UK Inflation Dynamics since the Thirteenth Century*

By James M. Nason and Gregor W. Smith[1]

CAMA, Australian National University and Virginia Center for Economic Policy; Queen's University, Canada

Historians have suggested there were waves of inflation or price revolutions in the UK (and earlier England) in the 13th, 16th, and 18th centuries, prior to the ongoing inflation since 1935. We study retail price inflation since 1251 and model its dynamics. The model is an AR(n) but allows for gradually evolving or drifting parameters and stochastic volatility. The long-horizon forecasts suggest only one inflation wave, that of the 20th century. We also use the model to measure inflation predictability and price-level instability from the beginning of the sample and to provide measures of real interest rates since 1695.

JEL classification: E31, E37.

Keywords: inflation, price revolutions, stochastic volatility, time-varying parameters.

∗ Manuscript received August 2021; revised December 2022.

Running Head: UK Inflation Dynamics

1: We thank the Social Sciences and Humanities Research Council of Canada and the Jenkins Family Economics Fund of North Carolina State University for support of this research. For helpful comments we thank the editor (Jesús Fernández-Villaverde) and two referees of this journal as well as Pooyan Amir-Ahmadi, Tim Cogley, Lee Craig, Thorsten Drautzburg, Stefano Giglio, Ian Keay, Laura Liu, Dong Lu, Christian Matthes, Elmar Mertens, seminar participants at the University of Colorado Boulder, Remin University, the Bundesbank, the University of British Columbia, Indiana University, and the Federal Reserve Bank of Philadelphia, and conference participants at the CEA, the IAAE, and the NBER-NSF SBIES. Corresponding author: James M. Nason, *jmnason87@gmail.com.*

> Then again, price history has not yet succeeded in acquiring its own tools of analysis. For better or worse, it must rely on those provided by economists and statisticians.
>
> —Braudel and Spooner (1967, p 374)

1. Introduction

In this study we statistically model inflation in the UK (and earlier England) since the thirteenth century. There are three aims. First, while the inflation of the 20th century is well known, historians have argued that there was an earlier wave of inflation (or price revolution) in the 16th century and that there were also such waves in the 13th and 18th centuries. In historical studies a wave is defined as a sustained increase in inflation, relative to the rate in adjacent time periods. We ask whether these two (or possibly four) waves appear in long-horizon inflation *forecasts*, within a unified statistical framework. Second, we report on inflation predictability (as defined by Cogley, Primiceri and Sargent, 2010), price-level uncertainty, and price-level instability (as defined by Cogley and Sargent, 2015 and Cogley, Sargent, and Surico, 2015) for this long span. Third, we provide measures of inflation expectations, based on a time-series model, that in turn allow new measurements of historical, real interest rates.

Price history as a field of research has three main components: (a) selecting and recording individual prices and constructing price indexes; (b) documenting patterns in prices or inflation rates; (c) attributing these patterns to causes such as changes in the money supply or real shocks. This study falls in category (b). We draw on recent research in category (a) that constructs consumer or retail price indexes. Research in category (c) may remain too challenging for this long time span, given the lack of data on some other macroeconomic indicators.

Modelling UK inflation over the past 750–800 years is itself challenging because inflation rates over most of this period appear to be very noisy. The noise might reflect actual changes caused by climate and harvests varying as well as unavoidable measurement error associated with collecting and aggregating prices before the advent of the Cost of Living Index (which began in 1914) and the Retail Price Index (which began in 1947). We first allow for the fact that reported inflation may have measurement error with a variance that depends on the data source. We then model underlying or true inflation using an nth-order autoregression, AR(n). The AR(n) grounds its forecast of inflation on recent inflation history. In practice our evidence shows the lag length is 3 years. The noise makes modelling heteroskedasticity in the regression error important. We estimate the AR(n) models with slowly evolving stochastic volatility (SV). The AR(n)s also feature time-varying parameters (TVPs) so that the intercept and autoregressive parameters also may slowly evolve over time. This gives an AR(n) the chance to fit data from the high-inflation period after 1935 and from earlier centuries, without our knowing turning points in advance.

The resulting TVP-SV-AR(n) optimally weights past inflation experiences to forecast future inflation. Forecasts (from h-year ahead to the long run) eliminate the noise making evident patterns in these predictions that are obscured in sample inflation yet may have mattered for setting interest rates or other contractual prices. We thus define inflation waves anew, as persistent fluctuations in long-horizon, expected inflation. The idea is to see whether these forecasts are persistently higher over certain long spans in addition to the post-1935 time period.

The forecasting model is univariate. Particularly as one moves into the 20th century one might think of forecasting with other variables in the information set, or, later still, using data from forecast surveys or financial markets. However, Faust and Wright (2013) find that slowly-evolving autoregressions do quite well in an inflation forecasting tournament for several countries including the UK. Survey data improve them but conditioning on activity variables generally does not. Still, this study certainly is not intended as a last word on post-1945 inflation forecasts for example. We focus on the TVP-SV-AR model for (a) some comparability with other studies of inflation dynamics such as Cogley, Sargent, and Surico (2015) and especially (b) comparability over time, so we can compare long-horizon forecasts, inflation predictability, and price-level instability across many centuries.

Section 2 briefly summarizes a number of historical studies, many of which documented commodity prices. We use more recently constructed consumer (or retail) price indexes assembled by Clark (2020). We thus focus on overall inflation rather than on grain prices for example. Section 3 describes this index. Section 4 describes the measurement-error model, the TVP-SV-AR(n), and prior densities for their parameters.

Section 5 presents the findings in the form of posterior densities for a range of statistics, including measures of persistence and stochastic volatility. We first report on estimates of the extent to which measurement error declined over time and then discuss the dynamics of inflation. The main finding is that there is only one wave in long-horizon inflation forecasts, that of the 20th century. The wave takes the form of a higher inflation intercept and especially higher persistence in inflation, though that second feature began much earlier. We also find that volatility tends to fall over time as the economy becomes more diversified. But this pattern of declining SV is by no means monotone. For example, volatility increases in the 1500s, a period often characterized as a price revolution, during the Napoleonic Wars, and during World War I.

Section 5 also describes the measurement and findings for inflation predictability and price-level instability. We find that inflation was predictable (measured by the posterior distribution of its R^2 statistic) at horizons of 1, 2, 3, and 5 years, for every year since the thirteenth century. Price-level instability generally trends down over time, but with local peaks in 1558, 1800, 1917–1922, and 1975.

Section 6 provides new measures of real interest rates at short and very long maturities for the periods since 1695 and 1701 respectively. For the 1695–2019 period we estimate that the short-term real interest rate peaked in 1922 while the long-term real interest rate peaked in 1974. Section 7 summarizes the findings. A variety of details concerning the data, algorithm, and formulas are collected in appendices in the online supplement.

2. Historical Studies

UK price history has a long history, including classic studies by Rogers (1866–1902), Abel (1935; translated 1980), Beveridge (1939), and Phelps Brown and Hopkins (1981). Fischer's (1996) book *The Great Wave* provided an overview of price history. He argued that there were four waves in inflation overall in Europe: roughly 1180–1350; 1480–1650; then from the late 1700s to the early 1800s; and from 1896 to the present. He also argued that inflation was higher and less volatile in each successive wave.

For the first of these hypothesized waves Mate (1975) discussed a burst of inflation during 1305–1310. She attributed it to an inflow of silver to England, in turn due to a change in France's mint ratio, an export surplus, and a prohibition on exporting English coins.

The second wave or price revolution is much more studied. Doughty (1975) aggregated prices collected by Rogers (1866–1902) and Beveridge (1939) to form an annual index of industrial prices in southern England for 1401–1640. He found the inflation rate in this index was particularly high during 1542–1560. Outhwaite (1982) used the indexes constructed by Phelps Brown and Hopkins (1981) to suggest an increase in inflation began around 1480 but certainly was clear by 1520.

Historians have long debated the causes of this inflation. Naturally a number of studies focus on changes in the money supply. Gould (1970) described the debasement that began in 1544, late in the reign of Henry VIII. Challis (1989) attributed the depreciation of the currency beginning in 1521 to an increased supply of silver to the mint, often from the crown itself (for example plate confiscated from monasteries). The debasements ended with the Elizabethan restoration of the coinage in 1560–1561, but Outhwaite suggested that inflation continued. Paper money was not introduced until late in the 17th century. But foreign coins and tokens made of copper or lead were used to solve the problem of small change (discussed by Sargent and Velde, 2003) in the Stuart period and so added to the money supply.

The idea that silver from South America via Spain was the cause of inflation did not become widespread until the 1600s. That idea was then popularized by Rogers (1866–1902) and then of course by Earl Hamilton in his books on Spain published in the 1930s. Other historians have emphasized physical explanations such as population changes, discussed by Phelps Brown and Hopkins (1957) for the 1500s or by Fischer. But Outhwaite suggested that data on both the money supply and population are unlikely to ever be refined enough retrospectively to decide among explanations.

Braudel and Spooner (1967, p 400) studied the three-hundred-year span from 1450 to 1750 in Europe, noting that the consensus was:

> The four long periods are: a fall, or rather stagnation, in the fifteenth century; a rise in the sixteenth continuing into the seventeenth century; then a fall until about 1720–50; finally a renewed inflation in the eighteenth century.

Fischer referred to the period from 1750 to 1820 as the third wave, with reference to commodity prices and for England to the index of Phelps Brown and Hopkins (1957). The British inflation of this period often is attributed to the suspension of gold convertibility from 1797 to 1821. Braudel and Spooner (1967) focused on the long term (*la longue durée*, to use the term of the Annales school), meaning the secular trend beyond cycles. We suggest measuring this feature with a long-horizon inflation forecast.

3. Measurement

We use the annual retail price index, constructed and described by Clark (2020) at `MeasuringWorth.com`. For 1209–1869 the locations are all in England. For simplicity we refer to the UK in the paper's title, following Clark (2020) and Thomas and Dimsdale (2017). For 1870–1914 the source is Feinstein (1995) and the series applies to the entire UK. For 1915–1946 the source is the UK cost-of-living index. After 1947 the price index is the series CDKO of the Office of National Statistics (ONS).

The largest component of the Clark price index is food, followed by fuel, lodging, light, services, and manufactured goods. These in turn are aggregates of more detailed sub-components. Clark describes how the weights vary over time to reflect changing patterns of expenditure (especially after the Industrial Revolution) and also the absence of some prices in the early centuries. As in a modern index some goods (like cider and tallow candles) have declining weights or disappear from the index while new goods (like potatoes, coffee, and gas lighting) are added. Observations are missing from the Clark series for 1213, 1215, 1222, 1228–1231, 1234, and 1238–1244. As a result, our estimation and tests use observations from 1245 on.

The upper panel of Figure 1 shows this (log) price index, $p_{uk,t}$, from 1250 to 2019. Here one can see the prolonged inflation beginning in 1935. But the series also drifts up over 1250–1934. The lower panel shows the corresponding inflation rate $\pi_{uk,t}$, which is the first difference of $p_{uk,t}$, beginning in 1251. The annual average inflation rate is 0.93% for the entire span and 0.42% for the period before 1935. Online Appendix A discusses and graphs the alternative indexes proposed by Allen (2001) and Thomas and Dimsdale (2017) as well as the earlier indexes collected by Mitchell (1988).

One might wonder whether to statistically model $p_{uk,t}$ or, instead, $\pi_{uk,t}$. To test for unit roots Table 1 shows the *DF-GLSu* and Q_T statistics derived by Elliott (1999). These offer higher power than the traditional ADF test. They draw the initial value from the unconditional distribution rather than using a value of zero as the Elliott-Rothenberg-Stock tests do. The lag length is selected by the BIC. Some critical values are given in the notes to the table.

The results are shown first for the entire 1245–2019 sample, then for four non-overlapping periods (1245–1439, 1440–1633, 1634–1827, and 1828–2019), and finally for the periods after World War I and after World War II. For $p_{uk,t}$ there is evidence of a unit root for the entire sample and for each quarter of the sample with the exception of 1245–1439, where both statistics suggest that this variable may be stationary. For $\pi_{uk,t}$ the *DF-GLSu* statistic is below the 1% critical value for the full sample, for each quarter, and for 1919–2019. It is below the 2.5% critical value for 1946–2019. The Q_T statistic is below its 1% critical value for each of these periods. Overall, then, Table 1 clearly suggests that we model the inflation rate as a stationary series.

Looking at $\pi_{uk,t}$ in the lower panel of Figure 1 then—so as to put the data all on a similar, stationary scale and allow comparisons—shows that it is not easy to detect patterns over time. High rates of inflation or deflation were common in earlier centuries: It is really the persistence of $\pi_{uk,t}$ that seems to have changed. Equivalently, the lower panel shows the noisiness in the early data, discussed in the introduction.

Historians therefore often have smoothed the commodity price or inflation rate series by averaging. Figure 2 illustrates the effects for $\pi_{uk,t}$. The three panels show its averages over the 10, 20, and 30 years prior to each year. The sample average is a random variable, so surrounding each average is a 68% confidence interval, constructed with a HAC (Newey-West) standard error with 1, 2, or 3 lags. As the window widens one can detect the waves discussed in section 2. The average was relatively high in the late 1200s, middle 1500s, and early 1800s and particularly low in the late 1800s.

Our idea in this paper is to instead base the smoother on the model with the best fit to the data or the best forecasts. This is in contrast to the averages in Figure 2, which of course use a pre-set window or span of years and fixed, equal weights on them. We test for the number of lags needed to statistically explain inflation or forecast its future values and we estimate the weights on those lags. And the weights are not fixed but may evolve slowly over time. This time-variation allows us to study this long series in a unified way, with no need for pre-set break dates or tests for structural breaks.

4. Time-Varying Autoregression

Our statistical framework begins with allowance for measurement error. The variance of the measurement error may depend on the span of data, because those data originate in several sources. These time spans are 1251–1869 (for the majority of the Clark data), 1870–1914 (for the Feinstein data), 1915–1946 (for the first, official cost-of-living index) and 1947–2019 (for the ongoing retail price index of the ONS). We index these four spans by $\tau = 1, 2, 3, 4$.

We assume that measured inflation, $\pi_{uk,t}$, is given by the sum of true inflation, π_t, and measurement error, m_t:

$$\pi_{uk,t} = \pi_t + m_t, \tag{1}$$

and

$$m_t = \sqrt{\beta_\tau}\sigma_u u_t \qquad u_t \sim N\big(0,\, 1\big), \tag{2}$$

so that $\beta_\tau \sigma_u^2$ is the variance of measurement error in span τ. As in Cogley and Sargent (2015), Cogley, Sargent, and Surico (2015), and Amir-Ahmadi, Matthes, and Wang (2016), the constancy of β_τ within a span, along with its discrete changes across spans, allows it to be identified separately from stochastic volatility in true inflation, which is described below.

In this model $\beta_1 = 1$ so that σ_u^2 is a baseline variance for 1251–1869. Also $\beta_4 = 0$ for 1947–2019 following Cogley, Sargent and Surico (2015), so inflation is measured without error in the most recent span. In between these values we estimate β_2 and β_3, which describe the extent to which the variance of measurement error may fall relative to the initial span. Hence equation (2) gives the variance of m_t as an element of the vector $[1\; \beta_2\; \beta_3\; 0]'\sigma_u^2$.

Table 2 describes priors for the three measurement-error parameters, σ_u^2, β_2, and β_3. For σ_u^2 the prior is inverse gamma with median 0.18 and a 5%–95% IQR of [0.007, 0.072]. For the parameters describing the decay in measurement-error variance over time we adopt $\mathcal{B}$eta distributions, because those are appropriate for proportions. For the 1870–1914 period the prior for β_2 is $\mathcal{B}$eta(5,3), which has a median of 0.635 and a 5%–95% IQR of [0.341, 0.871]. For the 1915–1946 period the prior for β_3 is $\mathcal{B}$eta(3,5), for which the corresponding statistics are 0.364 and [0.129, 0.659].

With these priors the weights β_2 and β_3 are restricted to lie in (0,1). In the early centuries of this study the consumption basket was no doubt simpler than it was in the 19th and 20th centuries for example. Nevertheless, there are good reasons to think measurement error was greater in the first span. First, some components are simply missing early in the Clark data. For example, there is no measure of lodging costs for much of the 13th century. Second, the number of locations and institutions from which data originate is also quite limited for the medieval and early modern periods. These facts motivate our prior to restrict the weights β_τ so that measurement error declines (or at least does not increase) over time. However, some alternate models of measurement error are described below and studied in Online Appendix B.

We model true inflation π_t as an AR(n) with TVPs and SV. The autoregression is:

$$(3) \qquad \pi_t = \alpha_{0,t} + \sum_{i=1}^{n} \alpha_{i,t}\pi_{t-i} + \xi_t\epsilon_t, \qquad \epsilon_t \sim N\big(0,1\big)$$

while SV evolves as a (log squared) random walk:

$$(4) \qquad \ln \xi_t^2 = \ln \xi_{t-1}^2 + \sigma_\phi \phi_t, \qquad \phi_t \sim N\big(0,\, 1\big).$$

The AR parameters are $\alpha_{i,t}$ for $i = 0, 1, 2, \ldots, n$. They follow a multivariate random walk with possibly correlated innovations. Let $\alpha_t \equiv \{\alpha_{0,t},\, \alpha_{1,t},\, \alpha_{2,t}, \ldots, \alpha_{n,t}\}$ and $\eta_t \equiv \{\eta_{0,t},\, \eta_{1,t},\, \eta_{2,t}, \ldots, \eta_{n,t}\}$. Then the TVPs evolve as:

$$(5) \qquad \alpha_t = \alpha_{t-1} + \eta_t, \qquad \eta_t \sim N(0_{n+1}, \Omega_\eta).$$

The innovations η_t to the multivariate random walk of the TVPs are not independent, according to equation (5), because the covariance matrix Ω_η is not restricted to be diagonal.

However, there is no correlation among either the elements of η_t and the innovation, ϵ_t, to the AR(n), $E(\eta_{i,t}\epsilon_t) = 0$, or the innovation to the (geometric) random walk, ϕ_t, of the SV ξ_t, $E(\eta_{i,t}\phi_t) = 0$, $i = 0, 1, \ldots, n$. The same holds for ϕ_t and ϵ_t, $E(\phi_t\epsilon_t) = 0$. Finally, the innovations to measurement error, u_t, also are uncorrelated with ϵ_t, $\eta_{i,t}$ and ϕ_t.

This time-series model has several appealing features. Using the recent history of inflation as a basis for forecasts seems natural for economists and for historical forecasters. But an AR with constant coefficients is unstable, if not also misspecified, when fitted to this long span of data. Imagine, then, fitting rolling ARs with a finite data window to track the evolution of a forecasting equation. In that case one would need to choose the length of the widow over which to estimate and any decay in the weights on observations within that span. Incorrect calibration of the window width and discount on past observation will induce bias and inefficiency in estimates of the rolling ARs.

Estimation of the TVP-SV-AR(n) automates these decisions, as weights on specific observations depend on the Kalman smoother, computed in the Bayesian sampler we use to construct the posterior of the model. Hence, changes in the underlying dynamics of UK inflation are addressed by time-variation in the scale of shocks to inflation, its unconditional mean, and persistence in the TVPs and SV of the AR(n).

We estimate the TVP-SV-AR(n) on the UK inflation series for 1251–2019 while conditioning on observations for 1245–1250, where $n = 1, \ldots, 6$. These choices set the estimation sample size to $T = 769$. Table 3 describes our priors for the TVP-SV-AR(n). Roughly speaking these are empirical Bayes priors, based on OLS regression for the estimation sample. The prior on the scale variance σ_ϕ in the innovation to SV is inverse gamma.

Table 4 describes how we select the autoregressive lag length, n. First, we calculate log marginal data densities (MDDs), as recommended by Geweke (2005). This criterion chooses the lag length n that the data most support, in terms of persistence and volatility, given equations (1) and (2). The first row of Table 4 shows the data favor $n = 3$ conditional on our priors.

Second, we calculate the Bayesian version of the Akaike information criterion (IC). Watanabe (2010) derives this IC and refers to it as the widely applicable information criterion (WAIC). The WAIC chooses the lag length to yield the best one-year-ahead forecasts (*i.e.* minimize predictive loss) along with estimating a penalty that approximates the number of (unconstrained) parameters using information in the data and priors. Gelman *et al* (2014, pp 173–174) label it the Watanbe-Akaike IC and recommend computing the WAIC $= -2\left[\widehat{\mathcal{L}}_n^T - \widehat{\mathcal{V}}_n^T\right]$, where $\widehat{\mathcal{L}}_n^T$ is the posterior mean of the log predictive likelihood of a TVP-SV-AR(n) and the penalty, $\widehat{\mathcal{V}}_n^T$, is the posterior variance of this log likelihood. The second row of Table 4 indicates the WAIC also chooses the lag length $n = 3$. Thus the choice of lag length is the same whether one looks for the best fit or the best forecasts.

Online Appendix B studies alternate models of measurement error and alternate priors, drawing on these two measures of fit. Online Appendix B.1 includes AR(1) persistence in the measurement error. However, estimates of the AR coefficient exhibited substantial uncertainty with Bayesian credible sets covering zero in all cases. Online Appendix B.2 then studies a model in which there are no restrictions on the three variances across $\tau \in \{1, 2, 3\}$, other than their being positive. This model is similar to the one studied by Cogley, Sargent, and Surico (2015) and Amir-Ahmadi, Mathes, and Wang (2016) except that there is no persistence in m_t. This alternate model of measurement error produces estimates of median volatility that increase across spans from 1251 to 1946, which replicates a puzzle documented by Cogley, Sargent, and Surico (2015). Otherwise, the two models lead to similar economic conclusions. However, the ln MDD and WAIC criteria support the baseline model compared with this alternate model. Finally, we also explored attaching measurement error to $p_{uk,t}$ rather than $\pi_{uk,t}$ but found that created explosive dynamics in the process.

We also studied an alternate model that allows for sudden changes in SV. To do this, we adopted a horseshoe prior that places time-varying heteroscedasticity on the variance of the innovation to the random walk of $\ln \xi_{t+1}^2$. Online Appendix B.3 sets $\pi_{uk,t} = \pi_t$ to establish a simpler environment in which to examine this issue. Online Appendix B.4 then applies the horsehoe prior in this environment. However, the two measures of fit suggest the data are indifferent to adopting this alternate prior.

The TVP-SV-SVAR was introduced by Cogley and Sargent (2005) and Primiceri (2005), who described a Markov chain Monte Carlo (MCMC) algorithm for estimation. Our code to run the Metropolis-Hastings (MH) in Gibbs sampler adapts the algorithm described by Canova and Pérez Forero (2015) to a univariate model. The MH step follows the Gibbs draw in the sampler. The Canova-Pérez Forero sampler uses the correction of Del Negro and Primiceri (2015). We estimate SV in the TVP-AR(n)s by incorporating into our version of the Canova-Pérez Forero sampler the ten-component mixture of normal distributions developed by Omori, Chib, Shephard, and Nakajima (2007). Online Appendix C describes our MCMC sampler in detail both for the baseline model and for the alternate models of Online Appendix B.

5. Posteriors

We next describe the findings in the form of posterior densities from the TVP-SV-AR(3) with our baseline model of measurement error. We begin with the three measurement-error parameters and then turn to TVPs and SV. The TVPs and SV allow us to study changes in inflation persistence and volatility over time and hence look for inflation waves. We also examine inflation predictability and price-level instability.

5.1. Measurement Error. Recall that the measurement error model begins with a baseline variance σ_u^2 for the first span. The posterior density for this value has a median of 0.00163 with a 5%–95% IQR of [0.00128, 0.00208]. This is a smaller scale of measurement error than found by Cogley and Sargent (2015) for example, in part because their early data were on more volatile commodity prices rather than a retail price index. The upper panels of Figure 3 graph the prior and posterior densities for the two fractions that scale down this variance for the next two time spans. For β_2, which applies to 1870–1914, the prior (dashed line) has a median of 0.636 while the posterior (solid line) has a median of 0.442 with a 5%–95% IQR of [0.196, 0.746]. For β_3, which applies to 1915–1946, the prior has a median of 0.364 while the posterior has a median of 0.290 with a 5%–95% IQR of [0.094, 0.587]. Thus both posteriors, although overlapping, have shifted left from their respective priors, implying somewhat larger reductions in measurement-error variance across time spans.

The lower left panel of Figure 3 shows the smoothed estimate of the measurement error year by year, with its 68% credible sets. The sets narrow over spans τ, until measurement error is zero after 1946. The lower right panel shows measured inflation, $\pi_{uk,t}$ along with the median of true inflation, π_t, and the associated 68% credible sets. Median π_t is of course less volatile than $\pi_{uk,t}$, while in some years its credible set extends to values comparable to $\pi_{uk,t}$.

5.2. Forecasts and Waves. Figure 4 presents posterior medians and credible sets for several hidden states. The upper left panel contains the posterior median of the time-varying intercept, $\alpha_{0,t}$. The upper right panel shows the posterior median of the sum of the lag TVPs, $\sum_{i=1}^{3} \alpha_{i,t}$. This is a measure of persistence. The lower left panel plots the median of the time-varying conditional mean or long-horizon forecast $\mu_{\pi,t}$. The changes in parameters are unpredictable, so this long-horizon forecast is:

$$\mu_{\pi,t} \equiv \lim_{j \to \infty} E_t \pi_{t+j} = \frac{\alpha_{0,t}}{1 - \sum_{i=1}^{3} \alpha_{i,t}}, \tag{6}$$

which is computed as the median of the ratio of the posterior draws of the numerator to the posterior draws of the denominator. This long-horizon forecast is the univariate version of the inflation trend estimated and studied using a TVP-SV-VAR by Cogley and Sbordone (2008). SV is depicted in the lower right panel. The four panels also display shadings that are 68% Bayesian credible sets (*i.e.* the 16% and 84% quantiles) of the TVPs and SV.

The intercept term $\alpha_{0,t}$ has considerable uncertainty. Its median is relatively high in the 13th and 16th centuries. But the credible sets contain zero from 1261 to 1923. After 1923 both the median and the 68% credible set rise sharply in the 20th century. The median of the posterior of $\alpha_{0,t}$ peaks at 1.89% in 1975.

The persistence measure—in the upper right panel of Figure 4—exhibits a very low frequency cycle. Its median falls until the 1350s, levels off from 1360 to 1560, rises for the next 100 years, levels off again from then to 1755, and then rises continuously until 1974 before falling after that date to the end of the sample. The median reaches its minimum of -0.460 in 1353 with a 68% Bayesian credible set [-0.768, -0.148]. It crosses zero in 1775–1776 moving from -0.004 [-0.428, 0.200] to 0.003 [-0.427, 0.197]. Its maximum is 0.623 in 1974 with a 68% Bayesian credible set [0.445, 0.789]. These differences *can* be distinguished statistically. Thus it is not the case that the early inflation data are simply so noisy that no patterns can be detected in them.

The key findings are in the lower left panel, which shows the posterior median and credible set for the long-run forecasts $\mu_{\pi,t}$ of equation (6). Prior to roughly 1920, the long-horizon forecasts $\mu_{\pi,t}$ do not vary significantly over time. One could draw a range of horizontal lines that would pass through these credible sets. Thus a series of (non-independent) tests of the hypothesis that each value was equal to this constant would not reject it one-by-one. The 68% credible sets of $\mu_{\pi,t}$ contain zero from 1261 to 1923. Note that we are not arguing that the forecasts might be zero, though that hypothesis also would not be rejected before 1924 except for early in the sample.

This lack of early waves in long-run forecasts contrasts with the historical views cited in section 2. For example, the traditional view of the price revolution of the 16th century is that inflation was low by the standards of the 20th century but persisted over many years to lead to a distinct inflation wave. In the UK inflation data, though, inflation (whether measured or true) was volatile and had negative persistence during this period. Long-horizon inflation forecasts were not noticeably different from those in the previous and subsequent centuries.

The one great wave in long-horizon forecasts begins in the 1920s and is associated with a rising intercept and rising persistence. The median of $\mu_{\pi,t}$ peaks at 5.055% in 1974 with a 68% Bayesian credible set [2.633, 8.362]. In the next year, 1975, the actual inflation rate (shown in the lower panel of Figure 1) also peaked, at 24.3%, the highest value since the mid 1500s. That was also the year the Wilson government began its incomes policy, described in the White Paper *Attack on Inflation*. The infamous IMF loan to the UK followed in 1976. After 1976 the median declines steadily until 2019. The decline features falls in the intercept and persistence measure. Median SV is also at an all-time low at the end of this period.

In the lower right panel of Figure 4, median SV falls on average over time, so that feature is central to the specification. The decline in SV could be due to better measurement from more sources or to genuine decreases in volatility with less dependence on weather and harvests. Either way, though, that fall is not monotonic. For example, median SV rises from 1477 to 1558, during the price revolution. During the first half of the sample, median SV has local peaks in 1251, 1319, 1371, 1549, and 1597. Median SV also rises during the Napoleonic Wars and during World War I. It also has local peaks in 1800, 1921, and 1975.

The one-year-ahead forecast or measure of expected inflation is:

$$E_t \pi_{uk,t+1} = \alpha_{0,t} + \sum_{i=0}^{n-1} \alpha_{i+1,t} \widehat{\pi}_{t-i}. \tag{7}$$

We next report the median of the ensemble of one-year-ahead expected inflation. These forecasts are created by running the Kalman filter on the posterior of a TVP-SV-AR(n) to produce $\widehat{\pi}_{t-i}$, $i = 0, \ldots, n-1$, multiplying these predictions by the posterior draws of $\alpha_{i+1,t}$ and to the result adding the associated posterior draws of $\alpha_{0,t}$.

Figure 5 shows the one-year-ahead forecasts of $\pi_{uk,t}$ and associated forecast errors from the TVP-SV-AR (3) and equations (1) and (2). It also shows 90% Bayesian credible sets. With our focus on the very long run we have so far described the 20th century as a period of inflation. But of course 1921–1934 was a period of deflation, as shown in Figure 1. Such fluctuations are one reason to adopt the TVP-SV-AR: It can adapt to such changes in the data. It is nevertheless not surprising that the forecast errors tend to mimic actual inflation during the early 1920s. It would perhaps be difficult for any time-series model to predict the sharp deflation of 1921–1923. The one-year-ahead forecasts tend to track $\pi_{uk,t}$ with a lag, as one would expect of an AR. Nevertheless, it is noteworthy that the long-horizon forecast (shown in Figure 4) does *not* fluctuate in response to this episode of deflation. The model thus features a steepening of the term structure of inflation expectations as the posterior density of $E_t \pi_{uk,t+1}$ drops while that of $\mu_{\pi,t}$ does not.

5.3. Predictability. To summarize the h-step-ahead predictability of sample inflation, we use the R^2_{ht} statistic proposed by Cogley, Primiceri, and Sargent (2010). This is 1 minus the ratio of the conditional variance to the unconditional variance. It lies between 0 and 1. For example, if $\pi_{uk,t}$ is completely unpredictable then the conditional variance and unconditional variances will coincide, their ratio will be 1, and $R^2 = 0$. Also note that the statistics will go to zero as $h \to \infty$. Online Appendix D in the online supplement develops the formula for these statistics for a TVP-SV-AR(n) and equations (1) and (2).

Figure 6 shows the median R^2_{ht} statistics for $\pi_{uk,t}$ at horizons of 1, 2, 3, and 5 years, with their 68% credible sets. The first, and central, finding is this: There is predictability at each horizon throughout the period. The Bayesian credible sets do not include zero at any horizon and date. (The 16th quantiles vary from 0.004 to 0.006 at $h = 3$ and from 0.001 to 0.002 at $h = 5$ from the mid 1500s to the end of the 1800s.) Historical inflation in the lower panel of Figure 1 appears very noisy, yet $\pi_{uk,t}$ exhibits predictability in every year since the thirteenth century, especially at horizons of 1 and 2 years.

Second, median predictability at each horizon peaks in the mid 1970s. These peaks are $R^2_{1,1974} = 0.394\,[0.230, 0.580]$, $R^2_{2,1975} = 0.164\,[0.051, 0.372]$, $R^2_{3,1975} = 0.097\,[0.019, 0.285]$, and $R^2_{5,1975} = 0.056\,[0.007, 0.209]$, where 68% credible sets appear in brackets.

5.4. Price-Level Instability. Recall that $p_{uk,t}$ denotes the log of the price level. We next measure price-level uncertainty at time t and horizon h, defined as the variance of accumulated inflation between t and $t + h$: $\text{var}_t(p_{uk,t+h} - E_t p_{uk,t+h})$. Formulating this measure involves first computing expected inflation and then finding the variance around that forecast. Price-level instability is then defined—following Cogley and Sargent (2015) and Cogley, Sargent, and Surico (2015)—as the total variation, and not just the unpredictable variation.

At horizon h, instability is given by the square root of the sum of the conditional variance and the squared, conditional mean:

$$\sqrt{\text{var}_t\left(p_{uk,t+h} - E_t p_{uk,t+h}\right) + \left(E_t p_{uk,t+h} - p_{uk,t}\right)^2}. \tag{8}$$

The TVPs make computing j-year ahead expected inflation difficult. We invoke a local approximation inspired by the anticipated utility model (AUM) of Kreps (1998) as implemented by Cogley and Sbordone (1998) and Cogley, Primiceri, and Sargent (2010). This holds the parameters at their current realizations to forecast inflation. We then use the TVP-SV-AR(n) of equations (1)–(5) to find the conditional mean of the log price level. Calculating the variance of the forecast error is more complicated. The details are in Online Appendix E.

Figure 7 reports on the posterior density of the instability measures (8) by graphing the median and 68% credible set at each year. We show results for $n = 3$ and at horizons h =1, 2, 3, and 5. (Results for the uncertainty measure on its own are similar to those for SV shown in the lower right panel of Figure 4 and so are not shown separately.) Instability behaves similarly over time at each horizon, but its value is greater year by year at longer horizons, as shown by the vertical axes in the four panels of Figure 7.

Instability is high during the first 300 years of the sample. At each horizon there are local peaks in 1318, 1371, 1438, 1528, and 1547. The last of the early modern peaks is in 1558, a year that saw a sharp deflation with $\pi_{uk,1558} = -26\%$. This was the year of Elizabeth I's accession to the throne (with Thomas Gresham as her advisor), bringing an end to the Tudor debasements, though debased coinage was not withdrawn until 1560.

Instability then trends down over time, but there are three noteworthy local peaks at each horizon. The first of these peaks is in 1800, during the War of the Second Coalition. In that year the inflation rate was 22%, the highest level of the era, as seen in the lower panel of Figure 1.

Second, at each horizon there are several elevated levels of median instability during World War I and in the early 1920s, a period of whipsawing inflation. Instability was high in 1917, when the wartime inflation rate peaked at 21%. It also was high in 1920 (when $\pi_{uk,1920} = 14\%$) and 1922, the second of two years of very sharp deflation with $\pi_{uk,1921} = -9\%$ and $\pi_{uk,1922} = -19\%$. This deflation followed the tightening of monetary policy in 1920–1921. Howson (1973, 1974) described the basis for this policy decision. Overall, this period from 1917 to 1922 exhibits the highest median instability of any year after 1558, at each horizon.

The third modern peak is in 1975, a year that saw the peak inflation of the 1970s as discussed earlier. While this peak in price-level instability also coincides with a local peak in uncertainty and SV, in this case there is a larger role for $(E_t p_{uk,t+h} - p_{uk,t})^2$ which also peaks in 1975. As already noted, the mid 1970s was an extraordinary period in UK economic history that saw the IMF loan and the government's attempt to control inflation with wage and price controls.

Of course we are modelling only inflation and could list events associated with any year. By listing some events that coincided with these peaks in price-level instability we are not implying that we necessarily understand their causes but rather leaving questions for future research.

Cogley, Sargent, and Surico (2015) study UK inflation from 1791 to 2011 using the unobserved components (UC) model of Stock and Watson (2007) complemented with an AR(1) measurement-error process. Their model thus features a stochastic trend in true inflation in the form of an unobserved, random walk. True inflation has two sources of SV, one in the permanent component and one in an unpredictable transitory component. Their model thus is not directly comparable to ours. (And we found it difficult to identify the modern, UC model for earlier centuries, given that we found evidence against a unit root in inflation in Table 1.) They also use the data assembled by Mitchell (1988) which display higher and more volatile inflation rates during the Napoleonic Wars than do the Clark data. Online Appendix A briefly describes the differences between these two series and graphs them. And they report uncertainty and instability for horizons of 5 and 10 years.

Despite these differences in time span, model, data, and horizons, we nevertheless can compare findings for the period after 1791. The two approaches both display local peaks in price-level instability in 1800 and 1975 at each horizon. There are some differences in the findings though, in both peaks and troughs. First, we find the highest post-1791 peak in 1917–1922 rather than in the mid-1970s, for each horizon. Overall we find much more evidence of price-level instability in the 1917–1922 period. Second, they find a trough at around 1890 that is never matched in the 20th century (by the median or the inter-quartile range). Possibly because our data now run to 2019, we find that median instability falls late in our sample to values comparable to those in the 1890s. We also find relatively wide 68% credible sets late in our sample, though, which makes comparisons less possible. A final difference of course is that we can estimate price-level instability since the 1250s to put the values of the post-1791 period in context. As noted, we can compare the instability of the 1920s or 1970s with that of the 13th to 17th centuries. Again, though, there is considerable dispersion in these distributions throughout this span, so that the credible sets sometimes overlap when comparing these periods.

6. Real Interest Rates

We next report implications for the history of real interest rates. We draw on the expertise of Thomas and Dimsdale (2017) by using their series for nominal interest rates, which begin in 1695. Thus our measures of real rates will differ from theirs only because of the measurement of inflation expectations.

The nominal short rate is series M9 from Thomas and Dimsdale (2017). They splice together annual averages of the Bank Rate 1695–1972, Minimum Lending Rate 1972–1981, Minimum Band 1 Dealing Rate 1981–1997, Repo Rate 1997–2006, and the Bank Rate, 2006–2017. We update it to 2019 from the Bank of England file `baserates.xls`. The long nominal interest rate is series M10, the yield on perpetual annuities/consols. It connects for 1703–1726 yields on new long-term issues, for 1727–1753 the yield on 3% perpetual annuities, for 1756–2015 the yield on Consols, and for 2016–2019 the 20-year zero-coupon yield. These nominal rates are shown in the top panels of Figure 8.

Thomas and Dimsdale's real short rate is based on NIESR inflation forecasts from 1959 and a range of measures from 1996 onwards. Prior to 1959 expectations are equal to realized inflation so the interest rate is an ex post one. Their measure of long-term inflation expectations applies the Hodrick-Prescott filter (with a parameter of 100) to actual inflation beginning in 1600. Notice also that their measure of inflation differs slightly from the Clark measure, as documented in Online Appendix A. We refer to these two real interest rates as the Bank of England series.

Our measures of inflation expectations are from the TVP-SV-AR(3) with equations (1) and (2). From the nominal interest rates we subtract the ensemble of values for the one-year ahead and long-horizon inflation forecasts, then plot the medians and the 90% credible sets of these ex ante real rates. In Figure 8 the lower left panel shows the short-term real rates and the lower right panel the long-term ones.

In comparing short-term, real interest rates in Figure 8, note first that the Bank of England series is much more volatile than ours, a fact that is not surprising since it is an ex post series for much of the sample. Until the 1880s it often lies outside the 90% credible sets for our ex ante real interest rate. The two series coincide more closely after that, except that the troughs in the Bank of England series are deeper than ours in the early 20th century. Both series reach their 20th-century peak in 1922, following the sharp monetary contraction of 1920–1921. In the Bank of England series comparably high values occur in the 18th and 19th centuries. A striking difference, therefore, is that 1922 marks the highest value for the entire 1695–2019 period for our ex ante, short-term, real rate. The median, real short rate rises from 0.11% in 1920 to 9.3% in 1921 and 11.7% in 1922 and does not fall below 2% until 1933. Notice also the median is negative during 1915–1918, 1935–1937, 1939–1952, at several times during the 1970s, and then also from 2009–2019.

Long-term real rates are shown in the lower left panel of Figure 8. It is interesting that the median of our series peaks in 1974, at the same time as the nominal rate and before the 1990 peak in the Bank of England series. This median long rate rises above 3% in 1968, peaks at 11.1% in 1974, and remains above 3% until 1998. But there is more uncertainty associated with the long-term real rate than the short-term real rate. For example, the Bank series lies within our 90% credible sets over the last 40 years while the 90% Bayesian credible sets of our long-term real rate cover zero from 1884 to 1973 and from 1997 to 2019. And of course there are other sources of information on long-horizon inflation expectations for the late 20th century, including surveys and market-based measures. Thus we emphasize earlier periods and those for which the differences between the two series are the largest.

For the late 18th and early 19th centuries the Bank of England long-term real rate is well below our credible sets, a pattern which is then reversed for several years after 1815. Our long-horizon inflation forecasts do not react as much as the HP-filtered inflation series does to the inflation during the Napoleonic Wars. Thus we do not estimate there to have been a sharp fall then rise in the long-term real rate during the early decades of the 19th century. A similar difference appears for the early 20th century. The inflation of World War I and deflation of the 1920s lead to large swings in the Bank of England series that lie outside our credible sets, and our median estimate is smoother. In particular, the Bank of England series is negative during both World Wars and again in the 1970s while our estimates of the median, long-term real interest rate are not.

Hamilton *et al.* (2016) study the short-term real interest rate for a large panel of countries that includes the UK since 1858. Borio *et al.* (2022) construct short-term and long-term (10-year) real interest rates for the UK and other countries since 1870. These studies model inflation expectations using an AR(1) model with a rolling multi-year window. Since the late 19th century they show a sawtooth pattern in the real interest rate, as it trends down, then up, then down again. Figure 8 shows the same pattern for our estimates of UK real rates at both maturities, concluding with very low or negative estimates at the end of the sample. The median estimate of the long real rate is negative for 1945–1947, 1949, and from 2015 to 2019.

These two studies just cited carefully examine a range of explanations for movements in real interest rates. Our paper extends real interest rate measures back in time a further 175 years to 1695 and uses a TVP-SV-AR to represent expected inflation, two features that we hope will make the results useful in further studies of secular trends in real interest rates.

7. Conclusion

We provide a statistical model that offers a window on the work of economic historians —principally Clark (2020)—who have assembled price indexes for the UK (and earlier England) since the thirteenth century. The model allows for measurement error that varies across spans corresponding to different underlying data sources. The TVP-SV-AR model describes the dynamics of inflation using a single model for the entire period, with slowly-evolving parameters and no need for structural break tests. This statistical method seems well-suited to studies in economic history. It allows us to estimate and compare the statistical properties of inflation for periods that are centuries apart.

A central finding is that there is a single, Hokusai-scale wave in historical, long-horizon inflation forecasts, occurring in the 20th century. When looking at the history of the price level or inflation in Figure 1, it is tempting to speculate about the social and economic changes that contributed to this wave, perhaps including the nature of World War I finance, the extension of the franchise after 1918, and/or a broader set of social insurance programs in the 1920s. But recall that Figure 4 shows inflation persistence rising throughout the 19th century and the long-horizon forecast turning up before 1914. We also found a tendency for stochastic volatility in inflation to fall, well before the 20th century.

Despite the noisiness that seems apparent in historical data, inflation is predictable at horizons of 1, 2, 3 and 5 years for each year since 1251, as measured by the R^2 statistic. Price-level instability tends to decline over time, which puts the instability of the 20th century in historical context. But there are notable, local peaks in median estimates of instability in 1558, 1800, 1917–1922, and 1975. For the 1695–2019 period we estimate that a distinct, overall peak in the short-term real interest rate occurred in 1922, following the monetary contractions of 1920–1921. In contrast, the median long-term real interest rate peaked in 1974, the same year in which inflation reached its highest value since the 1500s.

Supporting Information

Additional information may be found online in the Supporting Information section at the end of the article.

Online Appendix A: Alternative Price Indexes.

Figure A1: The Clark-, Bank of England-, and Allen-UK Price Levels on Several Subsamples.

Figure A2: The Clark- and Mitchell UK Price Level and Inflation Rates, 1781–2019.

Figure A3: UK and Statist Price Level and Inflation, 1847–1950.

Online Appendix B: Alternative Models.

Online Appendix B.1: Measurement Error with AR(1).

Online Appendix B.2: Alternative Models of Measurement Error.

Table B1: Alternative Model of Measurement Error.

Table B2: Evaluation of Fit of the TVP-SV-AR(n)s with Baseline and Alternative Models of Measurement Error.

Table B3: Summary of the Posterior Distributions of the Parameters of the Alternative Model of iid Measurement Error.

Table B4: Summary of the Posterior Distributions of the Parameters of the Baseline Model of Measurement Error.

Figure B1: Measurement Error Parameter Densities and Smoothed States.

Figure B2: Posterior Moments of the TVP-SV-AR(3) on UK Inflation, 1251–2019.

Figure B3: 1-Year Ahead Expected UK Inflation and Its Ex Post Forecast Error, 1251–2019.

Figure B4: Predictability of UK Inflation 1-, 2-, 3-, and 5-Years Ahead, 1251–2019.

Figure B5: UK Price-Level Instability, 1251–2019.

Figure B6: UK Nominal and Real Short- and Long-Term Interest Rates.

Online Appendix B.3: The TVP-SV-AR(n) minus Measurement Error.

Table B5: Comparison of Fit of the TVP-SV-AR(n)s without Measurement Error.

Figure B7: Posterior Moments of the TVP-SV-AR(4) on UK Inflation, 1251–2019.

References

Abel, W., *Agricultural Fluctuations in Europe from the Thirteenth to the Twentieth Centuries* (London: St. Martin's Press, 1980).
[Translated by O. Ordish from *Agrarkrisen und Agrarkonjunktur in Mittel Europa vom 13 bis zum 19 Jahrhundert* (Berlin, 1935; new eds. 1966, 1978).]

Allen, R. C., "The Great Divergence in European Wages and Prices from the Middle Ages to the First World War," *Explorations in Economic History* 38 (2001), 411–447.

Amir-Ahmadi, P., C. Matthes and M.-C. Wang, "Drifts and Volatilities under Measurement Error: Assessing Monetary Policy Shocks over the Last Century," *Quantitative Economics* 7 (2016), 591–611.

Beveridge, W., *Prices and Wages in England from the Twelfth to the Nineteenth Century* (New York: Longmans, Green and Company, 1939).

Borio, C., P. Disyatat, M. Juselius and P. Rungcharoenkitkul (2022) "Why So Low for So Long? A Long-Term View of Real Interest Rates," *International Journal of Central Banking* 18 (2022), 47–87.

Braudel, F. P. and F. Spooner, "Prices in Europe from 1450 to 1750," *The Cambridge Economic History of Europe* 4 (1967), 378–486.

Canova, F. and F. J. Pérez Forero, "Estimating Overidentified, Nonrecursive, Time-Varying Coefficients Structural Vector Autoregressions," *Quantitative Economics* 6 (2015), 359–384.

Challis, C. E., *Currency and the Economy in Tudor and Early Stuart England.* (London: The Historical Association, 1989).

Clark, G., "What Were the British Earnings and Prices Then?" (New Series) MeasuringWorth, 2020.

Cogley, T., G. Primiceri and T. J. Sargent, "Inflation-Gap Persistence in the U.S.," *American Economic Journal: Macroeconomics* 2 (2010), 43–69.

Cogley, T. and T. J. Sargent, “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.,” *Review of Economic Dynamics* 8 (2005), 262–302.

Cogley, T. and T. J. Sargent, “Measuring Price-Level Uncertainty and Instability in the United States, 1850–2012,” *Review of Economics and Statistics* 97 (2015), 827–838.

Cogley, T., T. J. Sargent and P. Surico, “Price-Level Uncertainty and Instability in the United Kingdom,” *Journal of Economic Dynamics & Control* 52 (2015), 1–16.

Cogley, T. and A. M. Sbordone, “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review* 98 (2008) 2101–2126.

Del Negro, M. and G. E. Primiceri, “Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum,” *Review of Economic Studies* 82 (2015), 1342–1345.

Doughty, R. A., “Industrial Prices and Inflation in Southern England, 1401–1640,” *Explorations in Economic History* 12 (1975), 177–192.

Elliott, G. “Efficient Tests for a Unit Root When the Initial Observation is Drawn from its Unconditional Distribution,” *International Economic Review* 40 (1999), 767–783.

Faust, J. and J. H. Wright, “Forecasting Inflation,” in G. Elliott and A. Timmermann, eds., *Handbook of Economic Forecasting*, vol.2A (Amsterdam: Elsevier, 2013), chapter 1.

Feinstein, C. H., “Changes in Nominal Wages, the Cost of Living and Real Wages in the United Kingdom over Two Centuries,” in P. Scholliers and V. Zamagni, eds., *Labour's Reward: Real Wages and Economic Change in 19th- and 20th-Century Europe* (Aldershot, Hants: Edward Elgar, 1995), 3–36.

Fischer, D. H., *The Great Wave: Price Revolutions and the Rhythm of History* (Oxford, UK: Oxford University Press, 1996).

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunsun, A. Vehtari and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. (New York: CRC Press, Taylor & Francis Group, 2014).

Geweke, J., *Contemporary Bayesian Econometrics and Statistics* (New York: Wiley, 2005).

Gould, J.D., *The Great Debasement* (Oxford, UK: Oxford University Press, 1970).

Hamilton, J. D., E. S. Harris, J. Hatzius and K. D. West, "The Equilibrium Real Funds Rate: Past, Present, and Future," *IMF Economic Review* 64 (2016), 660–707.

Howson, S., " 'A Dear Money Man'?: Keynes on Monetary Policy, 1920," *Economic Journal* 83 (1973), 456–464.

Howson, S. "The Origins of Dear Money, 1919–20," *Economic History Review* 27 (1974), 88–107.

Kreps, D., "Anticipated Utility and Dynamic Choice," in D.P. Jacobs, E. Kalai, M.I. Kamien, and N.L Schwartz, eds., *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997* (Cambridge, UK: Cambridge University Press, 1998), 242–274.

Mate, M., "High Prices in Early Fourteenth-Century England: Causes and Consequences," *Economic History Review* 28 (1975), 1–16.

Mitchell, B. R., *British Historical Statistics* (Cambridge, UK: Cambridge University Press, 1988).

Omori, Y., S. Chib, N. Shephard and J. Nakajima, "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference," *Journal of Econometrics* 140 (2007), 425–449.

Outhwaite, R. B., *Inflation in Tudor and Early Stuart England,* 2nd ed. (London: Macmillan, 1982)

Phelps Brown, H. and S. V. Hopkins, "Wage-Rates and Prices: Evidence for Population Pressure in the Sixteenth Century," *Economica* 24 (1957), 289–306.

Phelps Brown, H. and S. V. Hopkins, *A Perspective of Wages and Prices* London: Methuen, 1981).

Primiceri, G. E., "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies* 72 (2005), 821–852.

Rogers, J. E. T., *A History of Agriculture and Prices in England from the Year after the Oxford Parliament, 1259, to the Commencement of the Continental War, 1793* (Oxford, UK: Clarendon Press, 1866–1902).

Sargent, T. J. and F. R. Velde, *The Big Problem of Small Change* (Princeton, NJ: Princeton University Press, 2003).

Stock, J. H. and M. W. Watson, "Why has US Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (2007), 3–33.

Thomas, R. and N. Dimsdale, "A Millennium of UK data," Bank of England OBRA datset, 2017.

Watanabe, S., "Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory," *Journal of Machine Learning Research* 11 (2010), 3571–3594.

Table 1. Unit Root Tests

	Years	*DF-GLSu*		Q_T	
		$\pi_{uk,t}$	$p_{uk,t}$	$\pi_{uk,t}$	$p_{uk,t}$
Full Sample	1245–2019	-22.273	3.011	0.108	147.424
Four Quarters	1245–1439	-13.231	-3.513	0.338	4.430
	1440–1633	-12.921	0.147	0.375	92.241
	1634–1827	-11.342	-0.837	0.326	19.149
	1828–2019	-5.471	0.759	0.960	93.633
Post-WWI	1919–2019	-4.549	-0.425	1.542	84.246
Post-WWII	1946–2019	-2.977	-0.567	2.879	36.429

Notes: The table shows Elliott (1999) $DF{-}GLSu$ and Q_T test statistics for a unit root in the inflation rate $\pi_{uk,t}$ and the log price level $p_{uk,t}$ for the entire sample, for each quarter of the sample, and for subsamples after 1919 and after 1946. The lag length is selected by the BIC. A constant is included. The lag lengths, test statistics, and critical values are obtained using the RATS-Estima procedure erstest.src. Evidence against the unit root appears in the left tail of each distribution. For $DF{-}GLSu$ some asymptotic critical values are -3.28 (1%), -2.98 (2.5%), -2.73 (5%), and -2.46 (10%). For Q_T these are 3.06 (1%), 3.80 (2.5%), 4.65 (5%), and 5.94 (10%). Overall there is evidence of a unit root in the price level but not in the inflation rate, which supports modelling the inflation rate as a stationary series over these spans of annual data.

Table 2. Priors on the Measurement Error

$$\pi_{uk,t} = \pi_t + m_t$$

$$m_t = \sqrt{\beta_\tau}\sigma_u u_t \qquad\qquad u_t \sim \mathcal{N}(0, 1)$$

Measurement-Error Parameters	Prior Distribution	Parameters θ_1	θ_2
Measurement-error Volatility σ_u^2	$\mathcal{IG}$	$\psi_u T_u$	σ_{PR}^2
Adjustment β_2 (1870–1914)	$\mathcal{B}eta$	5	3
Adjustment β_3 (1915–1946)	$\mathcal{B}eta$	3	5

Notes: The subsample indicator τ defines $\mathcal{B}eta$ priors placed on two elements of the vector of adjustment factors, β_τ, attached to the scale volatility, σ_u^2, of measurement error. The vector is $\beta_\tau = \left[\beta_1\ \beta_2\ \beta_3\ \beta_4\right]'$. We normalize β_1=1 for 1251–1869 and set β_4 to zero for 1947–2019 to be consistent with Cogley, Sargent, and Surico (2015). Our priors on the remaining measurement error adjustment parameters are $\beta_2 \sim \mathcal{B}eta$ $(5.0, 3.0)$ for 1870–1914 $\beta_3 \sim \mathcal{B}eta$ $(3.0, 5.0)$ for 1915–1946. These priors imply coverage intervals for $\beta_2 =$ [0.341, 0.636, 0.871] and $\beta_3 =$ [0.129, 0.364, 0.659] at the 5th, 50th and 95th quantiles. The scale volatility parameter is endowed with an $\mathcal{IG}$ prior. The prior is parameterized as $\theta_1 = \psi_u T_u$ and $\theta_2 = \sigma_{PR}^2$, where $\psi_u = 0.430{\times}10^{-2}$, $T_u = 696 = T{-}(2019{-}1947{+}1)$, and $\sigma_{PR}^2 = 0.21114^2$. The scale parameter θ_2 is equated to the variance of UK inflation on the Bank of England OBRA dataset from 1209 to 1244; see Thomas and Dimsdale (2017). The median of this prior is 0.018 with 5th and 95th quantiles of 0.007 and 0.072.

Table 3. Priors on the TVP-SV-AR(n)s

$$\pi_t = \alpha_{0,t} + \sum_{i=1}^{n} \alpha_{i,t}\pi_{t-i} + \xi_t\epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0,1),$$

$$\alpha_t \equiv \{\alpha_{0,t},\ \alpha_{1,t},\ \alpha_{2,t}, \ldots, \alpha_{n,t}\},$$

$$\alpha_t = \alpha_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}\big(0_{n+1},\, \Omega_\eta\big),$$

$$\ln \xi_t^2 \ = \ \ln \xi_{t-1}^2 \ + \ \sigma_\phi \phi_t, \qquad \phi_t \ \sim \ \mathcal{N}\big(0,\, 1\big).$$

Initial Model Parameters	Prior Distribution	Parameters θ_1	θ_2
α_0, Initial Intercept and Lags of AR$\big(n\big)$	$\mathcal{MN}$	$\underline{\alpha}$	$\underline{\Omega}_{\underline{\alpha}}$
Ω_η, Covariance Matrix of Innovations to α_t	$\mathcal{IW}$	n+2	$\kappa_n\underline{\Omega}_{\underline{\alpha}}$
$\ln \xi_0^2$, Initial SV of π_t Regression Error	$\mathcal{LN}$	$\kappa_\xi \underline{\xi}_n^2$	1.0
σ_ϕ^2, Scale Variance of Innovations to $\ln \xi_1^2$	$\mathcal{IG}$	0.5ν	0.5ς

Notes: Columns under θ_1 and θ_2 are parameters of the prior distributions. A multivariate normal ($\mathcal{MN}$) prior is placed on the initial intercept and lag coefficients α_0 of the TVP-SV-AR(n)s, n = 1, ..., 6, of the UK inflation data. Fixed coefficient AR(n)s are estimated by ordinary least square (OLS) on the sample from 1251 to 2019 to set the prior mean, $\underline{\alpha}$. The OLS covariance matrix of these parameters is the source of $\underline{\Omega}_{\underline{\alpha}}$. Its initial draw is from an inverse-Wishart ($\mathcal{IW}$) distribution with n+2 degrees of freedom and a scale matrix $\kappa_n\underline{\Omega}_{\underline{\alpha}}$. The scalar κ_n is chosen to achieve acceptance rates for α_t between 50 and 60 percent across n = 1, ..., 6. The initial regression error SV, $\ln \xi_0^2$, is endowed with a log normal prior, $\mathcal{LN}\Big(\ln\Big(\kappa_\xi \underline{\xi}_n^2\Big) - 0.5,\ 1.0\Big)$, where $\underline{\xi}_n^2$ is set to the OLS estimates of the variances of the residuals of the AR(n)s and κ_ξ = 75.0. The prior on the scale variance, σ_ϕ^2, is inverse-gamma ($\mathcal{IG}$) with shape parameter equal to half of ν = 2 and a scale parameter equal to half of ς = 0.123. This prior is equivalent to a (univariate) random variable distributed $\mathcal{IW}$ with two degrees of freedom and a scale variance of ς. See Online Appendix C for further details.

Table 4. Evaluation of Fit of the TVP-SV-AR(n)s

	AR(1)	AR(2)	AR(3)	AR(4)	AR(5)	AR(6)
ln MDD	638.21	643.18	673.62	667.59	652.72	650.82
WAIC	-1431.95	-1457.55	-1487.87	-1483.30	-1487.75	-1465.82

Notes: The first row contains log marginal data densities, ln MDDs, of the TVP-SV-AR(n)s. The modified harmonic mean estimator of Geweke (2005) is employed to calculate the MDDs. The MDDs represent evidence the data have about the TVP-SV-AR(n)s. The second row reports the Widely Applicable Information Criterion (WAIC) developed by Watanabe (2010). The WAIC, also referred to as the Watanabe-Akaike-IC, is an estimate of the predictive loss of a TVP-SV-AR(n). This notion of predictive loss equals twice the difference between the sum of the posterior variances of the log predictive likelihoods and the mean of the log predictive likelihoods following the advice of Gelman *et al* (2014). Estimates of the likelihoods are obtained from the predictive steps of the Kalman filter. The sum of the posterior variances of the likelihood is an estimate of the effective dimension of the parameter vector of the TVP-SV-AR(n)s, which serves as the penalty term of the WAIC.

Figure 1: The UK Price Level and Inflation, 1251–2019

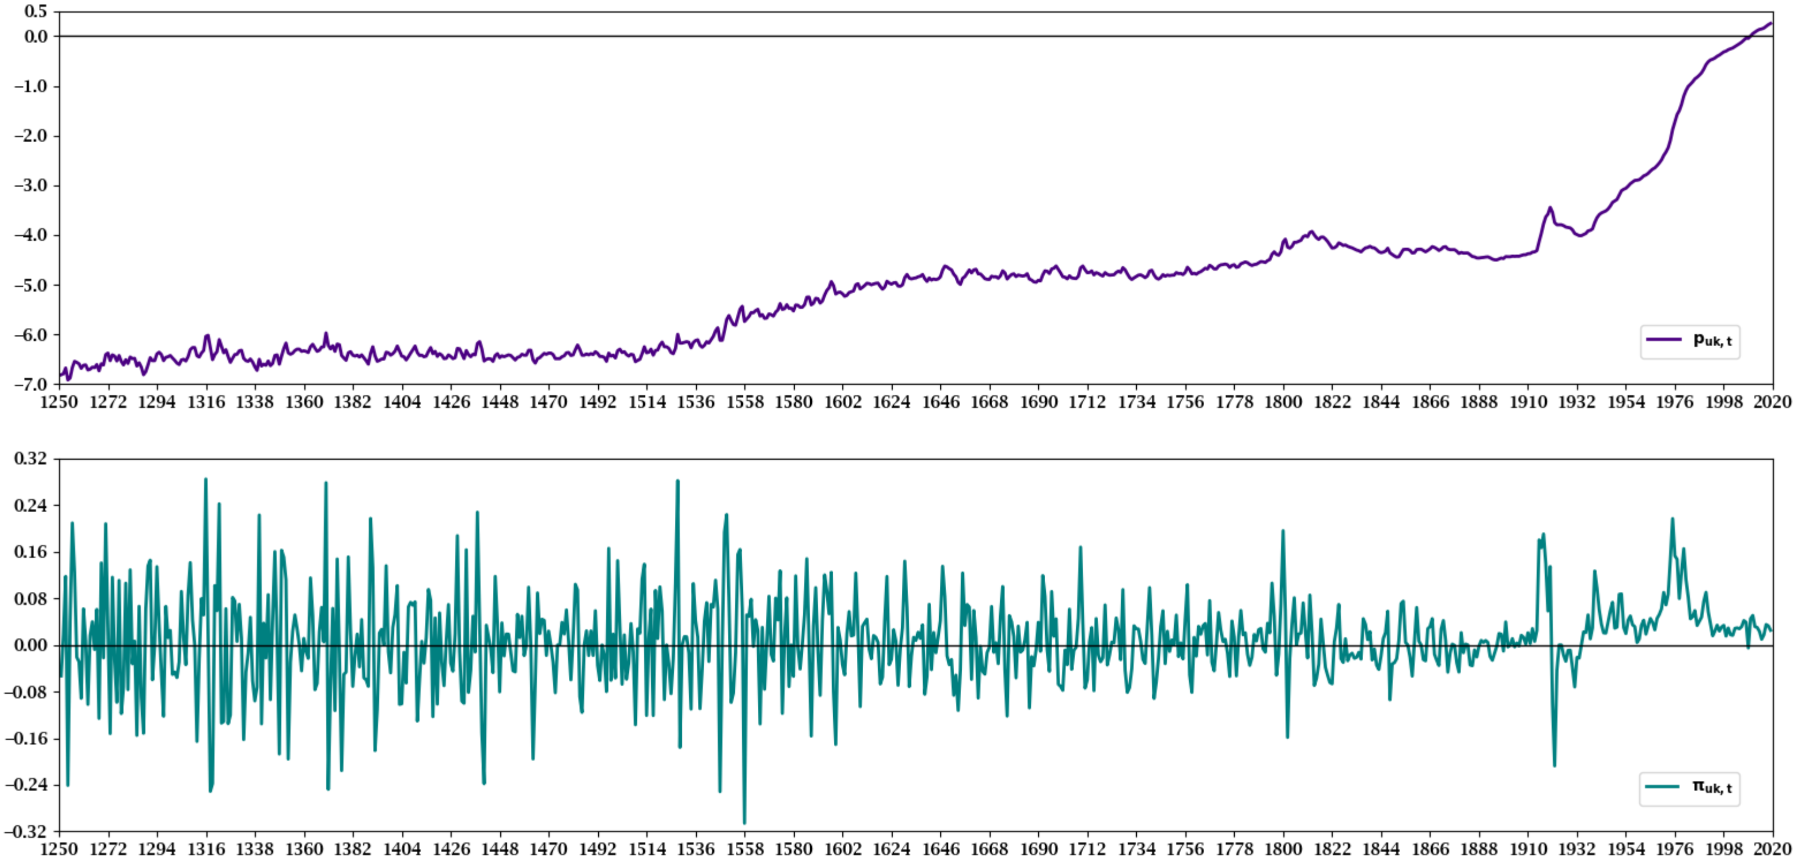

Notes: The top panel display the natural log of the UK price level, $p_{uk,t} = \ln(P_{uk,t}/100)$. The bottom panel plots the inflation rate of UK price series, $\pi_{uk,t} = p_{uk,t} - p_{uk,t-1}$.

Figure 2: Ten-, Twenty-, and Thirty-Year Moving Averages of UK Inflation

Notes: The top panel contains the 10-year moving average (MA) of UK inflation, $\pi_{10,uk,t}$, from 1254 to 2019 surrounded by shadings that are 68% Newey-West confidence bands computed with one lag. Two lags are used to construct the 68% Newey-West confidence bands of the 20-year MA of UK inflation, $\pi_{20,uk,t}$, that begins in 1264 and ends in 2019 and appear in the middle panel. The 30-year MA of UK inflation, $\pi_{30,uk,t}$, is on a 1274–2019 sample, are displayed in the bottom panel, and relies on three lags to calculate 68% Newey-West confidence bands.

Figure 3: Measurement Error Parameter Densities and Smoothed States

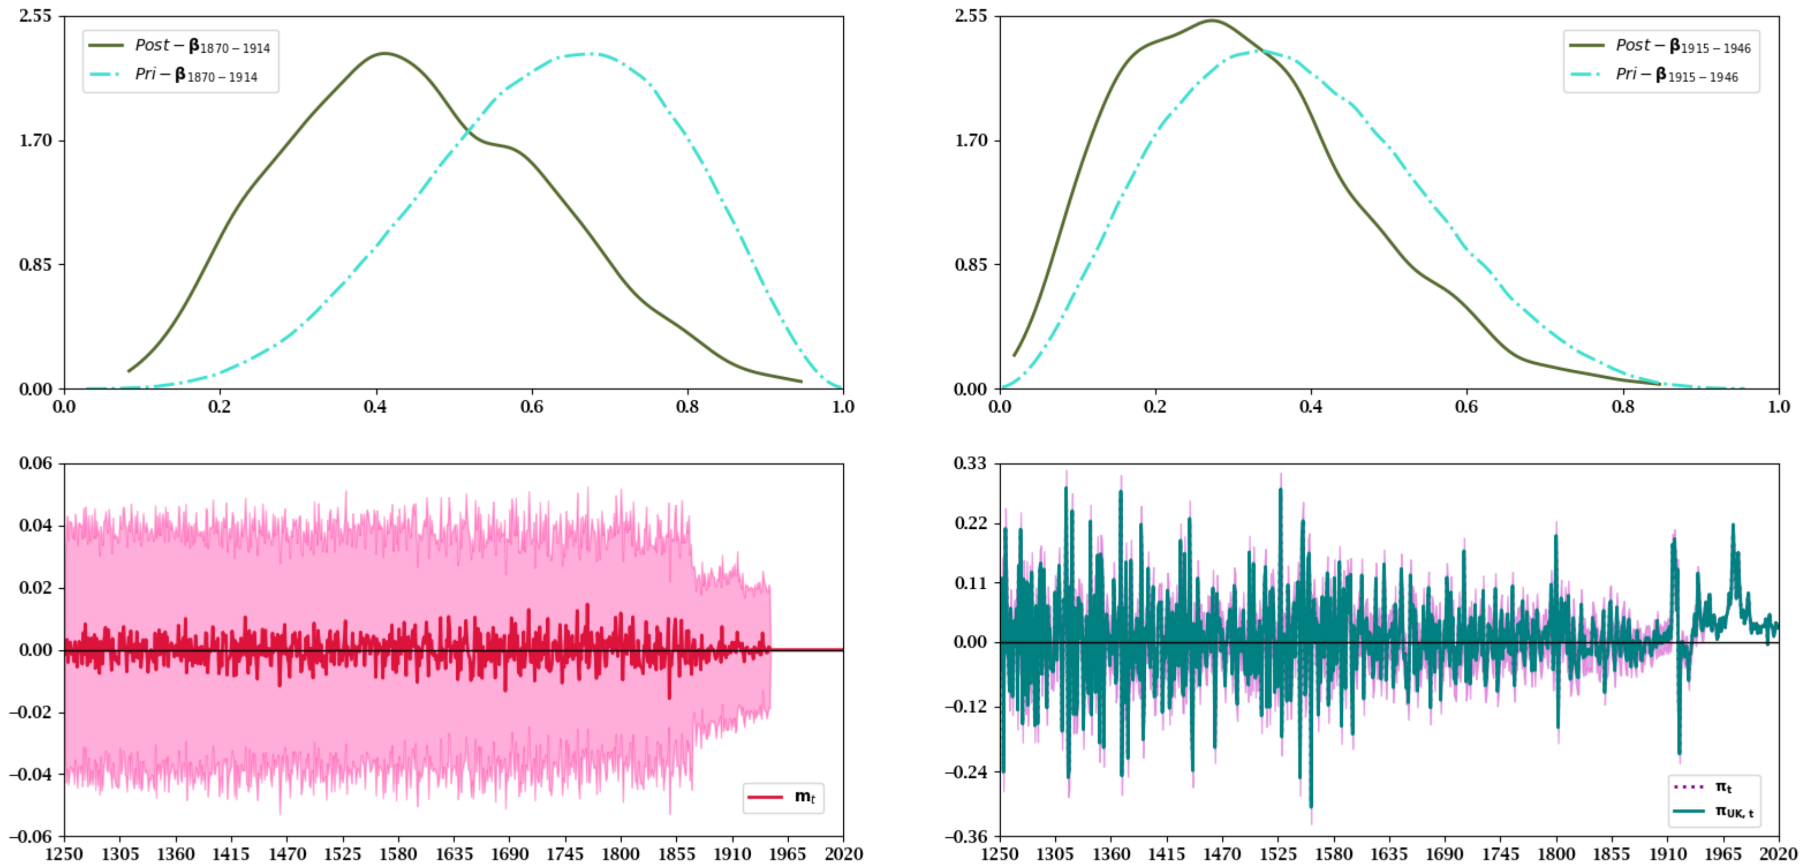

Notes: Prior and posterior densities of the β_τs appear in the top row of panels, τ = 1870-1914 and 1915-1946. The (olive) solid lines are posterior densities. The densities are constructed using a normal kernel and methods described by Silverman (1986). Posterior densities are the solid (olive) lines. The (turquoise) dot-dash lines are densities of the prior distributions. The latter distributions are simulated using the priors for the β_τs listed in table 2. The median of the posterior of smoothed measurement error, m_t, is the solid (red) line in the bottom left panel. Surrounding m_t are (pink) shadings, which are 68% uncertainty bands. The bottom right panel plots $\pi_{uk,t}$, as the solid green line, the median of the posterior of smoothed true inflation, π_t, is the dotted (magenta) line, and (orchid) shadings represent 68% confidence bands.

Figure 4: Posterior Moments of the TVP-SV-AR(3) on UK Inflation, 1251–2019

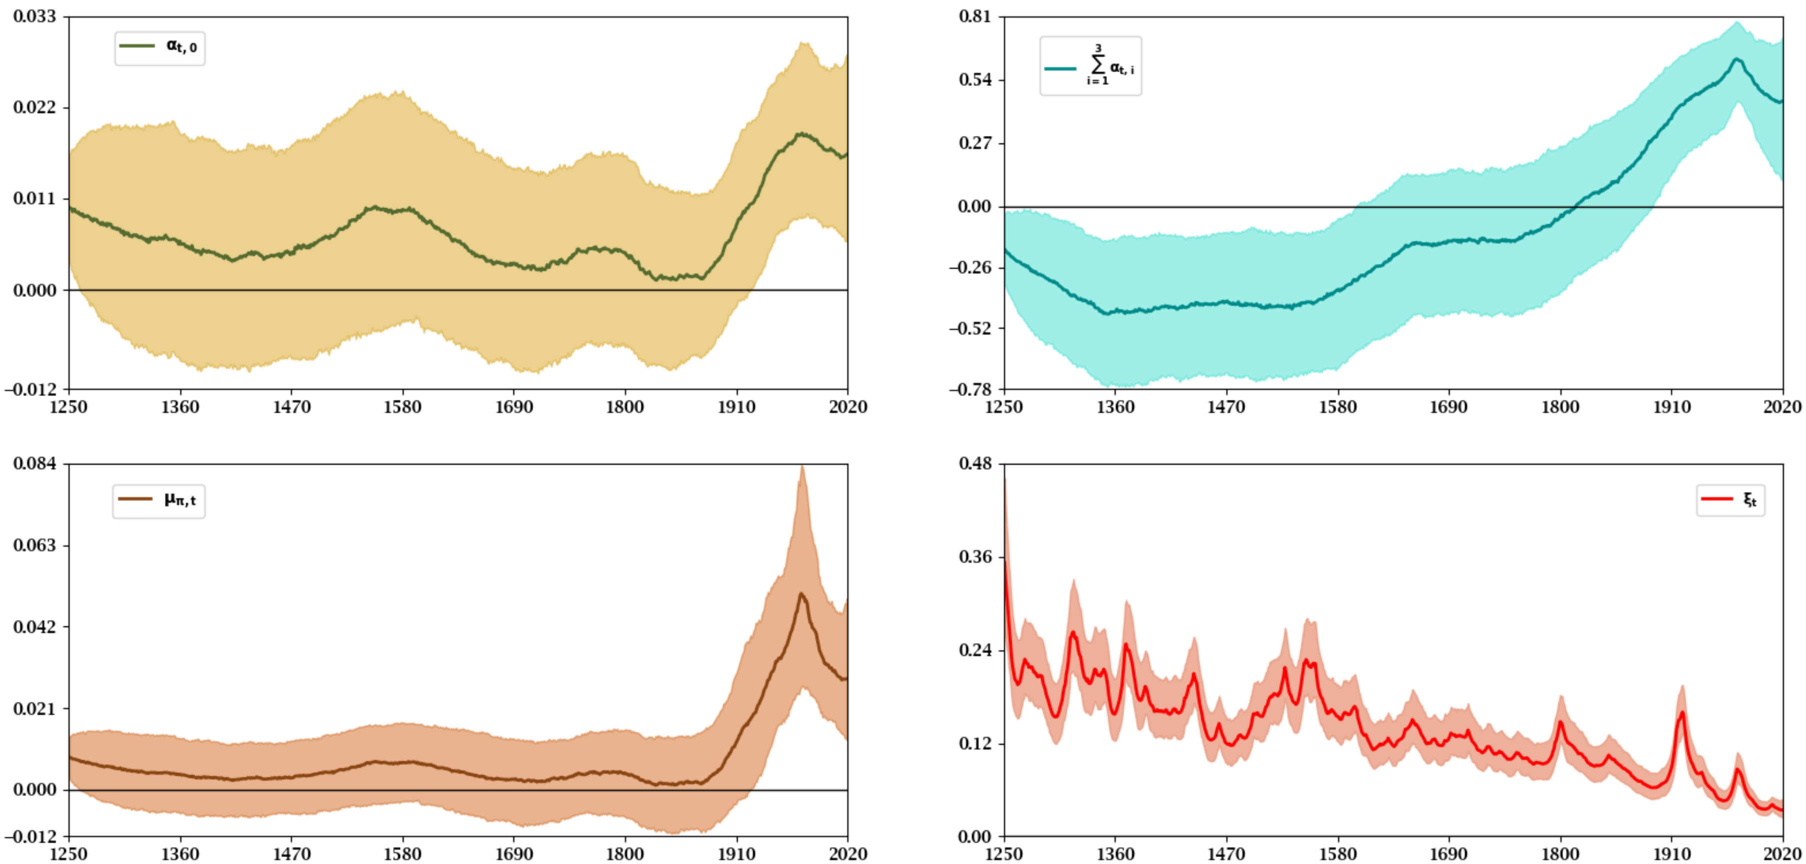

Notes: The top left panel contains the posterior median of the time-varying intercept, $\alpha_{0,t}$. The posterior median of the sum of the lag TVPs, $\sum_{\ell=1}^{3} \alpha_{\ell,t}$, is found in the top right panel. A plot of the time-varying conditional mean of UK inflation is depicted in the bottom left panel as the posterior median of $\mu_{\pi,t} = \alpha_{0,t}/(1 - \sum_{\ell=1}^{3} \alpha_{\ell,t})$. The SV of UK inflation is displayed in the bottom right panel. The four panels also display shadings that are 68% Bayesian credible sets (*i.e.*, 16% and 84% quantiles) of the TVPs and SV.

Figure 5: 1-Year Ahead Expected UK Inflation and Its Ex Post Forecast Error, 1251–2019

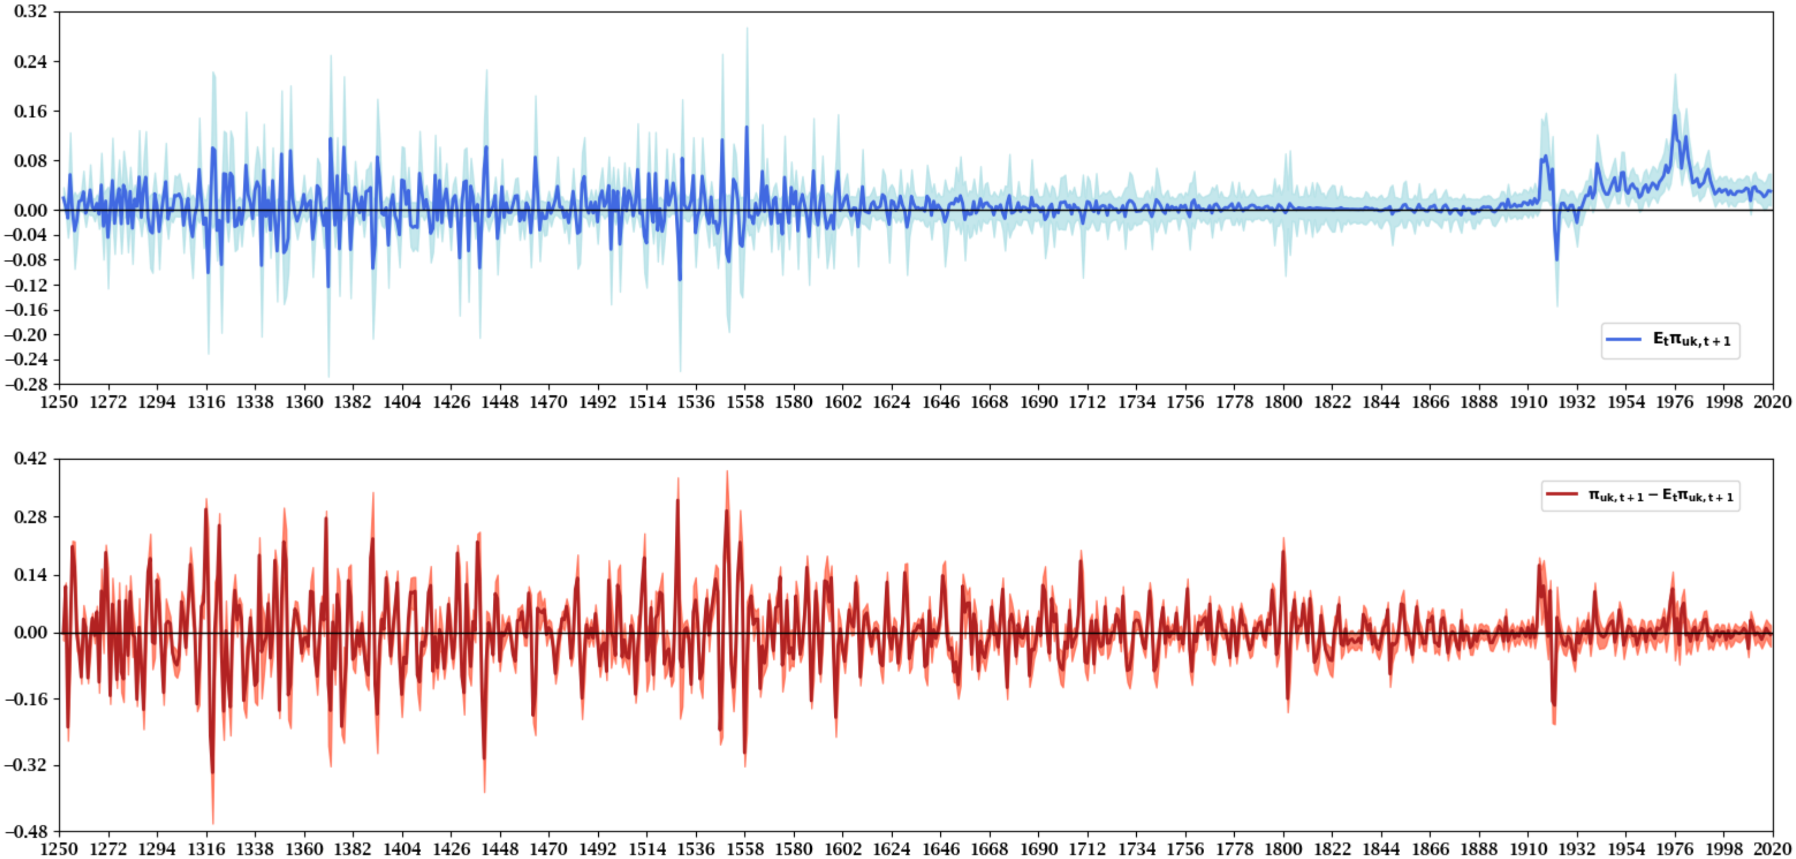

Notes: The top panel plots median 1-year ahead expected UK Inflation, $\mathrm{E}_t\pi_{uk,t+1}$. Expected inflation is estimated using the Kalman filter, $\pi_{uk,t}$, and the posterior distribution of the TVP-SV-AR(3) with measurement error. The bottom panel contains the ex post forecast error, $\pi_{uk,t+1} - \mathrm{E}_t\pi_{uk,t+1}$. The panels also contain shadings that are 90% Bayesian credible sets (*i.e.*, 5% and 95% quantiles).

Figure 6: Predictability of UK Inflation 1-, 2-, 3-, and 5-Years Ahead, 1251–2019

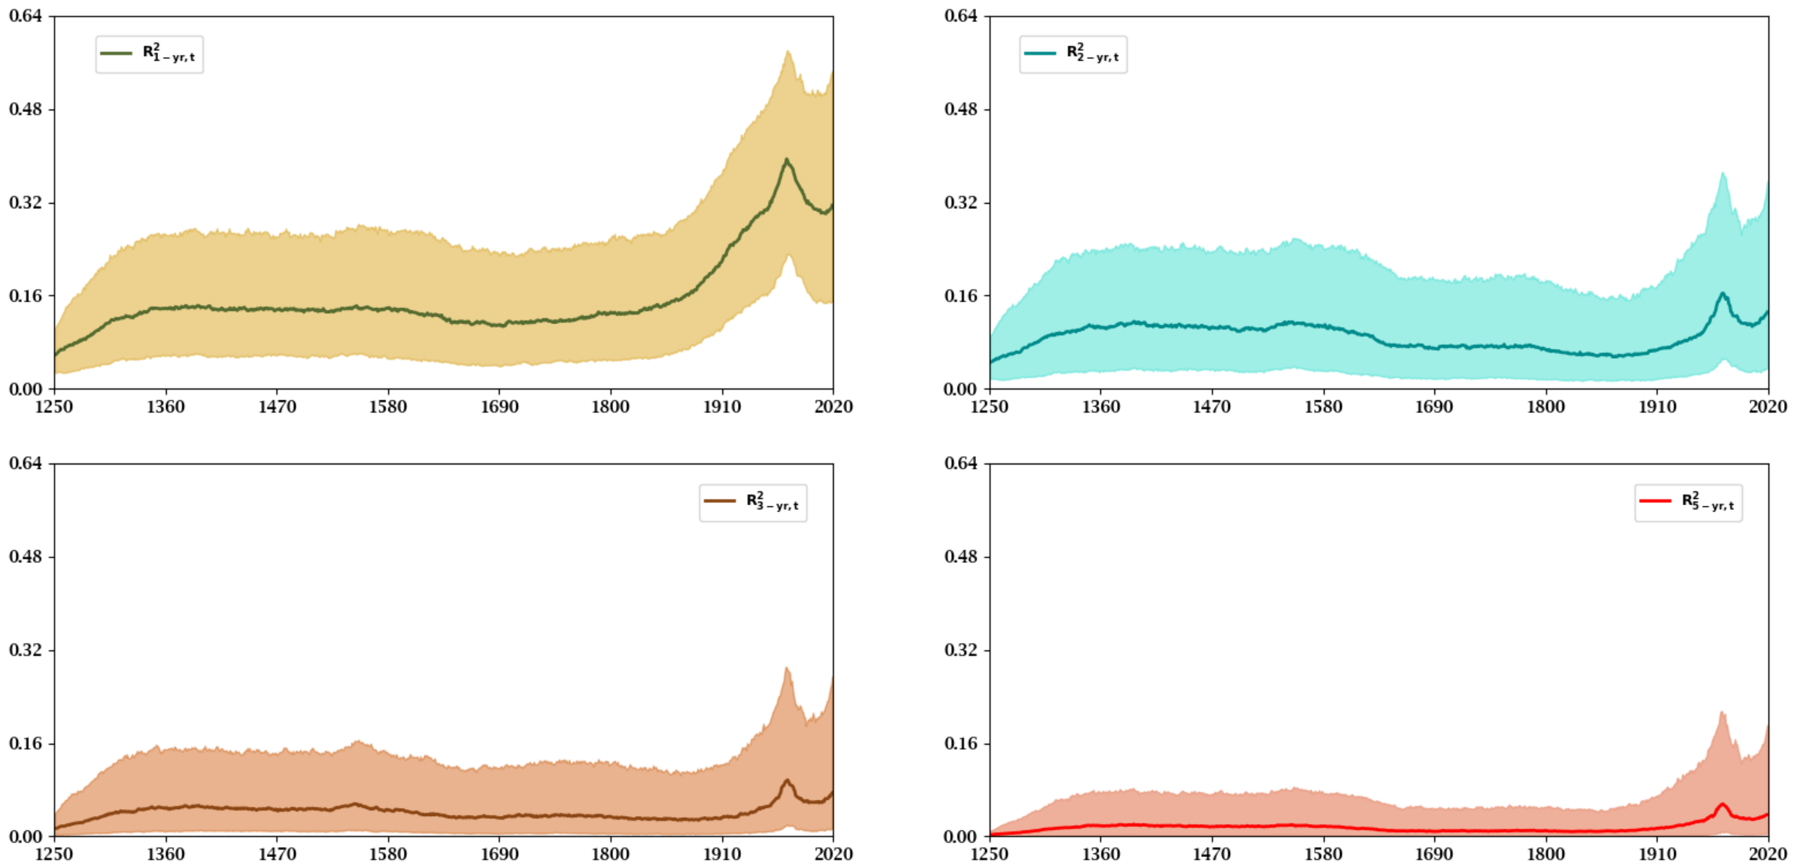

Notes: The top left panel plots the median 1-year ahead R-square statistic, R^2_{1t}, which is computed as described in appendix D, from 1251 to 2019. The shadings around R^2_{1t} are 68% Bayesian credible sets. Similarly, median R^2_{2t}, R^2_{3t}, and R^2_{5t} appear in the top right, bottom left, and bottom right panels along with 68% Bayesian credible sets as the shadings. The R^2_{ht} statistics are computed using the posterior distribution of the TVP-SV-AR(3) with measurement error.

Figure 7: UK Price-Level Instability, 1251–2019

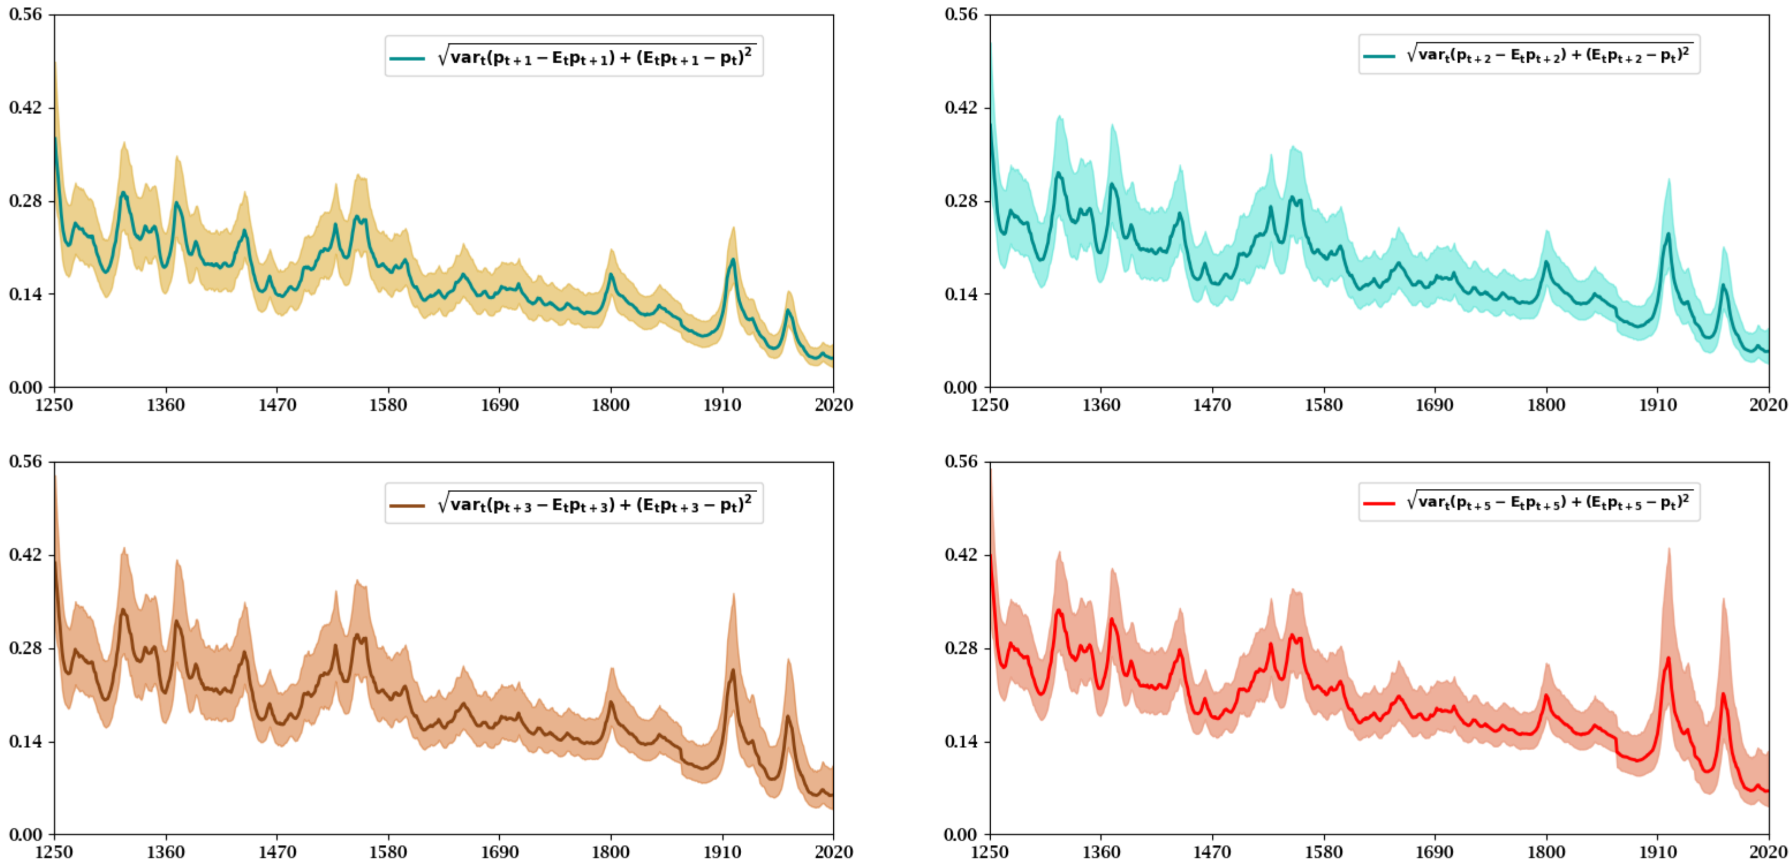

Notes: The top left panel plots the median of the 1-year ahead square root of the sum of the conditional variance and the squared, conditional mean, $\sqrt{\text{var}_t\left(p_{uk,t+1} - E_t p_{uk,t+1}\right) + \left(E_t p_{uk,t+1} - p_{uk,t}\right)^2}$, from 1251 to 2019; see section 5.4 and appendix E for details. The shadings around this statistic are 68% Bayesian credible sets. The median 2-, 3-, and 5-year ahead price-level stability statistics are displayed in the top right, bottom left, and bottom right panels along with shadings that are 68% Bayesian credible sets. The price-level stability statistics are computed using the posterior distribution of the TVP-SV-AR(3) with measurement error.

Figure 8: UK Nominal and Real Short- and Long-Term Interest Rates

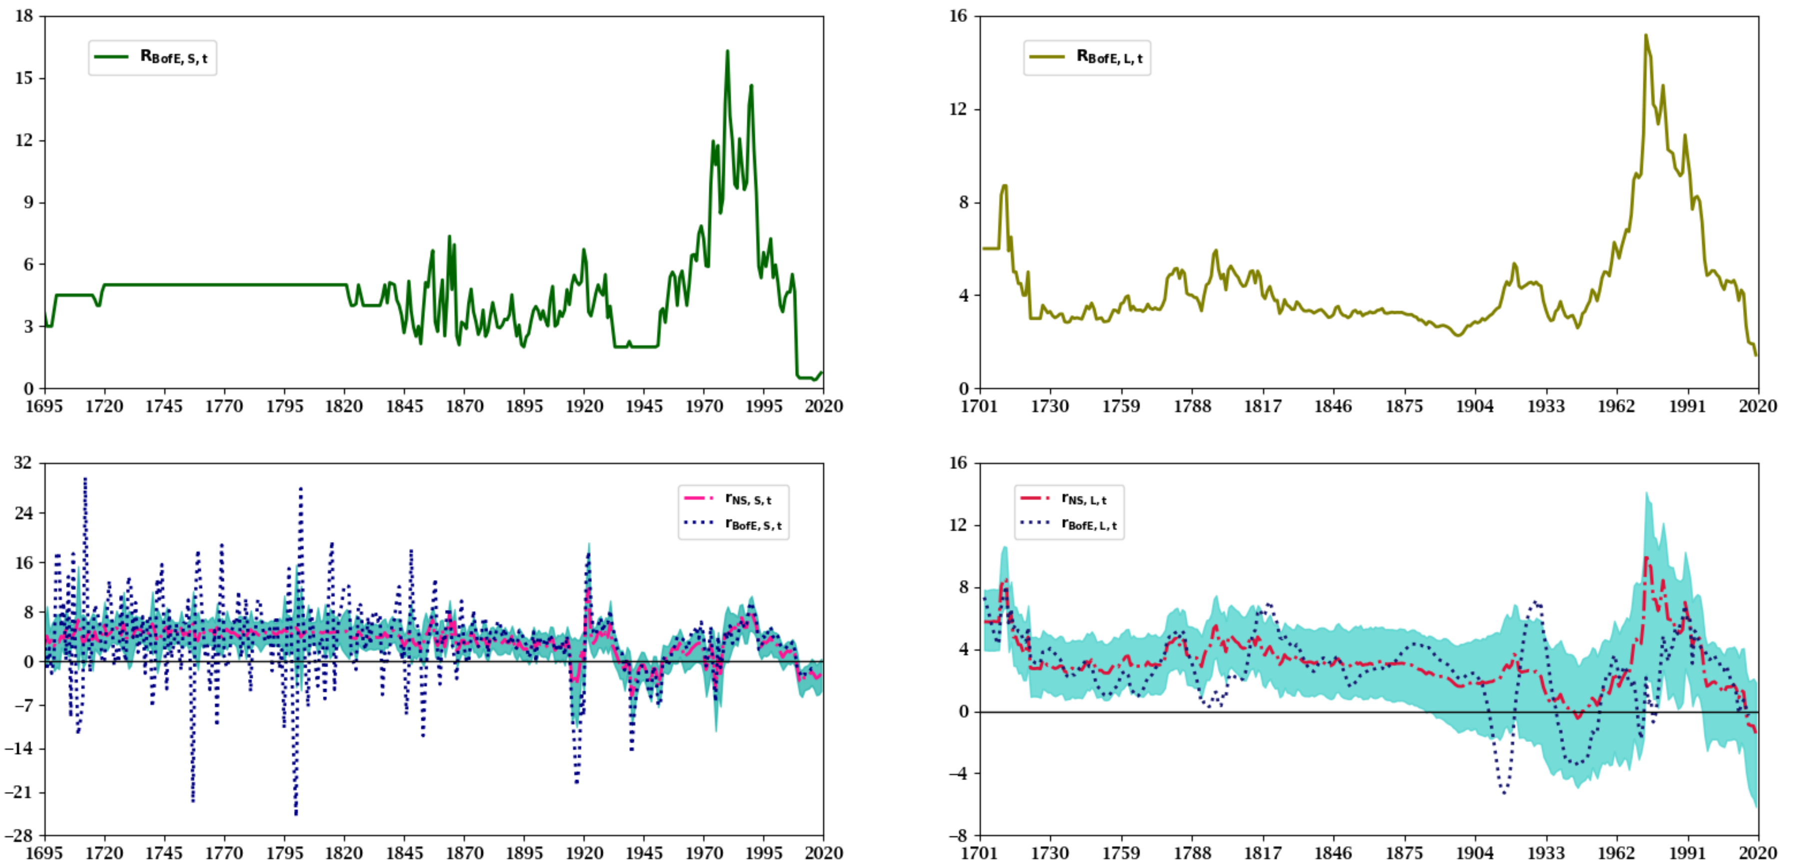

Notes: The top row plots UK nominal short- and long-term interest rates; see section 6 for details. The bottom row depicts ex ante short- and long-term real rates, $r_{NS,i,t} = R_{\mathrm{BofE},i,t} - E_t\pi_{uk,t+1}$, in the left and right panels, where $i = S$(short) and L(long). The $r_{NS,S,t}$ ($r_{NS,L,t}$) begins in 1695 (1703) and ends in 2019. One-year ahead expected UK inflation is computed on the posterior distribution of the TVP-SV-AR(3) with measurement error. In the bottom row of panels, the shadings are 90% Bayesian credible sets.