Noisy patterns: Bridging the gap between stochastics and dynamics

Hamster, C.H.S.

Author: Hamster, C.H.S
Title: Noisy patterns: Bridging the gap between stochastics and dynamics
Issue date: 2020-09-09

1 Introduction

Let us go back a few short moments in time, to the moment you decided to pick up this thesis to read it (or to just browse through, wondering if all those scribbles actually mean something). What happened between the moment you thought of picking it up and the moment your arm started to move? A small electric pulse travelled through your musculocutaneous nerve to bring the signal from your head to your arm. From our experience, we know that this always works in a healthy person, which is a small miracle if one thinks about the challenges such a pulse meets on its way. It has to hop from node of Ranvier to node of Ranvier, jumping over myelin sheaths, triggering ion pumps all along the way to boost the signal. Our intuition tells us that the signal propagation must be robust, as we know that the route from the brain to the muscle is not a perfectly straight and smooth highway. Not every cell the pulse encounters on the way is identical, there will be random external fluctuations in the potentials, or, in other words, the route will be noisy. Despite these challenges, the pulse will make it to the end.

Here we hit the core of the problems discussed in this thesis: How can a pulse, or more generally a pattern, travel in a noisy environment. Hence the title ‘Noisy Patterns’. Whenever I explain this story to non-mathematicians, the follow-up question is often ‘so you work together with biologists?’ The question to that answer goes along the line ‘not yet, but maybe, somewhere, once, in the future.’ The goal of this thesis is not to study pulses in human nerve cells, our goal is much more basic, but not less challenging! Our goal is the following: Can we develop the mathematical tools and machinery needed to rigorously study travelling waves in stochastic reaction-diffusion equations, in the same way we can study and understand their deterministic counterparts?

Over the past four years, we set several steps in this direction, resulting in multiple papers that form the core of this thesis. These results are certainly not the first results on stochastic reaction-diffusion equations, but what sets our results apart is the fact that they are firmly rooted in dynamical systems theory, or to be more specific, rooted in the language of travelling waves familiar to the nonlinear waves family that Leiden is a part of. This thesis will not contain any new theorems on the existence of solutions nor on the existence of invariant measures. No, the questions we will ask ourselves are the same as our community has always asked: can we show the existence of travelling waves, can we compute their shape and speed and can we study their stability? The answers to these questions are spread out over this thesis, but informally, we can formulate the following theorem:

Theorem. *Under certain technical assumptions, travelling waves or pulses in bistable reaction-diffusion equations persist on exponentially long timescales when the equation is forced by a small multiplicative noise term. The average wave speed and shape are close to the deterministic values and can be expanded systematically around the deterministic wave using Taylor series.*

We now give two main examples of stochastic partial differential equations (SPDEs) to which the theorem above applies. The first example is the stochastic Nagumo[1] equation,

$$u_t = \rho u_{xx} + u(1-u)(u-a) + \sigma u(1-u)\,\xi(x,t), \qquad x \in \mathbb{R},\ t \in \mathbb{R}^+, \tag{1.0.1}$$

and the second example is the two-component FitzHugh-Nagumo (FHN) equation

$$\begin{aligned} u_t &= \rho_1 u_{xx} + u(1-u)(u-a) + \sigma u\,\xi(x,t) \\ v_t &= \rho_2 v_{xx} + \varepsilon(u - \gamma v) \end{aligned} \tag{1.0.2}$$

for positive $\rho_1, \rho_2, \varepsilon, \gamma$ and $a \in (0,1)$. Here ξ is a Gaussian process with

$$\begin{aligned} E[\xi(x,t)] &= 0, \\ E[\xi(x,t)\xi(x',t')] &= \delta(t-t')q(x-x'). \end{aligned} \tag{1.0.3}$$

The last line means that the noise is uncorrelated in time and correlated in space with correlation function q, typically a Gaussian. The parameter σ indicates the noise strength and will be assumed to be small. For both equations, the existence of a stable travelling wave/pulse is known in the deterministic case [60, 113] and for the Nagumo equation the profile and speed can be computed explicitly, see Chapter 2. Remark that the noise term is multiplicative, i.e. the noise term ξ is multiplied by a function depending on u. Furthermore, note that these functions are chosen in such a way that the noise disappears in the background states of the wave (0 and 1 for the Nagumo equation and 0 for the FHN-equation). This is an assumption we will make throughout this thesis. There are two main advantages to this assumption, a technical and practical one. From a technical point of view, this type of multiplicative noise ensures that deviations from the travelling wave disappear at $\pm\infty$, hence making the deviations integrable. From a more practical viewpoint, additive noise can cause the wave to disappear or cause new waves to form [79]. This makes the study of a single travelling wave significantly more difficult.

We illustrate the practical questions we will answer in this thesis by showing two realizations of the equations above in Figure 1.1. For $\sigma = 0$, the waves are perfectly frozen in the $x - ct$ frame where c is the speed of the wave, but this is clearly not the case for $\sigma \neq 0$. Our main goal is to understand these pictures, which leads to the following natural questions:

[1] Nagumo is just one of many names attached to this equation. It is also known as the Allen-Cahn or Huxley equation [100], while in biology it is referred to as an extended Fisher-Kolmogorov equation with strong Allee effect [13] and in chemistry the term Schlögel model is used. The Nagumo equation is a specific example of a bistable reaction-diffusion equation.

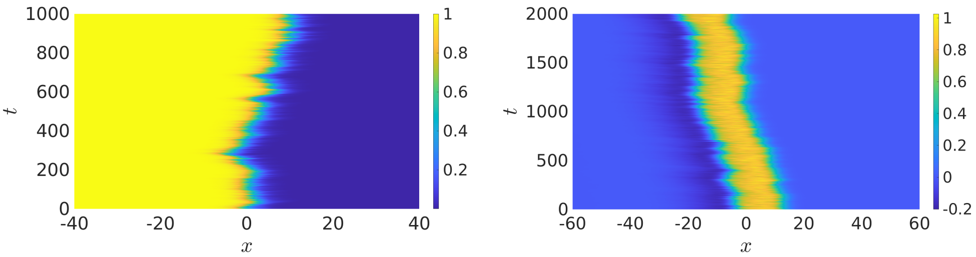

(a) (b)

Figure 1.1: Figure (a) shows a realization of equation (1.0.1) in the Stratonovich interpretation for $\rho = 1$, $a = 0.25$ and $\sigma = 0.5$, in a frame that moves with the deterministic speed. A clear drift to the right is visible. Figure (b) shows a realization of equation (1.0.2) for $\rho_1 = 1$, $\rho_2 = 0.01$, $a = 0.1$, $\varepsilon = 0.01$, $\gamma = 5$ and $\sigma = 0.1$ again in a frame that moves with the deterministic speed. This time a drift to the left is visible.

- Where is the pulse going? Is it drifting in a certain direction, and with which speed?
- What is the average shape of the pulse?
- How long can the pulse survive before the noise destroys it?

In order to answer these questions rigorously, we first need tools to understand equations such as (1.0.1) and (1.0.2). What is this object ξ? And how would we define a solution to these equations? We will introduce the necessary concepts to understand these equations in the next section. In the rest of this introduction, we will give a short introduction to the classic deterministic theory on travelling waves, followed by an overview of what is know for stochastic travelling waves. Then we will discuss the main difficulties in understanding stochastic travelling waves, which will explain why this thesis has become so heavy. We will end this introduction with a reading guide and an outlook.

1.1 Infinite dimensional stochastic integration

In the mathematical literature, equations such as (1.0.1) are often written[2] as

$$dU = [\rho U_{xx} + f(U)]dt + \sigma g(U) dW_t^Q. \tag{1.1.1}$$

This notation is a shorthand for

$$U(t) = U(0) + \int_0^t \rho U_{xx}(s) + f(U(s))ds + \sigma \int_0^t g(U(s)) dW_s^Q. \tag{1.1.2}$$

[2] We will write deterministic equations and stochastic equations from the physics literature with small letters, but equations in Itô or Stratonovich interpretation always with capitals.

The question that arises immediately is ‘what is W_t^Q and how is it related to $\xi(x,t)$?’ To answer this question we need the concept of a stochastic integral in an infinite dimensional Hilbert space, but that is not an easy task. The fact that this subject is hard, even if one understands ordinary Itô integrals, explains why Chapters 2 and 3 are written for a single Brownian motion. We first had to study the travelling waves in a simpler framework before we could lift it to the full infinite dimensional case. The integration theory that we will outline below is not new, but readable introductions are scarce. A classic work on the subject is the book[3] by Da Prato and Zabchyk [92], and a more detailed study of Gaussian processes, the construction of the stochastic integrals and applications to SPDEs can be found in [77, 93]. A slightly lighter introduction to start with is [23]. An extensive survey of different approaches to stochastic integrals is given in [29] and a description of the specific type of noise that will become important in Chapter 4 can be found in [90]. It is also worthwhile to study Martin Hairer’s lecture notes on SPDEs [44].

1.1.1 One dimensional Itô integral

The subject of stochastic integration theory basically evolves around the following question: Suppose I have a Brownian motion, or Wiener process, β_t and some function $f : [0,T] \times \Omega \to \mathbb{R}$. Under which conditions on f can we define

$$I(t) = \int_0^t f(s) d\beta_s \tag{1.1.3}$$

and what are the properties of $I(t)$? First we should notice is that β, f and I are all functions from $[0,T] \times \Omega$ to $\mathbb{R}$. This means that in principle we should write $I(t,\omega)$ and $f(s,\omega)$, but we will suppress the ω-dependence to avoid the notation becoming to cluttered as is common in the literature. There are multiple ways to define the stochastic integral above, most notably the Itô integral and the Stratonovich integral. Throughout this thesis we will use the Itô interpretation of the integral as it has better properties from a mathematical viewpoint, but the Stratonovich interpretation is often more physically relevant [110]. We will not construct the Itô integral, as there are many decent books on this subject [88], but we will record the most important results here.

Theorem 1.1.1. *Take a Brownian motion β_t on $[0,T]$ together with its natural filtration $\mathcal{F}_t$. Then, for every $f \in L^2([0,T] \times \Omega, \mathbb{R})$ that is adapted to $\mathcal{F}_t$, the Itô integral*

$$I(t) = \int_0^t f(s) d\beta_s \tag{1.1.4}$$

is well defined in $L^2([0,T] \times \Omega, \mathbb{R})$. Furthermore, $I(t)$ is a martingale with respect to $\mathcal{F}_t$, meaning that for all $0 \le s \le t \le T$ we have

$$E[I(t)|\mathcal{F}_s] = I(s) \tag{1.1.5}$$

[3] Be sure to get the updated second edition.

and $I(t)$ satisfies the Itô isometry:

$$E\left[\left(\int_0^t f(s)d\beta_s\right)^2\right] = E\left[\int_0^t f(s)^2 ds\right]. \tag{1.1.6}$$

The martingale property and the Itô isometry will later turn out to be essential tools in studying the properties of stochastic travelling waves. The first because it ensures us that the average of all the stochastic integrals we encounter in this thesis is always 0, and the second because it will allow us to compute the average of the norm of solutions to our SPDEs.

1.1.2 Infinite dimensional Wiener processes

When dealing with SPDEs, we must choose in which space the driving noise of the equation lives and, to begin simple, we start with $\mathbb{R}$, as in the previous paragraph. For a function $f : [0,T] \times \Omega \to H$ for some separable Hilbert space H, we can easily define

$$I(t) = \int_0^t f(s)d\beta_s \tag{1.1.7}$$

by mixing the Itô integral with the Bochner integral. $I(t)$ is still a martingale (in H) and the Itô isometry now becomes

$$E\left[\|\int_0^t f(s)d\beta_s\|_H^2\right] = E\left[\int_0^t \|f(s)\|_H^2 ds\right]. \tag{1.1.8}$$

We simply replaced the Euclidean norm in $\mathbb{R}$ with the norm on H. This is the type of integral we use in Chapter 2 and 3. It allowed us to develop our methods without being too distracted by complicated computations and more abstract integration theory.

The main problem is that in this framework we only deal with noise in time, not in space, while this is often relevant in applications. In order to build a concept of a stochastic integral where the noise depends on space and time, one needs an equivalent of Brownian motion in a function space. The theory we explain here following [93] works in general for separable Hilbert spaces, but we will for clarity stick with $L^2(\mathbb{R})$ or $L^2(D)$ for $D \subset \mathbb{R}$.

We now wish to define an equivalent of β_t on $L^2(D)$. This process, which we will call a Q-Wiener process W_t^Q, should be completely decorrelated in time (i.e. white in time) but can have correlation in space (i.e. coloured in space), where the operator $Q : L^2(D) \to L^2(D)$ in the notation W_t^Q describes the correlation. We call Q the covariance operator. Mathematically speaking, this means that W_t^Q should, for all $v, w \in L^2(D)$, satisfy

$$E[\langle W_t^Q, v\rangle_{L^2(D)}\langle W_s^Q, w\rangle_{L^2(D)}] = s \wedge t\langle Qv, w\rangle_{L^2(D)}. \tag{1.1.9}$$

The question is if we can find operators Q for which this process exists in $L^2(D)$. To investigate this question, we make the strong assumption that Q on $L^2(D)$ is Hilbert-Schmidt, linear, symmetric, and positive semi-definite[4]. Then we know that Q is

[4] Note that linear, symmetric and positive semi-definite are essential for Q to make sense as a covariance operator.

compact and of trace class and there is an orthonormal basis (e_k) in $L^2(D)$ such that $Qe_k = \lambda_k e_k$ and $\mathrm{Tr}(Q) = \sum_k \lambda_k < \infty$ [94, Ch. 6]. This allows us to define (i.e. we can prove that it converges) the following process in $L^2(D)$:

$$W_t^Q = \sum_{k=0}^{\infty} \sqrt{\lambda_k} e_k \beta_k(t), \tag{1.1.10}$$

where all the $\beta_k(t)$ are i.i.d. standard Brownian motions. On the other hand, we call a process W_t^Q a Gaussian process in $L^2(D)$ with mean 0 and covariance Q when for all $w, v \in L^2(D)$ we have

$$E[\langle W_t^Q, v\rangle_{L^2(D)} \langle W_s^Q, w\rangle_{L^2(D)}] = s \wedge t \langle Qv, w\rangle_{L^2(D)}. \tag{1.1.11}$$

We can check that W_t^Q as defined in (1.1.10) is a Gaussian process by simply plugging the basis expansion into the equation above and using the fact that $E[\beta_m(t)\beta_n(s)] = t\wedge s\, \delta_{mn}$. The converse is also true. By [93, Prop. 2.1.10] we know that for every Q of trace class, the process defined by (1.1.11) can be represented by a direct sum as in (1.1.10). Hence, all Gaussian processes on $L^2(D)$ are characterized by a covariance operator with finite trace and we have an explicit representation. Therefore, the definition of the integral of f over a Q-Wiener process is now straightforward:

$$I(t) = \int_0^t f(s) dW_s^Q = \sum_{k=0}^{\infty} \sqrt{\lambda_k} \int_0^t f(s) e_k d\beta_k(s). \tag{1.1.12}$$

The stochastic integral can hence be constructed as an infinite sum of well understood 1D stochastic integrals. Following [93], we can again show that $I(t)$ is a martingale and satisfies the following version of the Itô isometry:

$$E\left[\| \int_0^t f(s) dW_s^Q \|^2_{L^2(D)}\right] = E \sum_k \lambda_k \left[\int_0^t \|f(s)e_k\|^2_{L^2(D)} ds\right]. \tag{1.1.13}$$

This shows us that the integral is defined for a wide range of functions, for example the integral is well defined for $f = 1$, even though this is not an $L^2(D)$ function when $D = \mathbb{R}$. We can define the class of allowed integrands more precisely by rewriting the Itô isometry in the following way:

$$\begin{aligned} E\left[\| \int_0^t f(s) dW_s^Q \|^2_{L^2(D)}\right] &= E\left[\sum_k \int_0^t \|f(s)\sqrt{Q} e_k\|^2_{L^2(D)} ds\right] \\ &:= E \int_0^t \|f(s)\sqrt{Q}\|^2_{HS(L^2(D))} ds. \end{aligned} \tag{1.1.14}$$

Here $\|\cdot\|_{HS(L^2(D))}$ is the Hilbert-Schmidt operator norm on $L^2(D)$. In other words, the integral is well defined when $f(s)\sqrt{Q}$ is a Hilbert-Schmidt operator on $L^2(D)$ and $\|f(\cdot)\sqrt{Q}\|_{HS(L^2(D))}$ is square integrable.

Characterizing Q Now suppose Q has an integration kernel $q(x, y)$, i.e. for $v \in L^2(D)$ we have

$$Qv(x) = \int_D q(x, y)v(y)dy. \tag{1.1.15}$$

Then it must hold that $q \in L^2(D \times D)$ [94, Ch. 6]. When we (formally) replace v and w by $\delta(\cdot - x)$ and $\delta(\cdot - y)$ in the definition (1.1.11) of a Gaussian process, we find that

$$E[\langle W_t^Q, \delta(\cdot - x)\rangle_{L^2(D)} \langle W_s^Q, \delta(\cdot - y)\rangle_{L^2(D)}] = s \wedge t q(x, y). \tag{1.1.16}$$

In applications, this definition of the covariance is often written as

$$\langle dW^Q(x, t) dW^Q(y, s)\rangle = \delta(t - s)q(x, y). \tag{1.1.17}$$

Therefore, we need to make a distinction between the covariance Q (the covariance operator) and the covariance q (the kernel of the covariance operator). There is a clear reason why the applied literature chooses to work with the definition above instead of going for the more abstract definition via Gaussian processes on Hilbert spaces. The function q describes the correlation between two points in space and therefore has a clear interpretation and could in principle be measured. However, equation (1.1.17) will never be found in this form in the applied literature, as q is always chosen to be translation invariant, i.e. $q(x, y) = q(x - y)$. Hence, in applications it is assumed that the correlation depends only on the distance between two points.

A translation invariant kernel can only be an element of $L^2(D \times D)$ when D is finite. Therefore: A Gaussian process in $L^2(\mathbb{R})$ has a covariance operator Q with a kernel that is **not** translation invariant. The conclusion is that in order to model translation invariant noise on $\mathbb{R}$, and hence model physically relevant problems such as our two main examples (1.0.1) and (1.0.2), we need to study a bigger class of stochastic processes. More precisely, we need to study processes with a covariance operator that is not necessarily of trace class. These type of processes are known as a cylindrical Gaussian process. This implies that we need to study processes in larger spaces then $L^2(\mathbb{R})$, i.e. in distribution spaces or any type of abstract completion of $L^2(\mathbb{R})$.

1.1.3 Cylindrical Wiener processes and cylindrical integrals

In the previous section, we first defined Q to be an operator of trace class, constructed a Wiener process W_t^Q and then constructed the stochastic integral. The integrability condition, $f\sqrt{Q} \in L^2([0, T]; HS(L^2(\mathbb{R}))$, is then in general easy to check because Q itself is already of trace class. But what if we drop the assumption that Q is of trace class? The Itô isometry still makes sense under the same integrability condition, as long as we write $\sqrt{Q}e_k$ for $\sqrt{\lambda_k}e_k$ because an orthonormal basis for Q does not necessarily exist anymore. However, the construction of W_t^Q does not make sense without the basis representation for Q, nor does the stochastic integral.

A rather straightforward way to solve this problem would be the following procedure: Choose an explicit distribution space that contains $L^2(\mathbb{R})$ such that you can show that the sum (1.1.10) does converge, repeat the process of constructing and classifying the Gaussian process, constructing the integral, and then project the integral back into

$L^2(\mathbb{R})$. Such an approach can be found in [90], allowing for the explicit construction of a class of translation invariant processes on $L^2(\mathbb{R})$. However, we do not have to follow this procedure, as it turns out that the stochastic integral does not depend on the choice of completion of $L^2(\mathbb{R})$ we take, and hence we do not need to specify this abstract space. To make this precise, pick a non-negative symmetric operator $Q \in \mathcal{L}(L^2(\mathbb{R}), L^2(\mathbb{R}))$, possibly not of trace class. We then write[5]

$$L^2_Q(\mathbb{R}) = Q^{1/2}(L^2(\mathbb{R})), \tag{1.1.18}$$

which is again a separable Hilbert space with inner product

$$\langle v, w\rangle_{L^2_Q(\mathbb{R})} = \langle Q^{-1/2}v, Q^{-1/2}v\rangle_{L^2(\mathbb{R})}. \tag{1.1.19}$$

Furthermore, when (e_k) is an orthonormal basis of $L^2(\mathbb{R})$, $(\sqrt{Q}e_k)$ is a basis for $L^2_Q(\mathbb{R})$. Next, we define a Hilbert space $L^2_{\text{ext}}(\mathbb{R}) \supset L^2(\mathbb{R})$ such that the embedding $L^2_Q(\mathbb{R}) \subset L^2_{\text{ext}}(\mathbb{R})$ a Hilbert-Schmidt operator from $L^2_Q(\mathbb{R})$ to $L^2_{\text{ext}}(\mathbb{R})$. Such a larger Hilbert space always exists [93, §2.5], but is not necessarily unique. This extension space is the key ingredient that allows our noise process to be rigorously constructed. Following [66, eq. (2)], we introduce the sum below in $L^2_Q(\mathbb{R})$:

$$W^Q_{t,n} = \sum_{k=0}^{n} \sqrt{Q}e_k\beta_k(t). \tag{1.1.20}$$

This sum converges for $n \to \infty$ in $L^2_{\text{ext}}(\mathbb{R})$ for every $t \geq 0$. We will refer to this limiting process W^Q_t as a cylindrical Q-Wiener process in $L^2(\mathbb{R})$, but it is a regular Wiener process in $L^2_{\text{ext}}(\mathbb{R})$. Note that we call the process W^Q_t a cylindrical process in $L^2(\mathbb{R})$, even though it does not attain values in $L^2(\mathbb{R})$.

The stochastic integral over this cylindrical process can now be defined as the standard Itô integral over the Wiener process in $L^2_{\text{ext}}(\mathbb{R})$, and then be pulled back into $L^2(\mathbb{R})$. The important fact here is that the resulting integral in $L^2(\mathbb{R})$ does not depend on the choice of $L^2_{\text{ext}}(\mathbb{R})$! The take-home message in this section is therefore as follows: Even though W^Q_t and the integral over this process need to be defined in abstract spaces, in daily practice we can ignore this and treat the cylindrical integral as an ordinary integral, as long as we realise that we are cannot simplify Qe_k to $\lambda_k e_k$. A more detailed explanation of this subject can be found in §4.5.1

Some remarks on conventions Now suppose we have the Itô integral

$$\int_0^t f(s)dW^Q_s. \tag{1.1.21}$$

In the previous paragraph, we learned that the integral is well defined when $\|f\sqrt{Q}\|_{HS(L^2(\mathbb{R}))}$ is finite. In the literature, this is often written as $\|f\|_{L^0_2(L^2(\mathbb{R}))} < \infty$, which is defined as

$$L^0_2(L^2(\mathbb{R})) = L_2(L^2_0(\mathbb{R}), L^2(\mathbb{R})) := HS(L^2_Q(\mathbb{R}), L^2(\mathbb{R})). \tag{1.1.22}$$

[5] In the literature, the pair $(L^2_Q(\mathbb{R}), L^2(\mathbb{R}))$ is often denoted as (U_0, U), but in our setting this might be confusing with the solution $U(t)$.

The space $L_2(L^2(\mathbb{R}), L^2(\mathbb{R}))$ is the standard space of Hilbert-Schmidt operators, so the superscript 0 indicates the weighting of the norm with $\sqrt{Q}$.

There are many more ways in the literature to write down the condition $\|f\|^2_{L_2^0} < \infty$. Often, integral (1.1.21) is understood as

$$\int_0^t f(s) dW_s^Q = \int_0^t \tilde{f}(s) dW_s, \tag{1.1.23}$$

where W_t is the standard cylindrical Wiener process ($Q = I$) and $\tilde{f} = f\sqrt{Q}$ has to be an element of $L_2(L^2(\mathbb{R}), L^2(\mathbb{R}))$, which is the same requirement f being an element of $L_2^0(L^2(\mathbb{R}))$. Note that in the literature, stochastic terms are often introduced as BdW_t, i.e. a HS-operator B multiplied by spacetime white noise. The integral $\int_0^t BdW_s$ is a B^2-Wiener process, so BdW_t can be understood as $dW_t^{B^2}$, but B is still called the covariance [70]. Analogous, sometimes (see e.g. [72]) q is chosen to be the kernel of $\sqrt{Q}$. If we denote the kernel of $\sqrt{Q}$ by $p(x, y)$ we can relate q and p by

$$q(x, y) = \int_{\mathbb{R}} p(x, z)p(z, y)dz, \tag{1.1.24}$$

or, in the case of translational invariant kernels, simply as

$$q = p * p. \tag{1.1.25}$$

1.1.4 Stochastic PDEs

With the theory of the previous section, we can finally understand equations such as (1.0.1) and (1.0.2) in a rigorous framework. For example, the stochastic Nagumo equation (1.0.1)

$$u_t = \rho u_{xx} + f(u) + \sigma g(u)\xi(x, t), \tag{1.1.26}$$

can now be understood to be the following equation in Itô interpretation

$$dU = [\rho U_{xx} + f(U)]dt + \sigma g(U)dW_t^Q, \tag{1.1.27}$$

or, formulated in Stratonovich interpretation, to be

$$dU = [\rho U_{xx} + f(U)]dt + \sigma g(U) \circ dW_t^Q. \tag{1.1.28}$$

The Stratonovich equation can again be written in Itô formulation [108]:

$$dU = [\rho U_{xx} + f(U) + \frac{\sigma^2}{2} q(0)g'(U)g(U)]dt + \sigma g(U)dW_t^Q. \tag{1.1.29}$$

Note that W_t^Q is now a cylindrical Q-Wiener process in $L^2(\mathbb{R})$ with translation invariant kernel $q(x - y)$:

$$Qv(x) = \int_{\mathbb{R}} q(x - y)v(y)dy. \tag{1.1.30}$$

Hence, whenever one comes across an equation such as (1.1.26), it is important to realise that it is actually a pre-equation in the sense that we first must agree what we mean by the notation. If the equation is interpreted in Stratonovich sense, we can generally agree that equation (1.1.26) means (1.1.29), but note that in applications it is often not mentioned which interpretation is used, see e.g. [3].

Another way to give meaning to equation (1.1.26) is to replace ξ by a smooth approximation, solve the equation, and then take the limit back to the nowhere differentiable ξ. It can then be shown that the limiting solution is a solution of (1.1.29), and these kind of results are known as Wong-Zakai theorems [45]. Unfortunately, this also implies that it is not obvious how equation (1.1.26) should be interpreted in the case of space-time white noise (i.e. $q(x) = \delta(x)$), and indeed, interpreting space-time white noise in Stratonovich sense needs the significantly more abstract theory of Martin Hairer's regularity structures [45].

For equations in Itô interpretation such as (1.1.27), existence results are well established by now. For the existence results in §2.4, we apply theorems from [77, 93]. The specific shape of the assumptions for these theorems explains the technical setup in §2.2. However, using the techniques for the stability proofs, we can in the future also directly prove existence results without resorting to the variational framework from [77, 93].

1.2 Reaction-Diffusion equations and travelling waves

The basic equation this thesis evolves around is given by the following reaction-diffusion equation (RDE):

$$u_t = \rho u_{xx} + f(u). \tag{1.2.1}$$

We call a wave profile Φ_0 and a speed c_0 a travelling wave[6] solution when $\Phi_0(x - c_0 t)$ is a solution to the RDE above. This is equivalent to demanding that Φ_0 is a stationary solution of the RDE written in the travelling wave coordinate $\xi = x - c_0 t$:

$$u_t = \rho u_{\xi\xi} + c_0 u_\xi + f(u). \tag{1.2.2}$$

Hence, the pair (Φ_0, c_0) is a travelling wave solution when

$$\rho \Phi_0'' + c_0 \Phi_0' + f(\Phi_0) = 0, \qquad \Phi_0(-\infty) = u^-, \, \Phi_0(\infty) = u^+, \tag{1.2.3}$$

where u^+ and u^- are zeros of f. We now need to ask the following question: Does this equation have a solution, and if yes, is it unique? Of course, this question depends heavily on the function f. For the deterministic version of equations (1.0.1) and (1.0.2), the solution exists and is unique [60, 113], while for another important equation, the Fisher-KPP equation, there are infinitely many pairs (Φ, c). Over the past decades, many more nonlinearities have been studied, see for an overview e.g. [65, 99].

In order to investigate whether or not the travelling wave, if it exists, is stable, we follow the classic approach. We assume that the solution u in the ξ frame can be split

[6] Throughout this thesis we shall always write a subscript 0 for the deterministic wave.

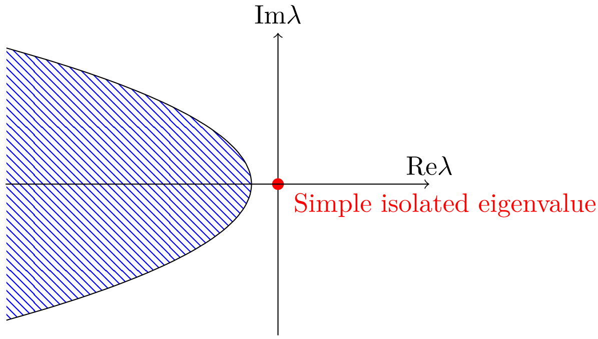

Figure 1.2: Sketch of the spectrum of the linear operator (1.2.5).

as $u(t) = \Phi_0 + v(t)$, which results in the following equation for $v(t)$:

$$v_t = \underbrace{\rho\Phi_0'' + c_0\Phi_0' + f(\Phi_0)}_{=0} + \mathcal{L}_{\mathrm{tw}}v + N(v), \tag{1.2.4}$$

where the linear operator $\mathcal{L}_{\mathrm{tw}} : H^2(\mathbb{R}) \to L^2(\mathbb{R})$ is given by

$$\mathcal{L}_{\mathrm{tw}}v = \rho\partial_{\xi\xi}v + c_0\partial_\xi v + f'(\Phi_0)v \tag{1.2.5}$$

and the nonlinearity N by

$$N(v) = f(\Phi_0 + v) - f(\Phi_0) - f'(\Phi_0)v. \tag{1.2.6}$$

In other words, we can say that (Φ_0, c_0) is chosen such that $v(t) = 0$ is a solution of equation (1.2.4). Furthermore, the travelling wave is a stable solution of the RDE when $v(t) = 0$ is a stable solution of (1.2.4).

We will now give a short overview of the main techniques we use to prove stability that will be extended to the stochastic equations in the rest of this thesis. These techniques, which we will refer to as phase tracking throughout this thesis, were developed by Zumbrun and Howard in the 90s [118], mainly for the purpose of studying a much harder problem, shock waves. Therefore, a full proof of the stability of travelling waves in bistable RDEs using their techniques is not readily available in the literature. We will write down the main Ansatz and do the calculations to describe the outline of the proof, but the technical details are omitted.

Spectral Properties The structure of the linear operator $\mathcal{L}_{\mathrm{tw}}$ as introduced in equation (1.2.5) is essential to understand the dynamics of the travelling wave. Especially, the shape of the spectrum is important to us. We always assume that the linear operator has a spectrum as depicted by Figure 1.2. This means that we always have an isolated eigenvalue in 0, corresponding to the translation invariance of the system. Indeed, a direct calculation shows that $\mathcal{L}_{\mathrm{tw}}\Phi_0' = 0$ so Φ_0' is the eigenfunction corresponding to the zero eigenvalue. We assume that the rest of the spectrum is in the left half plane, bounded away from the imaginary axis. This leads to a so called ‘spectral gap’ between the zero eigenvalue and the rest of the spectrum. The last important property of the spectrum is the fact that $\mathcal{L}_{\mathrm{tw}}$ is a sectorial operator. Using the techniques in [65], we can compute that the linearizations for equations (1.0.1) and (1.0.2) indeed have this structure.

From these properties we draw two important conclusions. First, the linear operator generates an analytic semigroup $S(t)$ [80], but more importantly, the semigroup can be split in two parts. The contour integral around the spectrum that defines $S(t)$ can be deformed into a contour integral around 0 and a contour integral around the rest of the spectrum. The integral around 0 results in the operator P, a projection operator onto Φ_0' given by $P(v) = \langle v, \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})} \Phi_0'$ for any $v \in L^2(\mathbb{R})$ with ψ_{tw} the normalized adjoint eigenfunction of Φ_0'. This means that ψ_{tw} is such that $\mathcal{L}_{\mathrm{tw}}^* \psi_{\mathrm{tw}} = 0$, $\langle \Phi_0', \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})} = 1$, where the formal adjoint $\mathcal{L}_{\mathrm{tw}}^*$ is defined as

$$\mathcal{L}_{\mathrm{tw}}^* v = \rho \partial_{\xi\xi} v - c_0 \partial_\xi v + f'(\Phi_0) v. \tag{1.2.7}$$

This has the following important consequence. Whenever a function $v \in L^2(\mathbb{R})$ is orthogonal to ψ_{tw} we find the following bound for the semigroup:

$$\|S(t)v\|_{L^2(\mathbb{R})} = \|S(t)(I-P)v\|_{L^2(\mathbb{R})} \leq M e^{-\beta t} \|v\|_{L^2(\mathbb{R})} \tag{1.2.8}$$

for some $M \geq 1$ and a $\beta > 0$ such that $-\beta$ lies in the spectral gap. All these computations can be made rigorous using [80].

Phase tracking Phase tracking is the key ingredient of this thesis. Essentially, we must be able to define a position for the travelling wave. There are many ways to define the location of the travelling wave, but it can be shown that they all result in the same qualitative behaviour [117]. The phase we choose is the one most practical for technical reasons, but also has an intuitive interpretation. We begin by introducing an Ansatz of the form

$$u\big(\cdot + \gamma(t), t\big) = \Phi_0(\cdot) + v(\cdot, t), \tag{1.2.9}$$

in which $\gamma(t)$ is the phase of u. We now demand that the evolution of the phase is governed by the ODE

$$\gamma(t) = \gamma_0 + c_0 t + \int_0^t a\big(v(s)\big) ds \tag{1.2.10}$$

for some (nonlinear) functional $a : L^2(\mathbb{R}) \to \mathbb{R}$ that we are still free to choose. The resulting equation for v is then given by

$$v_t(t) = \mathcal{L}_{\mathrm{tw}} v(t) + N\big(v(t)\big) + a\big(v(t)\big) \partial_\xi \big[\Phi_0 + v(t)\big], \tag{1.2.11}$$

which is an extended version of equation (1.2.4). This can be recast into the mild form

$$v(t) = S(t) v_0 + \int_0^t S(t-s) \Big[N\big(v(s)\big) + a\big(v(s)\big) \partial_\xi \big[\Phi_0 + v(s)\big] \Big] ds, \tag{1.2.12}$$

where $S(t)$ is the semigroup generated by $\mathcal{L}_{\mathrm{tw}}$, as discussed in the previous paragraph. To apply exponential bounds such as (1.2.8) to the semigroup S, we must avoid the neutral non-decaying part of the semigroup. In order to force the integrand to be orthogonal to the zero eigenspace, we recall that the integrand must be orthogonal to ψ_{tw} and choose

$$a(v) = -\frac{\langle N(v), \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}{\langle \partial_\xi(\Phi_0 + v), \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}. \tag{1.2.13}$$

This choice, together with an aptly chosen γ_0, defines the phase $\gamma(t)$ and ensures that $v(t)$ is orthogonal to ψ_{tw} for all time. By a standard bootstrapping procedure one can now establish the limits $\|v(t)\|_{L^2(\mathbb{R})} \to 0$ and $t^{-1}\gamma(t) \to c_0$ for $t \to \infty$, provided that the initial condition v_0 is sufficiently small. This allows us to conclude that the travelling wave is orbitally stable. By orbitally stable we mean that we have shown that perturbations around the wave die out, but they can induce a phase shift, so the solution converges to a shifted version of the wave we perturbed around.

At this point the definition for $\gamma(t)$ is rather technical. There is also a more intuitive interpretation to this definition. We need to control the growth in the direction of Φ_0'. Another way to do this is by just fixing the position of u. We can define the position of $u(\cdot, t)$ by the number $a(v(t))$ such that $\langle u(\cdot + c_0 t + a(v(t)), t) - \Phi_0, h \rangle_{L^2(\mathbb{R})} = 0$ for some reference function h. In other words, $\langle v(t), h \rangle_{L^2(\mathbb{R})} = 0$. This implies

$$\begin{aligned} 0 &= \frac{\partial}{\partial t} \langle v, h \rangle_{L^2(\mathbb{R})} \\ &= \langle \mathcal{L}_{\mathrm{tw}} v + a(v)\partial_\xi(v + \Phi_0) + N(v), h \rangle_{L^2(\mathbb{R})} \\ &= \langle v, \mathcal{L}_{\mathrm{tw}}^* h \rangle_{L^2(\mathbb{R})} + a(v) \langle \partial_\xi(v + \Phi_0), h \rangle_{L^2(\mathbb{R})} + \langle N(v), h \rangle_{L^2(\mathbb{R})}. \end{aligned} \tag{1.2.14}$$

Now to make sure that a is quadratic in v we have to choose $h = \psi_{\mathrm{tw}}$ and we find

$$a(v) = -\frac{\langle N(v), \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}{1 - \langle v, \psi_{\mathrm{tw}}' \rangle_{L^2(\mathbb{R})}}. \tag{1.2.15}$$

and this is exactly the same as we found before. Therefore, one can think of $\gamma(t)$ not just as a way to fix the position, but especially as **the** position that conveniently removes deviations, both constant and linear in v, from $c_0 t$.

Remark 1. *Suppose for the moment that we do use ψ_{tw} to define the position, but we do not know yet what (Φ_0, c_0) is. The equation for $a(v)$ would then become*

$$a(v) = -\frac{\langle \rho \Phi_0'' + c_0 \Phi_0' + f(\Phi_0), \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})} + \langle v, \mathcal{L}_{\mathrm{tw}}^* \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})} + \langle N(v), \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}{1 - \langle v, \psi_{\mathrm{tw}}' \rangle_{L^2(\mathbb{R})}}. \tag{1.2.16}$$

Hence, we could define (Φ_0, c_0) as the pair that ensures $a(0) = 0$. In the deterministic case this seems a rather cumbersome way to define (Φ_0, c_0), but in the stochastic case we a priori fix the position (be orthogonal to ψ_{tw}) and only after we did the computations decide what (Φ_σ, c_σ) must be, and then this definition becomes useful, see equation (2.2.46).

1.3 Stochastic travelling waves

Although there are not many results to be found in the mathematical literature on waves in stochastic bistable RDEs, the subject has a longer history in the physics and chemistry literature. Results for stochastic waves propagating into an unstable state (i.e. Fisher-KPP type dynamics) go back to 1989 [84] and results for bistable equations can already be found in 1991 [102]. In both cases, the driving motivation for studying stochastic waves was the desire to include thermal fluctuations into the system. In the second half of the 90s the subject of stochastic waves picked up speed and was centred around the group of García-Ojalvo in Barcelona, Schimansky-Geier in Berlin and Van Saarloos in Leiden. This resulted in a book called 'Noise in spatially extended systems' [39], which was published in 1999. In this book, basic techniques to study SPDEs are explained, and a range of specific examples is shown such as the Belousov-Zhabotinsky system, Ginzburg-Landau, Swift-Hohenberg and the bistable RDE. From a physics perspective, one could conclude that the subject has reached a certain maturity. However, mathematicians will agree with the authors that "many aspects need to be established in a more rigorous way." In this thesis we will develop techniques to study a large class of equations, including the equation studied [39, §6.1] for a limited set of parameters. The rest of the book is still open for every mathematician to explore, either by extending our techniques or by developing new ones.

1.3.1 Stochastic waves in the physics literature

The stochastic Nagumo equation as studied in [39] has the following form:

$$u_t = \rho u_{xx} + f(u) + \sigma g(u)\xi(x,t), \tag{1.3.1}$$

in which ξ is a Gaussian process that is white in time and coloured in space, i.e. it has the following properties:

$$\begin{aligned} E[\xi(x,t)] &= 0, \\ E[\xi(x,t)\xi(x',t')] &= \delta(t-t')q(x-x'). \end{aligned} \tag{1.3.2}$$

Here, q is a symmetric function describing the correlation in space. Typically, one can think of a spiked Gaussian. In this form, the equation is often referred to as a Langevin equation. In §1.1 we extensively discussed how this equation can be interpreted as an SPDE. In order to understand the observed wandering of the wavefront as in Figure 1.1, the following Ansatz is made. A phase $z(t)$ is introduced together with an unknown stochastic process $\Delta(t)$ and an average speed $\bar{c}$ such that

$$z(t) = z_0 + \bar{c}t - \Delta(t), \tag{1.3.3}$$

and then $u(t)$ is studied in the $x - z(t)$ frame. The equation is interpreted in the Stratonovich sense, so the classic chain rule still applies and the same computations as for the deterministic wave can be used. However, the Stratonovich integral over $\sigma g(u)\xi$ has an effective drift term of $\sigma^2/2q(0)g'(u)g(u)$ at lowest order, so the travelling wave equation at $\mathcal{O}(\sigma^2)$ becomes

$$\rho\Phi'' + \bar{c}\,\Phi' + f(\Phi) + \frac{\sigma^2}{2}q(0)g'(\Phi)g(\Phi) = 0, \tag{1.3.4}$$

which is an $\mathcal{O}(\sigma^2)$ perturbation of the deterministic equation (1.2.3). Now the key assumption under this approximation is that influence of the effective drift term on the front shape and the influence of the stochastic term $\Delta(t)$ on the front position can be separated at $\mathcal{O}(\sigma^2)$. In other words, the authors postulate that the $\Delta(t)$ dependence in $g(u(\cdot - \bar{c}t + \Delta(t), t))\xi$ does not influence the effective drift term at $\mathcal{O}(\sigma^2)$. Now this is not true as we will find out later in this thesis, but this does not take away the fact that equation (1.3.4) is a good approximation to the stochastic wave and this equation describes the numerical results quite accurately. See §4.3 for a more in depth comparison of equation (1.3.4) with our results.

Next, we would like to derive an equation for $\Delta(t)$. This can, at lowest order, be achieved by finding the $\mathcal{O}(\sigma)$ level of the equation for $u(x - z(t), t)$ by assuming that $\Delta(t)$ is of $\mathcal{O}(\sigma)$ and then tune $\Delta(t)$ such that the equation becomes solvable. This results in the following equation for $\Delta(t)$:

$$\dot{\Delta}(t) = \sigma \frac{\langle g(\Phi_0)\xi, \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}{\langle \Phi_0', \psi_{\mathrm{tw}} \rangle_{L^2(\mathbb{R})}}. \tag{1.3.5}$$

To characterize the wandering of the waves, a diffusion coefficient D is introduced:

$$D = \frac{E[\Delta(t)^2]}{2t}. \tag{1.3.6}$$

This coefficient can be computed using the equation for $\dot{\Delta}(t)$. As we will see in §4.2.4, this diffusion coefficient indeed describes the wandering of the wave at lowest order in σ. The diffusion coefficient is an important observable of the wave as it describes how, at lowest order, a single realization of the wave will wander about the deterministic speed. In recent work on understanding gene drift in populations, this coefficient again played an important role [13]. Also in the field of stochastic neural field equations this coefficient has been studied. See also §4.2.4 for a more detailed comparison of the drift coefficient with other results in the literature.

1.3.2 Stochastic waves in the mathematical literature

We now give a short overlook of the results in the mathematical literature. In the case of the monostable stochastic F-KPP equation, many of the main issues such as wave speed and existence of an invariant measure have been worked out by the group around Mueller. Especially, computing the equivalent of equation (1.3.6) and proving rigorously that the expansion in σ turned out to be a hard problem [85].

The first results on the bistable equation come from 2012 and are mainly numerical [79]. This result set the stage for further research, as it describes for equation (1.3.1) how the average wave speed and shape depend on the interpretation (Itô versus Stratonovich), on additive versus multiplicative noise, on the shape q of the correlation and on the intensity σ of the noise. Any successful attempt to study waves in the Nagumo equation must be able to reproduce their findings.

In the following years, Stannat and coworkers developed techniques to study travelling waves in bistable RDEs and neural field equations [72–74, 104, 105]. Building on this results Inglis and MacLaurin [57] developed techniques to study travelling waves in neural field equations. These works indeed present stability results for stochastic travelling waves, and certainly helped us gain insight into the important techniques needed for the subject. However, there are some technical and practical limitations. First, one needs an immediate contractivity condition, meaning that $M = 1$ in (1.2.8). With some effort, this can be shown for the Nagumo equation [105], but this is unknown for the FitzHugh-Nagumo equation.

Another major technical issue is that the covariance of the driving noise is assumed to be of trace class. In §1.1, we explained that this assumption destroys the translation invariance of the equation, and the physically relevant equation (1.3.1) is not included in these works. From a more practical viewpoint, we note that these approaches do not give an answer to some of our main questions, such as 'what is the average wave speed and profile' and 'how does the dynamics evolve over long timescales.' See also §2.1 for a more technical discussion of their methods.

A more recent formal approach came from Cartwright and Gottwald [19]. Using a collective coordinate approach, they came to the same stochastic wave as we will find in §2.2.4 for the Nagumo equation forced by a one dimensional Brownian motion. A major advantage of this approach is that it gives a quite direct way to compute the stochastic travelling wave, without any rigorous setup needed. However, up to now, their approach depends on the exact structure of the Nagumo equation, so it would be interesting to see if their approaches can be generalized to other equations. See for a more extensive discussion on stochastic travelling waves also the review by Kuehn [69].

1.3.3 Stochastic phase tracking

Our approach to studying stochastic travelling waves will be to extend the phase tracking approach from §1.2. Upon adding a stochastic term to the equation for $v(t)$, we directly see that we cannot choose $a(v)$ in such a way that the whole equation becomes orthogonal to ψ_{tw}. What we need is an extra degree of freedom to ensure that also the stochastic part of the equation can be chosen orthogonal to ψ_{tw}. Hence, we propose to replace the splitting

$$u(\cdot + \gamma(t), t) = v(t) + \Phi_0 \tag{1.3.7}$$

where $\gamma(t)$ satisfies the ODE

$$\gamma(t) = \gamma_0 + c_0 t + \int_0^t a(v(s))ds, \tag{1.3.8}$$

by the splitting

$$U(\cdot + \Gamma(t), t) = V(t) + \Phi_\sigma \tag{1.3.9}$$

where $\Gamma(t)$ satisfies the SDE

$$\Gamma(t) = \Gamma_0 + c_\sigma + \int_0^t a(V(s))ds + \sigma \int_0^t b(V(s))dW_s^Q. \tag{1.3.10}$$

Here we introduced a yet unknown stochastic travelling wave (Φ_σ, c_σ) and a function b that maps to a Hilbert-Schmidt operator from $L_Q^2(\mathbb{R})$ to $\mathbb{R}$. This will be made precise later, but for now, it is important to realise that, when we compute the equation for $V(t)$ using Itô calculus, the stochastic term for $V(t)$ will depend on b giving us the freedom to choose $V(t)$ orthogonal to ψ_{tw}. The advantage of splitting $U(t)$ this way is that we now, just as in the deterministic case, fully describe the dynamics of the solution along the family of travelling waves. Therefore, $\Gamma(t)$ gives us a meaningful way of defining the average speed of the wave. See §4.2.2 for an extensive discussion on choosing a, b, the wave (Φ_σ, c_σ) and what it means for the dynamics of the travelling wave.

1.4 Main difficulties

The first main difficulty is the following: If I define

$$V(t) = U(\cdot + \Gamma(t), t) - \Phi_\sigma \tag{1.4.1}$$

as we proposed in the previous section, what equation does $V(t)$ solve? Itô calculus is not suitable to directly calculate the time derivative of $U(\cdot + \Gamma(t), t)$. In §2.5 (for a single Brownian motion) and §4.5.3 (for a cylindrical Q-Wiener process) we explain how to derive an equation for $V(t)$. Especially computing the Itô correction term that describes the interaction between the stochastic terms from $U(t)$ and $\Gamma(t)$ is very subtle.

Unfortunately, the Itô calculus causes several delicate technical complications that are not observed in the deterministic setting. The first major issue comes from the Itô Isometry. In the deterministic setting, a semigroup can be used to lift a function from $L^2(\mathbb{R})$ to $H^1(\mathbb{R})$ at the cost of an integrable singularity: $\|S(t)\|_{\mathcal{L}(L^2,H^1)} \sim t^{-1/2}$. However, for stochastic integrals we run into the following problem when we apply the Itô Isometry:

$$E\|\int_0^t S(t-s)B(s)dW_s^Q\|_{H^1}^2 = E\int_0^t \|S(t-s)B(s)\|_{HS(L_Q^2,H^1)}^2 ds. \tag{1.4.2}$$

If we wish to estimate B on the right hand side in $L^2(\mathbb{R})$, we have to square the $(t-s)^{-1/2}$-singularity, which is unintegrable. This precludes us from estimating $E\|V(s)\|_{H^1}$ directly, but we are forced to study the integrated $H^1(\mathbb{R})$-norm. This extra integral gives us the possibility to get rid of the singularity, which is still anything but easy; see the sections on nonlinear stability §2.9, §3.5 and §5.5.

A second major complication is that stochastic phase-shifts lead to extra nonlinear diffusive terms. By contrast, deterministic phase-shifts such as (1.2.10) lead to extra convective terms, which are of lower order and hence less dangerous. As a consequence, we encounter quasi-linear equations in our analysis that do not immediately fit into a semigroup framework. Luckily, this extra nonlinear term is a scalar so we solve this problem by using a suitable stochastic time-transform to scale out this extra term.

However, scaling only works when the term in front of ∂_{xx} is a scalar, which is clearly not the case for equations with different diffusion coefficients such as (1.0.2). This causes the main divide between Chapter 2 and 3. In Chapter 2 we have to enforce that all diffusion coefficients are equal, excluding the FitzHugh-Nagumo equation. We solve this issue in Chapter 3 by scaling each component of the SPDE for V separately and carefully studying the off-diagonal elements of the semigroup.

A third major difficulty is proving results on long timescales. In the deterministic case, stability is a binary question: given an initial perturbation from the wave, you either converge back to (a translate of) the wave or not. In the stochastic case this is not true. At every moment in time the wave is pushed by the noise term, hence the probability of making a large excursion from the wave, i.e. becoming unstable, grows in time. The major question is if we can understand how fast this probability grows. Numerically it is straightforward to show that the supremum of deviations from the travelling wave grows logarithmically in time, but proving this cost us many grey hairs. A proof of the logarithmic growth bound can be found in Chapter 5.

1.5 Reading guide

One hardly ever reads a math paper from beginning to end in a linear fashion. Especially a thesis comprised of N papers ($N > 1$) should also not be attacked head on. The papers are presented here in a chronological way (of writing that is, not of publication) which makes sense because each new paper builds on the previous one. However, in the time that past between the first and the fourth paper, our capability to give a bird's-eye view on the subject grew, as well as our capability to write cleaner proofs. Therefore, Chapter 4 is a good place to start, as it directly describes the physically relevant equations such as (1.3.1) and gives an overview of our procedures and results. If you have just an hour to spare on this thesis, start with §4.1-§4.4.

After (or before?) Chapter 4 the nitty-gritty details start. Chapter 2 is the most detailed of all. This chapter is the starting point for all the other chapters. We deal with questions such as existence of solutions to the SPDE, existence of solutions to the SDE that describes the phase, existence and uniqueness of the stochastic travelling wave, we show how to compute the stochastic phase shift and how to compute the stochastic time transform needed for the stability analysis. After these preparatory computations, we prove the stability results.

There are two main restrictions on the results in Chapter 2. First, we needed that the diffusion coefficients were the same in all components. Lifting this strong restriction in order to study FitzHugh-Nagumo equations such as (1.0.2) is the main purpose of Chapter 3. Many of the details from the first chapter are not repeated and we focus on the techniques necessary to prove the results. The second restriction in Chapter 2 (which also applies to Chapter 3) is the fact that the equations are forced by a one-dimensional Brownian motion, and therefore these chapters can not directly deal with equations such as (1.0.1) and (1.0.2). As mentioned before, this was a deliberate choice because we had not yet developed the intuition needed to work with cylindrical Q-Wiener processes.

In Chapter 4 we develop the necessary tools to work with translation invariant processes, but also extensively discuss perturbation methods to approximate the stochastic wave and make comparisons with numerical approximations. These numerical results however also indicate that our results remain valid on exponentially long timescales which is a significantly stronger claim than what we prove in Chapters 2-4. Improving our results to include the exponentially long bounds is the subject of Chapter 5. It can be read directly after Chapter 4, but understanding the stability analysis in §2.9 will serve as a good basis to appreciate the delicacies in the proofs in Chapter 5. Studying the stability analysis for the FitzHugh-Nagumo equation in §3.5 also explains why we chose to set up the proof in Chapter 5 only for equations of Nagumo type.

Notation It is important to note that the four chapters to follow do not have a uniform notation and spelling. This introduction and Chapters 4 and 5 are written in British English while the other chapters are written in American English, although this separation is not always very strict. In Chapter 2, we write A_* for the operator $\rho\partial_{xx}$ to highlight the fact that we treat the operator as an operator from $H^1(\mathbb{R}) \to H^{-1}(\mathbb{R})$, in stead of $H^2(\mathbb{R}) \to L^2(\mathbb{R})$. Also note that in Chapters 2 and 3, Q is not the covariance operator, but defined as $Q = I - P$ where P is the projection onto the 0-eigenspace as defined in §1.2. As Chapters 2 and 3 are written for a single Brownian motion, there is within the chapters no chance of confusion.

1.6 Outlook

In this thesis we develop techniques for quite a broad class of RDEs, but our structural assumptions rule out two major classes of equations. In the first place, our techniques heavily depend on the spectral gap in the spectrum of the linear operator (see Figure 1.2). While this holds true for bistable equations, the essential spectrum typically touches the imaginary axis for waves in monostable equations such as the stochastic Fischer-KPP equation [85]:

$$du = [u_{xx} + u(1-u)]dt + \sigma\sqrt{u(1-u)}dW_t^Q \tag{1.6.1}$$

For the deterministic equation many waves exist, but compact initial conditions always converge to the wave with speed 2 [100]. Note that the multiplicative term $\sqrt{u(1-u)}$ is not Lipschitz in 0 and 1, which makes analysis much harder. In the stochastic setting, it was shown in [85], using a completely different class of techniques than ours, that the speed 2 is perturbed into

$$2 - \pi^2|\log(\sigma^2)|^{-2} + \mathcal{O}(\log|\log(\varepsilon)|\log(\varepsilon)|^{-3}). \tag{1.6.2}$$

This odd looking result, known as the Brunet-Derrida conjecture [17], deviates significantly from the perturbation results we will encounter in this thesis. It would be interesting to see if we could derive the same results. We could, for example, pose this equation in a weighted space to open up the spectral gap again. However, this would exclude perturbations with a heavy tail which might essentially contribute to the result above. Another option would be to split the semigroup not in an exponentially decaying and a nondecaying part (which is now impossible) but, inspired by [117], split the semigroup in an exponentially decaying part, a polynomially decaying part and a nondecaying part. This would probably require an extensive extension of our phase tracking Ansatz (1.3.10).

Another obstruction in our techniques is the fact that our methods rely on the smoothening properties of the semigroup. For example, we cannot yet directly deal with the classical FHN-equation (equation (1.0.2) with $\rho_2 = 0$) or with neural field equations. We expect for the FHN-equation that our perturbation results still hold for $\rho_2 \to 0$, but all our bounds will blow up.

Of course, there are many more RDEs that are well studied in the deterministic case but their stochastic counterparts not so much. Korteweg-de Vries, Gray Scott, Klausmeier, Burgers, to just name a few. This list shows that in the coming years many more exciting results can be expected!