Volume 2, Number 4

December 2006

IJCB

INTERNATIONAL JOURNAL OF CENTRAL BANKING

IJCB International Journal of Central Banking

Volume 2, Number 4 December 2006

The contents of this journal, together with additional materials provided by article authors, are available without charge at ***www.ijcb.org****.*

The *International Journal of Central Banking* is published quarterly (ISSN: 1815-4654). Online access to the publication is available free of charge at **www.ijcb.org**. Individual print subscriptions are available at an annual rate of $100 (USD).

Print subscription orders may be placed online at www.ijcb.org, by phone (202-452-3245), via fax (202-728-5886), or by e-mail (editor@ijcb.org).

Renewals, claims, address changes, and requests for permission to reprint material from this journal should be addressed to:

International Journal of Central Banking
Publications Fulfillment, MS-127
Federal Reserve Board
Washington, DC 20551

Phone: 202-452-3245
Fax: 202-728-5886
E-mail: editor@ijcb.org

ISSN: 1815-4654

Basel II and the Risk Management of Basket Options with Time-Varying Correlations*

Amy S. K. Wong
Tinbergen Institute
Erasmus University Rotterdam

The impact of jumps, regime switches, and linearly changing correlation term structures on the risk management of basket options has been examined in this paper. First, the results show that there is an asymmetric correlation effect on the value-at-risk of basket options. Second, the time at which a correlation shock occurs during the life of an option is particularly important for hedged basket options. Finally, the square-root-of-time rule can lead to severe underestimation of value-at-risk for basket options with time-varying correlations—for some cases, even by a factor exceeding the minimum regulatory stress factor.

JEL Codes: G15, G21.

1. Introduction

An important aspect of the "International Convergence of Capital Measurement and Capital Standards: A Revised Framework," also known as Basel II, is the management of the risk of financial products that cannot be entirely captured by value-at-risk (VaR); for example, nonlinear products or sudden correlation shifts (see Bank for International Settlements 2005). Basket options are derivatives that belong to the class of products that are subject to nonlinear and correlation risk. A basket option is an option on a portfolio of underlying assets, and the option price is highly dependent on the correlations between the underlying assets. This study examines

*I would like to thank the co-editor and anonymous referee for helpful comments. The usual disclaimer applies. Author contact: Tinbergen Institute, Erasmus University Rotterdam, The Netherlands; Home page: www.askwong.com.

basket options, because they are widely used across many financial markets, such as foreign exchange markets (see, e.g., Bennett and Kennedy 2004, Dammers and McCauley 2006), credit derivatives markets (e.g., Duffie and Singleton 2003), and equity markets (e.g., Pellizzari 2005). Many studies, such as Margrabe (1978), Curran (1994), Milevsky and Posner (1998), Brigo et al. (2004), and Deelstra, Liinev, and Vanmaele (2004), value basket options under the assumption of constant correlations between the processes of the underlying assets. However, recent empirical studies (e.g., Goetzmann, Li, and Rouwenhorst 2005), have shown that correlations of stock returns are considerably time varying.

Therefore, this paper examines the impact of time-varying correlation term structures on pricing and hedging of basket options as well as the implications for risk management. Empirical examination of the correlations between equity indices S&P 500, FTSE 100, and the Merrill Lynch government bond index shows that empirical features such as jumps, regime switches, and (nearly) linearly changing correlations can occur in practice. The main contribution of this study is to take these features into account in the correlation term structures of the basket options. To my knowledge, the impact of correlation jumps, regime switches, and linear correlation changes on the value-at-risk of basket options and the performance of the square-root-of-time rule have not yet been examined in the literature. Studies such as Skintzi, Skiadopoulos, and Refenes (2005) have analyzed the impact of estimation errors of constant correlation on value-at-risk of a portfolio of standard European options. Pellizzari (2005) examines a linearly increasing volatility structure for hedged basket options and the corresponding risk measures but does not look at time variation of correlations. Kupiec (1998) performs stress tests taking time-varying correlations into account for portfolios with linear exposure to underlying assets.

A second contribution of this study is an analysis on VaR of hedged basket options and the performance of the square-root-of-time rule for these positions when the correlation term structure contains the above-mentioned time variation. In practice, financial institutions often hedge their outstanding option positions to reduce the exposure to risk of the position and apply the square-root-of-time rule to the one-day VaR in order to obtain an estimate of the regulatory ten-day VaR. The performance of the square-root-of-time rule has been studied, for example, for GARCH processes by Diebold et al. (1997) and for jump diffusion processes by Danielsson and Zigrand (2005).

Moreover, this paper discusses the differences between the risk measure VaR and the coherent risk measure CVaR (conditional value at risk) for basket options, where the latter (more robust) measure can give additional information needed in some cases to give an adequate risk assessment of the derivatives position.

In this paper, a Monte Carlo simulation study has been performed for basket options with time-varying correlation term structures against the benchmark of constant correlations. The results are as follows.

First, there is an asymmetric correlation effect on the VaR of the basket option, where a change in negative (constant) correlations between the underlying assets has a greater impact on the VaR than a change in positive correlations of the same magnitude. This result is surprising: it is widely known that well-diversified asset portfolios (i.e., negative correlations) are less risky than portfolios with highly correlated assets, so one might expect that a basket option on a well-diversified portfolio is also relatively less risky. The results show that the potential loss *values* given by (C)VaR are indeed lower for basket options on negatively correlated assets, but at the same time the *changes* in (C)VaR are more subject to correlation risk as well. Ignoring this result can lead to serious underestimation of the VaR if sudden changes in market conditions occur. Another implication of this result is that VaR estimation of basket options on well-diversified baskets is relatively more prone to model risk, since in practice correlations have to be estimated under the assumption of a certain correlation model such as RiskMetrics^{TM}. For studies on the impact of estimation errors of constant correlations, see Fengler and Schwendner (2004) for basket options and Skintzi, Skiadopoulos, and Refenes (2005).

Second, the time at which correlation shocks occur during the life of the option is important for the VaR of basket options (especially if hedging is applied), even though the payoff of the basket option only depends on the value of the underlying assets at maturity. Compared with constant correlations at the average value of the time-varying correlations over the life of the option as a benchmark, the results show that the VaR estimates are dependent on the specific type of correlation time variation, even if the average correlation over the life of the option is the same. The estimates for the risk measures are also highly dependent on the hedge effectiveness for the option. This result is relevant for financial institutions, because option positions are often hedged in practice.

Third, the VaR estimate obtained by the square-root-of-time rule can lead to underestimation of the ten-day VaR for the unhedged option with time-varying correlations. The risk assessment of deep out-of-the-money derivatives plays an important role in the Basel II framework. This study shows that the square-root-of-time rule for the risk measure VaR underestimates the risks when the OTM basket option has a (time-varying) highly negatively correlated portfolio, and this underestimation even exceeds the minimum regulatory stress factor of value 3 for some cases. There is also underestimation of the VaR for basket options with constant correlations, but this underestimation remains below a factor of 3. When the time-varying correlations are relatively low at the start of the option, by using the square-root-of-time rule one implicitly assumes that the correlations remain this low for ten days. As a result, the risk implied by the ten-day VaR can be much larger than the estimate obtained from the square-root-of-time rule for time-varying correlations. For hedged options, the performance of the square-root-of-time rule is highly dependent on the difference in hedge effectiveness over a one-day and ten-day horizon and can lead to large deviations from the ten-day VaR.

Finally, VaR gives information about the potential loss of the option position corresponding to a certain confidence level but does not reveal the size of the loss if the VaR is exceeded. The difference in VaR and CVaR can be more than 40 percent for certain correlation term structures. Therefore, it is advisable to use VaR together with the coherent risk measure CVaR, because the CVaR can provide additional information needed to assess the risk of the basket options.

This paper is structured as follows. Section 2 contains an empirical examination of the correlations between equity and bond indices. Section 3 proceeds with the simulation framework, and in Section 4 the results of the simulation study are discussed. Finally, Section 5 concludes.

Figure 1. Correlations Generated by RiskMetrics^{TM} Model

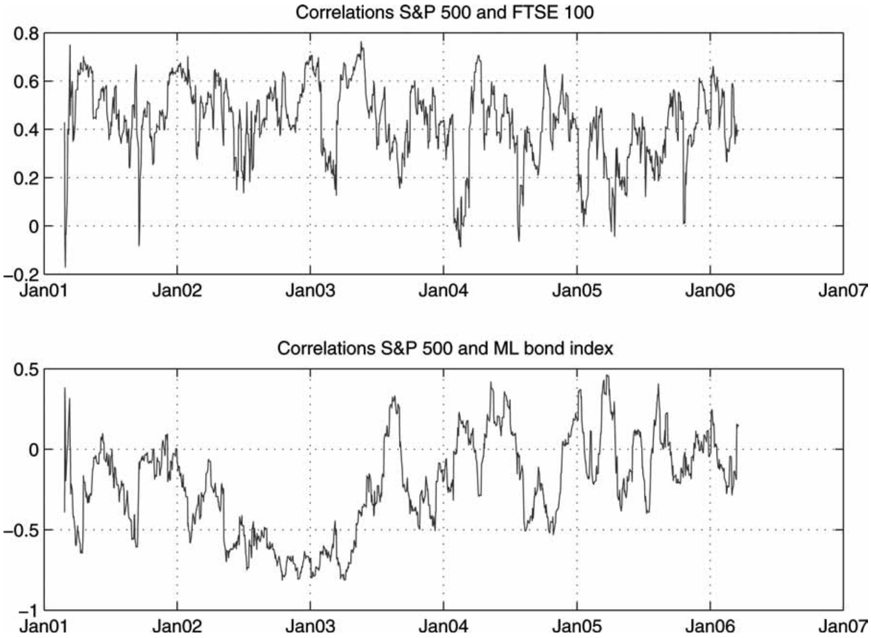

2. Empirical Motivation

This study examines the impact of time-varying correlations on the basket option and the implications for risk management. In order to see in what way this time variation could occur, the correlations between equity indexes S&P 500 (United States) and FTSE 100 (United Kingdom) as well as the Merrill Lynch U.S. Treasury one- to ten-year bond index (henceforth, ML bond index) are illustrated in figure 1. The data are from October 2000 to March 2006 and are obtained from Datastream. The correlations between the returns of N assets can be estimated in many different ways, but a widely used benchmark model is RiskMetrics^{TM} of J. P. Morgan (1996), specified as follows. Let r_t be the $1 \times N$ vector of asset returns at time t and $H_t = \{\sigma_{ij,t}\}_{i,j=1}^{N}$ be the conditional covariance matrix at time t. Then, RiskMetrics^{TM} estimates the covariance matrix by

$$H_t = \lambda H_{t-1} + (1 - \lambda) r'_{t-1} r_{t-1},$$

where the weighting parameter λ has value 0.94 for daily data.[1] Subsequently, the correlations can be easily calculated as $\rho_{ij,t} = \frac{\sigma_{ij,t}}{\sqrt{\sigma_{ii,t}\sigma_{jj,t}}}$.

Figure 1 shows that the correlations between S&P 500 and FTSE 100 as well as the correlations between S&P 500 and the ML bond index, are considerably time varying. The correlations between S&P 500 and FTSE 100 remain positive during most of the sample period, but there are several large jumps downward—for example, in 2005. The correlations between S&P 500 and the ML bond index are, during some periods, positive (the value stays below 0.5) but remain negative most of the time. One can also observe the following correlation patterns in figure 1: starting from the second quarter of 2003, the correlation appears to increase almost linearly. Moreover, one can see from the figure that the correlations from the second quarter of 2002 until the first quarter of 2003 are on average much more negative than the correlations before and after this period, which resembles a regime switch. Following these observations, the correlation term structures in the subsequent simulation study will contain jumps and regime switches as well as linear increase and decrease in correlations during the life of the option.

3. Simulation Methods

For simplicity, we will use a basket option on two underlying assets to analyze the impact of the correlation changes on basket options. The asset price dynamics will be simulated using the differential equations

$$\begin{aligned} dS_{1,t} &= \mu_1 S_{1,t}dt + \sigma_1 S_{1,t}dW_{1,t}, \\ dS_{2,t} &= \mu_2 S_{2,t}dt + \sigma_2 S_{2,t}dW_{2,t}, \\ dW_{1,t}dW_{2,t} &= \rho_t dt, \qquad \text{for } t = 1, \dots, T, \end{aligned} \tag{1}$$

where $W_{1,t}$ and $W_{2,t}$ are correlated Brownian motions with correlation ρ_t at time t. The value V_t of the basket option with strike price K

[1]The first 100 daily observations of the data sample are used to estimate the initial covariance matrix.

at time t is given by

$$C_T = (\omega_1 S_{1,T} + \omega_2 S_{2,T} - K)^+$$
$$= \max(\omega_1 S_{1,T} + \omega_2 S_{2,T} - K, 0), \tag{2}$$
$$V_t = e^{-r(T-t)} \mathbb{E}_Q \left[\sum_{m=1}^{M} C_T^{(m)} / M \right], \tag{3}$$

where C_T is the claim of the option at maturity T, ω_i are the portfolio weights of the underlying assets ($i = 1, 2$), Q is the risk-neutral martingale measure, and M is the number of simulations. The delta, Δ_i, is the sensitivity of the option value with respect to the price of asset i (used for delta hedging), and it is computed by the central difference method (for more details, see, e.g., Glasserman 2004):

$$\Delta_i(S_i) = \frac{\partial V}{\partial S_i} \approx \frac{V(S_i + h) - V(S_i - h)}{2h}, \quad i = 1, 2. \tag{4}$$

We will analyze the different correlation specifications for in-the-money (ITM), at-the-money (ATM), and out-of-the-money (OTM) basket options. Delta hedging will be done in the standard way (see, e.g., Hull 2006) with daily rebalancing. Each basket option contract is written on 100,000 underlying shares and has a maturity of three months in trading days ($T = \frac{63}{252}$). The simulated asset price paths and option values are computed using the parameters given in table 1 and the following correlation term structures for $t = \frac{1}{252}, \frac{2}{252}, \ldots, T$.

Table 1. Option Parameters

VaR horizon	10/252	$[\mu_1, \mu_2]$	$[0.1, 0.1]$
Maturity option T	63/252	r	0.05
dt	1/252	K_{ITM}	95
Nr. of assets	2	K_{ATM}	100
$[S_{1,0}, S_{2,0}]$	$[100, 100]$	K_{OTM}	105
$[\sigma_1, \sigma_2]$	$[0.35, 0.35]$	h in (4)	0.01
$[\omega_1, \omega_2]$	$[1/2, 1/2]$	M	5000

Constant correlations:

$$\text{C1 to C9:} \quad \rho_t = \rho,$$

$$\text{where } \rho \in \{-0.9, -0.7, -0.5, -0.2, 0, 0.2, 0.5, 0.7, 0.9\}$$

Negative correlations jump upward:

$$\text{T1:} \quad \rho_t = -0.9 + \left[0.9\left(t - \frac{1}{63}T\right)252\right]\mathbf{I}_{\left(\frac{1}{63}T < t \leq \frac{3}{63}T\right)} + \left[1.8 - 0.9\left(t - \frac{3}{63}T\right)252\right]\mathbf{I}_{\left(\frac{3}{63}T < t \leq \frac{5}{63}T\right)}$$

$$\text{T2:} \quad \rho_t = -0.9 + \left[0.9\left(t - \frac{59}{63}T\right)252\right]\mathbf{I}_{\left(\frac{59}{63}T < t \leq \frac{61}{63}T\right)} + \left[1.8 - 0.9\left(t - \frac{61}{63}T\right)252\right]\mathbf{I}_{\left(\frac{61}{63}T < t \leq T\right)}$$

Positive correlations jump downward:

$$\text{T3:} \quad \rho_t = 0.9 - 0.9\left[\left(t - \frac{1}{63}T\right)252\right]\mathbf{I}_{\left(\frac{1}{63}T < t \leq \frac{3}{63}T\right)} - \left[1.8 - 0.9\left(t - \frac{3}{63}T\right)252\right]\mathbf{I}_{\left(\frac{3}{63}T < t \leq \frac{5}{63}T\right)}$$

$$\text{T4:} \quad \rho_t = 0.9 - 0.9\left[\left(t - \frac{59}{63}T\right)252\right]\mathbf{I}_{\left(\frac{59}{63}T < t \leq \frac{61}{63}T\right)} - \left[1.8 - 0.9\left(t - \frac{61}{63}T\right)252\right]\mathbf{I}_{\left(\frac{61}{63}T < t \leq T\right)}$$

Correlation regime shifts:

$$\text{T5:} \quad \rho_t = -0.9\mathbf{I}_{\left(t \leq \frac{30}{63}T\right)} + -0.5\mathbf{I}_{\left(\frac{30}{63}T < t \leq T\right)}$$

$$\text{T6:} \quad \rho_t = 0.9\mathbf{I}_{\left(t \leq \frac{30}{63}T\right)} + 0.5\mathbf{I}_{\left(\frac{30}{63}T < t \leq T\right)}$$

Affine correlation term structure:

$$\text{T7:} \quad \rho_t = -0.9 + \frac{1.8}{62}(252t - 1)$$

$$\text{T8:} \quad \rho_t = 0.9 - \frac{1.8}{62}(252t - 1)$$

Figure 2. True Time-Varying Correlation Term Structures

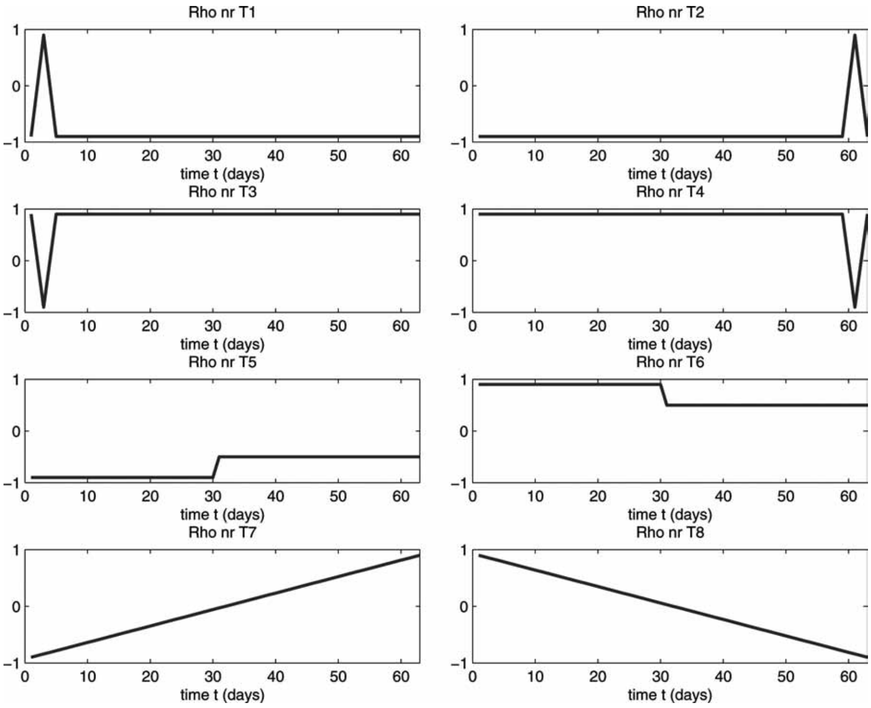

Figure 2 illustrates the time-varying correlation term structures, which can be interpreted as follows. The correlation term structures T1 and T2 correspond to a situation where the assets in the basket are usually well diversified (i.e., negative correlations), but suddenly the correlations jump upward at, respectively, the start and the end of the option life. This could happen, for example, during a financial crisis, when all stocks plummet simultaneously and the correlations between these stocks suddenly jump upward. Correlation specifications T3 and T4 correspond to a situation where the correlation is highly positive, but they suddenly jump downward at, respectively, the start and the end of the option life. T5 and T6 are regime-switching types of specifications, where the correlations shift, respectively, upward or downward, but are locally constant before and after the shift. Finally, T7 corresponds to a gradually linear increase of the correlations from −0.9 to +0.9, often associated with a bear market in which stocks become more correlated over time. T8 represents a linear decrease in correlations, which can be interpreted as a period in a bull market.

Eydeland and Wolyniec (2003) distinguish the instantaneous correlation ρ_t in (1) from the cumulative correlation defined by them as

$$\rho^*_{T,t} = \rho_{\ln S_{1,T} \ln S_{2,T}} = \frac{\mathbb{E}_t[\ln S_{1,T} \ln S_{2,T}] - \mathbb{E}_t[\ln S_{1,T}]\mathbb{E}_t[\ln S_{2,T}]}{\sqrt{\mathrm{var}_t[\ln S_{1,T}]}\sqrt{\mathrm{var}_t[\ln S_{2,T}]}} \tag{5}$$

$$= \frac{\int_t^T \sigma_{1,s}\sigma_{2,s}\rho_s ds}{\sqrt{\int_t^T \sigma^2_{1,s} ds}\sqrt{\int_t^T \sigma^2_{2,s} ds}}, \tag{6}$$

with the following properties

$$\lim_{t \to T} \rho^*_{T,t} = \rho_T \tag{7}$$

$$\text{if} \quad \sigma_{1,t} = \sigma_{2,t} = \sigma, \quad \forall t = 1, \cdots, T \quad \text{then} \quad \rho^*_{T,t} = \frac{\int_t^T \rho_s ds}{T - t}. \tag{8}$$

According to Eydeland and Wolyniec (2003), the cumulative correlation $\rho^*_{T,t}$ is more important than the instantaneous correlation ρ_t for option pricing and hedging as well as for VaR computation. When the VaR horizon is very short, by (7) the cumulative correlation becomes close to the instantaneous correlation. By (8) the cumulative correlation boils down to the average correlation over time period $T - t$, if the volatilities of both assets are the same and constant over time. So, in this special case one needs only an estimate of the average value of the correlation term structure instead of information on the entire term structure to compute the VaR of the option. This is a great simplification, and the results in the next section show to what extent this property holds for different time-varying correlation term structures.

Empirical correlations of asset returns are unobserved and have to be estimated using variance-covariance models, so it is very likely that the resulting correlations contain estimation errors. In this paper the focus is on the effects of changes in the true correlation values instead of correlation estimation errors. For this purpose, the same parameter values will be used both to generate the asset price data and to value the basket options. Hence, the effects of changes in the true correlations can be isolated without bothering about estimation errors of correlations (for more details on correlation estimation errors, see Fengler and Schwendner 2004).

3.1 Value-at-Risk

In the current Basel II framework, banks should develop a more forward-looking approach with respect to risk management. The risks of the basket option position are assessed using the risk measure value-at-risk. Value-at-risk (VaR) is a widely used measure for quantifying potential losses of asset portfolios at a certain confidence level α (conventionally, 95 percent or 99 percent). Let X be the profit and loss realizations of the simulations; then VaR is defined by

$$\mathbb{P}(X \leq VaR^{\alpha}) = 1 - \alpha. \tag{9}$$

The main methods of VaR computation are delta-normal, historical simulation, and Monte Carlo simulation (see, e.g., Jorion 2001). To determine which method is suitable, it is important to look at the characteristics of the portfolio for which the VaR needs to be computed. The asset portfolio here consists of a basket option with (time-dependent) correlations between the underlying assets; therefore, the risks involved in this asset position are highly nonlinear as well as time dependent. The Monte Carlo simulation method is used to compute the VaR in this study. The motivation for the choice of the Monte Carlo simulation method is that it can take time dependency and nonlinear risk into account. Moreover, with respect to the requirements of Basel II, the computation of VaR can be easily adapted to reflect current market conditions by changing the underlying parameters such that these coincide with market parameters.

The alternative methods are more often used by market practitioners, since they are less computationally intensive than Monte Carlo simulation. However, the delta-normal method and the historical simulation method are less adequate for computing the VaR of the basket options for the following reasons. The delta-normal method is based on the assumption of a normally distributed portfolio, thus it is only suitable for portfolios involving linear risk in the underlyings. The historical simulation method is often used by market practitioners, since this method does not involve any distributional assumptions of the portfolio and is relatively fast to compute. Historical returns are used to represent potential future losses, and this implies that all information of the historical data is retained. However, the resulting VaR estimate can only reflect the type of losses that already have occurred in the historical data sample used for estimation. Hence, this approach is not in line with the forward-looking approach required in the Basel II framework. Moreover, due to the use of a relatively large number of historical observations to avoid small-sample bias, the most recent market movements cannot be easily captured in the VaR computation. For more details on the historical simulation method, see Pritsker (2006) and the references therein.

3.2 CVaR and the Square-Root-of-Time Rule

This study will assess the risks of the option portfolio using VaR, because this risk measure is widely used for regulatory purposes as specified by the Basel Committee on Banking Supervision. However, it is well known that VaR is only an indication of the loss corresponding to a confidence level α and over a certain time horizon (usually ten days); it does not give an indication of the size of that loss if the VaR is exceeded. Moreover, it is not a coherent risk measure (for more details, see Artzner et al. 1999). Therefore, the conditional value-at-risk (CVaR) will also be computed, which satisfies the coherence properties. Let X be the profit and loss realizations of the simulations. The CVaR (also called expected shortfall) is defined for a given confidence level α as

$$CVaR^{\alpha} = \mathbb{E}[X|X \leq VaR^{\alpha}]. \tag{10}$$

The difference between the CVaR and the VaR can be expressed by the CVaR-to-VaR ratio. McNeil, Frey, and Embrechts (2005) describe that for the CVaR-to-VaR ratio of the normal distribution, it holds that $\lim_{\alpha\to 1} \frac{\mathrm{CVaR}^{\alpha}}{\mathrm{VaR}^{\alpha}} = 1$, whereas for the t-distribution, this ratio will go to a value greater than 1 in the limit. Hence, for a heavy-tailed distribution, the difference between the CVaR and the VaR will be larger than for a normal distribution. The 1996 Amendment of the 1988 Basel Capital Accord requires a capital charge for market risk (see, e.g., Jorion 2001 and Hull 2006 for more details). The market risk capital charge by the internal models approach can be computed as the maximum of the VaR of the previous day and the average VaR over the previous sixty days times a stress factor k. The stress factor k is determined by the local regulators and has at least an absolute value of 3. In practice, the square-root-of-time rule is often applied to compute the ten-day VaR and CVaR by multiplying the one-day VaR with $\sqrt{10}$, since the direct estimation of the ten-day VaR often requires too much historical data. This approach is valid for i.i.d. normally distributed observations; otherwise, the square-root-of-time rule gives an approximation of the true VaR (see Jorion 2001).

4. Results

In this section, the results obtained for basket options with time-varying correlation term structures will be discussed using the benchmark of constant correlation term structures. To gain a better understanding of the results of time-varying correlations, the results obtained from constant correlations are first discussed.

4.1 Constant Correlations

To illustrate how basket option prices of different moneyness levels change with respect to different constant correlation values between the underlyings, the option prices are plotted in figure 3 and are calculated with the parameters in table 1. Figure 3 shows the basket option prices that are computed for different moneyness levels using constant correlation values 0, ± 0.5, and ± 0.9 for the underlying assets. From this figure one can see that the differences between the option prices across correlation specifications are relatively pronounced for (near) ATM option prices. Moreover, the most striking observation from figure 3 is that there is an asymmetric correlation effect on the ATM option prices: changes in negative correlations have a greater impact on the option prices than changes of the same magnitude in positive correlations (for example, the difference between the option prices for ρ equal to -0.9 and zero is greater than the difference between the prices of ρ equal to $+0.9$ and zero). This asymmetric correlation effect can also be seen for the OTM

Figure 3. Basket Option Prices for Different Correlations for S_0 from 90 to 110

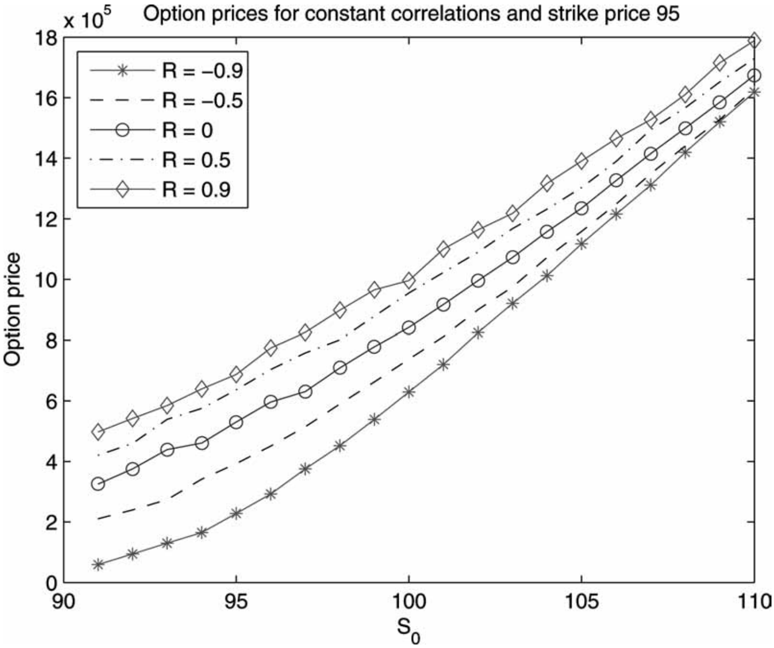

option prices, but to a lesser extent than for the ATM option prices. In figure 3 the option prices converge to the same value beyond a certain moneyness level regardless of the correlation value for the underlying (see, e.g., beyond S_0 of 100). However, the subsequent discussion of the simulation results shows that the risk of potential loss on the far ITM option over a certain time horizon does depend substantially on the correlation values of the underlying assets.

4.1.1 Asymmetric Correlation Effects

In this section the asymmetric effects of negative and positive correlation values on the option price as well as the risk measures VaR and CVaR of the basket option position will be discussed. The results of the simulation study for constant correlations are given in table 2 and turn out to be highly dependent on the initial moneyness level of the basket option. To discuss these results in more detail, let $\bar{V}_0$

Table 2. VaR and CVaR Results for ITM Options with Constant Correlations

	$\mathrm{VaR}^{0.99}$	$\mathrm{VaR}^{0.95}$	$\mathrm{CVaR}^{0.99}$	$\mathrm{CVaR}^{0.95}$	$\frac{\mathrm{CVaR}^{0.99}}{\mathrm{VaR}^{0.99}}$[a]
ITM C1 NH	$-3.66\cdot10^5$	$-2.68\cdot10^5$	$-4.24\cdot10^5$	$-3.29\cdot10^5$	1.16
ITM C1 HE	$-1.54\cdot10^4$	$-1.08\cdot10^4$	$-1.67\cdot10^4$	$-1.35\cdot10^4$	1.09
ITM C1 $\sqrt{10}$NH1	$-3.51\cdot10^5$	$-2.46\cdot10^5$	$-3.93\cdot10^5$	$-3.11\cdot10^5$	1.12
ITM C1 $\sqrt{10}$HE1	$-4.35\cdot10^4$	$-3.08\cdot10^4$	$-4.95\cdot10^4$	$-3.90\cdot10^4$	1.14
ITM C2 NH	$-5.60\cdot10^5$	$-4.14\cdot10^5$	$-6.63\cdot10^5$	$-5.10\cdot10^5$	1.19
ITM C2 HE	$-2.52\cdot10^4$	$-1.69\cdot10^4$	$-3.10\cdot10^4$	$-2.21\cdot10^4$	1.23
ITM C2 $\sqrt{10}$NH1	$-5.40\cdot10^5$	$-3.68\cdot10^5$	$-6.23\cdot10^5$	$-4.75\cdot10^5$	1.15
ITM C2 $\sqrt{10}$HE1	$-5.10\cdot10^4$	$-3.53\cdot10^4$	$-5.71\cdot10^4$	$-4.45\cdot10^4$	1.12
ITM C3 NH	$-6.98\cdot10^5$	$-5.06\cdot10^5$	$-8.11\cdot10^5$	$-6.28\cdot10^5$	1.16
ITM C3 HE	$-3.50\cdot10^4$	$-2.26\cdot10^4$	$-4.38\cdot10^4$	$-3.04\cdot10^4$	1.25
ITM C3 $\sqrt{10}$NH1	$-6.77\cdot10^5$	$-4.44\cdot10^5$	$-7.74\cdot10^5$	$-5.76\cdot10^5$	1.14
ITM C3 $\sqrt{10}$HE1	$-6.30\cdot10^4$	$-4.37\cdot10^4$	$-7.22\cdot10^4$	$-5.54\cdot10^4$	1.15
ITM C4 NH	$-8.48\cdot10^5$	$-6.07\cdot10^5$	$-9.86\cdot10^5$	$-7.67\cdot10^5$	1.16
ITM C4 HE	$-4.84\cdot10^4$	$-2.87\cdot10^4$	$-5.85\cdot10^4$	$-4.04\cdot10^4$	1.21
ITM C4 $\sqrt{10}$NH1	$-7.87\cdot10^5$	$-5.34\cdot10^5$	$-9.40\cdot10^5$	$-6.95\cdot10^5$	1.20
ITM C4 $\sqrt{10}$HE1	$-8.46\cdot10^4$	$-5.45\cdot10^4$	$-9.49\cdot10^4$	$-7.04\cdot10^4$	1.12
ITM C5 NH	$-9.42\cdot10^5$	$-6.60\cdot10^5$	$-1.10\cdot10^6$	$-8.43\cdot10^5$	1.17
ITM C5 HE	$-5.45\cdot10^4$	$-3.37\cdot10^4$	$-6.70\cdot10^4$	$-4.60\cdot10^4$	1.23
ITM C5 $\sqrt{10}$NH1	$-8.45\cdot10^5$	$-5.90\cdot10^5$	$-1.03\cdot10^6$	$-7.61\cdot10^5$	1.22
ITM C5 $\sqrt{10}$HE1	$-9.50\cdot10^4$	$-6.15\cdot10^4$	$-1.09\cdot10^5$	$-8.04\cdot10^4$	1.15
ITM C6 NH	$-1.05\cdot10^6$	$-7.05\cdot10^5$	$-1.20\cdot10^6$	$-9.12\cdot10^5$	1.15
ITM C6 HE	$-5.96\cdot10^4$	$-3.72\cdot10^4$	$-7.25\cdot10^4$	$-5.03\cdot10^4$	1.22
ITM C6 $\sqrt{10}$NH1	$-9.04\cdot10^5$	$-6.53\cdot10^5$	$-1.11\cdot10^6$	$-8.21\cdot10^5$	1.23
ITM C6 $\sqrt{10}$HE1	$-1.04\cdot10^5$	$-6.87\cdot10^4$	$-1.21\cdot10^5$	$-8.98\cdot10^4$	1.17
ITM C7 NH	$-1.19\cdot10^6$	$-7.60\cdot10^5$	$-1.35\cdot10^6$	$-1.01\cdot10^6$	1.13
ITM C7 HE	$-6.36\cdot10^4$	$-4.21\cdot10^4$	$-7.83\cdot10^4$	$-5.57\cdot10^4$	1.23
ITM C7 $\sqrt{10}$NH1	$-1.04\cdot10^6$	$-7.33\cdot10^5$	$-1.22\cdot10^6$	$-9.08\cdot10^5$	1.17
ITM C7 $\sqrt{10}$HE1	$-1.19\cdot10^5$	$-8.09\cdot10^4$	$-1.37\cdot10^5$	$-1.03\cdot10^5$	1.15
ITM C8 NH	$-1.21\cdot10^6$	$-8.12\cdot10^5$	$-1.43\cdot10^6$	$-1.07\cdot10^6$	1.19
ITM C8 HE	$-6.80\cdot10^4$	$-4.44\cdot10^4$	$-8.16\cdot10^4$	$-5.92\cdot10^4$	1.20
ITM C8 $\sqrt{10}$NH1	$-1.09\cdot10^6$	$-7.67\cdot10^5$	$-1.28\cdot10^6$	$-9.71\cdot10^5$	1.17
ITM C8 $\sqrt{10}$HE1	$-1.25\cdot10^5$	$-8.57\cdot10^4$	$-1.49\cdot10^5$	$-1.11\cdot10^5$	1.19
ITM C9 NH	$-1.30\cdot10^6$	$-8.67\cdot10^5$	$-1.51\cdot10^6$	$-1.14\cdot10^6$	1.16
ITM C9 HE	$-7.16\cdot10^4$	$-4.68\cdot10^4$	$-8.43\cdot10^4$	$-6.23\cdot10^4$	1.18
ITM C9 $\sqrt{10}$NH1	$-1.16\cdot10^6$	$-8.10\cdot10^5$	$-1.32\cdot10^6$	$-1.04\cdot10^6$	1.14
ITM C9 $\sqrt{10}$HE1	$-1.40\cdot10^5$	$-9.25\cdot10^4$	$-1.61\cdot10^5$	$-1.20\cdot10^5$	1.15

[a]NH = No hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\mathrm{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\mathrm{VaR}_{1-day}$(daily delta hedging)

denote the average basket option price of the simulation sample at time 0 and define[2]

$$\begin{aligned}
\delta\mathrm{VaR}^{\alpha}_{-} &= \mathrm{VaR}^{\alpha}(\rho = 0) - \mathrm{VaR}^{\alpha}(\rho = -0.9)\\
\delta\mathrm{VaR}^{\alpha}_{+} &= \mathrm{VaR}^{\alpha}(\rho = 0.9) - \mathrm{VaR}^{\alpha}(\rho = 0)\\
\delta\mathrm{VaR}^{\alpha}_{\mathrm{Total}} &= \delta\mathrm{VaR}^{\alpha}_{-} + \delta\mathrm{VaR}^{\alpha}_{+}.
\end{aligned}$$

The results for the unhedged basket option are as follows. First, as the constant correlation ρ increases from -0.9 to 0.9 (respectively, C1 and C9), the average initial option price $\bar{V}_0$ as well as the potential loss measured by VaR and CVaR increase for each moneyness level. For ITM options, this increase in the potential loss is unproportionally high relative to the increase in the average initial option price. For example, $\bar{V}_0^{ITM}$ of C1 and C9 are, respectively, $6.30 \cdot 10^5$ and $1.01 \cdot 10^6$, which is an increase with factor 1.61. The corresponding unhedged ten-day $\mathrm{VaR}^{0.95}$ estimates for C1 and C9 as given in table 2 are $-2.68 \cdot 10^5$ to $-8.67 \cdot 10^5$, respectively, so the potential losses increase with a factor of 3.23 (this factor is even larger for $\mathrm{VaR}^{0.99}$ and CVaR). As the correlations change from -0.9 to $+0.9$, the $(1-\alpha)$-quantiles of the option position's profits and losses increase to more than three times larger than the $\bar{V}_0$ for ITM options. In contrast to this result, higher correlation values affect $\bar{V}_0^{ATM}$ option prices and the corresponding VaR and CVaR estimates more proportionally, where $\bar{V}_0^{ATM}$ increases with factor 3.08 and (C)VaR estimates increase with about factor 4 when going from correlation value -0.9 to $+0.9$.

Secondly, the largest part of the increase in potential losses for ITM and ATM options can be found by increasing the negative correlations. For example, from table 2 the ITM $\delta\mathrm{VaR}^{0.95}_{-}$ has a value of $-3.92 \cdot 10^5$, which constitutes 65 percent of the total change $\delta\mathrm{VaR}^{0.95}_{\mathrm{Total}}$ of $-5.98 \cdot 10^5$. For the ATM and OTM option, this percentage can be obtained similarly and is 63 percent of the $\delta\mathrm{VaR}^{0.95}_{\mathrm{Total}}$. The relatively high sensitivity of basket options on well-diversified baskets (negative correlations) with respect to correlation changes seems counterintuitive. It is well known that well-diversified portfolios are

[2]The VaR here is the ten-day VaR^{α} of the unhedged option position given in table 2.

less risky than portfolios with highly correlated assets; therefore, investors might expect that options on a well-diversified portfolio are also less risky than options on highly correlated underlying assets. Although the results in table 2 show that the absolute *values* of the (C)VaR of the basket option increase with higher correlations, the *changes* in these (C)VaR values are relatively more subject to correlation risk for a negatively correlated portfolio. This observation is important for investors, because they may not be aware of this increased correlation risk when buying basket options on well-diversified ("safe") baskets of assets. Another implication of this result is that VaR estimation of basket options on well-diversified baskets is relatively more prone to model risk, since in practice correlations have to be estimated under the assumption of a correlation model (e.g., RiskMetrics^{TM}), which is in line with the results of Skintzi, Skiadopoulos, and Refenes (2005). They have found that the VaR measure becomes relatively more sensitive to correlation estimation errors with decreasing true correlations for linear portfolios as well as option portfolios containing plain-vanilla European options written on correlated underlying assets.

Finally, the option prices and risk measures of OTM options vary even more with the correlation of the underlying assets than those of ITM and ATM options. The $\bar{V}_0$ prices for ρ of -0.9 and $+0.9$ are, respectively, $5.45 \cdot 10^4$ and $5.23 \cdot 10^5$, and the corresponding $\mathrm{VaR}^{0.95}$ estimates are $-7.47 \cdot 10^4$ and $-6.11 \cdot 10^5$. So, the difference in option price and risk is extremely large across different correlation values for the OTM basket option position.

4.1.2 Delta-Hedged Portfolios

In practice, option positions of financial institutions are usually hedged to decrease the exposure to the full risks of the option position. Therefore, it is perhaps even more important to consider the effects of correlations on daily delta-hedged option positions. A delta-hedged option position accounts for the risk of changes in the price of the underlying assets. The estimates of the risk measures are given in table 2. The effectiveness of the delta hedge in reducing the risk of the option position will be measured by the no-hedge-to-hedge ratio of the risk measures given in table 3. If this ratio is very high, the risk of potential loss is reduced much better by hedging

Table 3. Risk Measure Ratios for ITM Options with Constant Correlations

	$VaR^{0.99}$	$VaR^{0.95}$	$CVaR^{0.99}$	$CVaR^{0.95}$[a]
ITM C1 NH/HE	23.86	24.89	25.33	24.48
ITM C1 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.08	8.00	7.93	7.96
ITM C1 NH/$\sqrt{10}$NH1	1.04	1.09	1.08	1.06
ITM C1 HE/$\sqrt{10}$HE1	0.35	0.35	0.34	0.35
ITM C2 NH/HE	22.18	24.50	21.38	23.10
ITM C2 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.59	10.44	10.90	10.68
ITM C2 NH/$\sqrt{10}$NH1	1.04	1.13	1.06	1.07
ITM C2 HE/$\sqrt{10}$HE1	0.50	0.48	0.54	0.50
ITM C3 NH/HE	19.96	22.41	18.53	20.69
ITM C3 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.75	10.14	10.72	10.40
ITM C3 NH/$\sqrt{10}$NH1	1.03	1.14	1.05	1.09
ITM C3 HE/$\sqrt{10}$HE1	0.56	0.52	0.61	0.55
ITM C4 NH/HE	17.51	21.12	16.86	18.96
ITM C4 $\sqrt{10}$NH1/$\sqrt{10}$HE1	9.30	9.80	9.91	9.88
ITM C4 NH/$\sqrt{10}$NH1	1.08	1.14	1.05	1.10
ITM C4 HE/$\sqrt{10}$HE1	0.57	0.53	0.62	0.58
ITM C5 NH/HE	17.28	19.60	16.41	18.35
ITM C5 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.89	9.59	9.41	9.46
ITM C5 NH/$\sqrt{10}$NH1	1.12	1.12	1.07	1.11
ITM C5 HE/$\sqrt{10}$HE1	0.57	0.55	0.61	0.57
ITM C6 NH/HE	17.57	18.93	16.60	18.14
ITM C6 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.73	9.51	9.20	9.15
ITM C6 NH/$\sqrt{10}$NH1	1.16	1.08	1.08	1.11
ITM C6 HE/$\sqrt{10}$HE1	0.58	0.54	0.60	0.56
ITM C7 NH/HE	18.74	18.07	17.24	18.14
ITM C7 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.72	9.06	8.90	8.83
ITM C7 NH/$\sqrt{10}$NH1	1.15	1.04	1.11	1.11
ITM C7 HE/$\sqrt{10}$HE1	0.53	0.52	0.57	0.54
ITM C8 NH/HE	17.72	18.28	17.59	18.15
ITM C8 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.73	8.96	8.59	8.75
ITM C8 NH/$\sqrt{10}$NH1	1.10	1.06	1.12	1.11
ITM C8 HE/$\sqrt{10}$HE1	0.54	0.52	0.55	0.53
ITM C9 NH/HE	18.13	18.52	17.87	18.25
ITM C9 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.31	8.75	8.23	8.67
ITM C9 NH/$\sqrt{10}$NH1	1.12	1.07	1.14	1.10
ITM C9 HE/$\sqrt{10}$HE1	0.51	0.51	0.53	0.52

[a]NH = No hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

than in the case of a low ratio. The results are very different across moneyness levels.

For ITM options, daily delta hedging reduces the risk by a factor of 23 or more for the highly negative correlation specification C1 for all risk measures. The hedge effectiveness reduces for higher correlations, but it is still considerable for positive correlations, as the no-hedge-to-hedge ratio of the risk measures has a value of more than 16 for different confidence levels.

The delta-hedge effectiveness deteriorates for ATM and OTM options for all values of correlations. Still, the risk reduction is substantial compared with the unhedged ATM and OTM option position in most cases, and the no-hedge-to-hedge ratios vary from around 7 to 15. In contrast to the result for ITM options, the reduction in risk achieved by delta hedging is the largest for highly positive correlations (especially for OTM options). Overall, the delta hedge is most effective for ITM options but still substantial for ATM and OTM options.

4.1.3 Risk Measures and the Square-Root-of-Time Rule

This section discusses the difference between the VaR and CVaR risk measures in this simulation experiment as well as the performance of the widely used square-root-of-time rule. Overall, the risk of potential loss increases for a higher value of the underlying constant correlation (see the VaR and CVaR estimates in table 2). From table 3 one can see that the $\mathrm{CVaR}^{0.99}$-to-$\mathrm{VaR}^{0.99}$ ratios of the ITM option are larger for the hedged option position than for the unhedged position in most cases, thus implying that the hedged profit and loss distribution has relatively heavier tails. However, for the ATM and OTM option, the profit and loss distribution of the unhedged option position exhibits heavier tails compared with the hedged option position. Define γ as

$$\gamma = \frac{\text{ten-day VaR}^{0.99}}{\sqrt{10}\cdot\text{one-day VaR}^{0.99}}, \tag{11}$$

where γ is used to measure the performance of the square-root-of-time rule. If γ is larger than one, the unhedged ten-day VaR is underestimated when applying the square-root-of-time rule. One would expect that the square-root-of-time rule performs worse for cases with high $\mathrm{CVaR}^{0.99}$-to-$\mathrm{VaR}^{0.99}$ ratios, since the validity of this rule critically depends on the assumption of i.i.d.-normal observations. This expectation is indeed confirmed by the results, since the $\mathrm{CVaR}^{0.99}$-to-$\mathrm{VaR}^{0.99}$ ratios for the unhedged OTM option are relatively high compared with the corresponding values for the unhedged ATM option, and this also holds for the underestimation of the ten-day $\mathrm{VaR}^{0.99}$. A notable exception is the unhedged ITM option position, where the square-root-of-time rule performs quite well. For the hedged option position, there is overestimation of the ten-day $\mathrm{VaR}^{0.99}$ by the square-root-of-time rule, since daily delta hedging has an offsetting effect on the profit and loss changes of the option position. The no-hedge-to-hedge ratios are quite high, with a minimum value of about 7. This overestimation declines as the difference of the hedge effectiveness between the one-day and ten-day horizon diminishes, i.e., as the difference in the no-hedge-to-hedge ratios decreases.

So, in most cases, the square-root-of-time rule leads to either underestimation or overestimation of the ten-day $\mathrm{VaR}^{0.99}$. However, none of the γ ratios exceed the minimum absolute stress factor value of 3. Hence, for the computation of market risk capital charge, the square-root-of-time rule used in combination with the stress factor k of at least 3 is reasonable here for basket options in case of constant correlations.

4.2 Time-Varying Correlations

In this section the results of time-varying correlation term structures containing jumps, regime switches, and affine term structures will be discussed for the basket option. By (8) the computation of the VaR for basket options with time-varying correlation term structures can be greatly simplified: one only needs to know an estimate of the average correlation over the option life instead of information on the entire future correlation term structure. The results in table 4 show that property (8) does not hold for the basket option with different correlation term structures. If the correlations are relatively high during the ten-day horizon over which the VaR is computed, then the potential loss will be higher than indicated by the VaR estimated using the average correlation over the life of the option

Table 4. VaR and CVaR Results for ITM Options with Time-Varying Correlations

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$	$\frac{\text{CVaR}^{0.99}}{\text{VaR}^{0.99}}$ [a]
ITM T1 NH	$-8.39 \cdot 10^5$	$-5.55 \cdot 10^5$	$-9.27 \cdot 10^5$	$-7.17 \cdot 10^5$	1.11
ITM T1 HE	$-5.96 \cdot 10^4$	$-3.19 \cdot 10^4$	$-8.58 \cdot 10^4$	$-5.12 \cdot 10^4$	1.44
ITM T1 $\sqrt{10}$NH1	$-3.40 \cdot 10^5$	$-2.35 \cdot 10^5$	$-3.78 \cdot 10^5$	$-2.97 \cdot 10^5$	1.11
ITM T1 $\sqrt{10}$HE1	$-4.29 \cdot 10^4$	$-3.24 \cdot 10^4$	$-4.84 \cdot 10^4$	$-3.89 \cdot 10^4$	1.13
ITM T2 NH	$-3.55 \cdot 10^5$	$-2.59 \cdot 10^5$	$-4.12 \cdot 10^5$	$-3.18 \cdot 10^5$	1.16
ITM T2 HE	$-1.49 \cdot 10^4$	$-1.10 \cdot 10^4$	$-1.68 \cdot 10^4$	$-1.35 \cdot 10^4$	1.12
ITM T2 $\sqrt{10}$NH1	$-3.37 \cdot 10^5$	$-2.36 \cdot 10^5$	$-3.78 \cdot 10^5$	$-2.98 \cdot 10^5$	1.12
ITM T2 $\sqrt{10}$HE1	$-4.52 \cdot 10^4$	$-3.08 \cdot 10^4$	$-5.02 \cdot 10^4$	$-3.96 \cdot 10^4$	1.11
ITM T3 NH	$-1.15 \cdot 10^6$	$-7.36 \cdot 10^5$	$-1.47 \cdot 10^6$	$-9.98 \cdot 10^5$	1.28
ITM T3 HE	$-6.82 \cdot 10^4$	$-4.24 \cdot 10^4$	$-8.08 \cdot 10^4$	$-5.80 \cdot 10^4$	1.19
ITM T3 $\sqrt{10}$NH1	$-1.16 \cdot 10^6$	$-8.11 \cdot 10^5$	$-1.33 \cdot 10^6$	$-1.04 \cdot 10^6$	1.14
ITM T3 $\sqrt{10}$HE1	$-1.38 \cdot 10^5$	$-9.13 \cdot 10^4$	$-1.60 \cdot 10^5$	$-1.19 \cdot 10^5$	1.16
ITM T4 NH	$-1.30 \cdot 10^6$	$-8.72 \cdot 10^5$	$-1.51 \cdot 10^6$	$-1.14 \cdot 10^6$	1.16
ITM T4 HE	$-7.08 \cdot 10^4$	$-4.72 \cdot 10^4$	$-8.50 \cdot 10^4$	$-6.19 \cdot 10^4$	1.20
ITM T4 $\sqrt{10}$NH1	$-1.16 \cdot 10^6$	$-8.07 \cdot 10^5$	$-1.32 \cdot 10^6$	$-1.04 \cdot 10^6$	1.14
ITM T4 $\sqrt{10}$HE1	$-1.32 \cdot 10^5$	$-8.98 \cdot 10^4$	$-1.55 \cdot 10^5$	$-1.18 \cdot 10^5$	1.18
ITM T5 NH	$-3.30 \cdot 10^5$	$-2.40 \cdot 10^5$	$-3.86 \cdot 10^5$	$-2.97 \cdot 10^5$	1.17
ITM T5 HE	$-1.66 \cdot 10^4$	$-1.22 \cdot 10^4$	$-1.83 \cdot 10^4$	$-1.49 \cdot 10^4$	1.10
ITM T5 $\sqrt{10}$NH1	$-3.15 \cdot 10^5$	$-2.15 \cdot 10^5$	$-3.52 \cdot 10^5$	$-2.73 \cdot 10^5$	1.12
ITM T5 $\sqrt{10}$HE1	$-4.82 \cdot 10^4$	$-3.24 \cdot 10^4$	$-5.42 \cdot 10^4$	$-4.16 \cdot 10^4$	1.12
ITM T6 NH	$-1.32 \cdot 10^6$	$-8.79 \cdot 10^5$	$-1.53 \cdot 10^6$	$-1.15 \cdot 10^6$	1.16
ITM T6 HE	$-7.36 \cdot 10^4$	$-4.78 \cdot 10^4$	$-8.60 \cdot 10^4$	$-6.32 \cdot 10^4$	1.17
ITM T6 $\sqrt{10}$NH1	$-1.18 \cdot 10^6$	$-8.17 \cdot 10^5$	$-1.34 \cdot 10^6$	$-1.05 \cdot 10^6$	1.14
ITM T6 $\sqrt{10}$HE1	$-1.29 \cdot 10^5$	$-8.60 \cdot 10^4$	$-1.50 \cdot 10^5$	$-1.13 \cdot 10^5$	1.17
ITM T7 NH	$-4.25 \cdot 10^5$	$-3.15 \cdot 10^5$	$-4.97 \cdot 10^5$	$-3.79 \cdot 10^5$	1.17
ITM T7 HE	$-2.89 \cdot 10^4$	$-2.06 \cdot 10^4$	$-3.31 \cdot 10^4$	$-2.55 \cdot 10^4$	1.15
ITM T7 $\sqrt{10}$NH1	$-2.86 \cdot 10^5$	$-1.97 \cdot 10^5$	$-3.24 \cdot 10^5$	$-2.47 \cdot 10^5$	1.13
ITM T7 $\sqrt{10}$HE1	$-8.14 \cdot 10^4$	$-5.63 \cdot 10^4$	$-9.15 \cdot 10^4$	$-7.12 \cdot 10^4$	1.12
ITM T8 NH	$-1.32 \cdot 10^6$	$-9.24 \cdot 10^5$	$-1.57 \cdot 10^6$	$-1.18 \cdot 10^6$	1.19
ITM T8 HE	$-7.90 \cdot 10^4$	$-5.09 \cdot 10^4$	$-1.02 \cdot 10^5$	$-7.00 \cdot 10^4$	1.29
ITM T8 $\sqrt{10}$NH1	$-1.25 \cdot 10^6$	$-8.54 \cdot 10^5$	$-1.40 \cdot 10^6$	$-1.10 \cdot 10^6$	1.12
ITM T8 $\sqrt{10}$HE1	$-1.24 \cdot 10^5$	$-7.45 \cdot 10^4$	$-1.55 \cdot 10^5$	$-1.04 \cdot 10^5$	1.25

[a]NH = No hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

(and vice versa). This result is in line with Eydeland and Wolyniec (2003), where they explain that for a short VaR horizon of ten days, the instantaneous correlation becomes more important. Hence, there are differences in the results depending on the specific type of time variation, even though the average correlation over the option life is the same, and this will be discussed in more detail later. This result implies that, in practice, the cumulative correlation might not be as relevant for VaR computations as expected. The reduction in risks achieved by daily delta hedging is often less for time-varying correlation term structures than for the constant (cumulative) correlations, but the risk reduction is still substantial for ITM or ATM options and, to a lesser extent, for OTM options. The results for the hedged option are highly dependent on the hedge effectiveness. The following discussion deals with the differences in the specific time variation of the correlation term structures T1 to T8.

4.2.1 Correlation Jumps

The correlation term structures T1 and T2 are initially highly negative and jump upward, respectively, at the start and the end of the option life. First, the time at which the correlation jump occurs is important for the VaR and CVaR estimates of the unhedged option but has an even larger impact on the effectiveness of the delta hedge. When no hedging is applied, a jump at the start of the option life is more than twice as risky as a jump near expiration for negative correlations; see table 4. For the delta-hedged option position this result is even stronger, since the risk measures indicate that the potential loss for a jump at the start is more than four times larger than for a jump near expiration. These results are robust for different moneyness levels. The average correlation for T1 and T2 is -0.84. The risk measures of only T2 give results that are comparable to those of C1 (constant correlation value of -0.9), but the risks for T1 are much higher than for C1. The no-hedge-to-hedge-ratios of the risk measures of T2 in table 5 are twice as large as for T1. So, the time of the temporary jump upward in the correlation has a large impact on the effectiveness of the delta hedge across all moneyness levels.

When the correlations are initially positive, a jump downward at the start (T3) is less risky than a jump near expiration (T4) for the unhedged basket option, regardless of the moneyness level. The hedge effectiveness of a jump near expiration is greater than a jump at the start for positive correlations. So, T2 and T4 have a relatively better hedging performance compared with, respectively, T1 and T3,

Table 5. Risk Measure Ratios for ITM Options with Time-Varying Correlations

	$\mathrm{VaR}^{0.99}$	$\mathrm{VaR}^{0.95}$	$\mathrm{CVaR}^{0.99}$	$\mathrm{CVaR}^{0.95}$[a]
ITM T1 NH/HE	14.08	17.43	10.81	13.99
ITM T1 $\sqrt{10}$NH1/$\sqrt{10}$HE1	7.93	7.27	7.80	7.62
ITM T1 NH/$\sqrt{10}$NH1	2.47	2.36	2.46	2.42
ITM T1 HE/$\sqrt{10}$HE1	1.39	0.99	1.77	1.32
ITM T2 NH/HE	23.75	23.44	24.53	23.57
ITM T2 $\sqrt{10}$NH1/$\sqrt{10}$HE1	7.47	7.64	7.52	7.52
ITM T2 NH/$\sqrt{10}$NH1	1.05	1.10	1.09	1.07
ITM T2 HE/$\sqrt{10}$HE1	0.33	0.36	0.34	0.34
ITM T3 NH/HE	16.85	17.36	18.16	17.23
ITM T3 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.46	8.88	8.30	8.78
ITM T3 NH/$\sqrt{10}$NH1	0.99	0.91	1.11	0.96
ITM T3 HE/$\sqrt{10}$HE1	0.50	0.47	0.51	0.49
ITM T4 NH/HE	18.39	18.46	17.81	18.46
ITM T4 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.80	8.98	8.52	8.83
ITM T4 NH/$\sqrt{10}$NH1	1.12	1.08	1.15	1.10
ITM T4 HE/$\sqrt{10}$HE1	0.54	0.53	0.55	0.53
ITM T5 NH/HE	19.89	19.68	21.11	19.96
ITM T5 $\sqrt{10}$NH1/$\sqrt{10}$HE1	6.53	6.63	6.49	6.57
ITM T5 NH/$\sqrt{10}$NH1	1.05	1.12	1.10	1.09
ITM T5 HE/$\sqrt{10}$HE1	0.34	0.38	0.34	0.36
ITM T6 NH/HE	17.95	18.38	17.77	18.25
ITM T6 $\sqrt{10}$NH1/$\sqrt{10}$HE1	9.12	9.49	8.88	9.25
ITM T6 NH/$\sqrt{10}$NH1	1.12	1.08	1.14	1.10
ITM T6 HE/$\sqrt{10}$HE1	0.57	0.56	0.57	0.56
ITM T7 NH/HE	14.71	15.26	15.01	14.87
ITM T7 $\sqrt{10}$NH1/$\sqrt{10}$HE1	3.51	3.50	3.54	3.47
ITM T7 NH/$\sqrt{10}$NH1	1.49	1.60	1.54	1.54
ITM T7 HE/$\sqrt{10}$HE1	0.36	0.37	0.36	0.36
ITM T8 NH/HE	16.76	18.15	15.44	16.93
ITM T8 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.08	11.45	9.01	10.60
ITM T8 NH/$\sqrt{10}$NH1	1.06	1.08	1.12	1.08
ITM T8 HE/$\sqrt{10}$HE1	0.64	0.68	0.65	0.67

[a]NH = No hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\mathrm{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\mathrm{VaR}_{1-day}$(daily delta hedging)

as shown by the no-hedge-to-hedge ratios in table 5. Hence, delta hedging of the basket option is more effective in reducing the risk of potential loss when the correlation jump does not occur within the VaR horizon of ten days. The mean and standard deviation of the initial option prices are nearly the same for T1 and T2, as well as for T3 and T4. Therefore, stress testing is quite important, since the valuation of the option does not show any of these differences in the potential loss of the position.

4.2.2 Correlation Regime Switches

The correlation term structure T5 has initially highly negative (constant) correlations but shifts upward in the second half of the life of the option. For the unhedged ITM and ATM option, a jump near expiration (T2) is riskier than a correlation shift (T5), although the jump only lasts for five days (see the risk measures in table 4).

Due to the relatively high hedge effectiveness of T2 for the ITM option compared with T5 (i.e., high no-hedge-to-hedge ratios), the (C)VaR estimates show that the potential loss of a jump near expiration is lower than a correlation shift for the hedged ITM option. For the hedged ATM option, the no-hedge-to-hedge ratios of T2 and T5 are comparable, and a jump near expiration still has a higher potential loss than a correlation shift for negative correlations. Regarding the OTM option, a correlation shift is riskier than a jump near expiration when no hedging is applied, and this also holds for the hedged option according to the (C)VaR estimates at the 5 percent significance level.

The risk measures do not differ largely between T4 and T6, which are correlation term structures with initially positive correlations and, respectively, a correlation jump and shift downward. The hedge effectiveness is much larger for T6 than for T5 in case of ATM and OTM options in terms of higher no-hedge-to-hedge ratios of risk measures.

4.2.3 Affine Correlation Term Structures

The linear increase and decrease of, respectively, correlations T7 and T8 have cumulative correlation of zero, which corresponds to C5. For C5, T7, and T8, the initial option prices have very similar mean and standard deviations. The unhedged (C)VaR results of T7 are much lower than those of C5, but T8 shows relatively higher risk of potential loss. However, the hedge effectiveness of T8 is relatively higher than for T7, as shown by the higher no-hedge-to-hedge ratios for T8. During bull markets, the correlations could decline; during bear markets or crisis, correlations often increase. Although the unhedged basket option is riskier for declining correlations (e.g., bull markets) in terms of higher potential loss estimates, the daily delta hedge is much more effective in reducing risk than for increasing correlations (e.g., bear markets).

4.2.4 Risk Measures and the Square-Root-of-Time Rule

When comparing the ten-day VaR to the ten-day CVaR, the results in table 4 show that the difference in the risk measures can be quite large. For example, the potential loss indicated by the $\mathrm{CVaR}^{0.99}$ is more than 40 percent larger than by the $\mathrm{VaR}^{0.99}$ for the hedged ITM option with correlation term structure T1. For the other term structures the difference is less pronounced but can still be substantial, especially at the 5 percent significance level. Therefore, it is advisable to look at both the VaR and CVaR estimates, since the CVaR could provide additional information on the risk exposure.

Many market practitioners use the square-root-of-time rule. Previously, the VaR and CVaR results for constant correlations have shown that the square-root-of-time rule provides reasonable estimates for the unhedged basket option. To my knowledge, there is not much evidence in the literature on how the square-root-of-time rule for VaR and CVaR performs for basket options if the correlation term structure is time varying. In this paper, simulation results show that the square-root-of-time rule (C)VaR estimates substantially underestimate the ten-day (C)VaR for the unhedged option with time-varying correlation term structures as shown in table 4. This underestimation is relatively severe for the cases T1 and T7, where the negative correlation between the underlying assets suddenly increases (as is often observed in financial crises). If the correlations at the first day of the VaR horizon are highly negative, then the VaR given by the square-root-of-time rule is computed based on this negative value and thus neglects the fact that the correlations increase over time.

The ratio γ of the ten-day $\text{VaR}^{0.99}$ to the square-root-of-time rule estimate is given in table 3 for the constant correlations and in table 5 for time-varying correlation term structures. The minimum of the stress factor k is 3, and in table 3 it is shown that the γ from our VaR results for the constant correlation specification do not exceed 3 for all constant correlation values and moneyness levels. The underestimation of the unhedged ten-day VaR by the square-root-of-time rule stays well below a factor of 1.5 for constant correlations. The results in table 5 show that the underestimation is much more severe for time-varying correlations, where the unhedged ten-day $\text{VaR}^{0.99}$ for T1 is underestimated with a factor of 2.47 for the unhedged ITM option and increases to 4.08 for the OTM option, thus even violating the minimum stress factor k of 3.

The performance of the square-root-of-time rule is worse for the hedged option, as can be seen in table 5. Depending on the difference between the hedge effectiveness for the one-day and ten-day horizon, the square-root-of-time rule could lead to either severe underestimation or overestimation of the ten-day (C)VaR. For example, the no-hedge-to-hedge ratios for the ten-day $\text{VaR}^{0.99}$ and the scaled one-day $\text{VaR}^{0.99}$ are, respectively, 14.08 and 7.93 for T1 in table 5. The T1 results for the unhedged ITM option show an underestimation of the ten-day $\text{VaR}^{0.99}$ by the square-root-of-time rule, and this also holds for the ten-day $\text{VaR}^{0.99}$ of the hedged option. For T7, the no-hedge-to-hedge ratios of the ten-day $\text{VaR}^{0.99}$ and the scaled one-day $\text{VaR}^{0.99}$ are, respectively, 14.71 and 3.51. Hence, the hedge effectiveness for the T7 scaled one-day $\text{VaR}^{0.99}$ is relatively much lower than for the ten-day $\text{VaR}^{0.99}$ in the case of T1. The hedged ten-day $\text{VaR}^{0.99}$ shows that the risk of potential loss for T1 is much larger than for T7, but the relatively low hedge effectiveness for the one-day horizon of T7 leads to the result that the scaled one-day $\text{VaR}^{0.99}$ of T7 is twice as large as the same risk measure for T1. Thus, the square-root-of-time rule leads to an overestimation of the T7 ten-day (C)VaR for the hedged option due to this large difference in hedge effectiveness between the one-day and ten-day horizon. The ten-day $\text{VaR}^{0.99}$ is also overestimated by the square-root-of-time rule for many other time-varying correlation term structures for the hedged option.

Since the performance of the square-root-of-time rule for hedged options is very sensitive to the hedge effectiveness, the minimum stress factor is likely to be violated in situations where hedging becomes relatively difficult—for example, for options that are out-of-the-money or near expiration. Hence, the square-root-of-time rule applied to a one-day (C)VaR to obtain a ten-day (C)VaR appears to be inadequate for time-varying correlation term structures, even in this simplified and ideal simulation environment.

4.3 Time Scaling for Time-Varying Correlations

So, the square-root-of-time rule gives reasonable results for unhedged VaR estimates in case of constant correlations but does not perform well for the hedged option and is highly dependent on the hedge effectiveness over different horizons. For time-varying correlations, the square-root-of-time rule can lead to serious under- and overestimation of the ten-day $\mathrm{VaR}^{0.99}$. Therefore, this section examines whether there is another scaling horizon adequate enough to account for the time-varying correlation term structure such that the under- or overestimation by the square-root-of-time rule remains within 10 percent of the ten-day $\mathrm{VaR}^{0.99}$. The square-root-of-time rule has been examined for the horizon of ten days with the following scalings of $\sqrt{10/j} \times \mathrm{VaR}_{j-day}$ for $j = 1, \ldots, 10$ to see whether any type of scaling is valid. The results show that for correlation term structures with initially positive correlations (T3, T4, T6, and T8), the square-root-of-time rule by scaling the one-day VaR is still reasonable for j of 7 or higher for both hedged and unhedged options.[3] However, the results for the correlation term structures with initially negative correlations (T1, T2, T5, and T7) show that the γ's become close to a value of 1 only for a scaling of j of 9 or more in most cases of the hedged option position. Hence, the scaling of the VaR can only be applied for a time horizon of one to three days in the future for the cases considered here.

4.4 Multivariate Baskets with Time-Varying Correlations

The simulation study in this paper considers a basket option with two underlying assets for simplicity, since there is only one correlation term structure involved in the analysis. However, in practice,

[3] Results available by request.

basket options are often traded on baskets consisting of more than two assets. Longin and Solnik (2001) find that during bear markets, two assets can become increasingly correlated, whereas these assets might have lower correlations in other market conditions. If many assets simultaneously become highly correlated during financial crises, not only do the pairwise correlations have an effect on a basket option on multiple assets, but an interaction effect between different correlation term structures could also exist. According to Fengler and Schwendner (2004), multiple assets in a basket can provide a diversification effect but could also introduce additional risks of unknown correlations, and as correlations between assets rise, the diversification effect will be diminished.

If the purpose of including many assets in a basket is to benefit from a diversification effect, the portfolio of the selected underlying assets should have low correlations in normal market conditions. In that case, the results in this paper for the correlation term structures with initially low correlations (such as T1, T2) are likely to be more relevant to basket options in practice than the other specifications considered. Moreover, the hedging of the basket options with many assets could be much more difficult. Since the performance of the square-root-of-time rule is highly dependent on the hedge effectiveness, this can have a negative impact on the accuracy of the scaled one-day VaR estimates as an approximation for the ten-day VaR. The basket option does not only react to changes in one correlation term structure as in the case of two assets, but will also react to changes in $\frac{1}{2}N(N-1)$ different correlation term structures and possibly to interaction effects between these correlation term structures for N assets ($N > 2$). To what extent opposite correlation movements in the basket will offset each other will also depend on the relative importance (i.e., basket weights) of the corresponding assets in the basket. A more detailed analysis of high-dimensional basket options is beyond the scope of this study and is left for further research.

5. Conclusions

The purpose of this paper is to analyze the effect of time-varying correlation term structures on pricing and hedging of basket options as well as on the risk measures VaR and CVaR. First, the benchmark with constant correlations has been used to examine the effect of the sign and size of correlations, and the results are as follows. A surprising result is that basket options on a well-diversified portfolio of assets (i.e., negatively correlated assets) are relatively much more sensitive to correlation changes than on a positively correlated portfolio of assets. This implies that sudden changes in the financial markets can lead to large losses for a basket option with well-diversified assets, which cannot be captured by the VaR when using the square-root-of-time rule. The ATM and OTM basket option prices react asymmetrically to positive and negative correlation changes, where a change in negative correlations has a higher impact on the option price than a change in positive correlations of the same magnitude. As a result, the corresponding VaR and CVaR estimates are also more sensitive to changes in negative correlations. The option price of a far ITM option does not differ largely for different correlation values. However, the risk measures of the far ITM basket show that the risk of potential loss increases substantially as the correlation increases. Thus, the risk of potential loss of basket options cannot be entirely captured by the value of the option, but extensive stress testing is needed to reveal these risks.

Second, dynamic delta hedging can reduce the VaR substantially. The specific type of time variation of the correlation term structure is essential to the effectiveness of the dynamic delta hedge of the basket option. Hence, one cannot use the average constant correlation for the computation of the VaR of a hedged basket option with time-varying correlations. Thus, the time of occurrence of a jump during the life of the option is very relevant for the risk of potential loss, even though the basket option payoff is based on the value of the underlying assets at the time of the expiration.

Third, the CVaR can differ from the VaR estimates substantially, especially for the time-varying negative correlations with a jump at the start of the option life and for estimates at the 5 percent significance level. Therefore, providing VaR estimates might not be sufficient, and the coherent measure CVaR can give the additional information needed in certain market conditions.

Finally, the square-root-of-time rule leads to underestimation of the unhedged ten-day VaR in most cases considered, but the underestimation ratio still remains below the regulatory stress factor minimum of 3 for constant correlation term structure. The scaled one-day

VaR can deviate heavily from the ten-day VaR for the hedged option when the difference in the hedge effectiveness for the one-day and ten-day horizon is large. The square-root-of-time rule performs relatively well for unhedged ITM basket options. However, applying the square-root-of-time rule can lead to large deviations from the ten-day VaR for (un)hedged ATM and OTM options. In general, the estimates do not improve very much by using different scaling horizons, unless the scaling is applied for a square-root-of-time factor close to 1. In these cases, stress testing is much more important for determining an accurate VaR estimate.

References

Artzner, P., F. Delbaen, J. M. Eber, and D. Heath. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9 (3): 203–28.

Bank for International Settlements. 2005. "International Convergence of Capital Measurement and Capital Standards: A Revised Framework." Technical Report (November).

Bennett, M. N., and J. E. Kennedy. 2004. "Quanto Pricing with Copulas." *Journal of Derivatives* (Fall):26–45.

Brigo, D., F. Mercurio, F. Rapisarda, and R. Scotti. 2004. "Approximated Moment-Matching Dynamics for Basket-Options Simulation. *Quantitative Finance* 4 (1): 1–16.

Curran, M. 1994. "Valuing Asian and Portfolio Options by Conditioning on the Geometric Mean Price." *Management Science* 40 (12): 1705–11.

Dammers, C. R., and R. N. McCauley. 2006. "Basket Weaving: The Euromarket Experience with Basket Currency Bonds." *BIS Quarterly Review* (March).

Danielsson, J., and J. P. Zigrand. 2005. "On Time-Scaling of Risk and the Square-Root-of-Time Rule." Working Paper, available at www.riskresearch.org.

Deelstra, G., J. Liinev, and M. Vanmaele. 2004. "Pricing of Arithmetic Basket Options by Conditioning." *Insurance, Mathematics and Economics* 34 (1): 55–77.

Diebold, F. X., A. Hickman, A. Inoue, and T. Schuermann. 1997. "Converting 1-Day Volatility to h-Day Volatility: Scaling by Root-h Is Worse than You Think." Wharton Working Paper Series, No. 97-34.

Duffie, D., and K. J. Singleton. 2003. *Credit Risk: Pricing, Measurement, and Management.* 1st ed. Princeton, NJ: Princeton University Press.

Eydeland, A., and K. Wolyniec. 2003. *Energy and Power Risk Management.* 1st ed. Hoboken, NJ: John Wiley and Sons, Inc.

Fengler, M. R., and P. Schwendner. 2004. "Quoting Multiasset Equity Options in the Presence of Errors from Estimating Correlations." *Journal of Derivatives* (Summer):43–54.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering.* 1st ed. New York: Springer-Verlag, Inc.

Goetzmann, W. N., L. Li, and K. G. Rouwenhorst. 2005. "Long-Term Global Market Correlations." *Journal of Business* 78 (1): 1–38.

Hull, J. 2006. *Options, Futures and Other Derivatives.* 6th ed. New York: Pearson Prentice Hall.

Jorion, P. 2001. *Value at Risk: The New Benchmark for Managing Financial Risk.* 2nd ed. McGraw Hill.

Kupiec, P. H. 1998. "Stress Testing in a Value at Risk Framework." *Journal of Derivatives* (Fall):7–24.

Longin, F., and B. Solnik. 2001. "Extreme Correlation of International Equity Markets." *Journal of Finance* 56 (2): 649–76.

Margrabe, W. 1978. "The Value of an Option to Exchange One Asset for Another." *Journal of Finance* 33 (1): 177–86.

McNeil, A. J., R. Frey, and P. Embrechts. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools.* 1st ed. Princeton, NJ: Princeton University Press.

Milevsky, M. A., and S. E. Posner. 1998. "A Closed-Form Approximation for Valuing Basket Options." *Journal of Derivatives* (Summer):54–61.

Morgan, J. P. 1996. "RiskMetrics—Technical Document." J. P. Morgan Bank. 4th ed. New York.

Pellizzari, P. 2005. "Static Hedging of Multivariate Derivatives by Simulation." *European Journal of Operational Research* 166 (2): 507–19.

Pritsker, M. G. 2006. "The Hidden Dangers of Historical Simulation." *Journal of Banking and Finance* 30 (2): 561–82.

Skintzi, V. D., G. Skiadopoulos, and A. P. N. Refenes. 2005. "The Effect of Misestimating Correlation on Value at Risk." *Journal of Alternative Investments* (Spring):66–82.

Firm-Specific Labor and Firm-Specific Capital: Implications for the Euro-Data New Phillips Curve*

Julien Matheron
Banque de France, Research Division

Standard GMM estimates of the New Phillips curve on euro-area data yield degrees of nominal rigidity that are not in accordance with recent microeconomic evidence. This paper studies whether similar conclusions are reached in a richer model where price setters face firm-specific capital and/or firm-specific labor. We find that combining these elements or considering firm-specific labor alone leads to statistically significant and economically reasonable estimates of the degree of nominal rigidity. In contrast, ignoring firm-specific labor yields estimates that are not supported by microeconomic evidence.

JEL Codes: E1, E3.

1. Introduction

In a set of influential papers, Galí and Gertler (1999) and Galí, Gertler, and López-Salido (2001) have shown that a hybrid New Keynesian Phillips curve (NKPC henceforth), based on the Calvo (1983) model, fit U.S. and European inflation data surprisingly well. Despite their attractiveness, these results have been criticized on the grounds that they imply unrealistic degrees of nominal rigidities.

*I thank Thibault Guyon, Eric Jondeau, Hervé Le Bihan, and Jean-Guillaume Sahuc for helpful discussions and comments during the course of this project. I also thank Michael Woodford (editor) and two anonymous referees for useful remarks and suggestions. The views expressed herein are those of the author and do not necessarily reflect those of the Banque de France. Author contact: Julien Matheron, Banque de France, DGEI-DIR, Service de Recherche en économie et finance (41-1391), 31 Rue Croix des Petits-Champs, 75049 Paris Cedex 01, France; Tel: +33-1-42-92-29-97; Fax: +33-1-42-92-62-92; e-mail: julien.matheron@banque-france.fr.

For example, Galí and Gertler (1999) report probabilities of price fixity for U.S. data almost always higher than 0.8, implying an average duration of price fixity higher than five quarters. Galí, Gertler, and López-Salido (2001) obtain an average price duration higher than ten quarters for euro-area data, based on the conventional NKPC.[1] Indeed, these results stand in sharp contrast with recent microeconomic evidence. For example, Bils and Klenow (2004), studying U.S. data, report that the average price duration is roughly two quarters. Dhyne et al. (2006) report an average price duration of ten months to four quarters on euro-area data. Thus, be it for U.S. or euro-area data, there seems to be a discrepancy between microeconomic data studies and macroeconomic estimates of the NKPC.

In response to these criticisms, a number of authors have argued in favor of including alternative supply-side refinements capable of diminishing the responsiveness of inflation to the real marginal cost. A fruitful direction of research, which has received attention in the recent literature, consists in rendering firms' real marginal cost increasing in their own output. In such a framework, when given the opportunity to reoptimize, a firm will change its price by a smaller amount than if its marginal cost were independent of its decisions. Everything else equal, this will translate into a smaller response of inflation to changes in the aggregate marginal cost.

Among the mechanisms working in this direction, firm-specific capital has been the subject of a significant number of recent papers, e.g., Altig et al. (2005), Christiano (2004), Eichenbaum and Fisher (2005), Sveen and Weinke (2004, 2005), and Woodford (2005). All suggest that this additional supply-side mechanism can have prominent effects on the dynamics of inflation and output in this kind of model.[2]

A key contribution to this literature is the paper by Eichenbaum and Fisher (2005). In that paper, the authors show that considering firm-specific capital and a variable demand elasticity allows an NKPC to better match postwar U.S. data. More precisely, they show

[1]However, in a specification allowing for firm-specific labor, they obtain much smaller average durations. See section 3 for similar results, though based on a slightly different econometric approach.

[2]See also de Walque, Smets, and Wouters (2004) for a similar analysis in the context of a model with Taylor (1980) contracts.

that the implied degree of nominal rigidity requested by the data is much smaller than obtained in conventional applications.

In this paper, we propose to estimate, via the generalized method of moments (GMM), NKPCs on euro data in the context of a refined supply-side environment, where firms face specific labor and capital markets. The parameters of interest are the probability of not reoptimizing a price and the degree of indexation to past inflation. We contrast the obtained estimates with those arising in (i) a model with aggregate markets for labor and capital, which yields the usual new Phillips curve; (ii) a model with fixed capital and firm-specific labor, as in Woodford's (2003) basic model; and (iii) a model with an aggregate labor market and firm-specific capital. All these specifications are observationally equivalent, yet they differ with respect to the implied degree of nominal rigidities.

Euro data aside, the present paper is a complementary study to Eichenbaum and Fisher (2005). Instead of assuming a variable demand elasticity, we explore the consequences of allowing for firm-specific labor in addition to firm-specific capital. As argued by Woodford (2003), allowing for firm-specific labor implies a higher degree of strategic complementarity between price setters, which, everything else equal, translates into a smaller partial elasticity of current inflation to the marginal cost. This in itself motivates the inclusion of firm-specific labor as an additional, possible channel of inflation persistence.

According to Woodford (2003), allowing for a variable demand elasticity and/or produced inputs strengthens the degree of strategic complementarity, but only to a marginal extent as long as labor is assumed to be firm specific. Specifically, the increase in the degree of strategic complementarity implied by the latter mechanism dwarfs that implied by a variable demand elasticity and/or the presence of produced inputs in price setters' production function.

Woodford (2005) shows that assuming firm-specific capital and capital adjustment costs plays a similar role, but it is unclear whether this additional mechanism contributes much to increasing the degree of strategic complementarity, compared to firm-specific labor. This is precisely the question under study in the present paper.

Obviously, there are many other supply-side refinements that can be included in a basic sticky-price model to help obtain sensible NKPC estimates. Recent research has followed this direction.

For example, in the context of euro-area data, McAdam and Willman (2004) consider a fairly disaggregated supply-side optimizing framework that allows them to treat possible nonstationarity in factor income shares and markups. Such a study is far beyond the scope of the present paper in which we seek to assess the usefulness of combining simple theoretical elements in order to obtain economically realistic parameter estimates of the euro-area NKPC.

Matheron and Maury (2004) is yet another paper investigating the robustness of Galí and Gertler (1999) conclusions to the inclusion of additional supply-side mechanisms. They show that when a sticky-price model features produced inputs and fixed production costs, the labor share is no longer an appropriate measure of the marginal cost. However, they demonstrate that this problem has no incidence on the message conveyed by Galí and Gertler (1999). In this paper, we abstract from these refinements and exclusively focus on firm-specific capital and/or firm-specific labor within the framework of a simple sticky-price model where the labor share is an appropriate measure of the marginal cost.

Our main results are as follows. First, considering firm-specific capital alone does not lead to reasonable degrees of nominal rigidities. This result is robust to various details of our empirical methodology. Overall, we find that this hypothesis barely improves the standard NKPC specification. Second, considering firm-specific labor yields very reasonable estimates of the degree of nominal rigidities. This result holds either with fixed capital or with firm-specific capital. Finally, abstracting from all these modeling refinements leads to statistically satisfying estimates of the NKPC, though at the cost of economically unacceptable average price durations.

The remainder of the paper is structured as follows. Section 2 details the sticky-price model with capital and labor both firm specific. The alternative specifications are introduced as particular cases of this general model. Section 3 lays out our GMM estimation procedure and then reports the results. The last section briefly concludes.

2. Model Specifications

In this section, we briefly sketch the reference model, which features capital and labor as both firm specific. All the details of the resolution are relegated in the appendix. The alternative specifications arise as specialized cases of the reference model. These are detailed in the following subsection.

2.1 A Sticky-Price Model with Capital and Labor Both Firm Specific

We consider a discrete-time economy populated by a large number of infinitely lived agents. The representative household's goal in life is to maximize

$$\mathrm{E}_0\left\{\sum_{t=0}^{\infty}\beta^t\left[U(c_t)-\int_0^1 V(h_t(\varsigma))\mathrm{d}\varsigma\right]\right\}, \tag{1}$$

where E_t is the expectation operator conditional on information available as of time t, β is the subjective discount factor, c_t is the consumption of final good, $h_t(\varsigma)$ is the supply of labor of type ς, $U(\cdot)$ is strictly concave and increasing, and $V(\cdot)$ is strictly convex and increasing.[3]

The household seeks to maximize (1) subject to the sequence of budget constraints

$$c_t \leq \int_0^1 w_t(\varsigma)h_t(\varsigma)\mathrm{d}\varsigma + \Phi_t, \tag{2}$$

where Φ_t is an aggregate variable summarizing all net sources of income other than wages (e.g., profits redistributed to the household, financial revenues, etc.), and $w_t(\varsigma)$ is the real wage paid to labor of type ς.

The final good y_t is produced by perfectly competitive firms according to the constant elasticity of substitution (CES) technology

$$y_t = \left(\int_0^1 y_t(\varsigma)^{(\theta-1)/\theta}\mathrm{d}\varsigma\right)^{\theta/(\theta-1)}, \quad \theta > 1, \tag{3}$$

where $y_t(\varsigma)$ is the input of intermediate good ς used in the production of y_t. The parameter θ is the elasticity of substitution between any

[3] Notice that to simplify our presentation, we have abstracted from preference and technology shocks.

two intermediate goods. The first-order condition associated with $y_t(\varsigma)$ is

$$y_t(\varsigma) = \left(\frac{P_t(\varsigma)}{P_t} \right)^{-\theta} y_t. \tag{4}$$

Notice that perfect competition ensures that the aggregate price level P_t must obey

$$P_t = \left(\int_0^1 P_t(\varsigma)^{1-\theta} \mathrm{d}\varsigma \right)^{1/(1-\theta)}. \tag{5}$$

Intermediate goods are produced by monopolistic firms, according to the constant-returns technology

$$y_t(\varsigma) = k_t(\varsigma) f\left(\frac{h_t(\varsigma)}{k_t(\varsigma)} \right), \tag{6}$$

where $f(x) = x^{1/\phi}$, $\phi > 1$. The capital stock $k_t(\varsigma)$ is assumed to be firm specific, i.e., we assume away the existence of an economy-wide capital market. Let $i_t(\varsigma)$ denote the real investment expenditures of firm ς in period t. Next-period's capital stock $k_{t+1}(\varsigma)$ and $i_t(\varsigma)$ are linked together through the relation

$$i_t(\varsigma) = k_t(\varsigma) I\left(\frac{k_{t+1}(\varsigma)}{k_t(\varsigma)} \right), \tag{7}$$

where $I(\cdot)$ is convex, with $I(1) = \delta \in (0, 1)$, $I'(1) = 1$, and $I''(1) = \epsilon > 0$. Here, δ is the depreciation rate. Eichenbaum and Fisher (2005) demonstrate that, in the context of the accumulation technology (7), ϵ is linked to the elasticity of the investment-capital ratio with respect to Tobin's q, denoted by ϱ, through the relation $\varrho = 1/(\delta\epsilon)$. Following Eichenbaum and Fisher (2005), we assume that firm ς selects the capital stock $k_{t+1}(\varsigma)$ based on information available as of time $t-1$, i.e., $i_t(\varsigma)$ is in firm ς's information set at $t-1$.

Additionally, as in Woodford (2003), we assume that in each period, a random fraction $1-\alpha$ of intermediate goods producers get to revise their price. The remaining firms simply rescale their price according to the simple rule $P_T(\varsigma) = x_{t,T} P_t(\varsigma)$, where

$$x_{t,T} = \begin{cases} \prod_{j=t}^{T-1} \pi^{1-\gamma} {\pi_j}^{\gamma} & \text{if } T > t \\ 1 & \text{otherwise} \end{cases}. \tag{8}$$

Here, non-reoptimizing firms partially index their prices to past levels of inflation and steady-state inflation π.[4] More precisely, the parameter $\gamma \in [0, 1]$ measures the degree of indexation to the most recently available inflation measure. This specification is an extension of the inflation indexation mechanism considered in Woodford (2003). While with the latter a hybrid New Phillips curve is only valid in the neighborhood of a zero-inflation steady state, the former enables us to consider strictly positive steady-state inflation rates. In the limiting case $\gamma = 1$, this specification reduces to that considered by Christiano, Eichenbaum, and Evans (2005).

Following Eichenbaum and Fisher (2005), we assume a delay in the implementation of a new price. The latter is chosen at $t - 1$ and becomes effective at date t. Thus, if drawn to reoptimize in period $t - 1$, firm ς will select its price $P_t^*(\varsigma)$ so as to maximize

$$\mathrm{E}_{t-1}^{\varsigma} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \frac{\lambda_T}{\lambda_t} \left\{ \left(\frac{x_{t,T} P_t(\varsigma)}{P_T} \right)^{1-\theta} y_T - s_T(\varsigma) \left(\frac{x_{t,T} P_t(\varsigma)}{P_T} \right)^{-\theta} y_T \right\}, \tag{9}$$

where λ_t is the Lagrange multiplier on constraint (2), $s_t(\varsigma)$ is firm ς's marginal cost in period t, and $\mathrm{E}_t^{\varsigma}\{\cdot\}$ is an expectation operator specific to firm ς that integrates over those future states of the world in which firm ς has not reset its price since t.

Our notations emphasize the fact that firm ς's real marginal cost depends on ς. There are two origins to this fact: (i) labor is firm specific, so that the wage rate paid by firm ς depends on firm ς's output, and (ii) firm ς accumulates its own capital stock that, consequently, depends on present and past price decisions. These mechanisms contribute to make $s_t(\varsigma)$ an increasing function of $y_t(\varsigma)$.

[4] In the context of our setup, π is a free parameter that we calibrate to inflation's observed average value.

The above model implies the following NKPC:

$$\mathrm{E}_{t-1}\left\{\frac{\gamma}{1+\beta\gamma}\hat{\pi}_{t-1}+\frac{\beta}{1+\beta\gamma}\hat{\pi}_{t+1}+\zeta^{-1}\frac{(1-\alpha\beta)(1-\alpha)}{\alpha(1+\beta\gamma)}\hat{s}_t-\hat{\pi}_t\right\}=0, \tag{10}$$

where a letter with a hat denotes the log-deviation of the associated variable, s_t is the average (economy-wide) real marginal cost, and

$$\zeta = 1 + (\phi(\nu+1)-1)\theta - \varkappa_0(\alpha,\beta,\delta,\epsilon,\theta,\phi,\nu),$$

where $\nu = V_{hh}h/V_h$ is the elasticity of labor's marginal disutility with respect to hours worked, evaluated in steady state, and $\varkappa_0(\cdot)$ is a positive, complicated function of the parameters listed, which we briefly describe in the appendix.[5]

2.2 Alternative Specifications

We consider three alternative specifications, which we briefly detail below.

2.2.1 Firm-Specific Labor and Fixed Capital

This specification is identical to that considered in Woodford (2003, chap. 3). We simply assume that there is no capital accumulation. The representative household still maximizes (1) subject to (2). Each firm possesses its own fixed stock of capital, and through a suitable normalization, they are assumed to operate the same technology given by

$$y_t(\varsigma) = f(h_t(\varsigma)).$$

Alternatively, one can interpret this specification as the limiting case in which capital adjustment costs are so important that firms simply prefer not to accumulate capital, i.e., the limiting case $\epsilon \to \infty$ (and $\delta = 0$).

Irrespective of the interpretation that one favors, what is important here is that firm ς faces a marginal cost that depends on ς.

[5]For further details, see also Christiano (2004), Eichenbaum and Fisher (2005), Sveen and Weinke (2004, 2005), and Woodford (2005).

This feature, as stressed out by Woodford (2003), generates a strong degree of strategic complementarity between price setters, which translates into a small elasticity of inflation with respect to the real marginal cost in the New Phillips curve.

This specification implies

$$\mathrm{E}_{t-1}\left\{\frac{\gamma}{1+\beta\gamma}\hat{\pi}_{t-1}+\frac{\beta}{1+\beta\gamma}\hat{\pi}_{t+1}+\zeta^{-1}\frac{(1-\alpha\beta)(1-\alpha)}{\alpha(1+\beta\gamma)}\hat{s}_t-\hat{\pi}_t\right\}=0, \tag{11}$$

$$\zeta = 1+(\phi(\nu+1)-1)\theta,$$

which is very close to that derived by Woodford (2003). Additionally, Galí, Gertler, and López-Salido (2001) and Sbordone (2002) study a similar specification.

Notice that in the case when firms face specific labor and capital markets, the implied ζ is necessarily smaller than in the present specification, as long as $\varkappa_0(\alpha,\beta,\delta,\epsilon,\theta,\phi,\nu)$ is positive, which is the case in the estimation reported below. This means that the degree of strategic complementarity is smaller when firms face specific labor and capital markets than in the present specification. As a consequence, we should expect a higher degree of nominal rigidity when firms face specific labor and capital markets than under the assumption of fixed capital and specific labor. However, the hypotheses underlying specification (10) sound more realistic than assuming that capital is fixed over the business cycle. It remains to be seen whether the implied degree of nominal rigidity is within the reasonable range described in the introduction.

2.2.2 Firm-Specific Capital, Aggregate Labor Market

This is a simplified version of the model considered by Eichenbaum and Fisher (2005). In this setting, we assume that firms accumulate their own capital stock and rent labor services on an economy-wide market. In this case, the representative household is assumed to maximize

$$\mathrm{E}_0\left\{\sum_{t=0}^{\infty}\beta^t[U(c_t)-V(h_t)]\right\}, \tag{12}$$

subject to the constraint

$$c_t \leq w_t h_t + \Phi_t,$$

where w_t is the real aggregate wage and Φ_t represents other net sources of income.

Intermediate goods producers still operate technology (6) and accumulate physical capital according to (7). If drawn to reoptimize in period t, firm ς will select its price $P_t^*(\varsigma)$ so as to maximize (9). Again, this specification allows the real marginal cost of firm ς to depend on ς, thus rendering the elasticity of inflation with respect to s_t smaller than in the model with aggregate labor and capital markets. Notice, however, that firm ς's real marginal cost $s_t(\varsigma)$ depends on ς only through the specificity of $k_t(\varsigma)$. We then obtain the NKPC

$$\mathrm{E}_{t-1}\left\{\frac{\gamma}{1+\beta\gamma}\hat{\pi}_{t-1} + \frac{\beta}{1+\beta\gamma}\hat{\pi}_{t+1} + \zeta^{-1}\frac{(1-\alpha\beta)(1-\alpha)}{\alpha(1+\beta\gamma)}\hat{s}_t - \hat{\pi}_t\right\} = 0, \tag{13}$$

$$\zeta = 1 + (\phi - 1)\theta - \varkappa_1(\alpha, \beta, \delta, \epsilon, \theta, \phi),$$

with $\varkappa_1(\alpha, \beta, \delta, \epsilon, \theta, \phi)$ another complicated function of the parameters listed, which we briefly describe in the appendix.

2.2.3 Aggregate Labor and Capital Markets

We now suppose that there exist aggregate capital and labor markets. Except for the automatic indexation that is considered here, this is the specification retained in Yun (1996). In this model, the representative household seeks to maximize (12) subject to the sequence of budget constraints

$$c_t + i_t \leq w_t h_t + \rho_t k_t + \Phi_t, \tag{14}$$

$$i_t = k_t I\left(\frac{k_{t+1}}{k_t}\right), \tag{15}$$

where w_t is the real wage rate, k_t is the stock of capital, which is now accumulated by the households, and ρ_t is the corresponding rental rate. Notice that in this case, neither h_t nor w_t depend on ς.

Intermediate goods producers still operate technology (6). However, since they have access to perfectly competitive inputs markets, and they operate with constant returns to scale, it can be shown that their marginal cost does not depend on ς.[6] In this case, they select $P_t^*(\varsigma)$ so as to maximize

$$\mathrm{E}_{t-1}^{\varsigma}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\frac{\lambda_T}{\lambda_t}\left\{\left(\frac{x_{t,T}P_t(\varsigma)}{P_T}\right)^{1-\theta}y_T - s_T\left(\frac{x_{t,T}P_t(\varsigma)}{P_T}\right)^{-\theta}y_T\right\}. \tag{16}$$

This specification implies the standard NKPC

$$\mathrm{E}_{t-1}\left\{\frac{\gamma}{1+\beta\gamma}\hat{\pi}_{t-1}+\frac{\beta}{1+\beta\gamma}\hat{\pi}_{t+1}+\frac{(1-\alpha\beta)(1-\alpha)}{\alpha(1+\beta\gamma)}\hat{s}_t-\hat{\pi}_t\right\}=0. \tag{17}$$

3. Estimating the New Phillips Curves

3.1 Estimation Strategy

Our goal is now to estimate the new Phillips curves implied by the different model specifications. Following Galí and Gertler (1999) and Eichenbaum and Fisher (2005), we estimate these equations by resorting to Hansen's (1982) generalized method of moments (GMM). Equations (10), (11), (13), and (17) can be interpreted as orthogonality conditions, which lend themselves to instrumental variable estimation. In particular, each of these equations can be generically written as

$$\mathrm{E}_{t-1}\{u_{t+1}(\boldsymbol{\psi})\}=0, \tag{18}$$

where we defined the parameter vector $\boldsymbol{\psi}=(\alpha,\gamma)'$. Recall that α is the probability of not reoptimizing prices in a given quarter and γ is the degree of price indexation to the most recently available inflation measure. These orthogonality conditions imply

$$\mathrm{E}\{\mathbf{f}_{t+1}(\boldsymbol{\psi})\}=0, \quad \mathbf{f}_{t+1}(\boldsymbol{\psi})\equiv u_{t+1}(\boldsymbol{\psi})\mathbf{Z}_{t-1} \tag{19}$$

for any ℓ-dimensional vector $\mathbf{Z}_t$ in the agents' time t information set. The fact that the vector of instruments belongs to date $t-1$'s

[6]See Chari, Kehoe, and McGrattan (2000) for a formal demonstration.

information set reflects our assumption regarding the timing of price decisions in the preceding section. Thus, our estimate of $\boldsymbol{\psi}$ is

$$\hat{\boldsymbol{\psi}} = \arg \min_{\boldsymbol{\psi} \in \boldsymbol{\Psi}} \mathbf{g}_T(\boldsymbol{\psi})' \mathbf{W}_T \mathbf{g}_T(\boldsymbol{\psi}), \tag{20}$$

where

$$\mathbf{g}_T(\boldsymbol{\psi}) = \frac{1}{T} \sum_{t=1}^{T} \mathbf{f}_{t+1}(\boldsymbol{\psi}).$$

Here, $\boldsymbol{\Psi}$ is the set of admissible values for $\boldsymbol{\psi}$, T is the sample size, and $\mathbf{W}_T$ is a symmetric positive definite weighting matrix. For later reference, it is convenient to define $J_T(\boldsymbol{\psi}) = \mathbf{g}_T(\boldsymbol{\psi})' \mathbf{W}_T \mathbf{g}_T(\boldsymbol{\psi})$.

The optimal choice for the weighting matrix consists in choosing the inverse of the spectral density matrix at frequency zero of $\mathbf{f}_{t+1}(\boldsymbol{\psi}_0)$, where $\boldsymbol{\psi}_0$ is $\boldsymbol{\psi}$'s true value. Let $\mathbf{S}$ denote this matrix. As argued by Eichenbaum and Fisher (2005), the theory predicts that $\mathbf{f}_{t+1}(\boldsymbol{\psi}_0)$ has an MA(1) representation, so that $\mathbf{S}$ obeys

$$\mathbf{S} = \sum_{k=-1}^{1} \mathrm{E}\big\{[\mathbf{f}_{t+1+k}(\boldsymbol{\psi}_0)][\mathbf{f}_{t+1+k}(\boldsymbol{\psi}_0)]'\big\}. \tag{21}$$

We could thus obtain a consistent estimator of $\mathbf{S}$ by directly implementing the empirical counterpart of equation (21)—the truncated kernel estimate of $\mathbf{S}$. This method has the obvious advantage of being completely consistent with our model and could be interpreted as a strict test of the theory expounded before. One important drawback, though, is that it can eventually lead to a semidefinite weighting matrix that is not positive in small samples. In practice, we encountered this problem in our application.

Following West (1997), a simple remedy to this problem is to fit an MA(1) to the GMM residuals $u_t(\boldsymbol{\psi})$ and then use the estimated auxiliary parameters to compute the spectral density matrix of $\mathbf{f}_{t+1}(\boldsymbol{\psi})$ evaluated at frequency zero.[7] More precisely, the model implies that $u_t(\boldsymbol{\psi})$ admits the representation

$$u_t(\boldsymbol{\psi}) = \epsilon_t + \eta \epsilon_{t-1}.$$

[7]See Jondeau and Le Bihan (2005) and Sahuc (2004) for early implementations of this method in a similar context.

Fitting this model to u_t yields an estimate $\hat{\eta}$ of the MA parameter. Let us define $\hat{\mathbf{d}}_{t+1} = (\mathbf{Z}_{t-1} + \mathbf{Z}_t\hat{\eta})\hat{\epsilon}_t$, where $\hat{\epsilon}_t$ is the residual obtained after the first GMM step. Then, a heteroskedasticity- and autocorrelation-consistent estimate of $\mathbf{S}$ is given by

$$\hat{\mathbf{S}} = \frac{1}{T-1}\sum_{t=1}^{T-1} \hat{\mathbf{d}}_{t+1}\hat{\mathbf{d}}'_{t+1}.$$

This is our preferred method since it imposes all the restrictions implied by our model.[8] Notice that, based on simulation experiments, West (1997) concludes that his estimator is efficient precisely when the truncated kernel yields an estimate of $\mathbf{S}$ that is not semidefinite positive. As a complementary remedy, we also used the Newey and West (1987) (NW) consistent estimate with a one-lag window. This approach is close in spirit to that advocated by Eichenbaum and Fisher (2005), though the practical details differ.

Alternatively, if one does not completely believe in the model, allowing for more serial correlation in the GMM errors might be desirable; thus we also use the NW consistent estimate with a twelve-lag window, as in Galí and Gertler (1999).

The vector of instruments $\mathbf{Z}_t$ contains a constant as well as inflation, the labor share, the output gap,[9] and the short-term nominal interest rate. We have thus five moment conditions and two estimated parameters, so that $TJ_T(\psi)$ is asymptotically distributed as a $\chi^2(3)$. Following Eichenbaum and Fisher, and in contrast with former studies relying on the GMM, we use a very small set of instruments. This is not innocent, since it has been shown that in small samples, using too many instruments can generate substantial biases.

To conclude this section, it is important to emphasize that equations (10), (11), (13), and (17) are observationally equivalent in terms of the J statistic. Yet they differ with respect to the implied degree of nominal rigidities. Thus, the question under study is not so much whether these equations pass the overidentification test but

[8] An alternative to this strategy would be to implement the procedure described in Eichenbaum, Hansen, and Singleton (1988).

[9] The output gap is defined as linearly detrended, logged real output. We also considered HP-filtered output and quadratically detrended output and obtained qualitatively similar results. See the next section.

whether they imply an economically reasonable degree of price stickiness. To judge this, we resort to the empirical evidence reported in Dhyne et al. (2006).

3.2 Calibrated Parameters and Data

Since we only estimate α and γ, we must calibrate the remaining parameters. As is conventional in the literature, we set $\beta = 0.99$. The capital depreciation rate δ is set to 2.5 percent. The elasticity of demand θ is set to 10, as is conventional in the literature. We set $\phi = 0.54^{-1}$, thus allowing for a labor share equal to its euro-data empirical counterpart. Notice that we implicitly assume that profits are redistributed proportionately to factor income, so that $1/\phi$ is indeed the steady-state labor share.

The elasticity of capital adjustment costs ϵ is set to 3, as in Woodford (2005). This, with the value set for δ, implies an elasticity of the investment-capital ratio of 13.33. As argued by Eichenbaum and Fisher (2005), this value is consistent with microeconomic evidence reported by Gilchrist and Himmelberg (1995) when it comes to the United States.

Finally, we set $\nu = 1$, as in Altig et al. (2005) and Christiano, Eichenbaum, and Evans (2005). This implies a larger elasticity of labor supply than that estimated by Avouyi-Dovi and Matheron (2004) or Smets and Wouters (2003) on euro-area data. For our purpose, this choice is conservative because setting $\nu = 2$, a value consistent with euro-area data, would greatly reinforce the importance of labor specificity. In the following sections, we explore the sensitivity of our results to this key parameter.

The calibrated parameters are summarized in table 1. Alternatively, we also investigate the calibration considered by Eichenbaum and Fisher (2005). In their benchmark case, they set $\phi = 1.5$ (a labor share equal to 2/3) and $\theta = 11$ (a steady-state markup equal to 10 percent).

Table 1. Calibrated Parameters

ϕ	θ	ν	β	δ	ϵ
1.855	10.000	1.000	0.990	0.025	3.000

The data used in this paper are extracted from the area-wide database compiled by Fagan, Henry, and Mestre (2005). The mnemonics are as follows: real output is YER, nominal output is YEN, the aggregate nominal wage bill is WIN, the GDP deflator is YED, and the short-term nominal interest rate is STN. We construct the labor share as the ratio WIN/YEN. Finally, the inflation rate is computed as the first difference of the logarithm of YED.

3.3 Results

The estimation results are reported in table 2. The table reports estimates for α and γ, as well as the implied price duration. In addition, the table shows the estimated value of TJ_T and the associated p-value. Recall that the four equations considered in the present paper are observationally equivalent in terms of J_T. Similarly, each estimated equation has the same γ. Consequently, table 2 reports TJ_T and $\hat{\gamma}$ only for the first equation.

Following Eichenbaum and Fisher (2005), the 95 percent confidence interval for the average duration is computed as follows. We first determine the 95 percent confidence interval of $\hat{\alpha}$ and then simply transform the latter using the function $\alpha \mapsto (1-\alpha)^{-1}$. Consequently, the obtained confidence band is not symmetric. Notice that in some cases, $\hat{\alpha} + 1.96\text{std}(\hat{\alpha})$ is higher than 1. For these cases, the upper bound is simply denoted by NA, since an infinite price duration is theoretically possible.

The top panel shows results pertaining to our preferred estimation method, i.e., that consisting of fitting an MA(1) to the GMM residuals $u_t(\boldsymbol{\psi})$, following West (1997). In this estimation, we use linearly detrended, logged output as a measure of the output gap. According to the overidentification test, the model cannot be rejected at conventional confidence levels, with a p value of 9.6 percent. The first point to mention is that γ is small, though not very well estimated. Apparently, a small degree of price indexation is required to match the data.[10] Second, as expected, the degree

[10]We investigated a constrained version of our setup, imposing $\gamma = 0$ or $\gamma = 1$. The first constraint is not rejected by the data. From a quantitative point of view, there are no substantial differences between our benchmark estimation and this constrained estimation. To the contrary, the second constraint is not supported and leads to statistical rejection of the model.

Table 2. Benchmark Estimation Results

MA(1) *Fit to GMM Residuals*					
	$\hat{\alpha}$	$\hat{\gamma}$	**Duration**	$\boldsymbol{TJ_T}$	$\boldsymbol{p}$
	Linearly Detrended Output				
Eq. (10)	0.7458 (0.0824)	0.1430 (0.1786)	3.93 [2.41,10.79]	6.35	9.59
Eq. (11)	0.7284 (0.0775)	—	3.68 [2.36,8.36]	—	—
Eq. (13)	0.8508 (0.0561)	—	6.70 [3.86,25.54]	—	—
Eq. (17)	0.9454 (0.0192)	—	18.31 [10.85,58.58]	—	—
	HP-Filtered Output				
Eq. (10)	0.7788 (0.0883)	0.2801 (0.1727)	4.52 [2.54,20.78]	7.68	5.31
Eq. (11)	0.7594 (0.0833)	—	4.16 [2.48,12.93]	—	—
Eq. (13)	0.8730 (0.0589)	—	7.87 [4.12,87.64]	—	—
Eq. (17)	0.9529 (0.0198)	—	21.24 [11.64,121.06]	—	—
	Quadratically Detrended Output				
Eq. (10)	0.7480 (0.0795)	0.1775 (0.1653)	3.97 [2.45,10.41]	6.87	7.60
Eq. (11)	0.7305 (0.0748)	—	3.71 [2.40,8.14]	—	—
Eq. (13)	0.8523 (0.0541)	—	6.77 [3.94,24.03]	—	—
Eq. (17)	0.9459 (0.0184)	—	18.49 [11.08,55.70]	—	—

Notes: The standard errors of α and γ are in parentheses below the empirical estimates. The 95 percent confidence intervals of the average price durations are below the empirical estimates, in brackets.

of nominal rigidity varies a lot depending on the specification of the NKPC. The lowest estimate is obtained in the case when capital is fixed and labor is firm specific (equation (11)). In this case, $\alpha = 0.728$, implying an average price duration of 3.7 quarters, in accordance with microeconomic evidence. Allowing for firm-specific variable capital in addition somewhat increases the degree of nominal rigidity, though to a small extent (equation (10)). In this case, we obtain $\alpha = 0.746$, implying an average price duration of 3.9 quarters. Thus, the model remains in the admissible range, at least when it comes to euro-area data. Our estimates also imply average price durations smaller than those reported by Galí, Gertler, and López-Salido (2001, 2003), even when they allow for firm-specific labor. These durations are always higher than four quarters.

In contrast, when an aggregate labor market is assumed, either with or without firm-specific capital, the probability of not reoptimizing prices appears too high when compared to results reported by Dhyne et al. (2006). When capital is firm specific (equation (13)), we obtain an average price duration of 6.70 quarters; when there are aggregate markets for both capital and labor (equation (17)), this duration reaches the astonishing level of 18.31 quarters.

The middle panel of table 2 reports results obtained when we use the same estimation strategy but use HP-filtered logged output as our measure of the output gap.[11] In this case, the model still passes the overidentification test, with a smaller p-value. The results are less encouraging. The average price durations are always higher than four quarters. Notice, however, that when sampling uncertainty is taken into account, the estimates are still consistent with the European microeconomic evidence. The last panel reports results when we use quadratically detrended output as our measure of the output gap. The model still passes the overidentification test. In this case, we obtain average price durations that resemble those derived in the top panel of the table.

Overall, our results suggest that including firm-specific labor in the NKPC yields estimates of α that deliver both a good statistical fit and an economically plausible degree of nominal rigidity, at least when compared to microeconomic studies on the euro area (Dhyne et al. 2006). This conclusion is similar to that emphasized by Galí, Gertler, and López-Salido (2001), though reached with a slightly different econometric framework. Additionally, we conclude

[11] Our benchmark measure of the output gap might be questioned on the grounds that a linear trend might not be the appropriate measure of potential output. It is thus important to assess the robustness of our results to alternative definitions of this variable.

Table 3. Results with Eichenbaum and Fisher's (2005) Calibration

MA(1) *Fit to GMM Residuals*					
	$\hat{\alpha}$	$\hat{\gamma}$	**Duration**	TJ_T	p
	Linearly Detrended Output				
Eq. (10)	0.7641 (0.0755)	0.1430 (0.1786)	4.24 [2.60,11.38]	6.35	9.59
Eq. (11)	0.7509 (0.0725)	—	4.01 [2.56,9.34]	—	—
Eq. (13)	0.8763 (0.0472)	—	8.08 [4.62,32.07]	—	—
Eq. (17)	0.9454 (0.0192)	—	18.31 [10.85,58.58]	—	—

Notes: The standard errors of α and γ are in parentheses below the empirical estimates. The 95 percent confidence intervals of the average price durations are below the empirical estimates, in brackets.

that allowing for firm-specific capital in addition to firm-specific labor increases moderately the required degree of nominal rigidity necessary to match the data. However, this amount of nominal rigidity remains economically reasonable. In contrast, a model featuring firm-specific capital and aggregate labor implies too high a probability of not reoptimizing prices.

3.4 Sensitivity Analyses

As a first robustness check, we investigate the consequences of adopting the calibration considered by Eichenbaum and Fisher (2005). The results are reported in table 3. To this end, we set $\phi = 3/2$ and $\theta = 11$. The key difference between what is investigated here and the results reported by Eichenbaum and Fisher is that in our context, the parameter ν affects our estimates when it comes to equation (10) and equation (11). This is not the case in Eichenbaum and Fisher's paper, because they abstract from labor specificity. The figure to keep in mind is their estimate of $\alpha = 0.72$ (see their table 3), which they obtain under the assumption of a constant demand elasticity. This framework is equivalent to that underlying equation (13).

In this case, using euro-area data, we obtain $\alpha = 0.88$ when we resort to the same calibration as theirs. This illustrates an important fact. With our benchmark calibration (irrespective of the chosen estimation strategy), we always obtain larger average price durations than they do with U.S. data. Obviously, this can originate either from our calibration or from the data. The present exercise suggests that the data are the correct suspect. This is reminiscent of the contrast arising between the results reported by Bils and Klenow (2004) and by Dhyne et al. (2006). We investigate further the sensitivity of our results below.

Before doing so, we also investigate to what extent our conclusions depend on our choice of a weighting matrix. Table 4 reports results obtained when we use the Newey-West estimator instead of the West estimator. In this case, we fall back to our benchmark calibration. The top panel enforces a one-lag window ($L = 1$). We obtain qualitatively similar conclusions. In this case, the overidentification test is even more supportive of our model, with a p-value of almost 43 percent. The estimated value of γ is almost four times as small as in the top panel and still not statistically different from 0. The ordering of α remains the same, though the estimated values appear somewhat smaller (with higher standard errors also). Equation (11) implies the smallest degree of nominal rigidity, followed by equation (10). Equations (13) and (17) are ranked third and fourth. Once again, the estimated average price durations seem reasonable for the first two specifications and at odds with microeconomic evidence for the last two.

Finally, the bottom panel of table 4 reports results obtained when using the Newey-West estimator with a twelve-lag window ($L = 12$ in table 2). Once again, the overidentification test does not allow us to reject the model, with a p-value slightly above 20 percent. Overall, we obtain very similar results in qualitative terms. Notice, however, that the estimates of α are now substantially higher than in the previous exercises, i.e., higher than 0.8 for all four specifications. Consequently, we obtain average price durations between 5.6 and 30.5 quarters. If we were to take these results seriously, we would conclude that none of the specifications considered can simultaneously match the data and be in accordance with euro-area microeconomic evidence. However, these estimates are obtained using auxiliary assumptions that are not supported by the theory

Table 4. Alternative Estimation Results

	$\hat{\alpha}$	$\hat{\gamma}$	**Duration**	TJ_T	p
	Newey-West HAC, $L = 1$				
Eq. (10)	0.7188 (0.1177)	0.0385 (0.1681)	3.56 [1.95,19.75]	2.78	42.65
Eq. (11)	0.7030 (0.1107)	—	3.37 [1.95,19.75]	—	—
Eq. (13)	0.8322 (0.0815)	—	5.96 [3.05,125.26]	—	—
Eq. (17)	0.9390 (0.0283)	—	16.39 [8.59,178.56]	—	—
	Newey-West HAC, $L = 12$				
Eq. (10)	0.8446 (0.1314)	0.1201 (0.1555)	6.43 [2.42,NA]	4.61	20.27
Eq. (11)	0.8220 (0.1264)	—	5.62 [2.35,NA]	—	—
Eq. (13)	0.9160 (0.0836)	—	11.90 [4.03,NA]	—	—
Eq. (17)	0.9673 (0.0280)	—	30.54 [11.42,NA]	—	—

Notes: The standard errors of α and γ are in parentheses below the empirical estimates. The 95 percent confidence intervals of the average price durations are below the empirical estimates, in brackets.

expounded in the previous section. It is thus unclear whether they should be granted much attention.

To complete this sensitivity analysis, we investigate to what extent our results depend on parameters calibrated prior to estimation. We focus our attention on four key parameters, namely ϵ, θ, ν, and ϕ in the model underlying equation (10).[12] The results are reported in table 5. Overall, we find that when these parameters increase, the probability of not reoptimizing prices decreases. These results are fairly intuitive. When ϵ increases, it becomes more and more difficult to adjust the capital stock, so that in the limit, capital remains fixed. In this case, the model converges to the specification underlying equation (11). Notice that the differences between the results under $\epsilon = 100$ and those under $\epsilon = 3$ are not very big.

[12]Notice that the parameters θ, ν, and ϕ also affect specification (11).

Table 5. Sensitivity Analysis

		ϵ			θ		
		0.01	**1**	**100**	**6**	**11**	**21**
Eq. (10)	$\hat{\alpha}$	0.79 (0.07)	0.76 (0.08)	0.73 (0.08)	0.80 (0.07)	0.73 (0.08)	0.65 (0.10)
	D	4.81 [2.82,16.40]	4.14 [2.48,12.60]	3.70 [2.26,8.51]	5.01 [2.99,15.54]	3.76 [2.31,10.08]	2.83 [1.81,6.44]
Eq. (11)	$\hat{\alpha}$				0.78 (0.06)	0.72 (0.08)	0.63 (0.10)
	D				4.56 [2.88,10.97]	3.54 [2.28,7.94]	2.74 [1.81,5.62]
Eq. (13)	$\hat{\alpha}$	0.91 (0.04)	0.87 (0.05)	0.84 (0.05)	0.88 (0.05)	0.84 (0.06)	0.79 (0.07)
	D	11.73 [6.29,86.36]	7.48 [4.13,39.98]	6.01 [3.71,15.72]	8.48 [4.81,35.78]	6.40 [4.81,35.78]	4.69 [2.80,14.34]
		ν			ϕ		
		0.01	**2**	**10**	$\mathbf{0.7^{-1}}$	$\mathbf{0.6^{-1}}$	$\mathbf{0.5^{-1}}$
Eq. (10)	$\hat{\alpha}$	0.85 (0.06)	0.68 (0.10)	0.45 (0.12)	0.78 (0.07)	0.76 (0.08)	0.73 (0.08)
	D	6.64 [3.82,25.11]	3.14 [1.97,7.67]	1.82 [1.27,3.23]	4.55 [2.79,12.41]	4.17 [2.55,11.45]	3.78 [2.31,10.32]
Eq. (11)	$\hat{\alpha}$	0.83 (0.05)	0.66 (0.09)	0.44 (0.11)	0.77 (0.07)	0.74 (0.07)	0.72 (0.08)
	D	5.92 [3.68,15.14]	2.98 [1.95,6.32]	1.78 [1.27,2.98]	4.31 [2.73,10.22]	3.92 [2.50,9.05]	3.53 [2.27,10.32]
Eq. (13)	$\hat{\alpha}$				0.89 (0.04)	0.87 (0.08)	0.84 (0.06)
	D				8.97 [5.12,36.41]	7.47 [4.28,29.34]	6.25 [3.61,23.29]

Notes: The standard errors of α are in parentheses below the empirical estimates. D stands for average duration. The 95 percent confidence intervals of the average price durations are below the empirical estimates, in brackets. A blank cell indicates that the corresponding equation is not affected by the parameter under study.

This is reminiscent of Eichenbaum and Fisher (2005). The other three parameters contribute to increasing the degree of strategic complementarity between price setters, as discussed by Woodford (2003). It thus comes as no surprise that the required degree of nominal rigidity decreases when these parameters are set to higher values.

4. Conclusion

The recent literature has emphasized the importance of assuming firm-specific capital in optimizing models with nominal rigidity based on the Calvo (1983) specification. Following this literature, we have sought to assess the fruitfulness of this hypothesis when one is concerned solely with obtaining economically realistic estimates of the probability of not reoptimizing prices in NKPC based on euro-area data. By “economically realistic,” we mean that this probability should imply an average price duration compatible with evidence drawn from microeconomic data (Dhyne et al. 2006).

Our objective was to compare the hypothesis of firm-specific capital with the complementary view that labor can also be firm specific. An important aspect of our analysis is our careful implementation of the generalized method of moments. First, following Eichenbaum and Fisher (2005), we select a very small number of instruments. Second, we exploit an important restriction implied by the theory, namely that the forecast error in the orthogonality condition should admit a first-order moving-average representation. Imposing this restriction during the course of the estimation has the merit of transparency.

We obtain two main conclusions. First, we confirm previous findings that allowing for firm-specific labor in a simple model with fixed capital yields a very reasonable estimate of the probability of not reoptimizing prices. The implied average price duration appears consistent with euro-area microeconomic evidence. Second, we found that allowing for firm-specific capital in addition to firm-specific labor results in a higher degree of nominal rigidity. This result is analytically derived by comparing equations (10) and (11). However, the implied probability of not reoptimizing remains acceptable when compared with microeconomic evidence. Thus, from the point of view of empirical realism, the model with labor and capital both firm specific fares well compared with the model with firm-specific labor only.

Appendix. Deriving the NKPC in the Benchmark Model

To make the paper completely self-contained, this appendix specializes the calculations in Woodford (2005) to the particular case under study here. Below, we outline, step by step, the calculations needed to obtain the NKPC.

Optimal Price Setting

The first step in deriving the New Keynesian Phillips curve is to obtain the first-order condition associated with the optimal choice of price. If drawn to reoptimize in period t, firm ς will select $P_t^*(\varsigma)$ so as to maximize

$$\mathrm{E}_{t-1}^{\varsigma}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\frac{\lambda_T}{\lambda_t}\left\{\left(\frac{x_{t,T}P_t^*(\varsigma)}{P_T}\right)^{1-\theta}y_T - s_T(\varsigma)\left(\frac{x_{t,T}P_t^*(\varsigma)}{P_T}\right)^{-\theta}y_T\right\}.$$

Here, $s_t(\varsigma)$ is the Lagrange multiplier associated with firm ς's production constraint. This term corresponds to the real marginal cost of producing an additional unit of intermediate good ς. Notice that firm ς discounts future cash flows according to $\beta^{T-t}\lambda_T/\lambda_t$, where λ_t is the Lagrange multiplier associated with the budget constraint in the representative household's program.

The associated first-order condition writes

$$\mathrm{E}_{t-1}^{\varsigma}\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\frac{\lambda_T}{\lambda_t}y_T\left(\frac{x_{t,T}P_t(\varsigma)}{P_T}\right)^{-\theta}\left\{\frac{x_{t,T}}{\pi_{t,T}}\frac{P_t(\varsigma)}{P_t} - \frac{\theta}{\theta-1}s_T(\varsigma)\right\} = 0.$$

Expressed in log-linear terms, the above equation rewrites

$$\sum_{T=t}^{\infty}(\alpha\beta)^{T-t}\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_t^*(\varsigma) + \hat{x}_{t,T} - \hat{\pi}_{t,T} - \hat{s}_T(\varsigma)\} = 0. \qquad (22)$$

Here, a letter with a hat refers to the log-deviation of the associated variable. We would like to solve this equation for $\hat{p}_t^*(\varsigma)$. Doing so requires that we know the behavior of $\hat{s}_t(\varsigma)$.

Real Marginal Cost

In the context of the present model, an expression for the real marginal cost can be obtained as follows. The first-order condition associated with the optimal choice of labor by firm ς is

$$w_t(\varsigma) = s_t(\varsigma)\frac{1}{\phi}\left(\frac{h_t(\varsigma)}{k_t(\varsigma)}\right)^{\frac{1}{\phi}-1}.$$

At the same time, the optimal supply of labor of type ς by the representative household obeys

$$\lambda_t w_t(\varsigma) = V_h(h_t(\varsigma)).$$

Combining these expressions, and making use of the production function (equation (6)), we arrive at

$$s_t(\varsigma) = \frac{\phi}{\lambda_t} V_h \left(\left(\frac{y_t(\varsigma)}{k_t(\varsigma)} \right)^{\phi} k_t(\varsigma) \right) \left(\frac{y_t(\varsigma)}{k_t(\varsigma)} \right)^{\phi - 1}.$$

Log-linearizing this expression yields

$$\hat{s}_t(\varsigma) = \omega \hat{y}_t(\varsigma) - (\omega - \nu)\hat{k}_t(\varsigma) - \hat{\lambda}_t, \tag{23}$$

where $\omega = \phi(\nu + 1) - 1$ and $\nu = V_{hh}h/V_h$. Using equation (4), we can eliminate $\hat{y}_t(i)$ and obtain

$$\hat{s}_t(\varsigma) = \omega \hat{y}_t - (\omega - \nu)\hat{k}_t(\varsigma) - \hat{\lambda}_t - \omega\theta\hat{p}_t(\varsigma),$$

where $\hat{p}_t(\varsigma)$ is the log-deviation of $P_t(\varsigma)/P_t$. Integrating the above equation yields

$$\hat{s}_t = \omega \hat{y}_t - (\omega - \nu)\hat{k}_t - \hat{\lambda}_t,$$

where $\hat{s}_t$ is interpreted as the average real marginal cost. Subtracting this equation from the previous one, we arrive at

$$\hat{s}_t(\varsigma) = \hat{s}_t - (\omega - \nu)\tilde{k}_t(\varsigma) - \omega\theta\hat{p}_t(\varsigma). \tag{24}$$

Here, $\tilde{k}_t(\varsigma)$ is the log-deviation of the ratio $k_t(\varsigma)/k_t$. Notice that in the case that $P_t(\varsigma)$ has not been reoptimized since period t, this equation rewrites

$$\hat{s}_T(\varsigma) = \hat{s}_T - (\omega - \nu)\tilde{k}_T(\varsigma) - \omega\theta[\hat{p}^*_t(\varsigma) + \hat{x}_{t,T} - \hat{\pi}_{t,T}]. \tag{25}$$

It is clear from either (24) or (25) that the real marginal cost depends on $\tilde{k}_t(\varsigma)$. Hence, the next step is to characterize the dynamic behavior of this variable.

Physical Capital Evolution

The first-order condition with respect to $k_{t+1}(\varsigma)$ is

$$I'\left(\frac{k_{t+1}(\varsigma)}{k_t(\varsigma)}\right) = \beta \mathrm{E}_{t-1}\left\{\frac{\lambda_{t+1}}{\lambda_t}\left[s_{t+1}(\varsigma)\left(1-\frac{1}{\phi}\right)\left(\frac{h_{t+1}(\varsigma)}{k_{t+1}(\varsigma)}\right)^{\frac{1}{\phi}} + \frac{k_{t+2}(\varsigma)}{k_{t+1}(\varsigma)} I'\left(\frac{k_{t+2}(\varsigma)}{k_{t+1}(\varsigma)}\right) - I\left(\frac{k_{t+2}(\varsigma)}{k_{t+1}(\varsigma)}\right)\right]\right\}.$$

Recall that we have assumed that $I(1) = \delta$, $I'(1) = 1$, and $I''(1) = \epsilon > 0$. Log-linearizing the above equation yields

$$\begin{aligned}
&\epsilon(\hat{k}_{t+1}(\varsigma) - \hat{k}_t(\varsigma)) + \hat{\lambda}_t \\
&\quad = \mathrm{E}_{t-1}\big\{\hat{\lambda}_{t+1} + (1-\beta(1-\delta))\big[\hat{s}_{t+1}(\varsigma) + \hat{y}_{t+1}(\varsigma) - \hat{k}_{t+1}(\varsigma)\big] \\
&\qquad + \beta\epsilon(\hat{k}_{t+2}(\varsigma) - \hat{k}_{t+1}(\varsigma))\big\},
\end{aligned}$$

where we made use of equation (6) and of the steady-state relation

$$s\left(1-\frac{1}{\phi}\right)\left(\frac{h}{k}\right)^{\frac{1}{\phi}} = \frac{1-\beta(1-\delta)}{\beta}.$$

Using equation (23) to eliminate $\hat{s}_{t+1}(\varsigma)$, we finally arrive at

$$\begin{aligned}
&\epsilon(\hat{k}_{t+1}(\varsigma) - \hat{k}_t(\varsigma)) + \hat{\lambda}_t \\
&\quad = \mathrm{E}_{t-1}\big\{\beta(1-\delta)\hat{\lambda}_{t+1} + (1-\beta(1-\delta))(1+\omega)\hat{y}_{t+1}(\varsigma) \\
&\qquad - (1-\beta(1-\delta))(1+\omega-\nu)\hat{k}_{t+1}(\varsigma) + \beta\epsilon(\hat{k}_{t+2}(\varsigma) - \hat{k}_{t+1}(\varsigma))\big\}.
\end{aligned}$$

Integrating this equation over the set of intermediate goods-producing firms and subtracting the results from the above equation yields

$$\xi \mathrm{E}_{t-1}\{\hat{p}_{t+1}(\varsigma)\} = \mathrm{E}_{t-1}\big\{\beta\tilde{k}_{t+2}(\varsigma) - \tau\tilde{k}_{t+1}(\varsigma) + \tilde{k}_t(\varsigma)\big\}, \qquad (26)$$

where we made use of equation (4) to eliminate $\hat{y}_{t+1}(\varsigma) - \hat{y}_t$ and we defined the auxiliary parameters

$$\begin{aligned}
\tau &= 1 + \beta + (1-\beta(1-\delta))(1+\omega-\nu)\epsilon^{-1}, \\
\xi &= \theta(1-\beta(1-\delta))\phi(1+\nu)\epsilon^{-1}.
\end{aligned}$$

Thus we are left with a serious problem: the optimal price depends on the real marginal cost. The latter is a function of $\tilde{k}_t(\varsigma)$, which in turn depends on next period's expected price. All this is an obvious consequence of capital specificity: we do not expect a firm that has reoptimized its price to select the same capital stock as a firm stuck with the same price as in the previous period.

Solving the Problem

Following Woodford (2005), we resort to the *undetermined coefficients method* to solve this problem. Let $\hat{p}_t^*$ denote the average value of $\hat{p}_t^*(\varsigma)$ across reoptimizing firms. We assume that

$$\hat{p}_t^*(\varsigma) = \hat{p}_t^* - \mu_{pk}\tilde{k}_t(\varsigma), \tag{27}$$

$$\tilde{k}_{t+1}(\varsigma) = \mu_{kk}\tilde{k}_t(\varsigma) - \mu_{kp}\mathrm{E}_{t-1}\{\hat{p}_t(\varsigma)\}, \tag{28}$$

where μ_{pk}, μ_{kk}, and μ_{kp} are the coefficients to be determined. Notice that the average value of $\tilde{k}_t(\varsigma)$ across reoptimizing firms is zero, since these firms are drawn with uniform probability over the entire set of firms.

First, let us use equation (28) to eliminate $\hat{k}_{t+2}(\varsigma)$ from (26). This yields

$$\xi\mathrm{E}_{t-1}\{\hat{p}_{t+1}(\varsigma)\} = \mathrm{E}_{t-1}\big\{(\beta\mu_{kk} - \tau)\tilde{k}_{t+1}(\varsigma) - \beta\mu_{kp}\hat{p}_{t+1}(\varsigma) + \tilde{k}_t(\varsigma)\big\}. \tag{29}$$

Now, notice that, according to the Calvo specification (augmented with the indexation mechanism considered here), the expected price of firm ς in $t+1$ conditional on information available as of time $t-1$ obeys

$$\mathrm{E}_{t-1}\{\hat{p}_{t+1}(\varsigma)\} = \alpha\mathrm{E}_{t-1}\{\hat{p}_t(\varsigma) + \pi_{t+1} - \gamma\pi_t\} + (1-\alpha)\mathrm{E}_{t-1}\{\hat{p}_{t+1}^*(\varsigma)\}.$$

Using equation (27), we then obtain

$$\begin{aligned}\mathrm{E}_{t-1}\{\hat{p}_{t+1}(\varsigma)\} &= \alpha\mathrm{E}_{t-1}\{\hat{p}_t(\varsigma) + \pi_{t+1} - \gamma\pi_t\} \\ &\quad + (1-\alpha)\mathrm{E}_{t-1}\{\hat{p}_{t+1}^*\} - (1-\alpha)\mu_{pk}\tilde{k}_{t+1}(\varsigma).\end{aligned}$$

Then, resorting to the indexation rule and equation (5), we obtain

$$(1-\alpha)\hat{p}_t^* = \alpha(\hat{\pi}_t - \gamma\hat{\pi}_t). \tag{30}$$

Plugging this equation into the previous one, we obtain

$$\mathrm{E}_{t-1}\{\hat{p}_{t+1}(\varsigma)\} = \alpha\mathrm{E}_{t-1}\{\hat{p}_t(\varsigma)\} - (1-\alpha)\mu_{pk}\tilde{k}_{t+1}(\varsigma).$$

Using this equation and (28) in equation (29), we finally arrive at

$$\begin{aligned}
&[\beta\mu_{kk} + \beta(1-\alpha)\mu_{kp}\mu_{pk} - \tau]\tilde{k}_{t+1}(\varsigma) \\
&\quad = [\beta\alpha\mu_{kp} + \xi(\alpha + (1-\alpha)\mu_{pk}\mu_{kp})]\mathrm{E}_{t-1}\{\hat{p}_t(\varsigma)\} \\
&\qquad - [1 + (1-\alpha)\mu_{pk}\mu_{kk}\xi]\tilde{k}_t(\varsigma).
\end{aligned}$$

Comparing this relation with equation (28), we obtain two restrictions on the μ's:

$$\mu_{kk} = -\frac{1 + (1-\alpha)\mu_{pk}\mu_{kk}\xi}{\beta\mu_{kk} + \beta(1-\alpha)\mu_{kp}\mu_{pk} - \tau} \tag{31}$$

and

$$\mu_{kp} = -\frac{\beta\alpha\mu_{kp} + \xi(\alpha + (1-\alpha)\mu_{pk}\mu_{kp})}{\beta\mu_{kk} + \beta(1-\alpha)\mu_{kp}\mu_{pk} - \tau}. \tag{32}$$

To complete the solution, we need an extra constraint that should derive from the optimal price-setting equation (22)—which we have not used up to now. To do so, let us plug equation (25) into equation (22). After rearranging a little, this yields

$$\begin{aligned}
\frac{1+\omega\theta}{1-\alpha\beta}\hat{p}_t^*(\varsigma) = \sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\big\{&\hat{s}_{t+j} - (1+\omega\theta)(\hat{x}_{t,t+j} - \hat{\pi}_{t,t+j}) \\
&- (\omega - \nu)\tilde{k}_{t+j}(\varsigma)\big\}.
\end{aligned} \tag{33}$$

What complicates this expression is the presence of $\tilde{k}_{t+j}(\varsigma)$. Notice, however, that according to (28),

$$\mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j+1}(\varsigma)\} = \mu_{kk}\mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j}(\varsigma)\} - \mu_{kp}\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+j}(\varsigma)\}, \quad j \geq 0.$$

Iterating over this equation yields

$$\mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j+1}(\varsigma)\} = \mu_{kk}^{j+1}\tilde{k}_t(\varsigma) - \mu_{kp}\sum_{i=0}^{j}\mu_{kk}^{j-i}\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+i}(\varsigma)\}.$$

Thus

$$\sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j}(\varsigma)\}$$
$$= \frac{1}{1-\alpha\beta\mu_{kk}}\tilde{k}_t(\varsigma) - \mu_{kp}\alpha\beta\sum_{j=0}^{\infty}(\alpha\beta)^j\sum_{i=0}^{j}\mu_{kk}^{j-i}\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+i}(\varsigma)\}.$$

Notice that

$$\sum_{j=0}^{\infty}(\alpha\beta)^j\sum_{i=0}^{j}\mu_{kk}^{j-i}\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+i}(\varsigma)\}$$
$$= \frac{1}{1-\alpha\beta\mu_{kk}}\sum_{i=0}^{\infty}(\alpha\beta)^i\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+i}(\varsigma)\},$$

thus

$$\sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j}(\varsigma)\}$$
$$= \frac{1}{1-\alpha\beta\mu_{kk}}\tilde{k}_t(\varsigma) - \frac{\mu_{kp}\alpha\beta}{1-\alpha\beta\mu_{kk}}\sum_{j=0}^{\infty}(\alpha\beta)^j\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+j}(\varsigma)\}.$$

Since we are only considering states of nature where firm ς is not allowed to reset its price after t, it must be the case that

$$\mathrm{E}_{t-1}^{\varsigma}\{\hat{p}_{t+j}(\varsigma)\} = \hat{p}_t^*(\varsigma) + \mathrm{E}_{t-1}\{\hat{x}_{t,t+j} - \hat{\pi}_{t,t+j}\}, \quad j \geq 0.$$

(Recall our notational convention: $\hat{x}_{t,t} = \hat{\pi}_{t,t} = 0$.) Inserting this into the previous equation, we obtain

$$\sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\{\tilde{k}_{t+j}(\varsigma)\}$$
$$= \frac{1}{1-\alpha\beta\mu_{kk}}\tilde{k}_t(\varsigma) - \frac{\mu_{kp}\alpha\beta}{(1-\alpha\beta\mu_{kk})(1-\alpha\beta)}\hat{p}_t^*(\varsigma)$$
$$- \frac{\mu_{kp}\alpha\beta}{1-\alpha\beta\mu_{kk}}\sum_{j=0}^{\infty}(\alpha\beta)^j\mathrm{E}_{t-1}\{\hat{x}_{t,t+j} - \hat{\pi}_{t,t+j}\}.$$

Plugging this relation into (33), we obtain

$$\zeta\hat{p}_t^*(\varsigma) = \sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\{(1-\alpha\beta)[\hat{s}_{t+j} - \zeta(\hat{x}_{t,t+j} - \hat{\pi}_{t,t+j})]\} - \frac{(\omega-\nu)(1-\alpha\beta)}{1-\alpha\beta\mu_{kk}}\tilde{k}_t(\varsigma), \tag{34}$$

where

$$\zeta = 1 + \omega\theta - \frac{(\omega-\nu)\mu_{kp}\alpha\beta}{(1-\alpha\beta\mu_{kk})}. \tag{35}$$

Integrating over the set of reoptimizing firms, we obtain

$$\zeta\hat{p}_t^* = \sum_{j=0}^{\infty}(\alpha\beta)^j \mathrm{E}_{t-1}^{\varsigma}\{(1-\alpha\beta)[\hat{s}_{t+j} - \zeta(\hat{x}_{t,t+j} - \hat{\pi}_{t,t+j})]\}. \tag{36}$$

Using this and equation (27), we finally obtain

$$(1+\omega\theta)\mu_{pk}(1-\alpha\beta\mu_{kk}) - (\omega-\nu)\mu_{kp}\alpha\beta\mu_{pk} = (\omega-\nu)(1-\alpha\beta). \tag{37}$$

Equations (31), (32), and (37) define a nonlinear system in $(\mu_{kk}, \mu_{kp}, \mu_{pk})$, the solution of which depends on $(\alpha, \beta, \delta, \epsilon, \theta, \phi, \nu)$. Once solved for $(\mu_{kk}, \mu_{kp}, \mu_{pk})$, this system implies

$$\varkappa_0(\alpha, \beta, \delta, \epsilon, \theta, \phi, \nu) = \frac{(\phi-1)(\nu+1)\mu_{kp}\alpha\beta}{(1-\alpha\beta\mu_{kk})}.$$

To obtain the NKPC, simply quasi-difference equation (33), and use equation (30) to eliminate $\hat{p}_t^*$.

In the case of firm-specific capital and aggregate labor markets, equation (37) rewrites

$$(1+(\phi-1)\theta)\mu_{pk}(1-\alpha\beta\mu_{kk}) - (\phi-1)\mu_{kp}\alpha\beta\mu_{pk} = (\phi-1)(1-\alpha\beta). \tag{38}$$

Equations (31) and (32) still hold, except that ξ and τ now obey

$$\xi = \theta(1-\beta(1-\delta))\phi\epsilon^{-1}, \quad \tau = 1 + \beta + (1-\beta(1-\delta))\phi\epsilon^{-1}.$$

In this case, we obtain

$$\zeta = 1 + (\phi - 1)\theta - \varkappa_1(\alpha, \beta, \delta, \epsilon, \theta, \phi),$$

with

$$\varkappa_1(\alpha, \beta, \delta, \epsilon, \theta, \phi) = \frac{(\phi - 1)\mu_{kp}\alpha\beta}{(1 - \alpha\beta\mu_{kk})}.$$

References

Altig, D., L. J. Christiano, M. Eichenbaum, and J. Linde. 2005. "Firm-Specific Capital, Nominal Rigidities and the Business Cycle." NBER Working Paper No. 11034.

Avouyi-Dovi, S., and J. Matheron. 2004. "Supply Shocks and Monetary Policy in an Estimated Sticky Price Model of the Euro Area." Banque de France Working Paper No. 126.

Bils, M., and P. J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947–85.

Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan. 2000. "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica* 68 (5): 1151–79.

Christiano, L. J. 2004. "Firm-Specific Capital and Aggregate Inflation Dynamics in Woodford's Model." Unpublished Manuscript, Northwestern University.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

de Walque, G., F. Smets, and R. Wouters. 2004. "Price Setting in General Equilibrium: Alternative Specifications." Mimeo, European Central Bank.

Dhyne, E., L. J. Álvarez, H. Le Bihan, G. Veronese, D. Dias, J. Hoffman, N. Jonker, P. Lünnemann, F. Rumler, and J. Vilmunen. 2006. "Price Changes in the Euro Area and in the United States: Some Facts from Individual Consumer Price Data." *Journal of Economic Perspectives* 20 (2): 171–92.

Eichenbaum, M., and J. M. D. Fisher. 2005. "Evaluating the Calvo Model of Sticky Prices." NBER Working Paper No. 10617, revised version available on M. Eichenbaum's home page, www.faculty.econ.northwestern.edu/faculty/eichenbaum/.

Eichenbaum, M., L. P. Hansen, and K. J. Singleton. 1988. "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty." *Quarterly Journal of Economics* 103 (1): 51–78.

Fagan, G., J. Henry, and R. Mestre. 2005. "An Area-Wide Model for the Euro Area." *Economic Modelling* 22 (1): 39–59.

Galí, J., and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (2): 195–222.

Galí, J., M. Gertler, and J. D. López-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45 (7): 1237–70.

———. 2003. "Erratum to 'European inflation dynamics'." *European Economic Review* 47 (4): 759–60.

Gilchrist, S., and C. P. Himmelberg. 1995. "Evidence on the Role of Cash Flows for Investment." *Journal of Monetary Economics* 36 (3): 541–72.

Hansen, L. P. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–54.

Jondeau, E., and H. Le Bihan. 2005. "Testing for the New Keynesian Phillips Curve. Additional International Evidence." *Economic Modelling* 22 (3): 521–50.

Matheron, J., and T. P. Maury. 2004. "Supply Side Refinements and the New Phillips Curve." *Economics Letters* 82 (3): 391–96.

McAdam, P., and A. Willman. 2004. "Supply, Factor Shares and Inflation Persistence: Re-examining Euro-Area New Keynesian Phillips Curves." *Oxford Bulletin of Economics and Statistics* 66 (September): 637–70.

Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–708.

Sahuc, J. G. 2004. "Partial Indexation and Inflation Dynamics: What Do the Data Say?" *Applied Economics Letters* 11 (13): 827–32.

Sbordone, A. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49 (2): 265–92.

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

Sveen, T., and L. Weinke. 2004. "Pitfalls in the Modelling of Forward-Looking Price Setting and Investment Behavior." Working Paper No. 2004-1, Norges Bank.

———. 2005. "New Perspectives on Capital and Sticky Prices." *Journal of Economic Theory* 123 (1): 21–39.

Taylor, J. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88 (1): 1–24.

West, K. D. 1997. "Another Heteroskedasticity- and Autocorrelation-Consistent Covariance Matrix Estimator." *Journal of Econometrics* 76 (1–2): 171–91.

Woodford, M. 2003. *Interest and Prices.* Princeton, NJ: Princeton University Press.

———. 2005. "Firm Specific Capital and the New Keynesian Phillips Curve." *International Journal of Central Banking* 1 (2): 1–46.

Yun, T. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37 (2): 345–70.

State-Dependent Stock Market Reactions to Monetary Policy*

Troy Davig[a] and Jeffrey R. Gerlach[b]
[a]Federal Reserve Bank of Kansas City
[b]College of William and Mary

This paper presents a test of the response of stock prices to Federal Reserve policy shocks using a Markov-switching framework. The framework endogenously identifies two distinct regimes. The first is a state where the S&P 500 index exhibits a significantly negative response to unexpected changes in the target federal funds rate in the thirty-minute window bracketing FOMC announcements, a result consistent with previous work. However, the model identifies a second regime from September 1998 to September 2002, in which the response of stock prices to policy shocks is insignificant and over ten times more volatile relative to the other regime.

JEL Codes: E44, G12, G14.

1. Introduction

The response of asset prices to Federal Reserve policy is a key component for analyzing the impact of monetary policy on the economy. As Blinder (1998) notes, "Monetary policy has important macroeconomic effects only to the extent that it moves financial market prices that really matter—like long-term interest rates, stock market values, and exchange rates." This paper presents a test of the response of stock prices to Federal Reserve policy shocks using an event-study Markov-switching framework. In the period following

*We thank the editor handling the paper and an anonymous referee for constructive comments. The views expressed herein are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Kansas City or the Federal Reserve System. Corresponding author: Davig: Federal Reserve Bank of Kansas City, 925 Grand Boulevard, Kansas City, MO 64198; e-mail: Troy.Davig@kc.frb.org.

the Federal Reserve's decision to announce changes in the target federal funds rate in 1994, the Markov-switching model identifies a separate state from September 1998 to September 2002. During this period, the volatility of the stock price response to Federal Reserve actions is more than ten times greater than in the remaining periods in the sample. Further, unexpected changes to the target federal funds rate in this high-volatility state have no statistically significant effect on the level of the S&P 500 index in the thirty-minute window bracketing the policy announcement. In the low-volatility state, the market response to unexpected changes in the target federal funds rate is significantly negative and reflects the recent body of work documenting this effect, including Rigobon and Sack (2004), Bernanke and Kuttner (2005), and Gurkaynak, Sack, and Swanson (2005).

A recent example using the event-study framework to assess the effects of monetary policy on the stock market is Bernanke and Kuttner (2005), who use daily CRSP value-weighted returns and a measure of unexpected changes to the target federal funds rate computed from federal funds futures contract prices the day prior to a change to the target rate. They find that an unexpected 25-basis-point cut in the target federal funds rate is associated with a 1 percent increase in equity prices. A similar approach is used by Gurkaynak, Sack, and Swanson (2005), who use high-frequency data to overcome the potential endogeneity and omitted-variables problems associated with using data covering a broad time window around policy announcements. By focusing on a narrow window around policy changes, they isolate the impact of unexpected moves in the target federal funds rate on equity prices, finding that the S&P 500 increases slightly more than 1 percent in response to a surprise 25-basis-point cut. Related literature measuring the response of equity returns to monetary policy using the event-study framework also includes Patelis (1997), Thorbecke (1997), D'Amico and Farka (2002), Bomfin (2003), Craine and Martin (2003), and Bentzen et al. (2004).[1]

[1] A broad literature exists assessing the reaction of bond prices to monetary policy using the event-study approach, such as Cook and Hahn (1989), Roley and Sellon (1998), Thornton (1998), Bomfin and Reihnart (2000), Kuttner (2001), Cochrane and Piazzesi (2002), and Poole, Rasche, and Thornton (2002).

Rigobon and Sack (2004) develop a heteroskedasticity-based technique for estimating the impact of monetary policy on asset prices and report that an unexpected 25-basis-point decrease in the three-month eurodollar futures rate results in a 1.7 percent increase in the S&P 500 index.

Jensen and Mercer (2002) allow for variation of asset returns across different states. They use three measures of monetary policy shocks—(i) changes in the Federal Reserve discount rate, (ii) changes in the target federal funds rate, and (iii) the Boschen and Mills (1995) monetary policy index—to separate Federal Reserve policy into expansive and restrictive periods. However, the different periods are exogenously specified and do not measure the response to unexpected monetary policy announcements. The framework in this paper endogenously detects different states and estimates the corresponding state-dependent response of stock prices to monetary policy.

The outline of the paper is as follows: Section 2 describes the data, section 3 discusses the Markov-switching framework, and section 4 presents the empirical results. Section 5 reports results from a set of robustness checks, and section 6 concludes.

2. Data

The sample consists of the eighty scheduled announcements associated with Federal Open Market Committee meetings from the beginning of 1994, when the Federal Reserve began to announce its policy decisions, through the end of 2003. Following Kuttner (2001), we use daily thirty-day federal funds rate futures, available from the Chicago Board of Trade, to measure the unexpected component of Federal Reserve policy decisions.[2] Our measure of shocks is identical to the daily monetary policy surprises in Gurkaynak, Sack, and Swanson (2005). Tick data provided the stock market data, which

[2]Krueger and Kuttner (1996) and Rudebusch (1998) confirm that federal funds futures prices are efficient. Söderström (2001) finds that federal funds futures rates have weak predictive power using daily data but are much more successful for predicting the average funds rate and funds rate changes around target changes and meetings of the FOMC. Poole, Rasche, and Thornton (2002) document that the futures market is better able to anticipate policy changes since 1994, when the Federal Reserve began to announce publicly its policy decisions.

consist of high-frequency observations of the S&P 500 index. The stock market returns are log differences of the S&P 500 index at the beginning and end of a thirty-minute window around announcements following FOMC meetings. The thirty-minute window is intended to be long enough to avoid market-microstructure issues but short enough to limit the endogeneity and omitted-variables problems associated with measuring the market response to Federal Reserve actions in daily data.

3. State-Dependent Market Reactions

The general approach follows an event-study framework, with the econometric specification deriving from the Markov-switching approach in Hamilton (1989). This framework allows for, but does not require, the market return to respond differently to unexpected changes in the target federal funds rate across different periods of the sample. The specification is

$$H_t = a + b^u(S_t)\Delta i_t^u + \varepsilon_t, \qquad (1)$$

where H_t is the market return, S_t is the unobserved state variable, Δi_t^u is the unexpected change in the target federal funds rate, and $\varepsilon_t \sim N(0, \sigma(S_t)^2)$. This specification allows the variance of the error term to vary with the state, requiring that $b^u(S_t)$ and $\sigma(S_t)$ switch synchronously.

A two-state Markov chain governs the evolution of the unobserved state,

$$\Pi = \begin{bmatrix} p_{00} & 1 - p_{00} \\ 1 - p_{11} & p_{11} \end{bmatrix}, \qquad (2)$$

where $p_{ij} = \Pr[S_t = j | S_{t-1} = i]$ for $i = 0, 1$ and $j = 0, 1$. A non-linear iterative filter formulates probabilistic estimates of the state and constructs the log-likelihood function, where estimation proceeds by maximizing the likelihood function (see Hamilton 1989; Kim and Nelson 1999). The model given by (1) and (2) specifies two states but provides for the possibility that only one state is in place over the sample period.[3] The data determine the properties and timing of each regime.

4. Results

4.1 High- and Low-Volatility States

Table 1 reports parameter estimates for the Markov-switching model using observations corresponding to scheduled FOMC announcements. Estimation endogenously delineates the sample between two regimes, one running from the beginning of the sample in February 1994 until November 1994 and then recurring from September 1998 through September 2002. Two primary factors characterize this regime: (i) the volatility of the market return is eleven times greater than in the other regime and (ii) the response of stock prices to

Table 1. Maximum-Likelihood Estimates

Parameter	**Scheduled FOMC Meetings**	**All Observations Including Intermeeting Moves**
a	$-.083^{**}$	$-.076^{*}$
	(.031)	(.031)
$b^u(0)$: Low Volatility	-1.914^{**}	-1.622^{**}
	(.635)	(.554)
$b^u(1)$: High Volatility	-1.547	-6.881^{**}
	(1.751)	(.675)
$\sigma(0)^2$: Low Volatility	$.028^{**}$	$.030^{**}$
	(.007)	(.008)
$\sigma(1)^2$: High Volatility	$.320^{**}$	$.423^{**}$
	(.066)	(.091)
ln Likelihood	-28.007	-37.5374

Note: * and ** denote significance at the 5 percent and 1 percent level, respectively.

[3]Due to data limitations, we do not consider more than two states as a reasonable description of the data.

an unexpected increase in the target federal funds rate is statistically insignificant. Thus, the direction of the stock price response to changes in Federal Reserve policy is less predictable and more volatile.

The low-volatility state corresponds to December 1994 through August 1998, recurring from 2002 to the end of the sample. Key attributes of this state are (i) low volatility of the market return and (ii) a significantly negative response to unexpected increases in the target federal funds rate. A 25-basis-point unexpected reduction in the target federal funds rate increases the S&P 500 index around 50 basis points, a response approximately half of that estimated by Bernanke and Kuttner (2005) and Gurkaynak, Sack, and Swanson (2005).

Figure 1 gives filtered and smoothed probabilities of being in the high-volatility state. The filtered probability of being in state j

Figure 1. Probability of High-Volatility State

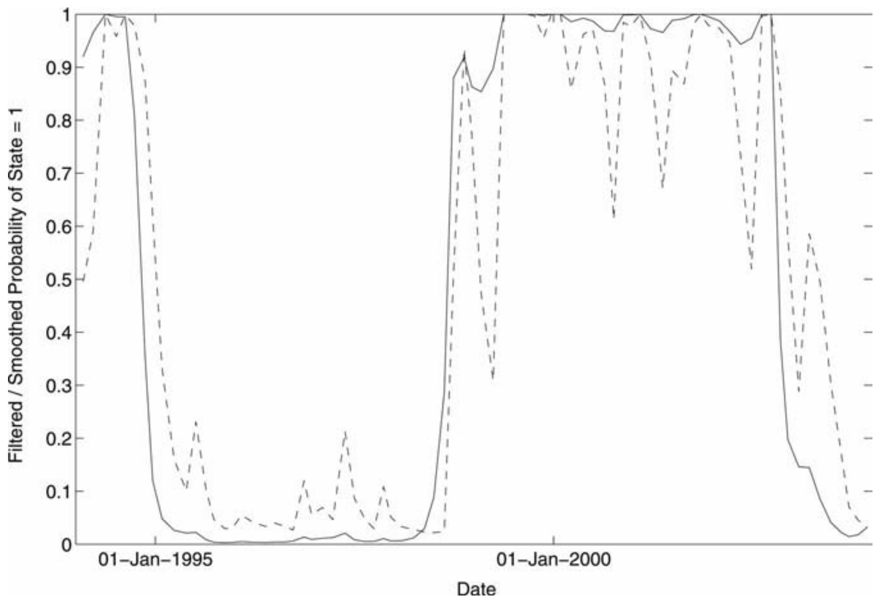

Note: The dashed line denotes filtered probabilities, while the solid line denotes smoothed probabilities.

at t is $\Pr[S_t = j|\Omega_t]$, where Ω_t is the information set that includes all past observations and those at time t. The smoothed probability is $\Pr[S_t = j|\Omega_T]$, where Ω_T includes information over the entire sample.

Ambiguity in the press releases for the first three public announcements of changes to the target federal funds rate likely causes four observations from 1994 to be placed in the high-volatility state. The Federal Reserve started officially announcing changes to the target federal funds rate at the beginning of 1994, but it wasn't until May 17, 1994, that it officially announced the exact target federal funds rate and provided a brief statement regarding the views of the FOMC about the state of the economy.[4] Thus, the market had less information regarding target federal funds rate changes prior to May 1994, resulting in a more uncertain and volatile market response. For this reason, the inclusion of these early observations into the high-volatility state appears reasonable.

Figures 2 and 3 are scatter plots of the unexpected component of changes in the target federal funds rate against the market return, broken out by state. Figure 2 is the low-volatility state, illustrating the negative relationship between the unexpected component of changes in the target federal funds rate and the market response, as well as the low variability of the response. Figure 3 is the high-volatility state, where the volatility of the response is more pronounced relative to the low-volatility state. In figure 3, the hollow markers denote intermeeting moves and are discussed in the following section.

Using data from Gurkaynak, Sack, and Swanson (2005) over the same sample period yields the same qualitative results for each regime. The difference in data used in their paper rests with how the unexpected change in the target federal funds rate is computed. They use the thirty-minute window surrounding the policy announcement, whereas our data follow Bernanke and Kuttner (2005), using the previous day's closing price. Both sets of data

[4]For example, the entire press release following the FOMC meeting on April 18, 1994, read as follows: "Chairman Alan Greenspan announced on April 18, 1994, that the Federal Reserve would increase slightly the degree of pressure on reserve positions. This action was expected to be associated with a small increase in short-term money market interest rates."

Figure 2. Low-Volatility State

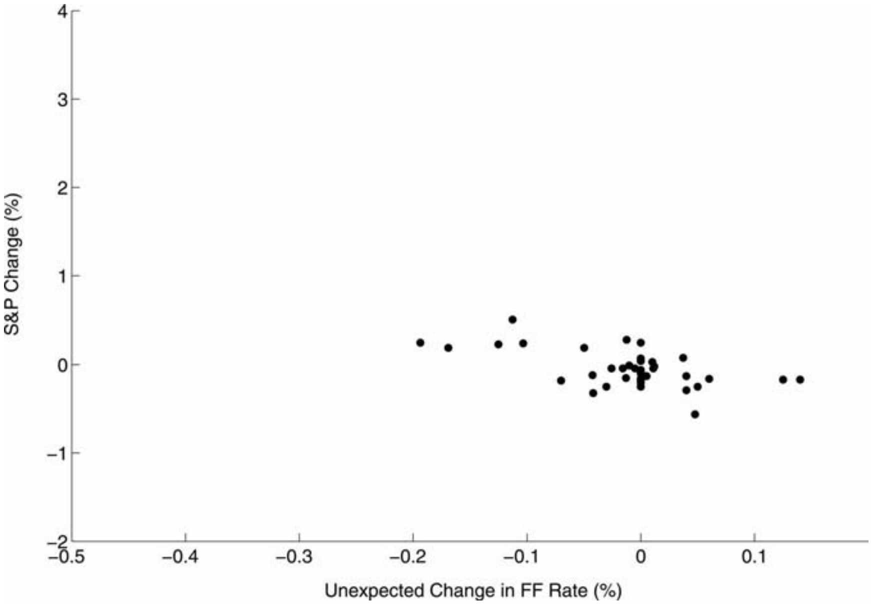

use the change in the S&P 500 index in the thirty-minute window surrounding the policy announcement. The one difference arising from the different data is the timing of the regimes, where the high-volatility regime extends to the end of the sample with data from Gurkaynak, Sack, and Swanson (2005), resulting in fifty-one of the eighty observations falling into this state. However, the probabilistic inference of being in the high-volatility state begins declining in September 2003. Otherwise, the timing of the states is in close agreement.

There are several possible explanations for changes in volatility in the sample. One possibility is that during the mid to late 1990s, traders simply faced increased noise in estimating current asset values. Thus, traders could not sort out the effects of monetary policy as well, causing a less-predictable market response to policy shocks.

A second possibility is that the market interpreted Federal Reserve actions more as a "signal" revealing private information during the high-volatility state, rather than altering underlying

Figure 3. High-Volatility State

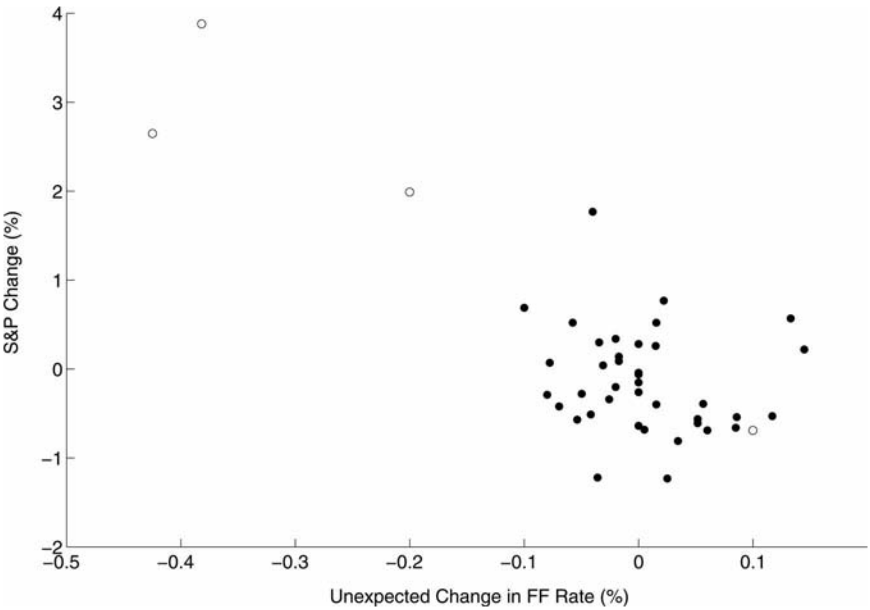

Note: Hollow markers denote intermeeting target federal funds rate changes.

fundamentals in financial markets.[5] For example, the market may have been more likely to interpret a rate increase in the high-volatility state as a signal from the Federal Reserve that the economy is stronger than expected. This interpretation suggests that the market simply interpreted information differently across the sample.[6]

[5]Market participants may perceive the Federal Reserve as possessing private information due to data gathering and analysis capabilities. Work addressing this issue includes Peek, Rosengren, and Tootell (1999, 2003), Romer and Romer (2000), and Faust, Swanson, and Wright (2004).

[6]Amato, Morris, and Shin (2002) provide an example of how agents may respond to noisy signals. In their model, agents have incentives to learn about economic fundamentals and coordinate their actions with those of other agents in the economy. Agents receive a private signal and a public signal, both containing information about the state of the economy and noise. The coordination incentive can cause agents to overreact to the public information.

4.2 Specification Testing

To establish the Markov-switching approach as an appropriate framework to address the stock market response to monetary policy, we first test the standard ordinary least squares (OLS) specification against the Markov-switching alternative. This test suffers from the problem in Davies (1977), in which some nuisance parameters are not identified under the null hypothesis. Consequently, the likelihood ratio statistic has a nonstandard distribution. Using the approach in Hansen (1992), the standardized likelihood test statistic is 3.16 with a 1 percent critical value of 2.36 and p-value less than 1e-6, indicating that the standard OLS model with time-invariant parameters is formally rejected in favor of the Markov-switching alternative.[7]

Given the close point estimates and large confidence interval around $b^u(1)$, a reasonable alternative model restricts the response of returns to be equal across regimes but allows the volatility to switch,

$$H_t = a + b^u \Delta i_t^u + \varepsilon_t, \tag{3}$$

where $\varepsilon_t \sim N(0, \sigma(S_t)^2)$. The restriction on (1) yielding (3) is given by

$$H_0 : b^u(0) = b^u(1).$$

We cannot reject this restriction, since the log-likelihood value for the model in (3) is -28.031, a value close to the unrestricted log-likelihood value of -28.007.[8] The estimate of b^u from (3) is -1.87 and significant, a value close to that in the low-volatility regime. Estimates of the variance in each regime in the restricted model are

[7]The critical values are actually bounds arising from simulating the asymptotic distribution and, as Hansen (1992) points out, are conservative estimates, which strengthens the result that the OLS model can be rejected in favor of the Markov-switching alternative.

The result of the test is invariant for $M = 1, 2, 3, 4$, where M is the bandwidth number in the Bartlett kernel used in simulating the covariance function (see Hansen 1996). The critical value reported is for $M = 4$.

[8]Given that not all parameters are restricted to be equal across both states, this approach to specification testing does not suffer from the problem of nuisance parameters not being identified under the null hypothesis (see Davies 1977).

nearly unchanged. This test, as well as the test against OLS, underscores the clear evidence for time variation in the variance of returns but raises questions regarding the response of returns in each regime. It is the case that returns in the unrestricted model do not respond significantly in the high-volatility regime. However, the above test reflects that the insignificant response in the high-volatility regime may be a consequence of increased noise instead of a breakdown in the relation between surprise changes in the target federal funds rate and returns.

Alternatively, we can specify a more-general model than (1) by allowing the intercept term to be state dependent and including the expected change to the target federal funds rate as an additional explanatory variable, also with a state-dependent coefficient. For example,

$$H_t = a(S_t) + b^u(S_t)\Delta i_t^u + b^e(S_t)\Delta i_t^e + \varepsilon_t, \tag{4}$$

where Δi_t^e is the expected change and $\varepsilon_t \sim N(0, \sigma(S_t)^2)$. The total change in the target federal funds rate is $\Delta i_t = \Delta i_t^e + \Delta i_t^u$.

The joint set of restrictions on the more-general model (4) that yield (1) is given by

$$H_0 : b^e(0) = b^e(1) = 0, \quad a(0) = a(1).$$

The null hypothesis tests whether the coefficients on the expected component in both states are zero and whether unconditional returns, governed by the intercept, are independent of the state. Setting the coefficient on Δi_t^e to zero in both states is justified on economic grounds, since the market should not respond to correctly anticipated changes to the target federal funds rate. The above restrictions cannot be rejected at conventional significance levels, supporting the specification given in (1). The statistically insignificant coefficients on expected changes to the target federal funds rate are consistent with estimates in Bernanke and Kuttner (2005) and the specification in Gurkaynak, Sack, and Swanson (2005). Also, the filtered and smoothed probabilistic estimates of the states under specifications (1) and (4) are in close agreement.

Further analysis suggests that the Markov-switching model is capable of quickly detecting a change between the low- and high-volatility states. The framework could have quickly detected the

Table 2. OLS Estimates by Regime

	Low Volatility	High Volatility
Observations	41	39
Intercept	$-.075^{**}$	$-.134$
	(.028)	(.093)
Unexpected Change	-1.821^{**}	-1.761
	(.421)	(1.638)
R^2	.335	.030

Note: ** denotes significance at the 1 percent level.

change in the relationship between unexpected changes in the target rate and stock prices. For example, using a truncated data set with only the first forty-nine observations from the sample, estimation of the model indicates a regime switch at approximately the same time as would occur if using the full sample. Since the data for such analysis are easily available and not subject to revision, the Markov-switching framework could possibly allow the Federal Reserve to learn within the course of only a few meetings if the market had entered a state where policy changes have a significantly more volatile effect on stock prices.

4.3 Comparison to OLS

To compare with other results, such as Bernanke and Kuttner (2005) and Gurkaynak, Sack, and Swanson (2005), table 2 reports OLS estimates for subsamples corresponding to the high- and low-volatility states. The coefficient estimates on the unexpected component for each state are similar to those from the Markov-switching model—significantly negative in the low-volatility state and negative, but not significant, in the high-volatility state. However, of primary interest is the R^2, which is .33 in the low-volatility state and .03 in the high-volatility state. The difference in R^2 is a clear indication of the higher volatility and unpredictability of the market response to monetary policy in the high-volatility state.

In Gurkaynak, Sack, and Swanson (2005), similar conclusions arise using high-frequency data in the thirty-minute window surrounding FOMC announcements. The coefficient on the unexpected component is significantly negative using observations in the low-volatility state and not significantly different than zero in the high-volatility state. The corresponding R^2 is .50 in the low-volatility state and .04 in the high-volatility state. The R^2 using the sample of observations from the beginning of 1994 to the end of 2003 corresponding only to scheduled FOMC meetings yields a significantly negative coefficient on the unexpected component and an R^2 of .07.

Comparing OLS with the Markov-switching estimates underscores that using OLS when regime changes occur yields estimates that are a weighted average of the responses in each regime. OLS estimates using the entire sample indicate that the market responds negatively, on average, to an unexpected increase in the target federal funds rate. However, OLS does not capture the changes in volatility that are apparent when using a Markov-switching framework.

5. Robustness Checks

5.1 Intermeeting Moves

The analysis thus far only considers the effects of changes to the target federal funds rate during scheduled FOMC meetings. Over the sample, there are five intermeeting changes to the target federal funds rate: April 18, 1994; October 15, 1998; January 3, 2001; April 18, 2001; and September 17, 2001.[9] These intermeeting rate changes differ from regularly scheduled FOMC meetings in that the unexpected component of the intermeeting moves is relatively large, since these changes are unexpected by their very nature.

The four intermeeting changes all occur during the high-volatility state and are represented by the hollow circles in figure 3. These observations include the three largest movements in both asset prices and the unexpected component of changes to the target federal funds rate.[10] Inspection reveals that the four observations line up precisely

[9]The move on September 17, 2001, the first day of trading on the New York Stock Exchange after September 11, is considered an anomaly and, following Bernanke and Kuttner (2005) and Gurkaynak, Sack, and Swanson (2005), is excluded from the analysis.

[10]Using influence analysis, Bernanke and Kuttner (2005) exclude these three intermeeting moves from their study.

Figure 4. Probability of High-Volatility State Including Intermeeting Target Changes

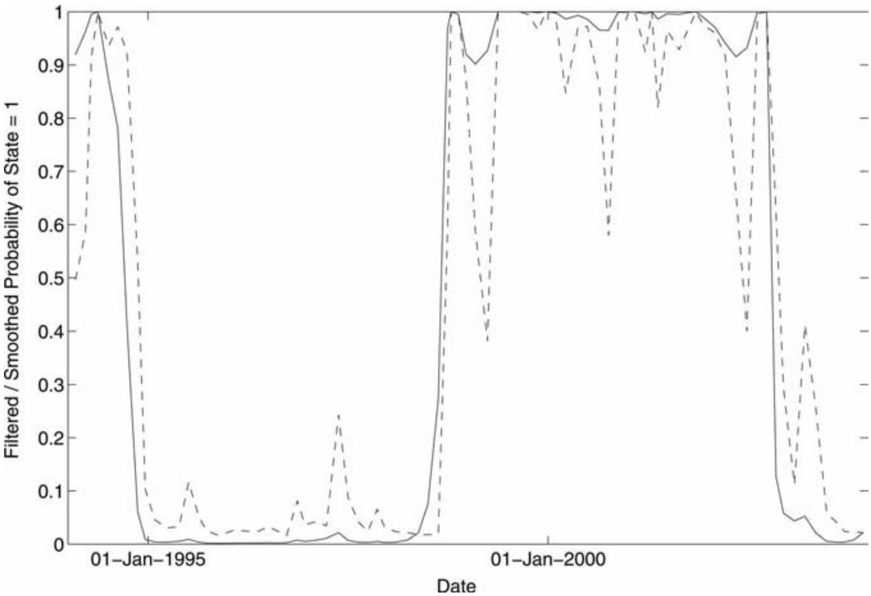

Note: The dashed line denotes filtered probabilities, while the solid line denotes smoothed probabilities.

and that including them in the sample strongly influences the coefficient estimate of the market response to the unexpected component in the high-volatility regime.

Table 1 provides estimates using the sample including intermeeting moves, and figure 4 provides the filtered and smoothed probabilistic estimates of the high-volatility state. The timing of the low- and high-volatility states is largely unaffected, as comparison of figures 1 and 4 indicates. The market response in the low-volatility state is −1.622 and significant, similar to the estimate using the sample excluding intermeeting moves. The primary difference arises in the high-volatility regime, where the market response is −6.881 and significant, contrasting to the estimate from the sample excluding intermeeting moves. Thus, unlike shocks that occur following regularly scheduled meetings, wholly unexpected rate changes appear to have a powerful effect on the market during the high-volatility state.[11]

5.2 Asymmetry, Discount Rate Changes, and Policy Reversals

As final robustness checks, we assess whether asymmetry, discount rate changes, or policy reversals have an effect on the results. Asymmetry exists if the market response depends on the sign of the unexpected change to the target federal funds rate. In other words, we test whether surprise increases affect the market differently from surprise decreases. Controlling for discount rate changes and policy reversals may also be important, given that such changes may convey additional information on the state of the economy to the public.

To account for each of these effects, we use the model given by (1) with the addition of a dummy variable,

$$H_t = a + b^u(S_t)\Delta i_t^u + b^d d_t + \varepsilon_t, \tag{5}$$

where d_t is either 0 or 1. Note that b^d is not state dependent. Results are given in table 3 for all the robustness checks. The coefficient on the dummy variable is insignificant and does not affect the timing of the low- and high-volatility states, indicating that asymmetry is not playing a role in the results.

To control for changes to the discount rate, $d_t = 1$ for the twenty observations when a change occurs and 0 otherwise.[12] The coefficient on the dummy is insignificant, and the timing of the states is unaffected. For policy reversals, $d_t = 1$ for the six observations when there was a change to the target federal funds rate in the opposite direction relative to the previous change. Again, the coefficient on the dummy is insignificant, and the timing of the states is unaffected.

[11] An important point to keep in mind, however, is that these results stem from adding only four observations to the sample, constituting a small subsample.

[12] We ignore May 19, 2000, and January 4, 2001, two dates when the discount rate changed, but the target federal funds rate did not.

Table 3. Robustness Checks

Parameter	Asymmetry	Discount Rate Changes	Policy Reversals
a	$-.101^{**}$	$-.090^{**}$	$-.095^{**}$
	(.032)	(.033)	(.037)
$b^u(0)$: Low Volatility	-2.304^{**}	-2.002^{**}	-2.097^{**}
	(.679)	(.660)	(.711)
$b^u(1)$: High Volatility	−1.838	−1.523	−1.601
	(1.839)	(1.836)	(1.694)
d	.139	.047	.150
	(.103)	(.082)	(.092)
$\sigma(0)^2$: Low Volatility	$.031^{**}$	$.030^{**}$	$.026^{**}$
	(.008)	(.008)	(.008)
$\sigma(1)^2$: High Volatility	$.302^{**}$	$.319^{**}$	$.312^{**}$
	(.075)	(.073)	(.072)
ln Likelihood	−27.155	−27.899	−26.985

Note: ** denotes significance at the 1 percent level.

6. Conclusion

Beginning in 1994, when the Federal Reserve began to announce publicly its policy decisions following FOMC meetings, through the end of 2003, there are two distinct regimes characterizing the response of the S&P 500 index to unexpected changes in the target federal funds rate—a high- and low-volatility regime. During the high-volatility regime, from September 1998 through September 2002, the market response to unexpected changes in the target federal funds rate following scheduled FOMC meetings is highly variable and not significantly different from zero. During the low-volatility regime, unexpected changes in the target federal funds rate are much less volatile and negatively related to stock prices. The results are robust when controlling for various factors, such as asymmetric responses, discount rate changes, and policy reversals.

References

Amato, J. D., S. Morris, and H. S. Shin. 2002. "Communication and Monetary Policy." *Oxford Review of Economic Policy* 18 (4): 495–503.

Bentzen, E., P. R. Hansen, A. Lunde, and A. Zebedee. 2004. "The Greenspan Effect on Equity Markets: An Intraday Examination of U.S. Monetary Policy Announcements." Mimeo, Copenhagen Business School.

Bernanke, B. S., and K. N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3): 1221–57.

Blinder, A. S. 1998. *Central Banking in Theory and Practice*. Cambridge, MA: MIT Press.

Bomfin, A. N. 2003. "Pre-announcement Effects, News Effects, and Volatility: Monetary Policy and the Stock Market." *Journal of Banking and Finance* 27 (1): 133–51.

Bomfin, A. N., and V. R. Reihnart. 2000. "Making News: Financial Market Effects of Federal Reserve Disclosure Practices." Finance and Economics Discussion Paper No. 2000/14, Board of Governors of the Federal Reserve System.

Boschen, J. F., and L. O. Mills. 1995. "The Relation Between Narrative and Money Market Indicators of Monetary Policy." *Economic Inquiry* 33 (1): 24–44.

Cochrane, J. H., and M. Piazzesi. 2002. "The Fed and Interest Rates—A High-Frequency Identification." *American Economic Review* 92 (2): 90–95.

Cook, T., and T. Hahn. 1989. "The Effect of Changes in the Federal Funds Rate Target on Market Interest Rates in the 1970s." *Journal of Monetary Economics* 24 (3): 331–51.

Craine, R., and V. Martin. 2003. "Monetary Policy Shocks and Security Market Response." Mimeo, University of California–Berkeley.

D'Amico, S., and M. Farka. 2002. "The Fed and Stock Market: A Proxy and Instrumental Variable Identification." Mimeo, Columbia University.

Davies, R. 1977. "Hypothesis Testing When a Nuisance Parameter Is Present Only Under the Alternative." *Biometrika* 64 (2): 247–54.

Faust, J., E. T. Swanson, and J. H. Wright. 2004. "Do Federal Reserve Policy Surprises Reveal Superior Information about the Economy?" *Contributions to Macroeconomics* 4 (1): 1–29.

Gurkaynak, R., B. Sack, and E. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1): 55–93.

Hamilton, J. D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57 (2): 357–84.

Hansen, B. E. 1992. "The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP." *Journal of Applied Econometrics* 7:S61–S82.

———. 1996. "Erratum: The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP." *Journal of Applied Econometrics* 11 (2): 195–98.

Jensen, G. R., and J. M. Mercer. 2002. "Monetary Policy and the Cross-Section of Expected Stock Returns." *The Journal of Financial Research* 25 (1): 125–39.

Kim, C.-J., and C. R. Nelson. 1999. *State-Space Models with Regime Switching*. Cambridge, MA: MIT Press.

Krueger, J. T., and K. N. Kuttner. 1996. "The Fed Funds Futures Rate as a Predictor of Federal Reserve Policy." *Journal of Futures Markets* 16 (8): 865–79.

Kuttner, K. N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.

Patelis, A. D. 1997. "Stock Return Predictability and the Role of Monetary Policy." *Journal of Finance* 52 (5): 1951–72.

Peek, J., E. S. Rosengren, and G. M. B. Tootell. 1999. "Is Bank Supervision Central to Central Banking?" *Quarterly Journal of Economics* 114 (2): 629–53.

———. 2003. "Does the Federal Reserve Possess an Exploitable Informational Advantage?" *Journal of Monetary Economics* 50 (4): 817–39.

Poole, W., R. H. Rasche, and D. L. Thornton. 2002. "Market Anticipations of Monetary Policy Actions." *Review* (Federal Reserve Bank of St. Louis) 84 (4): 65–94.

Rigobon, R., and B. Sack. 2004. "The Impact of Monetary Policy on Asset Prices." *Journal of Monetary Economics* 51 (8): 1553–75.

Roley, V. V., and G. H. Sellon. 1998. "Market Reaction to Monetary Policy Nonannouncements." Research Working Paper No. 98-06, Federal Reserve Bank of Kansas City.

Romer, C. D., and D. H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90 (3): 429–57.

Rudebusch, G. D. 1998. "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review* 39 (4): 907–31.

Söderström, U. 2001. "Predicting Monetary Policy with Federal Funds Futures Prices." *Journal of Futures Markets* 21 (4): 377–91.

Thorbecke, W. 1997. "On Stock Market Returns and Monetary Policy." *Journal of Finance* 52 (2): 635–54.

Thornton, D. L. 1998. "Tests of the Market's Reaction to Federal Funds Rate Target Changes." *Review* (Federal Reserve Bank of St. Louis) 80 (6): 25–36.

Monetary Policy Inertia: Fact or Fiction?*

Glenn D. Rudebusch
Federal Reserve Bank of San Francisco

Many interpret estimated monetary policy rules as suggesting that central banks conduct very sluggish partial adjustment of short-term policy interest rates. In contrast, others argue that this appearance of policy inertia is an illusion and simply reflects the spurious omission of important persistent influences on the actual setting of policy. Similarly, the real-world implications of the theoretical arguments for policy inertia are open to debate. However, empirical evidence on policy gradualism obtained by examining expectations of future monetary policy embedded in the term structure of interest rates is definitive and indicates that the actual amount of policy inertia is quite low.

JEL Codes: E44, E52.

1. Introduction

In recent years, there has been a clear shift in the focus of monetary policy research. While a decade or so ago, monetary aggregates were often used to model monetary policy, now the most common representation uses a short-term interest rate as the monetary policy instrument. Indeed, the literature on how central banks manipulate policy interest rates has grown very rapidly.[1] Especially since the introduction of the now-ubiquitous Taylor (1993) rule, many

*For helpful comments, I thank participants at the 2005 Riksbank Inflation Targeting conference, especially my discussant Axel Weber; colleagues throughout the Federal Reserve System, especially Bill English; and the editor Michael Woodford. Golnaz Motiey and Vuong Nguyen provided excellent research assistance. The views expressed in this paper are my own and should not necessarily be attributed to others at the Federal Reserve Bank of San Francisco. Author contact: Home page: www.frbsf.org/economists/grudebusch; e-mail: Glenn.Rudebusch@sf.frb.org.

[1]For some of the arguments underlying the shift away from monetary aggregates, see Rudebusch and Svensson (2002).

researchers have examined monetary policy rules or reaction functions that relate the policy interest rate to a small set of observables.[2] There has been voluminous research on the optimal design of such policy rules and on the empirical estimation of these rules using historical data. Many important normative and positive issues regarding the form of these rules have been considered—notably, the choice of the relevant argument variables in the rules and the nature of the dynamic adjustment embodied in the rules. This paper will examine the latter issue and broadly characterize the amount of monetary policy inertia or partial adjustment contained in optimal and empirical interest rate rules.

The dynamic adjustment process of monetary policy is a particularly interesting topic because of the lively debate about its nature. However, at the outset, it should be noted that this debate is largely limited to interest rate movements at a quarterly frequency, which is the relevant frequency for the empirical macroeconomic policy rules literature. In contrast, at a higher frequency—daily, weekly, or even monthly—the existence of a *short-run* smoothing of policy rates by central banks is widely acknowledged. Such short-term partial adjustment involves, for example, cutting the policy rate by two 25-basis-point moves in fairly quick succession, rather than reducing the rate just once by 50 basis points.[3] However, short-term partial adjustment *within a quarter* is essentially independent of whether there is monetary policy inertia *over the course of several quarters*, and this latter issue is the one that is relevant for estimated monetary policy rules and is discussed below.[4]

The debate about the dynamic adjustment of central bank policy rates focuses on the persistent quarterly cyclical fluctuations in

[2]See Svensson (2003) for a discussion of targeting rules as an alternative representation.

[3]For example, as described in Rudebusch (1995), central banks tend to adjust their policy interest rates in sequences of relatively small steps with only rare reversals of direction.

[4]Indeed, as described in Rudebusch (2002b), given their disparate time frames, a central bank could conduct short-run partial adjustment without quarterly inertia or vice versa. For example, a central bank could spread a desired change over several quarters but make only one rate adjustment per quarter. Alternatively, it could spread a desired change over a month or two but essentially hit its desired rate on a quarterly average basis. It is this latter scenario that is consistent with the evidence below.

Figure 1. U.S. Economic Data

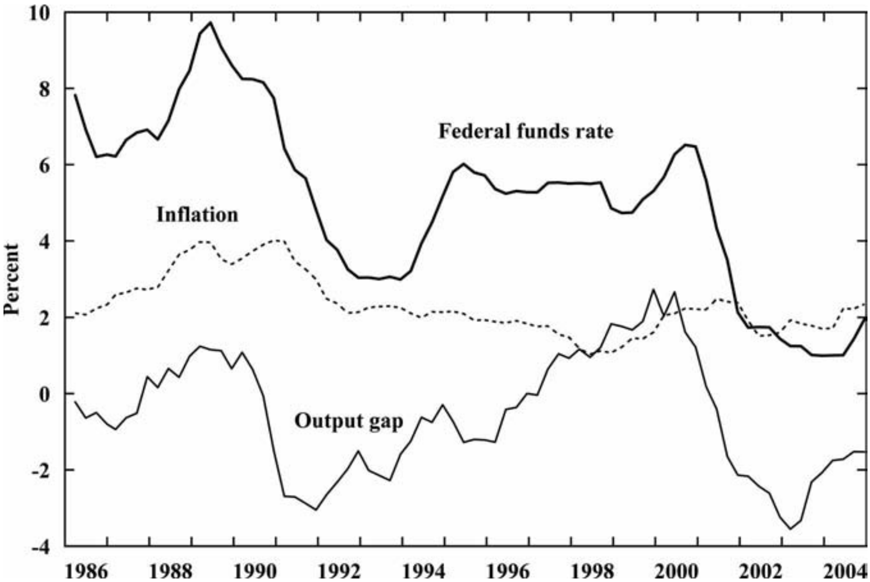

central bank policy rates—as illustrated in figure 1 for the United States.[5] The dispute is not about whether such slow adjustment exists but about its source. One school of thought, the partial-adjustment view, asserts that the persistence of policy rates reflects an inertia that is *intrinsic* or *endogenous* to the central bank. Under this view, there is a long, intentionally drawn-out adjustment of the policy rate in response to economic news. Such partial adjustment implies that the central bank knowingly distributes desired changes in the policy interest rate over an extended period of time; therefore, the smooth persistent policy rates reflect deliberate "interest rate smoothing" or "partial adjustment" or "gradualism" or "inertia" on the part of the central bank. For example, given typical empirical estimates, if a central bank knew it wanted to increase the policy rate by a percentage point, it would only raise it by about 20 basis

[5]Figure 1 shows the quarterly average federal funds rate as the U.S. monetary policy instrument. Figure 1 also displays two important indicators for policy: the four-quarter percent change in the price index for personal consumption expenditures excluding food and energy, labeled "inflation," and the output gap as estimated by the Congressional Budget Office (CBO).

points in the first three months and by about 60 basis points after one year. That is, there is a very slow convergence of the policy rate to its desired level.

The opposing view to partial adjustment is that the persistence of the policy rate simply reflects the response of the central bank to the slow cyclical fluctuations in the key macroeconomic driving variables of monetary policy, such as inflation and output, which are also illustrated in figure 1 for the United States. In this case, the persistence of the policy rate reflects an inertia that is *extrinsic* or *exogenous* to the central bank. Therefore, from this second perspective, the slow adjustment of the policy rate simply reflects the slow accretion of information relevant to the setting of the policy interest rate by policymakers, who then completely adjust the policy rate fairly promptly—typically within a few months—when confronted with new information.[6]

This disagreement is not just an academic debate about macroeconomic behavior but is highly relevant to the practical conduct of monetary policy. For example, as then–Federal Reserve Governor Larry Meyer noted at the February 1999 Federal Open Market Committee meeting (according to the now-public transcript): "I pay a lot of attention to the policy prescriptions from the Taylor rule. Sometimes the different rules that are in the standard packet yield quite different implications for policy." Dynamic adjustment was a key feature that differentiated the various rules supplied to Larry Meyer and other Federal Reserve governors.[7] Some of the rules assumed significant policy partial adjustment while others did not, and this difference led to alternative policy prescriptions. In particular, the crucial difference between, say, Taylor rules with and without endogenous inertia is evident in the policymaker's reaction to news about inflation and output. For example, when faced with a surprising economic recession or a jump in inflation, the inertial policymaker slowly changes the policy rate, while the non-inertial

[6]This same debate also occurs for other macroeconomic time series. For example, as many have noted, the infrequent adjustment of prices could reflect the inertial nature of price determination—menu costs, etc.—or it could indicate the sluggish economic determinants of completely flexible prices.

[7]This is evident in the now-public "standard packet," namely, the "Financial Indicators" FOMC material dated January 29, 1999.

policymaker responds to the news with immediate and sizable interest rate adjustments. (See the discussion in section 4 below.)

Policymakers themselves appear unclear about the source of the slow adjustment of policy interest rates. For example, Ben Bernanke (2004) was a proponent of the intrinsic view in which the slow adjustment of the policy rate reflects "partial adjustment and monetary policy inertia." In contrast, William Poole (2003) argued that there was no partial adjustment: "In my view of the world, future policy actions are almost entirely contingent on the arrival of new information. . . . Given information at the time of a meeting, I believe that the standing assumption should be that the policy action at the meeting is expected to position the stance of policy appropriately." A closely related policy debate, described in Rudebusch and Williams (2006), centers on how much information central banks can and should reveal about their future intentions for policy rate changes. Of course, a central bank that follows a partial-adjustment procedure typically will be able to communicate insights about likely future changes in the policy rate—namely, insights about the remaining policy partial adjustment. However, many policymakers vehemently deny that they are in a position to provide guidance about the future path of policy interest rates. As the Governor of the Bank of England (King 2006) recently noted: "The [Bank of England's monetary policy committee] reaches a new judgment each month, made afresh in the light of all the new information about the prospects for inflation. We don't decide in advance. So trying to give direct hints on the path of interest rates over the next few months risks deceiving financial markets into believing there are definite plans for the next few months when no such plans exist."[8]

Still, to be clear, the absence of central bank partial adjustment does not mean that central banks are not trying to influence long-term interest rates. Again, both sides of this debate agree that a change in the central bank policy rate is likely to persist, and both sides agree that such a change in the policy rate is likely to affect expectations of future short-term rates and hence long-term rates

[8]Goodhart (2005), a former member of the Bank of England's monetary policy committee, also relates how a central bank with no intrinsic inertia can still display an ex post track record with long sequences of small interest rate adjustments in the same direction.

as well. In order to influence the long rate, a central bank only must present a path for the policy rate that can shape expected future short rates. The partial-adjustment rule provides one such path, but it is not the only one. As noted by Goodfriend (1991) and Rudebusch (1995), an ex ante constant path, which is what some non-inertial rules approximate, is another obvious choice.

In the next section, in light of the clear theoretical and practical importance of the topic, I review the basic evidence for and against monetary policy inertia. The inertial view appears widely supported by estimated monetary policy rules. When such rules are estimated without policy inertia, the residuals indicate significant, persistent deviations of the rule recommendation from the historical policy rate. With the addition of partial adjustment (in the form of a lagged dependent variable), these deviations are greatly reduced. The alternative view, as noted above, is that the deviations represent persistent influences on central bank behavior that are not captured in a simple Taylor-type rule. These persistent influences may include, for example, responses to financial crises, judgmental adjustments, or differences between real-time and final revised data. Unfortunately, as is well known in econometrics, at least since Griliches (1967), the two dynamic representations of partial adjustment and persistent omitted variables can be very hard to distinguish in simple single-equation regressions. Indeed, this appears to be the case for the monetary policy rule regressions, especially since there is so much uncertainty about the exact arguments of the policy rules.

Therefore, section 3 turns to theory and examines whether a central bank would want to engage in sluggish partial adjustment from the perspective of optimal monetary policy prescriptions. There are three key rationales for inertial behavior, namely, to reduce interest rate volatility, to exploit the expectational channel for monetary policy, and to respond optimally to data and model uncertainty. While there appears to be some validity to each of these rationales, they do not appear to be able to justify the extremely slow monetary policy inertia suggested by the estimated monetary policy rules.

In contrast to the weak and inconclusive single-equation evidence and theoretical rationales in sections 2 and 3, a very powerful set of evidence on monetary policy inertia is introduced in section 4. This evidence is contained in the term structure of interest rates, which can bring a vast amount of information to bear on the appropriate monetary policy rule. Assuming financial market participants understand the policy rule that links short-term interest rates to the realizations of macroeconomic variables, they then will also use that rule in pricing forward interest rates. Accordingly, any deviations between expected future short-term rates and expected rule recommendations based on future macroeconomic conditions will be arbitraged away. Therefore, at any point in time, multiperiod interest rates, which embody expectations of future short rates, will contain much information about the properties of the monetary policy reaction function. Section 4 presents three different ways to use such yield-curve information—predictability regressions, macro-finance system estimates, and event studies based on macroeconomic data surprises. These procedures differ in the amount of economic structure imposed and also operate at three different frequencies—quarterly, monthly, and intraday. The resulting consistent set of results from these diverse methodologies appears to provide decisive evidence against the presence of significant monetary policy inertia.

Finally, section 5 concludes with some suggestions for future research.

2. Gradualism and Inertia in Policy Rules

An inertial view of monetary policy dynamic adjustment implies that the short-term policy rate is changed at a very sluggish pace, so a monetary policy reaction to new economic data is distributed over many quarters. I first clarify policy inertia as a general proposition (or, depending on your perspective, highlight some of the ambiguity involved with any such definition) and then survey some of the relevant empirical work.

2.1 Defining Policy Gradualism and Inertia

It is perhaps useful to discuss in general terms what is meant by the "inertial" and "non-inertial" hypotheses regarding the conduct of policy. In the literature, "inertial" rules follow the standard partial-adjustment form: $i_t = (1-\rho)\hat{\imath}_t + \rho i_{t-1}$, where i_t is the level of the policy interest rate set in quarter t, which is a weighted average of the current desired level, $\hat{\imath}_t$, and last quarter's actual value, i_{t-1}.

Based on historical data, estimates of ρ are often in the range of 0.8, so these empirical rules appear to imply a very slow speed of adjustment—about 20 percent per quarter—of the policy rate to its fundamental determinants. The large coefficient on the lagged dependent variable is widely interpreted as evidence for a "monetary policy inertia" behavior by central banks.[9]

In fact, at a general level, it does not seem that any logical distinction can be drawn between inertial and non-inertial rules as descriptions of policy. For example, by defining an "underlying" desired interest rate level as $\tilde{\imath}_t = \rho\tilde{\imath}_{t-1} + (1-\rho)\hat{\imath}_t$, the above inertial interest rate rule can be rewritten in an ostensibly non-inertial form as $i_t = \tilde{\imath}_t$.[10] That is, an inertial versus non-inertial designation makes sense only in conjunction with specific assumptions about the arguments of the rule. Of course, there is a natural set of arguments to consider, namely, the standard major macroeconomic data series—especially inflation and output, which are the arguments of the popular Taylor rule. Indeed, the hypothesis examined in this paper is not partial adjustment in all its generality, but partial adjustment toward a target that depends in a straightforward way on inflation and output (as exemplified by the Taylor rule). As described below, this is the case of overwhelming interest in the literature.

Therefore, to make progress, I will limit consideration to rules in which the desired rate is a simple function of a set of standard macroeconomic variables, formally, $\hat{\imath}_t = \beta' X_t$, where X_t is a vector of the variables influencing policy. The *inertial* rule can then be written as

$$i_t = (1-\rho)\beta' X_t + \rho i_{t-1}. \tag{1}$$

The corresponding *non-inertial* rule is

$$i_t = \beta' X_t. \tag{2}$$

[9]For example, Clarida, Galí, and Gertler (2000, 157–58) describe their U.S. estimates of various partial-adjustment policy rules as follows: "...the estimate of the smoothing parameter ρ is high in all cases, suggesting considerable interest rate inertia: only between 10 and 30 percent of a change in the [desired interest rate] is reflected in the Funds rate within the quarter of the change."

[10]This is just the observational equivalence of the information-smoothing and partial-adjustment models noted by Waud (1968).

The common finding in the empirical literature discussed below is that the inertial rule fits the data better (say, in an R^2 sense) than the non-inertial rule. An alternative view, however, is that the non-inertial rule is not misspecified in terms of dynamics but in terms of arguments, so there is an alternative non-inertial rule that could be formalized as

$$i_t = \beta' X_t + \phi' Z_t, \tag{3}$$

where Z_t is a vector of persistent omitted factors that also influence policy. The rest of this paper discusses the evidence for these varying specifications.

2.2 Gradualism and Inertia in Estimated Policy Rules

The belief in sluggish policy adjustment in the real world is based on estimated policy rules. The most commonly estimated inertial policy rules have been dynamic forms of the Taylor rule. In such rules, the actual interest rate partially adjusts to a desired interest rate that depends on inflation and the output gap; specifically,

$$i_t = (1-\rho)\hat{\imath}_t + \rho i_{t-1} + \xi_t \tag{4}$$

$$\hat{\imath}_t = k + g_\pi \bar{\pi}_t + g_y y_t, \tag{5}$$

where k is a constant incorporating an equilibrium real rate, r^*, and an inflation target, π^*, and g_π and g_y are the central bank response coefficients to (four-quarter) inflation ($\bar{\pi}_t$) and the output gap (y_t).[11]

To provide a benchmark for comparison, first consider an estimated *non-inertial* Taylor rule that assumes $\rho = 0$, as in Taylor (1999) and Yellen (2004). A least-squares regression on U.S. data

[11] The federal funds rate is a quarterly average rate. Inflation is defined using the price index for personal consumption expenditures excluding food and energy (denoted P_t, so $\pi_t = 400(\ln P_t - \ln P_{t-1})$ and $\bar{\pi}_t = \frac{1}{4}\Sigma_{j=0}^{3}\pi_{t-j}$), and the output gap is defined as the percent difference between actual real GDP (Q_t) and potential output (Q_t^*) estimated by the Congressional Budget Office (i.e., $y_t = 100(Q_t - Q_t^*)/Q_t^*$).

Figure 2. Actual and Fitted Federal Funds Rate

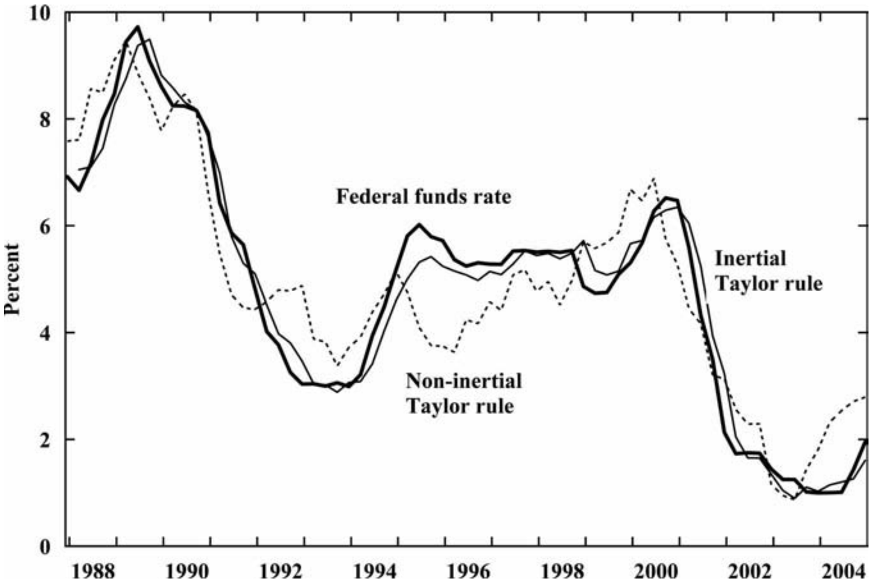

from 1987:Q4 to 2004:Q4 yields

$$\begin{aligned} i_t = \underset{(.28)}{2.04} + \underset{(.09)}{1.39}\,\bar{\pi}_t + \underset{(.06)}{.92}\,y_t + \xi_t^{NI} \equiv \hat{i}_t^{NI} + \xi_t^{NI}, \\ \sigma_\xi = .97, \quad \bar{R}^2 = .82, \quad DW = .34. \end{aligned} \tag{6}$$

The monetary policy response coefficients—namely, $g_\pi = 1.39$ for inflation response and $g_y = 0.92$ for output response—are not too far from the 1.5 and 0.5 that Taylor (1993) originally used. The fitted values from this non-inertial Taylor-rule regression, which will be denoted $\hat{i}_t^{NI}$, are shown as the dotted line in figure 2 and show a fairly good fit to the actual funds rate—the thick solid line. However, there are some large persistent deviations between the non-inertial rule and the historical funds rate, especially during 1992, 1993, 1999, and 2004 (when the actual rate was held below the rule) and during 1991, 1995, and 1996 (when the rate was pushed above the rule). As discussed below, the source of these deviations will be a critical element in interpreting the evidence and arguments for and against policy inertia.

A partial-adjustment mechanism is a standard econometric response to such persistent deviations, and a least-squares regression for an inertial policy rule on U.S. data from 1987:Q4 to 2004:Q4 yields

$$i_t = .22\,\hat{\imath}_t^I + \underset{(.04)}{.78}\,i_{t-1} + \xi_t^I, \tag{7}$$

$$\hat{\imath}_t^I = \underset{(.18)}{2.13} + \underset{(.18)}{1.33}\,\bar{\pi}_t + \underset{(.13)}{1.29}\,y_t \tag{8}$$

$$\sigma_\xi = .38, \quad \bar{R}^2 = .97.$$

In this regression, the estimated values of the response coefficients are not so different from the non-inertial rule; however, the estimate of the partial-adjustment coefficient ($\hat{\rho} = 0.78$) is economically and statistically significant. Such lagged dependence is an extremely robust empirical result in the literature.[12] Indeed, after taking into account the dynamic adjustment in equation (7), the fitted values in the inertial rule—which are shown as the thin solid line in figure 2—match the historical path of the funds rate much more closely than the non-inertial rule. This difference in fit is also apparent in figure 3, which charts the residuals (ξ_t^I and ξ_t^{NI}) from the inertial and non-inertial rules. The mean absolute residual for the non-inertial rule is .82 percentage point, which is almost three times larger than the .29-percentage-point mean absolute residual for the inertial rule.

The significance of ρ and the dramatic improvement in R^2 have been widely taken to be convincing evidence of monetary policy inertia. However, Rudebusch (2002b) argues that the monetary policy rule estimates are misleading and provide the illusion of monetary policy inertia. In particular, if the desired policy interest rate depends on persistent factors other than the current output and inflation in the Taylor rule, then such a misspecification could result in a spurious finding of partial adjustment. Accordingly, based only on these types of policy rule estimates, it would be very difficult

[12]Similar estimates are discussed by Kozicki (1999) and Rudebusch (2002b) for the United States and by Sauer and Sturm (2003), Gerdesmeier and Roffia (2004), and Castelnuovo (2006) for the euro area.

Figure 3. Residuals from Estimated Inertial and Non-Inertial Taylor Rules

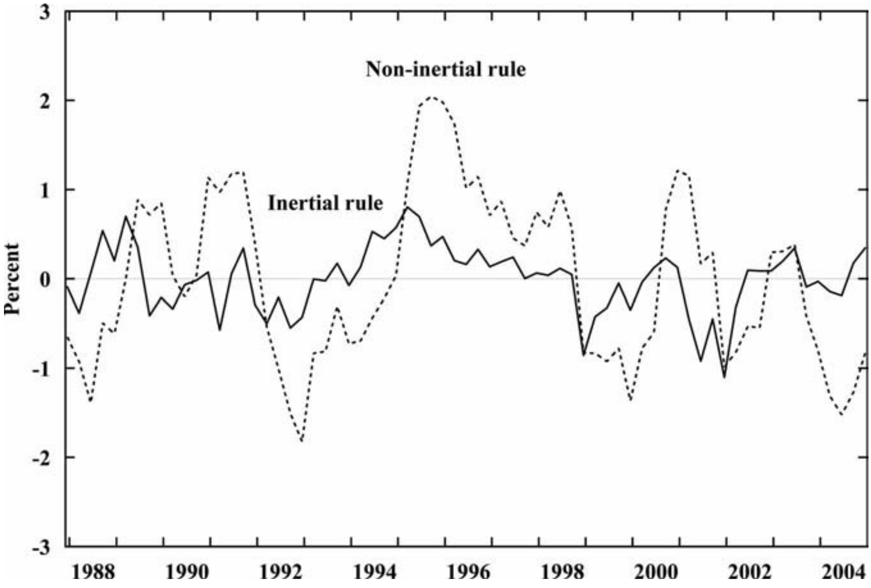

to distinguish whether the Federal Reserve's adjustment was sluggish or whether the Federal Reserve generally followed the Taylor rule with no policy inertia but sometimes deviated from the rule for several quarters at a time in response to other factors.

The intuition for this argument is illustrated in figures 4 and 5. Figure 4 displays the actual funds rate (thick solid line) and the "desired" funds rates from the two rules. The non-inertial rule desired rates are the fitted values $\hat{\imath}_t^{NI}$ from equation (6) (shown as the dotted line), and the inertial rule desired values are the $\hat{\imath}_t^{I}$ from equation (8) (shown as the thin solid line). The two desired levels generally move together, so deviations of these desired rates from the actual funds rate are similar across the two rules. Understanding the persistent deviations of the historical interest rate from the two Taylor-rule recommendations is key to interpreting the empirical evidence. Under the monetary policy inertia interpretation, these persistent deviations are the result of sluggish central bank responses to output and inflation gaps; that is, the central bank only gradually adjusts the policy rate to the level it would like to set in the

Figure 4. Actual and Desired Federal Funds Rate

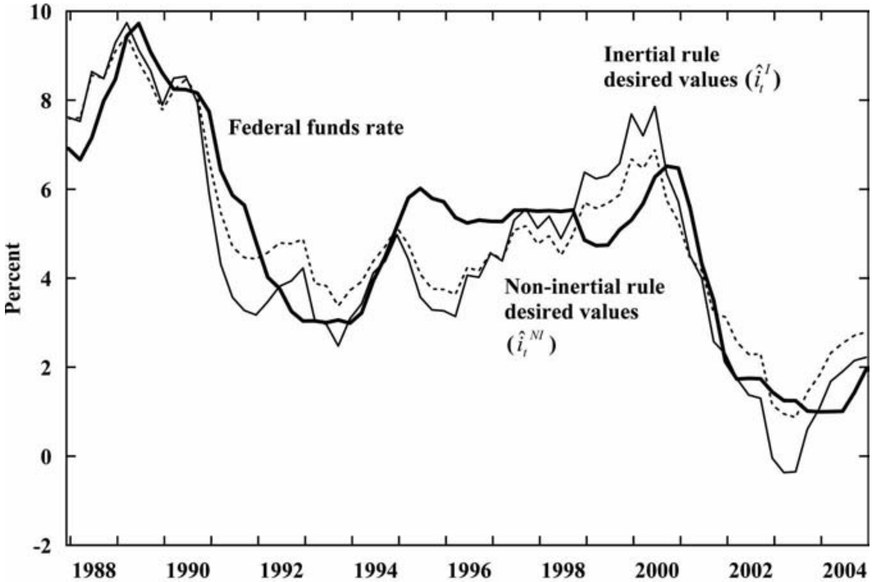

absence of some partial-adjustment constraint. However, there are several episodes evident in figure 4 that appear to contradict such an interpretation. For example, at the beginning of 1995, the actual funds rate matched the desired funds rate (as recommended by either rule), but over the rest of that year, the desired funds rate dropped almost 200 basis points, while the actual funds rate jumped 100 basis points. Conversely, after the third quarter of 1998, when the actual rate equaled the desired values, desired rates rose sharply for the next year, while the actual funds rate dropped. Adding a lagged funds rate to the equation will certainly improve the regression fit, but it appears misleading to characterize these episodes as central bank partial adjustment when the actual and desired funds rates moved so dramatically in opposite directions.

The deviations of the two desired funds rate series from the actual funds rate are shown in figure 5 (namely, $i_t - \hat{i}_t^{NI}$ as the solid line and $i_t - \hat{i}_t^I$ as the dotted line). Instead of a partial-adjustment explanation for these deviations, an alternative explanation is that the deviations in figure 5 reflect the incomplete description of monetary policy

Figure 5. Deviations of Actual Funds Rate from Desired Rule Value

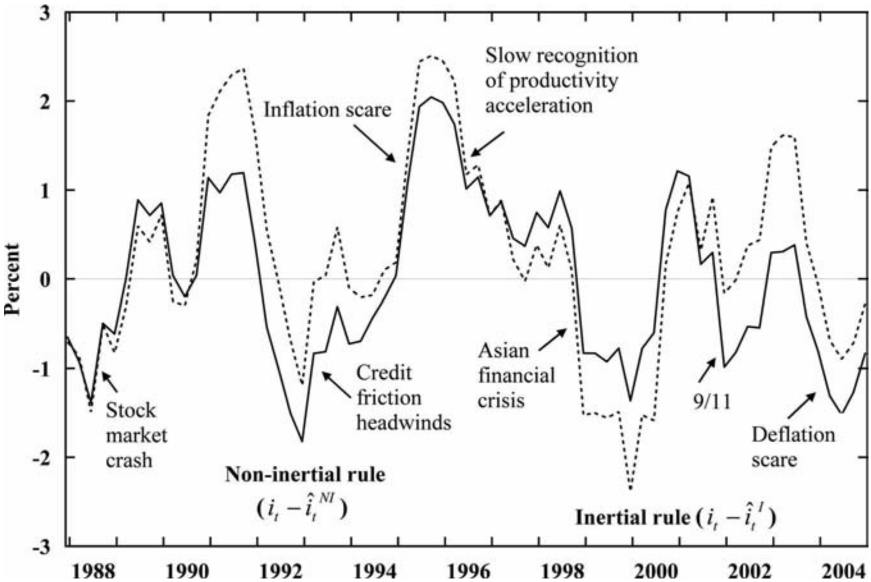

provided by the Taylor rule. Indeed, it is fairly straightforward to provide a basic narrative history of a variety of macroeconomic developments that the Federal Reserve appeared to respond to in addition to estimates of the contemporaneous output gap and inflation. Some of these developments are indicated in figure 5.[13] For example, relative to what the Taylor rule would have recommended, a response to the stock market crash may have lowered rates in 1988, and inflation worries—at least, as discussed below, when judged using real-time data—appear to have led to a greater-than-Taylor-rule tightening during 1989. The deviations toward looser monetary policy in 1992 and 1993 have been interpreted as the Federal Reserve's response to disruptions in the flow of credit or severe financial headwinds.[14]

[13]The original analysis of Taylor (1993) put forward a description of monetary policy that did not involve interest rate smoothing or partial adjustment. Taylor argued that deviations from the rule during various episodes were an appropriate response to special circumstances. Kozicki (1999) also makes this point.

[14]As then–Chairman of the Board of Governors Alan Greenspan testified to Congress on June 22, 1994: "Households and businesses became much more reluctant to borrow and spend and lenders to extend credit—a phenomenon often

An inflation scare at the end of 1994—evidenced by a rapid rise in long-term interest rates—preceded a sustained period of tight policy. Another factor that emerged during this period was the remarkable increase in the growth rate of productivity and potential output. At the time, most economists didn't recognize these changes and hence overestimated the degree of utilization in labor and product markets, which likely was reflected in tighter policy. In 1998 and 1999, a worldwide financial crisis following the Russian default and devaluation appears to have played a role in lowering rates.[15] Similarly, there was a rapid easing in response to events of September 11, 2001. Finally, 2003 and 2004 were dominated by fears of deflation, which would likely be reflected in lower rates than a simple Taylor rule would recommend, given potential concerns at the zero lower bound for the policy rate (as discussed in McGough, Rudebusch, and Williams 2005).

This narrative suggests that some Taylor-rule residuals reflect differences between policy judgments made with real-time data and Taylor-rule estimations conducted with final revised data—a topic that deserves special attention (see Rudebusch 1998, 2001, 2002a, 2002b, and Orphanides 2001, 2003). Figure 6 provides some evidence on the importance of these effects in the United States by showing the difference between real-time and current estimates of the output gap, which is denoted $y_{t|t} - y_{t|T}$, and the difference between real-time and current estimates of inflation, which is denoted $\bar{\pi}_{t|t} - \bar{\pi}_{t|T}$.[16]

referred to as the 'credit crunch.' In an endeavor to defuse these financial strains, we moved short-term rates lower in a long series of steps that ended in the late summer of 1992, and we held them at unusually low levels through the end of 1993—both absolutely and, importantly, relative to inflation."

[15] Federal Reserve Governor Larry Meyer (1999, 7) had this explanation for the easing of policy during late 1998: "There are three developments, each of which, I believe, contributed to this decline in the funds rate relative to Taylor Rule prescription. The first event was the dramatic financial market turbulence, following the Russian default and devaluation. The decline in the federal funds rate was, in my view, appropriate to offset the sharp deterioration in financial market conditions, including wider private risk spreads, evidence of tighter underwriting and loan terms at banks, and sharply reduced liquidity in financial markets."

[16] The output-gap series is Federal Reserve Board staff's real-time estimate—kindly supplied by David Small from the FOMC Secretariat—minus the current (as of 2005) CBO output-gap estimate. The inflation series is the real-time four-quarter GDP deflator inflation rate—obtained from the Federal Reserve Bank of Philadelphia real-time data website—minus the current release.

Figure 6. Differences between Real-Time and 2005 Data Vintages

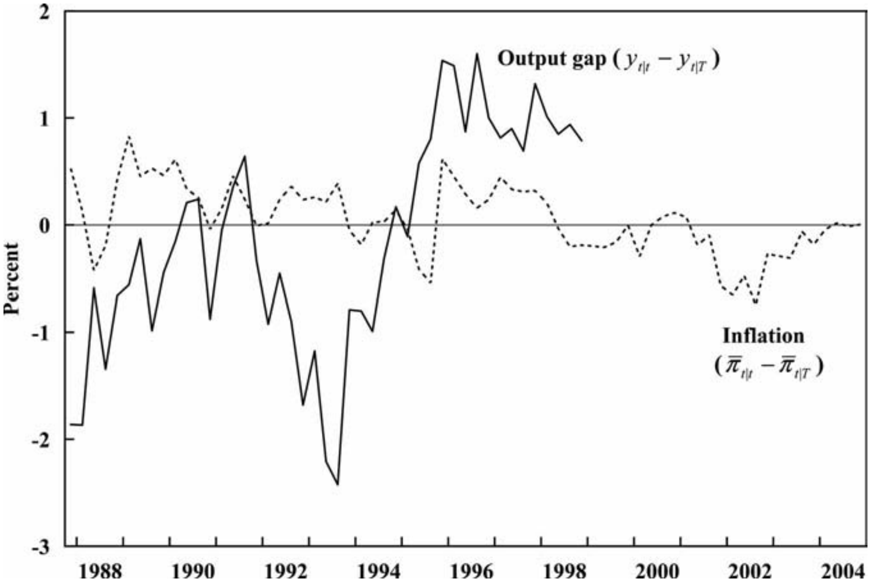

(The output-gap revisions end in 1998 because of data confidentiality.) For example, figure 6 shows that in real time, the output gap from 1996 through 1998 was estimated to be about a percentage point higher than the current estimate (because the estimated level of potential output was lower in real time). This underestimation of the degree of macroeconomic slack would be reflected in higher interest rates in real time than a Taylor rule estimated with current data would recommend. Similarly, during 1989, inflation was thought to be running about half of a percentage point faster than current estimates would indicate, so the actual policy rate would likely be higher than the final-data rule would recommend.

It is possible to provide a rough indication of the importance of the data revisions in accounting for the Taylor-rule residuals. The predicted Taylor-rule residuals based on the real-time to final-data revisions can be constructed under the assumption that the Federal Reserve followed the estimated non-inertial rule (6) in real time. Specifically, if the Federal Reserve used the estimated non-inertial

Figure 7. Matching Non-Inertial Rule Residuals with Real-Time Data Revisions

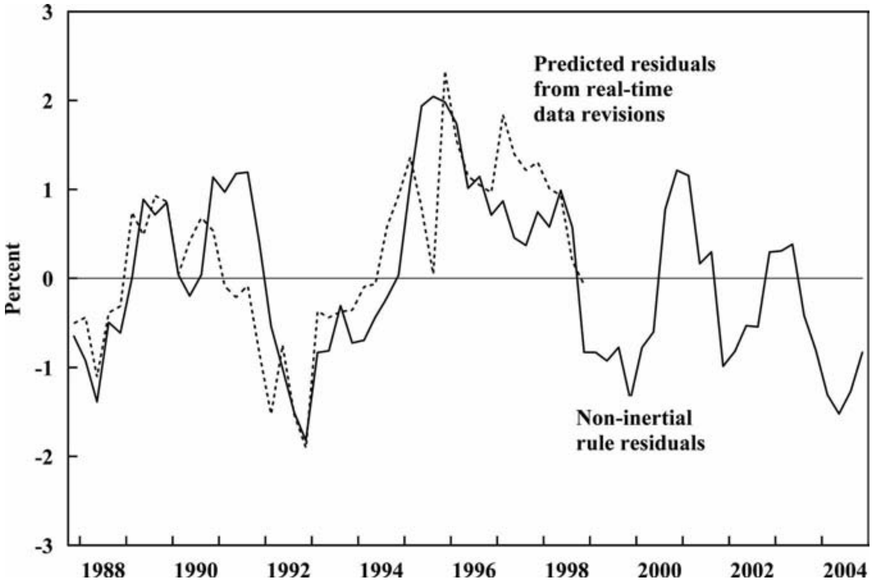

Taylor-rule coefficients to conduct policy, so $i_t = 2.04 + 1.39\bar{\pi}_{t|t} + .92y_{t|t}$, then the predicted residuals in equation (6) would equal $1.39(\bar{\pi}_{t|t} - \bar{\pi}_{t|T}) + .92(y_{t|t} - y_{t|T})$. These constructed residuals predicted by the data revisions are shown in figure 7, along with the non-inertial rule residuals from equation (6). The fairly close correlation between the predicted residuals from real-time data revisions and the actual residuals from the estimated non-inertial rule suggests that a substantial amount of the deviations of the actual rate from the rule estimated with the current vintage of data can be accounted for by the reactions to real-time data and not to central bank partial adjustment.[17]

[17]Lansing (2002) provides a careful simulation study that demonstrates the potential effectiveness of such real-time output-gap errors to account for spurious evidence of policy inertia. Also, see Mehra (2002) and Apel and Jansson (2005) for the United States and Sauer and Sturm (2003) for the euro area. In addition, given the large policy rule inflation response coefficient, inflation data revisions should not be ignored.

However, while real-time data revisions are undoubtedly part of the story, it is unlikely, as suggested in figure 5, that the Federal Reserve follows a Taylor rule in real time. Instead, like other central banks, it reacts in a less-simplistic fashion to a wide variety of macroeconomic developments; that is, the alternative to partial adjustment is the misspecification of the Taylor rule. This omitted-variables view of the non-inertial Taylor-rule residuals is supported by much contemporaneous press coverage and the narrative policy record. Still, it would be more satisfying to be able to provide econometric evidence distinguishing between the partial-adjustment and omitted-variables interpretations of the policy rule estimates.

Unfortunately, Rudebusch (2002b) argues that conclusive evidence from simple policy rule estimates on the extent of inertia is inherently difficult to obtain. For example, suppose that the non-inertial rule deviations, which presumably represent various persistent factors—credit crunches, financial crises, etc.—that a central bank might respond to, could be modeled as a simple first-order autoregressive process. Then, instead of the inertial model of central bank behavior in equations (4) and (5), the representation of policy would be the serially correlated shock model:

$$i_t = k + g_\pi \bar{\pi}_t + g_y y_t + \xi_t \tag{9}$$

$$\xi_t = \rho^e \xi_{t-1} + \omega_t. \tag{10}$$

The salient question is whether it is possible to distinguish between model (4) and (5) and model (9) and (10). Rudebusch (2002b) estimates a single equation that nests the inertial and serially correlated shocks rules and finds that the evidence distinguishing these two rules appears fragile to even modest changes in the sample period. His argument draws on a large literature in econometrics showing that estimates of partial-adjustment models commonly indicate an unrealistically slow adjustment—whether applied to inventory behavior (Blinder 1986) or money demand (Goodfriend 1985).[18] In particular, a standard policy rule with slow partial adjustment

[18]There is a large literature that argues that partial-adjustment models are difficult to identify and estimate empirically in the presence of serially correlated shocks (e.g., Griliches 1967; Hall and Rosanna 1991; and McManus, Nankervis, and Savin 1994).

and no serial correlation in the errors will be difficult to distinguish empirically from a policy rule that has immediate policy adjustment but highly serially correlated shocks. The choice between these two dynamic structures, which depends crucially on separating the influences of contemporaneous and lagged regressors, is especially difficult to untangle for empirical monetary policy rules for several reasons (also see Carare and Tchaidze 2005). First, the arguments of the rules—four-quarter inflation and the output gap—are highly serially correlated, so distinguishing the effect of, say, $\bar{\pi}_t$ from $\bar{\pi}_{t-1}$ is not easy. Second, the arguments of the rules are not exogenous. Third, only short data samples of plausibly consistent rule behavior are available with a limited amount of business-cycle variation in output and inflation. Fourth, there is some uncertainty about the appropriate arguments of the historical policy rule. Finally, as noted above, the actual interest rates are set on the basis of real-time data on output and inflation, which can also make it difficult to determine the correct dynamics.

There have been several interesting extensions of the analysis in Rudebusch (2002b). English, Nelson, and Sack (2003) and Gerlach-Kristen (2004b) provide two slightly different tests of the inertial and serially correlated shock interpretations that, unlike in Rudebusch (2002b), allow for both partial adjustment and serially correlated shocks to be jointly present. These authors find evidence that both features are significant elements in the data; therefore, the standard policy rule estimates in (7) and (8) are omitting important persistent factors (similar results for the euro area are provided by Castelnuovo 2006). However, considerable uncertainty remains, as illustrated by the insightful small-sample calculations conducted by English, Nelson, and Sack (2003). They investigate how much of the deviations of the actual rate from the desired rate (the $i_t - \hat{\imath}_t^I$ in figure 5) can be accounted for by partial adjustment. They find that a 95 percent confidence interval stretches from 8 percent to 88 percent; therefore, in the context of a single-equation regression, it is difficult to ascertain how economically important partial adjustment is for the policy rule. In addition, this wide range of uncertainty is only for a particular rule specification and estimation sample, so it ignores the broader uncertainty noted above.

Furthermore, the assumption that the persistent omitted variable is an AR(1) appears to be a gross simplification that may bias the results and boost the evidence for partial adjustment. Indeed, the narrative history summarized in figure 5 suggests a more-subtle reaction function than can be captured by equations (9) and (10). A few have tried to augment the estimated Taylor rule with other variables in order to capture directly the omitted persistent influences on policy that spuriously induce the appearance of policy inertia. For example, Gerlach-Kristen (2004b) and Driffill et al. (2006) find evidence that proxies for financial stability concerns, such as a private-public credit spread, have explanatory power in the Taylor rule. Also, expectations appear to play an important role in tempering the policy response to current readings on output and inflation, and Mehra (2002) suggests that expectations of future inflation—and, in particular, inflation scares in the bond market—are an important consideration for policy, which—when omitted—will appear as policy inertia. Finally, as shown by Trehan and Wu (2006), ignoring a true time-varying equilibrium real rate (r_t^*) can lead to finding policy inertia when there is really none.[19] Overall, however, the literature on augmenting the Taylor rule with the important determinants of policy other than current output and inflation appears to be incomplete at best.

3. Rationales for Sluggish Adjustment by Central Banks

The discussion above indicates that, given the distinct possibility of omitted persistent variables from the monetary policy rules, the usual single-equation evidence from estimated policy reaction functions is inconclusive regarding the empirical importance of policy inertia. In this section, I take a different tack and examine the normative case for interest rate smoothing. Presumably, if theory can provide a fairly compelling rationale for the existence of inertia as a feature of optimal monetary policy, then the case for real-world partial adjustment would be strengthened. Therefore, in this section, I consider the empirical relevance of the three most important

[19] In Europe, Gerlach and Schnabel (2000) find that a Taylor rule fits well without partial adjustment but with dummies for the period 1992:Q3–1993:Q3 to control for intra-European exchange market pressures. Gerlach-Kristen (2004a) finds that the long rate is significant in a euro-area Taylor rule, while Gerdesmeier and Roffia (2004) recommend inclusion of a money growth gap.

explanations for why central banks might find partial adjustment attractive.

3.1 Gradualism and Volatility Reduction

One consequence of policy inertia is to produce interest rates that are less volatile than would be suggested by the determinants of policy. As the speed of adjustment coefficient ρ increases, the variances of the level and changes in the policy instrument decline. Therefore, an obvious rationale for policy gradualism would be some desire on the part of the central bank to reduce the volatility in interest rates and, more generally, in asset prices. Such a desire can be modeled directly in the central bank's loss function, and then, together with a model of the economy, the optimal ρ coefficient can be calculated for an optimal simple Taylor rule (as in, for example, Rudebusch and Svensson 1999). If the optimal monetary policy partial-adjustment coefficient matched the high empirical estimates of ρ, then those estimates would have some greater credence.

The most common way to model a desire for smooth interest rates is to specify a loss function in which the central bank minimizes a weighted sum of the squared inflation gap, the squared output gap, and changes in the policy rate (see Clarida, Galí, and Gertler 1999 and Rudebusch and Svensson 1999):

$$L_t = 1/2\left[(\bar{\pi}_t - \pi^*)^2 + \lambda y_t^2 + \nu_{\Delta i}(\Delta i_t)^2\right], \tag{11}$$

where $\Delta i_t = i_t - i_{t-1.}$ The parameters $\lambda \geq 0$ and $\nu_{\Delta i} \geq 0$ are the relative weights on output and interest rate stabilization with respect to inflation stabilization. The intertemporal loss function in quarter t is the discounted sum of the expected future per-quarter losses,

$$E_t \sum_{\tau=0}^{\infty} \delta^\tau L_{t+\tau}, \tag{12}$$

with a discount factor δ ($0 < \delta < 1$). For $\delta = 1$, this loss function can be represented by the unconditional mean of the period loss function (Rudebusch and Svensson 1999)

$$E[L_t] = \mathrm{Var}[\bar{\pi}_t - \pi^*] + \lambda \mathrm{Var}[y_t] + \nu_{\Delta i}\mathrm{Var}[\Delta i_t], \tag{13}$$

which equals the weighted sum of the unconditional variances of the three goal variables and is the standard loss function in the literature.[20]

The presence of an interest rate smoothing motive in the loss function has some superficial plausibility, especially in light of the literature that analyzes changes in policy interest rates on a day-by-day basis. In the United States (e.g., Goodfriend 1991 and Rudebusch 1995) and many other countries (e.g., Goodhart 1997 and Lowe and Ellis 1997), central banks generally make changes in the policy rate at discrete intervals and in discrete amounts. Rudebusch (1995, 264), for example, describes a short-term interest rate smoothing in which the Federal Reserve adjusts interest rates "...in limited amounts ... over the course of several weeks with gradual increases or decreases (but not both)...." This smoothing likely reflects various institutional rigidities, such as a fixed monthly meeting schedule and perhaps certain sociological and political influences.[21] However, as noted in the introduction, short-term partial adjustment within a quarter is essentially independent of whether there is monetary policy inertia over the course of several quarters, and it is this latter issue that is relevant for empirical monetary policy rules. Indeed, if the underlying rationale for reducing interest rate volatility is to reduce instability in financial markets (as described by, for example, Goodfriend 1991, Rudebusch 1995, Cukierman 1996, and Lowe and Ellis 1997), then not wanting to move the policy rate by 50 basis points on a particular day is very different from not wanting to move it by 50 basis points on a quarterly average basis.

This issue is highlighted in trying to specify the weight $\nu_{\Delta i}$ on quarterly interest rate volatility relative to the variability of the output and inflation gaps. If λ and $\nu_{\Delta i}$ are both set equal to 1,

[20] However, the choice of δ is not innocuous. As shown in Dennis (2006), greater discounting may lead to less concern about the future and less interest rate smoothing.

[21] At a single meeting, large interest rate changes may be difficult to achieve politically because of the decision-making process (e.g., Goodhart 1997) or because such changes may be taken as an adverse signal of inconsistency and incompetence (e.g., Goodhart 1999). Indeed, many have noted an "aversion to reversals" in which raising (or lowering) the policy rate at one meeting precludes a lowering (raising) at the next. Again, it appears unlikely that such a meeting-by-meeting aversion would lead to quarterly inertia.

then the loss function equally penalizes a 1 percent output gap, a 1-percentage-point inflation gap, and a 1-percentage-point quarterly change in the funds rate. This penalty on interest rate volatility appears to be implausibly high, given the overwhelming emphasis among central banks on the first two objectives relative to the third (e.g., the "dual mandate" in the United States). Indeed, in practice, central banks have at times implemented large changes in policy rates, which contradicts the notion of a significant penalty. Perhaps the most extreme example occurred in September 1992, when the Swedish central bank raised its policy rate from 20 percent to 500 percent in one week in an attempt to maintain a fixed exchange rate. Also, during the 1979–82 monetary experiment, the United States had much greater interest rate volatility, which did not appear to impose, on its own, large costs. In the academic literature, $\nu_{\Delta i}$ is often set equal to 0.5 or 0.1. These loss functions equally penalize a 1 percent output gap, a 1-percentage-point inflation gap, and a 1.41- or 3.16-percentage-point quarterly change in the funds rate. Such weights still seem at the high end of the plausible range of penalties to reduce volatility, especially in a world with a wide variety of financial market instruments that allow for hedging against interest rate volatility.

Finally, I should note that even the specification of the interest rate smoothing objective in the loss function is unclear. Svensson (2003) notes that if the motive in reducing interest rate volatility is to avoid financial instability, then the loss function should be specified to minimize the *surprise* in the policy rate:

$$E[L_t] = \mathrm{Var}[\bar{\pi}_t - \pi^*] + \lambda \mathrm{Var}[y_t] + \nu_{Ei} \mathrm{Var}[E_{t-1}[i_t] - i_t], \quad (14)$$

where $\nu_{Ei} \geq 0$ is the relative weight on policy rate surprises. A third specification, advocated by Woodford (1999), penalizes the variability in the *level* of the policy rate:

$$E[L_t] = \mathrm{Var}[\bar{\pi}_t - \pi^*] + \lambda \mathrm{Var}[y_t] + \nu_i \mathrm{Var}[i_t - r^* - \pi^*], \quad (15)$$

where $\nu_i \geq 0$ is the weight on deviations of the nominal rate from a neutral level.[22]

[22]Woodford (1999, 2003) argues that smaller interest rate fluctuations reduce the likelihood of reaching the zero bound on nominal interest rates and the

On its own, motivating a large partial-adjustment coefficient through a central bank loss-function desire for interest rate smoothing appears unrealistic (e.g., Svensson 2003). This is particularly true in a model with no explicit forward-looking expectational terms, as in Rudebusch and Svensson (1999), where an optimal ρ in a dynamic Taylor rule of greater than .2 or .3 is difficult to obtain. However, results can be very different in forward-looking models, which are considered in the next subsection.

3.2 Central Bank Inertia as a Lever on Expectations

The most passionate advocates for optimal monetary policy partial adjustment base their case on the ability of such inertia to allow the central bank to influence the current state of the economy by promising future actions; that is, sluggish adjustment can be a lever to help move and manage expectations. In particular, partial adjustment can be optimal if the private sector is forward looking and the monetary policymaker is credibly committed to a gradual policy rule (see Levin, Wieland, and Williams 1999; Rotemberg and Woodford 1999; Woodford 1999, 2003; and Sack and Wieland 2000). In such a situation, the small inertial changes in the policy interest rate that are expected in the future can have a large effect on current supply and demand and can help the central bank control macroeconomic fluctuations.[23]

This argument can be elucidated and assessed within a simple expectational model. Rudebusch (2002b, 2005) describes an

---

associated adverse effects on macroeconomic stability; however, with a properly specified model, such concerns should be captured in the output and inflation stabilization concerns in the loss function. Woodford (2003) also tries to motivate this specification of the loss function by appealing to the transactions frictions underlying money demand (so-called shoe-leather costs).

[23]This argument can be thought of as a special case of the more general rationale that i_{t-1} is likely an important state variable given the dynamic structure of the economy, so the optimal instrument rule would include a response to it (e.g., Rudebusch and Svensson 1999). However, it should be noted that Woodford considers fully optimal policy, not an optimal simple rule of the form (1) and (2). Persistence of optimal policy under commitment arises because of the response of policy to previous promises through the lagged Lagrange multipliers (Dennis 2005). Some might interpret these lagged Lagrange multipliers as the unobserved persistent factors omitted from the simple policy rules.

empirical version of the New Keynesian model[24] suitable for quarterly data, where inflation and output are determined by future expectations and lags on the past:

$$\pi_t = \mu_\pi E_{t-1}\bar{\pi}_{t+3} + (1-\mu_\pi)\Sigma_{j=1}^{4}\alpha_{\pi j}\pi_{t-j} + \alpha_y y_{t-1} + \varepsilon_t, \quad (16)$$

$$y_t = \mu_y E_{t-1}y_{t+1} + (1-\mu_y)\Sigma_{j=1}^{2}\beta_{yj}y_{t-j} - \beta_r(r_{t-1} - r^*) + \eta_t, \ (17)$$

where $E_{t-1}\bar{\pi}_{t+3}$ represents the expectation of average inflation over the next year and $E_{t-1}y_{t+1}$ represents the expectation of period $t+1$ output conditional on a time $t-1$ information set. The real rate relevant for output, r_{t-1}, is defined as a weighted combination of an ex ante one-year rate and an ex post one-year rate:

$$r_{t-1} = \mu_r(E_{t-1}\bar{\imath}_{t+3} - E_{t-1}\bar{\pi}_{t+4}) + (1-\mu_r)(\bar{\imath}_{t-1} - \bar{\pi}_{t-1}), \quad (18)$$

where $\bar{\imath}_t$ is a four-quarter average of past interest rates, i.e., $\bar{\imath}_t = \frac{1}{4}\Sigma_{j=0}^{3}i_{t-j}$.

This model allows consideration of a wide range of explicit forward-looking behavior. At one extreme, the model with μ_π, μ_y, and μ_r set equal to zero matches the completely adaptive expectations model of Rudebusch and Svensson (1999) and Rudebusch (2001), which has had some success in approximating the time-series data in the manner of a small estimated vector autoregression (VAR) (see Estrella and Fuhrer 2002; Fuhrer and Rudebusch 2004). However, estimated forward-looking models also have had some success in fitting the data, as in Rotemberg and Woodford (1999) and Fuhrer (2000). The analysis below takes an eclectic view and conditions on a wide range of possible values for μ_π, μ_y, and μ_r.[25]

Table 1 summarizes the optimal amount of monetary policy inertia for various models and loss functions. The table displays the lag coefficients ρ from the *optimal* versions of the inertial Taylor rule in equations (4) and (5) across models with a range of forward-looking behavior and using the three different loss functions in

[24] Much of the appeal of the New Keynesian model lies in its foundations in a dynamic general equilibrium model with nominal price rigidities; see Walsh (2003) and Woodford (2003).

[25] In contrast, there is less contention regarding the values of the other parameters in the model, and these are set equal to the values given in table 1 of Rudebusch (2002b).

Table 1. Partial-Adjustment Coefficients for Optimal Inertial Taylor Rules

Model			Optimal ρ for Different Loss Functions					
μ_r	μ_π	μ_y	$\nu_{\Delta i} = .1$	$\nu_{\Delta i} = .5$	$\nu_{Ei} = .1$	$\nu_{Ei} = .5$	$\nu_i = .1$	$\nu_i = .5$
.0	.0	.0	−.12	.19	−.04	.34	−.57	−.51
.3	.3	.3	.18	.37	.27	.48	−.27	−.12
.5	.5	.5	.64	.70	.63	.68	.70	.80
.8	.8	.8	.90	.94	.90	.92	.93	.96
.0	.0	.5	.03	.17	.05	.26	−.34	−.23
.0	.5	.0	−.12	.16	−.04	.31	−.54	−.44
.5	.0	.0	.49	.61	.49	.67	.28	.30

Notes: The optimal lag coefficients for an inertial Taylor rule are reported for each of seven parameterizations of the model, which have varying μ_π, μ_y, and μ_r weights on expectational terms, and for six variations of the loss function. The loss functions have equal weight on output and inflation volatility ($\lambda = 1$) but a stronger or weaker interest rate smoothing motive—which may take the form of minimizing $\nu_{\Delta i}\text{Var}[\Delta i_t]$, $\nu_{Ei}\text{Var}[E_{t-1}[i_t] - i_t]$, or $\nu_i\text{Var}[i_t - r^* - \pi^*]$. The associated optimal g_π and g_y are not reported.

equations (13), (14), and (15). For each loss function, the weight on the interest rate smoothing ($\nu_{\Delta i}$, ν_{Ei}, or ν_i) is set equal to .5 or .1, while $\lambda = 1$.[26] Clearly in table 1, a large range of optimal lag coefficients—between –.6 and 1.0—can be rationalized for some combination of model and loss function. Most interesting, however, is how the expectational channel can magnify even a small cost of interest rate fluctuations in the central bank loss function to produce a sizable partial-adjustment coefficient in the policy rule. Also, note that the degree of optimal monetary policy inertia varies most strongly with the value of μ_r, which determines the degree to which interest rate expectations are forward looking. Such variation is consistent with the interpretation of Woodford (1999, 2003) and Levin,

[26] The results in table 1 are obtained by numerically minimizing the loss function over the parameters g_π, g_y, and ρ in the model. The results are obtained using the "AIM" algorithm (Anderson and Moore 1985), available at www.federalreserve.gov/pubs/oss/oss4/aimindex.html.

Wieland, and Williams (1999) that policy inertia is optimal when it alters expectations of future interest rates that are also important determinants of current demand.

While an expectational channel for optimal monetary policy inertia is valid in principle, it seems unlikely that such a channel is responsible for empirical monetary policy inertia, because its underlying assumption of a fully credible policy rule seems so unlikely historically. That is, even if economic agents were sufficiently forward looking (which is a separate, unresolved issue), the monetary policy rule must also be assumed to be perfectly credible, so agents know the rule and correctly assume that it will be followed.[27] This seems an unlikely description even for the relatively homogeneous 1987–2004 U.S. sample period underlying the above inertial policy rule estimates. In practice, the Federal Reserve may exhibit some transparency, but it does not appear to have a commitment technology.[28]

3.3 Uncertainty and Partial Adjustment

Uncertainty is the third general rationale often used to motivate optimal monetary policy inertia. The intuition appears clear: uncertainty breeds caution, and caution suggests a gradual adjustment of the policy rate. As noted by Bernanke (2004), "Because policymakers cannot be sure about the underlying structure of the economy or the effects that their actions will have on economic outcomes, and because new information about the economic situation arrives continually, the case for policymakers to move slowly and cautiously when changing rates seems intuitive." However, the implication that greater uncertainty produces greater inertia is not a general theoretical result, and the empirical evidence for this proposition appears weak as well.

[27]Still, this rationale may be a fruitful area for further research, particularly in examining cases, as in Kara (2003), of partial credibility and an intermediate amount of inertia.

[28]Informally, note that the Federal Reserve does not seem to have the requisite control over forward interest rates (as evidenced most recently by central banks' consternation regarding the "conundrum" of low long-term bond yields described in Rudebusch, Swanson, and Wu 2006). A formal commitment counterfactual is given in Dennis (2005).

Because economic data can be quite noisy, policymakers inevitably operate with imperfect knowledge about the current state of the economy. In addition, it may be the case that the noisier the economic data are, the less aggressive policymakers should be in responding to current readings on the economy (Rudebusch 2001; Orphanides 2003).[29] However, in empirical models, as noted by Rudebusch (2001), any such inducement toward timidity (that is, a low g_π and g_y) appears fairly modest and does not necessarily translate into greater sluggishness (that is, a high ρ).

Uncertainty about the model provides another rationale for caution. Indeed, ever since the classic Brainard (1967) analysis, uncertainty about the quantitative impact of policy and the dynamics of the economy has been widely cited as a rationale for damped policy action. However, in the general case, as Chow (1975, chap. 10) makes clear, *almost nothing can be said even qualitatively* about how the optimal rule under model uncertainty changes relative to the optimal rule under certainty. For example, the optimal policy response parameters are not necessarily reduced in the presence of uncertainty about several parameters. Thus, quantitative answers are required. Rudebusch (2001) provides some simple but instructive evidence that suggests that parameter uncertainty is not responsible for policy inertia. The policymaker is assumed to face an economy like (13), (14), and (15) *on average* (with μ_π, μ_y, and μ_r set equal to zero), but in any given quarter, the coefficients may take on a random value. These parameter shifts occur every quarter or every few years. The policymaker has to choose the g_π, g_y, and ρ parameters of the inertial Taylor rule (1) and (2), so that the loss function (10) is minimized. After allowing for uncertainty about all of the coefficients of the model, the optimal partial-adjustment coefficient actually falls a bit.[30]

[29] The general certainty-equivalence guideline is that optimal policy requires the same response under both partial and full information about the state of the economy. However, as discussed in Rudebusch (2001), the use of simple rules and inefficient output-gap estimates are two relevant exceptions for this analysis.

[30] This conclusion accords with much research on parameter uncertainty. Notably, in Estrella and Mishkin (1999), Peersman and Smets (1999), Shuetrim and Thompson (1999), and Tetlow and von zur Muehlen (2001), there is no significant attenuation of the rule parameters. Some attenuation is found in Salmon and Martin (1999), Söderström (1999), and Sack (2000).

Overall, although perhaps intuitive, the argument that uncertainty could account for the very gradual persistence in the data remains unproven.

4. Term-Structure Evidence on Inertial Policy Rules

To summarize the discussion so far, the single-equation estimation of policy rules has yielded inconclusive results regarding the existence of policy inertia, and the theoretical case for substantial interest rate smoothing appears unconvincing as well. To make some progress, this section turns to a vast and rich set of information about central bank reaction functions: the yield curve of interest rates. The yield curve contains such information because if financial market participants understand the policy rule that links short-term interest rates to the realizations of macroeconomic variables, then they will also use that rule in forming expectations of future short-term interest rates, which will be priced into long-term bonds.[31] In particular, any deviations from the policy rule embedded in expected future short-term rates and expected macroeconomic conditions would be arbitraged away. Therefore, at any point in time, multiperiod interest rates, which embody expectations of future short rates, contain much information about the properties of the reaction function (also see Ang, Dong, and Piazzesi 2005). In this section, I outline three different methods by which this information can be extracted to inform the debate on policy inertia. These methods differ primarily by the amount of economic structure imposed and by the frequency of data employed.[32]

[31]Note that the assumption is not one of credibility and commitment as in subsection 4.2 but one of transparency and learnability.

[32]For example, the three methodologies below use three different treatments of interest rate risk premiums. In the first one, a time-varying term premium is modeled in a simple ad hoc empirical fashion. In the second, a theoretical no-arbitrage consistency is enforced between the underlying factor dynamics and the term premium. In the final one, the term premium is assumed constant (i.e., the expectations hypothesis is assumed) over short thirty-minute windows.

4.1 Interest Rate Predictability at a Quarterly Frequency

Policy inertia has important implications for interest rate forecastability: in brief, the greater the delayed adjustment of the policy rate in reaction to current information, the greater the amount of forecastable future variation. Intuitively, if the funds rate typically is adjusted 20 percent toward its desired target in a given quarter, then the remaining 80 percent of the adjustment should be expected to occur in future quarters. Furthermore, assuming financial markets understand the inertial nature of monetary policy, they should anticipate the future partial adjustment of the funds rate and incorporate it into the pricing of longer-term maturities.

Rudebusch (2002b) shows that this general intuition is true in a wide variety of macroeconomic models. The amount of such forecastable variation in interest rates can be measured via a standard term-structure regression at a quarterly frequency such as

$$\Delta i_t = \delta + \gamma E_{t-2}(\Delta i_t) + \psi_t. \tag{19}$$

This equation regresses the realized change in the policy rate in quarter t (i.e., $\Delta i_t = i_t - i_{t-1}$) on the change that was expected two quarters earlier at the end of period $t-2$. Under rational expectations, this interest rate forecasting regression would yield in the limit an estimate of $\hat{\delta} = 0$ and $\hat{\gamma} = 1$. However, for assessing the forecastable variation in the interest rate and hence the degree of monetary policy inertia, the statistic of particular interest is the R^2 of this regression, which provides a natural measure of forecastability.

The theoretical relationship between the forecastable variation in the interest rate, as measured by the R^2 of the above prediction equation, and quarterly policy inertia, as measured by the ρ in the Taylor rule (1) and (2), is illustrated in figure 8. This figure graphs the implied (population) value of the R^2 of the regression (19) as a function of ρ for a representative case of the model described in section 3, namely, with $\mu_\pi = .3$, $\mu_r = .5$, and $\mu_y = 0$.[33] Note

[33] Also, g_π and g_y are set equal to 1.5 and 0.8, respectively. As in table 1, the unique stationary rational expectations solution for each specified policy rule and model is solved via AIM (see Anderson and Moore 1985 and Levin, Wieland, and Williams 1999). The reduced-form representation of the saddle-point solution is computed, the unconditional variance-covariance matrix of the

Figure 8. Implications for Interest Rate Forecastability from Policy Inertia

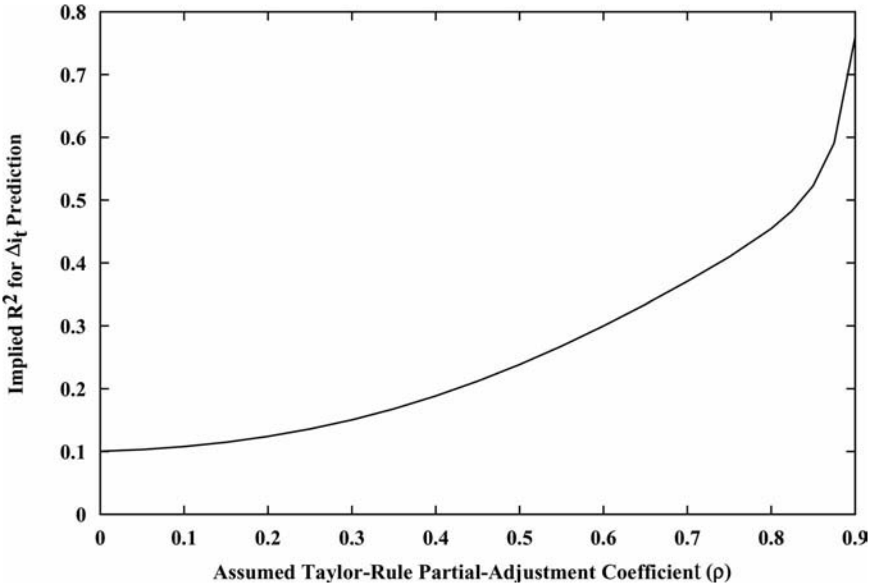

that even for the non-inertial policy rules, there is some predictable future movement in interest rates (with $R^2 = .10$ when $\rho = 0$), since there are predictable changes two quarters ahead in the output gap and in the four-quarter inflation rate, which partly determine future changes in interest rates. Even though the output gap and inflation are highly persistent in levels, the associated slow mean reversion implies only a modest predictability of future quarterly *changes* in these series and the desired funds rate. However, as ρ increases, the amount of predictable variation in Δi_{t+2} also increases, with an R^2 value of .45 at $\rho = 0.8$.

Rudebusch (2002b) shows that this theoretical relationship between partial adjustment and predictability is robust across a wide variety of models (and for forecast-based policy rules as well). This relationship can be empirically assessed by examining the extent of

---

model variables and the term spreads is obtained analytically, and the term-structure regression asymptotic R^2 is calculated using the appropriate variances and covariances.

forecastable future movements in the policy interest rate in the data. Specifically, if policy is highly inertial, as the single-equation reaction functions suggest, then financial markets should anticipate the future partial adjustment of the funds rate. In that case, a regression of actual changes in the funds rate on predicted changes embedded in the yield curve should provide a good explanatory fit and a fairly high R^2. In fact, researchers have found the opposite. They have estimated a variety of interest rate forecasting regressions and, using financial market expectations, have found little predictive information at quarterly frequencies with R^2s very close to zero.[34] For example, Rudebusch (2002b) shows that eurodollar futures from 1988 to 2000 have very little ability to predict the quarterly change in the funds rate two quarters ahead. The R^2 of such a regression is .11, which from figure 8 suggests that ρ is probably close to zero.[35]

This lack of predictive ability is well illustrated by the most recent episode of monetary policy *easing*. Figure 9 gives the actual target federal funds rate and various expected funds rate paths as of the middle of each quarter based on federal funds futures. Under quarterly policy inertia, the long sequence of target changes in the same direction in 2001 would be viewed as a set of gradual partial adjustments to a low desired rate. However, although the funds rate gradually fell in 2001, market participants actually anticipated few of these declines at a six- to nine-month horizon, as they would have if policy inertia were in place. Instead, markets assumed at each point in time that the Federal Reserve had adjusted the funds rate down to just about where it wanted the funds rate to remain based on current information available. Under this interpretation, the long sequence of declines is the result of a series of fairly prompt responses to new information that turned progressively more pessimistic. That is, the

[34] See, for example, Mankiw and Miron (1986) and Rudebusch (1995).

[35] Rudebusch (2002b) used a variety of structural models to show that the large estimated lag coefficients in the empirical inertial policy rules provided were inconsistent with the very low interest rate forecastability in the term structure of interest rates and that rules with highly serially correlated errors do not imply such forecastability. Söderlind, Söderström, and Vredin (2005) examine the former issue using a VAR model and survey data on macroeconomic forecasts and find evidence inconsistent with the standard inertial Taylor rule. In contrast, in a highly forward-looking empirical model, Berkelmans (2006) argues that a very inertial policy rule could be consistent with the interest rate predictability evidence.

Figure 9. Actual and Expected Federal Funds Rate

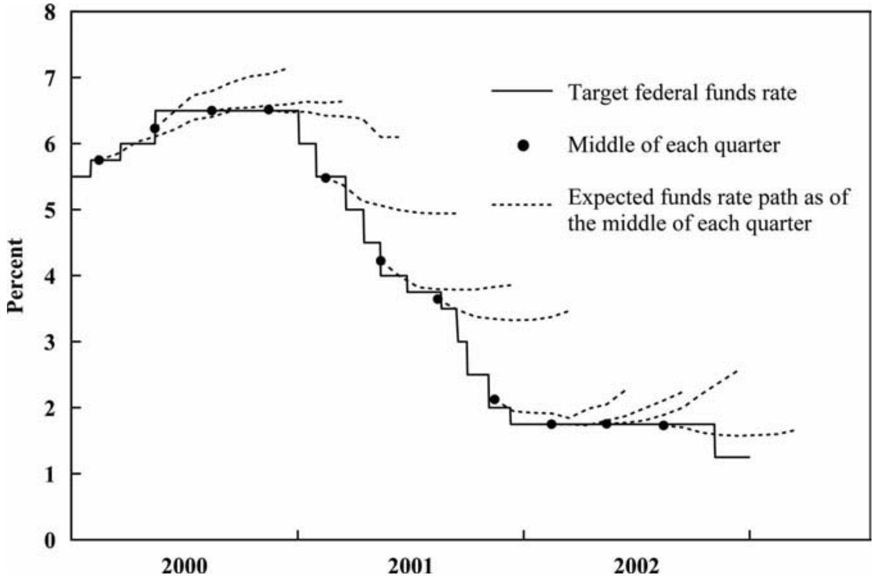

presence of quarterly partial adjustment or policy inertia is contradicted by the lack of forecastability of changes in the funds rate.

The latest episode of monetary policy *tightening* in the United States may at first glance seem to offer more support for gradualism and predictability in interest rates. During this episode, the FOMC raised the target federal funds rate by 25 basis points at each of the seventeen FOMC meetings that occurred during the two years from June 2004 through June 2006. Of course, the mere fact that the Federal Reserve engaged in a long series of interest rate increases is not informative regarding quarterly monetary policy inertia. Such persistent cyclical movements could reflect persistent changes in the determinants of policy rather than the gradual adjustment of the Federal Reserve to those determinants.[36] However, as described in

[36]Occasionally, the argument is made that long sequences of interest rate increases and decreases necessarily imply that changes in interest rates are predictable. This is the perennial argument of chartists and would suggest, for example, that equity prices, the dollar exchange rate, and commodity prices are all forecastable. Also, see the discussion in Goodhart (2005).

Figure 10. Actual and Expected Federal Funds Rate

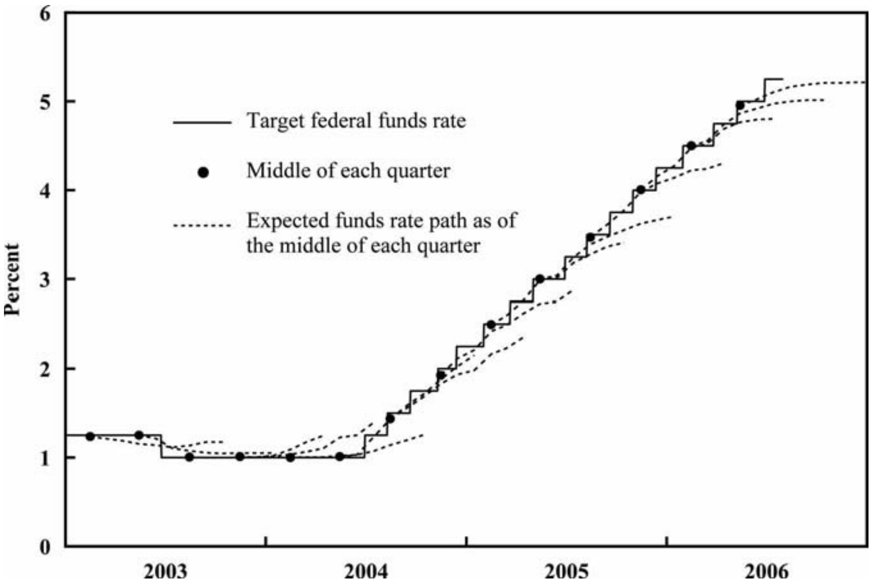

Rudebusch and Williams (2006), this latest episode was unprecedented in that the FOMC provided direct verbal signals about future policy rate changes. Starting in May 2004, the FOMC introduced the following language into its public statement: "The Committee believes that policy accommodation can be removed at a pace that is likely to be measured." This was a direct, though not unambiguous, indication that the FOMC anticipated that the policy interest rate could be gradually increased, and it was replaced in December 2005 by "some further policy firming is likely to be needed," and in January 2006 by "further policy firming may be needed."

These verbal signals of future policy intentions would seem likely to boost the predictability of interest rates, and, to a large extent, this appears to have occurred but—importantly—largely at very short horizons. This effect can be discerned in figure 10, which gives the actual target federal funds rate and various expected funds rate paths as of the middle of each quarter based on federal funds futures from 2003 to 2006. It is apparent that many of the expected interest rate paths are remarkably well aligned with the actual path for the first three or four months into the future; however, after about four

Figure 11. Actual and Expected Change in the Funds Rate

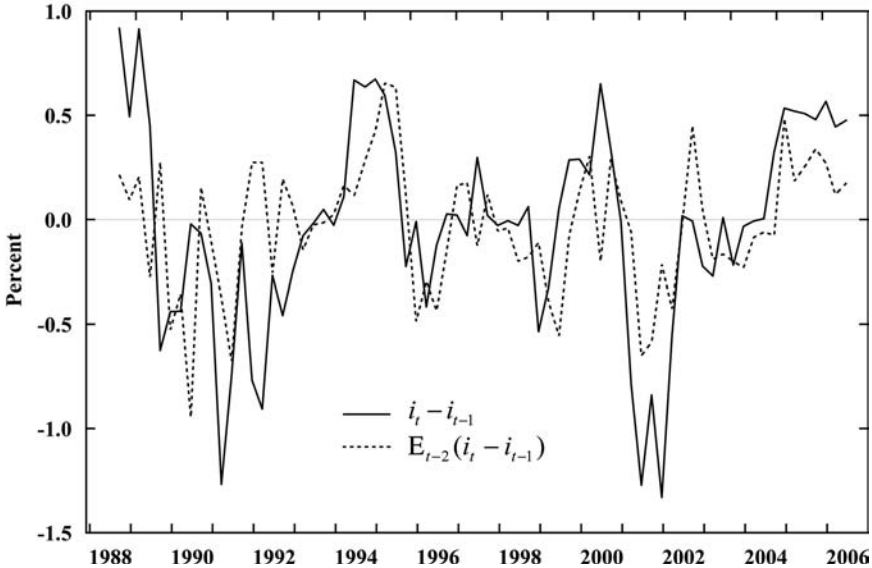

months, financial markets consistently underestimated the extent of the future tightening. That is, markets expected an even more gradual pace for the policy tightening than actually occurred. This is not too surprising, since FOMC members made it clear that future policy depended importantly on how the economic data unfolded in real time, and during much of this episode the economic recovery was not viewed as well established. For example, as the then–Vice Chairman of the Board of Governors noted (Ferguson 2004): "I believe it to be very important that the FOMC not go on a forced march to some point estimate of the equilibrium real federal funds rate. In my judgment, we should remove the current degree of accommodation at a pace that is importantly determined by incoming data and a changed outlook."

With respect to the forecasting regression (19), which is crucial for judging the extent of quarterly policy inertia, figure 11 displays the regression data, Δi_t and $E_{t-2}(\Delta i_t)$, updated through 2006:Q2. From this perspective, the past few years do not look that unusual. Indeed, the residuals from the forecasting regression, ψ_t, are plotted

Figure 12. Residuals from Funds Rate Forecast Regression

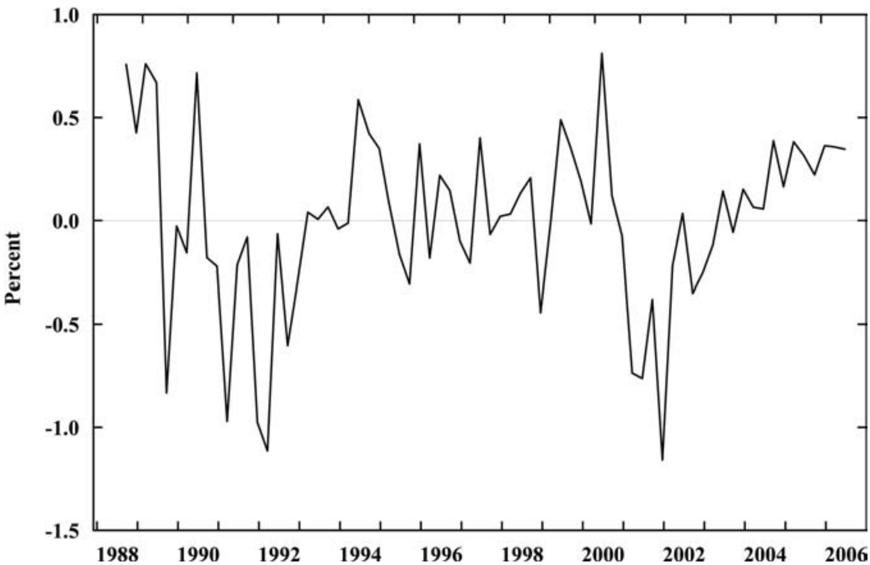

in figure 12, and the last two years of the sample are not notable for exhibiting extreme accuracy. Therefore, it appears that the recent tightening episode was an example of short-run smoothing of policy rates in the United States but is not inconsistent with the view that policymakers engage in a limited amount of quarterly policy inertia.

4.2 Term-Structure Model Estimation at a Monthly Frequency

While the evidence in section 4.1 on the predictability of interest rates is quite intuitive, it is somewhat indirect; that is, the absence of policy inertia is inferred from the lack of interest rate predictability evident in financial markets. More-direct estimates of the degree of interest rate smoothing would perhaps be more compelling, and this section considers direct estimates of ρ. However, these estimates of ρ differ from the single-equation ones given in section 2 because they are obtained in a complete system that combines key macroeconomic equations and information from the yield curve. The particular structure employed is from Rudebusch and Wu (2006). Their analysis uses monthly data to formally estimate a model that combines a fairly standard macroeconomic model with an off-the-shelf, no-arbitrage finance representation from the empirical bond-pricing literature. Again, it is the incorporation of yield-curve information that allows precise inference about the absence of monetary policy inertia.

Almost all movements in the yield curve can be captured in a no-arbitrage framework in which yields are linear functions of a few unobservable or latent factors (e.g., Duffie and Kan 1996; Dai and Singleton 2000). The Rudebusch-Wu macrofinance model employs such a framework: specifically, it features a constant factor volatility with state-dependent risk pricing of volatility, which implies conditionally heteroskedastic risk premiums. The one-month short rate is the sum of a constant and two unobserved term-structure factors,

$$i_t = \delta_0 + L_t^m + S_t^m, \tag{20}$$

where L_t^m and S_t^m are termed level and slope factors. The dynamics of these latent factors are given by

$$L_t^m = \rho_L L_{t-1}^m + (1-\rho_L)\pi_t + \varepsilon_{L,t} \tag{21}$$

$$S_t^m = \rho_S S_{t-1}^m + (1-\rho_S)\big[g_y y_t + g_\pi(\pi_t - L_t^m)\big] + u_{S,t} \tag{22}$$

$$u_{S,t} = \rho_u u_{S,t-1} + \varepsilon_{S,t}, \tag{23}$$

where π_t and y_t are inflation and the output gap.[37] These equations can be given a Taylor-rule interpretation, with the factor L_t^m interpreted as the inflation rate targeted by the central bank, as perceived by private agents. Private agents slowly modify their views about L_t^m as actual inflation changes, so L_t^m is associated with an interim or medium-term inflation target (as in Bomfim and Rudebusch 2000) with associated underlying inflation expectations over the next two to five years. The slope factor S_t^m captures the central bank's dual mandate to stabilize the real economy and keep inflation close to its target level. In addition, the dynamics of S_t^m allow for both partial adjustment and serially correlated shocks. If $\rho_u = 0$, the dynamics of S_t^m arise from monetary policy partial adjustment, as in an inertial Taylor rule. Conversely, if $\rho_S = 0$, the dynamics reflect the Federal Reserve's reaction to autocorrelated information or events not

[37] In this substitution with monthly data, π_t is the twelve-month inflation rate and y_t is capacity utilization.

captured by output and inflation, as in the Taylor rule with AR(1) shocks.

Appended to the above equations is a small macroeconomic model of inflation and output suitable for estimation with monthly data, which also has some New Keynesian justification:

$$\pi_t = \mu_\pi L_t^m + (1-\mu_\pi)[\alpha_{\pi_1}\pi_{t-1} + \alpha_{\pi_2}\pi_{t-2}] + \alpha_y y_{t-1} + \varepsilon_{\pi,t}. \quad (24)$$

$$y_t = \mu_y E_t y_{t+1} + (1-\mu_y)[\beta_{y1} y_{t-1} + \beta_{y2} y_{t-2}] - \beta_r(i_{t-1} - L_{t-1}^m) + \varepsilon_{y,t}. \quad (25)$$

That is, inflation in the current month is set as a weighted average of the public's expectation of the medium-term inflation target, identified as L_t^m, and two lags of inflation. Also, there is a one-month lag on the output gap to reflect adjustment costs and recognition lags. Current output is determined by expected future output, $E_t y_{t+1}$, lagged output, and the ex ante real interest rate, which is proxied by $i_{t-1} - L_{t-1}^m$ (that is, agents judge nominal rates against their view of the underlying future inflation rate, not just next month's inflation). Finally, the specification of longer-term yields follows the standard no-arbitrage formulation. For pricing longer-term bonds, the risk price associated with the structural shocks is assumed to be a linear function of L_t^m and S_t^m.

The above macrofinance model was estimated by maximum likelihood (ML) for the sample period from January 1988 to December 2000. The complete parameter estimates and details are in Rudebusch and Wu (2006); however, of particular interest for policy inertia are the estimates of ρ_S, which is a minuscule .026, and of ρ_u, which is .975. These estimates decisively dismiss the interest rate smoothing or monetary policy inertia interpretation of the Taylor rule. The persistent rule deviations occur not because the Federal Reserve was slow to react to output and inflation, but because the Federal Reserve responds to a variety of persistent determinants beyond current output and inflation. Some intuition for this result is given in figure 13, which displays the initial response of the entire yield curve to inflation and output shocks from the estimated macrofinance model. Positive shocks to inflation and output in this model are followed by immediate increases in short-term interest rates, and, for the inflation shock, these increases are more than one-for-one. These responses—shown as the solid lines—reflect the absence

Figure 13. Initial Yield-Curve Response to Output and Inflation Shocks

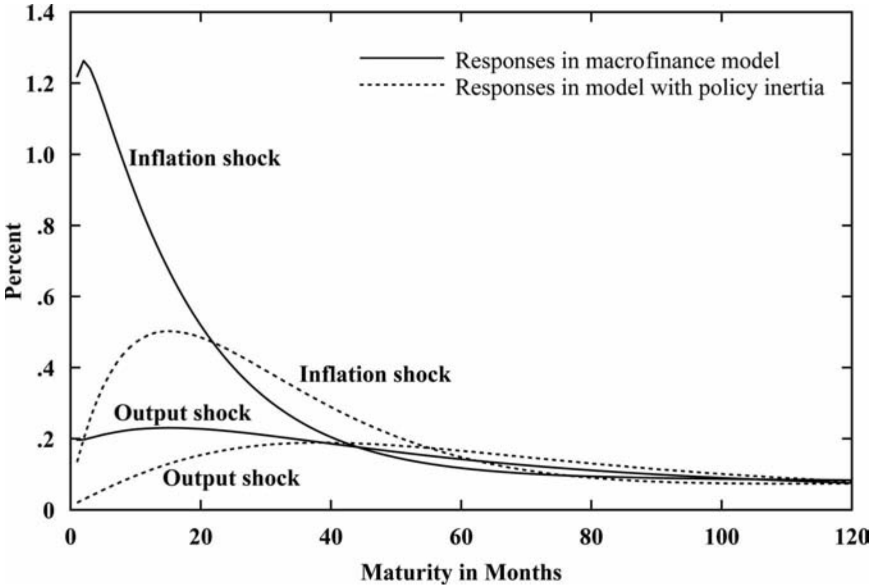

Note: The solid lines show the impact responses on the entire yield curve from a 1-percentage-point increase in inflation or output in the estimated macrofinance model in Rudebusch and Wu (2006). The dashed lines give similar responses in a macrofinance model that assumes substantial monetary policy inertia ($\rho_S = 0.9$) and serially uncorrelated policy shocks ($\rho_u = 0$).

of monetary policy partial adjustment or inertia. In contrast, the dashed lines in figure 13 display the yield-curve responses from a model that is identical to the estimated macrofinance model except that ρ_S is set equal to .9 and ρ_u equals 0. This hypothetical alternative model has substantial monetary policy inertia, and it displays markedly weaker responses to inflation and output shocks of yields that have maturities of less than two years. The two quite different responses of the yield curve in these models illustrate the potential importance of the information contained in the term structure for discriminating between the two models. Given the system ML estimates, it is clear that the data prefer the macrofinance model without policy inertia.

4.3 Intraday Interest Rate Reactions to Macroeconomic Data

As a third illustration of the power of the term structure to illuminate the nature of the monetary policy reaction function, I provide some new evidence on interest rate smoothing based on intraday movements of the yield curve. The underlying insight exploited here is similar to the one above: an inertial policy rule has important implications for the evolution of the entire term structure through time. Again, this approach is extremely powerful, because financial markets will enforce their understanding of the monetary policy reaction function at each point in time and across interest rates at all maturities. However, while the above results implement this idea with models estimated at a monthly or quarterly frequency and substantial economic structure, the results in this section are based on the intraday response of the yield curve to macroeconomic data releases and impose minimal structure. The resulting event study provides further compelling evidence against the existence of monetary policy inertia using very different data and information.

Intuitively, changes in the path of expected future interest rates following the release of news about the state of the economy should reveal the degree of interest rate smoothing, because financial markets will expect an inertial central bank to distribute the policy rate changes over several periods. To illustrate this mechanism in a simple formal structure, consider the policy inertia framework

$$i_t = (1-\rho)\beta\bar{\pi}_t + \rho i_{t-1}, \tag{26}$$

where i_t is the average short-term (daily) policy rate during quarter t, which is set by the central bank to respond gradually over time to the annual inflation rate, $\bar{\pi}_t$ (the four-quarter percent change). Also, annual inflation is assumed to be a simple AR(1) process,

$$\bar{\pi}_t = \delta\bar{\pi}_{t-1} + \varepsilon_{1,t} + \varepsilon_{2,t}, \tag{27}$$

with two sources of independent random variation. These two shocks are distinguished by the timing of their release dates during the quarter. News about inflation in $\varepsilon_{1,t}$ is revealed at the very beginning of quarter t, while the news in $\varepsilon_{2,t}$ is revealed sometime later in quarter t. This analysis just explores the effects of news in $\varepsilon_{1,t}$, while $\varepsilon_{2,t}$ is only included in the model to emphasize that knowledge of $\varepsilon_{1,t}$ does not determine $\bar{\pi}_t$. Also, one of the key elements of the methodology in this section is that δ can be pinned down by macroeconomic time-series data. In particular, for the inflation and output series shown in figure 1, which are the relevant policy determinants in the Taylor rule, the OLS estimates of δ, which have a well-known downward bias, are .97 for inflation and .95 for output. This evidence is consistent with the large literature examining the persistence of various macroeconomic series that indicates that δ is very close to 1.[38]

To calculate the immediate response of interest rates to the revelation of $\varepsilon_{1,t}$, note that at the end of period $t-1$, the expected value of the average interest rate over the next quarter is

$$E[i_t|e(t-1)] = \rho i_{t-1} + (1-\rho)\beta E[\pi_t|e(t-1)] \tag{28}$$

$$= \rho i_{t-1} + (1-\rho)\beta\delta\pi_{t-1}, \tag{29}$$

where $E[\cdot|e(t-1)]$ is the expectation conditional on the information set at the end of quarter $t-1$. Similarly, just after the revelation of $\varepsilon_{1,t}$ at the beginning of quarter t, the expected value of the quarter-t interest rate is

$$E[i_t|b(t)] = \rho i_{t-1} + (1-\rho)\beta E[\pi_t|b(t)] \tag{30}$$

$$= \rho i_{t-1} + (1-\rho)\beta(\delta\pi_{t-1} + \varepsilon_{1,t}), \tag{31}$$

where $E[\cdot|b(t)]$ is the expectation conditional on the information set at the beginning of quarter t. Therefore, the size of the revision to the expectation of i_t in response to $\varepsilon_{1,t}$ news about inflation equals

$$\Delta E[i_t|\Delta] \equiv E[i_t|b(t)] - E[i_t|e(t-1)] = (1-\rho)\beta\varepsilon_{1,t}. \tag{32}$$

That is, the change in the expectation of i_t equals the amount of inflation news multiplied by the policy response coefficient and reduced by a fraction for interest rate smoothing. Still, even with data on the change in interest rate expectations, it is difficult to determine the size of ρ from this equation, on its own, because $\beta\varepsilon_{1,t}$ must be measured in some way.[39]

[38] For evidence on this point and references to the voluminous literature, see Rudebusch (1992), Rudebusch and Svensson (1999), and Pivetta and Reis (2006).

[39] Macroeconomic data surprises relative to surveys of market participants may help but are clouded by information in revisions to earlier data.

However, combining the revisions in expectations of i_t with revisions of other expected future interest rates does allow the partial-adjustment coefficient to be determined. Specifically, note that at the end of quarter $t-1$, the expected value of i_{t+1} is

$$E[i_{t+1}|e(t-1)] = \rho E[i_t|e(t-1)] + (1-\rho)\beta E[\pi_{t+1}|e(t-1)] \quad (33)$$

$$= \rho^2 i_{t-1} + (\delta+\rho)(1-\rho)\beta\delta\pi_{t-1}. \quad (34)$$

At the beginning of quarter t, the expected value of i_{t+1} is

$$E[i_{t+1}|b(t)] = \rho E[i_t|b(t)] + (1-\rho)\beta E[\pi_{t+1}|b(t)] \quad (35)$$

$$= \rho^2 i_{t-1} + (\delta+\rho)(1-\rho)\beta(\delta\pi_{t-1} + \varepsilon_{1,t}), \quad (36)$$

so the revision to expectations of i_{t+1} in response to $\varepsilon_{1,t}$ is equal to

$$\Delta E[i_{t+1}|\Delta] \equiv E[i_{t+1}|b(t)] - E[i_{t+1}|e(t-1)] = (\delta+\rho)(1-\rho)\beta\varepsilon_{1,t}. \quad (37)$$

Finally, the ratio of the two revisions provides a straightforward expression:

$$\Delta E[i_{t+1}|\Delta]/\Delta E[i_t|\Delta] = \delta + \rho. \quad (38)$$

If, as noted above, the value of δ is pinned down by the well-known macroeconomic dynamics of inflation, then this ratio of revisions in expected future rates will identify ρ.

To estimate the ratio above, I use intraday data on yields of three- and six-month U.S. Treasury securities.[40] The revisions in these two yields are calculated over the half-hour period from five minutes before a release of macroeconomic data to twenty-five minutes after that release.[41] Changes in the three-month yield during this window

[40] I also obtained similar results using interest rate expectations from daily federal funds futures and eurodollar futures. However, an advantage to using the Treasury yields is that they enforce a consistent timing so that the macroeconomic news always occurs at the beginning of the monetary policy adjustment. In addition, Treasury markets are the most liquid ones.

[41] This thirty-minute window eliminates noise from extraneous sources, such as other data releases or monetary policy actions or communications. The data are discussed in Gürkaynak, Sack, and Swanson (2005a, 2005b) and were kindly supplied by the authors. They obtained tick-by-tick, on-the-run Treasury yield data back to 1991 from a consortium of interdealer brokers. They also show that a thirty-minute window is sufficiently wide to capture the full response of financial markets to news.

provide a reading on $\Delta E[i_t|\Delta]$, while changes in a combination of the two yields give $\Delta E[i_{t+1}|\Delta]$ via the expectations hypothesis—namely, twice the six-month yield minus the three-month yield.[42] For example, if the three-month rate increases by 5 basis points in response to a release of higher-than-expected consumer price inflation, and the three-month rate expected three months ahead increases by 9 basis points, then their ratio provides an estimate of $\delta + \rho$ equal to 1.8. Assuming inflation follows a unit-root process, so $\delta = 1$, then the monetary policy partial-adjustment coefficient is 0.8. That is, in response to news about persistently higher inflation, financial markets assume that an inertial central bank will boost the policy rate higher over the next few months but will also gradually raise it even higher in subsequent months. Alternatively, at the opposite end of the spectrum, if three- and six-month yields change by an identical amount in response to a persistent shock (so $\Delta E[i_t|\Delta] = \Delta E[i_{t+1}|\Delta]$), then $\delta + \rho = 1$ and financial markets assume that there is no monetary policy partial adjustment by central banks.

In fact, the data indicate quite clearly that the case of little or no inertia is the relevant one. I consider 315 macroeconomic data releases from July 1991 to September 2004 for the unemployment and CPI series, which are two of the most important and closely watched data releases. Of course, the formal structure outlined above applies to any persistent macroeconomic determinant of monetary policy, so the unemployment and CPI releases are pooled to increase the precision of the estimates. The median value of $\Delta E[i_{t+1}|\Delta]/\Delta E[i_t|\Delta]$ is 1.00; the mean value is 1.06 with a standard error of 0.15.[43] Again, with the assumption that macroeconomic time series are highly persistent, these results imply a central tendency for ρ of around 0 to .1 and a 95 percent confidence interval that lies entirely below .4.[44]

[42] This calculation ignores the time-varying term premium modeled in Rudebusch and Wu (2006) and discussed above; however, changes in the ratio of these premiums at these very short maturities are likely insignificant.

[43] The median expectational revision ratios for inflation and unemployment releases separately are also both equal to 1.0.

[44] These results also appear robust to consideration of longer maturities, as in $\Delta E[i_{t+k}|\Delta]/\Delta E[i_t|\Delta]$.

5. Conclusion

Does the persistence of the short-term policy interest rate reflect deliberate "partial adjustment" or "inertia" on the part of the central bank? As in many other areas of economics, understanding the nature of dynamic adjustment is a hard problem that simple regression estimates often cannot solve. However, in contrast to many other macrodynamic puzzles, interest rates have a rich set of term-structure information that can help provide answers. One of the key insights above is that although the short rate is a policy instrument, it is also a fundamental driver of long yields, so a joint macrofinance perspective can sharpen inference about the policy reaction function. The yield-curve results above—for quarterly predictability, monthly system estimation, and intraday responses to news—all point to fairly rapid central bank reactions to news and information and little real-world policy inertia. In essence, quarterly monetary policy partial adjustment does not appear to be consistent with the financial market's understanding of the monetary policy rule. This absence of intrinsic inertia appears in accord with the views of many central bankers, who often note that future policy actions will largely be contingent on incoming data and future changes in the economic outlook.

In terms of future research, much work can still be done to exploit yield-curve information about the monetary policy reaction function, especially in countries other than the United States. In addition, further policy rule estimation and investigation is recommended. The lagged policy rate in empirical monetary policy rules, although perhaps useful in mopping up residual serial correlation, should not be given a structural partial-adjustment interpretation with regard to central bank behavior. A better strategy is to identify and model the underlying persistent factors that influence central bank actions. This task will not be easy. As Svensson (2003, 467) argues, the missing elements may be largely judgmental in nature:

> Whereas simple instrument rules, like variants of the Taylor rule, may to some extent serve as very rough benchmarks for good monetary policy, they are very *incomplete* rules, because they don't specify when the central bank should or should not

> deviate from the simple instrument rule. Such deviations, by discretion and judgement, have been and will be frequent. . . .

In this case, policymakers should not be misled into viewing a Taylor rule, or any simple representation of policy, as a completely reliable guide to future actions.

References

Anderson, G. S., and G. Moore. 1985. "A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models." *Economics Letters* 17 (3): 247–52.

Ang, A., S. Dong, and M. Piazzesi. 2005. "No-Arbitrage Taylor Rules." Manuscript, University of Chicago (August).

Apel, M., and P. Jansson. 2005. "Some Further Evidence on Interest-Rate Smoothing: The Role of Measurement Errors in the Output Gap." Manuscript, Sveriges Riksbank (March).

Berkelmans, L. 2006. "Gradual Adjustment Versus Persistent Shocks: An Assessment of Predictability and Interest Rate Smoothing." Manuscript, Harvard University.

Bernanke, B. 2004. "Gradualism." Remarks at an economics luncheon, Seattle, Washington, May 20.

Blinder, A. S. 1986. "More on the Speed of Adjustment in Inventory Models." *Journal of Money, Credit, and Banking* 18 (3): 355–65.

Bomfim, A., and G. D. Rudebusch. 2000. "Opportunistic and Deliberate Disinflation Under Imperfect Credibility." *Journal of Money, Credit, and Banking* 32 (4): 707–21.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review Papers and Proceedings* 57 (2): 411–25.

Carare, A., and R. Tchaidze. 2005. "The Use and Abuse of Taylor Rules: How Precisely Can We Estimate Them?" IMF Working Paper No. 05/148 (July).

Castelnuovo, E. 2006. "Taylor Rules and Interest Rate Smoothing in the Euro Area." Forthcoming in *The Manchester School*.

Chow, G. C. 1975. *Analysis and Control of Dynamic Economic Systems*. New York, NY: John Wiley & Sons.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

———. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1): 147–80.

Cukierman, A. 1996. "Why Does the Fed Smooth Interest Rates?" In *Monetary Policy on the 75th Anniversary of the Federal Reserve System*, ed. M. T. Belongia, 111–47. Boston: Kluwer Academic Publishers.

Dai, Q., and K. J. Singleton. 2000. "Specification Analysis of Affine Term Structure Models." *Journal of Finance* 55 (5): 1943–78.

Dennis, R. 2005. "Inflation Targeting under Commitment and Discretion." *Economic Review* (Federal Reserve Bank of San Francisco) 1–13.

———. 2006. "The Policy Preferences of the US Federal Reserve." *Journal of Applied Econometrics* 21 (1): 55–77.

Driffill, J., P. Savona, Z. Rotondi, and C. Zazzara. 2006. "Monetary Policy and Financial Stability: What Role for the Futures Market?" *Journal of Financial Stability* 2 (1): 95–112.

Duffie, D., and R. Kan. 1996. "A Yield-Factor Model of Interest Rates." *Mathematical Finance* 6 (4): 379–406.

English, W. B., W. R. Nelson, and B. P. Sack. 2003. "Interpreting the Significance of the Lagged Interest Rate in Estimated Monetary Policy Rules." *Contributions to Macroeconomics* 3 (1): 1–16.

Estrella, A., and J. C. Fuhrer. 2002. "Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational-Expectations Models." *American Economic Review* 92 (4): 1013–28.

Estrella, A., and F. S. Mishkin. 1999. "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty." In *Monetary Policy Rules*, ed. J. B. Taylor, 405–30. Chicago: Chicago University Press.

Ferguson, R. W., Jr. 2004. "Equilibrium Real Interest Rate: Theory and Application." Remarks to the University of Connecticut School of Business.

Fuhrer, J. C. 2000. "Habit Formation in Consumption and Its Implications for Monetary-Policy Models." *American Economic Review* 90 (3): 367–90.

Fuhrer, J. C., and G. D. Rudebusch. 2004. "Estimating the Euler Equation for Output." *Journal of Monetary Economics* 51: 1133–53.

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., and G. Schnabel. 2000. "The Taylor Rule and Interest Rates in the EMU Area." *Economics Letters* 67 (2): 165–71.

Gerlach-Kristen, Petra. 2004a. "Interest-Rate Reaction Functions and the Taylor Rule in the Euro Area." Working Paper, Swiss National Bank (July).

———. 2004b. "Interest-Rate Smoothing: Monetary Policy Inertia or Unobserved Variables?" *Contributions to Macroeconomics* 4 (1): 1–17.

Goodfriend, M. 1985. "Reinterpreting Money Demand Regressions." *Carnegie-Rochester Conference Series on Public Policy* 22:207–42.

———. 1991. "Interest Rates and the Conduct of Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy* 34:7–30.

Goodhart, C. 1997. "Why Do the Monetary Authorities Smooth Interest Rates?" In *European Monetary Policy*, ed. S. Collignon, 119–78. Washington, DC: Pinter.

———. 1999. "Central Bankers and Uncertainty." *Bank of England Quarterly Bulletin* 39:102–15.

———. 2005. "The Monetary Policy Committee's Reaction Function: An Exercise in Estimation." *Topics in Macroeconomics* 5 (1): 1–40.

Griliches, Z. 1967. "Distributed Lags: A Survey." *Econometrica* 35:16–49.

Gürkaynak, R. S., B. Sack, and E. Swanson. 2005a. "Did Financial Markets Think Disinflation Was Opportunistic?" Manuscript, Federal Reserve Board.

———. 2005b. "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1): 55–93.

Hall, A., and R. J. Rosanna. 1991. "Estimating the Speed of Adjustment in Partial Adjustment Models." *Journal of Business and Economic Statistics* 9 (4): 441–53.

Kara, A. H. 2003. "Optimal Monetary Policy, Commitment, and Imperfect Credibility." Discussion Paper No. 61, Central Bank of Turkey.

King, M. 2006. Speech at the Lord Mayor's banquet for bankers and merchants, London, England, June 21.

Kozicki, S. 1999. "How Useful Are Taylor Rules for Monetary Policy?" *Economic Review* (Federal Reserve Bank of Kansas City) (Second Quarter):5–33.

Lansing, K. 2002. "Real-Time Estimation of Trend Output and the Illusion of Interest Rate Smoothing." In *Economic Review*, 17–34. San Francisco: Federal Reserve Bank of San Francisco.

Levin, A., V. Wieland, and J. C. Williams. 1999. "Robustness of Simple Monetary Policy Rules under Model Uncertainty." In *Monetary Policy Rules*, ed. J. B. Taylor, 263–99. Chicago: Chicago University Press.

Lowe, P., and L. Ellis. 1997. "The Smoothing of Official Interest Rates." In *Monetary Policy and Inflation Targeting: Proceedings of a Conference*, ed. P. Lowe, 286–312. Reserve Bank of Australia.

Mankiw, N. G., and J. A. Miron. 1986. "The Changing Behavior of the Term Structure of Interest Rates." *Quarterly Journal of Economics* 101 (2): 211–28.

McGough, B., G. D. Rudebusch, and J. C. Williams. 2005. "Using a Long-Term Interest Rate as the Monetary Policy Instrument." *Journal of Monetary Economics* 52 (5): 855–79.

McManus, D. A., J. C. Nankervis, and N. E. Savin. 1994. "Multiple Optima and Asymptotic Approximations in the Partial Adjustment Model." *Journal of Econometrics* 62 (2): 91–128.

Mehra, Y. 2002. "The Taylor Principle, Interest Rate Smoothing and Fed Policy in the 1970s and 1980s." Working Paper No. 02-03, Federal Reserve Bank of Richmond (August).

Meyer, L. 1999. "The Global Economic Outlook and Challenges Facing Monetary Policy around the World." Remarks at the annual dinner of the Society of Business Economists, London, England, February 25.

Orphanides, A. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91 (4): 964–84.

———. 2003. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics* 50 (3): 605–31.

Peersman, G., and F. Smets. 1999. "The Taylor Rule: A Useful Monetary Policy Benchmark for the Euro Area?" *International Finance* 2 (1): 85–116.

Pivetta, F., and R. Reis. 2006. "The Persistence of Inflation in the United States." Forthcoming in *Journal of Economic Dynamics and Control.*

Poole, W. 2003. "Fed Transparency: How, Not Whether." *Review* (Federal Reserve Bank of St. Louis) 85 (6): 1–8.

Rotemberg, J., and M. Woodford. 1999. "Interest Rate Rules in an Estimated Sticky Price Model." In *Monetary Policy Rules*, ed. J. B. Taylor, 57–119. Chicago: Chicago University Press.

Rudebusch, G. D. 1992. "Trends and Random Walks in Macroeconomic Time Series: A Re-Examination." *International Economic Review* 33 (3): 661–80.

———. 1995. "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure." *Journal of Monetary Economics* 35 (2): 245–74.

———. 1998. "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review* 39 (4): 907–31.

———. 2001. "Is the Fed Too Timid? Monetary Policy in an Uncertain World." *Review of Economics and Statistics* 83 (2): 203–17.

———. 2002a. "Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty." *Economic Journal* 112 (479): 402–32.

———. 2002b. "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia." *Journal of Monetary Economics* 49 (6): 1161–87.

———. 2005. "Assessing the Lucas Critique in Monetary Policy Models." *Journal of Money, Credit, and Banking* 37 (2): 245–72.

Rudebusch, G. D., and L. E. O. Svensson. 1999. "Policy Rules for Inflation Targeting." In *Monetary Policy Rules*, ed. J. B. Taylor, 203–46. Chicago: Chicago University Press.

———. 2002. "Eurosystem Monetary Targeting: Lessons from U.S. Data." *European Economic Review* 46 (3): 417–42.

Rudebusch, G. D., E. T. Swanson, and T. Wu. 2006. "The Bond Yield 'Conundrum' from a Macro-Finance Perspective." Forthcoming in *Monetary and Economic Studies.*

Rudebusch, G. D., and J. C. Williams. 2006. "Revealing the Secrets of the Temple: The Value of Publishing Interest Rate Projections." Manuscript, forthcoming in *Monetary Policy and Asset Prices*, ed. J. Campbell.

Rudebusch, G. D., and T. Wu. 2006. "Accounting for a Shift in Term Structure Behavior with No-Arbitrage and Macro-Finance Models." Forthcoming in *Journal of Money, Credit, and Banking.*

Sack, B. 2000. "Does the Fed Act Gradually? A VAR Analysis." *Journal of Monetary Economics* 46 (1): 229–56.

Sack, B., and V. Wieland. 2000. "Interest Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence." *Journal of Economics and Business* 52 (1–2): 205–28.

Salmon, C., and B. Martin. 1999. "Should Uncertain Monetary Policymakers Do Less?" In *Monetary Policy Under Uncertainty*, ed. B. Hunt and A. Orr, 226–56. Wellington, New Zealand: Reserve Bank of New Zealand.

Sauer, S., and J.-E. Sturm. 2003. "Using Taylor Rules to Understand ECB Monetary Policy." CESifo Working Paper No. 1110 (December).

Shuetrim, G., and C. Thompson. 1999. "The Implications of Uncertainty for Monetary Policy." In *Monetary Policy Under Uncertainty*, ed. B. Hunt and A. Orr, 259–92. Wellington, New Zealand: Reserve Bank of New Zealand.

Söderlind, P., U. Söderström, and A. Vredin. 2005. "Dynamic Taylor Rules and the Predictability of Interest Rates." *Macroeconomic Dynamics* 9 (3): 412–28.

Söderström, U. 1999. "Should Central Banks Be More Aggressive?" Manuscript, Sveriges Riksbank (May).

Svensson, L. E. O. 2003. "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules." *Journal of Economic Literature* 41 (2): 426–77.

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

———. 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, ed. J. B. Taylor, 319–41. Chicago: Chicago University Press.

Tetlow, R. J., and P. von zur Muehlen. 2001. "Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always

Call for Attenuated Policy?" *Journal of Economic Dynamics and Control* 25:911–49.

Trehan, B., and T. Wu. 2006. "Time-Varying Equilibrium Real Rates and Monetary Policy Analysis." Forthcoming in *Journal of Economic Dynamics and Control.*

Walsh, C. E. 2003. *Monetary Theory and Policy.* 2nd ed. Cambridge, MA: MIT Press.

Waud, R. N. 1968. "Misspecification in the 'Partial Adjustment' and 'Adaptive Expectations' Models." *International Economic Review* 9 (2): 204–17.

Woodford, M. 1999. "Optimal Monetary Policy Inertia." *The Manchester School* 67 (s1): 1–35.

———. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press.

Yellen, J. L. 2004. "Innovations and Issues in Monetary Policy: Panel Discussion." *American Economic Review* 94 (2): 45–7.

Expectations, Learning, and Discretionary Policymaking*

Christian Jensen
Southern Methodist University

When policy forecasts are based on the policymaker's present and past actions, current policy affects expectations of future policy, contrary to what happens when forecasters can replicate policymaking perfectly. We show that when forecasts are generated through any linear combination of present and past policy functions that produces expectations consistent with the implemented policy, the optimal discretionary policy exploiting learning converges toward the optimal commitment plan as we approach a situation where people do not discount the future. Since influencing expectations permits improving policy, successful policymakers need to know how policy expectations are formed and how they can affect these expectations.

JEL Codes: E52, E58, E61.

1. Introduction

In recent years, central banks appear to have been putting more effort into influencing people's expectations, not only by announcing explicit inflation targets, but also by incorporating predictions of future inflation and output, and in some cases even of future policy,

*Thanks for helpful comments are due to Takushi Kurozumi, Bennett McCallum, anonymous referees, and participants at the 2005 Missouri Economics Conference and the 2006 North American Econometric Society Summer Meetings. Author contact: Department of Economics, Box 750496, Southern Methodist University, Dallas, TX 75275; Tel: (214) 768-3736; e-mail: cjensen@alumni.cmu.edu.

in their policy announcements.[1] Furthermore, some of their statements indicate that they believe their actions affect the public's policy expectations and assert that they take this influence into account when deciding policy.[2] In the present paper, we assume that policy affects expectations due to the public's lack of knowledge about the policy decision process, in terms of the objectives and models policymakers use to choose policy, so that they cannot predict policy by putting themselves in the policymakers' shoes. Instead, their forecasts have to be based on present and past policy, which we assume is perfectly observable, and this way, present policy comes to affect expectations of future policy. Our objective is to show how the optimal discretionary policy can improve and become more similar to the optimal commitment plan, and even converge to it in the limiting case of no time discounting, when policymakers influence expectations through present policy. We do so in a setup where expectations are always consistent with the model and the implemented policy.

[1]See Bernanke and Mishkin (1997), McCallum (1996), and Svensson (1997, 1999) for discussions on inflation targeting. Some examples of announcements incorporating forecasts include the following: "Against this background, the Committee adopted a directive that was biased toward a possible firming of policy going forward. Committee members emphasized that such a directive did not signify a commitment to near-term action." (Federal Reserve System [FRS], October 5, 1999). "However, with inflation quite low and resource use slack, the Committee believes that policy accommodation can be maintained for a considerable period." (FRS, December 9, 2003). "The Committee judges that some further measured policy firming is likely to be needed to keep the risks to the attainment of both sustainable economic growth and price stability roughly in balance." (FRS, December 13, 2005). "While growth has been relatively modest so far, both external and domestic factors give reason to expect a strengthening of the recovery through 2004 and beyond." (European Central Bank [ECB], March 4, 2004). "All in all, our judgement remains that real economic growth will gradually improve over the period ahead." (ECB, June 2, 2005). "Euro area real GDP is projected to grow at rates of between 1.2% and 1.6% in 2005, and between 1.4% and 2.4% in 2006 and 2007.... The projections indicate average HICP inflation to lie between 2.1% and 2.3% in 2005, and between 1.6% and 2.6% in 2006.... For 2007 an average inflation rate of between 1.4% and 2.6% is projected." (ECB, December 1, 2005). "Looking ahead, given that our monetary policy continues to be accommodative, a progressive withdrawal of monetary accommodation will be warranted if our assumptions and baseline scenario are confirmed." (ECB, August 3, 2006). "The Executive Board's assessment is that the economic projections imply a sight deposit rate in the interval $1\frac{1}{4}$–$2\frac{1}{4}$% at the beginning of November 2004." (Norges Bank, July 1, 2004).

[2]"Furthermore, today's move by 50 basis points appeared to be the best way in which to avoid uncertainties regarding the future course of monetary policy." (ECB, November 4, 1999).

Without a credible commitment, the public expects policymakers to reconsider policy in every period and take whatever action is optimal from the perspective of that point in time. As a result, policymakers' present and past actions have no effect on the public's policy forecasts, since these are produced by going over the decision process that policymakers are expected to face in the future. However, predicting policy this way requires that forecasters be able to reproduce policymakers' decision process, which for monetary policy is difficult, if not impossible. In most countries, monetary policy is decided by a group of individuals relying on different models and beliefs about the workings of the economy and its contemporary state. In addition, there might not be agreement about the policy objectives. So how do people actually predict monetary policy? Instead of reconstructing policymakers' decision making, present and past policy is used to forecast future policy. Since the policy problem does not change much from one period to the next, recent policy is used as an indicator of the policy to be implemented in the future. An example from the literature is the Taylor rule, which is probably the most successful description of recent monetary policy in the United States.[3] While the rule is based on knowledge of monetary policymaking, it (probably) does not reproduce policymakers' decision process. Instead, it uses past policy to obtain a description of the behavior of the policy instrument. However, when individuals' policy expectations depend on policymakers' present and past actions, the discretionary solution changes, because, contrary to what is assumed under standard discretion, policy affects expectations.

Many authors question whether individuals are, in practice, able to reproduce policymaking to the extent required for this to be useful in forecasting. For instance, in the context of monetary policy, the learning literature—exemplified by Bullard and Mitra (2002), Carlstrom and Fuerst (2004), and Evans and Honkapohja (2001, 2003)—studies how equilibrium outcomes are affected when individuals learn policy, or equilibria, by applying least-squares methods to data from the past and whether these equilibria are stable and learnable. While these contributions assume that policy is exogenous and, therefore, unaffected by the learning process, later work—such

[3]See Judd and Rudebusch (1998), Orphanides (2003), and Taylor (1993, 1999).

as that by Gaspar, Smets, and Vestin (2005), Molnar and Santoro (2005), and Orphanides and Williams (2005)—not only lets policy adjust optimally to the public's expectations, but, in addition, chooses policy so as to shape these expectations in an optimal manner. A common observation in these papers is that the optimal discretionary policy exploiting adaptive learning can temporarily attain features of the optimal commitment plan. We show that when present and past policy functions are used to forecast future policy, the optimal discretionary policy exploiting learning only deviates from the optimal commitment plan due to time discounting. The reason is that the effect present policy has on expectations through these forecasts, when they are consistent with the implemented policy, is the same as the effect preannouncing policy has under commitment, apart from influencing expectations in different periods.

In the learning literature cited in the previous paragraph, it is common to assume that individuals know everything about the monetary policy problem except for having to estimate parameter values from noisy observations of past realizations. Instead, we assume that people know nothing about the policymaking process but can observe the implemented policy functions exactly. This allows people to predict policy perfectly in our model, which guarantees that our results are due to the exploitation of expectations and not because the expectations are inconsistent with actual policy. At the same time, such forecasts maximize the effect policy has on expectations, thereby bringing the discretionary solution as close as possible to the optimal commitment plan. We do this to illustrate the potential gains from exploiting expectations in a situation where these expectations are consistent with actual policy and, so, do not require that policymakers mislead the public. The actual gains will depend on exactly how people forecast policy, which is an empirical question. By focusing on policy functions instead of parameter values, our setup is closer to the Taylor-rule example than what is usual in the learning literature, in the sense that people use past policy to forecast future policy without reconstructing the policymaking process.

Since the optimal discretionary policy depends on how the public forecasts policy, it is important for policymakers to know how these forecasts are formed and how they are affected by the policymakers' actions. This knowledge can be used to influence expectations and thereby better achieve the policy objectives, much in the same way as a credible commitment to a plan can. Hence, the public's unawareness of the policymaking process, and use of present and past policy to forecast policymakers' future actions, can improve the discretionary solution. Moreover, if policymakers can affect the way in which policy expectations are formed, they can also influence the optimal discretionary, or time-consistent, policy. By determining how complicated and public the policy decision process is, policymakers can make it easier or harder for outsiders to forecast policy by reproducing this process. The less information people have about policymaking, the less likely that they can use it to forecast policy, and the more likely that their forecasts will be based on present and past policy. Thus, by affecting the knowledge the public has about the policy decision process and, thereby, how the public forecasts policy, policymakers can influence the discretionary solution.

The next section introduces a standard sticky-price model commonly used for modeling monetary policymaking. In section 3, this model is used to illustrate how the discretionary solution changes when policy affects expectations. We show that when people's next-period policy expectations are determined by the policy they see being implemented today, the discretionary solution becomes almost the same as the commitment plan, and we show that the two policies are identical when the discount factor is close enough to 1. Next, we explore how this result is modified when policy forecasts look further back in time, and when they are inconsistent with actual policy. In section 4, we argue that the optimal policy in our setup will, in general, differ from the one in a reputational model, even though policy affects expectations in both cases. We conclude that it is important for policymakers to know how the public's policy expectations are formed and, in particular, how they can influence these expectations.

2. Model

We employ a New Keynesian sticky-price model with monopolistic competition like the one used by, for example, Clarida, Galí, and Gertler (1999), McCallum and Nelson (2000), Svensson and Woodford (2003), and Woodford (1999a, 1999b, 2003). It assumes that it is costly, or impossible, to change prices frequently, so that present price setting needs to take into account expected future price movements. Assuming that the central bank can influence the rate of inflation through monetary policy, current prices will depend on both present and expected future monetary policy, so that present inflation depends on expected future inflation. The monetary policy problem is to respond to random cost-push shocks so as to keep inflation and output as close as possible to their flexible-price values, thereby minimizing the welfare loss due to price stickiness. For simplicity, and in line with current practice in the literature, we assume that the policymaker determines the rate of inflation directly and that individuals decide how much to produce, instead of the prices.

Letting π_t denote the rate of inflation and y_t the log of output, both in terms of deviations from flexible-price values, Calvo (1983), Rotemberg (1982), and Rotemberg and Woodford (1999) show that the policy problem in any period $t = 0$ is to minimize[4]

$$E \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \omega y_t^2 \right) \tag{1}$$

subject to

$$\pi_t = \beta E_t \pi_{t+1} + \alpha y_t + u_t \tag{2}$$

$$u_t = \phi u_{t-1} + \varepsilon_t . \tag{3}$$

The shock ε_t is white noise with variance σ_ε^2 and is assumed to be observed by the policymaker before setting the inflation rate π_t. The parameters α, ω, ϕ, and σ_ε are all strictly positive, while $\beta \in (0, 1)$.[5]

As is shown by Clarida, Galí, and Gertler (1999), Currie and Levine (1993), Svensson and Woodford (2003), and Woodford

[4]We minimize the unconditional expected value instead of the conditional one to avoid optimal policies being dependent on initial conditions. The time-inconsistency problem is the same with the conditional and unconditional objective functions, so our conclusions apply for both cases.

[5]$\phi > 0$ is needed to generate persistence in the model. If $\phi = 0$, expectations of future inflation, $E_t \pi_{t+1}$, would always be zero for some of the policy equations we consider, and the dependency of current inflation on expected future inflation would in practice disappear. Persistence can alternatively be introduced by including lagged values of variables in state equation (2) or in the policy function.

(1999a), the optimal commitment policy is

$$\pi_0 = -\frac{\omega}{\alpha} y_0 \tag{4}$$

$$\pi_t = -\frac{\omega}{\alpha} y_t + \frac{\omega}{\alpha} y_{t-1}, \quad t \geq 1. \tag{5}$$

This policy minimizes the objective in (1), but, as Kydland and Prescott (1977) argue, it is not time consistent, because if we reoptimized in any later period $\tau > 0$, the policymaker would want to deviate from the commitment policy in (5) and instead commit to implementing

$$\pi_\tau = -\frac{\omega}{\alpha} y_\tau \tag{6}$$

$$\pi_t = -\frac{\omega}{\alpha} y_t + \frac{\omega}{\alpha} y_{t-1}, \quad t > \tau. \tag{7}$$

When the policymaker reoptimizes in every period, the implemented policy is always

$$\pi_t = -\frac{\omega}{\alpha} y_t, \tag{8}$$

which—as is shown by Clarida, Gali, and Gertler (1999), McCallum and Nelson (2000), Svensson and Woodford (2003), and Woodford (1999a)—is the optimal discretionary policy. Combining (8) with the state equations in (2) and (3), one can show that the discretionary solution can equivalently be written as

$$\pi_t = \frac{\omega}{\alpha^2 + \omega(1 - \beta\phi)} u_t \tag{9}$$

in terms of the exogenous cost-push shock u_t.

3. Expectations

To study the role of policy expectations in the time-consistent solution, it is useful to be more explicit about these expectations. Assume that policy is given by

$$\pi_t = h_t u_t, \tag{10}$$

where $\{h_t\}_{t=0}^{\infty}$ is a sequence of policy parameters to be determined, so that the policy problem is now to choose h_t, instead of π_t directly.[6] State equation (2) summarizes individuals' optimal behavior. The reason it has current inflation, or output, depend on expected next-period inflation is that when prices are sticky, an individual's optimal present action depends on his or her inflation expectations for the next period. Hence, it is the public's inflation expectations that matter for the present state. This implies that when policy is as in equation (10), the expectations of next-period inflation, $E_t\pi_{t+1}$, in state equation (2) satisfy

$$E_t\pi_{t+1} = E_t h_{t+1} u_{t+1} = h^e_{t,t+1}\phi u_t, \tag{11}$$

where $h^e_{t,t+1}$ denotes the policy that the public at time t expects will be implemented in period $t+1$.[7] The reduced-form solution of the model described by equations (2), (3), (10), and (11) is given by (10) and

$$y_t = \frac{h_t - \beta\phi h^e_{t,t+1} - 1}{\alpha} u_t, \tag{12}$$

where the values of the endogenous variables, π_t and y_t, depend on the policy that is implemented in the contemporaneous period, h_t, and the policy that in period t is expected to be implemented in the period after, $h^e_{t,t+1}$. Inserting into the objective function in (1), we have

$$E\sum_{t=0}^{\infty}\beta^t\left(h_t^2 + \frac{\omega}{\alpha^2}\left(h_t - \beta\phi h^e_{t,t+1} - 1\right)^2\right)u_t^2. \tag{13}$$

The value of the objective function in any period $t = 0$ depends on both the policymaker's future actions, $\{h_t\}_{t=1}^{\infty}$, and the public's policy expectations, $\{h^e_{t,t+1}\}_{t=0}^{\infty}$. We distinguish between the two, because the policymaker is able to predict policy by resolving the policy problem, while the public might not be able to do so. The

[6] Writing policy as a function of y_t instead of u_t might be more in accordance with what policymakers actually do, but it makes the computation of the optimal policy more arduous.

[7] From the perspective of period t, individuals' policy expectations for the period after, $h^e_{t,t+1}$, are nonstochastic.

optimal time-consistent policy is found by minimizing the objective in (13) with respect to h_0, assuming that all later policy actions, $\{h_t\}_{t=1}^{\infty}$, will be chosen the same way.

With standard discretion, the present policy action, h_0, has no effect on policy expectations, so

$$\frac{\partial h_{t,t+1}^e}{\partial h_0} = 0, \quad \forall t, \tag{14}$$

and the first-order condition from minimizing (13) with respect to h_0 yields

$$h_0 = \omega \frac{\beta \phi h_{0,1}^e + 1}{\alpha^2 + \omega}. \tag{15}$$

Since future policy, $\{h_t\}_{t=1}^{\infty}$, and future policy expectations, $\{h_{t,t+1}^e\}_{t=1}^{\infty}$, are independent of present policy, h_0, the optimal present policy action only depends on present expectations of next-period policy, $h_{0,1}^e$. Because future policy will be determined by resolving exactly the same optimization problem in the future, the optimal policy will be the same in every period. Under standard discretion the public is assumed to forecast policy by replicating the policymaker's decision process, so they will be able to predict perfectly, and therefore we have

$$h_{0,1}^e = h_1 = h_0 = h_t, \quad \forall t. \tag{16}$$

Inserting this into the first-order condition in (15) and solving for h_t, we find that the time-consistent solution is

$$h_t = \frac{\omega}{\alpha^2 + \omega(1 - \beta\phi)}, \tag{17}$$

which gives the standard discretionary policy in (9).

Imagine now that the present policy action, h_0, affects the public's present or future policy expectations so that condition (14) does not hold. Then, the optimal time-consistent policy will differ from the standard discretionary solution in (9). When policy affects future policy expectations, $\left\{h_{t,t+1}^e\right\}_{t=1}^{\infty}$, the optimal present policy action will not only depend on the effect it has on expectations, $\left\{\frac{\partial h_{t,t+1}^e}{\partial h_0}\right\}_{t=1}^{\infty}$, but also on the expectations themselves, $\left\{h_{t,t+1}^e\right\}_{t=1}^{\infty}$, and the policy the policymaker will implement in the future, $\{h_t\}_{t=1}^{\infty}$. While the present policy action has no direct effect on the policymaker's future actions, it can affect these actions through the public's expectations when the policymaker reoptimizes in future periods.

For a simple, but explicit, example of policy affecting expectations in a context where these are consistent with the model and actual policy, imagine that policy expectations are

$$h^e_{t,t+1} = h_t, \quad \forall t, \tag{18}$$

so that the policy that is implemented today is expected to be implemented in the next period too.[8] While these expectations are plain, they should be reasonable in a situation where people have little knowledge about the policy decision process but in which they can observe the implemented policy, as we assume here. The public's policy expectations are not static in (18); they take into account that the policymaker will reoptimize in every period. Present policy, h_0, determines present expectations of next-period policy, $h^e_{0,1}$, but has no effect on future expectations, since $h^e_{1,2} = h_1$, $h^e_{2,3} = h_2, \ldots$. Hence, the policymaker takes into account that if the present policy action determines expectations of next-period policy today, future policy actions will determine expectations in the future. With a policy function like (10), this implies, in the present model, that the policy implemented in any period t will only affect inflation and output in period t. Inserting the expectations in (18) into the objective function in (13) and minimizing with respect to h_0, we find that the optimal discretionary policy of form (10) is now

$$\pi_t = \frac{\omega(1-\beta\phi)}{\alpha^2 + \omega(1-\beta\phi)^2} u_t, \tag{19}$$

no matter what policy was implemented in the past or will be implemented in the future. As Clarida, Galí, and Gertler (1999) show, this policy is the optimal commitment rule of the form $\pi_t = hu_t$.[9] Since

[8] We are not assuming that people expect the same inflation rate, but that they expect the same policy equation.

[9] The effect of exploiting learning is in this case the same as if the weight on output stabilization ω was lowered to $\omega(1-\beta\phi)$ in the standard discretionary

the optimal policy is the same in every period, using present policy to predict next-period policy, as in (18), should arguably be a reasonable way to forecast in this setup, especially given the assumption that individuals cannot reproduce the policymaker's decision process. In fact, these forecasts are exactly correct in terms of the policy function, and in terms of the inflation rate, they only miss due to the unpredictable white noise shock ε_t, the same as when forecasts are generated by reproducing the policymaker's decision process.

When people expect next-period policy to be the same as the present policy, the optimal discretionary policy will, in general, not be of the form in (10). For $\alpha = .05$, $\beta = .99$, $\omega = .25$, and $\phi = .8$, the policy

$$\pi_t = 3.251 u_t + .414 y_{t-1} \tag{20}$$

is a stationary discretionary equilibrium that minimizes the objective in (1) with a policy function of the same form as the optimal commitment policy,[10]

$$\pi_0 = 4.587 u_0 \tag{21}$$

$$\pi_t = 3.247 u_t + .455 y_{t-1}, \quad t \geq 1. \tag{22}$$

With a policy function of the form in (20), the policy implemented in period t does not only affect inflation and output in period t, since the lagged output term, y_{t-1}, in the policy function links periods together. Consequently, the optimal present policy action depends on what policy was implemented in the past. By a stationary equilibrium, we mean that the reported policy is optimal given that the same policy was implemented in the past. In other words, the reported policy is the one we would converge toward as the effect of policy implemented prior to the initial period, $t = 0$, dissipates.

Comparing equations (20) and (22), we see that when the present policy is expected to be implemented in the next period too, the

solution, and due to Rogoff's (1985) conservative central banker result, we know that lowering the weight on output stabilization can increase welfare.

[10]The commitment policy in (21)–(22) is equivalent to the one in (4)–(5). The policy in (20) can be computed algebraically, but the general solution is too extensive to reproduce here.

discretionary solution becomes very similar to the optimal commitment policy.[11] The reason is that the influence policymakers have on expectations with such forecasts is the same as the influence they have on expectations under commitment. The optimal commitment plan takes into account that the policy action to be implemented in any particular period will determine people's policy expectations for that period, since $h^e_{t,t+1} = h_{t+1}$ for $t = 1, 2, \ldots$, using the notation from above. When people forecast the next-period policy action to be the same as the present one, the action to be implemented in a period will determine people's expectations for the period immediately afterward, and $h^e_{t,t+1} = h_t$ for $t = 0, 1, \ldots$. Apart from the timing, the effects on expectations are the same. When the discount factor β is close enough to 1, periods are weighted almost the same, and the difference in timing becomes irrelevant, making the optimal policies under commitment and discretion the same. In fact, one can prove that as β converges toward 1, the optimal stationary discretionary policy and the optimal commitment policy converge, no matter what the values of α, ω, and ϕ (see the appendix).

Likewise, when expectations are determined by the policy action implemented $J > 0$ periods back in time, so that

$$h^e_{t,t+1} = h_{t-J}, \tag{23}$$

discretion will match commitment only if β^{J+1} is close to 1. The effect on policy expectations is the same, no matter what the value of J (including for commitment, which corresponds to the case with $J = -1$), but the larger J is, the more these effects are discounted away and ignored when determining the optimal discretionary policy. From this it follows that when expectations are a weighted average of past policy actions,

$$h^e_{t,t+1} = \sum_{j=0}^{J} \theta_j h_{t-j}, \tag{24}$$

[11] The parameters have been chosen somewhat arbitrarily from previous work. The key variable is the discount factor β, since, as shown below, the discretionary solution converges toward the optimal commitment plan as β approaches 1. A discount factor of .99 is fairly standard in the literature for a quarterly model.

where θ_j are arbitrary weights, the discretionary solution will be close to the optimal commitment rule only if

$$\sum_{j=0}^{J} \beta^{j+1}\theta_j \tag{25}$$

is close to 1.[12] For the expectations in (24) to be consistent with actual policy,

$$\sum_{j=0}^{J} \theta_j = 1 \tag{26}$$

needs to be satisfied, so the larger J is, the less similar discretion will be to commitment, given that $\beta \in (0, 1)$. Still, whatever effect policy has on expectations, it will be optimal to exploit this influence to improve policymaking.

It is much easier to match commitment if expectations are not required to be consistent with actual policy, and one can even outdo it when policy has no influence on expectations. Using state equation (2) to substitute for y_t in the standard discretionary policy in (8), we find that the optimal inflation rate under discretion when policy has no effect on expectations is

$$\pi_t = \frac{\omega}{\alpha^2 + \omega}(\beta E_t \pi_{t+1} + u_t), \tag{27}$$

while the associated output level is

$$y_t = -\frac{\alpha}{\alpha^2 + \omega}(\beta E_t \pi_{t+1} + u_t), \tag{28}$$

both in terms of the public's inflation expectations, $E_t \pi_{t+1}$.[13] By making these expectations equal $-\frac{1}{\beta}u_t$, one can make both inflation and output be zero at all times, which is much better than what can be achieved under commitment. Such expectations would, however, not be consistent with the implemented policy, defined in (27). To

[12] This can be verified numerically, or even algebraically for a small J.

[13] One can show that this is the optimal policy by substituting state equation (2) into the objective in (1) and minimizing with respect to π_t and $E_t \pi_{t+1}$ as two separate and independent variables.

guarantee that the gains from exploiting learning that we derive in the present study are not the result of inconsistent expectations or fooling the public, we assume throughout that policy is perfectly observable, which in our framework implies that policy forecasts are dead-on.

The learning literature is generally more optimistic than we have been above in terms of the knowledge the public is assumed to have about the policy decision process. In our setup, people are assumed to know nothing about this process, but, instead, we imagine that they can observe the implemented policy perfectly, which is also too optimistic. When policy is observed with error, after a lag, or when policymakers cannot exactly control what policy is implemented, policymakers' influence on expectations could be weakened, thus reducing the possible gains from its exploitation. Of course, such assumptions could also make it harder for forecasters to weed out expectations that are inconsistent with actual policy, facilitating a situation in which policymakers mislead the public.

Policy can also affect expectations when individuals are able to reproduce the policymaking process, as long as they cannot do so perfectly. An example is provided by Gaspar, Smets, and Vestin (2005), where forecasters are assumed to know everything about the model used to determine policy except the parameter values. Since the data that are used to estimate these parameters are observed with error, it can be difficult to distinguish policymakers' preferences, and parameter values in general, from strategic manipulation of the data, and policymakers should exploit this to influence the public's expectations. However, once policy forecasts are exactly correct because the true parameter values have been deduced, present policy has no effect on estimates and forecasts, and the optimal discretionary solution falls back to the standard one where expectations are independent of policy. In contrast, our forecasting scheme allows policy to affect expectations also when these are exactly correct, even if the magnitude of these effects depends on how far back forecasters look and how heavily people discount the future.

4. Discussion

We assume that there is no way for policymakers to credibly commit to implementing a policy. Barro (1986) and Rogoff (1987) argue that policymakers that have a reputation for implementing the policy they preannounce could achieve credible commitment.[14] As in our framework, policy affects expectations in their setup but in a different way. In a reputational model, the policymaker's past actions determine the credibility of the commitment and only influence policy expectations by affecting how likely the public believes it is that the announced policy will actually be implemented. If the policymaker's credibility is perfect, policy expectations will be determined solely by the announced policy. If the policymaker has no credibility, people will expect the standard discretionary policy. In our setup, there are no announcements, but policy affects expectations because people use it to learn what policy is, and will be, implemented. In one case, people learn whether they should trust the policymaker's announcements; in the other, they learn what policy is being implemented and never trust the policymaker's announcements. In a reputational model, the standard discretionary policy is not implemented because it would ruin the policymaker's good reputation and ability to credibly commit. In our setup, the standard discretionary policy is not implemented because when present policy is used to learn what policy is—and will be—employed, the optimal present action is not standard discretion.

Evans and Honkapohja (2006) and McCallum (1995) suggest that central banks should implement the optimal commitment policy even if they have no way of credibly committing to it, not even a reputation. This recommendation must be based on a belief that, sooner or later, people would learn that the standard discretionary policy is not being implemented and that they would therefore stop expecting this policy, since otherwise it would be optimal to implement ordinary discretion, not commitment. However, if this learning takes place, it must be because policy affects people's expectations of future policy, and, in this case, the optimal policy depends on the exact effect it has on expectations. If the effect is large enough, the optimal policy can be similar to the optimal commitment policy, but it is small; or if people discount the future too heavily, the optimal policy can be closer to the standard discretionary solution, and

[14]See Barro (1986), McCallum (1990), and Rogoff (1987) for discussions on why a reputation might be insufficient for the policymaker to achieve the desired commitment.

implementing the optimal commitment plan can give worse results in terms of achieving the policy objectives than standard discretion.[15]

5. Conclusions

The optimal time-consistent policy depends on the public's policy forecasts. We know how people should forecast policy assuming they are rational and have perfect and complete information; they should replicate the policymaker's decision process. When forecasts are formed this way, present and past policy is irrelevant for expected future policy, as is assumed in the standard discretionary solution. However, if forecasters do not have all the information required to reproduce the policymaker's decision process, they have to use present and past policy to learn about this process and predict policy. When they do so, time-consistent policy can reap the benefits of influencing expectations, just as in the commitment solution, to a degree determined by the exact learning/forecasting process and the level of time discounting in the economy. It is therefore important for policymakers to know how the public predicts policy and how their actions affect these predictions. This way, policymakers can influence policy expectations even if they are not able to commit and credibly preannounce their actions. If, in addition, policymakers can influence the way in which individuals forecast policy—for example, by providing more or less information about the policy decision process—they can affect the time-consistent solution and thereby how well discretionary policy achieves their objectives.

Appendix

This appendix outlines the proof of the proposition that the optimal (stationary) discretionary policy and the optimal commitment

[15] This discussion also applies to timeless-perspective commitment policies like the ones proposed by Svensson and Woodford (2003), Woodford (1999a, 2003), or Jensen and McCallum (2002), where the policymaker commits to a procedure for determining policy instead of a specific equation. The proposed timeless-perspective policies are only optimal if the policymaker can credibly commit and thereby dictate individuals' policy expectations. If policy, or commitment, is learned over time, the optimal timeless-perspective policies can differ from the ones proposed in the references above.

policy converge as β approaches 1, when expectations are such that the policy action implemented in the present period is expected to be implemented in the next period too. With a policy function

$$\pi_t = f_t u_t + g_t y_{t-1} \tag{29}$$

and expectations

$$f^e_{t,t+1} = f_t, \quad \forall t \tag{30}$$

$$g^e_{t,t+1} = g_t, \quad \forall t, \tag{31}$$

the reduced form of the model is given by (29) and

$$y_t = \frac{(1-\beta\phi)f_t - 1}{\alpha + \beta g_t} u_t + \frac{g_t}{\alpha + \beta g_t} y_{t-1}, \tag{32}$$

which we can use to write the policy objective in (1) as

$$E \sum_{t=0}^{\infty} \beta^t \left((f_t u_t + g_t y_{t-1})^2 + \omega \left(\frac{(1-\beta\phi)f_t - 1}{\alpha + \beta g_t} u_t + \frac{g_t}{\alpha + \beta g_t} y_{t-1} \right)^2 \right). \tag{33}$$

Minimizing with respect to f_0 and g_0, and then assuming stationarity ($\cdots = f_{-1} = f_0 = f_1 = \cdots$ and $\cdots = g_{-1} = g_0 = g_1 = \cdots$) and letting β converge toward 1, we find that the optimal discretionary policy converges toward

$$\pi_t = \lambda u_t + \rho y_{t-1}, \tag{34}$$

where

$$\lambda = \frac{\rho(2\omega(1-\phi^2) - \alpha^2\phi^2) + \alpha\omega(1+\phi^2)}{\rho(1+\phi)(\alpha^2\phi + (\alpha^2 + 2\omega)(\phi - 1)^2) + \alpha((3 - \phi - \phi^2 - \phi^3)\omega + \alpha^2)} \tag{35}$$

$$\rho = -\frac{1}{2}\alpha + \frac{1}{2}\sqrt{\alpha^2 + 4\omega}. \tag{36}$$

Combining equations (2), (3), and (5), one can show that when β converges toward 1, the optimal commitment policy converges toward

$$\pi_t = \frac{\omega}{\alpha^2 + \omega(1 - (r - 1 + \phi))} u_t + \frac{\omega}{\alpha}(1 - r) y_{t-1}, \tag{37}$$

where

$$r = \frac{\alpha^2 + 2\omega - \alpha\sqrt{\alpha^2 + 4\omega}}{2\omega}. \tag{38}$$

Showing that the discretionary policy given in (34)–(36) and the commitment policy given by (37)–(38) are identical is mere algebra.

References

Barro, R. J. 1986. "Reputation in a Model of Monetary Policy with Incomplete Information." *Journal of Monetary Economics* 17 (1): 3–20.

Bernanke, B. S., and F. S. Mishkin. 1997. "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives* 11 (2): 97–116.

Bullard, J., and K. Mitra. 2002. "Learning about Monetary Policy Rules." *Journal of Monetary Economics* 49 (6): 1105–29.

Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Carlstrom, C. T., and T. S. Fuerst. 2004. "Learning and the Central Bank." *Journal of Monetary Economics* 51 (2): 327–38.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Currie, D. A., and P. Levine. 1993. *Rules, Reputation and Macroeconomic Policy Coordination.* Cambridge University Press.

Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton, NJ: Princeton University Press.

———. 2003. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies* 70 (4): 807–24.

———. 2006. "Monetary Policy, Expectations and Commitment." *Scandinavian Journal of Economics* 108 (1): 1–182.

Gaspar, V., F. Smets, and D. Vestin. 2005. "Optimal Monetary Policy under Adaptive Learning." ECB Working Paper No. 644.

Jensen, C., and B. T. McCallum. 2002. "The Non-Optimality of Proposed Monetary Policy Rules Under Timeless-Perspective Commitment." *Economics Letters* 77 (2): 163–8.

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

Kydland, F. E., and E. C. Prescott. 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85:473–92.

McCallum, B. T. 1990. "Inflation: Theory and Evidence." In *Handbook of Monetary Economics*, Vol. 2, ed. B. M. Friedman and F. H. Hahn. North Holland.

———. 1995. "Two Fallacies Concerning Central-Bank Independence." *American Economic Review* 85 (2): 207–11.

———. 1996. "Inflation Targeting in Canada, New Zealand, Sweden, the United Kingdom, and in General." NBER Working Paper No. 5579.

McCallum, B. T., and E. Nelson. 2000. "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models." NBER Working Paper No. 7915.

Molnar, K., and S. Santoro. 2005. "Optimal Monetary Policy when Agents are Learning." Working Paper, Universitat Pompeu Fabra.

Orphanides, A. 2003. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50 (5): 983–1022.

Orphanides, A., and J. Williams. 2005. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation-Targeting Debate*, ed. B. Bernanke and M. Woodford. Chicago: University of Chicago Press.

Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

———. 1987. "Reputational Constraints on Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy* 26:141–82.

Rotemberg, J. J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (October): 517–31.

Rotemberg, J. J., and M. Woodford. 1999. "Interest Rate Rules in an Estimated Sticky Price Model." In *Monetary Policy Rules*, ed. J. B. Taylor. Chicago: University of Chicago Press.

Svensson, L. E. O. 1997. "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts." *American Economic Review* 87:98–114.

———. 1999. "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics* 43 (3): 607–54.

Svensson, L. E. O., and M. Woodford. 2003. "Indicator Variables for Optimal Policy." *Journal of Monetary Economics* 50 (3): 691–720.

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

———. 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, ed. J. B. Taylor, 319–41. Chicago: University of Chicago Press.

Woodford, M. 1999a. "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" In *New Challenges for Monetary Policy*, 277–316. Kansas City, MO: Federal Reserve Bank of Kansas City.

———. 1999b. "Optimal Monetary Policy Inertia." NBER Working Paper No. 7261.

———. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.

The Impact of Monetary Policy on the Exchange Rate: A Study Using Intraday Data*

Jonathan Kearns and Phil Manners
Reserve Bank of Australia

We investigate the impact of monetary policy on the exchange rate using an event study with intraday data for four countries. Carefully selecting the sample periods ensures that the policy change is exogenous to the exchange rate. An unanticipated tightening of 25 basis points leads to a rapid appreciation of around 0.35 percent. We also show that the impact depends on how the surprise affects expectations of future monetary policy. If expectations of future policy are revised by the full amount of the surprise, then the impact on the exchange rate is larger (0.4 percent) than if the surprise only brings forward an anticipated change in policy (0.2 percent).

JEL Codes: F31, G14.

1. Introduction

Recent studies have had some success in identifying the response of the exchange rate to macroeconomic variables by using high-frequency data.[1] This paper makes two important contributions to

*The authors are grateful to the editor and an anonymous referee, participants of a seminar at the Reserve Bank of Australia and the Monetary Shocks workshop at the University of Melbourne, and Ellis Connolly, Tim Hampton, and Christopher Kent for helpful comments. Responsibility for any errors rests with the authors. The views expressed in this paper are those of the authors and should not be attributed to the Reserve Bank. Corresponding author: Kearns: Reserve Bank of Australia, GPO Box 3947, Sydney NSW Australia; e-mail: kearnsj@rba.gov.au; Tel: +61-2-9551-8877; Fax: +61-2-9551-8833.

[1]For example, see Andersen et al. (2003), Zettelmeyer (2004), Faust et al. (forthcoming), and the references therein.

the literature on the response of the exchange rate to monetary policy. First, we use intraday data, which allows us to more precisely control for endogeneity and external factors that may influence both exchange rates and interest rates (such as macroeconomic data releases). With intraday data we can also examine the temporal response of the exchange rate. Our second contribution is to consider how changes in the expected path of future monetary policy that result from a monetary surprise influence the response of the exchange rate. Some interest rate changes may surprise with respect to the timing of the change; for example, the rate rise expected next month might occur this month. Others may surprise with respect to the expected path of monetary policy; for example, a surprise rate rise might be taken to indicate that a tightening phase is going to reach a higher maximum than previously anticipated. Because these surprises will have different effects on the expected future path of monetary policy, they are unlikely to have equivalent effects on the exchange rate.

A greater understanding of the impact of interest rates on exchange rates is of interest for several reasons. The theory of uncovered interest parity (UIP), which connects expected changes in the exchange rate to interest differentials, is central to almost all international macroeconomic models. Yet empirically, UIP is a resounding failure (Engel 1996). In addition, the response of the exchange rate to monetary policy is also an important monetary transmission channel in small, open economies (see, for example, Grenville 1995 and Thiessen 1995).

Our study includes four countries (Australia, Canada, New Zealand, and the United Kingdom) that are relatively small, and so changes in their interest rates are unlikely to affect global interest rates. This is important for isolating the impact of the change in one country's interest rate on the exchange rate. If the country studied was large, such as the United States, then markets might build the likely impact of changes in domestic monetary policy on foreign interest rates into the exchange rate's response. This would contaminate the measured response of the exchange rate. Further, these four countries have highly liquid financial markets, freely floating exchange rates, and similar monetary policy regimes.

We use an event-study methodology as has become common in the literature on asset prices. An event study is particularly useful because it can abstract from the joint determination of interest rates and exchange rates. The event is a monetary policy decision (either a surprise change in the policy interest rate or no change when a policy announcement was anticipated). We can be confident that we have isolated events in which causality runs in only one direction, from interest rates to exchange rates, for two reasons. Firstly, we use a narrow event window, only examining a short period around the policy change. Secondly, for the countries we study, the institutional structure of monetary policy decision making means that the decision is made well before the event window we use.

Several papers have recently used high-frequency data to examine the response of asset prices to macroeconomic shocks, including interest rates. This paper most closely follows that of Zettelmeyer (2004), who examines the response of exchange rates to interest rates using daily data (but not intraday). Unlike Zettelmeyer, we restrict our sample to a period in which the central banks we study did not explicitly respond to the exchange rate. We also use a more-accurate measure of the monetary surprise (based on one-month rather than three-month interest rates) and a larger sample, in part because we include decisions in which monetary policy does not change—that is, "no-change" surprises. We can include these observations because, under the monetary policy regimes we examine, the timing of the announcement of these no-changes was predetermined. Faust et al. (forthcoming) use intraday data to examine the response of exchange rates to macroeconomic announcements, including interest rate changes. But they only study surprises in U.S. interest rates, and so the exchange rate responses are potentially clouded by anticipated changes in foreign interest rates. Andersen et al. (2003) also examine the intraday response of the exchange rate to macroeconomic announcements but do not consider interest rate shocks. Bernanke and Kuttner (2005), studying the response of equity markets to interest rates using daily data, consider how the impact differs depending on the changes to the profile of anticipated future monetary policy, as we do in this study. A related literature has attempted to consider the longer-run impact of interest rates on the exchange rate. In an early study using a vector autoregression (VAR), Eichenbaum and Evans (1995) suggest that there exists a delayed overshooting. But by identifying surprise interest rate shocks using daily data, Faust et al. (2003) find that this result is not robust to allowing the foreign interest rate to respond. Faust and Rogers (2003) also fail to find evidence of delayed overshooting in a less-restricted VAR.

The remainder of the paper is structured as follows. Section 2.1 briefly outlines the application of the event-study methodology to monetary policy decisions. In section 2.2 we review the monetary policy operations of the four countries in the study and discuss how they influence the set of events that we consider. The data are described in more detail in section 2.3. In section 3.1 we present the results of the instantaneous impact and the timing of the response of the exchange rate. In section 3.2 we demonstrate how the response of the exchange rate depends on the effect of the monetary surprise on expectations. We examine the robustness of the results in section 3.3. Section 4 concludes.

2. The Estimation Framework

2.1 Event-Study Methodology

We use an event-study approach, estimating the change in the exchange rate around the announcement of "monetary policy decisions." Decisions include both announced changes to monetary policy and announcements of decisions to not change policy ("no-change" decisions), so long as the market knew for certain that a policy announcement would take place.[2] Further discussion of the events used can be found in section 2.2. In many cases, monetary policy decisions are widely anticipated by the market, and so their impact should already be incorporated into interest rates and exchange rates. In order to identify the impact of a monetary policy decision, we isolate the surprise component of the change in monetary policy by using changes in market interest rates rather than the change in the policy interest rate.[3] This technique, developed in

[2]On the event days in which there is no change in policy, markets may have given some probability to there being a change in policy, and so there may well have still been a surprise that contained news.

[3]This does not mean that anticipated changes in monetary policy have no effect on the exchange rate, but that the effect has been incorporated into the exchange rate at the same time as markets came to the conclusion that there would be a change in monetary policy.

Kuttner (2001), is commonly used in the literature. Market interest rates incorporate a risk premium, but the change in the market interest rate is a good proxy for the policy surprise, as the risk premium is unlikely to move in the short time periods used in the event study (Piazzesi and Swanson 2004).

For each of the events, we measure the movement in the exchange rate around the event using intraday data. We use a short, seventy-minute event window. This reduces the amount of information received by the market in the event window, reducing the number of events that would have to be discarded due to the exchange rate and interest rate jointly responding to other news, such as a macroeconomic data release. Because the interest rate surprise will be a more-dominant piece of information in a short event window, it should also result in more-accurate estimates.

One potential source of concern is that exchange rate movements could influence monetary policy decisions. However, this is not likely in this study because in each central bank, the main deliberation on policy changes occurs the day before the announcement. Given that at a daily frequency the exchange rate is typically considered to be a random walk, the event window will not contain exchange rate movements that influenced the policy decision.[4] The daily market interest rates may be affected by other events or an endogenous response to exchange rate movements, but this is minimized by the fact that events occur on days for which monetary policy is likely to be the most important shock to interest rates. Unfortunately, intraday interest rate data are not available for our sample of countries and time.[5] The events that are excluded from our sample for reasons of contamination are outlined in section 2.2.

[4]There is some evidence of weak serial correlation in exchange rates, which is often found using nonlinear models; for example, see Gencay et al. (2002). However, such serial correlation is so weak as to not be relevant from a policy perspective.

[5]We investigated using intraday exchange rate forwards to derive a measure of the intraday interest rate shock based on covered interest rate parity. But quotes for exchange rate forwards are not updated frequently, and so the length of the period used to measure the shock varied from one event to another, potentially in a way that correlates with the nature of the shock. This measure of the shock was then found to result in larger standard errors, though the point estimates were roughly equivalent.

2.2 Monetary Policy Operations

The monetary policy operations of the four economies used in this study have changed considerably over the past decade (Brown 1997; Archer, Brookes, and Reddell 1999; Parent, Munro, and Parker 2003; Zettelmeyer 2004). This section briefly outlines the current monetary policy regimes, how they have changed, and how these changes may influence this study. Using this information, section 2.3 explains how the set of events used in the analysis were determined.

The four countries currently have very similar monetary policy operations. In particular, they all have

- fixed announcement dates for monetary policy decisions, albeit with the option to make changes at other times in response to extreme events;
- a short-term interest rate as the policy instrument;
- a preference for not surprising the market; and
- an inflation target.

While all four countries have been inflation targeters for the full sample considered in this paper, institutional aspects of monetary policy operations have changed since the early 1990s in important ways. Some of these changes have been gradual, while others have been more abrupt. In Australia, monetary policy operations have changed progressively since the early 1990s, to resemble the current operational framework by about 1998. Prior to 1998, though the dates of the Board meetings were known (usually the first Tuesday of every month, with no meeting in January), monetary policy decisions were typically not announced or implemented immediately after a meeting. From 1998 onward, all changes in monetary policy have been announced the day after a Board meeting. Only since September 2002 has there been a public announcement on the day after the Board meeting in the event that policy was not being changed. However, market commentary in the period 1998 to 2002 suggests that if policy was not changed the day after a Board meeting, then no change was anticipated until the subsequent meeting. So for Australia we have included no-change decisions from the beginning of 1998, as well as all changes in monetary policy from mid-1993. Table 1 provides a summary of the sample periods and number of events for each country.

Table 1. The Set of Events

	Australia	**Canada**	**New Zealand**	**United Kingdom**
Sample	July 30, 1993–June 2, 2004	October 28, 1996–June 8, 2004	March 17, 1999 June 10, 2004	July 10, 1997–June 10, 2004
Number of Events Used	79	33[b]	42	82
Number of Changes	24	23	20	27
Number of No-Changes	55	10	22	55
Regime Change	1998	December 2000	March 1999	June 1997
Old Regime	Policy changes included (9 events)	Policy changes included (4 events)	None included	None included
New Regime	Changes and no-changes included			
Meetings per Year	11	8	8	12
Announcement Time[a]	9:30 a.m. on the day after the Board meeting	9:00 a.m. on fixed announcement days	9:00 a.m. on fixed announcement days	12:30 p.m. on the second day of the Monetary Policy Committee meeting

[a]Some events do occur at different times.
[b]Changes after the September 11, 2001, terrorist attacks and in response to the Russian crisis are excluded. Also, eight events that coincide with changes in the federal funds rate are excluded.

In the other countries, changes in monetary policy operations have been more distinct and, in some cases, substantial. At the start of 1999, the Reserve Bank of New Zealand (RBNZ) moved from focusing on a monetary conditions index (a combination of the overnight interest rate and the exchange rate) as the main intermediate target of policy to using the overnight cash rate to implement monetary policy. Accordingly, we begin our New Zealand sample in 1999, as prior to this the motivation for interest rate changes was inextricably linked to exchange rate movements.

Like New Zealand, Canada has implemented a system of eight fixed announcement dates (starting December 2000). However, unlike New Zealand, there has been little change in the framework and focus of policy. We therefore include most changes to policy from 1996 onward, when our intraday exchange rate data for the Canadian dollar begin. Eight changes are excluded, when they are on the day of, or the day after, a change in the federal funds rate, reflecting the likelihood of contamination of the measure of the surprise in policy. Two further policy changes are excluded—the one following the August 1998 Russian crisis and the one after the September 11, 2001, terrorist attacks—again, for reasons of possible contamination.

In the United Kingdom, the operational responsibility for monetary policy passed from the government to the Bank of England in June 1997, ensuring an independent monetary policymaker. From this date on, policy announcements of either a change or no change occurred according to a preannounced schedule. We exclude monetary policy decisions made prior to June 1997, owing to the large shift in the monetary policy regime and uncertainty about the exact time at which changes were announced prior to 1997.

For each event, we searched Bloomberg and other sources to ensure that there was no contaminating information in the event window. Because we use a narrow event window, we did not find cause to exclude any events other than those outlined for Canada.[6] We also record the exact time of the event in order to make the events completely comparable.

[6] A few events are excluded due to missing intraday exchange rate data.

Given the similarities of the current monetary policy regimes in the four countries, we also present results using a pooled sample. In order to keep the pooled sample as homogenous as possible, only those events that are part of the current regimes, in which monetary policy is implemented according to fixed announcement dates, are included. This is the full sample for New Zealand and the United Kingdom and the sample from 1998 for Australia and the end of 2000 for Canada.

2.3 The Data

The two data series used in the event study are interest rates and exchange rates. We use bank bill interest rates (one-month and three-month rates) and futures contracts on the three-month bank bill interest rate. Most of the literature for the United States has calculated monetary policy surprises using federal funds futures contracts (see, for example, Kuttner 2001, Faust et al. 2003, and Bernanke and Kuttner 2005). However, futures contracts over the policy instrument interest rates are not available for the countries we use over our sample period, and so we use bank bill interest rates to calculate monetary policy surprises. One advantage of bank bill rates is that, unlike futures, the horizon of the instrument does not vary from one event to another, thereby simplifying the calculation of the surprise. Piazzesi and Swanson (2004) find that eurodollar interest rate futures provide good measures of interest rate surprises for the United States and are only marginally outperformed by federal funds futures. The interest rate surprise is calculated as the change in the one-month or three-month bank bill interest rate from the close of the day prior to the monetary surprise to the close of the day of the monetary surprise. Note that the surprise can be nonzero even when the policy interest rate was not changed, if the market placed at least some probability on there being a change. The surprises are measured in percentage points (100 basis points). The interest rates we use, and their sources, are described in the appendix.

The bilateral exchange rates are the U.S. dollar price of the domestic currency, from the Reuters electronic trading system, at ten-minute intervals. At each ten-minute interval, the

Table 2. The Data

	Australia	**Canada**	**New Zealand**	**United Kingdom**
Number of Events Used	79	33	42	82
Average $\lvert\Delta i\rvert$	0.13	0.21	0.15	0.09
Average $\lvert\Delta i^s\rvert$	0.07	0.06	0.05	0.07
Ratio of Event to Nonevent Day Exchange Rate Volatility[a]	1.34	1.17	1.33	1.13
Average $\lvert\Delta e_{10m}\rvert$	0.049	0.041	0.059	0.052

Notes: $|\Delta i|$ is the absolute change in the official interest rate in percentage points. $|\Delta i^s|$ is the absolute change in the one-month interest rate in percentage points.
[a]The volatility is calculated as the average absolute change in the exchange rate over ten-minute intervals. Averages are taken over a window starting two hours before the event and ending six hours after. The sample of nonevent days is constructed by taking the day exactly one week prior to each event.
$|\Delta e_{10m}|$ is the absolute percent change over ten-minute intervals on event days.

exchange rate observation is the average of the closest active bid and ask quotes. Goodhart and Payne (1996) and Danielsson and Payne (2002) have found that at ten-minute intervals, quote data are good proxies for actual transaction prices in exchange rate markets.

Table 2 gives a brief summary of the data. The average absolute change in the policy rate, $|\Delta i|$, is based on change and no-change event days in the sample. The average absolute change is typically about twice as large as the average absolute surprise, $|\Delta i^s|$. Exchange rate volatility—measured as the average absolute change in the exchange rate over ten-minute intervals—is higher on event days than on nonevent days (from two hours before to six hours after the event), providing some initial evidence that monetary policy has an effect on the exchange rate.

3. Results

3.1 The Impact of Monetary Policy Shocks

To quantify the impact of monetary policy on the exchange rate, we regress the change in the exchange rate over the event window on the monetary policy surprise, as represented by equation (1),

$$\Delta e_{[t-10m,t+60m]} = \alpha + \beta \Delta i_t^s + \varepsilon_t, \tag{1}$$

where $\Delta e_{[t-10m,t+60m]}$ is the percentage change in the U.S. dollar bilateral exchange rate from ten minutes before the event to sixty minutes after, and Δi_t^s is the surprise move in policy measured by the daily change in market interest rates.[7] We use the exchange rate from ten minutes before the policy change, rather than at the time of the policy change, in case there are mismatches in the timing of our exchange rate data and policy implementation. We present results using surprises derived from both one-month and three-month interest rates.

These regressions suggest that a 100-basis-point surprise tightening of monetary policy is estimated to lead to an appreciation of the exchange rate in the range of 1–2 percent in the hour following the event (table 3). When we use the sample pooled across countries, the estimate is in the middle of this range, just under 1.5 percent.[8] In recent years, the countries used in this study have moved their policy rates in 25-basis-point increments.[9] A 25-basis-point surprise would lead to an appreciation of 0.25–0.50 percent. The surprise in monetary policy explains only about 10–20 percent of the movement in the exchange rate over the seventy-minute interval. The low

[7]Note that the daily interest rate change, which we use because intraday interest rate data are not available, is potentially a noisy indicator of the true interest rate surprise and so could lessen the explanatory power of our regressions. Equation (1) includes both change and no-change events in order to have a sufficiently large sample. In section 3.3, we examine whether change and no-change surprises have different impacts.

[8]The data do not reject the restriction that the coefficient on interest rates is constant across countries.

[9]For all countries, there are some larger policy moves earlier in the sample. However, because it is only the surprise component—not the change in the policy rate—that enters the regression, these changes are not necessarily larger values in the regression.

Table 3. Impact of a 100-Basis-Point Monetary Policy Surprise

Country	One Month	Three Month	$\overline{R}^2$	Observations
Australia	0.96 (0.00)		0.16	79
		1.88 (0.00)	0.32	79
Canada	1.56 (0.00)		0.22	33
		1.67 (0.00)	0.23	33
New Zealand	1.83 (0.02)		0.11	42
		1.97 (0.01)	0.15	42
United Kingdom	1.04 (0.00)		0.11	82
		1.58 (0.00)	0.17	82
Pooled Sample	1.45 (0.00)		0.13	222
		1.77 (0.00)	0.17	222

Notes: The dependent variable is the change in the exchange rate (relative to the U.S. dollar) from ten minutes before the event to sixty minutes after. P-values are in parentheses.

proportion of exchange rate movements explained by the interest rate surprise, even in such a short window with an important piece of information, is in line with other work on the exchange rate—for example, Andersen et al. (2003) and Faust et al. (forthcoming).

For all countries and the pooled sample, the coefficient on the interest rate surprise is larger when the surprise is measured using a three-month interest rate than in the equivalent regression using a one-month interest rate. Presumably this is because the three-month rate includes the impact of the decision on expectations of future monetary policy, at least over the next three months, an important issue that we explore in section 3.2. The point estimates for the three-month surprises for Australia and Canada are similar to those in Zettelmeyer (2004) using daily data. However, his result for New Zealand is larger, 2.7 percent, possibly because his sample, being mostly before 1999 when a monetary conditions index was being used, does not contain purely exogenous monetary shocks. The point estimates are also similar in magnitude to the 1.2 for the deutschemark/euro exchange rate response to changes in U.S. interest rates contained in Faust et al. (forthcoming). However, their estimate for the pound's response to U.S. interest rates of 0.66 is smaller, suggesting that their result may contain some bias in estimating the impact of a change in a large country's interest rate on the bilateral exchange rate with a smaller country.

The timing of the impact of a monetary policy surprise on the exchange rate can be determined by estimating equation (2) for k ranging from two hours before the event to six hours after (at ten-minute intervals).

$$\Delta e_{[t,t+k]} = \alpha + \beta \Delta i_t^s + \varepsilon_t \tag{2}$$

The results are shown in figures 1–4, where the surprise is measured using the one-month interest rate. In all four countries, there is a sharp spike in the impact in the ten minutes following the event, demonstrating that monetary policy announcements have a rapid impact on the exchange rate. The relative stability of the coefficients over the six hours after the event indicate that the surprise has little additional influence after its immediate impact. The standard errors, the dashed lines in the graphs, widen further from the event as the policy change becomes a smaller proportion of the information incorporated into the exchange rate. As a result, the statistical significance using daily data will be substantially weaker.

For Australia and Canada, there is significant movement in the exchange rate prior to the event in the same direction as the response following the event. Given that this is gradual for Canada, it is suggestive of late changes in market expectations of the policy announcement, perhaps as participants' expectations coalesce around a particular policy announcement. Such changes in market expectations would presumably also be reflected in intraday

Figure 1. Australian Dollar: Response to a 100-Basis-Point Interest Rate Surprise

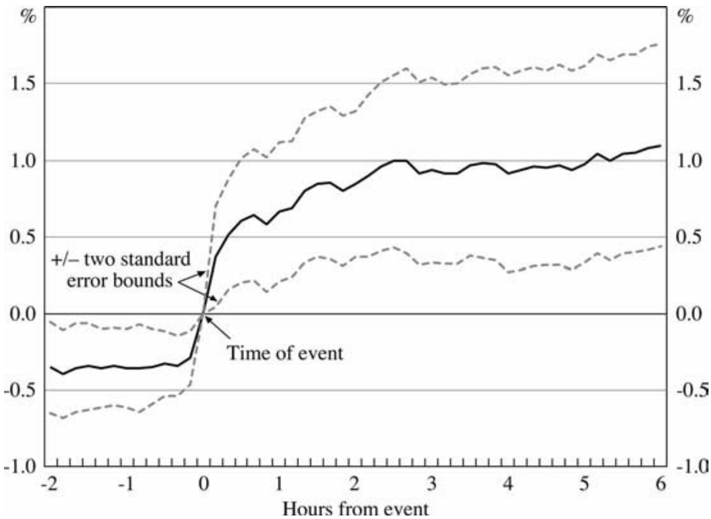

Figure 2. Canadian Dollar: Response to a 100-Basis-Point Interest Rate Surprise

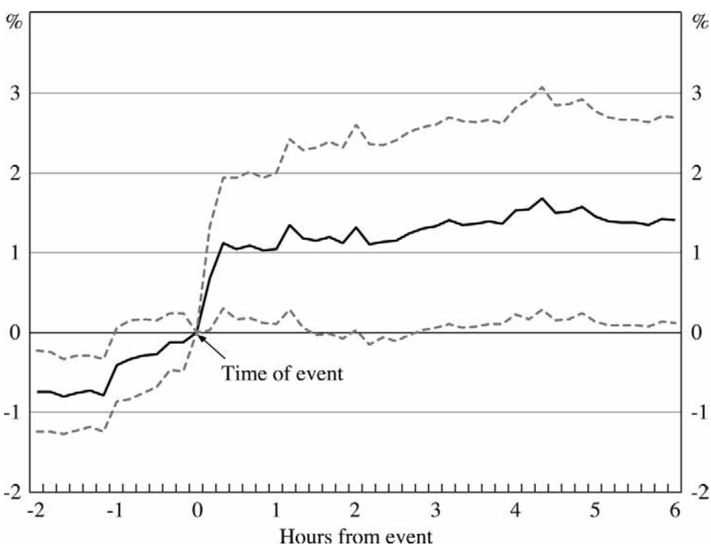

Figure 3. New Zealand Dollar: Response to a 100-Basis-Point Interest Rate Surprise

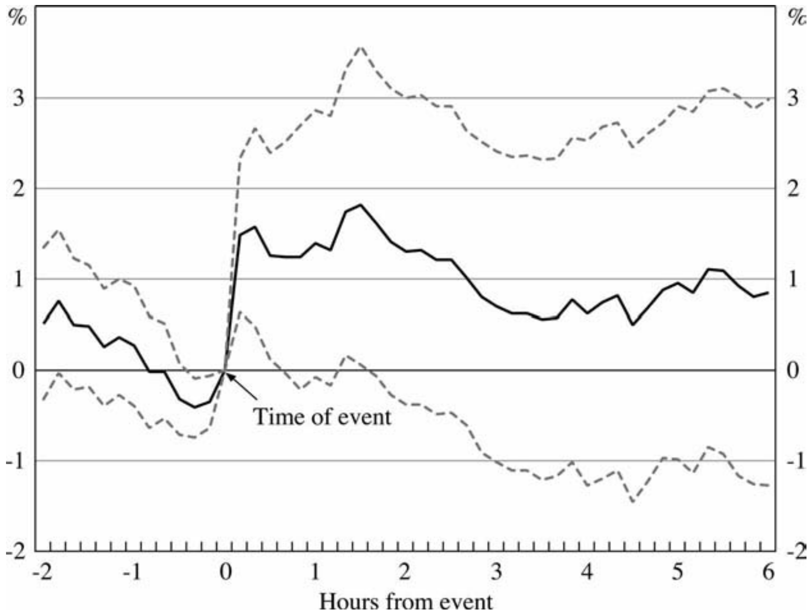

interest rate data were they available. For Australia, the sharp jump that occurs ten minutes before the event most likely reflects slight differences between the timing of the announcement and the exchange rate data. The significance of this change immediately prior to the event is not unduly influenced by any particular observations, and so the result does not appear to be the result of leaked information.[10] It is because of this possible timing mismeasurement that we base our main results on the exchange rate starting ten minutes before the event.

[10]This result is not sensitive to the exclusion of the only two events for Australia in which there is any suggestion of some participants seemingly having early access to the policy outcome: one in which the monetary policy decision was mistakenly released to some market participants six minutes early (February 2, 2000) and the other in which Bloomberg mistakenly released a report about one minute early, even though it did not yet know the outcome (July 3, 2002). Note that these do not affect our main results, because we use the exchange rate from ten minutes before the policy change.

Figure 4. British Pound: Response to a 100-Basis-Point Interest Rate Surprise

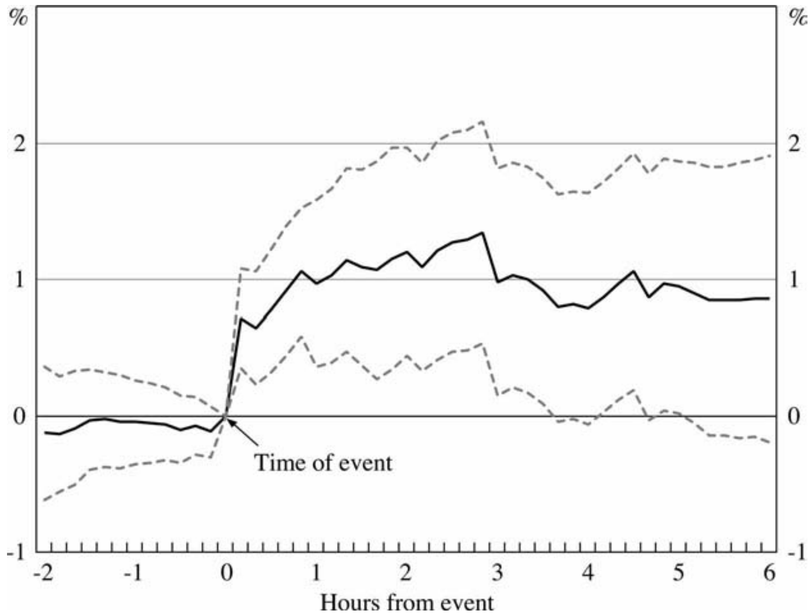

Given the importance for New Zealand of financial linkages with Australia, and New Zealand's smaller relative size, it is also interesting to examine the impact of monetary policy surprises in New Zealand on the New Zealand dollar/Australian dollar bilateral exchange rate. Figure 5 demonstrates that the response is more precisely estimated, and larger, than the response of the New Zealand dollar/U.S. dollar exchange rate shown in figure 3.

3.2 The Importance of Expectations

The impact of a policy decision on expectations of future policy may be important in determining the exchange rate's response. This is apparent from the different results obtained when measuring the monetary policy surprise using one-month and three-month interest rates. A monetary policy decision might simply surprise the market in its timing (a "timing" surprise), or it might be a surprise that shifts policy expectations at all horizons (a "level" surprise).

Figure 5. New Zealand Dollar/Australian Dollar Bilateral Rate: Response to a 100-Basis-Point Interest Rate Surprise

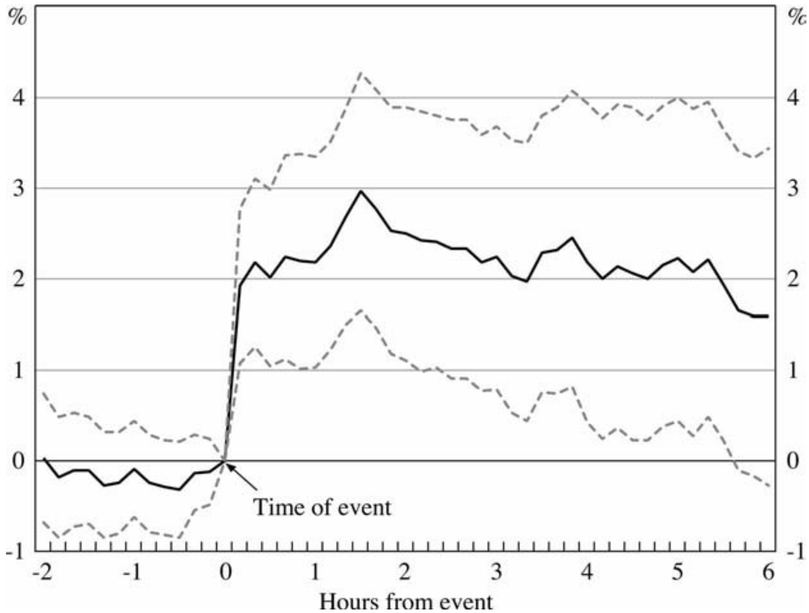

A change in current monetary policy could even shift expectations of future policy by more than the impact on one-month interest rates if market participants believe that it indicates that future changes in the same direction are likely.

To test whether expectations are important, we examine changes in three-month interest rate futures (details of these contracts are in the appendix). These are not perfect measures of policy expectations, but changes in these interest rates have been shown to be a reasonable guide to changes in policy expectations (Piazzesi and Swanson 2004). Typically, about eight futures contracts trade at any one time, giving a horizon of about two years in total. What is termed the "first" interest rate futures contract expires sometime in the next three months and is settled on the three-month interest rate prevailing on that expiration date. The "second" futures contract is settled on the three-month interest rate prevailing three months after the expiration of the first contract, and so on. Because the futures

Figure 6. Change in Spot Interest Rate versus Change in Futures Rate

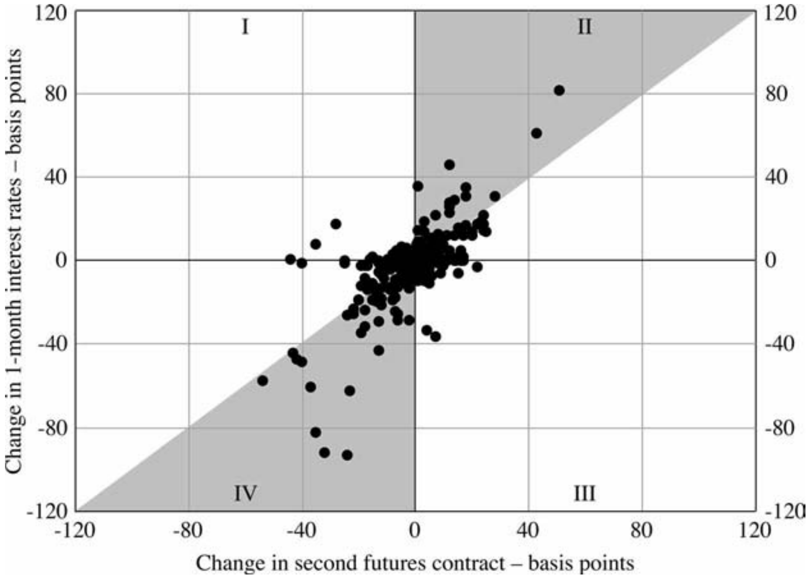

contracts expire at fixed points in time three months apart, but monetary policy events can happen at any time in a three-month period, the horizon until the contract expiration can differ from one event to another. While not ideal, this is unavoidable given that one-month futures contracts are not available for these countries over the relevant sample and the time between the fixed announcement dates can vary in Canada and New Zealand. We consider the sensitivity to this measurement issue in section 3.3.

Figure 6 shows a scatter plot of the change in the one-month interest rate against the change in the second futures contract, on days with a monetary policy event, for all four countries. We use the second futures contract, as it has the advantage that it does not expire until after at least two complete monetary policy decision cycles.

Given that monetary policy changes are not quickly reversed, surprises that lead to a positive (negative) change in the one-month interest rate are unlikely to also lead to a negative (positive) change in the futures rate. This suggests that there should be very few points in quadrants I and III. It would seem likely that most surprises would be somewhere between a level and timing surprise. Changes in the futures rate should then be of the same sign as, but smaller in magnitude than, the change in the one-month rate. In this case, most observations would lie in the shaded area in figure 6. The data broadly fit this pattern, especially for large monetary surprises, although there are still quite a number of points outside the shaded area. If all surprises had the same effect on expectations of future policy, then they would lie on a straight line. The scatter plot shows that this is clearly not the case, and regressions for each country indicate that the change in the one-month interest rate explains only about half of the change in the second futures rate. It thus appears that there is sufficient heterogeneity to test whether the effect that a monetary policy decision has on expectations of future policy is important in determining the decision's impact on the exchange rate.

We incorporate information on the change in the futures contract in equation (3):

$$\Delta e_{[t-10m,t+60m]} = \alpha + \beta \Delta i_t^s + \gamma \Delta i_t^{s,f} + \varepsilon_t. \tag{3}$$

A "timing" surprise—a surprise that does not change expectations of the level of future policy—is captured by β, since there is no change in the futures interest rate (that is, $\Delta i_t^{s,f} = 0$). A "level" surprise—a surprise that changes expectations of future interest rates by as much as the surprise in current policy, i.e., $\Delta i_t^s = \Delta i_t^{s,f}$—is measured by $\beta + \gamma$. The results are presented in table 4.

These results confirm that a surprise in the level of policy leads to a greater change in the exchange rate than does a surprise in the timing of policy. Estimates of the impact of a 100-basis-point surprise in timing, β, for the individual countries are imprecisely estimated but range from being negative to around 1 percent. However, in the larger pooled sample, the estimate is very precisely estimated to be 0.87. A 100-basis-point surprise increase in the level of the (current and future) policy instrument is estimated to lead to a 1.3–2.2 percent appreciation in the exchange rate, as seen by the estimates of $\beta + \gamma$. The pooled sample produces an estimate in the middle of this range, 1.68. These estimates of the impact of a level surprise are highly significant for all of the countries and the pooled sample. This indicates that a 25-basis-point timing (level)

Table 4. Timing and Level Surprises

	Timing Surprise β	γ	**Level Surprise** $\beta + \gamma$[a]	$\overline{R}^2$	**Observations**
Australia	0.46	0.80	1.26	0.21	79
	(0.14)	(0.02)	(0.00)		
Canada	0.84	0.85	1.69	0.23	33
	(0.26)	(0.20)	(0.00)		
New Zealand	1.06	0.74	1.80	0.11	42
	(0.35)	(0.35)	(0.02)		
United Kingdom	−0.21	2.44	2.22	0.29	82
	(0.59)	(0.00)	(0.00)		
Pooled Sample	0.87	0.81	1.68	0.17	222
	(0.00)	(0.00)	(0.00)		

Note: P-values are in parentheses.
[a]The p-value for $\beta + \gamma$ is calculated using a Wald test.

surprise would appreciate the exchange rate by around 0.2 percent (0.4 percent). Clearly, a level surprise has a much larger impact on the exchange rate than a timing surprise.

A key economic theory governing exchange rates is the prediction of UIP that if domestic interest rates are higher than foreign rates, the exchange rate should depreciate gradually in order to equalize returns. Macroeconomic models that incorporate UIP and rational expectations, such as Dornbusch (1976), typically predict a sharp appreciation in response to a surprise monetary tightening in order for the exchange rate to subsequently depreciate in line with UIP. While UIP fails empirically, our results show the exchange rate immediately appreciates in response to a monetary tightening, which accords with the prediction of these exchange rate models. Our work cannot say anything about whether the exchange rate subsequently depreciates as predicted by UIP, but it is interesting to compare the magnitude of our results to the initial jump that would be consistent with UIP.

To calculate the jump in the exchange rate in response to a surprise tightening in monetary policy that is consistent with UIP, we need to know both how long the change in monetary policy will be sustained and the level that the exchange rate is expected to return to after the change in the interest differential is eliminated. Using our estimates that separate the timing from the level effects, we can attempt to control for the first of these issues. But without knowing what caused the monetary surprise, and what impact that news had on the equilibrium value of the exchange rate, we cannot determine the level to which the exchange rate is expected to return. For the purpose of this calculation, we assume that the expected long-run value of the exchange rate did not change with the monetary surprise. If the monetary decision is a pure timing surprise, the interest rate given by the futures contract that expires in three to six months does not change. Assuming that the surprise lasts four and one-half months (that is, the midpoint of three and six months), in order for UIP to subsequently hold, a 100-basis-point surprise increase would require an immediate appreciation of less than 0.50 percent (0.375 percent).[11] In contrast, we estimate the response to a 100-basis-point timing surprise to be more than twice as large, 0.87 percent. So while UIP is found to fail empirically, our results suggest that the initial response of the exchange rate to a monetary policy surprise is in the direction predicted by macroeconomic models in which UIP holds, but it is seemingly larger than this theory would suggest. Of course, this interpretation is subject to the important caveat that we do not know how the long-run equilibrium level of the exchange rate has changed.

3.3 Robustness

To test the robustness of our findings, we include a range of other variables in the regressions. For brevity, we only report results using the pooled sample. The equivalent regressions for the individual countries produce similar results, though understandably with larger standard errors. Table 5 reports specifications using surprises based on one-month bank bill interest rates, while table 6 repeats the regressions using surprises based on three-month bank bill rates.

[11] This is the amount that the exchange rate would need to depreciate over the subsequent four and one-half months in order to equalize returns. This calculation assumes that the foreign interest rate remains constant.

Table 5. Pooled Results: One-Month Surprise

	Surprise Change β	Futures Contract γ 2nd	3rd	Expected Change	Maturity[a]	Change Dummy × Surprise Change	Change Dummy × 2nd Future	$\overline{R}^2$
I	1.45							0.13
	(0.00)							
II	0.87	0.81						0.17
	(0.00)	(0.00)						
III	0.95		0.74					0.17
	(0.00)		(0.00)					
IV	0.89	0.83		−0.21				0.18
	(0.00)	(0.00)		(0.00)				
V	0.94	0.76		−0.21	0.00			0.17
	(0.00)	(0.01)		(0.09)	(0.61)			
VI	0.61	0.72		−0.25	0.00	0.58		0.17
	(0.16)	(0.01)		(0.06)	(0.61)	(0.26)		
VII	0.74	0.47		−0.24	−0.01	0.15	0.63	0.18
	(0.10)	(0.19)		(0.06)	(0.46)	(0.82)	(0.26)	

Notes: The dependent variable is the change in the exchange rate (relative to the U.S. dollar) from ten minutes before the event to sixty minutes after. P-values are in parentheses. There are 222 observations in all regressions.
[a]The maturity variable is: (the change in the futures contract) × (the difference between the days to maturity and the average days to maturity).

The coefficients on monetary policy surprises are found to be robust and maintain their statistical significance across a range of specifications. Specifications I and II repeat the pooled results from tables 3 and 4. The estimates using the three-month interest rate surprise are slightly larger, reflecting their less-precise separation of timing and level surprises. Using the third rather than the second futures contract, specification III, does not change the results appreciably.

One surprising result is that the coefficient on the expected change (the change in monetary policy less the unexpected change) is always about −0.2 percent and marginally significant (specifications IV–VII). This runs counter to our priors that only unexpected changes in monetary policy should affect the exchange rate. This result owes a lot to one particular event in New Zealand on May 17, 2000. The tightening in monetary policy on this day was almost completely anticipated. But the particularly hawkish monetary policy statement released with the decision seemingly led to concerns

Table 6. Pooled Results: Three-Month Surprise

	Surprise Change β	Futures Contract 2nd γ	Futures Contract 3rd γ	Expected Change	Maturity[a]	Change Dummy × Surprise Change	Change Dummy × 2nd Future	$\overline{R}^2$
I	1.77 (0.00)							0.17
II	1.27 (0.00)	0.52 (0.08)						0.18
III	1.32 (0.00)		0.48 (0.07)					0.18
IV	1.32 (0.00)	0.51 (0.08)		−0.21 (0.08)				0.19
V	1.42 (0.00)	0.41 (0.22)		−0.20 (0.10)	−0.01 (0.49)			0.18
VI	0.94 (0.08)	0.39 (0.24)		−0.23 (0.06)	−0.01 (0.48)	0.72 (0.20)		0.19
VII	1.01 (0.08)	0.30 (0.45)		−0.23 (0.06)	−0.01 (0.45)	0.49 (0.54)	0.25 (0.71)	0.18

Notes: The dependent variable is the change in the exchange rate (relative to the U.S. dollar) from ten minutes before the event to sixty minutes after. P-values are in parentheses. There are 222 observations in all regressions.
[a]The maturity variable is: (the change in the futures contract) × (the difference between the days to maturity and the average days to maturity).

about the impact of the indicated course of policy on the growth of the economy and a sharp *depreciation* of the exchange rate. Excluding this observation, the expected change in policy does not have a statistically significant impact on the exchange rate.

A variable that controls for the changing number of days until maturity of futures contracts, used in specifications V–VII, is always economically and statistically insignificant. This suggests that our conclusions about the timing and level surprises are not unduly influenced by the fact that the horizon of interest rate futures is not constant across events.

The coefficient on a dummy for whether the decision was a change in monetary policy, multiplied by the monetary policy surprise, is always positive though not significant (specifications VI and VII). This suggests that there may be a slightly greater effect on the exchange rate when the surprise monetary decision is a change in the policy interest rate. Alternatively, this could simply reflect the fact that the proportion of the interest rate change that is caused by monetary policy is likely to be higher when the surprise is larger, which typically occurs when monetary policy is changing.

4. Conclusions

In this paper, we use an event study to isolate the impact of changes in monetary policy on the exchange rate. Two important aspects of our study enable us to abstract from endogeneity and the influence of other exogenous news. First, we use a sample period for four countries (Australia, Canada, New Zealand, and the United Kingdom) in which monetary policy does not focus on the exchange rate and the decision is predetermined when it is implemented. Second, we use intraday data with a narrow event window. The results indicate that the exchange rate appreciates on average by around 1.5 percent in response to an unanticipated 100-basis-point increase in the policy interest rate. The estimates for individual countries range from 1.0 percent to 1.8 percent. For a 25-basis-point surprise, this equates to an average appreciation of 0.35 percent (0.25–0.50 percent for individual countries). These results are slightly smaller than those in Zettelmeyer (2004) but, for the most part, are marginally larger than those in Faust et al. (forthcoming).

The impact of monetary policy changes on the exchange rate is found to occur virtually instantaneously. If we use an event window that ends well after the monetary policy decision, the estimates do not change, indicating that the news is rapidly incorporated into exchange rates, although the standard errors widen. Despite using a narrow event window in which no other identifiable events occurred, the monetary shock explains only 10–20 percent of the variation in the exchange rate in that short window. In general, the results suggest that monetary policy can account for only a small part of the observed volatility in the exchange rate. The small proportion explained by such high-profile news indicates that there is still much to learn in explaining exchange rate movements.

In the second part of the paper, we present new evidence that not all monetary surprises will have the same effect on the exchange rate. Those that cause a revision to expectations of future policy are found to have a much larger impact than surprises in the timing of a change in monetary policy. A 100-basis-point (25-basis-point) increase in current and future policy interest rates is found to appreciate the exchange rate by around 1.7 percent (0.4 percent). Estimates for individual countries range from 1.3 percent to 2.2 percent. In contrast, a monetary surprise that only brings forward an anticipated change in policy is found to appreciate the exchange rate by just 0.9 percent (0.2 percent).

Appendix

Table 7. Data Description and Sources

	Australia	**Canada**	**New Zealand**	**United Kingdom**
Interest Rates				
One-Month Interest Rate	Thirty-day bank bills (RBA)	One-month bankers acceptances (BoC)	One-month wholesale bill (RBNZ)	One-month LIBOR (Datastream: LDNIB1M)
Three-Month Interest Rate	Ninety-day bank bills (RBA)	Three-month bankers acceptances (BoC)	Three-month wholesale bill (RBNZ)	Three-month LIBOR (Datastream: LDNIB3M)
Futures Contracts	Ninety-day bank bills (Bloomberg: IR1 commodity)	Three-month bankers acceptances (Bloomberg: BA1 commodity)	Three-month bank bills (Bloomberg: ZB1 commodity)	Three-month LIBOR (Bloomberg: L1 commodity)
Futures Exchange	Sydney Futures Exchange	Montreal Exchange	Sydney Futures Exchange	London International Financial Futures Exchange
Settlement Months	March, June, September, and December for all countries			
Expiration Day	Second Friday	Third Tuesday	Thursday after first Wednesday after 9^{th} of month	Third Thursday
Exchange Rates	RBA/Reuters, ten-minute intervals, midpoint of two closest quotes			

Note: RBA is the Reserve Bank of Australia, BoC is the Bank of Canada, and RBNZ is the Reserve Bank of New Zealand.

References

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega. 2003. "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange." *American Economic Review* 93 (1): 38–62.

Archer, D., A. Brookes, and M. Reddell. 1999. "A Cash Rate System for Implementing Monetary Policy." *Bulletin* 62 (1): 51–61, Reserve Bank of New Zealand.

Bernanke, B., and K. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3): 1221–57.

Brown, G. 1997. "The New Monetary Policy Framework." Letter to the Governor of the Bank of England, available at www.hm-treasury.gov.uk/newsroom_and_speeches/press/1997/press_49_97letter.cfm.

Danielsson, J., and R. Payne. 2002. "Real Trading Patterns and Prices in Spot Foreign Exchange Markets." *Journal of International Money and Finance* 21 (2): 203–22.

Dornbusch, R. 1976. "Expectations and Exchange Rate Dynamics." *Journal of Political Economy* 84 (6): 1161–76.

Eichenbaum, M., and C. L. Evans. 1995. "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates." *Quarterly Journal of Economics* 110 (4): 975–1009.

Engel, C. 1996. "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence." *Journal of Empirical Finance* 3 (2): 123–92.

Faust, J., and J. H. Rogers. 2003. "Monetary Policy's Role in Exchange Rate Behavior." *Journal of Monetary Economics* 50 (7): 1403–24.

Faust, J., J. H. Rogers, E. T. Swanson, and J. H. Wright. 2003. "Identifying the Effects of Monetary Policy Shocks on Exchange Rates using High-Frequency Data." *Journal of the European Economic Association* 1 (5): 1031–57.

Faust, J., J. H. Rogers, S. B. Wang, and J. H. Wright. Forthcoming. "The High Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements." *Journal of Monetary Economics.*

Gencay, R., G. Ballocchi, M. Dacorogna, R. Olsen, and O. Pictet. 2002. "Real-Time Trading Models and the Statistical Properties of Foreign Exchange Rates." *International Economic Review* 43 (2): 463–91.

Goodhart, C. A. E., and R. G. Payne. 1996. "Microstructural Dynamics in a Foreign Exchange Electronic Broking System." *Journal of International Money and Finance* 15 (6): 829–52.

Grenville, S. A. 1995. "The Monetary Policy Transmission Process: What Do We Know? (And What Don't We Know?)" *Bulletin* (Reserve Bank of Australia) (September):15–33.

Kuttner, K. N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.

Parent, N., P. Munro, and R. Parker. 2003. "An Evaluation of Fixed Announcement Dates." *Bank of Canada Review* (Autumn):311.

Piazzesi, M., and E. T. Swanson. 2004. "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy." NBER Working Paper No. 10547.

Thiessen, G. 1995. "Uncertainty and the Transmission of Monetary Policy in Canada." Hermes-Glendon Lecture, York University, March 30.

Zettelmeyer, J. 2004. "The Impact of Monetary Policy on the Exchange Rate: Evidence from Three Small Open Economies." *Journal of Monetary Economics* 51 (3): 635–52.