MANIPAL COLLEGE
OF PHARMACEUTICAL SCIENCES
MANIPAL
(A constituent unit of MAHE, Manipal)

Poisson Distribution

Pharmacy Applications & Python Analysis

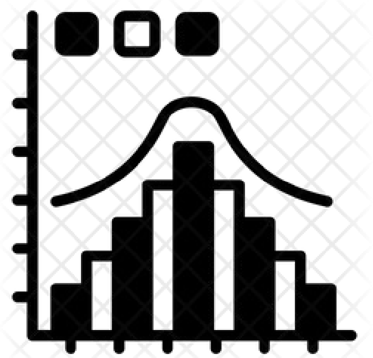

Dr. Pawan G. Nayak
Assistant Professor, Pharmacology
Manipal College of Pharmaceutical Sciences
Manipal Academy of Higher Education
Manipal – 576104, Karnataka, India.
Email: pawan.nayak@manipal.edu

Event counts

- A Poisson process models discrete events that occur randomly over time, area or volume.
- Individual event times are unpredictable, but a stable process has an approximately constant average rate.
- *Example:* dispensing-error reports received during each 8-hour hospital-pharmacy shift.
- Counts in disjoint intervals are assumed independent unless a shared cause creates clustering.

Random timing

- Illustrative ADR reports over 15 days.

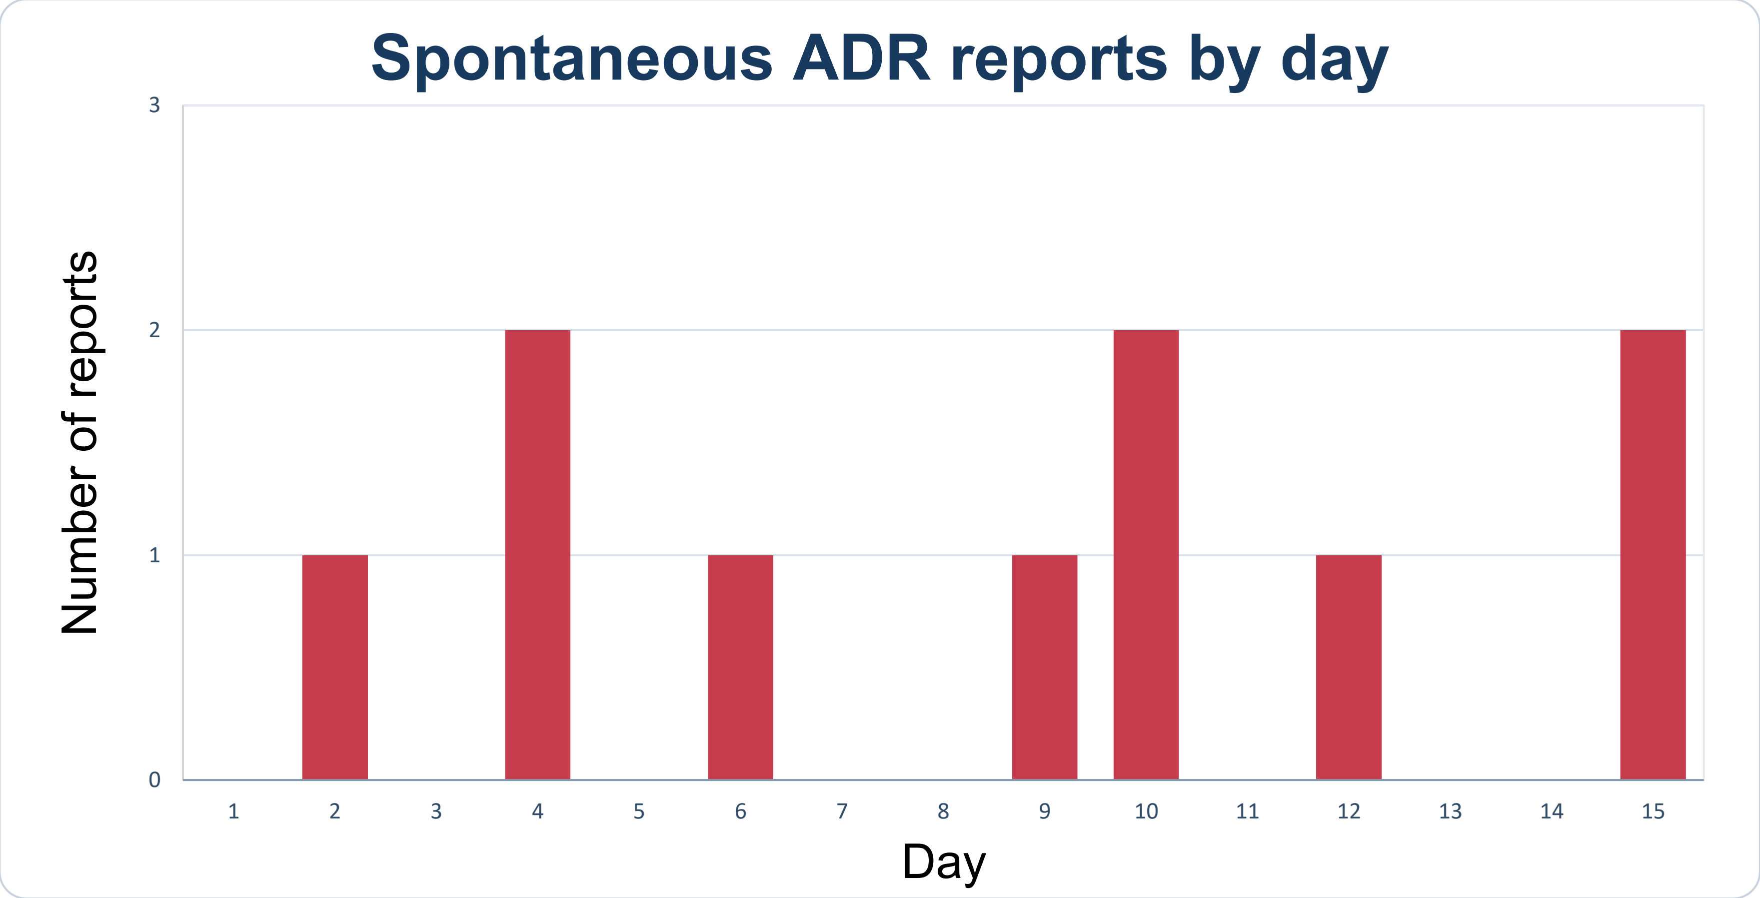

Assumptions

- Events are counted in a clearly defined interval.
- Counts in disjoint intervals are independent.
  - The average event rate is approximately constant within the interval.
  - The chance of more than one event in a very small sub-interval is negligible.
  - Look for clustering, seasonality, unequal exposure or overdispersion before using the model.

Count interval

- Define the exposure before counting: per shift, per batch, per 1,000 medication orders, per culture plate or per month.
- Each interval produces 0, 1, 2, … events.
- If exposure differs between intervals, use $\lambda = \text{rate} \times \text{exposure}$; raw counts should not be compared directly.

Pharmacy uses

- Prescriptions arriving per 15 minutes
- Dispensing incidents per work shift
- Spontaneous ADR reports per month
- Microbial colonies per plate or air sample
- Defective tablets in a sampled batch segment
- Equipment breakdowns per quarter

From process to probability

- The process describes how events occur. The Poisson distribution lets us calculate:
- probability of exactly k events;
- probability of at most, at least or more than a threshold;
- expected frequencies across many comparable intervals;
- and whether observed variation is compatible with a stable random process.

Poisson PMF

- Let X be the number of pharmacy events in a fixed interval.
- The Poisson probability mass function gives P(X = k) from the expected count λ, where k = 0, 1, 2, …

$$P(k \text{ events in time period}) = e^{-\frac{\text{events}}{\text{time}} * \text{time period}} * \frac{\left(\frac{\text{events}}{\text{time}} * \text{time period}\right)^{k}}{k!}$$

Rate × exposure = λ

- λ = average event rate × exposure in one interval.
- After this substitution, the Poisson probability depends on one parameter: λ, the expected count in that interval.

$$P(k \text{ events in interval}) = e^{-\lambda}\frac{\lambda^k}{k!}$$

Interpreting λ

- λ is the expected number of events in one defined interval; for a Poisson distribution, it is also the variance.
- Examples: 0.02 defects/tablet × 200 tablets gives λ = 4 defects; 0.625 reports/hour × 8 hours gives λ = 5 reports.

How λ changes the curve

- A larger λ shifts probability toward larger counts and increases the spread.
- Small λ values produce a strongly right-skewed distribution; as λ grows, the distribution becomes more symmetric.
- λ can change because the event rate changes, the exposure changes, or both.

Probability curves by λ

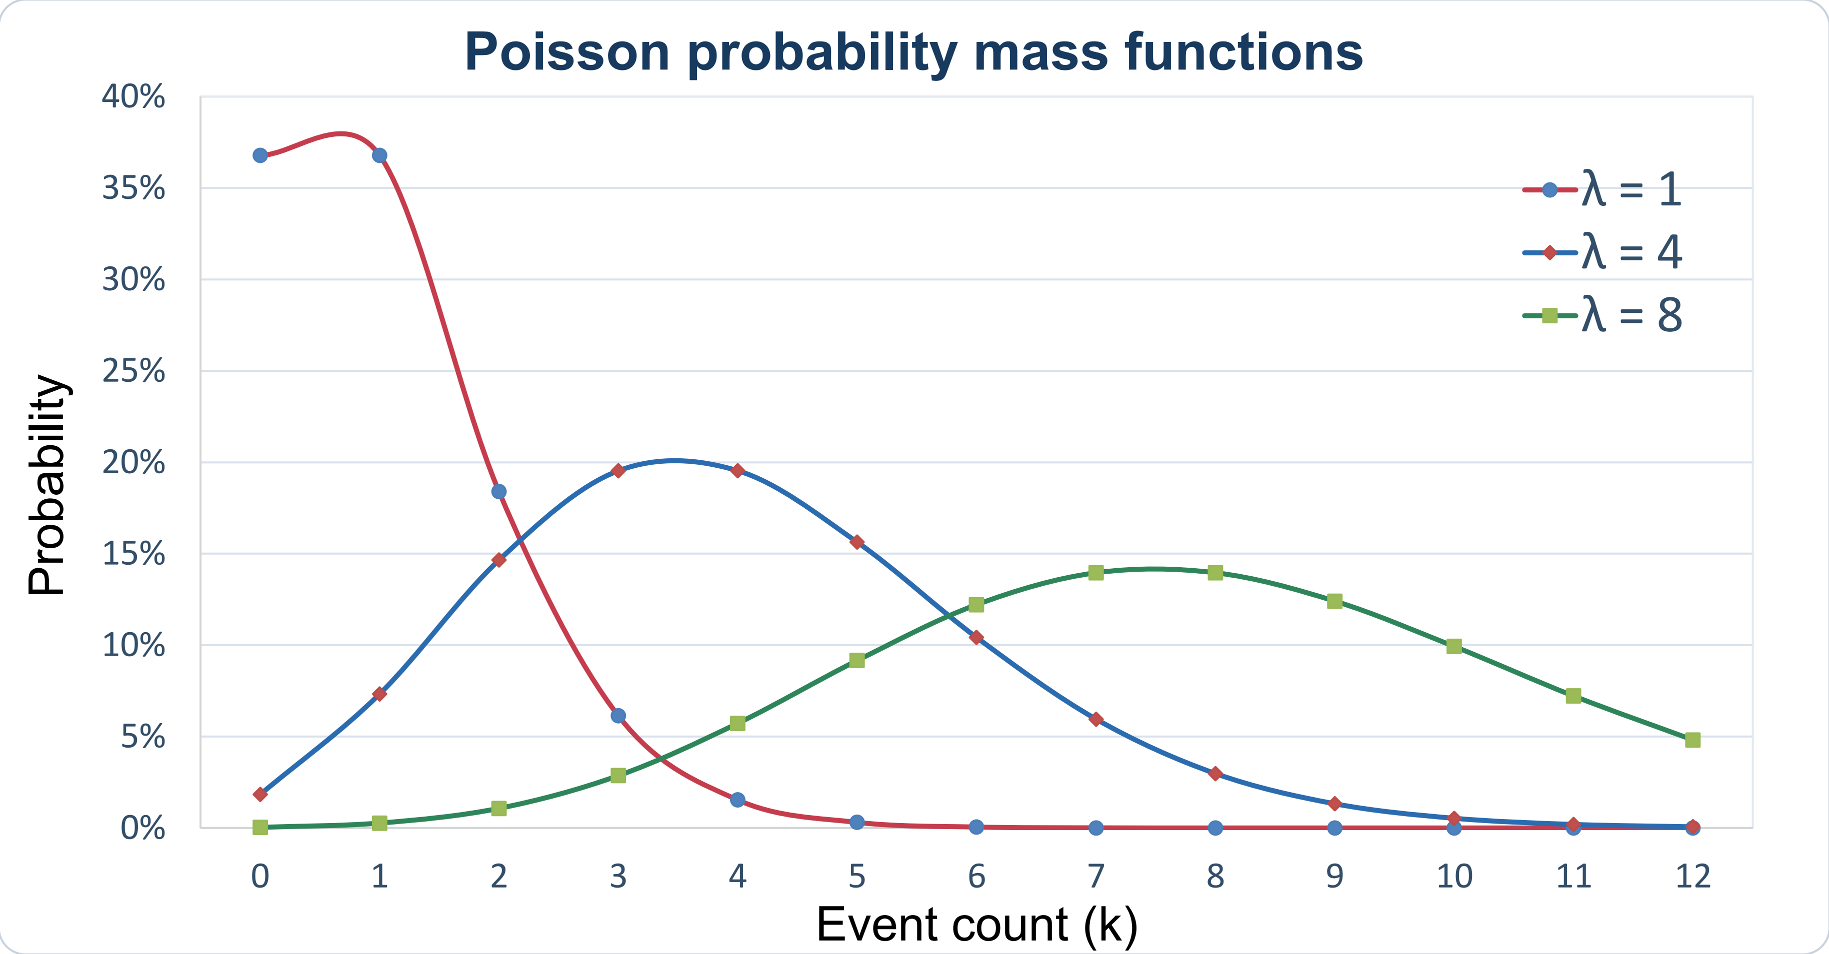

Probabilities occur only at integer counts; small λ is more right-skewed.

Mean, variance and mode

- For a Poisson random variable, $E(X) = \lambda$ and $Var(X) = \lambda$.
- If λ is not an integer, the mode is floor(λ).
- If λ is an integer, $\lambda - 1$ and λ are tied as the two modes.
- Observed counts remain integers even though λ need not be an integer.

Discrete Poisson counts

- Poisson uses a probability mass function: exact counts such as X = 3 have non-zero probability.
- The normal distribution uses a density function: probabilities are areas over intervals.
- Use Poisson when the outcome is a count, not when it is a continuous measurement.

Example: dispensing errors

- A hospital pharmacy records an average of 5 dispensing-error reports per 8-hour shift.
- For one comparable shift, what is the probability of observing exactly 3 reports?
- This is a synthetic teaching example—not hospital or patient data.

Is Poisson reasonable?

- The model is plausible if reports are independently counted, the process is stable during the shift, and all shifts have comparable exposure.
- A software change, staff shortage, common-cause incident or reporting campaign can create clustering and overdispersion; then Poisson may not fit.

Calculate λ for one shift

- Average rate = 5 reports per 8-hour shift.
- Exposure = 1 shift.
- λ = 5 reports/shift × 1 shift = 5 expected reports.

Rate parameter for the 8-hour interval

Calculate P(X = 3)

- Question: what is the probability of exactly 3 reports?
- Substitute: $P(X = 3) = e^{-5} \times 5^3 \div 3!$
- Result: $P(X = 3) = 0.1404 \approx 14.0\%$

Interpreting the 14% result

- Across many comparable shifts with $\lambda = 5$, about 14% would have exactly 3 reports.

- It does not mean every seventh shift must have 3 reports. Operational planning requires the complete probability distribution, not one exact count.

Full probability profile

With $\lambda = 5$, counts of 4 and 5 are tied as most likely, but lower and higher counts remain possible.

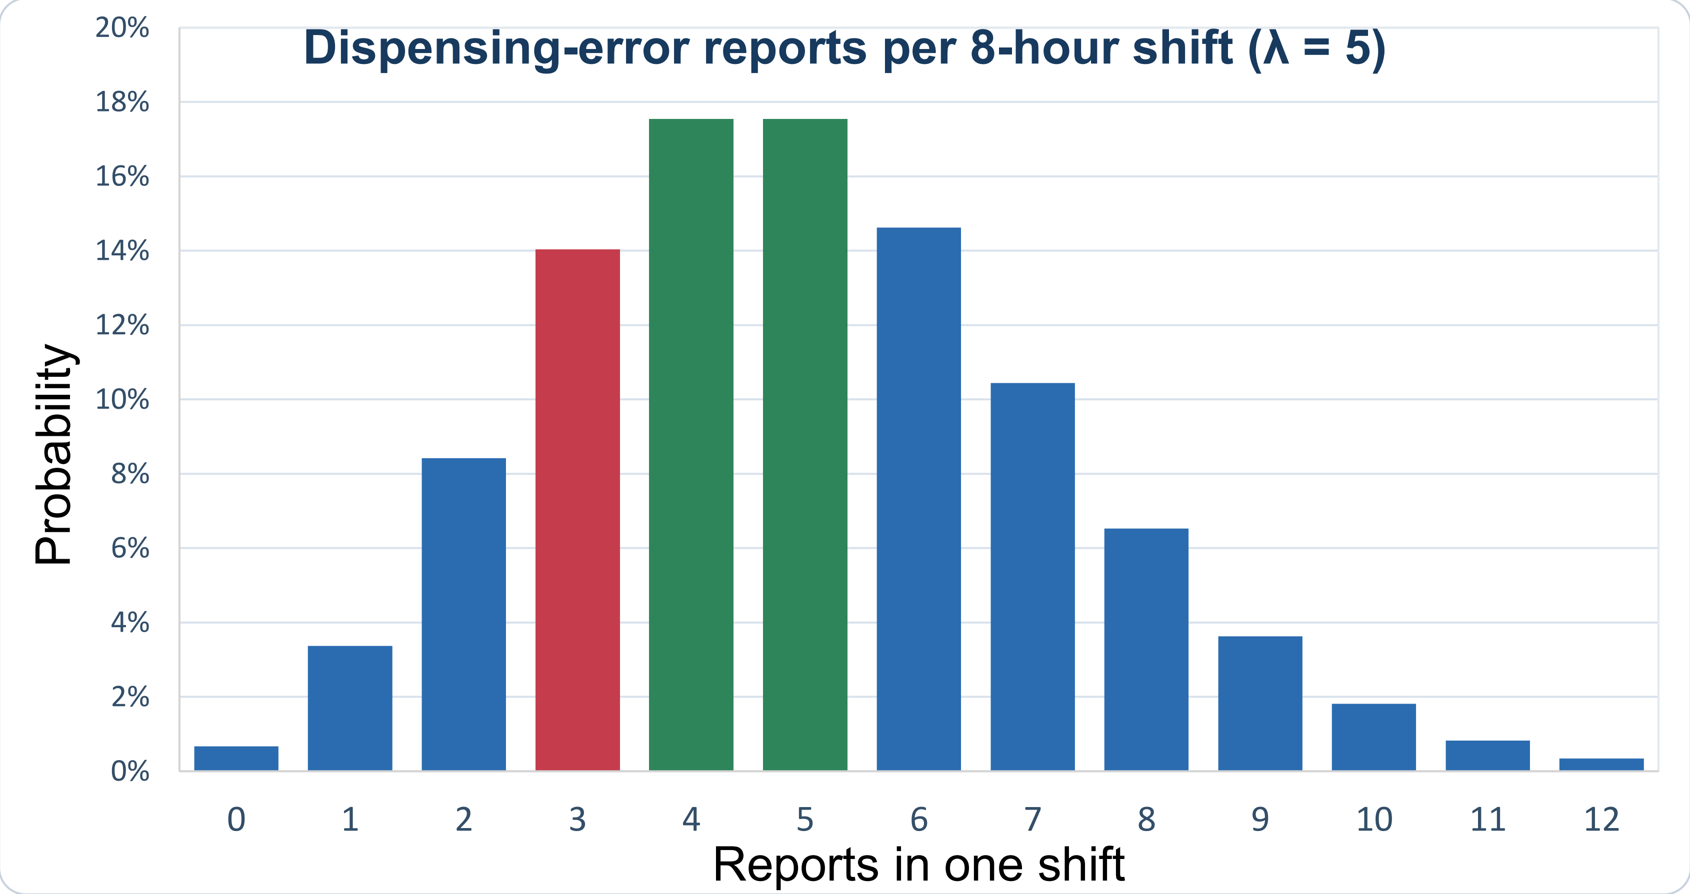

Exact and cumulative

- Exactly 3 reports: $P(X = 3) = 14.0\%$
- At most 3 reports: $P(X \leq 3) = 26.5\%$, so $P(X > 3) = 73.5\%$
- More than 5 reports: $P(X > 5) = 38.4\%$, so $P(X \leq 5) = 61.6\%$

Worked-example conclusions

- λ is an average–not a guaranteed count.
- Range probabilities are often more useful than a single most likely value.
- A sustained excess or clustering should trigger process investigation; probability does not define acceptable safety performance.

Simulation matches the PMF

A Python simulation of 10,000 shifts with λ = 5 produced mean 5.01 and variance 5.15; the histogram closely follows the Poisson pattern.

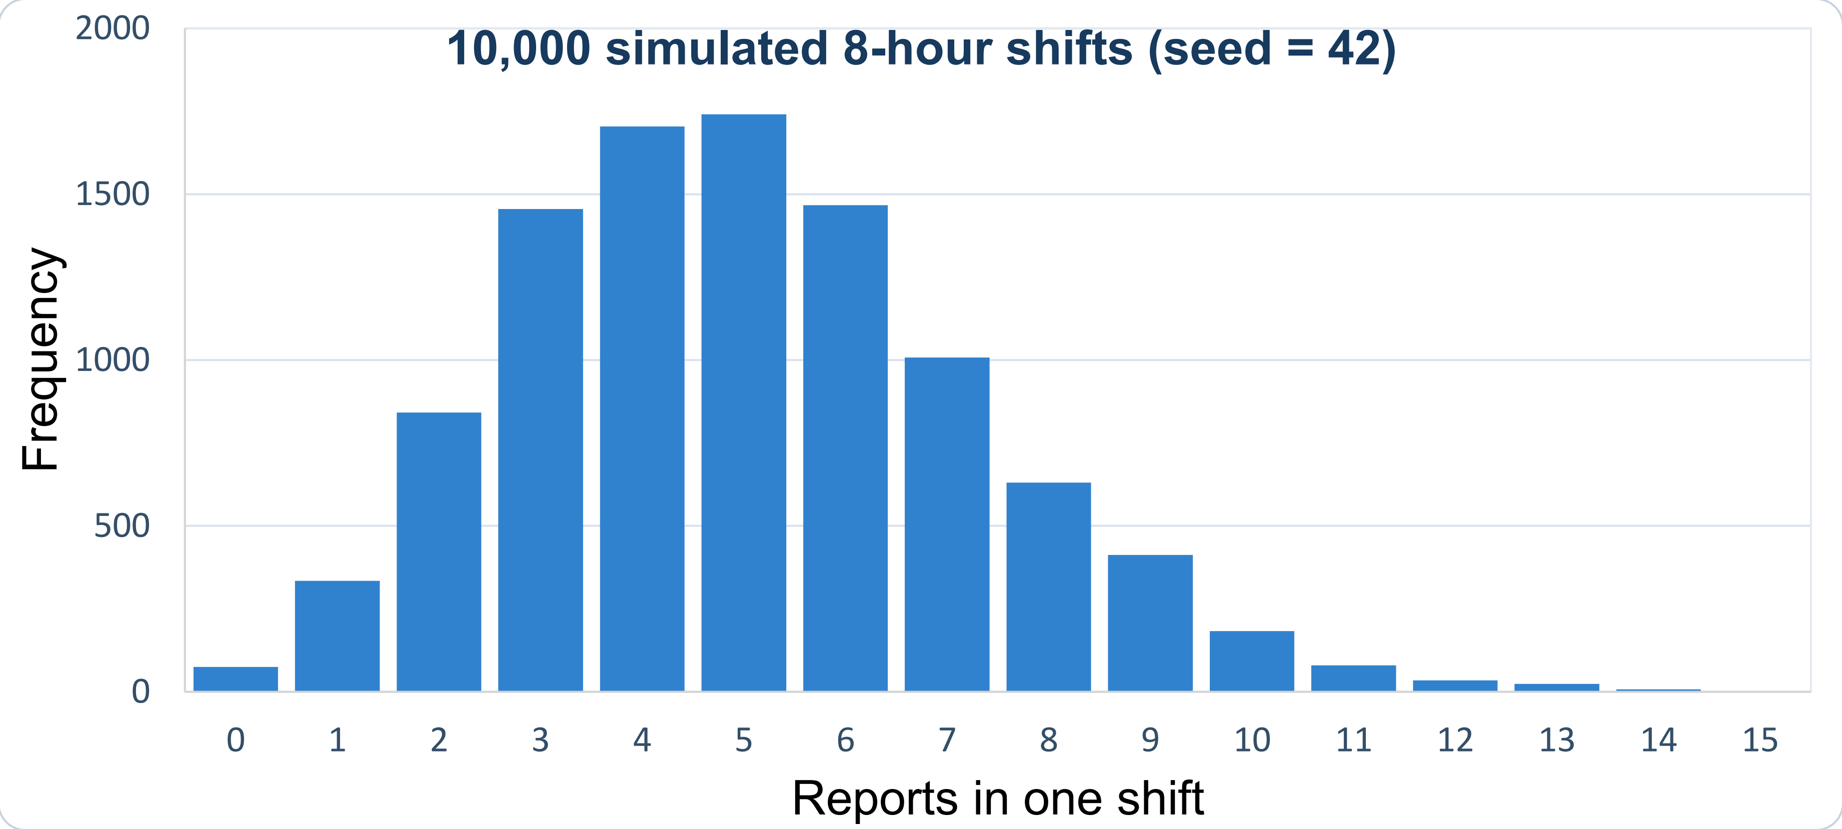

Poisson vs normal

Data: Poisson–whole-number counts | Normal–continuous measurements

Parameters: λ | μ and σ

Support: 0, 1, 2, … | all real values

Shape: often right-skewed at small λ | symmetric bell curve

Mean-variance: both equal λ | μ and σ^2 are separate

Pharmacy examples: errors/shift, colonies/plate | tablet weight, BP, concentration

Normal approximation

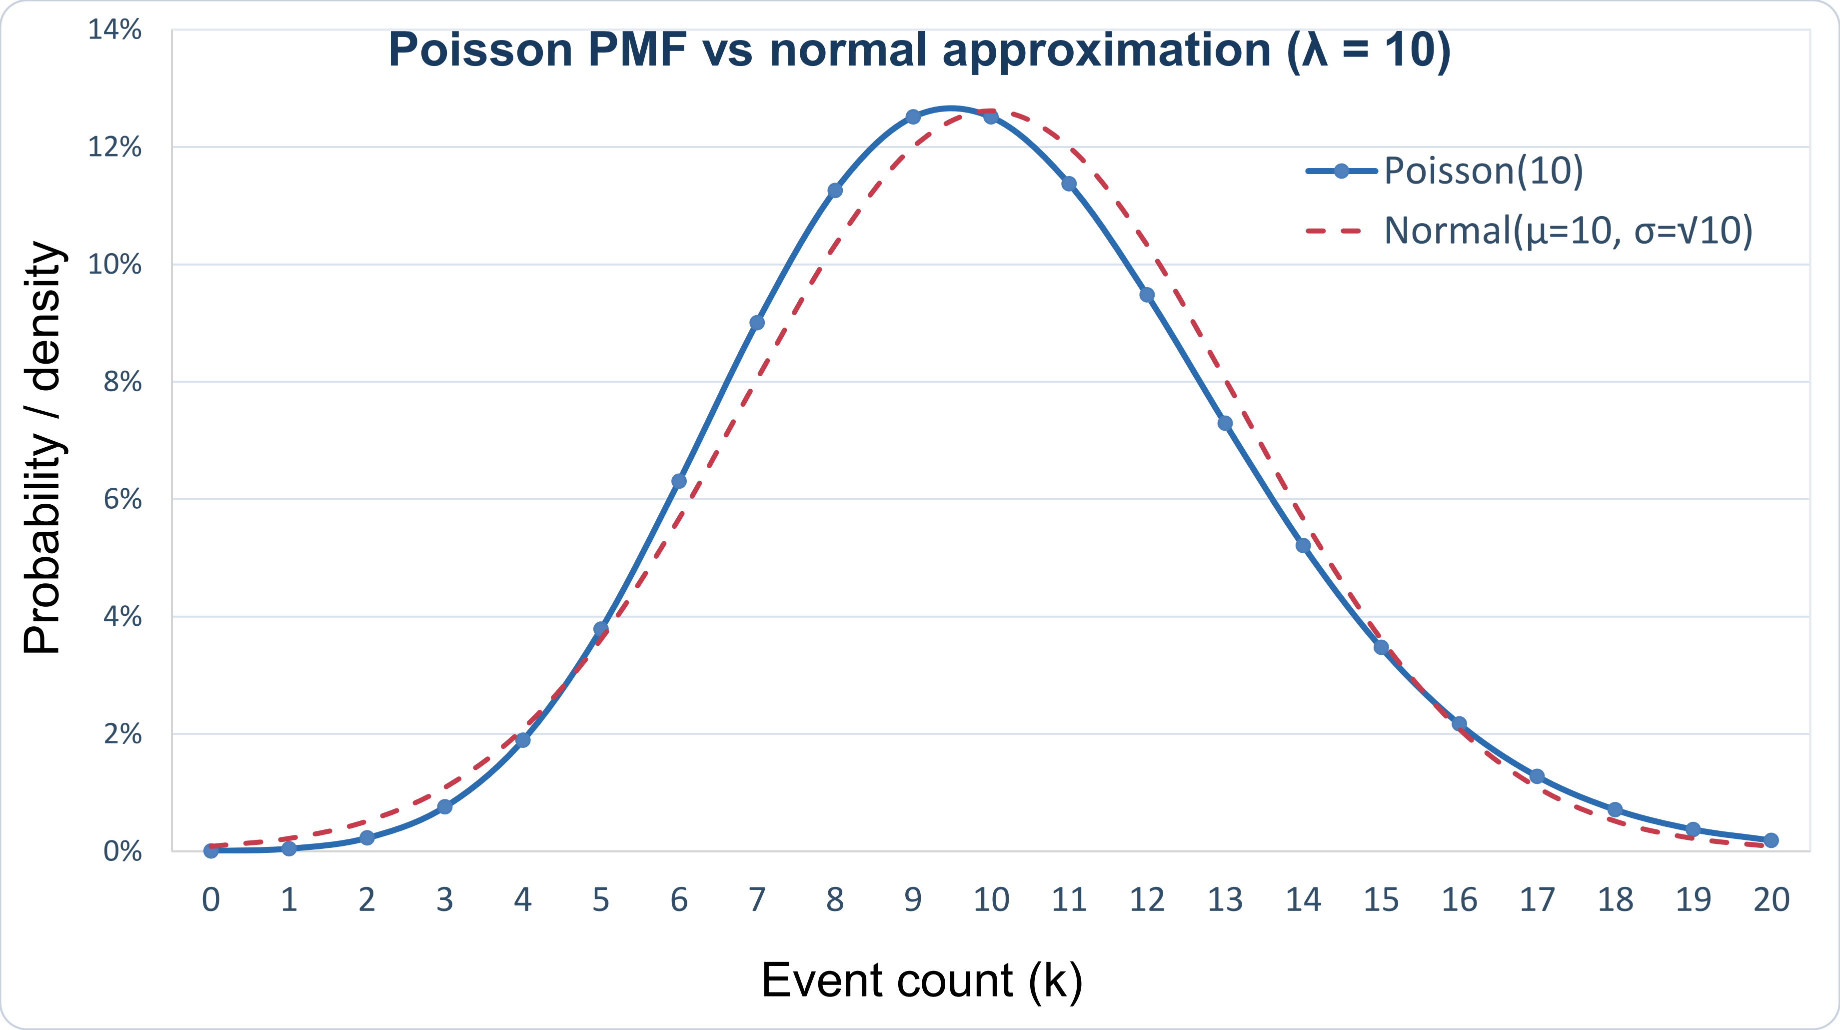

Python 1: arrange data

One row = one 8-hour shift; keep date, exposure and event count explicit.

```
df = pd.read_csv("dispensing_errors.csv")
df["date"] = pd.to_datetime(df["date"])
df = df.sort_values("date")
daily = df.groupby("date", as_index=False)["errors"].sum()
```

Python 2: analyse counts


```
from scipy.stats import poisson

lam = df["errors"].mean()          # 4.80
variance = df["errors"].var()      # 5.22
p_exact3 = poisson.pmf(3, mu=lam)
p_at_most3 = poisson.cdf(3, mu=lam)
p_more5 = poisson.sf(5, mu=lam)

# Mean and variance should be similar under Poisson.
```

Python 3: visualize fit


```
k = np.arange(0, 11); obs = df["errors"].value_counts().reindex(k,
fill_value=0)
exp = len(df) * poisson.pmf(k, mu=lam)
plt.bar(k - .2, obs, .4); plt.bar(k + .2, exp, .4)
```

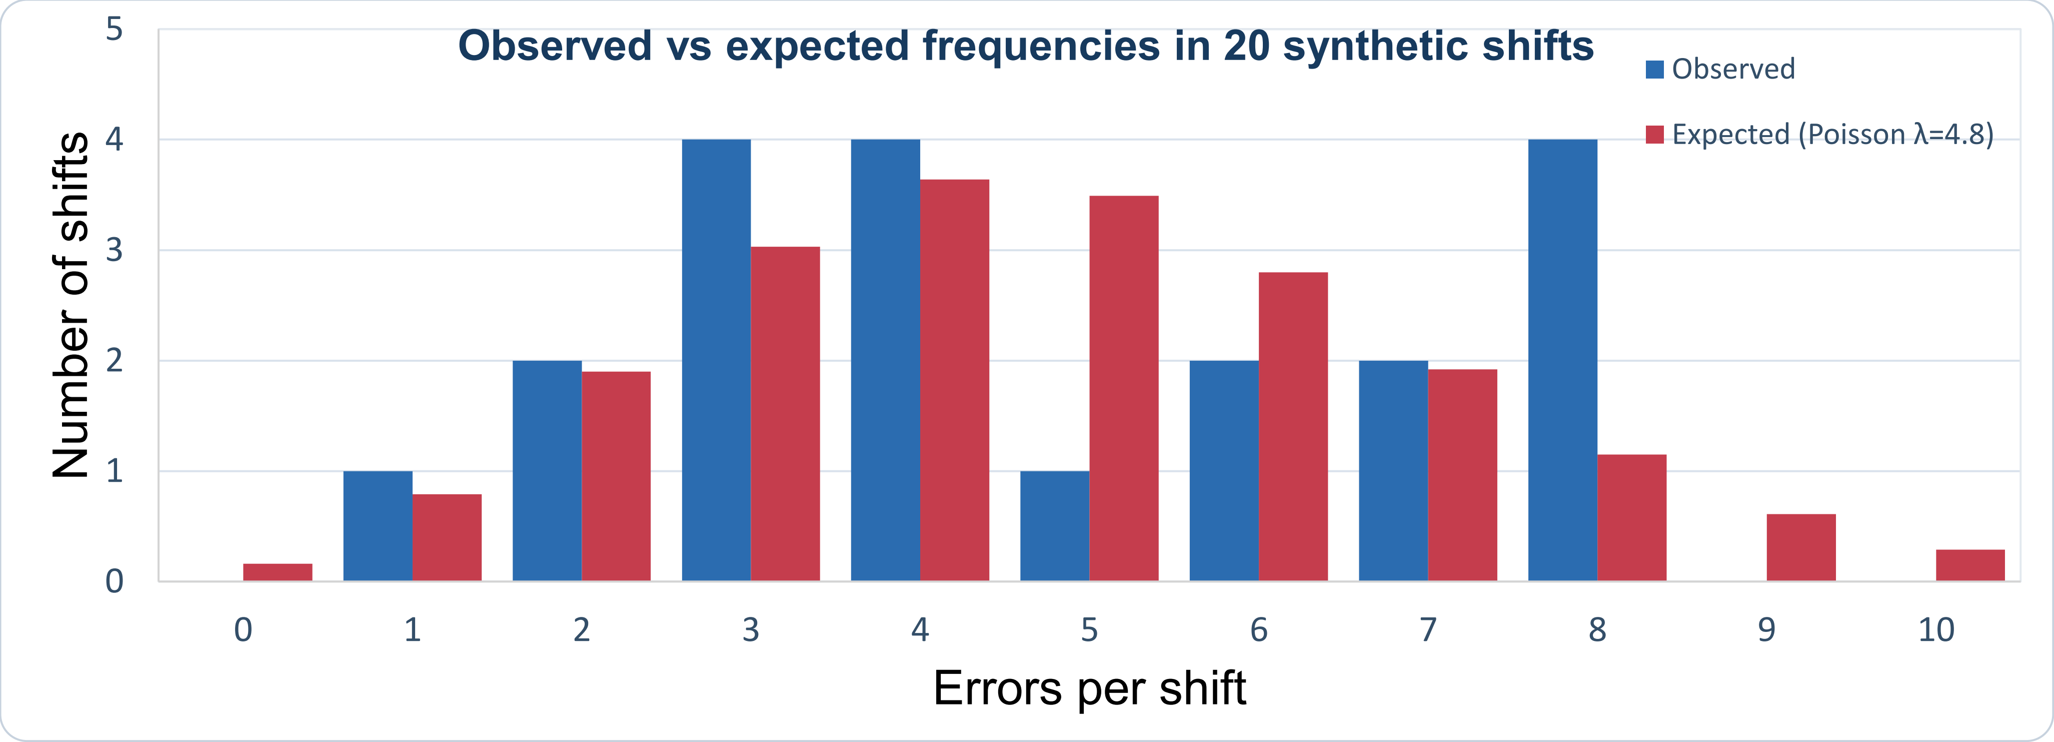

From counts to decisions

THANK YOU!