6. (10 points) HMMs and Particle Filtering

Consider a Markov Model with a binary state X (i.e., X_t is either 0 or 1). The transition probabilities are given as follows:

X_t	X_{t+1}	$P(X_{t+1} \mid X_t)$
0	0	0.9
0	1	0.1
1	0	0.5
1	1	0.5

(a) (2 pt) The prior belief distribution over the initial state X_0 is uniform, i.e., $P(X_0 = 0) = P(X_0 = 1) = 0.5$. After one timestep, what is the new belief distribution, $P(X_1)$?

X_1	$P(X_1)$
0	0.7
1	0.3

Since the prior of X_0 is uniform, the belief at the next step is the transition distribution from X_0:
$p(X_1 = 0) = p(X_0 = 0)p(X_1 = 0|X_0 = 0) + p(X_0 = 1)p(X_1 = 0|X_0 = 1) = .5(.9) + .5(.5) = .7$.
$p(X_1 = 1) = p(X_0 = 0)p(X_1 = 1|X_0 = 0) + p(X_0 = 1)p(X_1 = 1|X_0 = 1) = .5(.1) + .5(.5) = .3$.

Now, we incorporate sensor readings. The sensor model is parameterized by a number $\beta \in [0, 1]$:

X_t	E_t	$P(E_t \mid X_t)$
0	0	β
0	1	$(1 - \beta)$
1	0	$(1 - \beta)$
1	1	β

(b) (2 pt) At $t = 1$, we get the first sensor reading, $E_1 = 0$. Use your answer from part (a) to compute $P(X_1 = 0 \mid E_1 = 0)$. Leave your answer in terms of β.

$$p(X_1 = 0|E_1 = 0) = \frac{p(E_1 = 0|X_1 = 0)p(X_1 = 0)}{\sum_x p(E_1 = 0|X_1 = x)p(X_1 = x)} = \frac{\beta(0.7)}{\beta(0.7) + (1 - \beta)(0.3)}$$

(c) (2 pt) For what range of values of β will a sensor reading $E_1 = 0$ increase our belief that $X_1 = 0$? That is, what is the range of β for which $P(X_1 = 0 \mid E_1 = 0) > P(X_1 = 0)$?
$\beta \in (0.5, 1]$. Intuitively, observing $E_1 = 0$ will only increase the belief that $X_1 = 0$ if $E_1 = 0$ is more likely under $X_1 = 0$ than not. Note that $\beta > 0.5$; $\beta = 0.5$ is uninformative since the conditional distribution is uniform. This can be verified algebraically by setting $p(X_1 = 0) = p(X_1 = 0|E_1 = 0)$ and solving for β.

(d) (2 pt) Unfortunately, the sensor breaks after just one reading, and we receive no further sensor information. Compute $P(X_\infty \mid E_1 = 0)$, the stationary distribution *very many* timesteps from now.

X_∞	$P(X_\infty \mid E_1 = 0)$
0	$\frac{5}{6}$
1	$\frac{1}{6}$

The stationary distribution π for transition matrix P satisfies $\pi = \pi P$. It is purely a function of the transition matrix and not the prior or past observations.

Determine π by setting up the matrix equation $\pi = \pi P$ with the additional equation $\pi_0 + \pi_1 = 1$ from the sum-to-one constraint of the probabilities and solving.

(e) **(2 pt)** How would your answer to part (d) change if we never received the sensor reading E_1, i.e. what is $P(X_\infty)$ given no sensor information?

X_∞	$P(X_\infty)$
0	$\frac{5}{6}$
1	$\frac{1}{6}$

The stationary distribution does not depend on past observations, so the distribution is unchanged whether E_1 is observed or not.