Earth Cooling

Zachary del Rosario

2019-02-25

In 1863, Lord Kelvin carried out calculations that suggested the age of the earth was approximately 24 to 400 million years. [1] In this handout, you will re-analyze Lord Kelvin's analysis.

Learning Goals

Students will:

1. Identify sources of model-form uncertainty
2. Determine what information would be necessary to improve an analysis

Cooling Model

Lord Kelvin modeled the cooling of the Earth as diffusion on a semi-inifinite domain, with a Dirichlet boundary at the Earth's surface (here taken to be $T(x = 0, t) = 0$) and interior, and a uniform initial condition of $T(x > 0, t = 0) = T_0$. In full, this is

$$\begin{aligned} \partial_t T &= \alpha \partial_{xx} T, \\ T(x = 0, t) &= 0, \\ T(x \to \infty, t) &= T_0, \\ T(x > 0, t = 0) &= T_0. \end{aligned}$$

A time-varying solution to this model looks like the following.

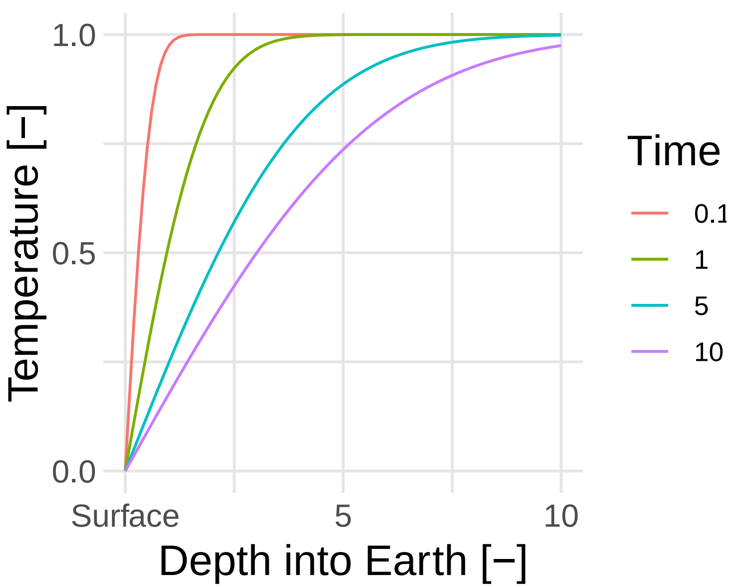

Q1: What assumptions does this model make?

Every equality in the statement above encodes one or more assumptions; what are these assumptions?

Just identify the assumptions for now; we'll assess how reasonable they are later.

A1:

1. The governing PDE is for conduction; this model encodes the assumption that conduction is the only relevant form of heat transfer.
2. The governing PDE also includes no source terms; this model encodes the assumption that there are no sources of heat generation in the Earth.
3. The left boundary condition encodes the assumption of a fixed Earth surface temperature.
4. The right boundary condition encodes the assumption of an infinite heat reservoir.
5. The initial condition encodes the assumption of a uniform temperature across the Earth.

Age Estimate

The conduction model above admits a solution given by

$$T(x,t) = T_o \text{erf}\left(\frac{x}{2\sqrt{\alpha t}}\right).$$

To approximate the age of the Earth, Lord Kelvin considered available temperature gradient data.

$$t = \frac{(KT_0/Q)^2}{\pi \alpha}.$$

Quantity	Symbol	Estimate	Units
Surface heat flux	Q	$[40 \times 10^{-3}, 80 \times 10^{-3}]$	Wm^{-2}
Initial temperature	T_0	1300	$°C$
Thermal conductivity	K	3	$Wm^{-1}K^{-1}$
Thermal diffusivity	α	1×10^{-6}	m^2s^{-1}

Rather than have you plug-and-chug, I've carried out the calculations for you.

Q	T0_C	K	a	t	years	Million Years
0.04	1300	3	0	3.025933e+15	95951717	95.952
0.08	1300	3	0	7.564833e+14	23987929	23.988

Lord Kelvin's model produces an age estimate of 24 to 96 million years. The commonly accepted estimate for the age of the Earth is `4.5` billion years.

Q2: Could uncertainties in the inputs be to blame?

We only have bounds for one of the inputs; could uncertainties in the others be to blame? How could you determine whether this was the case?

A2:

A very simple means to check would be to assume bounds for the estimated quantities, and re-compute the age. First, I re-compute using a fairly modest perturbation size.

```
fac <- 1.2
Q    = c(40e-3, 80e-3)
T0_C = c(1300 / fac, 1300 * fac)
K    = c(   3 / fac,    3 * fac)
A    = c(1e-6 / fac, 1e-6 * fac)

df_calcs2 <-
  map_dfr(
    Q,
    function(q) {map_dfr(
        T0_C,
        function(t) {map_dfr(
          K,
          function(k) {
            tibble(
              Q    = q,
              T0_C = t,
              K    = k,
              a    = A
            )}
          )
        })
    }) %>%
  mutate(
    t     = (K * T0_C / Q)^2 / (pi * a),
    years = t / (3600 * 24 * 365)
  )

df_calcs2 %>%
  ggplot(aes(years)) +
  geom_histogram(bins = 10) +
  geom_vline(xintercept = age_true) +
  scale_x_log10() +
  theme_common() +
  labs(
    x = "Estimated Age (Years)",
    y = "Counts",
    caption = "Perturbation factor 1.2"
  )
```

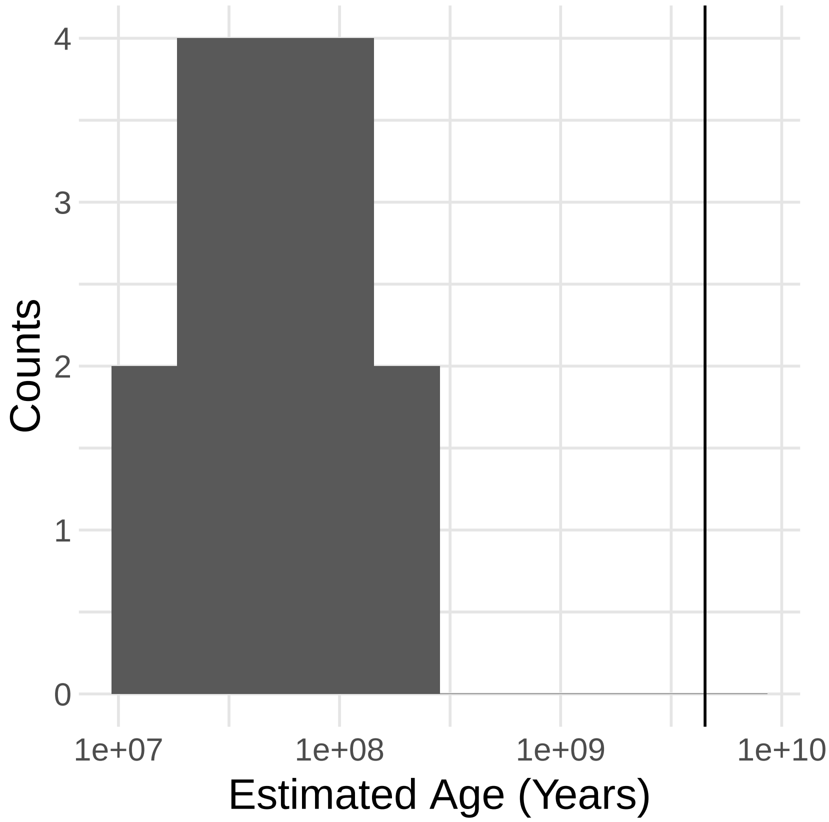

Perturbation factor 1.2

The ensemble of cases does not capture the now-accepted Earth age, so these uncertainty bounds are not compatible with the provided data. I next try larger bounds.

```
fac <- 3
Q    = c(40e-3, 80e-3)
TO_C = c(1300 / fac, 1300 * fac)
K    = c(   3 / fac,    3 * fac)
A    = c(1e-6 / fac, 1e-6 * fac)

df_calcs3 <-
  map_dfr(
    Q,
    function(q) {map_dfr(
        TO_C,
        function(t) {map_dfr(
          K,
          function(k) {
            tibble(
              Q    = q,
              TO_C = t,
              K    = k,
              a    = A
            )}
          )
        })
    }) %>%
  mutate(
```

```r
    t     = (K * T0_C / Q)^2 / (pi * a),
    years = t / (3600 * 24 * 365)
  )

df_calcs3 %>%
  ggplot(aes(years)) +
  geom_histogram(bins = 10) +
  geom_vline(xintercept = age_true) +
  scale_x_log10() +
  theme_common() +
  labs(
    x = "Estimated Age (Years)",
    y = "Counts",
    caption = str_c("Perturbation factor ", fac)
  )
```

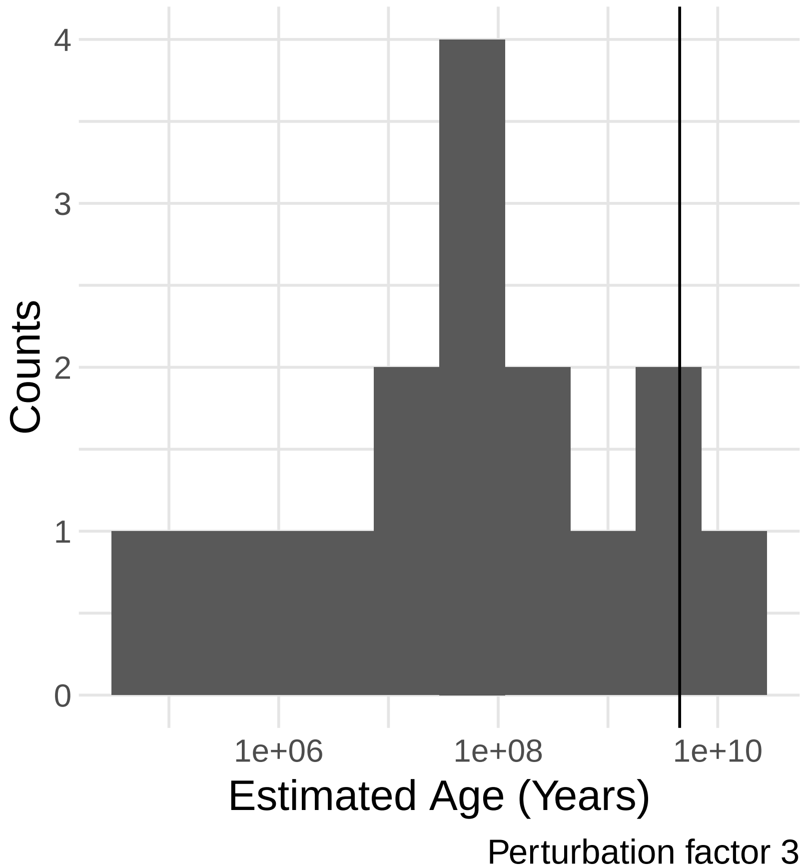

These uncertainty bounds are compatible with the provided Earth age. However, these are fairly wide bounds on the inputs – a factor of `fac` above and below the nominal estimates. One would need to argue that such large bounds are credible to accept this explanation for the discrepancy.

Note that we can *continue to inflate the perturbations to capture any proposed Earth age we want.* It is not enough to assume arbitrary uncertainties – we need data or principles when defining such bounds.

Q3: Could errors in the assumptions be to blame?

Return to the assumptions you listed above. What effect might these have on the estimated age? What information would you want to assess the credibility of these assumptions?

A3:

Let's revisit the assumptions:

1. The governing PDE is for conduction only.

This essentially assumes the Earth is solid throughout, or at least does not flow appreciably. Since conduction is one of the slowest forms of heat transfer, this could have a huge impact on our calculations.

To judge this assumption, I would want to know more about the composition of the Earth, particularly how it varies with depth.

2. The governing PDE includes no source terms.

This is potentially a faulty assumption. Additional sources of heat could increase the observed heat flux at the Earth's surface – not accounting for these sources could lead to an underestimate of the Earth's age.

3. The left boundary condition encodes the assumption of a fixed Earth surface temperature.

The surface temperature of Earth is certainly not fixed. However, the coldest and hottest temperatures recorded on Earth's surface are between $[-88°C, 58°C]$. This variation of about $100°C$ is unlikely to account for the discrepancy we've observed.

4. The right boundary condition encodes the assumption of an infinite heat reservoir.

This assumption ultimately depends on the relevant length and timescales associated with the Earth cooling, and interacts with the assumption of conduction above.

Assuming diffusion is the only relevant mechanism for heat transfer, this assumption is problematic if the diffusion distance $\sqrt{\pi\alpha t}$ is too large compared to relevant Earth lengthscales. At the nominal diffusivity and accepted earth age, the diffusion distance (in meters) is approximated below.

```
diffusion_distance <-
  sqrt(pi * 1e-6 * age_true * 356 * 24 * 3600)
print(sprintf("%4.3e", diffusion_distance))
```

```
## [1] "6.594e+05"
```

The radius of the Earth is approximately 6.4×10^6 meters, so this lengthscale is within an order of magnitude of being problematic.

However, this analysis is predicated on *diffusion only*; as we discussed above. As with Point 1. above, I would want to know more about the internal compostion of the Earth.

5. The initial condition encodes the assumption of a uniform temperature across the Earth.

This is almost surely untrue! This simple initial condition is closely related to the other assumptions. If the composition of the Earth varies greatly across its depth, then an initially uniform distribution of temperature could be incompatible with this inhomogeneity.

An Improved Analysis

John Perry was a former assistant of Lord Kelvin; he examined Kelvin's assumptions and proposed an alternative physical description. The key difference in Perry's model was the presence of liquid-rock below a thin band of solid Earth – a molten center for the planet. The figure below depicts the difference between their models graphically.

```
tribble(
   ~x, ~xend, ~y, ~yend, ~type,
    0,    10,  1,     1,  "Solid",
    0,   2.5,  0,     0,  "Solid",
  2.5,    10,  0,     0,  "Liquid"
) %>%
  ggplot(aes(x, y, color = type)) +
  geom_segment(aes(xend = xend, yend = yend), size = 2) +
  theme_common() +
  scale_color_discrete("Layer Type") +
  scale_x_continuous(
    breaks = c(
      "Surface" = 0,
      "L" = 2.5,
      10
    )
  ) +
  scale_y_continuous(
    breaks = c(
      "Kelvin" = 1,
      "Perry"  = 0
    )
  ) +
  coord_cartesian(ylim = c(-1, 2)) +
  labs(
    x = "Depth into Earth [-]",
    y = "Author",
    title = "Model Comparison"
  )
```

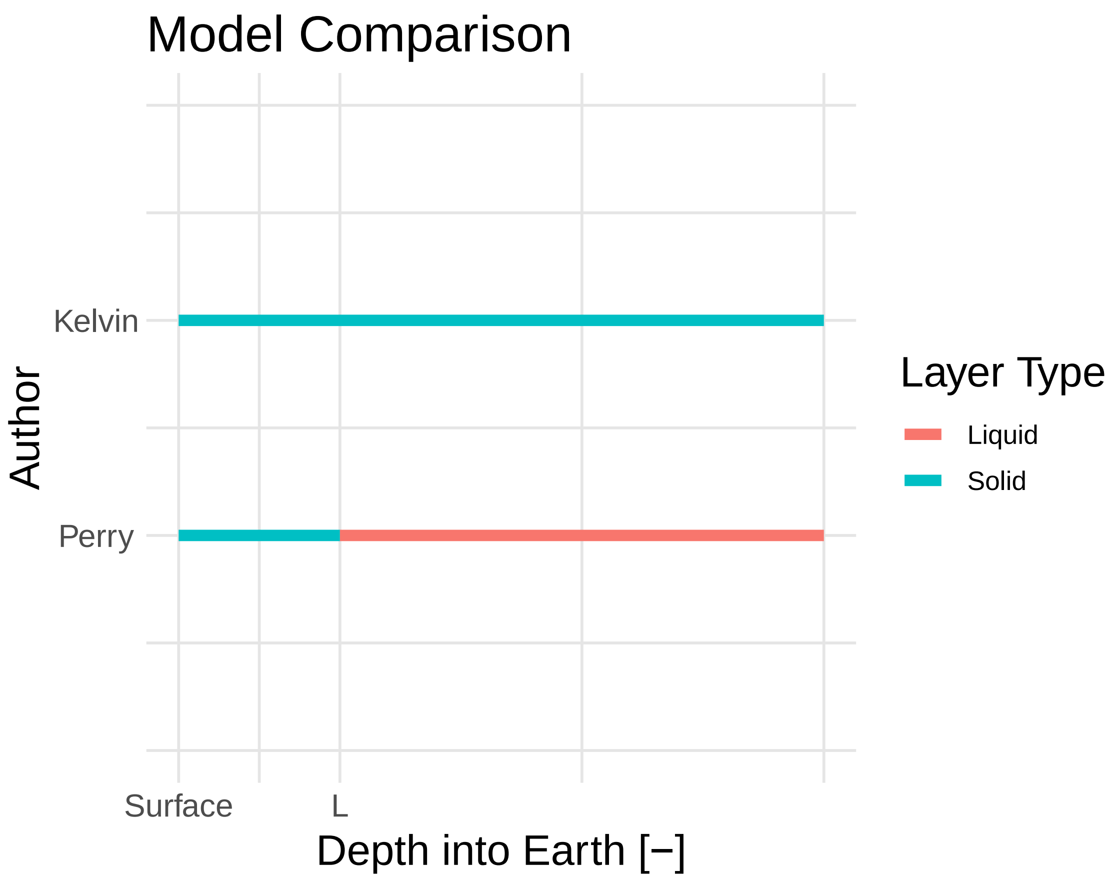

Perry decided to approximate the interior fluid as a well-mixed layer with uniform temperature – this leads to a thin *lid* of solid rock which dominates the heat transfer with conductive transport. This left Perry with another lengthscale L to estimate.

I reproduce a figure from the cited GSA article depicting the Earth's near-surface heat flux over time for Kelvin's and Perry's models.

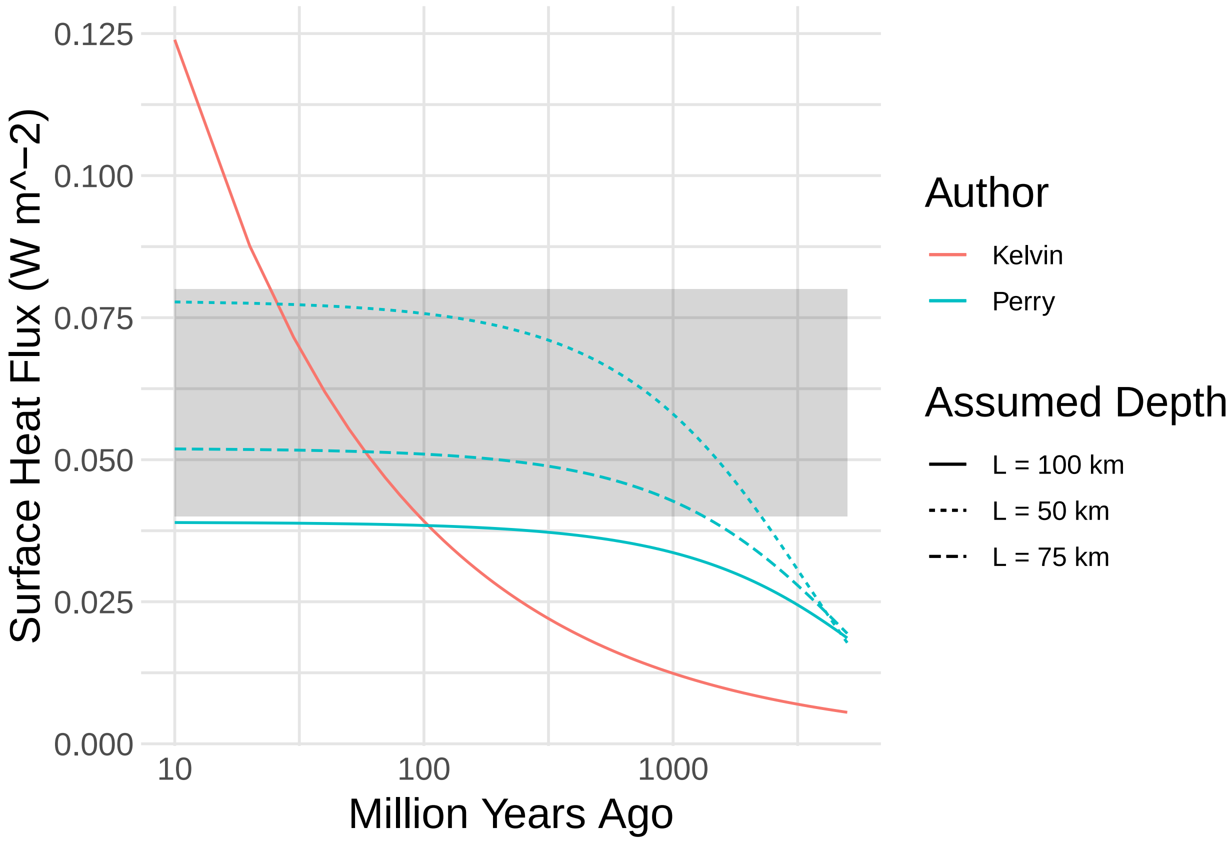

The shaded region above depicts the interval of measured surface heat flux values; we can see that Perry's model is compatible with the measured values, while Kelvin's overshoots. This overshoot results in an inaccurate estimate for the age of the Earth. While neither model agrees with more modern pictures of the Earth's interior, Perry's model captures an important feature of the Earth's composition.

Our modern understanding of Earth's temperature profile is more complicated still, with a number of different regions.[1] However, there is a thin 'lid' section similar to Perry's model, due to the boundary between the slow conducting solid region and the fast convecting liquid interior.

[1]By Bkilli1 - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=28934308

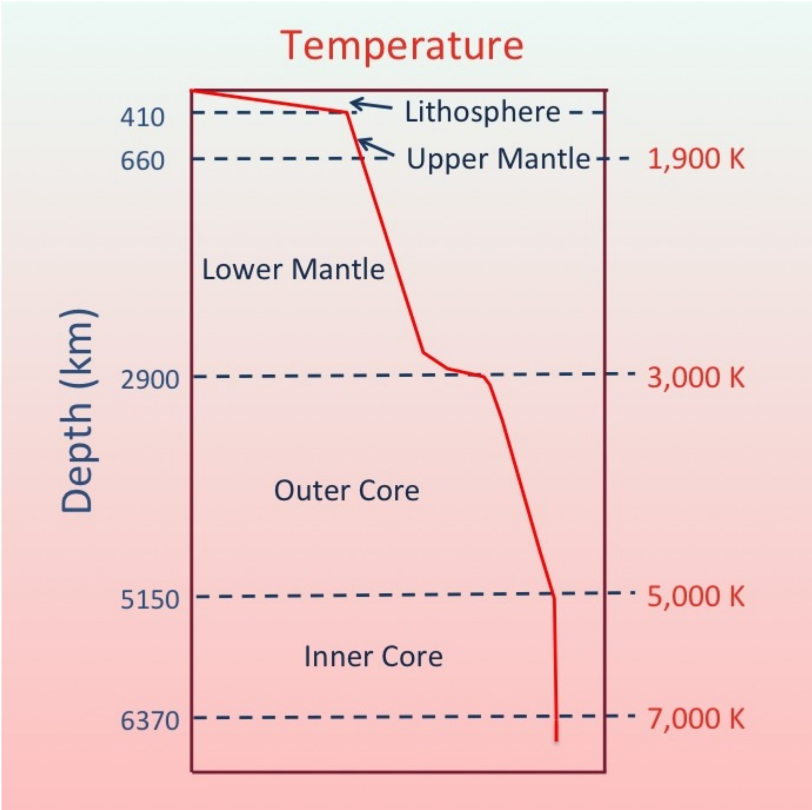

From this schematic, we can see that both Perry and Kelvin severly underestimated the interior temperature of the Earth. Modern approaches to estimating the Earth's age are based on radiometric dating, which relies on an entirely different principle.[2]

Main Takeaways

1. Quantifying uncertainty requires principles

It is misguided to apply arbitrary perturbations to inputs; we need either data (representative samples that inform numerical values) or principles (physics-based arguments for bounds, e.g. energy should be non-negative and conserved) to inform this sort of analysis.

We will learn techniques to learn from data

2. Model-form errors lie outside the model

Detecting model-form error requires that we criticize a model. We need to question the assumptions under which the model was constructed, and determine how those might affect our predictions.

The now-accepted answers to why Lord Kelvin's estimation was wrong are both of model-form error type.[3] and model-form error is often the most difficult to tackle and the most important source of uncertainty.

Schematic view for comparison against modern view:

[2] https://pubs.usgs.gov/gip/geotime/age.html

[3] Lord Kelvin neglected both the heat generation due to radioactivity (a source discovered some years later) and transport due to convection of molten rock. Even famed celebrated physicists can make mistakes!

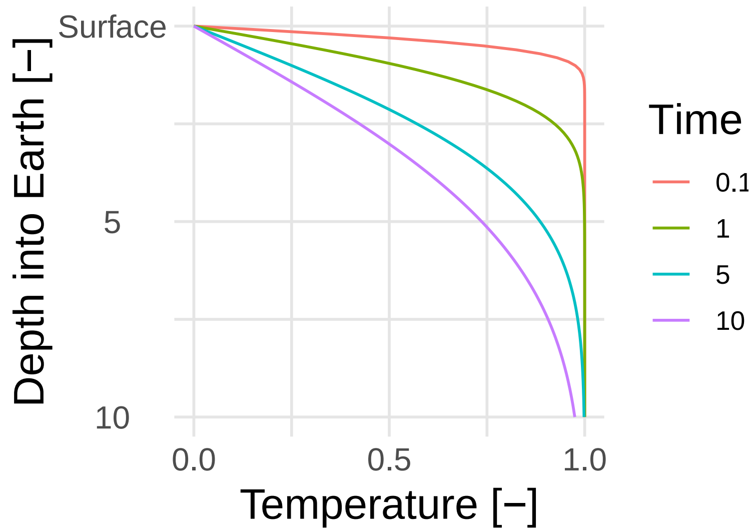


```
## Saving 4 x 3 in image
```

[1] P. England, P. Molnar, and F. Richter, Gsa Today **17**, (2007).