



First-Hitting Times Under Additive Drift

Timo Kötzing and Martin S. Krejca^(✉)

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
martin@krejca.de

Abstract. For the last ten years, almost every theoretical result concerning the expected run time of a randomized search heuristic used *drift theory*, making it the arguably most important tool in this domain. Its success is due to its ease of use and its powerful result: drift theory allows the user to derive bounds on the expected first-hitting time of a random process by bounding expected local changes of the process – the *drift*. This is usually far easier than bounding the expected first-hitting time directly.

Due to the widespread use of drift theory, it is of utmost importance to have the best drift theorems possible. We improve the fundamental additive, multiplicative, and variable drift theorems by stating them in a form as general as possible and providing examples of why the restrictions we keep are still necessary. Our additive drift theorem for upper bounds only requires the process to be nonnegative, that is, we remove unnecessary restrictions like a finite, discrete, or bounded search space. As corollaries, the same is true for our upper bounds in the case of variable and multiplicative drift.

1 Drift Theory

In the theory of randomized algorithms, the first and most important part of algorithm analysis is to compute the expected run time. A finite run time guarantees that the algorithm terminates almost surely, and, due to Markov's inequality, the probability of the run time being far larger than the expected value can be bounded, too. Thus, it is important to have strong and easy to handle tools in order to derive expected run times. The de facto standard for this purpose in the theory of randomized search heuristics is *drift theory*.

Drift theory is a general term for a collection of theorems that consider random processes and bound the expected time it takes the process to reach a certain value – the *first-hitting time*. The beauty and appeal of these theorems lie in them usually having few restrictions but yielding strong results. Intuitively speaking, in order to use a drift theorem, one only needs to estimate the expected change of a random process – the *drift* – at any given point in time. Hence, a drift theorem turns expected local changes of a process into expected first-hitting times. In other words, local information of the process is transformed into global information.

Drift theory gained traction in the theory of randomized search heuristics when it was introduced to the community by He and Yao [7, 8] via the *additive*

drift theorem. However, they were not the first to prove it. The result dates back to Hajek [6], who stated the theorem in a fashion quite different from how it is phrased nowadays. According to Lengler [13], the theorem has been proven even prior to that various times. Since then, many different versions of drift theorems have been proven, the most common ones being the *variable drift theorem* [9] and the *multiplicative drift theorem* [3]. The different names refer to how the drift is bounded other than independent of time: additive means that the drift is bounded by the same value for all states; in a multiplicative scenario, the drift is bounded by a multiple of the current state of the process; and in the setting of variable drift, the drift is bounded by any monotone function with respect to the current state of the process.

At first, the theorems were only stated over finite or discrete search spaces. However, these restrictions are seldom used in the proofs and thus not necessary, as pointed out, for example, by Lehre and Witt [12], who prove a general drift theorem without these restrictions. Nonetheless, up to date, all drift theorems require a bounded search space;¹ Semenov and Terkel [19] state a Theorem very much like an additive drift theorem for unbounded search spaces, but they require the process to have a bounded variance, as they also prove concentration for their result.

The area of randomized search heuristics is, in fact, in strong need of extended drift theorems and a careful discussion of what happens when restrictions are not met. While most search spaces are finite and, thus, the existing drift theorems sufficient, progress will be inhibited whenever search spaces are not naturally finite. Worse yet, the existing drift theorems might be applied where they are not applicable, as happened when in [4, Sect. 4] the additive drift theorem was applied on an unbounded search space.

While previously new drift theorems were proven on a need-to-have basis using whatever restrictions were present in the concrete application, we aim at providing the best possible theorem for any applications to come. For the restrictions that remain, we give examples that show that these restrictions are, in some sense, necessary. In this way, we want to further the understanding of random processes in general and not just for a concrete application; thus, this work should benefit a lot of future work in the area of randomized search heuristics.

Our most important results are the upper and lower bound of the classical additive drift theorem (Theorems 5 and 7, respectively), which we prove for unbounded² search spaces. These theorems are used as a foundation for all of our other drift theorems in other settings. Overall, our results can be summarized as follows:

¹ Lengler [13] briefly mentions infinite search spaces and also gives a proof for a restricted version of the additive drift theorem in the setting of an unbounded discrete search space.

² For the upper bound, we require the search space to be lower-bounded but not upper-bounded. We still refer to such a setting as unbounded.

For **additive drift**, we prove an upper bound for any nonnegative process (Theorem 5), and a lower bound for processes with bounded *expected* step size (Theorem 7).

For **multiplicative drift** and **variable drift**, we prove upper bounds for any nonnegative process (Corollaries 11 and 12; and Theorems 9 and 10, respectively).

The intention of this paper is to provide a fully-packed reference for very general yet easy-to-apply drift theorems. That is, we try to keep the requirements of the theorems as easy as possible but still state the theorems in the most general way, given the restrictions. Further, we discuss the ideas behind the different theorems and some of the proofs in order to provide insights into how and why drift works, we provide examples, and we discuss prior work at the beginning of each section.

We only consider bounds on the expected first-hitting time, as this is already a vast field to explore. However, we want to mention that drift theory has also brought forth other results than expected first-hitting times, namely, concentration bounds and negative drift, which are related. Both areas bound *the probability* of the first-hitting time taking certain values. Concentration bounds show how unlikely it is for a process to take much longer than the expected first-hitting time [2, 10]. On the other hand, negative drift bounds how likely it is for the process to reach the goal although the drift is going the opposite direction [10, 16, 17]. These results are also very helpful but out of the scope of this paper.

Our paper is structured as follows: in Sect. 2, we start by introducing important notation and terms, which we use throughout the entire paper. Further, we also discuss Theorem 1, which our proofs of the additive drift theorems rely on. In Sect. 3, we discuss additive drift and prove our main results. We then continue with variable drift in Sect. 4 as a generalization of additive drift. In this section, we introduce two different versions of first-hitting time that our results are based on. Last, we consider the scenario of multiplicative drift in Sect. 5.

Due to space limitations, many proofs are omitted in this paper. A version containing all of the proofs can be found on arXiv [11].

2 Preliminaries

We consider the expected *first-hitting time* T of a process $(X_t)_{t \in \mathbb{N}}$ over \mathbb{R} , which we call X_t for short. That is, we are interested in the expected time it takes the process to reach a certain value for the first time, which we will refer to as the *target*. Usually, our target is the value 0, that is, we will define the random variable $T = \inf\{t \mid X_t \leq 0\}$ (where we define that $\inf \emptyset := \infty$).

We provide bounds on $\mathbb{E}[T \mid X_0]$ with respect to the *drift* of X_t , which is defined as

$$X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t].$$

Note that $\mathbb{E}[T \mid X_0]$ as well as $\mathbb{E}[X_{t+1} \mid X_0, \dots, X_t]$ are both random variables. Because of the latter, the drift is a random variable, too. Further note that, if the drift is positive, X_t decreases its value in expectation over time when

considering positive starting values. This is why 0 will be our target most of the time.

We are only interested in the process X_t until the time point T . That is, all of our requirements only need to hold for all $t < T$ (since we also consider $t + 1$). While this phrasing is intuitive, it is formally inaccurate, as T is a random variable. We will continue to use it; however, formally, each of our inequalities in each of our requirements should be multiplied with the characteristic function of the event $\{t < T\}$. In this way, the inequalities trivially hold once $t \geq T$ and, otherwise, are the inequalities we state. This is similar to conditioning on the event $\{t < T\}$ but has the benefit of being valid even if $\Pr[t < T] = 0$ holds.

We want to mention that all of our results actually hold for a random process $(X_t)_{t \in \mathbb{N}}$ adapted to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$, where T is a stopping time defined with respect to \mathcal{F}_t .³ Since this detail is frequently ignored in drift theory, we phrase all of our results with respect to the natural filtration, making them look more familiar to usual drift results. For any time point $t \leq T$, we call X_0, \dots, X_{t-1} the *history* of the process.

Last, we state all of our results conditional on X_0 , that is, we bound $\mathbb{E}[T \mid X_0]$. However, by the law of total expectation, one can easily derive a bound for $\mathbb{E}[T] = \mathbb{E}[\mathbb{E}[T \mid X_0]]$.

2.1 Martingale Theorems

In this section, we state two theorems that we will use in order to prove our results in the next sections. Both theorems make use of *martingales*, a fundamental concept in the field of probability theory. A martingale is a random process with a drift of 0, that is, in expectation, it does not change over time. Further, a *supermartingale* has a drift of at least 0, that is, it decreases over time in expectation, and a *submartingale* has a drift of at most 0, that is, it increases over time in expectation.

The arguably most important theorem for martingales is the Optional Stopping Theorem (Theorem 1). We use a version given by Grimmett and Stirzaker [5, Chap. 12.5, Theorem 9] that can be extended to super- and submartingales.

Theorem 1 (Optional Stopping). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , and let T be a stopping time⁴ for X_t . Suppose that*

- (a) $\mathbb{E}[T] < \infty$ and that
 (b) there is some value $c \geq 0$ such that, for all $t < T$, it holds that
- $$\mathbb{E}[|X_{t+1} - X_t| \mid X_0, \dots, X_t] \leq c.$$

Then:

1. If, for all $t < T$, $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq 0$, then $\mathbb{E}[X_T] \leq \mathbb{E}[X_0]$.
2. If, for all $t < T$, $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \leq 0$, then $\mathbb{E}[X_T] \geq \mathbb{E}[X_0]$.

³ More information on filtrations can be found, for example, in *Randomized Algorithms* [15] in the section on martingales.

⁴ Intuitively, for the natural filtration, a stopping time T is a random variable over \mathbb{N} such that, for all $t \in \mathbb{N}$, the event $\{t \leq T\}$ is only dependent on X_0, \dots, X_t .

Theorem 1 allows us to bound $E[X_T]$ independently of its history, which is why our drift results are independent of the history of X_T as well.

Note that case (1) refers to supermartingales, whereas case (2) refers to submartingales. Intuitively, case (1) says that a supermartingale will have, in expectation, a lower value than it started with, which makes sense, as a supermartingale decreases over time in expectation. Case (2) is analogous for submartingales. For martingales, both cases can be combined in order to yield an equality.

Martingales are essential in the proofs of our theorems. We will frequently transform our process such that it results in a supermartingale or a submartingale in order to apply Theorem 1.

Another useful theorem for martingales is the following Azuma–Hoeffding Inequality [1]. This inequality basically is for martingales what a Chernoff bound is for binomial distributions.

Theorem 2 (Azuma-Hoeffding Inequality). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} . Suppose that*

(a) *there is some value $c > 0$ such that, for all $t \in \mathbb{N}$, it holds that $|X_t - X_{t+1}| < c$.*

If, for all $t \in \mathbb{N}$, $X_t - E[X_{t+1} | X_0, \dots, X_t] \geq 0$, then, for all $t \in \mathbb{N}$ and all $r > 0$,
 $\Pr[X_t - X_0 \geq r] \leq e^{-\frac{r^2}{2tc^2}}.$

3 Additive Drift

We speak of additive drift when the drift can be bounded by a value independent of the process itself. That is, the bound is independent of time and state.

When considering the first-hitting time T of a random process $(X_t)_{t \in \mathbb{N}}$ whose drift is *lower*-bounded by a value $\delta > 0$, then $E[T | X_0]$ is *upper*-bounded by X_0/δ . Interestingly, if the drift of X_t is *upper*-bounded by δ , $E[T | X_0]$ is *lower*-bounded by X_0/δ . Thus, if the drift of X_t is exactly δ , that is, we know how much expected progress X_t makes in each step, our expected first-hitting time is equal to X_0/δ . This result is remarkable, as it can be understood intuitively as follows: since we stop once X_t reaches 0, the distance from our start (X_0) to our goal (0) is exactly X_0 , and we make an expected progress of δ each step. Thus, in expectation, we are done after X_0/δ steps.

3.1 Upper Bounds

We give a proof for the Additive Drift Theorem, originally published (in a more restricted version) by He and Yao [7, 8]. We start by reproving the original theorem (which requires a bounded search space) but in a simpler, more elegant and educational manner. We then greatly extend this result by generalizing it to processes with a bounded step width. Finally, we lift also this restriction.

In all of these cases, we require our random process to only take nonnegative values. The intuitive reason for this is the following: when estimating an upper bound for the expected first-hitting time, we need a lower bound of the drift. This means the larger our bound of the drift, the better our bound for the first-hitting time. Since our process is nonnegative, the drift for values close to 0 provides a natural bound for the drift (which is uniform over the entire search space, since we look at *additive* drift). If our process could take values less than 0, we could artificially increase our lower bound of the drift for values that are now bounded by 0 and, thus, improve our first-hitting time. At the end of this section, we also give an example (Example 6), which shows how our most general drift theorem (Theorem 5) fails if the process can take negative values.

The proof of the following theorem transforms the process into a supermartingale and then uses Theorem 1. However, in order to apply Theorem 1, we have to make sure to fulfill its condition (a), which is the hardest part.

Theorem 3 (Upper Additive Drift, Bounded). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , and let $T = \inf\{t \mid X_t \leq 0\}$. Furthermore, suppose that,*

- (a) *for all $t \leq T$, it holds that $X_t \geq 0$, that*
- (b) *there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq \delta$, and that*
- (c) *there is some value $c \geq 0$ such that, for all $t < T$, it holds that $X_t \leq c$.*

Then $\mathbb{E}[T \mid X_0] \leq \frac{X_0}{\delta}$.

Note that condition (a) means that T can be rewritten as $\inf\{t \mid X_t = 0\}$, that is, we have to hit 0 exactly in order to stop. We show in Example 6 why this condition is crucial.

Condition (b) bounds the expected progress we make each time step. The larger δ , the lower the expected first-hitting time. However, due to condition (a), note that small values of X_t create a natural upper bound for δ , as the progress for such values can be at most $|X_t - 0| = X_t$.

Condition (c) means that we are considering random variables over the interval $[0, c]$. It is a restriction that all previous additive drift theorems have but that is actually not necessary, as we show with Theorem 5. In the following proof, we use this condition in order to show that $\mathbb{E}[T] < \infty$, which is necessary when applying Theorem 1.

Proof (Proof of Theorem 3). We want to use case (1) of the Optional Stopping Theorem in the version of Theorem 1. Thus, we define, for all $t < T$, $Y_t = X_t + \delta t$, which is a supermartingale, since

$$\begin{aligned} Y_t - \mathbb{E}[Y_{t+1} \mid Y_0, \dots, Y_t] &= X_t + \delta t - \mathbb{E}[X_{t+1} + \delta(t+1) \mid X_0, \dots, X_t] \\ &= X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] - \delta \geq 0, \end{aligned}$$

as we assume that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq \delta$ for all $t < T$. Note that we can change the condition Y_0, \dots, Y_t to X_0, \dots, X_t because the transformation from X_t to Y_t is injective.

We now show that $E[T | X_0] < \infty$ holds in order to apply Theorem 1. Let $r > 0$, and let a be any value such that $\Pr[X_0 \leq a] > 0$. We condition on the event $\{X_0 \leq a\}$, and we consider a time point $t' = (a + r)/\delta$ and want to bound the probability that $X_{t'}$ has not reached 0 yet, that is, the event $\{X_{t'} > 0\}$. We rewrite this event as $\{X_{t'} - a > -a\}$, which is equivalent to $\{Y_{t'} - a > -a + \delta t' = r\}$, by definition of Y and t' .

Note that, for all $t < T$, $|Y_t - Y_{t+1}| < c + \delta + 1$, as we assume that $X_t \leq c$. Thus, the differences of Y_t are bounded and we can apply Theorem 2 as follows, noting that $Y_0 = X_0 \leq a$, due to our condition on $\{X_0 \leq a\}$:

$$\Pr[Y_{t'} - a > r | X_0 \leq a] \leq \Pr[Y_{t'} - Y_0 \geq r | X_0 \leq a] \leq e^{-\frac{r^2}{2t'(c+\delta+1)^2}}.$$

If we choose $r \geq a$, we get $t' \leq 2r/\delta$ and, thus,

$$\Pr[Y_{t'} - Y_0 > r | X_0 \leq a] \leq e^{-\frac{r\delta}{4(c+\delta+1)^2}}.$$

This means that the probability that $X_{t'}$ has not reached 0 goes exponentially fast toward 0 as t' (and, hence, r) goes toward ∞ . Thus, the expected value of T is finite.

Now we can use case (1) of Theorem 1 in order to get $E[Y_T | X_0] \leq E[Y_0 | X_0]$. In particular, noting that $X_T = 0$ by definition,

$$\begin{aligned} X_0 &= E[X_0 | X_0] = E[Y_0 | X_0] \\ &\geq E[Y_T | X_0] = E[X_T + \delta T | X_0] = E[X_T | X_0] + \delta E[T | X_0] = \delta E[T | X_0]. \end{aligned}$$

Thus, we get the desired bound by dividing by δ . □

Note that the arguments in this proof only need the property of bounded differences in order to apply Theorem 2. Thus, we can relax the condition of a bounded state space into bounded step size, which can be seen in the following theorem.

Theorem 4 (Upper Additive Drift, Bounded Step Size). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , and let $T = \inf\{t \mid X_t \leq 0\}$. Furthermore, suppose that,*

- (a) *for all $t \leq T$, it holds that $X_t \geq 0$, that*
- (b) *there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - E[X_{t+1} | X_0, \dots, X_t] \geq \delta$, and that*
- (c) *there is some value $c \geq 0$ such that, for all $t < T$, it holds that $|X_{t+1} - X_t| \leq c$.*

Then $E[T | X_0] \leq \frac{X_0}{\delta}$.

Although the proof of Theorem 3 can be used for Theorem 4 as well, we provide a different proof strategy in the appendix, which we then generalize for our next theorem. This alternative strategy defines a process similar to X_t that behaves like X_t in the limit.

The proof of Theorem 4 makes use of Theorem 3 by artificially bounding the search space for a time that is sufficient in order to bound the expected first-hitting time. This approach can be used in order to let the restriction of the bounded step size fall entirely. Since we cannot make many assumptions about the process in this case anymore, we rely on Markov's inequality in order to show that our process will not leave, with sufficiently high probability, an interval large enough to properly bound the expected first-hitting time.

Theorem 5 (Upper Additive Drift, Unbounded). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , and let $T = \inf\{t \mid X_t \leq 0\}$. Furthermore, suppose that,*

- (a) *for all $t \leq T$, it holds that $X_t \geq 0$, and that*
- (b) *there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq \delta$.*

Then $\mathbb{E}[T \mid X_0] \leq \frac{X_0}{\delta}$.

As we already mentioned before, note that the condition of the process not being negative is important in order to get correct results. The following example highlights this fact.

Example 6. Let $n > 1$, and let $(X_t)_{t \in \mathbb{N}}$ be a random process with $X_0 = 1$ and, for all $t \in \mathbb{N}$, $X_{t+1} = X_t$ with probability $1 - 1/n$, and $X_{t+1} = -n + 1$ otherwise. Let T denote the first point in time t such that the event $X_t \leq 0$ occurs. We have, for all $t < T$, that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] = 1$ and, thus, $\mathbb{E}[T \mid X_0] \leq 1$ if we could apply any of the additive drift theorems. However, since T follows a geometric distribution with success probability $1/n$, we have $\mathbb{E}[T \mid X_0] = n$.

3.2 Lower Bound

In this section, we provide a lower bound for the expected first-hitting time under additive drift. In order to do so, we need an upper bound for the drift. Since we now lower-bound the first-hitting time, a large upper bound of the drift makes the result bad. Thus, we can allow the process to take negative values, as these could only increase the drift's upper bound. However, we need to have some restriction on the step size in order to make sure not to move away from the target. Again, we provide an example (Example 8) showing this necessity at the end of this section.

Theorem 7 (Lower Additive Drift, Expected Bounded Step Size). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , and let $T = \inf\{t \mid X_t \leq 0\}$. Furthermore, suppose that*

- (a) *there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \leq \delta$, and that*
- (b) *there is some value $c \geq 0$ such that, for all $t < T$, it holds that $\mathbb{E}[|X_{t+1} - X_t| \mid X_0, \dots, X_t] \leq c$.*

Then $E[T | X_0] \geq \frac{X_0}{\delta}$.

Proof. We make a case distinction with respect to $E[T | X_0]$ being finite. If $E[T | X_0]$ is infinite, then the theorem trivially holds. Thus, we now assume that $E[T | X_0] < \infty$.

Similar to the proof of Theorem 3, we define, for all $t < T$, $Y_t = X_t + \delta t$, which is a submartingale, since

$$\begin{aligned} Y_t - E[Y_{t+1} | Y_0, \dots, Y_t] &= X_t - \delta t - E[X_{t+1} - \delta(t+1) | X_0, \dots, X_t] \\ &= X_t - E[X_{t+1} | X_0, \dots, X_t] - \delta \leq 0, \end{aligned}$$

as we assume that $X_t - E[X_{t+1} | X_0, \dots, X_t] \leq \delta$ for all $t < T$ and because, again, the transformation of X_t to Y_t is injective.

Since we now assume that both $E[T | X_0] < \infty$ and, further, that $E[|X_{t+1} - X_t| | X_0, \dots, X_t] \leq c$ for all $t < T$, we can directly apply case (2) of Theorem 1 and get that $E[Y_T | X_0] \geq E[Y_0 | X_0]$. This yields, noting that $X_T \leq 0$,

$$\begin{aligned} X_0 &= E[X_0 | X_0] = E[Y_0 | X_0] \\ &\leq E[Y_T | X_0] = E[X_T + \delta T | X_0] = E[X_T | X_0] + \delta E[T | X_0] \leq \delta E[T | X_0]. \end{aligned}$$

Thus, we get the desired bound by dividing by δ . \square

Note that the step size has to be bounded in some way for a lower bound, as the following example shows.

Example 8. Let $\delta \in (0, 1)$, and let $(X_t)_{t \in \mathbb{N}}$ be a random process with $X_0 = 2$ and, for all $t \in \mathbb{N}$, $X_{t+1} = 0$ with probability $1/2$ and $X_{t+1} = 2X_t - 2\delta$ otherwise. Further, let T denote the first point in time t such that $X_t = 0$. Then T follows a geometric distribution with success probability $1/2$, which yields $E[T] = 2$. However, we have that $X_t - E[X_{t+1} | X_0, \dots, X_t] = \delta$. If Theorem 7 could be applied to this process (by neglecting the condition of the bounded step size), the theorem would yield that $E[T] \geq 2/\delta$, which is not true.

4 Variable Drift

In contrast to additive drift, *variable* drift means that the drift can depend on the current state of the process (while still being bounded independently of the time). Interestingly, these more flexible drift theorems can be derived by using additive drift. Intuitively, the reasoning behind this approach is to scale the search space such that the information relevant to the process's history cancels out.

It is important to note that variable drift theorems are commonly phrased such that the first-hitting time T denotes the first point in time such that the random process drops *below* a certain value (our target) – it is not enough to hit that value. However, this restriction is not always necessary. Thus, we also

consider the setting from Sect. 3, where T denotes the first point in time such that we *hit* our target. In this section, our target is no longer 0 but a value x_{\min} .

In all of our theorems in this section, we make use of a set D . This set contains (at least) all possible values that our process can take while not having reached the target yet. It is a formal necessity in order to calculate the bound of the first-hitting time (via an integral). However, when applying the theorem, it is usually sufficient to choose $D = \mathbb{R}$ or $D = \mathbb{R}_{\geq 0}$.

The first variable drift theorem was proven by Johannsen [9] and, independently in a different version, by Mitavskiy et al. [14]. It was later refined by Rowe and Sudholt [18]. In all of these versions, bounded search spaces were used. Due to Theorem 5, we can drop this restriction.

Going Below the Target. The following version of the theorem assumes that the process has to drop below the target, denoted by x_{\min} . We provide the other version afterward.

Theorem 9 (Upper Variable Drift, Unbounded, Below Target). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , $x_{\min} > 0$, and let $T = \inf\{t \mid X_t < x_{\min}\}$. Additionally, let D denote the smallest real interval that contains at least all values $x \geq x_{\min}$ that, for all $t \leq T$, any X_t can take. Furthermore, suppose that*

- (a) $X_0 \geq x_{\min}$ and, for all $t \leq T$, it holds that $X_t \geq 0$ and that
- (b) there is a monotonically increasing function $h: D \rightarrow \mathbb{R}^+$ such that, for all $t < T$, we have $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq h(X_t)$.

Then

$$\mathbb{E}[T \mid X_0] \leq \frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{X_0} \frac{1}{h(z)} dz.$$

Hitting the Target. As mentioned before, it is not always necessary to drop below the target. For the additive drift, for example, we are interested in the first time reaching the target. Interestingly, the proof for the following theorem is straightforward, as it is almost the same as the proof of Theorem 9. Intuitively, the waiting time for getting below the target, once it is reached, is eliminated from the expected first-hitting time. However, it is important to note that it is now not allowed to get below the target.

Theorem 10 (Upper Variable Drift, Unbounded, Hitting Target). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , $x_{\min} \geq 0$, and let $T = \inf\{t \mid X_t \leq x_{\min}\}$. Additionally, let D denote the smallest real interval that contains at least all values $x \geq x_{\min}$ that, for all $t \leq T$, any X_t can take. Furthermore, suppose that,*

- (a) for all $t \leq T$, it holds that $X_t \geq x_{\min}$ and that
- (b) there is a monotonically increasing function $h: D \rightarrow \mathbb{R}^+$ such that, for all $t < T$, we have $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq h(X_t)$.

Then

$$\mathbb{E}[T \mid X_0] \leq \int_{x_{\min}}^{X_0} \frac{1}{h(z)} dz.$$

5 Multiplicative Drift

A special case of variable drift is *multiplicative* drift, where the drift can be bounded by a multiple of the most recent value in the history of the process. As before, we provide upper bounds in the two versions of either dropping below the target or hitting it. In this setting, it can be intuitively argued why the version of dropping below the target is useful: consider a sequence of nonnegative numbers that halves its current value each time step. This process will never reach 0 within finite time. However, it drops below any value greater than 0.

Both upper bounds we state are simple applications of the corresponding variable drift theorems from Sect. 4.

Going Below the Target. Corollary 11 has first been stated by Doerr et al. [3] using finite state spaces. However, a closer look at the proof shows that this restriction is not necessary.

Corollary 11 (Upper Multiplicative Drift, Unbounded, Below Target). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , $x_{\min} > 0$, and let $T = \inf\{t \mid X_t < x_{\min}\}$. Furthermore, suppose that*

- (a) $X_0 \geq x_{\min}$ and, for all $t \leq T$, it holds that $X_t \geq 0$, and that
- (b) there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq \delta X_t$.

Then

$$\mathbb{E}[T \mid X_0] \leq \frac{1 + \ln\left(\frac{X_0}{x_{\min}}\right)}{\delta}.$$

Hitting the Target. By applying Theorem 10 instead of Theorem 9, we get the following theorem. As in the case of Theorem 10, the process now has to be lower-bounded by x_{\min} .

Corollary 12 (Upper Multiplicative Drift, Unbounded, Hitting Target). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , $x_{\min} > 0$, and let $T = \inf\{t \mid X_t \leq x_{\min}\}$. Furthermore, suppose that,*

- (a) for all $t \leq T$, it holds that $X_t \geq x_{\min}$, and that
- (b) there is some value $\delta > 0$ such that, for all $t < T$, it holds that $X_t - \mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \geq \delta X_t$.

Then

$$\mathbb{E}[T \mid X_0] \leq \frac{\ln\left(\frac{X_0}{x_{\min}}\right)}{\delta}.$$

Again, we provide an example that shows that the bounds above are as tight as possible, up to constant factors, for the range of processes we consider. The example describes a process that decreases deterministically, that is, it has a variance of 0.

Example 13. Let $\delta \in (0, 1)$ be a value bounded away from 1. Consider the process $(X_t)_{t \in \mathbb{N}}$, with $X_0 > 1$, that decreases each step deterministically such that $X_{t+1} = (1 - \delta)X_t$ holds. Let T denote the first point in time such that the process drops below 1. Thus, we get $T = \Theta(-\log_{(1-\delta)} X_0) = \Theta(-\ln(X_0)/\ln(1 - \delta)) = \Theta(\ln(X_0)/\delta)$, where the last equation makes use of the Taylor expansion of $\ln(1 - \delta) = \Theta(-\delta)$, as $1 - \delta$ does not converge to 0, by assumption.

References

1. Azuma, K.: Weighted sums of certain dependent random variables. *Tohoku Math. J.* **19**(3), 357–367 (1967)
2. Doerr, B., Goldberg, L.A.: Adaptive drift analysis. *Algorithmica* **65**(1), 224–250 (2013)
3. Doerr, B., Johannsen, D., Winzen, C.: Multiplicative drift analysis. *Algorithmica* **64**(4), 673–697 (2012)
4. Doerr, B., Kötzing, T., Lagodzinski, J.A.G., Lengler, J.: Bounding bloat in genetic programming. In: *Proceedings of the GECCO 2017*, pp. 921–928 (2017)
5. Grimmett, G.R., Stirzaker, D.R.: *Probability and Random Processes*. Oxford University Press, Oxford (2001)
6. Hajek, B.: Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Probab.* **14**(3), 502–525 (1982)
7. He, J., Yao, X.: Drift analysis and average time complexity of evolutionary algorithms. *Artif. Intell.* **127**(1), 57–85 (2001)
8. He, J., Yao, X.: A study of drift analysis for estimating computation time of evolutionary algorithms. *Nat. Comput.* **3**(1), 21–35 (2004)
9. Johannsen, D.: *Random combinatorial structures and randomized search heuristics*. Ph.D. thesis, Universität des Saarlandes (2010). http://scidok.sulb.uni-saarland.de/volltexte/2011/3529/pdf/Dissertation.3166_Joha_Dani_2010.pdf
10. Kötzing, T.: Concentration of first hitting times under additive drift. *Algorithmica* **75**(3), 490–506 (2016)
11. Kötzing, T., Krejca, M.S.: First-hitting times under additive drift. *CoRR abs/1805.09415* (2018). <https://arxiv.org/abs/1805.09415>
12. Lehre, P.K., Witt, C.: Concentrated hitting times of randomized search heuristics with variable drift. In: Ahn, H.-K., Shin, C.-S. (eds.) *ISAAC 2014*. LNCS, vol. 8889, pp. 686–697. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13075-0_54
13. Lengler, J.: Drift analysis. *CoRR abs/1712.00964* (2017). <http://arxiv.org/abs/1712.00964>
14. Mitavskiy, B., Rowe, J.E., Cannings, C.: Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *Int. J. Intell. Comput. Cybern.* **2**(2), 243–284 (2009)
15. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press, Cambridge (1995)
16. Oliveto, P.S., Witt, C.: Simplified drift analysis for proving lower bounds in evolutionary computation. *Algorithmica* **59**(3), 369–386 (2011)
17. Oliveto, P.S., Witt, C.: Erratum: simplified drift analysis for proving lower bounds in evolutionary computation. *CoRR abs/1211.7184* (2012). <http://arxiv.org/abs/1211.7184>

18. Rowe, J.E., Sudholt, D.: The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theor. Comput. Sci.* **545**, 20–38 (2014)
19. Semenov, M.A., Terkel, D.A.: Analysis of convergence of an evolutionary algorithm with self-adaptation using a Stochastic Lyapunov function. *Evol. Comput.* **11**(4), 363–379 (2003)