The missing month: US inflation in October 2025

Miguel Salema

Class note, November 2026

This note grew out of a graded exercise in *Time Series Analysis II* at the University of Bern, autumn semester 2026.

The US federal government shut down from October 1st to November 12th 2025, and the Bureau of Labor Statistics (BLS) collected no prices for the consumer price index (CPI) in that time. No CPI for October 2025 was ever published. In FRED the series `CPIAUCSL` has a hole.[1] This note fills it in three steps. The dynamics of the CPI series alone can only split the difference between September and November. A series that *was* observed in October, the CPI for gasoline, adds a lot of information. A prior centred on the BLS's own assumption then gives a posterior, computed by importance sampling. The answer is that prices barely moved in October.

[1] Data: FRED, `CPIAUCSL` (CPI) and `CUSR0000SETB01` (CPI for gasoline), monthly, seasonally adjusted, January 2000 to August 2026.

A hole with a known sum

Let P_t be the seasonally adjusted CPI and $y_t = 100\,\Delta \log P_t$ monthly inflation as a percentage, not annualised. Call October 2025 τ. Without P_τ, two inflation rates are missing, y_τ and $y_{\tau+1}$, but not their sum:

$$s \equiv y_\tau + y_{\tau+1} = 100 \log \frac{P_{\tau+1}}{P_{\tau-1}} = 0.252.$$

The unknown is $x \equiv y_\tau$. November is then $y_{\tau+1} = s - x$, and once x is fixed the data set is complete (Figure 1).

Two values of x are benchmarks. To compute the November index, the BLS carried the missing survey prices forward from September.[2] So $x = 0$. The other is the even split, $x = s/2 = 0.126$, the geometric mean of the September and November levels. The Bureau of Economic Analysis did the latter.[3]

[2] Bureau of Labor Statistics, "2025 federal government shutdown: impact on the Consumer Price Index", frequently asked questions, https://www.bls.gov/cpi/additional-resources/2025-federal-government-shutdown-impact-cpi-faq.htm.

[3] Bureau of Economic Analysis, "Personal income and outlays, October and November 2025", technical note, January 22nd 2026.

One series: split the difference

Model inflation as an AR(1),

$$y_t = \alpha + \rho\, y_{t-1} + \varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2),$$

and estimate $\phi = (\alpha, \rho, \sigma^2)$ by OLS on February 2000 to September 2025, before the hole. Given the first observation and normal errors, OLS is the maximum likelihood estimator: $\hat{\alpha} = 0.110$, $\hat{\rho} = 0.478$, $\hat{\sigma} = 0.266$. Now hold ϕ fixed and treat x as the only unknown.

Filling the hole with x changes three residuals and no others:

$$\begin{aligned} \varepsilon_\tau &= x - \alpha - \rho\, y_{\tau-1}, \\ \varepsilon_{\tau+1} &= (s - \alpha) - (1 + \rho)\, x, \\ \varepsilon_{\tau+2} &= (y_{\tau+2} - \alpha - \rho s) + \rho\, x, \end{aligned}$$

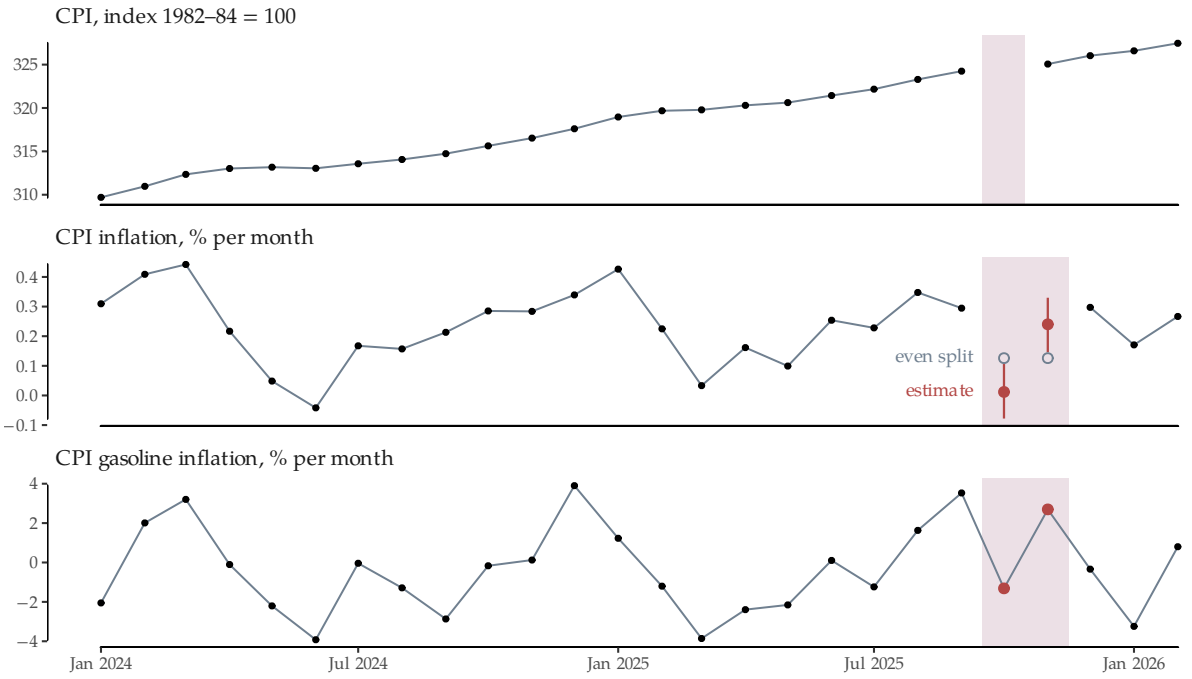


Figure 1: Monthly, seasonally adjusted. Top: the CPI, with one value missing, October 2025. Middle: CPI inflation, with two values missing (shaded) whose sum is known; hollow, the even split $s/2$; red, this note's estimate of October (posterior mean and 95% credible set) and the November it implies, $s - x$. Bottom: gasoline inflation, observed throughout. Source: FRED.

that is, $\varepsilon_t = \varepsilon_t^0 + r_t x$ with $r = (1, -(1+\rho), \rho)$. The log likelihood is a constant minus $\sum_t \varepsilon_t^2 / 2\sigma^2$, so it is quadratic in x. Completing the square gives $\mathcal{L}(x) \propto \exp\{-\frac{1}{2} P (x - \hat{x})^2\}$, with

$$P = \frac{\sum_t r_t^2}{\sigma^2} = \frac{2(1+\rho+\rho^2)}{\sigma^2},$$
$$\hat{x} = -\frac{\sum_t r_t \varepsilon_t^0}{\sum_t r_t^2} = \frac{s}{2} + \frac{\rho\,(y_{\tau-1} - y_{\tau+2})}{2(1+\rho+\rho^2)}.$$

So we have our likelihood. It's obvious that the MLE estimator will be $\hat{x}_{\text{MLE}} = \hat{x}$. Two things stand out. The intercept α drops out, because a change in α shifts all three residuals alike. And $\hat{x}$ leaves the even split iff September and December differ: with $\rho > 0$, persistence makes the path through the hole slope from $y_{\tau-1}$ towards $y_{\tau+2}$. In 2025 they hardly differ (0.295 and 0.297) so $\hat{x} = 0.126 = s/2$, with a standard deviation $P^{-1/2} = 0.144$. Just this one series cannot tell October from November.

Two series: let gasoline speak

Gasoline is one of the few indexes the BLS did publish for October. Let g_t be gasoline inflation, defined like y_t, and $w_t = (y_t, g_t)'$ a VAR(1),

$$w_t = \alpha + A\, w_{t-1} + e_t, \qquad e_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \Sigma),$$

again estimated by OLS on data up to September 2025:

$$\hat{A} = \begin{pmatrix} 0.721 & -0.015 \\ 2.597 & 0.213 \end{pmatrix},$$

with residual standard deviations of 0.262 and 5.43 and a correlation of 0.90 between the residuals of the two equations.

The algebra carries over. The unknown enters only the CPI element of w_τ and $w_{\tau+1}$: with $\iota = (1,0)'$ and w_t^0 the data completed with $x = 0$, $w_\tau = w_\tau^0 + \iota x$ and $w_{\tau+1} = w_{\tau+1}^0 - \iota x$, so

$$e_t = e_t^0 + r_t x, \quad r_\tau = \iota, \quad r_{\tau+1} = -(I + A)\,\iota, \quad r_{\tau+2} = A\,\iota,$$

and the likelihood is again normal in x, with

$$P = \sum_{t=\tau}^{\tau+2} r_t' \Sigma^{-1} r_t, \qquad \hat{x} = -P^{-1} \sum_{t=\tau}^{\tau+2} r_t' \Sigma^{-1} e_t^0.$$

In R the likelihood takes a few lines, with the parameters at their estimates (`W` holds y_t and g_t):

```
fill <- function(x) {      # October = x, November = s - x
  W[tau, "cpi"] <- x; W[tau + 1, "cpi"] <- s - x; W
}
loglik <- function(x) {    # up to a constant
  Wx <- fill(x)
  e  <- sapply(tau:(tau + 2), function(t)
          Wx[t, ] - alpha - A %*% Wx[t - 1, ])
  -0.5 * sum(diag(t(e) %*% solve(Sigma) %*% e))
}
optimize(loglik, c(-1, 1), maximum = TRUE)
```

The maximum is at $\hat{x} = 0.019$, and the curvature there gives a standard deviation of 0.058, as the closed form says. Gasoline moves the estimate from 0.126 to 0.019 and cuts its standard deviation by 60% (Figure 2).

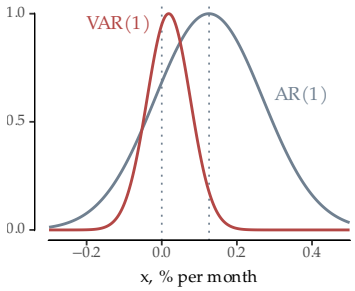


Figure 2: Likelihood of x, scaled to a maximum of one. Dotted: the carry-forward, $x = 0$, and the even split, $x = s/2$.

Why? The CPI and gasoline surprises move together within a month. At the estimate, gasoline inflation was 2.50 points lower than the VAR expected in October and 3.26 points higher in November. Consumer prices probably surprised in the same directions, so October gets much less than half of s. A price we observe is a noisy measurement of a shock that also hit the price we do not observe. There is a second, dynamic channel: lagged CPI inflation enters the gasoline equation with a coefficient of 2.6, so December's gasoline price says something about November's CPI growth.

A prior, and importance sampling

Now treat x as random, keep (α, A, Σ) at their estimates, and centre a prior on the carry-forward,

$$x \sim \text{Cauchy}(0, \kappa), \qquad \kappa = 0.08.$$

Half of a Cauchy's mass lies within κ of its centre, and before the shutdown inflation in half of all months was within about 0.08 of the average of that month and the next.[4] The heavy tails matter: had the data put October far from zero, a Cauchy prior would give way where a normal one would not. The log posterior is

$$\log p(x \mid \text{data}) = \text{const} + \ell(x) - \log\{1 + (x/\kappa)^2\},$$

with $\ell = \log \mathcal{L}$. A normal likelihood times a Cauchy prior is not a standard distribution, so its mean and quantiles have no closed form.

Importance sampling computes them. Draw $x_1, \ldots, x_R$ from a source density g, weight each draw by the posterior kernel over the source,

$$w_i = \frac{\mathcal{L}(x_i)\, p(x_i)}{g(x_i)}, \qquad \tilde{w}_i = \frac{w_i}{\sum_j w_j},$$

and average: $\mathrm{E}(x \mid \text{data}) \approx \sum_i \tilde{w}_i x_i$. The effective sample size, $1/\sum_i \tilde{w}_i^2$, says how many of the R draws count. Two sources suggest themselves. With the prior as source the weights are the likelihood; with $\mathcal{N}(\hat{x}, P^{-1})$, which is proportional to the likelihood, they are the prior. Both combine the same prior with the same data; only the computation differs.

[4] Half the interquartile range of $(y_t - y_{t+1})/2$, February 2000 to September 2025, is 0.078.

Source	Mean	95% credible set	ESS
Prior, Cauchy$(0, 0.08)$	0.012	$[-0.078,\ 0.103]$	5,400
$\mathcal{N}(\hat{x}, P^{-1})$	0.012	$[-0.078,\ 0.106]$	9,100
$\mathcal{N}(s/2,\ 0.02^2)$	0.086	$[\ 0.056,\ 0.146]$	96

Table 1: The posterior of x by importance sampling, $R = 10{,}000$ draws from each source. ESS: effective sample size.

```
x  <- rnorm(R, xhat, 1 / sqrt(P))            # source g
lw <- sapply(x, loglik) + dcauchy(x, 0, kappa, log = TRUE) -
      dnorm(x, xhat, 1 / sqrt(P), log = TRUE)   # log weights
w  <- exp(lw - max(lw)); w <- w / sum(w)
c(mean = sum(w * x), ESS = 1 / sum(w^2))
```

The normal source sits where the posterior is, so its weights, a Cauchy density over a narrow range, vary little, and almost every draw counts (Table 1). The prior is wider than the posterior and wastes about half of its draws where the likelihood is negligible. Both work because both cover the posterior. The third source, too narrow and in the wrong place, gives a confident wrong answer. Its effective sample size, 96 out of 10,000, gives it away. [5]

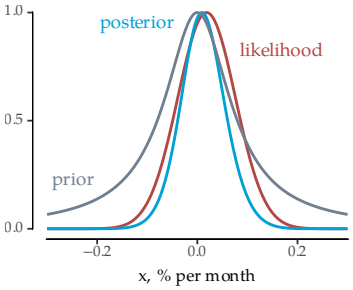


Figure 3: Prior, likelihood and posterior of x, each scaled to a maximum of one.

The answer

In October 2025 CPI inflation was about 0.01% (posterior mean 0.012, 95% credible set -0.08 to 0.11), about 0.1% at an annualised rate: October was flat, close to the carry-forward and far below the even split. November takes 0.24 of the two-month rise of 0.25 (Figures 1 and 3).

[5] For more about how to interpret the effective sample size, see A. B. Owen, *Monte Carlo theory, methods and examples*, 2013, section 9.3, free at https://artowen.su.domains/mc/. He also warns that a large effective sample size is no proof that the method worked: all the draws may miss part of the posterior. On where the formula comes from and how far to trust it: V. Elvira, L. Martino and C. P. Robert, "Rethinking the effective sample size", *International Statistical Review*, 2022.

The prior's scale hardly matters, because the likelihood itself sits close to zero. On a grid, the posterior mean is 0.000, 0.002, 0.013, 0.018 and 0.019 for $\kappa = 0.001$, 0.01, 0.1, 1 and 100. A tiny κ pins the posterior to the carry-forward whatever the data say; a large one leaves the normalised likelihood. A caveat is that the parameters are fixed at their estimates, so the credible set ignores the uncertainty about them..

Three lessons carry over to other holes in other series. A missing value is one more unknown. A series observed in the hole helps to the extent that its shocks are correlated with those of the missing one; Σ and A set the precision, the surprises the direction. And importance sampling is only as good as its source, which the effective sample size reveals.