Agenda

Multinomial regression

1. Categorical outcomes
2. Softmax link function
3. Interpreting coefficients
4. ***Hands on:***
   Multinomial logistic in R

Categorical outcomes

Are rich people more likely to be married?

- **Predictor:** Annual earnings
  $E_i \in [0, \infty)$
- **Outcome:** Marital status
  $M_i \in \{\text{Single}, \text{Married}, \text{Divorced}, \text{Widowed}\}$

The problem Outcome variable has multiple (>2) categories. *Binomial and Poisson models won't work.*

The solution Use a categorical / multinomial outcome distribution (and a new link function) to account for the data.

CATEGORICAL DISTRIBUTION

The ***categorical distribution*** is analagous to a Bernoulli distribution with *k* > 2 outcomes:

$$\text{Prob}(X = i) = p_i$$

$$\boldsymbol{p} = [p_1, p_2, \ldots, p_k]$$

$$\sum_{j=1}^{k} p_j = 1$$

$$\boldsymbol{p} = [0.20, 0.10, 0.45, 0.25]$$

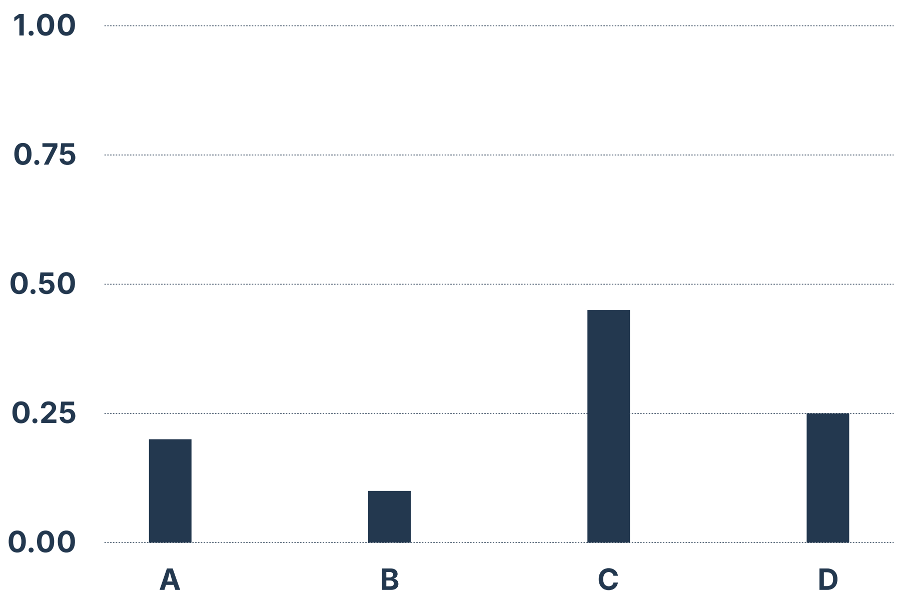

MULTINOMIAL DISTRIBUTION

The Binomial, Bernoulli, and categorical distributions are each special cases of the *multinomial distribution*.

Binomial distribution | $\text{Bin}(n, p) = \text{Multinom}(n, (1-p, p))$

Bernoulli distribution | $\text{Bernoulli}(p) = \text{Multinom}(1, (1-p, p))$

Categorical distribution | $\text{Cat}(p_1, p_2, \dots p_k) = \text{Multinom}(1, (p_1, p_2, \dots p_k))$

	1 trial	>1 trial
2 categories	**Bernoulli**	**Binomial**
>2 categories	**Categorical**	**Multinomial**

CATEGORICAL OUTCOME

$$M_i \sim \text{Categorical}(\boldsymbol{p}_i)$$

$$\boldsymbol{p}_i = f^{-1}\left([\mu_{\text{s}i}, \mu_{\text{m}i}, \mu_{\text{d}i}, \mu_{\text{w}i}]\right)$$

One category must be the *reference* category

$$\begin{aligned} \mu_{\text{s}i} &= 0 \\ \mu_{\text{m}i} &= \alpha_m + \beta_m E_i \\ \mu_{\text{d}i} &= \alpha_d + \beta_d E_i \\ \mu_{\text{w}i} &= \alpha_w + \beta_w E_i \end{aligned}$$

Each other category gets its own intercept and coefficient

(Can be thought of as a series of logistic regressions against the reference category)

But we still need a ***link function***

The *softmax* link function

Softmax is a straightforward generalization of the the _inverse logit_

$$[p_s, p_m, p_d, p_w] = \text{softmax}(\mu_s, \mu_m, \mu_d, \mu_w)$$
$$= \left[\frac{\exp(\mu_s)}{Z}, \frac{\exp(\mu_m)}{Z}, \frac{\exp(\mu_d)}{Z}, \frac{\exp(\mu_w)}{Z}\right]$$

$$\textit{where} \quad Z = \exp(\mu_s) + \exp(\mu_m) + \exp(\mu_d) + \exp(\mu_w)$$

Untransformed (*μ*)

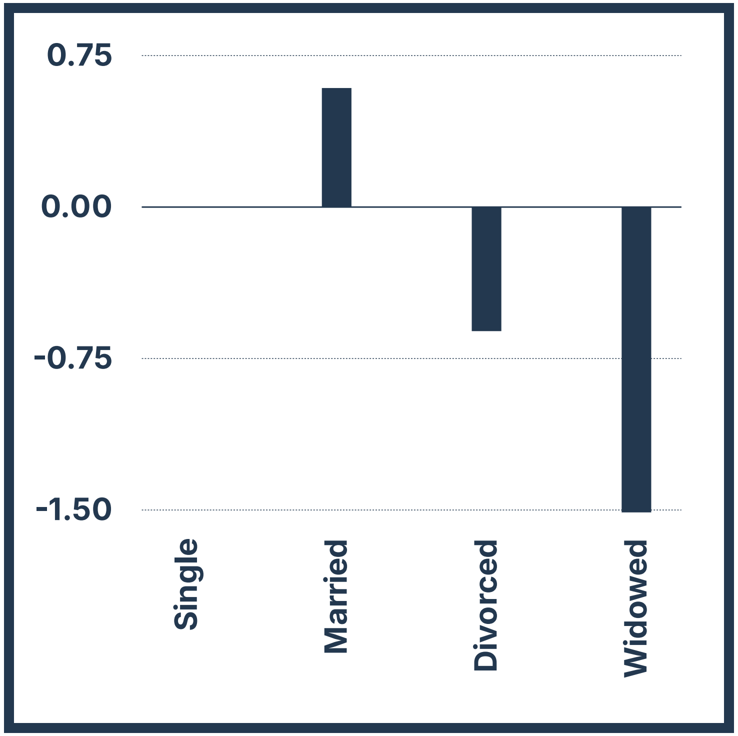

Transformed (*p*)

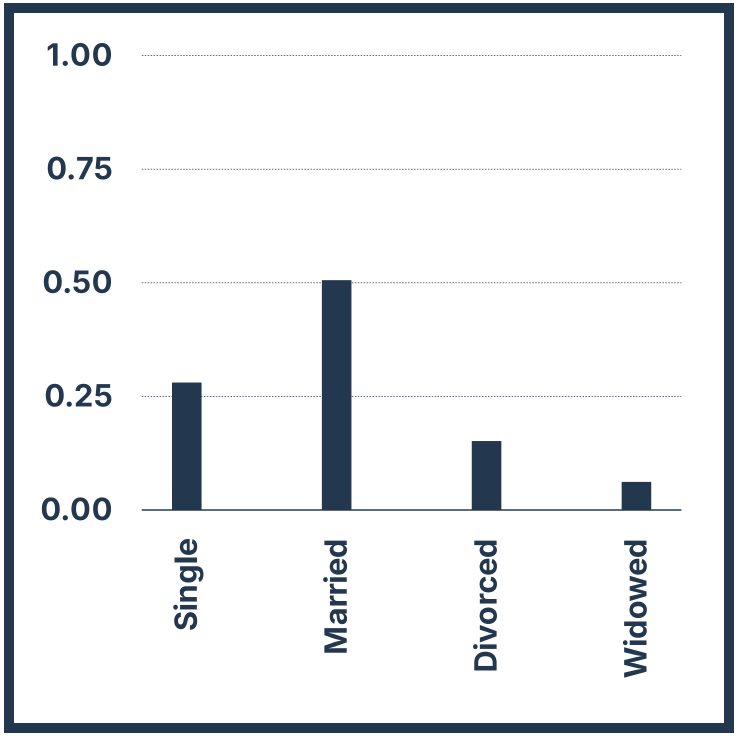

Putting it all together gives us the ***multinomial logistic regression model*** (a.k.a. ***categorical regression model***)

$$
\begin{aligned}
M_i &\sim \text{Categorical}(\boldsymbol{p}_i) \\
\boldsymbol{p}_i &= \text{softmax}\left([\mu_{\text{s}i}, \mu_{\text{m}i}, \mu_{\text{d}i}, \mu_{\text{w}i}]\right) \\
\mu_{\text{s}i} &= 0 \\
\mu_{\text{m}i} &= \alpha_m + \beta_m E_i \\
\mu_{\text{d}i} &= \alpha_d + \beta_d E_i \\
\mu_{\text{w}i} &= \alpha_w + \beta_w E_i \\
\alpha_m, \alpha_d, \alpha_w &\sim \text{Student}(3, 0, 2.5) \\
\beta_m, \beta_d, \beta_w &\sim \text{Norm}(0, 2)
\end{aligned}
$$

MULTINOMIAL MODEL

Note that with only two categories, the multinomial logistic regression reduces to the standard logistic regression:

$$M_i \sim \text{Categorical}(\boldsymbol{p}_i)$$

$$\boldsymbol{p}_i = \text{softmax}\left([\mu_{1i}, \mu_{2i}]\right)$$

$$\mu_{1i} = 0$$
$$\mu_{2i} = \alpha_2 + \beta_2 E_i$$

$$\begin{aligned} p_{1i} &= \frac{1}{1 + \exp(\mu_{2i})} \\ &= 1 - p_{2i} \end{aligned}$$

$$\begin{aligned} p_{2i} &= \frac{\exp(\mu_{2i})}{1 + \exp(\mu_{2i})} \\ &= \text{logit}^{-1}(\mu_{2i}) \end{aligned}$$

Interpreting coefficients

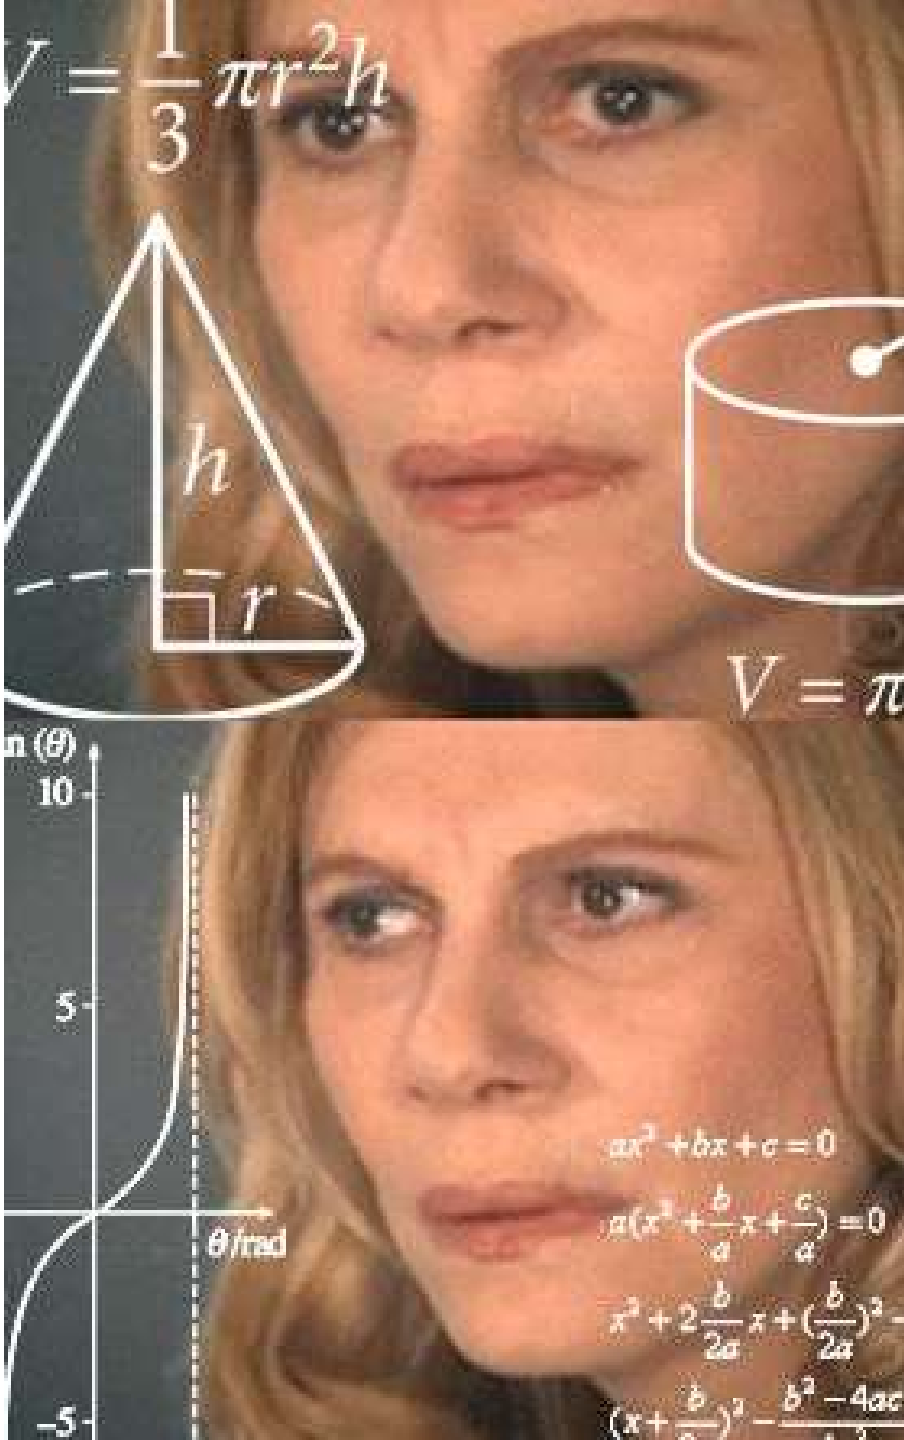

INTERPRETING COEFFICIENTS

$$M_i \sim \text{Categorical}(\boldsymbol{p}_i)$$

$$\boldsymbol{p}_i = \text{softmax}\left([\mu_{\text{s}i}, \mu_{\text{m}i}, \mu_{\text{d}i}, \mu_{\text{w}i}]\right)$$

$$\mu_{\text{s}i} = 0$$
$$\mu_{\text{m}i} = \alpha_m + \beta_m E_i$$
$$\mu_{\text{d}i} = \alpha_d + \beta_d E_i$$
$$\mu_{\text{w}i} = \alpha_w + \beta_w E_i$$

	Estimate	Q2.5	Q97.5
α_m	-5.40	-6.16	-4.67
β_m	0.59	0.52	0.67
α_d	-4.81	-5.80	-3.86
β_d	0.41	0.32	0.51
α_w	-3.43	-4.68	-2.23
β_w	0.19	0.07	0.32

Interpreting these cofficients on their own is complex — the results are analagous to a *series* of logistic regressions against the reference category, *conditional on* the other outcomes.

E.g. $\boldsymbol{\alpha_m}$ and $\boldsymbol{\beta_m}$ determine the probability of a person being ***married rather than single, assuming that they are neither divorced nor widowed***.

Assessing the sign of the α and β estimates is the only straightforward interpretation.

INTERPRETING COEFFICIENTS

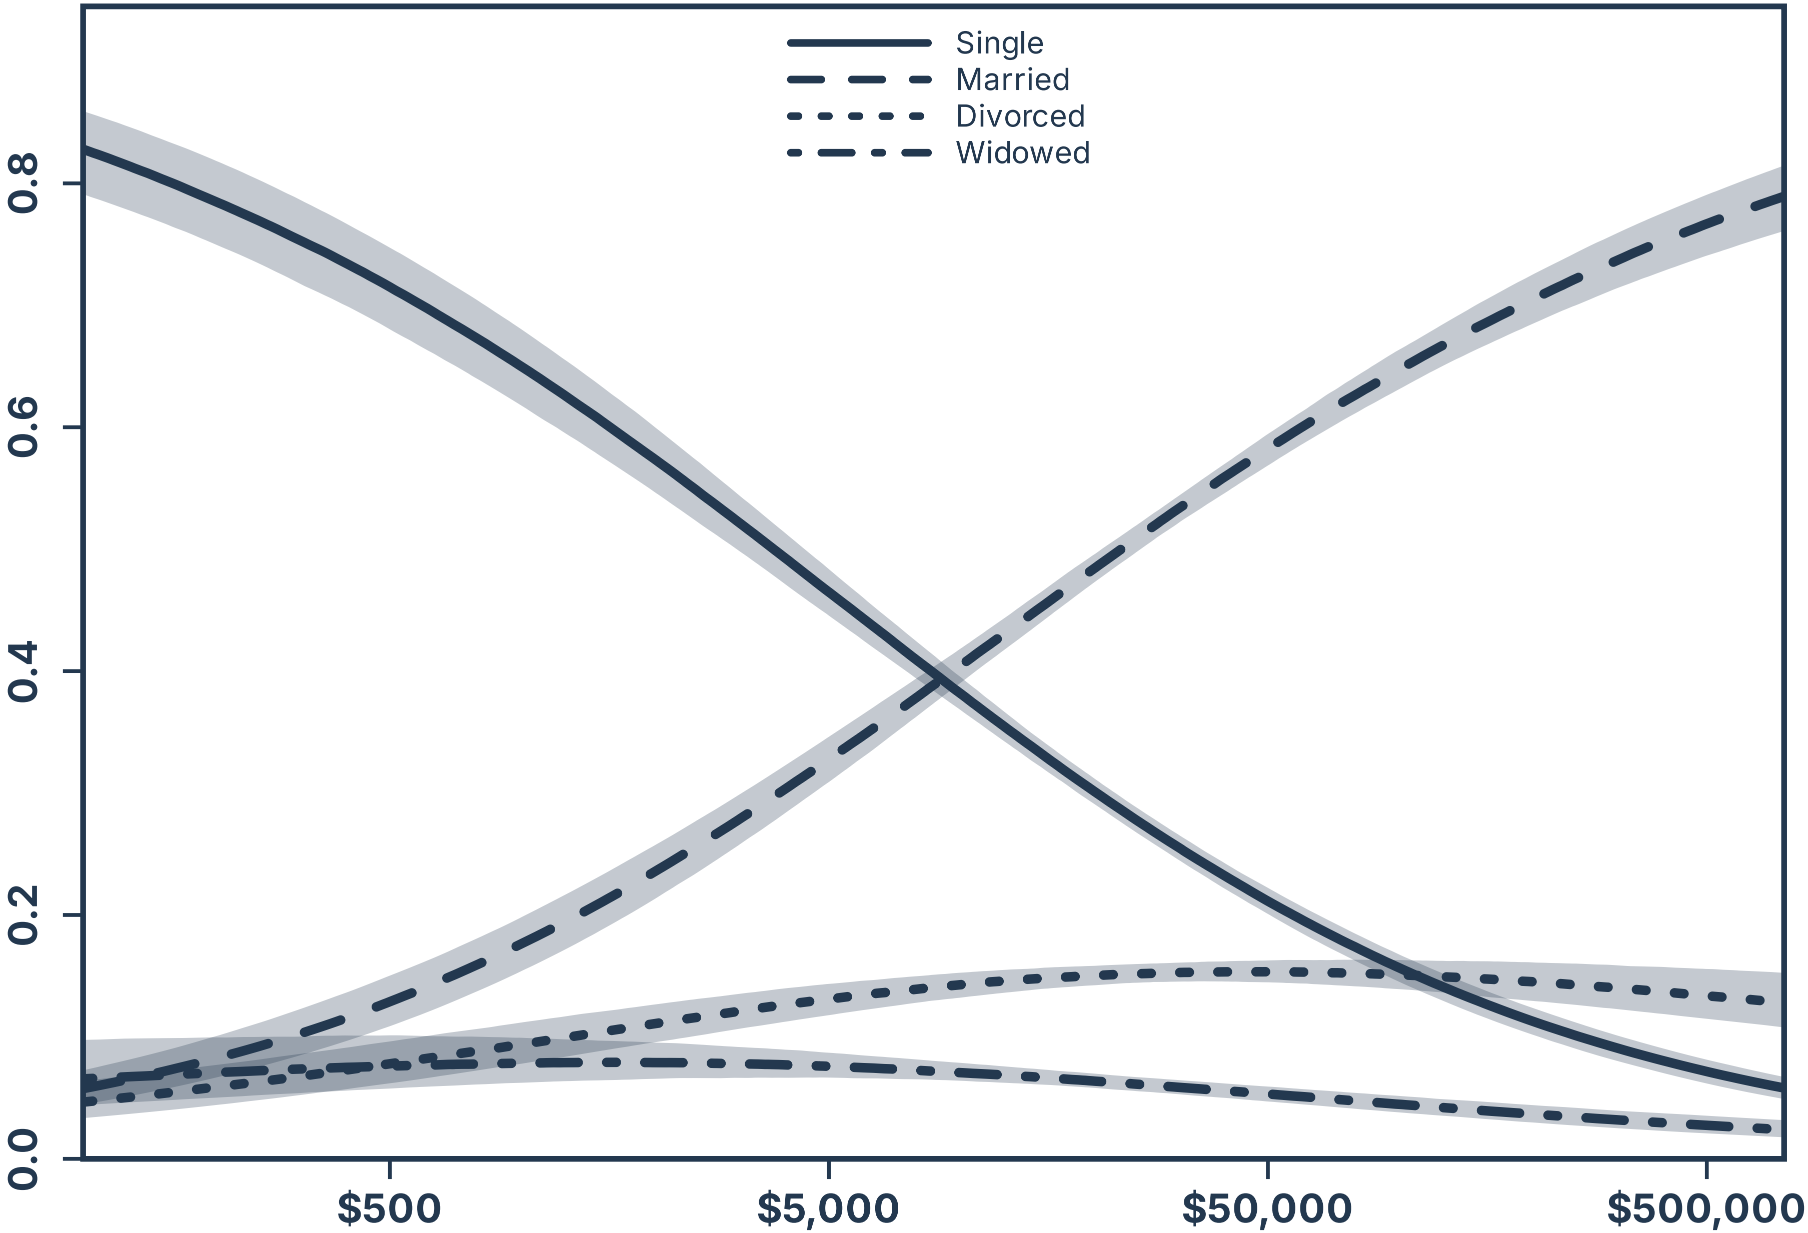

Image credit

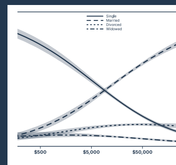

Figures by Peter McMahan (source code)

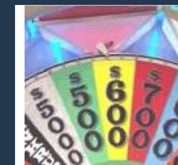

Clip from Wheel of Fortune (1983–)

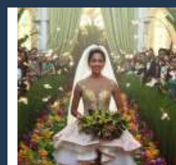

Still from Crazy Rich Asians (2018)

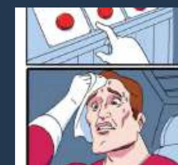

Adapted from image by Jake Clark

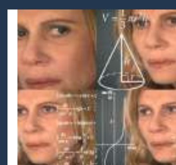

Stills adapted from Senhora do Destino (2004)