One-factor Hull-White Model Calibration for CVA Part II: Optimizing the Mean Reversion Parameter

Christoph M. Puetter*and Stefano Renzitti†

IHS Markit, 1066 West Hastings Street, Vancouver, BC V6E 3X1, Canada

July 9, 2020

This paper is the second of a multi-part series on the calibration of the one-factor Hull-White short rate model for the purpose of computing CVAs (and xVAs) with an xVA system. The first part introduces an atypical bootstrapping scheme for the calibration of the short rate volatility. The present second part focuses on the selection of the mean reversion parameter. In both expositions we present long-term time series results for EUR, JPY, and USD, covering the period from the beginning of 2009 (at the earliest) to spring 2020.

Valuation adjustment (xVA) systems build on a framework of inter-dependent models to facilitate pricing across asset classes and currencies.[1, 2] Simpler and more tractable models are typically used to keep the risk factor dimensionality and computational load at bay. At the heart of any xVA modeling framework are interest rates, whose dynamics, for the purpose of xVA calculations, are frequently derived from the one-factor Hull-White (HW 1F) short rate model.[3] The HW 1F model, despite its basic nature, is quite versatile when deployed accordingly. For example, in the first part of the paper series we showed that the HW 1F model is capable of simultaneously producing meaningful xVAs for a number of interest rate swaps that differ by maturity, provided that suitable swaptions – swaptions arranged in a chevron pattern that partially track the swaps' exposure peaks – are used for the volatility calibration. The calculations, without having gone into detail there, however, also took advantage of a fine-tuning procedure aimed at the mean reversion parameter.[4, 5] The mean reversion parameter is a crucial additional model degree of freedom that can significantly affect exposure simulations and xVA estimates. A deliberate selection therefore is advantageous and warrants a closer look in this second part of the series.

One of the goals of interest rate model calibration is to capture as much of a market-implied volatility surface as possible. Since the HW 1F model's volatility term structure is typically bootstrapped from at-the-money (ATM) swaptions, using one instrument per expiry, most of the ATM swaption volatility surface is merely implied by the model, but can be fitted to some extent, in a least square sense or else, by adjusting the mean reversion parameter. The mean reversion parameter's impact is exemplified in Figs. 1 and 2, where different choices for otherwise identical model and volatility calibration setups lead to distinctly different model-implied volatility surfaces. The volatility surfaces meet at the calibrating swaptions (white dots) but their orientations vary: a positive mean reversion parameter value tends to give the model-implied volatilities a negative slope as a function of swaption length and/or expiry and vice versa for a negative value.

A negative mean reversion parameter might seem an odd element in a mean reverting short rate model. But it is a pragmatic choice that helps modeling the dynamics of more involved

*christoph.puetter@ihsmarkit.com
†stefano.renzitti@ihsmarkit.com

quantities such as the swap rate. Past and recent developments of the major interest rate markets (EUR, JPY, and USD), which are dominated by a trend towards small and even negative bank rates and yields, all suggest that the use of negative mean reversion values is indispensable. A caveat is that a negative mean reversion parameter requires particular consideration when extrapolating the HW 1F model beyond the calibration regime (e.g., past the 30-year expiry mark in Figs. 1 and 2), which we discourage. The reason is that market-implied swaption volatilities tend to decay with swaption expiry and length; a HW 1F model with a negative mean reversion parameter on the other hand usually exhibits the opposite behavior and therefore can easily lead to overestimated exposures and xVAs for longer-dated trades and portfolios.

In trying to match the market-implied volatility surface by tuning the mean reversion parameter, one is free to choose the portion of the volatility matrix to put the main emphasis on. In an xVA system deployed to compute exposures and valuation adjustments for maturities of up to 30 years, a prudent, but not only, option is to minimize the root-mean-square relative error between model and market-implied volatilities of swaptions (Vol RMSRE) that approximately fall below the 30-year diagonal (or co-terminal).[1] Such a swaption fit metric can be defined and computed irrespective of the volatility calibration method. The resulting mean reversion parameter and swaption fit curves tend to be relatively stable over shorter time scales and are more variable over longer time horizons, see Figs. 3 and 4. The larger changes (of order $\sim 100\,\text{bp/yr}$ or more) often coincide with major market events such as central bank policy announcements. A case in point are the economic repercussions of the COVID-19 pandemic. Since March 2020, the market-implied swaption volatility surfaces of EUR, JPY, and USD have assumed very similar shapes. The optimized mean reversion parameter curve for USD has eventually entered negative territory, where it joins those for EUR and JPY. The mean reversion parameters for the three major currencies as of late spring / early summer 2020 roughly reside in the -100 to -200 bp/yr range.

The impact of the mean reversion parameter on the swaption fit is also clearly discernible from Figs. 3 and 4, with obvious ramifications for the accuracy of exposure and xVA calculations (see also first part of the series). The swaption fit errors vary by calibration method, currency, and date, and typically move in a range between 5 and 15% when the mean reversion parameter is purposely set. In the best cases, the swaption fit errors can be as small as three percentage points, see e.g., the USD curves at the end of 2019, but typically cannot match those derived from more sophisticated but related multi-factor Hull-White models (not shown).[3, 6] The graphs also suggest that a minor deviation of the mean reversion parameter from the optimal value by about 25-50 bp/yr is forgivable in terms. A large misspecification, on the other hand, as caused, e.g., by permanently using a fixed value of 500 bp/yr, can incur a large swaption fit penalty and lead to volatility calibration failures, and is better to be avoided.

From an implementation perspective, adding 25-50 bp/yr to the optimized value can help to configure a numerically more stable HW 1F model instance, without degrading exposures and xVA estimates too much. This can further be combined with a mapping to a discrete and piecewise constant mean reversion curve, effectively updating the mean reversion parameter with less than daily frequency, only when deemed necessary. Taking this route might be particularly beneficial when trying to reduce noise for sensitivity and profit-and-loss analyses.

Numerical instabilities are rare but occasionally do occur due to model limitations. Some of the spikes exhibited by the optimized mean reversion curves and marked by dots in Figs. 3 and 4 are artifacts thereof and in practice have to be replaced, e.g., by the previous day value.

[1] We correct for the density of the market-implied swaption volatilities in the expiry-length plane by using an (approximate) integral mean instead of a sample mean when computing the Vol RMSRE. We also constrain the mean reversion optimization procedure to find a value where the HW 1F model remains well defined and numerically stable, i.e. can price all market swaptions without failure. This can typically be achieved by choosing a sufficiently large (but not too large) mean reversion parameter. The optimization search range is furthermore limited to an interval between -500 and 1000 bp/yr.

Another artifact is that the mean-reversion-optimized swaption fit might appear sub-optimal in comparison to the fixed mean reversion alternatives as additional constraints are imposed on the optimized mean reversion parameter (see footnote 1 and, e.g., JPY curves in 2019 in Fig. 4).

Regardless of these details, calibrating the mean reversion parameter on a daily, weekly, monthly, or even less frequent basis is certainly beneficial. A more refined HW 1F modeling approach would involve an entire mean reversion term-structure instead of only a single parameter. For certain, one-sided portfolios, such an approach would likely be capable of providing exposures and xVAs on par with those from a multi-factor Hull-White model. A time-dependent mean reversion function allows for a highly-adaptable model-implied volatility surface, where each swaption strip in the length direction can effectively be fitted with an individual slope (see also [4]). Multi-tenor portfolios that are sensitive to decorrelation of interest rates, however, do require a multi-factor modeling approach. Details regarding the applicability and limits of multi-factor Hull-White short rate models for xVA calculations will be addressed in the future.

References

[1] Gregory, J. 2015. *The xVA Challenge: Counterparty Credit Risk, Funding, Collateral and Capital.* John Wiley & Sons.

[2] Green, A. 2016. *XVA: Credit, Funding and Capital Valuation Adjustments.* John Wiley & Sons.

[3] Andersen, L. B. G. and Piterbarg, V. V. 2010. *Interest Rate Modeling - Volume 2: Term Structure Models.* Atlantic Financial Press.

[4] Gurrieri, S., Nakabayashi, M. and Wong, T. 2009. Calibration Methods of Hull-White Model. Available at SSRN: `https://ssrn.com/abstract=1514192`

[5] Schlenkrich, S. 2012. Efficient Calibration of the Hull White Model, *Optim. Control. Appl. Meth.* 33, 352.

[6] Schrager, D. F. and Pelsser, A. A. J 2006. Pricing Swaptions and Coupon Bond Options in Affine Terms Structure Models. *Math. Finance* 16, 673-694.

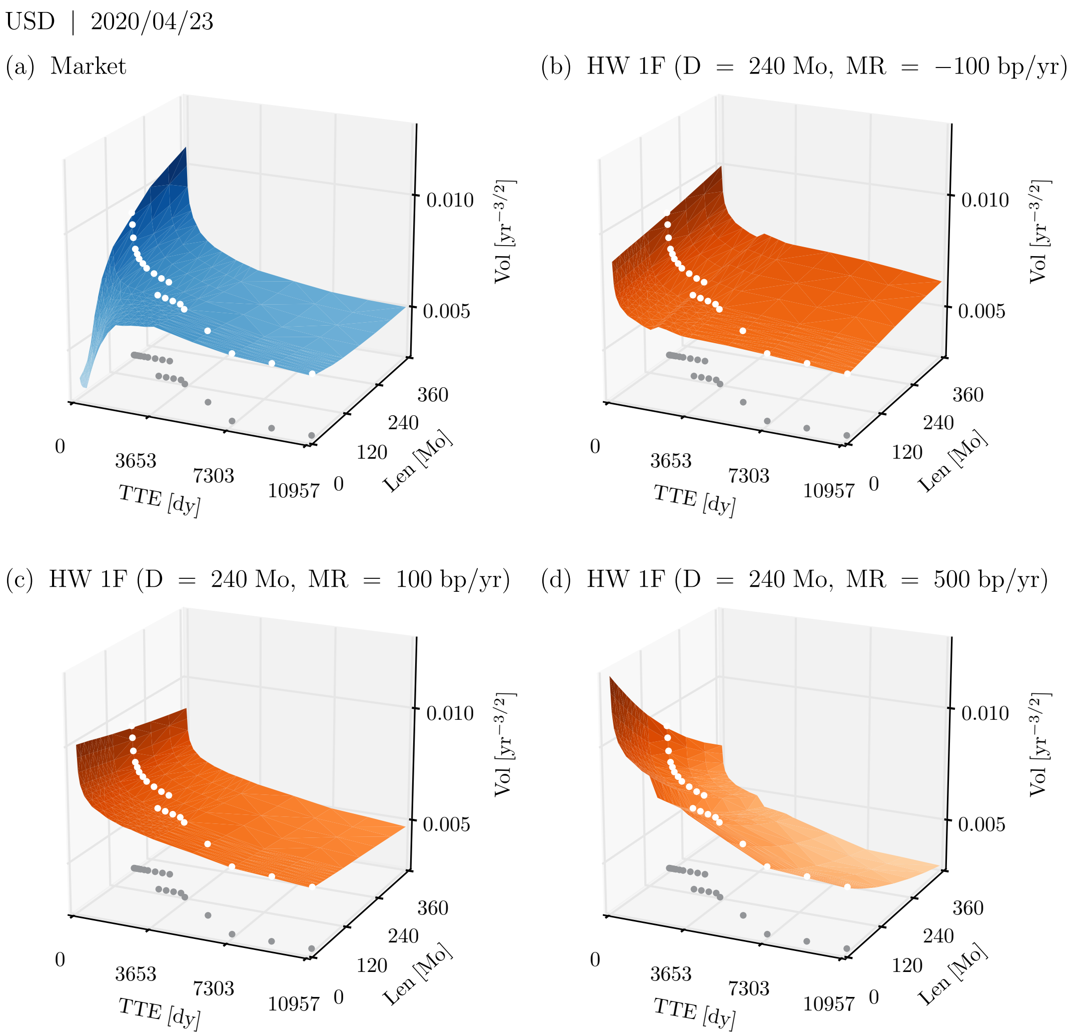

Figure 1: Market [(a)] and model-implied [(b)-(d)] normal ATM swaption volatility surfaces for USD as of 23 April 2020. The model-implied volatilities are based on calibration to the same 20-year diagonal (co-terminal) swaption set (white dots) but for different mean reversion parameters: $-100, 100,$ and 500 bp/yr. The mean reversion parameter choice clearly affects the fit to the market-implied volatility surface. A similar behavior transpires for other calibration options, see, e.g., Fig. 2, and the main text for more details.

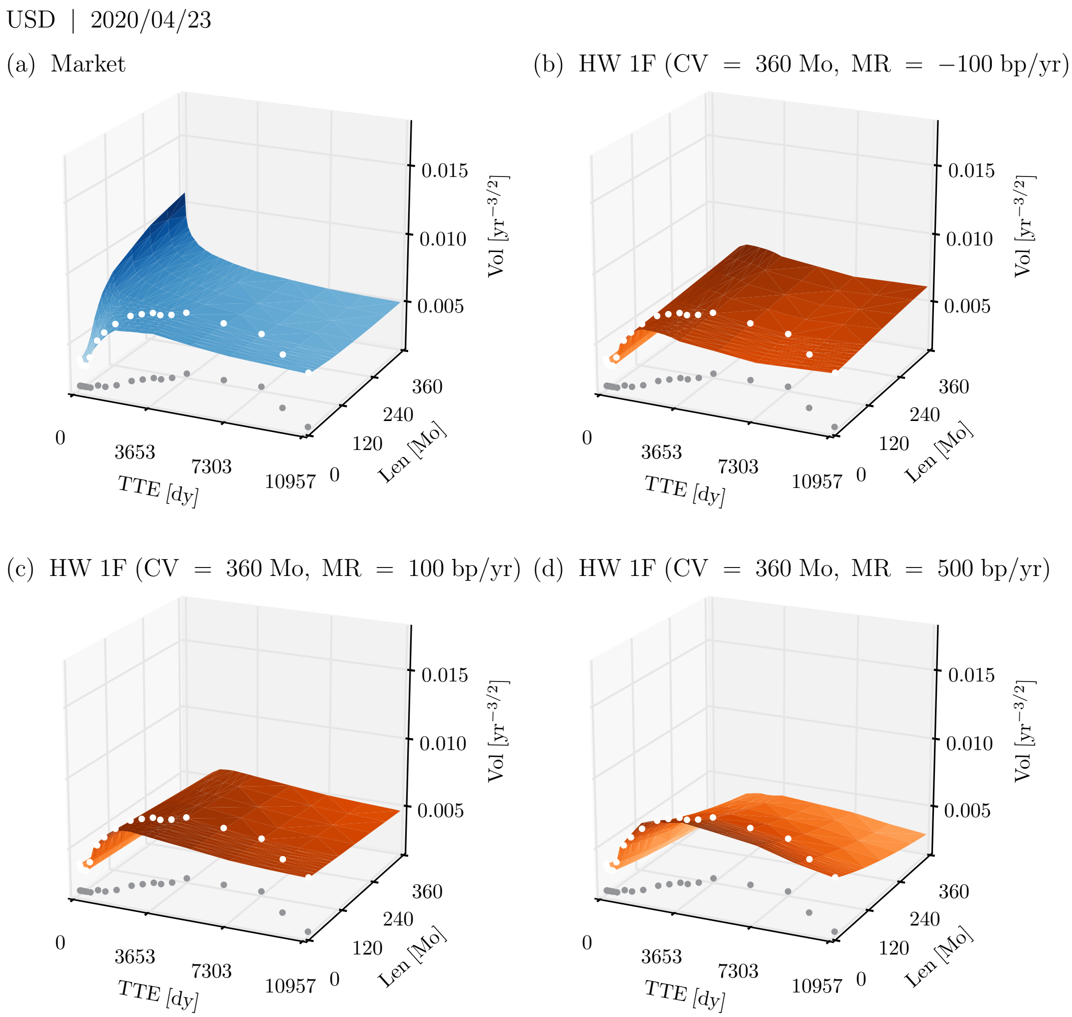

Figure 2: Market [(a)] and model-implied [(b)-(d)] normal ATM swaption volatility surfaces for USD as of 23 April 2020, similar to Fig. 1 but with HW 1F model instances calibrated to the 30-year swaption chevron (white dots). See caption of Fig. 1 and main text for more details.

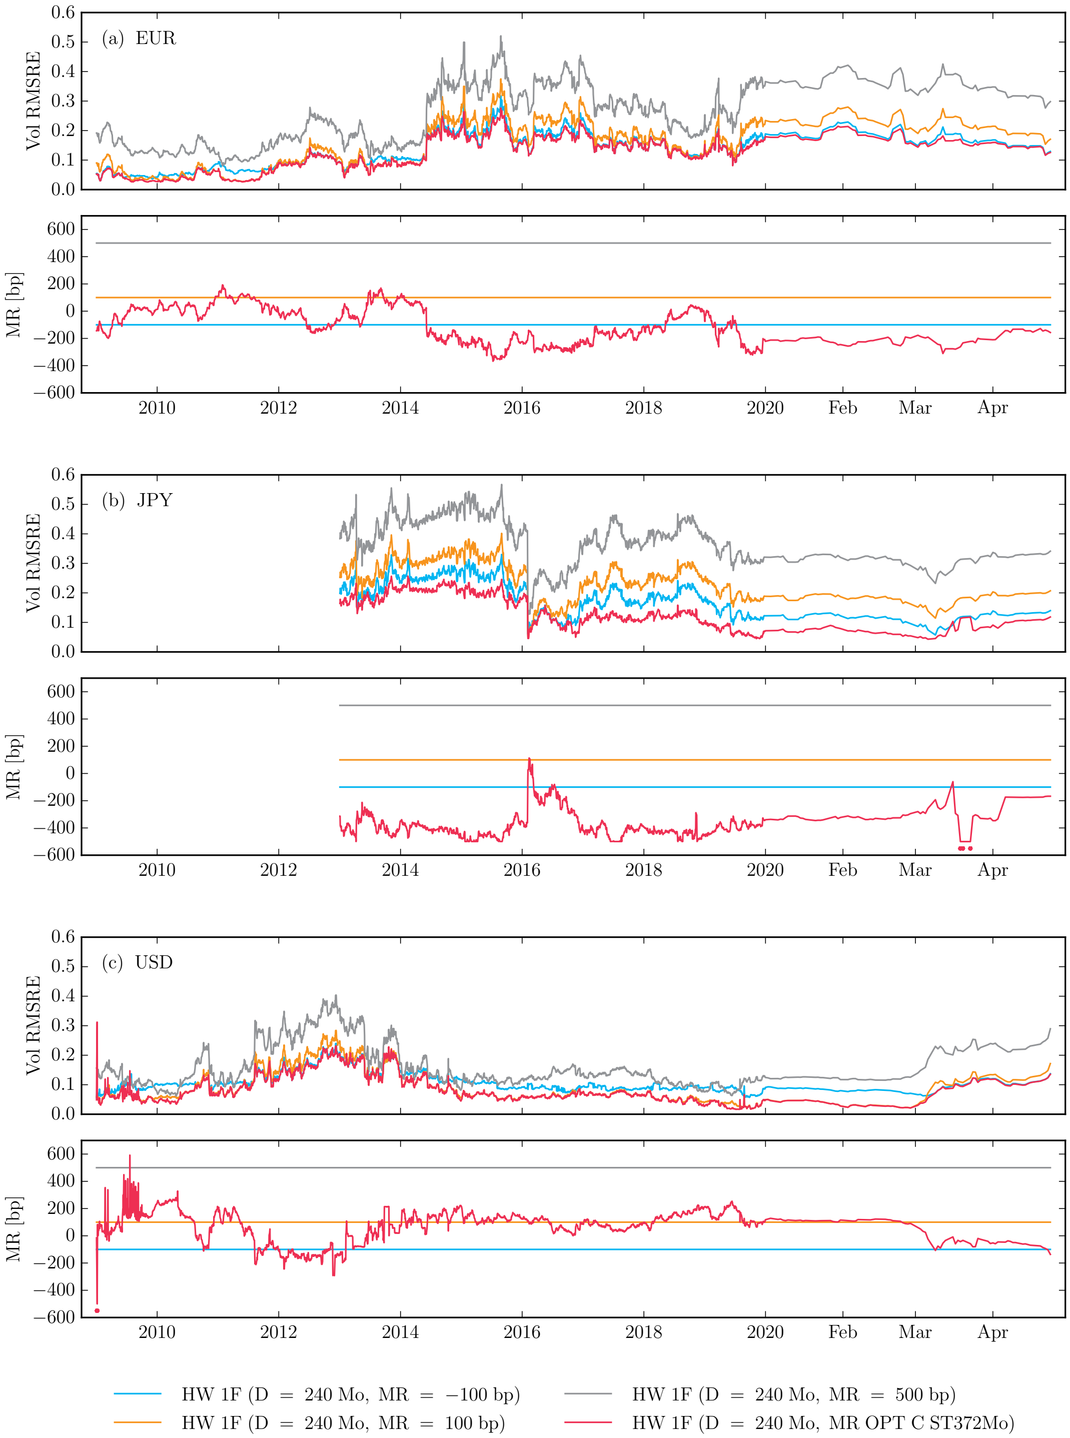

Figure 3: Swaption fit (Vol RMSRE) and mean reversion parameter (MR) time series for EUR, JPY, USD, based on HW 1F volatility calibration to the 20-year swaption diagonal (co-terminal). The dots under the calibrated mean reversion curves (red) indicate instances where a well-defined HW 1F model specification is elusive due to a failing mean reversion parameter optimization and/or instrument failures during the volatility calibration. Note also, that due to the constraints imposed on the optimization, the mean-reversion-optimized swaption fit can appear sub-optimal in comparison to the fixed mean reversion parameter alternatives. For more details see main text.

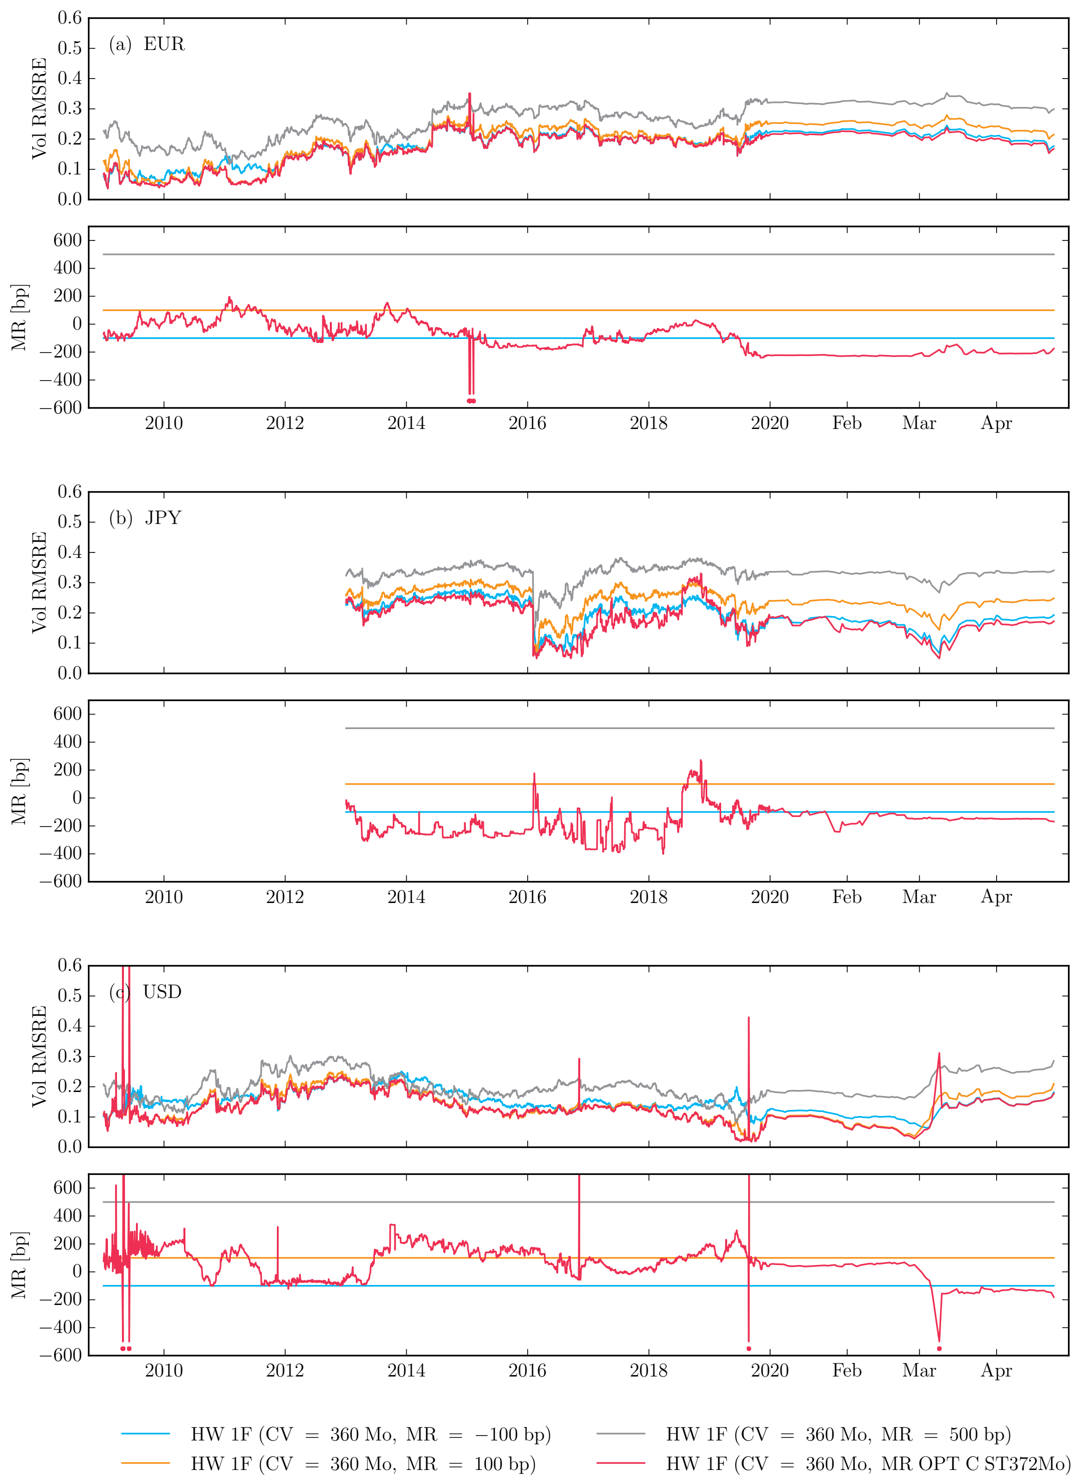

Figure 4: Swaption fit (Vol RMSRE) and mean reversion parameter (MR) time series for EUR, JPY, USD, based on HW 1F volatility calibration to the 30-year swaption chevron. See caption of Fig. 3 and main text for more details.