

Parameter Stability in Foreign Exchange Market Microstructure

Rafael Romeu*

(First Draft: Comments Welcome)

February 2001

Abstract

Microstructure models of foreign exchange markets emphasize two different channels of pressure in pricing dynamics; these are inventory and asymmetric information effects. Past empirical studies for foreign exchange markets have found evidence supporting the existence of both effects in the Dollar/DM market. This paper tests for the presence of within sample variation in the intensity of these effects and whether these are related to market conditions external to the pricing agent (the dealer). We test for and find multiple structural changes at conventional significance level. We incorporate the parameter instability into the estimated equations of foreign exchange. The new estimates reveal a pattern in which asymmetric information effects and inventory effects on the price of foreign exchange vary in intensity with external market conditions, and are dependent on these conditions.

* Department of Economics, University of Maryland 20742-7311; romeu@wam.umd.edu; 301.484.1588.

I. Introduction

Standard microstructure models of foreign exchange markets emphasize two different channels of pressure in pricing dynamics; these are inventory and asymmetric information effects (O'Hara 1995, Lyons 2000). The existence of inventory effects in pricing of foreign exchange implies that a dealer will adjust her price to counteract the accumulation of unwanted inventories. Similarly, the existence of asymmetric information effects implies that the dealer adjusts prices when trading to reflect the possibility that the opposite party has superior information about the future value of the asset. If either inventory effects or asymmetric information effects are present in the market, the dealer will increase her price when a buyer initiates a trade with her (i.e. transacted at the dealer's offer), and will lower prices when a seller initiates a trade with her (i.e. transacted at the dealer's bid).

The Lyons (1995) paper used transaction data to estimate a model of microstructure effects in the pricing decisions of a foreign exchange market dealer. This paper found both of the aforementioned effects present in the data. Building on this result, we look at whether these effects are muted or magnified in subsamples, and whether this is related to external market conditions. By looking at a variety of factors to determine what the market conditions may be external to the dealer, we find evidence robust to sample selection bias that the estimates of the inventory and asymmetric information effects vary with market conditions. Hence, we find evidence that the Lyons (1995) results confirming the microstructural hypothesis in foreign exchange exist in the market with

varying degrees of intensity. We find that the degree of intensity is related to the conditions that the dealer is facing in the overall market.

The rest of the paper is divided as follows. Section II gives a brief description of the theoretical framework used to describe the microstructure effects of the foreign exchange market. Section III describes the data used in the paper. In Section IV we model the effect of external conditions on price as being governed by a hidden state variable that is governed by a Markov process. In Section V we look for evidence of within sample variation of the data based on rolling regressions and other basic statistical evidence of parameter instability. In Section VI we look at statistical tests of breaks in the estimation based on the work of Bai and Perron (1998) and Bai (1999). These tests are very powerful in that they can reject among alternative hypotheses of multiple breaks, and select the optimal number of breaks, versus a null of no breaks. In Section VII we look for within sample variation in inventory and information effects based on parametric estimations. In Section VIII we look at non-linear estimation. In section IX we conclude. The appendix gives a brief description of the relevant economic events occurring at the time that the sample was recorded, as an extension of section II.

II. Theoretical Framework

This section presents a model of foreign exchange market trading that incorporates both inventory and asymmetric information effects. We follow the Lyons (1995) model, which explains and is used to test microstructure hypotheses in the foreign exchange market. We begin by assuming a world where the full information price of the risky asset is given by:

$$(1) V_t = \sum r_t = V_{t-1} + r_t \quad r_t \sim N(0, \sigma_r^2).$$

Dealers trading in this market quote prices by inferring the full information value of the asset (in this case the DM price of a Dollar) from the available signals. These signals are:

$$(2) S_t = V_t + \eta_t \quad \eta_t \sim N(0, \sigma_\eta^2).$$

$$(3) C_{jt} = V_t + \omega_t \quad \omega_t \sim N(0, \sigma_\omega^2).$$

$$(4) Q_t = V_t + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2).$$

In equation (2) S_t represents publicly available information about the asset. In equation (3) C_{jt} represents the opposite party's private information (we call the opposite party dealer j , and our dealer we call k). Finally, in equation (4) Q_t represents a signal of the market-wide orderflow as a whole.

Without any of these signals the dealer's best estimate of the value of the asset is

$E[V_t | \Omega_{t-1}] = V_t$. At the beginning of a period, the dealer observes S_t , then Q_t . Then the dealer quotes a schedule of prices for a possible Q_{jt} , where Q_{jt} is dealer j 's order to dealer k . In a real foreign exchange market transaction, a schedule would not be quoted because dealers trading with each other usually have a working relationship based on a \$10 Million standard order size. Here we allow for dealers to have different order sizes because each dealer will have a liquidity shock.

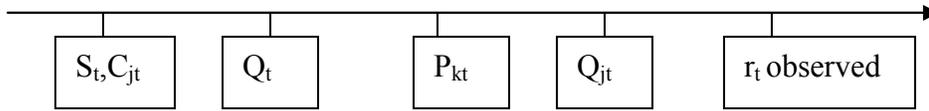


Figure 1. The timing of the market. This figure shows the signals in the sequence that our dealer observes them.

The diagram above shows the timing of the market. At each box, dealer k (our dealer) can update her information set based on the available incoming signals contained in the box. For example, in the first box, after observing S_t , dealer k 's estimate of V_t is:

$$E[V_t | \Omega_{t-1}, S_t] = S_t.$$

Notice that dealer k does not observe dealer j 's C_{jt} . This is because C_{jt} is dealer j 's private information. We do not include an extra step for incorporating dealer k 's own C_{kt} , as that could be suppressed into some new S_{tk}^* and would not add to the analysis. In the second box, after observing Q_t , dealer k 's inference about V_t becomes:

$$(5) \mu_t = E[V_t | \Omega_{t-1}, S_t, Q_t] = \rho S_t + (1-\rho)Q_t \quad \mu_t \sim N(V_t, \sigma_\mu^2).$$

$$(6) \rho = (\sigma_\varepsilon^2 + \sigma_\eta^2)^{-1} \sigma_\varepsilon^2$$

Equation (5) captures the inferred full information value of the asset, based solely on publicly available information. The intent of this model is want to see how dealer k sets prices. She does so by first looking at the publicly available information, which is equation (5). Having done this, she turns to trying to infer the private information of the opposite party, dealer j.

We assume that dealers have negative exponential utility, and as a result, their orders are linearly related to the difference between observed market prices and the private asset valuation.

$$(7) Q_{jt} = \theta(\mu_{jt} - P_{kt}) + X_{jt}.$$

Equation (7) is dealer j's order after observing dealer k's schedule of prices, P_{kt} . Dealer j orders Q_{jt} dollars, which is a linear combination of the difference between what he expected the dollar price to be (μ_{jt}) and what it is (P_{kt}), plus a liquidity shock (X_{jt}). Because dealers know the model, and each have the same utilities, they know the functional form of their utilities. Hence, we introduce the liquidity shock X_{jt} so that dealer k cannot invert Q_{jt} completely and learn the full private information of dealer j. This information is embedded in his best estimate of V_t , which we call μ_{jt} .

$$(8) \mu_{jt} = \lambda\mu_t + (1-\lambda)C_{jt}.$$

$$(9) \lambda = (\sigma_\mu^2 + \sigma_\omega^2)^{-1}\sigma_\omega^2$$

What dealer k will do now (he is the quoting dealer) is to isolate all the stochastic shocks that are not known to him around the value of V_t . Using (7) and (8) we have:

$$(10) \quad (1 - \lambda)^{-1}[(Q_{jt}/\theta) + P_{kt} - \lambda\mu_t] = V_t + \omega_{jt} + [(1 - \lambda)\theta]^{-1}X_{jt}.$$

Notice that everything on the left hand side of (10) is know to dealer k, and on the right hand side are all things that he does not know, and dealer j does. We will call the statistic on the left hand side Z_{jt} .

$$(11) \quad Z_{jt} \equiv \{ (1 - \lambda)^{-1}[(Q_{jt}/\theta) + P_{kt} - \lambda\mu_t] \} \sim N(V_t, \sigma_Z^2).$$

Z_{jt} contains the private shocks of dealer j, and dealer k (the quoting dealer) only knows the distribution. Hence, Z_{jt} is as close to unraveling dealer j's private information as the quoting dealer can get. So far we have public information summarized by μ_t and the information of the opposite party summarized by Z_{jt} . Dealer k's best estimate of the asset value then becomes:

$$(12) \quad \mu_{kt} = \kappa\mu_t + (1-\kappa)Z_{jt}$$

$$(13) \quad \kappa = (\sigma_\mu^2 + \sigma_Z^2)^{-1}\sigma_Z^2.$$

So that equation (12) is the best estimate of V_t by dealer k ¹. Assuming a simple inventory adjustment model, dealer k 's optimal price would be given by:

$$(14) \quad P_{kt} = \mu_{kt} - \alpha (I_{kt} - I^*) + \gamma D_t.$$

I^* indicates the optimal inventory level of the dealer.

Plugging in from the equations above, and redefining a parameter ϕ as:

$$(15) \quad \phi = [\kappa - \{\lambda(1-\kappa)/(1-\lambda)\}]$$

We arrive at the following pricing equation:

$$(16) \quad P_{kt} = \phi \mu_t + (1-\phi)[(Q_{jt}/\theta) + P_{kt}] - \alpha[I_{kt} - I^*] + \gamma D_t,$$

Or

$$(17) \quad P_{kt} = \phi \mu_t + [(1-\phi)/(\phi\theta)]Q_{jt} - (\alpha/\phi)[I_{kt} - I^*] + (\gamma/\phi)D_t.$$

¹ Note here that Lyons suppresses dealer k 's own private information, otherwise we would have something looking like

$$(12') \quad \mu_{kt}' = \gamma \mu_{kt} + (1-\gamma)C_{kt},$$

where dealer k is using his inference of dealer j 's information through Z_{jt} , public information through μ_t , as well as his own private information in C_{kt} . This, however, is probably suppressed since it would not add to the analysis, but would add more coefficients and variances.

Equation (17) is as far as the model can go without specifying a functional form for the signal of market-wide orderflow. In the FX markets brokered orderflow is the only available signal, however it is far from a complete picture of market-wide orderflow. Assuming that orderflow signals are the sum of m dealers' orders, plus a noise term, brokered orderflow becomes,

$$(18) \quad B_t = \sum_{i=1}^m Q_{it} + \psi_t.$$

Using this definition of brokered trading, the model is then extended to include a statistic Q_t , which will be centered on V_t with a random error. It can be shown that with (18) we can define this Q_t to be:

$$(19) \quad Q_{it} = B_t / [mp\theta(1-\lambda)] + S_t$$

Where S_t is the signal of the asset value defined in equation (2). The price equation then becomes:

$$(20) \quad P_{kt} = B_t / [mp\theta(1-\lambda)] + S_t + [(1-\phi)/(\phi\theta)]Q_{jt} - (\alpha/\phi)[I_{kt} - I^*] + (\gamma/\phi)D_t.$$

The only term left to specify in this equation is S_t , which is the marketwide signal of asset value. We assume that S_t is taken to be last period's value:

$$(21) \quad S_t = \mu_{k,t-1} + \varepsilon_{kt} = P_{k,t-1} + (\alpha/\phi)[I_{k,t-1} - I^*] - (\gamma/\phi)D_t.$$

Equation (21) uses the definition of $P_{k,t}$ to rewrite the definition of the 'random walk' type expectation of V_t specified by S_t . Plugging this into the price equation, we get:

$$(22) \quad \Delta P_{kt} = [\alpha - (\alpha/\phi)]I^* + [(1-\phi)/(\phi\theta)]Q_{jt} - (\alpha/\phi)I_{kt} + \alpha I_{kt-1} + (\gamma/\phi)D_t + \gamma D_{t-1} + B_t/[\mu\theta(1-\lambda)] + \varepsilon_{kt}.$$

Or

$$(23) \quad \Delta P_{kt} = \beta_0 + \beta_1 Q_{jt} + \beta_2 I_{kt} + \beta_3 I_{kt-1} + \beta_4 D_t + \beta_5 D_{t-1} + \beta_6 B_t + \varepsilon_{kt}.$$

Where the model predicts $\{\beta_1, \beta_3, \beta_4, \beta_6\} > 0$, and $\{\beta_2, \beta_5\} < 0$. Furthermore, using the definitions of S_t , μ_t , and Q_t , we can show that the error structure ε_{kt} is a moving average.

As a result,

$$(24) \quad \varepsilon_{kt} = \beta_7 v_{k,t-1} + v_{k,t}.$$

Hence, plugging into the price equation:

$$(25) \quad \Delta P_{kt} = \beta_0 + \beta_1 Q_{jt} + \beta_2 I_{kt} + \beta_3 I_{kt-1} + \beta_4 D_t + \beta_5 D_{t-1} + \beta_6 B_t + \beta_7 v_{k,t-1} + v_{k,t}$$

Finally, inventory can be controlled in three different ways. One is to adjust prices. The second is to sell to another dealer, that is, an out going order. The third is to sell through a broker, or brokered orders. We can write the inventory at any date as:

$$(26) \quad I_{kt} = I_{k,t-1} + Q_{k\Delta t}^o + Q_{k\Delta t}^b - Q_{jt}.$$

That is, (26) says that this trade's inventory (I_{kt}) is last trade's inventory ($I_{k,t-1}$) plus net orderflow from outgoing trades ($Q_{k\Delta t}^o$), as well as net orderflow from brokered trades ($Q_{k\Delta t}^b$), that have occurred between the last incoming trade, and the present incoming trade. We use Δt since a time period is defined as the elapsed time since the last incoming order ($Q_{j,t-1}$), and this one (Q_{jt}).

Hence, incorporating these extra inventory control measures brings us to the final estimable equation, after taking into account the power of the data in disentangling the elements of ΔI_{kt} ,

$$(27) \quad \Delta P_{kt} = \beta_0 + \beta_1 Q_{jt} + \beta_2 Q_{k\Delta t}^o + \beta_2' (Q_{k\Delta t}^b - Q_{jt}) + (\beta_2 + \beta_3) I_{k,t-1} + \\ + \beta_4 D_t + \beta_5 D_{t-1} + \beta_6 B_t + \beta_7 v_{k,t-1} + v_{k,t}.$$

Hence (27) represents the estimable equation. The moving average has a negative autocorrelation, so we would predict $\beta_7 < 0$. Furthermore, the model predicts that $\beta_2 = \beta_2'$.

III. Data

The data used to estimate the model in the previous section is unique in that it is direct transaction data for a dealer. As such it contains no aggregation error, and almost no measurement error. The data set is composed of the recorded transactions of a foreign exchange dealer based out of New York, trading in the Dollar/DM market for a large bank. The data is recorded in transaction time, not wall-clock time. This means that the time subscript t refers to the transaction number, not to the time at which the transaction occurred. For example, if $t=7$, this indicates that it is the seventh transaction of the week for the dealer, not that this transaction occurred at 7:00 a.m. or seven minutes after the opening of the market, or anything related to wall-clock times.

There are 843 observations for the dealer, which were recorded in the week of August 3-7, 1992. Monday had 181 observations, or $t=1$ to 181. Tuesday had 149 observations, or $t=182$ to 330. Wednesday had 110 observations, or $t=331$ to 440. Thursday had 152 observations, or $t=441$ to 592. Friday had 251 observations, or $t=593$ to 843. Excluding the first observation for the change in price component, as well as four overnight observations that are irrelevant for the purposes of measuring intraday price changes, we are left with a sample size of 838. For each observation we have measurements for the following variables:

P_t : The price of the dealer at which an incoming sale or purchase occurred.

Q_{jt} : The incoming quantity demanded by the opposite party. If the opposite party demands to buy, so that the quantity is transacted at the offer, this quantity is

recorded as positive. If the opposite party demands to sell, so that the quantity is transacted at the bid, this quantity is recorded as negative.

- I_t : This is the dealer's inventory at the time of (but not including) the incoming quantity Q_{jt} .
- D_t : This is an indicator that picks up the direction of trade. If the trade is buyer initiated, and occurs at the offer, $D_t=1$. If the trade is seller initiated, and occurs at the bid, $D_t=-1$. Hence, D_t and Q_{jt} always have the same sign. After controlling for asymmetric information via Q_{jt} , and for inventory effects via I_t , D_t picks up the size of the effective spread. Having this indicator, however, can be interpreted as *prima facie* evidence of non-linearities in the data.
- Q_t : This component of the data set is made up of the brokered deals of one of five DM/Dollar brokers in the New York market. As a result, it is a rough measure of the market-wide orderflow available to the dealer.

Beyond this data we consider external information relevant to the estimation of the model. We include a short summary of the events relevant to foreign exchange trading in Appendix B. A general overview of the week in question is as follows. In August of 1992 the US economy was in a sluggish economic recovery, and in the midst of an election between George Bush and William Clinton. The German economy was torn between two competing goals, fiscal and monetary. The fiscal authority, was concerned with financing the reunification of Germany. As a result, local, state and off-budget federal spending were increasing, as was inflation, and the money supply. The Bundesbank, headed by Helmut Schlesinger, had raised interest rates by seventy-five

basis points on July 16 of that year; German rates were six percentage points above U.S. rates, but the money supply was not contracting. In the week of August 3-7, 1992, the employment numbers for the U.S. economy were to be released. On Wednesday of that week, the Bundesbank announced that it would not increase its Lombard rate. On Friday, the employment numbers for July were released, and did not meet expectations. This sent the dollar lower. At the close of European trading, the Federal Reserve intervened multiple times to support the falling dollar.

One approach to measuring the effects of external conditions is to include this information as well as information about other market (e.g. equity or other currency markets beyond the Dollar/DM). The problem with this approach is that the strength of this data set lies in its precision in measuring the behavior of the dealer. This data contains the exact transactions of the dealer with no measurement or aggregation error. Hence, because the data shows such a precise picture, including anything that would depart from this dealer's behavior would contaminate the results. We can see turns in the dealer's pricing data that reflect the intervention of the Fed, for example. Hence, we can easily break the sample up in that sense. We cannot, however, introduce data outside of this data set, such as indicative quotes for other equity or currency markets, since its timing would not be synchronized with the transactions of our dealer, and in some cases is subject to measurement error (see Lyons 1995).

To get a feel for the data set, we include a graph over the week of trading for each of the aforementioned variables (for further descriptive statistics see Lyons 1995). In figure 2, we see the price and the change in price over the trading week. As is evident from the

upper panel, the price was generally falling over the week, until there is a sudden upturn in the market on the last day. This upturn is due to the repeated intervention of the Federal Reserve to boost the dollar after the release of the US employment data for July. The signal for orderflow, Q_{jt} , inventory I_t , and brokered trading, Q_t , are pictured in figure 3.

3. The empirical evidence for microstructure effects in foreign exchange markets obtained by Lyons 1995 were based on the data presented here, which was used to estimate an equation of the form of equation (25). We now turn to the question of parameter instability, and within sample variation of parameter value and significance in the estimation of this model.

IV. A First Pass: Markov Processes and Regime Switching

We reproduce here equation (25) for the reader.

$$(25) \quad \Delta P_{kt} = \beta_0 + \beta_1 Q_{jt} + \beta_2 I_{kt} + \beta_3 I_{kt-1} + \beta_4 D_t + \beta_5 D_{t-1} + \beta_6 B_t + \beta_7 v_{k,t-1} + v_{k,t}$$

In figure 4 we see the results from the baseline estimation of the model. The equation confirms the hypothesis of both inventory and information asymmetry effects being present in the process of determination of foreign exchange prices. All variables are significant and properly signed. Accepting that these are the variables that are relevant for estimating the true data generating process², we now turn to the question of whether the functional form of equation (25) accurately captures the relationship between these variables.

Notice the functional form of the equation (25). If we interpret the estimates of β as the response to changes in inventory, incoming orders, and so on, we are implicitly making two assumptions. The first is that the response of the dealer is linearly related to the variables observed. So, for example, we would assume that the dealer lowers his price by β_1 for the eleventh order to sell, after selling ten times in a row, just the same as if he had just received the first sell order. However, what if the dealer senses that the market is falling? We may hypothesize that the dealer may be inclined to sell into the falling market differently than he would be selling during a trending or 'normal' market. For example, looking back at figure two, we can see that the market price is more volatile on

² The question of whether a misspecification is present due to an omitted variable is, perhaps, relevant, but beyond the scope of this paper.

the last two days than on the first three. Thursday begins at observation 441 (recall that observations are on the abscissa), where the price moves more erratically. The question that we wish to analyze is whether by averaging the observations taken on these days with the observations on the other days, we are muting the potential effects of non-linearities in the market. For example, there may be non-linearities due to regime switching.

If we know there to be such non-linearities, one way to capture is by using a maximum likelihood estimation method with a model of regime switching coefficients. To model this, we define a state variable s_t , which defines the current state of the economy. Agents cannot observe s_t , but they try to infer where they are, and then update their beliefs taking the inferred state into account. An N-state Markov process is assumed to govern the movement of the variable s_t .

The model of regime switching with a hidden Markov state variable can be generalize to include a large class of processes, including state variables that are not Markov. After looking at various forms of the Markov regime-switching model, including regime switching autoregressive conditional heteroskedasticity, and three states instead of two. The more elaborate models reproduced the basic results of the simple two state model, so we chose this more parsimonious representation. We will assume that there are two basic states of nature. Let the low state be $s_t=1$, and the high state be $s_t=2$. The transition matrix is a matrix P where each entry $p_{a,b}$ describes the probability of going from state a

to state b^3 . If we believe that the probability of going from one regime this period, to another next period is Markov, then,

$$(28) \quad P(s_t=j \mid s_{t-1}=i, s_{t-2}=k, s_{t-3}=l, \dots) = P(s_t=j \mid s_{t-1}=i) = p_{ij}.$$

In this case the transition matrix would look like:

$$(29) \quad P = \begin{bmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{bmatrix}.$$

Hence in this model, depending on which regime he believes he is in, the trader will respond differently to the incoming signals in the market, e.g. inventory or orderflow changes. This is in contrast to Lyons (1995), which estimates uniform responses throughout the entire week of trading. If we look at figure 2, we would be saying that the trading on Monday was subject to the same responses as the falling markets of Thursday and Friday. With multiple regimes, we can capture the optimal response to different trading environments in the markets intraday. Furthermore, we can estimate these transition probabilities endogeneously through a grid search; we do not have to assign the transition probabilities.

The estimation of this model is based on the EM algorithm in Hamilton (1994),

³ Note that $p_{a,b}$ is **not** the $(a,b)^{\text{th}}$ entry, it is the $(b,a)^{\text{th}}$ entry. So the probability of going from state $s_t=1$ to $s_{t+1}=2$ is $p_{1,2}$, and this is entry $(2,1)$ in the P matrix, because the columns of the P matrix describe the distribution of going from state $s_t=1$ to all possible states.

Engle and Hamilton (1990), and Hamilton and Susmel (1996). We begin by rewriting equation (25) into vector form.

$$(30) \quad \Delta P_t = \beta_{s_t}' x_t + v_t.$$

Where we have a vector of parameters that depends on the hidden state variable governed by the transition matrix given in (29). For each state, there is a density particular to that state that describes the data generating process. We stack the densities into a vector in equation (31).

$$(31) \quad \eta_t = \begin{bmatrix} f(\Delta P_t | s_t = 1, x_t; \beta) \\ \vdots \\ f(\Delta P_t | s_t = N, x_t; \beta) \end{bmatrix} = \begin{bmatrix} (2\pi\sigma_{s=1}^2)^{-1} \exp\left\{\frac{-(\Delta P_t - \beta_{s=1}' x_t)^2}{2\sigma_{s=1}^2}\right\} \\ \vdots \\ (2\pi\sigma_{s=N}^2)^{-1} \exp\left\{\frac{-(\Delta P_t - \beta_{s=N}' x_t)^2}{2\sigma_{s=N}^2}\right\} \end{bmatrix}.$$

Equation (31) describes the vector of densities for the case of N states. For each state and each point in time, there is a density describing the data generating process. Note that to estimate this, we must estimate k parameters in the $\beta_{s_t=j}$ vector of parameters, for j=1 to N states. Hence, this is a very tall order. The estimation is carried out by defining an (Nx1) vector of estimated state probabilities $\hat{\epsilon}_{t|t}$, such that $\hat{\epsilon}_{t|t}$ is the distribution of probabilities of being in the N states, contingent on the information at t. Accordingly, $\hat{\epsilon}_{t+1|t}$ represents the estimated probability distribution of landing in the N states at time t+1, given the information set at time t. To estimate an equation of the form of (30), we iterate over the following algorithm:

$$(32) \quad \hat{\epsilon}_{t|t} = \frac{(\hat{\epsilon}_{t|t-1}(\bullet)\eta_t)}{1'(\hat{\epsilon}_{t|t-1}(\bullet)\eta_t)}.$$

$$(33) \quad \hat{\epsilon}_{t+1|t} = P\hat{\epsilon}_{t|t}.$$

For further details, see Hamilton (1994), and Engle and Hamilton (1990). What is important to note is the number of parameters estimated in the likelihood function embedded in the algorithm. There are k parameters by N states. Doing a grid search over a problem of this dimensionality is infeasible. Instead, we begin by looking for two state regime switching in the left hand side variable only. This way, we subsume all of the parameters of the right hand side into the estimated mean and variance of the densities on the right hand side of (31). Hence, for now we postulate that a better approximation to the data generating process is of the form:

$$(34) \quad \Delta P_t = \begin{cases} N(\mu_1, \sigma_1^2) \\ N(\mu_2, \sigma_2^2) \end{cases}$$

Where (30) shows that the change in price is in one of two normal distributions, and the probability of switching from one distribution to the other is governed by the transition matrix of equation (29). Hence for now we subsume all of the right-hand-side variables into the parameters describing the distributions of equation (30), and calculate this as a univariate model.

In figure five we see the results of this estimation. The data generating process seems to exhibit Markov-type regime changes. For example, in observation 401, the probability of being in regime one is almost zero. In observation 801, which is around where the Fed intervenes in the market, we see the same probability jump to 1. Hence, if we estimate our parameters using data points such as these, we would be forcing a linear average of these two distributions, and perhaps muting some parameter differences among the regimes. Because we have only included the left hand side variable in the Markov regime switching estimation, we cannot speculate on whether this regime changing is

affecting the estimated parameters. That is, we have no direct evidence that the right hand side is failing to pick up these regime changes, and hence making the parameters unstable. Also shown in figure five is a chart with the estimated moments of the two distributions. The chart shows the estimated means and variances for the two states. State one has both a higher (less negative, closer to zero) mean and a higher variance than state two. Hence, state two is a state in which the price is falling more sharply than state one, and fluctuate less than state one.

To address the problem of not including the possible effects of the regressors, we take the residuals from the estimated equation (25), and look for regime changes in them. Figure six shows the estimated residuals from equation (25), and an estimate of regime switching for these residuals.

From the figure we can see that the residuals display almost identical Markov switching behavior. Note what this implies. We have looked at the left-hand-side variable and found evidence of regime switching. We have then attempted to explain the variation in the left hand side variable using equation (25). We have taken the unexplained part of equation (25) and found the same regime switching. Hence, our model is not explaining some of the variation in the left-hand-side variable, particularly the variation picked up by Markov regime switching.

While we cannot take this to be clear evidence that the data generating process is in fact represented by a Markov model in the spirit of (29) and (30), we can take this as evidence

that there are non-linearities that are not captured by (25). We now turn to looking within the sample for some direction as to what these non-linearities may imply.

V. A Closer Look at the Data

To begin isolating some of the potential parameter instability that the specified equation may fail to capture, we roll a 150-observation window across the sample, and re-estimate equation (25) at each of the subsamples. So, for example, the first estimate would consist of estimating equation (25) on the subsample beginning with observation one and ending with observation 150. The idea here is to allow for piecewise variation in the parameter values within the sample. To give an example of why this may be helpful, suppose that we are interested in seeing whether the market is behaving differently on Tuesday than on Wednesday. There are on average 170 incoming observations per day. By allowing a window of 150 observations, we can look at an average of twenty parameter estimates just for each day. As the parameter evolves, we can get an idea of whether the values for Tuesday are very different than on Wednesday. With this approach we sacrifice the explanatory power of a greater sample size, but gain the within sample variation.

The figures below show these rolling estimates. If we look at figure seven, for incoming dealer orders, we can see that the estimates are highly significant in the middle of the week. Figure eight shows that inventory is less significant in the beginning of the week, but much more so in the middle and towards the end, when the market becomes more hectic. In fact, inventory becomes evermore significant, as the week draws on, until the point of intervention by the Fed, where its significance falls off.

Figures nine and ten show the piecewise variation in the lagged inventory coefficient, and the adjusted R^2 . Note that lagged inventory shows the same behavior in terms of growing significance throughout the sample. As the more volatile days are estimated, lagged inventory becomes a highly significant variable, whereas in the more tranquil markets of the beginning of the week, lagged inventory is not piecewise significant.

The adjusted R^2 shows itself higher than the total R^2 . The reader will recall that a reduction in sample-size should *reduce* the adjusted R^2 , not increase it⁴. In the figure, however, we can observe that the adjusted R^2 is consistently *higher* than the estimated R^2 for equation (25). Hence, the variation in the data generating process that this equation averages out by assuming a linear form is so significant, that taking a piecewise specification of eighteen percent of the sample (150 of 838 total) increases the explanatory power of the regressors throughout almost the entire sample.

⁴ The adjusted R^2 can never be greater than the R^2 , and can decrease as one adds regressors. Because the derivative is *positive*, as we *decrease* T, we would expect that the adjusted R^2 would also *decrease*. (Note that $k=7$ in equation (25).)

$$\bar{R}^2 = 1 - (1 - R^2) \left(\frac{T-1}{T-k} \right)$$

$$\frac{\partial \bar{R}^2}{\partial T} = (1 - R^2) \left(\frac{k-1}{(T-k)^2} \right) > 0 \Leftrightarrow k > 1$$

The rolling estimates for the direction of trade indicators, as well as for the moving average coefficient and brokered trading are reported in Appendix C.

The piecewise observed variation in the coefficients can be examined further with a rolling F-test for structural breaks. In figure 11 we report the results of a rolling breakpoint estimate for equation (25). The estimate involves picking a breakpoint transaction, and splitting the sample at that transaction. We estimate the resulting two subsamples, and compare the residual sum of squares using a standard F-test (sometimes called a Chow test). The advantage to this test is that it rolls the breakpoint along every possible transaction, so long as the two resulting subsamples contain sufficient observations to estimate equation (25). In this sense, it is fairly comprehensive over the entire range of possible breakpoints. The drawback is that this test has no power to reject the null hypothesis of no breaks versus any alternative hypothesis besides the hypothesis of one break. Hence, it is of limited value in detecting multiple breaks. This is a point to which we will return in section VI.

In the figure we can see that the test seems to find a break in the middle of the week, and then falls again towards the very end of the sample. Hence, it would seem that the model changes from the beginning of the week to the end. To further corroborate this hypothesis, we run a Wald test to see if the coefficients are equal at the beginning of the sample, versus at the end. We report the results in figure twelve.

So far we have explored the issue of whether there is within sample parameter stability and goodness of fit for equation (25). The evidence presented so far would support the hypothesis of parameter instability, however, no statistical test presented so far has the necessary power to reject the null of no breaks in favor of an alternative hypothesis of a more general character than just of one break. We now turn to examining this issue.

VI. Statistically Testing for Breaks

After describing a model of microstructure in foreign exchange in section II, section III documented economic factors that have the potential to affect the Dollar/DM market over the week in question. Section IV shows a two-state Markov process that finds market switching present in the data, and not picked up by the equation specified in previous work. Section V finds that rolling estimates indicate parameter instability within the sample. Particularly, that inventory effects and asymmetric information effects tend to offset each other within the sample. In this section we rigorously search for parameter instability by examining statistical tests for within sample structural breaks. We can test for exactly how many breaks are present in the sample without knowing total number of breaks, or the location, and then find them at better than 2.5% significance. We find evidence of multiple structural breaks at the aforementioned significance levels.

The statistical tests presented in this section are based on recent econometric developments regarding testing for and dating of multiple structural changes within samples⁵. We can estimate breaks in the model by estimating minimized sum of squared residuals over a set of potential breaks. We treat the number of breaks, and the dates of their occurrence as unknown variables to be estimated. This estimation strategy is robust to general forms of serially correlated and heteroskedastic errors, as well as partial structural changes, in which some parameters are allowed to stay fixed throughout the entire sample.

⁵ See Bai and Perron (1998), Bai (1999).

Consider the following model.

$$(35) \quad \begin{aligned} y_t &= x_t' \beta + z_t' \delta_j + u_t \\ x &= (T \times p), \beta = (1 \times p) \\ z &= (T \times q), \delta_j = (1 \times q) \end{aligned} \quad \text{for} \quad (t = T_{j-1} + 1, \dots, T_j)$$

In equation (35) we assume a generalized model of partial structural change, with m breaks (i.e. $m+1$ regimes). In particular, this model allows for some of the regressors (z_t') to change regimes, and consequently their corresponding coefficients to vary among regimes (i.e. δ_j). The model assumes that the other regressors (x_t) are not subject to structural breaks, and as such their corresponding coefficients (β_t) are constant throughout the entire sample. If $p=0$, we have the case of full structural change with every regime change.

We define a partition (T_1, \dots, T_m) to be a set of m dates where the sample changes regimes (i.e. T_j represents the last date of regime j). We denote a particular partition j as $\{T_j\}$, so that the coefficients δ_j and β are estimated for each partition $\{T_j\}$. Denoting the true value of the parameters with a 0 superscript, we can write the true data generating process as:

$$(36) \quad Y = X\beta^0 + \bar{Z}^0 \delta^0 + U$$

$$(37) \quad (\beta^0, \delta_1^0, \dots, \delta_{m+1}^0, T_1^0, \dots, T_m^0)$$

With (32) being the matrix form of the data generating process of (31), and (33) the vector of true parameter values to be estimated. For any given partition $\{T_j\}$, we first estimate β , and δ_j by minimizing the sum of squared residuals. That is,

$$(38) \quad \min_{\beta, \delta_j} \sum_{i=1}^{m+1} \sum_{t=T_{i-1}}^{T_i} \left[y_t - x_t' \beta - z_t' \delta_j \right]^2 = \hat{\beta}(\{T_j\}), \hat{\delta}(\{T_j\}).$$

Hence, (34) gives the optimal parameter estimates as a function of the partition. We then substitute these into the objective function and denote the sum of squared residuals as $S_T(T_1, \dots, T_m)$. The estimated breakpoints are the points at which the sum of squared residuals is minimized.

$$(39) \quad \begin{aligned} (\hat{T}_1, \dots, \hat{T}_m) &= \arg \min_{T_1, \dots, T_m} S_T(T_1, \dots, T_m); \\ \therefore (T_1, \dots, T_m) &\text{ satisfies} \\ T_i - T_{i-1} &\geq q \end{aligned}$$

A test of no break versus $m=k$ breaks would take the following form. If (T_1, \dots, T_k) is a partition such that $T_i = [\lambda_i T]$ for $(i=1, \dots, k)$, then

$$(40) \quad F_T(\lambda_1, \dots, \lambda_k; q) = \left(\frac{T - (k+1)q - p}{kq} \right) \frac{\hat{\delta}' R' \left(R(\bar{Z}' M_x \bar{Z})^{-1} R' \right)^{-1} R \hat{\delta}}{SSR_k}, \text{ where}$$

$$(41) \quad (R\delta)' = \left(\delta_1' - \delta_2', \dots, \delta_k' - \delta_{k+1}' \right), \text{ and}$$

$$(42) \quad M_x = I - X(X'X)^{-1} X'.$$

Here, SSR_k is the sum of squared residuals under the alternative. We must place restrictions on the break dates so that they are asymptotically distinct, and bounded from the boundaries of the sample. For $\varepsilon > 0$, we define the set of possible partitions as:

$$(43) \quad \Lambda_\varepsilon = \{(\lambda_1, \dots, \lambda_k) \mid \lambda_{i+1} - \lambda_i \geq \varepsilon, \lambda_1 \geq \varepsilon, \lambda_k \leq 1 - \varepsilon\},$$

where $\lambda_i = (T_i/T)$. It is important to note here that (39) imposes a crucial restriction on the last possible break date; the last break date is limited to being no closer than ε proportion of the sample away from the last observation. To find the asymptotic distribution of the test statistic, we adopt $\varepsilon = 0.05$, to ensure that the test has sufficient power. Hence the last possible break date can be no closer than five percent or no less than 42 observations from the end of the sample. With 838 observations, this restriction amounts to not allowing a breakpoint at or after observation 796. Note, however, that the Fed interventions on Friday, August 7, 1992 occurred at the end of the trading day, after European trading had ceased. They are picked up in the sample in the neighborhood of observation 796. On the day when the Fed intervened no less than three times to support the dollar, the prior would be that the point of intervention would represent a structural break in the sample, and indeed we observe a change in the price direction at that moment. Yet here we will not expect the test to pick this up.

The test of no breaks versus k breaks gives no feel for what is occurring between zero and k breaks, in the event of a rejection of the null hypothesis. Instead we consider a test of l versus $l+1$ breaks. We adopt the null hypothesis of l breaks. We can then estimate the breakpoints consistent with (35). Using the estimated partition $(\hat{T}_1, \dots, \hat{T}_m)$ we can test each of the $(l+1)$ segments for the presence of an additional break. At each step we conduct $(l+1)$ tests (one on each sub-segment) of no structural change versus the alternative of one structural change. This test occurs over each of the $l+1$ sub-segments, where a sub-segment spans $(\hat{T}_{i-1} + 1, \hat{T}_i)$ for $(i=1, \dots, l)$, and $(\hat{T}_0 = 0, \hat{T}_{l+1} = T)$. The null

hypothesis of l structural breaks is rejected in favor of the alternative of $(l+1)$ structural breaks when the sum of squared residuals for the smallest of the segments where $(l+1)$ breaks were found is sufficiently smaller than the minimum sum of squared residuals for the l breaks. That is:

$$(44) \quad F_T(l+1 | l) = \frac{\{S_T(\hat{T}_1, \dots, \hat{T}_l) - \min_{1 \leq i \leq l+1} \inf_{\tau \in \Lambda_{i,\eta}} S_T(\hat{T}_1, \dots, \hat{T}_{i-1}, \tau, \hat{T}_i, \dots, \hat{T}_l)\}}{\hat{\sigma}^2}$$

where,

$$(45) \quad \Lambda_{i,\tau} = \{\tau; \hat{T}_{i-1} + (\hat{T}_i - \hat{T}_{i-1})\eta \leq \tau \leq \hat{T}_i - (\hat{T}_i - \hat{T}_{i-1})\eta\}.$$

Here, $\hat{\sigma}^2$ is a consistent estimator of σ^2 under the null hypothesis. Equation (44) defines the test statistic, and equation (45) gives the area between the two breakpoints in question, where the new potential breakpoint is allowed to take on values. We begin by estimating this model for no breaks, and perform the test for one break. If we reject the null, we test for one break versus two breaks, and so on. At every rejection of the null hypothesis, we increase l sequentially until the test fails to reject the null hypothesis of no additional breaks.

The estimation of these test statistics are computationally intensive, and as a result, the maximum number of breaks initially permitted was limited to five. This limitation turned out to be of no consequence, since the test rejected five breaks in all cases at conventional significance levels.

In figure (13) we can see the results of the Bai Perron tests for multiple structural breaks. The data on brokered interdealer trading was excluded in this estimation for computational efficiency. The result of interest is that of column one. Here we see that allowing full structural change, and a maximum of five breaks, we reject the null hypothesis of no breaks in favor of one break occurring at observation number 449, at the one percent significance level. Furthermore, we cannot reject the null hypothesis of one break at observation 449 versus two breaks at any other point in the sample. Hence here we have statistical evidence of one break in the sample.

In the second panel of figure (13) we report the results of restricting the coefficients of the variables in the left column to stay constant within sample. This is solely for comparison purposes, because there is no theoretical justification for assuming that some of these coefficients vary while others do not. The purpose of fixing one coefficient throughout the sample is to increase power when the researcher **knows** this coefficient is invariant among (potential) structural breaks. This is not the case here. Withstanding this caveat, we confirm a breakpoint in the vicinity of observation 449 in three cases. Interestingly, we cannot reject the null of no breaks when we fix the coefficient on lagged inventory (I_{t-1}) and the coefficient on lagged orderflow indicator (D_{t-1}) across potential structural breaks. Finally, fixing inventory across potential breaks yields not one but three breaks at the one percent level, and four breaks at the 2.5% level.

Hence, we have found one break in the sample that is statistically significant, and corresponds to the vicinity of observation 449. This test has confirmed the results of

previous sections insofar as finding a break at approximately the beginning of Thursday, August 6, 1992 in our sample. Finally, this test is, by construction, unable to pick up a change in regime at the time of the Fed intervention.

VII. Parametric Estimation with Structural Breaks

In this section we estimate equation (25) taking into account the information available from previous sections on structural breaks present in the sample. By using several criteria for choosing subsamples and break dates, we avoid the issue of sample selection bias. We identify and characterize the asymmetric information and inventory effects on price setting behavior by foreign exchange traders in three different states within the sample. The first state can generally be thought of as the normal or sideways trending market state. Here, the trader is not facing a tumultuous market, or too much uncertainty. The second state can be thought of as one of intense trading, with market events changing rapidly, and prices and orderflow fluctuating more heavily. Here the trader is adapting to a faster moving market, where perhaps fundamentals are shifting, and uncertainty is greater. The third state is one in which the Fed intervenes. Here the monetary authority is ratcheting the market upwards, and as a result the force of the intervention changes the relationship between price changes and the right hand side variable. We first present each of the estimations resulting from subdividing the sample according to different regime selection criteria. We then present a generalized pattern that is common to all the results, and this pattern is robust to the different criteria used to select the regimes. This generalized result better characterizes the behavior of a dealer in the foreign exchange markets, in accordance with the microstructure theory of this market.

Markov Process

In this section, we use the estimated transition probabilities to segregate the areas of the sample that are in different regimes. The general idea is to segregate areas in the sample where the structural relationship of the data generating process is different from other areas. By breaking the sample into areas in which we believe that structural changes have occurred, we can test for breaks and look at the parameter changes in the different subsamples.

In looking at the transition probabilities (figure 5), the majority of the sample is in regime two. There are segments, however, where the probability of being in regime comes close to unity. After looking at several selection criteria, we defined three regimes within the full sample. The first is composed of three disjoint windows, which separate two disjoint parts of the second regime, and the third.

In figure (14) we see a diagram depicting these subsamples. Each regime is shaded differently, to show how the disjoint parts of the sample are combined based on the estimated probability of being in state 1 for both the change in price, and the residual of equation (25). As shown in the figure, regime one has three parts. The first part is composed of observations 1 to 49. The second part is composed of observations 96 to 549. The third is composed of observations 741 to 794. The idea behind the selection criteria used to segregate the sample into these subsamples was based on the following. We assume most of the observations are in the first regime. If the estimated probability of being in state one approached unity for an observation, we considered including that

observation to be in the second regime. If enough surrounding observations were close to unity as well, then the neighborhood of observations was included in the second regime (and hence excluded from the first). We defined approaching unity as having an estimated probability of being in state one of no less than 0.9. If an observation and its residual met this criterion, we then considered a window of five observations above and below it. If the surrounding observations also approached unity, we considered the same neighborhood around them. If three or more observations met the criteria, we included the neighborhood in regime two, i.e. not part of the first regime. We then selected out the last part of the sample, which had been a part of regime two, where the Fed intervened, because of the unusual continuum of points at which the probability of being in regime one was better than 0.98. Clearly, there was no switching in this part of the sample, and it was not like the other parts of regime two. Hence, we defined it as its own regime.

In figure (15) we depict the transition probabilities, with dark vertical lines indicating the regimes. In the top panel we see two pairs of dark vertical lines, and between each pair of lines, the estimated probability of being in regime one jumps up erratically. The areas between these dark lines define regime two. In the bottom panel, at the right end of the graph, we see another pair of dark lines. Between these lines are the observations that comprise regime three.

The model was estimated according to equation (25), with the sole addition of a dummy to indicate the jumps in the sample. The data was arranged in chronological (event time) order within regimes, and then arranged by regimes. The jump dummy takes on the

value of one when the data switches from one subsection of the sample to another, zero everywhere else. The brokered data was excluded to facilitate the estimation of regime three, as there are only forty-nine observations. With six right hand side variables, plus a dummy for the jumps in the sample, plus a constant and a moving average coefficient, there is not enough data for estimation. Since regime three has no jumps, we exclude the dummy for jumps in the sample. Since the brokered data is zero throughout most of the sample in general, and throughout all of regime three in particular, we exclude that as well.

Figure (16) shows the estimation results for the regimes selected according to the criteria based on the estimated state probabilities. In the upper panel we see the estimation results for regime one. This regime corresponds to observations that the Markov estimation found to come primarily from state two. This state has a lower (more negative, i.e. higher in absolute value) mean and a lower variance than state one. The observations from this state make up the majority of the sample, which is consistent with the estimated transition probability of 0.05 for exiting state two. The lower mean is consistent with the general bearish or falling market for the week in question.

Comparing the estimates from the upper panel, which correspond to regime one, with the overall estimation in the bottom panel, we see that the coefficients are all significant at conventional levels and signed as predicted. The adjusted R^2 is almost twice that of the baseline model, even though the sample size is reduced by over 35%. Looking at the asymmetric information effect, as measured by the coefficient on orderflow, Q_{jt} , we see that in regime one, the dealer does not shade his price against asymmetric information as

much as in the general model. This is consistent with a lower variance in the price as estimated by the Markov switching model. In the baseline model the dealer increases his spread by 2.8 pips (0.00028 DM) per \$10 million to protect against adverse selection (which is twice the coefficient on Q_{jt}). In this regime we predict that the dealer's spread will increase by 1.8 pips, or 33% less wide, for the same quantity. The coefficient on inventory reflects the amount of price shading that the dealer engages in to adjust inventory levels. Here we see that for a unit increase in inventories, a dealer lowers his price less in this regime than in the overall model. The coefficient on lagged inventory is also lower, since the dealer must raise his price less to adjust for the previous trade's inventory imbalance. The coefficients on the orderflow indicator (DEE_t in the chart) measure the effective spread after we take into account inventory and information effects. In this model, the dealer has a wider spread than in the baseline estimates; the spread is twice the coefficient on DEE , or 2.38 pips.

For the second regime, we are looking at or in the neighborhood of the estimated Markov state one observations. State one was estimated to have a higher negative mean (lower in absolute value) than state two, and a lower variance. The estimates for the region we characterized as regime two showed a much wider increase in the estimated spread to protect against adverse selection, as measured by the coefficient on Q_{jt} . Here the dealer's spread is 33% wider than the baseline model. More importantly, signed orderflow is **not** significant at the five percent level in this regime (with a p-value of .0721). Additionally, the inventory effect is magnified in this regime, so prices adjust more poignantly to uncovered inventory positions, and consequently adjust more quickly to lagged inventory

imbalances (as measured by the coefficients on I_t and I_{t-1} , respectively). The signed orderflow indicators (DEE_t and DEE_{t-1}) are significant at conventional levels, and while not as highly significant as in regime one, are of about the same magnitude. They also do not satisfy the model prediction of having a higher coefficient (in absolute value) on D_t than D_{t-1} . Finally, the adjusted R^2 is higher than the baseline model, even though the sample size for regime two is over 70% smaller. From all this we can characterize regime two as one in which inventory plays a more important role than orderflow in predicting price changes.

For regime three, given that we have only 49 observations, or about 5% of the sample, it is not surprising to see that most of the coefficients are insignificant at conventional levels, especially considering the heavy intervention of the Fed. What is surprising is that orderflow is **more** significant in regime three than regime two, and comes in at almost ten times the size of the baseline estimate. This regime represents the foreign exchange market activity as the Fed intervenes repeatedly, and in a short period of time, to support the dollar. Clearly, as the Fed forces the price upwards, our dealer reacts by increasing rapidly his price, and maintaining a wide spread. The coefficient on Q_{jt} shows a spread of 26 pips, and is significant at the five percent level, whereas the other coefficients are not remotely significant at any conventional level.

Finally, we report results for a traditional F-test for multiple joint structural breaks at the points predicted to be breaks in the sample, i.e. points 50, 96, 551, 741, and 795, in accordance with figure (14). The results are presented below in figure (17).

Bia and Perron Test for Multiple Structural Breaks

We now turn to segregating the sample according to the statistical method discussed in section VI. The test based on the Bai and Perron (1998) sup-F test statistic for multiple structural breaks found (exactly) one breakpoint at observation 449, at better than one percent significance level. The computation of the test statistic, however, required imposing assumptions about the location of the breakpoints. These assumptions required that the breakpoints be asymptotically distinct, and bounded away from the ends of the sample by a proportion of no less than 0.05 of the sample size. While the requirement that breakpoints be asymptotically distinct turned out to be inconsequential, the second requirement is of more concern. With 838 observations, five percent of the sample represents 42 observations. Hence, as stated in the previous section, we cannot expect the test to pick up a break after the observation that is 43rd from the end of the sample, or observation 795. Because our prior is that intervention by the Fed is a structural break in intraday foreign exchange activity, and because an intervention by the Fed occurred at around observation 795, we impose a breakpoint on that date. We report results for estimations both with and without the breakpoint at 795 in what follows.

In figure (18), we see results for the Bia and Perron sample split. We will call the observations prior to the first break (449) regime one. What is most striking about this regime is that inventory and lagged inventory are not even remotely significant. In this regime the adjusted R^2 of 0.32 is much higher than the baseline estimates, all coefficients are signed properly, the data is continuous, in chronological (event-time) order, and is over 53% of the available observations. All this withstanding, inventory plays no role in

the price determination. This, while signed orderflow (i.e. the asymmetric information effect) and the orderflow indicators are significant, and estimated at magnitudes similar to the baseline estimates, including an estimated spread of 2.6 pips per \$10 million to protect against adverse selection.

We refer to the estimates from the sub-sample 449 to 794 as the second regime, which are reported in the second panel. Here the estimates show orderflow to be inconsequential to the dealer, and the inventory components to be highly significant. Furthermore, with just 34.5% of the available observations, the adjusted R^2 is estimated at 0.29, higher than the adjusted R^2 of 0.22 of the baseline estimate, which is estimated with all available observations. All coefficients here are signed properly, however, the estimates fail to satisfy the prediction that the coefficient of DEE_t be greater than the absolute value of the coefficient of DEE_{t-1} . The change in price per \$10 million dollars of uncovered inventory is estimated at -2.04 pips, while the same for lagged inventory is 1.84. These magnitudes are more than double their counterparts in the baseline estimate, indicating much more aggressive inventory management, in sharp contrast to regime one.

The third regime comprises the sample ranging from observation 795 to 838. This break represents the market under the Fed intervention. While the sample size is small, there is enough explanatory power in the sub-sample to bring attention to orderflow, Q_{jt} . In similar fashion to the Markov results, the orderflow coefficient is significant at the 10% level (with a p-value of 0.064) and measures 14.1, which implies a spread of 28 pips. All other regressors are highly insignificant in this subsample, and have the opposite of their

predicted signs. With just five percent of the sample recording a market intervention by the Fed, it is not surprising to see a low adjusted R^2 .

In the third panel, we report the results of combining regimes two and three. Here we ignore the predicted break at observation 796, and hence the results, for a subsample with 46% of the available observations, are poor. The model does not fit this subsample well. The adjusted R^2 is lower than the baseline model, the orderflow indicators do not satisfy the predicted relationship, and while the coefficients are signed correctly, some are not significant at the five percent level (i.e. DEE_t and $MA(1)$).

Finally, we report the traditional F-test for structural breaks at observations 449 and 796 below. Note, however, that this is less powerful than the Bai and Perron test used to find these breaks.

F-Test

So far, the Markov and Sup-F tests have yielded results that would suggest parameter instability in the relationships estimated by equation (25), and a more dynamic relationship between the regressors and the left-hand-side variable. In the interest of guarding against sample selection bias, we include estimates of the baseline equation over subsamples where the breakpoints are determined where the p-value of the F-test reaches a local minimum in the rolling F-test estimation depicted in figure (11). This rolling test reaches three local minimums at conventional significance levels. These are (p-values in parenthesis) observations 505 (0.023), 699 (.009), and 722 (.031). As mentioned previously, the F-test has no power to discriminate between one break and

three. At these points the null of no break was rejected in favor of one break. The Sup-F test rejected the alternative of two breaks (within the adjusted sample), as well as that of three breaks, against a null of one break. Because, however, the Sup-F test statistic cannot find breaks that are within five percent of the end of the sample we cannot waive hands and eliminate the possibility of a break there. This and the search for a generalization of parameter stability that is robust to sample selection leads estimate the equation over these subsamples.

Figure (20) gives the results of the estimation based on the aforementioned breakpoints. The same general pattern emerges here as with the Markov and Sup-F breakpoints. At the beginning of the week, during the more tranquil periods, orderflow is significant at conventional levels, and inventory is not. As the market enters a sharper downturn at the beginning of Thursday, a breakpoint emerges in which orderflow becomes insignificant, and inventory becomes significant. Finally, at the end of the week, with the Fed intervention, orderflow again becomes significant, but generally speaking the model fails.

Following the convention of the previous sections, we can refer to regime one as the estimation of the top panel of figure (18). Here orderflow and the directional indicators are highly significant, whereas inventory is not. Furthermore, the adjusted R^2 is higher than baseline estimates. Regime two (the second panel) shows inventory becoming highly significant, and signed orderflow insignificant. Here, the directional indicators of orderflow are still significant, and the adjusted R^2 is at about the level of the baseline estimate (with 88% less observations used to fit the model). The third regime shows a

striking contrast with the baseline model. Both inventory and its lag are highly significant, and there is no role for orderflow, or its directional indicators. What is more remarkable is the adjusted R^2 of 0.65 with less than 3% of the sample used to fit the model. Finally, the last regime reports only orderflow as significant, and the usual poor fit of the model towards the end of the sample.

VIII. Non-Parametric Estimation

Given the extent of the variability in the model estimates reported in the previous section, a natural question that arises is that of whether the model itself is inappropriate for estimating the relationship between the stipulated variables. Our prior is that this is not the case. We attempt to look for alternatives that depart from the theory and do not impose the structure of the parametric model. The casual observation of the dprice series reveals its erratic nature. One non-parametric alternative is a kernel. A kernel estimate attempts to find a smooth non-linear function that can fit this series, and give the conditional expected value of the left-hand-side variable as a non-linear function of the right-hand side variables. Because it does not rely on a structural model, a kernel is a viable alternative for estimating the data generating process in the case where the structural model based on theory is unsuccessful.

We proceeded by fitting a multivariate gaussian kernel using the right hand side variables of equation (25) as the arguments of the function. We selected the bandwidth endogeneously based on the scaled comparison of the interquartile range of the datapoints to the sample standard deviation. The results are shown in figure (21). The estimated kernel fit the data poorly as the conditional expected value of the price change. The estimates of the derivatives were of unreasonable magnitudes, due in part to the poor fit of the kernel. The function was unable to pick up the majority of the dynamics in the price change series, and from the figure we can see that it does not look like the data.

Given these results, we maintain that the data generating process for the function is more piecewise linear, and we do not abandon the model of section two.

IX. Conclusion

This paper has looked at the issue of whether the microstructural hypotheses regarding the presence of asymmetric information and inventory effects in the process of price determination in the foreign exchange markets should be modified to take into account relative changes in the intensity of the effects. We find evidence not just of the presence of these effects in the process of price determination in these markets, but rather that there is a tradeoff between these effects. In subsamples where the price is more volatile, the market more hectic, we find inventory to be more important in the pricing decision of the dealer. We find that these effects are muted by the baseline regression. We find that orderflow is more important for the dealer in subsamples when trading is occurring over less dramatic price ranges. Improvements in the estimation and testing of the microstructural hypothesis can be achieved by a piecewise linear specification, and the fit of the model can be unambiguously improved by reducing sample size over fewer trades, where the aforementioned effects are not played off against each other (i.e. trades in the same regime). Finally, we find that the model fails in explaining pricing behavior when the Fed intervenes to force the price of the dollar upwards. In this case the only variable that has any explanatory power is orderflow. When the Fed intervenes, we find that orderflow is significant at much larger magnitudes than during other periods of trading. These results suggest that the dealer inventory and orderflow compete in terms of their influence over prices, and that changing market conditions dictate the relative benefits of using one over the other in pricing foreign exchange.

Previous theoretical work presents the issue in terms of either inventory effects or asymmetric information effects (i.e. orderflow) being the predominant driving force behind foreign exchange prices. Previous empirical work finds that both are present. This paper finds that while both are present, they are not necessarily present simultaneously within samples. Instead, their role as the driving force behind foreign exchange price dynamics adjusts to and is dictated by market conditions.

Appendix A

Figure 2

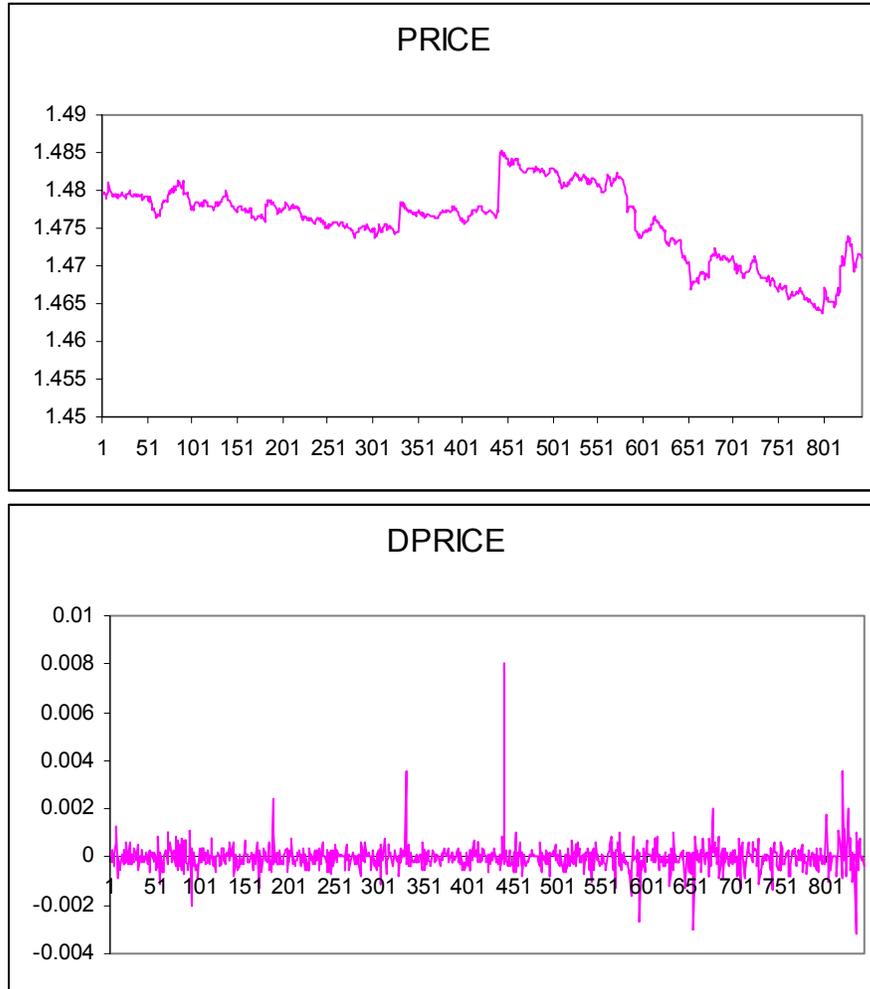


Figure 2. In the upper panel we see the price of one US Dollar in terms of Deutschemarks for the trading week of August 3-7, 1992. In the panel below we see change in price. For both panels, the abscissa gives the observation number, and the ordinate gives the price or price change in 0.0001DM's, also called pips. The discrete jumps in the price (and corresponding outliers in dprice) are overnight changes that we discard. The upturn in the price of the dollar at the end of the week corresponds to the intervention of the Fed.

Figure 3

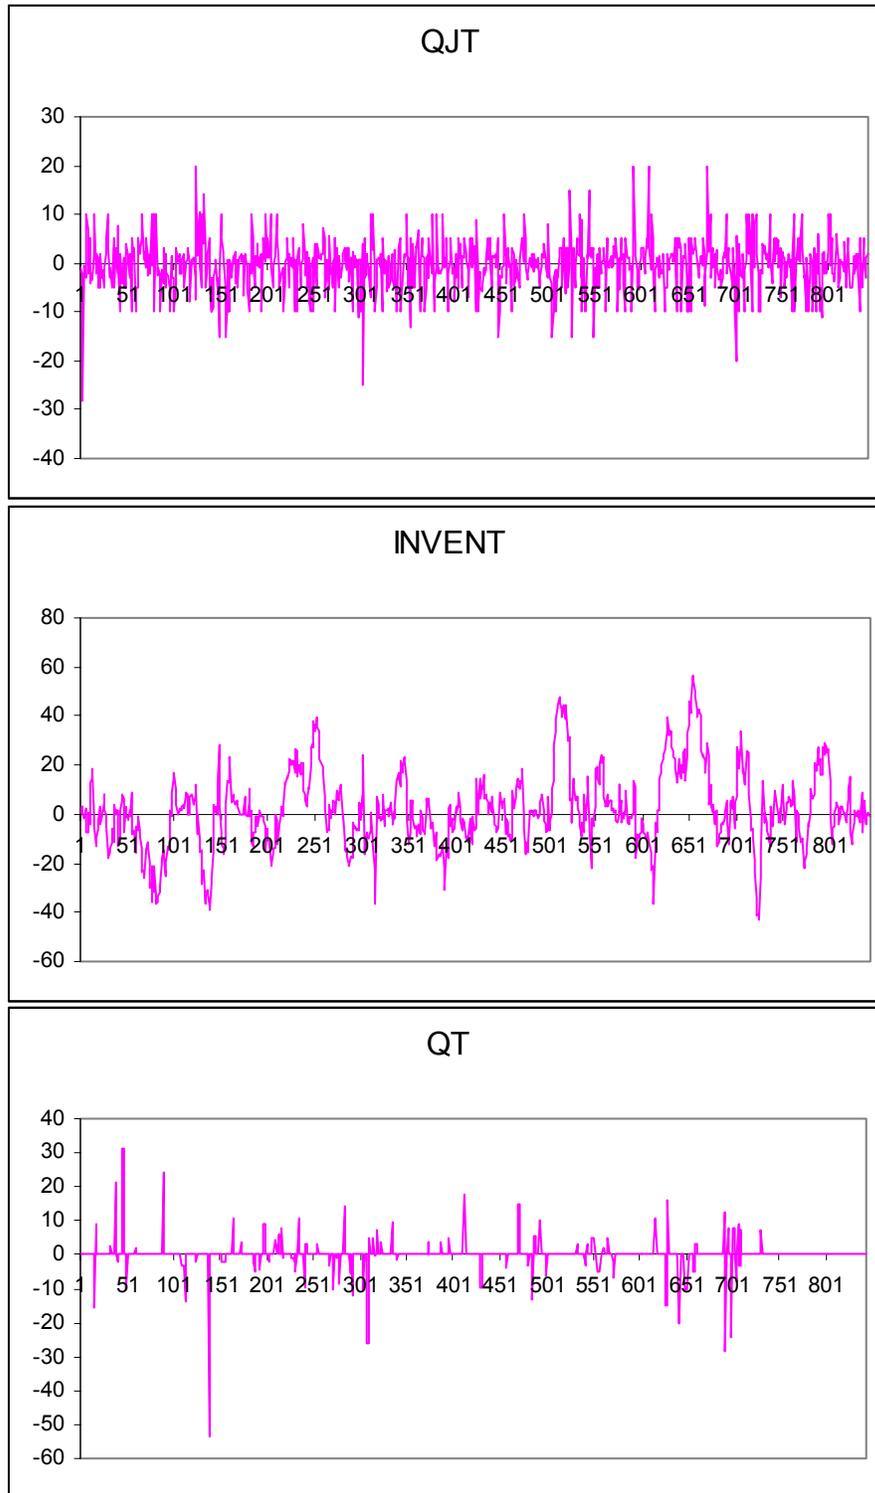


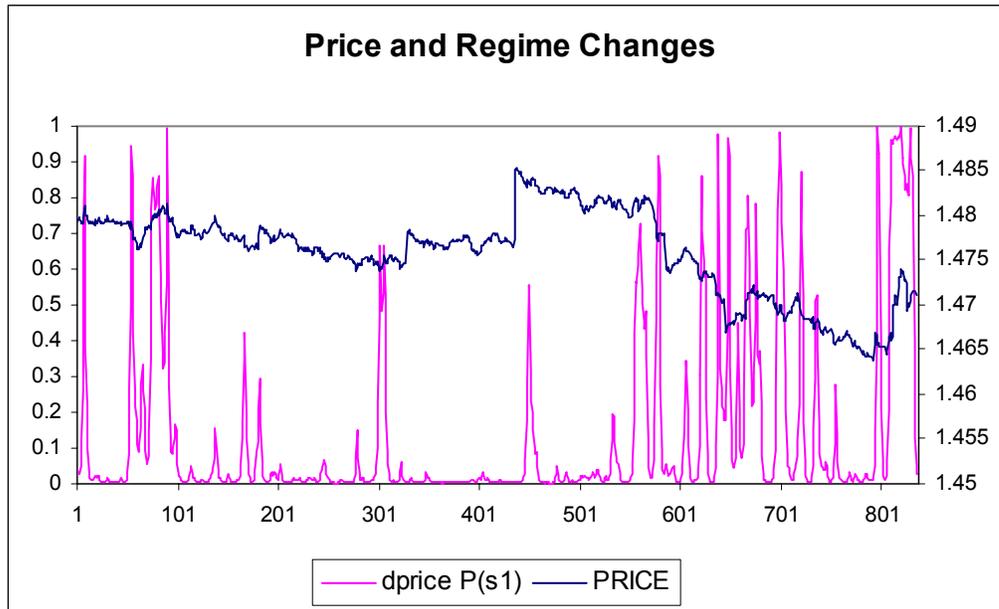
Figure 3. In the upper panel we see orderflow. In the middle panel, inventory, and in the lower panel, brokered trading, for the week of August 3-7, 1992. The abscissa gives the observation number, and the ordinate gives quantity in millions of US dollars.

Figure 4

Variable	Coefficient	Std. Error	t-Statistic	Prob.	Predicted
C	-1.21E-05	1.35E-05	-0.895035	0.371	
QJT	1.47E-05	4.65E-06	3.158934	0.0016	>0
INVENT	-9.22E-06	2.72E-06	-3.391972	0.0007	<0
INVENT(-1)	7.32E-06	2.62E-06	2.793635	0.0053	>0
DEE	0.000101	2.16E-05	4.695455	0	>0
DEE(-1)	-8.98E-05	1.47E-05	-6.120262	0	<0
QT	6.85E-06	3.11E-06	2.199042	0.0282	>0
MA(1)	-0.087047	0.034841	-2.49839	0.0127	<0
R-squared	0.22419	Mean dependent var		-2.13E-05	
Adjusted R-squared	0.217615	S.D. dependent var		0.00047	
S.E. of regression	0.000416	Akaike info criterion		-12.7227	
Sum squared resid	0.000143	Schwarz criterion		-12.6774	
Log likelihood	5313.376	F-statistic		34.09903	
Durbin-Watson stat	1.994299	Prob(F-statistic)		0	
Inverted MA Roots	0.09				

Figure 4. Above are estimated results for equation (25). The last column shows the predicted sign of the coefficients from the model. Below we see the adjusted R2 as well as other goodness of fit statistics.

Figure 5



	Mean	Variance
State 1	-0.000888	0.00358109
State 2	-0.001819	0.000475643
P11	0.785752	
P22	0.955316	

Figure 5. Above are estimated results for Markov switching in the change in price. The left axis measures the probability of a switch, and the right axis measures price. We see price superimposed onto the switching probability. Hence, we see that when the price is turbulent, the probability of being in regime one is very high.

Figure 6

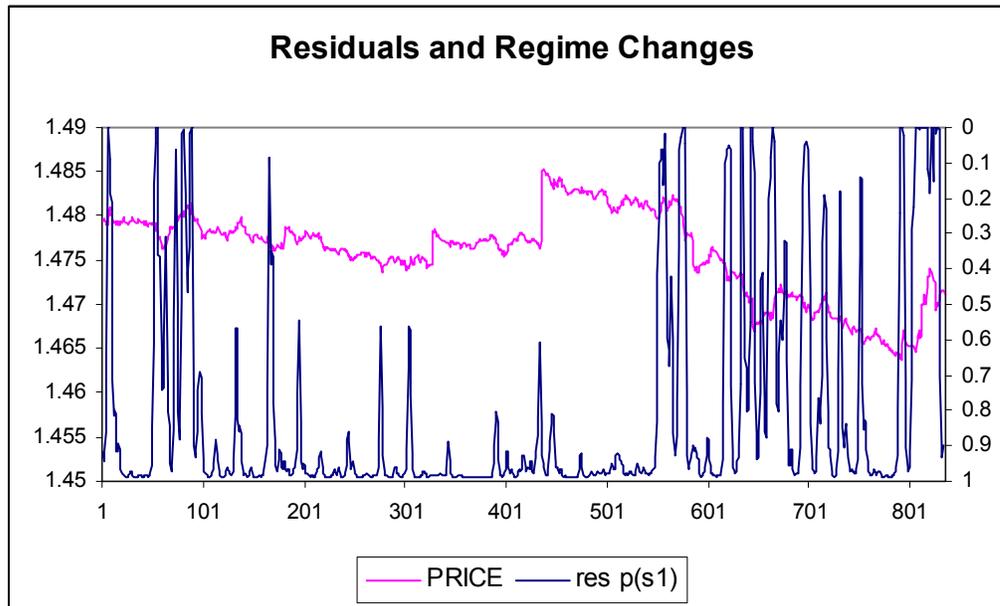


Figure 6. Above are estimated results for Markov switching in the residuals of equation (25) along with the price superimposed. We can observe that the residuals switch at the same time as the change in price. Hence the right hand side variables have not accounted for Markov switching (the right axis is inverted for presentation purposes).

Figure 7

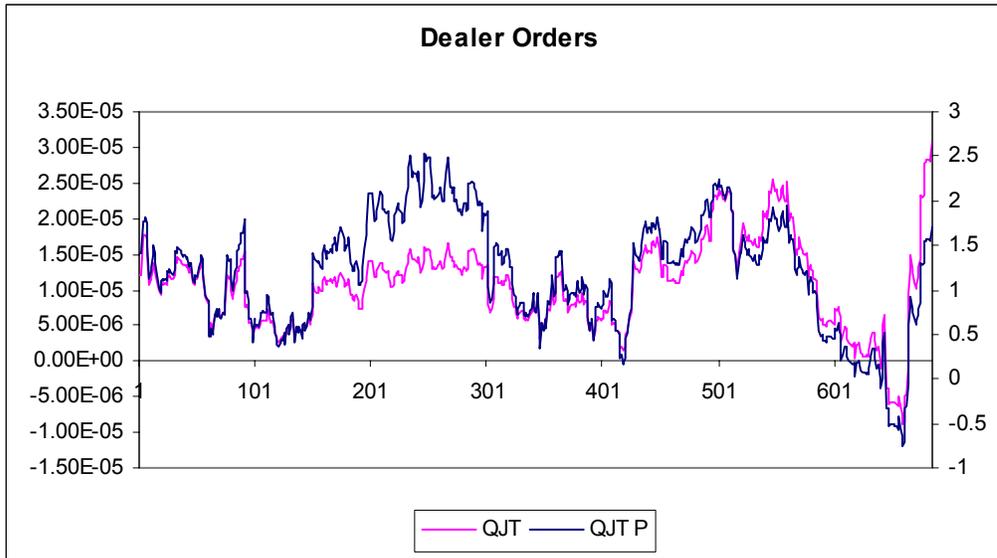


Figure 7. Above we see the coefficient of Q_{jt} from a window of 150 observations beginning with the observation indicated on the abscissa. The left axis measures the value of the coefficient, which is indicated by the Q_{JT} line, and the right axis measures the t -statistic of this coefficient, which is indicated by $Q_{JT P}$.

Figure 8

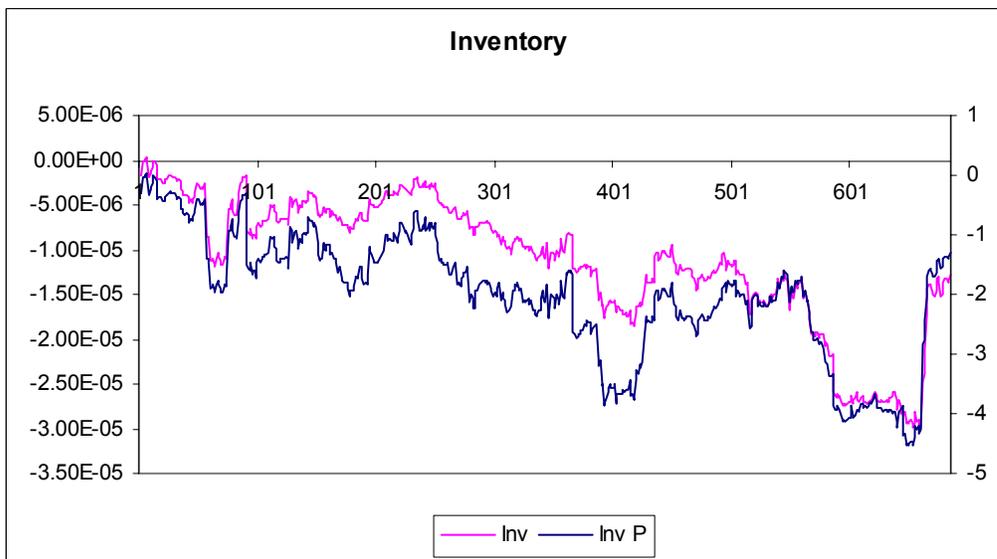


Figure 8. Here we see the coefficient on inventory for the 150 sample-size window rolling regression. INV indicates the coefficient value, which is measured on the left axis, $INV P$ measures its t -statistic, and the abscissa measures the first observation of the window.

Figure 9

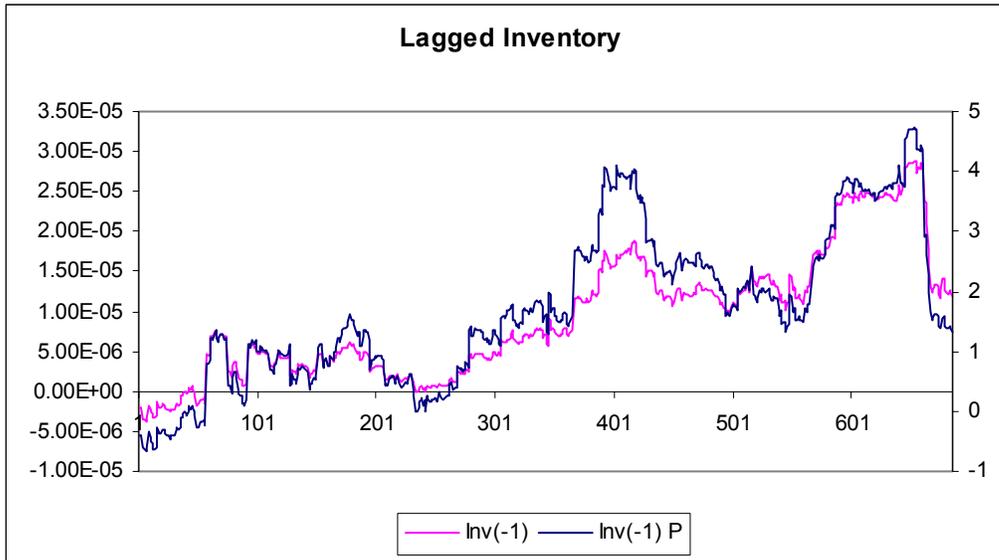


Figure 9. Here we see the coefficient on the rolling regression on the lagged inventory for the 150 sample-size. $INV(-1)$ indicates the coefficient value, which is measured on the left axis, $INV P(-1)$ measures its t-statistic, and the abscissa measures the first observation of the window.

Figure 10

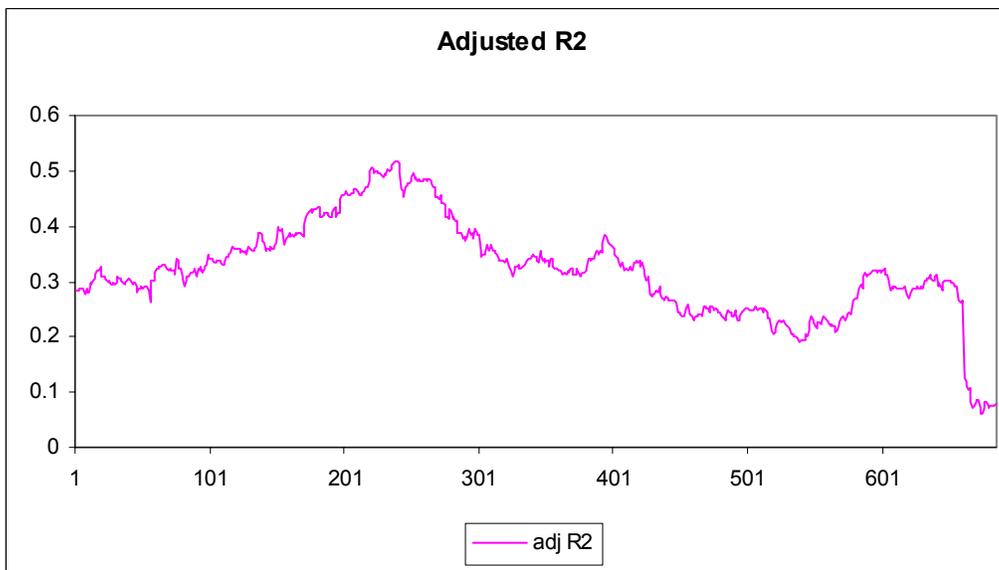


Figure 10. The adjusted R^2 for the rolling regression on a window of 150 observations beginning with the observation indicated in the abscissa. Compared with the original full-sample regression of $R^2=0.22$, the R^2 for the rolling regression is greater than 0.22 throughout most of the sample.

Figure 11

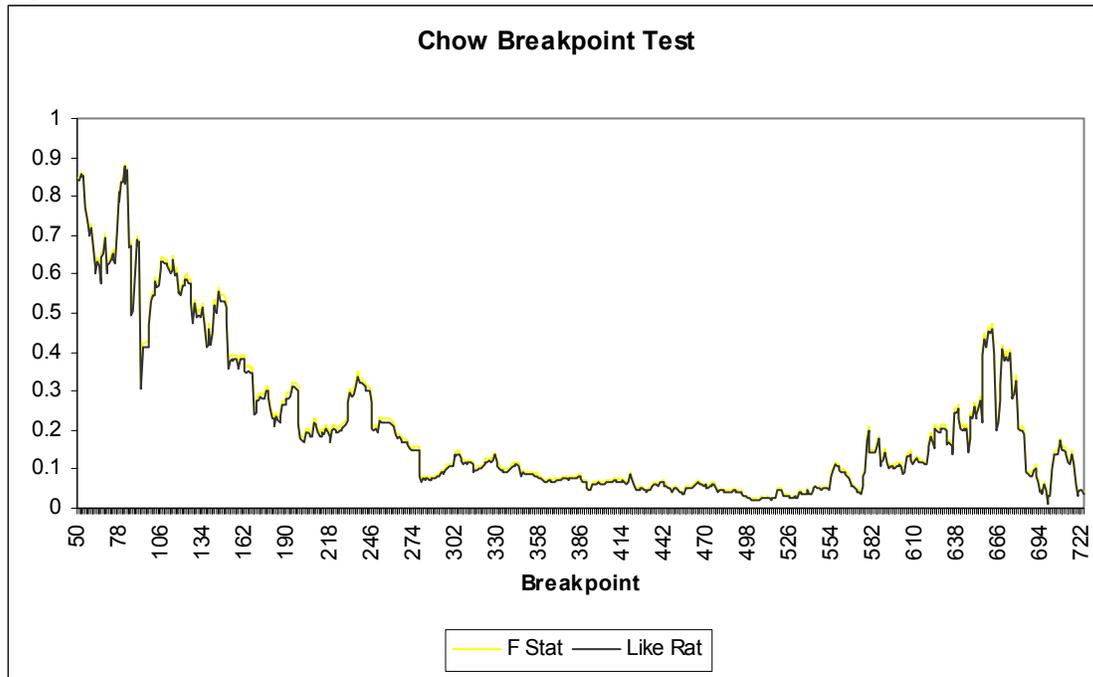


Figure 11. The rolling breakpoint test measures if the point indicated in the abscissa is a breakpoint, with P-value indicated in the ordinate. We calculate this for both the F-test and the Likelihood Ratio test. Above we can see one probable breakpoint in the middle of the sample, indicating a possible break around the end of Wednesday, beginning of Thursday in the one-week sample. Other probable breakpoints are toward the end of the sample at 699 and 722.

Figure 12

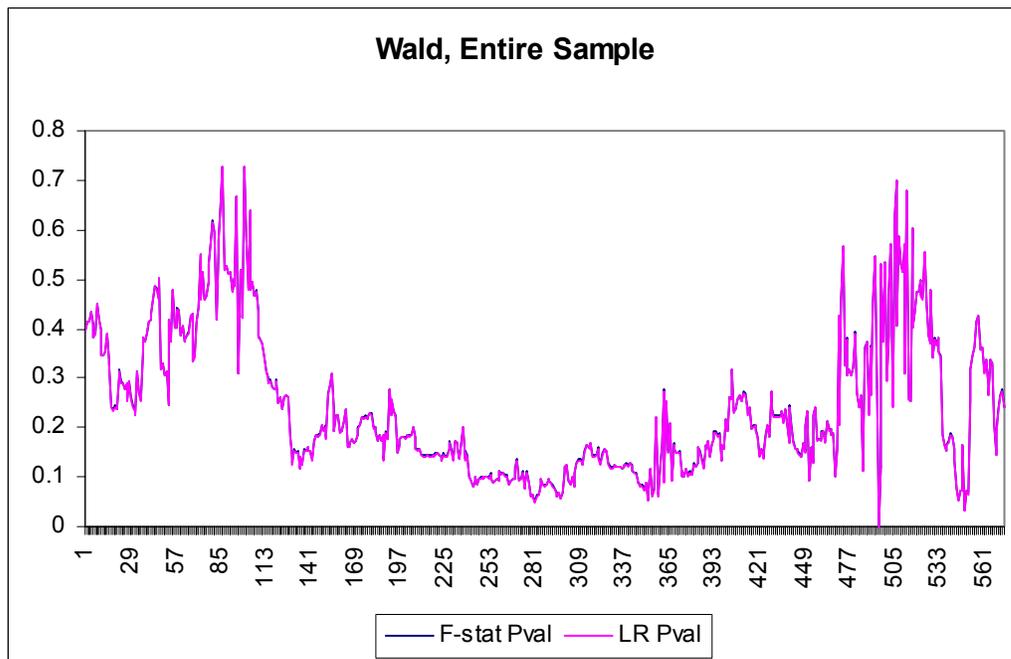


Figure 12. Above we give the specification for the Wald test. The breakpoint is recorded on the abscissa and the p-value on the ordinate. Again, this test serves as evidence supporting breaks in the coefficient values within sample.

Figure 13

Structural Breaks					
All Estimates are without Bt					
Fixed	Break(s)	Point(s)	Fixed	Break(s)	Point(s)
(p=0)			(p=1)		
None	1	449	Q(jt)	1	449
			I(t) (@1%)	3	167
					558
					691
			I(t) (@2.5%)		167
					416
Sample changes days on					558
mon-tue	182				691
tue-wed	331		D(t)	1	446
wed-thu	440		I(t-1)	0	na
thu-fri	593		D(t-1)	0	na

Figure 13. Above we give the results for the Bai and Perron tests for multiple structural breaks. In the left panel we report the test with no fixed regressors across the entire sample. The left column (of that panel) reports 'none' for no fixed regressors, the middle reports the number of break points found, and the right reports the location(s) in the sample of the breakpoint(s). In the right panel we report the results for one fixed regressor. The fixed regressor is given on the left (of the right panel), the number of breaks found is given in the middle, and the location of the break(s) is given on the right. For It results were found at the 2.5% and 1% significance level, and both are reported. For comparison, the dates of overnight changes are reported in column one, so as to clarify that the reported breaks do not (necessarily) correspond to overnight changes.

Figure 14

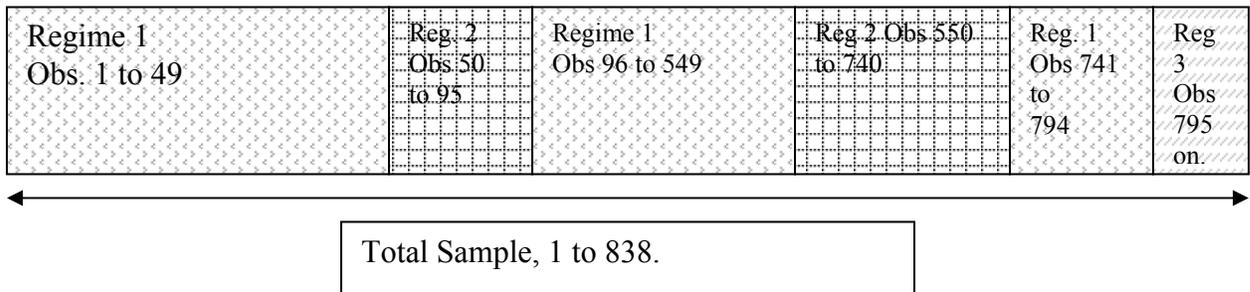


Figure 14. Above is depicted the regime selection based on the smoothed probability of being in state 1 on the observation in question. The first regime is made up of three disjoint segments; they are: (1,49), (96, 549), (741, 794). The second regime has two disjoint segments; they are (50, 95) and (550, 740). The third regime starts at observation 795, and goes to the end of the sample.

Figure 15

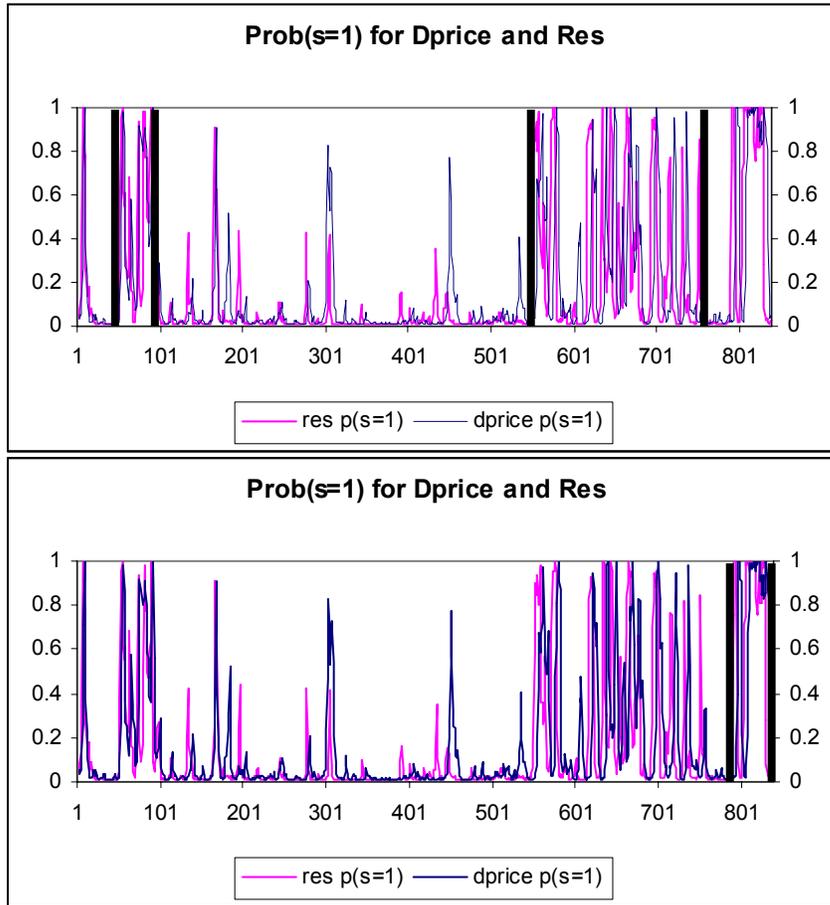


Figure 15. Above we see the graph for regime two in the top panel, and the graph for regime three in the bottom panel. The dark lines show the areas of the state probabilities that determined the regimes. In the upper panel, there are two pairs of dark lines that bound the sample area of regime two. Within these lines, we can see the probability of being in state one approach unity. In the lower panel, at the rightmost part of the graph, we see a pair of dark lines that bound regime three. This is the regime of the Fed intervention, and within we can see the probability jump to unity and maintain the level.

Figure 16

Sample(adjusted): 2 553 IF REGIMES=1									
<i>Variable</i>	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	JUMPS	MA(1)	Rbar2
<i>Coefficient</i>	-2.14	0.941	-0.833	0.697	11.9	-8.25	-79.5	-0.181077	0.383105
<i>Prob.</i>	0.0326	0.0111	0.0002	0.0012	0	0	0	0	
Sample(adjusted): 554 789 IF REGIMES=2									
<i>Variable</i>	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	JUMPS	MA(1)	Rbar2
<i>Coefficient</i>	-3.07	1.81	-1.45	1.18	11.4	-12.9	29	-0.126656	0.23974
<i>Prob.</i>	0.2772	0.0721	0.0095	0.0292	0.018	0.0001	0.5481	0.0585	
Sample(adjusted): 790 838 IF REGIMES=3									
<i>Variable</i>	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	JUMPS	MA(1)	Rbar2
<i>Coefficient</i>	13.9	13.8	1.79	-1.94	-30.9	-4.75		0.099181	-0.034627
<i>Prob.</i>	0.3936	0.0444	0.4228	0.3849	0.214	0.7508		0.5372	
Original Estimation of Equation (25)									
<i>Variable</i>	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	JUMPS	MA(1)	Rbar2
<i>Coefficient</i>	-1.25	1.43	-0.914	0.721	10.4	-9.21		-0.088931	0.219641
<i>Prob.</i>	0.3556	0.0021	0.0008	0.0062	0	0		0.0108	

Figure 16. Above se see the estimation results for the estimation of regimes one, two and three, and also for the original equation. Each panel shows the sample size, the coefficient values, and the p-value of the t-statistic. All coefficients are multiplied by $10exp4$ except the MA(1).

Figure 17

Chow Breakpoint Test: 50 96 551 741 795			
F-statistic	1.5924	Probability	0.017055
Log likelihood r	56.718	Probability	0.011538

Figure 17. F-test for parameter stability, we reject the null hypothesis of no breaks in favor of our stipulated breaks at 1% significance level.

Figure 18

Bai and Perron Sup-F Statistic								
Sample(adjusted): 2 448								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-1.7	1.3	-0.354	0.122	12.5	-8.85	-0.20472	0.324832
<i>Prob.</i>	0.1614	0.0046	0.2031	0.647	0	0	0	
Sample: 449 794								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-3.17	0.897	-2.04	1.86	11	-11.2	-0.1023	0.29919
<i>Prob.</i>	0.1375	0.1885	0	0	0.0006	0	0.0612	
Sample: 795 838								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	16.9	14.1	3.32	-2.66	-25.3	0.292	0.107834	-0.04893
<i>Prob.</i>	0.3459	0.0643	0.3796	0.4242	0.3678	0.9864	0.5313	
Sample: 449 838								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-0.621	1.73	-1.64	1.45	7.16	-10.1	-0.04547	0.174611
<i>Prob.</i>	0.8189	0.0431	0.0008	0.0018	0.0658	0.0002	0.375	
Original Estimation of Equation (25)								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-1.25	1.43	-0.914	0.721	10.4	-9.21	-0.08893	0.219641
<i>Prob.</i>	0.3556	0.0021	0.0008	0.0062	0	0	0.0108	

Figure 18. Above the sample is subdivided into three areas, with breaks at observations 449 and 795, per Sup-F test. The top panel reports the first segment, observations 1 to 448. The second panel reports from observations 449 to 794. The third panel reports from observation 795 to 838. The fourth panel reports results from observation 449 to 838, in effect skipping the stipulated break at 795. The final panel reports the original estimation of equation (25).

Figure 19

Chow Breakpoint Test: 449 795			
F-statistic	4.173755	Probability	0
Log likeliho	57.88762	Probability	0

Figure 19. F-test for parameter stability, we reject the null hypothesis of no breaks in favor of our stipulated breaks at 1% significance level.

Figure 20

F-Test Breakpoints								
Sample(adjusted): 2 508								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-1.91	1.27	-0.488	0.227	12.6	-8.67	-0.21988	0.344391
<i>Prob.</i>	0.0781	0.0029	0.0603	0.3636	0	0	0	
Sample: 509 699								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-2.15	1.28	-1.65	1.48	10.5	-12.5	-0.05374	0.227315
<i>Prob.</i>	0.5654	0.1924	0.0057	0.0102	0.0335	0.0003	0.4739	
Sample: 700 722								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-8.53	-4.58	-5.53	5.13	30.2	1.99	-0.98995	0.655636
<i>Prob.</i>	0.0076	0.2538	0.0002	0.0002	0.1367	0.8141	0	
Sample: 723 838								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	3.74	4.89	-0.457	0.463	-5.54	-6.62	0.048266	0.045433
<i>Prob.</i>	0.5995	0.0495	0.7288	0.7013	0.5909	0.323	0.6171	
Original Estimation of Equation (25)								
	C	QJT	INVENT	INVENT(-1)	DEE	DEE(-1)	MA(1)	Rbar2
<i>Coefficient</i>	-1.25	1.43	-0.914	0.721	10.4	-9.21	-0.08893	0.219641
<i>Prob.</i>	0.3556	0.0021	0.0008	0.0062	0	0	0.0108	

Figure 20. Breakpoints based on Rolling F-test for parameter stability. The breakpoints are given as the first point in each subsample.

Figure 21

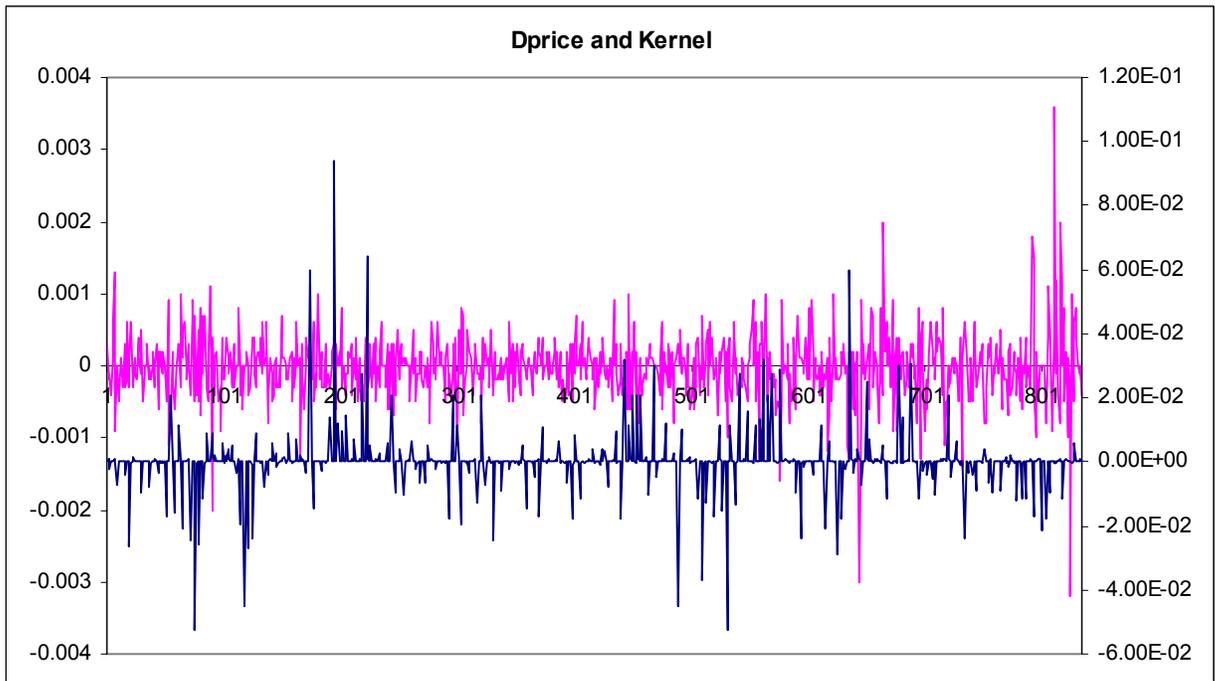


Figure 21. Above are is the Kernel with the series dprice. The higher series is the price change for the week of August 3-7, 1992. The lower series is the estimated kernel.

Appendix B

Here we present a brief summary of the economic conditions on the week of August 3 through August 7, 1992.

Monday, August 3, 1992: In early June, the Dow Jones Industrial Average had climbed to an all time high above 3400. The dollar was weak as the US economy was recovering from the 1990-1991 recession. On July 16, 1992 the Bundesbank raised its discount rate by 75 basis points to 8.75%. This forced the Federal Reserve to mount an intervention campaign. The Bundesbank was trying to fight inflation, reduce the money supply, and at the same time deal with the fiscal expansion resulting from the reunification of East Germany. It was reported in the financial press that traders and foreign exchange analysts predicted that the July 1992 US employment report would show improvement in the US economy, and hence push the dollar upward. It was clear that this report would determine the direction of the dollar. Japanese traders were concerned with Tokyo share prices, and the future direction of the US economy as well. Trading was confined to dealers, with many customers such as exporting companies absent from the market. Trading settled into a narrow range, and the market was described as dull.

Tuesday, August 4, 1992 – The dollar was quoted unchanged from the previous session. The bank of Italy cut the discount rate to 13.25%, but there was no significant move on the Dollar/DM market by day's end. Trading was described as lethargic in the traditionally slow month of August, as dealers marked time for the release of the

employment report. The currency was locked into a narrow trading range. The Japanese markets were in the same narrow trading range, and the financial press reiterated expectations of a dollar increase as the employment report for July was released. The summer holidays also contributed to the narrow and quiet market. The threat by an unnamed official from the Bank of Italy of intervention in the market to keep the dollar from going lower boosted the dollar in European trading. Later, when it was reported that US Construction spending fell by 1.5% in June, and the National Association of Purchasing Managers Index rose to 54.2% from 52.8%, with supplier deliveries the slowest since March 1989, the dollar lost ground in European trading, and again in New York. There was a question about whether central bankers, particularly the Fed, were willing to see the dollar fall past its historic low of DM1.4470. The threat of intervention is what kept short sales at bay.

Wednesday, August 5, 1992 – On Tuesday the dollar inched higher with most of the market awaiting the July employment report set to be released Friday. The markets were reported to have shrugged off the index of US leading indicators for June. The employment report was described as the watershed in market sentiment toward the dollar for the rest of the month. One report described the potential for disappointment as having the ability to bring the selling out and to test the central banks' resolve to defend the dollar. In Japan, trading was light, and the dollar closed almost unchanged as people optimistically awaited the employment numbers. Gains in the Tokyo stockmarket had no effect on the dollar. Dealers also awaited the meeting of Bundesbank officials on Thursday to see if they would continue to raise interest rates, however this was generally

viewed as unlikely. Optimism over the employment report seemed to be the overriding market sentiment.

Thursday, August 6, 1992 – The dollar drifted higher in New York trading on Wednesday, gaining against the DM, as well as the yen. On Thursday morning, Japanese trading was directionless, and good employment numbers were reported to be widely expected. The dollar ended up closing higher against the yen in Japanese trading. The British pound fell, and unemployment numbers for Germany were reported to be higher.

Friday, August 7, 1992 – Thursday's US stocks fell on disappointing earnings reports for the third time in three sessions. The dollar ended mostly lower after New York trading Thursday, as participants took profits ahead of the employment report that was released that Friday. Most economists were reported to be forecasting an improvement in the nation's employment situation. Dealers believed that if the report met expectations, the dollar could go to DM1.5. The dollar also received a boost from the decision by the Bundesbank to not raise interest rates on Thursday. On Friday, the dollar eased against the DM in Japanese trading.

Saturday, August 8, 1992 – Shortly after the London markets closed on Friday, the Federal Reserve bought dollars four or five times at levels ranging from DM1.4675 to DM1.473. The intervention lifted the dollar, but it drifted lower by the end of New York trading. This was due to a disappointing employment report released on Friday. The dollar was reported to have risen quickly when the Fed intervened, climbing one pfennig

in ten minutes, from 1.47DM to 1.475DM. Comments by Helmut Schlesinger suggested that the concerted central bank intervention in July had been a one-off event, and this caused traders to test the resolve of the central bank to avoid new lows for the Dollar.

References

- Bai, Jushan. "Likelihood ratio tests for multiple structural changes." *Journal of Econometrics* 91 (1999) 299-323.
- Bai, Jushan and Pierre Perron. "Estimating and testing linear models with multiple structural changes." *Econometrica* 66 (1998) 47-48.
- Business Week*; various articles.
- Campbell, John Y., Andrew W. Lo, A. Craig MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press (1997) Princeton, New Jersey.
- The Daily Telegraph*; various articles.
- Engle, Charles and James D. Hamilton. "Long Swings in the Dollar: Are They in the Data and Do Markets Know It?" *The American Economic Review* 80 (1990) 689-713.
- Engle, Charles. "Can the Markov Switching Model Forecast Exchange Rates?" *Journal of International Economics* 36 (1994) 151-165.
- The Financial Times (London); various articles.
- Fomby, Thomas B. and R. Carter Hill. **Applying Kernel and Nonparametric Estimation to Economic Topics**. Jai Press Inc. (2000) Stamford, CT.
- Hamilton, James D. "A Quasi-Bayesian Approach to Estimating Parameters for Mixtures of Normal Distributions." *Journal of Business and Economic Statistics* 9 (1991) 27-39.
- Hamilton, James D. and Raul Samusel. "Autoregressive Conditional Heteroskedasticity and Changes in Regime." *Journal of Econometrics* 64 (1994) 307-333.

Hamilton, James. **Time Series Analysis**. Princeton University Press (1994), Princeton, New Jersey.

Hansen, Bruce E. “The New Econometrics of Structural Change: Understanding and Dating Changes in US Productivity.” University of Wisconsin typescript, January 20001.

The Irish Times; various articles.

Jiji Press Ticker Service; various articles.

Koning, Ruud H. *Kernel: a Gauss Library for Kernel Estimation*.

www.xs4all.nl/~rhkoning/gauss.htm.

Kyodo News Service, Japan Economic Newswire; various articles.

Lyons, Richard K. “Tests of Microstructural Hypotheses in the Foreign Exchange Market.” *Journal of Financial Economics* 39 (1995) 321-351.

Lyons, Richard K. **The Microstructure Approach to Exchange Rates**. MIT Press, forthcoming.

The New York Times; various articles.

O’Hara, Maureen. **Market Microstructure Theory**. Blackwell Publishers, (1995) Malden, MA.

Tay, Anthony S. and Kenneth F. Wallis. “Density Forecasting: A Survey.” *Journal of Forecasting* 19 (2000) 235-254.

The Times; various articles.

The Toronto Star; various articles.

Wall Street Journal; various articles.