Repeated measurement analysis of binary responses[1]

Because language is discrete

[1] This slide show is an open-source document. See last slide for copyright information.

Modern tools

In R, the `glmer` function in the `lme4` package

- An extension of logistic regression (`glm` in R).
- $y = 0$ or 1.
- Random shock for subject again.
- So it's a mixed model.
- Technically a non-linear mixed model.

Odds

An indirect representation of probability

$$\text{Odds} = \frac{p}{1-p}$$

- If $P(Y = 1) = 1/2$, Odds $= .5/(1 - .5) = 1$ (to 1)
- If $P(Y = 1) = 2/3$, Odds $= 2$ (to 1)
- If $P(Y = 1) = 3/5$, Odds $= (3/5)/(2/5) = 1.5$ (to 1)
- If $P(Y = 1) = 1/5$, Odds $= .25$ (to 1)

The higher the probability, the greater the odds

$$\text{Odds} = \frac{p}{1-p}$$

$$0 \leq \text{Odds} < \infty$$

A linear model for the log odds

- Under the surface, a multiple regression equation
- Not for the probability of Yes
- And not for the odds of Yes
- But for the natural logarithm (base e, not 10) of the odds.
- $-\infty < \ln(\text{Odds}) < \infty$

The greater the probability, the greater the log odds

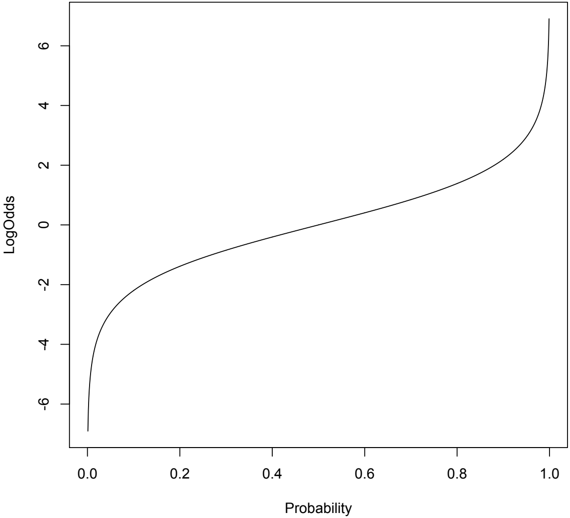

Log odds range from $-\infty$ to ∞

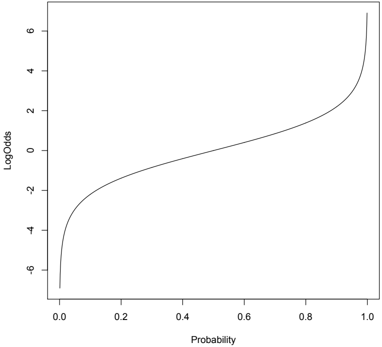

- Probabilities of zero and one are ruled out.
- Estimated probabilities of zero and one are ruled out too.

The odds ratio

- Conclusions tend to emerge in terms of odds.
- Instead of a treatment adding or subtracting something from the control group mean
- It might multiply the odds by something.
- If the something is greater than one, it's an increase in probability.
- Like the odds of a recursive structure are 1.4 times as great in the experimental condition.
- If the something is less than one, it's a decrease in probability.
- This is called the *odds ratio*.

Main effects: What are we testing?

- Tests for main effects are tests for differences of marginal mean log odds.

	Treatment			
	A	*B*	*C*	
Adult	1.0	1.7	0.9	1.2
Child	1.5	1.8	2.4	1.9
	1.25	1.75	1.65	1.55

- These correspond to ratios of marginal *geometric mean* odds.
  - Geometric mean $= (7 \cdot 2 \cdot 5 \cdot 6)^{1/4} = 4.53$ (Not 5)
  - Geometric mean $\leq$ Arithmetic mean.
- Simple statements about geometric means of odds translate into complicated statements about probabilities.
- Still, they often capture the reality that something is more likely under certain conditions.

Is there an interaction?

	Group		
Treatment	1	2	3
Experimental	0.5	0.7	0.9
Control	0.2	0.4	0.6

Interaction means different odds ratios

Not different differences between probabilities

		Group		
		1	2	3
Exper	Probability	0.50	0.70	0.90
	Odds	1.00	2.33	9.00
Control	Probability	0.20	0.40	0.60
	Odds	0.25	0.67	1.50
	Odds Ratio	4.00	3.50	6.00

Repeated measures

- Suppose subjects appear in more than one treatment condition.
- There is a baseline log odds of Yes for each treatment combination.
- A random shock for each subject is added to the baseline log odds.
- According to the usual model, the random shock is normally distributed.
- When a subject appears in more than one condition, her log odds of Yes are pushed up or down by the same amount.
- This produces correlated binary responses from the same subject.

A two-stage probability process

Implied by the random shocks model

- Subjects in a treatment combination are a random sample from some hypothetical population.
- The random sampling of subjects is Stage One.
- Within each treatment combination, log odds of Yes are normally distributed.
  - Population mean is determined by the treatment combination.
  - Population variance is the same for all treatment combinations.
- Log odds values from the same subject are correlated.
- For each observation, log odds of Yes are converted to a probability.
- Then a Yes or No is randomly produced.
- The "coin toss" to produce Yes or No is Stage Two.
- All we can observe are the Yes and No values.

Theoretical distribution of the probability of Yes

Based on a normal distribution of the log odds with standard deviation = 1

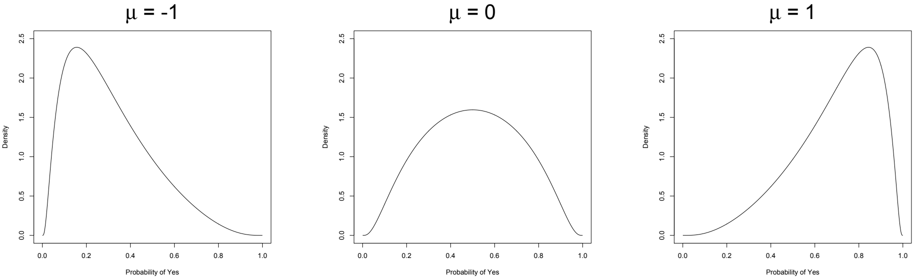

Main points

- Model implies a different log odds of Yes for each subject (in each condition).
- Log odds are normally distributed.
- Population mean log odds depends on the treatment combination.
- Tests of main effects and interactions have the usual meaning *in terms of population mean log odds*.
- You should continue to think in terms of probabilities (proportion of Yes responses in each condition).
- You can still believe pairwise differences between means, because mean log odds are different if and only if mean probability is different.
- You might want to look at estimated mean log odds to see what the tests of main effects and interactions are reflecting.

State of the art

Contemporary, not just modern

- The theory behind the tests is straightforward.
- But it's asymptotic.
- Meaning it's justified as the number of data points $\to \infty$.
- Computation is a bit bleeding edge.
- Methods for finding parameter estimates are iterative.
- Convergence problems are common.
- Software from different vendors does not always produce the same numbers.
- It's easy to produce examples with simulated data that fail the test of large-sample accuracy (yes, in R).
- Let's do it anyway.

The `glmer` function in the `lme4` package

- Syntax is like `lmer` for linear models.
- And like `glm` for generalized linear models with fixed effects.
- We are going to keep it simple.
- Just add `+(1|Subject)` for the random shock (intercept).
- And maybe sometimes `+(1|Item)`.
- Use effect coding (`contr.sum`) if there are interactions between factors.
- `Anova(model,type='III')` from the `car` package to test each effect controlling for all others.

Copyright Information