FlatPPL 0.1 / QUICK REFERENCE

Zero to hero

ONE RATE, UNEQUAL EXPOSURES

$$n_i \sim \text{Poisson}(rt_i)$$

COMPLETE EXAMPLE

One rate from three counts

```
1  %%%
2  A shared event rate in three exposures.
3  %%%
4  flatppl_compat = "0.1"
6  exposure = [1.0, 2.0, 1.0]
7  counts = [2, 3, 1]
8  rate ~ Gamma(shape = 2.0, rate = 1.0)
10 expected = rate .* exposure
11 n ~ Poisson.(expected)
13 prior = lawof(record(rate = rate))
14 K = kernelof(record(n = n), rate = rate)
15 L = likelihoodof(K, record(n = counts))
16 posterior = bayesupdate(L, prior)
18 point = record(rate = 1.5)
19 log_L = logdensityof(L, point)
20 predict = functionof(expected, rate = rate)
21 fitted = predict(rate = point.rate)
23 # Exact update for this Poisson model.
24 shape_post = 2.0 + sum(counts)
25 rate_post = 1.0 + sum(exposure)
26 mean_rate = shape_post / rate_post
28 # Predict the count in two new exposure units.
29 rate_law = Gamma(shape_post, rate_post)
30 r, state = rand(rate_law, rnginit(41))
31 next_n, state2 = rand(Poisson(2.0 * r), state)
```

A likelihood without a prior

`rate = elementof(posreals)`
For a likelihood fit, replace the Gamma draw on line 8 with this declaration and keep `K` and `L`. For Bayesian inference, restore the Gamma draw before constructing `prior` and `posterior`.

Check the log likelihood

At `rate = 1.5` the three log probabilities sum to `log_L` ≈ -3.97267. Each column pairs one count with its own exposure.

Exposure	1	2	1
Observed count	2	3	1
Poisson mean	1.5	3	1.5
Log probability	−1.38222	−1.49592	−1.09453

ANNOTATED LINES

A Set the data and the support 6–8

```
exposure = [1.0, 2.0, 1.0]
counts = [2, 3, 1]
rate ~ Gamma(shape = 2.0, rate = 1.0)
```

Entry `i` pairs a count with its exposure. The Gamma prior has shape 2 and rate 1, giving a prior mean rate of 2.

B Build the observation law 10–11

```
expected = rate .* exposure
n ~ Poisson.(expected)
```

Use `=` for a calculation and `~` for a draw. Multiplication by `rate` acts elementwise. The dotted Poisson constructor gives independent counts at the resulting means.

C Condition on the observed counts 13–16

```
prior = lawof(record(rate = rate))
K = kernelof(record(n = n), rate = rate)
L = likelihoodof(K, record(n = counts))
posterior = bayesupdate(L, prior)
```

At each rate, `K` gives a law for the record field `n`. Supply the observed counts to `likelihoodof`, then pass the likelihood and prior to `bayesupdate`.

D Score a point and predict 18–21

```
point = record(rate = 1.5)
log_L = logdensityof(L, point)
predict = functionof(expected, rate = rate)
fitted = predict(rate = point.rate)
```

Evaluate `logdensityof` at a record with the field `rate`. The function returned by `functionof` gives the expected counts `[1.5, 3.0, 1.5]` at that rate.

Analytic rate estimate

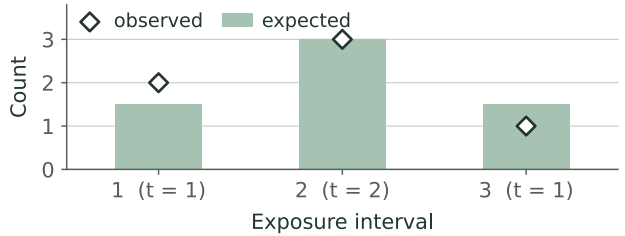
The counts sum to 6 over 4 exposure units. The likelihood peaks at $\frac{6}{4} = 1.5$. The exact posterior is `Gamma(8, 5)` with mean 1.6.

Check the file before fitting

Save the listing as `model.flatppl`. Run `flatppl infer model.flatppl model.flatpir` to check names, types and shapes. Then compare `log_L` with the hand calculation.

WRITING A MODEL

Expected and observed counts



At `rate = 1.5`, doubling the exposure doubles the expected count. The diamonds are the three observed counts.

Choose the data shape

A scalar parameter can serve every row. Use `iid(M, n)` for repeated draws from one law. Use a dotted constructor when its parameters vary across rows.

Change the observation family

```
y ~ Normal.(expected, 0.5)
hits ~ Binomial.([10, 20, 10], 0.3)
```

Use a Normal law for continuous readings. Use a Binomial law for successes out of known trial totals. Name the observed field in both `K` and `L`.

Accept data from the host

```
counts = external(cartpow(nonnegintegers, 3))
```

Declare the counts as external and supply a length-three vector before evaluating `L`. Use the same field name, `n`, in the kernel and the observation record.

Extend the predictor

```
beta = elementof(reals)
eta = log(rate) .+ beta .* [-1, 0, 1]
expected = exposure .* exp.(eta)
```

Replace `expected` on line 10. The log rate is the intercept and `beta` is the slope. For a likelihood fit, put both names in `K` and `point`.

Estimation and prediction

Maximize `L` to estimate the rate, or sample `posterior` to quantify its uncertainty. For prediction, draw a rate from the exact Gamma posterior. The mean count for the next exposure is $2 \times 1.6 = 3.2$.

Jessica Cox · bcox@mpp.mpg.de