PSC 202
SYRACUSE UNIVERSITY

INTRODUCTION TO POLITICAL ANALYSIS

MORE BIVARIATE HYPOTHESIS TESTING, HYPOTHESIS TESTING WITH SAMPLES

EXAM

- Next week Wednesday: Exam #2
  - Can bring a calculator (no phone etc.)
  - Allowed to bring one single-page letter-size (8.5x11) sheet with you. Front side only. What you put on it is up to you, but it has to be your own.
- Monday: Review
  - Email questions etc. by Sunday evening
- If you take exams at CDR, please sign up now!

PROBLEM SETS

- Problem set 6 due on Friday
- Problem set 7 will be posted tomorrow
  - Due Friday next week, but good idea to complete it before the exam

TODAY

- Finishing up bivariate hypothesis testing
- Hypothesis testing with samples

BIVARIATE RELATIONSHIPS

Dependent Variable	Independent Variable: Nominal/Ordinal	Independent Variable: Interval
Nominal/Ordinal	Cross-Tabulation	Not In This Class…
Interval	Mean Comparison	Correlation Coefficient, Linear Regression

EXAMPLE 1

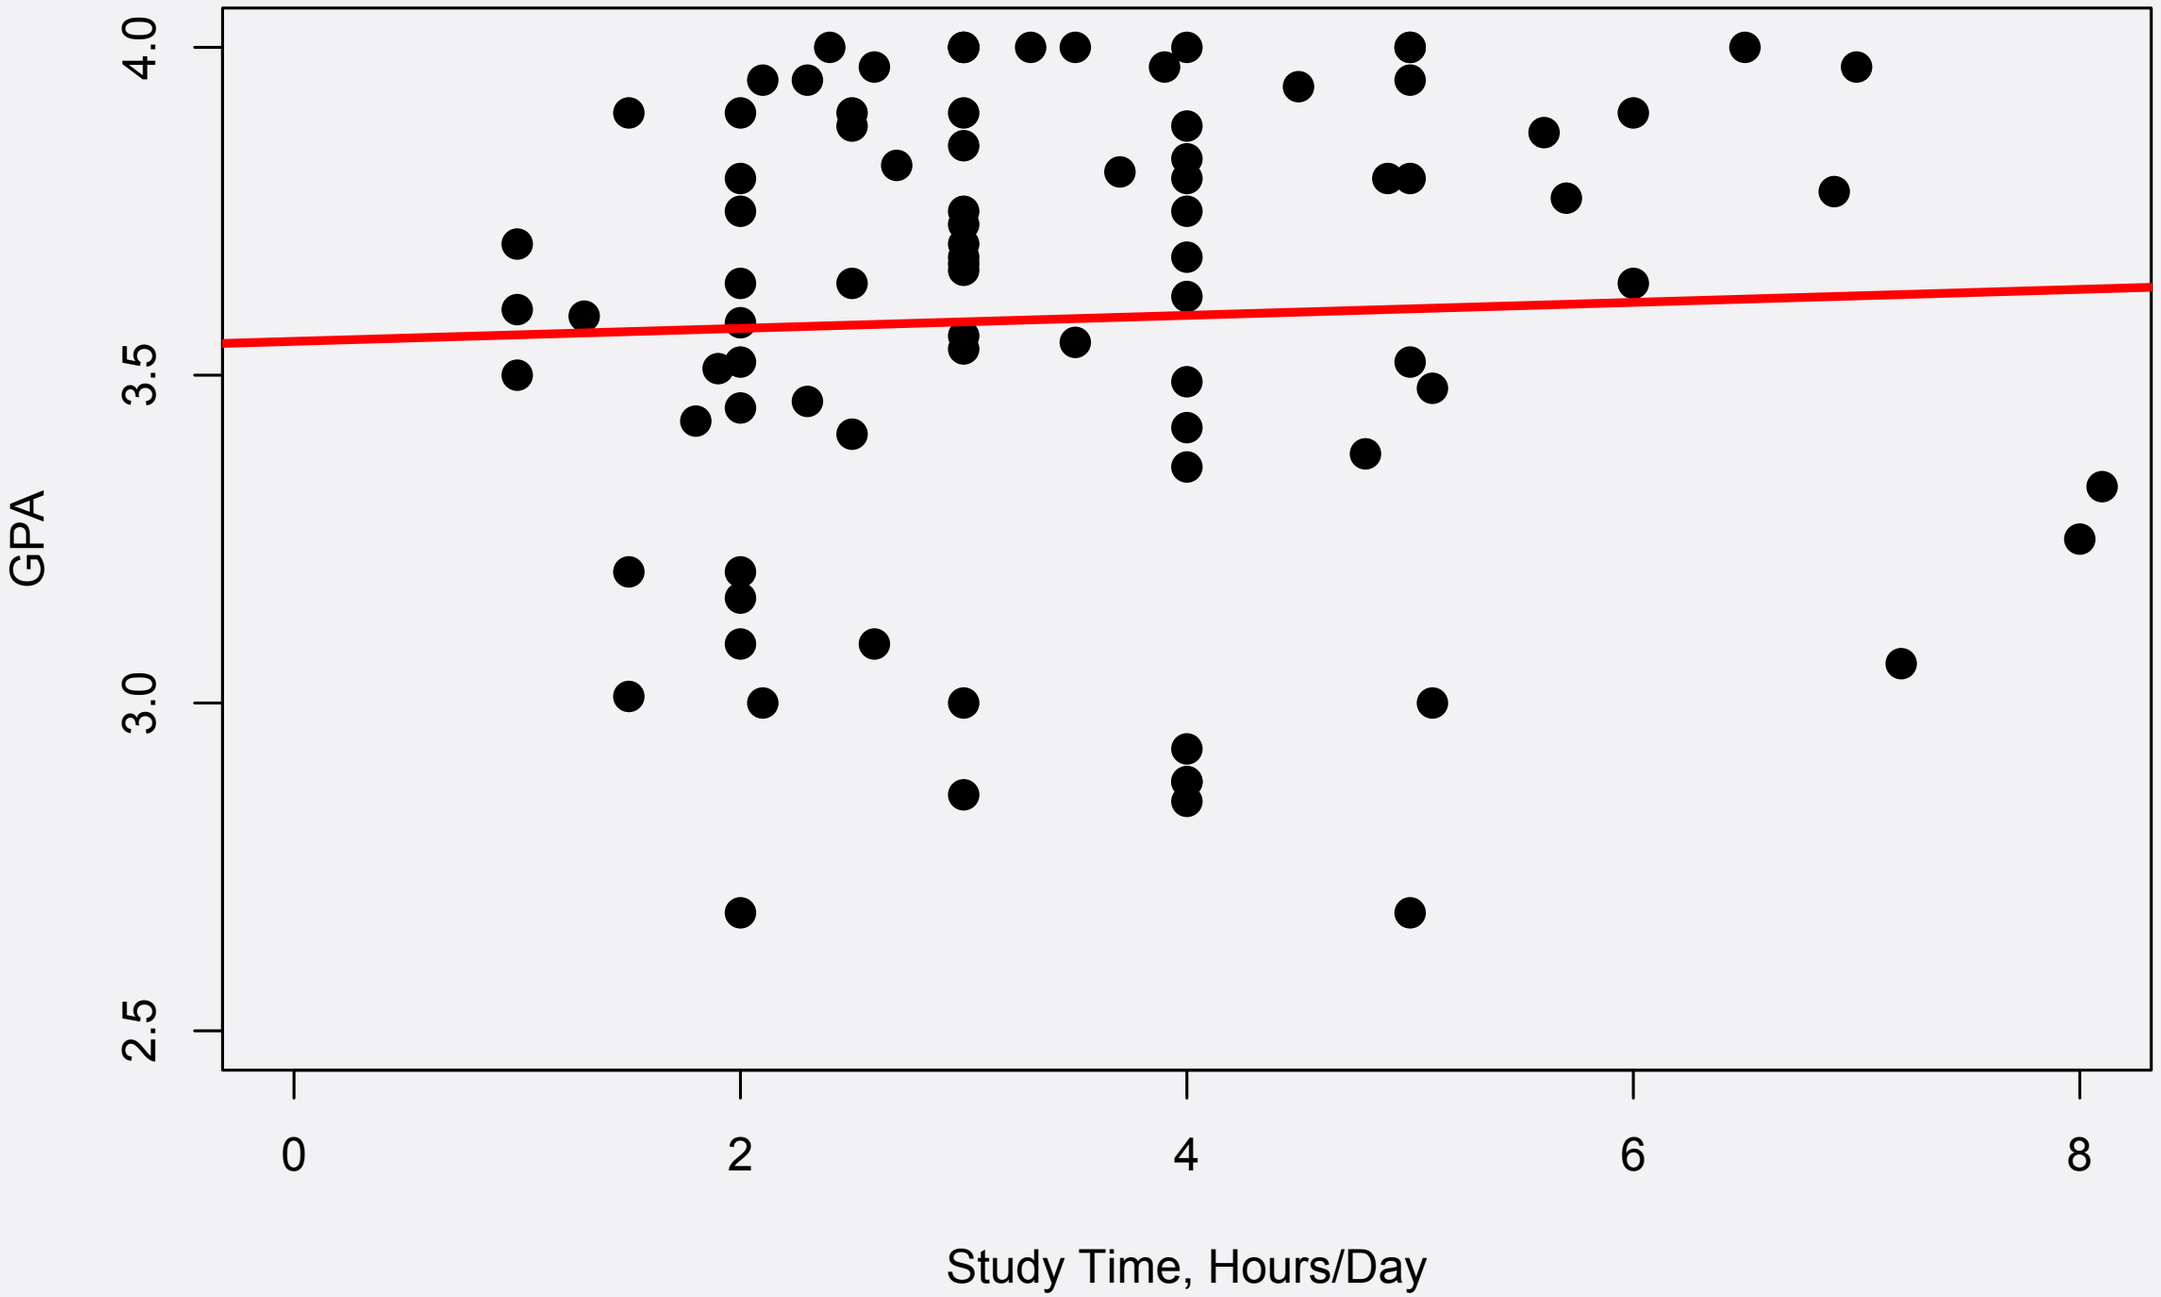

- GPA = 3.55 + 0.01 * Study Hours/Day

REGRESSION EQUATION

- GPA = 3.55 + 0.01 * Study Hours/Day
- General form: y = a + b * x
- y: dependent variable
- a: intercept
- b: slope
- x: independent variable

REGRESSION EQUATION

- Regression: $y = a + b * x$
- High School Math: $y = m * x + b$

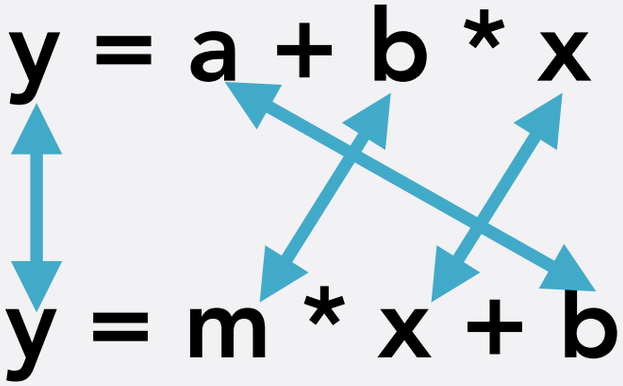

INTERPRETATION

- y = a + b * x
  - Interpretation of slope: For every one unit increase in x, y changes by b units
  - Interpretation of intercept: When x=0, y takes the value a

EXAMPLE 2

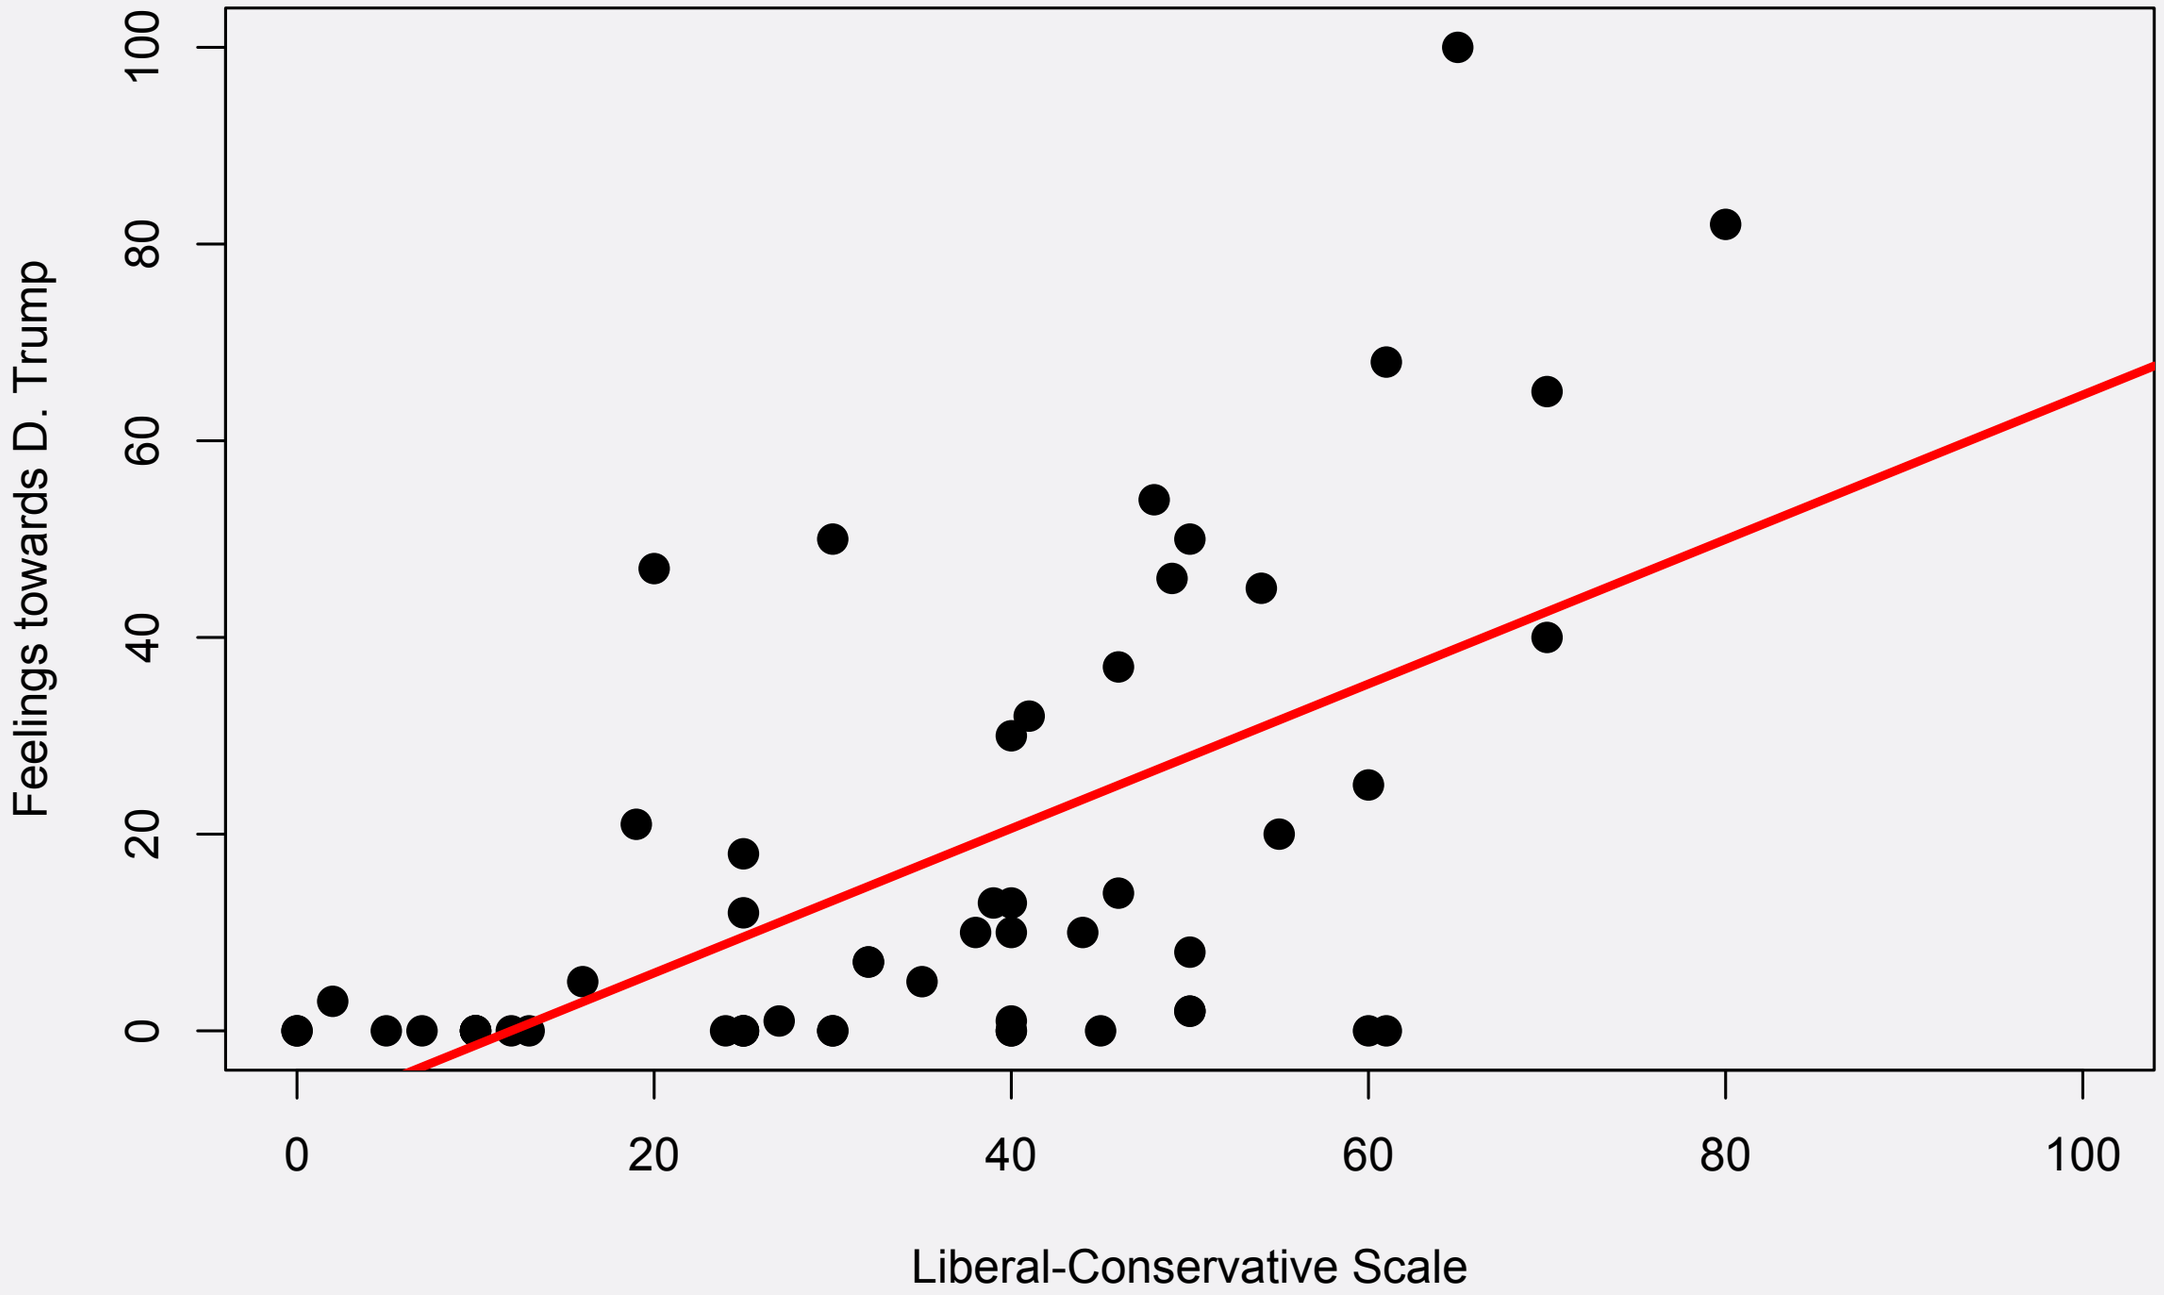

- **Trump Thermometer = -9 + 0.7 * Lib/Cons**

EXAMPLE 3

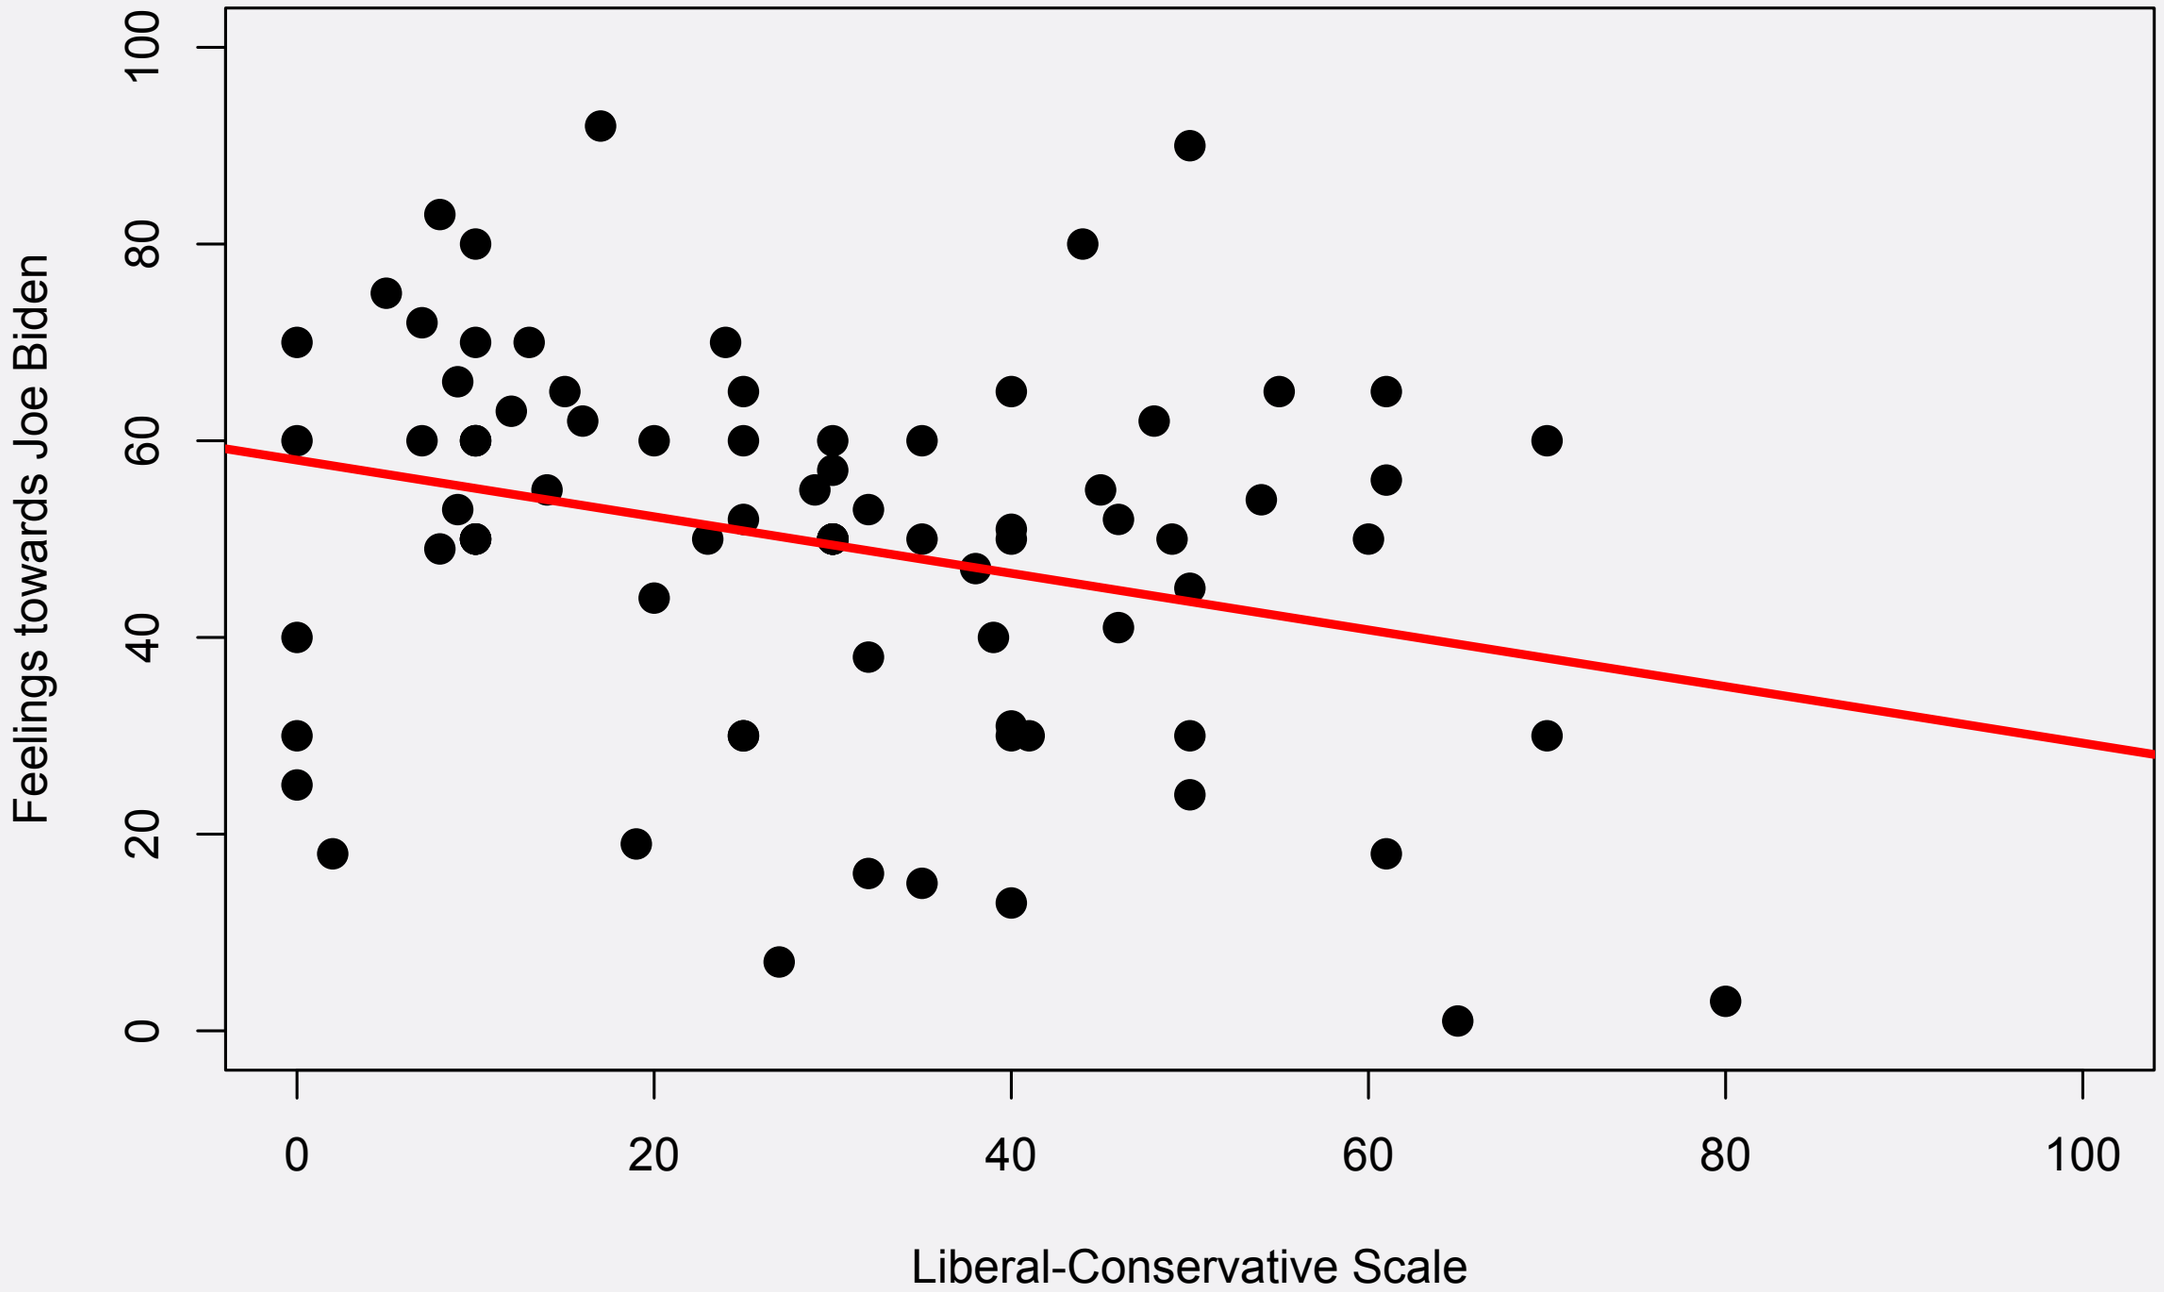

- **Biden Thermometer = 60 - 0.3 * Lib/Cons**

HOW TO PICK THE LINE

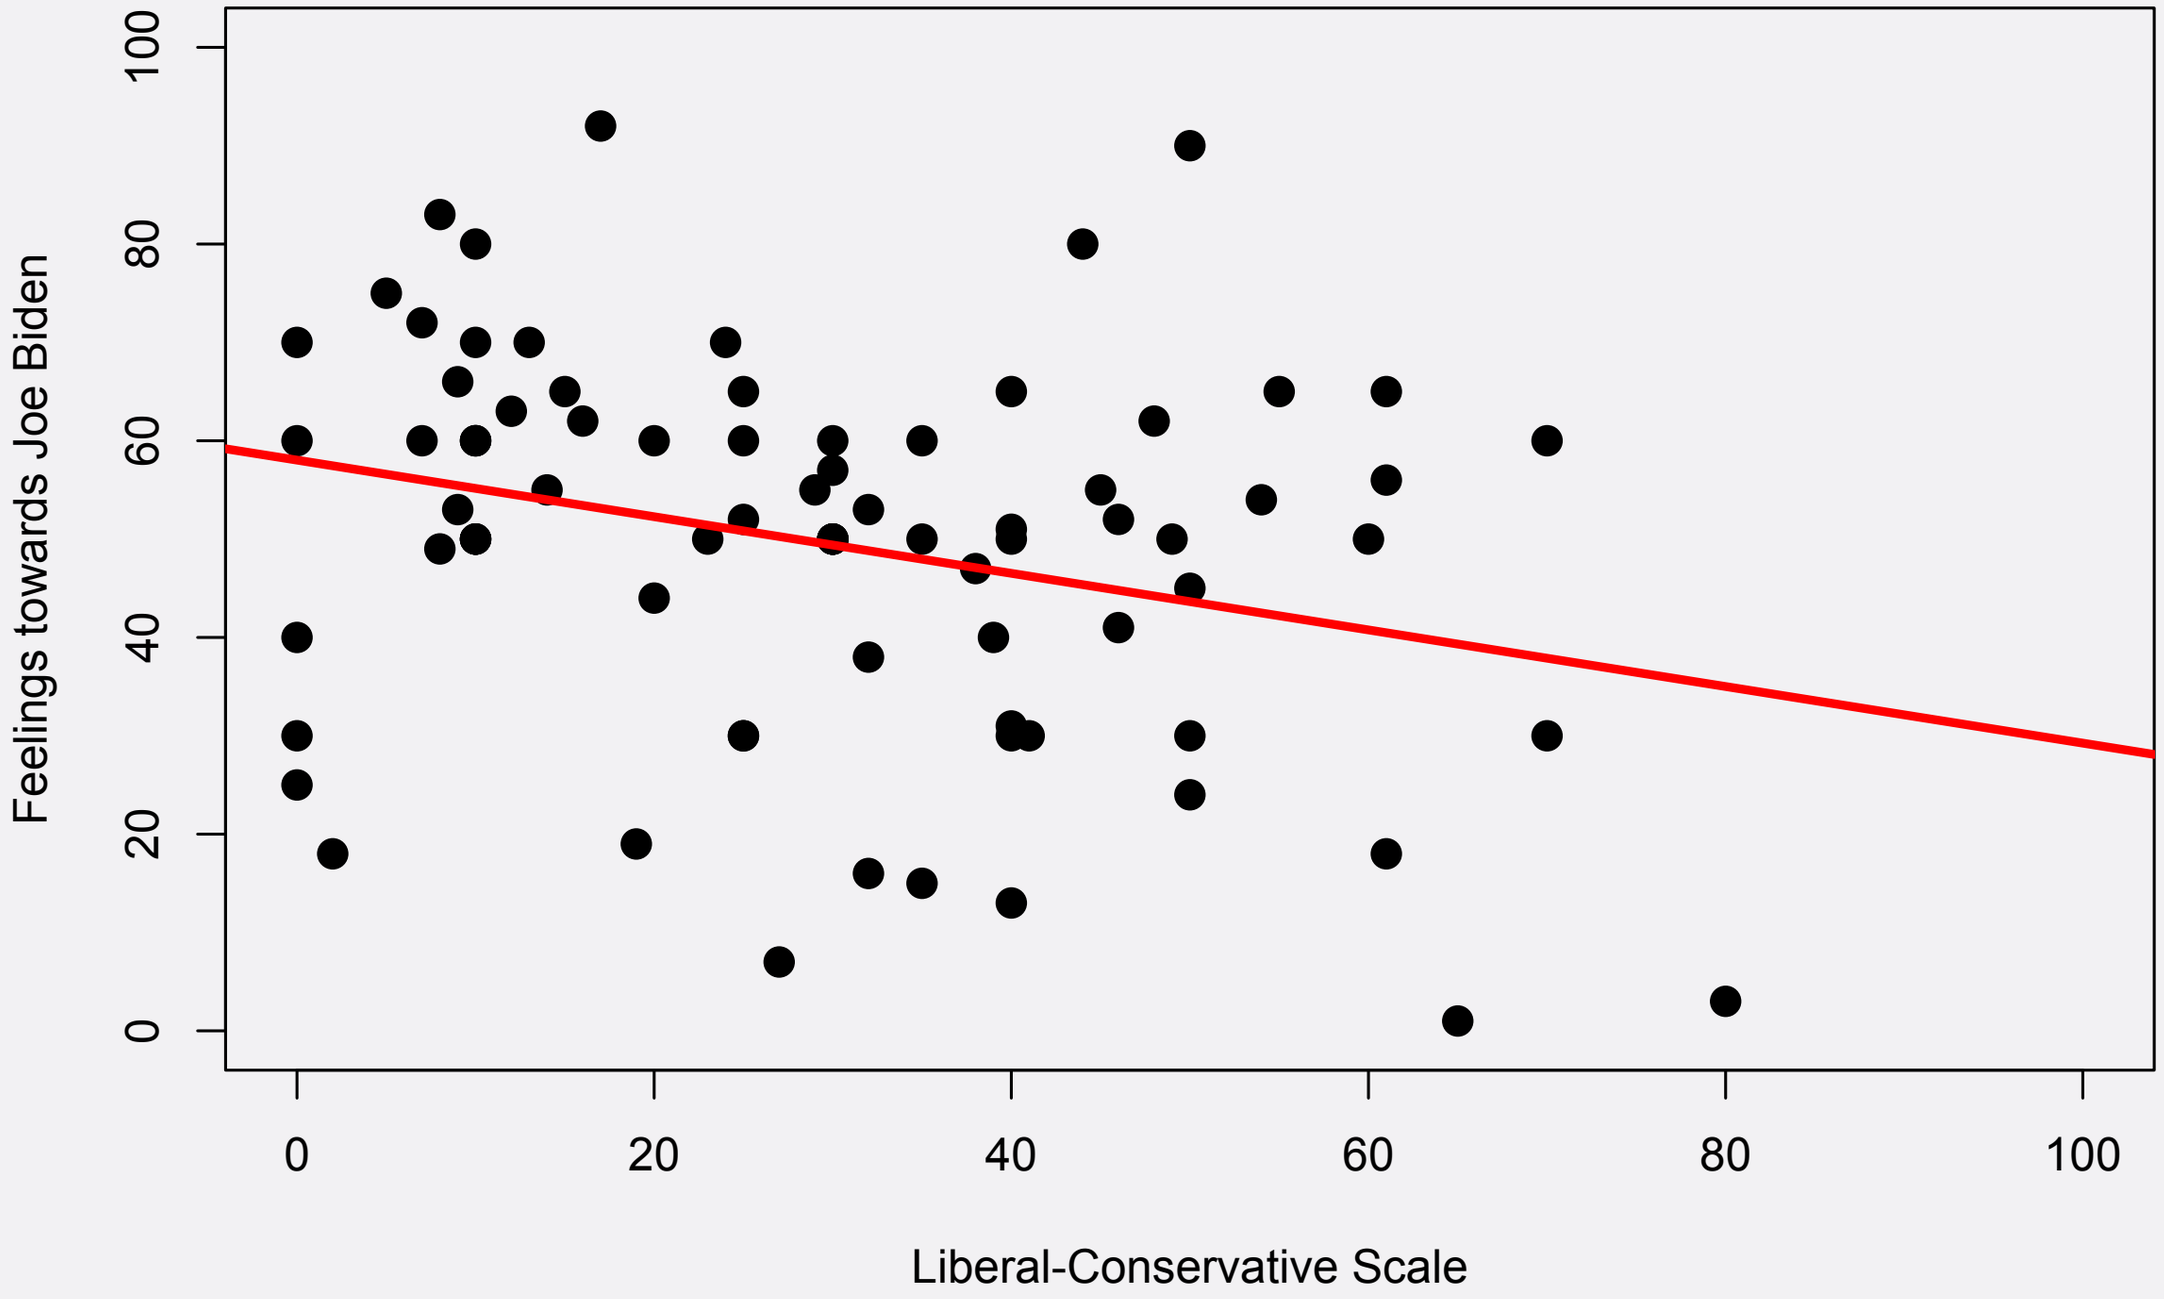

- **Why this line?**

HOW TO PICK THE LINE

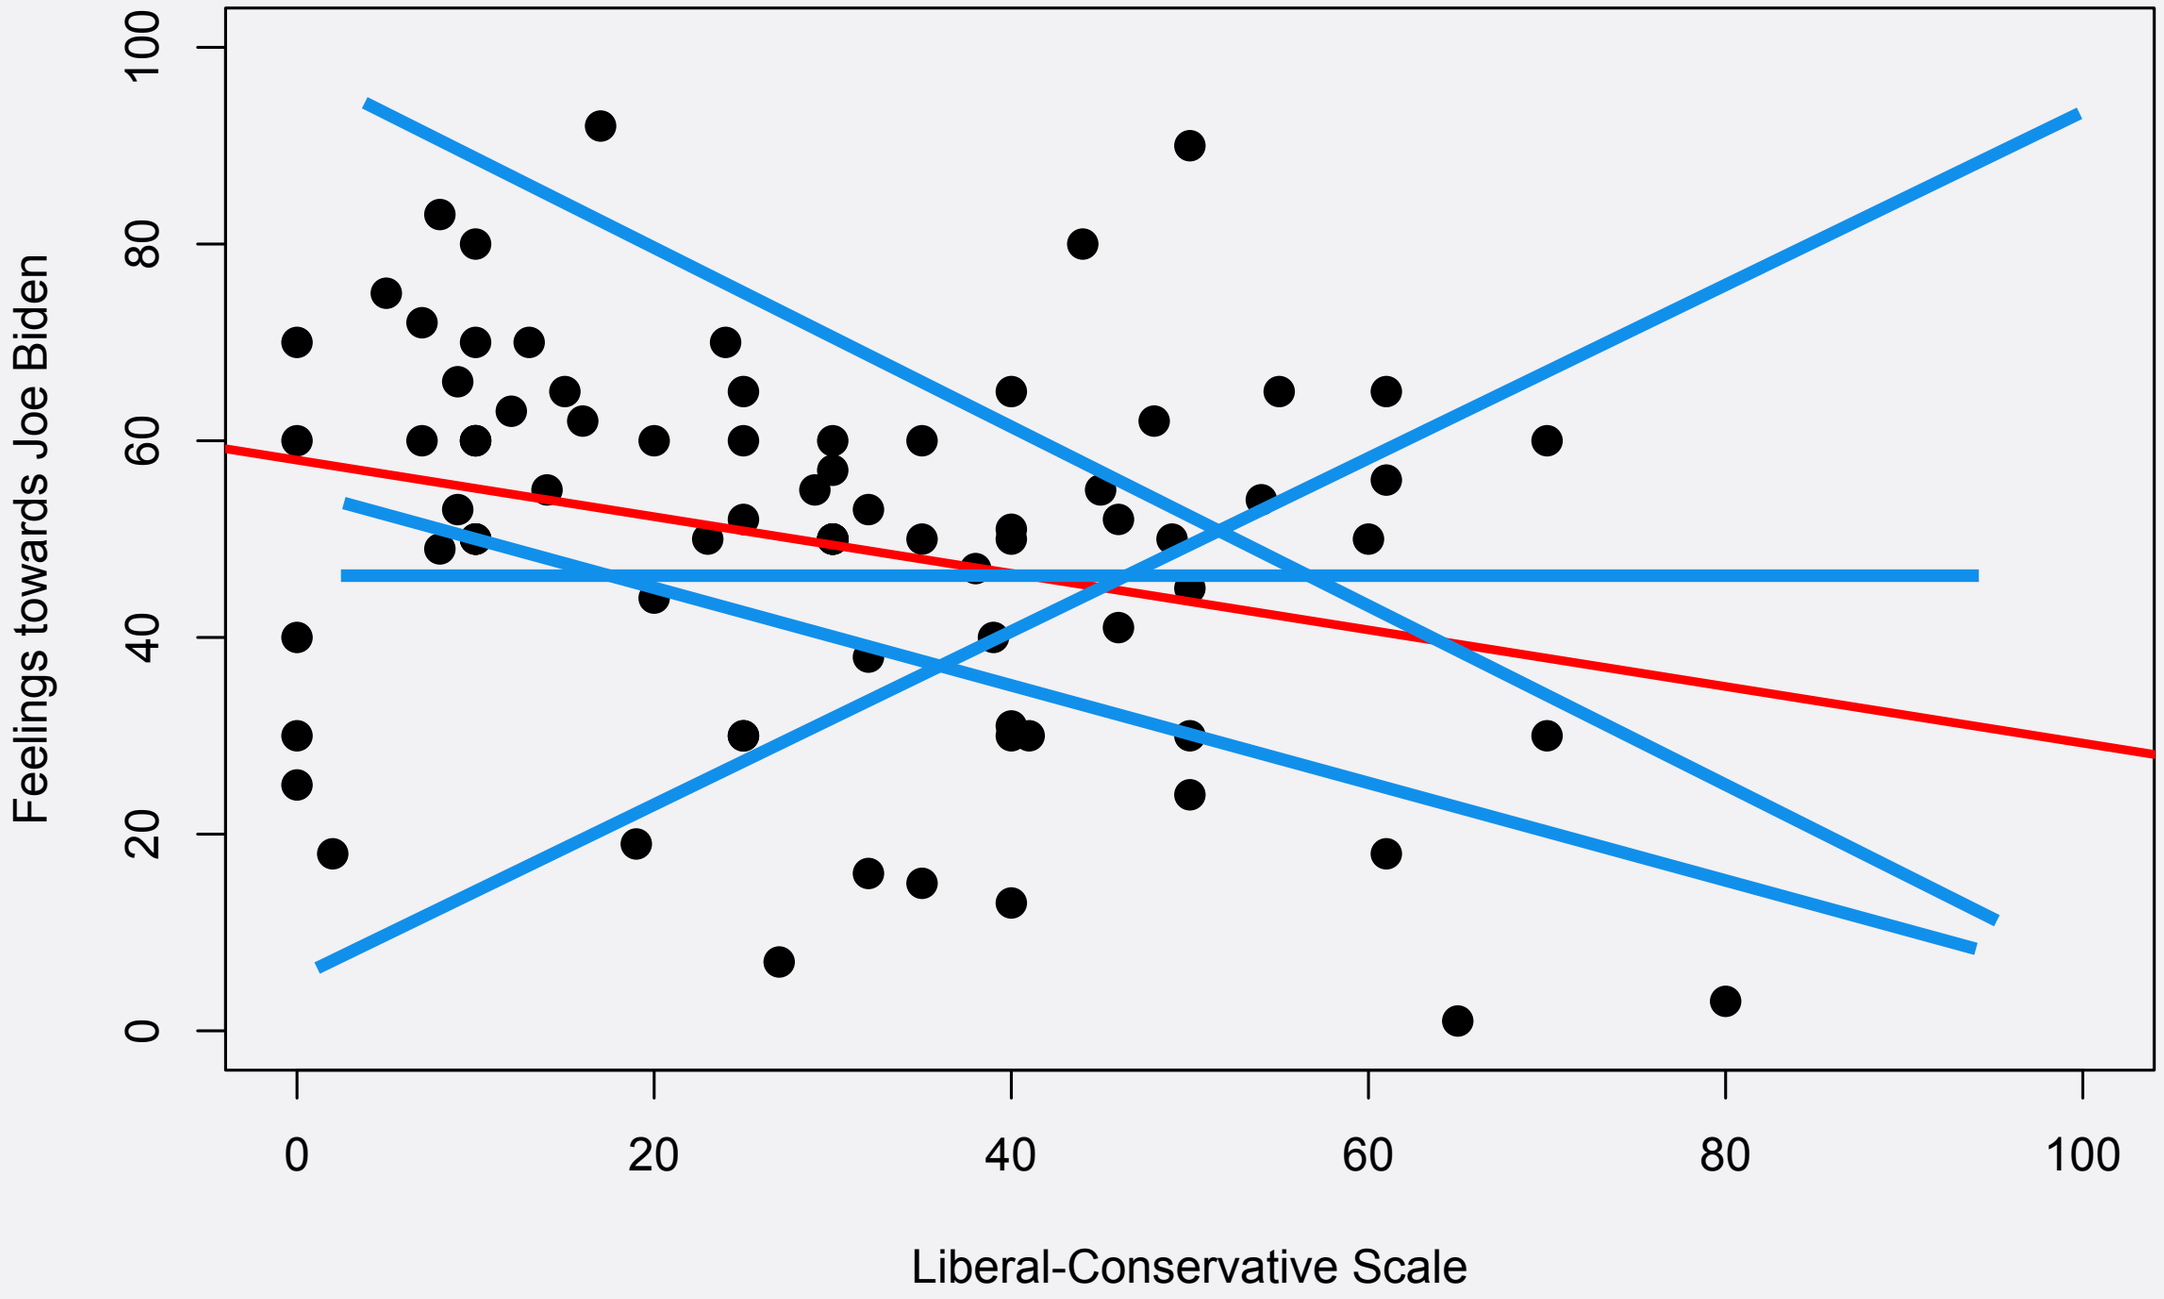

- **Why not these?**

HOW TO PICK THE LINE

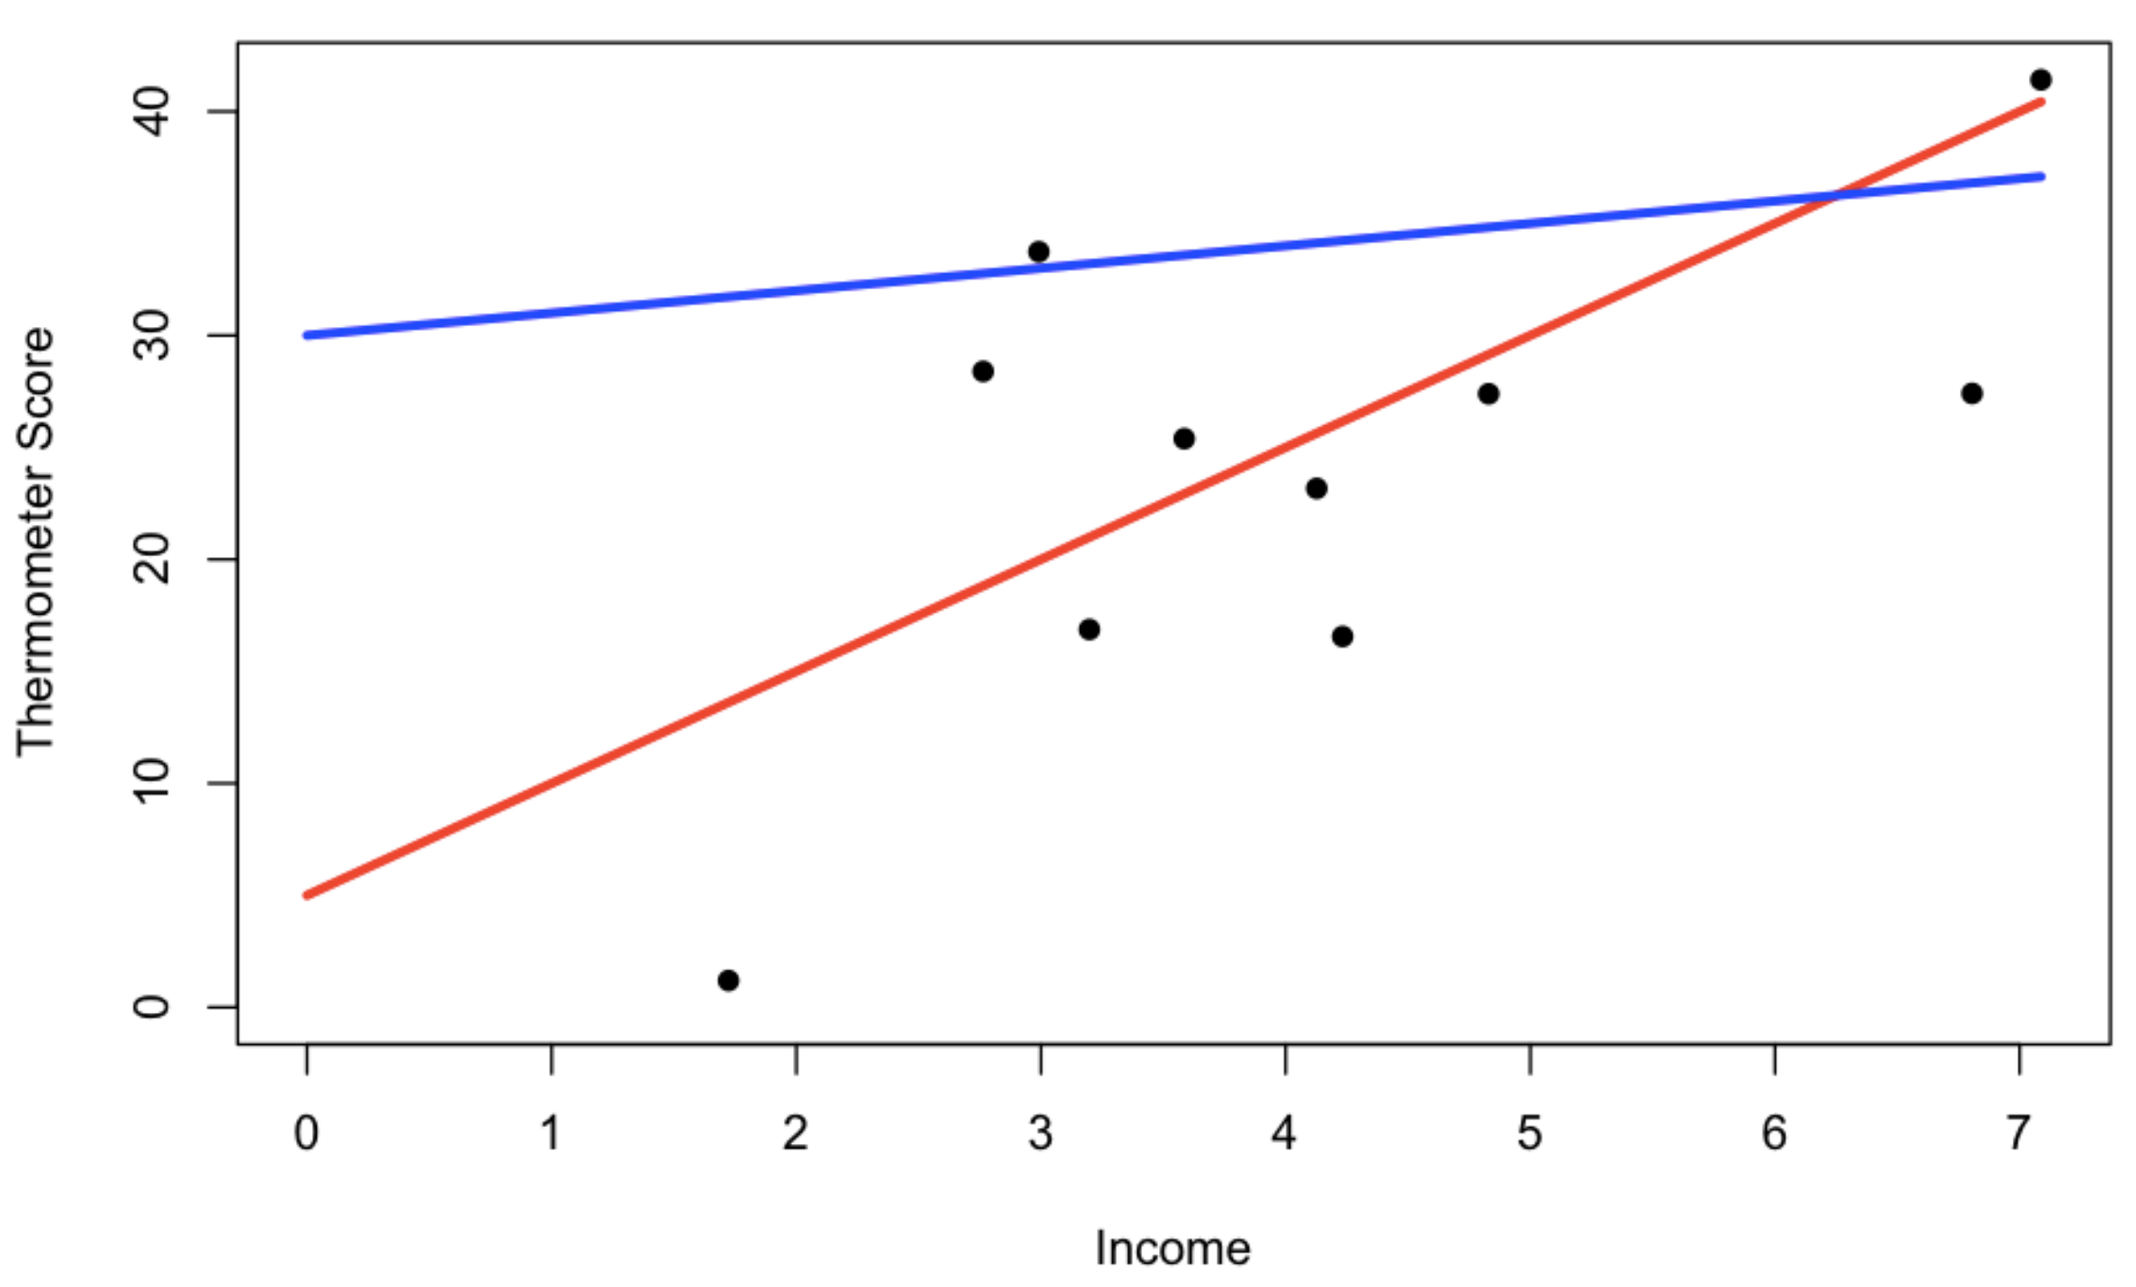

- **Which line is better?**

HOW TO PICK THE LINE

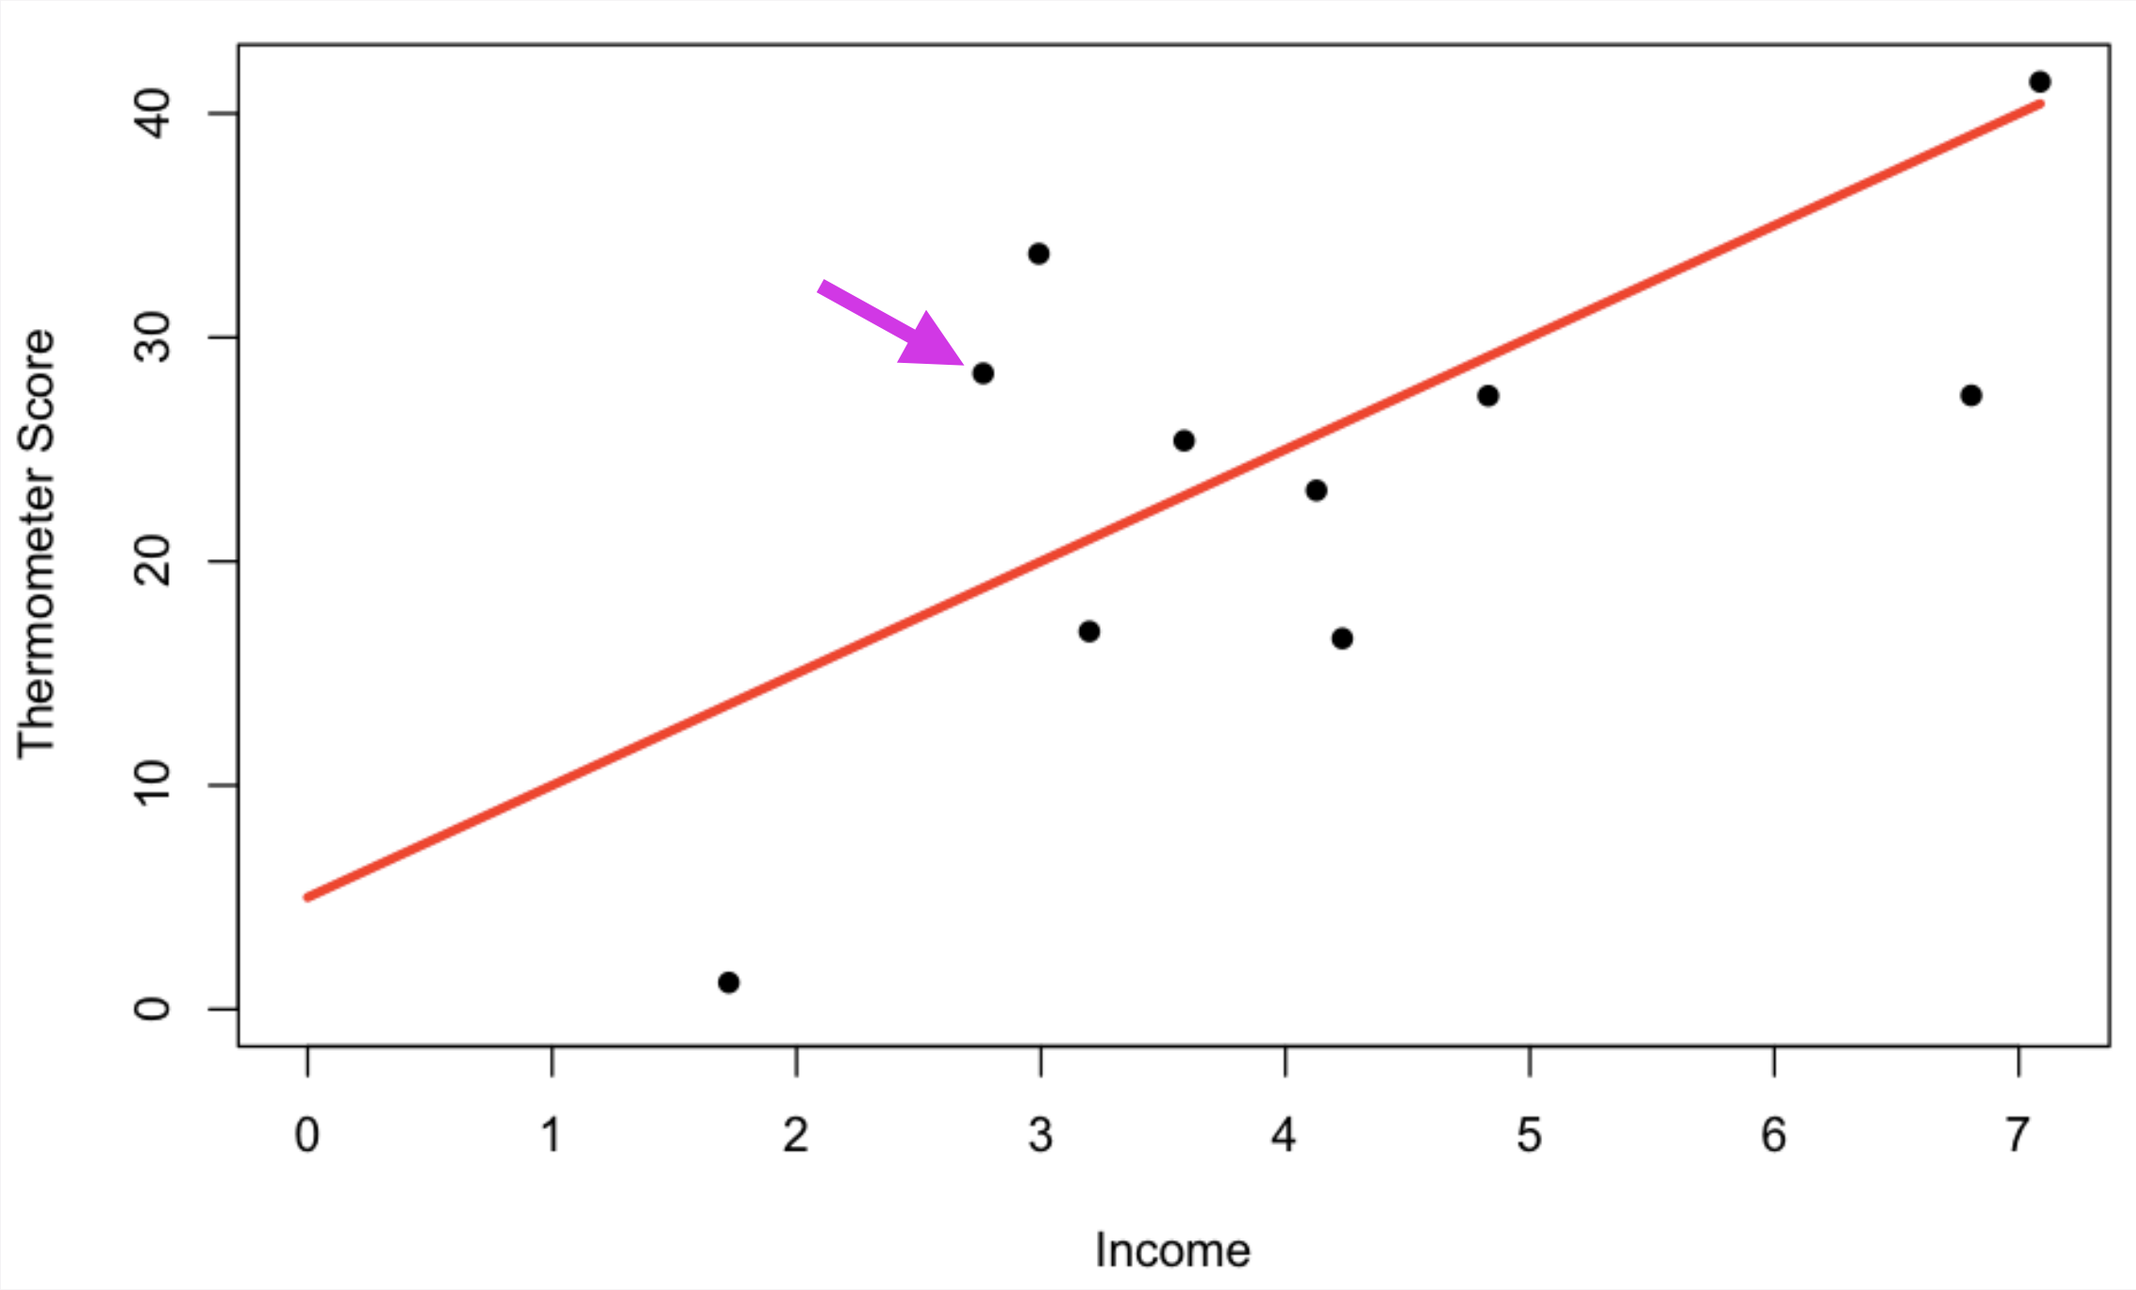

HOW TO PICK THE LINE

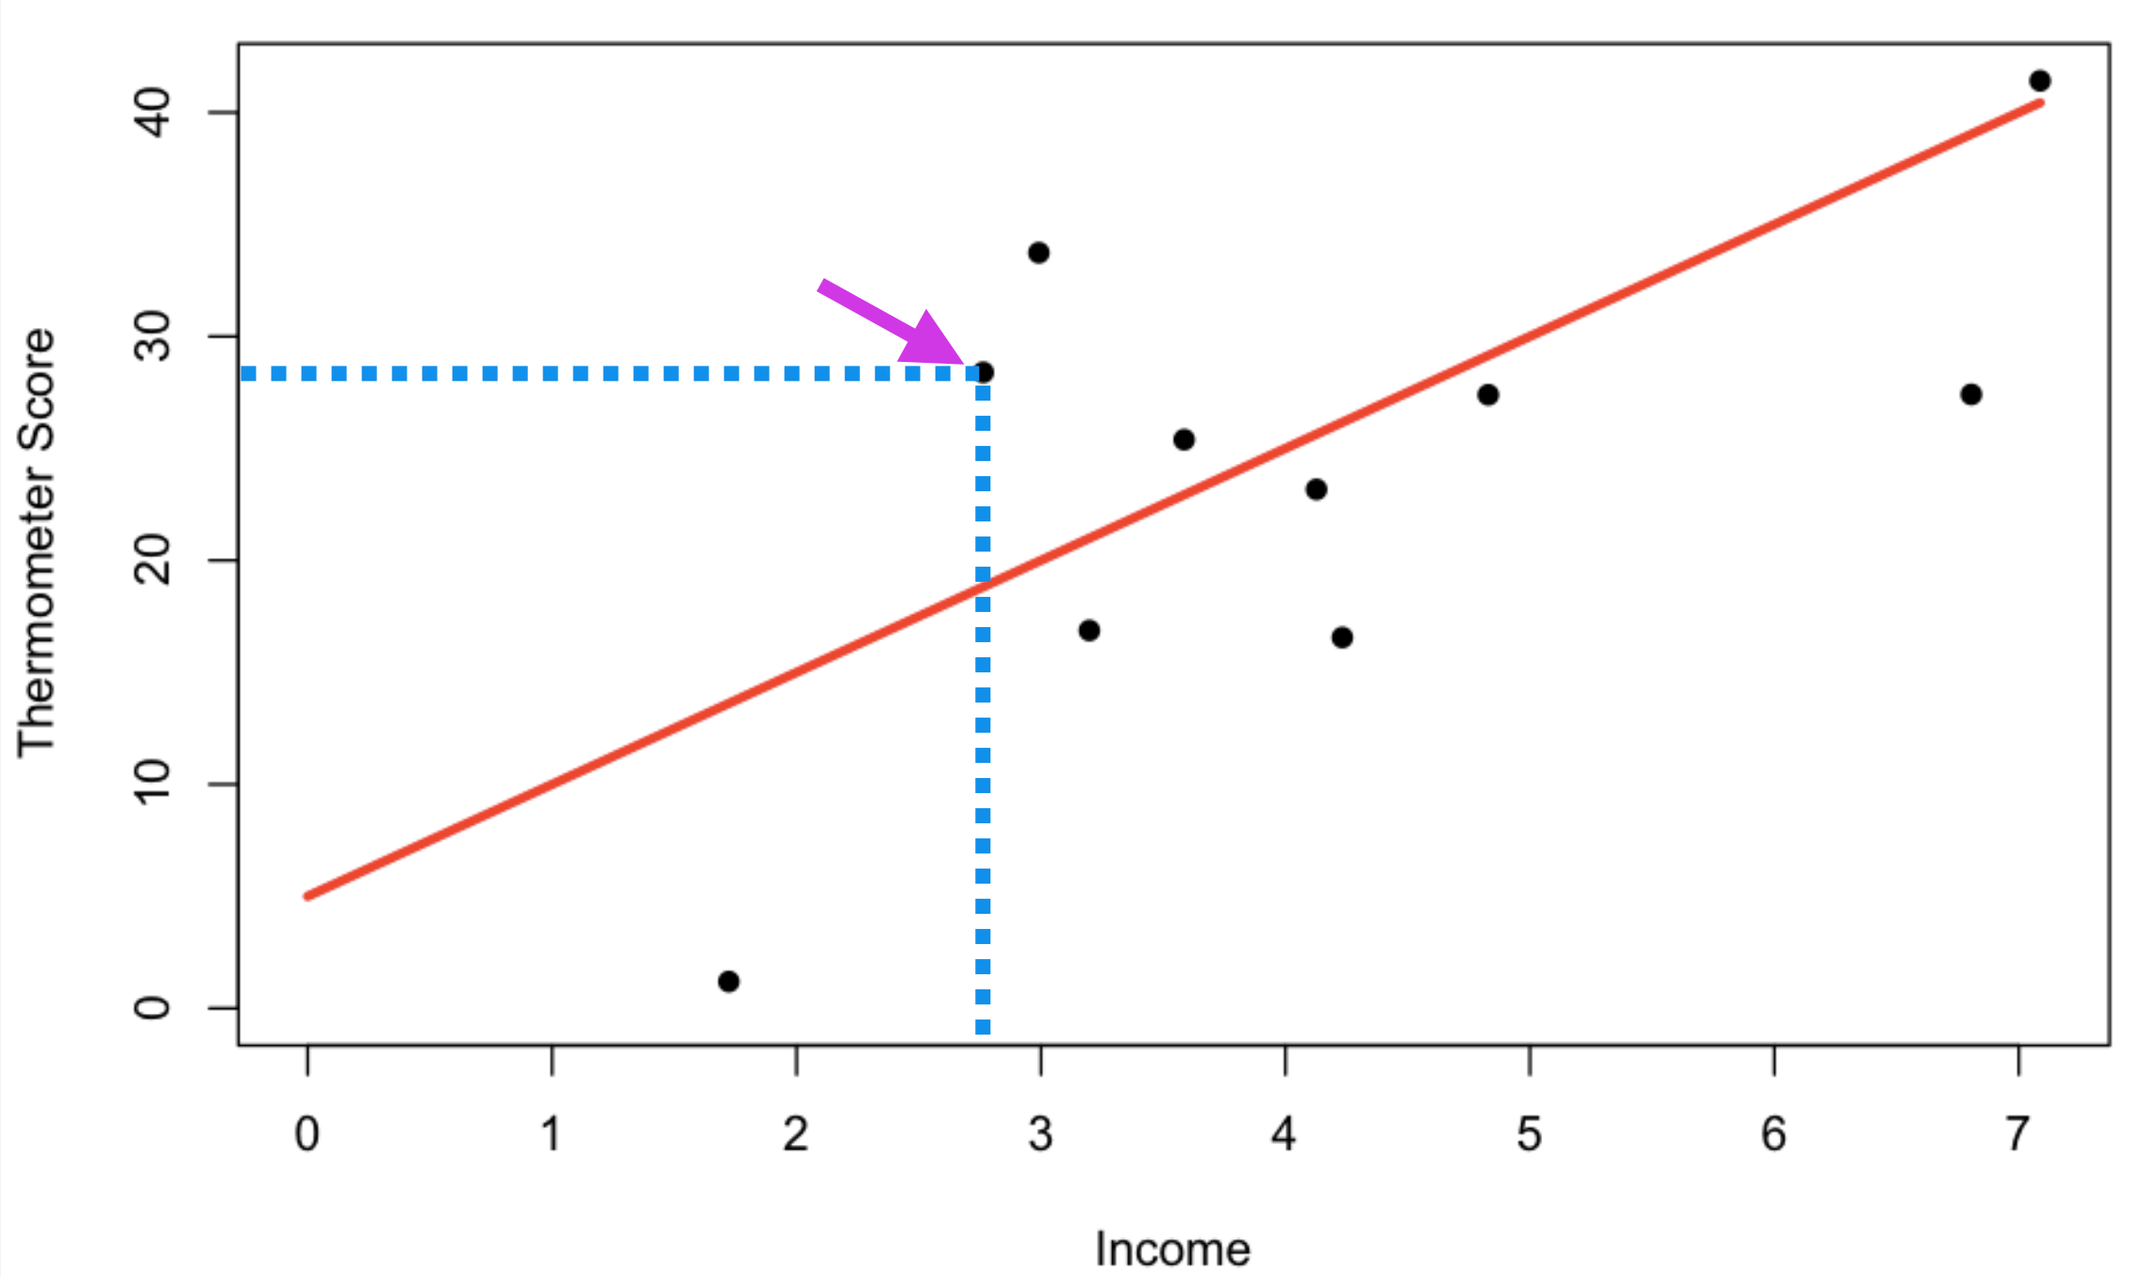

- **Actual y-value: y=28**

HOW TO PICK THE LINE

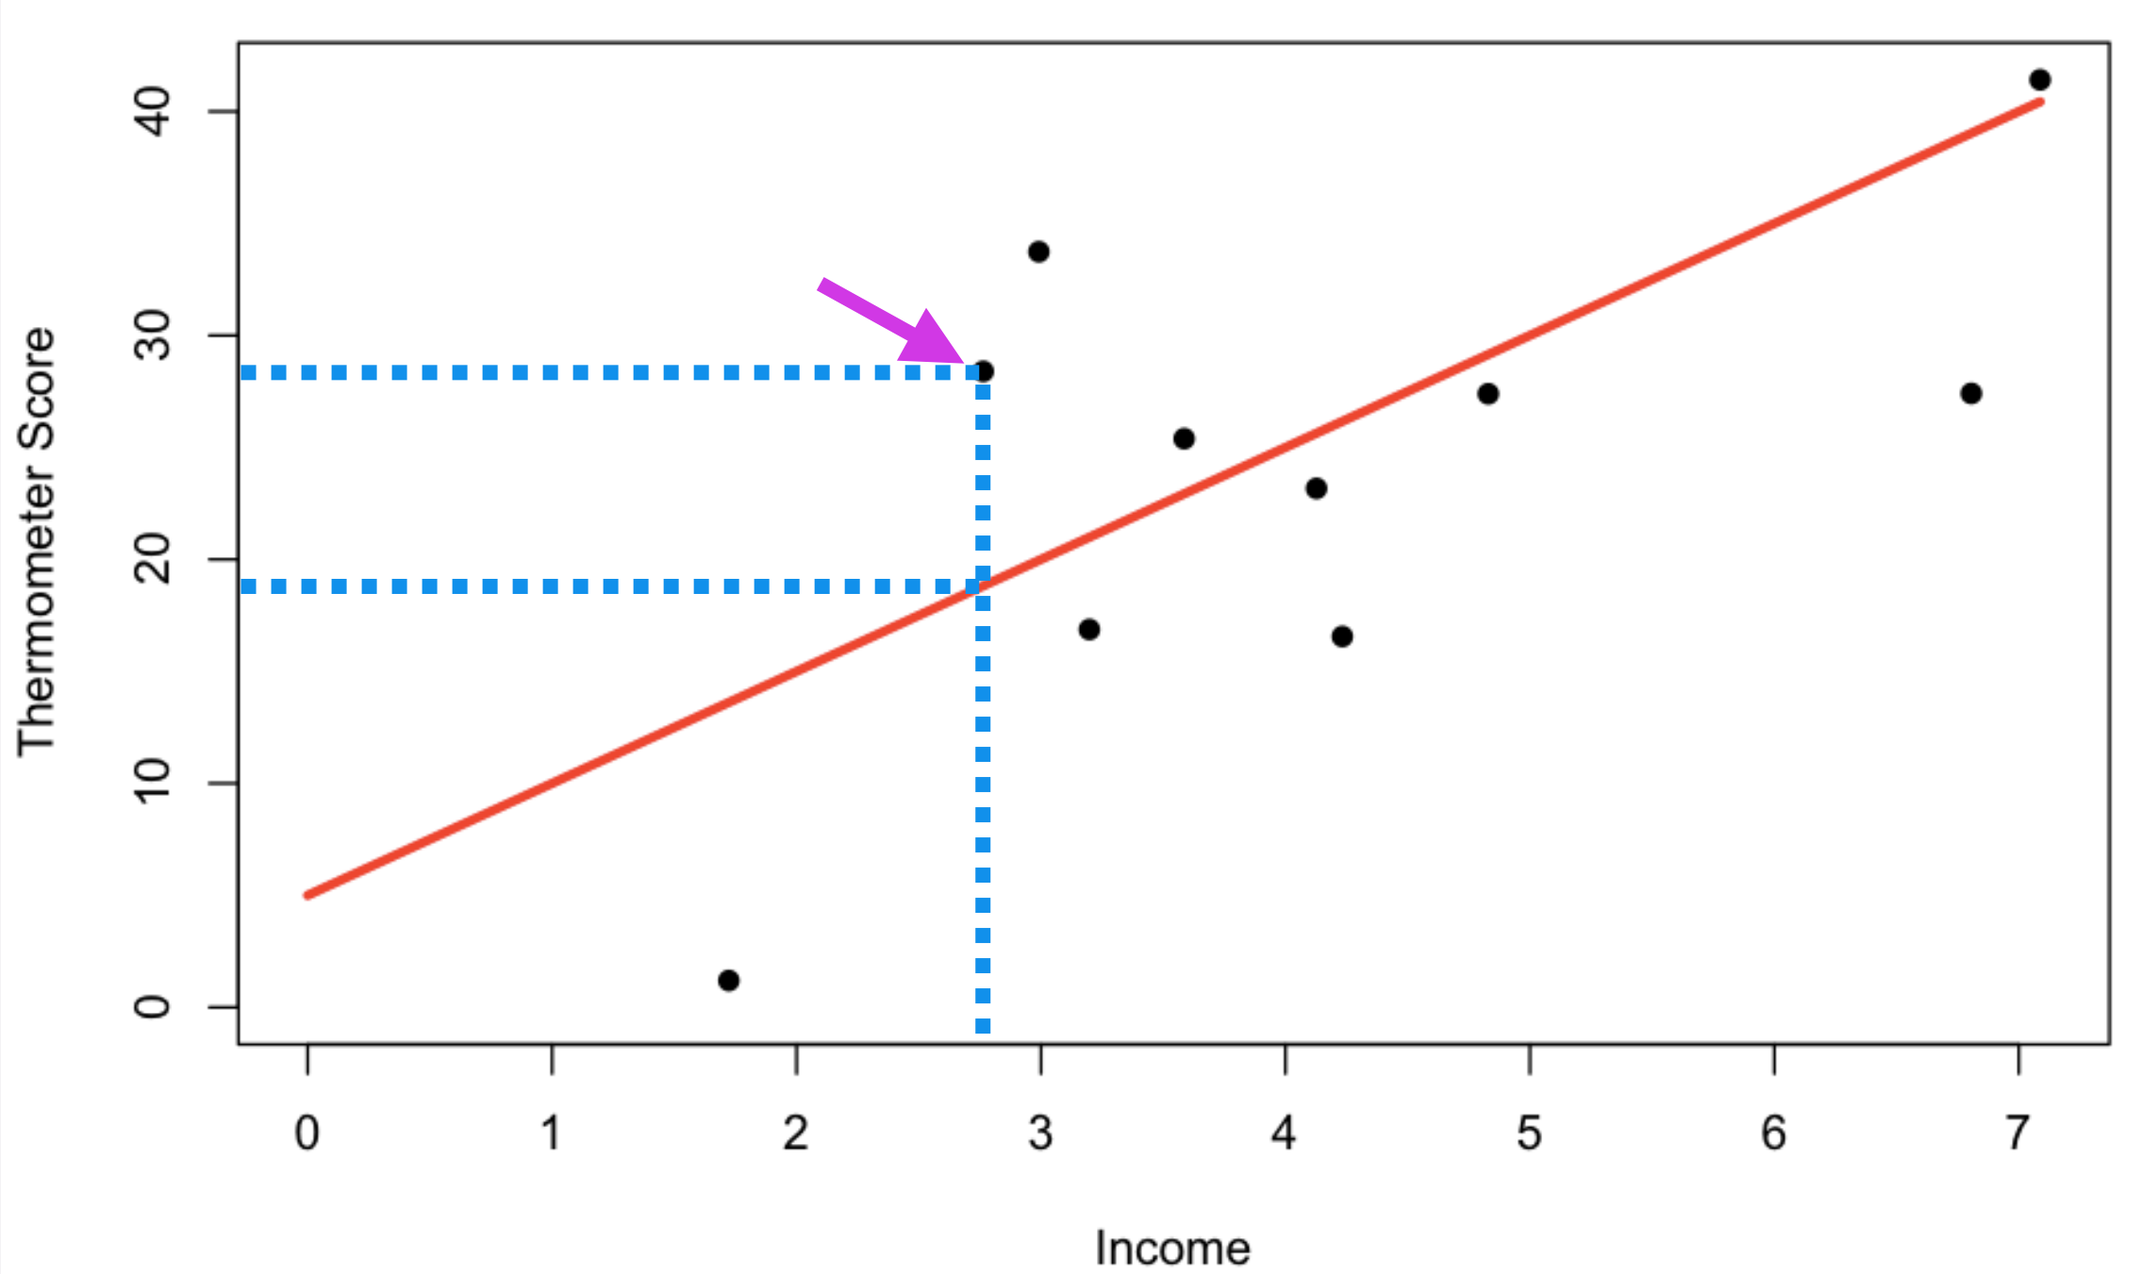

- **Predicted y-value: $\hat{y}=19$**

HOW TO PICK THE LINE

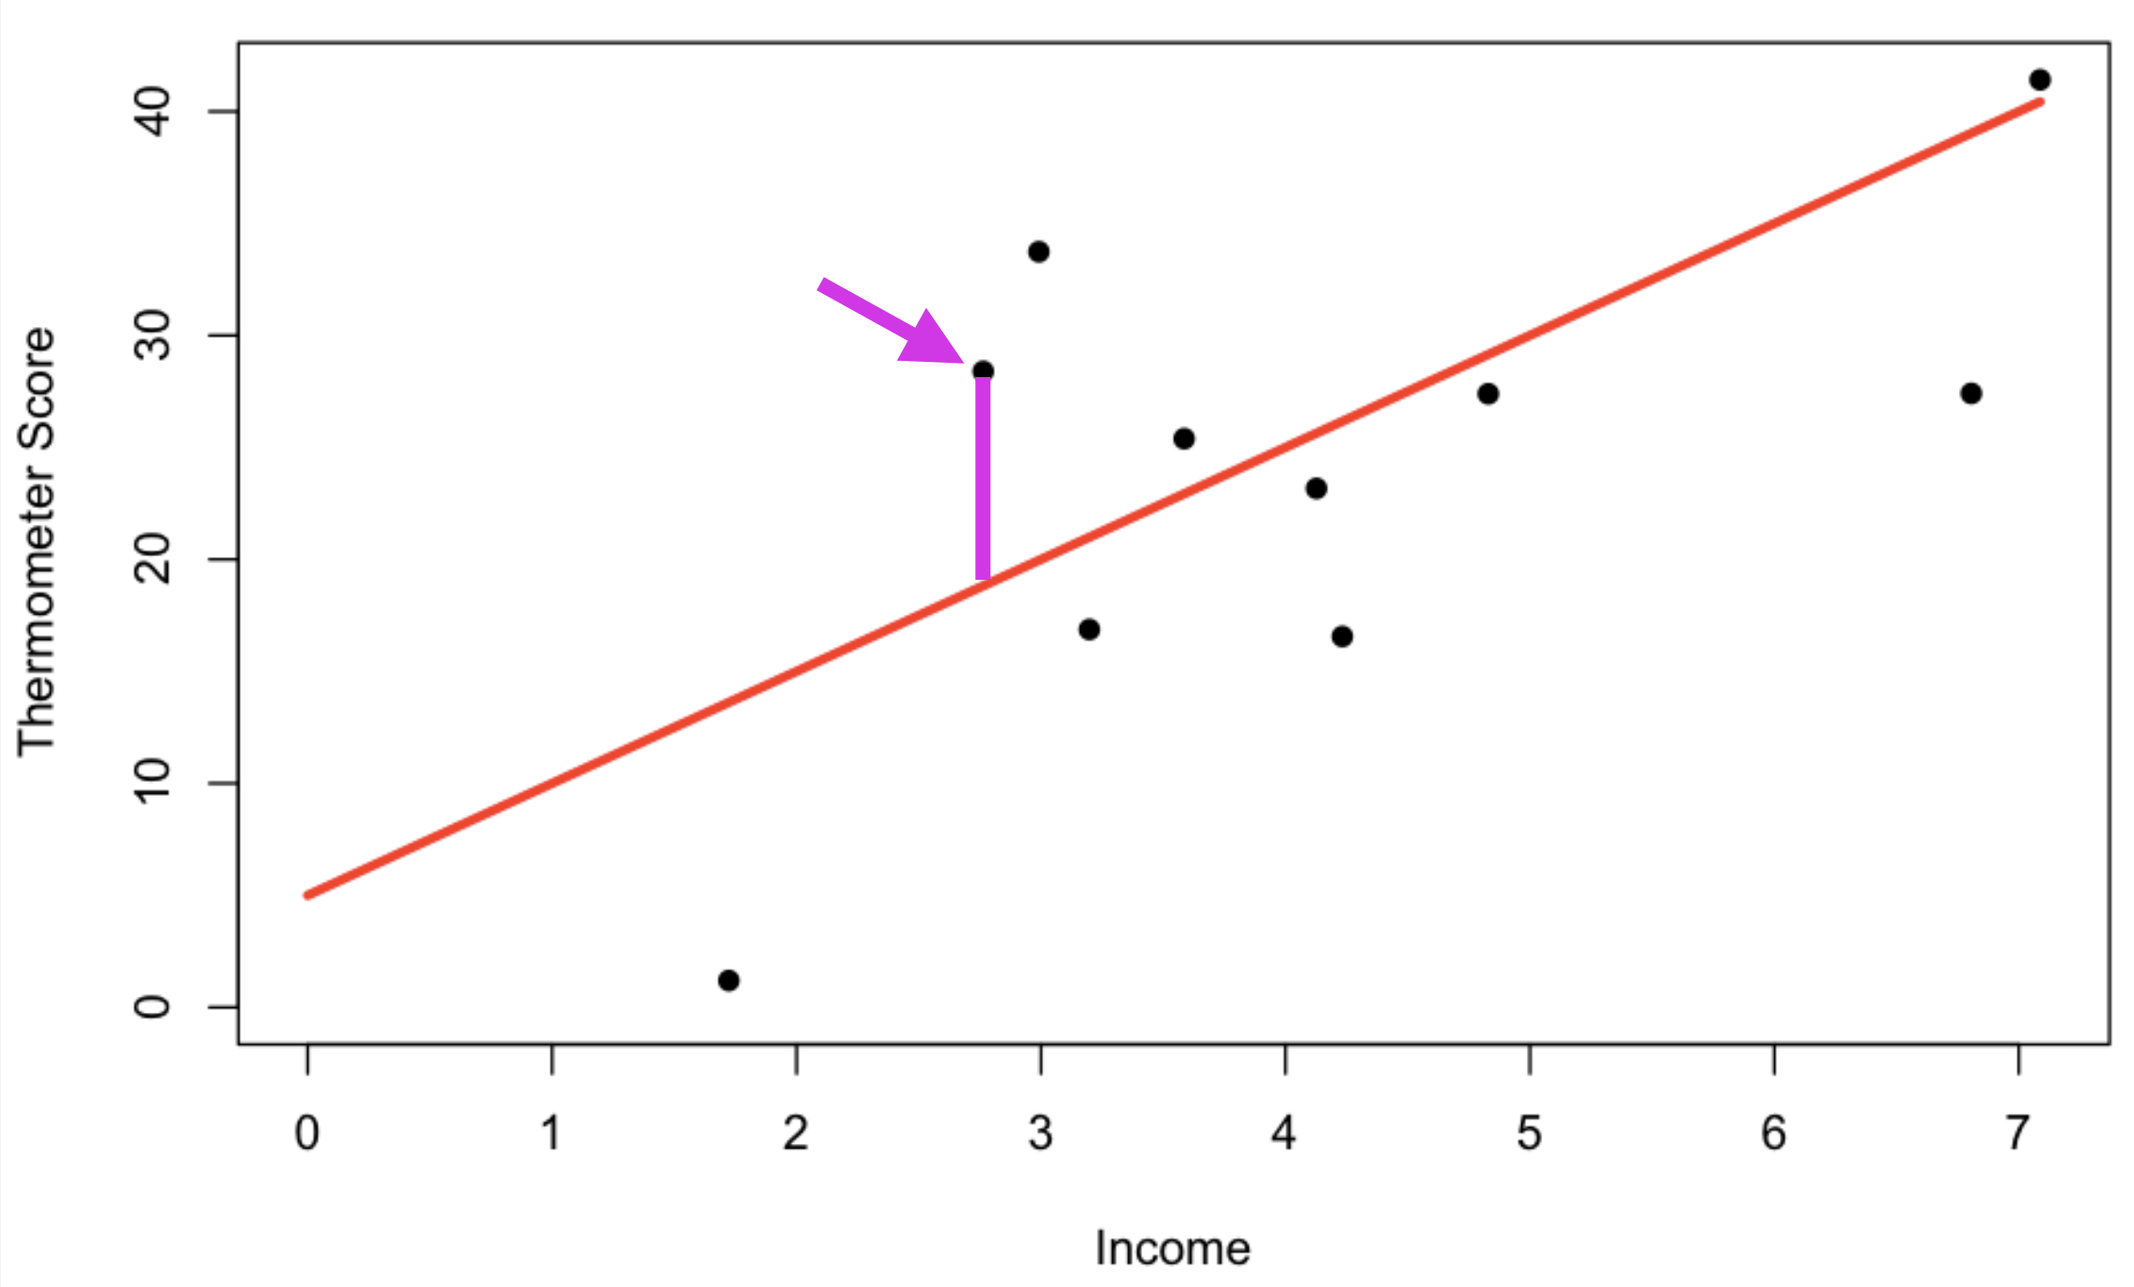

- **Prediction error: $y - \hat{y} = 28 - 19 = 9$**

HOW TO PICK THE LINE

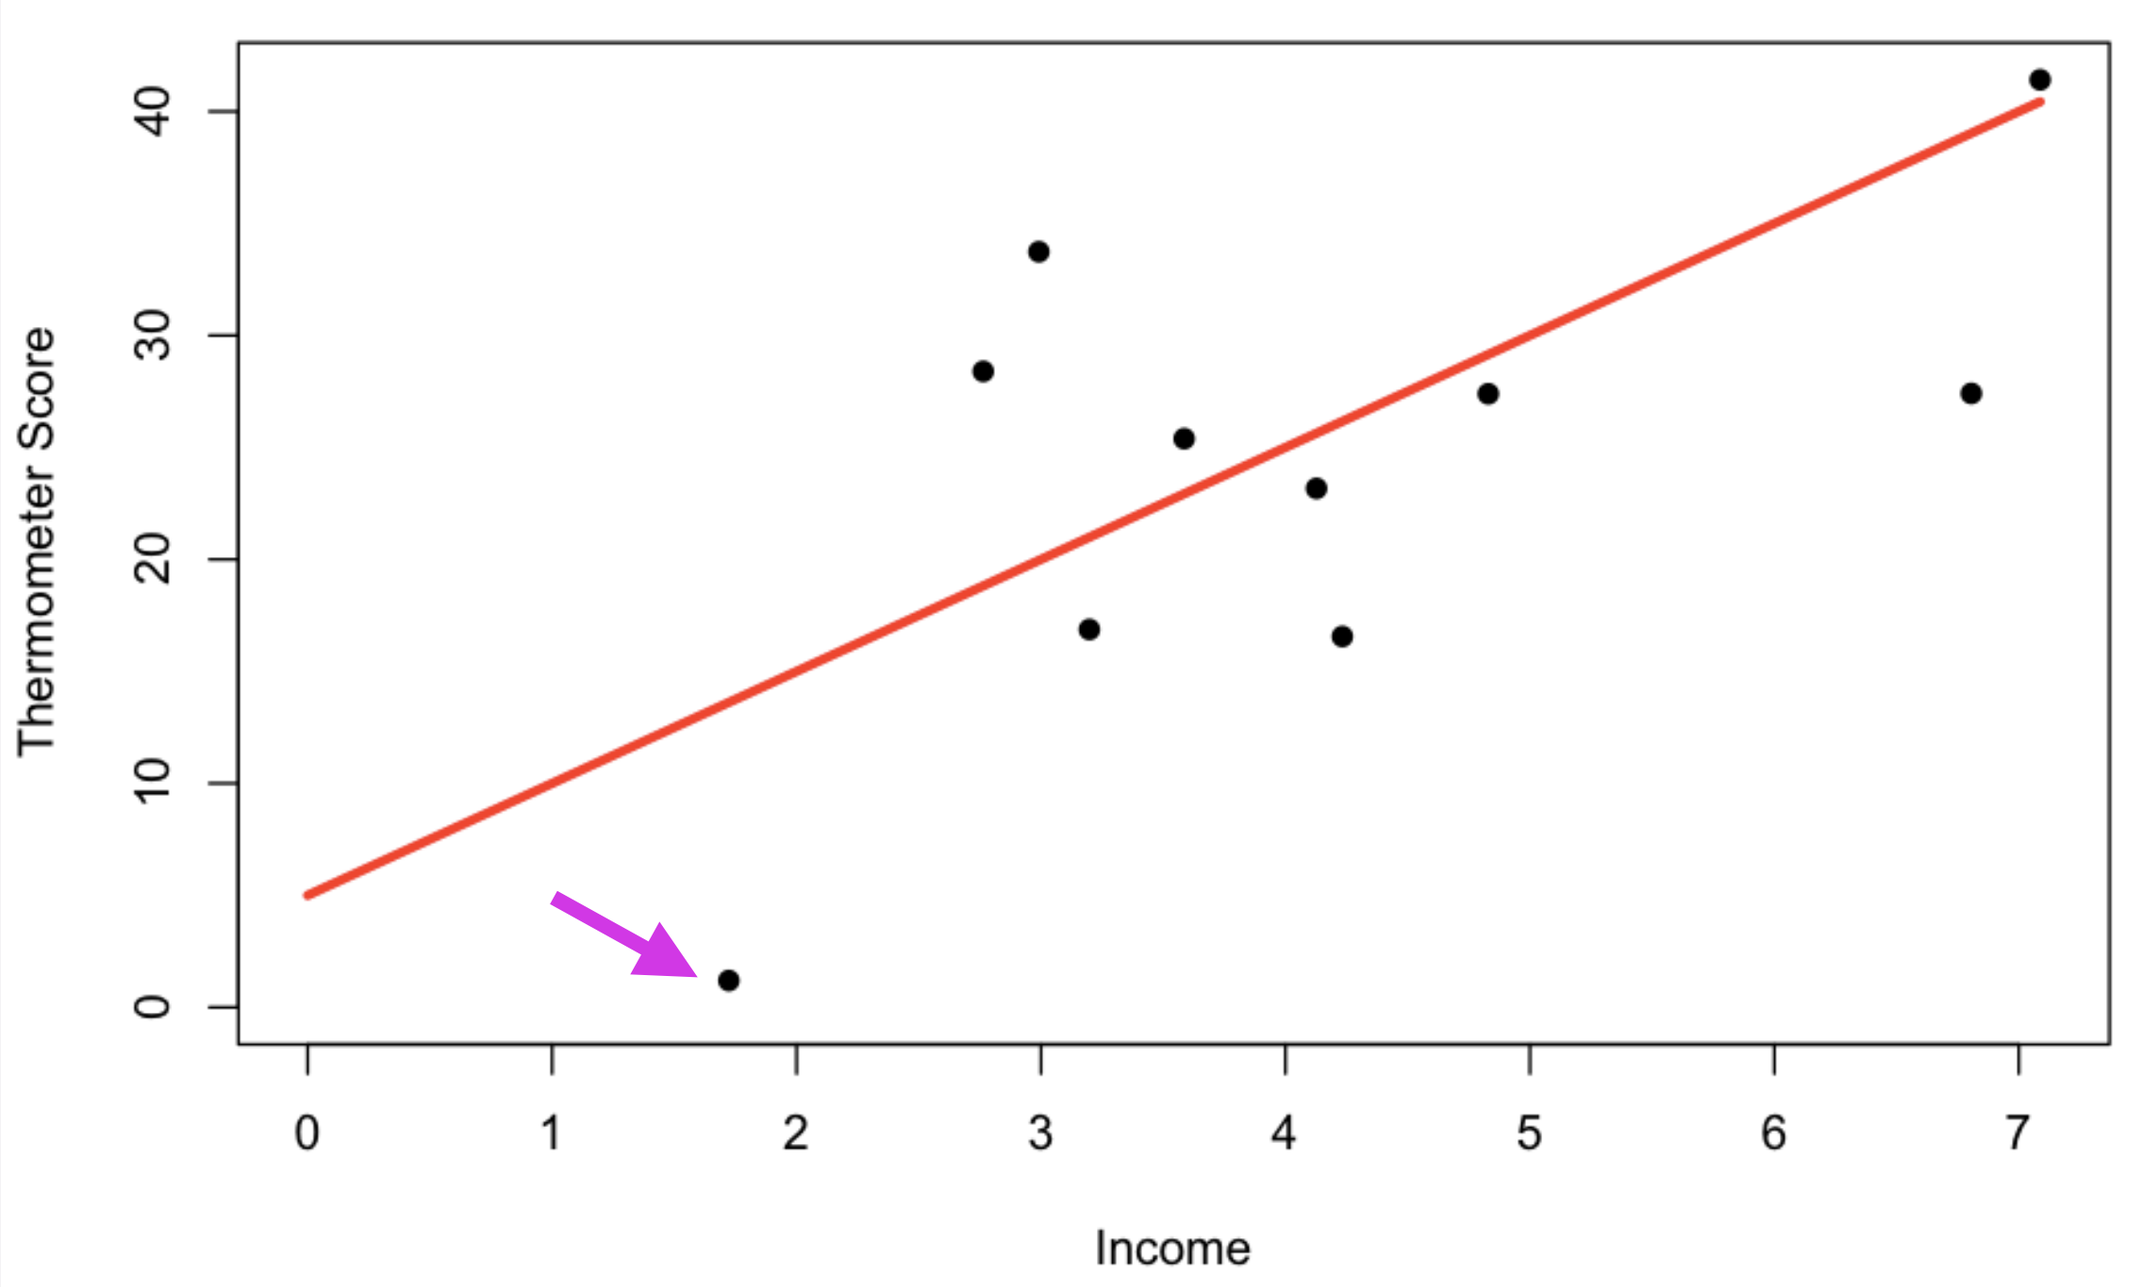

HOW TO PICK THE LINE

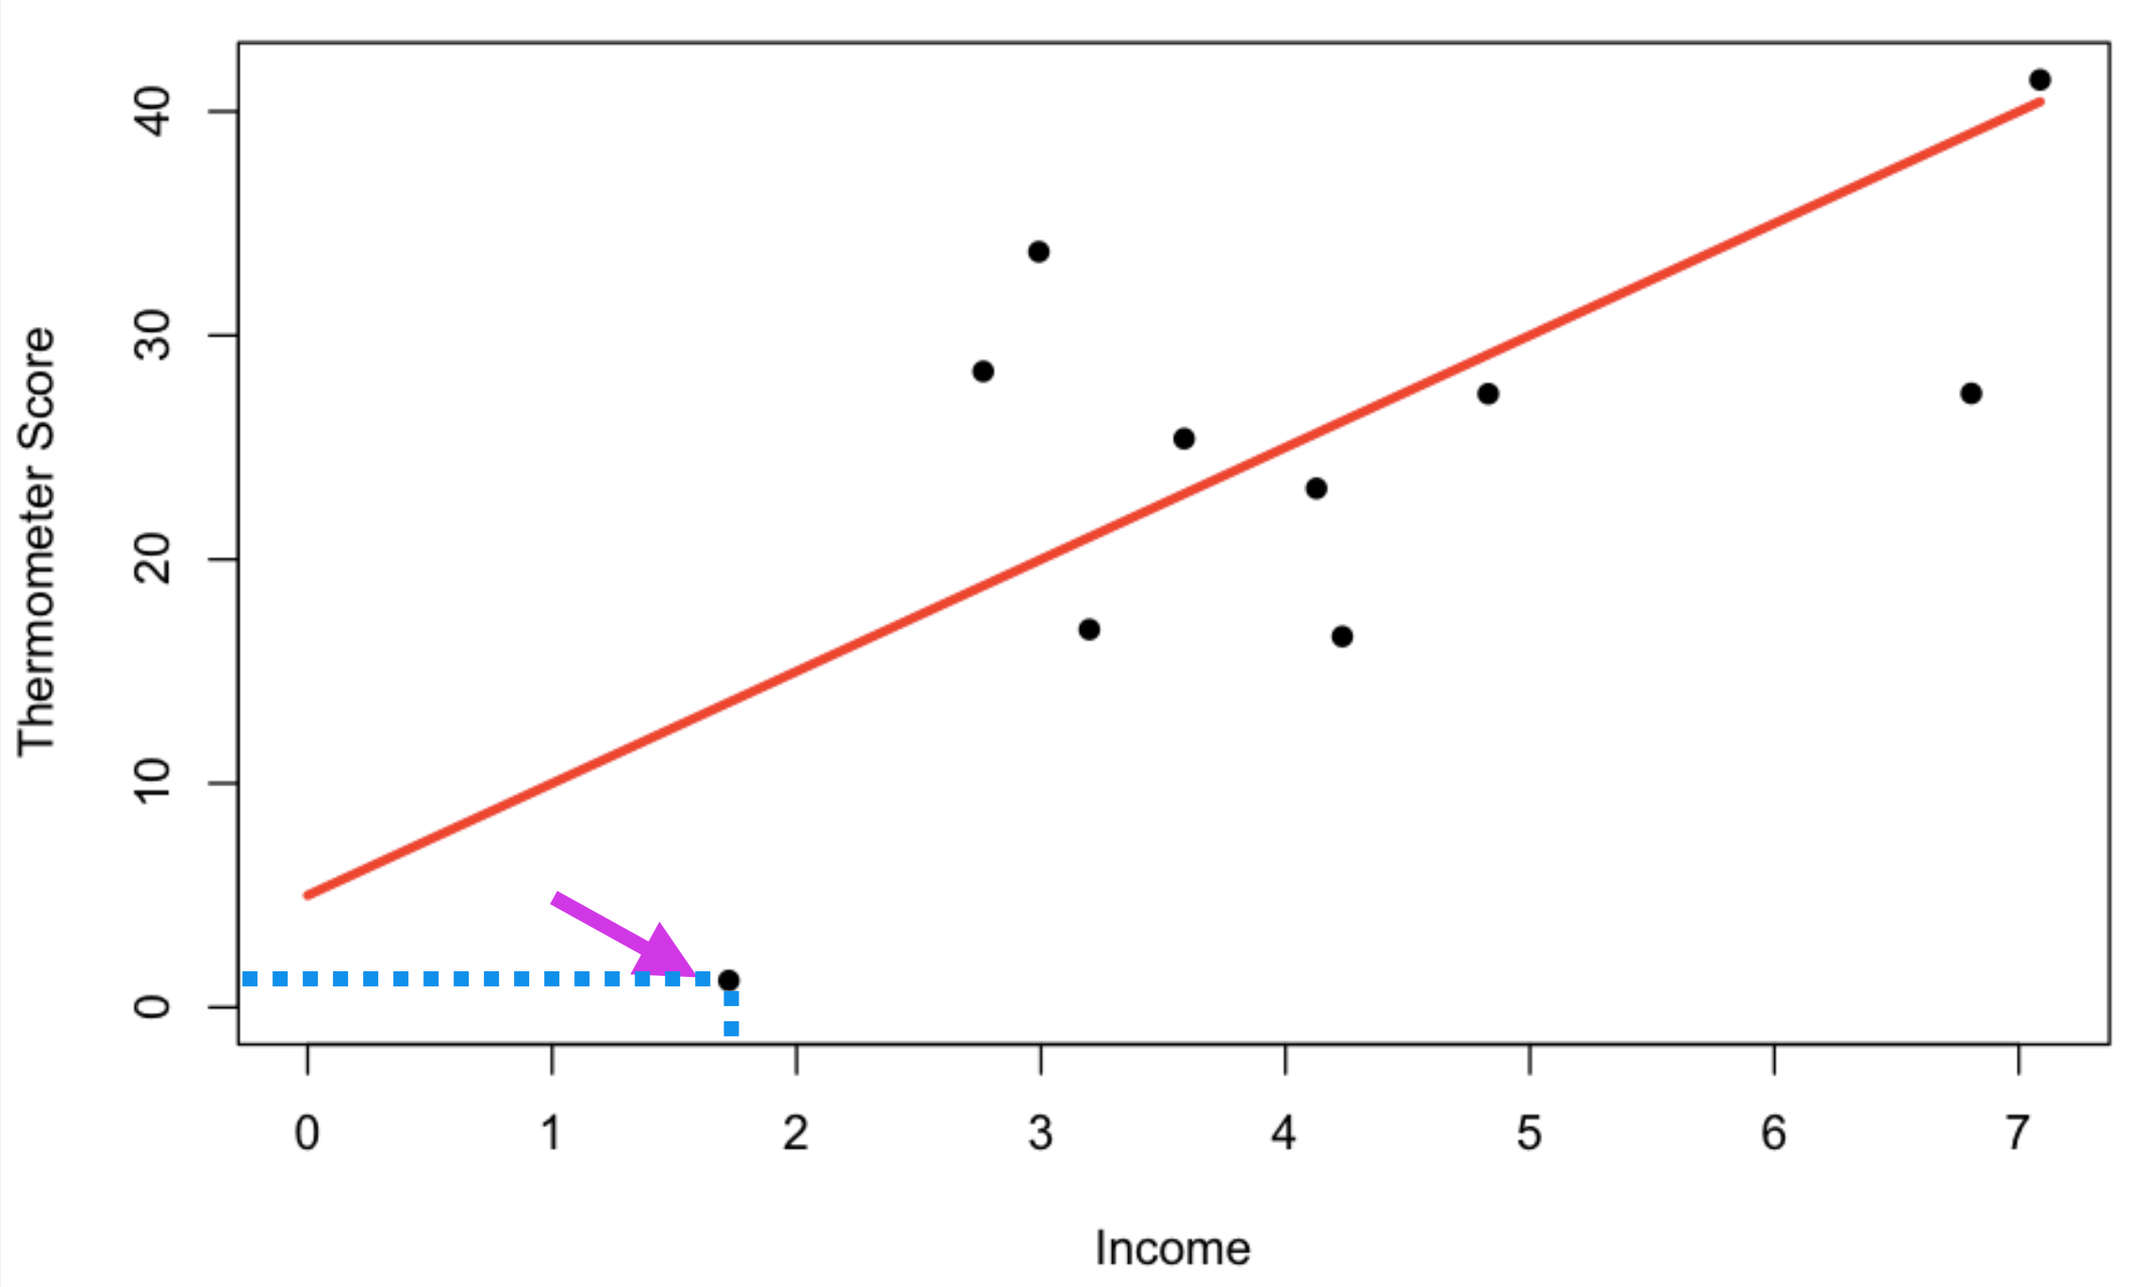

- **Actual y-value: y=1**

HOW TO PICK THE LINE

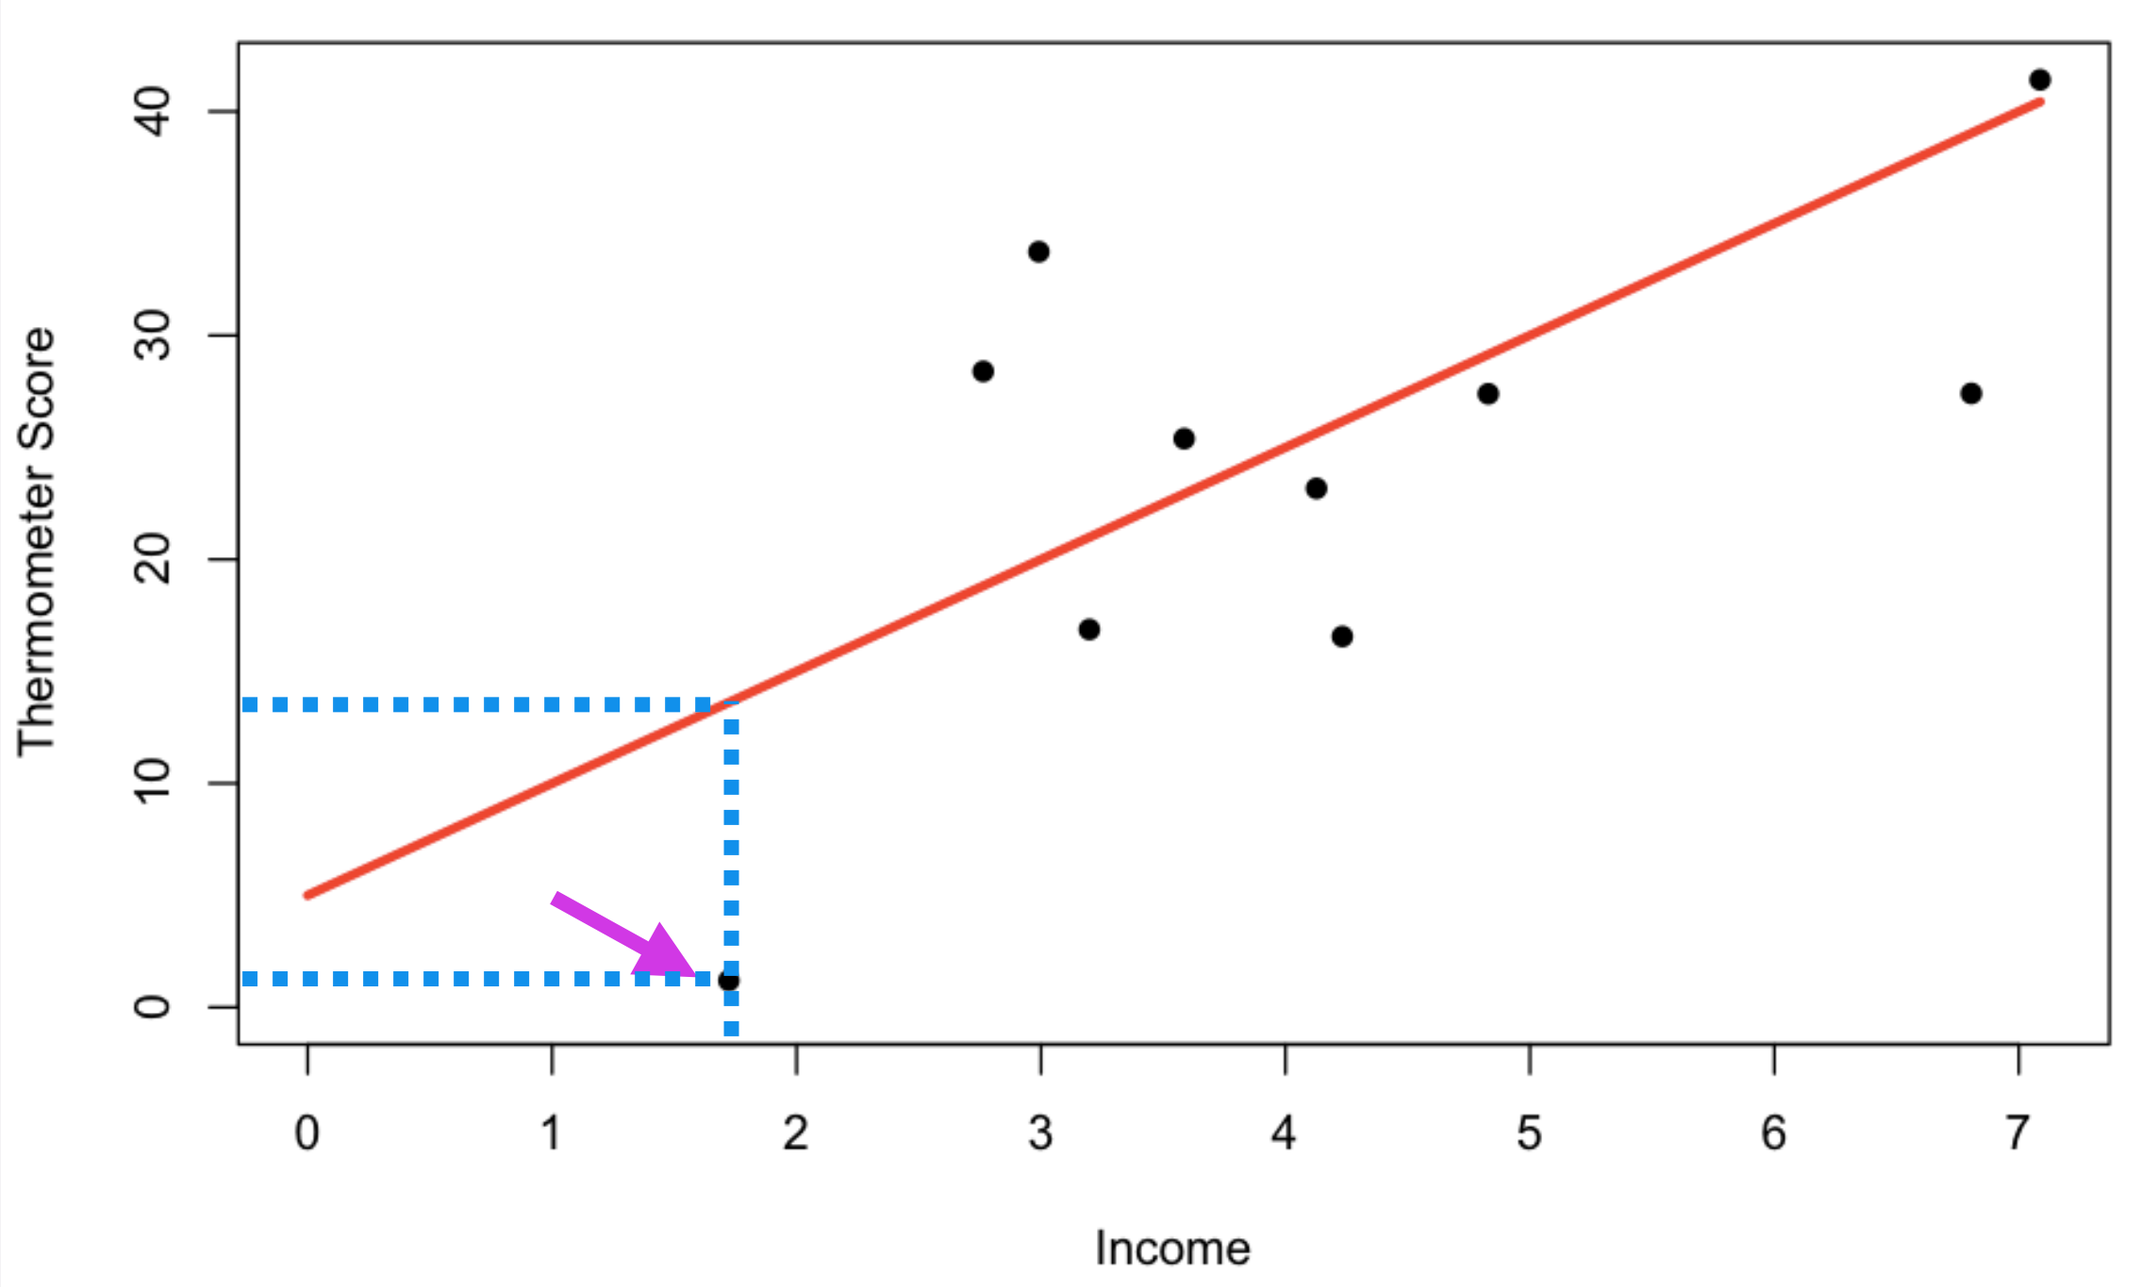

- **Predicted y-value: $\hat{y}=14$**

HOW TO PICK THE LINE

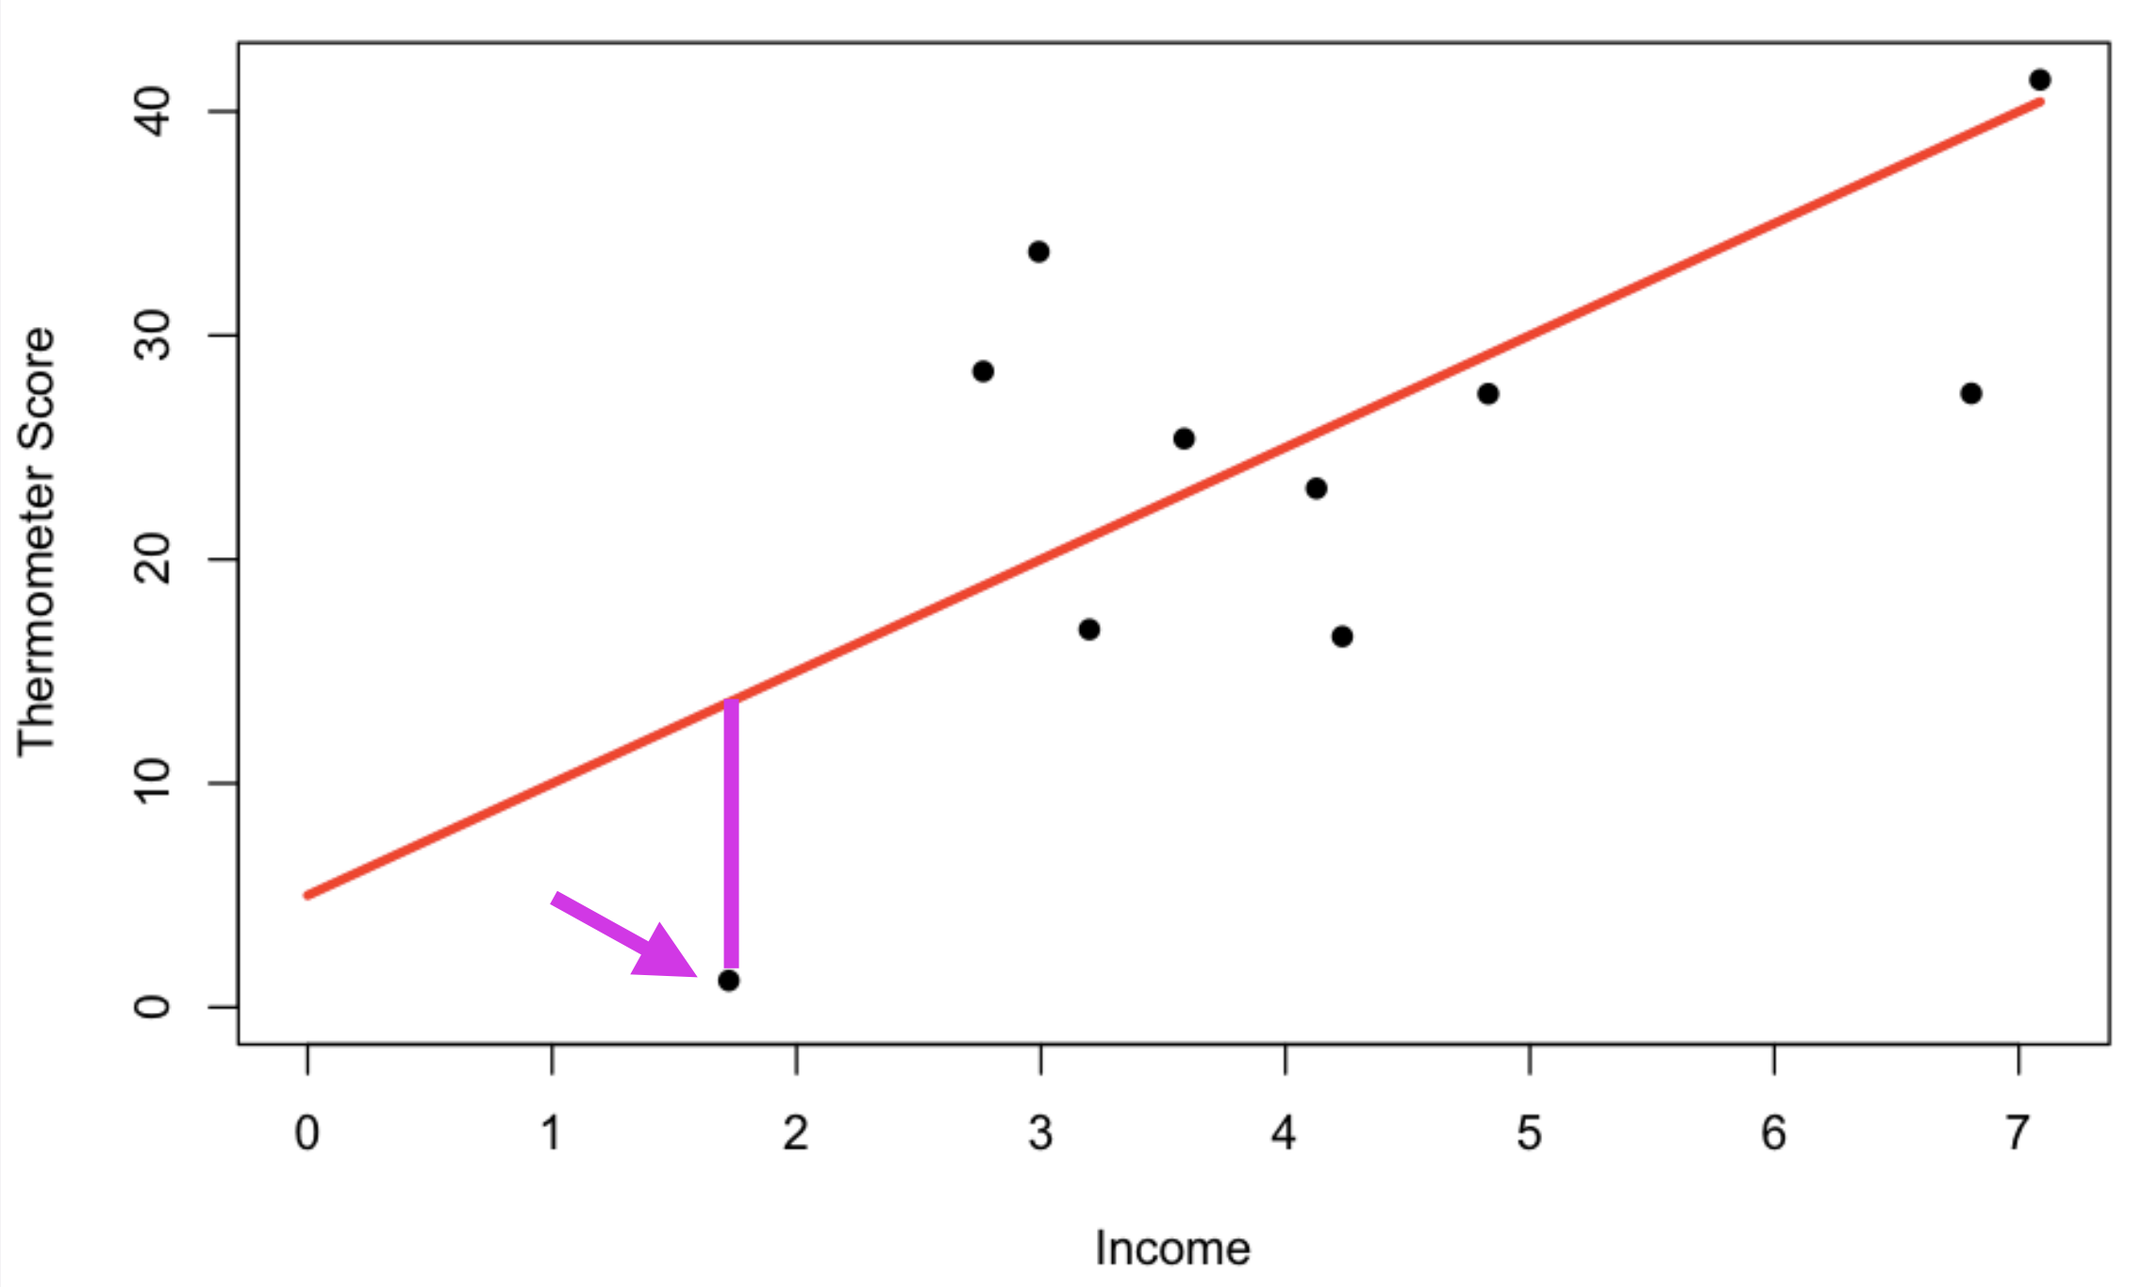

- **Prediction error: $y - \hat{y} = 1 - 14 = -13$**

HOW TO PICK THE LINE

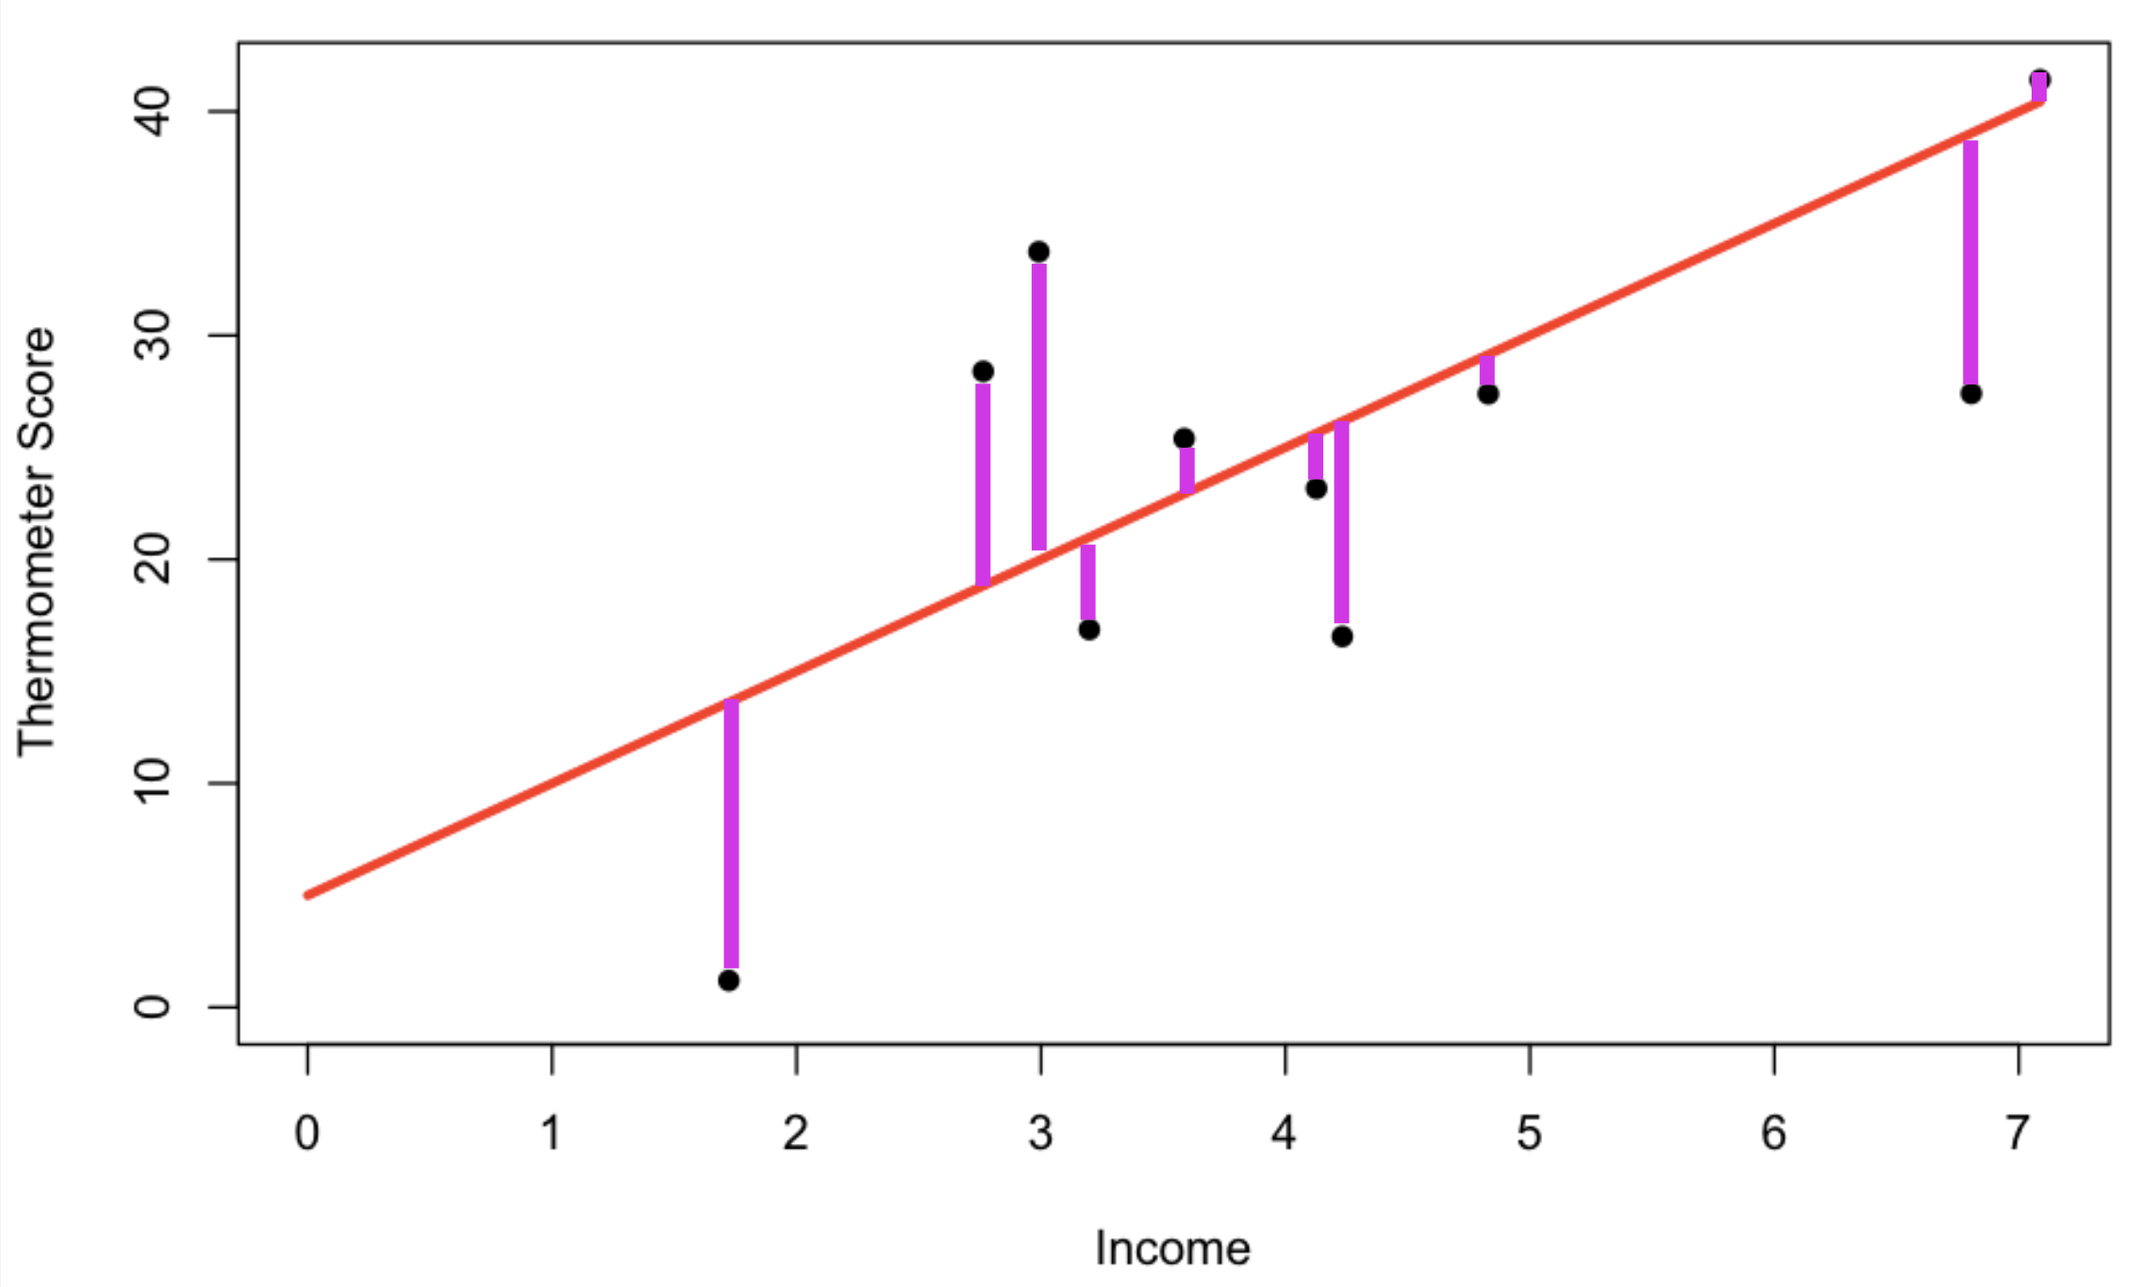

- Get prediction error for each observation

PREDICTION ERROR

- For each observation, we have a prediction error: $y - \hat{y}$
  - Some are positive, some are negative
- We square the prediction errors: $(y - \hat{y})^2$
  - Now all are positive

SQUARED PREDICTION ERROR

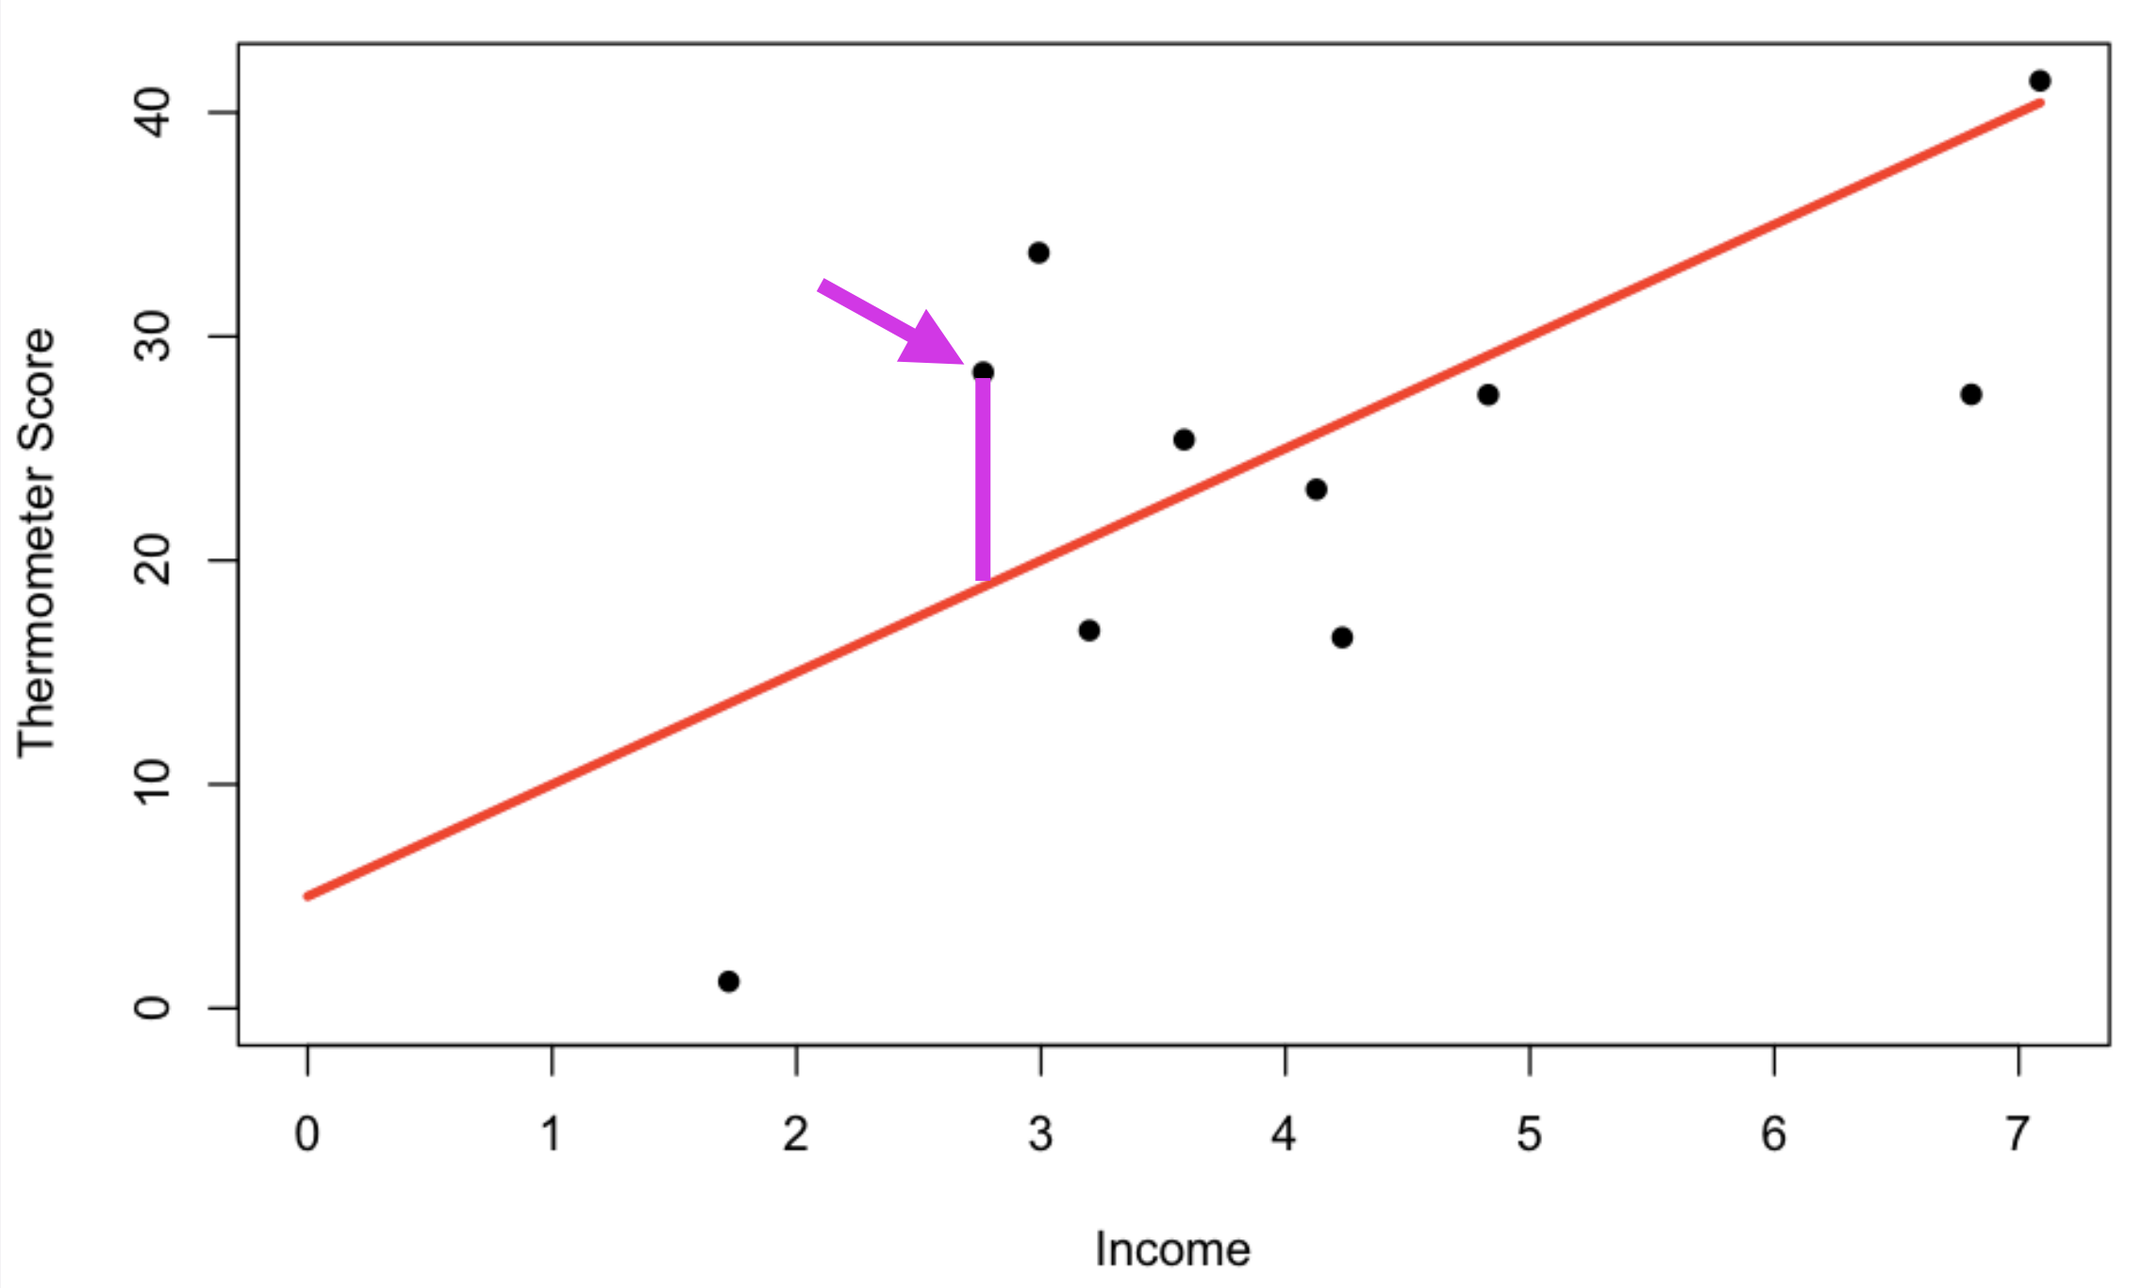

- Prediction error: $y - \hat{y} = 28 - 19 = 9$
- Squared prediction error: $9^2=81$

SQUARED PREDICTION ERROR

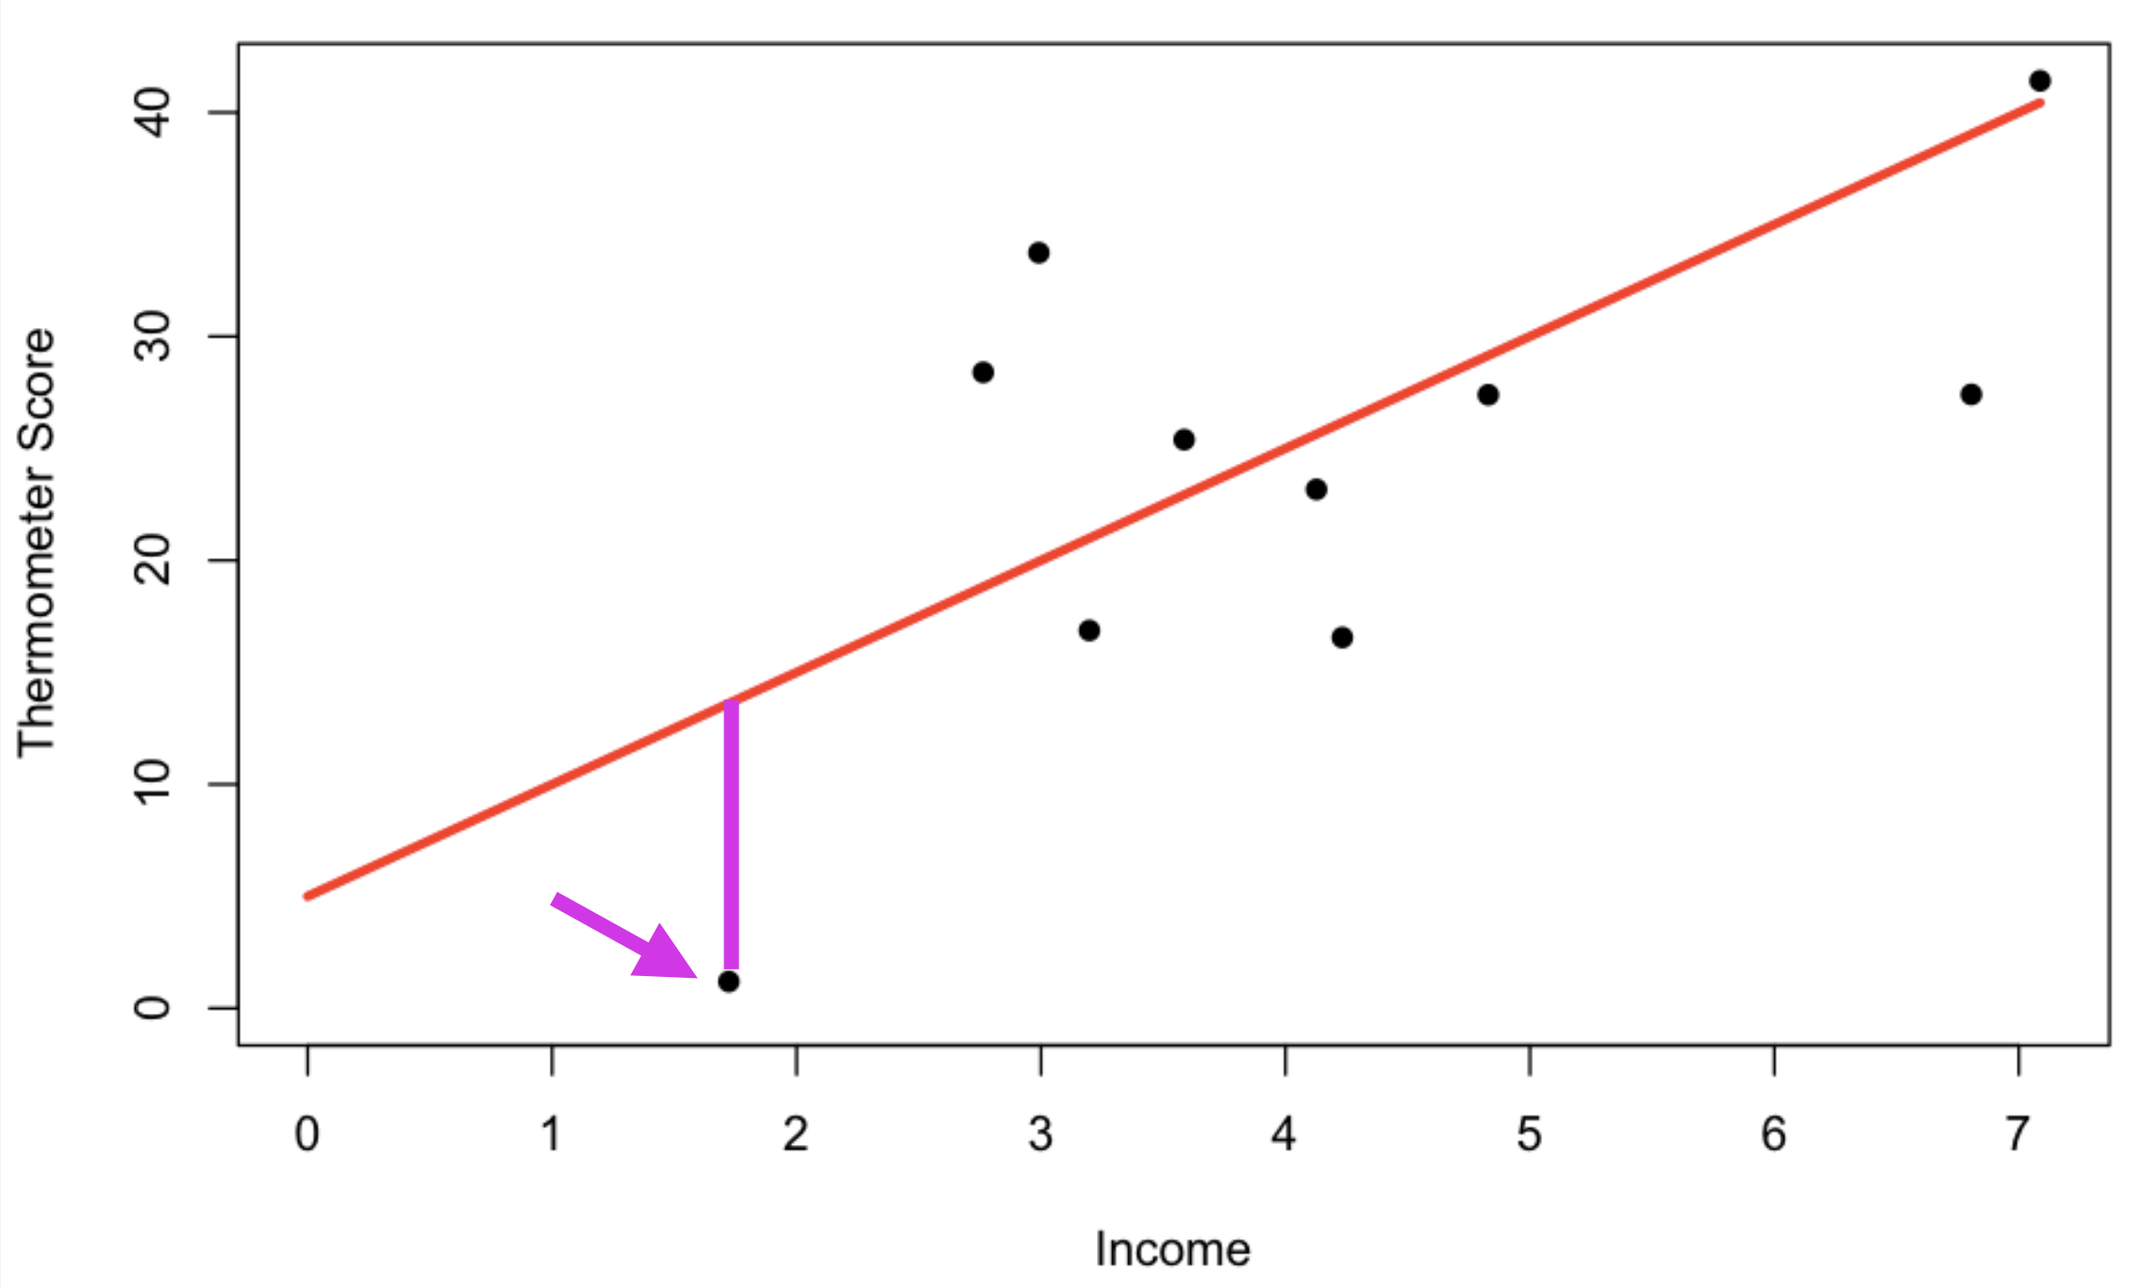

- Prediction error: $y - \hat{y} = 1 - 14 = -13$
- Squared prediction error: $(-13)^2 = 169$

SQUARED PREDICTION ERROR

- We sum squared prediction errors for all observations
- 81 + 169 + all the other observations = 696

SQUARED PREDICTION ERROR

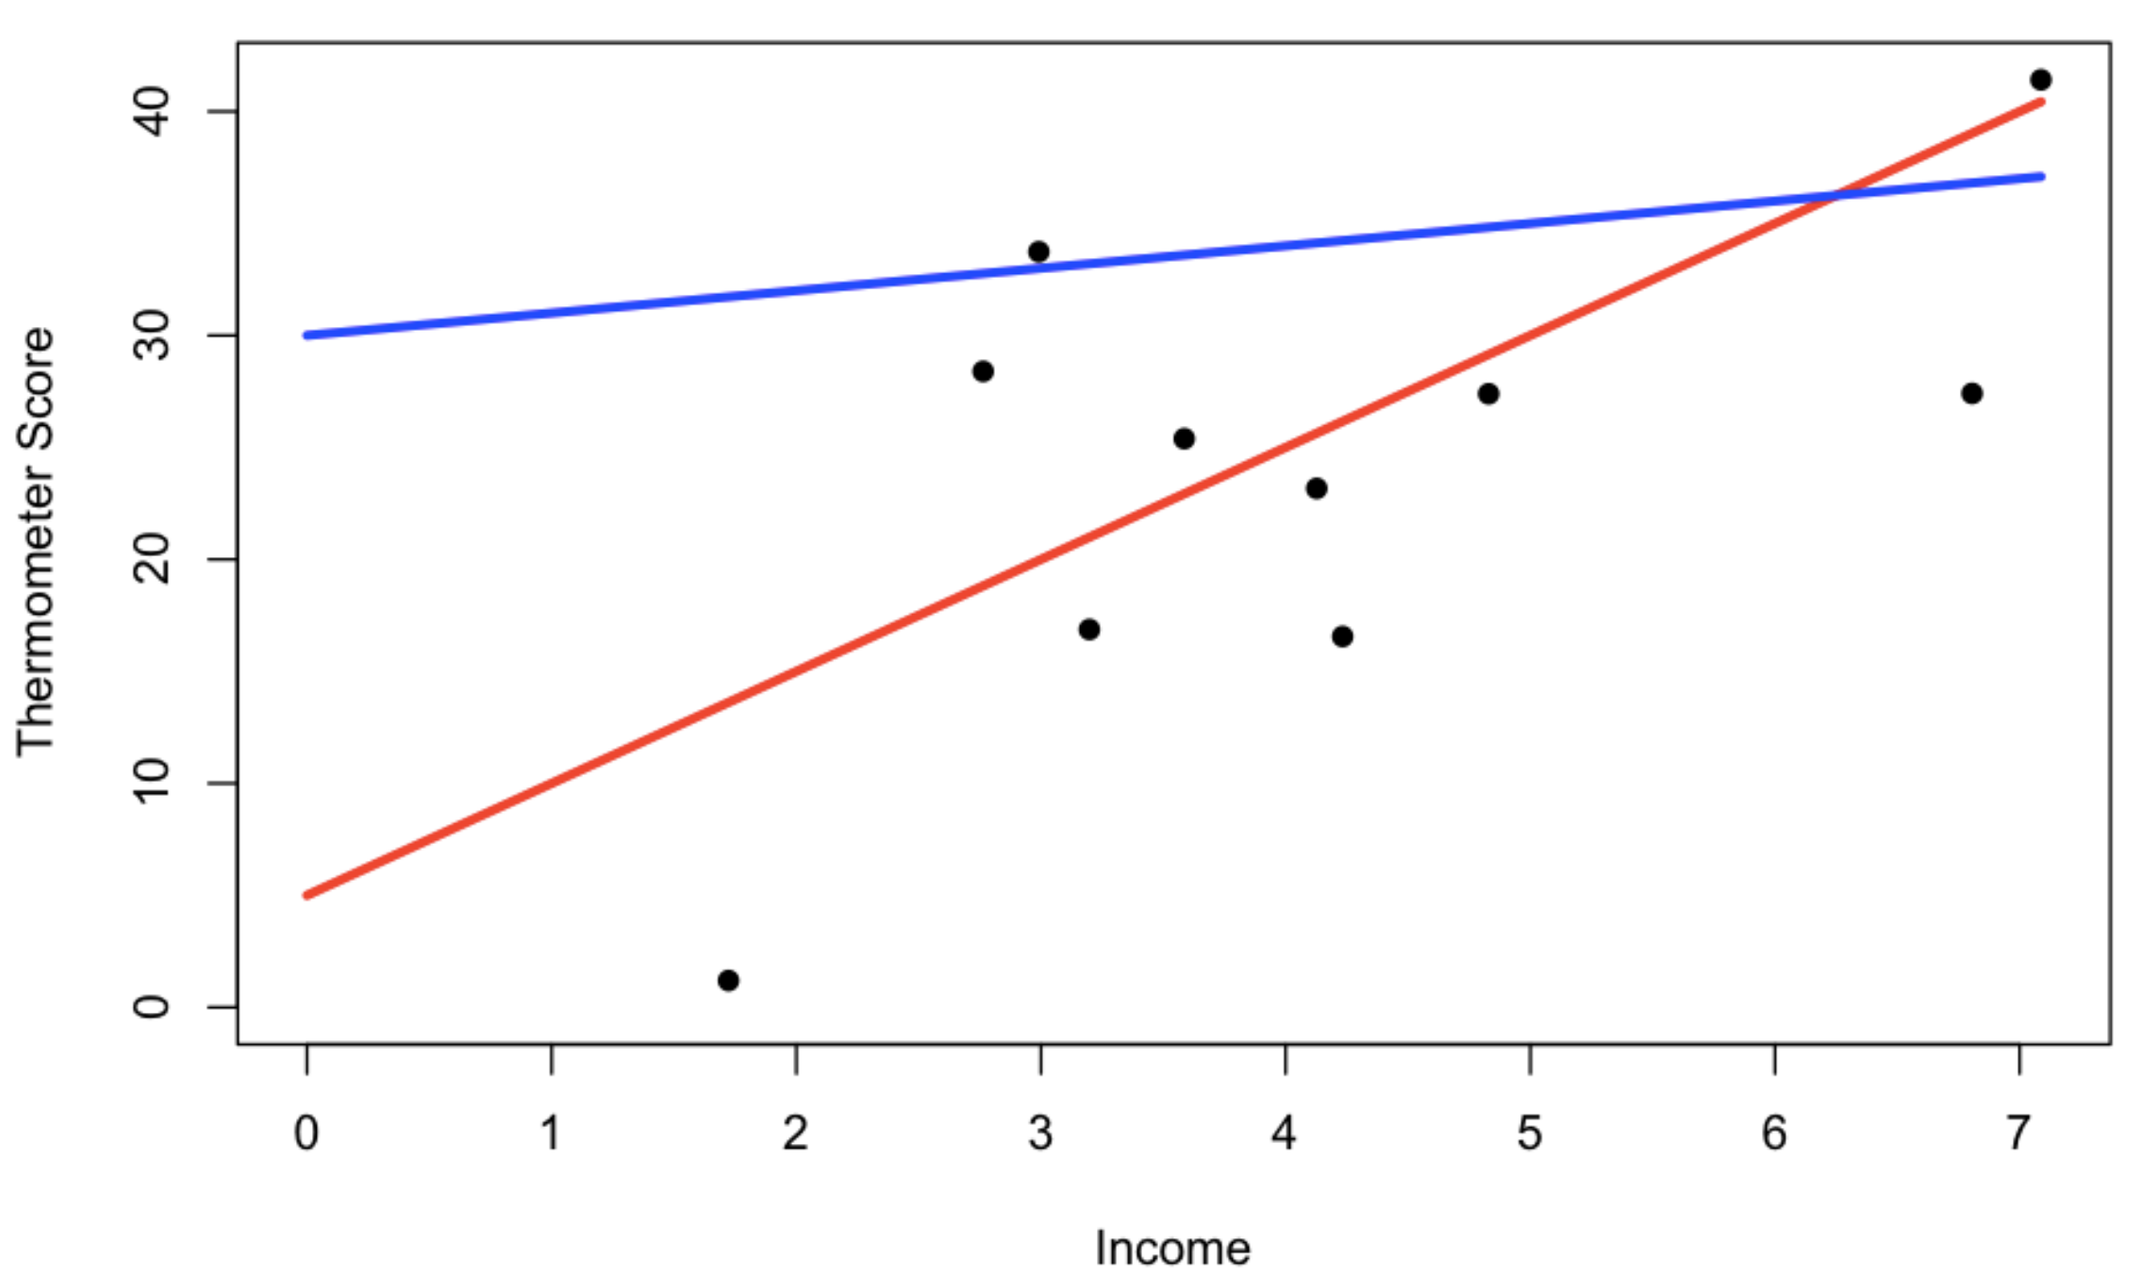

- Sum of squared prediction error red line: 696
- Sum of squared prediction error blue line: 1880

BEST LINE

- The best line is the one with the smallest sum of squared prediction errors
- “Ordinary Least Squares” (OLS) Linear Regression

BEST-FITTING LINE

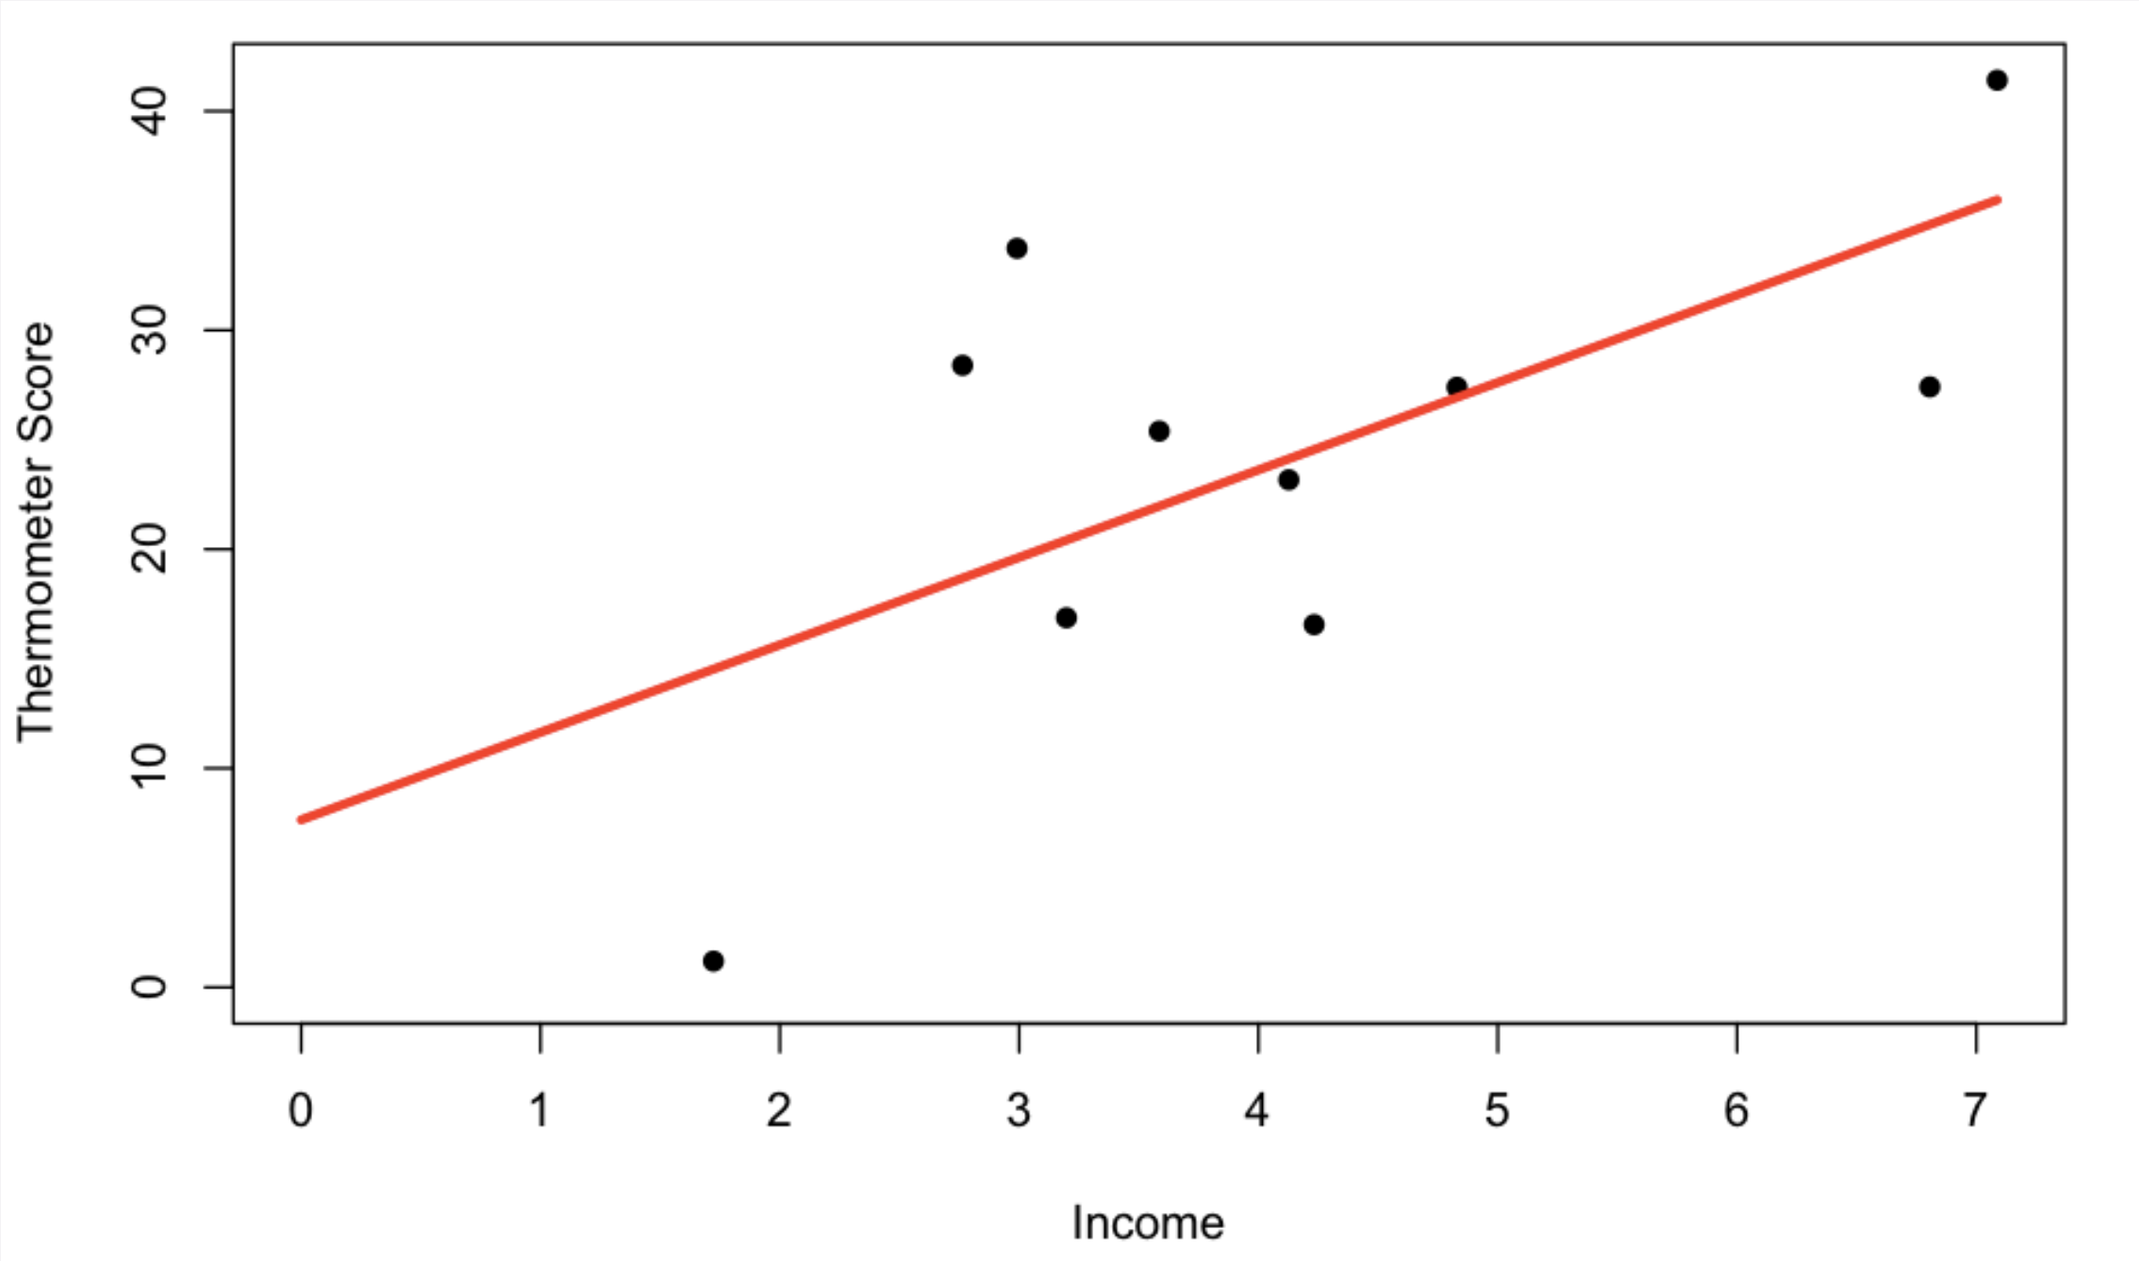

- **Sum of squared prediction errors: 646.3**

FINDINGS THE BEST LINE

- There is a lot of complicated math behind how to find the best line

$$\hat{\beta} = \frac{\sum x_i y_i - \frac{1}{n}\sum x_i \sum y_i}{\sum x_i^2 - \frac{1}{n}(\sum x_i)^2} = \frac{\operatorname{Cov}[x,y]}{\operatorname{Var}[x]}, \quad \hat{\alpha} = \bar{y} - \hat{\beta}\,\bar{x}\,.$$

- Thankfully there are computer programs like R or Stata that do this for us….

BACK TO BIDEN EXAMPLE

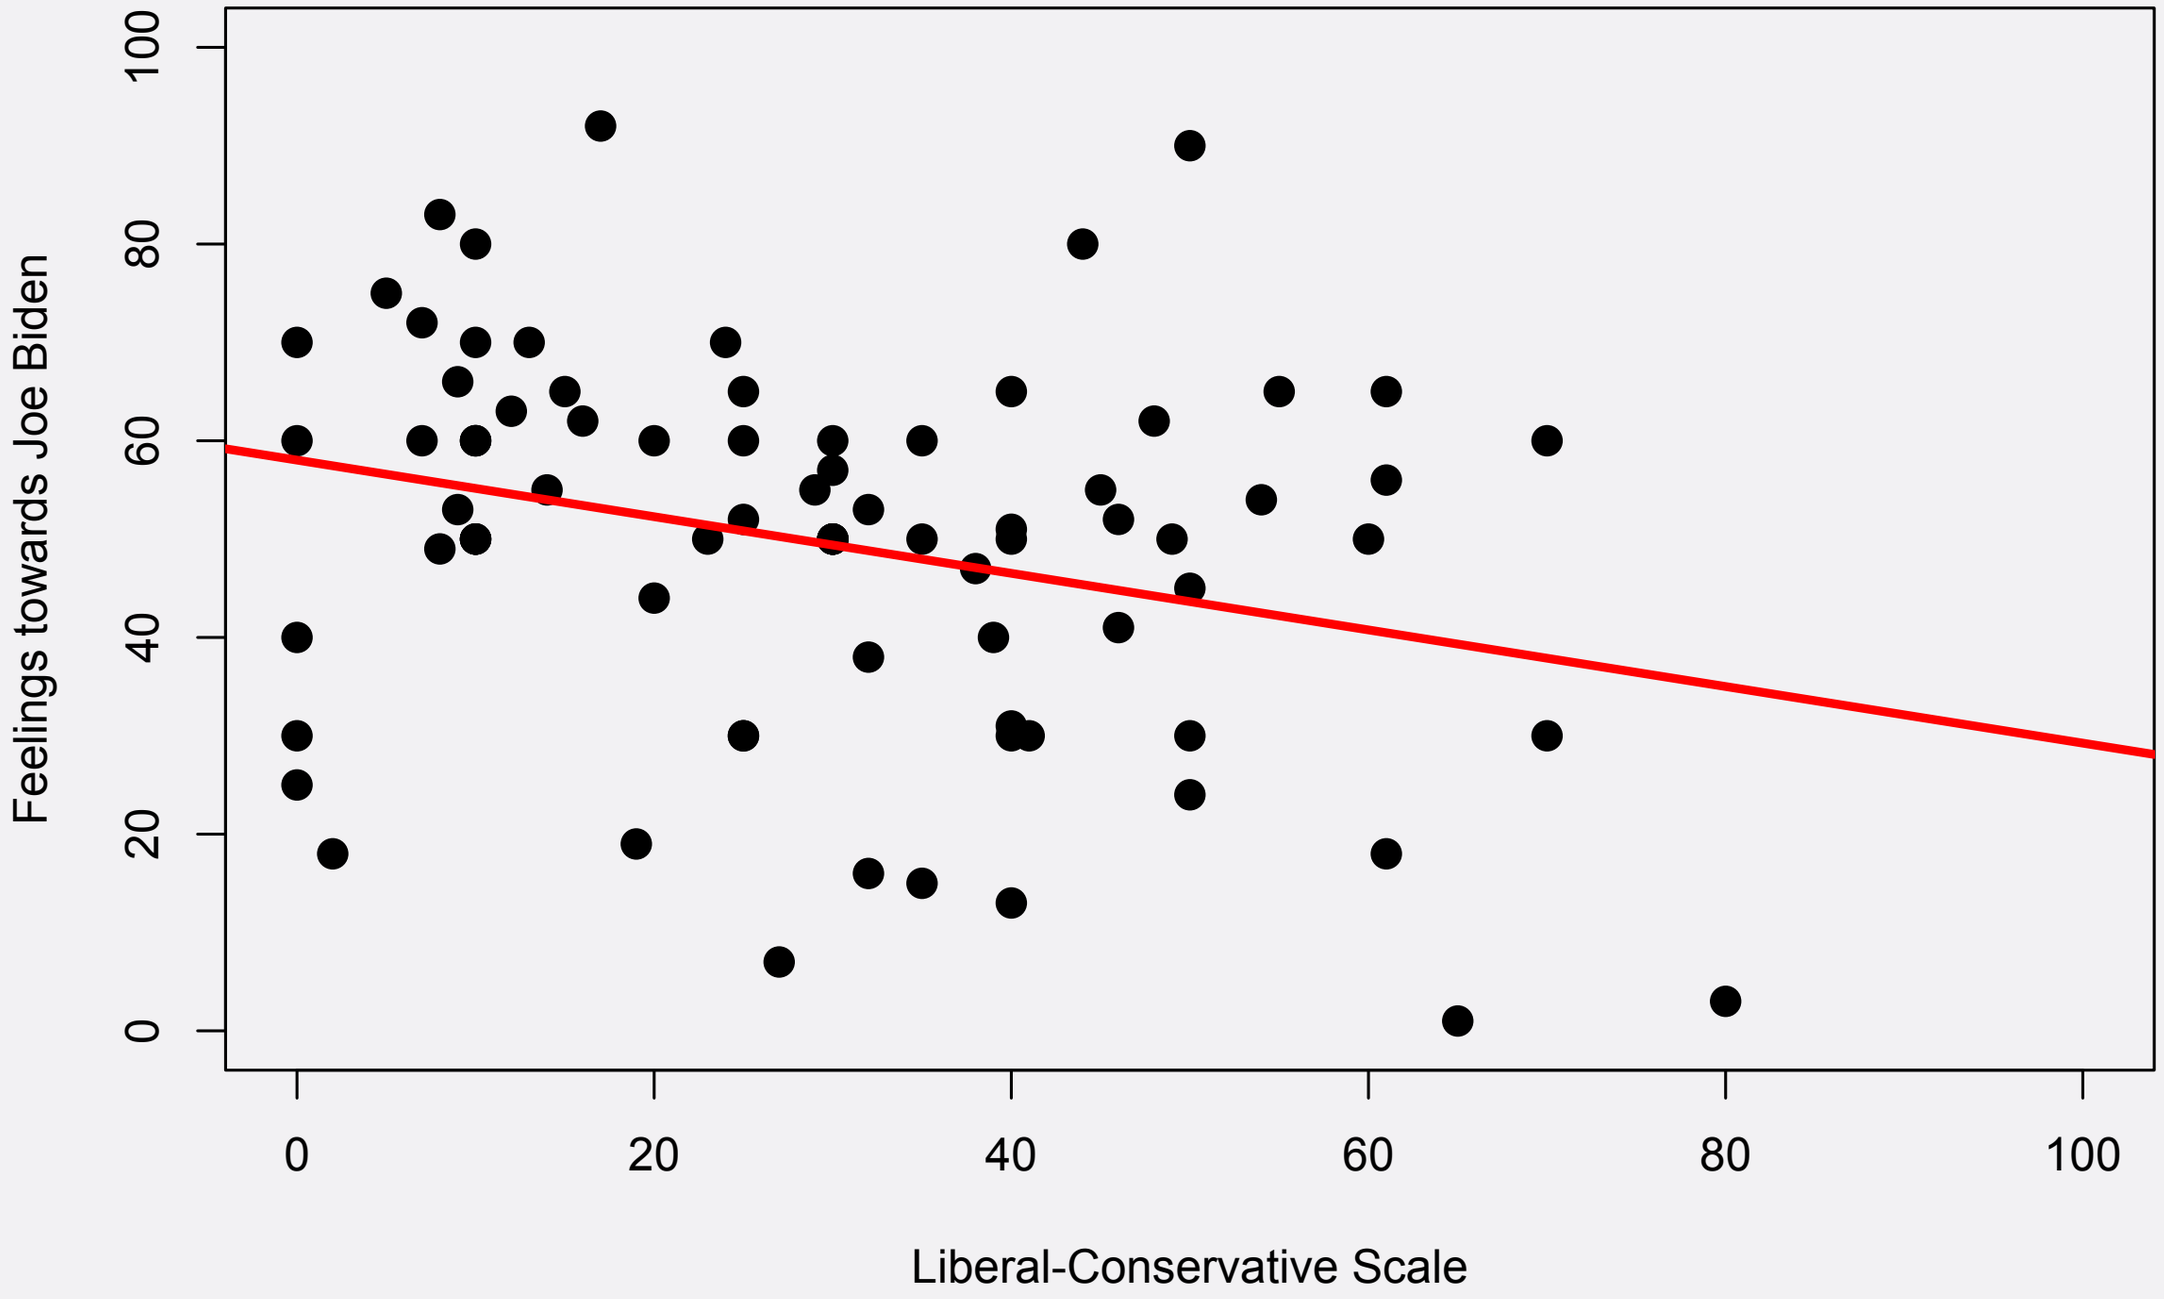

BACK TO OUR EXAMPLE

```
> m <- lm(therm_2 ~ libcons_1, data = data)
> summary(m)

Call:
lm(formula = therm_2 ~ libcons_1, data = data)

Residuals:
    Min      1Q  Median      3Q     Max
-43.261  -8.178   2.005  11.115  46.358

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  58.0308     3.7359  15.533  < 2e-16 ***
libcons_1    -0.2878     0.1065  -2.702  0.00842 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

- DV: Rating of J. Biden (therm_2)
- IV: Liberal-conservative scale (libcons_1)

BACK TO OUR EXAMPLE

Intercept

```
> m <- lm(therm_2 ~ libcons_1, data = data)
> summary(m)

Call:
lm(formula = therm_2 ~ libcons_1, data = data)

Residuals:
    Min      1Q  Median      3Q     Max
-43.261  -8.178   2.005  11.115  46.358

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  58.0308     3.7359  15.533  < 2e-16 ***
libcons_1    -0.2878     0.1065  -2.702  0.00842 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

BACK TO OUR EXAMPLE

```
> m <- lm(therm_2 ~ libcons_1, data = data)
> summary(m)

Call:
lm(formula = therm_2 ~ libcons_1, data = data)

Residuals:
    Min      1Q  Median      3Q     Max
-43.261  -8.178   2.005  11.115  46.358

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  58.0308     3.7359  15.533  < 2e-16 ***
libcons_1    -0.2878     0.1065  -2.702  0.00842 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Slope

BACK TO OUR EXAMPLE

```
> m <- lm(therm_2 ~ libcons_1, data = data)
> summary(m)

Call:
lm(formula = therm_2 ~ libcons_1, data = data)

Residuals:
    Min      1Q  Median      3Q     Max
-43.261  -8.178   2.005  11.115  46.358

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  58.0308     3.7359  15.533  < 2e-16 ***
libcons_1    -0.2878     0.1065  -2.702  0.00842 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

Slope

- Thermometer Score = 58.0 - 0.29 * Lib/Cons
- (I simplified numbers earlier to make math easier…)

NOW

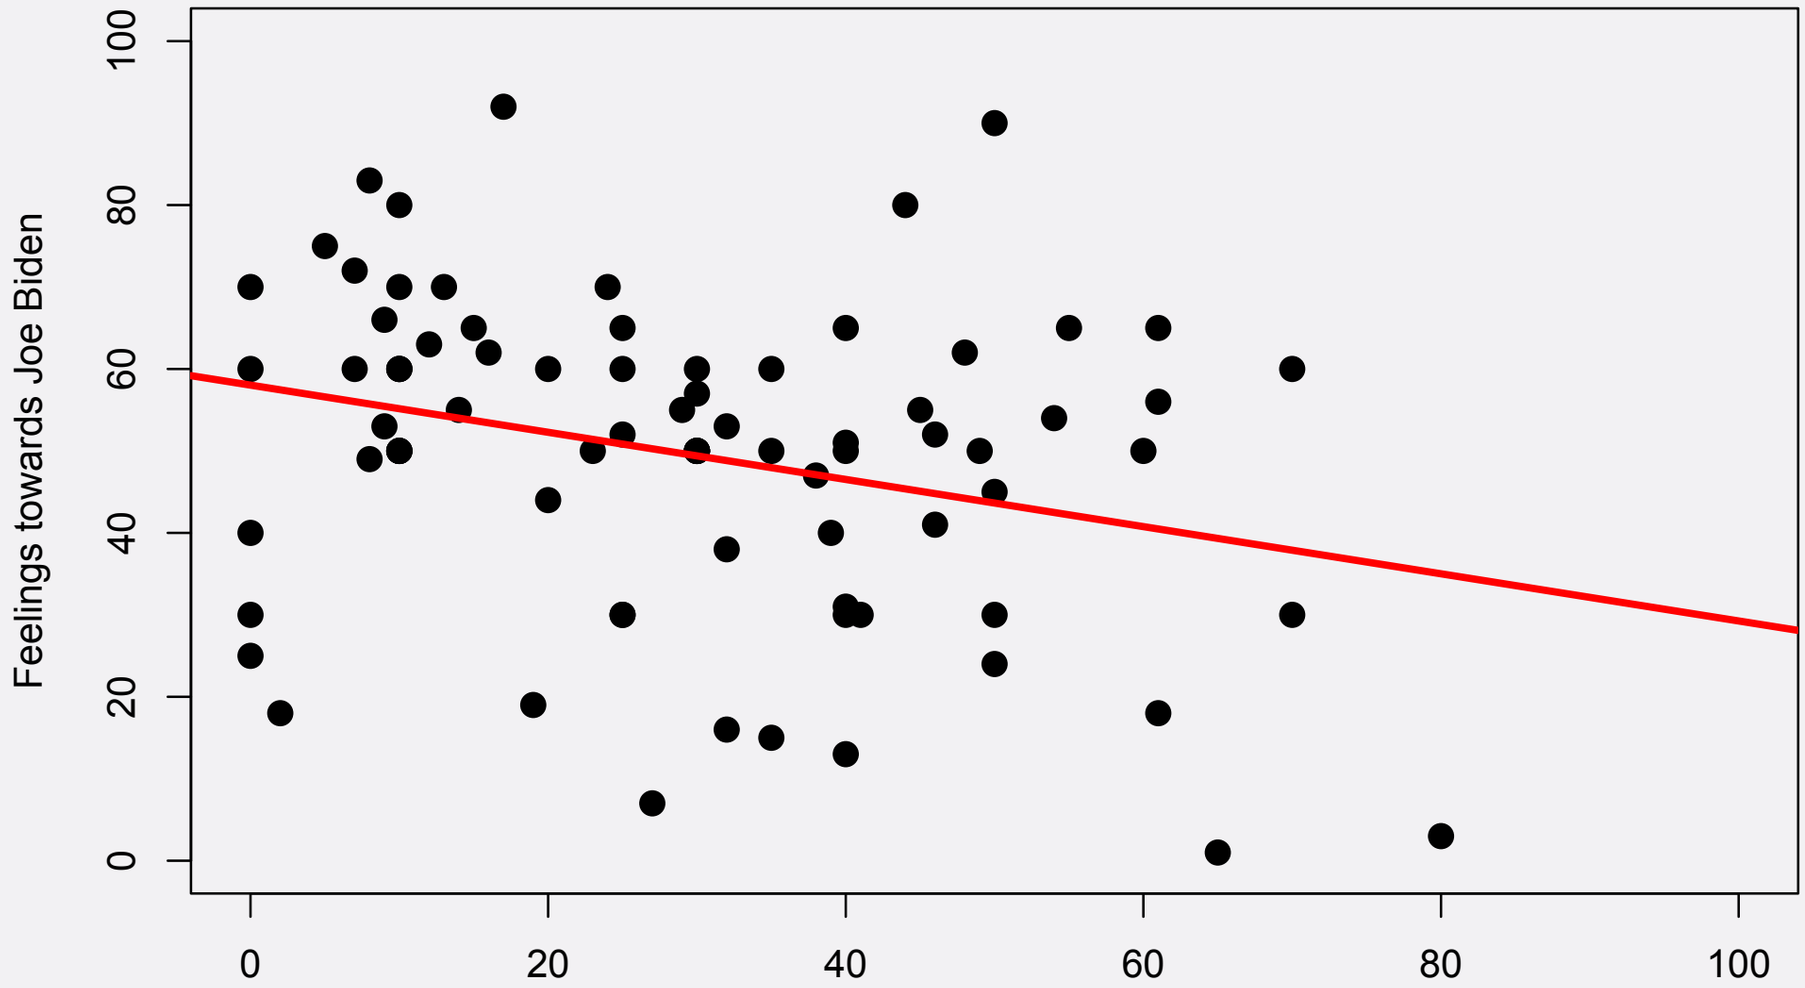

- Is this effect real?
- Or is this just something we found in our sample, but lib/cons actually has no effect on perceptions of Biden in the population?

TODAY

- Finishing up bivariate hypothesis testing
- Hypothesis testing with samples

REMEMBER

POLITICS JANUARY 25, 2023

Biden Averaged 41% Job Approval in His Second Year

Results for this Gallup poll are based on telephone interviews conducted Jan. 2-22, 2023, with a random sample of 1,011 adults, aged 18 and older, living in all 50 U.S. states and the District of Columbia. For results based on the total sample of national adults, the margin of sampling error is ±4 percentage points at the 95% confidence level. All reported margins of sampling error include computed design effects for weighting.

BIVARIATE RELATIONSHIP

Gender → Approval for J. Biden

- Hypothesis: In a comparison of individuals, women are more likely to approve of J. Biden than men
  - "gender gap"

BIVARIATE RELATIONSHIP

Biden Approval Ratings Diverge by Gender, Education, Race

Job Approval Ratings of President Biden, by Subgroup

	Approve	Disapprove	N
	%	%	
All U.S. adults	56	39	2,937
Gender			
Men	49	45	1,643
Women	62	34	1,294

PROBLEM

- Is the effect of gender on approval real?
  - Does it exist in the population?
  - Maybe gender actually has no effect in the population, and we just found one *by chance* in this sample?

PROBLEM

- We have a *random sample*
  - Men: 49% approval
  - Women: 62% approval
- Want to know: is mean approval rating of men and women in the *population* the same or not?

ALTERNATIVE HYPOTHESIS

- There *is* a relationship between the independent and dependent variable in the population

- H_A or H_1

NULL HYPOTHESIS

- In the population, there is *no relationship* between dependent and independent variable
  - If there is a difference in the sample, it is due to random sampling error

- H_0

IN OUR CASE

- H_0: In a comparison of individuals, there is *no difference* between men and women in approval of J. Biden
- H_A: In a comparison of individuals, there *is a difference* between men and women in approval of J. Biden

BACK TO MISTAKES

- Idea: Use relation between two variables in *sample* to make inference about relation between two variables in *population*
- Of course, means we can make mistakes

ERRORS

	There Is A Relation In The Population	There Is No Relation In The Population
We Conclude There Is A Relation	✔	✘
We Conclude There Is No Relation	✘	✔

ERRORS

	There Is A Relation In The Population	There Is No Relation In The Population
We Conclude There Is A Relation	✔	✘
We Conclude There Is No Relation	✘	✔

TYPE I ERROR

- We conclude there is a relationship between X and Y when in reality there is not
- Example: There is no difference between men and women in approval rating in the population, but we conclude that there is

TYPE I ERROR

- We conclude there is a relationship between X and Y when in reality there is not
- “Type I error”
- We falsely reject H_0

ERRORS

	There Is A Relation In The Population	There Is No Relation In The Population
We Conclude There Is A Relation	✔	✘
We Conclude There Is No Relation	✘	✔

TYPE II ERROR

- We conclude there is no relationship between X and Y when in reality there is
- Example: There is a difference between men and women in approval rating in the population, but we conclude that there is none

TYPE II ERROR

- We conclude there is no relationship between X and Y when in reality there is
  - "Type II error"
  - We falsely do not reject H_0

ERRORS

	There Is A Relation In The Population	There Is No Relation In The Population
We Conclude There Is A Relation	✔	✘ Type I
We Conclude There Is No Relation	✘ Type II	✔

DECISION

- It's bad if we conclude there is a relationship when in reality there is not (Type I error)
  - Type II error is also not great, but not as bad
- We privilege H_0

DECISION

- By default: We start out with assumption that there is no relationship in population (so H_0 is true)
- No difference between men and women in Biden approval in population

DECISION

- Ask: Is there enough evidence in the *sample* to reject H_0?
- Is the observed difference between mean and women in *sample* large enough to reject null hypothesis that no difference between them in population?

DECISION

Job Approval Ratings of President Biden, by Subgroup

	Approve	Disapprove	N
	%	%	
All U.S. adults	56	39	2,937
Gender			
Men	49	45	1,643
Women	62	34	1,294

- **The larger the difference in approval ratings between men and women in our samples, the less likely it is that the mean in the population is the same**

P-VALUE

- Q: When do we decide that we have "enough" evidence?
- A: When the chance of falsely rejecting H_0 is 5% or less
  - Equivalent: Change of Type I error less than 5%
  - Probability of falsely rejecting H_0 is called the "p-value"