This proves the displayed bound in (b).

(c) The optimal error cannot increase with more observations, because a rule can ignore the additional toss. For odd m, $E_m \geq E_{m+1}$, and the even-case bound with $m+1$ gives the stated exponent $2\lceil m/2 \rceil \varepsilon^2/(1-\varepsilon^2)$. The affinity proof in fact also works directly for odd m.

(d) If $E_m \leq \delta < 1/4$, the strict bound forces

$$2\lceil m/2 \rceil > \frac{1-\varepsilon^2}{\varepsilon^2} \log \frac{1}{8\delta(1-2\delta)}.$$

For fixed $\delta < 1/4$ this is $\Omega(\varepsilon^{-2})$; as $\delta \downarrow 0$ it is $\Omega(\varepsilon^{-2}\log(1/\delta))$ in the small-bias regime. This is a necessary bound, not sufficient. Hoeffding applied to the threshold rule gives $E_m \leq e^{-m\varepsilon^2/2}$, so $m \geq 2\varepsilon^{-2}\log(1/\delta)$ is sufficient. Together with (e), the lower bound applies to every decision rule.

Structural or analytic derivation

The sufficient statistic is the count of heads. Grouping the 2^m possible sequences by that count reduces exact error computation to a binomial tail with at most $m+1$ terms. The likelihood-ratio proof establishes that no more elaborate classifier improves the error for the declared equal-prior model.

Comparative cost and preferred method

Exact evaluation costs $O(m)$ recurrence steps for the binomial probabilities, versus 2^m sequence enumeration. Runtime to read the tosses is necessarily $\Omega(m)$ for an exact head count. The lower and upper sample-complexity bounds agree in order; they are not an exact least-integer sample threshold, which can be found from the binomial error.