ECB Reaction Functions and the Crisis of 2008*

Stefan Gerlach[a] and John Lewis[b]
[a]Central Bank of Ireland and CEPR
[b]Economics and Research Division, De Nederlandsche Bank

We estimate a reaction function for the European Central Bank (ECB) using forecasts of economic growth and inflation as regressors. We detect a shift after Lehman Brothers failed in September 2008 when the pre-crisis reaction function indicates that the zero lower bound may become a constraint. We detect a shift back in the second half of 2010, several months prior to the April 2011 rate increase. The interest rate cuts in 2008 were more aggressive than forecast by the pre-crisis reaction function. These findings are compatible with the literature on optimal monetary policy in the presence of a zero lower bound.

JEL Codes: C2, E52.

1. Introduction

The recent financial crisis illustrates that if a large, contractionary shock occurs, interest rates may be pushed to the zero lower bound (ZLB). Several authors have argued that this possibility should affect central banks' rate setting even before the constraint becomes binding. However, how it should do so remains an issue of debate. The literature on optimal policy suggests that central banks should be

*The views expressed in this paper are solely our own. We are grateful to seminar participants at De Nederlandsche Bank and at the KOF-ETH-UZH Seminar in International Economic Policy for useful comments, to Jakob de Haan and Fauve Stolwijk for the interpolated Consensus Economics Forecasts, and to Likun Wang for research assistance. Author contact: Stefan Gerlach, Central Bank of Ireland, PO Box No. 559, Dame Street, Dublin 2, Ireland, E-mail: stefan.gerlach@centralbank.ie, Tel: +353 1 224 6007, Fax: +353 1 671 6528; John Lewis, Economics and Research Division, De Nederlandsche Bank, PO Box 98 1000 AB, Amsterdam, The Netherlands, E-mail: j.m.lewis@dnb.nl, Tel: +31 20 524 2835, Fax: +31 20 524 2506.

more aggressive if there is a prospect that interest rates may reach zero in the near future (Orphanides and Wieland 2000; Reifschneider and Williams 2000; Adam and Billi 2006). But the risk of hitting an interest rate floor is sometimes seen as a reason to exercise greater caution in cutting rates—for example, to preserve the option to cut in the future or to avoid generating excessive pessimism (e.g., Bini Smaghi 2008).

In this paper we explore the ECB's interest rate setting behavior during the financial crisis using an econometric framework which allows the coefficients of the reaction function to change over time and where the timing and speed of any regime switch are estimated rather than imposed.

Several key results emerge. First, the ZLB did appear to be a significant constraint on policy. Second, we detect a rapid but not instantaneous shift of the reaction function in the final quarter of 2008, well before rates reached their apparent floor, and a switch back in the latter half of 2010. Third, under this alternative regime the ECB cut interest rates more rapidly in response to worsening economic conditions than the normal reaction function would have posited, consistent with theoretical work on optimal monetary policy in the vicinity of the ZLB.

Although the literature on ECB reaction functions is extensive,[1] few papers analyze how the crisis has affected ECB rate-setting behavior. Gorter et al. (2010) estimate reaction functions with rolling coefficients and detect parameter instability in reaction functions over the period 1999–2010, but do not propose an alternative specification which resolves the problem of unstable coefficients. Gerlach (2011) and Gerlach and Lewis (2011) utilize the smooth transition approach adopted here to examine how the ECB's responses to *actual macroeconomic data* changed during the crisis, and show that interest rate setting became more aggressive in late 2008.

However, Orphanides (2010) notes that a calibrated forward-looking reaction function under which the ECB sets monetary policy

[1] See, for example, Gerlach-Kristen (2003), Surico (2003), Gerdesmeier and Roffia (2004), Carstensen (2006), Castelnuovo (2007), Gerlach (2007), Sauer and Sturm (2007), and Gorter et al. (2010).

according to *macroeconomic forecasts* fits the ECB's observed interest rate setting well throughout the 2008–9 period, suggesting that the reaction function did not shift.[2] This raises the possibility that the apparent non-linearity detected by Gerlach (2011) and Gerlach and Lewis (2011) arises because the relationship between current data and forecasts, rather than the underlying reaction function, changed during the crisis. In this paper we therefore use forecasts of inflation and economic growth as regressors in the empirical reaction function in order to see whether the result that the ECB's reaction function changed still goes through. In addition, by extending the sample period to include the rate increases of April and July 2011, we are able to examine the ECB's behavior when interest rates rose from their apparent floor.

The remainder of the paper is organized as follows. The next section provides an overview of the literature on monetary policy and the ZLB. Section 3 details our empirical approach, section 4 presents the results, and section 5 concludes.

2. The Lower Bound on the Interest Rate

Before proceeding, it is useful to review the literature on the ZLB. The problems posed by a floor on the interest rate were revived in the modern literature by Summers (1991).[3] Subsequently, several papers used model-based simulations to gauge the frequency and extent to which the ZLB would constrain rate setting. They concluded that because the gap between the desired (negative) interest rate and the ZLB is small, the ZLB would have little effect (Reifschneider and Williams 2000; Coenen 2003; Coenen, Orphanides,

[2]It should be noted that there are some differences between our and Orphanides' choice of data and regressors. Thus, Orphanides uses the repo rate rather than the interbank rate as the measure of policy, and uses the difference in forecasts of actual and trend GDP growth as a measure of output conditions. Furthermore, he evaluates this function's performance using a one-step-ahead forecast of the interest rate rather than a longer out-of-sample forecast. The range of the forecast comes from fitting several alternative measures of GDP forecast growth to a rule with fixed parameters, rather than from uncertainty about the parameters themselves.

[3]The constraint was noted much earlier by authors such as Keynes (1936) and even Fisher (1896). See Ullersma (2002) for a comprehensive survey.

and Wieland 2004; and Billi 2005). But as Williams (2009) notes, these earlier studies were based on the probability distribution of disturbances during the Great Moderation and hence may underestimate the frequency and severity of tail events such as the crisis that followed the collapse of Lehman Brothers. Furthermore, he finds that the ZLB did not significantly accentuate the sharp decline in autumn 2008 in inflation and output in the United States, because developments were too fast for monetary policy to react preemptively, but argues that it has slowed the recovery considerably. Our results contribute to this literature by providing an empirical analysis of how long and by how much the ZLB may have constrained interest rates in the euro area over the 2008–11 period.

Other work has analyzed the related question of how monetary policy should respond when central banks are concerned that the ZLB may bind in the near future. Several papers have argued that faced with the possibility of the ZLB constraining rates in the future, policymakers should compensate by cutting interest rates today by more than their regular reaction function would suggest (Orphanides and Wieland 2000; Reifschneider and Williams 2000; Adam and Billi 2006). Accordingly, our paper seeks to contribute to this literature by providing an empirical exploration of the ECB's interest rate setting during the Great Recession.

3. Econometric Considerations

To set the stage for the empirical work that follows, in this section we first discuss the data and then discuss the econometric model.

3.1 Data Choice

Several contributions in the literature have argued that central banks respond to forecasts of future economic variables rather than current levels (Orphanides 2010). Using forecasts as explanatory variables has the advantage that they in principle incorporate all information available at the time to the policymaker, including "soft" information, other variables, and policymakers' expert judgment. This could be particularly important in crisis times, since central banks might look to a much broader set of variables (such as financial stress indicators or asset prices) to gauge the economic outlook than they would in normal times.

The ECB commissions its own Survey of Professional Forecasters (SPF), which is published in the middle month of each quarter and contains average forecasts for inflation, growth, and unemployment at the one- and two-year time horizons. Since this is the measure the ECB itself uses to assess the macroeconomic outlook, it represents a natural data choice for our purposes.

While Taylor (1993) used the output gap and inflation to characterize the interest rate setting of the Federal Reserve, a number of papers have argued that central banks in practice respond more strongly to the growth rate of real GDP than to the output gap. One reason for this is that it is difficult to estimate the trend level of output with any precision in real time (Orphanides and van Norden 2002). Indeed, as emphasized by Gerlach (2007), the ECB's own commentary on its policy decisions emphasizes the role of business sentiment indicators which are much more closely correlated with real GDP growth rather than the deviation of real GDP from trend. In modeling the policy choices of the ECB's Governing Council, of which he was a member, Orphanides (2010) also uses the differential between (forecast) real and potential GDP growth rather than the output gap.[4]

Since we wish to analyze monetary policy at the monthly frequency at which it is set and the SPF data is only available quarterly, we need to interpolate their values in the months between surveys. Another option would be simply to take the latest SPF data available in a given month. However, that would be tantamount to assuming that no new information is received by the ECB in the three months between SPF rounds. This is improbable, especially during the recent crisis. The simplest alternative would be to use linear interpolation for the intermediate months, but this procedure relies on the use of data (the next round of SPF figures) which was not available in real time.

[4]Using inflation and economic growth in a reaction function can be grounded theoretically in the literature on "speed limit" policy rules. This finds that a monetary policy rule based on economic growth and inflation is a close approximation to the unconstrained optimal rule (Walsh 2003; McCallum and Nelson 2004).

Figure 1. SPF Data and Interpolated Data, Inflation Forecasts (% change, year-on-year)

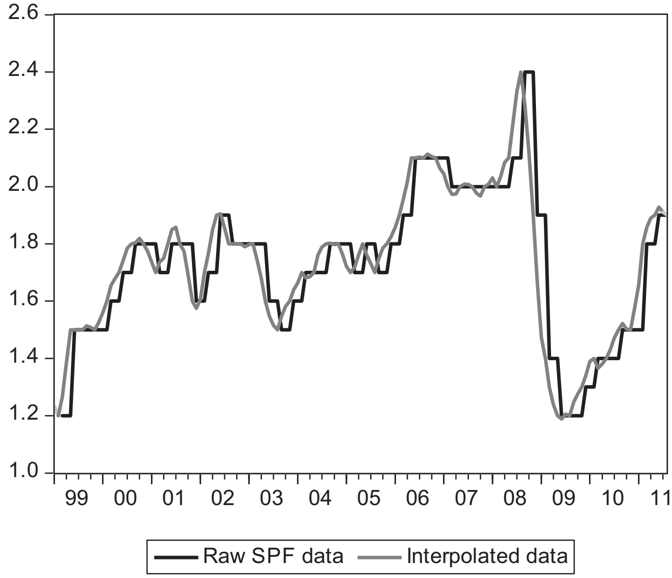

We therefore use the Chow-Lin (1971) method to interpolate the quarterly SPF series using the monthly expectations data supplied by Consensus Economics. This too is a survey, but unlike the SPF, it asks participants to forecast what inflation and growth will be at the end of the current and the following calendar year. Since the forecast horizon is not constant, we use linear interpolation to derive an estimate for forecast inflation twelve months ahead, which we then use to interpolate monthly values for the SPF data.

The raw SPF data together with the interpolated data is shown in figures 1 and 2. These show that inflation forecasts started to rise in the second half of 2008, before abruptly falling after the collapse of Lehmann Brothers in September. By the summer of 2009 the fall was arrested, and inflation forecasts started to rise again. Similarly, the collapse of Lehman Brothers had a large impact on forecast real GDP growth, which fell to –2 percent in the summer of 2009 before rebounding.

An important question that arises when estimating reaction functions for the ECB during the crisis is what interest rate to use as

Figure 2. SPF Data and Interpolated Data, Real GDP Growth Forecasts (% change year-on-year)

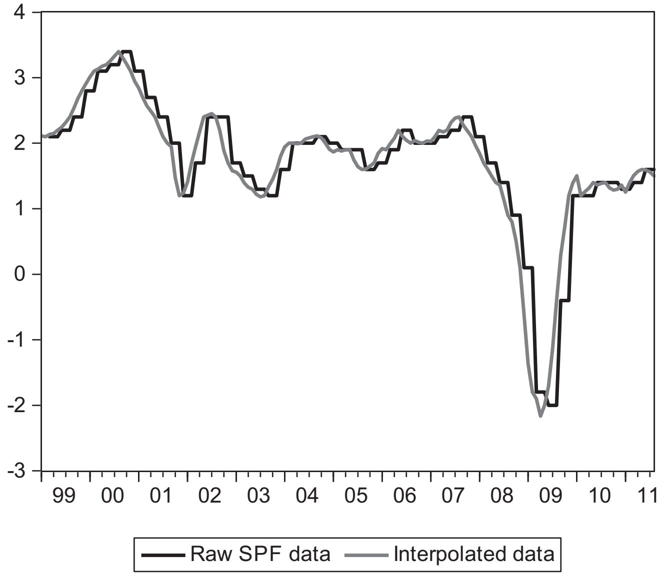

dependent variable. While the literature to date has tended to use the ECB's repo rate, during the crisis the overnight rate (Eonia) fell substantially below the repo rate (Gerlach and Lewis 2011). The ECB had strong incentives to reduce interest rates as much as possible to stimulate the economy and could have prevented the overnight rate from declining below the repo rate if it had felt that it fell too low. We therefore view the rapid decline of Eonia relative to the repo rate as an expression of policy and thus use Eonia as our dependent variable in the econometric analysis below.

3.2 Modeling Interest Rate Setting

We base our approach on a modified version of the model of Judd and Rudebusch (1998). Let i_t denote the Eonia overnight rate, i_t^T the ECB's "target" for the overnight rate, y_{t+12}^f the forecast of economic growth twelve months ahead, and π_{t+12}^f the forecast of inflation twelve months ahead.[5] The target level for the interest rate is given by

$$i_t^T = \alpha_0 + \alpha_y y_{t+12}^f + \alpha_\pi \pi_{t+12}^f, \quad (1)$$

where all coefficients are expected to be positive.[6] The overnight rate is allowed to move gradually towards the target according to

$$i_t - i_{t-1} = \beta_0 (i_t^T - i_{t-1}) + \beta_1 \Delta i_{t-1}, \quad (2)$$

where β_0 governs the gradualism with which this is done and where β_1 captures the extent to which there is inertia in the change in rates.

Combining (1) and (2), we have the following reaction function:

$$i_t = \tilde{\alpha}_0 + \tilde{\alpha}_y y_{t+12}^f + \tilde{\alpha}_\pi \pi_{t+12}^f + \tilde{\alpha}_i + (1 - \beta_0) i_{t-1} + \beta_1 \Delta i_{t-1}, \quad (3)$$

where $\tilde{\alpha}_j = \alpha_j \beta_0$.[7] For ease of exposition, we rewrite this equation more compactly as

$$i_t = \Theta Z_t + e_t, \quad (4)$$

where Θ is a row vector of parameters and Z_t is a column vector of data.

Estimating a single reaction function for the whole period may be problematic because interest rates may have been constrained by the ZLB during the crisis. First and as discussed above, it is eminently plausible that the ways in which policymakers respond to economic conditions may change if it is plausible that the ZLB might

[5] We use the term "target" here for brevity. It should not be thought of as an "optimal rate" in the sense that any delayed adjustment towards it represents "bad policy."

[6] We also experimented with a more general form of (1), based on Smets and Wouters (2003), which adds core inflation, the change in expected future inflation, and the current output gap as additional terms in the target rate function. However, the coefficients on these extra variables were all insignificant at the 5 percent level in both regimes, and the Aikaike, Schwarz Bayes, and Hannan-Quinn criteria all preferred our benchmark model.

[7] Elsewhere in the literature, other papers including a lagged dependent variable use a slightly different terminology: $\tilde{\alpha}$ is sometimes referred to as the "short-run reaction" and α as the "long-run reaction."

be reached in the near future. Second, the fitted reaction function, which is typically linear, must satisfy the non-negativity constraint on the interest rate. As economic conditions worsen, it is otherwise possible that the target interest rate implied by the normal reaction function would be negative. While the ZLB constrains the actual interest rate, which adjusts gradually to the target rate, such a situation can only persist for a limited period of time. At some point the reaction function must shift to an alternative which respects the ZLB constraint.[8] This shift could be either gradual or discrete, and may happen at or before the moment that the ZLB binds. Similarly, when rates rise from their floor, that also implies a switch in reaction function.

In estimating the reaction function, we must therefore take into account the possibility that there may have been two regimes in operation during the sample period. The full model can thus be written

$$i_t = \Theta_I Z_t + e_t \text{ where } e_t \sim N(0, \sigma_I^2) \tag{5a}$$

$$i_t = \Theta_{II} Z_t + e_t \text{ where } e_t \sim N(0, \sigma_{II}^2). \tag{5b}$$

The simplest way to model the structural change would simply be to assume breaks at some point in the sample at dates suggested by the chronology of the crisis such as during the money-market tensions in August 2007 or the collapse of Lehman Brothers in September 2008 for the first break, and the first ECB rate hike in April 2011 for the second. Alternatively, we could employ an Andrews test to select the appropriate break date. In both cases, however, the break would be assumed to be discrete, which rules out the possibility of a smooth transition between regimes that seems eminently plausible.

We therefore employ the smooth transition model of Mankiw, Miron, and Weil (1987). In any given period, the interest rate is a weighted average of the two regimes:

$$i_t = (1 - \omega_t)\Theta_I Z_t + \omega_t \Theta_{II} Z_t + \varepsilon_t, \tag{6}$$

[8]Of course, if the deterioration of the economy is not too extreme, in the sense that equation (1) does not predict a negative interest rate, it is possible that a linear reaction function will fit the data.

where the variance of the errors is also a weighted average of the two regime-specific variances: $\sigma^2 = (1-\omega_t)^2\sigma_I^2 + \omega_t^2\sigma_{II}^2$. Note that this means the error terms are heteroskedastic and so this will be an additional source of specification error if one simply estimates a single reaction function.

The weights follow a logistic function, $L(\bullet)$ of time:

$$\omega_t = L(\theta, \lambda, \tau_t) \equiv \frac{-\exp(\theta(\tau_t - \lambda_1)(\tau_t - \lambda_2))}{1 - (\exp(\theta(\tau_t - \lambda_1)(\tau_t - \lambda_2))}, \tag{7}$$

where τ_t is a time trend.

Since a time trend is deterministic and increases monotonically, this specification permits only two changes of regimes. However, this seems appropriate for the current sample in which there appears to have been a switch into a crisis reaction function and then a shift back as interest rates were raised.

The speed and timing of the transition are captured by the parameters θ, λ_1, and λ_2, respectively. The λ parameters give the value of the trend at the midpoint of the switch (i.e., when $\omega = 0.5$) and the parameter θ describes the speed (a larger negative value implies a faster switch). This functional form also nests the case of a discrete break. In that case, the speed of the change tends towards negative infinity and the midpoint of the change occurs between months. Testing the restrictions implied allows us to formally test whether the switch was discrete.

It should be noted that the model is not identified as it currently stands. Since $L(\theta_0; \ldots) \equiv 1 - L(-\theta_0; \ldots)$, a model with $\theta = \theta_0$ will be identical to a model with $\theta = -\theta_0$. However, the model we estimate below implies restrictions on parameters in the second period that are sufficient to ensure identification.

4. Empirical Results

As suggested by Mankiw, Miron, and Weil (1987), the model can be estimated with maximum likelihood.[9] We begin by estimating an unrestricted version which included forecast growth, forecast

[9] Interestingly, they do not estimate the location and speed of the switch but use a grid search procedure to determine these parameters.

inflation, a lagged interest rate, and the lagged change in the interest rate, specification 1 in table 1.

The first switching point is located in November 2008. (We therefore refer to the regime in force in the first part of the sample as the normal period and that in the second part as the crisis regime.) The failure of Lehman Brothers in September will first be felt in the October forecasts and (since the regressors are lagged by one month to capture reporting lags) will only influence policy in November.

Turning to the coefficients in the reaction function, the pre-crisis estimates have the expected signs and, with the exception of the lagged change in the interest rate, are significant. In the crisis period, both forecast inflation and real GDP growth are insignificant. This suggests that the appropriate restricted form of the reaction function in the crisis is a model with only a constant and the lagged interest rate, which implies that the target interest rate was constant in the crisis period and that the actual interest rate converged gradually towards and then fluctuated around this constant target.

We next dropped all the insignificant parameters and show the results for the restricted model as specification 2 in table 1. In the pre-crisis period, the coefficients are consistent with the earlier literature on ECB reaction functions. The lagged rate is significant, as are the reactions to forecasted inflation and economic growth. Importantly, the long-run response to forecast inflation and real GDP growth are large (4.55 and 1.74, respectively), indicating that the ECB has responded strongly to inflation and to economic growth, which is an important driver of inflation.

The speed and switching parameters are both quite precisely estimated and locate the midpoint of the first switch in the autumn of 2008 and the midpoint of the second switch in mid-2010. To quantify this uncertainty, figure 3 shows the point estimate for the weighting variable over time and together with a 95 percent confidence interval. The confidence interval was obtained by drawing 10,000 times from the joint distribution of the speed and switching date parameters and computing the weighting function in each case. Discarding the top and bottom 2.5 percent of realized outcomes then yields a confidence interval.

The solid line shows the median of the estimates of ω_t, which suggests that the first switch occurred around November 2008 and that it took only four months for the weighting parameter to go from 0.1 to 0.9. The vertical distance between the dashed lines gives a

Table 1. Maximum-Likelihood Estimates, Sample Period: January 1999–August 2011

	1. Unrestricted Model		2. Restricted Model		3. Discrete Break	
Regime:	**Normal**	**Crisis**	**Normal**	**Crisis**	**Normal**	**Crisis**
Constant	−0.708	−0.617	−0.589	0.106	−0.570	
	(0.118) *0.000*	(1.017) *0.733*	(0.100) *0.000*	(0.068) *0.122*	(0.080) *0.000*	
Lagged Interest Rate	0.930	0.575	0.940	0.693	0.937	
	(0.010) *0.000*	(0.132) *0.000*	(0.001) *0.000*	(0.049) *0.000*	(0.011) *0.000*	
Lagged Change in Interest Rate	−0.142	−0.044				
	(0.079) *0.131*	(0.084) *0.605*				
Expected GDP Growth	0.169	−0.047	0.147		0.157	
	(0.030) *0.000*	(0.114) *0.829*	(0.021) *0.046*		(0.009) *0.000*	
Expected Inflation	0.319	−0.047	0.270		0.248	
	(0.055) *0.000*	(0.326) *0.884*	(0.051) *0.000*		(0.042) *0.001*	
St. Dev. of Error Term	0.118	0.071	0.118	0.072	0.169	0.127
	(0.007) *0.000*	(0.014) *0.000*	(0.007) *0.000*	(0.013) *0.000*	(0.062) *0.042*	(0.007) *0.000*
Speed (K)	−2.900		−2.777			
	(0.585) *0.000*		(0.041) *0.000*			
Switching Date (λ_1)	226.220		226.178			
	(1.029) *0.000*		(0.431) *0.000*			
Switching Date (λ_2)	247.185		244.923			
	(2.022) *0.000*		(2.627) *0.000*			
Log-Likelihood	122.141		120.225		92.18	
Akaike Info Criterion	−1.410		−1.437		−1.134	
Schwarz Criterion	−1.111		−1.218		−1.015	
Hannan-Quinn Criterion	−1.289		−1.348		−1.086	

Notes: Standard errors in parentheses; p-values in italics. Likelihood-ratio test of implied restrictions in specification 2 vs. 1 yields p-value of 0.429.

Figure 3. Estimated Weighting Function and 95% Confidence Band

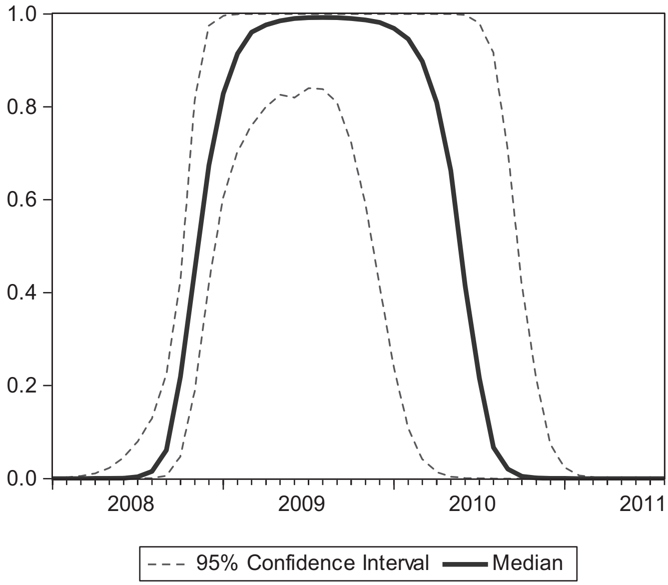

95 percent empirical confidence interval for ω_t. Overall, the switching between regimes seems quite tightly identified. In early 2008, the upper bound of the confidence interval is close to zero, indicating that monetary policy is quite well characterized by the normal reaction function. By September 2008, when the point estimate is below 0.1, even the upper bound is around 0.25, suggesting that the normal reaction function continued to explain the ECB's policy decisions quite well.

Similarly, 90 percent of the second switch was complete by the end of 2010. That suggests that the ECB had switched back into its normal reaction function some months before the interest rate increase of April 2011 and hence the holding of rates in the early months of that year was consistent with its regular reaction function. However, the second switch is somewhat less tightly identified, as the greater vertical span of the confidence interval indicates.

To check whether the transition was smooth rather than discrete, we first estimate an alternative version of the model, with a different specification of the weighting variable. Specifically, ω equals zero when the interest rate is above 0.5, and equals unity otherwise. This implies the following representation of the interest rate:[10]

$$i_t = \max[0.5, \Theta_I Z_t] + \varepsilon_t. \tag{8}$$

The results of this are reported as specification 3 of table 1. Although the coefficients in the reaction function for the pre-crisis period are similar, all three information criteria reject this discrete break specification.

To explore this issue further, we estimated a series of alternative specifications where a discrete break is imposed in a given month. The log-likelihood of these is shown in figure 4 (solid line), together with the log-likelihood of the smooth transition model (dashed line) reported in table 1. The horizontal axis records the month where the break is imposed is in the discrete model. In the upper panel, the second switch is fixed and the first is allowed to vary; in the lower panel, the first switch is fixed and the second is allowed to vary. In both cases, we search six months either side of midpoint identified by the smooth transition model.

Holding the second switch date fixed, the highest log-likelihood for the first switch is the case where October 2008 is the break point, although other break points in the following months yield log-likelihoods which are almost as high. Holding the first switch date fixed in October 2008 and allowing the second to vary gives a local maximum for the log-likelihood function in September 2010. Since the discrete break in October 2008 is a restricted form of our benchmark smooth transition model, the restrictions implied can be tested formally. A likelihood-ratio test of the implied restrictions gives a p-value of 0.02 and thus provides formal support for modeling the regime change as a smooth transition.[11]

As a robustness check, we also investigated several other alternative forms of the model similar to the alternatives considered by Gerlach and Lewis (2011). These included a single reaction function

[10]This reaction function differs from a Tobit model since the interest rate is not constant at the censoring value but is subject to an error term in both regimes. Consequently, the error variance has the form $\sigma^2 = (1 - \omega_t)^2\sigma_I^2 + \omega_t^2\sigma_{II}^2$.

[11]Of course, this test disregards the fact that the discrete break date was selected on the basis of highest log-likelihood amongst alternative models rather than being exogenously given. Correcting for this would further strengthen the rejection of the discrete break model.

Figure 4. Log-Likelihood Function for Discrete Breaks

A. First Break Varies, Second Break Fixed (May 2010)

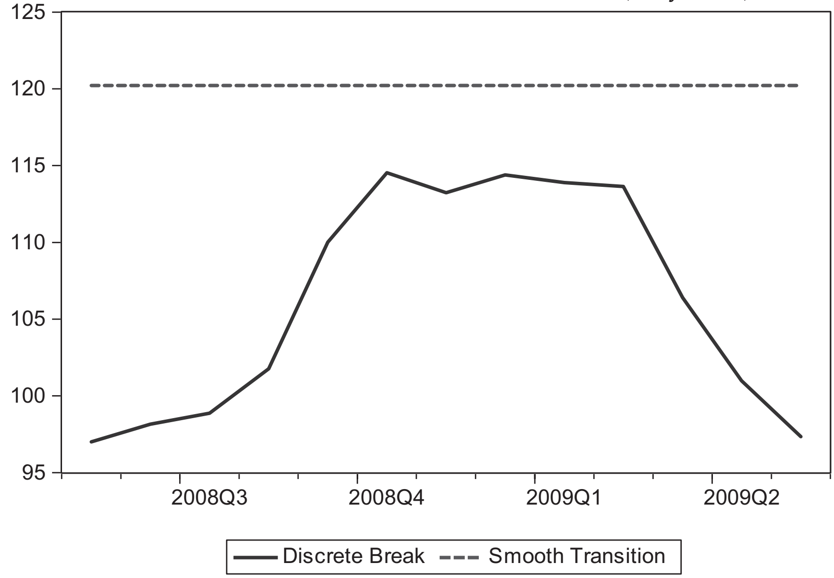

B. First Break Fixed (November 2008), Second Break Varies

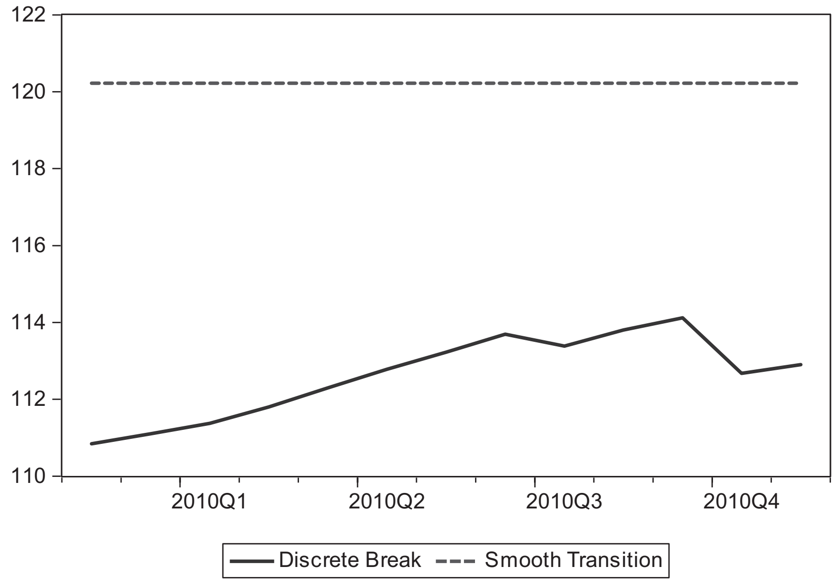

Figure 5. Target Interest Rate with 95% Confidence Band

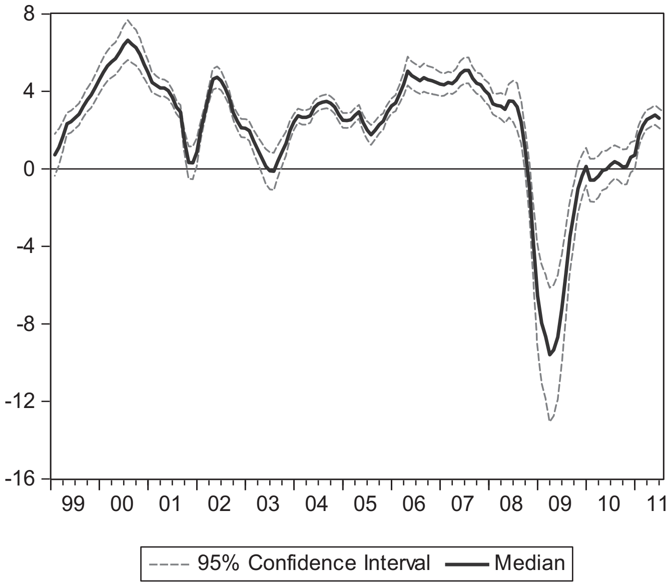

where the coefficients are constant throughout, and one where only the gradualism parameter changes across regime. Log-likelihood-based tests strongly rejected these alternatives.

The above results provide evidence for a change in the ECB's reaction function as the crisis progressed, but they do not say anything directly about how far actual interest rates were from those posited by the normal reaction function. It is to this issue that we now turn.

We begin with the estimates of the target interest rate, i_t^T. This is the interest rate towards which the Eonia rate would converge if the forecast rate of inflation and real GDP growth remained constant at their current levels. Figure 5 shows the target interest rate, assuming that the normal reaction function applied throughout, together with a 95 percent confidence interval, which is constructed by sampling from the joint distribution of the parameters.[12]

[12] As before, we draw 10,000 from the joint distribution of parameters, calculate the path of the interest rate implied by each one, and then discard the upper and lower 2.5 percent in each period.

Several features stand out. First, from the start of EMU until the crisis starts, the point estimate of the target rate is always positive, although in the early part of the 2000s the lower bound of the confidence interval went briefly below zero. Second, following the collapse of Lehman Brothers on September 15, 2008, the target rate declines sharply and by November that year is significantly below zero and remains well below zero for a considerable time. Thus, the pre-crisis coefficients suggest that, given the ECB's past way of setting interest rates, the target rate was negative, and hence it was possible that the overnight rate would reach the ZLB in the future, although the actual rate was some way above zero at the time. Third, the rate rose substantially in the second half of 2009, but the confidence interval does not fully rise clear of zero until the end of 2010.

We next turn to the overnight rate. In order to see how much of the sharp interest rate reduction after the collapse of Lehman Brothers can be attributed to deteriorating macroeconomic forecasts, we construct a dynamic out-of-sample forecast over the period September 2008 to December 2009 using coefficients from the normal reaction function. This forecast is shown in figure 6 (solid thin line), together with a 95 percent confidence interval (the dashed line) and the actual path of the interest rate (thick line).[13]

The figure shows that the predicted policy rate falls sharply after the failure of Lehman Brothers and turns negative in the middle of 2009. However, the cuts seen in the final quarter of 2008 represent a significantly larger and more rapid relaxation of monetary policy than predicted by normal reaction function. By January 2009 this discrepancy was over 100 basis points. Furthermore, the observed policy rate lay outside the 95 percent confidence interval, indicating both a statistically and economically significant difference. As the interest rates level off, this divergence closed, and it then reversed as the forecast rate goes below zero.[14]

This figure is particularly relevant in light of the theory on optimal monetary policy in the vicinity of the lower bound. As noted earlier, theory suggests that as the economic outlook worsens, central banks should cut interest rates faster than one would expect on the

[13]This is constructed using the same method as our confidence intervals for the target rate.

[14]We repeated this experiment using November 2008 and January 2009 as alternative starting points for the forecast, and also obtained similar results.

Figure 6. Dynamic Out-of-Sample Forecasts Starting in September 2008

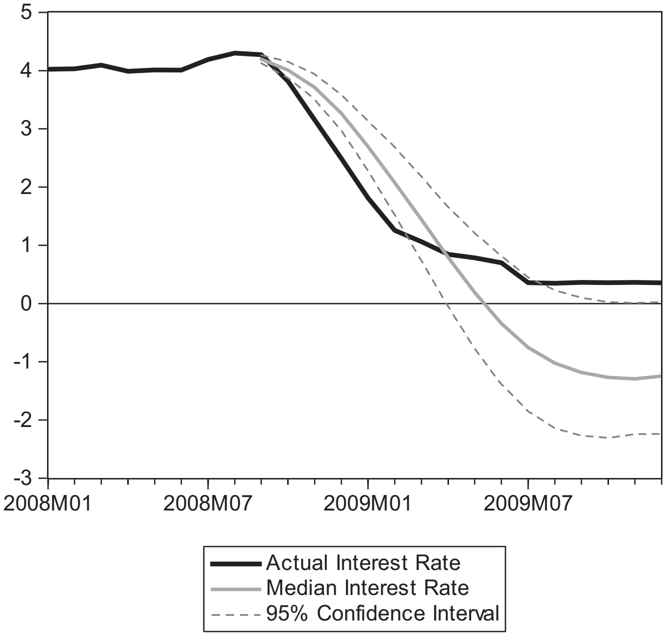

basis of their regular reaction function before. The ECB's behavior appears consistent with the predictions of the theoretical literature which holds that central banks should cut interest rates faster in response to a contractionary shock when there is a likelihood that the ZLB will bind in the future. Some caution needs to be applied here, however, as our findings depend on the assumption that the SPF is a good proxy for the ECB's own forecasts at the time. In the absence of monthly point forecasts from the ECB, this cannot be tested directly, although both the range and midpoint of quarterly ECB staff forecasts appeared to be revised downwards in a similar fashion to SPF data during the period of sharp rate cuts.

5. Conclusions

In this paper we have studied the ECB's interest rate setting during the global financial crisis of 2007–8. Our main findings can be summarized as follows.

First, we find evidence that the ZLB did constrain monetary policy during the financial crisis. We draw this conclusion because the estimated target rate was significantly negative for an extended period from the autumn of 2008 onwards and did not rise significantly above zero during the rest of the sample (which ends in December 2010). Furthermore, the dynamic forecast of the interest rate based on the estimated normal reaction coefficients posits a negative interest rate, which furthermore is statistically significant, for a prolonged period.

Second, the ECB's responses to forecasts of inflation and growth shifted shortly after the collapse of Lehman Brothers. The midpoint of this switch happens at almost exactly the same moment as the target interest rate (computed on the basis of the parameters estimate from data from before the crisis) turns negative. Taken together, these two findings suggest the ZLB was indeed a binding constraint on policy.

Third, there is a switch back which occurs in late 2010. Although this is difficult to tie down very precisely, it appears that the switch was complete several months before the ECB's rate increase of April 2011.

Fourth, as evidenced by the shift of the reaction function, during the crisis interest rates were reduced significantly faster than suggested by pre-crisis estimates of the reaction function. Thus, the shift in the reaction function is not only statistically but also economically significant.

Fifth, the change in the ECB's interest rate setting behavior is consistent with the findings in the theoretical literature on optimal monetary policy in the presence of the ZLB, which suggests that central banks will cut interest rates rapidly when the ZLB may bind in the near future.

References

Adam, K., and R. Billi. 2006. "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates." *Journal of Money, Credit and Banking* 38 (6): 1877–1905.

Billi, R. M. 2005. "The Optimal Inflation Buffer with a Zero Bound on Nominal Interest Rates." *Computing in Economics and Finance* 25.

Bini Smaghi, L. 2008. "Careful with the 'D' Words!" Speech given in the European Colloquia Series, Venice, November 25.

Carstensen, K. 2006. "Estimating the ECB Reaction Function." *German Economic Review* 7 (2): 1–34.

Castelnuovo, E. 2007. "Taylor Rules and Interest Rate Smoothing in the Euro Area." *The Manchester School* 75 (1): 1–16.

Chow, G., and A. Lin. 1971. "Best Linear Unbiased Interpolation, Distribution and Extrapolation of Time Series by Related Series." *Review of Economics and Statistics* 53 (4): 372–75.

Coenen, G. 2003. "Zero Lower Bound: Is It a Problem in the Euro Area?" In *Background Studies for the ECB's Evaluation of Its Monetary Policy Strategy,* ed. O. Issing, 139–55. Frankfurt am Main: European Central Bank.

Coenen, G., A. Orphanides, and V. Wieland. 2004. "Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates are Bounded at Zero." *Advances in Macroeconomics* 4 (1): 1–23.

Fisher, I. 1896. *Appreciation and Interest.* New York: McMillan.

Gerdesmeier, D., and B. Roffia. 2004. "Empirical Estimates of Reaction Functions for the Euro Area." *Swiss Journal of Economics and Statistics* 140 (1): 37–66.

Gerlach, S. 2007. "Interest Rate Setting by the ECB, 1999–2006: Words and Deeds." *International Journal of Central Banking* 3 (3): 1–45.

———. 2011. "ECB Repo Rate Setting during the Financial Crisis." *Economics Letters* 112 (2): 186–88.

Gerlach, S., and J. Lewis. 2011. "The Zero Lower Bound, ECB Interest Rate Policy and the Financial Crisis." Forthcoming in *Empirical Economics.*

Gerlach-Kristen, P. 2003. "Interest Rate Reaction Functions and the Taylor Rule in the Euro Area." ECB Working Paper No. 258.

Gorter, J., F. Stolwijk, J. Jacobs, and J. de Haan. 2010. "ECB Policy-Making and the Financial Crisis." DNB Working Paper No. 272.

Judd, J., and G. Rudebusch. 1998. "Taylor's Rule and the Fed: 1970–1997." *Economic Review* (Federal Reserve Bank of San Francisco) 3: 3–16.

Keynes, J. 1936. *The General Theory of Employment, Interest and Money.* Cambridge: Cambridge University Press.

Mankiw, N., J. Miron, and D. Weil. 1987. "The Adjustment of Expectations to a Change in Regime: A Study of the Founding of the Federal Reserve." *American Economic Review* 77 (3): 358–74.

McCallum, B. T., and E. Nelson. 2004. "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models." *Review* (Federal Reserve Bank of St. Louis) 86 (2): 43–56.

Orphanides, A. 2010. "Monetary Policy Lessons from the Crisis." Working Paper No. 2010-1, Central Bank of Cyprus.

Orphanides, A., and S. van Norden. 2002. "The Unreliability of Output-Gap Estimates in Real Time." *Review of Economics and Statistics* 84 (4): 569–83.

Orphanides, A., and V. Wieland. 2000. "Efficient Monetary Policy Design Near Price Stability." *Journal of the Japanese and International Economies* 14 (4): 327–65.

Reifschneider, D., and J. Williams. 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit and Banking* 32 (4): 936–66.

Sauer, S., and J.-E. Sturm. 2007. "Using Taylor Rules to Understand European Central Bank Monetary Policy." *German Economic Review* 8 (3): 375–98.

Smets, F., and R. Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1 (5): 1123–75.

Summers, L. 1991. "How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit and Banking* 23 (3): 625–31.

Surico, P. 2003. "Asymmetric Reaction Functions for the Euro Area." *Oxford Review of Economic Policy* 19 (1): 44–57.

Taylor, J. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39 (1): 195–214.

Ullersma, C. 2002. "The Zero Bound on Nominal Interest Rates and Monetary Policy Effectiveness: A Survey." *De Economist* 150 (3): 273–97.

Walsh, C. 2003. "Speed Limit Policies: The Output Gap and Optimal Monetary Policy." *American Economic Review* 93 (1): 265–78.

Williams, J. 2009. "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound." *Brookings Papers on Economic Activity* 40 (2): 1–49.