A Calibration of the Shadow Rate to the Euro Area Using Genetic Algorithms

Eric McCoy and Ulrich Clemens

EUROPEAN ECONOMY

Economic and Financial Affairs

***Europe Direct is a service to help you find answers
to your questions about the European Union.***

Freephone number (*)**:**
00 800 6 7 8 9 10 11

(*) The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

A Calibration of the Shadow Rate
to the Euro Area Using Genetic Algorithms

Eric McCoy, Ulrich Clemens

Abstract

In the face of the lower bound on interest rates, central banks have relied on unconventional policy tools such as large-scale asset purchases and forward guidance to try to affect long-term interest rates and provide monetary stimulus to the economy. Assessing the impact of these measures and summarising the overall stance of monetary policy in this new environment has proven to be a challenge for academics and central banks. As a result, researchers have worked on modifying current term structure models and have adapted them to the current situation of close to zero or even negative interest rates. The paper begins by providing a non-technical overview of Leo Krippner's two-factor shadow rate model (K-ANSM2), explaining the underlying mechanics of the model through an illustrative example. Thereafter, the paper presents the results obtained from calibrating Krippner's K-ANSM2 shadow rate model to the euro area using genetic algorithms and discusses the pros and the cons of using genetic algorithms as an alternative to the optimisation method currently used (Nelder-Mead optimisation routine). Finally, the paper ends by analysing the strengths and weaknesses of using the shadow short rate as a tool to illustrate the stance and the dynamics of monetary policy.

CONTENTS

1. INTRODUCTION

Traditionally, the most important indicator for a central bank's monetary policy stance is the policy rate. As such, it can be assessed against historical or normative benchmarks such as the Taylor rule to gauge the appropriateness of the level of monetary policy accommodation. The convergence of policy rates towards a lower bound that is assumed to be close to zero or even in negative territory in combination with the adoption of unconventional measures targeting the longer end of the yield curve (e.g. large scale asset purchases or various forms of forward guidance on the expected policy rate path), however, have impaired the function of policy rates as a summary indicator of a central bank's monetary policy stance. In light of this, alternative measures of the central bank's policy stance that can help quantify those monetary stimulus measures that go beyond steering the policy rates have been proposed. For example, the size of the central bank balance sheet should reflect quantitative easing measures, but cannot capture policy measures that are targeted at influencing markets' expectations regarding future policy rates and announcement effects.[1] By contrast, measures built on financial market prices can capture those expectation effects. To capture the monetary policy induced flattening of the yield curve, for instance one can simply look at the spread between long- and short-term market rates. Nevertheless, this rather rough measure does not provide a continuous stance measure that reflects monetary policy in both a conventional and unconventional environment. Alternatively, monetary conditions indices (MCI), which typically combine information on real interest rates and exchange rates, can be calculated. However, the inclusion of variables such as the exchange rate implies that the MCI can be substantially driven by factors outside the control of the central bank, for example foreign central banks' monetary policies. Lastly, and most prominently, term structure-based shadow short rate (SSR) models such as the one presented here have been proposed to capture monetary policy at the lower bound.[2] The advantages and disadvantages of this particular measure of the monetary policy stance are discussed in the following sections.

Term structure modelling used to be rather straightforward in the non-zero lower bound (ZLB) economic context that prevailed before the global financial crisis. In fact, Gaussian affine term structure models (GATSMs), including the popular subclass of arbitrage-free Nelson and Siegel models (ANSM), were widely favoured by many academics, researchers, central bankers and private-sector practitioners. Notwithstanding their popularity and widespread use, GATSMs have one well-acknowledged downside which was deemed small relative to their benefits in the pre-crisis period. That is, the mathematical specifications underlying the stochastic process for the short rate imply non-zero probabilities of interest rates of any maturity evolving to negative values. However, allowing for interest rates to evolve to negative levels in a GATSM model becomes untenable in a ZLB environment as the model would no longer fit the observed market yield curve (which is constrained by the ZLB). Therefore, the market-calibrated GATSM cannot provide anymore a valid representation of the term structure and its dynamics.

Recently developed models for the shadow rate offer a solution in the sense that on the one hand the Gaussian diffusion process for the "shadow" short rate can freely adopt negative values, while on the other hand the "ZLB" short rate diffusion process will remain constrained by the "lower bound". These academic developments are essential, as there is no longer a misspecification of the model relative to the data being modeled (i.e. the observed market yield curve) and at the same time by virtue of the fact that the SSR can take on negative values, the dynamics of the shadow short rate can now reflect the effects of non-standard monetary policy measures (like the asset purchase programmes as well as

---

[1] Furthermore, the mapping from the central banks' balance sheet size to market prices and interest rates is uncertain and depends on various factors.

[2] A somewhat hybrid, purely data-oriented approach that incorporates both price and quantity measures has been proposed by Lombardi and Zhu (2014), who estimate a shadow short rate for the US by pooling together a large dataset of both price and quantity variables and summarising this information by extracting common factors through a dynamic factor model. These latent factors are then used to construct a shadow federal funds rate, such that this shadow rate is a weighted average of all monetary information contained in the original dataset, with weights determined on the basis of the historical correlations.

forward guidance). The SSR is one of the outputs from the shadow rate term structure model which suggests itself as a means of quantifying the stance of monetary policy in ZLB environments.

While it is a priori not clear exactly how low the effective lower bound on nominal interest rates is, it is widely accepted that interest rates cannot be set arbitrarily low as market participants always have the 'cash option' of holding cash at a zero nominal interest rate instead of accepting negative interest. The strike price of said option then constitutes the lower bound to the interest rate, whereby the existence of transaction and storage costs of cash holdings might explain the fact that the actual lower bound is below zero. Shadow short rates derived from term structure models rely on the assumption of the presence of such a cash option. The shadow short rate then represents the interest rate that would prevail in a hypothetical world without this cash option, e.g. a world where cash is not available. As will be explained later on in the paper, the K-ANSM (2) model simultaneously estimates the shadow rate (based on GATSM principles) and the option effect such that both combined fit the observed market data.

In this paper we contribute to the existing literature on the shadow rate by testing an alternative method to calibrate shadow rate models to market data. More specifically we use genetic algorithms to calibrate Krippner's (2016) two-factor shadow rate model to the euro area OIS swap curve, showing that overall the results obtained corroborate those of Krippner (with even slightly improved calibration results). We also investigate and confirm Krippner's finding that the iterated extended Kalman filter (IEKF) outperforms the extended Kalman filter (EKF) by providing much more reliable and accurate parameter estimations from different starting points, which compensates for the fact that the Nelder-Mead optimisation routine is known to often result in a "local" optimum. Our empirical tests also demonstrate that this compensation is not perfect and thus this is where genetic algorithms may prove to be useful when facing new datasets (new regions, other countries or yield curves etc …), where one has no idea about the starting values for the parameters.

The paper is organised as follows: after providing a non-technical overview of Krippner's term structure based shadow rate model in Section 2, we discuss how to calibrate the model to yield curve data using genetic algorithms in Section 3; we then go on to discuss in Section 4 results obtained from calibrating Krippner's two-factor shadow rate model, K-ANSM (2), to the Euro area (EA) using genetic algorithms as well as the pros and cons of using genetic algorithms as compared to currently/commonly used calibration methods. In Section 5 we summarise the strengths and weaknesses of using the shadow short rate as a tool to illustrate the stance and the dynamics of monetary policy. Finally, Section 6 concludes the paper and Section 7 (technical appendix) provides a succinct technical review of the K-ANSM (2) model.

The paper is designed in such a way as to appeal to two groups of readers: (1) those with prior knowledge of term-structure modelling and shadow rate models and (2) readers without such technical background knowledge. The latter group might primarily be interested in the practical application of the shadow rate model to the fields of macro-economics, monetary policy and financial markets (bond option pricing, interest rate risk management …). This group of readers is invited to start by reading Section 2 which provides a non-technical overview of the shadow rate model, and then subsequently jump to Sections 4 and 5 which in turn discuss the calibration results for the euro area (EA) as well as the pros and cons of the shadow rate as a tool to illustrate monetary policy stance. The first group of readers, with a background in term-stucture modelling and knowledge of the shadow rate literature, will be essentially interested in the new calibration method proposed (based on genetic algorithms) in this paper. This group of readers is thus invited to skip the non-technical review and directly start with Section 3, which discusses the calibration algorithm in more detail in addition to the yield curve dataset used for the calibration, followed by Section 4, which provides the calibration results for the Euro Area (EA) and which also discusses the advantages/disadvantages of calibrating with genetic algorithms. Section 7 in addition offers a succinct technical refresher of the K-ANSM (2) shadow rate model. As the paper is designed to address both readers with and without knowledge of term-structure modelling and shadow rate models, reiterations of certain key aspects throughout the paper are inevitable.

2.   A NON-TECHNICAL OVERVIEW OF THE SHADOW RATE

The two-factor model of Krippner, commonly referred to as K-ANSM (2), which is used in this paper is a term-structure based shadow rate model that has received much attention in the literature. To be more specific, Krippner uses an affine term structure model (GATSM) to capture yield curve dynamics through time. The term structure of interest rates refers to the relationship between the yields-to-maturity of a set of bonds (or swaps) and their times-to-maturity. It is a simple descriptive measure of the cross-section of bond/swap prices observed at a single point in time. An affine term structure model hypothesises that the term structure of interest rates at any point in time is a mathematical function of a small set of common state variables or factors. Once assumptions regarding the dynamics of the state variables are specified, the dynamics of the term structure are also determined. However, in times of negative interest rates, these affine term structure models are not able to take into account the lower bound on interest rates. Fischer Black (1995) first introduced the notion of a "zero lower bound" on interest rates by explaining that so long as people can hold cash, they would prefer to keep money in their mattresses rather than holding financial instruments bearing negative interest rates. In his original formulation Black explains that interest rates cannot become negative because of the existence of this option for people to switch to holding cash when rates fall below zero. Since then, the methodology has been refined and today it allows for a lower bound which is negative. Nevertheless, the basic principle remains: the term structure model has to somehow incorporate this "option effect" to account for the lower bound on interest rates. In this section, we provide a non-technical description of the shadow rate model by illustrating how it works in the context of a specific example. To set the scene for the example, we will focus on the period between March-2016 and September-2016 and the specific question we will look at is whether and how the shadow rate model picks up on the observed flattening of the OIS yield curve during this 6-month period?

Before starting with a description of the shadow rate model itself, we briefly discuss the interest rate context during this 6-month period. Graph 1 below illustrates the market OIS yield curve for a selected number of maturities (1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y) as observed in March-16 (red-squares) and in Sep-16 (green-diamonds). We immediately notice that while the OIS curve during this 6-month period does not move much at the short-end (below 1Y), it flattened for the longer maturities (beyond 3Y). The EONIA in this same period basically remained stable (going from -0.30% in March-16 to -0.33% in Sep-16) as the ECB's deposit facility rate also remained unchanged at -0.40% throughout this 6-month period. If we look at the 10Y zero-coupon rate, it dropped from 0.31% in March-16 to 0.04% in Sep-16 thus confirming the observed flattening of the yield curve. Given the rigidity of interest rates at the short-end of the yield curve due to the assumed presence of a lower bound[3], the ECB opted to implement so-called non-standard monetary policy measures which are explicitly targeted at the longer end of the yield curve. In the next paragraphs, we will explain how the shadow rate model is able to incorporate the information from the yield curve flattening between March-2016 and Sep-2016 and how it translates into a shadow rate estimate which falls due to this yield curve flattening.[4]

The first step towards acquiring an understanding of the shadow rate model is to grasp the basic idea underlying term structure models of the yield curve (commonly referred to as General Affine Term Structure Models or GATSM). Coming back to our example in Graph 1, the red squares represent the observed yield curve rates for maturities 1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y in March-16. The question is how do we obtain the red-dotted smooth line which seems to connect (although not perfectly) the red-squares (i.e. observed market zero-coupon rates)? GATSMs allow us to calculate any point on the yield curve depicted by the red-dotted smooth line and even extrapolate below the 1M rate for example to obtain the 1week rate or the overnight rate or beyond the 10Y rate to calculate the 14Y or 20Y interest rates for example. So how do GATSMs enable us to calculate the interest rate for any maturity? GATSMs look back through the historical evolution of the observed yield curve points

---

[3] The interested reader can find in Section 5 a more detailed discussion of the choice of the lower bound.

[4] Whether this flattening was monetary policy induced or not will not be discussed in this didactic example.

(1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y) to extract two underlying factors which are sufficient to model the dynamics of the entire yield curve though time: they are referred to as the "level" and "slope" parameters. Accordingly, the zero-coupon interest rate for any time to maturity and at any given point in time is a function of these two underlying factors.

Graph 1 — Evolution of the EA OIS and shadow yield curves (Mar-16 to Sep-16).
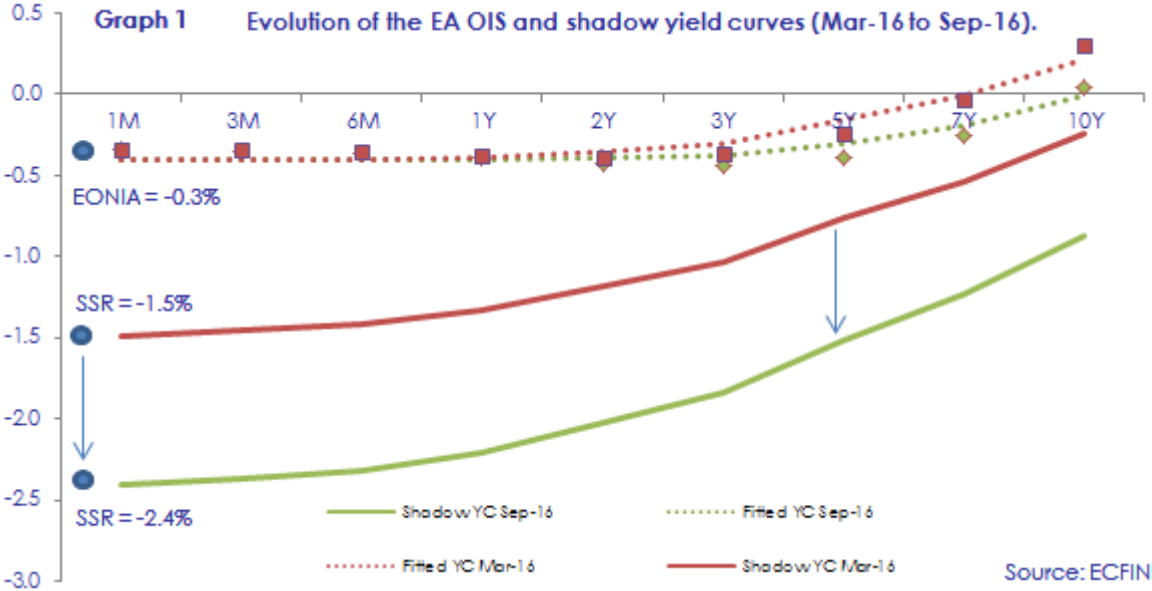
In essence, the GATSM captures the dynamics of the entire term structure of interest rates by extracting the "level" and "slope" components and thereby generates a time series of "level" and "slope" coefficients[5]. In order for the GATSM to be able to extract these two factors from the historical evolution of the yield curve, some assumptions need to be made regarding the behavior of the short rate (i.e. overnight rate) that drives the entire yield curve. Loosely speaking, the evolution of the short rate in time is assumed to be composed of a long-term trend term and a variance term which accounts for random market shocks. Thus, the short rate evolves over time according to some kind of general trend, but the model refines this even further by adding two dimensions to this trend: a long-term mean parameter as well as a mean-reversion parameter. In other words, when the level of the short rate deviates too much from its long-term mean, it will revert back to this mean at a speed governed by the mean-reversion parameter. This mean-reversion process is hampered in its ability to get back to its long-term mean level due to the diffusion or shock term[6].

The second step towards acquiring an understanding of the shadow rate model is to grasp how we go from the market yield curves to the "shadow" yield curves. This is a key step and once again we will illustrate it with the help of our specific example. In Graph 1, we see that there exists a hidden "shadow" yield curve in March-16 (red solid line) which lies well below the market yield curve (red dotted line). This "shadow" yield curve is based on the previously discussed GATSM principles and thus we can easily extrapolate (i.e. slide along the red solid curve to the blue dot) to obtain the shadow short rate (SSR) which according to the model stands at -1.5% in March-16. Before explaining how we obtain this "shadow" yield curve, first let us define the shadow short rate (SSR or equivalently shadow

---

[5] In other words, the estimated coefficients are time varying; at each point in time (i.e. each end-of-month) where we observe a market yield curve, the GATSM generates "level" and "slope" components specific to that point in time.

[6] For the interested reader, the appendix provides a succinct overview of mathematical assumptions underlying the arbitrage-free Nelson Siegel model (ANSM), a specific subclass of GATSMs, which serves as the basis for the shadow rate model of this paper.

rate). The shadow rate is commonly defined as the economic interest rate which would prevail in the absence of the lower bound on interest rates.[7] In our specific example in Graph 1, we have set the lower bound at -0.40% (the value of the deposit facility rate) to run the model calculations. That is, we assume that the central bank cannot steer money market rates below their current levels by adjusting downwards its key policy rate to below -0.40% in order to provide additional monetary stimulus because market participants would rather hold cash than accept an interest rate below that lower bound. Thus, as the yield curve is constrained at the short end, the central bank needs to resort to non-standard measures (affecting the longer-end of the yield curve) to implement monetary policy easing. If on the other hand, the lower bound constraint was absent, the central bank would have more control over the short end of the yield curve and could thus drive money market rates well below their current levels by further cutting its policy rates. The shadow rate is hence a theoretical unobserved interest rate which by definition has no direct bearing on the short-term financing conditions for companies. However, as we will see in the next paragraphs, the shadow rate – although theoretical – is affected by and captures the central bank's non-standard measures that affect the long end of the yield curve.
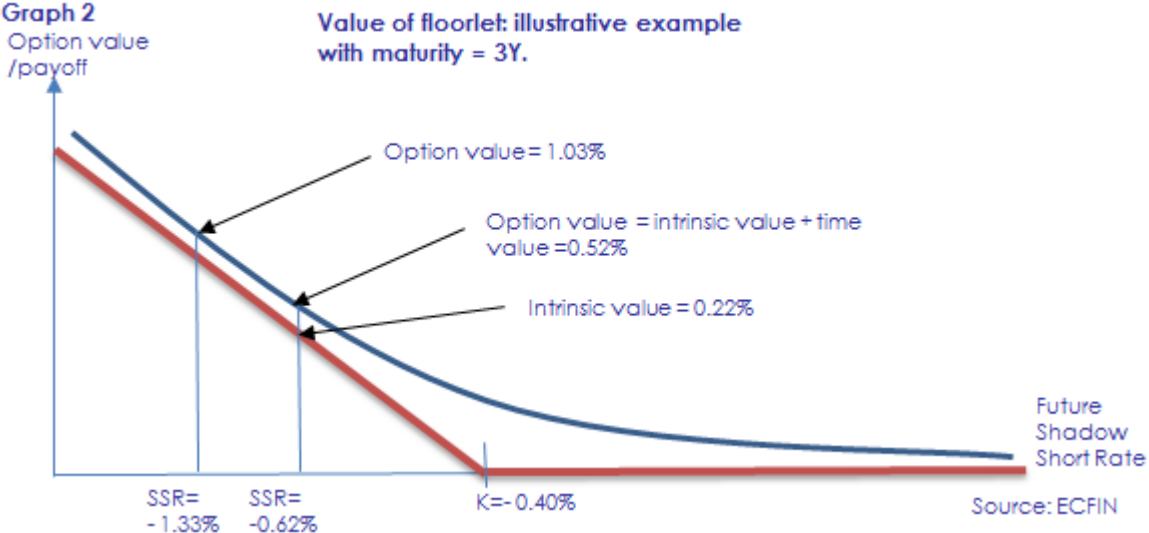
In order to understand why the "shadow" yield curve in March-16 (solid red line in Graph 1) lies below the market yield curve (dotted red line), we first need to take a short detour and introduce a commonly traded financial instrument called the "floor". An interest rate floor is a derivative contract in which the buyer receives payments at the end of each period in which the interest rate is below an agreed "strike" interest rate. An example of such a floor would be an agreement to receive a payment for each quarter that the EURIBOR 3M rate falls below 1%. In this example, the strike rate is 1% and the instrument underlying the floor is EURIBOR 3M. In financial markets, we find that the most commonly traded floors are on underlyings corresponding to the EURIBOR 1M, 3M or 6M. The floor itself also has a tenor or maturity; for example, a 5Y floor on EURIBOR 3M with a strike rate of 1% provides a positive payment at the end of each 3M period (known as the interest rate fixing or reset period) over the next 5 years whenever the 3M interest rate is below 1%. So the buyer of a 5Y floor on EURIBOR 3M with a strike of 1% could potential get 20 (4 quarters * 5 years) positive payments if EURIBOR 3M fixings all fall below 1%. Using financial jargon, we say that the buyer of the floor in effect holds a series of 20 "put" options on the interest rate with a strike rate of 1%. Each put option in this example covers the 3-month period in between two EURIBOR 3M fixings and is commonly referred to as a "floorlet". If interest rates fall substantially over the next 5 years then this implies that each put option or floorlet will automatically be exercised and will result every 3 months in a positive payoff to the buyer of the floor. Having given a brief description of the interest rate floor, let us know transpose this over to our shadow rate example.

Let us now imagine the following purely hypothetical example: we buy a 10Y-floor with the underlying being the EONIA and a strike rate of -0.40%. Notice that we chose the maturity of the floor as being equal to 10Y which is the longest maturity depicted on the yield curve in Graph 1, and the strike rate of -0.40% which corresponds to the ECB's deposit facility rate (which we used to set the lower bound on interest rates when running the shadow rate model). The buyer of this floor in effect has bought a series of floorlets or equivalently put options with a strike rate of -0.40%. Every day, the EONIA's fixing is observed and if this fixing falls below -0.40% then the buyer of the floor receives a positive payment which is equal to the difference between the strike rate of -0.40% and the fixing. If the fixing on any given day is above -0.40%, then there is no payment from the floorlet. Given the negative interest rate environment prevailing in March-16, with rates hovering not very far from -0.40% (the strike rate), this floor clearly has value for the buyer and hence the buyer would need to pay a premium to purchase this floor. The total value of the floor is in fact equal to the discounted sum of the values of the individual floorlets.. Moreover, it is important to note that in the shadow rate model itself, the underlying instrument of a given put option is not the future observed EONIA fixing but is rather the model's calculated "shadow" forward short rate. Loosely speaking, in the shadow rate model, the value of each individual floorlet/put option is added to the model's calculated "shadow"

---

[7] Correspondingly, the shadow yield curve can also be interpreted as the one that would prevail in the absence of the lower bound on interest rates.

forward short rate such that the sum is equal to the market's forward EONIA rate[8]. So one of the cornerstones of the shadow rate model is to understand how the value of a given individual floorlet changes in function of time to maturity and interest rates levels.
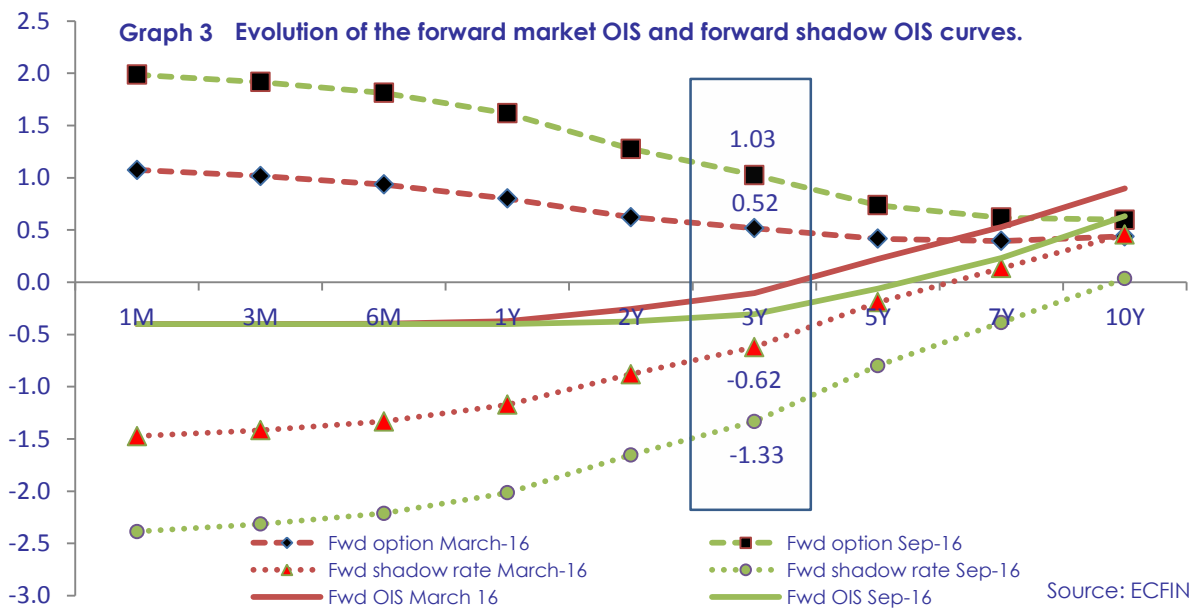
Pursuing our didactical example, suppose we focus on the following particular floorlet: expiry in 3Y with strike rate K = -0.40% which is a put option on the "shadow" short rate in 3 years. Graph 2 illustrates the value of the floorlet in function of the underlying (i.e. "shadow" short rate) and shows for example that if the "shadow" short rate in 3 years turns out to be -0.62%, then the payoff to exercising the floorlet is 0.22% (that is, the difference between the strike rate of the option which is -0.40% and the future value of the underlying which is -0.62%). But the option expires only in 3 years, and in the meantime the "shadow" short rate can still evolve to even more negative levels; this implies that the option value today is worth even more than the option's "intrinsic value" of 0.22%. The additional value is known as the option's "time value" and reflects the fact that until the floorlet's expiry in 3 years, "shadow" interest rates could potentially fall even lower which would imply that the future payoff of the option would be higher.[9] Conversely, a floorlet that expires e.g. in one day would have virtually no additional time value as the probability for an unexpected downward change in the time remaining is very low.

Graph 2
Option value /payoff

Value of floorlet: illustrative example with maturity = 3Y.

Option value= 1.03%

Option value = intrinsic value + time value =0.52%

Intrinsic value = 0.22%

Future Shadow Short Rate

SSR= -1.33%   SSR= -0.62%   K=-0.40%   Source: ECFIN

As previously mentioned, the example of the floorlet is purely didactical and serves to only illustrate in a non-technical manner the so-called "option effect" embedded in the shadow rate model. Moving on to the model itself, Graph 3 below illustrates the evolution of the market forward OIS curve between March-16 and Sep-16 (the solid red and green lines respectively). Lying underneath these forward market OIS curves are the "shadow" forward OIS curves for the same period (the dotted red and green lines respectivey). The difference between the forward market OIS and "shadow" forward OIS curves corresponds to the model's calculated forward option effect. For example and as illustrated in Graph 3, for the 3Y maturity the forward market OIS rate fell from -0.10% in March-16 to -0.30% in Sep-16 while the "shadow" forward OIS rates fell from -0.62% to -1.33% during the same period. In March-16, the difference between the forward market OIS rate (-0.10%) and the "shadow" forward OIS rate (-0.62%) was 0.52% which is exactly equal to the forward option value. Using a similar reasoning for Sep-16, we can observe that the gap between between the forward market OIS rate (-0.30%) and the "shadow" forward OIS rate (-1.33%) was 1.03%, exactly the value of the forward option effect.

---

[8] The mathematical details regarding how the option effect is embedded into the shadow rate model are covered in the technical appendix (Section 7).

[9] Of course, the shadow short rates could also rise in the meantime. However, the payoff risks are asymmetric, as the option value in the worst case becomes zero, whereas the positive payoffs are theoretically unlimited.
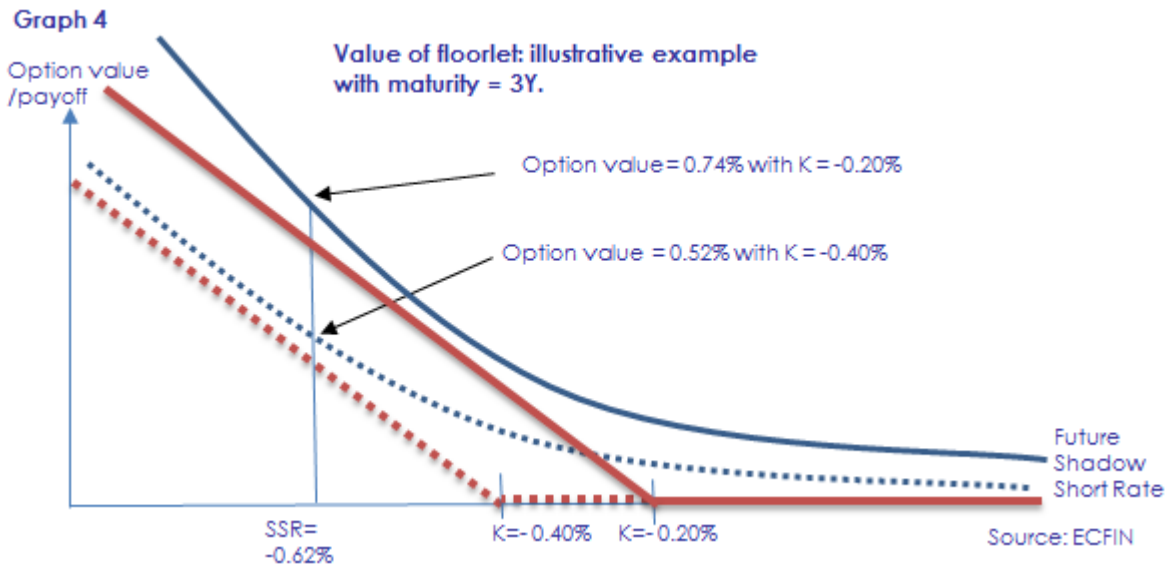
Graph 3   Evolution of the forward market OIS and forward shadow OIS curves.

2.5
2.0
1.5
1.0
0.5
0.0
-0.5
-1.0
-1.5
-2.0
-2.5
-3.0

1M   3M   6M   1Y   2Y   3Y   5Y   7Y   10Y

1.03
0.52
-0.62
-1.33

Fwd option March-16   Fwd option Sep-16
Fwd shadow rate March-16   Fwd shadow rate Sep-16
Fwd OIS March 16   Fwd OIS Sep-16

Source: ECFIN

The option effect has thus increased from March-16 to Sep-16 and this is due to the fact that the model's calculated "shadow" forward EONIA rate has dropped from -0.62% to -1.33% during this same period leading to a higher expected value of the put option (see Graph 2 to visualise the impact using once again our didactial example of the floorlet where we observe the option value increasing from 0.52% to 1.03%). The same reasoning applies to the other points or maturities along the yield curve; for example, the model's calculated "shadow" forward rate for the 3M maturity also dropped substantially from March-16 (-1.42%) to Sep-16 (-2.32%) even though the 3M spot market rate has not changed. This is because the calculated 3M "shadow" forward rates are computed based on the model's extracted "level" and "slope" factors in March-16 and Sep-16 respectively, which take into account the observed flattening on the longer end of the yield curve during this period. The option value for this 3M maturity increases because as the "shadow" forward EONIA drops, the option gains in intrinsic value. The increase in the option value compensates for the drop in the "shadow" forward rate. This leads us to the conclusion that the area between the forward market OIS curve and the "shadow" forward OIS curve is explained by the option effect. This is a key aspect of the shadow rate model, which is worth summarising at this point: the model's forward "option effect" when added to the "shadow" forward OIS curve enables us to get back to the market forward OIS curve. Given that the shadow rate model is calibrated to actual market data, we can therefore safely add that the model is calibrated in such a way that the calculated forward option effect when added to the "shadow" forward OIS curve (which is itself based on GATSM extracted "level" and "slope" parameters) enables us to obtain a fit to the observed market OIS data[10], i.e. the model simultaneously estimates the shadow yield curve and the option effect such that both combined fit the observed data. This also implies that when observed market rates are far above the lower bound – that is when the lower bound is "not binding" – the shadow yield curve coincides with the observed market yield curve as the option value is zero. Thus, in times when the lower bound is not binding, the shadow short rate closely tracks the observed EONIA, making the shadow short rate a measure that is applicable both in conventional and unconventional monetary policy environments. Focusing now only on the "shadow" short rate (i.e. the two blue dots in Graph 1 representing the shortest maturity on the yield curve), the workaround which enables the shadow rate model to be calibrated to observed market yield curve data (which is constrained by the lower bound) while at the same time allowing the "shadow" short rate to take on negative values is illustrated by the following equation:

---

[10] To be specific, Krippner has opted to calibrate the shadow rate model to the zero-coupon yield curve rates instead of the forward rates. The interested reader will find more details about the calibration dataset and algorithm in Section 3 as well as in the technical appendix in Section 7.

$$\underline{r}(t) = max\{r_L, r(t)\} = r(t) + max\{r_L - r(t), 0\}$$

where $\underline{r}(t)$ is the actual market-observed short rate which is subject to the lower bound constraint (denoted by r_L) and $r(t)$ is the non-observable "shadow" short rate (i.e. the two blue dots in Graph 1). Thus, the shadow short rate $r(t)$ can be viewed as the short rate that would prevail in the absence of the lower bound on nominal interest rates. The shadow short rate $r(t)$ is estimated using GATSM "level" and "slope" factors while the value to the option-component $max\{r_L - r(t), 0\}$ reflects the difference between the lower bound and the "shadow" short rate[11]. What else can affect the option's value and thus the shadow rate besides a flattening of the observed market yield curve? As the graph below illustrates, the option's value is strongly affected by its strike rate. Coming back once again to our didactical example of the floorlet, and as illustrated in Graph 4 below, setting the strike rate of the option to -0.20% instead of -0.40% will lead to a higher option value for a given level of the underlying rate. In fact, the pricing/value curve of the option shifts upwards from the blue dotted curve (corresponding to a strike rate of -0.40%) to the solid blue curve (corresponding to the strike rate of -0.20%). For example, if the underlying "shadow" rate is currently at -0.62%, then it makes intuitive sense that the option value increases because it has more intrinsic value for this level of the underlying "shadow" rate and thus we say that the option becomes further "in-the-money" (where the strike price becomes more favorable in relation to the current underlying price). If we now apply the above reasoning to our actual shadow rate model then it becomes clearer why the choice of the lower bound will have a significant impact on the model's forward "option effect": as the lower bound is set to a less negative level (i.e. from -0.40% to -0.20%) then the forward "option effect" increases in value across all maturities which in turn generates a lower forward shadow yield curve. But if the forward "shadow" yield curve shifts downwards, then this implies that the shadow short rate (SSR) falls as well.

Graph 4

Option value /payoff

Value of floorlet: illustrative example with maturity = 3Y.

Option value = 0.74% with K = -0.20%

Option value = 0.52% with K = -0.40%

Future Shadow Short Rate

SSR = -0.62%

K = -0.40%   K = -0.20%

Source: ECFIN

Graph 5 below illustrates how setting the lower bound at -0.20% instead of -0.40% has resulted in lower forward "shadow" rates across all maturities in Sep-16. As a consequence, the shadow short rate

---

[11] For points/maturities other than the spot "shadow" short rate (i.e the forward "shadow" rates), the option effect/component is a more complex mathematical function (see Section 7 for the details).

(SSR) estimate itself in Sep-16 is estimated at -3.51% when using a lower bound of -0.20% as opposed to -2.42% when setting the lower bound to -0.40%.

Graph 5 Impact of lower bound on forward shadow OIS curves in Sep-16.
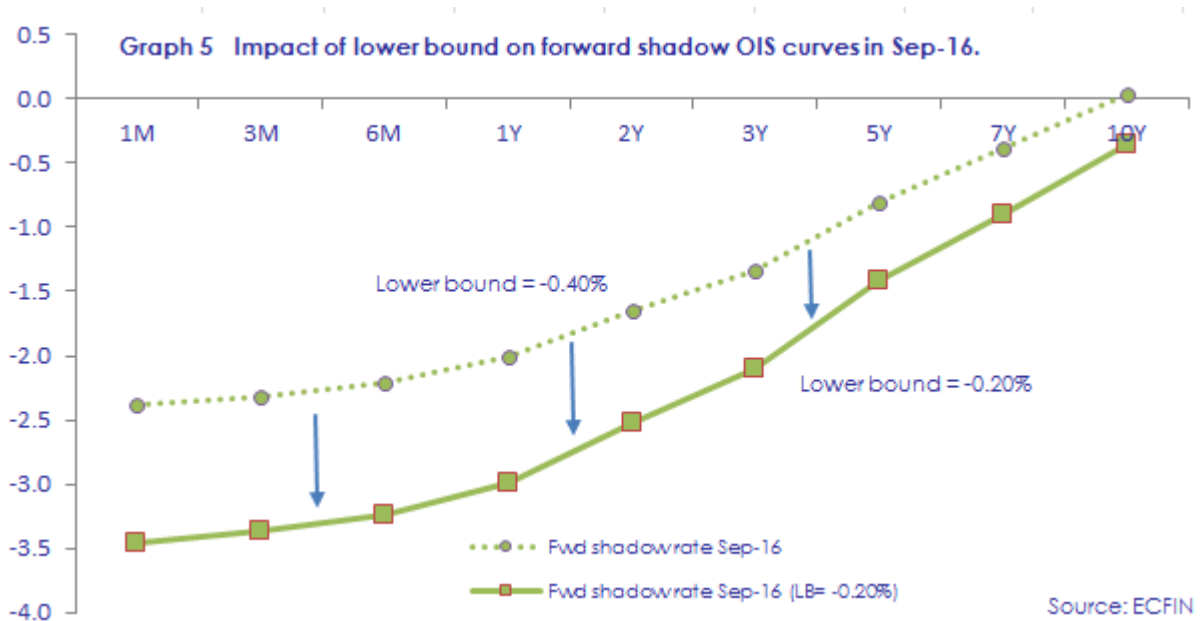
In fact, one of the fundamental characteristics of Krippner's two factor model is that when the lower bound is set to be less negative (i.e. closer to zero) it tends to capture earlier and reflect more strongly the observed movements of the market yield curve as compared to when running the model with a more negative lower bound whereby the option effect is less pronounced. Graph 6 below illustrates this point and it shows the evolution of the euro area shadow rate since Dec-11 (when shadow rates first started to become negative) with a lower bound of 0% and also with a lower bound of -0.40%. The graph also plots the evolution of the 10Y – 3M OIS market yield spread (a possible proxy for the "slope" of the market yield curve) and the 3M OIS market yield (taking this as a proxy for the "level" of the market yield curve). The graph is further broken down into different periods characterised by either a flattening/steepening of the yield curve and/or a level change in the yield curve. One clearly sees that when the shadow rate model is run with a lower bound of 0% instead of -0.40%, it is more sensitive to and better reflects changes to the market yield curve's observed level and slope "proxy" parameters. For an assumed lower bound of -0.40%, the estimated shadow rate does not deviate meaningfully from EONIA until late 2015, as the lower bound only becomes binding afterwards, implying a positive value of the cash option. Conversely, for an assumed lower bound of 0% the cash option is "in the money" already by mid-2012 and the estimated shadow rate can pick up on yield curve movements relatively early.

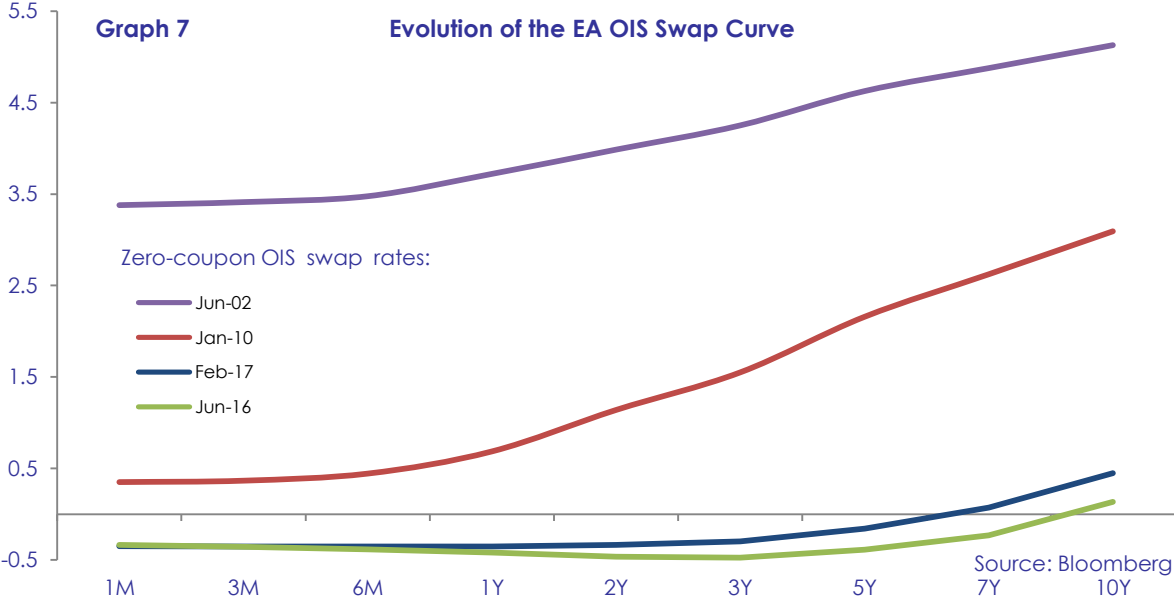Graph 6   Shadow rate estimates with different lower bounds.

Source: ECFIN
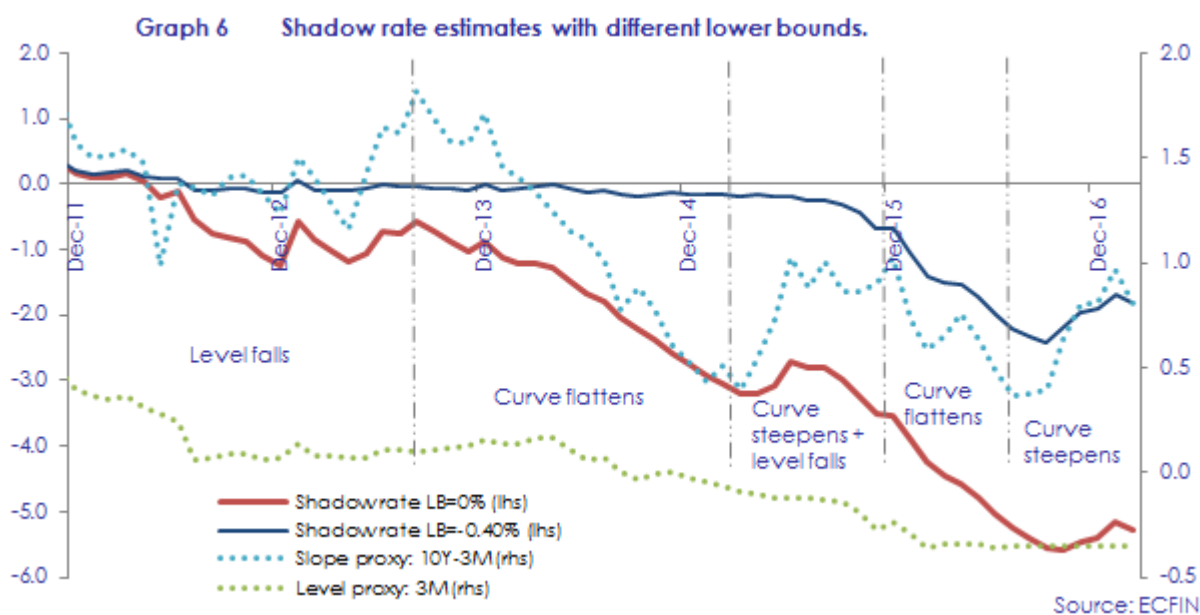
Although we chose to provide a non-technical summary of the shadow rate model in this section by focusing on a specific period (namely from March-16 to Sep-16) during which we observed a flattening of the OIS market yield curve and explained how the model produces a decrease in the shadow rate, the same conceptual framework also applies in the case whereby the OIS market yield curve experiences a downwards parallel shift. In the case of a downwards parallel shift or even a combination of a downwards parallel shift and flattening, the shadow rate should also in principle fall as a result of the yield curve movement. Conversely, in periods where we observe a steepening, or even a combination of an upwards parallel shift and steepening of the OIS market yield curve, the shadow short rate (SSR) should consequently increase. We therefore conclude this section by acknowledging that the shadow short rate (SSR) is reactive to changes in both the level and the slope (flattening/steepening) of the OIS market yield curve. Section 5 discusses in more detail the link between such yield curve changes and monetary policy, namely the advantages/disadvantages of the shadow rate as a stance indicator of monetary policy.

3. CALIBRATING THE SHADOW RATE MODEL

3.1   THE CALIBRATION DATASET

In this sub-section, we provide a succinct description of the dataset used to calibrate Krippner's two factor K-ANSM (2) shadow rate model. Graph 7 below illustrates the movements of the EA OIS swap curve (zero-coupon rates) from June 2002 (the beginning of the sample period) to February 2017 (the end of the sample period). As euro area OIS markets were generally not well developed in the early years of the monetary union, reliable data at least for the longer maturities is not available pre-2006. We therefore only use OIS data from 2006 onwards and apply the percentage evolution of Euribor swap rates before January 2006 to back-fill the OIS swaps dataset to June 2002. Generally, as the overnight swap contracts only involve the exchange of net payments and no principal, OIS rates can be seen as risk-free and thus can better reflect monetary policy than for example government bond yields, which might mirror other developments, e.g. with regard to countries' credit risk. Over the sample period, the EA OIS curve has shifted down materially, resulting in a relatively flat curve up to the two year maturity in February 2017 with most yields (up to 7 years maturity) in negative territory. A

clearly inverted swap curve with swap rates as low as -0.47% at the 3 year maturity was observed at the end of June 2016. The K-ANSM (2) model aims to capture the movements of the OIS swap curve through time by extracting two hidden factors, commonly referred to as the "level" and "slope" components of the yield curve. A mathematical algorithm (the Kalman filter) extracts these hidden yield curve components/factors, which vary through time, in such a way so as to obtain the best fit with the observed market yield curve data.

Graph 7   **Evolution of the EA OIS Swap Curve**

Zero-coupon OIS swap rates:

- Jun-02
- Jan-10
- Feb-17
- Jun-16

Source: Bloomberg

(1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, 10Y)

THE CALIBRATION ALGORITHM

The model's underlying state variables/factors are estimated using a Kalman filter approach, which is useful in situations such as this one, where the underlying state variables are not directly observable. In the particular case of the two factor model, the first state variable can be interpreted as the "level" component of the yield curve while the second state variable represents the yield curve's "slope" component. The key to estimating these "hidden" factors lies in the relationship between the bond prices and the underlying state variables. Indeed, the calibration algorithm begins with an observed system of equations called the measurement system; this system represents exactly this affine relationship between market zero-coupon rates (which is a simple logarithmic transformation of the bond price function) and the unobserved state variables. A second, unobserved system of equations called the transition system describes the dynamics of the state variables as they were originally formulated in the model (i.e. the process for the short rate). Together, the measurement and transition equations represent what is called the state-space form of the model. Once the initial conditions (for the state variables and state variance matrix) have been specified (the so-called "priors") and given a set of starting values for the model parameters (which define the stochastic process for the short rate) to be optimised, the Kalman filter then uses this state-space formulation to recursively make inferences about the unobserved values of the state variables (transition system) by conditioning on the observed market zero-coupon rates (measurement system). These recursive inferences are then used to construct and maximise a log-likelihood function to find an optimal parameter set for the system of equations [see Section 7 for a mathematical description of the shadow rate model]. The addition of the "option" component introduces non-linearity in the Kalman filter's measurement equations, and hence the calibration of the K-ANSM (2) model requires a non-linear Kalman filter.

To take into account non-linearities in the Kalman filter's measurement equation, a commonly used method is to use the extended Kalman filter (EKF) which has been employed namely by Christensen and Rudebusch (2013a,b, 2014) and Wu and Xia (2013, 2014). The EKF effectively calculates a first-order Taylor approximation of the non-linear interest rate function around the best available estimate of the state variable vector. Krippner (2015a) points out that using the EKF within an optimisation algorithm often results in a likelihood value and parameter estimates that can be greatly improved on by using the iterated extended Kalman filter (IEKF) within the same optimisation algorithm. The iterated extended Kalman filter improves the linearisation of the extended Kalman filter by recursively modifying the center point of the Taylor expansion. This reduces the linearisation error at the cost of increased computational requirements. Krippner (2015a) demonstrates that the IEKF-based optimisation obtains a materially higher value of the likelihood function along with materially different parameter estimates as compared to the EKF-based optimisation, which hence justifies Krippner's choice of the iterated extended Kalman filter (IEKF) for calibrating the shadow rate model.

As described above, in order to compute the "optimal" model parameters, the Kalman filter algorithm is embedded within an optimisation routine which seeks to obtain those parameters which provide the best possible fit to the observed yield curve data. A very commonly used algorithm, which is also employed by Krippner, to maximise the model's fit to the observed yield curve, is the Nelder-Mead simplex search method of Lagarias et al (1998). One known weakness of this algorithm is that it provides a "local" solution (for the optimal model parameters) meaning that the algorithm returns a local optimum of the mathematical function near its starting point (for the model parameters) set at the beginning of the algorithm's run. As mentioned previously, Krippner (2015a) finds that in practice the IEKF outperforms the EKF in that it provides much more reliable and accurate parameter estimations from different starting points. However, it is not clear to what extent this increased stability stemming from using the IEKF compensates for the fact that the optimisation routine being used is known to often result in a "local" optimum. This is where the use of genetical optimisation algorithms may prove to be useful.

A genetic algorithm (GA) is a method for solving both constrained and unconstrained optimisation problems based on a natural selection process that mimics biological evolution. Genetic algorithms embed a probabilistic search which is founded on and mimics the idea of an evolutionary process. The GA procedure is based on the Darwinian principle of survival of the fittest. An initial population is created containing a predefined number of individuals (or solutions), each represented by a genetic string (incorporating the variable's information). Each individual has an associated fitness measure, typically representing an objective value. The concept that fittest (or best) individuals in a population will produce fitter offspring is then implemented in order to reproduce the next population. Selected individuals are chosen for reproduction (or crossover) at each generation, with an appropriate mutation factor to randomly modify the genes of an individual, in order to develop the new population. The result is another set of individuals based on the original subjects leading to subsequent populations with better (min. or max.) individual fitness. Therefore, the algorithm identifies the individuals with the best optimismaximising fitness values, and those with lower fitness will naturally get discarded from the population. Ultimately this search procedure finds a set of variables that optimises the fitness of an individual and/or of the whole population. As a result, the GA technique has advantages over traditional non-linear optimisation techniques which tend to be sensitive to the initial starting point and which may end up being trapped in local minima/maxima. The next section presents the resuls of calibrating Krippner's two-factor shadow rate model to the Euro area using genetic algorithms.

Before explaning the various steps that a typical genetic algorithm goes through to optimise a mathematical function, some preliminary terminology needs to be introduced. The function that the algorithm tries to optimise is called the "fitness function" and in our case it is the log-likelihood function which is given by Log_LKL= $f\{\phi, k_{11}, k_{12}, k_{21}, k_{22}, \theta_1, \theta_2, \sigma_1, \sigma_2, \rho_{12}, \sigma_\eta\}$ (see the technical appendix in Section 7 for more details). An "individual" is any point (input vector in our case) to which you can apply the fitness function and the value of the fitness function for a given individual is known as its "score". An individual is also commonly referred to as a "chromosome" and the vector entries of an individual as the "genes". A "population" is an array of individuals and as will be explained below during each iteration of the genetic algorithm a series of computations are made on the current population to produce a new population. Each successive population is commonly referred to as a new "generation".

Typically, a genetic algorithm goes through the following steps to optimise the fitness function (in our case this means finding the 11 input parameters that maximise the log-likelihood function):

1) The algorithm begins by creating a random initial population.

2) It then scores each individual in the current population by computing its fitness value. It ranks them by fitness value and subsequently retains only a select sub-group comprised of the fittest individuals, known as the "parents", to be able to produce the next generation.

3) The genetic algorithm creates three types of "children" which will form the next generation:

a) "Elite children" are the individuals/parents in the current generation with the best fitness values. These individuals automatically survive to the next generation.
b) "Crossover children" are created by combining the vector entries of a pair of parents.
c) "Mutation children" are created by introducing random changes, or mutations, to a single parent's vector entries or genes.

The algorithm then proceeds by replacing the current population with the children to form the next generation.

In an iterative manner, the algorithm creates a sequence of new populations/generations and it finally stops when one of the stopping criteria is met. Most often, the genetic algorithm is set up such that the algorithm stops if the relative change in the best fitness function value is less than or equal to a predefined user tolerance level (usually set at 1e-6).

What makes genetic algorithms (GAs) a good global optimisation tool?

(1) Most other algorithms are serial and can only explore the solution space to a problem in one direction at a time wheras GAs can explore the solution space in multiple directions at once. If one path turns out to be unfruitful, they can easily eliminate it and continue work on more promising avenues, giving them a greater chance at each step of finding the global optimal.
(2) A notable strength of genetic algorithms is that they perform well in problems for which the solution space is complex - ones where the fitness function is discontinuous, noisy, changes over time, or has many local optima. GA has proven to be effective at escaping local optima and discovering the global optimum in even a very rugged and complex fitness landscape.
(3) Another aspect in which genetic algorithms prove particularly efficient is in their ability to manipulate many parameters simultaneously (when the dimension of the solution space is high).
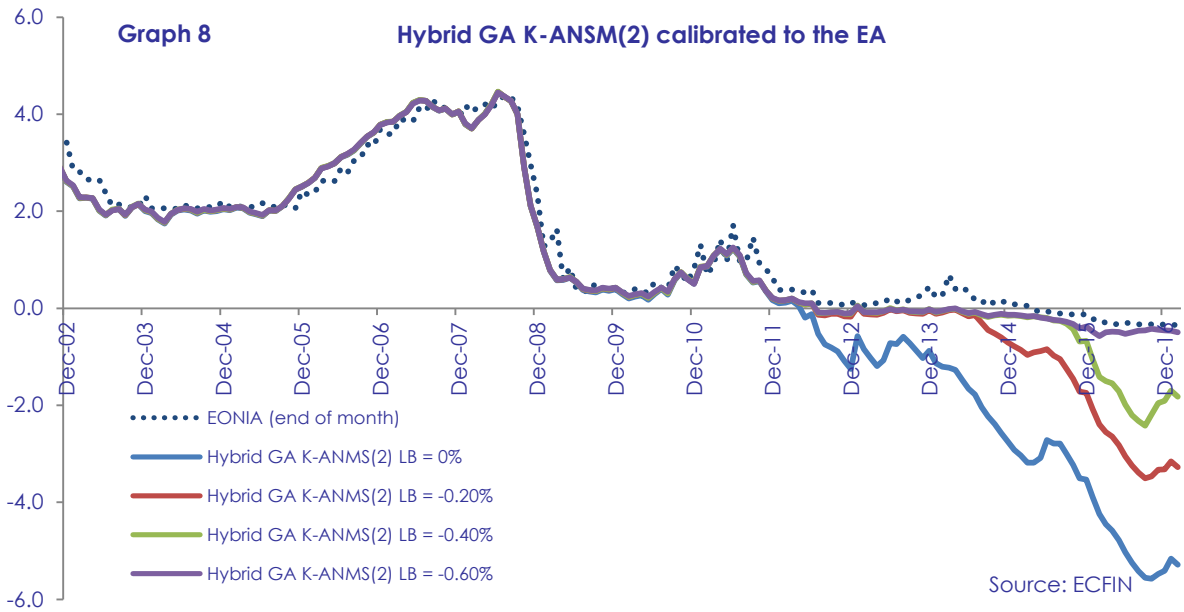
4. EMPIRICAL RESULTS FOR THE EURO AREA

Graph 8 below plots the shadow rate for the EA over the period Dec-02 to Feb-17 when calibrating the model using genetic algorithms with different scenarios for the lower bound: 0%, -0.20%, -0.40%, and -0.60% (for a discussion of the choice of the lower bound, see Section 5). The model was calibrated to the euro area OIS swap curve[12] using monthly data on the following yield curve tenors: 1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y. Two stylised facts emanate from the genetically calibrated shadow rate model, GA K-ANSM (2), as illustrated in Graph8. First, the GA K-ANSM (2) model seems to track the EONIA rate fairly well for all chosen lower bound calibrations in a non-ZLB environment, i.e. until the second quarter of 2012, when the lower bound for the first specification (LB = 0%) becomes binding.[13] EONIA is often used to proxy the ECB's interest rate policy since it conveniently includes the switch of the ECB's effective policy rate from the interest rate on its main refinancing operations (MRO) to the deposit facility rate (DFR) that followed the increase in the central banks' excess liquidity post October 2008, when tender operations with fixed rate full allotments were introduced.[14] Second, Graph 8 shows that the choice of the lower bound used to calibrate the GA K-ANSM (2) model results in substantially different estimates of the shadow rate. In fact, the shadow rate estimates in Feb-2017 vary from -0.5% (when using a lower bound of -0.60%) to as low as -5.3% (when using a lower bound of 0%). Choosing an appropriate level for the lower bound is thus a key decision resulting in potentially very different estimates of the shadow rate. This raises the question of what explains the model's sensitivity to the choice of the lower bound.

The answer to this question does not rest with the type of calibration algorithm used, but is rather linked to an inherent characteristic of the model: the value of the cash option. To illustrate this point, Graph 9 below plots the value of the cash option (using a forward start/tenor of one year as an example) for each of the different lower bounds chosen. Even by simple visual inspection of Graphs 8 and 9, one can easily deduce that the higher is the lower bound, the higher is the value of the cash option effect. As explained in Section 2, given that the market-calibrated short rate is the sum of the shadow rate and the cash option's value, in the case of a lower bound set at 0%, the highly positive intrinsic value of the cash option compensates for the model's estimated (based on the "level" and "slope" parameters) highly negative shadow short rate (-5.3% in Feb-2017) such that the sum equals the observed market short rate.This also makes sense intuitively especially if one thinks of the model's lower bound as the strike price of the cash option, i.e. the interest rate that this option guarantees upon exercise (e.g. 0% in the classical framework of switching to cash without any transaction and storage costs). The value of the option at the time of exercise should then be the difference of the option's strike rate and the "shadow" short rate (or zero if the "shadow" short rate lies above the strike rate). From a forward looking perspective, for any given probability distribution that assigns non-zero probabilities to OIS rates falling below the strike price, the value of the option will increase with the strike price, as a higher cumulative probability will be assigned to an outcome in which rates fall below the strike price (which would trigger exercising this option).

---

[12] We use Eonia based OIS rates for the period after January 2006 and apply the percentage evolution of Euribor swap rates before January 2006 to back-fill the OIS swaps dataset to June 2002 (see Section 3 for more details on the dataset).

[13] That is, although EONIA was still trading above 0%, non-zero probabilities are assigned to OIS rates below the lower bound, which implies a positive value of the cash option resulting in negative values of the respective shadow rate. This was intensified by the deposit facility rate cut to 0% in July 2012, which shifts a considerable part of the probability distribution into negative territory.

[14] While the Eurosystem had been operating under a structural liquidity deficit pre-October 2008, the introduction of tender operations with fixed rate full allotments created an environment of excess liquidity in which banks are incentivised to lend overnight to other banks at a small spread over the DFR. Excess liquidity increased considerably with the allotment of longer-term refinancing operations and in partular with the start of the asse purchases under the expanded asset purchase programme.

Graph 8   **Hybrid GA K-ANSM(2) calibrated to the EA**

Legend:
- EONIA (end of month)
- Hybrid GA K-ANMS(2) LB = 0%
- Hybrid GA K-ANMS(2) LB = -0.20%
- Hybrid GA K-ANMS(2) LB = -0.40%
- Hybrid GA K-ANMS(2) LB = -0.60%

Source: ECFIN

A key question emerging from this analysis is thus: what is an appropriate level to choose for the lower bound? Without further analysis, one can already notice that by simply setting the lower bound at the minimum observed value of the DFR over the calibration sample history, -0.40%, the resulting shadow rate estimate stands at -1.8% in Feb-2017. Although subjective, this appears to be a more reasonable estimate of the shadow rate from a monetary policy perspective. At the same time, this shadow rate estimate shows little variation before the start of the ECB's expanded asset purchase programme and thus seems to capture earlier unconventional monetary policies insufficiently (e.g. forward guidance, TLTROs).

Graph 9   **Hybrid GA K-ANSM(2) option effect for the EA**

Legend:
- Hybrid GA Option Effect LB = 0%
- Hybrid GA Option Effect LB = -0.20%
- Hybrid GA Option Effect LB = -0.40%
- Hybrid GA Option Effect LB = -0.60%

Option effect's tenor/forward start = 1 year

Source: ECFIN
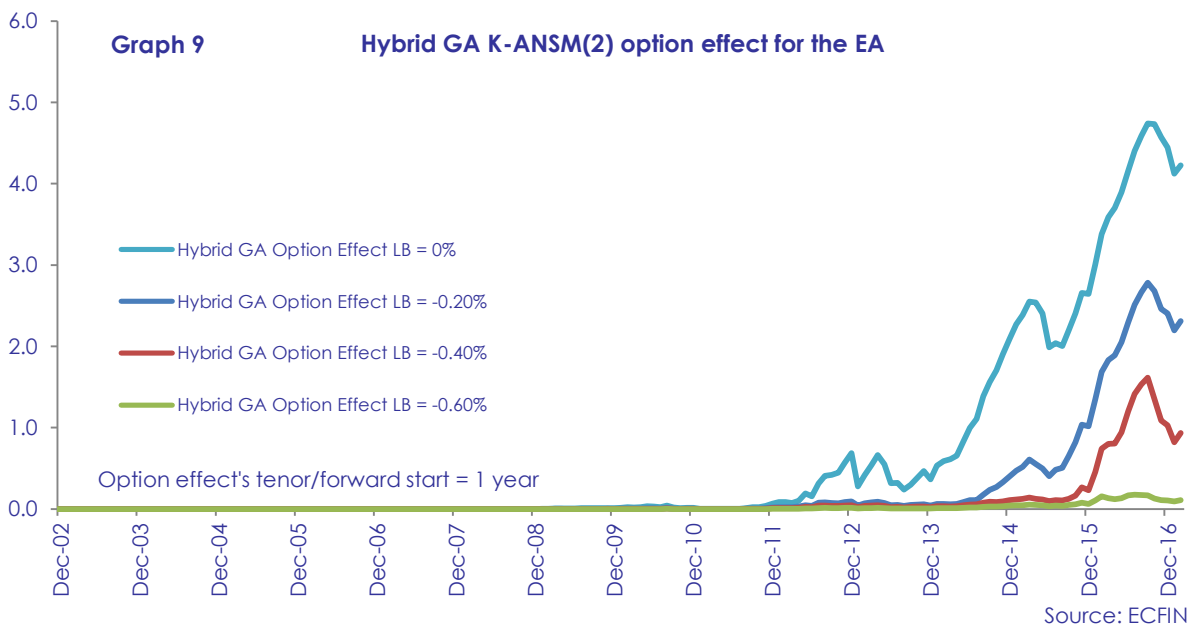
It is important at this point to exclude any possible distorting effects that the choice of the calibrating algorithm could have on the cash option value and thus also on the shadow rate estimate itself. In other words, do the more commonly used Nelder-Mead simplex algorithm implemented by Krippner and the genetic algorithm tested in this paper provide similar results? Table 1 below provides a comparison of the quality of the calibration between the two approaches. As can be seen, both the optimised log-

17

likelihood measure for the whole model and the RMSEs (root mean-squared errors) for the different maturities suggest that the results indeed are fairly similar.

Table 1	LB = 0%		LB = -0.20%		LB = -0.40%		LB = -0.60%	
	Nelder-Mead	Genetic Algorithm	Nelder-Mead	Genetic Algorithm	Nelder-Mead	Genetic Algorithm	Nelder-Mead	Genetic Algorithm
Log-likelihood	7768	7788	7953	7968	7968	8000	7889	7983
RMSE (*)								
1M	17.9	17.6	16.3	15.8	15.5	15.2	16.1	15.0
3M	14.5	14.1	11.2	10.9	10.0	9.9	10.9	9.7
6M	13.1	12.9	8.2	8.1	6.1	6.4	7.4	6.4
1Y	15.9	15.9	11.5	11.7	10.1	10.4	10.8	10.4
2Y	18.1	18.1	15.6	15.7	15.1	15.0	15.4	14.9
3Y	17.0	16.9	15.6	15.6	15.5	15.4	16.2	15.2
5Y	12.0	11.6	10.7	10.6	10.9	10.7	12.3	10.4
7Y	8.8	8.1	6.5	6.3	6.6	6.0	8.1	5.3
10Y	11.9	11.2	12.5	12.6	13.2	13.2	14.5	13.5

(*) Values expressed in basis points (bps)   Source: ECFIN

The results shown under the column heading "Nelder-Mead" use as a starting point for the parameters to be calibrated an already optimised set of values for the euro area[15]. The results shown under the column heading "Genetic Algorithm" use an arbitrary starting point value which assumes no prior knowledge of the dataset to which the shadow rate model is being calibrated. In fact, we have set the starting value for all parameters as being equal to 0.001 except for the 2x2 matrix k (which represents the mean-reversion of state variables – see Section 8 for a more technical description). For the mean-reversion matrix k, we randomly generated a positive-definite matrix, which in our case gave us k $= \begin{bmatrix} 0.14842 & 0.09308 \\ 0.09308 & 0.08504 \end{bmatrix}$.

Overall, the genetic algorithm provides a slightly superior calibration across all scenarios for the lower bound as the value of the optimised log-likelihood is higher in all cases. The value of the optimised log-likelihood function is the yardstick commonly used to ascertain the quality of the calibration; in other words, the higher the log-likelihood value resulting from the optimisation algorithm, the better the model captures the movements of the yield curve over time (in addition to its fit at any given point in time). In addition, the average quality of the fit to the observed yield curve can be measured by the root mean-squared error (RMSE) statistic. As Table 1 clearly illustrates, both the Nelder-Mead and the genetic calibration algorithms provide a similar fit overall (with a slightly better fit obtained nevertheless by the genetic algorithm across the different lower bounds and yield curve tenors). Thus, to summarise, the alternative calibration technique (genetic algorithms) tested in this paper corroborates overall the empirical results obtained by Krippner and at the same time also confirms Krippner's stated advantage of using the iterated extended Kalman filter (Krippner 2015a).

Having demonstrated the relative convergence of both calibration algorithms, we can safely conclude that the highly positive value obtained for the option effect (and thus the highly negative shadow rate)

---

[15] The starting values come from the already optimised set of parameters for the euro area (daily dataset) as published on Krippner's website: EA_BB_Govt_BB_OIS_rL125_Daily_20160429.

when the lower bound is set to zero is not linked to the calibration of the model but rather to intrinsic economic phenomena being captured by the cash option effect (see Section 5). This being said, it is nevertheless worthwhile diving a bit more deeply into the pros and cons of calibrating the shadow rate model using genetic algorithms versus the traditional Nelder-Mead algorithm. While the Nelder-Mead algorithm is known to reduce the computational burden as it converges quickly to the local optimum solution, it is in principle prone to estimation errors as it often tends to result in local optimum solutions and is thus relatively more sensitive to the starting point assumptions. The genetic algorithm, by contrast, searches the entire solution space testing far out points to exclude that the global optimum maybe lying somewhere in a far out corner of the complex solution space. Therefore it also guarantees that the end result is not sensitive to the choice of the starting parameters set prior to launching the calibration algorithm. However, there are two documented disadvantages to using genetic algorithms. The first weakness relates to the fact that although they are efficient in searching a complex, multi-dimensional solution space to find the region near the most likely global optimum point, once they reach this region they have difficulties in converging to the exact final global optimal point within this sub-region. Second, genetic algorithms present a significant computational burden and thus are very time consuming to run.

To overcome these shortcomings, while keeping the favourable properties of both algorithms, we decided to embed a hybrid scheme into the genetic algorithm which works as follows: the genetic algorithm first searches through the entire solution space (even in far out regions of the solution space) to seek the region which should contain the global optimum (thereby testing and excluding possible distant local optimum points); once the genetic algorithm reaches the region near the global optimum point, we switch to the Nelder-Mead algorithm which converges more quickly to the optimum solution. Thus, while making use of the Nelder-Mead algorithm's faster convergence, we also eliminate its main weakness, i.e. its sensitivity to the starting point assumptions. To confirm this, we also tested the calibration of the K-ANSM (2) model using the traditional Nelder-Mead algorithm and also using the same "arbitrary" starting point values as for the genetic algorithm described previously; the calibration results turned out significantly poorer this time: the values of the log-likelihood function obtained were 7707, 7223, 6901, 6857 respectively for the various lower bounds 0%, -0.20%, -0.40%, and -0.60%. For these same "arbitrary" starting values, the hybrid genetic algorithm converged well to obtain a good calibration fit and relatively high log-likelihood values (as demonstrated in Table 1). Hence, although our tests confirm that the IEKF outperforms the EKF by providing much more reliable and accurate parameter estimations from different starting points (compensating for the fact that the Nelder-Mead optimisation routine is known to often result in a "local" optimum), our tests have also shown that there are limits to this compensation. This is where using a hybrid-genetic algorithm approach may prove to be very useful: in cases where one needs to calibrate the shadow rate model to new datasets (new regions, other countries or yield curves etc …) when the starting point assumptions are unknown or in the presence of substantial market events and movements. In such cases, one only needs to set arbitrary starting values for the parameters to be optimised and the hybrid-genetic algorithm should converge to obtain a good fit to the market dataset. Unfortunately, the computational burden of the hybrid-genetic algorithm still remains high compared to the simple Nelder-Mead algorithm.

Finally, another useful and complementary way of establishing the reliability of calculated model parameters and the resulting estimated level of the shadow rate is to look at the size of the confidence intervals produced by the model. Graph 10 below illustrates the 95% confidence intervals for the estimated shadow rate when calibrated to the euro area using the hybrid-genetic algorithm and a lower bound of -0.40%. As the graph clearly illustrates, the confidence band around the estimated shadow rate is quite narrow in the period from Dec-02 to Dec-15; afterwards, it starts to widen as the shadow rate dips more deeply into negative territory, i.e. when the cash option's value becomes more positive and its effect on the shadow rate relatively more important. In other words, the fact that the cash option that has to be taken into account in the measurement equation of the Kalman filter adds another layer of model complexity (and non-linearity) and thus broadens the confidence band. Nonetheless, overall, it is reassuring to see that the shadow rate estimates obtained from running the hybrid-genetic

calibration algorithm and using a lower bound of -0.40% (reflecting the minimum observed value of the DFR in the calibration sample history) lie within relatively tight confidence bands.

Graph 10   Hybrid GA K-ANSM(2) calibrated to the EA (lower bound =-0.40%)
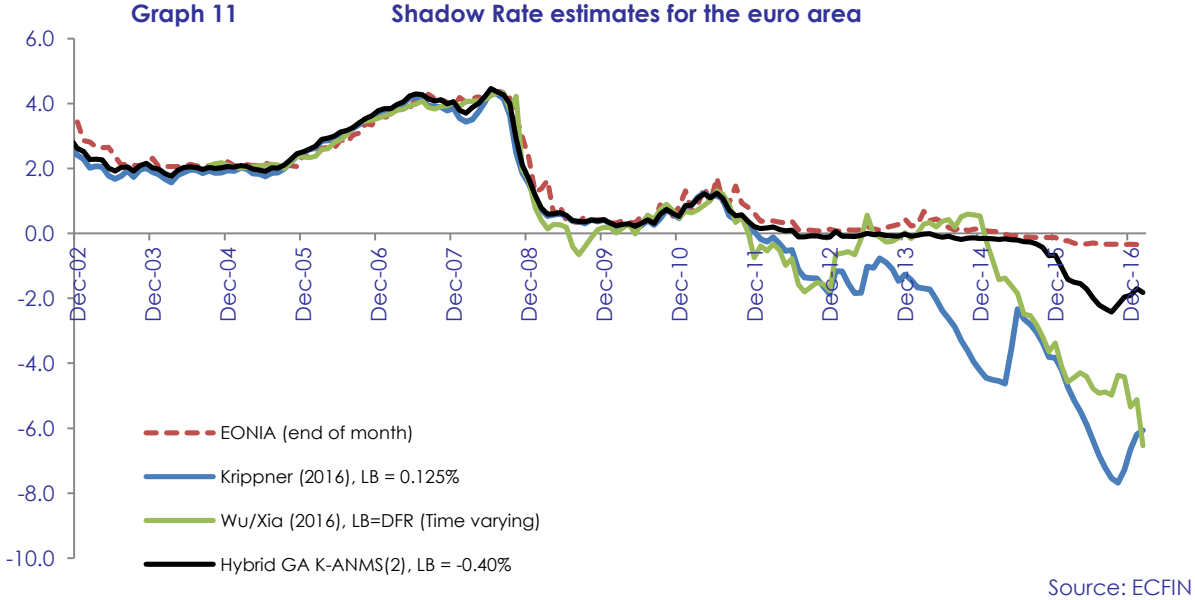
Source: ECFIN
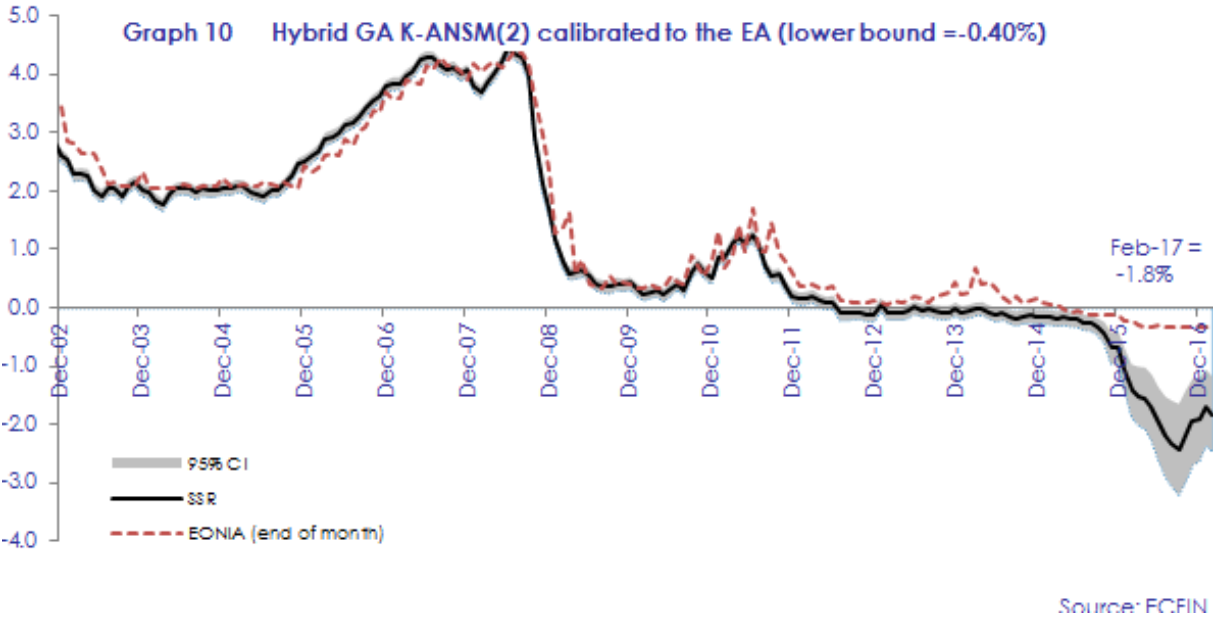
5. THE SHADOW RATE AS A MEASURE OF THE MONETARY POLICY STANCE

In the presence of the lower bound of interest rates, central banks have increasingly turned to unconventional monetary policies targeting the long end of the term structure. Shadow rates derived from term-structure models can capture those effects along the yield curve and try to quantify the monetary stimulus implied by measures beyond variations of the policy rates as they boil down monetary policy actions into one estimated number. Generally, this intuitiveness also represents a major advantage of the SSR concept. The shadow rate can also be used as a time series in further analysis instead of the restricted policy rate.[16] Nonetheless, the fact that the SSR is a purely hypothetical rate implies limits for its use in economic inference, as for example a decline in the shadow rate has no obvious impact on the short-term financing conditions of enterprises, which still face the lower bound constrained rates.

The most important drawback of the shadow rate, however, is its sensitivity to the chosen model-specification. For illustrative purposes, Graph 11 plots shadow rate estimates for the euro area by Wu and Xia (2016), who use a three-factor model and a time-varying lower bound that follows the ECB's deposit facility rate, and from the two-factor model by Krippner (2016) (with a lower bound of 0.125%) as well as from our two-factor Hybrid-GA model (with an assumed lower bound of -0.40%). The results vary considerably, in particular depending on the number of factors included in the model and the assumed or endogenously estimated lower bound. Further discrepancies can arise from the estimation method, the maturity spectrum included in the estimations and the yield curve data to which

---

[16] To this end, Wu and Xia (2014) estimate a monetary policy VAR in which they replace the Federal Funds rate in the years that the lower bound was binding with their SSR estimates and show, using impulse response analyses, that the effects of the shadow rate on macroeconomic variables are similar to those of the federal funds rate.

the model is calibrated.[17] Using three factors, as is conventional in the term-stucture literature, allows for more flexibility in the model and thus leads to a better model-fit for the short- and mid-term maturity spectrum, thereby improving the overall model-fit. However, this might impair the shadow rate's ability to capture unconventional monetary policies that are predominantly targeted at the long end of the yield curve. To this end, Krippner (2015b) compares his model specification with the Wu/Xia (2015) shadow rate for the United States and concludes that the two-factor model's dynamics are more consistent with the Federal Reserve's unconventional monetary policies.

Graph 11   **Shadow Rate estimates for the euro area**

- - - EONIA (end of month)
——— Krippner (2016), LB = 0.125%
——— Wu/Xia (2016), LB=DFR (Time varying)
——— Hybrid GA K-ANMS(2), LB = -0.40%

Source: ECFIN

The second major source of sensitivity is the choice of the lower bound. To illustrate this, Graph 12 plots again the shadow rate estimates derived from the hybrid GA model for a range of lower bounds between 0% and -0.60%. Evidently, the higher the lower bound is set, the more negative does the shadow rate get in response to unconventional monetary policies in times when the lower bound is binding. That is, the higher the lower bound, the more probability is attributed cet. par. to an outcome below the lower bound, which increases the value of the cash option. Equivalently, a higher lower bound assigns more weight to movements at the long end of the term structure, as a larger portion of implied forward rates at the short end (i.e. below the specified lower bound) is ignored in the estimation process.[18] This means at the same time that a high lower bound might ignore monetary policy induced movements at the short end of the yield curve. The evident question is thus whether there is an optimal choice for the lower bound parameter.

The answer to this question depends to a large extent on which interpretation we want to give the lower bound and thus our understanding of the shadow rate. In its original *economic interpretation,* we assume that the lower bound represents the level of interest rates at which investors switch to holding cash (i.e. exercising the cash option). This is then also consistent with the shadow rate being interpreted as the rate that would materialise in the absence of cash. The 'extreme' choice of a lower bound of -0.60% as shown in Graph 12, which as explained above shows little reaction to unconventional monetary policies but stays at more economically reasonable levels, can be justified

---

[17] For instance, Krippner uses zero coupon government rates spliced with OIS rates once they are available, while Wu and Xia are calibrating their model to forward rates. Furthermore, Krippner uses for his estimations daily yield curve data and also includes the 30 year maturity (as opposed to the maximum 10 year maturity used by Wu/Xia and ourselves).

[18] That is, the measurement equation errors at the short end of the yield curve (where the lower bound applies) are not taken into account. The estimation then maximises the model-fit to the remaining longer maturities.

economically: it represents the minimum of the *observed and expected* (i.e. based on forward rates) OIS spot rates as derived from OIS curves. While Eonia has stayed at above -0.40% throughout the sample horizon, the inverted OIS swap curve up to the 3 year maturity in July 2016 implies expected spot rates that were considerably below. The argument then can be made that markets' expectation of a lower OIS rate in the future also implies their willingness to participate in the market at such rates, which means they don't exercise the cash option that constitutes the lower bound at this rate. Thus, the effective lower bound of interest rates must be at or below that minimum point. In that sense, the chosen lower bound represents an upper limit to the 'true' lower bound, which might be even more negative. As mentioned above, while this lower bound calibration yields shadow rate estimates whose levels seem more reasonable from a historical perspective, they show little correlation with the ECB's unconventional monetary policies, as illustrated in Graph 12. Consequently, the estimated shadow rate with a lower bound at -0.60% by definition hardly deviates from EONIA, as a very low cumulated probability is assigned to OIS rates falling below the lower bound such that the cash option value is diminished considerably.

Alternatively, and more prominently in the literature, one can understand the lower bound rather as an econometric parameter choice to provide the desired features of the model. For example, if the goal is to map unconventional monetary policies that affect the long end of the yield curve, a higher lower bound than what can be justified under a purely economic interpretation might be desireable. As Graph 6 illustrates, setting the lower bound to -0.20% or even at 0% – both materially above realised OIS spot rates over the sample horizon – produces shadow rate estimates that correlate highly with the ECB's unconventional monetary policy measures. Furthermore, standardisation of those estimates can still provide a useful measure for a monetary policy variable in time series models.[19] However, at the same time, those shadow rates are also more prone to pick up other movements in long-term rates that are not directly related to the monetary policy stance. For instance, lower long-term growth- and inflation expectations should lead to a flattening of the yield curve and thus to a decrease in the shadow rate, signaling a further easing of the monetary policy stance when measured by the shadow rate. While the decrease at the long end of the term structure also reflects expectations of ensuing lower future policy rates, in our opinion it can hardly be interpreted as part of the monetary policy stance in the sense of an active steering of interest rate expectations. The uptick of the shadow rate in the second quarter of 2015, just after the start of the public sector asset purchase programme (PSPP), for instance seems to be driven by increasing long-term inflation expectations on the back of increasing oil prices with realised inflation also picking up. Similarly, the increasing shadow rates post September 2016 may reflect to a large exent a higher growth and inflation outlook. Macroeconomic developments thus introduce some noise in the time series that does not necessarily reflect the ECB's monetary policy measures.

---

[19] As the two-factor shadow rate estimates show material variations in magnitude but are similar with respect to their profiles and dynamics, they could be scaled, e.g. to match the Taylor rule during the unconventional monetary policy period, in order to obtain an estimate that could be used as a policy rate analogue in unconventional monetary policy periods.

Graph 12   EA Shadow Rate and ECB Monetary Policy Events

Source: ECFIN
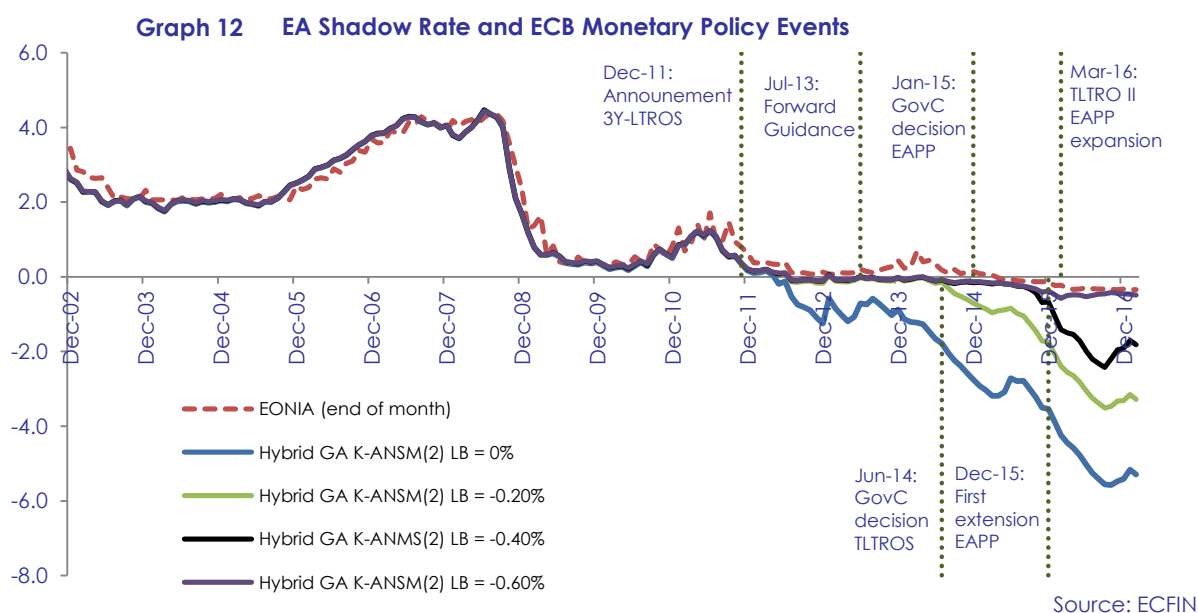
Finally, an elegant way to ensure that the shadow rates capture unconventional monetary policies while avoiding overly unrealistic lower bound specifications from an economic perspective is to introduce a time-varying lower bound, e.g. by setting it to zero up to July 2012 and to the level of the deposit facility rate thereafter.[20] This way, the lower bound stays close to realised rates even when they are not in deeply negative territory, thereby ensuring that the cash option value is non-zero, while avoiding that realised spot rates fall below the lower bound at any time. However, the idea of a time-varying lower bound is rather econometrically than theoretically motivated, as it is not clear why economic agents' lower bound should change with changes in monetary policy.

All in all, there is a trade-off in the calibration of shadow rate models between an economically reasonable interpretation and the usefulness of the shadow rate estimates as a stance indicator. While shadow rates can thus be a useful indicator of the monetary policy stance, it should not be confused for a proxy for the policy rate. In light of the model-sensitivity to the parameter choices, other outputs from shadow rate models have been proposed as stance indicators, such as the Effective Monetary Stimulus (EMS) measure, which sets the expected path of short term rates embedded in the yield curve in relation to an estimated long-term neutral rate, or lift-off measures, which measure the expected time until the policy rate crosses the lower bound.[21] While more abstract than the shadow short rate, both measures have the advantage that they are more robust to different model specifications as shadow yield curves tend to be more similar across model specifications (Bauer and Rudebusch, 2013).

---

[20] For example, Kortela (2016) and Lemke and Vladu (2016) use time-varying lower bounds.

[21] See for example Halberstadt and Krippner (2016). However, the lift-off measure is only defined for an unconventional monetary policy framework, i.e. it does not exist when the lower bound is not binding. It can thus not provide a continuous measure of the monetary policy stance over time.

6. CONCLUSION

The convergence of monetary policy rates towards a lower bound and the introduction of unconventional monetary policy measures targeting the long end of the yield curve have impaired the function of policy rates as a summary indicator for a central bank's policy stance. In this context, shadow short rates derived from term structure models have been proposed as a stance measure that captures movements along the yield curve and boils it down into a single indicator. One such shadow rate model that has received much attention in the literature is the two factor shadow rate model by Krippner, which we draw on in this paper.

In this paper, we test an alternative method to calibrate shadow rate models to market data. More specifically, we use hybrid-genetic algorithms to calibrate Krippner's two factor shadow rate model to the euro area OIS swap curve, showing that overall the results obtained corroborate those of Krippner (with even slightly improved calibration results as compared to those obtained when using Krippner's already optimised starting values for the parameters). We also confirm Krippner's finding that the IEKF outperforms the EKF by providing much more reliable and accurate parameter estimations from different starting points, which compensates for the fact that the Nelder-Mead optimisation routine is known to often result in a "local" optimum. But as our tests have also shown, this compensation is not perfect and in the case where the starting values of the parameters are set arbitrarily, the calibration quality obtained using hybrid-genetic algorithms was higher than the one obtained using the Nelder-Mead algorithm for those same arbitrary starting values. Thus, one noticeable advantage of using this hybrid-genetic algorithm approach is that it may prove itself to be very useful in cases where one needs to calibrate the shadow rate model to new datasets (new regions, other countries or yield curves etc …) when the starting point assumptions are unknown or in the presence of substantial market events and movements.

While employing genetic algorithms can thus improve the overall calibration quality, it does not eradicate the general caveats associated with the shadow rate concept such as its overall high uncertainty arising from its sensitivity to the chosen model specification. Estimation results vary considerably, in particular depending on the number of factors included in the model and the assumed or endogenously estimated lower bound. While the case can be made for choosing two factors over three in order to better reflect unconventional monetary policy measures, it is not obvious which level of the lower bound should be chosen, as there is a trade-off between an economically reasonable interpretation of the shadow rate and its usefulness as a stance indicator at the lower bound. While shadow rates can thus be a useful indicator of the monetary policy stance, the uncertainy regarding their overall level and the multitude of measures they comprise warrant a careful interpretation of the estimation results.

7. TECHNICAL ANNEX

7.1   A TERM STRUCTURE-BASED SHADOW RATE

Krippner uses a popular subclass of GATSMs (Gaussian affine term structure) models called arbitrage-free Nelson and Siegel (1987) models (ANSMs) to represent the shadow term structure. An affine term structure model is a financial model that relates zero-coupon bond prices (i.e. the discount curve) to a model for short rate. It is particularly useful for deriving the zero-coupon yield curve from quoted bond prices. Thus the starting point for the development of the affine class is the postulation of a stochastic process for the short rate and the related state variables, or factors, which drive the dynamics of the term structure. These factors are the underlying source of uncertainty in the model of the term structure.Under a GATSM, the short rate $r(t)$ at time t is a linear function of the state variables $x(t)$ at time t:

$$r(t) = a_0 + b_0' x_n(t)$$

where r(t) is a scalar, a_0 is a constant scalar, b_0 is a constant Nx1 vector containing the weights for the N state variables $x_n(t)$. Under the risk-neutral $\mathbb{Q}$-measure, x(t) evolves as a correlated vector Ornstein-Uhlenbeck process[22]:

$$dx(t) = \tilde{k}[\tilde{\theta} - x(t)]dt + \sigma d\widetilde{W}(t)$$

where $\tilde{\theta}$ is a constant Nx1 vector representing the long-run level of x(t), $\tilde{k}$ is a constant NxN matrix that governs the deterministic mean reversion of x(t) to $\tilde{\theta}$, σ is a constant NxN matrix representing the correlated variance of innovations to x(t), and $d\widetilde{W}(t)$ is an Nx1 vector with independent risk-neutral Wiener components. The main advantages working with continuous-time specifications are that they are more amenable to mathematical manipulation, and they lead to closed-form analytic solutions for all representations of the term structure.

The dynamics for the state variables vector x (t) are given by the solution to the above stochastic differential equation [for a more complete discussion see Krippner 2015, "Zero Lower Bound Term Structure Modeling", p48]. The resulting closed-form expressions for short rate and the zero-coupon interest rates in the particular case of the two-factor arbitrage-free Nelson Siegel model, ANSM (2), are presented below:

$$r(t) = x_1(t) + x_2(t)$$

$$R(t,\tau) = a(\tau) + [b(\tau)]' \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}$$

where $r(t)$ is the short rate and $R(t,\tau)$ is the zero-coupon interest rate for time to maturity τ.

The expressions $a(\tau)$ and $b(\tau)$ are themselves functions of τ and of the parameters $\tilde{k}$, $\tilde{\theta}, \sigma$ which define the stochastic process for the short rate [see Krippner 2015, "Zero Lower Bound Term Structure Modeling", p65]:

$$a(\tau) = -\sigma_1^2 \cdot \frac{1}{6}\tau^2 - \sigma_2^2 \cdot \frac{1}{2\phi^2}\left[1 - \frac{1}{\tau}G(\phi,\tau) - \frac{1}{2\tau}\phi[G(\phi,\tau)]^2\right] - \rho\sigma_1\sigma_2$$
$$\cdot \frac{1}{\phi^2}\left[1 - \frac{1}{\tau}G(\phi,\tau) + \frac{1}{2}\phi\tau - \phi G(\phi,\tau)\right]$$

$$[b(\tau)]' = \left[1, \frac{1}{\tau}G(\phi,\tau)\right]$$

where $G(\phi,\tau) = \frac{1}{\phi}[1 - exp(-\phi\tau)]$. A key aspect is that the interest rate factor loadings are defined by $[b(\tau)]'$ and this particular loading structure implies that the first factor is responsible for parallel yield curve shifts, since the effect of this factor is identical for all maturities; the second factor represents slope parameter of the yield curve slope. The key step to go from a yield curve model to the shadow rate model is to derive the closed-form analytic expression for the forward rate option effect z(t, τ).

Because the shadow rate framework uses a GATSM specification for shadow short rates, the dynamics of r(t + τ), the short rate in τ years from now, in terms of the state variables x(t), also define the dynamics for the option component $\max\{r_L - r(t + \tau), 0\}$. In fact, the ZLB forward rate $\underline{f}(t,\tau)$ is obtained by calculating the expectation of ZLB short rates $\underline{r}(t + \tau)$ under the t + τ forward $\mathbb{Q}$-measure:

$$\underline{f}(t,\tau) = \tilde{E}_{t+\tau}\left[\underline{r}(t+\tau)|x(t)\right]$$

$$= \tilde{E}_{t+\tau}\left[r(t+\tau)|x(t)\right] + \tilde{E}_{t+\tau}[\max\{r_L - r(t+\tau), 0\}|x(t)]$$

---

[22] In mathematics, the Ornstein–Uhlenbeck process is a stochastic process that can be considered to be a modification of the random walk in continuous time in which the properties of the process have been changed so that there is a tendency of the walk to drift towards its long-term mean, with a greater mean-reversion when the process has drifted further away from it's long-term mean.

$$= f(t, \tau) + z(t, \tau)$$

where $f(t, \tau)$ is the shadow forward rate and $z(t, \tau)$ is the forward option effect. The forward option effect is derived by evaluating directly the expectation $\tilde{E}_{t+\tau}[\max\{r_L - r(t + \tau), 0\}|x(t)]$ which results in the following analytical expression for the forward option effect [see Krippner 2015, "Zero Lower Bound Term Structure Modeling", p111]:

$$z(t, \tau) = [r_L - f(t, \tau)] \cdot \left(1 - \Phi\left[\frac{f(t, \tau) - r_L}{\omega(\tau)}\right]\right) + \omega(t) \cdot \emptyset\left[\frac{f(t, \tau) - r_L}{\omega(\tau)}\right]$$

where $\Phi(\cdot)$ is the cumulative unit normal density function, $\emptyset(\cdot)$ is the unit normal probability density function, and $\omega(\tau)$ is the shadow short rate standard deviation. Substituting the result for $z(t, \tau)$ above into the expression $\underline{f}(t, \tau) = f(t, \tau) + z(t, \tau)$ yields the ZLB forward rate expressed in terms of the shadow forward rate and the option effect:

$$\underline{f}(t, \tau) = f(t, \tau) + z(t, \tau)$$

$$= r_L + [f(t, \tau) - r_L] \cdot \Phi\left[\frac{f(t, \tau) - r_L}{\omega(\tau)}\right] + \omega(\tau) \cdot \emptyset\left[\frac{f(t, \tau) - r_L}{\omega(\tau)}\right]$$

In the particular case of the two-factor shadow rate model, K-ANSM (2), the ZLB forward rate as derived from the above theoretical foundations has the following closed-form analytic form [see Krippner 2015, "Zero Lower Bound Term Structure Modeling", p129]:

$$\underline{f}(t, \tau) = r_L + [f(t, \tau) - r_L] \cdot \Phi\left[\frac{f(t,\tau) - r_L}{\omega(\tau)}\right]$$

$$+ \omega(\tau) \cdot \frac{1}{\sqrt{2\pi}} \cdot \exp\left(-\frac{1}{2}\left[\frac{f(t, \tau) - r_L}{\omega(\tau)}\right]^2\right)$$

Furthermore, in the particular case of the two-factor model, the shadow short rate volatility function is derived to be:

$$\omega(\tau) = \sqrt{\sigma_1^2 \cdot \tau + \sigma_2^2 \cdot G(2\phi, \tau) + 2\rho_{12}\sigma_1\sigma_2 \cdot G(\phi, \tau)}$$

where $G(\phi, \tau) = \frac{1}{\phi}[1 - exp(-\phi\tau)]$.

In a final step, ZLB zero-coupon interest rates are derived from forward rates using the usual term structure relationship:

$$\underline{R}(t, \tau) = \frac{1}{\tau}\int_0^\tau \underline{f}(t, u)du$$

where $\underline{R}(t, \tau)$ is the ZLB zero-coupon interest rate for time to maturity τ.

The integral does not have a closed-form analytic solution, because $\underline{f}(t, u)$ contains the cumulative Gaussian distribution, but univariate numerical integration of $\underline{f}(t, u)$ over time to maturity τ may be used to calculate the integral to arbitrary precision. Now that the theoretical foundations for the shadow rate model have been established, the next section will discuss how to calibrate the shadow rate to market data.

7.2   CALIBRATING SHADOW RATE MODELS

The two-factor shadow rate model, K-ANSM (2), has the following 11parameters to estimate, of which the first 10 originate from the definition of the stochastic process for the short rate, that is

$$\text{Parameter set B} = \{\phi, k_{11}, k_{12}, k_{21}, k_{22}, \theta_1, \theta_2, \sigma_1, \sigma_2, \rho_{12}, \sigma_\eta\}$$

plus the variable σ_η which represents the measurement equation (Kalman filter) standard deviation.

The state variables and the parameters when expressed in matrix form can be linked directly to the expression defining the stochastic process for the short rate (a correlated vector Ornstein-Uhlenbeck process):

$$x(t) = \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}; \ a_0 = 0; \ b_0 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}; \ k = \begin{bmatrix} k_{11} & k_{12} \\ k_{21} & k_{22} \end{bmatrix}; \ \theta = \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix}$$

$$\sigma = \begin{bmatrix} \sigma_1 & 0 \\ \rho_{12}\sigma_2 & \sigma_2\sqrt{1-\rho_{12}^2} \end{bmatrix}; \ \tilde{k} = \begin{bmatrix} 0 & 0 \\ 0 & \phi \end{bmatrix}; \ \tilde{\theta} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Calibrating the K-ANSM (2) involves embedding the Kalman filter into an optimisation algorithmso as to estimate the parameters for the specified model, and the state variables associated with those parameters and calculated by the Kalman filter are also an output of the optimisation. The Kalman filter is based on a state equation, which specifies how the state variables evolve over time, and a measurement equation, which specifies how the state variables explain the observed data at each point in time. In our particular case, the state variables are in the vector x (t), the measurement equation is the GATSM yield curve expression as a function of x (t), and the data is the observed yield curve data at each point in time. The objective function of the optimisation algorithm is to maximise the log-likelihood function given by the expression

$$Log\_LKL\big(B, \sigma_\eta, \{ZC_1, \dots, ZC_T\}\big) = -\frac{1}{2}\sum_{t=1}^{T}[K \cdot log(2\pi) + log(|\mathcal{M}_t|) + \eta_t'\mathcal{M}_t^{-1}\eta_t]$$

where $\{\eta_1 \dots \eta_T\}$ is the time series of Kx1 vectors containing the unexplained component of the yield curve data at time t relative to the K-ANSM (2) model (obtained using the measurement equation) and $\{\mathcal{M}_1 \dots \mathcal{M}_T\}$ is the time series of KxK matrices obtained at each time step of the Kalman filter algorithm. The constant K refers to the number of yield curve tenors used in the calibration sample, which in our case refers to the following yield curve 9 tenors: 1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y.

The addition of the "option" component introduces non-linearity in the Kalman filter's measurement equations, and hence the calibration of the K-ANSM (2) model requires a nonlinear Kalman filter: Krippner opted for the iterated extended Kalman filter (IEKF). The IEKF effectively calculates a first-order Taylor approximation[23] of the non-linear interest rate function around the best available estimate of the state variable vector. The use of the IEKF allows for the non-linearity of $\underline{R}(t, \tau)$ with respect to the state variables. The state equation for the K-ANSM (2) is a first-order vector autoregression of the following form:

$$x_t = \theta + exp(-k\Delta t)(x_{t-1} - \theta) + \varepsilon_t$$

where the subscripts t are an integer index to represent the progression of time in steps of Δt between observations (in the case of monthly data $\Delta t = 1/12$), $exp(-k\Delta t)$ is the matrix exponential of $-k\Delta t$, and ε_t is the vector of innovations to the state variables. The variance of ε_t is:

$$var(\varepsilon_t) = \int_0^{\Delta t} exp(-ku)\sigma\sigma'exp(-k'u)\, du$$

which is a 2x2 matrix. The measurement equation for the K-ANSM (2) is given by:

$$\begin{bmatrix} R_t(\tau_1) \\ \vdots \\ R_t(\tau_K) \end{bmatrix} = \begin{bmatrix} \underline{R}(x_t, \tau_1, B) \\ \vdots \\ \underline{R}(x_t, \tau_K, B) \end{bmatrix} + \begin{bmatrix} \eta_t(\tau_1) \\ \vdots \\ \eta_t(\tau_K) \end{bmatrix}$$

---

[23] In mathematics, a Taylor series is a representation of a function as an infinite sum of terms that are calculated from the values of the function's derivatives at a single point.

where K is the index for the yield curve tenor τ_K, $R_t(\tau_K)$ is the observed interest rate at time t for time to maturity τ_K, $\underline{R}(x_t, \tau_K, B)$ are the the zero-coupon model-based interest rate functions evaluated at τ_K, and $\eta_t(\tau_K)$ is the component of $R_t(\tau_K)$ that is unexplained by the K-ANSM(2) model. The variance of η_t is specified to be homoskedastic and diagonal:

$$\Omega_\eta = diag\left[\{\sigma_\eta^2, \dots, \sigma_\eta^2\}\right]$$

where Ω_η is a KxK matrix with entries σ_η^2. Furthermore, reflecting standard practice, the vectors η_t and ε_t are assumed to be uncorrelated over time.

In a previous version of Krippner's K-ANSM (2) model (prior to July-2016), the variance of η_t was specified as heteroskedastic and diagonal which translates to:

$$\Omega_\eta = diag\left[\{\sigma_\eta^2(\tau_1), \dots, \sigma_\eta^2(\tau_K)\}\right]$$

That specification led to larger variances in the residuals for the short and long-maturity data, which had the practical effect of a less close fit to the short-maturity data and more volatile shadow rate estimates in both the non ZLB and ZLB periods. The homoskedastic specification enforces similar sized residuals across the yield curve data, which results in less volatile shadow rate estimates.

7.3   THE OPTIMISATION ALGORITHM

A very commonly used derivative-free algorithm to maximise the log-likelihood function is the Nelder-Mead simplex search method of Lagarias et al. This is a direct search method that does not use numerical or analytic gradients (FMINSEARCH function in MATLAB). It is also the method employed by Krippner to calibrate the shadow rate model (minimising the negative of the log-likelihood function). In a nutshell, it works as follows: consider the function $f(x)$ to be minimised where $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is called the objective function and n the dimention. A simplex is a geometric figure in n dimensions which is the convex hull of $n+1$ vertices (denote the simplex with vertices $x_1, \dots, x_{n+1}$ by Δ). The method iteratively generates a sequence of simplices to seek the minimum point and at each iteration the vertices are ordered according to the objective function values $f(x_1) \leq f(x_2) \leq \cdots \leq f(x_{n+1})$. We refer to x_1 as the best vertex and to x_{n+1} as the worst vertex. The algorithm uses four possible operations: reflection α, expansion β, contraction γ and shrink δ, each being associated with a scalar parameter. Let $\bar{x}$ be the centroid of the n best vertices, which implies that $\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$. Any one given iteration of the Nelder-Mead algorithm follows the following scheme:

1. Sort: evaluate f at the $n+1$ vertices of Δ and sort the vertices (as described above).

2. Reflection: compute the reflection point x_r from $x_r = \bar{x} + \alpha(\bar{x} - x_{n+1})$. Then evaluate $f_r = f(x_r)$. If $f_1 \leq f_r < f_n$ then replace x_{n+1} with x_r.

3. Expansion: if $f_r < f_1$ then compute the expansion point x_e from $x_e = \bar{x} + \beta(x_r - \bar{x})$ and evaluate $f_e = f(x_e)$. If $f_e < f_r$ then replace x_{n+1} with x_e, otherwise replace x_{n+1} with x_r.

4. Outside contraction: if $f_n \leq f_r < f_{n+1}$ then compute the outside contraction point $x_{oc} = \bar{x} + \gamma(x_r - \bar{x})$ and evaluate $f_{oc} = f(x_{oc})$. If $f_{oc} \leq f_r$ then replace x_{n+1} with x_{oc}, otherwise go to step 6.

5. Inside contraction: if $f_r \geq f_{n+1}$ then compute the inside contraction point $x_{ic} = \bar{x} - \gamma(x_r - \bar{x})$ and evaluate $f_{ic} = f(x_{ic})$. If $f_{ic} < f_{n+1}$ then replace x_{n+1} with x_{ic}, otherwise go to step 6.

6. Shrink: for $2 \leq i \leq n+1$ then define $x_i = x_1 + \delta(x_i - x_1)$. Replace all points except the first/best x_1 and go to step 1.

This scheme is repeated until the diameter of the simplex is less than the specified tolerance.

8. REFERENCES

Bauer, M. and G. Rudebusch (2013). Monetary policy expectations at the zero lower bound. Working Paper, Federal Reserve Bank of San Francisco 18.

Black, F. (1995). Interest rates as options. Journal of Finance 50(7), 1371-1376.

Grewal, M. and A. Andrews (2008). Kalman Filtering, Third Edition. Wiley and Sons.

Gürkaynak, R., B. Sack, and J. Wright (2007). The U.S. Treasury yield curve: 1961 to the present. Journal of Monetary Economics 54(8), 2291-2304.

Halberstadt, A. and L. Krippner (2016). The effect of conventional and unconventional euro area monetary policy on macroeconomic variables. Deutsche Bundesbank Discussion Paper 49/2016.

Kortela, T. (2016). A shadow rate model with time-varying lower bound of interest rates. Bank of Finland Research Discussion Paper, 19.

Krippner, L. (2011). Modifying Gaussian term structure models when interest rates are near the zero lower bound. Discussion paper, Centre for Applied Macroeconomic Analysis 36/2011.

Krippner, L. (2012a). Measuring the stance of monetary policy in zero lower bound environments. Discussion Paper, Reserve Bank of New Zealand DP 2012/04.

Krippner, L. (2012b). Modifying Gaussian term structure models when interest rates are near the zero lower bound. Discussion paper, Centre for Applied Macroeconomic Analysis 5/2012.

Krippner, L. (2012c). Modifying Gaussian term structure models when interest rates are near the zero lower bound. Discussion Paper, Reserve Bank of New Zealand DP 2012/02.

Krippner, L. (2012d). A theoretical foundation for the Nelson and Siegel class of yield curve models. Discussion Paper, Centre for Applied Macreconomic Analysis 11/2012.

Krippner, L. (2013a). Faster solutions for Black zero lower bound termstructure models. Working Paper, Centre for Applied Macroeconomic Analyis 66/2013.

Krippner, L. (2013b). Measuring the stance of monetary policy in zero lower bound environments. Economics Letters 118(1), 135-138.

Krippner, L. (2013c). A tractable framework for zero-lower-bound Gaussian term structure models. Discussion Paper, Centre for Applied Macreconomic Analysis 49/2013.

Krippner, L. (2013d). A tractable framework for zero-lower-bound Gaussian term structure models. Discussion Paper, Reserve Bank of New Zealand DP 2013/02.

Krippner, L. (2014b). Measuring the stance of monetary policy in conventional and unconventional environments. Working Paper, Centre for Applied Macroeconomic Analysis 6/2014.

Krippner, L. (2014c). Measuring the stance of monetary policy in conventional and unconventional environments. Working Paper. http://conference.nber.org/confer/2014/EASE14/Krippner.pdf.

Krippner, L. (2014d). A theoretical foundation for the Nelson-Siegel class of yield curve models. Journal of Applied Econometrics (forthcoming).

Krippner, L. (2015a). Term Structure Modeling at the Zero Lower Bound: A Practitioner's Guide. Palgrave-Macmillan.

Krippner, L. (2015b). A comment on Wu and Xia (2015), and the case for two-factor Shadow Short Rates. CAMA Working Paper 48/2015

Lagarias, J., J. Reeds, M. Wright, and P. Wright (1998). Convergence properties of the Nelder–Mead simplex method in low dimensions. SIAM Journal on Optimisation, 9(1):112–147, 1998.

Lefebvre, T., H. Bruyninckx, and J. De Schutter (2004). Kalman filters for nonlinear systems: A comparison of performance. International Journal of Control 77 (7), 639-653.

Lemke, W. and A. Vladu (2016). Below the zero lower bound – a shadow-rate term structure model for the euro area, Deutsche Bundesbank, Discussion Paper 32-2016

Lombardi, M. and F. Zhu (2014). A shadow policy rate to calibrate US monetary policy at the zero lower bound. BIS Working Papers 452, Bank for International Settlements.

Nelson, C. and A. Siegel (1987). Parsimonious modelling of yield curves. Journal of Business 60 (4), 473-489.

Wu, J. and F. Xia (2014). Measuring the macroeconomic impact of monetary policy at the zero lower bound. Working Paper, National Bureau of Economic Research 2017.

EUROPEAN ECONOMY DISCUSSION PAPERS

European Economy Discussion Papers can be accessed and downloaded free of charge from the following address:
https://ec.europa.eu/info/publications/economic-and-financial-affairs-publications_en?field_eurovoc_taxonomy_target_id_selective=All&field_core_nal_countries_tid_selective=All&field_core_date_published_value[value][year]=All&field_core_tags_tid_i18n=22617

Titles published before July 2015 under the Economic Papers series can be accessed and downloaded free of charge from:
http://ec.europa.eu/economy_finance/publications/economic_paper/index_en.htm

Alternatively, hard copies may be ordered via the "Print-on-demand" service offered by the EU Bookshop: http://publications.europa.eu/bookshop.

HOW TO OBTAIN EU PUBLICATIONS

Free publications:

- one copy:
  via EU Bookshop (http://publications.europa.eu/bookshop);

- more than one copy or posters/maps:
  – from the European Union's representations (http://ec.europa.eu/represent_en.htm);
  – from the delegations in non-EU countries (https://eeas.europa.eu/headquarters/headquarters-homepage/area/geo_en);
  – by contacting the Europe Direct service (http://europa.eu/europedirect/index_en.htm) or
   calling 00 800 6 7 8 9 10 11 (freephone number from anywhere in the EU) (*).

  (*)   The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

Priced publications:

- via EU Bookshop (http://publications.europa.eu/bookshop).

Publications Office