Indefinitely Oscillating Martingales

Jan Leike Marcus Hutter

August 14, 2014

Abstract

We construct a class of nonnegative martingale processes that oscillate indefinitely with high probability. For these processes, we state a uniform rate of the number of oscillations and show that this rate is asymptotically close to the theoretical upper bound. These bounds on probability and expectation of the number of upcrossings are compared to classical bounds from the martingale literature. We discuss two applications. First, our results imply that the limit of the minimum description length operator may not exist. Second, we give bounds on how often one can change one's belief in a given hypothesis when observing a stream of data.¹

Keywords. martingales, infinite oscillations, bounds, convergence rates, minimum description length, mind changes.

1 Introduction

Martingale processes model fair gambles where knowledge of the past or choice of betting strategy have no impact on future winnings. But their application is not restricted to gambles and stock markets. Here we exploit the connection between nonnegative martingales and probabilistic data streams, i.e., probability measures on infinite strings. For two probability measures P and Q on infinite strings, the quotient Q/P is a nonnegative P-martingale. Conversely, every nonnegative P-martingale is a multiple of Q/P P-almost everywhere for some probability measure Q.

One of the famous results of martingale theory is Doob's Upcrossing Inequality [Doo53]. The inequality states that in expectation, every nonnegative martingale has only finitely many oscillations (called *upcrossings* in the martingale literature). Moreover, the bound on the expected number of oscillations is inversely proportional to their magnitude. Closely related is Dubins' Inequality [Dub62] which asserts that the probability of having many oscillations decreases exponentially with their number. These bounds are given with respect to oscillations of fixed magnitude.

In Section 4 we construct a class of nonnegative martingale processes that have infinitely many oscillations of (by Doob necessarily) decreasing magnitude. These martingales satisfy uniform lower bounds on the probability and the expectation of the number of upcrossings. We prove corresponding upper bounds in Section 5 showing that these lower bounds are asymptotically tight. Moreover, the construction of the martingales is agnostic regarding the underlying

¹ This is the extended technical report. The conference version can be found at [LH14].

probability measure, assuming only mild restrictions on it. We compare these results to the statements of Dubins' Inequality and Doob's Upcrossing Inequality and demonstrate that our process makes those inequalities asymptotically tight. If we drop the uniformity requirement, asymptotics arbitrarily close to Doob and Dubins' bounds are achievable. We discuss two direct applications of these bounds.

The Minimum Description Length (MDL) principle [Ris78] and the closely related Minimal Message Length (MML) principle [WB68] recommend to select among a class of models the one that has the shortest code length for the data plus code length for the model. There are many variations, so the following statements are generic: for a variety of problem classes MDL's predictions have been shown to converge asymptotically (predictive convergence). For continuous independently identically distributed data the MDL estimator usually converges to the true distribution [Grü07, Wal05] (inductive consistency). For arbitrary (non-i.i.d.) countable classes, the MDL estimator's predictions converge to those of the true distribution for single-step predictions [PH05] and ∞ -step predictions [Hut09]. Inductive consistency implies predictive convergence, but not the other way around. In Section 6 we show that indeed, the MDL estimator for countable classes is *inductively inconsistent*. This can be a major obstacle for using MDL for prediction, since the model used for prediction has to be changed over and over again, incurring the corresponding computational cost.

Another application of martingales is in the theory of mind changes [LS05]. How likely is it that your belief in some hypothesis changes by at least $\alpha > 0$ several times while observing some evidence? Davis recently showed [Dav13] using elementary mathematics that this probability decreases exponentially. In Section 7 we rephrase this problem in our setting: the stochastic process

$P(\text{hypothesis} \mid \text{evidence up to time } t)$

is a martingale bounded between 0 and 1. The upper bound on the probability of many changes can thus be derived from Dubins' Inequality. This yields a simpler alternative proof for Davis' result. However, because we consider nonnegative but unbounded martingales, we get a weaker bound than Davis.

2 Strings, Measures, and Martingales

We presuppose basic measure and probability theory [Dur10, Chp.1]. Let Σ be a finite set, called *alphabet*. We assume Σ contains at least two distinct elements. For every $u \in \Sigma^*$, the *cylinder set*

$$\Gamma_u := \{ uv \mid v \in \Sigma^\omega \}$$

is the set of all infinite strings of which u is a prefix. Furthermore, fix the $\sigma\text{-algebras}$

$$\mathcal{F}_t := \sigma\left(\{\Gamma_u \mid u \in \Sigma^t\}\right)$$
 and $\mathcal{F}_\omega := \sigma\left(\bigcup_{t=1}^\infty \mathcal{F}_t\right).$

 $(\mathcal{F}_t)_{t\in\mathbb{N}}$ is a *filtration*: since $\Gamma_u = \bigcup_{a\in\Sigma} \Gamma_{ua}$, it follows that $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ for every $t \in \mathbb{N}$, and all $\mathcal{F}_t \subseteq \mathcal{F}_{\omega}$ by the definition of \mathcal{F}_{ω} . An *event* is a measurable set $E \in \mathcal{F}_{\omega}$. The event $E^c := \Sigma^{\omega} \setminus E$ denotes the complement of E. See also the list of notation in Appendix A.1.

Definition 1 (Stochastic Process). $(X_t)_{t \in \mathbb{N}}$ is called (\mathbb{R} -valued) stochastic process iff each X_t is an \mathbb{R} -valued random variable.

Definition 2 (Martingale). Let P be a probability measure over $(\Sigma^{\omega}, \mathcal{F}_{\omega})$. An \mathbb{R} -valued stochastic process $(X_t)_{t \in \mathbb{N}}$ is called a *P*-supermartingale (*P*-submartingale) iff

- (a) each X_t is \mathcal{F}_t -measurable, and
- (b) $\mathbb{E}[X_t \mid \mathcal{F}_s] \leq X_s \ (\mathbb{E}[X_t \mid \mathcal{F}_s] \geq X_s)$ almost surely for all $s, t \in \mathbb{N}$ with s < t.

A process that is both P-supermartingale and P-submartingale is called P-martingale.

We call a supermartingale (submartingale) process $(X_t)_{t\in\mathbb{N}}$ nonnegative iff $X_t \geq 0$ for all $t \in \mathbb{N}$.

A stopping time is an $(\mathbb{N} \cup \{\omega\})$ -valued random variable T such that $\{v \in \Sigma^{\omega} \mid T(v) = t\} \in \mathcal{F}_t$ for all $t \in \mathbb{N}$. Given a supermartingale $(X_t)_{t \in \mathbb{N}}$, the stopped process $(X_{\min\{t,T\}})_{t \in \mathbb{N}}$ is a supermartingale [Dur10, Thm. 5.2.6]. If $(X_t)_{t \in \mathbb{N}}$ is bounded, the limit of the stopped process, X_T , exists almost surely even if $T = \omega$ (Martingale Convergence Theorem [Dur10, Thm. 5.2.8]). We use the following variant on Doob's Optional Stopping Theorem for supermartingales.

Theorem 3 (Optional Stopping Theorem [Dur10, Thm. 5.7.6]). Let $(X_t)_{t\in\mathbb{N}}$ be a nonnegative supermartingale and let T be a stopping time. The random variable X_T is almost surely well defined and $\mathbb{E}[X_T] \leq \mathbb{E}[X_0]$.

For two probability measures P and Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$, the measure Q is called absolutely continuous with respect to P on cylinder sets iff $Q(\Gamma_u) = 0$ for all $u \in \Sigma^*$ with $P(\Gamma_u) = 0$. We exploit the following two theorems that state the connection between probability measures on infinite strings and martingales. For two probability measures P and Q the quotient Q/P is a nonnegative Pmartingale if Q is absolutely continuous with respect to P on cylinder sets. Conversely, for every nonnegative P-martingale there is a probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that the martingale is P-almost surely a multiple of Q/Pand Q is absolutely continuous with respect to P on cylinder sets.

Theorem 4 (Measures \rightarrow Martingales [Doo53, II§7 Ex. 3]). Let Q and P be two probability measures on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that Q is absolutely continuous with respect to P on cylinder sets. Then the stochastic process $(X_t)_{t \in \mathbb{N}}$,

$$X_t(v) := \frac{Q(\Gamma_{v_{1:t}})}{P(\Gamma_{v_{1:t}})}$$

is a nonnegative P-martingale with $\mathbb{E}[X_t] = 1$.

Theorem 5 (Martingales \rightarrow Measures). Let P be a probability measure on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ and let $(X_t)_{t \in \mathbb{N}}$ be a nonnegative P-martingale with $\mathbb{E}[X_t] = 1$. There is a probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ that is absolutely continuous with respect to P on cylinder sets and for all $v \in \Sigma^{\omega}$ and all $t \in \mathbb{N}$ with $P(\Gamma_{v_{1,t}}) > 0$,

$$X_t(v) = \frac{Q(\Gamma_{v_{1:t}})}{P(\Gamma_{v_{1:t}})}.$$

For completeness, we provide proofs for Theorem 4 and Theorem 5 in Appendix A.2.

Remark 17 (Absolute continuity and absolute continuity on cylinder sets). A measure Q is called *absolutely continuous with respect to* P iff Q(A) = 0 implies P(A) = 0 for all measurable sets $A \in \mathcal{F}_{\omega}$. Absolute continuity trivially implies absolute continuity on cylinder sets. However, the converse is not true: absolute continuity on cylinder sets is a strictly weaker condition than absolute continuity.

Let P be a Bernoulli(2/3) and Q be a Bernoulli(1/3) process. Formally, we fix $\Sigma = \{0, 1\}$ and define for all $u \in \Sigma^*$,

$$P(\Gamma_u) := \left(\frac{2}{3}\right)^{\operatorname{ones}(u)} \left(\frac{1}{3}\right)^{\operatorname{zeros}(u)},$$
$$Q(\Gamma_u) := \left(\frac{1}{3}\right)^{\operatorname{ones}(u)} \left(\frac{2}{3}\right)^{\operatorname{zeros}(u)},$$

where ones(u) denotes the number of ones in u and zeros(u) denotes the number of zeros in u. Both measures P and Q are nonzero on all cylinder sets: $Q(\Gamma_u) \ge 3^{-|u|} > 0$ and $P(\Gamma_u) \ge 3^{-|u|} > 0$ for every $u \in \Sigma^*$. Therefore Q is absolutely continuous with respect to P on cylinder sets. However, Q is not absolutely continuous with respect to P: define

$$A := \left\{ v \in \Sigma^{\omega} \mid \limsup_{t \to \infty} \frac{1}{t} \operatorname{ones}(v_{1:t}) \leq \frac{1}{2} \right\}.$$

The set A is \mathcal{F}_{ω} -measurable since $A = \bigcap_{n=1}^{\infty} \bigcup_{u \in U_n} \Gamma_u$ with $U_n := \{u \in \Sigma^* \mid |u| \ge n \text{ and ones}(u) \le |u|/2\}$, the set of all finite strings of length at least n that have at least as many zeros as ones. We have that P(A) = 0 and Q(A) = 1, hence Q is not absolutely continuous with respect to P.

While Theorem 4 trivially also holds if Q is absolutely continuous with respect to P, Theorem 5 does not imply that Q is absolutely continuous with respect to P. Consider the process $X_0(v) := 1$,

$$X_{t+1}(v) := \begin{cases} 2X_t, & \text{if } v_{t+1} = 0 \text{ and} \\ \frac{1}{2}X_t, & \text{if } v_{t+1} = 1. \end{cases}$$

The process $(X_t)_{t\in\mathbb{N}}$ is a nonnegative *P*-martingale since every X_t is \mathcal{F}_t -measurable and for $u = v_{1:t}$ we have

$$\mathbb{E}[X_{t+1} \mid \mathcal{F}_t](v) = P(\Gamma_{u0} \mid \Gamma_u) 2X_t(v) + P(\Gamma_{u1} \mid \Gamma_u) \frac{1}{2} X_t(v) = \frac{1}{3} 2X_t(v) + \frac{2}{3} \cdot \frac{1}{2} X_t(v) = X_t(v).$$

Moreover,

$$Q(\Gamma_u) = \left(\frac{1}{3}\right)^{\operatorname{ones}(u)} \left(\frac{2}{3}\right)^{\operatorname{zeros}(u)}$$

= $\left(\frac{2}{3}\right)^{\operatorname{ones}(u)} \left(\frac{1}{3}\right)^{\operatorname{zeros}(u)} 2^{-\operatorname{ones}(u)} 2^{\operatorname{zeros}(u)} = P(\Gamma_u) X_t(v).$

Hence $X_t(v) = Q(\Gamma_{v_{1:t}})/P(\Gamma_{v_{1:t}})$ *P*-almost surely. The measure *Q* is uniquely defined by its values on the cylinder sets, and as shown above, *Q* is not absolutely continuous with respect to *P*.

3 Martingale Upcrossings

Fix $c \in \mathbb{R}$ and $\varepsilon > 0$, and let $(X_t)_{t \in \mathbb{N}}$ be a martingale over the probability space $(\Sigma^{\omega}, \mathcal{F}_{\omega}, P)$. Let $t_1 < t_2$. We say the process $(X_t)_{t \in \mathbb{N}}$ does an ε -upcrossing between t_1 and t_2 iff $X_{t_1} \leq c - \varepsilon$ and $X_{t_2} \geq c + \varepsilon$. Similarly, we say $(X_t)_{t \in \mathbb{N}}$ does an ε -downcrossing between t_1 and t_2 iff $X_{t_1} \geq c + \varepsilon$ and $X_{t_2} \leq c - \varepsilon$. Except for the first upcrossing, consecutive upcrossings always involve intermediate downcrossings. Formally, we define the stopping times

$$T_0(v) := 0,$$

$$T_{2k+1}(v) := \inf\{t > T_{2k}(v) \mid X_t(v) \le c - \varepsilon\}, \text{ and }$$

$$T_{2k+2}(v) := \inf\{t > T_{2k+1}(v) \mid X_t(v) \ge c + \varepsilon\}.$$

The $T_{2k}(v)$ denote the indexes of upcrossings. We count the number of upcrossings with the random variable $U_t^X(c-\varepsilon, c+\varepsilon)$, where

$$U_t^X(c-\varepsilon, c+\varepsilon)(v) := \sup\{k \ge 0 \mid T_{2k}(v) \le t\}$$

and $U^X(c - \varepsilon, c + \varepsilon) := \sup_{t \in \mathbb{N}} U_t^X(c - \varepsilon, c + \varepsilon)$ denotes the total number of upcrossings. We omit the superscript X if the martingale $(X_t)_{t \in \mathbb{N}}$ is clear from context.

The following notation is used in the proofs. Given a monotone decreasing function $f : \mathbb{N} \to [0, 1)$ and $m, k \in \mathbb{N}$, we define the event $E_{m,k}^{X,f}$ that there are at least k-many f(m)-upcrossings:

$$E_{m,k}^{X,f} := \left\{ v \in \Sigma^{\omega} \mid U^X(1 - f(m), 1 + f(m))(v) \ge k \right\}.$$

For all $m, k \in \mathbb{N}$ we have $E_{m,k}^{X,f} \supseteq E_{m,k+1}^{X,f}$ and $E_{m,k}^{X,f} \subseteq E_{m+1,k}^{X,f}$. Again, we omit X and f in the superscript if they are clear from context.

4 Indefinitely Oscillating Martingales

In this section we construct a class of martingales that has a high probability of doing an infinite number of upcrossings. The magnitude of the upcrossings decreases at a rate of a given summable function f (a function f is called summable iff it has finite L_1 -norm, i.e., $\sum_{i=1}^{\infty} f(i) < \infty$), and the value of the martingale X_t oscillates back and forth between $1 - f(M_t)$ and $1 + f(M_t)$, where M_t denotes the number of upcrossings so far. The process has a monotone decreasing chance of escaping the oscillation. We need the following condition on the probability measure P.

Definition 18 (Perpetual Entropy). A probability measure P has *perpetual* entropy iff there is an $\varepsilon > 0$ such that for every $u \in \Sigma^*$ and $v \in \Sigma^{\omega}$ with $P(\Gamma_u) > 0$ there is an $a \in \Sigma$ and a $t \in \mathbb{N}$ with $1 - \varepsilon > P(\Gamma_{uv_{1:t}a} | \Gamma_{uv_{1:t}}) > \varepsilon$.

This condition states that after seeing some string $u \in \Sigma^*$, there is always some future time point where there are two symbols that both have conditional probability greater than ε . In other words, observing data distributed according to P, we almost surely never run out of symbols with significant entropy. This is stronger than demanding that the observed string is nonconstant with high probability, because we get a single lower bound ε for all observed strings u. **Theorem 6** (An indefinitely oscillating martingale). Let $0 < \delta < 1/2$ and let $f : \mathbb{N} \to [0, 1)$ be any monotone decreasing function such that $\sum_{i=1}^{\infty} f(i) \leq \delta/2$. For every probability measure P with perpetual entropy there is a nonnegative martingale $(X_t)_{t\in\mathbb{N}}$ with $\mathbb{E}[X_t] = 1$ and

$$P[\forall m. U(1 - f(m), 1 + f(m)) \ge m] \ge 1 - \delta.$$

Proof. By grouping symbols from Σ into two groups, we can without loss of generality assume that $\Sigma = \{0, 1\}$. Since $P(\Gamma_{u0} | \Gamma_u) + P(\Gamma_{u1} | \Gamma_u) = 1$, we can define a function $a : \Sigma^* \to \Sigma$ that assigns to every string $u \in \Sigma^*$ a symbol $a_u := a(u)$ such that $p_u := P(\Gamma_{ua_u} | \Gamma_u) \leq \frac{1}{2}$. In Claim 7 we show that without loss of generality, we can group such that $p_u > \varepsilon$ infinitely often for some $\varepsilon > 0$.

In the following we define the stochastic process $(X_t)_{t \in \mathbb{N}}$. This process depends on the random variables M_t and γ_t , which are defined below. Let $v \in \Sigma^{\omega}$ and $t \in \mathbb{N}$ be given and define $u := v_{1:t}$. For t = 0, we set $X_0(v) := 1$, and if $p_u = 0$, we set $X_{t+1} = X_t$. Otherwise we distinguish the following three cases.

(i) For $X_t(v) \ge 1$:

$$X_{t+1}(v) := \begin{cases} 1 - f(M_t(v)) & \text{if } v_{t+1} \neq a_u, \\ X_t(v) + \frac{1 - p_u}{p_u} (X_t(v) - (1 - f(M_t(v)))) & \text{if } v_{t+1} = a_u. \end{cases}$$

(ii) For $1 > X_t(v) \ge \gamma_t(v)$:

$$X_{t+1}(v) := \begin{cases} X_t(v) - \gamma_t(v) & \text{if } v_{t+1} \neq a_u, \\ 1 + f(M_t(v)) & \text{if } v_{t+1} = a_u. \end{cases}$$

(iii) For
$$X_t(v) < \gamma_t(v)$$
 and $X_t(v) < 1$:
let $d_t(v) := \min\{\frac{p_u}{1-p_u}X_t(v), \frac{1-p_u}{p_u}\gamma_t(v) - 2f(M_t(v))\}$;

$$X_{t+1}(v) := \begin{cases} X_t(v) + d_t(v) & \text{if } v_{t+1} \neq a_u, \\ X_t(v) - \frac{1-p_u}{p_u} d_t(v) & \text{if } v_{t+1} = a_u. \end{cases}$$

The random variables M_t and γ_t are defined as

$$\begin{split} \gamma_t(v) &:= \frac{p_u}{1 - p_u} \Big(1 + f(M_t(v)) - X_t(v) \Big) \\ M_t(v) &:= 1 + \operatorname*{arg\,max}_{m \in \mathbb{N}} \left\{ \forall k \le m. \ U_t^X (1 - f(k), 1 + f(k)) \ge k \right\}, \end{split}$$

i.e., M_t is 1 plus the number of upcrossings completed up to time t.

We give an intuition for the behavior of the process $(X_t)_{t\in\mathbb{N}}$. For all m, the following repeats. First X_t increases while reading a_u 's until it reads one symbol that is not a_u and then jumps down to 1 - f(m). Subsequently, X_t decreases while not reading a_u 's until it falls below γ_t or reads an a_u and then jumps up to 1 + f(m). If it falls below 1 and γ_t , then at every step, it can either jump up to 1 - f(m) or jump down to 0, whichever one is closest (the distance to the closest of the two is given by d_t). See Figure 1 for a visualization.



Figure 1: An example evaluation of the martingale defined in the proof of Theorem 6.

For notational convenience, in the following we omit writing the argument v to the random variables X_t , γ_t , M_t , and d_t .

Claim 1: $(X_t)_{t\in\mathbb{N}}$ is a martingale. Each X_{t+1} is \mathcal{F}_{t+1} -measurable, since it uses only the first t + 1 symbols of v. Writing out cases (i), (ii), and (iii), we get

$$\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] \stackrel{(i)}{=} (1 - f(M_t))(1 - p_u) + \left(X_t + \frac{1 - p_u}{p_u}(X_t - (1 - f(M_t)))\right)p_u = X_t,$$

$$\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] \stackrel{(ii)}{=} \left(X_t - \frac{p_u}{1 - p_u}((1 + f(M_t)) - X_t)\right)(1 - p_u) + (1 + f(M_t))p_u = X_t,$$

$$\mathbb{E}[X_{t+1} \mid \mathcal{F}_t] \stackrel{(iii)}{=} (X_t + d_t)(1 - p_u) + (X_t - \frac{1 - p_u}{p_u}d_t)p_u = X_t.$$

Claim 2: If $X_t \ge 1 - f(M_t)$ then $X_t > \gamma_t$. In this case

$$\gamma_t = \frac{p_u}{1 - p_u} (1 + f(M_t) - X_t) \le 2 \frac{p_u}{1 - p_u} f(M_t),$$

and thus with $p_u \leq \frac{1}{2}$ and $f(M_t) \leq \sum_{k=1}^{\infty} f(k) \leq \frac{\delta}{2} < \frac{1}{4} < \frac{1}{3}$,

$$X_t - \gamma_t \ge 1 - f(M_t) - 2\frac{p_u}{1 - p_u}f(M_t) = 1 - \frac{1 + p_u}{1 - p_u}f(M_t) \ge 1 - 3f(M_t) > 0.$$

Claim 3: If $p_u > 0$, $X_t < \gamma_t$, and $X_t < 1$ then $d_t \ge 0$. We have $\frac{p_u}{1-p_u}X_t \ge 0$ since $p_u > 0$ and $X_t \ge 0$. Moreover, $\frac{1-p_u}{p_u}\gamma_t - 2f(M_t) = 1 - f(M_t) - X_t > 0$ by the contrapositive of Claim 2.

Claim 4: The following holds for cases (i), (ii), and (iii).

- (a) In case (i): $X_{t+1} \ge X_t$ or $X_{t+1} = 1 f(M_t)$.
- (b) In case (ii): $X_{t+1} \leq X_t$ or $X_{t+1} = 1 + f(M_t)$.
- (c) In case (iii): $X_t < 1 f(M_t)$ and $X_{t+1} \le 1 f(M_t)$.

If $p_u = 0$ then $X_{t+1} = X_t$, so (a) and (b) hold trivially. Otherwise, for (a) we have $\frac{1-p_u}{p_u} > 0$ and $X_t \ge 1 - f(M_t)$. For (b) we have $\gamma_t > 0$ since $X_t < 1 + f(M_t)$. For (c), $X_t < 1 - f(M_t)$ follows from the contrapositive of Claim 2. If $p_u > 0$ then by Claim 3 we have $d_t \ge 0$ and hence $X_{t+1} \le X_t + d_t \le X_t + (1 + f(M_t) - X_t) - 2f(M_t) = 1 - f(M_t)$. Claim 5: $X_t \ge 0$ and $\mathbb{E}[X_t] = 1$. The latter follows from

arm 5:
$$X_t \ge 0$$
 and $\mathbb{E}[X_t] = 1$. The latter follows from

$$\mathbb{E}[X_t] = \mathbb{E}[\mathbb{E}[X_t \mid \mathcal{F}_{t-1}]] = \mathbb{E}[X_{t-1}] = \dots = \mathbb{E}[X_0] = 1.$$

Regarding the former, we use $0 \le f(M_t) < 1$ to conclude

- (i≠) $1 f(M_t) \ge 0$, (i=) $\frac{1-p_u}{p_u}(X_t - (1 - f(M_t))) \ge 0$ for $X_t \ge 1$,
- (ii \neq) $X_t \gamma_t \ge 0$ for $X_t \ge \gamma_t$,
- (ii=) $1 + f(M_t) \ge 0$,
- (iii \neq) $X_t + d_t \ge 0$ since $d_t \ge 0$ by Claim 3, and
- (iii=) $X_t \frac{1-p_u}{p_u} d_t \ge 0$ since $d_t \le \frac{p_u}{1-p_u} X_t$.

Claim 6: $X_t \leq 1 - f(M_t)$ or $X_t \geq 1 + f(M_t)$ for all $t \geq T_1$. We use induction on t: the induction start holds with $X_{T_1} \leq 1 - f(M_t)$ and the induction step follows from Claim 4.

Claim 7: $P(\{v \in \Sigma^{\omega} \mid p_{v_{1:t}} > \varepsilon \text{ for infinitely many } t\}) = 1 \text{ for some } \varepsilon > 0.$ By assumption P has perpetual entropy; let ε be as in Definition 18.

$$A := \{ v \in \Sigma^{\omega} \mid P(\Gamma_{v_{1:t}}) > 0 \text{ for all } t \}$$

Its complement $A^c = \bigcup_{u \in \Sigma^*: P(\Gamma_u)=0} \Gamma_u$ is the countable union of null sets and therefore P(A) = 1. Let $v \in A$ be some outcome, let $t \in \mathbb{N}$ be the current time step, and define $u := v_{1:t}$. Because P has perpetual entropy and $P(\Gamma_u) > 0$ since $v \in A$, there exists $u' \in \Sigma^*$, $a \in \Sigma$, and $v' \in \Sigma^\omega$ such that v = uu'av' and $1 - \varepsilon > P(\Gamma_{uu'a} \mid \Gamma_{uu'}) > \varepsilon$. If $P(\Gamma_{uu'a} \mid \Gamma_{uu'}) \leq 1/2$ we can select $a_{uu'} := a$; if $P(\Gamma_{uu'a} \mid \Gamma_{uu'}) > 1/2$ then, with abuse of notation, for the symbol group $b := \Sigma \setminus \{a\}$ we have $\varepsilon < P(\Gamma_{uu'b} \mid \Gamma_{uu'}) \leq 1/2$ and hence we can select $a_{uu'} := b$. In either case $p_{uu'} > \varepsilon$ for a suitable grouping of symbols.

Claim 8: $(X_t)_{t\in\mathbb{N}}$ converges almost surely to a random variable $X_{\omega} \in \{0, 1\}$. According to the Martingale Convergence Theorem [Dur10, Thm. 5.2.8], the process $(X_t)_{t\in\mathbb{N}}$ converges almost surely to a random variable X_{ω} . Assume that X_{ω} attains some value x_{ω} other than 0 and 1. Pick an $\varepsilon' > 0$ such that $|x_{\omega}| > 2\varepsilon'$ and $|1 - x_{\omega}| > 2\varepsilon'$. Since $X_t \to x_{\omega}$ we have $|x_{\omega} - X_t| < \varepsilon'$ for all but finitely many t, and hence there is a $t_0 \in \mathbb{N}$ such that $|X_t| > \varepsilon'$ and $|1 - X_t| > \varepsilon'$ for all $t \ge t_0$. Recall that $\varepsilon > 0$ is fixed and depends only on P. Below we show for cases (i), (ii), and (iii) that $|X_{t+1} - X_t| > \min\{\varepsilon \cdot \varepsilon', \varepsilon', \frac{1}{8}\}$ if $p_u > \varepsilon$. By Claim 7 we almost surely have infinitely many $t \ge t_0$ with $p_u > \varepsilon$, which is a contradiction to the fact that $(X_t)_{t\in\mathbb{N}}$ converges almost surely.

- (i) Assume $X_t \ge 1$, then $X_t > 1 + \varepsilon'$ by assumption. Either $X_{t+1} = 1 f(M_t) \le 1 < X_t \varepsilon'$ or $X_{t+1} = X_t + \frac{1-p_u}{p_u}(X_t 1 + f(M_t)) > X_t + (X_t 1 + f(M_t)) \ge X_t + (X_t 1) > X_t + \varepsilon'$ because $p_u \le \frac{1}{2}$ implies $\frac{1-p_u}{p_u} \ge 1$.
- (ii) Assume $\gamma_t \leq X_t < 1$, then $\varepsilon' < X_t < 1 \varepsilon'$. Either $X_{t+1} = 1 + f(M_t) \geq 1 > X_t + \varepsilon'$ or $X_{t+1} = X_t \gamma_t$ and thus $X_t X_{t+1} = \gamma_t = \frac{p_u}{1 p_u} (1 + f(M_t) X_t) > \varepsilon (1 + f(M_t) X_t) \geq \varepsilon (1 X_t) > \varepsilon \varepsilon'$.
- (iii) Assume $X_t < \gamma_t$ and $X_t < 1$, then since $0 \le X_t$ by Claim 5, $\varepsilon' < X_t < \gamma_t$ and $X_t < 1 - \varepsilon'$. Either $d_t = \frac{p_u}{1 - p_u} X_t > \varepsilon \varepsilon'$ and we are done, or $d_t = \frac{1 - p_u}{p_u} \gamma_t - 2f(M_t)$. If $X_t \ge \frac{5}{8}$ then $d_t > \frac{1 - p_u}{p_u} X_t - 2f(M_t) > X_t - \frac{1}{2} \ge \frac{1}{8}$, since $f(M_t) \le \frac{\delta}{2} < \frac{1}{4}$. If $X_t < \frac{5}{8}$ then $d_t = 1 - f(M_t) - X_t > \frac{3}{4} - X_t > \frac{1}{8}$. Hence either $X_{t+1} - X_t = d_t > \min\{\varepsilon \varepsilon', \frac{1}{8}\}$ or $X_t - X_{t+1} = \frac{1 - p_u}{p_u} d_t > d_t > \min\{\varepsilon \varepsilon', \frac{1}{8}\}$.

Claim 9: For all $m \in \mathbb{N}$, if $E_{m,m-1} \neq \emptyset$ then $P(E_{m,m} \mid E_{m,m-1}) \geq 1 -$ 2f(m). Let $v \in E_{m,m-1}$ and let $t_0 \in \mathbb{N}$ be a time step such that exactly m-1 upcrossings have been completed up to time t_0 , i.e., $M_{t_0}(v) = m$. The subsequent downcrossing is completed eventually with probability 1: we are in case (i) and in every step there is a chance of $1 - p_u \ge \frac{1}{2}$ of completing the downcrossing. Therefore we assume without loss of generality that the downcrossing has been completed, i.e., that t_0 is such that $X_{t_0}(v) = 1 - f(m)$. We will bound the probability $p := P(E_{m,m} \mid E_{m,m-1})$ that X_t rises above 1 + f(m) after t_0 to complete the *m*-th upcrossing.

Define the stopping time $T: \Sigma^{\omega} \to \mathbb{N} \cup \{\omega\},\$

$$T(v) := \inf\{t \ge t_0 \mid X_t(v) \ge 1 + f(m) \lor X_t(v) = 0\},\$$

and define the stochastic process $Y_t = 1 + f(m) - X_{\min\{t_0+t,T\}}$. Because $(X_{\min\{t_0+t,T\}})_{t\in\mathbb{N}}$ is martingale, $(Y_t)_{t\in\mathbb{N}}$ is martingale. By definition, X_t always stops at 1 + f(m) before exceeding it, thus $X_T \leq 1 + f(m)$, and hence $(Y_t)_{t\in\mathbb{N}}$ is nonnegative. The Optional Stopping Theorem yields $\mathbb{E}[Y_{T-t_0} \mid \mathcal{F}_{t_0}] \leq$ $\mathbb{E}[Y_0 \mid \mathcal{F}_{t_0}]$ and thus $\mathbb{E}[X_T \mid \mathcal{F}_{t_0}] \geq \mathbb{E}[X_{t_0} \mid \mathcal{F}_{t_0}] = 1 - f(m)$. We show that $X_T \in \{0, 1 + f(m)\}$ almost surely. If T is finite then this holds by definition of T. If $T = \omega$ then the random variable X_T is defined as the limit $\lim_{t\to\infty} X_t$. By Claim 8 the limit $X_T \in \{0, 1\}$ and according to Claim 6 we have $X_t \leq 1 - f(M_t)$ for all $t \in \mathbb{N}$, so X_t cannot converge to 1. We conclude that

$$1 - f(m) \le \mathbb{E}[X_T \mid \mathcal{F}_{t_0}] = (1 + f(m)) \cdot p + 0 \cdot (1 - p),$$

hence $P(E_{m,m} \mid E_{m,m-1}) = p \ge 1 - f(m)(1+p) \ge 1 - 2f(m)$.

Claim 10: $E_{m+1,m} = E_{m,m}$ and $E_{m+1,m+1} \subseteq E_{m,m}$. By definition of M_t , the *i*-th upcrossings of the process $(X_t)_{t\in\mathbb{N}}$ is between 1-f(i) and 1+f(i). The function f is monotone decreasing, and by Claim 6 the process $(X_t)_{t\in\mathbb{N}}$ does not assume values between 1 - f(i) and 1 + f(i). Therefore the first m f(m+1)upcrossings are also f(m)-upcrossings, i.e., $E_{m+1,m} \subseteq E_{m,m}$. By definition of $E_{m,k}$ we have $E_{m+1,m} \supseteq E_{m,m}$ and $E_{m+1,m+1} \subseteq E_{m+1,m}$. Claim 11: $P(E_{m,m}) \ge 1 - \sum_{i=1}^{m} 2f(i)$. For $P(E_{0,0}) = 1$ this holds trivially.

Using Claim 9 and Claim 10 we conclude inductively

$$P(E_{m,m}) = P(E_{m,m} \cap E_{m,m-1}) = P(E_{m,m} | E_{m,m-1})P(E_{m,m-1})$$
$$= P(E_{m,m} | E_{m,m-1})P(E_{m-1,m-1})$$
$$\ge (1 - 2f(m))\left(1 - \sum_{i=1}^{m-1} 2f(i)\right) \ge 1 - \sum_{i=1}^{m} 2f(i).$$

From Claim 10 follows $\bigcap_{i=1}^{m} E_{i,i} = E_{m,m}$ and therefore $P(\bigcap_{i=1}^{\infty} E_{i,i}) = \lim_{m \to \infty} P(E_{m,m}) \ge 1 - \sum_{i=1}^{\infty} 2f(i) \ge 1 - \delta$.

Theorem 6 gives a *uniform* lower bound on the probability for many upcrossings: it states the probability of the event that for all $m \in \mathbb{N}$, $U(1-f(m), 1+f(m)) \geq m$ holds. This is a lot stronger than the nonuniform bound $P[U(1-f(m), 1+f(m)) \ge m] \ge 1-\delta$ for all $m \in \mathbb{N}$: the quantifier is inside the probability statement.

As an immediate consequence of Theorem 6, we get the following uniform lower bound on the *expected* number of upcrossings.

Corollary 7 (Expected Upcrossings). Let $0 < \delta < 1/2$ and let $f : \mathbb{N} \to [0, 1)$ be any monotone decreasing function such that $\sum_{i=1}^{\infty} f(i) \leq \delta/2$. For every probability measure P with perpetual entropy there is a nonnegative martingale $(X_t)_{t\in\mathbb{N}}$ with $\mathbb{E}[X_t] = 1$ and for all $m \in \mathbb{N}$,

$$\mathbb{E}[U(1 - f(m), 1 + f(m))] \ge m(1 - \delta).$$

Proof. From Theorem 6 and Markov's inequality.

By choosing a specific slowly decreasing but summable function f, we get the following concrete results.

Corollary 8 (Concrete lower bound). Let $0 < \delta < 1/2$. For every probability measure P with perpetual entropy there is a nonnegative martingale $(X_t)_{t \in \mathbb{N}}$ with $\mathbb{E}[X_t] = 1$ such that

$$P\left[\forall \varepsilon > 0. \ U(1-\varepsilon, 1+\varepsilon) \in \Omega\left(\frac{\delta}{\varepsilon \left(\ln \frac{1}{\varepsilon}\right)^2}\right)\right] \ge 1-\delta \text{ and}$$
$$\mathbb{E}[U(1-\varepsilon, 1+\varepsilon)] \in \Omega\left(\frac{1}{\varepsilon \left(\ln \frac{1}{\varepsilon}\right)^2}\right).$$

Moreover, for all $\varepsilon < 0.015$ we get $\mathbb{E}[U(1-\varepsilon, 1+\varepsilon)] > \frac{\delta(1-\delta)}{\varepsilon(\ln \frac{1}{\varepsilon})^2}$ and

$$P\left[\forall \varepsilon < 0.015. \ U(1-\varepsilon, 1+\varepsilon) > \frac{\delta}{\varepsilon \left(\ln \frac{1}{\varepsilon}\right)^2}\right] \ge 1-\delta.$$

Proof. Define

$$g: (0, e^{-2}] \to [0, \infty), \qquad \varepsilon \mapsto 2\delta \left(\frac{1}{\varepsilon (\ln \varepsilon)^2} - \frac{e^2}{4} \right).$$

We have $g(e^{-2}) = 0$, $\lim_{\varepsilon \to 0} g(\varepsilon) = \infty$, and

$$\frac{dg}{d\varepsilon}(\varepsilon) \ = \ 2\delta\left(\frac{-1}{\varepsilon^2(\ln\varepsilon)^2}+\frac{-2}{\varepsilon^2(\ln\varepsilon)^3}\right) \ = \ -\frac{2\delta(2+\ln\varepsilon)}{\varepsilon^2(\ln\varepsilon)^3} \ < \ 0 \ {\rm on} \ (0,e^{-2}).$$

Therefore the function g is strictly monotone decreasing and hence invertible. Choose $f := g^{-1}$. Using the substitution $t = g(\varepsilon)$, $dt = \frac{dg}{d\varepsilon}(\varepsilon)d\varepsilon$,

$$\begin{split} \sum_{t=1}^{\infty} f(t) &\leq \int_{0}^{\infty} f(t) dt = \int_{g^{-1}(0)}^{g^{-1}(\infty)} f(g(\varepsilon)) \frac{dg}{d\varepsilon}(\varepsilon) d\varepsilon \\ &= 2\delta \left(\int_{e^{-2}}^{0} \frac{-1}{\varepsilon (\ln \varepsilon)^{2}} d\varepsilon + \int_{e^{-2}}^{0} \frac{-2}{\varepsilon (\ln \varepsilon)^{3}} d\varepsilon \right) \\ &= 2\delta \left(\left[\frac{1}{\ln \varepsilon} \right]_{e^{-2}}^{0} + \left[\frac{1}{(\ln \varepsilon)^{2}} \right]_{e^{-2}}^{0} \right) = 2\delta \left(\frac{1}{2} - \frac{1}{4} \right) = \frac{\delta}{2} \end{split}$$

Now we apply Theorem 6 and Corollary 7 to $m := g(\varepsilon)$ and get

$$\begin{split} P\left[U(1-\varepsilon,1+\varepsilon) \geq 2\delta\left(\frac{1}{\varepsilon(\ln\varepsilon)^2} - \frac{e^2}{4}\right)\right] \geq 1-\delta, \text{ and} \\ \mathbb{E}[U(1-\varepsilon,1+\varepsilon)] \geq 2\delta(1-\delta)\left(\frac{1}{\varepsilon(\ln\varepsilon)^2} - \frac{e^2}{4}\right). \end{split}$$
 For $\varepsilon < 0.015$, we have $\frac{1}{\varepsilon(\ln\varepsilon)^2} > \frac{e^2}{2}$, hence $g(\varepsilon) > \frac{\delta}{\varepsilon(\ln\varepsilon)^2}$.

The concrete bounds given in Corollary 8 are *not* the asymptotically optimal ones: there are summable functions that decrease even more slowly. For example, we could multiply the function g with the factor $\sqrt{\ln(1/\varepsilon)}$ (which still is not optimal).

5 Martingale Upper Bounds

In this section we state upper bounds on the probability and expectations of many upcrossings (Dubins' Inequality and Doob's Upcrossing Inequality). We use the construction from the previous section to show that these bounds are asymptotically tight. Moreover, with the following theorem we show that the uniform lower bound on the probability of many upcrossings guaranteed in Theorem 6 is also asymptotically tight.

Every function f is either summable or not. If f is summable, then we can scale it with a constant factor such that its sum is smaller than $\frac{\delta}{2}$, and then apply the construction of Theorem 6. If f is not summable, the following theorem implies that there is no *uniform* lower bound on the probability of having at least *m*-many f(m)-upcrossings.

Theorem 9 (Upper bound on upcrossing rate). Let $f : \mathbb{N} \to [0, 1)$ be a monotone decreasing function such that $\sum_{t=1}^{\infty} f(t) = \infty$. For every probability measure P and for every nonnegative P-martingale $(X_t)_{t\in\mathbb{N}}$ with $\mathbb{E}[X_t] = 1$,

$$P[\forall m. U(1 - f(m), 1 + f(m)) \ge m] = 0.$$

Proof. Define the events $D_m := \bigcup_{i=1}^m E_{i,i}^c = \{ \forall i \leq m. \ U(1 - f(i), 1 + f(i)) \geq i \}$. Then $D_m \subseteq D_{m+1}$. Assume there is a constant c > 0 such that $c \leq P(D_m^c) = P(\bigcap_{i=1}^m E_{i,i})$ for all m. Let $m \in \mathbb{N}, v \in D_m^c$, and pick $t_0 \in \mathbb{N}$ such that the process $X_0(v), \ldots, X_{t_0}(v)$ has completed *i*-many f(i)-upcrossings for all $i \leq m$ and $X_{t_0}(v) \leq 1 - f(m+1)$. If $X_t(v) \geq 1 + f(m+1)$ for some $t \geq t_0$, the (m+1)-st upcrossing for f(m+1) is completed and thus $v \in E_{m+1,m+1}$. Define the stopping time $T : \Sigma^{\omega} \to (\mathbb{N} \cup \{\omega\})$,

$$T(v) := \inf\{t \ge t_0 \mid X_t(v) \ge 1 + f(m+1)\}.$$

According to the Optional Stopping Theorem applied to the process $(X_t)_{t \ge t_0}$, the random variable X_T is almost surely well-defined and $\mathbb{E}[X_T | \mathcal{F}_{t_0}] \le \mathbb{E}[X_{t_0} | \mathcal{F}_{t_0}] = X_{t_0}$. This yields $1 - f(m+1) \ge X_{t_0} \ge \mathbb{E}[X_T | \mathcal{F}_{t_0}]$ and by taking the expectation $\mathbb{E}[\cdot | X_{t_0} \le 1 - f(m+1)]$ on both sides,

$$1 - f(m+1) \ge \mathbb{E}[X_T \mid X_{t_0} \le 1 - f(m+1)]$$

$$\ge (1 + f(m+1))P[X_T \ge 1 + f(m+1) \mid X_{t_0} \le 1 - f(m+1)]$$

by Markov's inequality. Therefore

$$P(E_{m+1,m+1} \mid D_m^c) = P[X_T \ge 1 + f(m+1) \mid X_{t_0} \le 1 - f(m+1)]$$

$$\cdot P[X_{t_0} \le 1 - f(m+1) \mid D_m^c]$$

$$\le P[X_T \ge 1 + f(m+1) \mid X_{t_0} \le 1 - f(m+1)]$$

$$\le \frac{1 - f(m+1)}{1 + f(m+1)} \le 1 - f(m+1).$$

Together with $c \leq P(D_m^c)$ we get

$$P(D_{m+1} \setminus D_m) = P(E_{m+1,m+1}^c \cap D_m^c)$$

= $P(E_{m+1,m+1}^c | D_m^c) P(D_m^c) \ge f(m+1)c.$

This is a contradiction because $\sum_{i=1}^{\infty} f(i) = \infty$:

$$1 \ge P(D_{m+1}) = P\left(\biguplus_{i=1}^{m} (D_{i+1} \setminus D_i)\right) = \sum_{i=1}^{m} P(D_{i+1} \setminus D_i) \ge \sum_{i=1}^{m} f(i+1)c \to \infty.$$

Therefore the assumption $P(D_m^c) \ge c$ for all m is false, and hence we get $P[\forall m. U(1 - f(m), 1 + f(m)) \ge m] = P(\bigcap_{i=1}^{\infty} E_{i,i}) = \lim_{m \to \infty} P(D_m^c) = 0.$

By choosing a specific decreasing non-summable function f for Theorem 9, we get that $U(1-\varepsilon, 1+\varepsilon) \notin \Omega(\frac{1}{\varepsilon \log(1/\varepsilon)})$ *P*-almost surely.

Corollary 10 (Concrete upper bound). Let P be a probability measure and let $(X_t)_{t \in \mathbb{N}}$ be a nonnegative martingale with $\mathbb{E}[X_t] = 1$. Then for all a, b > 0,

$$P\left[\forall \varepsilon > 0. \ U(1-\varepsilon, 1+\varepsilon) \ge \frac{a}{\varepsilon \log(1/\varepsilon)} - b\right] = 0.$$

Proof. We proceed analogously to the proof of Corollary 7. Define

$$g:(0,c]\to [g(c),\infty),\qquad \varepsilon\mapsto \frac{a}{\varepsilon\ln\frac{1}{\varepsilon}}-b$$

with c < 1 and $g(c) \ge 1$. We have $\lim_{\varepsilon \to 0} g(\varepsilon) \to \infty$ and

$$\frac{dg}{d\varepsilon}(\varepsilon) = \frac{-a}{\varepsilon^2 \ln \frac{1}{\varepsilon}} + \frac{-a}{\varepsilon^2 (\ln \frac{1}{\varepsilon})^2} < 0 \text{ on } (0,c].$$

Therefore the function g is strictly monotone decreasing and hence invertible. Choose $f := g^{-1}$. Using the substitution $t = g(\varepsilon)$, $dt = \frac{dg}{d\varepsilon}(\varepsilon)d\varepsilon$,

$$\begin{split} \sum_{t=1}^{\infty} f(t) &\geq \int_{g(c)}^{\infty} f(t)dt = \int_{c}^{g^{-1}(\infty)} f(g(\varepsilon)) \frac{dg}{d\varepsilon}(\varepsilon)d\varepsilon \\ &= \int_{c}^{0} \frac{-a}{\varepsilon \ln \frac{1}{\varepsilon}} d\varepsilon + \int_{c}^{0} \frac{-a}{\varepsilon (\ln \frac{1}{\varepsilon})^{2}} d\varepsilon = \int_{-\ln c}^{-\ln 0} \frac{a}{u} du + \int_{c}^{0} \frac{-a}{\varepsilon (\ln \frac{1}{\varepsilon})^{2}} d\varepsilon \\ &= [a \ln u]_{-\ln c}^{+\infty} + \left[\frac{a}{\ln \frac{1}{\varepsilon}}\right]_{c}^{0} = \infty - a \ln(-\ln c) + 0 - \frac{a}{\ln \frac{1}{\varepsilon}} = \infty. \end{split}$$

Now we apply Theorem 9 to $m := g(\varepsilon)$.

Theorem 11 (Dubins' Inequality [Dub62, Thm. 13.1]). For every nonnegative P-martingale $(X_t)_{t\in\mathbb{N}}$ and for every c > 0 and every $\varepsilon > 0$,

$$P[U(c-\varepsilon, c+\varepsilon) \ge k] \le \left(\frac{c-\varepsilon}{c+\varepsilon}\right)^k \mathbb{E}\left[\min\left\{\frac{X_0}{c-\varepsilon}, 1\right\}\right]$$

Dubins' Inequality immediately yields the following bound on the probability of the number of upcrossings.

$$P[U(1 - f(m), 1 + f(m)) \ge k] \le \left(\frac{1 - f(m)}{1 + f(m)}\right)^k$$

The construction from Theorem 6 shows that this bound is asymptotically tight for $m = k \to \infty$ and $\delta \to 0$: define the monotone decreasing function $f : \mathbb{N} \to [0, 1)$,

$$f(t) := \begin{cases} \frac{\delta}{2m}, & \text{if } t \le m, \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$
(1)

Then the martingale from Theorem 6 yields the lower bound

$$P[U(1 - \frac{\delta}{2k}, 1 + \frac{\delta}{2k}) \ge k] \ge 1 - \delta,$$

while Dubins' Inequality gives the upper bound

$$P[U(1 - \frac{\delta}{2k}, 1 + \frac{\delta}{2k}) \ge k] \le \left(\frac{1 - \frac{\delta}{2k}}{1 + \frac{\delta}{2k}}\right)^k = \left(1 - \frac{2\delta}{2k + \delta}\right)^k \xrightarrow{k \to \infty} \exp(-\delta).$$

As δ approaches 0, the value of $\exp(-\delta)$ approaches $1 - \delta$ (but exceeds it since exp is convex). For $\delta = 0.2$ and m = k = 3, the difference between the two bounds is already lower than 0.021.

The following theorem places an upper bound on the rate of *expected* upcrossings. In Appendix A.3 we discuss different versions of this inequality and prove this inequality tight.

Theorem 12 (Doob's Upcrossing Inequality [Xu12]). Let $(X_t)_{t\in\mathbb{N}}$ be a submartingale. For every $c \in \mathbb{R}$ and $\varepsilon > 0$,

$$\mathbb{E}[U_t(c-\varepsilon, c+\varepsilon)] \le \frac{1}{2\varepsilon} \mathbb{E}[\max\{c-\varepsilon - X_t, 0\}].$$

Asymptotically, Doob's Upcrossing Inequality states that with $\varepsilon \to 0$,

$$\mathbb{E}[U(1-\varepsilon,1+\varepsilon)] \in O\left(\frac{1}{\varepsilon}\right).$$

Again, we can use the construction of Theorem 6 to show that these asymptotics are tight: Let f be as in (1). Then for $\delta = \frac{1}{2}$, Corollary 7 yields a martingale fulfilling the lower bound

$$\mathbb{E}[U(1-\frac{1}{4m},1+\frac{1}{4m})] \ge \frac{m}{2}$$

and Doob's Upcrossing Inequality gives the upper bound

$$\mathbb{E}[U(1-\frac{1}{4m},1+\frac{1}{4m})] \le 2m$$

which differs by a factor of 4. In Theorem 23 we show that Doob's Upcrossing Inequality can also be made exactly tight.

The lower bound for the expected number of upcrossings given in Corollary 7 is a little looser than the upper bound given in Doob's Upcrossing Inequality. Closing this gap remains an open problem. We know by Theorem 9 that given a non-summable function f, the uniform probability for many f(m)-upcrossings goes to 0. However, this does not necessarily imply that expectation also tends to 0; low probability might be compensated for by high value. So for expectation there might be a lower bound larger than Corollary 7, an upper bound smaller than Doob's Upcrossing Inequality, or both.

If we drop the requirement that the rate of upcrossings be uniform, Doob's Upcrossing Inequality is the best upper bound we can give: using the little-o notation, assume there is a smaller upper bound $g(m) \in o(m)$ such that for every martingale process $(X_t)_{t\in\mathbb{N}}$,

$$\mathbb{E}\left[U(1-\frac{1}{m},1+\frac{1}{m})\right] \in o(g(m)).$$
(2)

In the following we sketch how to construct a martingale that violates this bound. Define f(m) := g(m)/m, then $f(m) \to 0$ as $m \to \infty$, so there is an infinite sequence $(m_i)_{i \in \mathbb{N}}$ such that $\sum_{i=0}^{\infty} f(m_i) \leq 1$. We define the martingale process $(X_t)_{t \in \mathbb{N}}$ such that it picks an $i \in \mathbb{N}$ with probability $f(m_i)$, and then becomes a martingale that makes Doob's Upcrossing Inequality tight for upcrossings between $1 - 1/m_i$ and $1 + 1/m_i$: for every *i*, we apply the construction of Theorem 23. This would give the following lower bound on the expected number of upcrossings for each *i*:

$$\forall i \ \mathbb{E}\left[U(1-\frac{1}{m_i}), 1+\frac{1}{m_i})\right] \ge m_i f(m_i) = g(m_i).$$

Since there are infinitely many m_i , we get a contradiction to (2). Using a similar argument, we can show that nonuniformly, Dubins' bound is also the best we can get.

6 Application to the MDL Principle

Let \mathcal{M} be a countable set of probability measures on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$, called *environ*ment class. Let $K : \mathcal{M} \to [0, 1]$ be a function such that $\sum_{Q \in \mathcal{M}} 2^{-K(Q)} \leq 1$, called *complexity function on* \mathcal{M} . Following notation in [Hut09], we define for $u \in \Sigma^*$ the minimal description length model as

$$\mathrm{MDL}^{u} := \operatorname*{arg\,min}_{Q \in \mathcal{M}} \left\{ -\log Q(\Gamma_{u}) + K(Q) \right\}.$$

That is, $-\log Q(\Gamma_u)$ is the (arithmetic) code length of u given model Q, and K(Q) is a complexity penalty for Q, also called *regularizer*. Given data $u \in \Sigma^*$, MDL^u is the measure $Q \in \mathcal{M}$ that minimizes the total code length of data and model.

The following corollary of Theorem 6 states that in some cases the limit $\lim_{t\to\infty} \text{MDL}^{v_{1:t}}$ does not exist with high probability.

Corollary 13 (MDL may not converge). Let P be a probability measure on the measurable space $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ with perpetual entropy. For any $0 < \delta < 1/2$, there is a set of probability measures \mathcal{M} containing P, a complexity function $K : \mathcal{M} \to [0, 1]$, and a measurable set $Z \in \mathcal{F}_{\omega}$ with $P(Z) \ge 1 - \delta$ such that for all $v \in Z$, the limit $\lim_{t\to\infty} MDL^{v_{1:t}}$ does not exist.

Proof. Fix some positive monotone decreasing summable function f (e.g., the one given in Corollary 8). Let $(X_t)_{t\in\mathbb{N}}$ be the *P*-martingale process from Theorem 6. By Theorem 5 there is a probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that

$$X_t(v) = \frac{Q(\Gamma_{v_{1:t}})}{P(\Gamma_{v_{1:t}})}$$

P-almost surely. Choose $\mathcal{M} := \{P, Q\}$ with K(P) := K(Q) := 1. From the definition of MDL and Q it follows that

$$X_t(u) < 1 \iff Q(\Gamma_u) < P(\Gamma_u) \implies \text{MDL}^u = P$$
, and
 $X_t(u) > 1 \iff Q(\Gamma_u) > P(\Gamma_u) \implies \text{MDL}^u = Q.$

For $Z := \bigcap_{m=1}^{\infty} E_{m,m}^{X,f}$ Theorem 6 yields

$$P(Z) = P[\forall m. U(1 - f(m), 1 + f(m)) \ge m] \ge 1 - \delta.$$

For each $v \in Z$, the measure $\text{MDL}^{v_{1:t}}$ alternates between P and Q indefinitely, and thus its limit does not exist.

Crucial to the proof of Corollary 13 is that not only does the process Q/P oscillate indefinitely, it oscillates around the constant $\exp(K(Q) - K(P)) = 1$. This implies that the MDL estimator may keep changing indefinitely, and thus it is inductively inconsistent.

7 Bounds on Mind Changes

Suppose we are testing a hypothesis $H \subseteq \Sigma^{\omega}$ on a stream of data $v \in \Sigma^{\omega}$. Let $P(H \mid \Gamma_{v_{1:t}})$ denote our belief in H at time $t \in \mathbb{N}$ after seeing the evidence $v_{1:t}$. By Bayes' rule,

$$P(H \mid \Gamma_{v_{1:t}}) = P(H) \frac{P(\Gamma_{v_{1:t}} \mid H)}{P(\Gamma_{v_{1:t}})} =: X_t(v).$$

Since X_t is a constant multiple of $P(\cdot | H)/P$ and $P(\cdot | H)$ is a probability measure on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ that is absolutely continuous with respect to P on cylinder sets, the process $(X_t)_{t\in\mathbb{N}}$ is a P-martingale with respect to the filtration $(\mathcal{F}_t)_{t\in\mathbb{N}}$ by Theorem 4. By definition, $(X_t)_{t\in\mathbb{N}}$ is bounded between 0 and 1.

Let $\alpha > 0$. We are interested in the question how likely it is to often change one's mind about H by at least α , i.e., what is the probability for $X_t = P(H \mid \Gamma_{v_{1:t}})$ to decrease and subsequently increase m times by at least α . Formally, we define the stopping times $T'_{0,\nu}(v) := 0$,

$$T'_{2k+1,\nu}(v) := \inf\{t > T'_{2k,\nu}(v) \mid X_t(v) \le X_{T'_{2k,\nu}(v)}(v) - \nu\alpha\},\$$

$$T'_{2k+2,\nu}(v) := \inf\{t > T'_{2k+1,\nu}(v) \mid X_t(v) \ge X_{T'_{2k+1,\nu}(v)}(v) + \nu\alpha\}.$$

and $T'_k := \min\{T'_{k,\nu} \mid \nu \in \{-1,+1\}\}$. (In Davis' notation, $X_{T'_{0,\nu}}, X_{T'_{1,\nu}}, \ldots$ is an α -alternating W-sequence for $\nu = 1$ and an α -alternating M-sequence for $\nu = -1$ [Dav13, Def. 4].) For any $t \in \mathbb{N}$, the random variable

$$A_t^X(\alpha)(v) := \sup\{k \ge 0 \mid T_k'(v) \le t\},\$$

is defined as the number of α -alternations up to time t. Let $A^X(\alpha) := \sup_{t \in \mathbb{N}} A_t^X(\alpha)$ denote the total number of α -alternations.

Setting $\alpha = 2\varepsilon$, the α -alternations differ from ε -upcrossings in three ways: first, for upcrossings, the process decreases below $c - \varepsilon$, then increases above $c + \varepsilon$, and then repeats. For alternations, the process may overshoot $c - \varepsilon$ or $c + \varepsilon$ and thus change the bar for the subsequent alternations, causing a 'drift' in the target bars over time. Second, for α -alternations the initial value of the martingale is relevant. Third, one upcrossing corresponds to two alternations, since one upcrossing always involves a preceding downcrossing. See Figure 2.

To apply our bounds for upcrossings on α -alternations, we use the following lemma by Davis. We reinterpret it as stating that every bounded martingale process $(X_t)_{t\in\mathbb{N}}$ can be modified into a martingale $(Y_t)_{t\in\mathbb{N}}$ such that the probability for many α -alternations is not decreased and the number of alternations



Figure 2: This example process has two upcrossings between $c-\alpha/2$ and $c+\alpha/2$ (completed at the time steps of the vertical orange bars) and four α -alternations (completed when crossing the horizontal blue bars).

equals the number of upcrossings plus the number of downcrossings. A sketch of the proof can be found in Appendix A.4.

Lemma 14 (Upcrossings and alternations [Dav13, Lem. 9]). Let $(X_t)_{t\in\mathbb{N}}$ be a martingale with $0 \leq X_t \leq 1$. There exists a martingale $(Y_t)_{t\in\mathbb{N}}$ with $0 \leq Y_t \leq 1$ and a constant $c \in (\alpha/2, 1 - \alpha/2)$ such that for all $t \in \mathbb{N}$ and for all $k \in \mathbb{N}$,

$$P[A_t^X(\alpha) \ge 2k] \le P[A_t^Y(\alpha) \ge 2k] = P[U_t^Y(c - \alpha/2, c + \alpha/2) \ge k].$$

Theorem 15 (Upper bound on alternations). For every martingale process $(X_t)_{t \in \mathbb{N}}$ with $0 \leq X_t \leq 1$,

$$P[A(\alpha) \ge 2k] \le \left(\frac{1-\alpha}{1+\alpha}\right)^k$$

Proof. We apply Lemma 14 to $(X_t)_{t\in\mathbb{N}}$ and $(1-X_t)_{t\in\mathbb{N}}$ to get the processes $(Y_t)_{t\in\mathbb{N}}$ and $(Z_t)_{t\in\mathbb{N}}$. Dubins' Inequality yields

$$P[A_t^X(\alpha) \ge 2k] \le P[U_t^Y(c_+ - \frac{\alpha}{2}, c_+ - \frac{\alpha}{2}) \ge k] \le \left(\frac{c_+ - \frac{\alpha}{2}}{c_+ + \frac{\alpha}{2}}\right)^k =: g(c_+) \text{ and}$$
$$P[A_t^{1-X}(\alpha) \ge 2k] \le P[U_t^Z(c_- - \frac{\alpha}{2}, c_- - \frac{\alpha}{2}) \ge k] \le \left(\frac{c_- - \frac{\alpha}{2}}{c_- + \frac{\alpha}{2}}\right)^k = g(c_-)$$

for some $c_+, c_- \in (\alpha/2, 1 - \alpha/2)$. Because Lemma 14 is symmetric for $(X_t)_{t \in \mathbb{N}}$ and $(1-X_t)_{t \in \mathbb{N}}$, we have $c_+ = 1-c_-$. Since $P[A_t^X(\alpha) \ge 2k] = P[A_t^{1-X}(\alpha) \ge 2k]$ by the definition of $A_t^X(\alpha)$, we have that both are less than $\min\{g(c_+), g(c_-)\} = \min\{g(c_+), g(1-c_+)\}$. This is maximized for $c_+ = c_- = 1/2$ because g is strictly monotone increasing for $c > \alpha/2$. Therefore

$$P[A_t^X(\alpha) \ge 2k] \le \left(\frac{\frac{1}{2} - \frac{\alpha}{2}}{\frac{1}{2} + \frac{\alpha}{2}}\right)^k = \left(\frac{1 - \alpha}{1 + \alpha}\right)^k.$$

Since this bound is independent of t, it also holds for $P[A^X(\alpha) \ge 2k]$.

The bound of Theorem 15 is the square root of the bound derived by Davis [Dav13, Thm. 10 & Thm. 11].

$$P[A(\alpha) \ge 2k] \le \left(\frac{1-\alpha}{1+\alpha}\right)^{2k} \tag{3}$$

This bound is tight [Dav13, Cor. 13]. A similar bound for upcrossings was proved by Dubins [Dub72, Cor. 1].

Because $0 \leq X_t \leq 1$, the process $(1-X_t)_{t\in\mathbb{N}}$ is also a nonnegative martingale, hence the same upper bounds apply to it. This explains why the result in Theorem 15 is worse than Davis' bound (3): Dubins' bound applies to all nonnegative martingales, while Davis' bound uses the fact that the process is bounded from below and above. For unbounded nonnegative martingales, downcrossings are 'free' in the sense that one can make a downcrossing almost surely successful (as done in the proof of Theorem 6). If we apply Dubins' bound to the process $(1-X_t)_{t\in\mathbb{N}}$, we get the same probability bound for the downcrossings of $(X_t)_{t\in\mathbb{N}}$ (which are upcrossings of $(1-X_t)_{t\in\mathbb{N}}$). Multiplying both bounds yields Davis' bound (3); however, we still require a formal argument why the upcrossing and downcrossing bounds are independent.

The following corollary to Theorem 15 derives an upper bound on the *expected* number of α -alternations.

Theorem 16 (Upper bound on expected alternations). For every martingale $(X_t)_{t \in \mathbb{N}}$ with $0 \leq X_t \leq 1$, the expectation $\mathbb{E}[A(\alpha)] \leq \frac{1}{\alpha}$.

Proof. By Theorem 15 we have $P[A(\alpha) \ge 2k] \le \left(\frac{1-\alpha}{1+\alpha}\right)^k$, and thus

$$\mathbb{E}[A(\alpha)] = \sum_{k=1}^{\infty} P[A(\alpha) \ge k]$$

= $P[A(\alpha) \ge 1] + \sum_{k=1}^{\infty} \left(P[A(\alpha) \ge 2k] + P[A(\alpha) \ge 2k+1] \right)$
 $\le 1 + \sum_{k=1}^{\infty} 2P[A(\alpha) \ge 2k] \le 1 + 2\sum_{k=1}^{\infty} \left(\frac{1-\alpha}{1+\alpha}\right)^k = \frac{1}{\alpha}.$

We now apply the technical results of this section to the martingale process $X_t = P(\cdot | H)/P$, our belief in the hypothesis H as we observe data. The probability of changing our mind k times by at least α decreases exponentially with k (Theorem 15). Furthermore, the expected number of times we change our mind by at least α is bounded by $1/\alpha$ (Theorem 16). In other words, having to change one's mind a lot often is unlikely.

Because in this section we consider martingales that are bounded between 0 and 1, the lower bounds from Section 4 do not apply here. While for the martingales constructed in Theorem 6, the number of 2α -alternations and the number of α -up- and downcrossings coincide, these processes are not bounded. However, we can give a similar construction that is bounded between 0 and 1 and makes Davis' bound asymptotically tight.

8 Conclusion

We constructed an indefinitely oscillating martingale process from a summable function f. Theorem 6 and Corollary 7 give uniform lower bounds on the probability and expectation of the number of upcrossings of decreasing magnitude. In Theorem 9 we proved the corresponding upper bound if the function f is not summable. In comparison, Doob's Upcrossing Inequality and Dubins' Inequality give upper bounds that are not uniform. In Section 5 we showed that for a certain summable function f, our martingales make these bounds asymptotically tight as well.

Our investigation of indefinitely oscillating martingales was motivated by two applications. First, in Corollary 13 we showed that the minimum description length operator may not exist in the limit: for any probability measure P we can construct a probability measure Q such that Q/P oscillates forever around the specific constant that causes $\lim_{t\to\infty} MDL^{v_{1:t}}$ to not converge.

Second, we derived bounds for the probability of changing one's mind about a hypothesis H when observing a stream of data $v \in \Sigma^{\omega}$. The probability $P(H \mid \Gamma_{v_{1:t}})$ is a martingale and in Theorem 15 we proved that the probability of changing the belief in H often by at least α decreases exponentially.

A question that remains open is whether there is a *uniform* upper bound on the *expected* number of upcrossings tighter than Doob's Upcrossing Inequality.

References

- [Dav13] Ernest Davis. Bounding changes in probability over time: It is unlikely that you will change your mind very much very often. Technical report, 2013. https://cs.nyu.edu/davise/papers/dither.pdf.
- [Doo53] Joseph L. Doob. Stochastic Processes. Wiley, New York, 1953.
- [Dub62] Lester E Dubins. Rises and upcrossings of nonnegative martingales. Illinois Journal of Mathematics, 6(2):226–241, 1962.
- [Dub72] Lester E Dubins. Some upcrossing inequalities for uniformly bounded martingales. Symposia Mathematica, IX:169–177, 1972.
- [Dur10] Rick Durrett. Probability: Theory and Examples. Cambridge University Press, 4th edition, 2010.
- [Grü07] Peter D. Grünwald. *The Minimum Description Length Principle*. The MIT Press, Cambridge, 2007.
- [Hut09] Marcus Hutter. Discrete MDL predicts in total variation. In Advances in Neural Information Processing Systems 22 (NIPS'09), pages 817– 825, Cambridge, MA, USA, 2009. Curran Associates.
- [LH14] Jan Leike and Marcus Hutter. Indefinitely oscillating martingales. In Proc. 25th International Conf. on Algorithmic Learning Theory (ALT'14), pages 321–335. Springer, 2014.
- [LS05] Wei Luo and Oliver Schulte. Mind change efficient learning. In Proc. 18th Annual Conference on Learning Theory (COLT'05), volume 3559 of LNAI, pages 398–412, Bertinoro, Italy, 2005. Springer.

- [PH05] Jan Poland and Marcus Hutter. Asymptotics of discrete MDL for online prediction. *IEEE Transactions on Information Theory*, 51(11):3780–3795, November 2005.
- [Ris78] Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.
- [RW94] L Chris G Rogers and David Williams. Diffusions, Markov Processes, and Martingales: Volume 1, Foundations. Cambridge University Press, 2nd edition, 1994.
- [Wal05] Christopher S. Wallace. Statistical and Inductive Inference by Minimum Message Length. Springer, Berlin, 2005.
- [WB68] Christopher S. Wallace and David M. Boulton. An information measure for classification. *Computer Journal*, 11(2):185–194, August 1968.
- [Xu12] Weijun Xu. Martingale convergence theorems. Technical report, 2012. http://people.maths.ox.ac.uk/xu/Martingale_ convergence.pdf.

A Appendix

A.1 Notation

- $\bullet \ := \ denotes \ a \ definition.$
- $A^c := \Sigma^{\omega} \setminus A$ denotes the complement of a measurable set $A \subseteq \Sigma^{\omega}$.
- For a set X, the power set of X is denoted by 2^X .
- $\mathbb{1}_X$ is the characteristic function for a set X, i.e., $\mathbb{1}_X(x) = 1$ if $x \in X$ and 0 otherwise.
- ω is the smallest infinite ordinal.
- \mathbb{N} is the set of natural numbers.
- \mathbb{R} is the set of real numbers.
- For $a, b \in \mathbb{R}$, [a, b] denotes the closed interval with end points a and b; (a, b] and [a, b) denote half-open intervals and (a, b) denotes an open interval.
- The set Σ denotes a finite alphabet. The set of all finite strings of length n is denoted Σ^n , the set of all finite strings is denoted Σ^* , and the set of all infinite strings is denoted Σ^{ω} .
- For a string $u \in \Sigma^*$, |u| denotes the length of u.
- For $v \in \Sigma^{\omega}$, $v_{1:k}$ denotes the first k characters of v.
- $f \in \Omega(g)$ denotes $g \in O(f)$, i.e., $\exists k > 0 \ \exists x_0 \ \forall x \ge x_0$. $g(x) \cdot k \le f(x)$.
- $f \in o(g)$ denotes $\lim_{x \to \infty} \frac{f(x)}{g(x)} = 0.$

A.2 Measures and Martingales

In this section we prove Theorem 4 and Theorem 5, establishing the connecting between measures on infinite strings and martingales.

Proof of Theorem 4. X_t is only undefined if $P(\Gamma_{v_{1:t}}) = 0$. The set

$$\{v \in \Sigma^{\omega} \mid \exists t. \ P(\Gamma_{v_{1:t}}) = 0\}$$

has P-measure 0 and hence $(X_t)_{t \in \mathbb{N}}$ is well-defined almost everywhere.

 X_t is constant on Γ_u for all $u \in \Sigma^t$, and \mathcal{F}_t is generated by a collection of finitely many disjoint sets:

$$\Sigma^{\omega} = \biguplus_{u \in \Sigma^t} \Gamma_u.$$

(a) Therefore X_t is \mathcal{F}_t -measurable.

(b) $\Gamma_u = \biguplus_{a \in \Sigma} \Gamma_{ua}$ for all $u \in \Sigma^t$ and $v \in \Gamma_u$, and therefore

$$\mathbb{E}[X_{t+1} \mid \mathcal{F}_t](v) = \frac{1}{P(\Gamma_u)} \sum_{a \in \Sigma} X_{t+1}(ua) P(\Gamma_{ua}) = \frac{1}{P(\Gamma_u)} \sum_{a \in \Sigma} \frac{Q(\Gamma_{ua})}{P(\Gamma_{ua})} P(\Gamma_{ua})$$
$$\stackrel{(*)}{=} \frac{1}{P(\Gamma_u)} \sum_{a \in \Sigma} Q(\Gamma_{ua}) = \frac{Q(\Gamma_u)}{P(\Gamma_u)} = X_t(v).$$

At (*) we used the fact that Q is absolutely continuous with respect to P on cylinder sets. (If Q were not absolutely continuous with respect to P on cylinder sets there are cases where $P(\Gamma_u) > 0$, $P(\Gamma_{ua}) = 0$, and $Q(\Gamma_{ua}) \neq 0$. Therefore $X_{t+1}(ua)$ does not contribute to the expectation and thus $X_{t+1}(ua)P(\Gamma_{ua}) = 0 \neq Q(\Gamma_{ua})$.)

 $P \ge 0$ and $Q \ge 0$ by definition, thus $X_t \ge 0$. Since $P(\Gamma_{\epsilon}) = Q(\Gamma_{\epsilon}) = 1$, we have $\mathbb{E}[X_0] = 1$.

The following lemma gives a convenient condition for the existence of a measure on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$. It is a special case of the Daniell-Kolmogorov Extension Theorem [RW94, Thm. 26.1].

Lemma 19 (Extending measures). Let $q: \Sigma^* \to [0,1]$ be a function such that $q(\epsilon) = 1$ and $\sum_{a \in \Sigma} q(ua) = q(u)$ for all $u \in \Sigma^*$. Then there exists a unique probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that $q(u) = Q(\Gamma_u)$ for all $u \in \Sigma^*$.

To prove this lemma, we need the following two ingredients.

Definition 20 (Semiring). A set $\mathcal{R} \subseteq 2^{\Omega}$ is called *semiring over* Ω iff

- (a) $\emptyset \in \mathcal{R}$,
- (b) for all $A, B \in \mathcal{R}$, the set $A \cap B \in \mathcal{R}$, and
- (c) for all $A, B \in \mathcal{R}$, there are pairwise disjoint sets $C_1, \ldots, C_n \in \mathcal{R}$ such that $A \setminus B = \bigcup_{i=1}^n C_i$.

Theorem 21 (Carathéodory's Extension Theorem [Dur10, Thm. A.1.1]). Let \mathcal{R} be a semiring over Ω and let $\mu : \mathcal{R} \to [0, 1]$ be a function such that

- (a) $\mu(\Omega) = 1$ (normalization),
- (b) $\mu(\biguplus_{i=1}^{n} A_{i}) = \sum_{i=1}^{n} \mu(A_{i})$ for pairwise disjoint sets $A_{1}, \ldots, A_{n} \in \mathcal{R}$ such that $\biguplus_{i=1}^{n} A_{i} \in \mathcal{R}$ (finite additivity), and
- (c) $\mu(\bigcup_{i\geq 0} A_i) \leq \sum_{i\geq 0} \mu(A_i)$ for any collection $(A_i)_{i\geq 0}$ such that each $A_i \in \mathcal{R}$ and $\bigcup_{i\geq 0} A_i \in \mathcal{R}$ (σ -subadditivity).

Then there is a unique extension $\overline{\mu}$ of μ that is a probability measure on $(\Omega, \sigma(\mathcal{R}))$ such that $\overline{\mu}(A) = \mu(A)$ for all $A \in \mathcal{R}$.

Proof of Lemma 19. We show the existence of Q using Carathéodory's Extension Theorem. Define $\mathcal{R} := \{\Gamma_u \mid u \in \Sigma^*\} \cup \{\emptyset\}.$

- (a) $\emptyset \in \mathcal{R}$.
- (b) For any $\Gamma_u, \Gamma_v \in \mathcal{R}$, either

- u is a prefix of v and $\Gamma_u \cap \Gamma_v = \Gamma_v \in \mathcal{R}$, or
- v is a prefix of u and $\Gamma_u \cap \Gamma_v = \Gamma_u \in \mathcal{R}$, or
- $\Gamma_u \cap \Gamma_v = \emptyset \in \mathcal{R}.$

(c) For any $\Gamma_u, \Gamma_v \in \mathcal{R}$,

- Γ_u \ Γ_v = ⋃_{w∈Σ^{|v|-|u|}\{x}} Γ_{uw} if v = ux, i.e., u is a prefix of v, and
 Γ_u \ Γ_v = Ø otherwise.
- $-u (-v) \rightarrow \cdots \rightarrow \cdots$

Therefore \mathcal{R} is a semiring. By definition of \mathcal{R} , we have $\sigma(\mathcal{R}) = \mathcal{F}_{\omega}$.

1

The function $q: \Sigma^* \to [0,1]$ naturally gives rise to a function $\mu: \mathcal{R} \to [0,1]$ with $\mu(\emptyset) := 0$ and $\mu(\Gamma_u) := q(u)$ for all $u \in \Sigma^*$. We will now check the prerequisites of Carathéodory's Extension Theorem.

- (a) (Normalization.) $\mu(\Sigma^{\omega}) = \mu(\Gamma_{\epsilon}) = q(\epsilon) = 1.$
- (b) (Finite additivity.) Let $\Gamma_{u_1}, \ldots, \Gamma_{u_k} \in \mathcal{R}$ be pairwise disjoint sets such that $\Gamma_w := \biguplus_{i=1}^k \Gamma_{u_i} \in \mathcal{R}$. Let $\ell := \max\{|u_i| \mid 1 \le i \le k\}$, then $\Gamma_w = \biguplus_{v \in \Sigma^\ell} \Gamma_{wv}$. By assumption, $\sum_{a \in \Sigma} q(ua) = q(u)$, thus $\sum_{a \in \Sigma} \mu(\Gamma_{ua}) = \mu(\Gamma_u)$ and inductively we have

$$\mu(\Gamma_{u_i}) = \sum_{s \in \Sigma^{\ell - |u_i|}} \mu(\Gamma_{u_i s}), \tag{4}$$

and

$$\mu(\Gamma_w) = \sum_{v \in \Sigma^{\ell}} \mu(\Gamma_{wv}).$$
(5)

For every string $v \in \Sigma^{\ell}$, the concatenation $wv \in \Gamma_w = \bigcup_{i=1}^k \Gamma_{u_i}$, so there is a unique *i* such that $wv \in \Gamma_{u_i}$. Hence there is a unique string $s \in \Sigma^{\ell-|u_i|}$ such that $wv = u_i s$. Together with (4) and (5) this yields

$$\mu\left(\biguplus_{i=1}^{k}\Gamma_{u_{i}}\right) = \mu(\Gamma_{w}) = \sum_{v\in\Sigma^{\ell}}\mu(\Gamma_{wv}) = \sum_{i=1}^{k}\sum_{s\in\Sigma^{\ell-|u_{i}|}}\mu(\Gamma_{u_{i}s}) = \sum_{i=1}^{k}\mu(\Gamma_{u_{i}}).$$

(c) (σ -subadditivity.) We will show that each Γ_u is compact with respect to the topology \mathcal{O} generated by \mathcal{R} . σ -subadditivity then follows from (b) because every countable union is in fact a finite union.

We will show that the topology \mathcal{O} is the product topology of the discrete topology on Σ . (This establishes that $(\Sigma^{\omega}, \mathcal{O})$ is a Cantor Space.) Every projection $\pi_k : \Sigma^{\omega} \to \Sigma$ selecting the k-th symbol is continuous, since $\pi_k^{-1}(a) = \bigcup_{u \in \Sigma^{k-1}} \Gamma_{ua}$ for every $a \in \Sigma$. Moreover, \mathcal{O} is the coarsest topology with this property, since we can generate every open set $\Gamma_u \in \mathcal{R}$ in the base of the topology by

$$\Gamma_u = \bigcap_{i=1}^{|u|} \pi_i^{-1}(\{u_i\})$$

The set Σ is finite and thus compact. By Tychonoff's Theorem, Σ^{ω} is also compact. Therefore Γ_u is compact since it is homeomorphic to Σ^{ω} via the canonical map $\beta_u : \Sigma^{\omega} \to \Gamma_u, v \mapsto uv$.

From (a), (b), and (c) Carathéodory's Extension Theorem yields a unique probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that $Q(\Gamma_u) = \mu(\Gamma_u) = q(u)$ for all $u \in \Sigma^*$.

Using Lemma 19, the proof of Theorem 5 is now straightforward.

Proof of Theorem 5. We define a function $q: \Sigma^* \to \mathbb{R}$, with

$$q(u) := X_{|u|}(v)P(\Gamma_u)$$

for any $v \in \Gamma_u$. The choice of v is irrelevant because $X_{|u|}$ is constant on Γ_u since it is \mathcal{F}_t -measurable. In the following, we also write $X_t(u)$ if |u| = t to simplify notation.

The function q is non-negative because X_t and P are both non-negative. Moreover, for any $u \in \Sigma^t$,

$$1 = \mathbb{E}[X_t] = \int_{\Sigma^{\omega}} X_t dP \ge \int_{\Gamma_u} X_t dP = P(\Gamma_u) X_t(u) = q(u).$$

Hence the range of q is a subset of [0, 1].

We have $q(\epsilon) = X_0(\epsilon)P(\Gamma_{\epsilon}) = \mathbb{E}[X_0] = 1$ since P is a probability measure and $\mathcal{F}_0 = \{\emptyset, \Sigma^{\omega}\}$ is the trivial σ -algebra. Let $u \in \Sigma^t$.

$$\sum_{a \in \Sigma} q(ua) = \sum_{a \in \Sigma} X_{t+1}(ua) P(\Gamma_{ua}) = \int_{\Gamma_u} X_{t+1} dP$$
$$= \int_{\Gamma_u} \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] dP = \int_{\Gamma_u} X_t dP = P(\Gamma_u) X_t(u) = q(u).$$

By Lemma 19, there is a probability measure Q on $(\Sigma^{\omega}, \mathcal{F}_{\omega})$ such that $q(u) = Q(\Gamma_u)$ of all $u \in \Sigma^*$. Therefore, for all $v \in \Sigma^{\omega}$ and for all $t \in \mathbb{N}$ with $P(\Gamma_{v_{1:t}}) > 0$,

$$X_t(v) = \frac{q(v_{1:t})}{P(\Gamma_{v_{1:t}})} = \frac{Q(\Gamma_{v_{1:t}})}{P(\Gamma_{v_{1:t}})}.$$

Moreover, Q is absolutely continuous with respect to P on cylinder sets since $P(\Gamma_u) = 0$ implies

$$Q(\Gamma_u) = q(u) = X_{|u|}(u)P(\Gamma_u) = 0.$$

A.3 Different Upcrossing inequalities and their tightness

There are different versions of the upcrossing inequality in circulation. Let a < b and let $(X_t)_{t \in \mathbb{N}}$ be a martingale process. Doob [Doo53, VII§3 Thm. 3.3] states

$$\mathbb{E}[U_t(a,b)] \le \frac{1}{b-a} \mathbb{E}[\max\{X_t - a, 0\}].$$
(6)

Durrett [Dur10, Thm. 5.2.7] gives a slightly stronger version:

$$\mathbb{E}[U_t(a,b)] \le \frac{1}{b-a} \Big(\mathbb{E}[\max\{X_t - a, 0\}] - \mathbb{E}[\max\{X_0 - a, 0\}] \Big).$$
(7)

We will prove tight the version stated in Theorem 12 [Xu12, Thm. 1.1]:

$$\mathbb{E}[U_t(a,b)] \le \frac{1}{b-a} \mathbb{E}[\max\{a - X_t, 0\}].$$
(8)

For nonnegative martingales we can estimate $\mathbb{E}[\max\{a - X_t, 0\}] \leq a$ to get a bound independent of t from the upcrossing inequality (8). To get a bound independent of t from (6) or (7), we look at the upcrossings of the martingale process $(-X_t)_{t\in\mathbb{N}}$, which are the downcrossings of $(X_t)_{t\in\mathbb{N}}$. The number of downcrossings differs from the number of upcrossings by at most 1, so we can conclude from (7),

$$\mathbb{E}[U_t^X(a,b)] \le \mathbb{E}[U_t^{-X}(-b,-a)] + 1$$

$$\le \frac{1}{b-a} \Big(\mathbb{E}[\max\{a - X_t, 0\}] - \mathbb{E}[\max\{a - X_0, 0\}] \Big) + 1.$$

The origin of the diversity in upcrossing inequalities stems from the details of their proofs. When we start betting every time the process $(X_t)_{t\in\mathbb{N}}$ falls below aand stop every time it rises above b, our gain at time t is at least $(b-a)U_t(a,b)$ plus some amount R that we gained or lost since we started betting last time in case the last upcrossing has not yet completed. Because we are betting on a martingale, our expected gain is zero, hence $(b-a)\mathbb{E}[U_t(a,b)] = \mathbb{E}[-R]$. The right hand sides of the equations (6), (7), and (8) arise from the way we estimate R from below. The inequality (8) estimates R by taking into account any possible losses ignoring gains since we last started betting at a. Contrarily, (6) estimates R by taking into account any possible gains ignoring losses since we started betting at a. In (7) we additionally suppose that we are betting starting at time 0 and take into account any losses before X_t falls below a for the first time.

Lemma 22 (Tightness Criterion for (8)). Let a < b and let $(X_t)_{t \in \mathbb{N}}$ be a martingale such that

- (a) X_t does not assume any values between a and b, and
- (b) all upcrossings are completed at b and all downcrossings are completed at a:

$$X_{T_{2k}} = b$$
 and $X_{T_{2k+1}} = a$ $\forall k \in \mathbb{N}.$

Then the inequality (8) is tight, i.e.,

$$\mathbb{E}[U_t(a,b)] = \frac{1}{b-a} \mathbb{E}[\max\{a - X_t, 0\}].$$

Proof. This proof essentially follows the proof of Doob's Upcrossing Inequality given in [Xu12]. Define the process

$$D_t(v) := \sum_{k=1}^{\infty} \left(X_{\min\{t, T_{2k}\}}(v) - X_{\min\{t, T_{2k-1}\}}(v) \right).$$

Since all but finitely many terms in the infinite sum are zero, D_t is well-defined. The process $(D_t)_{t\in\mathbb{N}}$ is martingale:

$$\mathbb{E}[D_{t+1} \mid \mathcal{F}_t] = \sum_{k=1}^{\infty} \left(\mathbb{E}[X_{\min\{t+1, T_{2k}\}} \mid \mathcal{F}_t] - \mathbb{E}[X_{\min\{t+1, T_{2k-1}\}} \mid \mathcal{F}_t] \right)$$

Fix some $i \in \mathbb{N}$. Conditioning on \mathcal{F}_t , we know whether $T_i > t$ or $T_i \leq t$ since T_i is a stopping time. In case $T_i > t$ we have $t + 1 \leq T_i$, implying $X_{\min\{t+1,T_i\}} = X_{t+1}$ and thus $\mathbb{E}[X_{\min\{t+1,T_i\}} \mid \mathcal{F}_t] = \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] = X_t = X_{\min\{t,T_i\}}$ because $(X_t)_{t\in\mathbb{N}}$ is martingale. In case $T_i \leq t$ we have $X_{\min\{t+1,T_i\}} = X_{T_i}$ and hence $\mathbb{E}[X_{\min\{t+1,T_i\}} \mid \mathcal{F}_t] = \mathbb{E}[X_{T_i} \mid \mathcal{F}_t] = X_{T_i} = X_{\min\{t,T_i\}}$. In both cases we get $\mathbb{E}[X_{\min\{t+1,T_i\}} \mid \mathcal{F}_t] = X_{\min\{t,T_i\}}$, therefore $\mathbb{E}[D_{t+1} \mid \mathcal{F}_t] = D_t$.

Let $t \in \mathbb{N}$ be some time step, and fix $v \in \Sigma^{\omega}$. Let $U_t := U_t(a, b)$ denote the number of upcrossings that have been completed up to time t. We distinguish the following two cases.

- (i) There is an incomplete upcrossing, $T_{2U_t+1} \leq t < T_{2U_t+2}$.
- (ii) There is no incomplete upcrossing, $T_{2U_t} \leq t < T_{2U_t+1}$.

In case (i) we have $X_t < b$ and therefore $X_t \leq a$ by assumption (a). With assumption (b) we get

$$D_{t} = \sum_{k=1}^{U_{t}} (X_{T_{2k}} - X_{T_{2k-1}}) + X_{t} - X_{T_{2U_{t+1}}}$$

$$= \sum_{k=1}^{U_{t}} (b-a) + X_{t} - a = (b-a)U_{t} + X_{t} - a.$$
(9)

In case (ii) we have $X_t > a$. With assumption (b) we get

$$D_t = \sum_{k=1}^{U_t} (X_{T_{2k}} - X_{T_{2k-1}}) = \sum_{k=1}^{U_t} (b-a) = (b-a)U_t.$$
 (10)

From (9) and (10) follows that

$$D_t = (b - a)U_t + \min\{X_t - a, 0\}.$$

We have $D_0 = 0$ and since $(D_t)_{t \in \mathbb{N}}$ is martingale it follows that $\mathbb{E}[D_t] = 0$. Hence

$$(b-a)\mathbb{E}[U_t] = \mathbb{E}[-\min\{X_t - a, 0\}] = \mathbb{E}[\max\{a - X_t, 0\}].$$

Theorem 23 (Tightness of Doob's Upcrossing Inequality). Let P be a probability measure with perpetual entropy. For all b > a > 0 there is a nonnegative martingale $(X_t)_{t \in \mathbb{N}}$ with $X_0 = a$ that makes Doob's Upcrossing Inequality tight for all t > 0:

$$0 < \mathbb{E}[U_t^X(a,b)] = \frac{1}{b-a} \mathbb{E}[\max\{a - X_t, 0\}].$$

We added the requirement $\mathbb{E}[U_t^X(a,b)] > 0$, because otherwise the constant process $X_t = a$ would trivially make the inequality tight.

Proof. Fix c := (a + b)/2 and set f(t) := (b - a)/(b + a) = (b - a)/(2c); then f(t) < 1 because a > 0. Define $X_0 := X_1 := a/c$. The function f is not summable, but we nonetheless apply the same construction as in Theorem 6: for t > 1 let X_t be defined as in the proof of Theorem 6. We prove that the scaled process $Y_t := c \cdot X_t$ makes the inequality (8) tight. Since $(X_t)_{t \in \mathbb{N}}$ does upcrossings between $1 - f(M_t) = a/c$ and $1 + f(M_t) = b/c$, the scaled process $(Y_t)_{t \in \mathbb{N}}$ does upcrossings between a and b.

By Claim 1 $(X_t)_{t\in\mathbb{N}}$ is a martingale process, and by Claim 5 $X_t \ge 0$, hence this also applies to the scaled process $(Y_t)_{t\in\mathbb{N}}$. We check the criterion given in Lemma 22.

- (a) This holds for X_t for t = 0 and t = 1 according to the definition of $(X_t)_{t \in \mathbb{N}}$. For t > 1 this follows from Claim 6 since $T_1 = 1$ because $X_1 = a/c$.
- (b) From Claim 4, this is fulfilled in cases (i) and (ii). In case (iii) we have $X_{t+1} \leq 1 f(M_t)$, so the process cannot do an (up-)crossing.

It remains to show that $\mathbb{E}[U_t^Y(a,b)] > 0$. By Claim 2 $X_0 > \gamma_0$, so for all $v \in \Sigma^{\omega}$ with $v_0 = a_{\epsilon}$ we have $X_1 = 1 + f(M_t)$, therefore $U_1^X(a/c, b/c)(v) \ge 1$. Since $P(\Gamma_{a_{\epsilon}}) > 0$ by assumption, this yields $\mathbb{E}[U_1^X(a/c, b/c)] > 0$ and hence $\mathbb{E}[U_t^Y(a,b)] > 0$ for all t > 1.

The process from Theorem 23 also gives a tightness result as $t \to \infty$. A weaker lower bound $\mathbb{E}[U^X(a,b)] \geq \frac{a+b}{8(b-a)} - \frac{1}{2}$ can be derived directly from Corollary 7 using $\delta := 1/2$ and

$$f(i) := \begin{cases} \frac{b-a}{b+a}, & \text{if } i \le \frac{b+a}{4(b-a)}, \\ 0, & \text{otherwise}, \end{cases}$$

and scaling the process with (b+a)/2.

Corollary 24 (Asymptotic tightness of Doob's Upcrossing Inequality). Let P be a probability measure with perpetual entropy. For all b > a > 0 there is a nonnegative martingale $(X_t)_{t \in \mathbb{N}}$ with $X_0 = a$ such that

$$\mathbb{E}[U^X(a,b)] = \frac{a}{b-a}$$

Proof. Consider the process $(Y_t)_{t\in\mathbb{N}}$ from the proof of Theorem 23. Since $(X_t)_{t\in\mathbb{N}}$ is a nonnegative martingale, the Martingale Convergence Theorem [Dur10, Thm. 5.2.8] implies that $(X_t)_{t\in\mathbb{N}}$ converges almost surely to a limit $X_{\omega} \geq 0$. This limit can only be 0 or 1 by Claim 8. Since f(t) = (b-a)/(2c) > 0 for all t, $(X_t)_{t\in\mathbb{N}}$ does not converge to 1 by Claim 6 $(T_1 = 1 \text{ by construction})$. Thus $X_{\omega} = 0$ almost surely, but this generally does not imply $\lim_{t\to\infty} \mathbb{E}[X_t] = 0$ [Dur10, Ex. 5.2.3]. However, $\max\{a - Y_t, 0\} = \max\{a - cX_t, 0\}$ is bounded, therefore uniformly integrable. By [Dur10, Thm. 5.5.2] (a generalization of the dominated convergence theorem),

$$\lim_{t \to \infty} \mathbb{E}[\max\{a - Y_t, 0\}] = \mathbb{E}[\max\{a - cX_\omega, 0\}] = a.$$
(11)

By Dubins' Inequality,

$$\begin{split} \mathbb{E}[U_t^Y(a,b) \cdot \mathbbm{1}_{U_t^Y(a,b) \ge k}] &= \sum_{i=k}^{\infty} P[U_t^Y(a,b) \ge i] \le \sum_{i=k}^{\infty} a^i b^{-i} \\ &= a^k b^{-k} \frac{b}{b-a} \to 0 \text{ as } k \to \infty, \end{split}$$

hence $U_t^Y(a, b)$ is also uniformly integrable and by the same theorem [Dur10, Thm. 5.5.2] and (11),

$$(b-a)\mathbb{E}[U^Y(a,b)] = \lim_{t \to \infty} (b-a)\mathbb{E}[U^Y_t(a,b)] = \lim_{t \to \infty} \mathbb{E}[\max\{a-Y_t,0\}] = a. \quad \Box$$

The same process can also be used to show that Dubins' Inequality is tight. For a specific underlying probability measure a proof of this is sketched by Dubins [Dub62, Thm. 12.1]. We prove a version that is agnostic with respect to the probability measure P. **Corollary 25** (Tightness of Dubins' Inequality). Let P be a probability measure with perpetual entropy. For all b > a > 0 there is a nonnegative martingale $(X_t)_{t \in \mathbb{N}}$ with $X_0 = a$ that makes Dubins' Inequality tight:

$$P[U^X(a,b) \ge k] = \frac{a^k}{b^k}$$

Proof. We use Dubins' Inequality on the process $(X_t)_{t\in\mathbb{N}}$ from Corollary 24;

$$\mathbb{E}[U^X(a,b)] = \sum_{k=1}^{\infty} P[U^X(a,b) \ge k] \le \sum_{k=1}^{\infty} \frac{a^k}{b^k} = \frac{a}{b-a} = \mathbb{E}[U^X(a,b)],$$

so the involved inequalities must in fact be equalities.

A.4 Davis' Lemma

We do not reproduce Davis' proof in detail. It needs to be adapted to the martingale setting, which is quite cumbersome to do. Below we give an outline of the proof.

Proof sketch for Lemma 14. This proof relies on the observation that the probability that X_t rises (falls) by at least α does not decrease as α decreases. Formally, we argued in the proof of Theorem 9 that $P(X_T \ge x + \alpha \mid X_t = x) \le \frac{x}{x+\alpha}$ using the Optional Stopping Theorem. Since $(X_t)_{t\in\mathbb{N}}$ is bounded by 1 from above, the same argument can be carried out for the process $(1 - X_t)_{t\in\mathbb{N}}$, giving an analogous bound $P(X_T \le x - \alpha \mid X_t = x) \le \frac{1-x}{1-x+\alpha}$ when $(X_t)_{t\in\mathbb{N}}$ decreases. These bounds are tight.

The idea of the proof is to define a martingale process $(Z_t)_{t\in\mathbb{N}}$; the process defined by $Y_t := X_t + Z_t$ is then a martingale. We need to show that $(Y_t)_{t\in\mathbb{N}}$ has the desired properties: $0 \leq Y_t \leq 1$ and the probability of having at least 2k 2α -alternations of $(X_t)_{t\in\mathbb{N}}$ does not exceed the probability of having at least k α -upcrossings of $(Y_t)_{t\in\mathbb{N}}$.

There are two sources of misalignment between 2ε -alternations and ε upcrossings; we consider them in turn.

First, drift: if the martingale $(X_t)_{t\in\mathbb{N}}$ overshoots the target and becomes larger than $X_{T'_{2k+1}} + \alpha$ or smaller than $X_{T'_{2k}} - \alpha$, it changes the target in subsequent alternations. Without loss of generality, consider the first case. Suppose we are in time step t, have observed $u \in \Sigma^t$ and 2k + 1 alternations have been completed, i.e., $T'_{2k+1} \leq t < T'_{2k+2}$. By observing a symbol $a \in \Sigma$, we would have $X_{t+1}(ua) \geq X_{T'_{2k+1}}(u) + \alpha$ with a possible overshoot $\gamma := X_{t+1}(ua) - (X_{T'_{2k+1}} + \alpha)$. To compensate, we set $Z_{t+1}(ua) = Z_t(u) - \gamma$ and $Z_{t+1}(ub) \geq Z_t$ appropriately for $b \in \Sigma \setminus \{a\}$ such that Z_t fulfills the martingale condition (b) of Definition 2. Removing the overshoots from the martingale makes upcrossings and alternations coincide, i.e. $A_t^Y(\alpha) = 2U_t^Y(c-\alpha/2, c+\alpha/2)$ for a suitable constant c, which we will discuss below. According to the aforementioned observation, the new martingale is at least as likely to complete the alternation as the old one, since we have reduced the distance needed to be traveled.

Second, the initial value Y_0 . Let $c \in (\alpha/2, 1-\alpha/2)$ and define $Z_0 := c+\alpha/2-X_0$. The constant c denotes the center of the alternations, i.e., Y_n alternates between $c-\alpha/2$ and $c+\alpha/2$, since $Y_0 = X_0+Z_0 = c+\alpha/2$. What value should we

assign to c? Since we only care about cases where the number of alternations is even, c = 1/2 maximizes the probability of successful upcrossings [Dav13, Lem. 7]. This intuitively makes sense: there is an equal number of up- and downcrossings and the probability of each of them being successful depends on the process' distance from 0 or 1 respectively.

At this point we have a martingale process $(Y_t)_{t\in\mathbb{N}}$ that is bounded between 0 and 1, and upcrossings and alternations coincide: $A_t^Y(\alpha) = 2U_t^Y(\frac{1-\alpha}{2}, \frac{1+\alpha}{2})$. It remains to show that the probability of at least k alternations has not decreased compared to the process $(X_t)_{t\in\mathbb{N}}$. By construction, this is already to case for single down- and upcrossings. However, there could be cases where the drift that we removed from the process would cause us to move to a region where successful alternations are more likely. But since we centered the process optimally, this is not possible.

There is one other technical problem that we glossed over: we have to make sure that the process $(Y_t)_{t\in\mathbb{N}}$ exceeds neither 0 nor 1; We have to stop the process at these points. Moreover, if the process Y_t reaches $1 + Z_t$ or Z_t but its value is in (0, 1) instead of stopping it, we switch to a random walk until we 'get back on track'.