Do Institutions Cause Prosperity? Settler Mortality as an Instrument for Institutional Quality

Example Paper — Gov 51: Data Analysis and Politics

2026-01-01

1 Introduction

Why are some countries rich and others poor? This question has occupied economists and political scientists for centuries, and the answers matter: if poverty is caused by geography, then policy can do little. If it is caused by institutions, then reform is possible.

In 2001, Daron Acemoglu, Simon Johnson, and James Robinson published a paper that changed how economists think about this question. Their argument was simple: the quality of institutions — in particular, how well a country protects private property rights and limits government expropriation — is a fundamental cause of long-run economic prosperity. Countries with strong institutions tend to be rich; countries with weak institutions tend to be poor.

But this correlation does not identify causation. Rich countries can *afford* better institutions; countries with strong institutions may also have other advantages (geography, climate, human capital) that drive prosperity independently. The relationship between institutions and income is plagued by reverse causation and omitted variable bias. Any regression of income on institutional quality is likely to be biased.

Acemoglu, Johnson, and Robinson (AJR) proposed an elegant solution: an instrumental variable. They used **settler mortality** — the death rates faced by European colonizers in the 19th century — as an instrument for institutional quality today. Their logic:

1. Where European settlers died at high rates (from malaria, yellow fever), colonizers extracted resources rather than settling. They built exploitative institutions.
2. Where settlers survived in large numbers, they built inclusive institutions modeled on their home countries.
3. These institutions persisted: the "reversal of fortune" is documented in their data.
4. Settler mortality affects income today *only through* the institutions it shaped — not through any direct path.

This paper replicates AJR's core analysis, explains the instrumental variables strategy in detail, and interprets the results.

2 Data and Measurement

2.1 The AJR Dataset

```r
ajr <- read_csv("data/ajr2001.csv", show_col_types = FALSE)

glimpse(ajr)
```

```
Rows: 64
Columns: 8
$ country  <chr> "Angola", "Argentina", "Australia", "Burkina Faso", "Banglade~
$ shortnam <chr> "AGO", "ARG", "AUS", "BFA", "BGD", "BHS", "BOL", "BRA", "CAN"~
$ africa   <dbl> 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 1, 1, 1, 0, 0, 0, 1, 0, 1, 1, 1~
$ latitude <dbl> 0.13666667, 0.37777779, 0.30000001, 0.14444445, 0.26666668, 0~
$ euro1900 <dbl> 0, 90, 99, 0, 0, 10, 30, 55, 98, 50, 0, 0, 0, 25, 20, 25, 0, ~
$ avexpr   <dbl> 5.363636, 6.386364, 9.318182, 4.454545, 5.136364, 7.500000, 5~
$ logpgp95 <dbl> 7.770645, 9.133459, 9.897972, 6.845880, 6.877296, 9.285448, 7~
$ logem4   <dbl> 5.634789, 4.262680, 2.145931, 5.634789, 4.268438, 4.442651, 4~
```

The dataset contains 64 former European colonies with data on three key variables:

- **`logpgp95`**: Log GDP per capita in 1995 (outcome, Y). The log transformation compresses the enormous variation in income across countries.
- **`avexpr`**: Average protection against expropriation risk, scored 0-10 (endogenous treatment, D). Higher values indicate stronger property rights and rule of law. This is AJR's measure of "institutional quality."
- **`logem4`**: Log settler mortality, measured in deaths per thousand per year among European military and naval personnel, clergy, and laborers in the 19th century (instrument, Z).

The dataset also contains `africa` (indicator for sub-Saharan Africa) and `latitude` (absolute latitude), which AJR use as controls in robustness checks.

2.2 Descriptive Statistics

```r
ajr |>
  summarise(
    `N countries`                = n(),
    `Mean log GDP (logpgp95)`    = round(mean(logpgp95, na.rm = TRUE), 2),
    `SD log GDP`                 = round(sd(logpgp95, na.rm = TRUE), 2),
    `Mean institutions (avexpr)` = round(mean(avexpr, na.rm = TRUE), 2),
    `SD institutions`            = round(sd(avexpr, na.rm = TRUE), 2),
    `Mean log settler mortality` = round(mean(logem4, na.rm = TRUE), 2),
    `SD log settler mortality`   = round(sd(logem4, na.rm = TRUE), 2)
  ) |>
  pivot_longer(everything(), names_to = "Variable", values_to = "Value") |>
  kable(caption = "Descriptive statistics, AJR (2001) dataset (N = 64 former colonies)")
```

Table 1: **?(caption)**

(a) Descriptive statistics, AJR (2001) dataset (N = 64 former colonies)

Variable	Value
N countries	64.00
Mean log GDP (logpgp95)	8.06
SD log GDP	1.04
Mean institutions (avexpr)	6.52
SD institutions	1.47
Mean log settler mortality	4.64
SD log settler mortality	1.25

```r
# Drop rows with missing values on key variables
ajr_clean <- ajr |>
  filter(!is.na(logem4), !is.na(avexpr), !is.na(logpgp95))

cat("Complete cases:", nrow(ajr_clean), "\n")
```

```
Complete cases: 64
```

3 Why OLS Fails

The obvious starting point is a regression of log income on institutional quality:

$$\log(\text{GDP})_i = \alpha + \beta \cdot \text{institutions}_i + \varepsilon_i$$

```r
ols_fit <- lm_robust(logpgp95 ~ avexpr, data = ajr_clean)

pred_line <- tibble(
  avexpr   = seq(min(ajr_clean$avexpr), max(ajr_clean$avexpr), length.out = 200),
  logpgp95 = coef(ols_fit)["(Intercept)"] + coef(ols_fit)["avexpr"] * avexpr
)

ggplot(ajr_clean, aes(x = avexpr, y = logpgp95)) +
  geom_point(color = "#A51C30", size = 2.5, alpha = 0.75) +
  geom_text(aes(label = shortnam), size = 2.2, hjust = -0.2, alpha = 0.6,
            check_overlap = TRUE) +
  geom_line(data = pred_line, color = "#1E3C72", linewidth = 1.0) +
  annotate("text", x = 8.5, y = 7.3,
           label = sprintf("OLS slope = %.3f\n(SE = %.3f)",
                           coef(ols_fit)["avexpr"],
                           ols_fit$std.error["avexpr"]),
           color = "#1E3C72", size = 3.2, hjust = 0) +
  labs(
```

```
  title    = "Strong Correlation Between Institutions and Income...",
  subtitle = "...but is this causal?",
  x        = "Institutional quality (avg. protection against expropriation, 0-10)",
  y        = "Log GDP per capita, 1995",
  caption  = "Source: Acemoglu, Johnson & Robinson (2001)"
) +
theme_minimal(base_size = 11) +
theme(panel.grid.minor = element_blank())
```

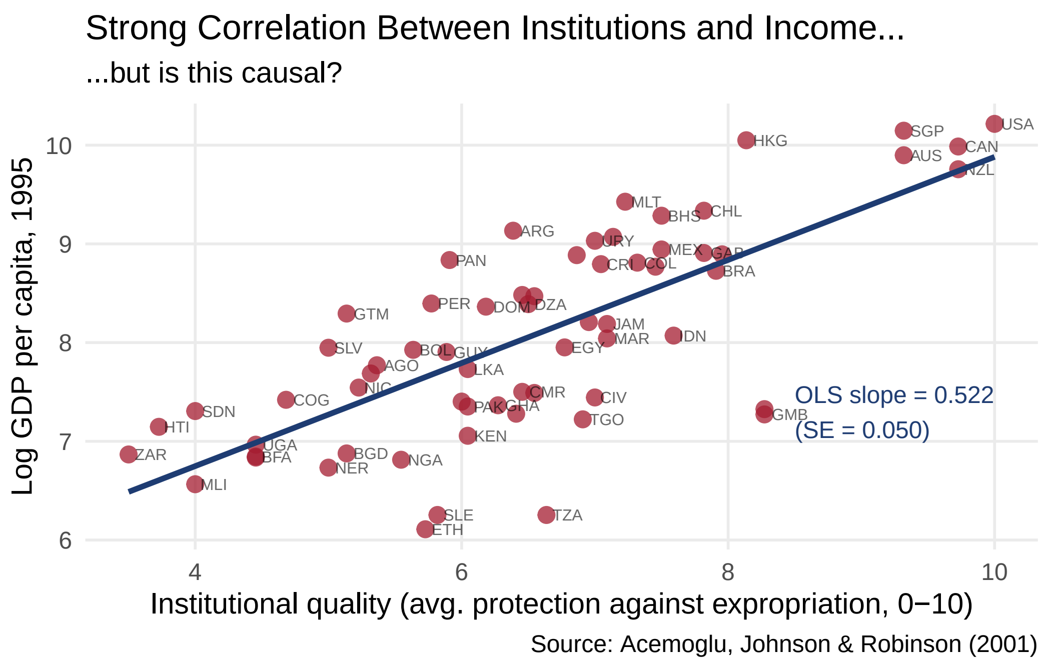

Figure 1: OLS relationship between institutional quality and log GDP per capita. Each point is a country. The upward slope reflects the strong positive correlation — but correlation is not causation.

The OLS estimate of $\hat{\beta} = 0.522$ says: a one-unit improvement in institutional quality is associated with a 0.52-log-point increase in income. This is a large effect — it implies that moving from the 25th to the 75th percentile of institutional quality (a 2.5-unit shift) would multiply income by roughly $e^{2.5\times 0.52} \approx 3.7$.

But this estimate is almost certainly biased. Three threats to causal identification are severe:

1. **Reverse causation**: Rich countries can invest more in institutional infrastructure. Income $\rightarrow$ institutions as much as institutions $\rightarrow$ income.
2. **Omitted variables**: Geography (latitude, disease environment), human capital, culture, and colonial legacy all affect both institutions and income.
3. **Measurement error in institutions**: The `avexpr` variable is measured in the 1980s–90s. Countries that are prosperous may score higher on expropriation protection because their prosperity gives them the resources to maintain strong rule of law.

All three biases push the OLS estimate upward — the true effect of institutions on income is likely *smaller* than 0.522. But the question is not just whether OLS overstates the effect; it is whether institutions have *any* causal effect at all.

4 The Instrument: Settler Mortality

4.1 Why Settler Mortality Works

AJR's key insight is that the disease environment faced by European colonizers in the 18th and 19th centuries was exogenous to current economic outcomes. Yellow fever and malaria were determined by ecology, not by the economic potential of the territory. And the disease environment directly shaped what kind of colonization occurred:

- **High-mortality environments** (coastal West Africa, parts of South America): colonizers died quickly; only small garrisons remained; extraction-oriented institutions dominated.
- **Low-mortality environments** (North America, Australia, New Zealand): settlers could survive and reproduce; European-style inclusive institutions were transplanted.

These institutions persisted through independence. The colonizer's approach to property rights, legal systems, and political accountability left a legacy that shaped development trajectories centuries later.

The instrument `logem4` is the log of settler mortality per 1,000 people per year. Negative values of the predicted relationship with `avexpr` are expected: higher mortality → less settlement → weaker institutions → lower income today.

4.2 The Two Conditions

For settler mortality to be a valid instrument for institutional quality, it must satisfy two conditions:

Relevance: Settler mortality must be correlated with institutional quality. If mortality was high, colonizers built extractive institutions; if low, they built inclusive ones. We can test this directly.

Exclusion restriction: Settler mortality must affect income *only through* institutional quality — not through any other pathway. This is an assumption that cannot be directly tested. The main concern: could settler mortality affect income through channels other than institutions? For instance, areas with high malaria burden might have lower productivity today through direct effects on labor and health. AJR argue that controlling for latitude and Africa indicators addresses this, and that the bulk of the long-run effect runs through institutions.

5 The IV Analysis

5.1 First Stage: Settler Mortality → Institutions

The first stage regresses institutional quality on the instrument:

$$\text{institutions}_i = \alpha_0 + \alpha_1 \cdot \log(\text{settler mortality})_i + u_i$$

```
fs_fit <- lm_robust(avexpr ~ logem4, data = ajr_clean)
```

```r
fs_line <- tibble(
  logem4 = seq(min(ajr_clean$logem4), max(ajr_clean$logem4), length.out = 200),
  avexpr = coef(fs_fit)["(Intercept)"] + coef(fs_fit)["logem4"] * logem4
)

ggplot(ajr_clean, aes(x = logem4, y = avexpr)) +
  geom_point(color = "#A51C30", size = 2.5, alpha = 0.75) +
  geom_text(aes(label = shortnam), size = 2.2, hjust = -0.2, alpha = 0.6,
            check_overlap = TRUE) +
  geom_line(data = fs_line, color = "#1E3C72", linewidth = 1.0) +
  annotate("text", x = 6.8, y = 8.8,
           label = sprintf("Slope = %.3f (SE = %.3f)\nF-stat = %.1f",
                           coef(fs_fit)["logem4"],
                           fs_fit$std.error["logem4"],
                           fs_fit$fstatistic[1]),
           color = "#1E3C72", size = 3.2, hjust = 0) +
  labs(
    title    = "High Settler Mortality Predicts Weaker Institutions Today",
    subtitle = "First stage: Z -> D",
    x        = "Log settler mortality (deaths per 1,000 per year)",
    y        = "Institutional quality (avexpr, 0-10)",
    caption  = "Source: Acemoglu, Johnson & Robinson (2001)"
  ) +
  theme_minimal(base_size = 11) +
  theme(panel.grid.minor = element_blank())
```

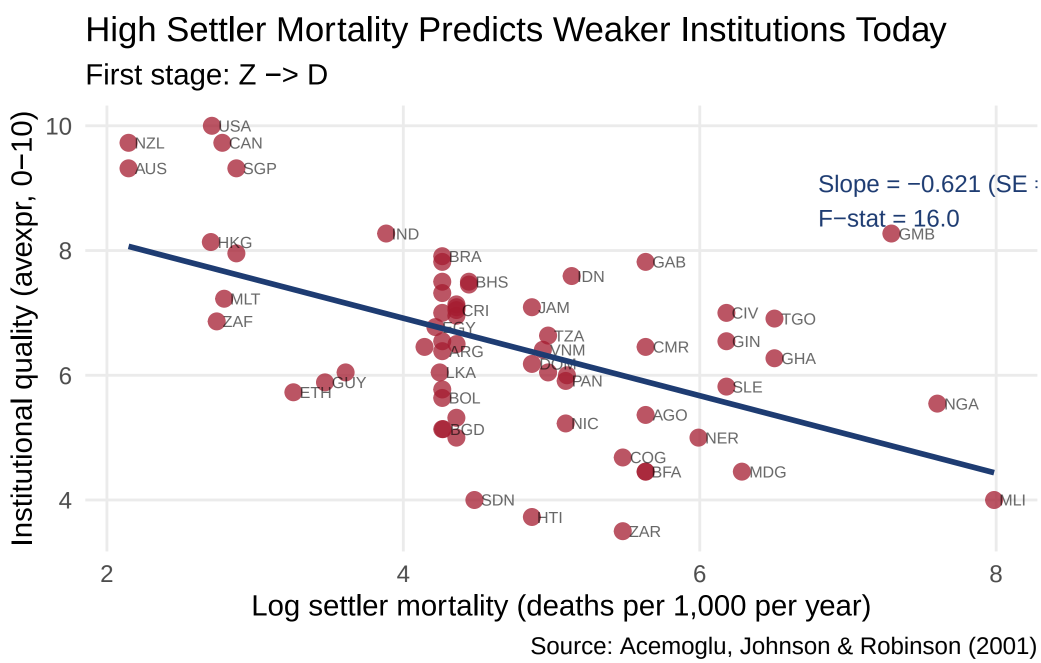

Figure 2: First stage: log settler mortality predicts institutional quality. Each point is a country. Countries where European settlers faced high mortality (right side) ended up with weaker institutions (lower avexpr). The negative slope is strong: F-statistic 16.

```
cat(sprintf("First stage coefficient on logem4: %.3f (SE = %.3f)\n",
            coef(fs_fit)["logem4"], fs_fit$std.error["logem4"]))
```

```
First stage coefficient on logem4: -0.621 (SE = 0.155)
```

```
cat(sprintf("F-statistic on instrument: %.1f\n", fs_fit$fstatistic[1]))
```

```
F-statistic on instrument: 16.0
```

```
cat(sprintf("Interpretation: 1-unit increase in log settler mortality → %.3f-unit decrease i
            coef(fs_fit)["logem4"]))
```

```
Interpretation: 1-unit increase in log settler mortality → -0.621-unit decrease in institutiona
```

The first-stage coefficient is negative (-0.621) and highly significant. The F-statistic of 16 exceeds the conventional threshold of 10 for a strong instrument (Stock and Yogo, 2005). The instrument is relevant: settler mortality systematically predicts institutional quality across these 64 former colonies.

5.2 Reduced Form: Settler Mortality → Income

The reduced form regresses the outcome directly on the instrument, bypassing the endogenous treatment entirely:

$$\log(\text{GDP})_i = \pi_0 + \pi_1 \cdot \log(\text{settler mortality})_i + e_i$$

```r
rf_fit <- lm_robust(logpgp95 ~ logem4, data = ajr_clean)

rf_line <- tibble(
  logem4   = seq(min(ajr_clean$logem4), max(ajr_clean$logem4), length.out = 200),
  logpgp95 = coef(rf_fit)["(Intercept)"] + coef(rf_fit)["logem4"] * logem4
)

ggplot(ajr_clean, aes(x = logem4, y = logpgp95)) +
  geom_point(color = "#2980B9", size = 2.5, alpha = 0.75) +
  geom_text(aes(label = shortnam), size = 2.2, hjust = -0.2, alpha = 0.6,
            check_overlap = TRUE) +
  geom_line(data = rf_line, color = "#1E3C72", linewidth = 1.0) +
  annotate("text", x = 6.8, y = 9.5,
           label = sprintf("Slope = %.3f (SE = %.3f)",
                           coef(rf_fit)["logem4"],
                           rf_fit$std.error["logem4"]),
           color = "#1E3C72", size = 3.2, hjust = 0) +
  labs(
    title    = "Settler Mortality Predicts Income Today",
    subtitle = "Reduced form: Z -> Y",
    x        = "Log settler mortality (deaths per 1,000 per year)",
    y        = "Log GDP per capita, 1995",
    caption  = "Source: Acemoglu, Johnson & Robinson (2001)"
  ) +
  theme_minimal(base_size = 11) +
  theme(panel.grid.minor = element_blank())
```

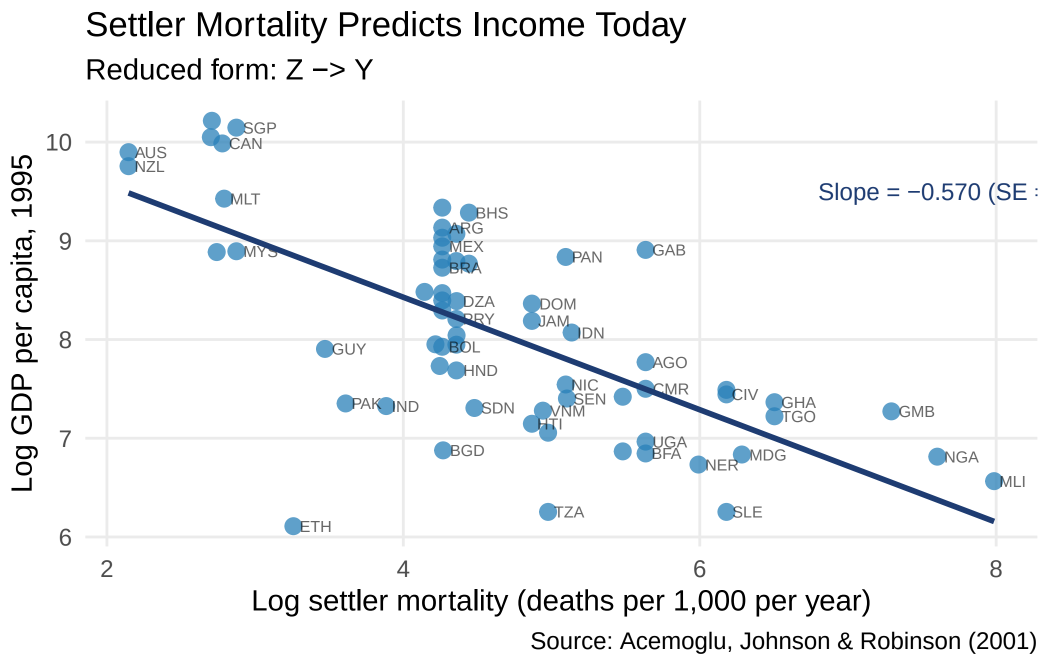

Figure 3: Reduced form: log settler mortality predicts log income. The instrument affects income through the institutional channel it created. The negative slope is clearly visible — colonies with healthier disease environments for settlers have higher incomes today.

```
cat(sprintf("Reduced form coefficient on logem4: %.3f (SE = %.3f)\n",
            coef(rf_fit)["logem4"], rf_fit$std.error["logem4"]))
```

```
Reduced form coefficient on logem4: -0.570 (SE = 0.074)
```

The reduced form slope is -0.57. The instrument moves income: a one-unit increase in log settler mortality is associated with a 0.57-log-point *decrease* in current income. This is the total effect of the instrument on the outcome, running through the institutional channel.

5.3 The Wald Estimator

With one instrument and one endogenous variable, the IV estimator reduces to a simple ratio:

$$\hat{\beta}^{\text{Wald}} = \frac{\hat{\pi}_1}{\hat{\alpha}_1} = \frac{\text{Reduced form slope}}{\text{First stage slope}}$$

The intuition: the instrument moves institutions by $\hat{\alpha}_1$ and income by $\hat{\pi}_1$. Dividing the income effect by the institution effect gives the causal effect of institutions on income.

```
wald_estimate <- coef(rf_fit)["logem4"] / coef(fs_fit)["logem4"]

cat(sprintf("Wald = RF / FS = %.3f / %.3f = %.3f\n",
            coef(rf_fit)["logem4"],
            coef(fs_fit)["logem4"],
            wald_estimate))
```

```
Wald = RF / FS = -0.570 / -0.621 = 0.917
```

The Wald estimate is 0.917. This is the causal effect of institutional quality on log income — specifically, the effect identified by the variation in institutions that is driven by settler mortality.

5.4 Two-Stage Least Squares (2SLS)

The Wald estimator is exact with a single instrument and no controls. The more general method is two-stage least squares (2SLS), which extends naturally to multiple instruments and covariates.

Stage 1 — Project the endogenous variable onto the instrument:

$$\hat{D}_i = \hat{\alpha}_0 + \hat{\alpha}_1 Z_i$$

Stage 2 — Regress the outcome on the projected values $\hat{D}_i$:

$$Y_i = \beta_0 + \beta^{2SLS} \hat{D}_i + e_i$$

The key insight: $\hat{D}_i$ contains only the variation in institutions that comes from the instrument — variation that is exogenous by assumption. By using $\hat{D}_i$ rather than D_i, we remove the endogenous component of institutions and estimate only the causal effect.

```r
iv_fit <- iv_robust(logpgp95 ~ avexpr | logem4, data = ajr_clean)

cat(sprintf("2SLS coefficient on avexpr: %.3f\n", coef(iv_fit)["avexpr"]))
```

```
2SLS coefficient on avexpr: 0.917
```

```r
cat(sprintf("SE: %.3f\n",                              iv_fit$std.error["avexpr"]))
```

```
SE: 0.173
```

```r
cat(sprintf("95%% CI: [%.3f, %.3f]\n",
            iv_fit$conf.low["avexpr"],
            iv_fit$conf.high["avexpr"]))
```

```
95% CI: [0.572, 1.262]
```

```r
cat(sprintf("Note: 2SLS = Wald = %.3f (confirmed)\n", wald_estimate))
```

```
Note: 2SLS = Wald = 0.917 (confirmed)
```

The 2SLS estimate (0.917) matches the Wald estimate exactly, as it must with a single instrument and no additional controls. The 95% confidence interval is [0.572, 1.262].

Table 2: **?(caption)**

(a) OLS vs. 2SLS estimates of the effect of institutional quality on log GDP per capita. Outcome: logpgp95. Treatment: avexpr. Instrument (2SLS only): logem4.

Estimator	Coefficient	Std. Error	95% CI
OLS	0.522	0.050	[0.422, 0.623]
2SLS (IV)	0.917	0.173	[0.572, 1.262]

6 OLS vs. 2SLS: What the Difference Tells Us

```r
tibble(
  Estimator   = c("OLS", "2SLS (IV)"),
  Coefficient = round(c(coef(ols_fit)["avexpr"], coef(iv_fit)["avexpr"]), 3),
  `Std. Error` = round(c(ols_fit$std.error["avexpr"], iv_fit$std.error["avexpr"]), 3),
  `95% CI`    = c(
    sprintf("[%.3f, %.3f]", ols_fit$conf.low["avexpr"], ols_fit$conf.high["avexpr"]),
    sprintf("[%.3f, %.3f]", iv_fit$conf.low["avexpr"],  iv_fit$conf.high["avexpr"])
  )
) |>
  kable(caption = "OLS vs. 2SLS estimates of the effect of institutional quality on log GDP
```

The 2SLS estimate (0.917) is *larger* than the OLS estimate (0.522). This finding is counterintuitive if you expected OLS to be biased upward — and it requires an explanation.

Why is 2SLS > OLS? Several mechanisms are possible:

1. **Measurement error in institutions**: The `avexpr` variable is measured with error. OLS attenuates coefficients toward zero when the right-hand-side variable is measured imprecisely (attenuation bias). IV corrects for measurement error by using an external source of variation.

2. **LATE vs. ATE**: The 2SLS estimator recovers a **Local Average Treatment Effect** — the effect for countries whose institutional quality was affected by settler mortality (the "compliers"). If the marginal effect of institutions is larger for countries close to the institutional quality margin induced by the instrument, 2SLS will exceed the population average effect.

3. **Reverse causation biasing OLS downward**: In equilibrium, very poor countries may also face so much political instability that their formal property rights protections are weak *and* they attract less investment — but the direction of this effect is ambiguous.

The most compelling interpretation, offered by AJR themselves, is that measurement error in institutions biases OLS toward zero, and the IV corrects this. The true causal effect of institutions on income is larger than OLS suggests.

7 What "Local" Means in the LATE

The 2SLS with settler mortality as the instrument estimates the causal effect of institutions for a specific subset of countries: those whose institutional quality was changed by the settler mortality environment. These are the **compliers** in Angrist-Imbens language.

In this context, compliers are countries for which: - Low settler mortality led to higher institutional quality (they would have had weaker institutions if mortality had been high) - High settler mortality led to lower institutional quality (they would have had stronger institutions if mortality had been low)

Countries that would have had strong institutions regardless of settler mortality (always-takers) and countries that would have had weak institutions regardless (never-takers) do not contribute to the 2SLS estimate.

The LATE interpretation: **a one-unit increase in institutional quality causes a 0.92-log-point increase in income per capita for countries whose institutional development was shaped by the settler mortality environment.** This is a large effect: it implies moving from the 25th to the 75th percentile of institutional quality (roughly 2.5 units) would multiply income by $e^{2.5\times0.917} \approx 10$.

8 Limitations and the Albouy Critique

AJR's paper generated significant debate, the most prominent challenge from David Albouy (2012). Albouy argued that roughly one-third of the settler mortality data were imputed, often from regions with different ecological conditions than the country itself. He showed that excluding or correcting suspicious observations substantially weakens the first stage and the IV estimate.

AJR responded that their main results are robust to various alternative sample restrictions and coding choices. The debate illustrates a general principle: the exclusion restriction — the assumption that the instrument affects the outcome only through the treatment — cannot be tested, and its plausibility must be assessed through robustness checks and qualitative arguments.

Other concerns:

- **Institutional persistence mechanisms**: The paper assumes that institutions established during colonization persist largely unchanged. The channel through which 19th-century mortality affects 21st-century prosperity is long and complex.
- **Definition of institutions**: `avexpr` is a composite measure combining several dimensions of property rights protection. Different institutional dimensions (legal origins, constitutional constraints, bureaucratic quality) tell somewhat different stories.
- **Sample selection**: The dataset covers 64 former European colonies. The findings may not generalize to countries that were never colonized or that had very different colonial experiences.

9 Conclusion

AJR's analysis provides compelling evidence that institutions causally affect economic prosperity. The OLS estimate of 0.522 — large as it is — likely understates the true effect due to measurement error attenuation. The 2SLS estimate of 0.917 recovers a local average treatment effect that is nearly twice as large, suggesting that for countries whose institutional quality was shaped by the European settler mortality environment, institutional quality has enormous economic consequences.

Four numbers summarize the analysis:

- **OLS**: 0.522 — biased (likely toward zero due to measurement error)

- **First stage F**: ~16 — strong instrument (exceeds Stock-Yogo threshold of 10)
- **Wald = RF/FS**: 0.917 — causal effect by instrumental variables
- **2SLS**: 0.917 — confirming the Wald calculation

The elegance of the strategy lies in exploiting a source of variation — 19th-century settler mortality — that is plausibly unrelated to current economic conditions except through the institutional channel it created. Whether you find this argument fully convincing depends on how seriously you take the exclusion restriction. That is the fundamental question every IV analysis must answer.

This paper was produced using Quarto with R. All code is included above. Data are from Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2001. "The Colonial Origins of Comparative Development: An Empirical Investigation." American Economic Review 91(5): 1369–1401. The `estimatr` package (Blair et al., 2022) is used for heteroskedasticity-robust standard errors throughout.