Lesson Objectives

1. Frame. A random intercept on the log-odds scale: `glmer` is `lmer` plus a `family`, the way `glm` is `lm` plus a `family`.
2. Apply. Fit a model with fixed and random effects and read each piece of the output.
3. Mean. Small groups shrink toward the average pass rate; a perfect record on a few tries is not evidence of the best group.
4. Communicate. Rank groups, and report fixed effects as odds ratios, without being fooled by small samples.

Bottom Line Up Front

1. **Same idea as Lesson 23, new scale.** A random intercept for a group shifts the log-odds instead of the mean.
2. **Fixed effects break on small perfect groups.** Everyone passed, so `glm` chases a log-odds of $+\infty$.
3. **Random effects borrow strength.** A group's own rate gets weight $w_j = n_j/(n_j + \lambda)$; the rest comes from the average.

Key Definitions

Fixed effect: a factor whose levels are the only ones you care about, and whose effect you want to estimate. Distance and firing position.

Random effect: a factor whose levels are a sample from many. You care how much they differ, and you want to predict a level not in the data. Companies, shooters.

Mixed model: fixed and random effects in one model.

Three Models for a 0/1 Response

Cadet (or round) i in group j; p_{ij} is the chance of a 1.

Fixed $\log \dfrac{p_{ij}}{1-p_{ij}} = \beta_0 + \beta_j$

```
glm(pass ~ company, family = binomial)
```

Random $\log \dfrac{p_{ij}}{1-p_{ij}} = \mu + u_j, \qquad u_j \sim N(0, \sigma_u^2)$

```
glmer(pass ~ 1 + (1 | company), family = binomial)
```

Mixed $\log \dfrac{p_{ij}}{1-p_{ij}} = \beta_0 + \beta_1 x_{ij} + \cdots + u_j, \qquad u_j \sim N(0, \sigma_u^2)$

```
glmer(hit ~ x + (1 | shooter), family = binomial)
```

Probability $\hat{p}_{ij} = \dfrac{e^{\text{log-odds}}}{1 + e^{\text{log-odds}}}$ **Latent ICC** $= \dfrac{\hat{\sigma}_u^2}{\hat{\sigma}_u^2 + \pi^2/3}$

Which Company Is Best at the IOCT?

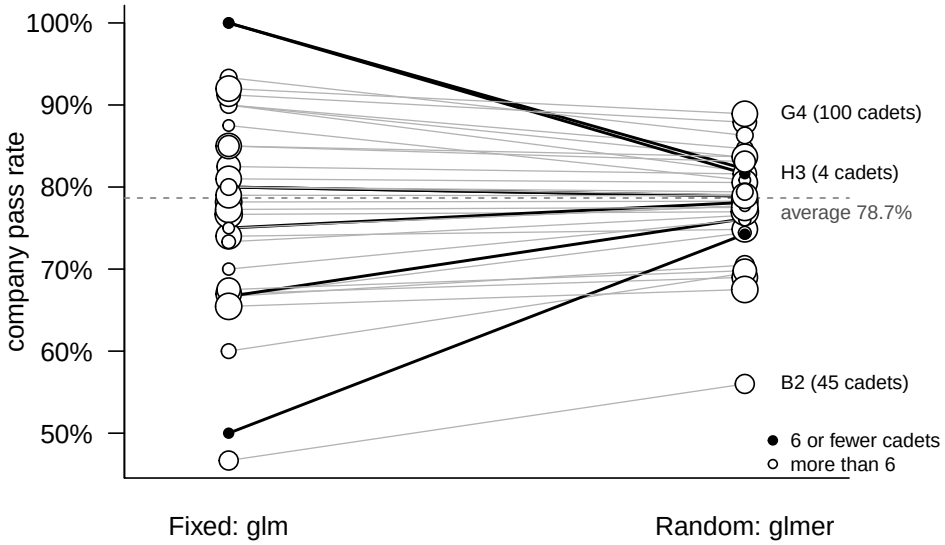
Indoor Obstacle Course Test, pass or fail: 1,684 cadets in 36 companies; the average company passes 78.7%. G4 passed 92 of 100. H3 passed 4 of 4, A2 passed 5 of 5.

```
fitF <- glm(pass ~ company, family = binomial)

            Estimate Std. Error z value
(Intercept)    1.099      0.816   1.346
companyA2     13.467    394.776   0.034
companyG4      1.344      0.896   1.500
companyH3     13.467    441.372   0.031
```

```
fitR <- glmer(pass ~ 1 + (1 | company), ...)

Random effects:
 Groups  Name        Variance Std.Dev.
 company (Intercept) 0.2532   0.5032
Fixed effects:
            Estimate Std. Error z value
(Intercept)   1.3048     0.1183   11.03
```



1. H3's coefficient is 13.47 with a standard error of 441. What is `glm` trying to do with a company where everyone passed, and why does the standard error blow up?

2. The `glmer` summary has a `company` row but no `Residual` row. Why not? What plays the role of σ^2?

3. `glm` ranks H3 and A2 first at 100%. `glmer` gives H3 81.6%, A2 82.2%, G4 88.9%. Use the figure: which company do you tell the commander is best, and why did the order flip?

How Much Does Each Company Shrink?

`glmer` has no closed form for $\hat{u}_j$ (the logistic curve puts u_j on both sides). Lesson 23's formula, on pass rates, gets within 1.7 points of `glmer` for every company.

The Blend

$$\hat{p}_j \approx \bar{p} + w_j \left(\frac{k_j}{n_j} - \bar{p} \right), \qquad w_j = \frac{n_j}{n_j + \lambda}, \qquad \lambda = \frac{\text{noise}}{\text{signal}} = \frac{\bar{p}(1-\bar{p})}{\hat{\sigma}_u^2 \left[\bar{p}(1-\bar{p})\right]^2}$$

k_j of n_j cadets passed. Here $\bar{p} = 0.787$, $\hat{\sigma}_u^2 = 0.253$, so $\lambda = 23.5$.

4. Rearranged, the blend is $\hat{p}_j \approx \dfrac{k_j + \lambda \bar{p}}{n_j + \lambda}$. Say what $\lambda = 23.5$ means in terms of cadets.

5. Two new companies: one passes 9 of 10, the other 90 of 100. Same raw rate. Without computing, which gets the higher $\hat{p}$ from `glmer`? Then check with the blend.

6. Suppose companies differed twice as much: $\hat{\sigma}_u^2$ doubles. What happens to λ, to every w_j, and to the amount of shrinkage? Why does that make sense?

7. B2 passed 21 of 45 (46.7%). The blend says 57.7%; `glmer` says 56.0%, the biggest miss of all 36 companies. Why is the blend worst for B2?

Rifle Qualification

48 shooters, one row per round: `shooter`, `position` (prone supported, prone unsupported, kneeling), `distance` (50 to 300 m), `hit` (1 or 0). Most fired 40 rounds; lightning closed the range before 6 finished, and they fired 10. The OIC asks: how much do distance and position matter, and who are the best shooters?

8. Which of `distance`, `position`, `shooter` should be fixed and which random? Why?

`dist100` is distance in hundreds of meters.

```
fit_rifle <- glmer(hit ~ dist100 + position + (1 | shooter), family = binomial)

Random effects:
 Groups  Name        Variance Std.Dev.
 shooter (Intercept) 0.4503   0.671
Fixed effects:
                           Estimate Std. Error z value Pr(>|z|)
(Intercept)                 2.87243    0.19717  14.568  < 2e-16
dist100                    -0.91371    0.07187 -12.714  < 2e-16
positionprone unsupported -0.44369    0.13906  -3.191  0.00142
positionkneeling          -0.95333    0.13769  -6.924  4.4e-12
```

9. $e^{-0.914} = 0.40$ and $e^{-0.953} = 0.39$. Tell the OIC what each means. Whose odds are being compared?

10. The latent ICC is $0.450/(0.450 + 3.29) = 0.12$. What does it say about shooters?

`ranef(fit_rifle)` top five, with each shooter's record

shooter	u_hat	rounds	hits	avg distance (m)
S22	1.182	40	37	185.0
S15	0.965	40	35	185.0
S13	0.837	40	35	172.5
S17	0.763	10	10	170.0
S29	0.703	40	34	171.3

11. S17 went 10 for 10, a perfect record. Why does S22, 37 of 40, rank above S17?

12. S13 and S15 both hit 35 of 40. Why does S15 rank higher?

13. Fit without the shooter, `glm` gives -0.821 per 100 m; `glmer` gives -0.914. Why is `glm`'s effect smaller, and which one answers "how much harder is 300 m for my shooter"?

14. In two sentences, answer the OIC.

Key Takeaway

A random effect on a 0/1 response works the way it did in Lesson 23: each group gets weight $n_j/(n_j+\lambda)$ on its own record and the rest on the average. Rank groups by the model, not by raw rates on a handful of tries.