



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

**A New Approach to Estimating the Natural
Rate of Interest**

Luca Benati

16-10

August 2016

DISCUSSION PAPERS

Schanzeneckstrasse 1
Postfach
CH-3001 Bern, Switzerland
<http://www.vwi.unibe.ch>

A New Approach to Estimating the Natural Rate of Interest*

Luca Benati
University of Bern[†]

Abstract

I illustrate a novel, and straightforward method to extract the natural rate of interest from the joint dynamics of inflation and the term structure of nominal interest rates, based on cointegrated structural VARs. The key identifying assumption is that the unit root component of nominal interest rates is driven by two permanent shocks, an inflation shock, and a shock to the natural rate, which can be disentangled *via* standard SVAR techniques.

I estimate the natural rate for the Zero Lower Bound (ZLB) period by simulating the pre-ZLB estimated system *conditional* on the actual dynamics of the unconstrained variables at the ZLB, and then imposing the identifying restrictions upon the conditional projections. Evidence suggests that since the beginning of the financial crisis the natural rate has decreased, in the United States, by 0.5-0.7 percentage points, and it stands, in the second half of 2016, at about 0.5 per cent.

Keywords: Natural rate of interest; structural VARs; unit roots; cointegration.

*I wish to thank James Hamilton for useful suggestions. Usual disclaimers apply.

[†]Department of Economics, University of Bern, Schanzeneckstrasse 1, CH-3001 Bern, Switzerland. Email: luca.benati@vwi.unibe.ch

1 Introduction

Since the outbreak of the financial crisis, the natural rate of interest has been one of the most intensely discussed issues in either policymaking circles, academia, or the financial press. Two key pieces of evidence suggest that the natural rate has been progressively decreasing over the last three decades: *first*, in both the United States, and several other countries, *ex post* real interest rates have been steadily declining since the early 1980s;¹ *second*, over the same period, trend GDP growth has been slowing down in most advanced countries.² Since, as discussed by Laubach and Williams (2003, henceforth, LW), standard growth models imply a linear relationship between the natural rate of interest and trend output growth, the fact that the latter has been decreasing provides strong *prima facie* evidence that the former might also have been declining.

There are currently two main approaches to estimating the natural rate of interest. In the former, which was originally introduced by LW (2003), the natural rate is modelled as an I(1) process; it is embedded within a semi-structural framework which typically also features a Phillips curve-type relationship; and it is then extracted from the data *via* Kalman filtering methods.³ The second approach—for a recent example, see e.g. Barsky *et al.* (2014)—is instead based on the notion of extracting the natural rate from an estimated, fully-specified DSGE model.

In this paper I illustrate a novel, and straightforward method to extract the natural rate of interest from the joint dynamics of inflation and the term structure of nominal interest rates based on cointegrated structural VARs.⁴ The key identifying assumption is that the unit root component of nominal interest rates is driven by two permanent orthogonal shocks, an inflation shock (*via* the Fisher effect), and a shock to the natural rate, which can be disentangled *via* standard SVAR techniques. I estimate the cointegrated system based on U.S. data up to October 2008, after which the Zero Lower Bound (ZLB) started being *de facto* binding. Estimates suggest that the natural rate had been declining since the first half of the 1980s. I then estimate

¹See, e.g., Chapter 3 (‘Perspectives on Global Real Interest Rates’) of the April 2014 International Monetary Fund’s *World Economic Outlook*.

²For the United States, see e.g. Antolin-Diaz, Drechsel, and Petrella (2016).

³For recent applications of this methodology, see Laubach and Williams (2016) and Holston *et al.* (2016).

⁴In conceptually related work, Hamilton *et al.* (2016, Section 5) postulate a cointegration relationship between the U.S. and world real interest rates, and extract estimates of the U.S. real FED Funds rate from a bivariate VECM for the two series. The *only* element of commonality between Hamilton *et al.* (2016) and the present work, however, is that they both use cointegration methods. Beyond that, the two papers have else in common. By the same token, Holston *et al.* (2016) estimate the natural rate for the United States, the Euro area, Canada, and United Kingdom based on LW’s (2003) methodology, and then set up a VECM for the four estimated natural rates in order to explore the existence of a common stochastic trend in individual countries’ natural rates. Again, the only point of contact between Holston *et al.* (2016) and the present work is the use of cointegration methods.

the natural rate for the period since November 2008 by simulating (i.e., bootstrapping) the system I previously estimated up to October 2008 *conditional* on the actual dynamics of the variables unconstrained by the ZLB (i.e., all variables except for the nominal interest rates, for which—conceptually in line with Swanson and Williams (2014)—I produce evidence that, since November 2008, they have likely been constrained by the ZLB). Finally, I estimate the natural rate by imposing the identifying restrictions upon the cointegrated system’s conditional projections. Evidence suggests that, since the beginning of the financial crisis, the U.S. natural rate has decreased by 0.5-0.7 percentage points, and it stands, in the second half of 2016, at about 0.5 per cent.

The paper is organized as follows. The next section discusses results from unit root and cointegration tests, which strongly suggest that—in line with basic economic intuition—(i) both inflation, and all nominal rates are I(1); (ii) all of the interest rate spreads are stationary, thus implying that the seven nominal interest rates I use are cointegrated, with the six cointegration vectors being the theoretical ones (i.e, featuring a 1, a -1, and five zeros); and (iii) when the system is augmented with other I(1) series (hours and GDP *per capita*) no additional cointegration relationship is present. Section 3 discusses the proposed methodology, which is based on the notion of identifying the cointegrated structural VARs *via* long-run restrictions. In particular, the key identifying assumption is that the permanent component of nominal interest rates is driven by two, and only two orthogonal shocks: a permanent inflation shock, and a permanent shock to the natural rate of interest. Section 4 presents Classical point estimates for either the period up to September 2008, or the subsequent ZLB period, whereas Section 5 characterizes uncertainty around the point estimates *via* bootstrapping methods, and presents statistical evidence on the natural rate’s progressive decline since the early 1980s. As I show, uncertainty around the point estimates is, unsurprisingly, quite substantial. I therefore discuss a simple approach to shrinking the extent of uncertainty, based on the intuitive notion of rejecting ‘implausible’ paths for the natural rate, where implausibility is defined in terms of the implied average gap for the real FED Funds rate. Since, from a conceptual point of view, imposing prior knowledge upon the data is best implemented within a Bayesian framework, in Section 6 I turn to Bayesian estimation of the cointegrated systems. As I show, however, implementing the analysis within a Bayesian framework suffers from the crucial shortcoming that estimates of the natural rate end up depending, in a material way, on the priors which are being imposed in order to shrink the extent of uncertainty. In the end, my own conclusion is therefore that—for practical purposes—the best thing to do is probably to adopt a Classical approach, and to simply ignore ‘implausible’ paths as in Section 5.

2 Integration and Cointegration Properties of the Data

The benchmark system for the United States I use throughout the paper features core CPI inflation; the nominal Federal Funds rate; the nominal 3- and 6-month Treasury bill rates; the nominal 1-, 3-, 5-, and 10-year government bond yields; the logarithm of total hours worked in private industries *per capita*; the unemployment rate; and the rate of capacity utilization in the industrial sector. (All of the data are discussed in detail in Appendix A.) The sample period used for *estimation* of the cointegrated VECM extends from January 1964 (when hours worked in private industries first becomes available) to October 2008 (in November 2008 the Federal Funds rate, taking a value of 0.39 per cent, reached *de facto* the ZLB).

In order to check for (i) robustness of the estimates, and (ii) what the minimal system producing robust estimates in fact is, I also consider either smaller systems, or larger ones featuring the logarithm of real GDP *per capita*, and the vacancy rate. To anticipate, smaller systems produce somehow different and—as I will argue—manifestly less reliable estimates. On the other hand, augmenting the benchmark system with GDP and/or the vacancy rate does not produce materially different results.

As the next sub-section will show, all of the previously mentioned variables are I(1) with the single exception of the unemployment and vacancy rates, and of capacity utilization. From an estimation standpoint this does not cause any problem,⁵ because all of the cointegrated systems will be estimated in the VECM (i.e., stationary) form by imposing, in estimation, the six theoretical cointegration vectors for the seven nominal interest rates. This allows to expand the VECM by adding I(0) series such as the unemployment rate and capacity utilization.

2.1 Evidence from unit root tests

Table 1 reports bootstrapped *p*-values for Elliot, Rothenberg, and Stock (1996) unit root tests.⁶ For the logarithm of real GDP *per capita*, which exhibits an obvious trend, the tests are based on models including an intercept and a time trend.⁷ For log hours *per capita*, for which evidence of a trend is not clear-cut, I consider tests with an intercept, and either with or without a time trend. For all other series—which, as a matter of economic logic, should *not* possess any trend—I only consider tests with

⁵I wish to thank James Hamilton for an helpful email exchange on this issue.

⁶For either series, *p*-values have been computed by bootstrapping 10,000 times estimated ARIMA(*p*,1,0) processes. In all cases, the bootstrapped processes are of length equal to the series under investigation. As for the lag order, *p*, since, as it is well known, results from unit root tests may be sensitive to the specific lag order which is being used, for reasons of robustness we consider four alternative lag orders, either 3, 6, 9, or 12 months.

⁷The reason for including a time trend is that, as discussed e.g. by Hamilton (1994, pp. 501), the model used for unit root tests should be a meaningful one also under the alternative.

Table 1 Bootstrapped p -values for Elliot, Rothenberg, and Stock unit root tests ^a				
	$p=3$	$p=6$	$p=9$	$p=12$
I: Tests with an intercept and no time trend				
Inflation	0.00	0.02	0.15	0.12
Federal Funds rate	0.17	0.28	0.14	0.28
3-month Treasury bill rate	0.31	0.41	0.31	0.29
6-month Treasury bill rate	0.35	0.42	0.33	0.26
1-year government bond yield	0.39	0.44	0.40	0.30
3-year government bond yield	0.53	0.57	0.54	0.40
5-year government bond yield	0.56	0.55	0.54	0.42
10-year government bond yield	0.60	0.57	0.52	0.42
Log aggregate hours in private industries <i>per capita</i>	0.43	0.20	0.21	0.26
Unemployment rate	0.30	0.05	0.05	0.06
Vacancy rate	0.16	0.01	0.01	0.00
Rate of capacity utilization	0.12	0.07	0.05	0.01
<i>Spreads</i>				
3-month Treasury bill rate minus Federal Funds rate	0.00	0.00	0.00	0.00
6-month Treasury bill rate minus 3-month Treasury bill rate	0.00	0.00	0.00	0.00
1-year government bond yield minus 6-month Treasury bill rate	0.00	0.01	0.04	0.04
3-year government bond yield minus 1-year government bond yield	0.00	0.01	0.01	0.01
5-year government bond yield minus 3-year government bond yield	0.06	0.12	0.06	0.02
10-year government bond yield minus 5-year government bond yield	0.02	0.06	0.06	0.05
II: Tests with an intercept and a time trend				
Log aggregate hours in private industries <i>per capita</i>	0.69	0.28	0.30	0.30
Log real GDP <i>per capita</i>	0.74	0.53	0.26	0.36
^a Based on 10,000 bootstrap replications of estimated ARIMA processes.				

Table 2 Bootstrapped p-values for Elliot, Rothenberg, and Stock unit root tests^a				
	$p=3$	$p=6$	$p=9$	$p=12$
Tests with an intercept and no time trend				
First difference of inflation	0.00	0.00	0.00	0.00
First difference of Federal Funds rate	0.00	0.00	0.00	0.00
First difference of 3-month Treasury bill rate	0.00	0.00	0.00	0.00
First difference of 6-month Treasury bill rate	0.00	0.00	0.00	0.00
First difference of 1-year government bond yield	0.00	0.00	0.00	0.00
First difference of 3-year government bond yield	0.00	0.00	0.00	0.00
First difference of 5-year government bond yield	0.00	0.00	0.00	0.00
First difference of 10-year government bond yield	0.00	0.00	0.00	0.00
Log-difference of aggregate hours in private industries <i>per capita</i>	0.00	0.00	0.00	0.00
Log-difference of real GDP <i>per capita</i>	0.00	0.00	0.00	0.00
^a Based on 10,000 bootstrap replications of estimated ARIMA processes.				

an intercept, but no time trend. The results in the table can thus be summarized. At the 10 per cent significance level, which I take as the benchmark throughout the entire paper,

(i) either the unemployment rate, the vacancy rate, or the rate of capacity utilization, are $I(0)$.

(ii) All of the spreads are $I(0)$,⁸ thus implying that either of the six pairs of nominal interest rates I consider (the 3-month Treasury bill rate and the Federal Funds rate; the 6- and 3-month Treasury bill rates; ...; and the 10- and 5-year government bond yields) is—as it should logically be expected—cointegrated with cointegration vector $[1 \ -1]'$. This implies that either permanent inflation shocks, or permanent shocks to the natural rate of interest, have an identical impact on all of the seven nominal rates. Throughout the entire paper, I will therefore impose such theoretical cointegration vectors in estimation.

(iii) All other series—inflation; all nominal interest rates; and the logarithms of either hours or real GDP *per capita*—are $I(1)$.

Finally, a necessary condition for performing Johansen’s tests is that the series under investigation contain a unit root, but that their order of integration is not greater than one. Table 2 therefore reports bootstrapped p -values for Elliot *et al.*’s (1996) unit root tests with an intercept, but no time trend, for the log-differences of hours and real GDP *per capita*, and for the first differences of all other $I(1)$ series. For all series, and for either of the four lag orders I am here considering, the bootstrapped p -values are uniformly equal to 0.00.

2.2 Evidence from Johansen’s cointegration tests

Table 3 reports results from Johansen’s cointegration tests for three alternative systems. All of them feature the seven nominal interest rates and inflation, which, based on the proposed methodology, constitute the minimal system which allows for the identification of the shocks to the natural rate. The second system also features hours, whereas the third also features hours and GDP. I bootstrap Johansen’s test statistics⁹ *via* the procedure proposed by Cavaliere *et al.* (2012; henceforth, CRT).¹⁰ Benati, Lucas, Nicolini, and Weber (2016, appendix C.3) contain extensive Monte Carlo ev-

⁸At the 10 per cent level, out of the 24 overall unit root tests performed on the spreads, only in a *single* case—for the spread between the 5- and the 3-year government bond yields, based on $p=6$ —I cannot reject the null of a unit root (with a p -value of 0.12). In *all* other cases a unit root can be rejected, most of the times comfortably so.

⁹The rationale for bootstrapping critical and p -values for Johansen’s tests was provided by Johansen (2002) himself, who showed how, in small samples, trace and maximum eigenvalue tests based on asymptotic critical values typically tend to perform poorly.

¹⁰For tests of the null of no cointegration against the alternative of one or more cointegrating vectors the model which is being bootstrapped is a simple, non-cointegrated VAR in differences. For the maximum eigenvalue tests of h versus $h+1$ cointegrating vectors, on the other hand, the model which is being bootstrapped is the VECM estimated under the null of h cointegrating vectors.

Table 3 Results from Johansen's cointegration tests^a[illegible]

idence on the excellent performance of CRT's bootstrapping procedure. I select the VAR lag order as the maximum¹¹ between the lag orders chosen by the Schwartz and the Hannan-Quinn criteria¹² for the VAR in levels.

Based on the previously discussed results for the spreads, we would expect either of the three systems to feature *at least* six cointegration vectors. In fact, the results in Table 3 suggest that either system does indeed feature six, and *only* six cointegration vectors. At the 10 per cent significance level, this is very clear for either the second or the third system, and based on either the trace or the maximum eigenvalue test statistics. As for the system only featuring the seven nominal interest rates and inflation, the trace test of the null of no cointegration against the alternative of 7 or more cointegration vectors has a p -value of 0.01, but the maximum eigenvalue test of 6 versus 7 cointegration vectors has a p -value of 0.33, which safely does not reject the null of 6 vectors.

In what follows, I will therefore proceed under the assumptions that

(I) either system features six, and only six cointegration vectors, and

(II) in line with the previously discussed results from unit root tests for the spreads, the cointegration vectors are the theoretical ones pertaining to the six pairs of nominal interest rates I am considering, that is, for either pair, $[1 \ -1]'$.

3 Methodology

3.1 Estimation

I estimate the cointegrated VAR in the VECM form, imposing in estimation the six theoretical cointegration vectors. The fact that I estimate the cointegrated system in the VECM form allows, as needed, to expand it with a vector of stationary series. Specifically, let Y_t be the vector of I(1) series, comprising, in the smaller system, the seven nominal interest rates and inflation, and in progressively larger systems also comprising hours, and then also GDP; and let S_t be the vector of I(0) series, comprising the unemployment rate and capacity utilization, and, in the largest system, also comprising the vacancy rate. Let Y_t and S_t be of dimension $(N_1 \times 1)$ and $(N_0 \times 1)$, respectively, with $N = N_1 + N_0$. The general form of the systems I estimate is then given by

$$\begin{bmatrix} \Delta Y_t \\ S_t \end{bmatrix} = B_0 + B_1 \begin{bmatrix} \Delta Y_{t-1} \\ S_{t-1} \end{bmatrix} + \dots + B_{p-1} \begin{bmatrix} \Delta Y_{t-(p-1)} \\ S_{t-(p-1)} \end{bmatrix} - BA'Y_{t-1} + \epsilon_t \quad (1)$$

¹¹I consider the maximum between the lag orders chosen by the SIC and HQ criteria because the risk associated with selecting a lag order smaller than the true one (model mis-specification) is more serious than the one resulting from choosing a lag order greater than the true one (over-fitting).

¹²On the other hand, I do not consider the Akaike Information Criterion since, as discussed (e.g.) by Luetkepohl (1991), for systems featuring I(1) series the AIC is an inconsistent lag selection criterion, in the sense of not choosing the correct lag order asymptotically.

where p is the lag order, A is a $(6 \times N_1)$ matrix whose columns are the six theoretical cointegration vectors, and B is the $(N \times 6)$ matrix of the loadings (i.e. of the coefficients translating the deviations from long-run equilibrium captured by the cointegration residuals into an adjustment towards such equilibrium). I select the VAR lag order as the maximum between the lag orders chosen by the Schwartz and the Hannan-Quinn criteria for the VAR in levels.

3.2 Identification

I identify the estimated system *via* long-run restrictions. My identifying assumptions are that

(1) the I(1) component of nominal interest rates is driven by two, and only two, permanent shocks: a permanent inflation shock (*via* the Fisher effect), and a permanent shock to the natural rate of interest.

(2) The permanent shock to the natural rate does not have any long-run impact on inflation. On the other hand, it is allowed to have a long-run impact on all other I(1) series (although, as we will see, none of these impacts is significantly different from zero).

The fact that, beyond the permanent inflation shock, the natural rate shock is defined as the only other shock having a permanent impact on nominal rates requires some discussion. The implication is that *all* other disturbances impacting upon nominal rates are transitory, including (e.g.) shocks to the risk premium. Now suppose—just for the sake of the argument—that shocks to the risk premium were partly transitory, and partly permanent (the argument applies to any other shock impacting upon nominal rates). Under these circumstances, the proposed identification strategy would interpret permanent risk premium shocks as natural rate shocks. The key issue here, however, is that—especially for practical (i.e., policymaking) purposes—this is *exactly* what we would want. The fact that, for a given equilibrium inflation rate, the equilibrium nominal Federal Funds rate increases by x per cent because of (say) a permanent risk premium shock—as opposed to a supposedly ‘authentic’ shock to the natural rate (due, e.g., to an acceleration of productivity growth)—is, for all practical purposes, completely irrelevant. What matters is that the equilibrium FED Funds rate has increased by x per cent. So, although in principle it is possible to draw fine distinctions about what, exactly, is behind the identified natural rate shocks, such distinctions are, for monetary policy purposes, essentially irrelevant.

3.3 Computing the natural rates corresponding to individual interest rates series

For each individual nominal interest rate R_t —with R_t being either the FED Funds rate, the 3- or 6- month Treasury bill rate, etc.—I compute the corresponding natural rate as the component of $R_t - \pi_t$ (with π_t being the inflation rate) which is uniquely

driven by natural rate shocks.¹³ In fact, inflation fluctuations due to natural rate shocks are close to being negligible,¹⁴ so that, in practice, the estimated natural rates are very close (up to a constant) to the components of the R_t 's which are uniquely driven by natural rate shocks.

3.4 Characterizing the extent of uncertainty *via* bootstrapping

In Section 5 I will characterize the extent of uncertainty surrounding the point estimates by bootstrapping the estimated system (1). The best way to think of this is simply as an extension of CRT's (2012) methodology for bootstrapping cointegrated VARs to a system which also comprises the stationary series collected in the vector S_t . Apart from this, there is no difference between CRT's methodology and what I am doing there.

I now turn to discussing the natural rate estimates for both the pre-ZLB and the ZLB periods; the impulse-response functions (henceforth, IRFs) to natural rate shocks; and the fractions of forecast error variance (henceforth, FEV) explained by these shocks.

4 Classical Point Estimates

4.1 Natural rate estimates for the pre-ZLB period

The first panel of Figure 1 shows, for the period up to October 2008, the *ex post* real Federal Funds rate, together with the point estimates of the *natural FED Funds rate*—i.e., the component of the *ex post* real FED Funds rate uniquely driven by natural rate shocks—produced by two of the models considered herein: the benchmark system (in red), and the one only including the nominal rates, inflation, and hours. The second panel, on the other hand, shows the point estimates of the natural FED Funds rate produced by four of the systems considered herein, which are representative of the

¹³To be precise, the way I implement this is the following. First, I re-run history only conditional on transitory interest rate shocks (that is: killing off both permanent inflation shocks, and natural rate shocks). From the difference between the actual series, $[Y_t, S_t]'$, and the thus-computed counterfactual series, I therefore obtain the components of $[Y_t, S_t]'$ which are uniquely driven by either of the permanent interest rate shocks. Finally, I run a counterfactual for these components in which I kill off the permanent inflation shocks, thus obtaining the components of $[Y_t, S_t]'$ which are uniquely driven by natural rate shocks. Another way of doing this would simply have been to re-run the history of the actual series by just killing off all shocks other than the natural rate shocks. This produces natural rate estimates which are near-identical to my benchmark estimates after a couple of years at the beginning of the sample, for which there is some minor difference.

¹⁴I do not report this evidence for reasons of space, but it is available upon request.




Figure 1 *Ex post* real Federal Funds rate, and Classical estimates of the natural rate of interest based on alternative models

overall set of estimates produced by all of the models I estimate.¹⁵

The main findings emerging from Figure 1 can be summarized as follows:

first, conceptually in line with the evidence produced (e.g.) by Laubach and Williams (2016) and Holston *et al.* (2016), the natural FED Funds rate is estimated to have decreased from about 3.4-3.5 per cent towards the end of 1964, to between 1.2 and 1.6 per cent (depending on the model) in October 2008. In particular, based on the benchmark model with nominal rates, inflation, hours, the unemployment rate, and capacity utilization, the natural rate is estimated to have fallen slightly below 3 per cent at the beginning of the 1970s, and to have remained mostly there until 1984; to have then progressively decreased until the early 1990, reaching a trough of about 2.2 per cent; to have temporarily surged beyond 2.5 per cent in the first half of the 1990s; and to have then fallen all the way down to about 1.2 per cent in October 2008, with only a temporary increase around the mid-2000s.

Second, the different systems often produce materially different natural rate estimates. In particular, the estimate produced by the smallest system—i.e., the one only featuring the nominal rates and inflation (see the brown line)—exhibits a sharp fall from the mid-1960s to the mid-1970s, and a much gentler decline since then. Augmenting this system with hours (blue line) has no material impact until the mid-1980s, but over the following period it produces a somehow higher path, which is very broadly constant between the second half of the 1980s and the early 2000s. Further augmenting the system with unemployment and capacity utilization (red line: the benchmark system) produces the completely different, previously-mentioned path. Finally, augmenting the benchmark system with either GDP (black line) or GDP and the vacancy rate (not shown, but available upon request), does not produce any material difference compared with the benchmark model.

4.2 Which system produces the more reliable estimates?

The fact that different systems tend to produce materially different estimates raises two questions:

- (i) which estimates are more reliable? And
- (ii) what is the minimal system producing reliable estimates?

Figure 2 and Table 4 show evidence on this. Specifically, Figure 2 shows, based on either of the four systems reported in the second panel of Figure 1, the estimated natural rates—defined, as in Section 3.3, as the components of the $(R_t - \pi_t)$'s uniquely driven by natural rate shocks—for the FED Funds rate and the 3-, 5-, and 10-year government bond yields.¹⁶ Table 4, on the other hand, reports the simple contemporaneous correlation between the first differences of the natural FED Funds rate, and

¹⁵In order to make the figure more intelligible, all series have been smoothed with a 3-month rolling moving average.

¹⁶As In figure 1, in order to make the figure more intelligible, all series have been smoothed with a 3-month rolling moving average.

Estimates based on model with:




Figure 2 Classical estimates of the natural rates of interest for the Federal Funds rate, and for 1-, 3-, 5-, and 10-year maturities, based on alternative models

Table 4 Simple correlation between the first differences of the natural FED Funds rate, and of natural rates for alternative maturities

[illegible]

of the natural rates for alternative maturities.

Since the natural rates at the various maturities are all uniquely driven by the very same permanent natural rate shocks, their month-to-month changes should be, in principle, perfectly correlated, and their levels should be identical up to a constant capturing (e.g.) maturity-specific risk premia. To put it differently, a positive (negative) natural rate shock should shift the entire set of natural rates at the various maturities upwards (downwards) by an identical amount.¹⁷

Based on this metric, it is apparent from both the table, and especially the figure, that systems smaller than the benchmark one produce manifestly less reliable estimates. For example, based on the smallest system (with just nominal rates and inflation) the correlations between the first differences of the natural FED Funds rate, and of natural rates for alternative maturities, range between 0.825 and 0.934. Adding hours improves things quite significantly, with the correlations now ranging between 0.959 and 0.992. These correlations, however, are still somehow lower—although by comparatively small amounts—than the corresponding ones for the benchmark system, which range between 0.986 and 0.997. Finally, adding GDP or even (not reported) the vacancy rate to the benchmark system only leads to negligible decreases in the correlations.

The evidence in Figure 2 is even clearer than in Table 4, with the time-profiles of the estimated natural rates produced by the benchmark system (or by the benchmark system augmented with GDP) for the various maturities being very strongly correlated, and pretty much a scaled-up, or scaled-down version of one another.¹⁸ Smaller systems, on the other hand, tend to produce natural rates' paths which clearly are not a scaled-up or scaled-down version of one another, and exhibit, most of the times, a markedly different profile. This is already apparent for the system with nominal rates, inflation, and hours, and it becomes dramatically so for the smallest model with just inflation and nominal rates.

Based on this evidence, from now on I will exclusively focus on the results produced by the benchmark system.

I now turn to estimating the natural rate for the ZLB period.

4.3 Estimating the natural rate for the ZLB period

As previously mentioned, throughout the entire paper estimation is performed uniquely up to October 2008, the last month for which we can reasonably be sure that the ZLB was not binding.¹⁹ In order to estimate the natural FED Funds rate for the period starting in November 2008, I therefore proceed as follows.

¹⁷The fact that the impact should be identical for all rates originates from the fact that—as discussed in Section 3.1—we are here imposing the theoretical cointegration vectors, which for any rates' pair is $[1 \ -1]'$.

¹⁸With the exception of the first 3-4 years up to 1967-1968, for which the correlations appear weaker.

¹⁹In that month, the FED Funds rate was equal, on average, to 0.97 per cent.

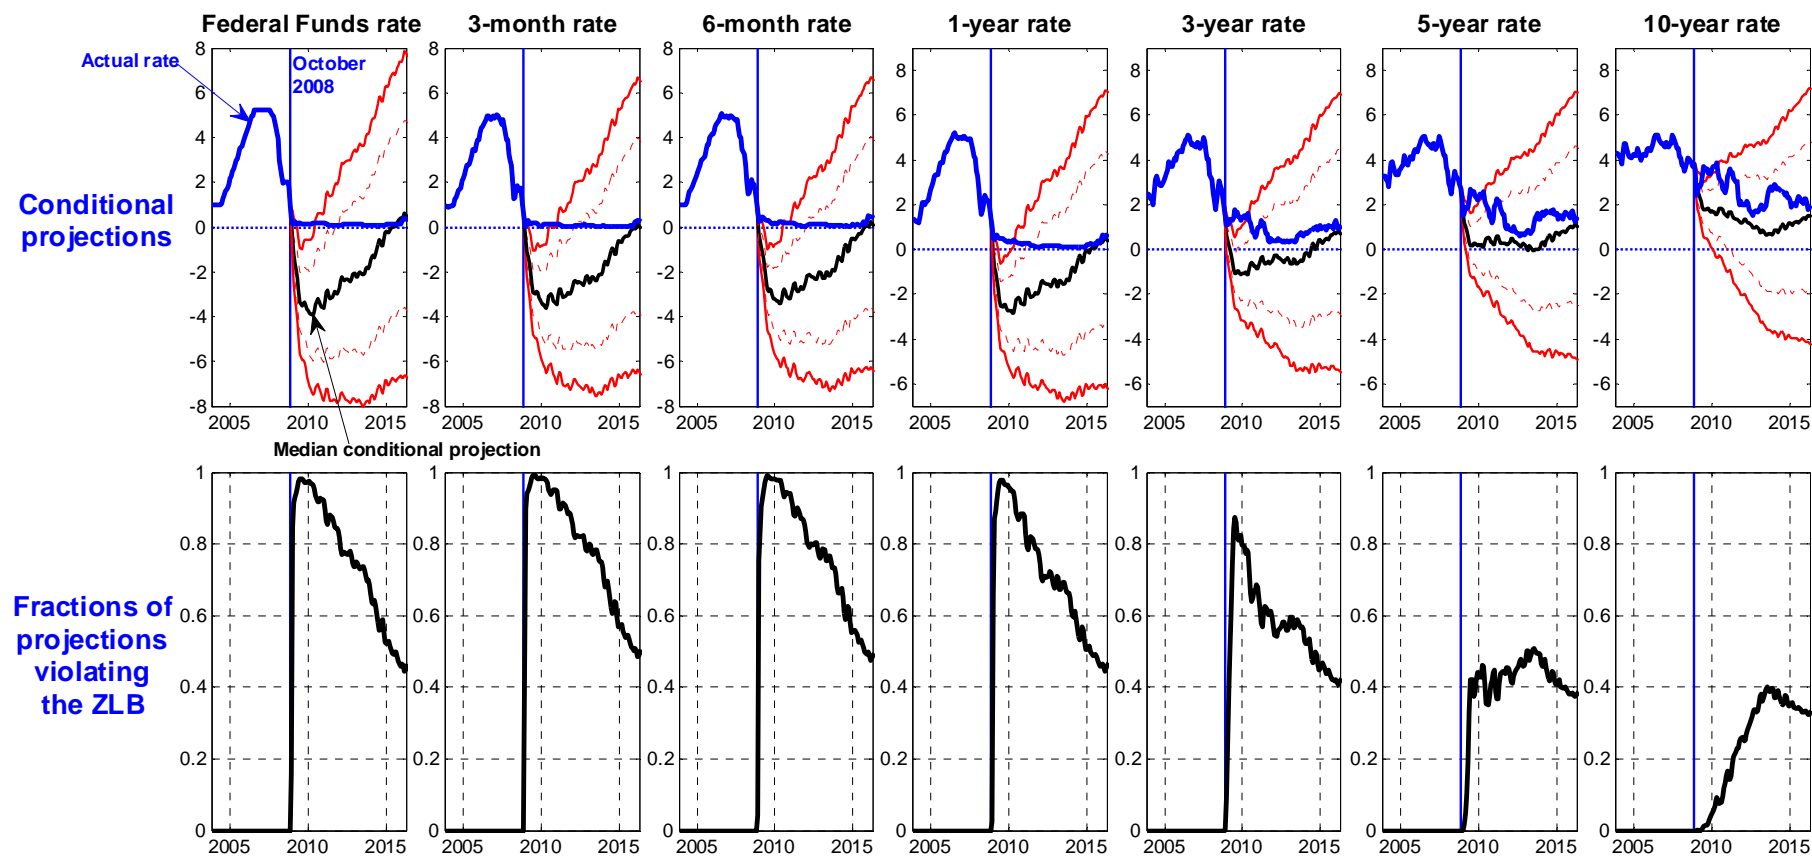


Figure 3 Classical unconstrained interest rates' projections for the period since October 2008 conditional on the actual dynamics of inflation and real activity series (with 16-84 and 5-95 bootstrapped percentiles)

Taking as given the data up to October, starting in November I bootstrap the estimated system into the future²⁰ up until the end of the available sample (May 2016) *conditional*²¹ on the actual dynamics of all variables other than the nominal interest rates.²² The rationale for doing so is that we have strong reasons to believe that, since November 2008, the ZLB has been constraining not only the dynamics of the FED Funds rate and of short-term rates (such as the 3- and 6-month Treasury bill rate), but also rates at longer maturities (see in particular some of the evidence reported in Swanson and Williams (2014)). Figure 3 provides support to this conjecture. The first row of the figure shows, for either of the seven nominal rates, the median conditional projection, together with the 16-84 and the 5-95 percentiles of the bootstrapped distribution, whereas the second row shows, for each individual month, the fraction of bootstrapped projections which violate the ZLB. Evidence that the ZLB may have been binding is very strong for all maturities up to 3 years ahead, whereas—in line with Swanson and Williams (2014)—it is weaker for longer maturities. In particular, for the 10-year government bond yield, evidence that the ZLB may have been binding is virtually non-existent for the first 3-4 years since November 2008, but it steadily creeps up going forward, peaking at about 40 per cent in 2014. Because of this, in order to be absolutely sure that in *no way* natural rate estimates for the period since November 2008 have been distorted by nominal rates being constrained by the ZLB, in what follows I will work with the just-described conditional projections, thus allowing for the possibility that, in fact, *all* nominal rates may have been constrained, at some time, by the ZLB.

For each bootstrapped replication j , for $j = 1, 2, \dots, 10,000$, I then augment the associated conditional projections for the nominal rates with the actual historical values taken by the remaining series, thus obtaining, for each replication, a system identical to (1), with the *only* difference that the nominal rates' paths have been simulated, rather than being actual data. I then impose the same identifying restrictions I previously imposed upon the actual data up to October 2008 upon each of the resulting, partially forecasted systems, thus obtaining a bootstrapped distribution for either of the natural rates' point estimates at the various maturities.²³ Figure 4

²⁰Based on 10,000 bootstrap replications.

²¹Specifically, the projections for the nominal rates conditional on the actual historical dynamics of inflation, hours, the unemployment rate, and the rate of capacity utilization have been computed based on the estimated 'conditional version' of the benchmark system, that is: the version of (1) in which both S_t and all of the variables in ΔY_t other than the nominal rates, have been moved to the right-hand side.

²²To be precise: In doing so, I reject all bootstrapped paths for which the unemployment rate turns out to have been negative for at least one month. In fact, out of 10,000 bootstrapped paths, the fraction of rejected paths is negligible, for the simple reason that this period was dominated by the Great Recession. By the same token, I imposed rejection of all paths for which the unemployment rate was greater than 1, or the capacity rate was outside the $[0, 1]$ interval. These constraints were never violated.

²³To be absolutely clear: the fact that we are here dealing with *distributions* of point estimates originates from the fact that nominal rates' paths have been forecasted, rather than being actual

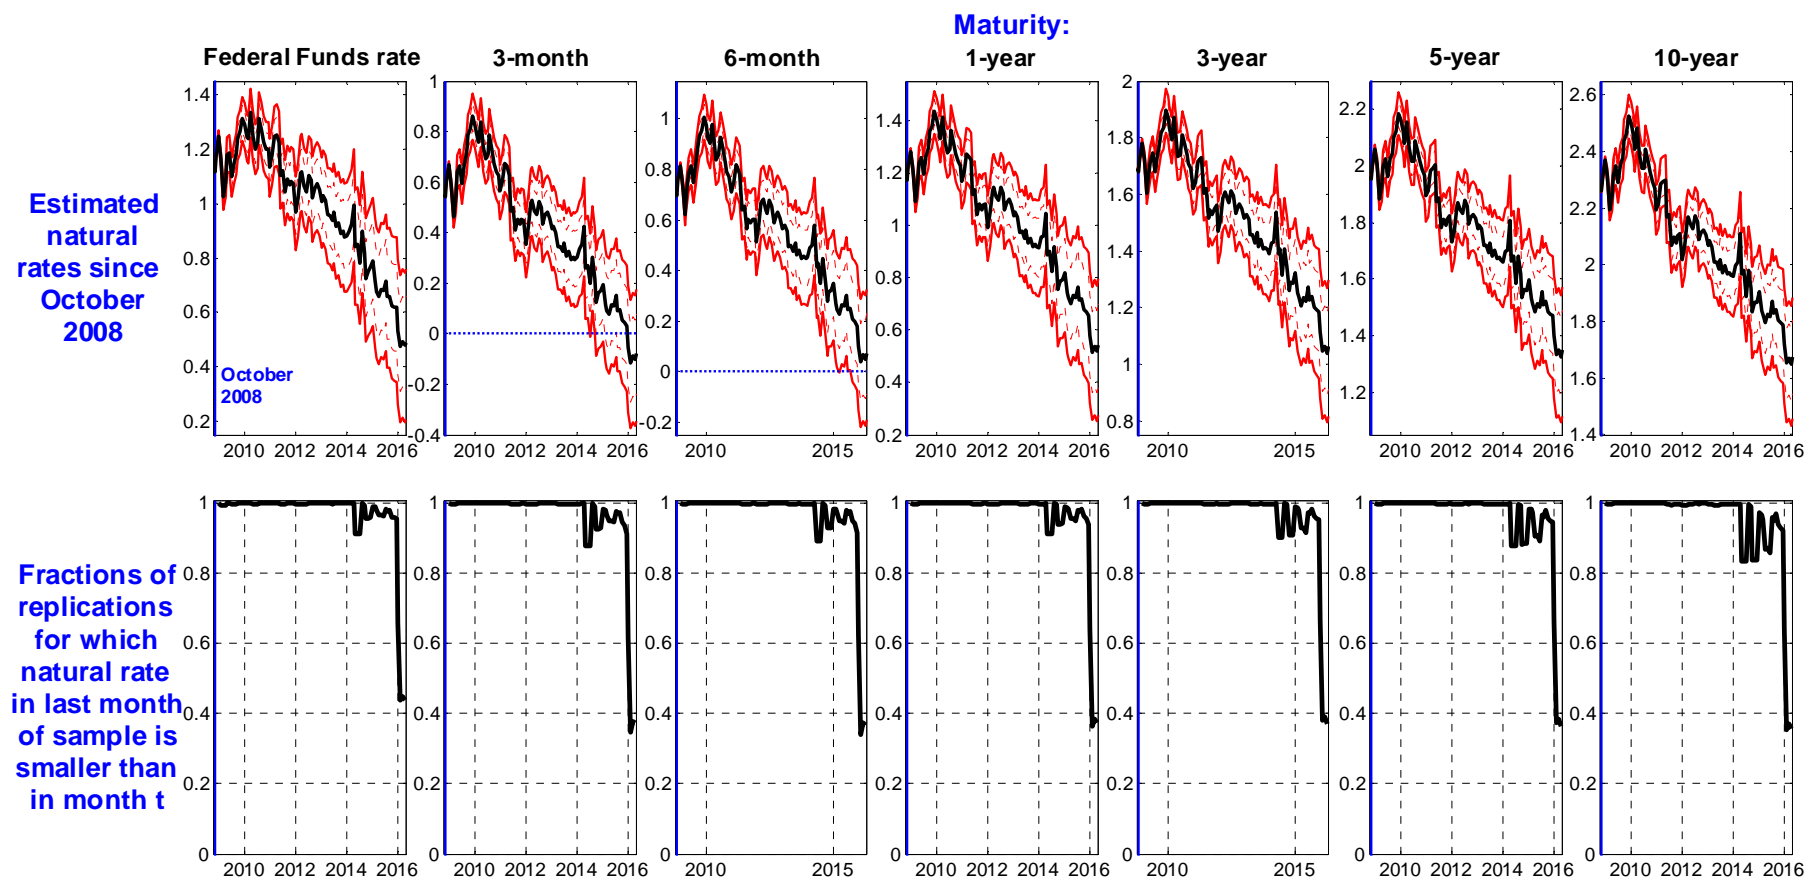


Figure 4 Classical estimates of the natural rates at the various maturities for the period since October 2008 (medians, and 16th-84th and 5th-95th percentiles of bootstrapped distributions, abstracting from model uncertainty)

shows the results. The first row shows, for either maturity from the FED Funds rate to the 10-year government bond yield, the median point estimate of the corresponding natural rate, together with the 16-84 and the 5-95 percentiles of the bootstrapped distributions of the point estimates. The second row provides statistical evidence on whether the *point estimate*²⁴ of the natural rate has been decreasing since November 2008, by showing the fractions of replications for which the natural rate’s point estimate in the last month of the sample (May 2016) was lower than it had been in month t . Evidence of a decrease in natural rates is uniformly very strong, with the fractions in the second row being very close to one up until the end of 2013. As for the point estimates, the natural FED Funds rate is estimated to have fallen from about 1.2 per cent in November 2008 to about 0.5 at the end of the sample.

4.4 Impulse-response functions and variance decompositions

Figure A.1 in the appendix shows, for either of the variables in the benchmark model, bias-corrected²⁵ IRFs to a normalized one per cent shock to the natural rate of interest, together with the simple estimates, and with the 16-84 and 5-95 percentiles of the bootstrapped distribution. All IRFs are exactly as expected based on economic logic, that is, natural rate shocks *only* have a long-run unitary impact on all of the nominal rates (by construction), but other than that they do not have any statistically significant impact on any other variable.

Figure A.2 in the appendix shows, for either variable, the fractions of FEV explained by natural rate shocks, together with the 16-84 and 5-95 percentiles of the bootstrapped distribution. *Point estimates* uniformly point towards natural rate shocks explaining minor-to-negligible fractions of the FEV of all variables—including nominal rates—at all horizons. For variables other than nominal rates the extent of econometric uncertainty is small-to-negligible. For nominal rates, on the other hand, uncertainty is quite substantial: for the FED Funds rate, in particular, the bootstrapped 84th and 95th percentiles at the 10-year horizon are just shy of 20 per cent, and beyond 50 per cent, respectively, thus pointing towards a potentially non-negligible role played by natural rate shocks in driving the FED funds (and other nominal rates).

I now turn to characterizing uncertainty around the estimates *via* bootstrapping methods.

data. So these distributions do *not* encode the extent of sample uncertainty, which we will characterize in Section 5 *via* bootstrapping.

²⁴So, again, this evidence does not take into account of sample uncertainty.

²⁵Bias-correction has been implemented as in Kilian (1998).




Figure 5 Characterizing uncertainty around the Classical benchmark estimates *via* bootstrapping

5 Characterizing Uncertainty Around the Classical Estimates *via* Bootstrapping

Figure 5 shows, for the FED Funds rate and for the 1-, 5-, and 10-year government bond yields, the respective estimated natural rates²⁶ together with the 16th and 84th percentiles of the bootstrapped distributions (bootstrapping has been implemented exactly as before; see Section 3.4). The main finding emerging from the four panels is the uniformly sizeable extent of econometric (i.e., sample) uncertainty surrounding the point estimates. For the natural FED Funds rate, in particular, a 68 per cent-coverage confidence interval stretches from 0.01 to 2.29 per cent for October 2008, whereas in the middle of the sample (in May 1986) uncertainty is even larger, with the corresponding confidence interval going from -1.35 to 6.05 per cent. Intuitively, the reason for such a large extent of uncertainty is that, as discussed in Section 4.4, point estimates suggest that natural rate shocks explain uniformly small fractions of the FEV of any variable. For the FED Funds rate and all other nominal rates, in particular, the fractions are uniformly below 5 per cent, which logically implies that estimating *anything* which pertains to these ‘small’ shocks is going to be a challenge.²⁷

Such a sizeable extent of uncertainty raises the question of whether there might be a way to decrease it, thus providing policymakers with more precise estimates of the natural rate. Before addressing this issue, in Section 5.2, I next discuss statistical evidence of a decrease in the natural FED Funds rate.

5.1 Statistical evidence of a decrease in the natural FED Funds rate

The bottom-right panel of Figure 5 reports the fraction of bootstrapped replications for which the natural rate in October 2008 was lower than in month t (the black line); and the corresponding fractions of replications for which the natural rate in month t was lower than it had been either 5 or 10 years before (the blue and red lines, respectively). Evidence of a decrease in the natural rate since the start of the sample is strong, with the fraction of replications for which in October 2008 it was lower than in month t being, at the very beginning, close or even beyond 90 per cent. Going forward, evidence becomes progressively weaker, but it remains almost uniformly beyond 60 per cent of the replications, and, up until the end of 1983, beyond 70 per cent. Evidence based on the other two measures is uniformly weaker, and most of the times just barely above the 50 per cent threshold.

In assessing such comparative weakness of the statistical evidence of a decrease it is important to keep in mind what I previously pointed out about natural rate

²⁶All of the series shown in Figure 6 have been smoothed, as before, *via* a 3-month rolling window, in order to make them more intelligible.

²⁷Canova and Paustian (2011) contain a conceptually related discussion of the problems posed by ‘small’ shocks within the context of identification *via* sign restrictions.

shocks being ‘small’: given the limited role they most likely play in driving the system, it is not surprising that not only natural rate estimates are characterized by a sizeable extent of uncertainty, but also evidence of a decrease in the natural rate is not overwhelming.

This brings us to the issue of whether it might be possible to shrink uncertainty, to which I now turn.

5.2 Is there a way to decrease the extent of uncertainty?

The key to shrink the uncertainty surrounding point estimates lies in recognizing that some of the paths produced by bootstrapping are manifestly implausible. For example, in the first panel of Figure 5, any reasonable observer would likely deem the 84th percentile of the bootstrapped distribution as an implausible estimate of the natural rate. The key issue here is how to go beyond such ‘I know it when I see it’²⁸ approach, thus making the notion of ‘implausibility’ operational. The reason why I regard the just-mentioned 84th percentile in Figure 5 implausible is because it cannot possibly be a trend for the real *ex post* FED Funds rate. Another way of seeing this is that the 84th percentile automatically implies an almost uniformly negative transitory component of the real *ex post* FED Funds rate: in the same way as an output gap estimate which had consistently been negative over the entire post-WWII period should be regarded as implausible, we should here regard as equally unlikely paths for the natural rate which map into implausible paths for the transitory component of the real *ex post* FED Funds rate.

Figure 6 shows the bootstrapped distribution and cumulative distribution of the *absolute value* taken by the *average real FED Funds rate gap* over the entire sample period, where the real FED Funds rate gap is defined, for each individual bootstrapped replication, as the difference between the real *ex post* FED Funds rate and the estimated natural FED Funds rate for that replication. I focus on the absolute value because (say) an average real FED Funds rate gap of either 10 or -10 per cent should be regarded as equally implausible. As the figure shows, a non-negligible, and in fact large fraction of bootstrap replications is associated with comparatively large values (either positive or negative) of the average real FED Funds rate gap, with (e.g.) about one-third of the replications being associated with values beyond 2 per cent.

In order to gauge an idea of how (im)plausible these figures are, let’s consider a standard estimate of the U.S. output gap, i.e. the one implied by the *Congressional Budget Office’s* (CBO) estimate of potential GDP. Over the period used here for estimation purposes (1964Q1-2008Q3), the CBO output gap estimate had been equal, on average, to -0.38 per cent, with a standard deviation of 2.40 per cent. Over the same period, the real *ex post* FED Funds rate (converted to the quarterly frequency) had been equal on average to 1.91 per cent, with a standard deviation of 2.28 per

²⁸As in U.S. Supreme Court Justice Potter Stewart’s definition of obscenity.

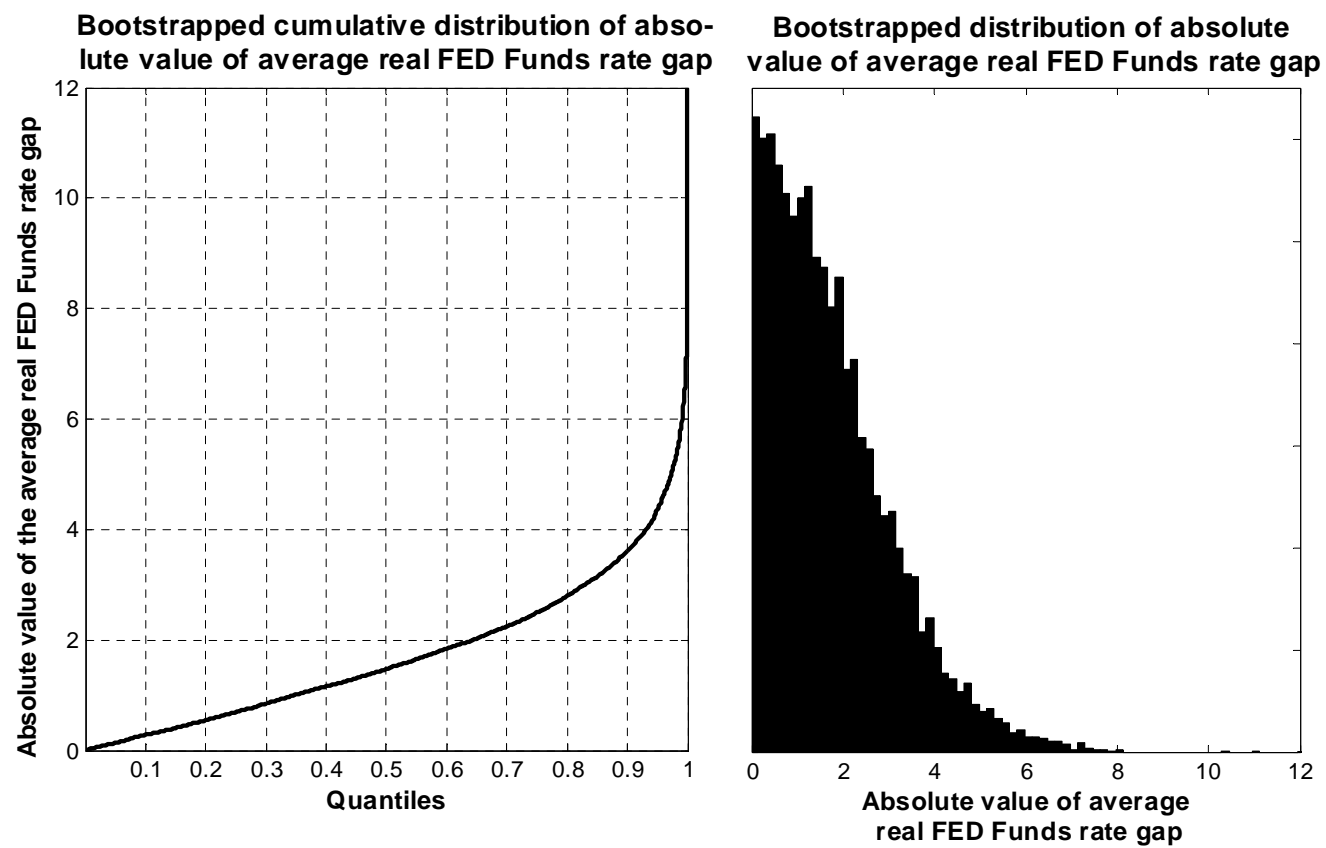


Figure 6 Bootstrapped distribution of the absolute value of the average natural FED Funds rate gap

cent. Since any plausible estimate of the natural FED Funds rate should behave like a slow-moving time-varying mean of the real *ex post* FED Funds rate, it logically follows that the standard deviation of the real FED Funds rate gap should not be too different from the standard deviation of the real *ex post* FED Funds rate, that is, it should be roughly around 2 per cent.²⁹ Since this is somehow smaller than the standard deviation of the CBO estimate of the output gap, it logically follows that the 0.38 per cent (in absolute value) average value for the output gap should be regarded as an *upper bound* for the corresponding average real FED Funds rate gap.³⁰

What precedes suggests that a simple way to make the notion of ‘(im)plausibility’ of a specific natural FED Funds rate estimate operational is to rule out (i.e., reject) natural rate paths implying average real FED Funds rate gaps which are (in absolute value) ‘way beyond’ 0.4 per cent. Although, from a conceptual point of view, imposing prior information³¹ upon the data is best implemented within a Bayesian framework, as I will show in the next Section a Bayesian approach suffers from the crucial drawback that median estimates of the natural FED Funds rate turn out to depend, in a material way, on the prior which is being imposed upon the absolute value of the average real FED Funds rate gap. This is why I will mostly discuss this issue within a Classical context.

Figure 7 shows the Classical natural FED Funds rate estimates, together with the 16-84 and 5-95 percentiles of the bootstrapped distributions which are obtained by imposing progressively larger upper bounds on the absolute value of the average real FED Funds rate gap. Either of the six bootstrapped distributions whose percentiles are being shown there has been obtained by retaining, out of the bootstrapped distribution whose 16-84 percentiles were shown in the first panel of Figure 5, only the replications satisfying the constraint. As the figure shows, the imposition of plausible prior information on the average real FED Funds rate gap allows to shrink the extent of uncertainty by substantial amounts. In particular, imposing an upper bound of 0.5 per cent—which, based on the previous discussion, should be regarded as entirely reasonable—allows to ‘compress’ the 16-84 percentiles by a sizeable amount compared to the first panel of Figure 5. Once again, however, it is to be stressed how uncertainty is still large, with 16-84 percentiles being typically about 4 percentage points wide.

²⁹Since time-variation in the natural FED Funds rate will cause the standard deviation of the real FED Funds rate gap to be necessarily smaller than the standard deviation of the real *ex post* FED Funds rate.

³⁰The intuition here is that, when the economy is in equilibrium, the transitory components of GDP, the real *ex post* FED Funds rate, and the unemployment rate should all be zero, whereas deviations of GDP from potential, unemployment from the natural rate, and the real *ex post* FED Funds rate from the natural rate should all be strongly correlated.

³¹It is to be noticed that such prior information has been elicited from the CBO’s output gap, i.e. from a series which is here *not* being used in estimation.

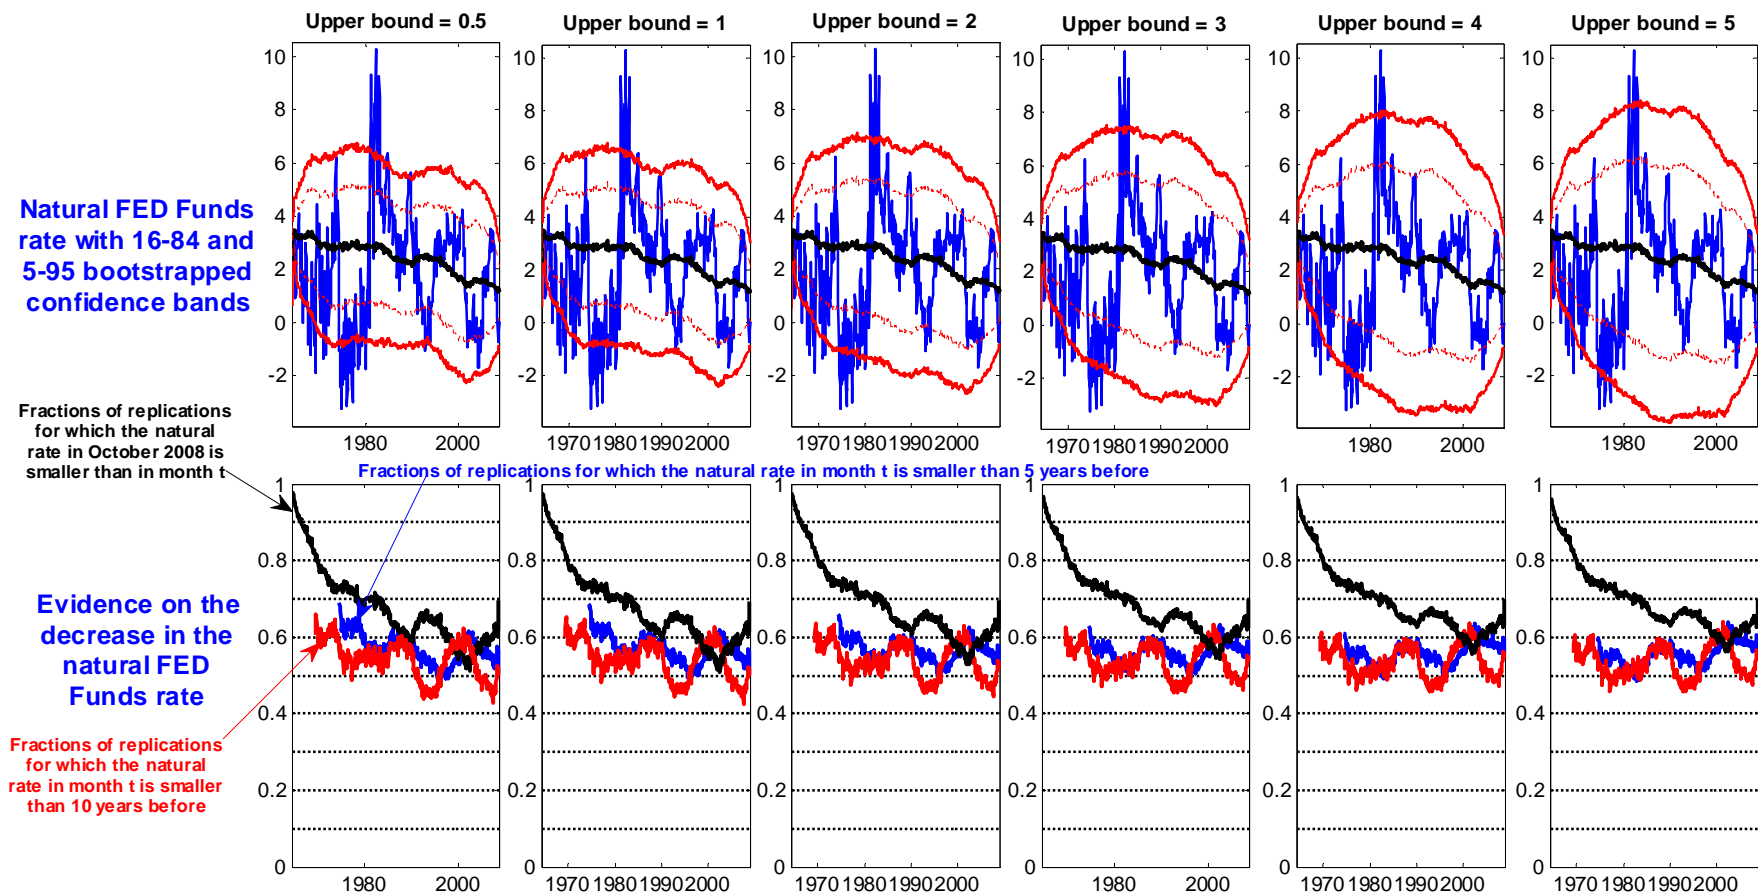


Figure 7 Classical results based on alternative upper bounds on the absolute value of the average transitory FED Funds rate gap

6 Bayesian Estimates

I finally turn to the Bayesian estimates, which, as previously mentioned, perform in an unsatisfactory way, as in general they depend in a non-negligible way on the priors. I estimate the reduced-form system (1) *via* the Bayesian estimator detailed in Uhlig (1998, 2005), and I consider several alternative types of priors:

(i) priors which (as in the previous section) rule out models producing average real FED Funds rate gaps whose absolute value exceeds a certain threshold.

(ii) ‘Tunnel priors’, according to which the natural FED Funds rate path should lie inside a pre-specified ‘tunnel’. I consider two alternative ways of setting up such a tunnel. First, I simply define intervals $[K_{\text{lower}}, K_{\text{upper}}]$, for alternative values of K_{lower} and K_{upper} , and I reject all draws associated with natural FED Funds rate paths which wander outside the interval. Second, I consider arguably more plausible specifications, in which, for each month t , the lower and upper bounds of the tunnel are given by $K_t^{\text{lower}} = \gamma_t - K$ and $K_t^{\text{upper}} = \gamma_t + K$, where γ_t is the value taken in month t by a linear time trend for real GDP growth³² (which is estimated *via* OLS), and K is a scalar. The rationale for doing so is that, as mentioned in the introduction, and as discussed e.g. by LW (2003), standard growth models imply a linear relationship between trend GDP growth and the natural rate. Further, under plausible assumptions, the natural rate should be not too different from trend GDP growth. The ‘tunnel prior’ is designed to capture both features.

(iii) Finally, I further add, on top of the ‘tunnel priors’, priors on the extent of time-variation in the natural rate, which I implement by imposing upper bounds on the standard deviation of the month-on-month change in the natural rate.

In a nutshell, none of these approaches produces robust results, as the estimates of the natural FED Funds rate turn out to depend, in a non-negligible way, on the priors which are being imposed in order to rule out implausible paths for the real FED Funds rate gap (I do not report these results for reasons of space, but all of them are available upon request).

My own conclusion is therefore that for practical (i.e., policymaking) purposes, the best approach is probably the one I detailed in Section 5.2, based on the notion of just rejecting all of the bootstrapped replications associated with implausibly large values of the real FED Funds rate gap. Intuitively, a Classical approach does not suffer from the just-mentioned problem of dependence on the priors which instead plagues the Bayesian approach for the very simple reason that, within this approach, first you compute the point estimates, and then you characterize uncertainty around them *via* bootstrapping. As a consequence, within a Classical approach the point estimates are independent of the way uncertainty is characterized and somehow ‘limited’. Within a Bayesian approach, on the other hand, the estimate of the natural rate is the median (or mode, or mean) of the posterior distribution, and it therefore may depend in a

³²So, once again, the prior here is a proper one, since it has been set up based on a series (real GDP growth) which is not being used in estimation.

material way on how the priors are set up. Therefore, as previously mentioned, in practice the priors which are being imposed in order to limit the extent of uncertainty end up affecting natural rate estimates in a non-negligibly way.

7 Conclusions

In this paper I have illustrated a novel, and straightforward method to extract the natural rate of interest from the joint dynamics of inflation and the term structure of nominal interest rates, based on cointegrated structural VARs. The key identifying assumption is that the unit root component of nominal interest rates is driven by two permanent shocks, an inflation shock, and a shock to the natural rate, which can be disentangled via standard SVAR techniques. I estimate the natural rate for the Zero Lower Bound (ZLB) period by simulating the pre-ZLB estimated system conditional on the actual dynamics of the unconstrained variables at the ZLB, and then imposing the identifying restrictions upon the conditional projections. Evidence suggests that since the beginning of the financial crisis the natural rate has decreased, in the United States, by 0.5-0.7 percentage points, and it stands, in the second half of 2016, at about 0.5 per cent.

References

- ANTOLIN-DIAZ, J., T. DRECHSEL, AND I. PETRELLA (2016): “Tracking the Slow-down in Long-Run GDP Growth,” Hanyang University and Chinese University of Hong Kong, mimeo.
- BARSKY, R., A. JUSTINIANO, AND L. MELOSI (2014): “The Natural Rate of Interest and Its Usefulness for Monetary Policy,” *American Economic Review*, 104(5), 37–43.
- BENATI, L., R. E. LUCASJR., J.-P. NICOLINI, AND W. WEBER (2016): “International Evidence on Long-Run Money Demand,” *NBER Working Paper 22475*, July 2016.
- CANOVA, F., AND M. PAUSTIAN (2011): “Measurement with Some Theory: Using Sign Restrictions to Evaluate Business Cycle Models,” *Journal of Monetary Economics*, 58, 345–361.
- CAVALIERE, G., A. RAHBK, AND R. TAYLOR (2012): “Bootstrap Determination of the Co-Integration Rank in Vector Autoregressive Models,” *Econometrica*, 80(4), 1721–1740.
- ELLIOT, G., T. ROTHENBERG, AND J. H. STOCK (1996): “Efficient Tests for an Autoregressive Unit Root,” *Econometrica*, 64(4), 813–836.
- HAMILTON, J. (1994): *Time Series Analysis*. Princeton, NJ, Princeton University Press.
- HAMILTON, J. D., E. S. HARRIS, J. HATZIUS, AND K. D. WEST (2016): “The Equilibrium Real Funds Rate: Past, Present and Future,” Universities of California at San Diego and Wisconsin at Madison, Bank of America Merrill Lynch, and Goldman Sachs, mimeo.
- HOLSTON, K., T. LAUBACH, AND J. C. WILLIAMS (2016): “Measuring the Natural Rate of Interest: International Trends and Determinants,” *Finance and Economics Discussion Series 2016-073*. Washington: Board of Governors of the Federal Reserve System.
- JOHANSEN, S. (2002): “A Small Sample Correction of the Test for Cointegrating Rank in the Vector Autoregressive Model,” *Econometrica*, 70, 1929–1961.
- KILIAN, L. (1998): “Small-Sample Confidence Intervals for Impulse-Response Functions,” *Review of Economics and Statistics*, pp. 218–230.
- LAUBACH, T., AND J. WILLIAMS (2003): “Measuring the Natural Rate of Interest,” *The Review of Economics and Statistics*, 85(4), 1063–1070.

- LAUBACH, T., AND J. C. WILLIAMS (2016): “Measuring the Natural Rate of Interest Redux,” *Business Economics*, 51(2), 57–67.
- LUETKEPOHL, H. (1991): *Introduction to Multiple Time Series Analysis, II edition*. Springer-Verlag.
- SWANSON, E. T., AND J. C. WILLIAMS (2014): “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates,” *American Economic Review*, 104(10), 3154–3185.
- UHLIG, H. (1998): “Comment On: The Robustness of Identified VAR Conclusions About Money,” *Carnegie-Rochester Conference Series on Public Policy*, 49, 245–263.
- (2005): “What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure,” *Journal of Monetary Economics*, 52(2), 381–419.
- WU, J. C., AND F. D. XIA (2016): “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” *Journal of Money, Credit and Banking*, 48(2-3), 253–291.

A The Data

A monthly seasonally unadjusted series for the Federal Funds rate (the acronym is FEDFUNDS) is from the Board of Governors of the Federal Reserve System, and it is available since July 1954. Monthly seasonally unadjusted series for the 3- and the 6-month Treasury bill rates (the acronyms are TB3MS and DTB6, respectively), available since January 1934 and January 1959, respectively, are from the St. Louis FED. Monthly seasonally unadjusted series for the 1-, 3-, 5-, and 10-year Treasury constant maturity rates (the acronyms are GS1, GS3, GS5, and GS10, respectively), all available since April 1953, are from the St. Louis FED. A monthly seasonally unadjusted population series (CNP16OV), available since January 1948, is from the U.S. Department of Labor: Bureau of Labor Statistics. A monthly seasonally adjusted series for the civilian unemployment rate (UNRATE), available since January 1948, is from the U.S. Department of Labor: Bureau of Labor Statistics. A monthly seasonally adjusted series for the core CPI ('CPILFESL: Consumer Price Index for All Urban Consumers: All Items Less Food and Energy'), available since January 1947, is from the US. Bureau of Labor Statistics. The monthly seasonally unadjusted series for the Wu-Xia shadow Federal Funds rate is from the Federal Reserve Bank of Atlanta's website. A monthly seasonally adjusted series for the aggregate weekly hours of production and nonsupervisory employees in total private industries (AWHI), available since January 1964, is from the U.S. Department of Labor: Bureau of Labor Statistics. A monthly seasonally adjusted series for the rate of capacity utilization in the manufacturing sector (CUMFNS), available since January 1948, is from the Board of Governors of the Federal Reserve System. A monthly seasonally adjusted interpolated series for real GDP is from Mark Watson's homepage for the period January 1959-June 2010, and it is from *Macroeconomic Advisers* since then. Over the period of overlapping (January 1992-June 2010), the contemporaneous correlation between the log-differences of the two series is equal to 0.9094, which justifies their linking. A monthly seasonally adjusted series for the civilian labor force (CLF16OV), available since January 1948, is from the U.S. Department of Labor: Bureau of Labor Statistics. A monthly seasonally adjusted help wanted index, available for the period January 1951-December 2014, is from the Conference Board until 1995, from Regis Barnichon's website after that. The vacancy rate has been computed as the ratio between the help wanted index and civilian labor force. A quarterly seasonally adjusted series for potential GDP (GDPPOT) is from the *Congressional Budget Office* (CBO), and it is available since 1949Q1. The output gap estimate mentioned in Section 5.2 has been computed as the percentage difference, for each quarter, between real GDP ('GDPC96, Gross Domestic Product, Seasonally Adjusted Annual Rate, Quarterly, Billions of Chained 2009 Dollars') from the U.S. Department of Commerce: Bureau of Economic Analysis, and the CBO's potential GDP estimate.

Online Appendix

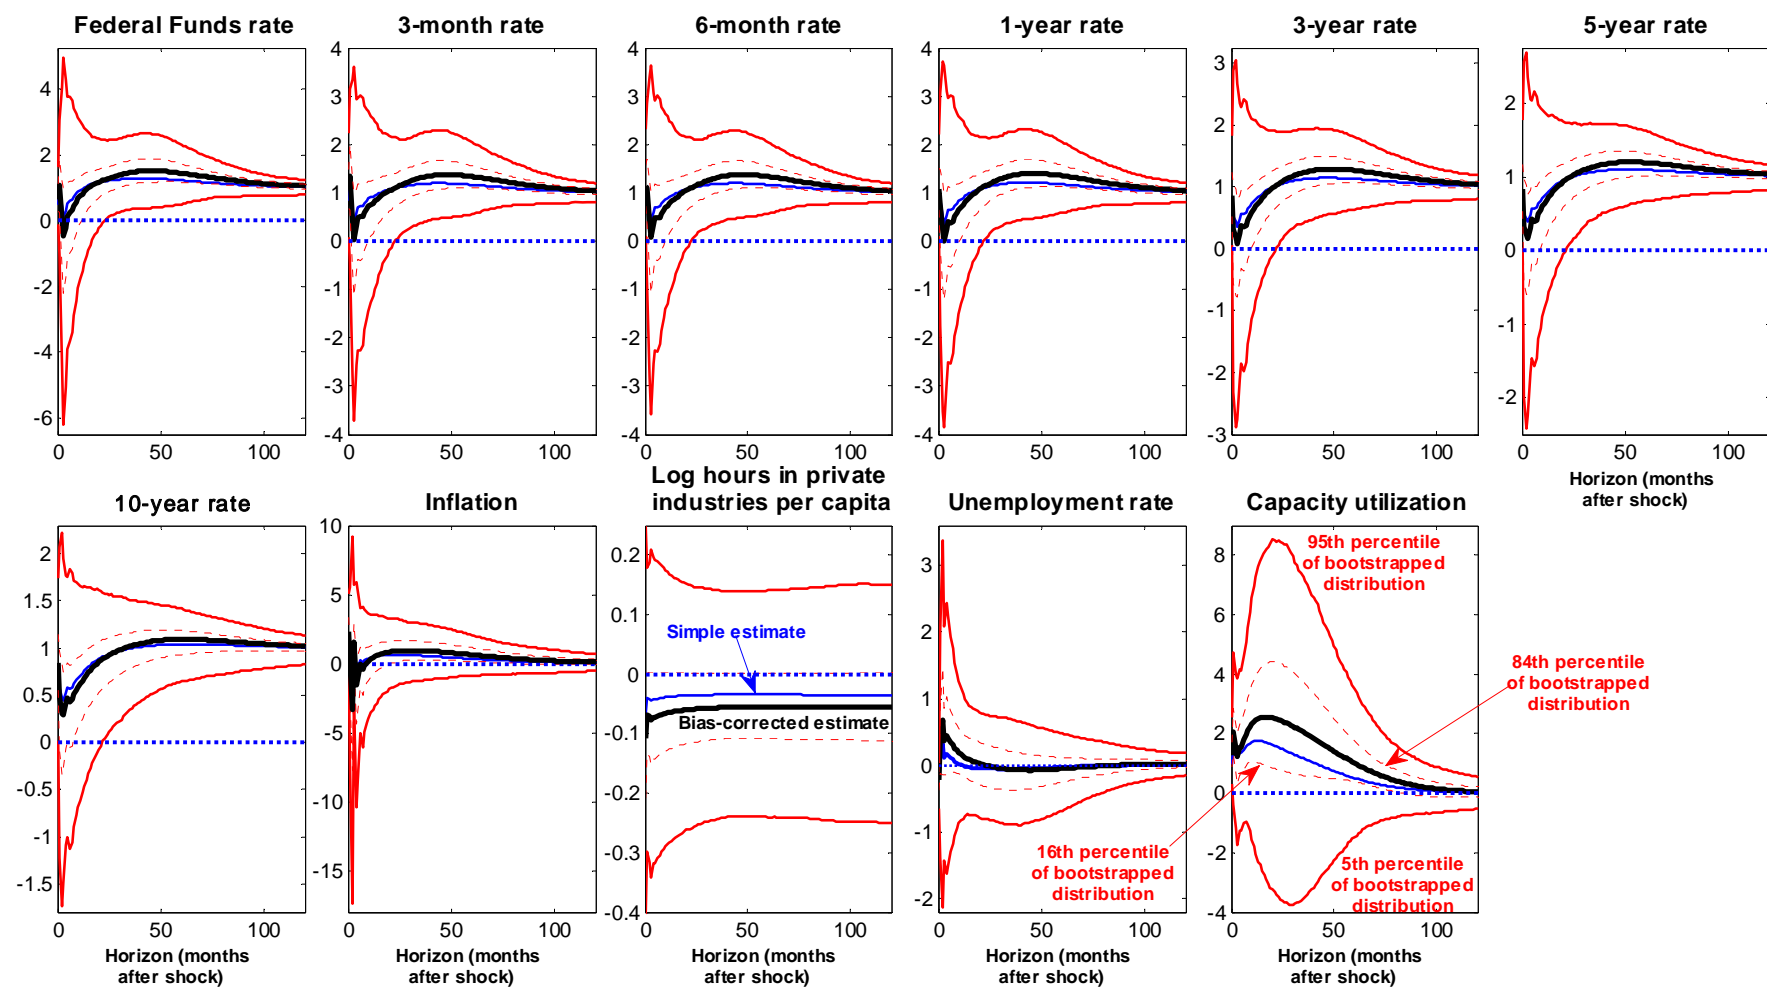


Figure A.1 Classical estimates based on the benchmark model: Impulse-response functions to a normalized 1% permanent shock to the natural rate of interest

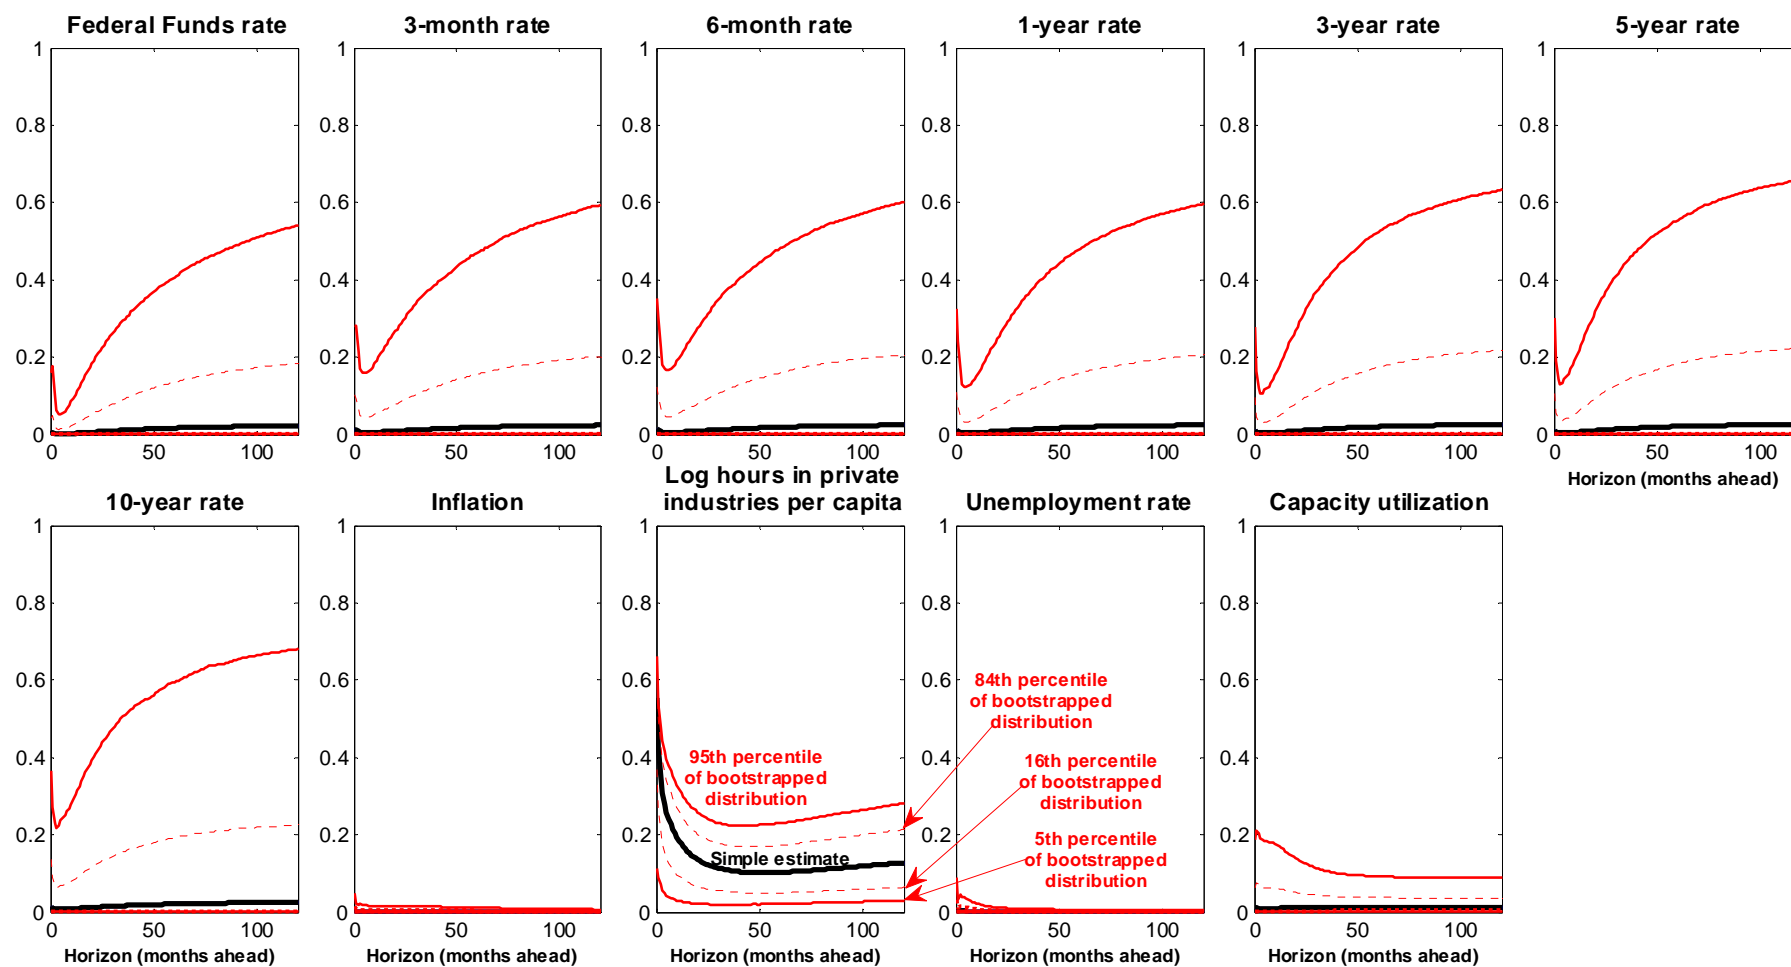


Figure A.2 Classical estimates based on the benchmark model: Fractions of forecast error variance explained by shocks to the natural rate of interest