

IDENTIFYING THE EFFECTS OF MONETARY POLICY SHOCKS ON EXCHANGE RATES USING HIGH FREQUENCY DATA

Jon Faust

John H. Rogers

Eric Swanson

Jonathan H. Wright

Board of Governors,

Federal Reserve System

Abstract

This paper proposes a new approach to identifying the effects of monetary policy shocks in an international vector autoregression. Using high-frequency data on the prices of Fed Funds futures contracts, we measure the impact of the surprise component of the FOMC-day Federal Reserve policy decision on financial variables, such as the exchange rate and the foreign interest rate. We show how this information can be used to achieve identification without having to make the usual strong assumption of a recursive ordering. (JEL: C32, E52, F30)

1. Introduction

The role of monetary policy in explaining the dynamics and volatility of exchange rates is a central theme in empirical international finance. The current predominant approach to identifying structural monetary policy shocks, in both closed- and open-economy settings, involves using a vector autoregression (VAR). This approach relies on making identifying assumptions relating structural shocks to the reduced form errors of the VAR. While many identification approaches have been proposed for identifying VARs, most often short-run restrictions are used. These specify that some structural shock has no contemporaneous effect on one or more variables. In an open-economy setting, such identifying assumptions are used by Eichenbaum and Evans (1995), Kim and Roubini (2000), and Kim (2001).

Identification of structural monetary policy shocks in VARs is contentious because, as the authors generally acknowledge, there are few highly credible

Acknowledgments: We are grateful to Jonathan Halket and Shing-Yi Wang for superlative research assistance. We are grateful to Jim Stock, Glenn Rudebusch, Chris Sims, Harald Uhlig, and two anonymous referees for their very helpful comments on earlier drafts of this manuscript. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

E-mail addresses: Faust: jon.faust@frb.gov; Rogers: john.h.rogers@frb.gov;
Swanson: eric.t.swanson@frb.gov; Wright: jonathan.h.wright@frb.gov

identifying assumptions. Open economy VAR applications raise particularly thorny simultaneity issues. For example, most closed economy applications involve a single financial market variable, a short-term interest rate; long-term rates are generally excluded due to the identification problems that arise when they are included.¹ To be minimally credible, the open economy analogs must include three financial market variables: a short-rate in each country and the exchange rate. Satisfactory identifying restrictions for sorting out the contemporaneous movements of these variables simply have not been found. For example, some papers assume that U.S. monetary policy shocks have no effect on foreign interest rates until a month after the policy move (Eichenbaum and Evans 1995, Kim and Roubini 2000). This is at odds with the fact that foreign central banks regularly change policy in the wake of Federal Reserve policy decisions. Other authors assume that the Fed ignores any surprise movements in exchange rates and/or short-term interest rates that have occurred during the month in which decisions on the policy variable are made (Eichenbaum and Evans 1995 and Kim and Roubini 2000). If true, these assumptions would call into question why the Federal Reserve Board staff invest tremendous effort in providing the Board with minute-by-minute information about surprising movements in financial markets.

Aware that the assumptions are not entirely credible, authors typically discuss results from a few alternative identifications, indicating that the published results are robust to changes. Such robustness checks are of course indispensable. Nonetheless, in cases where the alternative identifications are each recursive, a sense of dissatisfaction lingers since we expect simultaneity among asset market variables.

Motivated by these considerations, Faust and Rogers (2003) apply an approach to identification, originally developed by Faust (1998). This is an approach that allows one to do inference in partially identified models. Using such methods, one can test whether the answers to key questions are robust to dropping implausible identifying assumptions. Using a standard open-economy VAR, Faust and Rogers find that some key results are highly sensitive to the assumed recursive structure of money market variables, while other results are robust. For example, the “delayed overshooting” response of the nominal exchange rate commonly found under the assumption that foreign interest rates do not respond contemporaneously to U.S. monetary policy shocks vanishes when even a slight response of foreign rates is allowed. On the other hand, the assumption that monetary policy shocks generate large deviations from uncovered interest rate parity is not sensitive to loosening the recursive structure.

The approach of Faust and Rogers can show which answers are sensitive to allowing simultaneity among financial market variables. When sensitivity is found, additional identifying information is needed to sharpen our inferences.

In this paper, we bring high-frequency financial market data to bear in

1. See Leeper, Sims, and Zha (1996) for a thorough description of this issue and examples of VARs with long and short rates.

identifying the monetary policy shock following the approach of Faust, Swanson, and Wright (2002a). We assume that the change in the Fed Funds target rate on the days of Federal Open Market Committee (FOMC) meetings that was not anticipated by futures markets represents a monetary policy shock. We then regress changes in exchange rates and spot and future interest rates in a narrow window around the FOMC decision on the surprise change in the target rate, and then impose that the impulse responses of the exchange rate and U.S. and foreign short-term interest rates in a standard open-economy VAR match the responses we have estimated from the high frequency financial market data.

Our key results are these:

1. Most of the impulse responses of the system to U.S. policy shocks under the new identification are consistent with those from the recursive identification. However, the effect of the U.S. policy shock on foreign output and interest rates lasts longer than with the recursive identification. The price puzzle in the recursive identification is avoided with the new identification. For Germany, we formally reject the recursive identification, but not for the United Kingdom.
2. The peak timing of the exchange rate response is imprecisely estimated as in Faust and Rogers (2003). Whereas the recursive identification suggests strong evidence of delayed overshooting, the confidence interval for the peak timing in the new identification includes immediate peaks and delays of several years.
3. All the approaches agree that monetary policy shocks generate large UIP deviations. The movements of the exchange rate following U.S. policy shocks do not seem to be driven by UIP.
4. The confidence interval for the variance share of the exchange rate due to the policy shock in the new identification is somewhat larger than in the recursive identification, but is bounded by about $\frac{1}{3}$. This is somewhat tighter than the estimates of Faust and Rogers, reflecting the fact that additional information has been brought to bear.

Other authors have also used high-frequency financial market data to help identify the monetary policy shock in an otherwise conventional VAR. Bagliano and Favero (1999) take a monetary policy shock identified by interest rate moves around policy decisions and use it to identify the effects of a policy shock in closed and open economy VARs. Cochrane and Piazzesi (2002) use a similar approach in a closed economy VAR. The primary difference in our method is that we also exploit futures market data and high-frequency spot exchange rate data. In particular, we require that the VAR replicate the effect of the policy shock on expected future home and foreign rates as measured from futures markets and spot exchange rates. We also focus on a different set of questions—exchange rate effects of U.S. policy shocks—than these other papers.

Section 2 discusses the approach to identification. Section 3 presents our approach and results from the high-frequency data exercise. Section 4 contains

the VAR results. Section 5 contains some tests of our identifying assumptions, and Section 6 concludes.

2. Identification

2.1 The Simplest Case

Consider the reduced form VAR,

$$A(L)Y_t = u_t \quad (1)$$

where Y_t is $G \times 1$, $A(L) = \sum_{j=0}^{\infty} A_j L^j$ and $A_0 = I$. Following the literature we assume that $A(L)$ is invertible so that the system can be written as,

$$Y_t = B(L)u_t \quad (2)$$

where $B(L) = A(L)^{-1}$.

The identified VAR literature makes the assumption that the G reduced form errors u_t are related to structural errors ε_t by the relation: $u_t = S\varepsilon_t$, where S is full rank. One of the structural shocks is assumed to be the monetary policy shock of interest. We can order things such that this is the first structural shock. The VAR can be written in terms of the structural shocks as,

$$Y_t = B(L)S\varepsilon_t \quad (3)$$

Call the first column of S , α ; this is the column corresponding to the policy shock. The impulse response of all variables in the VAR to the policy shock is,

$$B(L)\alpha = \sum_{j=0}^{\infty} B_j \alpha L^j$$

This is a $G \times 1$ vector of lag polynomials and the coefficients of the g^{th} element trace out the response of the g^{th} variable to the policy shock.

The B s are given by the reduced form estimates and so identifying the impulse response requires picking the G elements of α . One restriction is a normalization, choosing the sign and units of the policy shock. In most work, one normalizes the standard deviation of the shock to be 1. In our work, the VAR includes the three-month eurodollar interest rate and we normalize the shock to have a contemporaneous -25 basis point effect on this interest rate.²

2. The choice of normalizing the impact effect or the standard deviation of the shock is innocuous in the point estimates, although not in confidence intervals. Suppose a one-standard-deviation shock has a -25 basis point effect in the point estimates. A 95 percent confidence interval for the effects of a one-standard-deviation shock need not be a 95 percent confidence interval for a 25 basis point shock. This is because the impact effect of a one-standard-deviation shock is stochastic. Our normalization is chosen for two reasons. It is technically convenient, and it leads to confidence intervals for something we want to learn: the effects of a given size shock. We are less interested in the effects of a one-standard-deviation shock, where the value of the standard deviation is not stipulated.

We complete the identification by requiring that certain impulse responses match values given from the high-frequency data. For this section, simply take it as given that we have some restrictions saying that the impulse response of the j^{th} variable to the policy shock at lag h is r_{jh} . This restriction can be written,

$$B_{h,j}\alpha = r_{jh} \quad (4)$$

where $B_{h,j}$ is the j^{th} row of B_h . If we have G such restrictions, we can stack them to form

$$R\alpha = r$$

Clearly, if R and r are taken as known and R is full rank, α is uniquely identified as $R^{-1}r$.

2.2 Factors Complicating Inference

In the previous discussion, we treated R and r as known. In practice R will be implied by the reduced-form estimates of the VAR and r will be estimated from the futures market data. We must take account of uncertainty in each when doing our inference. More problematically, full identification rests on the condition that R is of rank G . When we test the rank of our estimated R s below, we cannot reject rank deficiency. We discuss reasons for this below. Thus our restrictions $R\alpha = r$ leave the system only partially identified—some linear combinations of α may be well identified while others are not. We must use methods appropriate for partially identified systems. An alternative would be to find more restrictions. We argue below that we are exploiting all the restrictions from high-frequency data that we could identify.

We take a classical approach to inference in partially identified systems. First, we use economic reasoning to bound the magnitude of each element of α above and below. Remember that the elements of α are the contemporaneous effect of the policy shock on each variable in the VAR and that we normalize the effect on the short-rate in the United States to 25 basis points. Thus, one can interpret our bounds as limits on the relative effect on other variables of a shock that lowers the U.S. interest rate by 25 basis points. We choose bounds that are largely uncontroversial, but these limits remain a substantive part of the identification. It is important to note that several of these restrictions are strictly looser than the restrictions imposed in recursive identification. In that work, certain contemporaneous responses are set to zero.

These bounds are required because confidence intervals for linear functions of α (such as structural impulse responses) would typically be unbounded when α is only partially identified. Even with these bounds, the failure of the rank condition means that we must give up on point estimation and only consider confidence intervals constructed in a way that is robust to the failure of the rank condition.

2.3 Confidence Intervals Under Partial Identification

Suppose we want to learn about some scalar parameter f . This could be the share of the forecast error variance of output at horizon 48 due to the policy shock or the impulse response of prices to the policy shock at some horizon. Calling all the reduced form parameters of the VAR θ , f is a function of θ and α : $f(\theta, \alpha)$.

To form confidence intervals for f , first, we form a $v_1\%$ confidence set for α by a method that takes account of uncertainty in R and r and that does not rely on assumptions about the rank of R . The construction of this confidence set follows the work of Stock and Wright (2000) and is discussed in detail in Appendix A1. Call this confidence set A .³

For any fixed α in A , we can use a conventional bootstrap to construct a $v_2\%$ confidence interval for $f(\theta, \alpha)$. Let this confidence interval be $[\underline{c}(\alpha), \bar{c}(\alpha)]$. Next form the outer envelope of all of these intervals across all α s in A , as $[\inf_{\alpha \in A} \underline{c}(\alpha), \sup_{\alpha \in A} \bar{c}(\alpha)]$. This confidence interval has asymptotic coverage of at least $v_1 + v_2 - 100\%$, from the Bonferroni inequality, because asymptotically (i) the true α is included in A with probability $v_1\%$, and (ii) the bootstrap confidence interval has $v_2\%$ coverage for any fixed α . The technique is conservative in that coverage may asymptotically be higher than $v_1 + v_2 - 100\%$.⁴ The resulting confidence interval may be wide, reflecting in part its construction as a conservative confidence interval using the Bonferroni inequality. Henceforth in this paper, we set $v_1 = 95$ and $v_2 = 73$ ensuring that the asymptotic coverage is at least 68 percent, a coverage rate commonly applied in VAR work.

An alternative approach to inference in unidentified systems is provided by a Bayesian framework. While the Bayesian approach might be simpler in some respects, the results in underidentified systems may be highly sensitive to the prior, even in large samples. Our classical confidence intervals will (asymptotically) have at least the stated coverage so long as the bounds we impose on α are correct. Thus, so long as ones prior for α is not inconsistent with our bounds, the stated results should be of interest.

3. High-Frequency Asset Price Data and Impulse Responses to Policy Shocks

This section develops the claim, taken as given in the last section, that some structural impulse responses to a monetary policy shock can be measured from high-frequency data on interest rates, interest rate futures, and exchange rates.

3. We also form confidence intervals reflecting uncertainty regarding α but not θ . Fixing θ at the reduced form point estimate $\hat{\theta}$, we can find the range of $f(\hat{\theta}, \alpha)$ consistent with α in our confidence set A : $[\inf_{\alpha \in A} f(\hat{\theta}, \alpha), \sup_{\alpha \in A} f(\hat{\theta}, \alpha)]$. We do not report these to save space. They are substantially smaller than the confidence intervals we report later.

4. For example, even when the true α is not in A , the confidence interval may contain the true f .

Our method relies on several important assumptions. These are to some degree testable and we present evidence in support of the assumptions in Section 5.

3.1 The Principal Assumptions

Since March 1994, the FOMC has made a public announcement about its target for the Federal Funds rate at 2:15 P.M. Eastern time on each of its eight regularly scheduled meeting dates every year. We follow Kuttner (2001) and other papers⁵ in assuming that the unexpected change in the target rate in the FOMC announcement can be measured from Federal Funds futures contracts (described precisely below) and that this unexpected target change reflects an exogenous monetary policy shock. This assumption will fail, for example, if the FOMC's policy move at 2:15 P.M. reveals private information of the Fed's about the state of the economy. In Section 5, we discuss our tests of this assumption. This assumption does not require that there are no policy shocks on other days—say, days the Chairman testifies in Congress.

For three variables in the VAR (domestic and foreign short-term interest rates and the exchange rate), we have high-frequency data. We estimate the contemporaneous effect of the policy shock on these variables by regressing the change in each variable in a narrow window around the FOMC meeting on the unexpected change in the target rate at the FOMC meeting, as measured from the Federal Funds futures market (which we henceforth refer to as the FFT shock). We assume that the relative movements in these variables in response to the FFT shock measure the effect of a monetary policy shock on these variables in the VAR. We gain efficiency in estimating the contemporaneous effects of the policy shock on other variables by picking the narrowest possible window around the policy shock.⁶

We gain four additional restrictions by assuming that we can measure the change in expectations of future spot rates due to the policy shock from interest rate futures markets. In particular, we use high-frequency three-month and six-month futures data on the home and foreign short rate. We take the change in these rates due to the FFT shock in a narrow window around the FOMC announcement as a measure of the effect of the policy shock on expected interest rates at the relevant horizon. In addition to the assumptions stated so far, this step requires that risk premia embedded in the futures rates do not change over the window of time when we measure the change. We test some implications of this assumption next.

One might wonder why we do not use additional high-frequency data in

5. Cochrane and Piazzesi (2002) use changes in the spot thirty-day eurodollar interest rate around the FOMC announcement instead.

6. Picking the narrowest possible window for the left-hand side variables in this regression is purely a matter of efficiency, not consistency.

order to gain further restrictions. We think we have exploited much of the relevant data. We use high-frequency data for every variable in the VAR that is measured at high frequency. There are other futures or forward data for the interest rates and exchange rates in the VAR. Unfortunately, our tests indicate that the risk premia in these data appear to be too variable to reliably treat these as measures of expectations.

3.2 *Asset Price Data*

The Federal Funds futures contract, traded on the CBOT, settles on the last day of the month for the average effective Federal Funds rate over all days in that month. Following Kuttner (2001), we measure the surprise change in the target Federal Funds rate on the FOMC day by taking the difference between the closing price of the Federal Funds contract that day (3:00 P.M. Eastern time) and the closing price the previous day and multiplying this change by (days in month/days left in month).⁷ We would prefer to have the change in the Fed Funds futures rate in a narrower window of time around the FOMC announcement (2:00 P.M.–2:30 P.M.), but we have these data only for the last two years of the sample.

We have observations at five-minute intervals for the sterling and the mark/euro exchange rates against the dollar, obtained from Olsen Associates (Zurich, Switzerland). We measure the change around the FOMC announcement from 2:00 P.M. and 2:30 P.M. The data on spot and future three-month interest rates are observed only at the daily frequency, so we use the change in daily quotes. Since these quotes are taken at different times for different assets, the main issue is whether we want the close from the day before to the day of the announcement or the change from the day of to the day after. We measure U.S., U.K., and German interest rates using the spot three-month eurodollar, spot three-month sterling LIBOR, and spot three-month fibor/euribor deposit rates, respectively (mark and euro rates are spliced). These rates are directly comparable to each other, the associated assets are very actively traded by international market participants, and there are very well developed futures markets corresponding to each. We measure expected future interest rates three-months ahead and six-months ahead using these eurodollar, LIBOR, and fibor/euribor futures contracts that trade in Chicago, London, and Frankfurt,

7. Near the end of the month, this scaling factor is quite large. Unfortunately, our Fed Funds futures rate data are recorded only to the nearest basis point (to the nearest half basis points since 1995). Thus, our measured changes involve measurement error that is greatly exacerbated when the time-of-month scaling factor is large. For this reason, we take the target surprise to be the change in the next month's contract whenever the change is after the 22nd of the month. Whenever the FOMC meeting occurs this late in the month, there is no FOMC meeting the next month.

TABLE 1. MEASURES OF THE IMPULSE RESPONSE TO A POLICY SHOCK

Variable	Horizon	UK VAR		German VAR	
		Rel. response	St. error	Rel. response	St. error
<i>i</i>	0	-0.25		-0.25	
	3	-0.233	0.042	-0.233	0.042
	6	-0.206	0.060	-0.206	0.060
<i>i</i> *	0	0.028	0.071	-0.040	0.022
	3	-0.016	0.076	-0.113	0.043
	6	-0.016	0.089	-0.127	0.067
<i>s</i>	0	0.352	0.156	0.611	0.271

Notes: The results are for a least squares regression of the change in the spot/future interest rate or exchange rate on the unexpected change in the target Federal Funds rate, with no intercept, around the FOMC meeting. The coefficient on the spot U.S. interest rate is normalized to -25 basis points. There are 62 observations; the standard errors are conventional OLS standard errors using the delta method to adjust for the normalization.

respectively.⁸ Our Eurodollar spot rate is the British Banker's Association trimmed mean of market quotes at 11 A.M. London time each day, well before the FOMC announcement. Libor and Fibor/Euribor interest rate futures prices are closing prices in London and Frankfurt, and these markets close before the FOMC announcement. Thus, for all these series we take the change from the day of to the day after the announcement. The eurodollar futures prices are taken at the Chicago close which is after the FOMC announcement. Thus, in this case, we use the change from the day before to the day of the announcement.

3.3 Estimating the Structural Impulse Responses

We run the regression of the changes in the spot three-month interest rates, three-month and six-month interest rate futures for the home and foreign country, and the exchange rate on the FFT shock (the unexpected target rate change) for all of the 62 FOMC meetings from March 1994 to October 2001, inclusive. We then normalized the coefficient on the spot U.S. interest rate to -25 basis points to get our estimate of structural impulse responses.⁹

The results, using the United Kingdom and Germany as the foreign countries are reported in Table 1. The coefficients on the three-month- and six-month-ahead U.S. interest rates are negative and highly significant. The coefficients on the spot and future interest rates are not significantly different from zero for the United Kingdom. The coefficient on the three-month-ahead German interest rate is significantly negative and is about half the size of the 25 basis

8. These contracts all settle based on the spot eurodollar, Libor or Fibor/Euribor interest rate on the last day of the contract. Liquid contracts exist settling in March, June, September, and December of each year. We use linear interpolation to compute the implied three- and six-month-ahead forecast interest rates.

9. The Fed had a Federal Funds rate target before March 1994, but we use only data since March 1994 because we need to know the exact time at which the public learns of the FOMC target rate decision. A public announcement has been made at 2:15 P.M. since March 1994 only.

point decline in the spot U.S. rate. The coefficients on the spot and six-month ahead German interest rates are negative and borderline significant. The coefficient on the exchange rate is positive (a surprise loosening of monetary policy depreciates the dollar), and the coefficient is significant for both the United Kingdom and Germany. The magnitudes of the coefficients are, in our view, quite reasonable. The effect on U.S. rates decays slightly over the six-month horizon, but is nearly constant. The identification procedure takes these point estimates as the r_{jt}^s in (4). We take account of uncertainty in these estimates using the conventional variance-covariance matrix. We next combine these results with the VAR to gain identification.¹⁰

4. Results on the Identified VAR

In this section we apply the methodology to a benchmark seven-variable VAR of Eichenbaum and Evans (1995). Our dataset consists of monthly observations from January 1974 through October 2001. The variables are domestic and foreign output (y and y^*) measured as industrial production, U.S. prices (p) measured as the CPI, the three-month U.S. and foreign interest rates (i and i^*) described previously, the ratio of nonborrowed reserves to total reserves in the U.S. ($nbrx$), and the exchange rate (s) measured as the dollar price of foreign currency. The two foreign countries are the United Kingdom and Germany. The details of the data sources are in the data appendix. All of the variables, except the i and i^* are in logs, and the VAR includes six lags and a constant.

Eichenbaum and Evans (1995; EE) estimate a recursive VAR with the data ordered as y , p , y^* , i^* , $nbrx$, i , and s , calling the shock in the NBRX equation the monetary policy shock. Figures 1 and 2 show the estimated impulse responses and 68 percent bootstrap confidence intervals for the recursive identification, for both countries. The results are generally reasonable by the standards of the literature and generally consistent with what EE find using slightly different data and a sample ending in May 1990. The surprise 25 basis point loosening of U.S. policy persists for about six months and then decays rapidly. The U.K. interest rate falls about half as much but is more persistent. Home output rises gradually to a peak effect of nearly a percentage point after two years and then decays. Foreign output follows a similar pattern, but at about half the magnitude. There is a “price puzzle” in that the home price level initially falls significantly following a monetary policy loosening. The exchange rate response is quite different from that in EE, however. It initially rises and then has a second mode at a horizon of about three years. The German results show roughly the same pattern.

We are particularly interested in three questions concerning the exchange

10. We assume there is no covariance between these estimates and the estimated VAR coefficients discussed in the next section. This assumption strikes us as reasonable.

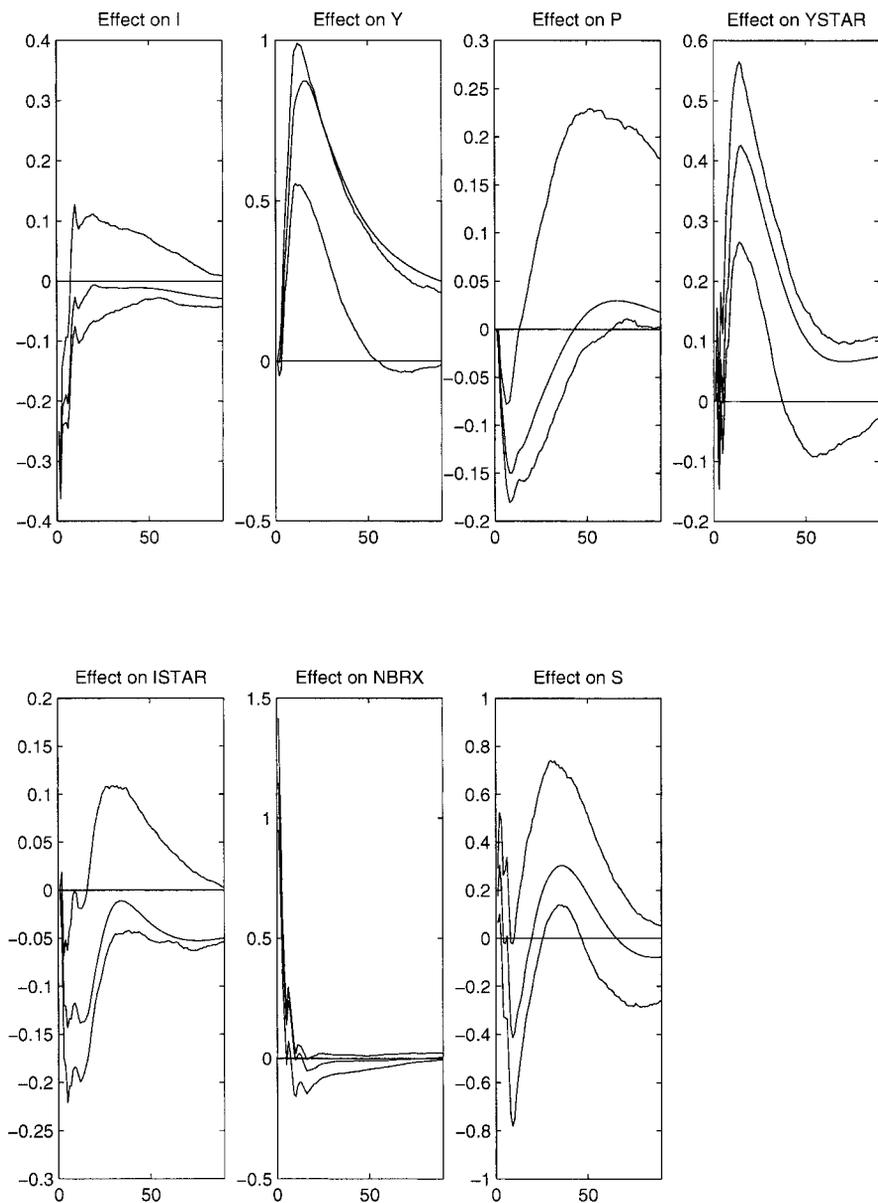


FIGURE 1. Recursive Identification Impulse Responses for the United Kingdom (with 68% Bootstrap Intervals)

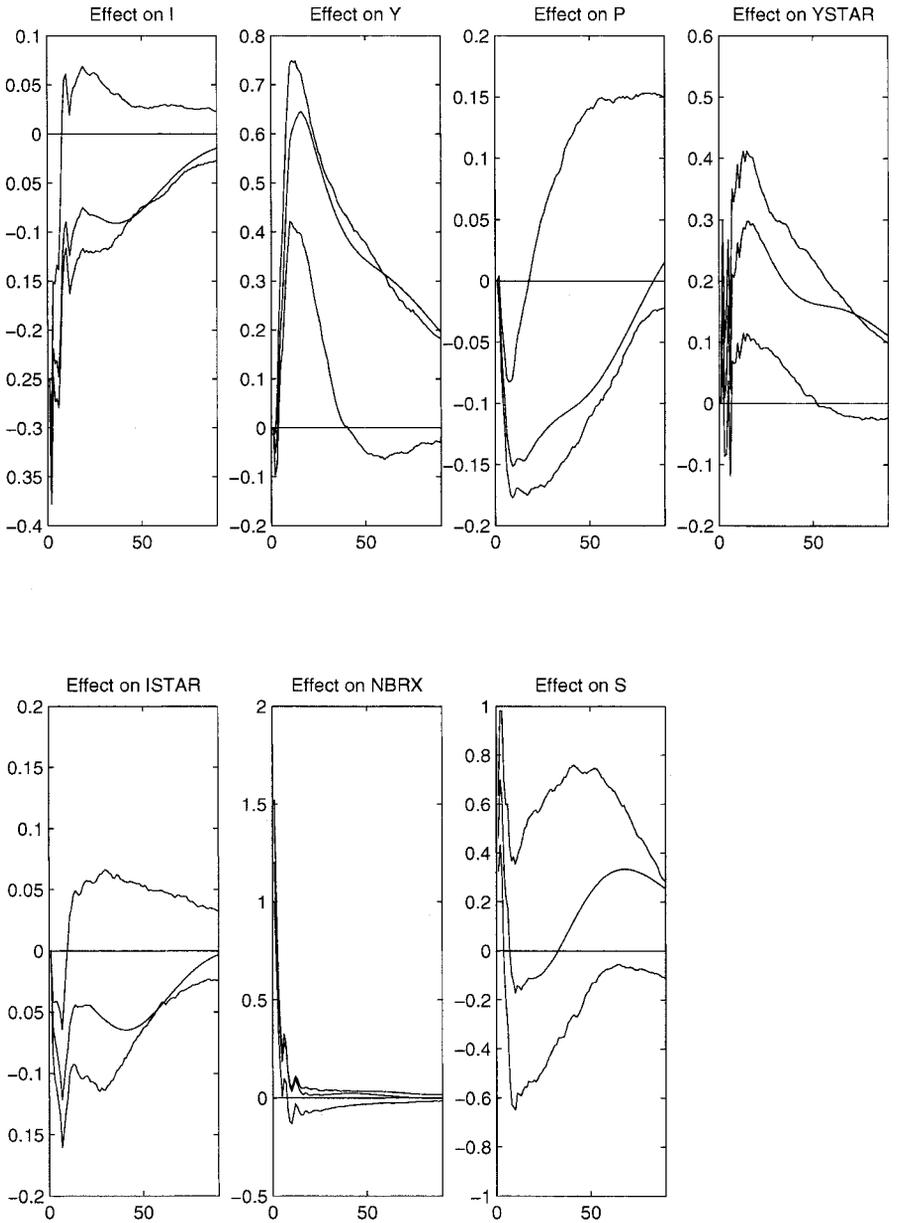


FIGURE 2. Recursive Identification Impulse Responses for Germany (with 68% Bootstrap Intervals)

rate response: 1) What is the timing of the peak exchange rate effect? 2) What share of the variance of exchange rates is due to monetary policy shocks? 3) Is the response to policy shocks consistent with uncovered interest rate parity (UIP)? These questions can all be motivated by Dornbusch's classic work on overshooting (1976). This model was designed to help explain the high volatility of exchange rates relative to macroeconomic fundamentals. In Dornbusch-style overshooting, the peak exchange rate effect should come contemporaneously with the shock, and the dynamics of the exchange rate are consistent with UIP.

With regard to the question of UIP, we know that UIP does not hold unconditionally in the data. The deviation from UIP is interpreted as a time varying risk premium and called the forward premium bias puzzle (see, e.g., Engel 1996). It remains conceptually possible, however, that UIP holds conditionally in response to money shocks. In this case, the monetary policy shock does not drive the variance of the risk premium or equivalently, monetary policy shocks do not contribute to the forward premium bias. Most prior work finds that conditional UIP does not hold.¹¹

To assess this issue we calculate the implied root mean square UIP deviation (UIPD) over forty-eight months following the money shock. The expected UIPD deviation at $t + h$ of a shock at t is given by,¹²

$$c(i, l) - c(i^*, l) - 400[c(s, l + 3) - c(s, l)].$$

where $c(x, l)$ is the response of variable x at lag l to the policy shock. The RMSE of the UIPD comes from summing the squared deviations over the forty-eight-month horizon, and taking the square root of this object.¹³ A large RMSE UIPD implies either large absolute deviations or highly variable deviations, or both.

The top panel of Table 2 shows the estimates and 68 percent bootstrap confidence intervals for various parameters relevant to answering our three questions: (1) the fraction of the variance of exchange rates at horizons 12, 24, 36, 48, and 60 months that are due to the monetary policy shock, (2) the time of the peak effect of the monetary policy shock on exchange rates, and (3) the RMSE UIPD.

The EE model draws mixed conclusions at best regarding Dornbusch overshooting as an explanation for exchange rate movements. For both countries and all horizons, the confidence interval for the variance share of the exchange

11. Eichenbaum and Evans (1995), Cushman and Zha (1997), and Kim and Roubini (2000) report that policy shocks generate deviations from UIP that are several times larger than the generated interest rate differential. Cushman and Zha note that the pointwise coverage intervals on the UIP deviations cover zero, but do not report a joint statistic on the statistical significance of the UIP deviations.

12. This is annualized, presumes monthly data, and three-month interest rates in annual percentage rate units.

13. Some tricky timing and definition questions arise. We use monthly average data for exchange rates and interest rates. If the identification is correct, then the calculated UIP deviations should be interpreted as the expected path of the monthly average UIP deviation in response to a money shock.

TABLE 2. SUMMARY OF THE RESPONSE OF THE EXCHANGE RATE TO THE MONETARY POLICY SHOCK

	Variance share at horizon										Peak time		UIPD	
	12	24	36	48	60									
Recursive identification														
UK pt. est.	0.01	0.01	0.01	0.02	0.02	1		0.45						
UK CI	0.01	0.05	0.01	0.06	0.02	0.09	0.02	0.11	0.02	0.11	30	34	0.47	0.98
Germany pt. est.	0.01	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01	0.01	1	1	0.47	1.05
Germany CI	0.01	0.05	0.01	0.06	0.01	0.07	0.01	0.08	0.01	0.09	1	1	0.47	1.05
New identification														
UK CI	0.00	0.35	0.01	0.28	0.01	0.25	0.02	0.24	0.02	0.26	0	64	0.34	0.99
Germany CI	0.00	0.34	0.01	0.31	0.01	0.33	0.02	0.34	0.02	0.33	1	67	0.31	0.96

Notes: The confidence intervals are conservative 68 percent bootstrap intervals as discussed in the text. The peak time and variance share horizons are in months. UIPD is the root mean square UIP deviation at horizon 48.

rate accounted for by the policy shock is 11 percent or less. The confidence interval for the United Kingdom shows the peak exchange rate effect occurring more than two years after the shock; the German peak is much earlier. Finally, for both countries the RMSE UIPD is quite large.

4.1 Results for the Identified VAR

Remember that we can view the identification problem as choosing a vector α and that the seven elements of α give the impact effect of the policy shock on the seven variables in the VAR. The element of α corresponding to the domestic interest rate is normalized to -0.25 (a surprise 25 basis point easing). As discussed in Section 2, we bound all other elements of α above and below: we require that (1) the elements of α corresponding to p , y , and y^* are between 0 and 0.05, (2) the element corresponding to i^* is between -0.25 and 0, (3) the element corresponding to $nbrx$ is between 0 and 0.25, and (4) the element corresponding to s is between 0 and 2.5. We therefore require that a surprise loosening of monetary policy cannot lower output (foreign or domestic), prices, or NBRX contemporaneously, that it cannot cause the dollar to appreciate contemporaneously and that it cannot cause foreign interest rates to rise, contemporaneously.

Such assumptions are commonly applied either formally or informally in the literature (e.g., Faust 1998). We also set fairly weak bounds on the magnitude of these contemporaneous effects. We think larger contemporaneous effects are implausible. Recursive identifications make the stronger restriction that there is no contemporaneous effect on variables such as output and prices that are higher in the ordering. While we view our restrictions as quite reasonable, others may disagree. One of the nice features of our approach is that any

TABLE 3. TEST OF THE RANK OF R IN $R\alpha = r$, TEST STATISTIC, AND (P -VALUE)

Null	United Kingdom	Germany
$\rho = 3$	226.62 (0.00)	244.86 (0.00)
$\rho = 4$	98.45 (0.00)	108.54 (0.00)
$\rho = 5$	11.48 (0.32)	8.54 (0.58)
$\rho = 6$	1.40 (0.84)	1.19 (0.88)

Notes: See Appendix A2 for details on this test.

restrictions that are viewed as implausible may be loosened as much as one likes: the cost of removing restrictions is simply wider confidence intervals. We discuss some modifications of this variety next.

We use the results from Table 1 to obtain an estimate of r with an associated variance-covariance matrix. If the matrix R were of Rank 7, then α would be just identified. We test hypotheses about the rank of the matrix R using the method described in Appendix A2. We know that the matrix R has rank of at least 3, since one restriction normalizes the monetary policy shock to lower interest rates by 25 basis points and the contemporaneous effects of the monetary policy shocks on exchange rates and foreign interest rates are also imposed. For both countries, the hypotheses that R has Rank 3 or 4 are clearly rejected (Table 3). The hypotheses that it has Rank 5 or 6 are not rejected. Thus α is not fully identified, and this partial identification means that we will not have any point estimates and must construct confidence intervals as described above.¹⁴

Figures 3 and 4 show pointwise confidence intervals on the impulse response of the variables in the system to the monetary policy shock.¹⁵ These are conservative confidence intervals in the sense that they will have coverage of at least 68 percent asymptotically. In these confidence intervals, we have substantially weakened the restrictions of the recursive identification, allowing simultaneity among all the variables. We have instead achieved identification with restrictions taken from the high-frequency financial market data and with the interval restrictions on α . Since we effectively have only five identifying restrictions and the interval restrictions on α , one might suppose that our confidence intervals will be very wide. In practice, our confidence intervals are quite similar (both in width and shape) to those found for the recursive identification.

Part of the precision comes from our a priori bounds on α . This is necessary because each impulse response is linear in α , which is unidentified in some directions, so that our confidence intervals for impulse responses will typically

14. An intuition for the rank deficiency is that the rows of R associated with the U.S. interest rate response at horizons 3 and 6 are nearly proportional—using the notation in (4), B_{3j} is nearly proportional to B_{6j} where the U.S. interest rate is the j th variable in the VAR. This is not surprising: similarly dated impulse response estimates are typically highly collinear. The same argument holds for the foreign interest rate.

15. As noted above, these take account of both uncertainty in our estimates of α and the reduced form parameters.

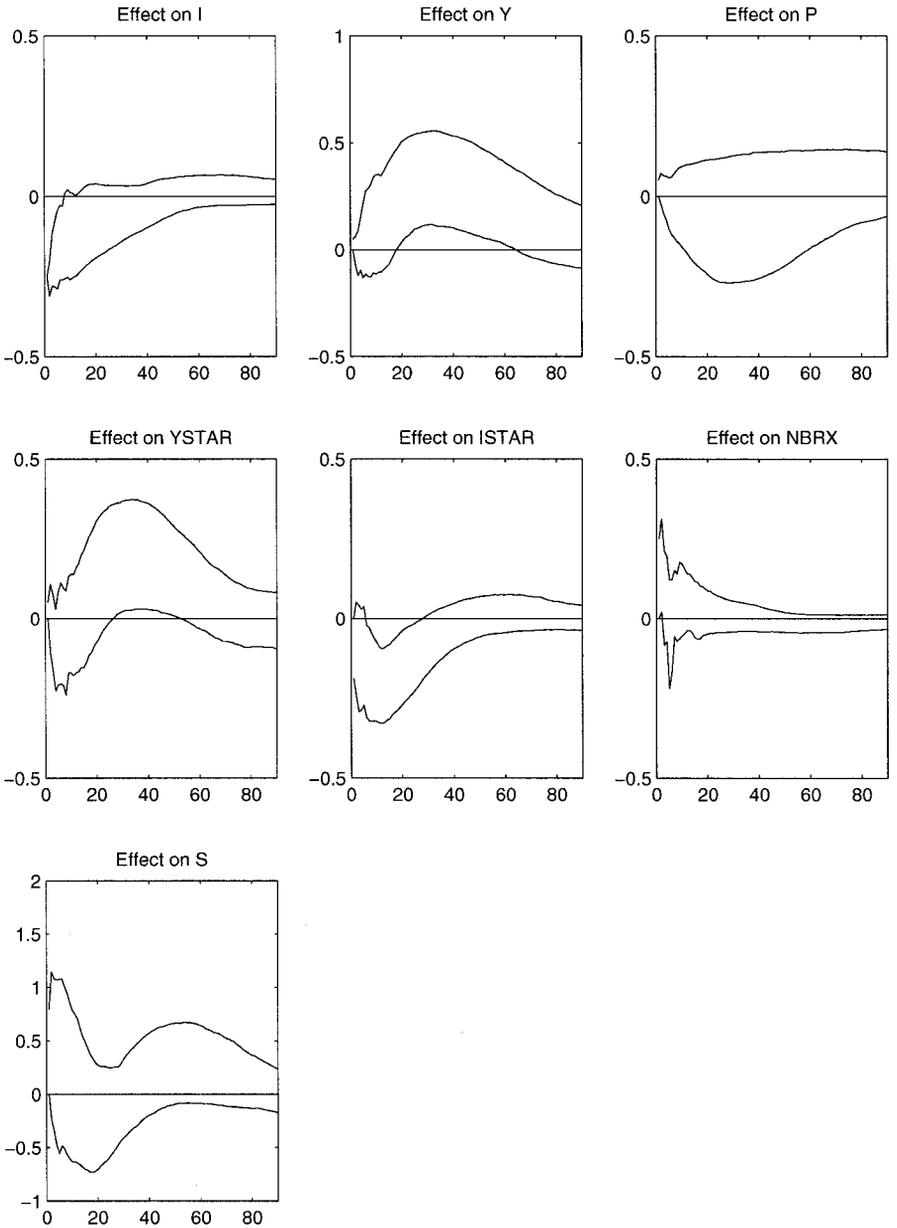


FIGURE 3. New Identification Confidence Intervals for U.K. Impulse Responses

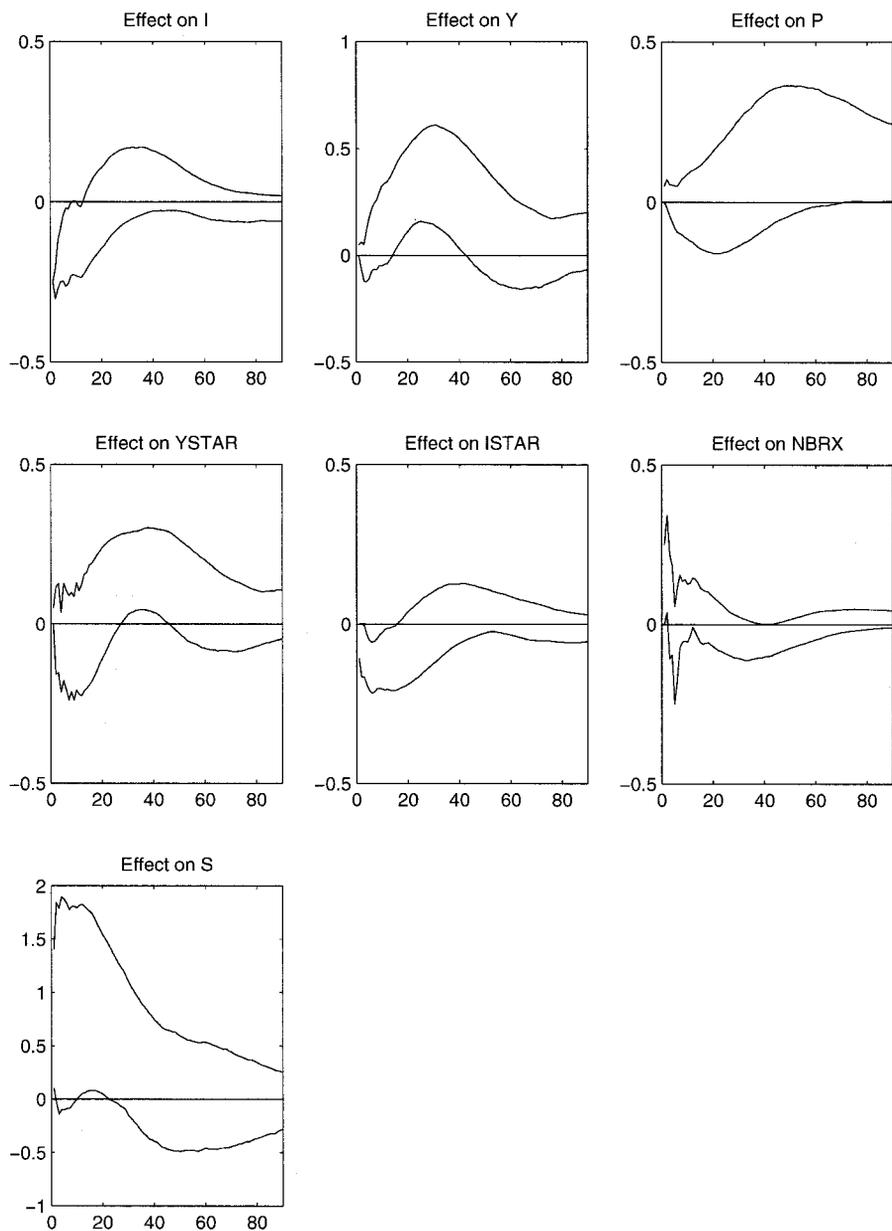


FIGURE 4. New Identification Confidence Intervals for German Impulse Responses

be unbounded otherwise. It is however natural to ask if the high-frequency financial data are contributing anything over and above these bounds on α . We have redone some of the results relying only on these bounds on α , and find that the confidence intervals are substantially wider.¹⁶

While the general character of the impulse response to the policy shock matches the recursive identification, there are some differences. For example, the effect on output is somewhat delayed and somewhat moderated relative to the recursive identification. The effect of the U.S. policy shock on foreign output and interest rates lasts longer than with the recursive identification. The confidence interval for prices is shifted up so that at no horizon is there a pointwise significant fall in prices following the policy loosening.

The confidence intervals for the variance share of the exchange rate due to the policy shock are considerably wider than those from the recursive identification, going from about 0 to 30 percent (Table 2, bottom panel). But they are still considerably narrower than those reported by Faust and Rogers (2003) who drop the strict recursiveness assumption but do not use the financial market data. Thus, while there are other differences among the three approaches, it appears that the very small confidence intervals in the top panel of Table 2 rely on the strict recursiveness assumption. Dropping that, as in Faust and Rogers, leads to the possibility that policy shocks are the main source of exchange rate variation. Adding the restrictions implied by the financial market data, reduces the maximal share to under one-third.

Consistent with Faust and Rogers, we find that the peak timing is not tightly identified. Our confidence interval goes from an immediate peak to a peak at a horizon over five years.¹⁷ Thus, the delayed overshooting found in EE seems to rely on strict recursiveness. For those who find strict recursiveness implausible, further information will have to be brought to bear to further reduce our uncertainty on this point.

The recursive identification, the approach of Faust and Rogers, and the new approach in this paper concur that UIP deviations following money shocks are quite large. The new identification actually narrows the confidence interval some relative to the recursive identification.

Recently, there has been some interest in the possibility that monetary policy loosening represents cost-shocks that could boost aggregate supply and lower prices in the short-run (see, for example, Christiano and Eichenbaum 1992, Christiano, Eichenbaum, and Evans 1997, and Barth and Ramey 2000). In addition, it would be possible to argue that a monetary policy loosening could cause the dollar to appreciate. In order to allow for these possibilities we also

16. The lag zero impulse responses are just the elements of α . For example, the upper bound on our confidence interval for the contemporaneous exchange rate effect, incorporating high-frequency financial data, is about 0.5 percent; our a priori upper bound for this parameter is 2.5 percent.

17. Specifically, here we are referring to confidence intervals on the lag time of the maximum impulse response over lag horizons zero through 90.

considered relaxing our bounds on α to specify that the element of α corresponding to p is between -0.05 and 0.05 and the element corresponding to s is between -2.5 and 2.5 . The results are very similar and our key conclusions emphasized above are not altered by this modification.¹⁸

4.2 Testing the Validity of the Recursive Identification

Our method drops some strong restrictions implied by the recursive identification. The benefit is that we do not have to be concerned about robustness of our results to minor and plausible changes in assumptions such as allowing small simultaneous interactions where recursion imposes no response. The cost is that confidence intervals for some items are quite wide. Thus, it is worth checking whether the recursive identification can be maintained in the face of the information from financial markets.

The α implied by the Eichenbaum and Evans (1995) recursive identification is simply the fifth column of the Cholesky factor of the covariance matrix of the reduced form errors, using their ordering of the variables (in which $nbrx$ comes fifth). For the United Kingdom this choice of α is included in the confidence set A , but for Germany it is not. In other words, the recursive identification is rejected by our identification for Germany but not the UK (the p -values for the test of $R\alpha = r$ are 0.42 and 0.00 for the United Kingdom and Germany, respectively).¹⁹

5. Support for the Identifying Assumptions

Our approach to identification relies on the following principal assumptions.

1. The surprise change in the target rate on FOMC day is a monetary policy shock.
2. The change in the interest rate futures gives an accurate measure of a change in the expectation of future spot rates.

We take up these assumptions in this section.

18. We have also rerun the original exercise but limiting the VAR estimation sample to begin in 1984:02, as a stability check. Once again the results are quite similar. Note that since the high-frequency data start in 1994, that portion of the estimation is unchanged when we estimate the VAR reduced form over the shorter sample.

19. We also tested the recursive ordering in which i^* comes after $nbrx$; the $nbrx$ shock is still the policy shock. This recursive ordering is not rejected for either country (p -values of 0.40 and 0.94 for the United Kingdom and Germany, respectively).

5.1 Is the FOMC Day Surprise Strictly Due to a Monetary Policy Shock?

There are two ways that this assumption could fail. First, other important information could hit the market on the day of announced target changes. Second, the Fed's decision on FOMC day could reveal private information of the Fed about the state of the economy.

The issue of other important information hitting the market on the day of announced target changes could be effectively circumvented by using the change in the Fed Funds futures price from 2:00 P.M. to 2:30 P.M. on FOMC day, instead of the daily close-to-close change. We have intradaily data on the Fed Funds futures prices at 2:00 P.M. and 2:30 P.M. for the last two years of the sample only. The correlation between the change in the target surprise measured using the 2:00 P.M.–2:30 P.M. intradaily data, and using the daily close-to-close data is 0.955. This extremely high correlation would be surprising for other futures markets, but in the case of the Fed Funds rate, we are dealing with a rate relatively tightly targeted by the Federal Reserve. On FOMC days very little happens to change the expectation of this rate over the remainder of the month except the FOMC announcement.

Nonetheless, we checked whether any important pieces of macro data were announced on the day of FOMC meetings. We find that on the sixty-two FOMC days in our sample, durable goods and GDP were released once each, PPI was released twice, industrial production was released three times and CPI was released five times. There were no FOMC meeting days in our sample on which retail sales were released. Deleting the FOMC days on which there is one of these macro releases does not change the regression estimates in Table 1 by much.

The Federal Reserve might, however, have an information advantage through earlier access to data (especially data that are produced by the Federal Reserve, such as industrial production) or through superior economic analysis provided by the Fed's staff economists. In short, the Fed announcement itself might effectively release macroeconomic data. Faust, Swanson, and Wright (2002a, 2002b) test this hypothesis by regressing nine macroeconomic data releases on survey expectations for those data taken before the FOMC meeting, and the FOMC-day target surprise measured from Federal Funds futures contracts. We find that the coefficient on the target surprise is not significantly different from zero at the 1 percent level for any data release, and is significant at the 5 percent level for only one of the nine releases (industrial production). We interpret this as, at most, weak evidence against our assumption that the Fed releases no macro information through the FOMC-day target surprise.

If the FOMC day surprise is strictly due to a monetary policy shock, we would expect FOMC days to be special in some way. If the behavior of the Fed Funds futures price or its association with other variables is not different on these days than other days, it would be difficult to argue that we are capturing a policy shock.

TABLE 4. MEASURES OF THE IMPULSE RESPONSE TO THE NONFARM PAYROLLS DATA SURPRISE

Variable	Horizon	U.S.–U.K. Data		U.S.–German Data	
		Rel. response	St. error	Rel. response	St. error
<i>i</i>	0	–0.25		–0.25	
	3	–0.463	0.058	–0.463	0.058
	6	–0.647	0.088	–0.647	0.088
<i>i</i> *	0	0.036	0.057	0.003	0.028
	3	–0.176	0.057	–0.094	0.029
	6	–0.263	0.077	–0.135	0.043
<i>s</i>	0	0.641	0.181	0.121	0.273

Notes: The results are for a least squares regression of the change in the spot/future interest rate or exchange rate on the unexpected component of the nonfarm payrolls data release, with no intercept, around the data release time. For the U.S.–U.K. and U.S.–German data, the home country is the United States and the foreign country is the United Kingdom and Germany, respectively. The coefficient on the spot U.S. interest rate is normalized to –25 basis points. There are 62 observations; the standard errors are conventional OLS standard errors using the delta method to adjust for the normalization.

In recent years the unexpected component of the FOMC decision on the target rate, measured from the futures market, is usually less than 10 basis points. Although this seems like a small surprise, the standard deviation of the price of the same-month Federal Funds contract is about 3 times larger on FOMC days than on non-FOMC days.²⁰

If we are capturing the response to a monetary policy shock, then we would expect the relative movements of asset prices around a macroeconomic data release, in response to the unexpected component of that data release, to be quite different from our estimated response to policy shocks because macroeconomic data releases are not monetary policy shocks. To investigate this, we regressed the change in the exchange rate from 8:15 A.M. to 8:45 A.M. and the daily changes in spot interest rates and interest rate futures²¹ on the surprise component of the nonfarm payrolls release. This is a data release that comes out at 8:30 A.M. Eastern time, and was found by Fleming and Remolona (1997) to be the most important monthly macroeconomic release. The surprise component is measured as the deviation between the actual nonfarm payrolls data release and the Money Market Services median forecast. We then computed the relative effects of this shock, normalizing the effect on the spot U.S. interest rate to –25 basis points. The results are shown in Table 4, and are the direct analog of the results in Table 1, except using the unexpected components of the nonfarm payrolls data release instead of the unexpected component of the FOMC decision. The relative movements of asset prices in response to the nonfarm payrolls data release are indeed quite different from those in response to FOMC decision. In particular, the effect of the macroeconomic release on U.S. interest

20. The FOMC-day standard deviation is 6 basis points, versus a non-FOMC-day standard deviation of 2.1 basis points. This does not necessarily mean that the monthly monetary policy shock in the VAR is small. Again, our identification assumes that the FFT is purely a monetary policy shock, not that it is the only monetary policy shock.

21. The spot and future interest rate series are the same as described in Subsection 3.2.

TABLE 5. FORECAST EFFICIENCY TESTS FOR INTEREST RATE FUTURES

	$\hat{\alpha}$	$\hat{\beta}$
Eurodollar	-0.16	-0.01
1 quarter ahead	(0.52)	(0.90)
Eurodollar	0.00	-0.08
2 quarters ahead	(1.00)	(0.51)
Sterling LIBOR	-0.08	0.00
1 quarter ahead	(0.94)	(0.97)
Sterling LIBOR	-0.53	0.04
2 quarters ahead	(0.37)	(0.55)
Euribor	-0.12	0.02
1 quarter ahead	(0.32)	(0.37)
Euribor	-0.40	0.06
2 quarters ahead	(0.14)	(0.19)

Notes: These results refer to the standard efficiency test evaluating the forecast of 1 and 2 quarter ahead spot interest rates implicit in interest rate futures markets. The forecast error is regressed on a constant and the forecast. There are four observations per year, corresponding to the settlement days of the interest rate futures contracts. The p -values associated with coefficient estimates are shown in parentheses. For one-quarter-ahead forecasts, conventional OLS standard errors are used. For two-quarter-ahead forecasts, Hansen-Hodrick standard errors are used.

rates builds over time; the effect at a six-month horizon is larger than at a three-month horizon, which is larger than the contemporaneous effect, in contrast to the results in Table 1. Also, the effects on the six-month ahead expected future interest rate in the United Kingdom, and on the sterling exchange rate, are much larger in response to the nonfarm payrolls release than in response to the FOMC decision.

5.2 Do the Interest Rate Futures Data Accurately Reflect Expected Future Spot Rates?

Eurodollar, LIBOR and euromark/euribor futures all settle in the middle of March, June, September and December. We assess the efficiency of the interest rate forecasts from each of these markets as predictors of the actual interest rate on the settlement day²² one or two quarters later by the standard forecast rationality regression. Specifically, we regress the forecast error on a constant and the forecast interest rate. If there is no time-varying term premium, then the slope coefficient should not be significantly different from 0. The results are reported in Table 5. In all cases the hypothesis that the slope coefficient is 0 is not rejected, so that we can think of the term premia in interest rate futures as being time invariant.²³

Interestingly, if we redo this exercise using the forecast of interest rates from the futures market four or eight quarters ahead, then the slope coefficient is significantly below 0. This indicates that the term premia vary over time, and

22. This is implied by the settlement price of the contract.

23. These findings are also consistent with Favero and Mosca's (2001) results that the expectations theory cannot be rejected in the post-1993 data.

may therefore be affected by a monetary policy shock. This, combined with the lower liquidity on longer dated contracts, are the reasons why we do not use future interest rates more than six months ahead.

6. Conclusions

Structural inference about the effects of monetary policy shocks on exchange rates suffers from the normal problems in identifying structural models and more. In the open economy context one must sort out the simultaneous interaction of at least three financial market variables: home and foreign interest rates and the exchange rate. No recursive relation among these variables is very plausible. Nonetheless, various recursive identifications have been proposed and generally plausible answers have emerged from this work.

In this paper, we bring high frequency financial market information to bear in identifying the reaction of financial market variables to a policy shock. Essentially, we require that the impulse response of the VAR match the high frequency response of financial market variables around the time of FOMC announcements. Using this new approach, we find support for the general characteristics of the impulse response of the system to policy shocks.

We find this quite reassuring. We drop all recursiveness assumptions and use instead very different restrictions coming from financial market data. The basic pattern of most of the responses is little changed in the face of large changes in the approach to identification. However, the effect of the U.S. policy shock on foreign output and interest rates lasts longer than with the recursive identification. There is a price puzzle in the recursive identification, which is avoided with the new identification. With specific regard to the exchange rate response, our results are between those of Eichenbaum and Evans (1995) and Faust and Rogers (2003). We find that the peak timing of the exchange rate effect is quite imprecisely estimated: it may come nearly immediately as in Dornbusch overshooting or come several years later. The estimated variance share of exchange rate movements due to the policy shock—bounded at about $\frac{1}{3}$ —is between the Eichenbaum and Evans and Faust and Rogers estimates. Like both previous studies, we find added support for the view that policy shocks generate large UIP deviations.

Appendices

A1 Partial Identification

Here we describe how to construct the confidence set A for the vector α when the restrictions $R\alpha = r$ must be satisfied, R is estimated by \hat{R} , r is estimated by

\hat{r} , R may be rank deficient, $T^{1/2}(\text{vec}(\hat{R}) - \text{vec}(R)) \rightarrow_d N(0, V_R)$ and $T^{1/2}(\hat{r} - r) \rightarrow_d N(0, V_r)$. Consider the GMM objective function

$$S(\alpha) = T(\hat{R}\alpha - r)'[(\alpha \quad \mathbb{E} \quad I_K)\hat{V}_R(\alpha' \quad \mathbb{E} \quad I_K) + \hat{V}_r]^{-1}(\hat{R}\alpha - r).$$

In standard GMM terminology, this is the continuous updating GMM objective function. The estimator $\hat{\alpha}$ that minimizes this objective function is not consistent for the true α because of the rank deficiency of the matrix R . However $S(\alpha_0)$ has a χ^2 null distribution regardless of the rank of R where α_0 denotes the true value of the vector α . Define the confidence set

$$A = \{\alpha \in A^+ : S(\alpha) \leq F_{\chi^2}\}$$

where F_{χ^2} denotes the 95th percentile of a χ^2 distribution (degrees of freedom equal to the number of elements in r) and A^+ is our parameter space for α , each element of which is constrained to lie between a lower bound and an upper bound.

The use of such confidence sets in models that are not fully identified was proposed by Stock and Wright (2000), where they are referred to as S-sets. If the matrix R is rank deficient, then there exists a subspace of vectors α that are observationally equivalent to α_0 . Any vector in this subspace must be included in A with probability 95 percent, asymptotically. Any other vector α will be excluded from A with probability 1, asymptotically. This is a correct statement of what we do and do not know about α , when R is rank deficient. More formally, the confidence set A is unbounded with probability 0.95, asymptotically: this must be the case for any confidence set for an unidentified parameter if the confidence set is to have 95 percent asymptotic coverage uniformly in the parameter space (Dufour 1997).

Concretely, we proceed by forming a grid with 20 million points in A^+ . For each point in this grid, we calculate the objective function $S(\alpha)$. If this is above the critical value, we compute nothing else and simply proceed to the next point in the grid. If the objective function $S(\alpha)$ is below the critical value, we include that value of α in the confidence set A . For each such α , we then compute the lower and upper bounds of the bootstrap confidence intervals for all the parameters of interest (notably variance shares and impulse responses) conditional on that α . Each bootstrap replication involves calculating a new θ from the bootstrap sample, while holding α fixed—the confidence interval then consists of the $(100 - v_2)/2$ and $(100 + v_2)/2$ percentiles of the parameters over a total of 500 bootstrap replications. Having cycled through all points in the grid, our confidence intervals for the objects of interest are given by the smallest and largest values of these percentiles, respectively. This completes the algorithm we use, for which MATLAB code is available from the authors, on request.

For each point α in the grid such that $S(\alpha)$ is below the critical value, we have used the Runkle (1987) bootstrap to form a confidence interval for the parameters of interest, given that θ is uncertain. We could instead take draws

from the posterior for θ that corresponds to the RATS prior and interpret the percentiles of this posterior as a classical confidence interval. This turns out to give similar results with the model and data considered in this paper. Implementing the bias-adjusted bootstrap of Kilian (1998) pushes the largest root of the VAR inside the unit circle except that Kilian's algorithm then calls for the bias-adjustment to be scaled back so as to induce a unit root but not an explosive root. We are not aware of evidence that the bias-adjustment works well under these specific circumstances, even though it too gives similar results except at long horizons. We accordingly report results using just the bootstrap.

A2 Testing the Rank of R

We wish to test the hypothesis that $\rho(R) = L$ against the alternative that $\rho(R) > L$, where $\rho(\cdot)$ denotes the rank of the argument. Assume that $T^{1/2}(\hat{\theta} - \theta) \rightarrow_d N(0, V_\theta)$. See Hamilton (1994) for primitive conditions for these convergence results and \hat{V}_θ , a consistent estimator of V_θ . The matrix R is a nonlinear function of θ and can be estimated by \hat{R} , where this denotes this same nonlinear function of $\hat{\theta}$. By the delta method, $T^{1/2}(\text{vec}(\hat{R}) - \text{vec}(R)) \rightarrow_d N(0, V_R)$ where $V_R = (d\text{vec}(R)'/d\theta) V_\theta (d\text{vec}(R)/d\theta)$.

To test the hypothesis about that rank of R , we use the test statistic

$$T \min_{P \in \pi(L)} (\text{vec}(\hat{R}) - \text{vec}(P))' \hat{V}_R^{-1} (\text{vec}(\hat{R}) - \text{vec}(P))$$

where \hat{V}_R is $(d\text{vec}(\hat{R})'/d\theta) \hat{V}_\theta (d\text{vec}(\hat{R})/d\theta)$ and $\pi(L)$ is the space of all conformable matrices of rank L . By Theorem 1 of Cragg and Donald (1997), under the null hypothesis, this test statistic has a χ^2 null limiting distribution.

A3 Data

High-Frequency Data: The spot and futures interest rate data were acquired from Datastream and CBOT and consist of daily closing prices, as described in the text. The exchange rate data consist of 2 P.M. and 2:30 P.M. Eastern Time quotes (midpoint of bid and ask) obtained from Olsen and Associates (Zurich, Switzerland).

VAR Data. The data were acquired from the Federal Reserve Board's International Finance and the IMF's International Financial Statistics (IFS) databases. All series are expressed in natural logarithms except interest rates, which are expressed in percentage points. The series definitions are as follows:

$y(y^*)$ = index of U.S. (foreign) industrial production;

p = U.S. CPI—all urban, all items;

nbr = nonborrowed reserves plus extended credit, seasonally adjusted, monthly average;

- tr = total reserves, seasonally adjusted, monthly average;
 $nbrx = nbr/tr$;
 s = spot exchange rate; monthly average; US\$/foreign currency;
 i, i^* = for the United States, 90-day T-bill rate, monthly average (line 60c, IFS); for the United Kingdom, 90-day T-bill rate, monthly average (line 60cs, IFS), for Germany, 90-day Fibor/Euribor rate.

References

- Bagliano, Fabio, and Carlo A. Favero (1999). "Information from Financial Markets and VAR Measures of Monetary Policy." *European Economic Review*, 43, pp. 825–837.
- Barth, Marvin J. and Valerie A. Ramey (2000). "The Cost Channel of Monetary Transmission." Working paper, University of California at San Diego.
- Christiano, Lawrence J. and Martin Eichenbaum (1992). "Liquidity Effects and the Monetary Transmission Mechanism." *American Economic Review*, 82, pp. 346–353.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (1997). "Sticky Price and Limited Participation Models of Money: A Comparison." *European Economic Review*, 41, pp. 1201–1249.
- Cochrane, John H. and Monika Piazzesi (2002). "The Fed and Interest Rates: A High Frequency Identification." *American Economic Review*, 92, pp. 90–95.
- Cragg, John G. and Stephen G. Donald (1997). "Inferring the Rank of a Matrix." *Journal of Econometrics*, 76, pp. 223–250.
- Cushman, David O., and Tao Zha (1997). "Identifying Monetary Policy in a Small Open Economy under Flexible Exchange Rates." *Journal of Monetary Economics*, 39, pp. 433–448.
- Dornbusch, Rudiger (1976). "Expectations and Exchange Dynamics." *Journal of Political Economy*, 84, pp. 1161–1176.
- Dufour, Jean-Marie (1997). "Some Impossibility Theorems in Econometrics with Applications to Structural and Dynamic Models." *Econometrica*, 65, pp. 1365–1387.
- Eichenbaum, Martin and Charles L. Evans (1995). "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates." *Quarterly Journal of Economics*, 110, pp. 975–1010.
- Engel, Charles (1996). "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence." *Journal of Empirical Finance*, 3, pp. 123–192.
- Faust, Jon (1998). "The Robustness of Identified VAR Conclusions about Money." *Carnegie-Rochester Series on Public Policy*, 49, pp. 207–244.
- Faust, Jon and John H. Rogers (2003). "Monetary Policy's Role in Exchange Rate Behavior." *Journal of Monetary Economics*, forthcoming.
- Faust, Jon, Eric T. Swanson, and Jonathan H. Wright (2002a). "Identifying VARs Based on High Frequency Futures Data." International Finance Discussion Paper 720, Federal Reserve Board of Governors, Washington, D.C.
- Faust, Jon, Eric T. Swanson, and Jonathan H. Wright (2002b). "Do Federal Reserve Policy Surprises Reveal Inside Information About the Economy?," mimeo. Washington, D.C.: Federal Reserve Board.
- Favero, Carlo A. and Federico Mosca (2001). "Uncertainty on Monetary Policy and the Expectations Model of the Term Structure of Interest Rates." *Economics Letters*, 71, pp. 369–375.
- Fleming, Michael and Eli Remolona (1997). "What Moves the Bond Market?" *Federal Reserve Bank of New York Economic Policy Review*, 3, pp. 31–50.
- Hamilton, James D. (1994). *Time Series Analysis*. Princeton, New Jersey: Princeton University Press.

- Killian, Lutz (1998). "Small Sample Confidence Intervals for Impulse Responses." *Review of Economics and Statistics*, 80, pp. 218–230.
- Kim, Soyoung (2001). "International Transmission of U.S. Monetary Policy Shocks: Evidence from VARs." *Journal of Monetary Economics*, 48, pp. 339–372.
- Kim, Soyoung and Nouriel Roubini (2000). "Exchange Rate Anomalies in the Industrial Countries: A Solution with a Structural VAR Approach." *Journal of Monetary Economics*, 45, pp. 561–586.
- Kuttner, Kenneth N. (2001). "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics*, 47, pp. 523–544.
- Leeper, Eric M., Christopher A. Sims, and Tao Zha (1996). "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity*, pp. 1–63. Washington, D.C.: Brookings Institute.
- Runkle, David E. (1987). "Vector Autoregressions and Reality." *Journal of Business and Economic Statistics*, 5, pp. 437–442.
- Stock, James H. and Jonathan H. Wright (2000). "GMM with Weak Identification." *Econometrica*, 68, pp. 1055–1096.