{B}(\mathbf{H}_1).$$

$\Gamma(U)$ is called the second quantization of U . Now let U_t be a one-parameter strongly continuous unitary group on \mathbf{H}_1 with infinitesimal generator A ; then $\Gamma(U_t)$ defines a one-parameter strongly continuous unitary group on $\Gamma(\mathbf{H}_1)$ (once its domain is extended from $\mathbf{E}(\mathbf{H}_1)$ to all of $\Gamma(\mathbf{H}_1)$). Hence by Stone's theorem, there exists a self-adjoint operator $\Lambda(A)$ such that $\Gamma(U_t) = \exp(it\Lambda(A))$.

Having generated a bunch of abstract nonsense, let us turn to the probabilistic side of things. Let us begin by introducing some notation.

DEFINITION B.2.1. The field quantum probability space (\mathcal{W}, φ) is defined by $\mathcal{W} = \mathcal{B}(\Gamma)$ and $\varphi(X) = \langle \Phi, X\Phi \rangle$, where $\Phi = e(0)$ is called the vacuum vector and φ is called the vacuum state.

DEFINITION B.2.2. We denote by $\mathcal{W}_{[s,t]} = \mathcal{B}(\Gamma_{[s,t]})$, $\mathcal{W}_t = \mathcal{W}_{[0,t]}$, and $\mathcal{W}_{[t, \infty)} = \mathcal{W}_{[t, \infty)}$, so that $\mathcal{W} = \mathcal{W}_s \otimes \mathcal{W}_{[s,t]} \otimes \mathcal{W}_t$. A stochastic process is a family of self-adjoint operators $\{X_t\}_{t \in \mathbb{R}_+}$. A stochastic process is called adapted if $X_t \eta \mathcal{W}_t$ for every t , where the filtration \mathcal{W}_t is defined as $\mathcal{W}_t = \vee \mathbf{N}\{X \otimes I : X \in \mathcal{W}_t\} \subset \mathcal{W}$.

DEFINITION B.2.3. A classical or commutative stochastic process is an adapted process $\{X_t\}_{t \in \mathbb{R}_+}$ such that $\vee \mathbf{N}\{X_t : t \in \mathbb{R}_+\}$ is a commutative Von Neumann algebra. Such a process generates a filtration $\mathcal{X}_t = \vee \mathbf{N}\{X_s : 0 \leq s \leq t\} \subset \mathcal{W}_t$.

A commutative process X_t becomes a stochastic process in the classical sense $x_t = \iota(X_t)$ when the spectral theorem is applied to the Von Neumann algebra generated by $\{X_t\}_{t \in \mathbb{R}_+}$. Applying the spectral theorem to the filtration of subalgebras, $\iota(\mathcal{X}_t) = L^\infty(\mathcal{F}_t)$, gives rise to the classical filtration \mathcal{F}_t generated by x_t .

Let us now investigate some interesting stochastic processes in (\mathcal{W}, φ) .

DEFINITION B.2.4. Define the following commutative stochastic processes: the quadratures $Q_t = B(iI_{[0,t]})$, $P_t = B(-I_{[0,t]})$, and the gauge process $\Lambda_t = \Lambda(M_{[0,t]})$, where $M_{[0,t]} : L^2(\mathbb{R}_+) \rightarrow L^2(\mathbb{R}_+)$ is defined by $M_{[0,t]}f = I_{[0,t]}f$. Define also the annihilation process $A_t = (Q_t + iP_t)/2$ and the creation process $A_t^\dagger = (Q_t - iP_t)/2$. The operators A_t, A_t^\dagger and Λ_t (or Q_t, P_t and Λ_t) are called the fundamental noises.

The fact that Q_t, P_t and Λ_t are adapted and commutative is easily verified. Beware, however, that these processes do not commute with each other.

REMARK B.2.5. A_t and A_t^\dagger are not self-adjoint; we use these chiefly for historical reasons (quantum stochastic calculus is defined in terms of A_t and A_t^\dagger , rather than Q_t and P_t). It can be shown that A_t, A_t^\dagger and Λ_t can all be restricted to an invariant domain, which includes the exponential domain \mathbf{E} as a subset, such that $A_t + A_t^\dagger, i(A_t^\dagger - A_t)$ and Λ_t are essentially self-adjoint on that domain.

We can finally give a concrete probabilistic result.

LEMMA B.2.6. $\iota(Q_t)$ and $\iota(P_t)$ both define a Wiener process in the vacuum state φ (but not on the same probability space, as Q_t and P_t do not commute).

PROOF. Let us prove this for Q_t ; the proof for P_t is identical. It suffices to prove that $q_t = \iota(Q_t)$ has independent increments, and that $q_t - q_s$ is a Gaussian random variable with mean zero and variance $t - s$, under the state \mathbf{P} obtained from φ through the spectral theorem. To this end, let us calculate the joint characteristic function of two nonoverlapping increments $q_t - q_s$ and $q_v - q_u$, $u \leq v \leq s \leq t$:

$$\begin{aligned} \xi(\alpha, \beta) &= \mathbf{P}(e^{i\alpha(q_t - q_s) + i\beta(q_v - q_u)}) = \langle \Phi, W(i\alpha I_{[s,t]} + i\beta I_{[u,v]}) \Phi \rangle \\ &= \exp\left(-\frac{1}{2}\|\alpha I_{[s,t]} + \beta I_{[u,v]}\|^2\right) = \exp\left(-\frac{1}{2}\alpha^2(t-s) - \frac{1}{2}\beta^2(v-u)\right). \end{aligned}$$

Hence $q_t - q_s$ and $q_v - q_u$ are independent (as the characteristic function factorizes), $q_t - q_s$ is Gaussian distributed with mean zero and variance $t - s$, and $q_v - q_u$ is Gaussian distributed with mean zero and variance $v - u$. \square

REMARK B.2.7. Technically a Wiener process should have continuous sample paths. In the current context this does not make any sense: it is easy to break continuity of the sample paths by changing $\iota(X_t)$, for each t , only on a set of measure zero. Such properties are discarded when we consider a classical stochastic process as a map $t \rightarrow L^0(\Omega, \mathcal{F}, \mu)$ (where L^0 is the set of a.s. finite random variables, up to a.s. equivalence). Hence in the quantum setting, we should generally consider stochastic processes up to modification. Of course, we are always free to choose a continuous modification of these Wiener processes (by the Kolmogorov-Čentsov theorem) after the spectral theorem has been applied.

The process $\iota(\Lambda_t)$ defines a rather boring process under the vacuum state: it is a.s. zero (this follows immediately from $\Gamma(U)\Phi = \Phi$). This process is more interesting in a coherent state $\varphi_f(X) = e^{-\|f\|^2} \langle e(f), X e(f) \rangle$, where Λ_t can be shown to be a Poisson process with time-dependent intensity $|f(t)|^2$ by proceeding in the same way as in the previous proof. In quantum optics, Q_t and P_t can be observed by using a homodyne detector to measure the vacuum (Q_t and P_t correspond to orthogonal quadratures of the field), while Λ_t can be observed using a photon counter. In this thesis we are chiefly interested in the former type of detection, though photodetection is easily treated using the same techniques.

B.2.2. Quantum stochastic calculus on the exponential domain. We have seen that there are at least three interesting processes in our field probability space—the Wiener processes Q_t , P_t and the Poisson process Λ_t . It is now natural to ask, following the classical case, whether we can build stochastic integrals based on these processes, which would in turn allow us to develop a whole class of interesting processes through stochastic differential equations. The motivation for this stems from the fact that we would like to model physical systems that interact with the field—i.e., we need to introduce Schrödinger equations which are driven by the noises Q_t , P_t and Λ_t . As these noises are singular, these quantum stochastic Schrödinger equations must be defined as quantum stochastic differential equations (QSDEs). In this section we will develop the corresponding integration theory, chiefly following [Hud03, HP84]. Section B.2.3 is devoted to QSDEs.

The difficulty in defining quantum stochastic integrals stems from the unboundedness of the fundamental noises A_t , A_t^\dagger and Λ_t (we switch to using A_t and A_t^\dagger for

historical reasons). As we will see, it is not difficult to construct stochastic integrals for simple processes, provided that they are adapted. However, we are faced at this point with the problem of extending the integral to general adapted integrands through a suitable limiting procedure. Taking limits of unbounded operators is a touchy business—how can such a limit be meaningfully defined when the various terms in the sequence may have different domains? Can we guarantee that the limit operator is even well defined on a suitable domain, let alone that it is closable or even self-adjoint? A general discussion of these questions appears, e.g., in [RS72, sec. VIII.7]. We are going to be faced with even more unpleasant problems, as a useful theory of stochastic integrals should provide a stochastic calculus that allows us to manipulate such integrals. It is clear that any quantum stochastic calculus is going to be riddled with domain problems.

To resolve these issues and develop a viable stochastic calculus, R. L. Hudson and K. R. Parthasarathy [HP84] adopted a very simple policy: they simply fix a nice domain of exponential vectors, and define all the operators involved (integrands, integrators, integrals) only on this domain. This allows one to take strong limits on this fixed domain; i.e., if $\{A_n\}$ is a sequence of operators on a fixed domain D , and A is an operator on the same domain, then we say that $A_n \rightarrow A$ strongly on D if $\|(A_n - A)\psi\| \rightarrow 0$ for all $\psi \in D$. The theory now becomes fairly transparent. On the other hand, important information is lost in this process. For example, suppose we would like to generate a commutative stochastic process as the solution of a QSDE. Using the Hudson-Parthasarathy theory, we can at most obtain operators that are symmetric with respect to the domain D ; it is almost always unclear, however, whether or not such operators are actually essentially self-adjoint. Hence we can hardly expect to be able to apply the spectral theorem directly to processes obtained from the Hudson-Parthasarathy theory.

It would thus appear that the Hudson-Parthasarathy theory is not very useful; what is the point in defining operators on a fixed domain, if this prohibits us from interpreting them as observables? Things are not as bad as they seem, however. It is very well possible to obtain processes from the Hudson-Parthasarathy theory that are bounded. If an operator is bounded on the dense domain D , then it is uniquely extended to a bounded operator on all of Γ . There are no domain problems for such operators, and the Hudson-Parthasarathy theory is very successful in defining and manipulating these processes. In particular, we are mainly interested in using the Hudson-Parthasarathy theory to define quantum stochastic Schrödinger equations, whose solutions are unitary and hence necessarily bounded. Hence the theory we are about to present is extremely useful, but should be handled with care.

As we are generally interested in coupling a quantum system to the field, let us introduce the full model now (previously we only considered the field).

DEFINITION B.2.8. Let \mathfrak{h} denote the initial system Hilbert space, set $\mathcal{B} = \mathcal{B}(\mathfrak{h})$, and let ρ be a state on \mathcal{B} (the initial state). The standard quantum probability space for use in quantum stochastic analysis is defined as $(\mathcal{A}, \mathbb{P})$, where $\mathcal{A} = \mathcal{B} \otimes \mathcal{W}$ is a Von Neumann algebra on $\mathfrak{H} = \mathfrak{h} \otimes \Gamma$, and $\mathbb{P} = \rho \otimes \varphi$.

The terminology “initial system” will be clarified in section B.2.3.

Let us now introduce a suitable domain of exponential vectors.

DEFINITION B.2.9. The restricted exponential domain is defined by $D = \{e(f) : f \in L^2(\mathbb{R}_+) \cap L^{\infty, \text{loc}}(\mathbb{R}_+)\} \subset E$, i.e., D is the set of $e(f)$ with locally bounded f .

\mathbf{D} is a dense domain in Γ , and we will define the Hudson-Parthasarathy theory exclusively on the dense domain $\mathbf{h} \otimes \mathbf{D}$ (\otimes denotes the algebraic tensor product). The reason for using \mathbf{D} rather than \mathbf{E} is the difficulty in defining integrals with respect to Λ_t for arbitrary $\psi \in \mathbf{E}$; a heuristic explanation can be sought in the fact that under the state φ_f with $f \notin L^{\infty, \text{loc}}(\mathbb{R}_+)$, the intensity of the Poisson process Λ_t blows up at a finite time. This could make it difficult to control the integrals.

DEFINITION B.2.10. An operator X on $\mathbf{h} \otimes \Gamma$ is called allowable if $\mathbf{D}(X) = \mathbf{h} \otimes \mathbf{D}$ and $\mathbf{D}(X^*) \supset \mathbf{D}(X)$. For an allowable operator X , the operator X^\dagger is defined as the restriction of X^* to $\mathbf{h} \otimes \mathbf{D}$.

Note that, as mentioned before, $X = X^\dagger$ does not imply that X is essentially self-adjoint. However, allowable operators admit at least a little regularity.

LEMMA B.2.11. *Any allowable X is closable, X^\dagger is allowable, and $X^{\dagger\dagger} = X$.*

PROOF. The fact that X is closable follows from the fact that X^* has a dense domain, and moreover this means that $(X^*)^* = \overline{X}$ (see [RS72, Thm. VIII.1]). But $X^\dagger \subset X^*$ (X^* is an extension of X^\dagger), so $(X^\dagger)^* \supset (X^*)^* \supset X$. \square

DEFINITION B.2.12. An admissible process is a family $\{X_t\}_{t \in \mathbb{R}_+}$ of allowable operators. An admissible process is called adapted if there exists an operator $X_{[t]}$ on $\mathbf{h} \otimes \mathbf{D}_{[t]}$ such that $X_t \psi \otimes e(f) = (X_{[t]} \psi \otimes e(f_{[t]})) \otimes e(f_{[t]})$ for any $\psi \in \mathbf{h}$, $f \in L^2 \cap L^{\infty, \text{loc}}$.

Note that the fundamental noises A_t , A_t^\dagger , and Λ_t are admissible adapted processes once we restrict their domains to $\mathbf{h} \otimes \mathbf{D}$, and that it is indeed the case that $A_t^\dagger = (A_t)^\dagger$. We note also that if M_t is one of the fundamental noises, then the increments $M_t - M_s$ are adapted to the future, i.e., there exists an operator $\Delta M_{[s]}$ on $\mathbf{D}_{[s]}$ such that $(M_t - M_s) \psi \otimes e(f) = \psi \otimes e(f_{[s]}) \otimes \Delta M_{[s]} e(f_{[s]})$. This important property makes it easy to define stochastic integrals with respect to simple integrands. In particular, note that if X_s is admissible and adapted, then $X_s(M_t - M_s)$ is an allowable operator—there are no domain problems in the operator multiplication, as $X_s(M_t - M_s) = X_{[s]} \otimes \Delta M_{[s]}$ is just a tensor product in disguise.

DEFINITION B.2.13. An admissible process X_t is called simple if there exists an increasing sequence $t_n \nearrow \infty$ with $t_0 = 0$, such that $X_t = \sum_{n=0}^\infty X_{t_n} I_{[t_n, t_{n+1})}(t)$. If X_t is an adapted simple process and M_t is a fundamental noise, we define

$$\int_0^t X_t dM_t = \sum_{n=0}^\infty X_{t_n} (M_{t_{n+1} \wedge t} - M_{t_n \wedge t}) \quad \text{on the domain } \mathbf{h} \otimes \mathbf{D}.$$

We are now faced with extending this integral to a more general class of processes. Classically, the Itô isometry allows us to extend the Itô integral by taking limits in L^2 . In the quantum case, something similar happens, only the Itô isometry is replaced by a suitable estimate of the following form [HP84, Cor. 1]: defining

$$I_t = \int_0^t (E_s d\Lambda_s + F_s dA_s + G_s dA_s^\dagger + H_s ds),$$

we can estimate for any $\psi \otimes e(f) \in \mathbf{h} \otimes \mathbf{D}$

$$\begin{aligned} \|I_t \psi \otimes e(f)\|^2 &\leq C(t, f) \times \\ &\int_0^t (\|E_s \psi \otimes e(f)\|^2 + \|F_s \psi \otimes e(f)\|^2 + \|G_s \psi \otimes e(f)\|^2 + \|H_s \psi \otimes e(f)\|^2) ds, \end{aligned}$$

where $C(t, f) < \infty$ and (E_s, F_s, G_s, H_s) are admissible adapted simple processes. It is now evident how to define the stochastic integral for more general processes.

DEFINITION B.2.14. Let (E_s, F_s, G_s, H_s) be admissible adapted processes such that there exists a sequence $(E_s^n, F_s^n, G_s^n, H_s^n)$ of simple processes where

$$\int_0^t (\|(E_s - E_s^n) \psi \otimes e(f)\|^2 + \|(F_s - F_s^n) \psi \otimes e(f)\|^2 + \|(G_s - G_s^n) \psi \otimes e(f)\|^2 + \|(H_s - H_s^n) \psi \otimes e(f)\|^2) ds \xrightarrow{n \rightarrow \infty} 0$$

for any $\psi \otimes e(f) \in \mathfrak{h} \otimes \mathbb{D}$. Then (E_s, F_s, G_s, H_s) is said to be quantum stochastically integrable, and the quantum Itô integral I_t of this quadruple is uniquely defined as the strong limit of the corresponding simple integrals on the domain $\mathfrak{h} \otimes \mathbb{D}$.

An admissible adapted process X_t is called square-integrable if for any $t < \infty$

$$\int_0^t \|X_s \psi \otimes e(f)\|^2 ds < \infty \quad \text{for all } \psi \otimes e(f) \in \mathfrak{h} \otimes \mathbb{D}.$$

By [HP84, Prop. 3.2], any square-integrable process is stochastically integrable.

Let us list some simple properties of the quantum Itô integral. First, I_t is an admissible, adapted process. Moreover, I_t is easily shown to be strongly continuous in t on its domain; hence I_t is square-integrable, and in particular I_t is itself quantum stochastically integrable. Moreover, if (E_t, F_t, G_t, H_t) are stochastically integrable, then $(E_t^\dagger, F_t^\dagger, G_t^\dagger, H_t^\dagger)$ are also stochastically integrable and

$$I_t^\dagger = \int_0^t (E_s^\dagger d\Lambda_s + F_s^\dagger dA_s^\dagger + G_s^\dagger dA_s + H_s^\dagger ds).$$

We conclude this section with two key Theorems in the Hudson-Parthasarathy theory. The first gives the matrix elements of the quantum Itô integral. The second is the quantum Itô formula—in disguise.

THEOREM B.2.15. Let $\psi \otimes e(f), \phi \otimes e(g) \in \mathfrak{h} \otimes \mathbb{D}$, and let I_t be as above. Then

$$\langle \psi \otimes e(f), I_t \phi \otimes e(g) \rangle = \int_0^t \langle \psi \otimes e(f), \{f(s)^* E_s g(s) + F_s g(s) + f(s)^* G_s + H_s\} \phi \otimes e(g) \rangle ds.$$

THEOREM B.2.16. Let $\psi \otimes e(f), \phi \otimes e(g) \in \mathfrak{h} \otimes \mathbb{D}$, let I_t be the quantum Itô integral of the quantum stochastically integrable quadruple (E_t, F_t, G_t, H_t) , and let I_t' the quantum Itô integral of the integrable quadruple (E_t', F_t', G_t', H_t') . Then

$$\begin{aligned} \langle I_t' \psi \otimes e(f), I_t \phi \otimes e(g) \rangle = & \int_0^t \langle I_t' \psi \otimes e(f), \{f(s)^* E_s g(s) + F_s g(s) + f(s)^* G_s + H_s\} \phi \otimes e(g) \rangle ds \\ & + \int_0^t \langle \{g(s)^* E_s' f(s) + F_s' f(s) + g(s)^* G_s' + H_s'\} \psi \otimes e(f), I_t \phi \otimes e(g) \rangle ds \\ & + \int_0^t \langle \{E_s' f(s) + G_s'\} \psi \otimes e(f), \{E_s g(s) + G_s\} \phi \otimes e(g) \rangle ds. \end{aligned}$$

To see that this is indeed a quantum Itô rule, let us pretend that the products $(I_t')^\dagger I_t$, $(I_t')^\dagger E_t$, $(E_t')^\dagger G_t$, etc., are well defined. Then we can proceed to bring all the operators from the left-hand sides of the inner products to the right-hand sides. Using Theorem B.2.15 it is now easy to see that $(I_t')^\dagger I_t$ is again a quantum

Itô integral with coefficients that can be read off from the expression in Theorem B.2.16. In particular, the first line of that expression is precisely $(I_t')^\dagger dI_t$, the second line is $d(I_t')^\dagger I_t$, and the third line is the Itô correction term $d(I_t')^\dagger dI_t$.

Of course, the fact of the matter is that it is usually not so easy to establish that all these products are well defined—after all, this requires I_t to map $\mathfrak{h} \otimes \mathbb{D}$ into itself, which is rarely the case. The cleverness of Theorem B.2.16 is that it allows us to retain a quantum stochastic calculus, even though true operator multiplication is usually difficult to establish. Nonetheless, these results are most useful in precisely those cases where we can establish a “true” Itô rule in the sense of operator multiplication (by extending, if necessary, the domain of the leftmost operator). The following corollary is extremely useful, and has an obvious proof.

COROLLARY B.2.17. *For a pair of bounded allowable operators X and Y , define the multiplication $X \cdot Y = \overline{XY}$. Let*

$$\begin{aligned} dI_t &= E_t d\Lambda_t + F_t dA_t + G_t dA_t^\dagger + H_t dt, \\ dI_t' &= E_t' d\Lambda_t + F_t' dA_t + G_t' dA_t^\dagger + H_t' dt. \end{aligned}$$

Assume additionally that $I_t, E_t, F_t, G_t, H_t, I_t', E_t', F_t', G_t', H_t'$ are all bounded processes in the sense that $\sup_{s \leq t} \|I_s\| < \infty$ for all $t < \infty$, etc. Then

$$d(I_t \cdot I_t') = I_t \cdot dI_t' + dI_t \cdot I_t' + dI_t \cdot dI_t',$$

where the correction term $dI_t \cdot dI_t'$ is evaluated according to the quantum Itô table

$dI \setminus dI'$	dA_t	$d\Lambda_t$	dA_t^*	dt
dA_t	0	dA_t	dt	0
$d\Lambda_t$	0	$d\Lambda_t$	dA_t^*	0
dA_t^*	0	0	0	0
dt	0	0	0	0

The notation which we have used here should be obvious, but for concreteness let us be a little more precise: $I_t \cdot dI_t'$ contains terms such as $I_t \cdot F_t' dA_t$; $dI_t \cdot I_t'$ contains terms such as $F_t \cdot I_t' dA_t$; and $dI_t \cdot dI_t'$ contains terms such as $F_t \cdot G_t' dt$.

REMARK B.2.18. We have only considered the case of a Fock space with single multiplicity $\Gamma = \Gamma(L^2(\mathbb{R}_+))$. When the field has multiple degrees of freedom, this is usually introduced by considering the Fock space $\Gamma(L^2(\mathbb{R}_+) \otimes \mathbb{C})$, where \mathbb{C} is called the “color space” and $\dim \mathbb{C}$ is the number of degrees of freedom of the field. The theory here is much the same: one can introduce an orthonormal basis in \mathbb{C} , and for every basis element one obtains an independent copy of the fundamental noises. Now, however, there are additional fundamental noises (the exchange operators) which can scatter excitations between the different degrees of freedom. In practice very little changes in the theory as we have presented it, both conceptually and technically; however, the large number of degrees of freedom does bring with it a significant notational burden (though this is somewhat alleviated by the elegant Evans notation). The reader is referred to [Par92, Mey93] for further details. For notational simplicity and optimal clarity I have decided to stick with the simple Fock space Γ for the purposes of this thesis. The reader should keep in mind, however, that everything we will do extends almost trivially to the general case.

B.2.3. Quantum stochastic differential equations. To complete our discussion of the Hudson-Parthasarathy theory, we will briefly discuss quantum stochastic differential equations. We will only discuss linear equations of this type, and these are indeed the most useful. Let us first give a general existence and uniqueness result, see, e.g., [Mey93, page 173-175].

THEOREM B.2.19. *Consider the quantum stochastic differential equation*

$$V_t = I + \int_0^t \{L_1 d\Lambda_s + L_2 dA_s + L_3 dA_s^* + L_4 ds\} V_s,$$

where L_1, \dots, L_4 are fixed bounded operators of the form $L \otimes I$ on $\mathfrak{h} \otimes \Gamma$. Then there exists a unique admissible adapted process V_t that solves this equation.

This result admits a straightforward proof through Picard iteration; see also [HP84, sec. 7] for such a proof. Note that Hudson and Parthasarathy use the left form of the equation, i.e., where V_s is placed on the left of the L_i rather than on the right. This is essentially equivalent to the right equation, however, as it simply corresponds to solving the equation for V_t^\dagger in our notation.

We can now use the method of quantum stochastic differential equations to model the time evolution of a quantum system in interaction with the field. Hitherto we have been dealing purely with quantum *probability*; it is useful at this point to recall some quantum *mechanics*. Quantum mechanics is the theory that models the dynamics of physical systems within the framework of quantum probability. The observables in the theory are defined as self-adjoint operators on some quantum probability space. To model how these random variables evolve in time, one introduces a two-parameter unitary group $U_{s,t}$. In particular, $U_{s,t}$ is a unitary transformation that defines an automorphism of the quantum probability space, and $U_{r,t}U_{s,r} = U_{s,t}$ for all $0 \leq s \leq r \leq t$. The fact that the time evolution is an automorphism of the algebra emphasizes the idea that the algebra defines a physical model of a “closed system”; no information is lost from the system, and hence the time evolution is always reversible. This is a physical idea, not a probabilistic one; we will not dwell on its implications and accept it as physical fact.

REMARK B.2.20. In the time-homogeneous case treated in most quantum mechanics textbooks, a one-parameter unitary group U_t suffices to describe the time evolution. Here, however, we will be defining time evolutions that are driven in some sense by the fundamental quantum noises, and we are thus necessarily in the time-nonhomogeneous case. Note that this situation parallels closely the classical theory of stochastic flows generated by stochastic differential equations.

How does the time evolution work? Suppose that we have somehow defined the time evolution $U_{s,t}$. If a certain physical quantity is described by the observable X at time s , then this quantity will be described by the time evolved observable $U_{s,t}^* X U_{s,t}$ at time t . (The fact that the observable corresponding to the same physical quantity changes with time should be compared to the classical case where, e.g., the position of a particle can be described by a stochastic process; the same physical quantity is now described by a different random variable at every time t .) We can now understand the reason for the terminology “initial system” and “initial state” for the algebra \mathcal{B} and state ρ , which we have adjoined to the field probability space (\mathcal{W}, φ) for the purpose of quantum stochastic calculus. The observables affiliated to \mathcal{B} describe the physical quantities associated with the quantum system at the initial time $t = 0$; similarly, ρ gives their expectations at that time. After

some time t , the system will have interacted with the field and the associated observables will thus have evolved out of the initial algebra.

We are going to generate time evolutions by solving a particular quantum stochastic differential equation. First, we should figure out how to ensure that the solution V_t of the QSDE introduced above is in fact unitary (more precisely, that its closure is unitary). Necessary conditions are easily obtained by assuming that V_t is unitary, then calculating $V_t^\dagger \cdot V_t$ and $V_t \cdot V_t^\dagger$ using the quantum Itô rules and requiring that the coefficients vanish (as $V_t^* V_t = V_t V_t^* = I$ for a unitary operator). In fact these conditions are also sufficient; see [HP84, Thm. 7.1].

THEOREM B.2.21. *The unique solution of the QSDE of Theorem B.2.19 extends to a unitary process if and only if it is of the form*

$$dU_t = \left\{ (W - I) d\Lambda_t + L dA_t^\dagger - L^* W dA_t - \frac{1}{2} L^* L dt - iH dt \right\} U_t, \quad U_0 = I,$$

where W, L, H are bounded operators of the form $X \otimes I$ on $\mathfrak{h} \otimes \Gamma$, W is unitary, and H is self-adjoint. This QSDE is called the Hudson-Parthasarathy equation.

We have now almost constructed a suitable $U_{s,t}$. Indeed, if we define

$$dU_{s,t} = \left\{ (W - I) d\Lambda_t + L dA_t^\dagger - L^* W dA_t - \frac{1}{2} L^* L dt - iH dt \right\} U_{s,t}, \quad U_{s,s} = I,$$

then it is not difficult to verify that $U_{s,t}$ satisfies the desired group property. Beside the fact that this evidently defines an acceptable physical model, the Hudson-Parthasarathy equation is known to be an extremely good model for real physical phenomena, particularly in quantum optics [GZ04].

Finally, it should be mentioned that we have only scratched the surface of the theory of quantum stochastic differential equations (and quantum stochastic analysis in general). The current model is quite general already, and will suffice for our purposes. Particularly interesting, from the author's point of view, is the extension of these models to incorporate feedback controls. This leads to the notion of controlled quantum flows [BV06] and more generally to controlled quantum Markov processes, a general study of which still remains to be performed.

Bibliography

- [AA99] J. L. Anderson and S. L. Anderson, *A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts*, Monthly Weather Rev. **127** (1999), 2741–2758.
- [AAS⁺02] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi, *Adaptive homodyne measurement of optical phase*, Phys. Rev. Lett. **89** (2002), 133602.
- [AHSV06] A. Apte, M. Hairer, A. M. Stuart, and J. Voss, *Sampling the posterior: An approach to non-Gaussian data assimilation*, 2006, Physica D, at press.
- [AM06] M. A. Armen and H. Mabuchi, *Low-lying bifurcations in cavity quantum electrodynamics*, Phys. Rev. A **73** (2006), 063801.
- [AS82] E. Abraham and S. D. Smith, *Optical bistability and related devices*, Rep. Prog. Phys. **45** (1982), 815–885.
- [ASK04] B. Azimi-Sadjadi and P. S. Krishnaprasad, *A particle filtering approach to change detection for nonlinear systems*, EURASIP J. Appl. Signal Process. (2004), no. 15, 2295–2305.
- [ASK05] ———, *Approximate nonlinear filtering and its application in navigation*, Automatica **41** (2005), 945–956.
- [AZ97a] R. Atar and O. Zeitouni, *Exponential stability for nonlinear filtering*, Ann. Inst. H. Poincaré Probab. Statist. **33** (1997), 697–725.
- [AZ97b] ———, *Lyapunov exponents for finite state nonlinear filtering*, SIAM J. Control Optim. **35** (1997), 36–55.
- [BCL04] P. Baxendale, P. Chigansky, and R. Liptser, *Asymptotic stability of the Wonham filter: Ergodic and nonergodic signals*, SIAM J. Control Optim. **43** (2004), 643–669.
- [Bel80] V. P. Belavkin, *Quantum filtering of Markov signals with white quantum noise*, Radiotekhnika i Elektronika **25** (1980), 1445–1453.
- [Bel88] ———, *Nondemolition stochastic calculus in Fock space and nonlinear filtering and control in quantum systems*, Proceedings XXIV Karpacz winter school (R. Guellerak and W. Karwowski, eds.), Stochastic methods in mathematics and physics, World Scientific, Singapore, 1988, pp. 310–324.
- [Bel92] ———, *Quantum stochastic calculus and quantum nonlinear filtering*, J. Multivar. Anal. **42** (1992), 171–201.
- [Ben81] V. E. Beneš, *Exact finite-dimensional filters for certain diffusions with nonlinear drift*, Stochastics **5** (1981), 65–92.
- [Ben92] A. Bensoussan, *Stochastic control of partially observable systems*, Cambridge University Press, 1992.
- [Bes86] J. Besag, *On the statistical analysis of dirty pictures*, J. Roy. Statist. Soc. Ser. B **48** (1986), 259–302.
- [Bez71] Z. I. Bezaeva, *Limit theorems for conditional Markov chains*, Teor. Veroyatnost. i Primenen. **16** (1971), 437–445.
- [BGM04] L. Bouten, M. Guță, and H. Maassen, *Stochastic Schrödinger equations*, J. Phys. A: Math. Gen. **37** (2004), 3189–3209.
- [BHL99] D. Brigo, B. Hanson, and F. Le Gland, *Approximate nonlinear filtering by projection on exponential manifolds of densities*, Bernoulli **5** (1999), 495–534.
- [Bia95] P. Biane, *Calcul stochastique non-commutatif*, Lectures on probability theory (Saint-Flour, 1993), Lecture Notes in Math., vol. 1608, Springer, Berlin, 1995, pp. 1–96.
- [Bil99] P. Billingsley, *Convergence of probability measures*, second ed., Wiley Series in Probability and Statistics, John Wiley & Sons, New York, 1999.

- [Bis81] J.-M. Bismut, *Mécanique aléatoire*, Lecture Notes in Mathematics 866, Springer-Verlag, 1981.
- [BJ87] R. S. Bucy and P. D. Joseph, *Filtering for stochastic processes with applications to guidance*, second ed., Chelsea Publishing Co., New York, 1987.
- [BJP02] R. K. Boel, M. R. James, and I. R. Petersen, *Robustness and risk-sensitive filtering*, IEEE Trans. Automat. Control **47** (2002), 451–461.
- [BK98] A. Budhiraja and H. J. Kushner, *Robustness of nonlinear filters over the infinite time interval*, SIAM J. Control Optim. **36** (1998), 1618–1637.
- [BK99] ———, *Approximation and limit results for nonlinear filters over an infinite time interval*, SIAM J. Control Optim. **37** (1999), 1946–1979.
- [BKK95] A. G. Bhatt, G. Kallianpur, and R. L. Karandikar, *Uniqueness and robustness of solution of measure-valued equations of nonlinear filtering*, Ann. Probab. **23** (1995), 1895–1938.
- [BKK99] ———, *Robustness of the nonlinear filter*, Stochastic Process. Appl. **81** (1999), 247–254.
- [BL76] H. J. Brascamp and E. H. Lieb, *On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation*, J. Funct. Anal. **22** (1976), 366–389.
- [BM82] J.-M. Bismut and D. Michel, *Diffusions conditionnelles. II. Générateur conditionnel. Application au filtrage*, J. Funct. Anal. **45** (1982), 274–292.
- [Bor00] C. Borell, *Diffusion equations and geometric inequalities*, Potential Anal. **12** (2000), 49–71.
- [BP03] J. Bröcker and U. Parlitz, *Analyzing communication schemes using methods from nonlinear filtering*, Chaos **13** (2003), 195–208.
- [BR87] O. Bratteli and D. W. Robinson, *Operator algebras and quantum statistical mechanics. 1*, second ed., Texts and Monographs in Physics, Springer-Verlag, New York, 1987.
- [Buc65] R. Bucy, *Nonlinear filtering theory*, IEEE Trans. Automat. Control **10** (1965), 198.
- [BV06] L. Bouten and R. Van Handel, *On the separation principle of quantum control*, 2006, preprint, math-ph/0511021.
- [BV07] ———, *Controlled quantum stochastic processes*, 2007, in preparation.
- [BVJ06a] L. Bouten, R. Van Handel, and M. R. James, *A discrete invitation to quantum filtering and feedback control*, 2006, preprint, math.PR/0606118.
- [BVJ06b] ———, *An introduction to quantum filtering*, 2006, preprint, math.OC/0601741.
- [Car93] H. J. Carmichael, *An open systems approach to quantum optics*, Springer-Verlag, 1993.
- [CC05] J. M. C. Clark and D. Crisan, *On a robust version of the integral representation formula of nonlinear filtering*, Probab. Theory Related Fields **133** (2005), 43–56.
- [CCB90] D. A. Coast, G. G. Cano, and S. A. Briller, *Use of hidden Markov models for electrocardiographic signal analysis*, J. Electrocardiol. **23** (1990), 184–191.
- [CDMS97] H. Carvalho, P. Del Moral, A. Monin, and G. Salut, *Optimal nonlinear filtering in GPS/INS integration*, IEEE Trans. Aerospace Electr. Syst. **33** (1997), 835–850.
- [Che00] A. M. Chebotarev, *Lectures on quantum probability*, Aportaciones Matemáticas: Textos [Mathematical Contributions: Texts], vol. 14, Sociedad Matemática Mexicana, México, 2000.
- [Chi06] P. Chigansky, *Stability of nonlinear filters: A survey*, 2006, Mini-course lecture notes, Petropolis, Brazil. Available at <http://www.wisdom.weizmann.ac.il/~pavel/>.
- [CHL97] M. Capitaine, E. P. Hsu, and M. Ledoux, *Martingale representation and a simple proof of logarithmic Sobolev inequalities on path spaces*, Electron. Comm. Probab. **2** (1997), 71–81.
- [CKM91] S. H. Chung, V. Krishnamurthy, and J. B. Moore, *Adaptive processing techniques based on hidden Markov models for characterizing very small channel currents buried in noise and deterministic interferences*, Phil. Trans. Roy. Soc. Lond. B **334** (1991), 357–384.
- [CL06] P. Chigansky and R. Liptser, *On a role of predictor in the filtering stability*, Electron. Comm. Probab. **11** (2006), 129–140.
- [Cla78] J. M. C. Clark, *The design of robust approximations to the stochastic differential equations of nonlinear filtering*, Communication systems and random process theory (Proc. 2nd NATO Advanced Study Inst., Darlington, 1977), Sijthoff & Noordhoff,

- Alphen aan den Rijn, 1978, pp. 721–734. NATO Advanced Study Inst. Ser., Ser. E: Appl. Sci., No. 25.
- [CLR06] J. Cvitanic, R. Liptser, and B. Rozovskii, *A filtering approach to tracking volatility from prices observed at random times*, Ann. Appl. Probab. **16** (2006), 1633–1652.
- [CM84] M. Chaleyat-Maurel and D. Michel, *Des résultats de non existence de filtre de dimension finie*, Stochastics **13** (1984), 83–102.
- [CMX⁺90] S. H. Chung, J. B. Moore, L. Xia, L. S. Premkumar, and P. W. Gage, *Characterization of single channel currents using digital signal processing techniques based on hidden Markov models*, Phil. Trans. Roy. Soc. Lond. B **329** (1990), 265–285.
- [COC99] J. M. C. Clark, D. Ocone, and C. Coumarbatch, *Relative entropy and error bounds for filtering of Markov processes*, Math. Control Signals Systems **12** (1999), 346–360.
- [Con85] J. B. Conway, *A course in functional analysis*, Graduate Texts in Mathematics, vol. 96, Springer-Verlag, New York, 1985.
- [CV06] P. Chigansky and R. Van Handel, *Model robustness of finite state nonlinear filtering over the infinite time horizon*, 2006, Ann. Appl. Probab., at press, math.PR/0604295.
- [CZ95] P. E. Caines and J. F. Zhang, *On the adaptive control of jump parameter systems via nonlinear filtering*, SIAM J. Control Optim. **33** (1995), 1758–1777.
- [Dav76] E. B. Davies, *Quantum theory of open systems*, Academic Press, 1976.
- [Dav80] M. H. A. Davis, *On a multiplicative functional transformation arising in nonlinear filtering theory*, Z. Wahrsch. Verw. Gebiete **54** (1980), 125–139.
- [Dav82] ———, *A pathwise solution of the equations of nonlinear filtering*, Teor. Veroyatnost. i Primenen. **27** (1982), 160–167.
- [DE97] P. Dupuis and R. S. Ellis, *A weak convergence approach to the theory of large deviations*, Wiley Series in Probability and Statistics, John Wiley & Sons, New York, 1997.
- [DJ99] A. C. Doherty and K. Jacobs, *Feedback control of quantum systems using continuous state estimation*, Phys. Rev. A **60** (1999), 2700–2711.
- [DKW01] H. Deng, M. Krstić, and R. J. Williams, *Stabilization of stochastic nonlinear systems driven by noise of unknown covariance*, IEEE Trans. Automat. Control **46** (2001), 1237–1253.
- [DM82] C. Dellacherie and P.-A. Meyer, *Probabilities and potential. B*, North-Holland Mathematics Studies, vol. 72, North-Holland Publishing Co., Amsterdam, 1982, Theory of martingales, Translated from the French by J. P. Wilson.
- [Dos77] H. Doss, *Liens entre équations différentielles stochastiques et ordinaires*, Ann. Inst. H. Poincaré sec. B **13** (1977), 99–125.
- [Dun68] T. E. Duncan, *Evaluation of likelihood functions*, Information and Control (1968), no. 13, 62–74.
- [DZ91] B. Delyon and O. Zeitouni, *Lyapunov exponents for filtering problems*, Applied stochastic analysis (London, 1989), Stochastics Monogr., vol. 5, Gordon and Breach, New York, 1991, pp. 511–521.
- [EAM95] R. J. Elliott, L. Aggoun, and J. B. Moore, *Hidden Markov models*, Applications of Mathematics (New York), vol. 29, Springer-Verlag, New York, 1995.
- [EK86] S. N. Ethier and T. G. Kurtz, *Markov processes: Characterization and convergence*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, New York, 1986.
- [Ell82] R. J. Elliott, *Stochastic calculus and applications*, Applications of Mathematics (New York), vol. 18, Springer-Verlag, 1982.
- [Ell86] ———, *Reverse-time Markov processes*, IEEE Trans. Inform. Theory **32** (1986), 290–292.
- [EM02] Y. Ephraim and N. Merhav, *Hidden Markov processes*, IEEE Trans. Inf. Th. **48** (2002), 1518–1569.
- [Eph92] Y. Ephraim, *Statistical-model-based speech enhancement systems*, Proc. IEEE **80** (1992), 1526–1555.
- [FGDW02] W. Fong, S. J. Godsill, A. Doucet, and M. West, *Monte Carlo smoothing with application to audio signal enhancement*, IEEE Trans. Signal Proc. **50** (2002), 438–449.
- [Fit05] P. J. Fitzsimmons, *Functional monotone class theorem*, 2005, lecture notes.

- [FK71] P. Frost and T. Kailath, *An innovations approach to least-squares estimation—Part III: Nonlinear estimation in white Gaussian noise*, IEEE Trans. Automat. Control **16** (1971), 217–226.
- [FKK72] M. Fujisaki, G. Kallianpur, and H. Kunita, *Stochastic differential equations for the non linear filtering problem*, Osaka J. Math. **9** (1972), 19–40.
- [FM83] W. H. Fleming and S. K. Mitter, *Optimal control and nonlinear filtering for nondegenerate diffusion processes*, Stochastics **8** (1982/83), 63–77.
- [FR75] W. H. Fleming and R. W. Rishel, *Deterministic and stochastic optimal control*, Springer-Verlag, Berlin, 1975.
- [FR92] D. R. Fredkin and J. A. Rice, *Bayesian restoration of single-channel patch clamp recordings*, Biometrics **48** (1992), 427–448.
- [FR01] R. Frey and W. J. Runggaldier, *A nonlinear filtering approach to volatility estimation with a view towards high frequency data*, Int. J. Theor. Appl. Finance **4** (2001), 199–210, Information modeling in finance (Évry, 2000).
- [Fri75] A. Friedman, *Stochastic differential equations and applications. Vol. 1*, Academic Press, New York, 1975, Probability and Mathematical Statistics, vol. 28.
- [FS06] W. H. Fleming and H. M. Soner, *Controlled Markov processes and viscosity solutions*, second ed., Stochastic Modelling and Applied Probability, vol. 25, Springer, New York, 2006.
- [GGB⁺02] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund, *Particle filters for positioning, navigation and tracking*, IEEE Trans. Signal Proc. **50** (2002), 425–437.
- [Gri75] D. Griffeath, *Uniform coupling of non-homogeneous Markov chains*, J. Appl. Probability **12** (1975), 753–762.
- [GSDM03] J. M. Geremia, J. K. Stockton, A. C. Doherty, and H. Mabuchi, *Quantum Kalman filtering and the Heisenberg limit in atomic magnetometry*, Phys. Rev. Lett. **91** (2003), 250801.
- [GSM05] J. M. Geremia, J. K. Stockton, and H. Mabuchi, *Suppression of spin projection noise in broadband atomic magnetometry*, Phys. Rev. Lett. **94** (2005), 203002.
- [GV06] J. Gough and R. Van Handel, *Singular perturbation of quantum stochastic differential equations with coupling through an oscillator mode*, 2006, J. Stat. Phys., at press, math-ph/0609015.
- [GY06] X. Guo and G. Yin, *The Wonham filter with random parameters: Rate of convergence and error bounds*, IEEE Trans. Automat. Control **51** (2006), 460–464.
- [GZ04] C. Gardiner and P. Zoller, *Quantum noise*, third ed., Springer, 2004.
- [Hau85] U. G. Haussmann, *On the drift of a reversed diffusion*, Stochastic differential systems (Marseille-Luminy, 1984), Lecture Notes in Control and Inform. Sci., vol. 69, Springer, Berlin, 1985, pp. 164–175.
- [HC99] J. L. Hibey and C. D. Charalambous, *Conditional densities for continuous-time nonlinear hybrid systems with applications to fault detection*, IEEE Trans. Automat. Control **44** (1999), 2164–2169.
- [Hol91] A. S. Holevo, *Quantum stochastic calculus*, J. Soviet Math. **56** (1991), 2609–2624, Translation of Itogi Nauki i Tekhniki, Ser. Sovr. Prob. Mat. 36, 3–28, 1990.
- [HP84] R. L. Hudson and K. R. Parthasarathy, *Quantum Ito’s formula and stochastic evolutions*, Comm. Math. Phys. **93** (1984), 301–323.
- [HP86] U. G. Haussmann and É. Pardoux, *Time reversal of diffusions*, Ann. Probab. **14** (1986), 1188–1205.
- [HS87] C.-R. Hwang and S. J. Sheu, *A generalization of Chernoff inequality via stochastic analysis*, Probab. Theory Related Fields **75** (1987), 149–157.
- [HSV06] M. Hairer, A. M. Stuart, and J. Voss, *Analysis of SPDEs arising in path sampling. II. The nonlinear case*, 2006, preprint.
- [Hud03] R. L. Hudson, *An introduction to quantum stochastic calculus and some of its applications*, Quantum probability communications, vol. XI (Grenoble, 1998), QP-PQ, XI, World Scientific, River Edge, NJ, 2003, pp. 221–271.
- [HUL93] J.-B. Hiriart-Urruty and C. Lemaréchal, *Convex analysis and minimization algorithms. II*, Grundlehren der Mathematischen Wissenschaften, vol. 306, Springer-Verlag, Berlin, 1993, Advanced theory and bundle methods.

- [HUL01] ———, *Fundamentals of convex analysis*, Grundlehren Text Editions, Springer-Verlag, Berlin, 2001.
- [JB88] M. R. James and J. S. Baras, *Nonlinear filtering and large deviations: A PDE-control theoretic approach*, *Stochastics* **23** (1988), 391–412.
- [Kad56] R. V. Kadison, *Operator algebras with a faithful weakly-closed representation*, *Ann. Math.* **64** (1956), 175–181.
- [Kal60] R. E. Kalman, *A new approach to linear filtering and prediction problems*, *Trans. ASME ser. D* **82** (1960), 35–45.
- [Kal80] G. Kallianpur, *Stochastic filtering theory*, Springer, Berlin, 1980.
- [Kal97] O. Kallenberg, *Foundations of modern probability*, Probability and its Applications (New York), Springer-Verlag, 1997.
- [KR97a] R. V. Kadison and J. R. Ringrose, *Fundamentals of the theory of operator algebras. Vol. I*, Graduate Studies in Mathematics, vol. 15, American Mathematical Society, Providence, RI, 1997.
- [KR97b] ———, *Fundamentals of the theory of operator algebras. Vol. II*, Graduate Studies in Mathematics, vol. 16, American Mathematical Society, Providence, RI, 1997.
- [Kri05] V. Krishnan, *Nonlinear filtering and smoothing: An introduction to martingales, stochastic integrals and estimation*, Dover, 2005.
- [KS68] G. Kallianpur and C. Striebel, *Estimation of stochastic systems: Arbitrary system process with additive white noise observation errors*, *Ann. Math. Statist.* **39** (1968), 785–801.
- [KS69] ———, *Stochastic differential equations occurring in the estimation of continuous parameter stochastic processes*, *Teor. Verojatnost. i Primenen* **14** (1969), 597–622.
- [KS91] I. Karatzas and S. E. Shreve, *Brownian motion and stochastic calculus*, second ed., Graduate Texts in Mathematics, vol. 113, Springer-Verlag, New York, 1991.
- [KSH00] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear estimation*, Prentice Hall, New York, 2000.
- [Kun84] H. Kunita, *Stochastic differential equations and stochastic flows of diffeomorphisms*, École d'été de probabilités de Saint-Flour, XII—1982, Lecture Notes in Math., vol. 1097, Springer, Berlin, 1984, pp. 143–303.
- [Kun90] ———, *Stochastic flows and stochastic differential equations*, Cambridge Studies in Advanced Mathematics, vol. 24, Cambridge University Press, 1990.
- [Kus64] H. J. Kushner, *On the differential equations satisfied by conditional probability densities of Markov processes, with applications*, *J. Soc. Indust. Appl. Math. Ser. A Control* **2** (1964), 106–119.
- [Kus67] ———, *Dynamical equations for optimal nonlinear filtering*, *J. Diff. Eq.* **3** (1967), 179–190.
- [Kus72] ———, *Stochastic stability*, Stability of stochastic dynamical systems (Proc. Internat. Sympos., Univ. Warwick, Coventry, 1972), Springer, Berlin, 1972, pp. 97–124. Lecture Notes in Math., vol. 294.
- [Kus79] ———, *A robust discrete state approximation to the optimal nonlinear filter for a diffusion*, *Stochastics* **3** (1979), 75–83.
- [LB98] A. V. Lukashin and M. Borodovsky, *GeneMark.hmm: New solutions for gene finding*, *Nucl. Acids Res.* **26** (1998), 1107–1115.
- [Lin02] T. Lindvall, *Lectures on the coupling method*, Dover Publications, Mineola, NY, 2002.
- [Liv95] C. Liverani, *Decay of correlations*, *Ann. Math.* **142** (1995), 239–301.
- [LO03] F. Le Gland and N. Oudjane, *A robustification approach to stability and to uniform particle approximation of nonlinear filters: The example of pseudo-mixing signals*, *Stochastic Process. Appl.* **106** (2003), 279–316.
- [LO04] ———, *Stability and uniform approximation of nonlinear filters using the Hilbert metric and application to particle filters*, *Ann. Appl. Probab.* **14** (2004), 144–187.
- [LS97] C. Lemaréchal and C. Sagastizábal, *Practical aspects of the Moreau-Yosida regularization: theoretical preliminaries*, *SIAM J. Optim.* **7** (1997), 367–385.
- [LS01] R. S. Liptser and A. N. Shiryaev, *Statistics of random processes. I*, Applications of Mathematics (New York), vol. 5, Springer-Verlag, Berlin, 2001.
- [LSBS96] R. S. Liptser, Y. Steinberg, B. Z. Bobrovsky, and Z. Schuss, *Fixed lag smoothing of scalar diffusions. I. The filtering-smoothing equation*, *Stochastic Process. Appl.* **64** (1996), 237–255.

- [Maa03] H. Maassen, *Quantum probability applied to the damped harmonic oscillator*, Quantum probability communications, vol. XII (Grenoble, 1998), QP-PQ, XII, World Scientific, River Edge, NJ, 2003, pp. 23–58.
- [Mab06] H. Mabuchi, 2006, private communication.
- [Mey93] P.-A. Meyer, *Quantum probability for probabilists*, Lecture Notes in Mathematics, vol. 1538, Springer-Verlag, Berlin, 1993.
- [Mit79] S. K. Mitter, *On the analogy between mathematical problems of nonlinear filtering and quantum physics*, Ricerche Automat. **10** (1979), 163–216.
- [Mit81] ———, *Nonlinear filtering and stochastic mechanics*, Stochastic systems: the mathematics of filtering and identification and applications (Les Arcs, 1980), NATO Adv. Study Inst. Ser. C: Math. Phys. Sci., vol. 78, Reidel, Dordrecht, 1981, pp. 479–503.
- [Mit82] ———, *Lectures on nonlinear filtering and stochastic control*, Nonlinear filtering and stochastic control (Cortona, 1981), Lecture Notes in Math., vol. 972, Springer, Berlin, 1982, pp. 170–207.
- [MN03] S. K. Mitter and N. J. Newton, *A variational approach to nonlinear estimation*, SIAM J. Control Optim. **42** (2003), 1813–1833.
- [Mor66] R.E. Mortensen, *Optimal control of continuous-time stochastic systems*, Ph.D. thesis, Univ. California, Berkeley, 1966.
- [MV06] M. Mirrahimi and R. Van Handel, *Stabilizing feedback controls for quantum systems*, 2006, SIAM J. Control Optim., at press, math-ph/0510066.
- [MW95] L. Mandel and E. Wolf, *Optical coherence and quantum optics*, Cambridge University Press, 1995.
- [MW98] H. Mabuchi and H. M. Wiseman, *Retroactive quantum jumps in a strongly coupled atom-field system*, Phys. Rev. Lett. **81** (1998), 4620–4623.
- [MYK99] H. Mabuchi, J. Ye, and H. J. Kimble, *Full observation of single-atom dynamics in cavity QED*, Appl. Phys. B **68** (1999), 1095–1108.
- [NP88] D. Nualart and É. Pardoux, *Stochastic calculus with anticipating integrands*, Probab. Theory Related Fields **78** (1988), 535–581.
- [Nua95] D. Nualart, *The Malliavin calculus and related topics*, Probability and its Applications (New York), Springer-Verlag, New York, 1995.
- [OK91] D. Ocone and I. Karatzas, *A generalized Clark representation formula, with application to optimal portfolios*, Stochastics Stochastics Rep. **34** (1991), 187–220.
- [OP89] D. Ocone and É. Pardoux, *A generalized Itô-Ventzell formula. Application to a class of anticipating stochastic differential equations*, Ann. Inst. H. Poincaré Probab. Statist. **25** (1989), 39–71.
- [OP96] ———, *Asymptotic stability of the optimal filter with respect to its initial condition*, SIAM J. Control Optim. **34** (1996), 226–243.
- [OR05] N. Oudjane and S. Rubenthaler, *Stability and uniform particle approximation of nonlinear filters in case of non ergodic signals*, Stoch. Anal. Appl. **23** (2005), 421–448.
- [Par81] É. Pardoux, *The solution of the nonlinear filtering equation as a likelihood function*, Proc. 20th IEEE Conf. Decision and Control, 1981, pp. 316–319.
- [Par86] ———, *Grossissement d’une filtration et retournement du temps d’une diffusion*, Séminaire de Probabilités, XX, 1984/85, Lecture Notes in Math., vol. 1204, Springer, Berlin, 1986, pp. 48–55.
- [Par91] ———, *Filtrage non linéaire et équations aux dérivées partielles stochastiques associées*, École d’Été de Probabilités de Saint-Flour XIX—1989, Lecture Notes in Math., vol. 1464, Springer, Berlin, 1991, pp. 67–163.
- [Par92] K. R. Parthasarathy, *An introduction to quantum stochastic calculus*, Monographs in Mathematics, vol. 85, Birkhäuser Verlag, Basel, 1992.
- [Par82] É. Pardoux, *Équations du filtrage non linéaire, de la prédiction et du lissage*, Stochastics **6** (1981/82), 193–231.
- [Ped79] G. K. Pedersen, *C*-algebras and their automorphism groups*, London Mathematical Society Monographs, vol. 14, Academic Press, London, 1979.
- [PR02] Z. Palmowski and T. Rolski, *A technique for exponential change of measure for Markov processes*, Bernoulli **8** (2002), 767–785.
- [Pro04] P. E. Protter, *Stochastic integration and differential equations*, Applications of Mathematics (New York), vol. 21, Springer-Verlag, Berlin, 2004.

- [Rab89] L. R. Rabiner, *A tutorial on hidden Markov models and selected applications in speech recognition*, Proc. IEEE **77** (1989), 257–286.
- [Roz90] B. L. Rozovskii, *Stochastic evolution systems*, Mathematics and its Applications (Soviet Series), vol. 35, Kluwer Academic Publishers Group, Dordrecht, 1990, Linear theory and applications to nonlinear filtering, Translated from the Russian by A. Yarkho.
- [RS72] M. Reed and B. Simon, *Methods of modern mathematical physics. I. Functional analysis*, Academic Press, New York, 1972.
- [Rud73] W. Rudin, *Functional analysis*, McGraw-Hill Book Co., New York, 1973.
- [RW00] L. C. G. Rogers and D. Williams, *Diffusions, Markov processes, and martingales. Vol. 2: Itô calculus*, Cambridge University Press, 2000.
- [Sak98] S. Sakai, *C^* -algebras and W^* -algebras*, Classics in Mathematics, Springer-Verlag, Berlin, 1998.
- [SH04] J. Sass and U. G. Haussmann, *Optimizing the terminal wealth under partial information: The drift process as a continuous time Markov chain*, Finance Stoch. **8** (2004), 553–577.
- [She83] S. J. Sheu, *Solution of certain parabolic equations with unbounded coefficients and its application to nonlinear filtering*, Stochastics **10** (1983), 31–46.
- [She85] ———, *Stochastic control and exit probabilities of jump processes*, SIAM J. Control Optim. **23** (1985), 306–328.
- [She91] ———, *Some estimates of the transition density of a nondegenerate diffusion Markov process*, Ann. Probab. **19** (1991), 538–561.
- [Shi63] A. N. Shiryaev, *Optimal methods in quickest detection problems*, Teor. Veroyatnost. i Primenen. **8** (1963), 26–51.
- [Shi73] ———, *Statistical sequential analysis: Optimal stopping rules*, American Mathematical Society, Providence, R.I., 1973, Translations of Mathematical Monographs, vol. 38.
- [Shi06] ———, *Proof of the Poincaré-Chernoff inequality and the logarithmic Sobolev inequality by the methods of stochastic calculus for Brownian motion*, Russ. Math. Surv. **61** (2006), 571–573.
- [Sta04] W. Stannat, *Stability of the pathwise filter equation on \mathbb{R}^d* , 2004, Technical report, Univ. Bielefeld, <http://www.mathematik.tu-darmstadt.de/~stannat/>.
- [Sta05] ———, *Stability of the filter equation for a time-dependent signal on \mathbb{R}^d* , Appl. Math. Optim. **52** (2005), 39–71.
- [Sta06] ———, *Stability of the optimal filter via pointwise gradient estimates*, Stochastic partial differential equations and applications—VII, Lect. Notes Pure Appl. Math., vol. 245, Chapman & Hall/CRC, Boca Raton, FL, 2006, pp. 281–293.
- [Str60] R. L. Stratonovich, *Conditional Markov processes*, Teor. Veroyatnost. i Primenen. **5** (1960), 172–195.
- [Str66] ———, *A new representation for stochastic integrals and equations*, SIAM J. Control **4** (1966), 362–371.
- [Str68] ———, *Conditional Markov processes and their application to the theory of optimal control*, Translated from the Russian by R. N. and N. B. McDonough for Scripta Technica. Modern Analytic and Computational Methods in Science and Mathematics, No. 7, American Elsevier Publishing Co., New York, 1968.
- [Sus78] H. J. Sussmann, *On the gap between deterministic and stochastic ordinary differential equations*, Ann. Probab. **6** (1978), 19–41.
- [SVM04] J. K. Stockton, R. Van Handel, and H. Mabuchi, *Deterministic Dicke-state preparation with continuous measurement and control*, Phys. Rev. A **70** (2004), 022106.
- [SVP+04] I. Siddiqi, R. Vijay, F. Pierre, C. M. Wilson, M. Metcalfe, C. Rigetti, L. Frunzio, and M. H. Devoret, *RF-driven Josephson bifurcation amplifier for quantum measurement*, Phys. Rev. Lett. **93** (2004), 207002.
- [SVW04] A. M. Stuart, J. Voss, and P. Wiberg, *Conditional path sampling of SDEs and the Langevin MCMC method*, Commun. Math. Sci. **2** (2004), 685–697.
- [Tak02] M. Takesaki, *Theory of operator algebras. I*, Encyclopaedia of Mathematical Sciences, vol. 124, Springer-Verlag, Berlin, 2002.
- [Thi03] M. Thill, *A non-commutative monotone selection principle*, 2003, preprint, math.FA/0312061 and math.FA/0311438.

- [Van06] R. Van Handel, *Almost global stochastic stability*, SIAM J. Control Optim. **45** (2006), 1297–1313.
- [Van07] ———, *Stopping times in quantum control*, 2007, in preparation.
- [VC03] M. H. Vellekoop and J. M. C. Clark, *A nonlinear filtering approach to changepoint detection problems: Direct and differential-geometric methods*, SIAM J. Control Optim. **42** (2003), 469–494.
- [VM05] R. Van Handel and H. Mabuchi, *Quantum projection filter for a highly nonlinear model in cavity QED*, J. Opt. B: Quant. Semiclass. Opt. **7** (2005), S226–S236.
- [VM06] ———, *Optimal error tracking via quantum coding and continuous syndrome measurement*, 2006, preprint, [quant-ph/0511221](#).
- [Von29] J. Von Neumann, *Zur algebra der funktionaloperationen und theorie der normalen operatoren*, Math. Ann. **102** (1929), 370–427.
- [VSM05a] R. Van Handel, J. K. Stockton, and H. Mabuchi, *Feedback control of quantum state reduction*, IEEE Trans. Automat. Control **50** (2005), 768–780.
- [VSM05b] ———, *Modelling and feedback control design for quantum state preparation*, J. Opt. B: Quant. Semiclass. Opt. **7** (2005), S179–S197.
- [Wil91] D. Williams, *Probability with martingales*, Cambridge University Press, Cambridge, 1991.
- [WM94] D.F. Walls and G.J. Milburn, *Quantum optics*, Springer Verlag, Berlin Heidelberg, 1994.
- [Won65] W. M. Wonham, *Some applications of stochastic differential equations to optimal nonlinear filtering*, J. Soc. Indust. Appl. Math. Ser. A Control **2** (1965), 347–369 (1965).
- [YFSJ03] M. F. Yanik, S. Fan, M. Soljagic, and J. D. Joannopoulos, *All-optical transistor action with bistable switching in a photonic cross-waveguide geometry*, Opt. Lett. **28** (2003), 2506–2508.
- [Zak69] M. Zakai, *On the optimal filtering of diffusion processes*, Z. Wahrsch. Verw. Geb. **11** (1969), 230–243.