JSPS Grants-in-Aid for Creative Scientific Research

Understanding Inflation Dynamics of the Japanese Economy

Working Paper Series No.73

A New Method for Identifying the Effects of Foreign Exchange Interventions

Chih-nan Chen
Tsutomu Watanabe
and
Tomoyoshi Yabu

First draft: October 2005
This version: May 18, 2011

Research Center for Price Dynamics
Institute of Economic Research, Hitotsubashi University
Naka 2-1, Kunitachi-city, Tokyo 186-8603, JAPAN
Tel/Fax: +81-42-580-9138
E-mail: sousei-sec@ier.hit-u.ac.jp
http://www.ier.hit-u.ac.jp/~ifd/

A New Method for Identifying the Effects of Foreign Exchange Interventions*

Chih-nan Chen† — Harvard University

Tsutomu Watanabe‡ — Hitotsubashi University

Tomoyoshi Yabu§ — Keio University

First draft: October 2005

This version: May 18, 2011

Abstract

Central banks react even to intraday changes in the exchange rate; however, in most cases, intervention data is available only at a daily frequency. This temporal aggregation makes it difficult to identify the effects of interventions on the exchange rate. We apply the Bayesian MCMC approach to this endogeneity problem. We use "data augmentation" to obtain intraday intervention amounts and estimate the efficacy of interventions using the augmented data. Applying this new method to Japanese data, we find that an intervention of one trillion yen moves the yen/dollar rate by 1.7 percent, which is more than twice as much as the magnitude reported in previous studies applying OLS to daily observations. This shows the quantitative importance of the endogeneity problem due to temporal aggregation.

JEL Classification Numbers: C11, C22, F31, F37

Keywords: foreign exchange intervention; intraday data; Markov-chain Monte Carlo method; endogeneity problem; temporal aggregation

---

*We would like to thank Rasmus Fatum, Simon Gilchrist, Hidehiko Ichimura, Yuko Imura, Takatoshi Ito, Daisuke Nagakura, Jouchi Nakajima, Masao Ogaki, Paolo Pesenti, Toshiaki Watanabe, Hung Yin-Ting, and seminar participants at the Bank of Japan, Boston University, and NBER Japan Project Meeting for their helpful comments and suggestions. This research forms part of the project on "Understanding Inflation Dynamics of the Japanese Economy" funded by a JSPS Grant-in-Aid for Creative Scientific Research (18GS0101).

†Center for Multicultural Mental Health Research, Harvard University (`cchen@chareresearch.org`).

‡Hitotsubashi University and University of Tokyo (`tsutomu.w@r.hit-u.ac.jp`).

§Corresponding author. Faculty of Business and Commerce, Keio University (`tyabu@fbc.keio.ac.jp`).

1 Introduction

Are foreign exchange interventions effective? This issue has been debated extensively since the 1980s, but no conclusive consensus has emerged.[1] A key difficulty faced by researchers in answering this question is the endogeneity problem: the exchange rate responds "within the period" to foreign exchange interventions *and* the central bank reacts "within the period" to fluctuations in the exchange rate. This difficulty would not arise if the central bank responded only slowly to fluctuations in the exchange rate, or if the data sampling interval were sufficiently fine.

As an example, consider the case of Japan. The central bank of Japan, which is known to be one of the most active interveners, started to disclose intervention data in July 2001, and this has rekindled researchers' interest in the effectiveness of interventions. Studies using this recently disclosed data include Ito (2003), Fatum and Hutchison (2006), Dominguez (2003), Chaboud and Humpage (2005), Galati et al. (2005), Fratzscher (2005), Watanabe and Harada (2006), and Fatum (2008), among others. However, the information disclosed is limited: only the total amount of interventions on a day is released to the public at the end of a quarter, and no detailed information, such as the time of the intervention(s), the number of interventions over the course of the day, and the market(s) (Tokyo, London, or New York) in which the intervention(s) was/were executed, is disclosed.[2] Most importantly, the low frequency of the disclosed data poses a serious problem for researchers because, as is well known, the Japanese central bank often reacts to intraday fluctuations in the exchange rate.[3]

In this paper, we propose a new methodology based on Gibbs sampling to eliminate this endogeneity problem due to temporal aggregation. Consider a simple two-equation system. Hourly changes in the exchange rate, Δs_h, satisfy $\Delta s_h = \alpha I_h +$ disturbance, where I_h is the hourly amount of interventions. On the other hand, the central bank policy reaction function is given by $I_h = \beta \Delta s_{h-1} +$ disturbance. Suppose that this two-equation system represents the true structure of the economy, and that s_h is observable at the hourly frequency but I_h is not - researchers are able to observe only the daily sum of I_h, and in that sense, intervention data suffers from temporal

[1]See Edison (1993), Dominguez and Frankel (1993), Sarno and Taylor (2003), and Neely (2005) for surveys on this topic.

[2]This also applies to monetary authorities in other industrialized countries. Important exceptions are the Bank of Canada, Denmark's National Bank, and the Swiss National Bank, which disclose information regarding intraday transactions to researchers. Studies using the Swiss data to evaluate the efficacy of intraday interventions include Fischer and Zurlinden (1999) and Payne and Vitale (2003), while Fatum and Pedersen (2009) use the Danish data.

[3]Chang and Taylor (1998), for example, counting the number of reports by Reuters about Japanese central bank interventions from October 1, 1992 to September 30, 1993, find that there were reports of 154 interventions on 69 days, implying that, when it intervenes, the Japanese central bank intervenes, on average, two or three times a day.

aggregation. Given this environment, our task is to estimate the unknown parameters (i.e., α, β, and the variance of each disturbance term).

The key idea of the methodology we propose is as follows. Suppose we have a guess for the values of the unknown parameters. Then the exchange rate equation and the policy reaction function allow us to recover the hourly amount of intervention, subject to the constraint that the sum of hourly amounts equals the daily amount, which is observable. In the extreme case where the variance of the disturbance term in the first equation is zero, we estimate I_h as $I_h = \alpha^{-1}\Delta s_h$ using the first equation. In the other extreme case where the variance of the disturbance term in the second equation is tiny, we have $I_h = \beta\Delta s_{h-1}$ from the second equation. In more general cases, one can guess (and we will verify this later) that the estimate of I_h is a weighted average of the two, with the weights being determined by the relative importance of the two disturbance terms. Once we obtain an estimate for the hourly amount of intervention in this way, we can estimate the unknown parameters without encountering an endogeneity problem. By repeating this procedure, we are able to estimate the unknown parameters as well as the hourly amount of intervention.

Our method can be seen as an application of data augmentation techniques based on Markov-chain Monte Carlo (MCMC) methods to the endogeneity problem. The idea of using data augmentation to cope with various problems due to temporal aggregation goes back to Liu (1969), who proposed a simple method to convert low frequency (say, quarterly) observations into high frequency (say, monthly) observations. Chow and Lin (1971) developed a best linear unbiased method to convert low frequency observations into high frequency observations. Our paper is most closely related to Hsiao (1979) and Palm and Nijman (1982), whose models consist of two equations in which an endogenous variable y is determined by an explanatory variable x (i.e., $y_t = bx_t + u_t$) and x is determined by an exogenous variable z (i.e., $x_t = az_t + v_t$). They consider a setting in which researchers have access to semi-annual observations for y and z but only annual observations for x, and obtain an ML-estimator for b by integrating out the missing observations.

The model we seek to estimate in this paper differs from those of Hsiao (1979) and Palm and Nijman (1982) in some important respects. First, the extent to which data is aggregated is much higher than in these previous studies. Specifically, it is assumed in this paper that the intervention amount is decided by the central bank on an hourly basis but is observable only at the daily frequency. Thus, it is necessary to integrate out twenty four missing observations, which is more difficult to implement. Second, the model to be estimated in this paper is non-linear. Recent studies on the central bank policy reaction function, including Almekinders and Eijffinger (1996) and Ito and Yabu (2007), emphasize that the policy reaction function has the nature of an Ss rule due to the presence of fixed costs associated with policy changes. In this case, I_h no longer depends linearly on Δs_{h-1}, so that the resulting likelihood function is more complicated. Given such a non-linear structure of the model, it is very hard or no longer possible to compute likelihood functions in the way suggested by Hsiao (1979) and Palm and Nijman (1982). To overcome this difficulty, we employ the Bayesian MCMC method.

The idea of applying MCMC methods to data augmentation was first proposed by Tanner and Wong (1987), and MCMC methods have been employed in several studies, including Eraker (2001) who used it in the context of estimating parameters in continuous diffusion processes when only discrete, and sometimes low-frequency, data are available. However, to the best of our knowledge, this paper is the first application of the Bayesian MCMC approach to the endogeneity problem due to temporal aggregation.[4]

The rest of the paper is organized as follows. Section 2 provides a detailed explanation of our methodology to address the endogeneity problem, while Section 3 presents simulation results to demonstrate how the methodology works. In Sections 4 and 5, we apply our method to Swiss and Japanese intervention data. A unique feature of the Swiss data is that it records the amount of intraday intervention with a time stamp up to the minute. Using the Swiss data, we conduct an experiment in which we first apply our method to aggregated daily intervention data to estimate the efficacy of intervention and then compare the estimate with the one obtained using the hourly intervention data. We find that the two estimates are very close to each other, implying that endogeneity bias due to temporal aggregation is successfully eliminated by our method. Applying our method to the Japanese data, we find that an exchange rate intervention (e.g., a sale) of one trillion yen leads to a 1.7 percent change in the value of the yen (depreciation). This is more than twice as large as the magnitude reported in previous studies such as Ito (2003) and Fratzscher (2005), which apply ordinary least squares to daily intervention and exchange rate data. This result is consistent with the prediction that endogeneity creates a bias toward zero for the intervention coefficient as long as the central bank follows a leaning-against-the-wind policy. Section 6 concludes

[4]The issue we discuss in this paper is related to the macroeconomic argument that if agents' decision intervals do not coincide with the sampling interval, then inferences made about the behavior of economic agents from observed time series can be distorted (see Christiano and Eichenbaum 1987 and Sims 1971). McCrorie and Chambers (2006) investigate the problem of spurious Granger causality relationships that arise from temporal aggregation.

the paper, while the Appendix provides the technical details of our methodology.

2 Methodology

2.1 The endogeneity problem in identifying the effects of foreign exchange interventions

In this section, we present a detailed description of our methodology to address the endogeneity problem in identifying the effects of foreign exchange interventions on the exchange rate. Consider a simple model of the following form:

$$s_{t,h} - s_{t,h-1} = \alpha I_{t,h} + \epsilon_{t,h} \tag{1}$$

$$I_{t,h} = \beta(s_{t,h-1} - s_{t-1,h-1}) + \eta_{t,h} \tag{2}$$

where $s_{t,h}$ is the log of the exchange rate at hour h of day t ($t = 1, ..., T$ and $h = 1, ..., 24$), $I_{t,h}$ is the purchase of domestic currency (and the selling of foreign currency) implemented by a central bank between $h-1$ and h of day t, and $\epsilon_{t,h}$ and $\eta_{t,h}$ are disturbance terms satisfying $\epsilon_{t,h} \sim i.i.d.\ N(0, \sigma_\epsilon^2)$ and $\eta_{t,h} \sim i.i.d.\ N(0, \sigma_\eta^2)$. Equation (1) represents the exchange rate dynamics, while equation (2) is the central bank's policy reaction function. We assume that α is negative, implying that intervention consisting of the selling of domestic currency ($I_{t,h} < 0$) leads to a depreciation of the domestic currency ($s_{t,h} - s_{t,h-1} > 0$) and vice versa. An important assumption is that the exchange rate is observable at the hourly frequency, while interventions are observable only at the daily frequency: namely, we observe $I_t \equiv \Sigma_{h=1}^{24} I_{t,h}$. Note that if we were able to observe $I_{t,h}$ at the hourly frequency, we could obtain unbiased estimators of α and β by applying OLS to each of the two equations separately.

Taking partial sums of both sides of the equations leads to a daily model of the following form:

$$s_{t,24} - s_{t-1,24} = \alpha I_t + \epsilon_t \tag{3}$$

$$I_t = \beta \sum_{h=1}^{24} (s_{t,h-1} - s_{t-1,h-1}) + \eta_t \tag{4}$$

where $s_{t,24} - s_{t-1,24} = \sum_{h=1}^{24}(s_{t,h} - s_{t,h-1})$, $I_t = \sum_{h=1}^{24} I_{t,h}$, $\epsilon_t \equiv \Sigma_{h=1}^{24} \epsilon_{t,h}$, and $\eta_t \equiv \Sigma_{h=1}^{24} \eta_{t,h}$. This shows that the endogeneity problem arises in this daily model, so that a simple application of OLS to each of the two equations separately no longer works. To illustrate this, suppose that the central bank adopts a leaning-against-the-wind policy, so that β takes a positive value. Then an increase in $\epsilon_{t,h}$ leads to an increase in $s_{t,h} - s_{t,h-1}$ through equation (1), and to an increase in $I_{t,h+1}$ through equation (2). This means that I_t and ϵ_t in equation (3) are positively correlated, so that an OLS estimator of α has an upward bias. On the other hand, an increase in $\eta_{t,h}$ increases $I_{t,h}$ through equation (2), thereby creating an appreciation of the yen as long as α is negative. This implies that the error term in equation (4), η_t, and the regressor, $\sum(s_{t,h} - s_{t-1,h-1})$, are negatively correlated and, as a result, an OLS estimator of β has a downward bias.

2.2 MCMC method

We propose a method for estimating equations (1) and (2) using the daily data for interventions and the hourly data for the exchange rate. The set of parameters to be estimated is α, β, σ_ϵ^2, and σ_η^2. We first introduce an auxiliary variable, $I_{t,h}$, to substitute missing observations. Then we obtain a conditional distribution of each parameter, given the other parameters and the values of the auxiliary variable. Similarly, we obtain a conditional distribution of the auxiliary variable, given the parameters. Finally, we use the Gibbs sampler to approximate joint and marginal distributions of the entire parameters and the auxiliary variable from these conditional distributions. See Kim and Nelson (1999) for more on Gibbs sampling.

2.2.1 Prior distributions

We choose the following priors for the unknown parameters. We adopt a flat prior for α and β. On the other hand, we assume that the priors for σ_ϵ^2 and σ_η^2, are more informative than the flat ones but still relatively diffused. Specifically, we assume that the prior of σ_ϵ^2 is given by

$$IG\left(\frac{\nu_1}{2}, \frac{\delta_1}{2}\right)$$

with $\nu_1 = 10$ and $\delta_1 = 0.0002$, implying that the mean of σ_ϵ is 0.0015 and that the 95 percent confidence interval is 0.0010 to 0.0025. The prior of σ_η^2 is given by

$$IG\left(\frac{\nu_2}{2}, \frac{\delta_2}{2}\right)$$

with $\nu_2 = 10$ and $\delta_2 = 0.35$, implying that the mean of σ_η is 0.2031 and that the 95 percent confidence interval is 0.1315 to 0.3291.[5]

[5]The mean of σ_ϵ and the mean of σ_η are chosen using the Japanese data. As for the mean of σ_ϵ, we use the standard error of the hourly log difference in the yen-dollar rate, which is equal to 0.0015. As for the mean of σ_η, we guess this not from the standard error of $I_{t,h}$, which is not observable, but from the standard error of I_t. Specifically, the standard error of I_t calculated using the entire observations is 0.0888, while the standard error of I_t calculated using only non-zero observations is 0.3174. The average of the two values, 0.2031, is used as the mean of σ_η.

2.2.2 Computational algorithm

The above assumptions about the priors and the data generating process provide us with posterior conditional distributions that are needed to implement Gibbs sampling. The following steps 1 through 5 are iterated to obtain the joint and marginal distributions of the parameters and the values of the auxiliary variable. The summations are taken from $(t,h) = (1,1)$ to $(T,24)$, unless otherwise stated.

Step 1 Generate α conditional on $s_{t,h}$, $I_{t,h}$ and σ_ϵ^2. We have the regression $s_{t,h} - s_{t,h-1} = \alpha I_{t,h} + \epsilon_{t,h}$. Hence, the posterior distribution is $\alpha \sim N(\phi_s, \omega_s)$, where $\phi_s = \sum I_{t,h}(s_{t,h} - s_{t,h-1}) / \sum I_{t,h}^2$ and $\omega_s = \sigma_\epsilon^2 / \sum I_{t,h}^2$.

Step 2 Generate σ_ϵ^2 conditional on $s_{t,h}$, $I_{t,h}$ and α. The posterior is $\sigma_\epsilon^2 \sim IG\left(\frac{\nu_s}{2}, \frac{\delta_s}{2}\right)$ where $\nu_s = \nu_1 + T$ and $\delta_s = \delta_1 + RSS_s$ with $RSS_s = \sum(s_{t,h} - s_{t,h-1} - \alpha I_{t,h})^2$.

Step 3 Generate β conditional on $s_{t,h}$, $I_{t,h}$ and σ_η^2. We have the regression $I_{t,h} = \beta(s_{t,h-1} - s_{t-1,h-1}) + \eta_{t,h}$. Hence, the posterior distribution is $\beta \sim N(\phi_I, \omega_I)$, where $\phi_I = \sum I_{t,h}(s_{t,h-1} - s_{t-1,h-1}) / \sum(s_{t,h-1} - s_{t-1,h-1})^2$ and $\omega_I = \sigma_\eta^2 / \sum(s_{t,h-1} - s_{t-1,h-1})^2$.

Step 4 Generate σ_η^2 conditional on $s_{t,h}$, $I_{t,h}$ and β. The posterior distribution is $\sigma_\eta^2 \sim IG\left(\frac{\nu_I}{2}, \frac{\delta_I}{2}\right)$ where $\nu_I = \nu_2 + T$ and $\delta_I = \delta_2 + RSS_I$ with $RSS_I = \sum(I_{t,h} - \beta(s_{t,h-1} - s_{t-1,h-1}))^2$.

Step 5 Generate $I_{t,h}$ conditional on $s_{t,h}$, I_t, α, β, σ_ϵ^2 and σ_η^2. Consider the case in which the aggregated intervention amount is not known. Then, the posterior distribution is as follows:

$$(I_{t,1}, ..., I_{t,24})' \sim N(\Xi_t, \Psi)$$

where $\Xi_t = (\xi_{t,1}, ..., \xi_{t,24})'$ and $\Psi = diag(\varphi, ..., \varphi)$ with $\varphi = (\frac{1}{\sigma_\eta^2} + \frac{\alpha^2}{\sigma_\epsilon^2})^{-1}$ and $\xi_{t,h} = (\varphi \frac{1}{\sigma_\eta^2})[\beta(s_{t,h-1} - s_{t-1,h-1})] + (\varphi \frac{\alpha^2}{\sigma_\epsilon^2})[\alpha^{-1}(s_{t,h} - s_{t,h-1})]$. Note that the expectation of $I_{t,h}$ is a weighted average of the two components, $\beta(s_{t,h-1} - s_{t-1,h-1})$ and $\alpha^{-1}(s_{t,h} - s_{t,h-1})$, with the weights being determined by σ_ϵ^2, σ_η^2, and α.[6] We consider the posterior distribution of $(I_{t,1}, ..., I_{t,23}, I_t)$. Note that when we know $(I_{t,1}, ..., I_{t,23}, I_t)$, the intervention in the last hour, $I_{t,24}$, is already determined. The posterior distribution is as follows:

$$(I_{t,1}, ..., I_{t,23}, I_t)' \sim N(\Xi_t^*, \Psi^*)$$

[6] Note that our method to deduct intraday timing works even if intervention is not effective at all. In an extreme case in which intervention is not effective at all, i.e., $\alpha = 0$, the expectation of $I_{t,h}$ is simply equal to $\beta(s_{t,h-1} - s_{t-1,h-1})$, implying that $I_{t,h}$ is estimated only from the reaction function.

where $\Xi_t^* = B\Xi_t$ and $\Psi^* = B\Psi B'$ with

$$B = \begin{bmatrix} 1 & 0 & ... & 0 \\ 0 & 1 & ... & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & ... & 1 \end{bmatrix}. \tag{5}$$

We can partition the matrices Ξ_t^* and Ψ^* as follows:

$$\Xi_t^* = \begin{bmatrix} \Xi_{t,1}^* \\ \Xi_{t,2}^* \end{bmatrix}, \ \Psi^* = \begin{bmatrix} \Psi_{11}^* & \Psi_{12}^* \\ \Psi_{21}^* & \Psi_{22}^* \end{bmatrix},$$

where $\Xi_{t,1}^*$ is 23×1, $\Xi_{t,2}^*$ is 1×1, Ψ_{11}^* is 23×23, Ψ_{12}^* is 23×1, Ψ_{21}^* is 1×23, and Ψ_{22}^* is 1×1. Finally, we can construct the posterior distribution of $(I_{t,1}, ..., I_{t,23})$ conditional on I_t as follows:

$$\left(I_{t,1}, ..., I_{t,23} | \ I_t\right)' \sim N\left(\Xi_{t,1}^* + \Psi_{12}^*(\Psi_{22}^*)^{-1}(I_t - \Xi_{t,2}^*), \Psi_{11}^* - \Psi_{12}^*(\Psi_{22}^*)^{-1}\Psi_{21}^*\right).$$

By generating the values of the auxiliary variable $I_{t,1}, ..., I_{t,23}$ from this posterior distribution conditional on the parameters, the hourly exchange rate, and the aggregated intervention, we can construct the intervention in the last hour as $I_{t,24} = I_t - \Sigma_{h=1}^{23} I_{t,h}$.

We iterate steps 1 through 5 $M + N$ times and discard the realizations of the first M iterations but keep the last N iterations to form a random sample of size N on which statistical inference can be made. M must be sufficiently large so that the Gibbs sampler converges. Also, N must be large enough to obtain the precise empirical distribution. In our simulations, we set $M = 2000$ and $N = 2000$ and run 3 independent Markov chains.

3 Simulation Analysis

In this section we conduct Monte Carlo simulations to evaluate the performance of our methodology. We start by assuming that the data generating process is given by equations (1) and (2) with $s_{0,24} = \ln(100)$, $\alpha = -0.015$, $\beta = 3.2$, $\sigma_\epsilon = 0.0015$, and $\sigma_\eta = 0.2031$. We borrow the estimates of α and β from the study on intervention in Japan by Kearns and Rigobon (2005):[7] $\alpha = -0.015$ implies that a 1 trillion yen intervention by the Japanese monetary authorities moves the yen/dollar rate by 1.5 percent; on the other hand, $\beta = 3.2$ implies that a one percent deviation of the exchange

[7] We divide their estimate of β by 24 to convert their estimate, which is based on a daily frequency, to one based on an hourly frequency.

rate from its target level causes the Japanese monetary authorities to intervene with 32 billion yen. Note that our framework is a classical one in that parameters are treated as unknown constants to be estimated.

We generate bivariate time series $\{s_{t,h}, I_{t,h}\}$ by (1) and (2). The length of the time series is set to 100 days ($T = 100$), and 500 replications of this length are generated. We repeat this for $T = 250$ and 500. We then estimate the unknown parameters under the following three cases. The first case is what we refer to as the "infeasible estimator." We assume that the hourly amount of intervention, $I_{t,h}$, is observable to researchers, and we simply apply OLS to the hourly data of intervention and exchange rates. This estimator can be seen as the best one (although it is infeasible), and will be used as a benchmark. The second case is what we refer to as the "naive OLS estimator," where we assume that intervention data is available only at the daily frequency, and we apply OLS to the daily intervention and exchange rate data. Specifically, we estimate equations (3) and (4) separately. This estimator suffers from the endogeneity problem, as explained earlier. The third case is what we refer to as the "MCMC estimator," where we assume that exchange rate data is available at the hourly frequency, but intervention data is available only at the daily frequency, and we apply our MCMC method to these data. The MCMC method provides us with a posterior distribution for each of the parameters. We use the mean of the distribution as a point estimate.

Table 1 presents the simulation results. We evaluate the performance of the three estimators in terms of the *Mean*, which is defined as the mean of estimated values of α and β over 500 replications, as well as the corresponding root mean squared error, which is denoted by $\sqrt{MSE}$. We see from the table that the infeasible estimators for α and β are close to the true values (i.e., $\alpha = -0.015$ and $\beta = 3.2$), and that the corresponding root mean squared errors are small. On the other hand, the naive OLS estimators perform badly; the estimate of α is of the opposite sign, and so is the estimate of β. Importantly, we see no clear sign of improvement in the performance of these estimators as sample size T increases. In contrast, the MCMC estimator performs as well as the infeasible estimator: the means of α and β are almost the same as those of the infeasible estimator, and although the root mean squared errors are slightly larger, the difference tends to become smaller as T increases.[8]

The MCMC estimators reported in Table 1 are obtained based on the assumption that the

[8] As a robustness check, we repeated simulations using different true parameter values for α and β, including the case of $\alpha = 0$. We confirmed that the MCMC estimator performs as well as the infeasible estimator.

disturbance terms, $\varepsilon_{t,h}$ and $\eta_{t,h}$, are both normally distributed. One may wonder to what extent the results depend on this assumption. To check this, we conduct the following experiment. We use the same data generating process as before, but we now assume that the disturbance terms follow t distributions with different degrees of freedom, or ARCH innovations with different degrees of persistence.[9] The results are shown in Table 2, in which the number in each cell represents the estimate of α for different specifications of $\varepsilon_{t,h}$ and $\eta_{t,h}$. We see that the estimate of α is quite close to the true value even when the disturbance terms are not Gaussian, suggesting that the MCMC estimator is robust to changes in the specification of the disturbance terms.

4 An Experiment Using Swiss Data

In the simulation analysis conducted in the previous section, we artificially generated hourly intervention data; however the monetary authorities in some countries (see footnote 2) disclose information about intraday interventions. Therefore, we are able to conduct a similar exercise using actual (not artificial) intraday intervention data. This is what we do in this section. Specifically, we conduct two different estimations regarding the efficacy of intervention: one with (actual) intraday intervention data, and the other one with aggregated daily intervention data. We then compare the two estimates in order to see how our method performs.

We use intraday intervention data disclosed by the Swiss National Bank. The Swiss National Bank discloses the amount of intervention with an up-to-the-minute time stamp for various currency pairs, including the Swiss franc (CHF) versus the US dollar (USD), the Swiss franc versus the German mark, and the US dollar versus the German mark, for the period of October 1986 to August 1995. We produce hourly intervention data for the pair CHF/USD for the period of January 1991 to August 1995.[10]

In applying our method to the Swiss data, we modify the policy reaction function described by equation (2) in the following way. Equation (2) implies that interventions are every-day events: namely, the central bank intervenes (by a small amount) even on quiet days when the exchange rate is fairly stable. But this is not consistent with the fact that interventions were carried out only on 1.2 percent of the total business days (namely, 20 out of 1735 business days), during the

[9]Note that the variance of the t distribution with d degrees of freedom is given by $d/(d-2)$ for $d > 2$. To adjust the variance, we multiply the disturbance term generated from the t distribution by $\sqrt{\sigma_\varepsilon^2(d-2)/d}$ or $\sqrt{\sigma_\eta^2(d-2)/d}$.

[10]Payne and Vitale (2003) suggest that there may exist a structural break somewhere around the year 1990, so we decided not to use the entire sample but a subsample for the period after 1990. The data we use is downloaded from the website operated by the Federal Reserve Bank of St. Louis (http://research.stlouisfed.org/fei/).

sample period. In this sense, Swiss interventions have an "all or nothing" property, suggesting that we need to incorporate some form of transaction costs associated with the conduct of interventions. Specifically, following Almekinders and Eijffinger (1996) and Ito and Yabu (2007), we assume that the Swiss monetary authorities have to pay some fixed costs on intervention days in the form of political costs. These political costs may include, for example, the costs incurred by the Swiss government in conducting negotiation with governments of relevant countries, as pointed out by Ito and Yabu (2007). The Swiss monetary authorities are assumed to compare the benefits of intervention (greater stability in the exchange rate) and the fixed costs they have to incur in implementing interventions. As is well known, the solution to this type of optimization with fixed costs is characterized by a state-dependent rule: namely, the monetary authorities carry out interventions only when the optimal level of intervention for that day exceeds a certain threshold. Specifically, we use a state-dependent rule of the form:

$$I_{t,h} = 1(|I^*_{t,1} - \mu_I| > c) I^*_{t,h} \tag{6}$$

$$I^*_{t,h} = \mu_I + \beta(s_{t,h-1} - s_{t-1,h-1}) + \eta_{t,h} \tag{7}$$

Equation (7) describes how the optimal level of intervention, $I^*_{t,h}$, is determined, while equation (6) represents a state-dependent policy reaction function, where $1(\cdot)$ is an indicator variable, which is equal to unity if the condition stated in the parentheses is satisfied and zero otherwise. In equation (7), we assume that the optimal level of intervention depends on the change in the exchange rate over the last 24 hours. In equation (6), we assume that intervention is carried out if the optimal level of intervention at the beginning of a day (i.e., 9am Swiss time), $I^*_{t,1}$, exceeds a prespecified threshold c, which is determined by the size of the political costs. Note that $I_{t,h}$ equals $I^*_{t,h}$ for any h as long as $I^*_{t,1}$ exceeds the threshold. In other words, once the monetary authorities decide to intervene on day t at the beginning of that day, they are allowed to intervene for every hour of day t without incurring any extra political costs. In this sense, the monetary authorities' decision on whether to intervene or not is made only once a day, although the amount of intervention for every hour of the day is decided during the day depending on fluctuations in the exchange rate over the course of the day.

We estimate the effect of intervention on the exchange rate using the model, consisting of equations (1), (6), and (7), as well as daily intervention data and hourly exchange rate data.[11] The result

[11]We employ an estimation procedure different from the one described in Section 2, which is for the model without

is presented in Table 3. We see from the table that the infeasible estimator of α is -0.198, implying that intervention of 10 million CHF moves the CHF/USD rate by 0.0198 percent.[12] On the other hand, the naive OLS estimator of α is -0.122, which is closer to zero than the infeasible estimator. Turning to the MCMC estimator, the estimate of α is -0.221, which is quite close to the infeasible estimator. Also note that the 95 percent confidence intervals for the infeasible estimator and the MCMC estimators overlap. This result confirms the finding in the previous section that the estimate produced by the MCMC method using daily observations of intervention is almost as precise as the one obtained by using hourly intervention data. As for the other parameters, the estimates of β and c are consistent with the theoretical prediction. In particular, the point estimate of β is positive and, as shown in the column labeled "Pr(< 0)," the probability that the estimate of β falls below zero is small, indicating that intraday intervention by the Swiss National Bank is characterized by a leaning-against-the-wind policy.

To investigate the precision of the MCMC estimator in greater detail, we look at whether the model correctly predicts the presence of intervention for hours in which the Swiss National Bank actually intervened, as well as whether the model correctly predicts the absence of intervention for hours in which the Swiss National Bank did not intervene. The result is presented in Table 4, in which we say that "the model predicts intervention" for a particular hour h on a particular day t if the 99 percent posterior interval of $I_{t,h}$ does not include zero, and otherwise we say "the model predicts no intervention." The number of days on which the Swiss National Bank intervened is 20 in the sample period, and the number of hours in which intervention took place in these 20 days is 29 (29 hours out of the 480 hours). We see from the table that the model successfully predicts the presence of intervention for 12 hours out of the 29 hours in which intervention actually took place, while it fails to do so for the remaining 17 hours. In other words, the ratio of *correct* signals turns out to be 0.413 (= 12/29). On the other hand, the model correctly predicts the absence of intervention for 427 hours out of the 451 hours in which intervention did not take place, and fails to do so for the remaining 24 hours. In other words, the ratio of *false* signals is 0.053 (= 24/427). Thus, the noise-to-signal ratio, which is calculated by dividing the ratio of false signals by the ratio

political costs. The estimation procedure for the model with political costs is given in the Appendix. As for the mean of σ_ϵ and the mean of σ_η, they are chosen in the same way as in Section 2 (see footnote 5) except that the Swiss data, rather than the Japanese data, is now used.

[12] Payne and Vitale (2003) estimate the effect of Swiss intervention using intraday intervention data and report that the *immediate* impact of an intervention of 50 million USD on the exchange rate is 0.15 percent, which is larger than the value implied by our estimate (i.e., 0.08 percent). A possible reason for this difference is that while Payne and Vitale (2003) estimate the immediate impact, our result in Table 3 shows the hourly impact.

of correct signals, turns out to be 0.128 (= 0.053/0.413).

Ito and Yabu (2007) estimate a policy reaction function for Japan using the daily intervention data and then calculate a similar noise-to-signal ratio to see how well their estimated reaction function tracks the actual intervention behavior. Specifically, they consider a situation in which a researcher, using the estimated policy reaction function, forecasts whether intervention will occur on day t or not at the beginning of that day. Note that the exercise Ito and Yabu (2007) conducted is to predict the presence or absence of intervention at the *daily* frequency, while the exercise we conduct in Table 4 is to predict the presence or absence of intervention at the *hourly* frequency, which doubtlessly is a more difficult exercise. They report that the ratio of correct signals and the ratio of false signals are 0.143 and 0.0094 respectively, so that the noise-to-signal ratio is 0.066. Comparing the results, our model outperforms theirs in terms of the ratio of correct signals, but it performs worse in terms of the ratio of false signals. As a result, the noise-to-signal ratio for our model is slightly higher than the one for their model.

5 Application to Japanese Data

5.1 Baseline specification

In this section, we apply our method to the Japanese data with daily observations of intervention and hourly observations of the yen/dollar exchange rate. Figures 1 and 2 show the hourly movement of the yen/dollar exchange rate and the daily amount of intervention implemented by the Japanese monetary authorities, both for the period from April 1991 to December 2002.[13] An important thing to note from Figure 2 is that there is a structural break somewhere around 1995: interventions are small in size but frequent during the period before 1995, while they are larger in size but less frequent during the period after 1995. As noted by, among others, Ito (2003), this break coincides with a change in the person in charge of the conduct of interventions in June that year.[14] Kearns and Rigobon (2005) make use of this shift in Japanese intervention policy as a key piece of information in identifying the effects of Japanese intervention on the yen/dollar rate.

[13] Note that our sample period does not include the period of "Great Intervention" in 2003 and 2004, during which the Japanese monetary authorities aggressively purchased US dollars and sold yen as part of their "quantitative easing" policy. We deliberately exclude this period, because, as shown by previous studies, the motivation for interventions was quite different from that in preceding periods. See Taylor (2006), Ito (2007), and Watanabe and Yabu (2009) for more on the intervention policy during this period.

[14] In Japan, decisions on exchange rate interventions fall under the aegis of the Internal Finance Bureau of the Ministry of Finance. On June 21, 1995, Eisuke Sakibara (subsequently known as "Mr. Yen") was appointed as Director General of the International Finance Bureau.

To incorporate this structural change in the policy reaction fucntion, we modify equation (6) as follows:

$$I_{t,h} = \begin{cases} 1(|I^*_{t,1} - \mu_I| > c_1)I^*_{t,h} & \text{for } t < T_B \\ 1(|I^*_{t,1} - \mu_I| > c_2)I^*_{t,h} & \text{for } t \geq T_B \end{cases} \tag{8}$$

where T_B is the break date (namely, June 1995), and c_1 and c_2 are different thresholds for the two subperiods. Here we assume that the change in the Japanese policy reaction function can be represented solely by a change in threshold c, or the size of political costs, and that the other parameters are identical across the two subperiods. We make this assumption simply to obtain empirical results that are comparable to those of Kearns and Rigobon (2005), whose identification method requires such an assumption. Note that our identification method does not require imposing this assumption.

5.2 Baseline result

In our baseline regressions, we use equations (1), (7) and (8). Table 5 presents the results. We run regressions with and without the lagged intervention term I_{t-1}, with the left half of the table showing the result without that term, and the right half showing that with that term. We see from the left hand side of the table that the coefficient on the intervention variable, α in equation (1), is negative and significantly different from zero. Note that the frequency of finding negative values, $\Pr(<0)$, equals unity, indicating that we never find positive values in 10,000 draws. The estimated value of α is equal to -0.0164, implying that a yen-selling (yen-buying) intervention of one trillion yen leads to a 1.64 percent depreciation (appreciation) of the yen. The result for the specification with the lagged intervention variable, which is reported on the right hand side of the table, is almost the same.

Our estimate regarding the impact of foreign exchange interventions is more than twice as large as that obtained in previous studies. Ito (2003), for example, applying OLS to daily data of Japanese interventions and the yen/dollar rate, arrived at a corresponding change of 0.6 percent for the sample period of April 1991 to March 2001 and 0.9 percent for the subperiod from June 1995 to March 2001. Similarly, Fratzscher (2005), applying a similar regression as Ito (2003) using daily data for the period 1990-2003, found that Japanese interventions of ten billion dollars, which is approximately equal to one trillion yen, moves the yen/dollar rate by 0.8 percent. Our much larger estimation result suggests that these previous studies suffer from the endogeneity problem, so that their estimates of the effectiveness of interventions on the exchange rate was biased toward zero.

Kearns and Rigobon (2005), who identified the effects of intervention by making use of the structural change in the policy reaction function, report that an intervention of one billion dollars moves the yen/dollar rate by 1.5 percent, which is relatively close to our estimate, although it is still outside our 95 percent posterior interval.

Turning to the coefficients in the policy reaction function, we find that the coefficient on the change in the exchange rate, β, is positive and significantly different from zero, indicating that a leaning-against-the-wind policy was adopted by the Japanese monetary authorities. We also find that the estimates of c_1 and c_2 are both positive as predicted, and more importantly, c_2 is significantly larger than c_1, providing an explanation of the fact that interventions during the latter sample period were larger but less frequent.

Our MCMC approach gives us a posterior distribution for the auxiliary variable, $I_{t,h}$, for each t and h. Figure 3 shows the estimate of this variable and the yen/dollar rate for each hour on April 10, 1998, when the Japanese monetary authorities purchased 2.6 trillion yen, the largest yen-buying intervention in our sample period. The solid line represents the mean of the posterior distribution of $I_{t,h}$, while the dotted lines indicate the 99 percent confidence interval. We see from the figure that the estimated hourly amount of intervention is almost always positive (i.e., almost all interventions were yen-buying interventions). The estimated hourly amount takes the largest value, 0.5 trillion yen, at 6-7am GMT (i.e., 2-3pm in Tokyo), and this is exactly the time when the yen exhibits a sharp appreciation and records its highest level on this day. This concurrence can be interpreted as evidence of aggressive yen-buying intervention during this hour causing a sharp appreciation.

Figure 4 shows the movement of the yen/dollar rate before and after the hour that a yen-selling intervention is carried out. To construct this figure, we collected the estimates of $I_{t,h}$ for 148 business days when yen-selling interventions were implemented. We then identified h when the estimate of $I_{t,h}$ is significantly different from zero (i.e., the 99 percent confidence interval of $I_{t,h}$ does not include zero). Note that $\tau = 0$ in the figure represents the hour of intervention and that the yen/dollar rates for other hours are divided by the exchange rate levels at the hour of intervention for normalization. The solid line represents the 50th percentile of the distribution of the exchange rate, while the two dotted lines represent the 40th and 60th percentiles, respectively. We see from the figure that there is a trend of yen appreciation prior to the hour of yen-selling intervention. The yen falls very quickly in response to the intervention and stays there for at least twelve hours after the intervention, indicating that interventions have a persistent effect on the level of the yen/dollar rate.

Finally, we decompose the yen-amount of intervention per business day into an extensive margin (i.e., the probability of intervention for a given day) and an intensive margin (i.e., the yen-amount per intervention day). Furthermore, we decompose the yen amount per intervention day into an extensive margin (the probability of intervention in a given hour on a day that interventions were conducted) and an intensive margin (the yen-amount per intervention hour). The results are shown in Table 6. We see from the first three rows of the table that the post-1995 period is characterized by a lower extensive margin and a higher intensive margin; this confirms what we saw in Figure 2. More importantly, we see from the last two rows that the larger yen-amount per intervention day in the post-1995 period comes partly from the larger extensive margin, but mostly from the larger intensive margin. These results indicate that the post-1995 period is characterized by a higher intensive margin not only at the daily frequency, but also at the hourly frequency.

5.3 Robustness to changes in the specification of the policy reaction function

An important feature of our MCMC approach is that it makes use of knowledge about the structure of the economy, which is represented by the equation for exchange rate dynamics and the equation for the policy reaction function. This implies that the performance of the entire estimation process crucially depends on whether the structure of the economy is properly specified or not. In this subsection we check the sensitivity of the baseline results to various changes in the specification of the policy reaction function.

5.3.1 Alternative assumptions regarding the policy lag

In the baseline case, we assumed that the central bank can respond to exchange rate changes that occurred at least an hour earlier, but not to changes less than an hour earlier. The baseline result might still suffer from the endogeneity problem if the central bank can respond in less than one hour.

To address this issue, we examine whether the baseline results are sensitive to changes in the frequency of observations. We consider the same structure of the model as in the baseline case (the model without the lagged intervention term) but change the frequency h at which the exchange rate is observed. Specifically, the frequencies considered here are: 12 hours, 8 hours, 6 hours, 4 hours, 2 hours, 1 hour, 30 minutes, and 15 minutes. For example, for the 30 minute frequency, we assume that the central bank is able to respond to exchange rate changes that occurred more than 30 minutes earlier, but not to change that took place less than 30 minutes earlier. Figure 5 shows the estimate of α and its 95 percent confidence interval for different frequencies chosen in the estimation. The figure shows that when the frequency is as low as 12 hours, the estimated value of α is -0.0076, which is very close to the estimate obtained by Ito (2003) using the daily data. As the frequency increases, the value of α converges to the baseline estimate. Importantly, the value of α does not differ much for $h = 1$ hour, $h = 30$ minutes, and $h = 15$ minutes, suggesting that the Japanese monetary authorities did not respond in less than one hour. Therefore, the frequency of one hour is sufficiently fine to eliminate the endogeneity problem.

5.3.2 Higher political costs at night

What time of the day were interventions conducted? Neely (2001) asked this question to 44 central banks and collected 22 responses. He provided the following options: "prior to normal business hours," "morning of the business day," "afternoon of the business day," and "after normal business hours." One of the interesting features we learn from the responses to this question is that about 56 percent of central banks answered that they *never* intervene "prior to normal business hours," and similarly about 35 percent answered that they *never* intervene "after normal business hours." Various pieces of anecdotal evidence regarding the intervention behavior of Japan's monetary authorities suggest that they are active during hours in which the Tokyo market is open, while they are much less active during other hours, which is more or less similar to what Neely's (2001) survey results indicate.

The fact that central banks seldom intervene during night hours may be interpreted as reflecting that the political costs are much higher at night than during the daytime, so that central banks hesitate to intervene at night even if the optimal level of intervention is not zero. Based on this line of reasoning, we assume that $I_{t,h}$ is equal to zero at night ($h = 9, \ldots, 24$, or between 5pm and 9am Tokyo time). Specifically, we replace equation (8) by:

$$I_{t,h} = \begin{cases} 1(|I^*_{t,1} - \mu_I| > c_1)I^*_{t,h} & \text{for} \quad h = 1, \ldots, 8, \text{ and} \quad t < T_B \\ 1(|I^*_{t,1} - \mu_I| > c_2)I^*_{t,h} & \text{for} \quad h = 1, \ldots, 8, \text{ and} \quad t \geq T_B \\ 0 & \text{for} \quad h = 9, \ldots, 24 \end{cases} \tag{9}$$

and repeat the same exercise as before. The regression result is presented in Table 7, showing that the baseline result is not sensitive to this change in the policy reaction function. That is, the coefficient associated with the effectiveness of intervention, α, is negative and significantly different from zero as before, although it is now a little smaller; the coefficient on the change in the exchange rate in the policy reaction function, β, is positive and significantly different from zero, suggesting again a leaning-against-the-wind policy; the coefficients related to the size of the political costs, c_1 and c_2, are both positive and significantly different from zero as before, and the political costs are significantly larger in the latter sample period.

5.3.3 Ito-Yabu (2007) specification

Equation (7), which is basically identical to the intervention function adopted by Kearns and Rigobon (2005), may be too simple to capture the details of Japanese intervention policy. Ito and Yabu (2007) derive a policy reaction function that can be regarded as a better approximation to the Japanese policy reaction function. Specifically, the optimal amount of intervention depends on the deviation of the actual exchange rate from its target level, which is determined by the weighted average of $s_{t-1,h-1}$, $s_{t-21,h-1}$, and $s^{MA}_{t-1,h-1}$, where $s^{MA}_{t,h}$ is defined as the moving average of the exchange rate over the last one year. To incorporate this idea into our model, we replace equation (7) by:

$$I^*_{t,h} = \mu_I + \beta_1(s_{t,h-1} - s_{t-1,h-1}) + \beta_2(s_{t,h-1} - s_{t-21,h-1}) + \beta_3(s_{t,h-1} - s^{MA}_{t,h-1}) + \eta_{t,h} \qquad (10)$$

Note that equation (7) is a special case of the above equation, with both β_2 and β_3 being equal to zero.

We conduct the same exercise as before and the results are presented in Table 8. We confirm that the coefficient associated with the effectiveness of intervention, α, is negative and significantly different from zero; the coefficients related to political costs, c_1 and c_2, are both positive and significantly different from zero. Turning to the new coefficients in the policy reaction function, β_1 and β_2 are both positive, but β_3 is very close to zero, indicating that Japan's monetary authorities take a leaning-against-the-wind policy with respect to changes in the exchange rate at the daily and monthly frequencies, but not at the annual frequency.

6 Conclusion

Estimating the effects of foreign exchange interventions is not an easy task because central banks react even to intraday changes in the exchange rate, while intervention data is usually available only at the daily frequency. In this paper, we therefore proposed a new methodology based on Markov Chain Monte Carlo simulation to cope with this endogeneity problem. We first conduct "imputation" or "data augmentation" to obtain intraday amounts of intervention and then estimate the efficacy of interventions using the augmented data. Although a number of previous studies have pursued the idea of augmenting observed low-frequency data with simulated high-frequency data by applying MCMC methods, especially in the area of finance, this paper is the first application of the MCMC approach to the endogeneity problem due to temporal aggregation.

We applied this method to Swiss and Japanese intervention data. The Swiss data is unique in that the amounts of intraday interventions with up-to-the-minute time stamps are available. Using the Swiss data, we conducted an experiment in which we first applied our method to aggregated daily intervention data to estimate the efficacy of intervention and then compared it with the estimate obtained using the hourly intervention data. We found that the two estimates are very close to each other, implying that endogeneity bias due to temporal aggregation is successfully eliminated by employing our method. Applying our method to Japanese data, we found that an intervention of one trillion yen moves the yen/dollar exchange rate by 1.7 percent, which is more than twice as much as the magnitude reported in previous studies applying OLS to daily observations. We interpreted this difference as highlighting the quantitative importance of the endogeneity problem due to temporal aggregation.

A Estimation procedure of the model with political costs

The methodology for estimating the model without political costs given by equations (1) and (2) is presented in Section 2.2. The purpose of this appendix is to provide details regarding the estimation of the model in which the policy reaction function is given by equations (6) and (7). The parameters to be estimated are μ_s, α, μ_I, β, c_1, c_2, σ_ϵ^2, and σ_η^2. In addition to these parameters, we estimate auxiliary variables, $I_{t,h}$ and $I_{t,h}^*$, for each h and t. A flat prior is adopted for μ_s, α, μ_I, β, c_1, and c_2. As for σ_ϵ^2 and σ_η^2, priors are the same as those used in the main text.

The posterior conditional distributions, which are needed to implement Gibbs Sampling, are obtained from the priors and the assumptions of the data generating process. The following steps 1 through 6 are iterated to obtain joint and marginal distributions of the parameters and the auxiliary variables.

Step 1 Generate μ_s and α conditional on $s_{t,h}$, $I_{t,h}$ and σ_ϵ^2. We have the regression $s_{t,h} - s_{t,h-1} = \mu_s + \alpha I_{t,h} + \epsilon_{t,h}$. Hence, the posterior distribution is $(\mu_s, \alpha)' \sim N(\phi_s, \omega_s)$ where $\phi_s = (X_s' X_s)^{-1} X_s' Y_s$ and $\omega_s = (X_s' X_s)^{-1} \sigma_\epsilon^2$ with the matrices $X_s = \{1, I_{t,h}\}$ and $Y_s = \{s_{t,h} - s_{t,h-1}\}$.

Step 2 Generate σ_ϵ^2 conditional on $s_{t,h}$, $I_{t,h}$, μ_s and α. The posterior is $\sigma_\epsilon^2 \sim IG\left(\frac{\nu_s}{2}, \frac{\delta_s}{2}\right)$ where $\nu_s = \nu_1 + T$ and $\delta_s = \delta_1 + RSS_s$ with $RSS_s = \sum (s_{t,h} - s_{t,h-1} - \mu_s - \alpha I_{t,h})^2$.

Step 3 Generate μ_I and β conditional on $s_{t,h}$, $I_{t,h}^*$ and σ_η^2. We have the regression $I_{t,h}^* = \mu_I + \beta(s_{t,h-1} - s_{t-1,h-1}) + \eta_{t,h}$. Hence, the posterior distribution is $(\mu_I, \beta)' \sim N(\phi_I, \omega_I)$ where $\phi_I = (X_I' X_I)^{-1} X_I' Y_I$ and $\omega_I = (X_I' X_I)^{-1} \sigma_\eta^2$ with the matrices $X_I = \{1, s_{t,h-1} - s_{t-1,h-1}\}$ and $Y_I = \{I_{t,h}^*\}$.

Step 4 Generate σ_η^2 conditional on $s_{t,h}$, $I_{t,h}^*$, μ_I and β. The posterior is $\sigma_\eta^2 \sim IG\left(\frac{\nu_I}{2}, \frac{\delta_I}{2}\right)$ where $\nu_I = \nu_2 + T$ and $\delta_I = \delta_2 + RSS_I$ with $RSS = \sum (I_{t,h}^* - \mu_I - \beta(s_{t,h-1} - s_{t-1,h-1}))^2$.

Step 5 Generate $I_{t,h}$ and $I_{t,h}^*$ conditional on $s_{t,h}$, I_t, μ_s, α, μ_I, β, c_1, c_2, σ_ϵ^2 and σ_η^2. Consider the case without the political costs. The posterior distribution without knowing I_t is as follows:

$$(I_{t,1}, ..., I_{t,24})' \sim N(\Xi_t, \Psi)$$

where $\Xi_t = (\xi_{t,1}, ..., \xi_{t,24})'$ and $\Psi = diag(\varphi, ..., \varphi)$ with $\varphi = (\frac{1}{\sigma_\eta^2} + \frac{\alpha^2}{\sigma_\epsilon^2})^{-1}$ and $\xi_{t,h} = (\varphi \frac{1}{\sigma_\eta^2}) [\mu_I + \beta(s_{t,h-1} - s_{t-1,h-1})] + (\varphi \frac{\alpha^2}{\sigma_\epsilon^2}) \left[\alpha^{-1}(s_{t,h} - s_{t,h-1} - \mu_s)\right]$. Hence, the posterior distribution of $(I_{t,1},...,I_{t,23},I_t)$ is as follows:

$$(I_{t,1},...,I_{t,23},I_t)^{'} \sim N\left(\Xi^*,\Psi^*\right)$$

where $\Xi_t^* = B\Xi_t$ and $\Psi^* = B\Psi B'$ with B defined by (5). We can partition the matrices Ξ_t^* and Ψ^* as follows:

$$\Xi_t^* = \begin{bmatrix} \Xi_{t,1}^* \\ \Xi_{t,2}^* \end{bmatrix}, \quad \Psi^* = \begin{bmatrix} \Psi_{11}^* & \Psi_{12}^* \\ \Psi_{21}^* & \Psi_{22}^* \end{bmatrix}$$

where $\Xi_{t,1}^*$ is 23×1, $\Xi_{t,2}^*$ is 1×1, Ψ_{11}^* is 23×23, Ψ_{12}^* is 23×1, Ψ_{21}^* is 1×23, and Ψ_{22}^* is 1×1. Then we can construct the posterior distribution of $(I_{t,1},...,I_{t,23})$ conditional on I_t as follows:

$$(I_{t,1},...,I_{t,23}|I_t)^{'} \sim N\left(\widehat{\Xi},\hat{\Psi}\right)$$

where $\widehat{\Xi} = \Xi_1^* + \Psi_{12}^*(\Psi_{22}^*)^{-1}(I_t - \Xi_2^*)$ and $\hat{\Psi} = \Psi_{11}^* - \Psi_{12}^*(\Psi_{22}^*)^{-1}\Psi_{21}^*$. We can partition the matrices $\widehat{\Xi}$ and $\hat{\Psi}$ as follows:

$$\widehat{\Xi}_t = \begin{bmatrix} \widehat{\Xi}_{t,1} \\ \widehat{\Xi}_{t,2} \end{bmatrix}, \quad \hat{\Psi} = \begin{bmatrix} \hat{\Psi}_{11} & \hat{\Psi}_{12} \\ \hat{\Psi}_{21} & \hat{\Psi}_{22} \end{bmatrix}$$

where $\widehat{\Xi}_{t,1}$ is 1×1, $\widehat{\Xi}_{t,2}$ is 22×1, $\hat{\Psi}_{11}$ is 1×1, $\hat{\Psi}_{12}$ is 1×22, $\hat{\Psi}_{21}$ is 22×1, and Ψ_{22}^* is 22×22. Then the posterior distribution of $I_{t,1}$ conditional on I_t is $I_{t,1}|\ I_t \sim N(\widehat{\Xi}_{t,1},\hat{\Psi}_{11})$. The posterior distribution of $(I_{t,2},...,I_{t,23})^{'}$ conditional on I_t and $I_{t,1}$ is as follows:

$$(I_{t,2},...,I_{t,23}|\ I_t,I_{t,1})^{'} \sim N\left(\widehat{\Xi}_{t,2} + \hat{\Psi}_{21}(\hat{\Psi}_{11})^{-1}(I_t - \widehat{\Xi}_{t,1}),\hat{\Psi}_{22} - \hat{\Psi}_{21}(\hat{\Psi}_{11})^{-1}\hat{\Psi}_{12}\right) \tag{11}$$

Since we have the political costs, $I_{t,h}$ and $I_{t,h}^*$ are generated from the following:

$t \in \{\ I_t \neq 0,\ t < T_B\}$: Generate $I_{t,1}$ from a truncated normal distribution such as $N(\widehat{\Xi}_{t,1},\hat{\Psi}_{11})$ conditional on $|I_{t,1} - \mu_I| > c_1$. Then generate $(I_{t,2},...,I_{t,23})^{'}$ from (11) and construct $I_{t,24} = I_t - \sum_{h=1}^{23} I_{t,h}$. Set $I_{t,h}^* = I_{t,h}$ for $h = 1,...,24$.

$t \in \{I_t = 0,\ t < T_B\}$: Generate $I_{t,1}^*$ from a truncated normal distribution such as $N(\mu_I + \beta_1(s_{t-1,24} - s_{t-2,24}),\sigma_\eta^2)$ conditional on $|I_{t,1}^* - \mu_I| < c_1$. Then generate $I_{t,h}^*$ from $N(\mu_I + \beta_1(s_{t,h-1} - s_{t-1,h-1}),\sigma_\eta^2)$ for $h = 2,...,24$. Set $I_{t,h} = 0$ for $h = 1,...,24$.

$t \in \{I_t \neq 0,\ t \geq T_B\}$: Generate $I_{t,1}$ from a truncated normal distribution such as $N(\widehat{\Xi}_{t,1},\hat{\Psi})$ conditional on $|I_{t,1} - \mu_I| > c_2$. Then generate $(I_{t,2},...,I_{t,23})^{'}$ from (11) and construct $I_{t,24} = I_t - \sum_{h=1}^{23} I_{t,i}$. Set $I_{t,h}^* = I_{t,h}$ for $h = 1,...,24$.

$t \in \{I_t = 0,\ t \geq T_B\}$: Generate $I^*_{t,1}$ from a truncated normal distribution such as $N(\mu_I + \beta_1(s_{t-1,24} - s_{t-2,24}), \sigma^2_\eta)$ conditional on $|I^*_{t,1} - \mu_I| < c_2$. Then, generate $I^*_{t,h}$ from $N(\mu_I + \beta_1(s_{t,h-1} - s_{t-1,h-1}), \sigma^2_\eta)$ for $h = 2, ..., 24$. Set $I_{t,h} = 0$ for $h = 1, ..., 24$.

Step 6 Generate c_1 and c_2 conditional on μ_s, α, μ_I, β, σ^2_ϵ and σ^2_η. Define the cumulative distribution functions of $N(\widehat{\Xi}_{t,1}, \widehat{\Psi}_{11})$ and $N(\mu_I + \beta_1(s_{t-1,24} - s_{t-2,24}), \sigma^2_\eta)$ as $\Phi^{I_t \neq 0}_t$ and $\Phi^{I_t = 0}_t$, respectively. The posterior distribution of c_1 is

$$\Pi_{t<T_B} \left[\Phi^{I_t=0}_t(c_1 + \mu_I) - \Phi^{I_t=0}_t(-c_1 + \mu_I)\right]^{1(I_t=0)} \left[1 - \Phi^{I_t\neq 0}_t(c_1 + \mu_I) + \Phi^{I_t\neq 0}_t(-c_1 + \mu_I)\right]^{1(I_t\neq 0)}.$$

Similarly, the posterior distribution of c_2 is

$$\Pi_{t\geq T_B} \left[\Phi^{I_t=0}_t(c_2 + \mu_I) - \Phi^{I_t=0}_t(-c_2 + \mu_I)\right]^{1(I_t=0)} \left[1 - \Phi^{I_t\neq 0}_t(c_2 + \mu_I) + \Phi^{I_t\neq 0}_t(-c_2 + \mu_I)\right]^{1(I_t\neq 0)}.$$

These densities are intractable and hence we implement the Metropolis-Hastings algorithm to draw from them.

We iterate steps 1 through 6 $M + N$ times and discard the realizations of the first M iterations but keep the last N iterations to form a random sample of size N on which statistical inference can be made. We set $M = 10,000$ and $N = 10,000$.

References

[1] Almekinders, G. J. and S. C. W. Eijffinger (1996), "A Friction Model of Daily Bundesbank and Federal Reserve Intervention," *Journal of Banking and Finance* 20(8), 1365-1380.

[2] Chaboud, A. P. and O. F. Humpage (2005), "An Assessment of the Impact of Japanese Foreign Exchange Intervention: 1991-2004," *International Finance Discussion Paper* No. 824, Board of Governors of the Federal Reserve System.

[3] Chang, Y. and S. J. Taylor (1998), "Intraday Effects of Foreign Exchange Intervention by the Bank of Japan," *Journal of International Money and Finance* 17(1), 191-210.

[4] Chow, G. C. and A. Lin (1971), "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series," *Review of Economics and Statistics* 53(4), 372-375.

[5] Christiano, L. J. and M. Eichenbaum (1987), "Temporal Aggregation and Structural Inference in Macroeconomics," *Carnegie-Rochester Series on Public Policy* 26(1), 63-130.

[6] Dominguez, K. and J. Frankel (1993), *Does Foreign Exchange Intervention Work?* Washington DC: Institute for International Economics.

[7] Dominguez, K. (2003), "Foreign Exchange Intervention: Did It Work in the 1990s?" in F. Bergsten and J. Williamson (eds.), *Dollar Overvaluation and the World Economy*, Washington, D.C.: Institute for International Economics.

[8] Edison, H. J. (1993), *The Effectiveness of Central-Bank Intervention: A Survey of the Literature After 1982*, Special Papers in International Economics, no. 18, Princeton, N.J.: Princeton University.

[9] Eraker, B. (2001), "MCMC Analysis of Diffusion Models with Application to Finance," *Journal of Business and Economic Statistics* 19(2), 177-191.

[10] Fatum, R. (2008) "Bank of Japan Intervention in the JPY/USD Exchange Rate Markets: Is It Effective and Through Which Channel Does It Work?" University of Alberta.

[11] Fatum, R. and M. Hutchison (2003), "Is Sterilized Foreign Exchange Intervention Effective After All? An Event Study Approach," *Economic Journal* 113(487), 390-411.

[12] Fatum, R. and M. Hutchison (2006), "Effectiveness of Official Daily Foreign Exchange Market Intervention Operations in Japan," *Journal of International Money and Finance* 25(2), 199-219.

[13] Fatum, R. and J. Pedersen (2009), "Real-Time Effects of Central Bank Intervention in the Euro Market," *Journal of International Economics* 78, 11-20.

[14] Fischer, A. M. and M. Zurlinden (1999), "Exchange Rate Effects of Central Bank Interventions: An Analysis of Transaction Prices," *Economic Journal* 109(458), 662-676.

[15] Fratzscher, M. (2005), "How Successful Are Exchange Rate Communication and Interventions? Evidence from Time-Series and Event-Study Approaches," *ECB Working Paper* No. 528.

[16] Galati, G., W. Melick, and M. Micu (2005), "Foreign Exchange Market Intervention and Expectations: The Dollar/Yen Exchange Rate," *Journal of International Money and Finance* 24(6), 982-1011.

[17] Gelman, A, J. B. Carlin, H. S. Stern, and D. B. Rubin (2003), *Bayesian Data Analysis*, CRC Press.

[18] Ito, T. (2003), "Is Foreign Exchange Intervention Effective? The Japanese Experiences in the 1990s," in Paul Mizen (ed.), *Monetary History, Exchange Rates and Financial Markets, Essays in Honor of Charles Goodhart*, Volume 2, Cheltenham, UK: Edward Elgar.

[19] Ito, T. (2007), "Myths and Reality of Foreign Exchange Interventions: An Application to Japan," *International Journal of Finance and Economics* 12, 133-154.

[20] Ito, T. and T. Yabu (2007), "What Prompts Japan to Intervene in the Forex Market: A New Approach to a Reaction Function," *Journal of International Money and Finance* 26, 193-212.

[21] Kearns, J. and R. Rigobon (2005), "Identifying the Efficacy of Central Bank Interventions: Evidence from Australia and Japan," *Journal of International Economics* 66(1), 31-48.

[22] Kim, C. J. and C. Nelson (1999), *State-Space Models with Regime Switching*, Cambridge: MIT Press.

[23] McCrorie, J. R. and M. J. Chambers (2006), "Granger Causality and the Sampling of Economic Processes," *Journal of Econometrics* 132(2), 311-336.

[24] Neely, C. J. (2001), "The Practice of Central Bank Intervention: Looking Under the Hood," Federal Reserve Bank of St. Louis *Review*, May/June 2001, 83(3), 1-10.

[25] Neely, C. J. (2005), "An analysis of recent studies of the effect of foreign exchange intervention," Federal Reserve Bank of St. Louis Review November/December 2005, 87(6), 685–717.

[26] Payne, R. and P. Vitale (2003), "A Transaction Level Study of the Effects of Central Bank Intervention on Exchange Rates," *Journal of International Economics* 61(2), 331-352.

[27] Sarno, L. and M. P. Taylor (2003), *Economics of Exchange Rates*, Cambridge: Cambridge University Press.

[28] Sims, C. A. (1971), "Discrete Approximations to Continuous Time Lag Distributions in Econometrics," *Econometrica* 39(3), 545-564.

[29] Tanner, M. A. and W. H. Wong (1987), "The Calculation of Posterior Distributions by Data Augmentation," *Journal of the American Statistical Association* 82(198), 528-540.

[30] Taylor, J. B. (2006), "Lessons from the Recovery from the "Lost Decade" in Japan: The Case of the Great Intervention and Monetary Injection," Paper presented at the International Conference of the Economic and Social Research Institute, Cabinet Office, Government of Japan.

[31] Watanabe, T. and K. Harada (2006), "Effects of the Bank of Japan's Interventions on Yen/Dollar Exchange Rate Volatility," *Journal of the Japanese and International Economies* 20(1), 99-111.

[32] Watanabe, T. and T. Yabu (2009), "The Great Intervention and Massive Money Injection: The Japanese Experience 2003-2004," Hitotsubashi University.

Table 1: Finite Sample Properties of the Three Estimators

		Infeasible estimator		Naive OLS estimator		MCMC estimator	
		$Mean$	$\sqrt{MSE}$	$Mean$	$\sqrt{MSE}$	$Mean$	$\sqrt{MSE}$
$T = 100$	α	-0.0150	0.0013	0.0082	0.0021	-0.0150	0.0017
	β	3.2199	0.0450	-0.9848	6.9027	3.2416	0.0777
$T = 250$	α	-0.0151	0.0011	0.0088	0.0012	-0.0149	0.0015
	β	3.2114	0.0376	-0.7797	3.9145	3.2251	0.0626
$T = 500$	α	-0.0152	0.0006	0.0091	0.0008	-0.0155	0.0006
	β	3.2094	0.0237	-1.0747	1.7552	3.2050	0.0378

Note: The data generating process is given by equations (1) and (2) with $\alpha = -0.015$ and $\beta = 3.2$. "Mean" is defined as the mean of estimators over 500 replications. "$\sqrt{MSE}$" represents the root mean squared error for each estimator. We estimate three chains from independent starting points in each replication. Each chain runs $4,000$ draws and the first $2,000$ are discarded as the burn-in-phase.

Table 2: MCMC Estimators of α for Different Error Distributions

		Distribution of $\eta_{t,h}$					
		$N(0, \sigma_\eta^2)$	t_{10}	t_5	t_3	$ARCH(0.3)$	$ARCH(0.85)$
	$N(0, \sigma_\varepsilon^2)$	-0.0150	-0.0153	-0.0157	-0.0152	-0.0157	-0.0154
	t_{10}	-0.0154	-0.0144	-0.0149	-0.0148	-0.0152	-0.0152
Distribution	t_5	-0.0150	-0.0159	-0.0150	-0.0154	-0.0143	-0.0152
of $\varepsilon_{t,h}$	t_3	-0.0146	-0.0144	-0.0151	-0.0153	-0.0162	-0.0150
	$ARCH(0.3)$	-0.0151	-0.0148	-0.0146	-0.0150	-0.0152	-0.0151
	$ARCH(0.85)$	-0.0159	-0.0152	-0.0165	-0.0154	-0.0149	-0.0170

Note: The data generating process is given by equations (1) and (2) with $\alpha = -0.015$ and $\beta = 3.2$. We consider various distributions to generate the disturbance terms. We set σ_ε and σ_η to 0.01 and 0.1, respectively. We multiply the disturbance terms generated from a t distribution by $\sqrt{\sigma_\varepsilon^2(1-d)/d}$ or $\sqrt{\sigma_\eta^2(1-d)/d}$. On the other hand, if $\varepsilon_{t,h}$ is ARCH(λ), $\varepsilon_{t,h}$ follows $N(0, \sigma_{t,h}^2)$, where $\sigma_{t,h}^2 = (1-\lambda)\sigma_\varepsilon^2 + \lambda\varepsilon_{t,h-1}^2$ with $\lambda = 0.3, 0.85$. Each chain runs $4,000$ draws and the first $2,000$ are discarded as the burn-in-phase. We estimate three chains from independent starting points in each replication.

Table 3: Interventions by the Swiss National Bank

	Infeasible estimator	Naive OLS estimator	MCMC estimator	
			Mean	Pr(< 0)
α	-0.198	-0.122	-0.221	1.000
	[-0.234, -0.152]	[-0.278, 0.034]	[-0.269, -0.175]	
β			0.050	0.053
			[-0.010, 0.105]	
c			0.023	0.000
			[0.020, 0.028]	

Note: The values in brackets are the 95 percent confidence intervals of the parameters. The column labeled "Mean" shows the mean of the marginal distribution of a parameter. The column labeled "Pr(< 0)" shows the frequency of finding negative values for a parameter. The MCMC estimation is conducted by five chains from independent starting points. Each chain runs $20,000$ draws and the first half is discarded as the burn-in-phase.

Table 4: Does the Model Predict Interventions Correctly?

	Hours in which the SNB intervened	Hours in which the SNB did not intervene
The model predicts intervention	12	24
The model predicts no intervention	17	427
Total	29	451

Note: We say that "the model predicts intervention" in a particular hour h if the 99 percent posterior interval of $I_{t,h}$ does not include zero. Otherwise, we say that "the model does not predict intervention" in that hour.

Table 5: Japanese Intervention: Baseline Results

	Without lagged intervention term				With lagged intervention term			
	Mean	Std. Dev.	Pr(< 0)	$\hat{R}$	Mean	Std. Dev.	Pr(< 0)	$\hat{R}$
	Equation for exchange rate dynamics							
α	-0.0164	0.0008	1.000	1.18	-0.0167	0.0006	1.000	1.09
	[-0.0180, -0.0148]				[-0.0179, -0.0154]			
	Equation for policy reaction function							
β	0.2556	0.0801	0.000	1.01	0.2414	0.0714	0.000	1.00
	[0.1045, 0.4168]				[0.1000, 0.3852]			
ρ					0.0143	0.0087	0.050	1.00
					[-0.0027, 0.0315]			
c_1	0.1089	0.0059	0.000	1.16	0.1070	0.0046	0.000	1.07
	[0.0983, 0.1210]				[0.0974, 0.1163]			
c_2	0.1820	0.0100	0.000	1.16	0.1789	0.0076	0.000	1.07
	[0.1633, 0.2021]				[0.1656, 0.1975]			

Note: Constants are estimated but not reported. The columns labeled "Mean" and "Std. Dev." refer to the mean and standard deviation of the marginal distribution of a parameter. The columns labeled "Pr(< 0)" refer to the frequency of finding negative values. The columns labeled $\hat{R}$ refer to the Gelman-Rubin statistic to monitor the convergence of the Markov chains. $\hat{R} < 1.1$ is considered as a sign of convergence. The values in brackets are the 95 percent posterior bands of the parameter. We estimate five chains from independent starting points. Each chain runs $20,000$ draws and the first half is discarded as the burn-in-phase.

Table 6: Intensive and Extensive Margins of Japanese Interventions

	1991-2002	1991-1995 (A)	1995-2002 (B)	B/A
Yen-amount per business day [trillion]	0.010	0.007	0.012	1.84
Probability of intervention day	0.070	0.149	0.025	0.16
Yen-amount per intervention day [trillion]	0.155	0.047	0.519	11.06
Probability of intervention hour	0.084	0.064	0.151	2.35
Yen-amount per intervention hour [trillion]	0.077	0.030	0.143	4.69

Note: "Yen-amount of intervention per business day" is defined as the total amount of intervention during the observation period divided by the number of business days. "Probability of intervention day" is defined as the number of intervention days divided by the number of business days. "Yen-amount per intervention day" is defined as the total amount of intervention during the observation period divided by the number of intervention days. "Probability of intervention hour" is defined as the number of intervention hours divided by the number of intervention days multiplied by 24. "Yen-amount per intervention hour" is defined as the total amount of intervention during the observation period divided by the number of intervention hours.

Table 7: Intervention Only During the Daytime

	Without lagged intervention				With lagged intervention			
	Mean	Std. Dev.	Pr(< 0)	$\hat{R}$	Mean	Std. Dev.	Pr(< 0)	$\hat{R}$
				Equation for exchange rate dynamics				
α	-0.0114	0.0006	1.000	1.10	-0.0114	0.0005	1.000	1.18
	[-0.0125, -0.0104]				[-0.0124, -0.0105]			
				Equation for policy reaction function				
β	0.5727	0.1916	0.001	1.02	0.5220	0.1855	0.000	1.00
	[0.1935, 0.9481]				[0.1756, 0.9170]			
ρ					0.0428	0.0216	0.028	1.00
					[-0.0008, 0.0850]			
c_1	0.1534	0.0083	0.000	1.11	0.1534	0.0071	0.000	1.26
	[0.1375, 0.1698]				[0.1405, 0.1672]			
c_2	0.2621	0.0132	0.000	1.12	0.2618	0.0115	0.000	1.24
	[0.2368, 0.2882]				[0.2395, 0.2851]			

Note: Constants are estimated but not reported. The columns labeled "Mean" and "Std. Dev." refer to the mean and standard deviation of the marginal distribution of a parameter. The columns labeled "Pr(< 0)" refer to the frequency of finding negative values. The columns labeled $\hat{R}$ refer to the Gelman-Rubin statistic to monitor the convergence of the Markov chains. $\hat{R} < 1.1$ is considered as a sign of convergence. The values in brackets are the 95 percent posterior bands of the parameter. We estimate five chains from independent starting points. Each chain runs $20,000$ draws and the first half is discarded as the burn-in-phase.

Table 8: Ito-Yabu (2007) Specification

	Without lagged intervention				With lagged intervention			
	Mean	Std. Dev.	Pr(<0)	$\hat{R}$	Mean	Std. Dev.	Pr(<0)	$\hat{R}$
	Equation for exchange rate dynamics							
α	-0.0160 [-0.0173, -0.0150]	0.0006	1.000	1.06	-0.0161 [-0.0180, -0.0148]	0.0008	1.000	1.08
	Equation for policy reaction function							
β_1	0.1923 [0.0257, 0.3571]	0.0856	0.012	1.00	0.2020 [0.0241, 0.4027]	0.0932	0.015	1.00
β_2	0.0849 [-0.0232, 0.1937]	0.0566	0.072	1.00	0.0633 [-0.0557, 0.1721]	0.0586	0.137	1.00
β_3	-0.0109 [-0.0655, 0.0373]	0.0253	0.665	1.00	-0.0053 [-0.0548, 0.0460]	0.0251	0.567	1.02
ρ					0.0128 [-0.0062, 0.0317]	0.0096	0.092	1.00
c_1	0.1117 [0.1039, 0.1194]	0.0041	0.000	1.05	0.1107 [0.0963, 0.1217]	0.0062	0.000	1.07
c_2	0.1868 [0.1717, 0.1996]	0.0075	0.000	1.06	0.1862 [0.1631, 0.2024]	0.0095	0.000	1.08

Note: Constants are estimated but not reported. The columns labeled "Mean" and "Std. Dev." refer to the mean and standard deviation of the marginal distribution of a parameter. The columns labeled "Pr(<0)" refer to the frequency of finding negative values. The columns labeled $\hat{R}$ refer to the Gelman-Rubin statistic to monitor the convergence of the Markov chains. $\hat{R} < 1.1$ is considered as a sign of convergence. The values in brackets are the 95 percent posterior bands of the parameter. We estimate five chains from independent starting points. Each chain runs $20,000$ draws and the first half is discarded as the burn-in-phase.

Figure 1: Hourly Fluctuations in the Yen-Dollar Rate

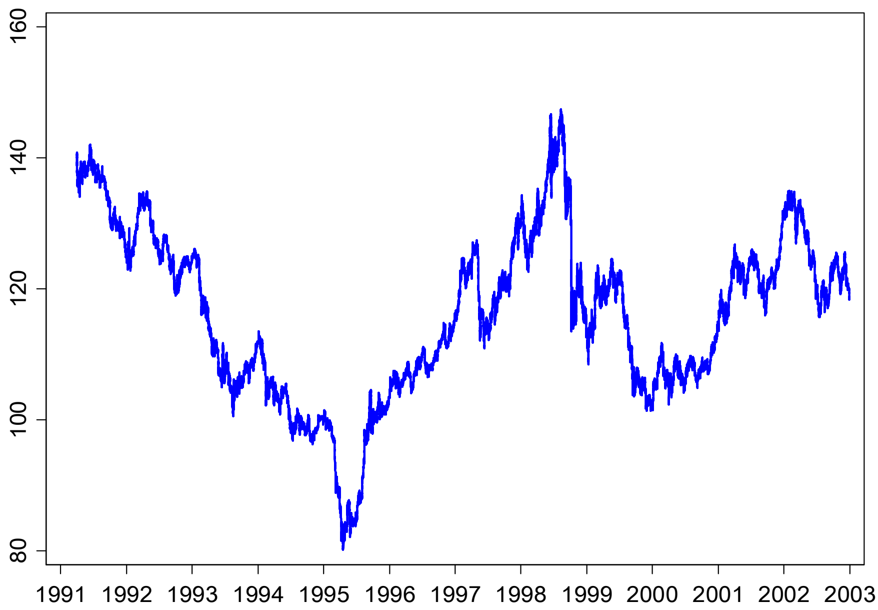

Figure 2: Daily Amounts of Intervention by Japan's Monetary Authorities

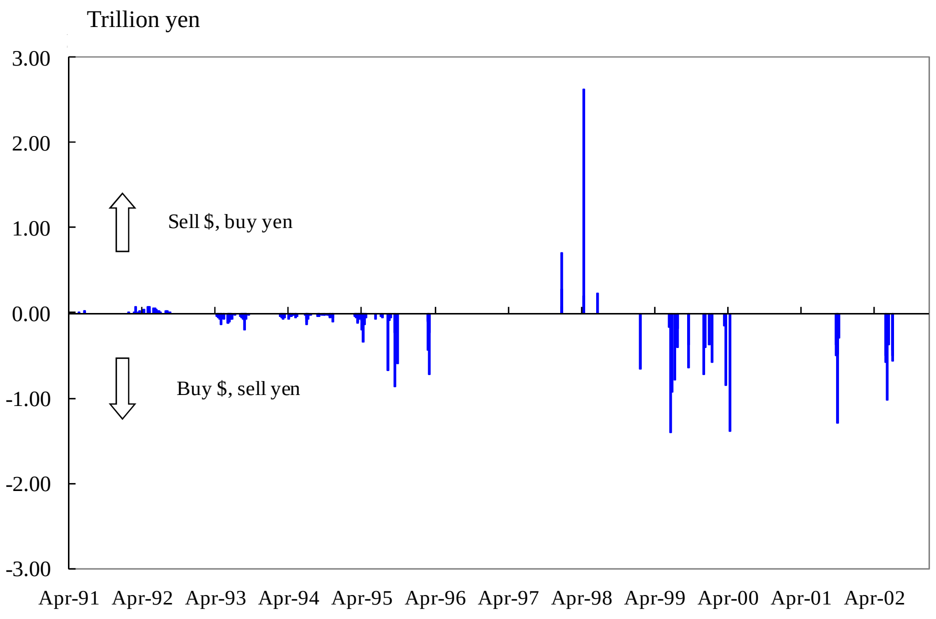

Figure 3: Estimated Hourly Amounts of Intervention on April 10, 1998

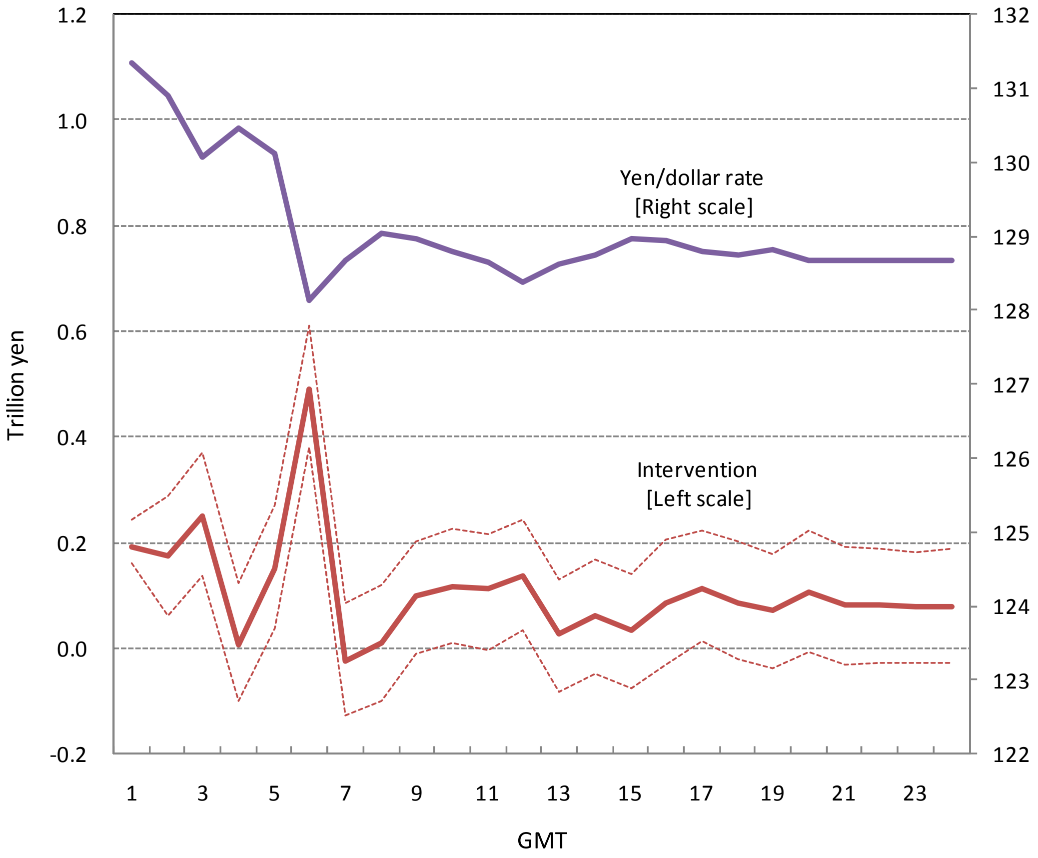

Figure 4: Exchange Rates Before and After Intervention

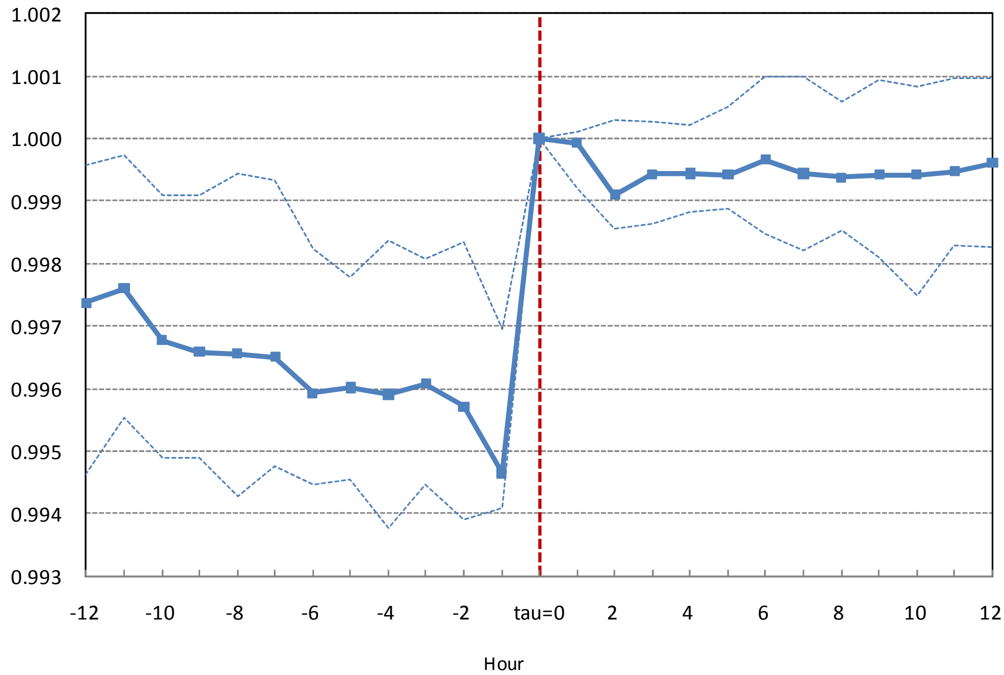

Figure 5: The Estimate of α for Different Data Frequencies