

A control version of the Freidlin-Wentzell Theorem

Pratik Worah

Google Research

New York, USA

PWORAH@GOOGLE.COM

Abstract

The Freidlin-Wentzell theorem gives the probability of large deviations under small noise. Very roughly, it gives the probability that a particle undergoing diffusion inside a box, under very small but positive noise, will eventually hit the boundary of the box in a given time. Exit always happens at a particular point irrespective of the starting position, and this point of exit and time to exit can be calculated using the drift and volatility terms. We present an analogous control version of the theorem, where the goal of the controller is to keep the particle inside the box, as long as possible.

1. Introduction

Typically stochastic control problems aim to place bounds on the expectations of linear or quadratic functions of some underlying stochastic process. Theoretical progress in non-linear control has been limited, although much has been done in the second half of last century (see for example the book by Bensoussan (1992)). The use of expectations of linear and quadratic functions as the objective is more natural as opposed to optimizing the underlying tail probabilities directly. The latter problem seems theoretically more difficult. Moreover, diffusions tend to spread out rather quickly so there is little to gain, as far as applications go, in investigating incremental gains by using control variables and optimizing the tail probability, especially since the events themselves are rare events.

However, in the small noise case i.e., when the diffusion looks like:

$$dx_t^\varepsilon = b(x_t^\varepsilon)dt + \sqrt{\varepsilon}dW_t, \tag{1}$$

where W_t is standard Brownian motion in the plane and $\varepsilon \rightarrow 0+$ i.e., ε small but remains positive, the diffusion (mostly) stays in a narrow tube for a long time i.e., it does not spread out much. Moreover, the probability that such a diffusion exits its bounding box is closely related to where it exits the box, and surprisingly there is a unique point of exit, irrespective of the starting point. Furthermore, this unique point of exit is used to explicitly calculate the time to exit from the given domain (our box). This is the well-known Wentzell-Freidlin theory (see Freidlin and Wentzell (1991) for a formal treatment). In this short paper, we investigate a very simple control version of the large deviations theorem in Wentzell-Freidlin theory (Theorem 1), and its implications to the point of exit and exit time (Lemma 7 and 8).

2. Background and Definitions

We refer the reader to the book Fleming and Rishel (1975) for a background in control theory and the book Varadhan (2017) for a quick background on large deviations theory and

Wentzell-Freidlin theory. The book Giaquinta and Hilderbrandt (1996) is a good reference on variational methods.

Optimal control in the stochastic setting has been extensively studied (see for example the books Fleming and Rishel (1975) and Bensoussan (1992)). The problem is simply stated by starting with an Itô process,

$$dx_t = A(t)x_t + B(t)\theta_t + C(t)dt + adW_t, \quad t \in [0, T]$$

where W_t is Brownian motion in \mathbb{R}^d with covariance matrix a , θ_t is adapted to the natural filtration. The objective is typically to minimize a quadratic cost function of the form:

$$\mathbb{E} \left[\int_0^T (x_t^T M x_t + \theta_t^T N \theta_t) dt + x_T^T M_T x_T + \theta_T^T N_T \theta_T \right].$$

An explicit solution to the problem is known via dynamic programming and by now it's standard in text books. The situation for non-linear control is very different, little is known systematically. The book Bensoussan (1992) and the references therein, do investigate some eclectic cases which permit a solution. In particular, they also study a selected case where the signal has large noise but the observation is required to have small noise. However, their objective is to approximate the filtered process and is therefore different from ours, which is to control exit times and bound the probability of exit in the small noise case using variational methods.

3. Our Result

Notation: Throughout we will let $\dot{f}(t)$ denote derivative with respect to time and $f'(g(x))$ will be used to denote derivative of f with respect to its argument.

We make the following assumptions, some are necessary and others serve to simplify the presentation. First, we need some smoothness assumptions for our theorem to hold.

Assumption 1 *In what follows we will assume that the domain D is smooth, convex and bounded. The paths $u(t)$, $u : [0, T] \rightarrow D$ are Lipschitz continuous with constant c .¹*

The following assumption will ensure concavity of a and convexity of the rate functional in Theorem 1. Although, these assumptions can be relaxed a bit, something similar will be necessary absent new ideas to prove the required minimax inequality.

Assumption 2 *In the following a will denote a 2×2 diagonal matrix in \mathbb{R} . For $i \in \{1, 2\}$, $a_{ii}^{-1} : D \rightarrow \mathbb{R}$ are C^∞ with derivatives bounded by c in D , and they are positive, bounded away from zero and finite. Furthermore, assume that:*

1. $(a_{ii}^{-1})''$ are not positive.
2. a^{-1} is large enough, i.e., for λ being the minimal eigenvalue of $W^{1,2}(D)$:

$$\min_D a_{ii}^{-1} \geq c^2 + c^3 \lambda + c^2 \lambda. \quad (2)$$

1. As opposed to just Hölder continuous

The following strong assumption simplifying the drift term b is not conceptually necessary, it can be replaced with assumptions on the norm of derivatives of b , but it simplifies our presentation greatly.

Assumption 3 *We will assume that $b \equiv (b_1, b_2)$, $b_i : D \rightarrow D$, is simply the drift of an OU process with $b_i(x) := -c'x_i$, for some suitably small and positive c' .*

Our main result is an extension of the large deviation principle underlying the bounds in Wentzell-Freidlin theory.

Theorem 1 *For any positive ε , given an Itô process:*

$$dx_t^\varepsilon = -c'x_t^\varepsilon dt + \sqrt{\varepsilon}a^{\frac{1}{2}}(\theta(x_t^\varepsilon))dW_t, \quad (3)$$

where W_t is standard Brownian motion in \mathbb{R}^2 . Define

$$I_T(u, \theta) := \frac{1}{2\varepsilon} \int_0^T \sum_{i \in \{1,2\}} a_{ii}^{-1}(\theta(u(t))) (\dot{u}_i(t) + c'u_i(t))^2 dt, \quad (4)$$

where we assume $\theta \equiv (\theta_1, \theta_2)$, $\theta_i : \mathbb{R} \rightarrow (0, 1)$, is Lipschitz continuous with coefficient c . Then, under Assumptions 1, 2 and 3, we have:

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log \mathbb{P}(x_T^\varepsilon \notin D) = - \inf_{u(T) \notin D} \sup_{\theta} I_T(u, \theta). \quad (5)$$

The idea behind the above theorem is quite simple: Show that the functional $I(u, \theta)$ is convex in u and concave in θ . This would imply the existence of a saddle point using Ky Fan's theorem from game theory. However, proving that a functional is convex is not easy, especially when it consists of a mix of terms with \dot{u} and u – this is why the convexity method in calculus of variations is not as widely useful – most such functionals do not have much structure in them to work with. In our case, we do have some structure and we make simplifying assumptions (Assumptions 2 and 3 above) to bring it out. Using Poincaré inequality we are able to show that our functionals are indeed convex and concave. This leads to the above large deviation bound. Using this bound and referring to existing arguments from Wentzell-Freidlin theory, we are able to show bounds in Lemma 7 and 8 that such a controlled process would have a larger exit time. When equilibrium exists, the expected exit increases from: $\inf_{y \in \partial D} \inf_{\substack{u, u(0)=x_0 \\ u(T)=y}} I(u)$ to $\inf_{y \in \partial D} \inf_{\substack{u, u(0)=x_0 \\ u(T)=y}} \sup_{\theta} I(u, \theta)$.²

4. Proofs and Calculations

4.1 Existence of Nash equilibrium

We begin by recalling a large deviation inequality implied by the Freidlin-Wentzell theorem.

2. We assume that once the particle penetrates D , it stays there. This is because we need to assume here and throughout that T is large, larger than the exit time i.e., $T \rightarrow \infty$.

Theorem 2 *Freidlin and Wentzell (1991), Varadhan (2017)* For any positive ε , given an Itô process:

$$dx_t^\varepsilon = b(x_t^\varepsilon)dt + \sqrt{\varepsilon}a^{\frac{1}{2}}(x_t^\varepsilon)dW_t, \quad (6)$$

where a (a is positive definite) and $b : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are uniformly Lipschitz continuous and W_t is standard Brownian motion in \mathbb{R}^2 , define the rate function:

$$I_T(u) := \frac{1}{2} \int_0^T \frac{1}{\varepsilon} \langle a^{-1}(u(t))(\dot{u}(t) - b(u(t))), (\dot{u}(t) - b(u(t))) \rangle dt, \quad (7)$$

where \dot{u} is Hoelder continuous, then on the Banach space of smooth functions with bounded sup-norm, for every open set G and closed set F ,

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log P(x_t^\varepsilon \in F, t \in [0, T]) \leq - \inf_{u \in F} I_T(u), \quad (8)$$

$$\liminf_{\varepsilon \rightarrow 0} \varepsilon \log P(x_t^\varepsilon \in G, t \in [0, T]) \geq - \inf_{u \in G} I_T(u). \quad (9)$$

In the case of a smooth closed bounded domain $D \subset \mathbb{R}^2$, suppose b is such that the dynamical system $\dot{u}(t) = b(u(t))$ has a stable fixed point in D . Theorem 2 then says that the probability of penetration through D in time T asymptotically equals $e^{-\inf_u I(u)}$ (where we have dropped the time subscript from I for brevity). Moreover, as $T \rightarrow \infty$, exit always happens with probability 1, for any positive ε , as long as determinant of a is bounded away from 0 (see for example Varadhan (2017)).

Consider now a simple parametrization of a in Theorem 2 with function $\theta : D \rightarrow (0, 1)^2$, so that $a(x_t^\varepsilon)$ in Equation 6 is replaced with $a(\theta(x_t^\varepsilon))$ i.e., θ will be our control function.

Remark 3 *Note that we restrict our control here to just the volatility term as opposed to the drift term or both. However, it is possible to consider controlling both the drift and volatility terms in this framework, as long as some restrictions and trade-offs are followed. The calculations for the second variation and the inequalities become longer and trickier in that case and we plan to investigate them in a full version of the paper.*

If we wanted to pick a θ so that the probability to exit in time T is minimized and therefore the time to exit maximized then we would immediately obtain from Theorem 2 an expression for the probability of exit as:

$$P(x \notin D) \simeq e^{-\sup_\theta \inf_u I_\theta(u)}. \quad (10)$$

However, if one closely follows the proofs in Freidlin and Wentzell (1991) or Varadhan (2017), then this naive bound implies that the particle penetrates D at an expected exit time $\sim e^{\frac{1}{\varepsilon} \sup_\theta \inf_u I_\theta(u)}$ as opposed to $e^{\frac{1}{\varepsilon} \inf_u I(u)}$. However, there is room for improvement in these bounds, and in particular the question is: Can we reverse the supinf to obtain a better guarantee under reasonable conditions?

A natural way to reverse the sup inf is to use Ky Fan's minimax theorem which gives sufficient conditions for the existence of Nash equilibria in two-person zero-sum games. A version of the theorem is stated below.

Theorem 4 *Raghavan (1994)* Let X, Y be compact convex subsets of locally convex topological vector spaces. Let $K : X \times Y \rightarrow \mathbb{R}$ be continuous. For any fixed \bar{x}, \bar{y} , let $K(\bar{x}, y) : Y \rightarrow \mathbb{R}$ be a convex function and $K(x, \bar{y}) : X \rightarrow \mathbb{R}$ be a concave function. Then

$$\sup_{x \in X} \inf_{y \in Y} K(x, y) = \inf_{y \in Y} \sup_{x \in X} K(x, y).$$

In order to use Theorem 4, we will need to show that our valuation functional $I(u, \theta)$ is convex with respect to u and concave with respect to θ in D , under the assumptions in Theorem 1 on the process parameters a and b .

We begin by investigating the convexity of $I_\theta(u)$ when θ is kept fixed and u is varied. Recall from the calculus of variations (see for example Giaquinta and Hilderbrandt (1996)) that the non-negativity of the second variation with respect to u is enough to show convexity of I . In other words, we need to show:

$$\delta^2 \left(\int_0^T \sum_{i \in \{1,2\}} a_{ii}^{-1}(\theta_i(u_i(t))) (\dot{u}_i(t) + c' u_i(t))^2 dt \right) \geq 0. \quad (11)$$

By definition of the second variation, checking convexity then equals verifying that the second variation of the Lagrangian is non-negative i.e.

$$\int_0^T \delta^2 \left(\sum_{i \in \{1,2\}} a_{ii}^{-1}(\theta_i(u_i(t))) (\dot{u}_i(t) + c' u_i(t))^2 \right) dt \geq 0. \quad (12)$$

We will verify the case for one coordinate dimension ($i = 1$), the other case is proved similarly. Therefore, we have to compute the second variation with respect to u for the following three types of terms:

1. $\delta^2(a_{11}^{-1} \cdot \theta_1)(u_1) \dot{u}_1^2$,
2. $\delta^2(a_{11}^{-1} \cdot \theta_1)(u_1) b_1(u_1)^2$, and
3. $\delta^2 2(a_{11}^{-1} \cdot \theta_1)(u_1) \dot{u}_1 \cdot b_1(u_1)$.

In all three cases, despite our simplified control setting the expressions involved are tedious to work with. However, assuming that $b_i(u_i) \equiv -c' u_i$ where c' is very small allows us to ignore cases 2 and 3 and concentrate on term 1.

The second variation from term 1 can be written as:

$$(a_{11}^{-1} \cdot \theta_1)(u_1) \dot{\phi}^2 + (a_{11}^{-1} \cdot \theta)''(u_1) \dot{u}_1^2 \phi^2 + 2(a_{11}^{-1} \cdot \theta)'(u_1) \dot{u}_1 \phi \dot{\phi}, \quad (13)$$

where \cdot denotes composition and $\phi \in C^1([0, T], D)$ (functions with continuous first derivatives).

In order to show that the integral corresponding to Equation 13 is non-negative, observe that the first term is positive by positivity of a . We will upper bound the contribution of the middle term and the last term using Poincaré inequality.

Since $\phi(0) = x_0$ and ϕ is assumed to have continuous first derivatives and its domain is bounded, so $\phi \in W_0^{1,2}$, and we can apply Poincaré inequality, to get:

$$\int \phi^2 dt \leq \lambda \int \dot{\phi}^2 dt, \quad (14)$$

where λ is the minimal eigenvalue for the Laplacian in $W_0^{1,2}([0, T], D)$. Therefore, we can use it to bound the middle-term in Equation 13.

Moreover, observe that $2\phi\dot{\phi} \leq (\phi^2 + \dot{\phi}^2)$, and so we can bound the last term in Equation 13 using a combination of Poincaré inequality together with our assumptions on a and θ .

Under the assumption that a is bounded away from zero and is uniformly Lipschitz continuous with bounded first derivatives i.e. Assumption 2,

$$\|(a_{11}^{-1} \cdot \theta)''\|_{\infty} \leq c, \quad (15)$$

$$\|(a_{11}^{-1} \cdot \theta)'\|_{\infty} \leq c. \quad (16)$$

Since u is Lipschitz continuous with coefficient c , under Assumptions 1,2 and 3 we have:

$$\int a_{11}^{-1}(\theta(u_1))\dot{\phi}^2 \geq \int (a_{11}^{-1} \cdot \theta)''(u_1)\dot{u}_1^2\phi^2 + 2(a_{11}^{-1} \cdot \theta)'(u_1)\dot{u}_1\phi\dot{\phi}. \quad (17)$$

This is because the coefficients of ϕ^2 and $\dot{\phi}^2$ in the RHS integral can be bounded by $c^3\lambda$ and $c^2 + c^2\lambda$ respectively.

Therefore, we have shown the convexity of the second variation of I with respect to u_1 .

In order to show concavity of the second variation with respect to θ_1 , we merely need to show that

$$\int (a^{-1})''(\theta_1(u_1))\phi^2(\dot{u}_1 + c'u_1)^2 dt \leq 0, \quad (18)$$

and this follows immediately from our assumption that $(a^{-1})''_{11} \leq 0$.

A similar set of required inequalities holds in the other coordinate as well, and since we have assumed that the correlations are zero we have verified the concavity and convexity conditions of Ky Fan's minimax theorem. Since our family of functions u and θ are assumed to be Lipschitz continuous, we also have compactness of underlying spaces, therefore:

$$\sup_{\theta} \inf_u I_{\theta}(u) = \inf_u \sup_{\theta} I_{\theta}(u). \quad (19)$$

Remark 5 *Although we have assumed the correlation to be zero for simplicity, the calculations work out when the correlation a_{12}^{-1} is assumed to be positive and constant. When the correlation is negative and large enough, the second variation is no longer non-negative – an equilibrium may not exist!*

Remark 6 *Note that in principle, we have still have excess convexity and concavity under appropriate assumptions; so we can add a control to b also. It would be an interesting although lengthier calculation, left for a full version of the article.*

Let θ^* and u^* be the equilibrium point, as obtained above. Note that we can find these quantities by solving two *linear* Euler-Lagrange equations obtained from the optimization problems:

$$\inf_u \int_0^T \delta^2 \left(\sum_{i \in \{1,2\}} a_{ii}^{-1}(\theta_i^*(u_i(t))) (\dot{u}_i(t) + c' u_i(t))^2 \right) dt, \quad (20)$$

$$\sup_\theta \int_0^T \delta^2 \left(\sum_{i \in \{1,2\}} a_{ii}^{-1}(\theta_i(u_i^*(t))) (\dot{u}_i^*(t) + c' u_i^*(t))^2 \right) dt. \quad (21)$$

We have now shown Theorem 1.

4.2 Point of exit

Finally, we compare the point of exit with a control under the assumption of existence of a Nash equilibrium, to the point of exit without any control. The proof is a straight-forward modification of the usual proof of the case without control (see for example Varadhan (2017) or Freidlin and Wentzell (1991)), the only difference being that now our large deviation rate function and therefore quasi-potential is different. The probability of penetrating D in time T now becomes:

$$\mathbb{P}(x \notin D) \simeq e^{-\inf_u \sup_\theta I(u, \theta)}. \quad (22)$$

Lemma 7 *Let*

$$V(x) := \inf_{\substack{u, u(0)=x_0, \\ u(T)=x}} \sup_\theta I(u, \theta), \quad (23)$$

where $V(x)$ is such that it has a unique minimum at y_0 on the boundary ∂D . Then irrespective of the starting point x_0 , exit will take place at y_0 with probability almost 1.

Similarly, following Freidlin and Wentzell (1991) one can bound the expected time to exit D as follows.

Lemma 8 *The expected time to exit D for our controlled process x_t^ε with start position x_0 is*

$$\varepsilon \log \mathbb{E}(\tau_D) \simeq \inf_{y \in \partial D} V(y, x_0), \quad (24)$$

where $V(x)$ is given by Equation 23.

5. Applications

We provide context for two applications. Further details will appear in a longer version of this paper.

First, in medicine: Often diseases constitute of interacting families of pathogens. For example, tumours consist of several kinds of cells. In particular, there is a class of cells which exhibits the Warburg effect (lactate producers) and those that exhibit the Reverse Warburg effect (lactate consumers). Each cell type has its own specific drug and the cell types may switch among themselves slowly. The goal of drugs i.e., the control, is to keep the

population in check i.e., keep the particle in a box! In theory, one can model the growth of malignant cells using in vitro data. The existence of a Nash equilibrium, as discussed above, in such a model would imply that one can keep the population of malignant cells in check for much longer, if one uses the two specific drugs based on the optimal control strategy from Equations 20 and 21. We note that optimal control in medical oncology has been studied in context of other scenarios (see Schaettler and Ledzewicz (2010)). However, the case of modelling the effect of different malignant cell types has not been dealt with before. But it seems relevant and we intend to revisit this case in a longer version of the paper.

Second, in financial economics: Often pairs trading strategies involve continuously trading a combination two co-integrated securities such that the resulting trade maintains very small (but non-zero) variance around a mean for a period of time. One wants this period to be long because it's in this period that one buys low and sells high (hopefully!). Problems arise when the portfolio moves too far away from its mean. One solution is to reweigh the portfolio holdings and therefore variance based on some control variables (see for example Fouque et al. (2015) for a related example of applications of optimal control to finance) to keep its value near the mean – again an instance of keeping a particle in a box!

Surely there are many other such problems and working out the above calculations in more cases may lead to more real world applications as well.

Acknowledgements

I am grateful to Raghu Varadhan for many helpful discussions. Without his gentle guidance, I wouldn't have enjoyed learning about Wentzell-Freidlin theory (among many other things), as much as I did.

References

- Alain Bensoussan. Stochastic Control of Partially Observable Systems. *Cambridge University Press*, 1992.
- Wendell Fleming and Raymond Rishel. Deterministic and Stochastic Optimal Control. *Springer*, 1975.
- Jean-Pierre Fouque, Andrew Papanicolaou, and Ronnie Sircar. Filtering and Portfolio Optimization with Stochastic Unobserved drift in Asset Returns. *Communications in Mathematical Sciences*, 2015.
- Mark Freidlin and Alexander Wentzell. Random Perturbations of Dynamical Systems. *Springer*, 1991.
- Mariano Giaquinta and Stefan Hildebrandt. Calculus of Variations 1. *Springer*, 1996.
- T.E.S. Raghavan. Zero-Sum Two-Person Games. *Handbook of Game Theory*, 1994.
- Heinz Schaettler and Urszula Ledzewicz. Optimal Control for Mathematical Models of Cancer Therapies. *Springer*, 2010.
- S.R.S. Varadhan. Large Deviations. *AMS Courant Lecture Notes*, 2017.