Supplementary Material: Internet Appendices

for

Are Shadow Rate Models of the Yield Curve Structurally Stable?

IA.I. Innovation Statistics

This appendix proves invariance to model rotation and asymptotic distribution of the innovation-based statistics discussed in Section III.C and Section IV.A.2.

Consider a general first-order Gaussian VAR. That is, suppose the N-dimensional state vector x_t follows the transition equation $x_t = m_0 + M_1 x_{t-1} + \varepsilon_t$, where the ε_t are i.i.d. $N(0, \Omega)$, with $\Omega > 0$. Choose any matrix decomposition such that $SS' = \Omega$ (for example, the Cholesky decomposition) and let $\eta_t = S^{-1}\varepsilon_t = S^{-1}(x_t - m_0 - M_1 x_{t-1})$. Then by construction, the η_t are i.i.d. $N(0, I)$. Next, take an arbitrary invariant transformation, that is, an N-vector l and invertible $N \times N$-matrix L, and let $\tilde{x}_t = l + Lx_t$. Then $\tilde{x}_t = \tilde{m}_0 + \tilde{M}_1 \tilde{x}_{t-1} + \tilde{\varepsilon}_t$, where the $\tilde{\varepsilon}_t$ are i.i.d. $N(0, \tilde{\Omega})$, with $\tilde{m}_0 = (I - LM_1L^{-1})l + Lm_0$, $\tilde{M}_1 = LM_1L^{-1}$, and $\tilde{\Omega} = L\Omega L'$. If we choose any $\tilde{S}$ such that $\tilde{S}\tilde{S}' = \tilde{\Omega}$ and let $\tilde{\eta}_t = \tilde{S}^{-1}\tilde{\varepsilon}_t = \tilde{S}^{-1}(\tilde{x}_t - \tilde{m}_0 - \tilde{M}_1\tilde{x}_{t-1})$, then, denoting the Frobenius matrix norm by $||\cdot||_F$, for $u = 1, 2, \ldots$,

$$
\begin{aligned}
||T^{-1}\textstyle\sum_t \tilde{\eta}_t \tilde{\eta}'_{t-u}||_F^2 &= ||T^{-1}\textstyle\sum_t \tilde{S}^{-1}\tilde{\varepsilon}_t \tilde{\varepsilon}'_{t-u}\tilde{S}^{-1'}||_F^2 \\
&= T^{-2}\text{tr}(\tilde{S}^{-1}\textstyle\sum_t \tilde{\varepsilon}_{t-u}\tilde{\varepsilon}'_t \tilde{S}^{-'}\tilde{S}^{-1}\textstyle\sum_t \tilde{\varepsilon}_t\tilde{\varepsilon}'_{t-u}\tilde{S}^{-1'}) \\
&= T^{-2}\text{tr}(\tilde{S}^{-1}L\textstyle\sum_t \varepsilon_{t-u}\varepsilon'_t L'(\tilde{S}\tilde{S}')^{-1}L\textstyle\sum_t \varepsilon_t\varepsilon'_{t-u}L'(\tilde{S}^{-1'}) \\
&= T^{-2}\text{tr}(\Omega^{-1}\textstyle\sum_t \varepsilon_{t-u}\varepsilon'_t\Omega^{-1}\textstyle\sum_t \varepsilon_t\varepsilon'_{t-u})
\end{aligned}
$$

$$= T^{-2}\text{tr}(S^{-1}\textstyle\sum_t \varepsilon_{t-u}\varepsilon_t' S^{-1'} S^{-1}\sum_t \varepsilon_t \varepsilon_{t-u}' S^{-1'})$$

$$= ||T^{-1}\textstyle\sum_t \eta_t \eta_{t-u}'||_F^2.$$

The statistic $||T^{-1}\sum_t \eta_t\eta_{t-u}'||_F^2$ is thus invariant both to the original choice of S (say, whether the Cholesky decomposition or a symmetric matrix square root is used), and to observationally equivalent model rotations. By an analogous argument, the same is true for $||T^{-1}\sum_t \eta_t\eta_t' - I||_F^2$.

Moreover, since the elements of η_t are independent both contemporaneously and across time, $E(\eta_{t,i}\eta_{t-u,j})$ is equal to 1 if $u = 0$ and $i = j$, and is equal to 0 otherwise. Similarly, $Cov(\eta_{t,i}\eta_{t-u,j}, \eta_{t,p}\eta_{t-v,q})$ is 2 if $u = v = 0$ and $i = j = p = q$; is equal to 1 if $u = v = 0$ and either $i = p \neq j = q$ or $i = q \neq j = p$, or if $u = v > 0$ and either $i = p, j = q$ or $i = q, j = p$; and, is equal to 0 otherwise. Thus, by the Law of Large Numbers for random vectors,

$$\sqrt{T}\begin{pmatrix} \text{vech}\left[(T^{-1}\sum_t \eta_t\eta_t' - I) \otimes (I + J)^{\circ -1/2}\right] \\ \text{vec}\left(T^{-1}\sum_t \eta_t\eta_{t-1}'\right) \\ \vdots \\ \text{vec}\left(T^{-1}\sum_t \eta_t\eta_{t-u}'\right) \end{pmatrix} \xrightarrow{d} N(0, I)$$

as $T \to \infty$. Note $\eta_t\eta_{t-u}'$ is symmetric only when $u = 0$, which is why we use half-vectorization in that case; furthermore, the diagonal elements in that matrix have variance 2, which is accounted for by the Hadamard operation $\otimes(I + J)^{\circ -1/2}$ which divides the diagonal elements by $\sqrt{2}$. Using this

convergence result, the Continuous Mapping Theorem implies that

$$\frac{1}{2}T||T^{-1}\textstyle\sum_t \eta_t\eta_t' - I||_F^2 \xrightarrow{d} \chi^2_{(N^2+N)/2}$$

$$T||T^{-1}\textstyle\sum_t \eta_t\eta_{t-u}'||_F^2 \xrightarrow{d} \chi^2_{N^2} \qquad u = 1, 2, \ldots$$

as $T \to \infty$. Furthermore, as the individual statistics are independent under their asymptotic distribution, they can be added to obtain a joint statistic with asymptotic χ^2 distribution and degrees of freedom equal to the sum of the individual degrees of freedom.

IA.II. Parameter Estimates

This appendix provides the parameter estimates, in original normalization (equation equation (13)), based on the pre-ELB sample ($\hat{\theta}_{pre}$), post-ELB sample ($\hat{\theta}_{post}$), and full sample ($\hat{\theta}_{full}$). The PC-rotated parameter estimates $\hat{\theta}^{\dagger}_{pre}$ and $\hat{\theta}^{\dagger}_{post}$ in Section V.B are invariant transformations of $\hat{\theta}_{pre}$ and $\hat{\theta}_{post}$, respectively. Recall that, to achieve econometric identification, we impose the normalization restrictions $k_0^{\mathbb{P}} = 0_{3\times1}$, $\Sigma = 0.01I$, $[K_1^{\mathbb{P}}]_{i,j(i<j)} = 0$.

A few properties of our estimates are of note: In $\hat{\theta}_{pre}$, the estimate of $[K_1^{\mathbb{P}}]_{33}$ does not have a standard error because the estimate is at the boundary ($[K_1^{\mathbb{P}}]_{33} = [K_1^{\mathbb{P}}]_{22}$); that is, the matrix has a repeated eigenvalue. In $\hat{\theta}_{post}$, the estimate of δ_y does not have a standard error as it is at the minimum bound (4 bps) that we imposed. If this parameter is left unconstrained in the post-ELB sample estimation, the estimated $\mathbb{Q}$ parameters display strong signs of overfitting. In all three models, we impose a lower bound of 50 bps on the measurement error estimate for the 5-to-10-year survey expectation.

	$\hat{\theta}_{pre}$	$\hat{\theta}_{post}$	$\hat{\theta}_{full}$
ρ_0	0.0450 (0.0023)	0.0374 (0.0097)	0.0465 (0.0028)
$[\rho_1]_1$	0.4051 (0.1051)	0.2541 (0.0824)	0.3512 (0.0315)
$[\rho_1]_2$	-0.1366 (0.0803)	0.1596 (0.0600)	0.0449 (0.0485)
$[\rho_1]_3$	0.6915 (0.0711)	0.4295 (0.0615)	0.7472 (0.0533)
$[K_1^{\mathbb{P}}]_{11}$	-0.0637 (0.0192)	-0.1109 (0.4311)	-0.0917 (0.0167)
$[K_1^{\mathbb{P}}]_{21}$	0.0584 (0.1125)	0.1390 (0.3667)	0.1235 (0.0603)
$[K_1^{\mathbb{P}}]_{31}$	-0.0744 (0.2951)	-0.2366 (0.7887)	-0.0151 (0.0582)
$[K_1^{\mathbb{P}}]_{22}$	-1.1215 (0.0538)	-0.6351 (0.0956)	-0.8391 (0.2096)
$[K_1^{\mathbb{P}}]_{32}$	2.5296 (0.3192)	2.8466 (0.9364)	1.8229 (0.2069)
$[K_1^{\mathbb{P}}]_{33}$	-1.1215	-0.6784 (0.0576)	-0.9855 (0.0375)
$[k_0^{\mathbb{Q}}]_1$	0.0073 (0.0015)	-0.0069 (0.0098)	0.0049 (0.0002)
$[k_0^{\mathbb{Q}}]_2$	-0.0068 (0.0016)	0.0063 (0.0086)	-0.0022 (0.0014)
$[k_0^{\mathbb{Q}}]_3$	0.0048 (0.0038)	0.0124 (0.0127)	0.0045 (0.0012)
$[K_1^{\mathbb{Q}}]_{11}$	0.1194 (0.1783)	-0.3428 (0.1242)	0.0138 (0.0066)
$[K_1^{\mathbb{Q}}]_{21}$	0.1212 (0.1129)	0.3423 (0.1345)	0.1212 (0.0089)
$[K_1^{\mathbb{Q}}]_{31}$	0.3998 (0.3284)	0.0872 (0.4703)	0.1682 (0.0074)
$[K_1^{\mathbb{Q}}]_{12}$	1.4357 (0.2536)	-0.5689 (0.3800)	0.6421 (0.0099)
$[K_1^{\mathbb{Q}}]_{22}$	-0.9906 (0.1593)	0.3044 (0.2306)	-0.4941 (0.0166)
$[K_1^{\mathbb{Q}}]_{32}$	2.5270 (0.3326)	2.1732 (0.6383)	1.7591 (0.0097)
$[K_1^{\mathbb{Q}}]_{13}$	-0.6656 (0.0794)	0.0465 (0.0371)	-0.2775 (0.0072)
$[K_1^{\mathbb{Q}}]_{23}$	0.1049 (0.0692)	-0.0081 (0.0272)	-0.0464 (0.0081)
$[K_1^{\mathbb{Q}}]_{33}$	-1.3073 (0.1178)	-0.4557 (0.0813)	-0.8964 (0.0097)
δ_y	0.0006 (0.0000)	0.0004	0.0006 (0.0000)
$\delta_{z,6m}$	0.0017 (0.0001)	0.0011 (0.0001)	0.0018 (0.0001)
$\delta_{z,12m}$	0.0023 (0.0001)	0.0017 (0.0001)	0.0025 (0.0007)
$\delta_{z,5-10y}$	0.0050	0.0050	0.0050
$\underline{r}$		0.0006 (0.0000)	0.0007 (0.0000)

Table IA.1

Parameter estimates for the three models used in the paper.

IA.III. Yields-Only Model

As a robustness check, we estimate a full-sample, single-regime shadow rate model based only on yield data, omitting surveys. We find, as expected, that the parameter estimates differ somewhat from those reported for the analogous model with surveys in Table IA.1. Particularly notable are a lower estimate of the long-term short rate ρ_0, at 0.0381, albeit with substantially larger standard error (0.0133), as well as a more negative eigenvalue $[K_1^{\mathbb{P}}]_{11}$, estimated to be -0.1755, implying a faster rate of convergence to the long-term mean, though equally with a relatively large standard error (0.1175). Both observations are consistent with the issues raised by Kim and Orphanides (2012) which motivate the use of survey data in the first place.

Based on this yields-only model, we repeat the LM tests from Section IV.B. As seen in Figure IA.1, the time series of statistics differs minimally from that based on the model with surveys (consistent with the small difference seen in the Ghysell-Hall statistics with and without surveys reported in Table 5). Note, however, that the close match between the LM statistics does *not* imply that the two models must have nearly identical parameter estimates, nor that survey data by itself does not reflect a structural break; it merely suggests that survey data contains little *incremental* information about a structural break, over and above the yield data. We conclude from this exercise that the evidence in favor of the *existence* of a structural break based on single-regime models (which is derived from information in the neighborhood of a single estimated parameter vector) is not meaningfully affected by the presence of survey data. Nevertheless, the addition of survey information may help characterize the *nature* of the structural break (i.e., identify the different estimated pre- and post-break parameter vectors); indeed, we have not been able to

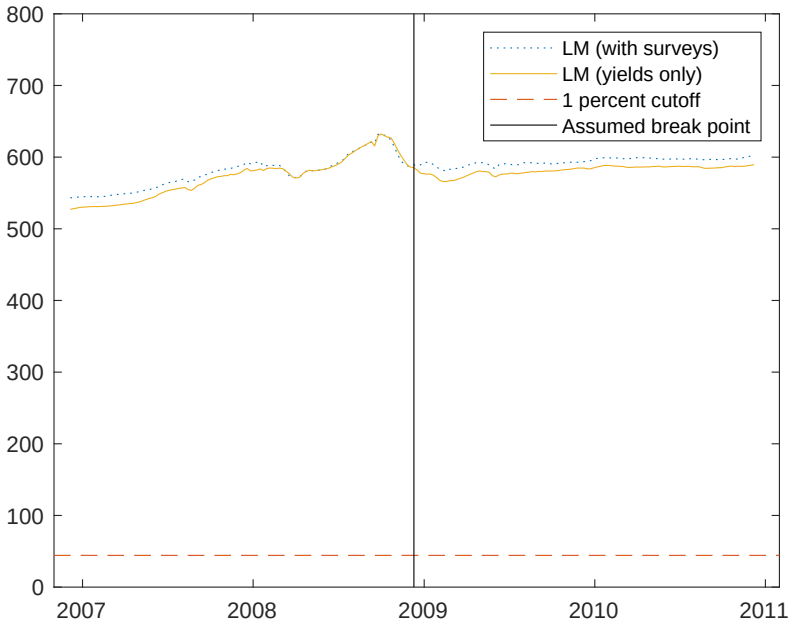

Figure IA.1

Time series of Lagrange multiplier statistics for the yields-only model (solid yellow line), with the dotted blue line replicating the statistics based on our model with surveys from Figure 3.

estimate a structurally broken model without survey data analogous to Section V due to the limited length of the post-ELB sample.[38]

IA.IV. $\mathbb{Q}$-Stable Model

In this appendix, we present results from a model in which the $\mathbb{P}$ parameters are permitted to change at the end of 2008 while the $\mathbb{Q}$ parameters remain stable. Such a model is motivated by

[38] With yields-only estimation of the post-ELB period data, we were not able to arrive at a global maximum of the likelihood function with any reasonable degree of confidence. The estimate with the highest likelihood value among the estimates we found had unreasonable properties; e.g., its implied time series of the 5-to-10-year-ahead expected shadow rate had negative values throughout the post-ELB sample period.

the desire to specify the structural break more parsimoniously than we do in Section V, where we allow all model parameters to change at the time of the structural break. While we are thus able to analyze the characteristics of the structural break in the greatest possible generality, this flexibility comes at the cost of doubling the number of free parameters.

Here, we follow an approach inspired by studies such as Andreasen et al. (2019) and Giacoletti et al. (2021), who argue that structural changes in the economic environment can be captured by allowing $\mathbb{P}$ parameters to vary while keeping $\mathbb{Q}$ parameters fixed. In particular, we estimate a model in which we normalize $\rho_1 = (1, \ldots, 1)'$ and $k_0^{\mathbb{Q}} = 0$, and keep ρ_0, the eigenvalues of $K_0^{\mathbb{Q}}$, as well as the instantaneous covariance matrix Σ fixed throughout the sample. Meanwhile, $k_0^{\mathbb{P}}$ and $K_1^{\mathbb{P}}$ are allowed to change in 2008.

As shown in Figure IA.2, this model fits the cross section of yields comparably as well as the fully flexible structurally broken model, with no readily discernible deterioration of fit in the post-ELB period. At the same time, as shown in Figure IA.3, the shadow rate implied by the model with stable $\mathbb{Q}$ parameters is more closely aligned with that implied by the structurally stable model; this is particularly notable in 2014 and 2015. Thus, while the model with stable $\mathbb{Q}$ parameters seems to fit the cross section of yields better than the structurally stable model in the post-ELB period, it does so at inferred state variables that are closer to the structurally stable model than the fully flexible structurally broken model. As seen in the left panel of Figure IA.4, the two-year yield term premium implied by the model with stable $\mathbb{Q}$ parameters periodically deviates from the term premium implied by the fully flexible structurally broken model.

Analogously to Section IV.B, we conduct an LM test for structural stability, here with respect to only the $\mathbb{Q}$ parameters. The statistic is 287.25, about half the value we obtain in the model in which we test for a break in all parameters (see Figure 3), but still very highly significant

as the test here has fewer degrees of freedom; the 1% cutoff for 10 d.f. is 23.21. In other words, from a purely statistical perspective, the null hypothesis of structural stability in the $\mathbb{Q}$ parameters is strongly rejected in a model in which the $\mathbb{P}$ parameters are allowed to change. Interestingly, this is the case even though such a model appears to be able to fit the cross section of yields adequately (as seen in Figure IA.2). Two related factors likely contribute to this finding: First, even modest improvements in yield fit tend to translate into large gains in likelihood. Second, such improvements in cross-sectional fit might be achieved by compromising the time-series properties of the filtered state variables. Indeed, recall from Figure IA.3 that the implied shadow rate of the model with stable $\mathbb{Q}$ parameters—a linear function of its filtered state variables—tracks more closely the shadow rate implied by the structurally stable model than that implied by the fully flexible structurally broken model.

IA.V. Additional Materials for Section V

V.A. Decomposition of Shadow Yields

Here we further examine the shadow yields, $y^s_{t,\tau}$, defined as

$$y^s_{t,\tau} = -\frac{1}{\tau}\log P^s_{t,\tau}, \quad P^s_{t,\tau} = E^{\mathbb{Q}}_t\left[\exp\left(-\int_t^{t+\tau} s_u du\right)\right],$$

with the shadow rate s_t. As in the decomposition of (actual) yield into expectations component and term premium component, we can decompose shadow yield into "shadow yield expectations"

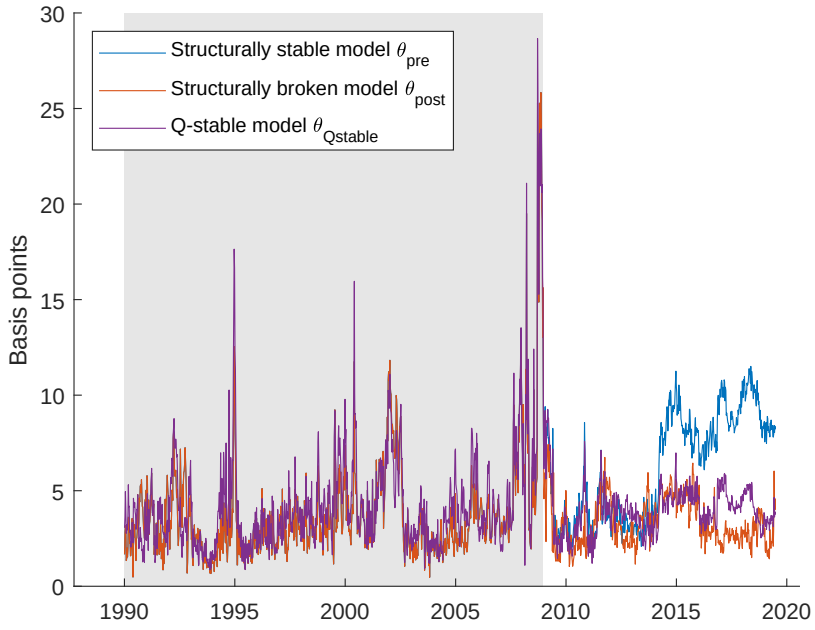

Figure IA.2

Time series of mean-root-squared yield fitting errors in different models. The pre-ELB period is shaded.

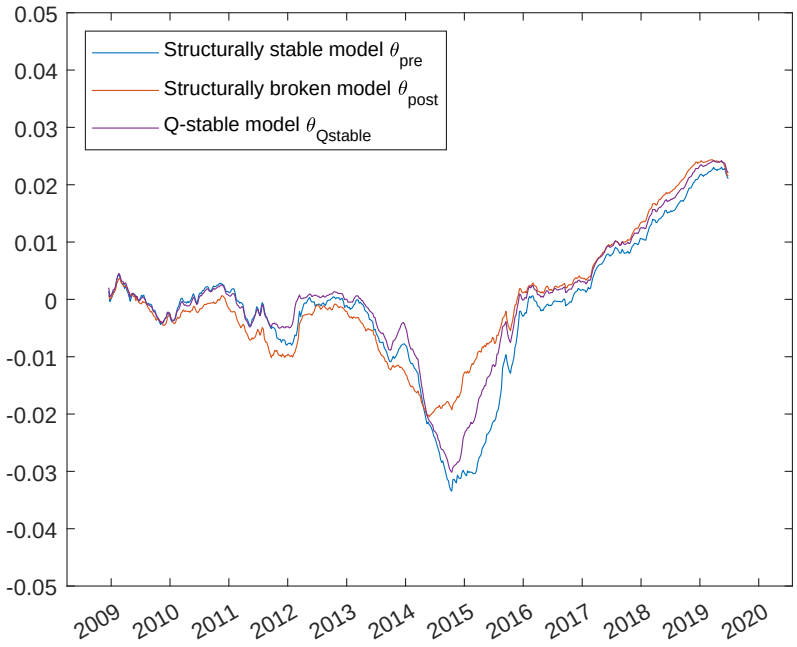

Figure IA.3

Time series of estimated shadow rates.

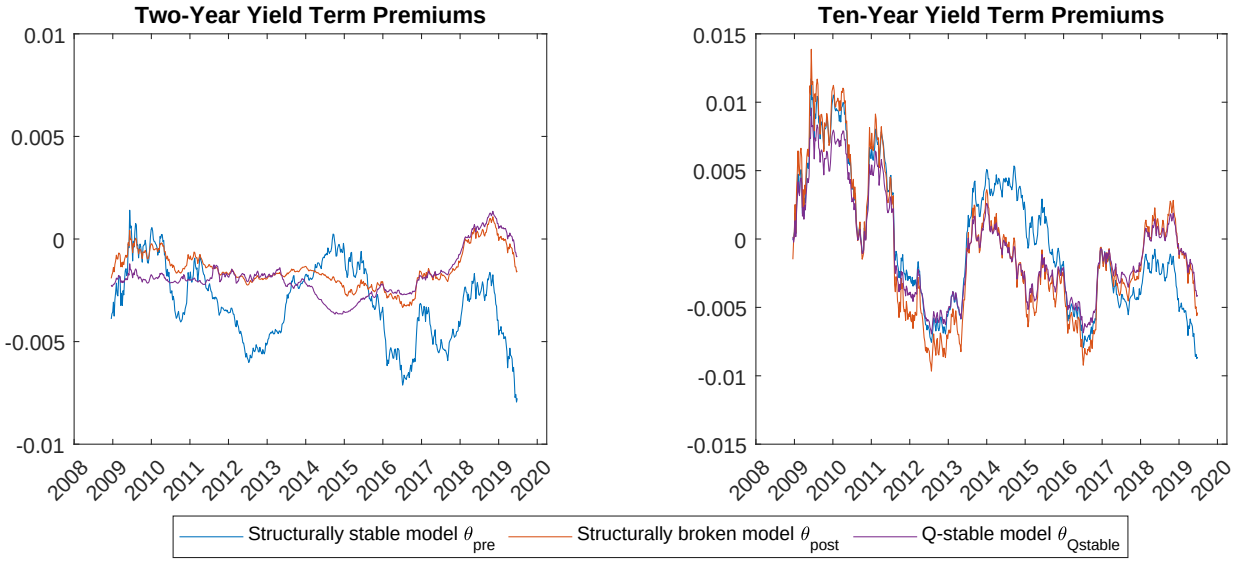

Figure IA.4

Two- (left panel) and ten-year (right panel) yield term premiums implied by our different models, shown for the post-ELB subsample.

component YE^s and "shadow yield term premium" component YTP^s, i.e.,

$$y^s_{t,\tau} \quad = \quad \mathrm{YE}^s_{t,\tau} + \mathrm{YTP}^s_{t,\tau}.$$

Note that y^s, YE^s, and YTP^s are all affine in the state vector x_t.

Figure IA.5 displays the time series of the shadow yields, shadow expectations components, and shadow term premiums for the 2-year maturity, implied by the structurally broken model and the structurally stable model. It can be seen that the 2-year shadow yields implied by the two models (structurally stable vs. broken) are fairly similar, but there is considerable difference in the behavior of 2-year shadow term premiums. In particular, the shadow term premium in the structurally stable model is more volatile and lower in level than the shadow term premium in the structurally broken model during much of the post-ELB sample period; this seems to be consistent with the feature in Figure 5 of the paper that the actual (i.e. not shadow) term premium in the 2-year yield is more negative and volatile in the structurally stable model than in the structurally broken model, and also with the findings in Section V.C of the paper that the shadow yields in the structurally broken model are close to the expectations hypothesis for short maturities in the post-ELB sample. In the case of longer maturities such as the 10-year horizon shown in Figure IA.6, the structurally broken model implies a greater gap between actual and shadow yields, relative to the structurally stable model. Our interpretation is that after the Financial Crisis (i.e., in the post-ELB period), permanently higher risk-adjusted odds of the scenario of returning to ELB are priced into yields, and this is manifesting itself as a greater gap between actual yields and shadow yields (lower than actual yields) in the structurally broken model. The decomposition of 10-year shadow yields into shadow EH and shadow term premium components indicate that the shadow

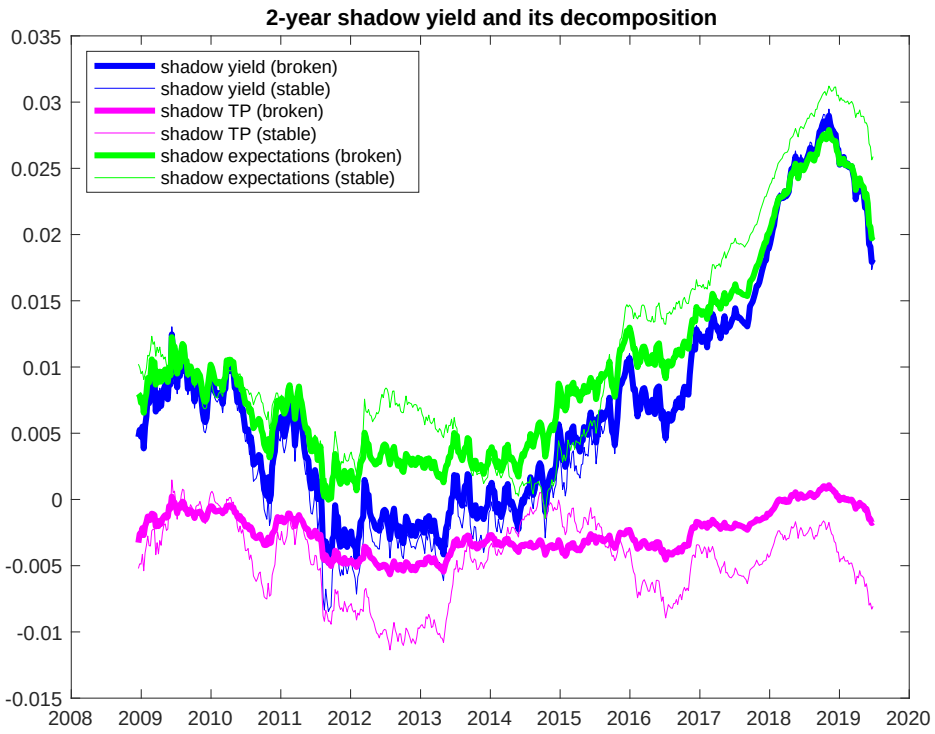

Figure IA.5

Decomposition of the 2-year shadow yields into shadow expectations components and shadow term premium components, based on structurally stable and broken models.

term premium in the structurally broken model tends to be lower than that in the structurally stable model in the post-ELB period, significantly contributing to the greater gap between shadow yields and actual yields in the structurally broken model.

V.B. Model-implied Shadow Rate Paths

Here we examine the time series of $\mathbb{P}$-measure expectations of the shadow rate implied by the structurally broken and structurally stable models. Figure IA.7 displays the time series of the shadow rate itself and the shadow rate expected at horizons of 1 year, 2 years, and 3 years, for both models. Interestingly, while the level of the shadow rate was much lower for the structurally stable

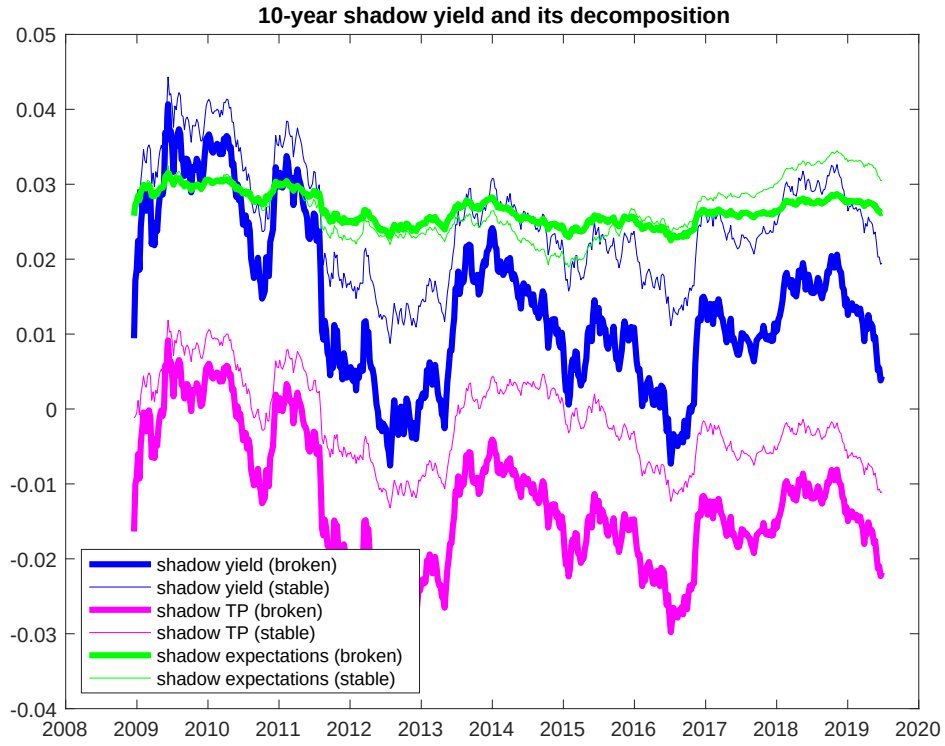

Figure IA.6

Decomposition of the 10-year shadow yields into shadow expectations components and shadow term premium components, based on structurally stable and broken models.

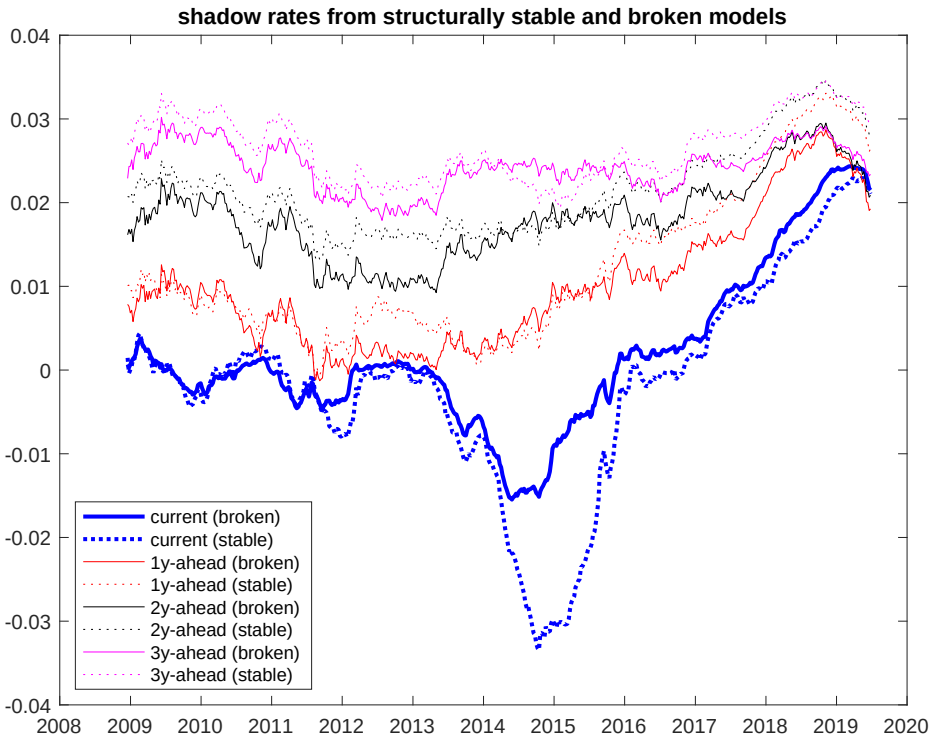

Figure IA.7

Times series of shadow rates and their $\mathbb{P}$-measure expectations 1-, 2-, and 3-years ahead, based on the structurally stable model and structurally broken model.

model during 2014–15, the *expected* shadow rate for horizons of 1 year and beyond are roughly similar during that period (i.e., the expected path of the shadow rate for horizons below 1 year was much steeper for the structurally stable model during 2014–15), highlighting that a lower level of shadow rate does not necessarily imply a lower path for the shadow rate beyond very near-term horizons. It can be also seen that for much of the post-ELB (2008–19) period, the 1∼2-year horizon shadow rate expectations based on the structurally stable model are persistently higher than those of the structurally broken model (especially during periods of 2010–2012 and 2016–2018). This is also consistent with the persistently more negative 2-year shadow yield term premium seen in Figure IA.5.

V.C. Economic Variables as Functions of PC Factors

We also examine model-implied economic variables of interest from the perspective of PC components, to discern how their roles might have changed between the structurally stable and the structurally broken models. Figure IA.8 shows the filtered estimates of these PC factors for the structurally stable model and the structurally broken model. While all three PC factors show some differences between the structurally stable and broken models, the question of how much a specific PC factor contributes to the difference between the structurally stable and broken models will depend on the specific object under consideration. For example, the "shadow yields" (of various maturities) can be written as affine functions of PC factors (i.e., $const. + \sum_{i=1}^{N} b_i \mathrm{PC}_{it}$), and therefore the impact of a structural break can show up as either changes in the PC factors (PC_i), in the factor loadings (b_i), or both. Therefore, to compare a PC factor's overall contribution to an object, we examine the product of the factor with its factor loading for the object ($b_i \mathrm{PC}_i$). In Figure IA.9, we plot this product for each of the PC factors (PC1, PC2, PC3), for the 2-year shadow yield.[39] This decomposition indicates that the slope factor (PC2) has a notably different role in the variation of the 2-year shadow yield, with PC2 contributing little in the case of the structurally stable model, and PC2 contributing substantially in the case of the structurally broken model. Similarly, we can examine the contribution of each of the PC factors to the expectations component of the 2-year shadow yield ($\mathrm{YE}^{s}_{t,2y}$); this is shown in Figure IA.10. It can be seen that there are visible differences in the contribution to the 2-year shadow expectations component between

[39]Here we focus on the 2-year maturity, as the results in Section V.D indicated important qualitative differences between structurally stable and broken models for relatively short-term maturities such as 2 years. Note also that we focus on shadow yields instead of actual yields, because the actual yields are nonlinear in the PC factors (or state variables more generally) and hence it is ambiguous how to define the contribution of each factor to actual yields.

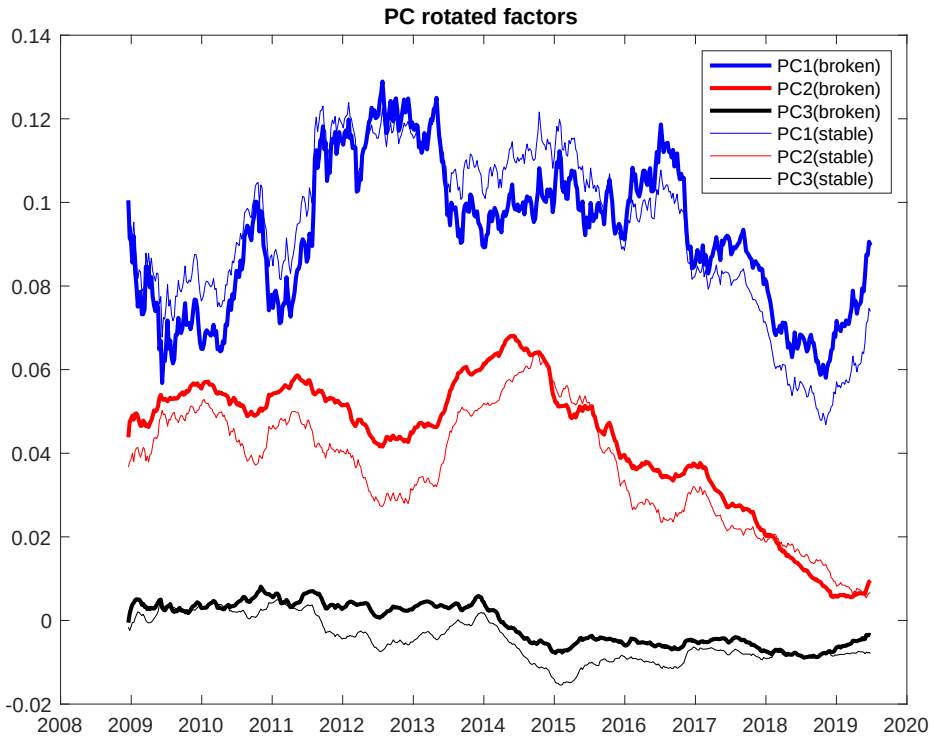

Figure IA.8

Time series of PC factors (level, slope, curvature) based on structurally stable and broken models.

structurally stable and structurally broken models for all three PC factors. Moreover, the differences are not confined to a relatively narrow period, say, for example, 2013–15.

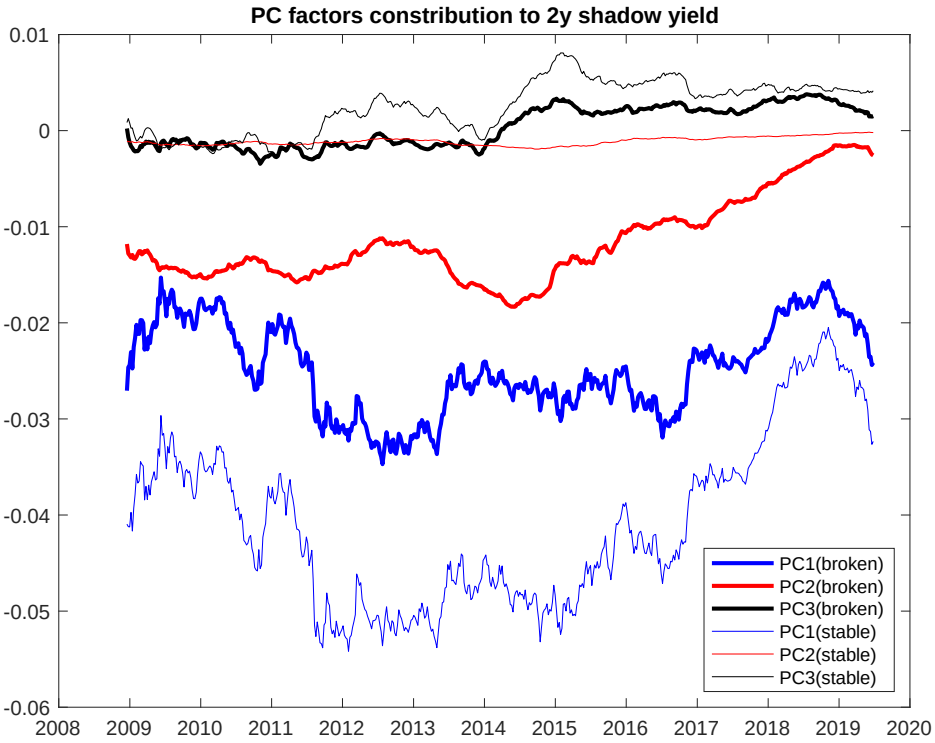

Figure IA.9

PC factors (PC1, PC2, PC3)'s contribution to the 2-year shadow yield ($y^s_{t,2y}$), based on structurally stable and broken models.

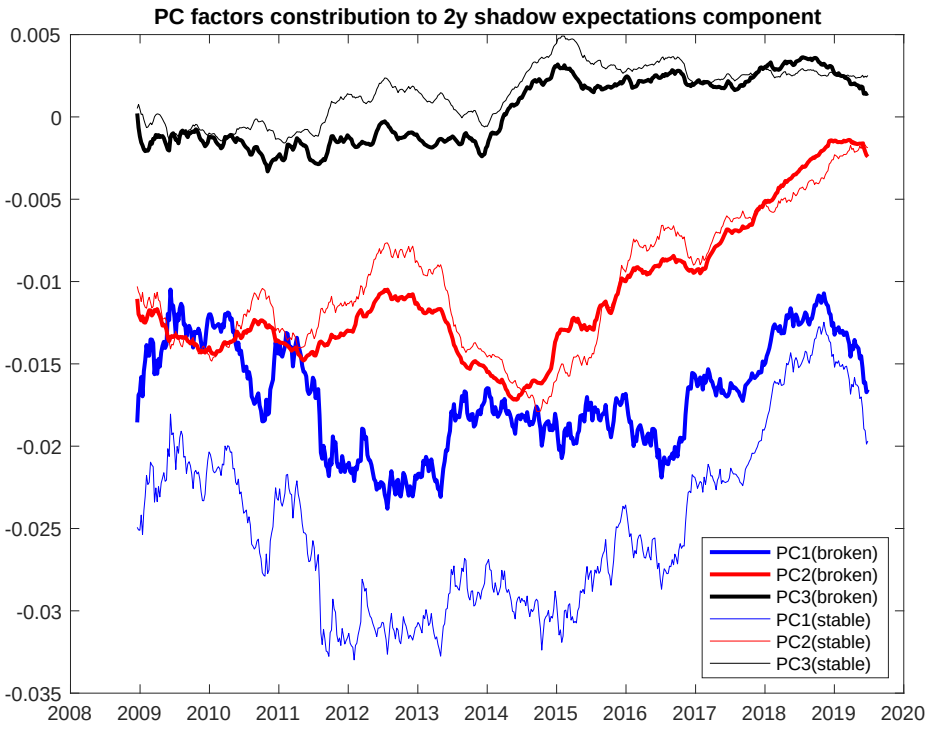

Figure IA.10

PC factors (PC1, PC2, PC3)'s contribution to the 2-year shadow yield expectations component ($\mathrm{YE}^{s}_{t,2y}$).