

Measuring Price-Level Uncertainty and Instability in the U.S., 1850-2012*

Timothy Cogley[†] and Thomas J. Sargent[‡]

August 4, 2014

Abstract

We measure price-level uncertainty and instability in the U.S. over the period 1850-2012. Major outbreaks of price-level uncertainty and instability occur both before and after World War II, alternating with three price-level moderations, one near the turn of 20th century, another under Bretton Woods, and a third in the 1990s. There is no evidence that the price level was systematically more stable or less uncertain before or after the Second World War. Moderations sometimes involved links to gold, but the experience of the 1990s proves that a well-managed fiat regime can achieve the same outcome.

JEL CLASSIFICATION: E31, C22

KEY WORDS: Price stability, inflation uncertainty, deflation risk, gold standard, Bretton Woods, stochastic volatility, measurement error, nonlinear signal extraction

1 Introduction

Figure 1 portrays annual data on the wholesale price level and inflation rate in the United States for the period 1798-2012. We study wholesale price indices because Christopher Hanes (2006) says they are more reliable than consumer price indices for the period before World War I. According to Hanes, statisticians collected more data on wholesale prices before 1914 than on retail prices or rents, thus compiling better raw materials for price indices. Since comparability of data across periods is important for us, we examine wholesale price indices for the period after 1914 as well.

*For comments and suggestions, we thank Christopher Sims, Nicholas Bloom, Gabriel Mathy, Emi Nakamura, François Velde, Mark Watson, two referees, and seminar participants at the Bank of France, Boston College, and the summer 2013 SITE workshop on uncertainty and volatility.

[†]New York University, tim.cogley@nyu.edu

[‡]New York University and Hoover Institution, thomas.sargent@nyu.edu




Figure 1: Wholesale Price Level and Inflation

The behavior of the wholesale price level, shown in the left panel, changed dramatically after the Second World War. Between 1798 and 1940, average inflation was essentially zero, and there was no sustained growth in the price level. One very rough stylized summary of pre-World War II US monetary history is that the price level surged during wars but fell afterward, thus providing long-run price stability. In contrast, after 1945 there was no rollback of prices, average inflation became positive, and the price level ascended. Because of that, one might suspect that price-level uncertainty was greater after World War II.

But to the extent that average inflation was predictable after World War II, a positive average would contribute to sustained growth in the price level without increasing uncertainty. The visual impression conveyed by an upward drift might therefore exaggerate the degree of uncertainty. Predictable variation would undermine price level *stability*, however, so that the price level might have been less stable after 1945, in the sense of total variability, even if were no more uncertain.

In addition, although average inflation was close to zero before 1940, there was a lot of volatility. Notice in particular that the rollbacks after the War of 1812 and the Civil War lasted for decades and eventually generated more cumulated postwar deflation than wartime inflation.¹ Uncertainty about the extent and duration of a rollback – especially about the degree of overshooting – could at least partially offset any presumed benefits. Furthermore, data on inflation, exhibited in the right panel,

¹ For instance, during the deflation following the Civil War, the price level overshot its 1860 value on the way down and eventually fell 30 percent below the prewar mark. After reaching a trough in 1896, another 10 years passed before the price level returned to its 1860 value.

suggest that volatility might have been lower *after* 1950, despite the Great Inflation and the 2008 financial crisis. For these reasons, whether price-level uncertainty increased or decreased after 1945 is unclear.

In this paper, we estimate a flexible statistical model for wholesale price inflation and use it to measure price-level uncertainty and instability before and after World War II. Our starting point is a simple and elegant model for inflation developed by Stock and Watson (2007). A version of their model presents state equations for a nonlinear state-space representation. Motivated by the work of Christina Romer (1986a,b), we also append measurement error to acknowledge noise in pre-World War II data. As Romer pointed out, one of the chief obstacles to assessing volatility across historical eras is that data are not directly comparable. Long historical time series such as the one shown in figure 1 are typically spliced together from a variety of sources that define and measure objects differently. Older data are usually noisier, and measurement error amplifies volatility.

To level the playing field, Romer posed and solved a ‘noise-creation’ problem, constructing a comparably noisy series for the postwar era by imitating prewar measurement procedures using postwar data sources. After estimating the variances of comparably noisy measures, she concluded that much of the apparent decline in U.S. unemployment and GDP volatility is a consequence of better measurement. Similarly, Hanes (1998) solved a noise-creation problem for U.S. wholesale price indices, constructing a noisy post-World War II series whose stochastic properties are consistent with prewar data.

This paper extends Hanes’s analysis by solving a signal extraction problem for prewar data, thus improving estimates of historical volatility. We begin by describing a nonlinear state-space model in which prewar inflation is measured with error. So that Romer’s conclusion is not hardwired, our model allows true but hidden volatilities to vary over time, as in Stock and Watson (2007). We solve the signal extraction problem by applying Bayes’s theorem. If the state-space model were linear and conditionally Gaussian, this would reduce to an operation closely related to Kalman smoothing. Because the model is nonlinear in state variables, however, we use a Markov Chain Monte Carlo (MCMC) algorithm to simulate the posterior distribution for the model’s parameters and hidden variables. From that posterior, we deduce the posterior distribution for conditional variances and second moments of future price levels, and then compare our measures of uncertainty and instability at various dates.

Specification of the measurement-error process and identification of its parameters are key elements of our analysis. Because measurement error is a latent variable weakly

restricted by theory, how to specify it and how to identify its parameters are not obvious. Here we exploit Hanes’s consistent postwar measure of wholesale price inflation. For the period 1948-1990, Hanes’s noisy measure overlaps with the putatively noise-free² measure constructed by the Bureau of Labor Statistics. The availability of two inflation measures – one noisy and one clean – sharply identifies measurement error for years in which they overlap, allowing precise estimation of its parameters.

After solving this signal extraction problem, we find that uncertainty was highest and the price level least stable during the Civil War, the two World Wars and interwar years, and during the Great Inflation and Great Recession. There were also three great price-level moderations, one near the turn of the 20th century, another under Bretton Woods, and a third during the 1990s. Because periods of high and low volatility occur both before and after the Second World War, we find no compelling evidence that the price level was systematically more stable or less uncertain in either era. There was less uncertainty and more stability during moderations than at other times, but moderations occur both before and after the Second World War.

2 Data

Our data are annual and are assembled from four sources, Warren and Pearson (1933), Hanes (1998), and two series from the Bureau of Labor Statistics (*BLS*). Hanes’s (2006) chapter in the Millennial edition of the *Historical Statistics of the United States (HSUS)* describes the historical data and documents their sources, and the data themselves can be found in the accompanying tables.³ Postwar data were downloaded from the *Federal Reserve Economic Database (FRED)*.

Data for the period 1797-1890 come from Warren and Pearson (1933). The dotted line in figure 2 depicts their measure of inflation. We use this series primarily as a training sample to calibrate aspects of the prior.

Hanes (1998) replicated and extended Warren and Pearson’s series for the periods 1860-1941 and 1947-1990 (see the solid line in figure 2). His objective was not to construct the best measure of wholesale price inflation for these periods but by following Warren and Pearson as closely as possible to construct a series consistent with theirs. For years in which their measures overlap, the correlation is 0.997, confirming that Hanes’s procedures are faithful to those of Warren and Pearson.⁴ Among other things,

² At least this is the conventional assumption in macroeconomics.

³ See <http://hsus.cambridge.org/HSUSWeb/toc/showChapter.do?id=Cc>.

⁴ Indeed, for the overlapping years, Warren and Pearson’s dotted line in figure 2 hides behind Hanes’s solid line.




Figure 2: US Wholesale Price Inflation: Overlapping Sources

we interpret ‘consistency’ to mean that the stochastic properties of measurement errors in Hanes’s data are the same before and after the Second World War.

The United States government began constructing wholesale price indices in 1902 and produced retrospective measures back to 1890. We make use of two data series from the *BLS*, one from *HSUS* covering the period 1890-1997 and another from *FRED* covering the period 1913-2013. The corresponding inflation measures are shown as dashed and dashed-dotted lines, respectively, in figure 2. The two *BLS* series are very highly correlated for the period in which they overlap (correlation 1914-1997 = 0.986), but because *FRED* data were revised more recently, we believe they are more accurate for the period after the Second World War. Differences are slight in any case.

Although the Warren-Pearson-Hanes series is our primary noisy measure, there is a gap from 1942-1947 that we fill with *BLS-HSUS* data. That the *BLS-HSUS* measure is comparable to the Warren-Pearson-Hanes series is not obvious, but the following defense can be offered. Warren and Pearson’s objective was to construct a retrospective measure comparable to the *BLS* index begun in 1890, and Hanes’s goal was to mimic Warren and Pearson (see Hanes 2006). To the extent that both succeeded, one might therefore expect early *BLS-HSUS* data to correlate highly with Hanes’s measure. Indeed, for the period 1914-1941, the point estimate in a regression of Hanes’s inflation measure onto that of *BLS-HSUS* is 1.02, its standard error is 0.04, and the R^2 is 0.989. Thus the hypothesis that $y_{BLS-HSUS}$ is an unbiased predictor of y_{Hanes} (i.e. that the regression coefficient is 1) cannot be rejected, and the standard deviation




Figure 3: US Wholesale Price Inflation: Spliced Sources

of the approximation error $y_{Hanes} - y_{BLS-HSUS}$ is 11 percent of the standard deviation of y_{Hanes} . Extrapolating this close association to 1942-1947 rationalizes bridging the gap between Hanes's two series with *BLS-HSUS* data. To be on the safe side, we also checked the influence of this patch by estimating a version of the model that treats 1942-1947 as missing, and our conclusions were the same.

After discarding redundant data, we are left with the data shown in figure 3 that we will use to calibrate priors and estimate the model. We construct a noisy inflation measure for 1798-1990 by splicing data from Warren-Pearson-Hanes and *BLS-HSUS*.⁵ Our clean post-World War II measure (1948-2012) comes from *BLS-FRED*. Thus, two inflation measures – one noisy and one clean – are available for 1948-1990, and a single measure – either noisy or clean – is available for all other years. That a noise-ridden series comparable to that for the period before World War II overlaps with a clean measure for much of the postwar period underlies our strategy for recognizing measurement error.

For guidance on how to represent measurement error, we examine the difference between the noisy and clean measures $y_{Hanes} - y_{BLS-FRED}$ for the period 1948-1990. On the assumption that *BLS-FRED* data are error free, this difference identifies noise in Hanes's measure of inflation, and its partial sum identifies noise in his measure of

⁵ In principle, precision could be enhanced by using all the noisy data, but since the various noisy measures are highly correlated we suspect the gain would be slight. Thus, we simplify by using a single noisy series.

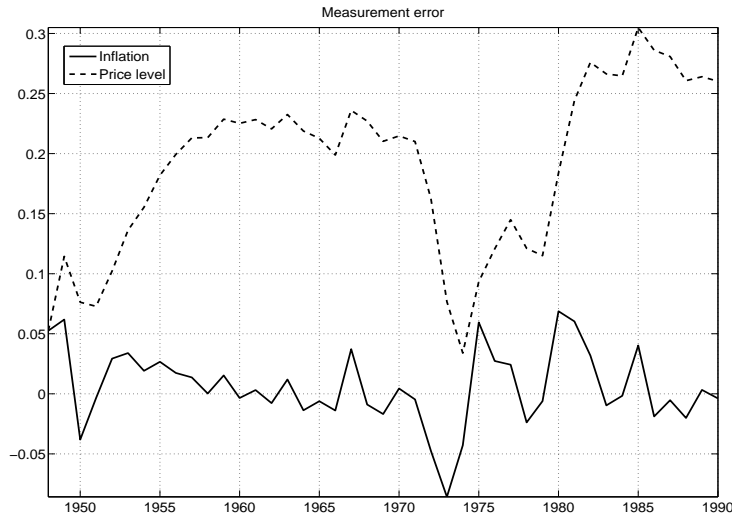


Figure 4: Measurement error in inflation and the log price level, 1948-1990

the log price level. As shown in figure 4, noise in inflation (depicted by the solid line) appears to be stationary, centered near zero, and weakly autocorrelated, while noise in the log price level (shown as a dashed line) appears to have a unit root. Thus, in what follows, we assume that our noisy measure of inflation is the sum of true inflation plus a mean zero, stationary measurement error.

3 A nonlinear state-space model with measurement error and time-varying volatilities

Following Stock and Watson (2007) and Shephard (2013), the state equations for our statistical model are

$$\begin{aligned}
 \pi_t &= \mu_t + \sqrt{r_t} \varepsilon_{\pi t}, \\
 \mu_t &= \mu_{t-1} + \sqrt{q_t} \varepsilon_{\mu t}, \\
 \ln r_t &= \ln r_{t-1} + \eta_{rt}, \\
 \ln q_t &= \ln q_{t-1} + \eta_{qt},
 \end{aligned}
 \tag{1}$$

where π_t is inflation, μ_t is trend inflation, and r_t and q_t are stochastic volatilities that evolve as geometric random walks. The innovations $\varepsilon_{\pi t}, \varepsilon_{\mu t}$, are standard normal, serially uncorrelated, and independent of all the other shocks in the model. Similarly, the log volatility innovations η_{rt} and η_{qt} are iid normal with mean zero and covariance

matrix

$$W = \begin{bmatrix} \sigma_r^2 & \sigma_{rq} \\ \sigma_{rq} & \sigma_q^2 \end{bmatrix}. \quad (2)$$

Introduced by Shephard (2013), this parameterization extends that of Stock and Watson (2007) by allowing correlated increments to $\ln q_t$ and $\ln r_t$. Shephard reported that this extension improved the model’s fit to quarterly post World War II US inflation.⁶

This model has three features that we believe an adequate statistical representation of inflation requires. First, to fit both the 19th century, when average inflation was close to zero, and the period after World War II, when average inflation was positive and variable, the model must include a stochastic trend. Second, there must also be a transient component to fit the short-term volatility seen at various times throughout the sample. Third, so that volatility can change, innovations to the two components must have time-varying variances.

Our main extension of the Stock-Watson-Shephard model aims to confront the measurement issues raised by Romer. To address her concerns, we append an additional state equation for a serially correlated measurement error m_t ,

$$m_t = \rho_m m_{t-1} + \sigma_m \varepsilon_{mt}, \quad (3)$$

where ε_{mt} is iid standard normal and independent of the other shocks in the model. The measurement equation for our state-space model is

$$y_t = C_{mt} \begin{bmatrix} \pi_t \\ m_t \end{bmatrix}, \quad (4)$$

where y_t represents observations on inflation. Our identifying assumptions are that the post-1948 *BLS-FRED* data are noise free and that the measurement error parameters ρ_m and σ_m are constant throughout Hanes’s sample. This is how we interpret the statement that Hanes’s prewar and postwar measurement procedures are consistent. Under these assumptions, the measurement arrays y_t and C_{mt} are defined as

- $y_t = y_{nt}$, $C_{mt} = [1 \ 1]$ when a single noisy observation y_{nt} is available (before 1948);
- $y_t = y_{ct}$, $C_{mt} = [1 \ 0]$ when a single clean observation y_{ct} is available (after 1990);

⁶ As we shall see, we find less evidence of correlation in log-volatility innovations, perhaps because our data are annual, measure wholesale prices instead of consumer prices, cover a longer time span, and are noisier than Shephard’s. We also estimated a version based on Stock and Watson’s original parameterization that constrains $\sigma_{rq} = 0$ and obtained results that are similar to those reported below.

- $y_t = [y_{nt} \ y_{ct}]'$, $C_{mt} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$ when both noisy and clean observations are available (1948-1990).

4 Priors

Our first task is to estimate the latent states π_t , μ_t , r_t , q_t , and m_t , the log-volatility innovation variance W , and the measurement error parameters ρ_m and σ_m . We do this via Bayesian methods.

Priors are wanted for the initial states μ_0 , π_0 , r_0 , q_0 , and m_0 as well as the parameters W , ρ_m , and σ_m . The transition equations then imply priors for the remaining states μ_t , π_t , q_t , r_t , and m_t . We assume priors for μ_0 , π_0 , r_0 , q_0 , m_0 , W , ρ_m , and σ_m to be mutually independent, and specify marginal priors for each. Many aspects of these priors are calibrated using a training sample consisting of measured inflation for 1798-1849.

For instance, the prior for the initial state (μ_0, π_0) is normal with mean equal to the training sample average (-0.90 percent per annum) and variance

$$P_0 = \begin{bmatrix} 0.15^2 & 0 \\ 0 & 0.025^2 \end{bmatrix}. \quad (5)$$

Since prior credible sets for π_0 and μ_0 are roughly (-0.3,0.3) and (-0.05,0.05), respectively, the prior is weakly informative about initial inflation.

We also adopt normal priors for $\ln r_0$ and $\ln q_0$, the logs of the initial innovation variances for the transitory and persistent components of inflation, respectively. The prior mean for $\ln r_0$ is the log of the training sample variance for measured inflation (-4.97), while that for $\ln q_0$ is the log of the training sample variance divided by 25 (-8.19). We set the prior standard deviation for both to 5, a value that is huge on a log scale. This makes the prior on the initial values $\ln r_0$ and $\ln s_0$ very weakly informative.

For the measurement error parameters ρ_m and σ_m , we adopt weakly informative priors portrayed by dashed lines in figure 5. The prior for ρ_m is normal with mean zero and standard deviation 0.45, thus centering on a white-noise specification and concentrating the preponderance of prior mass in the stationary region.⁷ For σ_m , we adopt an inverse-gamma prior whose mode equals 50 percent of the training sample standard deviation for measured inflation (70.71 percent of the variance), thereby expressing an initial belief that the Warren-Pearson-Hanes data are very noisy. Not wanting to hardwire this prior belief into a posterior, however, we set the prior degrees of freedom

⁷ Truncating the prior at the boundary of the stationary region was unnecessary because all the posterior draws were in the interior.




Figure 5: Priors and posteriors for measurement-error parameters

to 2, so that a centered 95 credible set ranges from 28 percent to more than 100 percent of the training sample standard deviation. The result is an $IG_1(0.051, 2)$ specification.

A prior for W , the covariance matrix for log-volatility innovations, is harder to calibrate using the training sample. Instead, we adopt an informative inverse-Wishart prior that is inspired by Stock and Watson’s calibration. We start with their parameter for the variance of log-volatility innovations in quarterly data, adjust for time aggregation to an annual sampling frequency, and then set the diagonal elements of the prior scale matrix so that the prior modes for σ_r^2 and σ_q^2 equal the adjusted value. Lacking a strong prior view about the covariance σ_{rq} , we set the off-diagonal elements of the prior scale matrix to zero, thus centering the prior covariance on zero. After centering the prior in this way, we set the degree of freedom parameter to deliver plausible prior credible sets. After some experimentation, we settled on an $IW(0.05 \cdot I_2, 10)$ prior for W . Dashed lines in figure 6 portray prior histograms for the standard deviations σ_r and σ_q and the correlation $\sigma_{rq}/\sigma_r\sigma_q$. A centered prior 95 percent credible set for the standard deviations ranges from 0.158 to 0.380, while that for the correlation covers the interval ± 0.6 .

5 Posteriors

The joint posterior distribution for parameters and hidden states is simulated via a MCMC algorithm. Except for the measurement error and assumptions about the prior, the model is very similar to that of Stock and Watson (2007) and is closely related to that of Cogley and Sargent (2005). The MCMC algorithm is therefore also similar to

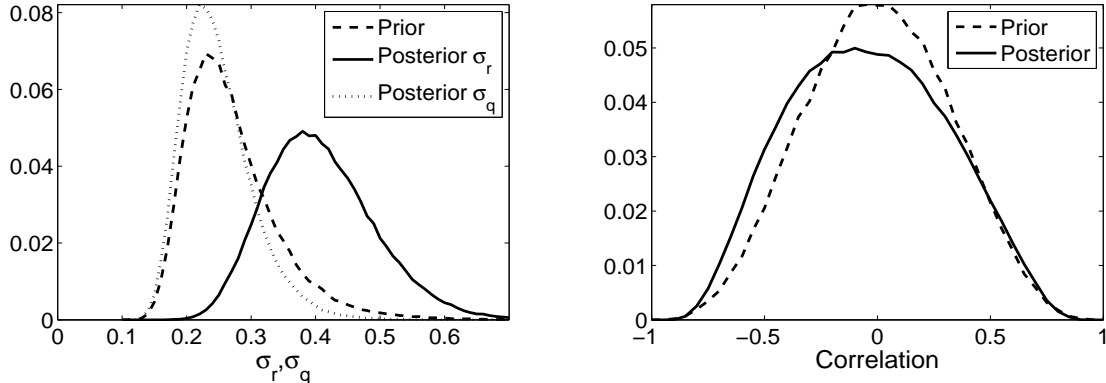


Figure 6: Priors and posteriors for variances and covariances of log-volatility innovations

Cogley and Sargent’s. Appendix A describes the details.⁸

5.1 Parameters and latent states

As we shall see, our measures of uncertainty and instability are functions of model parameters and latent states. We therefore begin by summarizing aspects of the posterior distribution.

Also shown in figure 5 are posterior distributions for the measurement-error parameters ρ_m and σ_m . As expected, measurement-error parameters are sharply identified from overlapping clean and noisy observations on inflation, and their posteriors are much more concentrated than the priors (compare the solid and dashed lines in figure 5). Measurement errors are autocorrelated (the mode for ρ_m is 0.33, and a centered 95 percent credible set ranges from 0.03 to 0.6), and there is little support for large values of σ_m (the posterior mode is 0.035, and a centered 95 percent credible set ranges from 0.028 to 0.045). At the posterior mode, the unconditional standard deviation of m_t is 0.037. Since the standard deviation of the Warren-Pearson-Hanes measure of inflation is 0.092, this amounts to 40 percent of measured inflation volatility, a substantial proportion but considerably less than some values entertained by the prior.

Posterior distributions for the standard deviations and correlation of log volatility innovations appear in figure 6. Perhaps less obviously, precise estimation of measurement-error parameters helps identify pre-1948 movements in r_t and therefore also improves

⁸ Shephard (2013) describes an alternative simulation algorithm based on an auxiliary particle filter. Because his algorithm intensively exploits parallel processing, our guess is that it is computationally more efficient than ours. Nevertheless, our algorithm delivers a valid sample from the joint posterior distribution for parameters and hidden states.

estimation of σ_r . The noise-ridden measure of inflation has two transient components, one $\sqrt{r_t}\varepsilon_{\pi t}$ that is genuine and another m_t due to measurement error. Since parameters governing the latter are precisely estimated, the amount of transient variation attributable to measurement error is restricted. Transient variation over and above an amount consistent with m_t must therefore be genuine. Since pre-1948 variation in r_t is more strongly identified, so is σ_r . Thus, in contrast to our earlier work (Cogley, et al. 2012), the posterior for σ_r differs from its prior, with the data wanting higher values of σ_r (compare the solid and dashed lines figure 6). The posterior mode for σ_r increases to 0.38, and a centered 95 percent credible set ranges from 0.26 to 0.59.

Alas, the other innovation variance parameters are weakly identified: their posteriors are essentially the same as their priors (compare the dotted and dashed lines in the left panel of figure 6 and the solid and dashed lines in the right). Because of the extent to which the prior and posterior for σ_q overlap, we also examined a model with a less informative $IW(0.05 \cdot I_2, 2)$ prior for W . This has long upper tails for σ_q and σ_r and thus promotes faster drift in q_t and r_t . Despite that, the posteriors were similar to those shown in figure 6, and our overall conclusions were much the same. For this reason, we doubt that our benchmark prior has undue influence on the results.

Aspects of the posterior for the hidden states μ_t , $\pi_t - \mu_t$, $\sqrt{q_t}$, and $\sqrt{r_t}$ are shown in figure 7. In each panel, a solid line depicts the posterior median at each date, and dotted lines portray the interquartile range.

Although the persistent component μ_t was centered around zero in the 19th century, it was not strongly anchored near zero (see the upper left panel). On the contrary, the posterior median wandered between +80 and -200 basis points per annum until the early 20th century, when it settled in the vicinity of 1.5 to 2 percent. It remained there until the Great Inflation of the 1970s, when the median estimate peaked at 7.4 percent per annum. After the Volcker disinflation, the median estimate fell to around 2.5 percent, and it fluctuated between 1.4 and 3.3 percent from then until the end of the sample. There is also substantial uncertainty about the location of μ_t at any given date. Before 1914, the interquartile range was typically ± 130 basis points around the median with a spike to ± 180 basis points during the Civil War. With better data after 1948, the interquartile range narrowed to ± 85 basis points around the median but spiked to ± 260 basis points during the Great Inflation.

Movements in q_t , the innovation variance of the persistent component μ_t , follow a similar pattern (see the lower left panel of figure 7). Volatility was high during the third quarter of the 19th century and declined gradually, with the median estimate of $\sqrt{q_t}$ falling from 1 percent in 1865 to 0.5 percent in 1945. As the economy's nominal

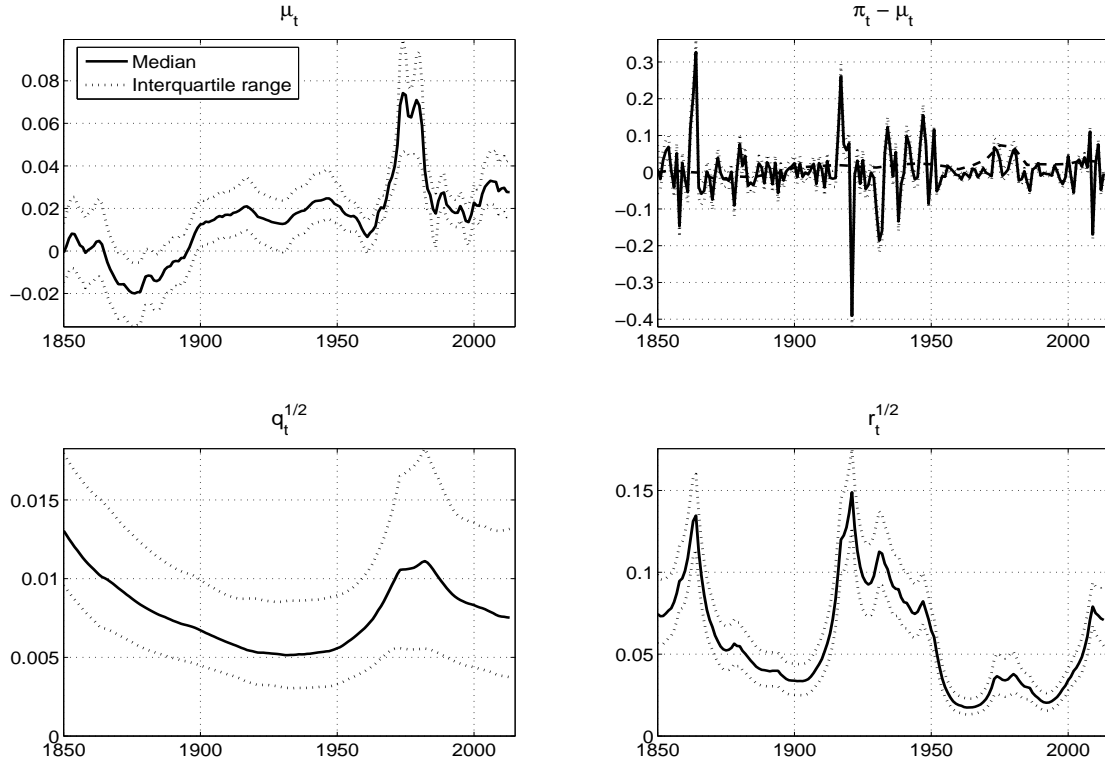


Figure 7: Posteriors for hidden states

anchor began to drag during the run-up to the Great Inflation, our estimate of $\sqrt{q_t}$ rose, reaching a local peak of 1.1 percent at the time of the Volcker disinflation. Thus, the post-World War II peak in q_t was about the same as at the end of the Civil War. Furthermore, the median estimate fell after 1983 and by the end of the sample had returned to roughly the same level as at the end of the 19th century. We therefore find no strong evidence that increments to the persistent component of inflation were systematically more variable after World War II. On the contrary, the volatility of $\Delta\mu_t$ was highest at the beginning of our sample.

The right column of figure 7 depicts the transient component of inflation $\pi_t - \mu_t$ and its innovation variance r_t . Although perhaps not obvious at first glance, the figure attests that most bursts in U.S. inflation were transitory. To see this, notice the difference in scale between the left and right columns. For instance, while the interquartile range for μ_t varies between -0.03 and 0.09, that for $\pi_t - \mu_t$ ranges from -0.4 to 0.4. Similarly, estimates of $\sqrt{r_t}$ are an order of magnitude greater than those

for $\sqrt{q_t}$. To highlight this difference in scale, we reproduce in the top right panel the median estimate of μ_t , which we portray as a dashed line. Except for the Great Inflation, fluctuations in μ_t are small in comparison with variations in $\pi_t - \mu_t$. As we shall see, transient variation dominates not only at the one-year horizon but also at longer horizons.

The largest bursts of transient inflation occurred during and after the Civil War, in and between the two World Wars, and during the recent financial crisis. A lesser burst of transient volatility also appeared during the Great Inflation, but this was comparatively small. Troughs in r_t occurred shortly after the turn of the 20th century, in the early 1960s under Bretton Woods, and in the 1990s. Thus, three ‘great moderations’ appear in our sample, one in the last quarter of the 19th century, another between World War II and the Great Inflation, and a third beginning in the mid 1980s and ending after the turn of the 21st century. Because measurement error has been purged, the model asserts that these moderations were authentic.

5.2 Orders of integration

Our model implies that inflation is $I(1)$ and that the log price level is $I(2)$, a common specification for post-World War II data. However, it is not obvious whether these orders of integration are consistent with data from before the First World War data. To examine this issue, we estimate an augmented Dickey-Fuller regression,

$$y_t = \mu + \rho y_{t-1} + \sum_{j=1}^2 \zeta_j \Delta y_{t-j} + u_t, \quad (6)$$

where y_t is measured inflation for the period 1850-1913, and we calculate the t -statistic for $\rho - 1$. This augmented Dickey-Fuller statistic is -4.56, and its 1-percent asymptotic critical value is -3.43. The test therefore seems strongly to reject a unit root in inflation.

However, the estimates shown in figure 7 suggest that the random-walk component of inflation was small prior to 1914 and that the transient component plus measurement error was much more volatile. It is tenuous whether an augmented Dickey-Fuller test can detect a small random-walk component hidden under substantial noise. To check the size of the test for our specification, we simulate the state-space model, generating artificial data on measured inflation by drawing from the posterior for π_t , ρ_m , and σ_m and then calculating the implied distribution for the augmented Dickey-Fuller statistic. It turns out that the null distribution is shifted well to the left of the asymptotic distribution and that the correct 10 percent critical value is -5.13, implying that a unit root in inflation is not rejected. Indeed, the p -value for a sample statistic of -4.56 is 0.39.

None of this proves that inflation was $I(1)$, but it does establish that our $I(1)$ representation is not grossly at odds with the data. We are sufficiently reassured that our representation for inflation is good enough for pre-World War I data to allow us to proceed.

5.3 Price-level uncertainty

To measure price-level uncertainty, we calculate the standard deviation of cumulative h -year inflation, $p_{t+h} - p_t$, conditioned on the log-volatility innovation variance W and histories of the time-varying states $\omega^t = (\pi^t, \mu^t, r^t, q^t)$. As shown in appendix B, the conditional mean and standard deviation are

$$\begin{aligned} E(p_{t+h} - p_t | \omega^t, W) &= h\mu_t, \\ \sigma(p_{t+h} - p_t | \omega^t, W) &= \sqrt{q_t \sum_{j=1}^h (h-j+1)^2 \exp(j\sigma_q^2/2) + r_t \sum_{j=1}^h \exp(j\sigma_r^2/2)}, \end{aligned} \tag{7}$$

respectively. Because the information set contains the current price level p_t , the conditional standard deviation of the future price level p_{t+h} is the same as that for cumulative inflation, $\sigma(p_{t+h} | \omega^t, W) = \sigma(p_{t+h} - p_t | \omega^t, W)$. Both measure uncertainty about price-level paths emanating forward from a given (ω^t, W) pair.⁹

The conditioning variables (ω^t, W) are unknown, but our MCMC algorithm delivers a sample from their joint posterior distribution. By calculating the conditional standard deviation $\sigma(p_{t+h} - p_t | \omega^t, W)$ for every (ω^t, W) pair in this sample, we obtain a sample from the posterior distribution $p(\sigma(p_{t+h} - p_t) | y^T)$, where y^T denotes the entire sample of observations and T is the length of the sample. We call the resulting values *smoothed conditional volatilities* because they are derived from a posterior for (ω^t, W) that conditions on the full sample y^T .^{10,11}

⁹ Notice that the permanent and transient components both contribute terms that make this conditional variance increase with the forecast horizon h . The increment to the term involving volatility in the transient component is $r_t \exp(h\sigma_r^2/2)$, while the increment to the term involving volatility in the permanent component consists of the sum of $q_t \exp(h\sigma_q^2/2)$ plus a number of other nonnegative terms. If the values associated with permanent and transient components were similar (i.e. if $q_t \approx r_t$ and $\sigma_q \approx \sigma_r$), the permanent component would grow more rapidly and dominate at long horizons. However, since we estimate $q_t \ll r_t$ and $\sigma_q < \sigma_r$, the dominance of the permanent component is not a foregone conclusion, at least not at horizons that interest us.

¹⁰ We draw from the smoothed density $p(\omega^t, W | y^T)$ rather than the filter densities $p(\omega^t, W | y^t)$ because filter densities for $t < 1948$ fail to exploit the identifying power of overlapping clean and noisy measures. Hence filter densities before and after World War II would be on a different footing vis a vis measurement error.

¹¹ For $t < T$, a smoothed conditional volatility is not a posterior conditional-on- y^T standard deviation, but we believe it is a better choice because the conditional-on- y^T standard deviation would tell us more about measurement error than about price-level volatility.




Figure 8: Posteriors for smoothed conditional volatilities 5 and 10 years ahead

Smoothed conditional volatilities for forecast horizons of 5 and 10 years are shown in figure 8. Solid lines again depict the posterior median at each date, and dotted lines portray the interquartile range. At least for these horizons, the evolution of price-level uncertainty closely resembles the movements in $\sqrt{r_t}$ shown in figure 7. Price level uncertainty was biggest near the end of the Civil War, during the interwar period and two World Wars, and near the end of the sample. A smaller bump appears during the Great Inflation, but this again seems minor. The periods of lowest volatility occurred around the turn of the 20th century, in the late 1950s and early 1960s under Bretton Woods, and during the Great Moderation of the 1990s. The close resemblance between $\sqrt{r_t}$ and smoothed conditional volatilities attests that transient bursts of inflation remain the dominant source of price-level uncertainty at horizons of 5 and 10 years. The dominance of transient variation follows from the fact that $q_t \ll r_t$ and $\sigma_q < \sigma_r$.

Because periods of high and low volatility appear both before and after 1945, there is no unequivocal answer to whether there was greater price-level uncertainty before or after World War II. On the contrary, the answer depends on the particular years being compared. We think it is sensible to focus on local peaks and troughs. As shown in table 1, median estimates of smoothed volatilities 10 years ahead peaked in 1864, 1921, 1980, and 2009 and reached troughs in 1904, 1959, and 1994. Dates of peaks and troughs for the 5-year horizon differ only slightly, so we concentrate on these 7 years.

Table 2 compares smoothed conditional volatilities in selected prewar and postwar years. The columns and rows refer, respectively, to particular prewar and postwar base years. Thus, the top left panel compares 1864 with 1959, while the bottom right compares 1921 with 2009. The top entry in each panel is the ratio of the postwar mean smoothed conditional standard deviation relative to that in the prewar base

Table 1: Peaks and Troughs in Smoothed Conditional Volatilities

Peaks	1864	1921	1980	2009
Troughs	1904	1959	1994	

Note: Median estimates, $h = 10$ years.

year. Entries in parentheses record the proportion of posterior sample paths on which conditional standard deviations are lower after the Second World War.

We begin by comparing peaks with peaks, viz. 1864 and 1921 v. 1980 and 2009. At postwar peaks, mean smoothed volatilities were approximately 30-50 percent lower than before the war, and the probability that uncertainty was greater at prewar peaks exceeds 80 percent in all cases and 90 percent in a few. Thus the evidence weakly suggests that price-level uncertainty was greater at prewar peaks. But since uncertainty was higher after the war on a significant minority of paths (8-17 percent), the evidence is not compelling.

When comparing troughs with troughs (1904 v. 1959 and 1994), the evidence favoring a postwar stabilization is even weaker. Mean smoothed volatilities at the 1959 trough were 11-17 percent lower at the 5 and 10 years horizons, and uncertainty was lower on roughly two-thirds of posterior sample paths. At the 1994 moderation there was about as much uncertainty as at the 1904 trough.

One can of course also compare peaks with troughs. Not surprisingly, there was

Table 2: Relative Conditional Standard Deviations

	1864		1904		1921	
	5 years	10 years	5 years	10 years	5 years	10 years
1959	0.350 (0.982)	0.283 (0.996)	0.887 (0.627)	0.829 (0.674)	0.361 (0.986)	0.276 (0.997)
1980	0.640 (0.831)	0.533 (0.915)	1.620 (0.205)	1.560 (0.188)	0.659 (0.826)	0.520 (0.924)
1994	0.466 (0.909)	0.370 (0.965)	1.180 (0.462)	1.082 (0.503)	0.480 (0.908)	0.361 (0.970)
2009	0.729 (0.844)	0.685 (0.884)	1.848 (0.133)	2.003 (0.078)	0.751 (0.834)	0.668 (0.894)

Note: The top entry in each panel is the ratio of the postwar mean smoothed conditional standard deviation relative to that in the prewar base year. Entries in parentheses record the proportion of sample paths on which conditional standard deviations are lower in the postwar year.

more uncertainty at prewar peaks than at postwar troughs. For instance, mean smoothed volatilities for 1959 and 1994 are approximately 60 percent lower than those for 1864 and 1921, and the probability that price levels were more predictable at postwar troughs exceeds 95 percent. What to make of this comparison is unclear, however, for there is also evidence – albeit a bit weaker – that uncertainty was greater at postwar peaks than prewar troughs. For instance, mean smoothed volatilities in 2009 were double those in 1904, and the probability that uncertainty was greater during the financial crisis is roughly 10 percent. Similarly, mean smoothed volatilities were about 60 percent higher in 1980, and conditional standard deviations were greater than in 1904 on 80 percent of posterior sample paths.

In any case, no simple answer emerges to whether the price level was more predictable before or after 1945. Episodes of high and low volatility occur in both periods, precluding a blanket characterization. The trough-to-trough calculations confirm the promises of Ricardo and Keynes that a well-managed fiat regime can deliver outcomes comparable to those of a commodity standard.¹² Absent big shocks such as wars, depressions, or financial panics, price level uncertainty can be moderated either by means of a link to gold, as under Bretton Woods or the classical gold standard, or via a well-managed fiat regime, as in the 1990s.

5.4 Price-level instability

While conditional variances are suitable for measuring uncertainty, the concept of price stability seems different. We interpret ‘stability’ as *total* variation, not just *unpredictable* variation. For assessing price instability, we therefore compare second moments across dates. Our measure of instability is the square root of the conditional second moment (aka the conditional root mean square) for cumulative inflation,

$$crms(p_{t+h} - p_t | \omega^t, W) = \sqrt{\sigma^2(p_{t+h} - p_t | \omega^t, W) + h^2 \mu_t^2}. \quad (8)$$

This measures total price-level variation emanating forward from a given (ω^t, W) pair.^{13,14}

¹² See Sargent (2008) for accounts of their ideas and proposals.

¹³ When central bankers speak of price stability, however, they sometimes seem to mean stability around an inflation target. This concept is closer to our notion of predictability. Readers who prefer this interpretation should therefore pay more attention to the results reported in section 5.3.

¹⁴ While the conditional variance for the log price level is the same as that for cumulative inflation, conditional second moments are not, because the conditional means differ. The conditional mean for the log price level h periods ahead is $p_t + h\mu_t$, so $(p_t + h\mu_t)^2$ must be added to the conditional variance. Conditional second moments of the log price level therefore penalize both past and future cumulative

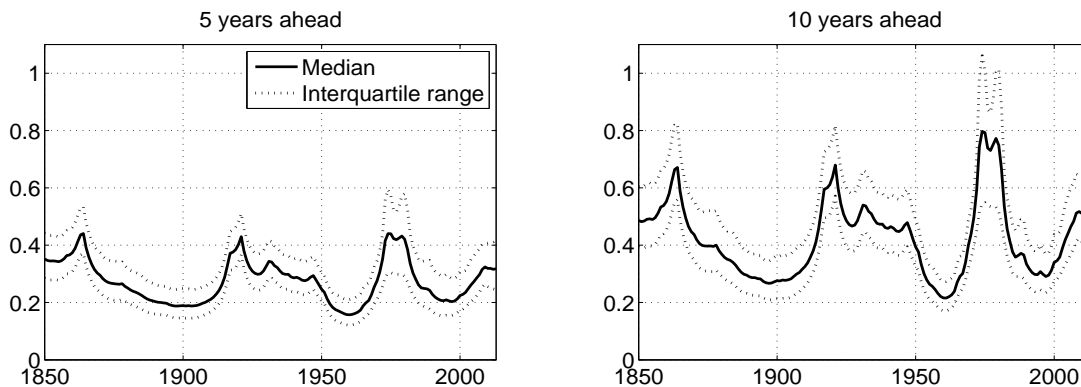


Figure 9: Posteriors for smoothed conditional root mean square statistics 5 and 10 years ahead

Smoothed conditional root mean square statistics are shown in figure 9. Since the conditional second moment is the sum of the conditional variance plus the square of the conditional mean, a difference between uncertainty and instability materializes only when the conditional mean $h\mu_t$ is sufficiently far from zero. The only times this happened in our sample were the Great Inflation of the 1970s and, to a lesser extent, the financial crisis of 2009. The primary difference between the conditional root mean square statistics shown in figure 9 and the conditional standard deviations shown in figure 8 is that there are larger bumps in our instability measure in the 1970s. Indeed, median estimates of price instability in late 1970s are larger than prewar peaks. Otherwise, the evolution of conditional root mean square statistics resembles that of conditional standard deviations.

Among other things, this means that our trough-to-trough comparisons of uncertainty extend to our measure of instability. It follows that no compelling evidence emerges that the price level was more stable during the two postwar moderations than at the turn of the 20th century.

The main difference concerns peak-to-peak comparisons. For instance, while there was weak evidence of less price-level *uncertainty* in 1980 than at prewar peaks, there was about as much price-level *instability*. Mean smoothed *crms* statistics for 1980 are approximately the same as those for 1864 and 1921, and the odds that prices were more stable in one era over the other are not far from even (see the first and third

inflation. Why past cumulative inflation should be penalized is not obvious, however, so we prefer to measure price instability by conditional second moments for cumulative inflation. Since variation in p_t dominates everything else, conditional second moments for the log price level increase with time and are greatest at the end of the sample.

Table 3: Relative Conditional Root-Mean-Square Statistics

	1864		1904		1921	
	5 years	10 years	5 years	10 years	5 years	10 years
1959	0.371 (0.984)	0.325 (0.995)	0.849 (0.646)	0.774 (0.692)	0.386 (0.988)	0.326 (0.996)
1980	0.967 (0.558)	1.092 (0.455)	2.216 (0.070)	2.599 (0.040)	1.007 (0.535)	1.095 (0.454)
1994	0.505 (0.911)	0.447 (0.960)	1.157 (0.425)	1.064 (0.441)	0.526 (0.910)	0.448 (0.966)
2009	0.795 (0.778)	0.796 (0.755)	1.821 (0.120)	1.896 (0.097)	0.828 (0.767)	0.798 (0.770)

Note: The top entry in each panel is the ratio of mean smoothed *CRMS* relative to that in the prewar base year. Entries in parentheses record the proportion of sample paths on which *CRMS* statistics are lower in the postwar year.

columns of the 1980 row in table 3). Comparing the prewar peaks with 2009 yields similar results. Mean smoothed *crms* statistics are about 20 percent lower than at the prewar peaks, but price levels 5 and 10 years ahead were less stable in 2009 on 20-25 percent of paths (see the first and third columns of the 2009 row in table 3).

5.5 Deflation risk

Another aspect of price level dynamics that changed after the Second World War was the risk of deflation. To develop evidence on this feature of the data, we calculate smoothed deflation probabilities and compare them across dates.

As an analogue to our measures of uncertainty and instability, we want the probability of cumulative deflation $p_{t+h} - p_t < 0$ going forward from a given (ω^t, W) pair. Toward this end, let ξ_t^h be a sequence of potential future shocks $\{\varepsilon_{\pi s}, \varepsilon_{\mu s}, \eta_{qs}, \eta_{rs}\}_{s=t+1}^{t+h}$. According to the model, ξ_t^h has a normal unconditional distribution that we denote $p_N(\xi_t^h)$. Given a pair (ω^t, W) along with a draw of potential future shocks ξ_t^h , a sequence of potential future price level realizations $\{p_s(\omega^t, W, \xi_t^s)\}_{s=1}^h$ can be calculated by iterating on the state equations. We define an indicator variable

$$\begin{aligned}
 I_{deflation}\left(p_h(\omega^t, W, \xi_t^h), p_t(\omega^t, W)\right) &= 1 \text{ if } p_{t+h} < p_t, \\
 &= 0 \text{ otherwise,}
 \end{aligned}
 \tag{9}$$

that records whether cumulative inflation is positive or negative. Conditional on

(ω^t, W) , the probability of deflation is

$$\begin{aligned} dpr(\omega^t, W) &\equiv pr(p_{t+h} - p_t < 0 | \omega^t, W), \\ &= \int I_{deflation}(p_h(\omega^t, W, \xi_t^h), p_t(\omega^t, W)) p_N(\xi_t^h) d\xi_t^h. \end{aligned} \tag{10}$$

It is easy to approximate this probability by Monte Carlo integration. By calculating dpr for every (ω^t, W) pair in the posterior sample, we obtain its posterior distribution. We call the resulting values *smoothed deflation probabilities* because they are derived from a posterior for ω^t and W that conditions on the full sample y^T .

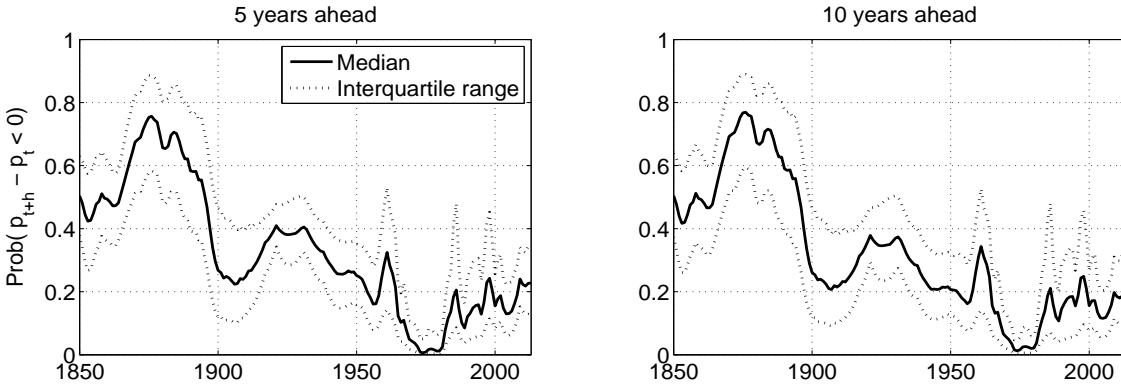


Figure 10: Smoothed deflation probabilities 5 and 10 years ahead

Figure 10 depicts the median and interquartile range for dpr for forecast horizons of 5 and 10 years. Not surprisingly, low values of μ_t and high values of $\sigma(p_{t+h} - p_t)$ both contribute to high deflation risk. Thus, the probability of deflation was highest – above 70 percent – after the Civil War when μ_t was negative and σ_t was high. The opposite scenario, periods with high μ_t and low σ_t , did not occur in our sample, but deflation risk was lowest – below 5 percent – during the Great Inflation when high values of μ_t trumped intermediate values of σ_t . The three price level moderations lie in between, with deflation probabilities on the order of 10-20 percent.¹⁵ Thus, low values of μ_t trump low values of σ_t during moderations.

Table 4 compares selected years before and after the Second World War. Peaks in deflation risk were reached in 1876, 1921, and 1998, and troughs occurred in 1909 and 1979. We also include 2009 to measure deflation risk during the financial crisis. As

¹⁵ Because stochastic volatilities are represented as geometric random walks, our model generates very thick tails for future price levels even in tranquil times.

Table 4: Relative Deflation Probabilities

	1876		1909		1921	
	5 years	10 years	5 years	10 years	5 years	10 years
1979	0.082 (0.995)	0.071 (0.995)	0.198 (0.887)	0.181 (0.885)	0.136 (0.993)	0.125 (0.992)
1998	0.396 (0.918)	0.381 (0.922)	0.958 (0.532)	0.971 (0.518)	0.656 (0.746)	0.671 (0.714)
2009	0.368 (0.951)	0.320 (0.954)	0.889 (0.528)	0.814 (0.553)	0.609 (0.840)	0.562 (0.828)

Note: The top entry in each panel is the ratio of the postwar mean probability of deflation relative to that in the prewar base year. Entries in parentheses record the proportion of sample paths on which the probability of deflation was lower in the postwar year.

before, the top row records the ratio of mean deflation risks before and after the war, and the bottom depicts the probability that $dpr(\omega^t, \theta)$ was lower after the war.

At the 1998 postwar peak, mean deflation risk was roughly 35-60 percent lower than at the prewar peaks of 1921 and 1876 (see the first and third columns of the row corresponding to 1998). Conditional deflation probabilities were also lower than in 1876 on more than 90 percent of posterior sample paths. Relative to 1921, however, the evidence is less decisive: $dpr(\omega^t, \theta)$ was higher in 1921 on 75 percent of posterior paths but was lower on the remaining 25 percent. Comparing 2009 with the prewar peaks yields similar results, with slightly stronger evidence of lower postwar deflation risk.

The postwar peak in dpr is comparable to the prewar trough (see the 1909-1998 panel), but the postwar trough is negligible in comparison with either of the prewar peaks (see the first and third columns of the 1979 row). A trough-to-trough comparison yields weak evidence of a postwar decline (see the 1909-1979 panel). At the 1979 postwar trough, mean deflation risk was approximately one-fifth as large as at the 1909 prewar trough, but the probability of decline fell short of 90 percent. Thus, on a nontrivial minority of posterior sample paths, there was less deflation risk at the prewar trough.

6 Discussion of related literature

Our paper is closely related to research on changing inflation persistence. For instance, Barsky (1987) documents that U.S. inflation was well approximated by a martingale-

difference process before World War II and became strongly autocorrelated only during the Great Inflation. Benati (2008) corroborates these findings and shows more generally that inflation is weakly persistent whenever a monetary regime has a clearly defined nominal anchor. Our results and those of Stock and Watson (2007) are broadly consistent with their findings.

While our statistical model stays close to Stock and Watson's (2007) representation, a number of other papers extend it in various directions to capture other features of postwar US inflation. For instance, Cogley, et al. (2010) investigate autocorrelation in the inflation gap, $\pi_t - \mu_t$, and Cecchetti, et al. (2007) model log-volatility innovations as a two-state Markov process, thus allowing jumps in r_t and q_t . Kozicki and Tinsley (2012) and Mertens (2011) introduce survey data on inflation forecasts to aid identification,¹⁶ Mertens (2011) and Stella and Stock (2012) develop multivariate forecasting models by incorporating data on other macroeconomic outcomes such as unemployment and the term structure of interest rates,¹⁷ and Chan, et al. (2013) introduce a priori bounds on the random walk component μ_t .¹⁸ Another worthy extension, especially for modeling 19th century inflation dynamics, would be to consider adding a stationary component to the log price level.

All of these extensions merit further research, but we chose not to undertake them here because we wanted to draw out the implications of a well-known and relatively simple statistical model for pre-World War II inflation volatility. We believe we have captured the main contours of the rise and fall in inflation volatility, and we encourage others to join us by studying how alternative model specifications affect the results.

7 Concluding remarks

Our analysis of US wholesale price inflation evinces recurring episodes of rising and falling price-level uncertainty and instability. Major outbreaks of price level instability and unpredictability are associated with the Civil War, the two World Wars and Great Depression, and the Great Inflation and Great Recession. In each instance, a crisis disrupted pre-existing monetary arrangements and created considerable uncertainty about the future. In each case, policy makers eventually found a path back to price stability, but that took a long time: the average time from peak to trough was 30 years.

¹⁶ Surveys do not extend far enough back in time to help us.

¹⁷ Coping with measurement error is a bigger challenge in a multivariate environment, one well worth undertaking.

¹⁸ In the present context, estimates of μ_t seem reasonable, making explicit bounds unnecessary. Whether bounds that are reasonable for an inflation-targeting regime would remain valid in the event of a crisis such as a great and dangerous war is open to doubt.

Whether prices were more predictable or stable before or after the Second World War seems to rest on trough-to-peak comparisons. When comparing peaks with peaks and troughs with troughs, no compelling evidence emerges that the price level was more stable or less uncertain in one era than the other.

Evidence emerges for three ‘great price-level moderations,’ one near the turn of 20th century, another under Bretton Woods, and a third in the 1990s. Well-managed systems sometimes involved links to gold – as under the classical gold standard and Bretton Woods – but the experience of the 1990s proves that a well-managed fiat regime can achieve the same outcomes, as Ricardo and Keynes promised. Thus, a link to gold or some other commodity seems not to have been necessary for moderating price level volatility.

Nor does a commodity standard provide a guarantee of permanent price stability. The major challenge is coping with big shocks such as wars, depressions, and financial panics. Neither a commodity nor a fiat regime is immune to big shocks, and convertibility has typically been suspended in times of crisis. During the US Civil War, policy makers in both the North and the South stopped backing their currencies with gold or silver because issues of war finance became their paramount concerns. Convertibility was restored in the North only in 1879, and never in the South. The classical gold standard and Bretton Woods system ended not because the authorities thought they had discovered better methods for maintaining price stability but because other economic objectives supervened. Because maintaining convertibility during a crisis was unlikely, there never was a believable unconditional commitment to a commodity standard. Before the breakdown of Bretton Woods we saw mixed standards, with convertibility during tranquil times and suspensions and reversions to fiat money after big shocks.

Our analysis is cast in terms of a purely *statistical* model, so we have nothing to say about why someone might *care* about the uncertainty and/or the stability of inflation and how they have changed over time. But plenty of economic models do suggest that we should care. For example, models in the tradition of Lucas (1973) make social welfare depend on the uncertainty of the log price level, but not its stability. Other models in the tradition of Lucas (2000) focus on welfare effects of the stability of inflation, not its uncertainty. In other more general macroeconomic models, people might care about both price level stability and price level uncertainty. Our statistical characterizations of price level stability and price level uncertainty set the stage for a quantitative economic analysis of their welfare consequences.

8 References

Barsky, Robert, “The Fisher Hypothesis and the Forecastability and Persistence of Inflation,” *Journal of Monetary Economics* 19 (1987), 3-24.

Benati, Luca, “Investigating Inflation Persistence Across Monetary Regimes,” *Quarterly Journal of Economics* 123 (2008), 1005-1060.

Cecchetti, Stephen, Peter Hooper, Bruce Kasman, Kermit Schoenholtz, and Mark Watson, “Understanding the Evolving Inflation Process,” *U.S. Monetary Policy Report* No. 1 (2007).

Chan, Joshua, Gary Koop, and Simon Potter, “A New Model of Trend Inflation,” *Journal of Business and Economic Statistics* 34 (2013), 91-106.

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

Cogley, Timothy, Girogio Primiceri, and Thomas J. Sargent, “Inflation-Gap Persistence in the U.S.,” *American Economic Journal: Macroeconomics* 2 (2010), 43-69.

Cogley, Timothy, Thomas J. Sargent, and Paolo Surico, “Price-Level Uncertainty and Instability in the U.K.,” unpublished manuscript, New York University and London Business School, (2012)

Hanes, Christopher, “Consistent Wholesale Price Series for the United States, 1860-1990,” in Trevor J. O. Dick, editor, *Business Cycles Since 1820: New International Perspectives from Historical Evidence*, Edward Elgar (1998).

Hanes, Christopher, “Prices,” in S.B. Carter, S.G. Gartner, M.R. Haines, A.L. Olmstead, R. Sutch, and G. Wright, editors, *Historical Statistics of the United States: Earliest Times to the Present, Millennial Edition*, Cambridge University Press (2006).

Kozicki, Sharon and Peter Tinsley, “Effective Use of Survey Information in Estimating the Evolution of Expected Inflation,” *Journal of Money, Credit, and Banking* 44 (2012), 145-169.

Lucas, Robert E., Jr, “Some International Evidence on Output-Inflation Tradeoffs,” *American Economic Review*, 63 (1973), 326-34.

Lucas, Robert E., Jr, “Inflation and Welfare,” *Econometrica*, 68 (2000), 247-274.

Mertens, Elmar, “Measuring the Level and Uncertainty of Trend Inflation,” Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, Paper 2011-42 (2011).

Romer, Christina, “Is the Stabilization of the Postwar Economy a Figment of the Data?,” *American Economic Review* 76 (1986a), 314-334.

Romer, Christina, “Spurious Volatility in Historical Unemployment Data,” *Journal of Political Economy* 94 (1986b), 1-37.

Sargent, Thomas J., “Evolution and Intelligent Design,” *American Economic Review*, 98 (2008), 5-37.

Shephard, Neil, “Martingale Unobserved Components Models,” unpublished manuscript, University of Oxford (2013).

Stella, Andrea and James Stock, “A State-Dependent Model for Inflation Forecasting,” Board of Governors of the Federal Reserve System, International Finance Discussion Paper No. 1062 (2012).

Stock, James and Mark Watson, “Why Has Inflation Become Harder to Forecast?,” *Journal of Money Credit and Banking* 39 (2007), 3-33.

Warren, George and Frank Pearson, *Prices*, Wiley (1933).