

Explain the components of this code:

with pm.Model() as model:

```
# ---- Population means ----
mu_beta = pm.Normal(
    "mu_beta",
    mu=[0.0, 0.0, 0.0],
    sigma=[10.0, 5.0, 5.0],
    shape=3
)

# ---- Covariance of random effects ----
chol, corr, sigmas = pm.LKJCholeskyCov(
    "chol",
    n=3,
    eta=2.0,
    sd_dist=pm.Exponential.dist(1.0),
    compute_corr=True
)

# ---- Non-centered random effects ----
z = pm.Normal(
    "z",
    mu=0,
    sigma=1,
    shape=(n_players, 3)
)

beta = pm.Deterministic(
    "beta",
    mu_beta + pt.dot(z, chol.T)
)

beta0 = beta[:, 0]
beta1 = beta[:, 1]
beta2 = beta[:, 2]
```

```
# ---- Linear predictor ----
mu = (
    beta0[player_idx]
    + beta1[player_idx] * age_c
    + beta2[player_idx] * age_c**2
)

# ---- Observation noise ----
sigma = pm.Exponential("sigma", 1.0)

# ---- Likelihood ----
points_obs = pm.Normal(
    "points_obs",
    mu=mu,
    sigma=sigma,
    observed=points
)

trace = pm.sample(
    2000,
    tune=2000,
    target_accept=0.9,
    chains=4
)
```

This is a **hierarchical (multilevel) Bayesian regression model** written in PyMC. I'll explain it section by section, focusing on *what each component represents statistically* and *why it's used*.

1. Model context

python

 Copy code

```
with pm.Model() as model:
```

This creates a PyMC probabilistic model. Everything inside this block defines random variables and relationships that make up the Bayesian model.

2. Population-level means (hyperpriors)

python

 Copy code

```
mu_beta = pm.Normal(
```

 ChatGPT ▾

Log in

Sign up for free

?

```
    sigma=[10.0, 5.0, 5.0],  
    shape=3  
)
```

What this is:

- `mu_beta` represents the **population-level mean coefficients** for three regression terms:
 1. Intercept
 2. Linear age effect
 3. Quadratic age effect

Why it's needed:

- Each player has their own regression coefficients, but those coefficients are assumed to come from a shared population distribution.
- `mu_beta` is the **mean of that population distribution**.

Interpretation:

- On average across all players:
 - Intercept ≈ 0
 - Linear age effect ≈ 0
 - Quadratic age effect ≈ 0
- The large `sigma` values make these **weakly informative priors**.

3. Covariance structure of random effects

python

Copy code

```
chol, corr, sigmas = pm.LKJCholeskyCov(
    "chol",
    n=3,
    eta=2.0,
    sd_dist=pm.Exponential.dist(1.0),
    compute_corr=True
)
```

What this is:

This defines the **covariance matrix** for the 3 random effects using an LKJ prior.

It returns:

- chol : Cholesky factor of the covariance matrix
- corr : Correlation matrix between coefficients
- sigmas : Standard deviations of each coefficient

Why it's needed:

- Player-specific intercepts and slopes may be **correlated**.
 - Example: Players with higher baseline performance may age differently.
- The LKJ prior allows flexible correlation while remaining stable for sampling.

Interpretation:

- eta=2.0 mildly favors lower correlations (but does not forbid strong ones).
- sd_dist=Exponential(1.0) controls how much players can deviate from the population mean.

4. Non-centered parameterization

python

Copy code

```
z = pm.Normal(
    "z",
    mu=0,
    sigma=1,
```

```
    shape=(n_players, 3)
)
```

What this is:

- `z` is a matrix of **standard normal latent variables**.
- One row per player, one column per coefficient.

Why it's used:

- This is a **non-centered parameterization**, which:
 - Improves sampling efficiency
 - Reduces divergences
 - Is especially important in hierarchical models

5. Player-specific coefficients

python

 Copy code

```
beta = pm.Deterministic(
    "beta",
    mu_beta + pt.dot(z, chol.T)
)
```

What this does:

- Transforms `z` into actual player-level coefficients.
- Each player gets:
 - Their own intercept
 - Their own linear age effect
 - Their own quadratic age effect

Mathematically:

$$\beta_i \sim \mathcal{N}(\mu_\beta, \Sigma)$$

where `$\Sigma = chol @ chol.T$` .

6. Split coefficients for readability

python

Copy code

```
beta0 = beta[:, 0]
beta1 = beta[:, 1]
beta2 = beta[:, 2]
```

These are:

- `beta0` : player-specific intercepts
- `beta1` : player-specific linear age slopes
- `beta2` : player-specific quadratic age slopes

7. Linear predictor

python

Copy code

```
mu = (
    beta0[player_idx]
    + beta1[player_idx] * age_c
    + beta2[player_idx] * age_c**2
)
```

What this is:

- The expected value of the outcome (`points`) for each observation.
- `player_idx` maps each observation to the correct player.

Interpretation:

This is a **player-specific quadratic regression**:

$$\mu_{ij} = \beta_{0,i} + \beta_{1,i} \cdot \text{age}_j + \beta_{2,i} \cdot \text{age}_j^2$$

where:

- i = player
- j = observation

8. Observation noise

python

 Copy code

```
sigma = pm.Exponential("sigma", 1.0)
```

What this is:

- Residual standard deviation of the observation model.

Why:

- Allows for unexplained variability in points not captured by age or player effects.

9. Likelihood (data model)

python

 Copy code

```
points_obs = pm.Normal(  
    "points_obs",  
    mu=mu,  
    sigma=sigma,  
    observed=points  
)
```

What this does:

- Connects the model to the observed data.
- Assumes points are normally distributed around the predicted mean.

10. Sampling

python

 Copy code

```
trace = pm.sample(  
    2000,  
    tune=2000,  
    target_accept=0.9,  
    chains=4  
)
```

What happens here:

- Draws posterior samples using NUTS (Hamiltonian Monte Carlo).
- `tune=2000` : warm-up steps
- `target_accept=0.9` : more conservative step size (helps avoid divergences)
- `chains=4` : multiple independent chains for convergence diagnostics

11. Big-picture summary

This model is:

- ✓ A hierarchical Bayesian regression
- ✓ With player-specific intercepts and age curves
- ✓ Using a non-centered parameterization
- ✓ With correlated random effects
- ✓ Estimated using modern HMC sampling

If you want, I can also:

- Translate this into math notation
- Explain how to interpret the posterior output
- Suggest diagnostics or model improvements

