

Daily Changes in Fed Funds Futures Prices

James D. Hamilton, Professor of Economics

University of California, San Diego

jhamilton@ucsd.edu

December 19, 2006

Revised: July 6, 2007

JEL codes: E40, E50, G13

Keywords: federal funds, futures, monetary policy

Manuscript number 07-009

This paper has benefitted from helpful suggestions by anonymous referees.

ABSTRACT

This paper explores the properties of daily changes in the prices for near-term fed funds futures contracts. The paper finds these contracts to be excellent predictors of the fed funds rate, and shows that the claim of a nonzero term premium in the short-horizon contracts is more sensitive to outliers than previous research appears to have recognized. I find some statistically significant evidence of serial correlation in the daily changes, but this accounts for only a tiny part of the one-day movements and there is essentially zero predictability for horizons longer than one day. Settlement futures prices for each day appear to incorporate the information embodied in that day's term structure of longer-horizon Treasury securities. Previous employment growth makes a statistically significant contribution to predicting futures price changes, though again this could only account for a tiny part of the daily variance. The paper concludes that futures prices provide a very useful measure of the daily changes in the market's expectation of near-term changes in Fed policy.

1 Introduction.

The federal funds rate is of considerable interest in economics and finance, both because it defines the shortest end of the term structure (the overnight rate being the shortest-maturity U.S. asset traded) and because it is the rate directly targeted and controlled by the Federal Reserve. Futures contracts based on the fed funds rate have come to be regarded as useful indicators of what the market expects future monetary policy to be (Krueger and Kuttner, 1996; Gürkaynak, Sack, and Swanson, 2007). Changes in futures prices hold particular promise for trying to assess the daily economic consequences of changes in Fed policy (Kuttner, 2001; Gürkaynak, 2005; Hamilton, 2006).

However, there are potential drawbacks to using these numbers for such purposes. A number of researchers have documented the existence of what appear to be time-varying term premia in longer-horizon fed funds futures contracts (Sack, 2004; Piazzesi and Swanson, 2006). To my knowledge, there is no systematic investigation of how big a contribution such effects may make to daily changes in futures prices.¹ This paper addresses this gap, reviewing the time-series properties of daily changes in fed funds futures prices at 1- to 3-month horizons.

Section 2 begins with a description of these contracts and a review of why one might expect daily changes in the prices of near-horizon contracts to be an approximate martingale difference sequence whose innovations are dominated by changes in the expected value of the fed funds rate over the contracted month. Section 3 documents that, particularly in the most recent data, the futures prices are remarkably good predictors of what the fed funds

rate actually turns out to be. Section 4 reviews the evidence for a nonzero term premium in the short-horizon contracts, and finds that it is more sensitive to outliers than previous research appears to have recognized. Section 5 looks for evidence of time-varying premia in the form of violations of the martingale difference condition. I find some statistically significant violations, though any predictable components appear to comprise a very small part of the daily movements. I conclude that a martingale is a good approximation for these data, and that it is quite reasonable to interpret daily changes in fed funds futures as primarily signalling changes in the market's assessment of future changes in Fed policy.

It should be noted that both the theoretical and the empirical analysis here are limited to very short-horizon futures contracts. The strongest empirical evidence against the martingale hypothesis from previous studies using monthly data has come from longer-horizon futures prices. Such longer-horizon contracts may also pose a particular problem for daily policy analysis, being thinly and erratically traded in the earlier part of the sample. However, for purposes of measuring the daily changes in the market's expectation of very near-term changes in Fed policy, the results of this paper suggest that futures prices are an excellent measure.

2 Fed funds futures data.

In a typical macroeconomic study using monthly data, the measure of the fed funds rate for month m is based on the average value of the daily effective fed funds rate over all the calendar days of the month. In theoretical discussions, it will be convenient to measure \tilde{r}_m

as a fraction of unity, whereas, because the daily changes in these magnitudes are so small, empirical results in this paper are all reported for r_m measured in basis points. Thus if the average effective fed funds rate for month m would be quoted at an annual interest rate as 5%, this will be represented in the notation of this paper as $\tilde{r}_m = 0.05$ and $r_m = 500$.

Since October 1988, the Chicago Board of Trade has hosted daily trade in futures contracts with cash settlement based on what the value of \tilde{r}_m actually turns out to be. Let \tilde{F}_{dm} denote the implied interest rate² for a contract based on the value \tilde{r}_m as traded on day d , where d denotes some banking day during or prior to month m . If \tilde{r}_m (which will not be known until the end of month m) turns out to be less than \tilde{F}_{dm} (which is known as of day d), the seller of the contract has to compensate the buyer by a certain amount (namely, \$41.67 on a standard contract) for every basis point by which \tilde{F}_{dm} exceeds \tilde{r}_m . If $\tilde{r}_m > \tilde{F}_{dm}$, the buyer must compensate the seller.

Let $d^*(m)$ denote the first business day of month m . Cash settlement of a contract based on \tilde{r}_m held to completion would take place on day $d^*(m+1)$. Let $1 - \tilde{\lambda}_{dm}$ denote the pricing kernel relating the current day d to a security paying off on $d^*(m+1)$. For example, with additively separable consumer preferences based on a daily utility function $U(c_d)$,

$$1 - \tilde{\lambda}_{dm} = \frac{\beta^{d^*(m+1)-d} U'(c_{d^*(m+1)}) / P_{d^*(m+1)}}{U'(c_d) / P_d} \quad (1)$$

for P_d the dollar price of a unit of consumption on day d .

If futures contracts were pure forward contracts in which no money changes hands until this final settlement, the cost of entering into the futures contract on day d would be zero. If you take the long side of the contract, the payoff on the first day of the month following

month m would be proportional to $\tilde{F}_{dm} - \tilde{r}_m$. Standard finance theory would then require the agreed-upon price \tilde{F}_{dm} to satisfy

$$E_d \left[(1 - \tilde{\lambda}_{dm})(\tilde{F}_{dm} - \tilde{r}_m) \right] = 0 \quad (2)$$

where $E_d[\cdot]$ denotes an expectation formed on the basis of information available on day d . Since \tilde{F}_{dm} is known with certainty on day d , it follows from (2) that

$$\tilde{F}_{dm} - \tilde{F}_{d-1,m} = E_d(\tilde{r}_m) - E_{d-1}(\tilde{r}_m) + \tilde{h}_{dm} - \tilde{h}_{d-1,m} \quad (3)$$

for

$$\tilde{h}_{dm} = \tilde{F}_{dm} E_d(\tilde{\lambda}_{dm}) - E_d(\tilde{\lambda}_{dm} \tilde{r}_m).$$

Recall that $E_d(1 - \tilde{\lambda}_{dm})$ is the reciprocal of the risk-free gross interest rate between day d and $d^*(m+1)$. As the length of calendar time separating these days shrinks, this gross discount rate converges to unity and $E_d(\tilde{\lambda}_{dm})$ converges to zero. For example, if the annual interest rate is 6% and we are looking at a payoff 1 month ahead, $E_d(\tilde{\lambda}_{dm}) \simeq 0.005$, which is an order of magnitude smaller number than \tilde{F}_{dm} or \tilde{r}_m . Likewise, as d and $d^*(m+1)$ get closer in time, the uncertainty about the future marginal utility of consumption becomes resolved and the variance of $\tilde{\lambda}_{dm}$ has to go to zero as well. For example, for a diffusion process, $E_d \left[\tilde{\lambda}_{dm} - E_d(\tilde{\lambda}_{dm}) \right]^2 = O_p\{[d^*(m+1) - d]\}$. Hence for the very near-term contracts, the contribution of the risk and term premium (\tilde{h}_{dm}) should be significantly smaller than the first two variables in (3), in which case, \tilde{F}_{dm} would approximately follow a martingale, with the innovation in this martingale corresponding to new information that market participants receive on day d about the value of \tilde{r}_m . The futures value \tilde{F}_{dm} would follow an exact

martingale in the special case of risk neutrality.

Because of margin requirements and the daily marking of the contract to market, these are not true forward contracts, and the pricing of these contracts is in theory more complicated than the above simple formulas. However, Piazzesi and Swanson (2006) demonstrated that adjustments to returns based on marking to market are likely to make very little difference in practice.

The basic data used in this study are the daily changes in the prices of contracts within a few months of settlement over the period October 3, 1988 through June 30, 2006. Specifically, let f_{1d} denote the change (in basis points) between day $d - 1$ and d in the implied interest rate for the “spot-month contract,” that is, the contract that will settle at the end of the current month,³

$$f_{1d} = F_{d,m^*(d)} - F_{d-1,m^*(d)}$$

where $m^*(d)$ is the month within which day d falls. Let f_{2d} denote the change on day d of a contract for settlement at the end of the following month,

$$f_{2d} = F_{d,1+m^*(d)} - F_{d-1,1+m^*(d)}$$

and f_{3d} the change for the month after next:

$$f_{3d} = F_{d,2+m^*(d)} - F_{d-1,2+m^*(d)}.$$

If daily changes in \tilde{h}_{dm} are negligible, then from (3), f_{id} would follow a martingale difference sequence reflecting new information that market participants receive on day d about the likely value of r_m . The following sections review the evidence on this hypothesis.

3 Usefulness of futures for predicting fed funds rate.

As documented by Krueger and Kuttner (1996), Rudebusch (1998), Kuttner (2001), and Gürkaynak, Sack, and Swanson (2007), one can obtain an excellent forecast of r_m on the basis of F_{dm} for near-term contracts. Let e_{1m} denote the difference between the average fed funds rate for month m and the futures rate as of the last day of the preceding month, measured in basis points,

$$e_{1m} = r_m - F_{d^*(m)-1,m}.$$

Let e_{im} denote the corresponding forecast error using the futures rate i months ahead. The second column of Table 1 gives the average squared value of the forecast error e_{im} based on the futures forecast looking $i = 1, 2$, or 3 months ahead. For example, for a 2-month ahead forecast the futures MSE is

$$(T_M - 2)^{-1} \sum_{m=3}^{T_M} e_{2m}^2 = (T_M - 2)^{-1} \sum_{m=3}^{T_M} [r_m - F_{d^*(m-1)-1,m}]^2 = 392.$$

One natural basis for comparison is provided by a random walk, in which the predicted funds rate for month $m + s$ is the value obtaining in month m . This no-change forecast is often extremely hard to beat out-of-sample for financial data. These baseline MSEs are reported in the first column in Table 1. For example, forecasting 2 months ahead with a random walk model would result in a mean squared forecast error over the full sample of

$$(T_M - 2)^{-1} \sum_{m=3}^{T_M} (r_m - r_{m-2})^2 = 1248,$$

or a root mean squared error of 35 basis points. Thus the futures prices represent a 69% improvement over the no-change forecast in terms of the 2-month-ahead MSE. Similar

improvements in forecasts are obtained at 1-month and 3-month horizons.

There has been a substantial improvement over time in the accuracy of the predictions embodied in fed funds futures, as noted by Poole and Rasche (2000), Poole, Rasche and Thornton (2002), Lang, Sack and Whitesell (2003), Carlson, et. al. (2006), and Swanson (2006), among others. In part this reflects the fact that fed funds changes themselves have become more modest (see the bottom panel, first column of Table 1). Even so, the quality of futures forecasts have improved far more than proportionally (column 2). Over the last 3-1/2 years, even looking 3 months ahead, the fed funds futures have an astonishing root MSE of 6.9 bp and an average absolute error of only 5.4 bp, this for predicting a series r_m whose average absolute 3-month change is 36 bp.⁴

Of particular interest is how futures prices have responded to the news of changes in the target for the fed funds rate as announced in periodic news releases from the FOMC. Hamilton (2006) noted that on each of the 15 most recent occasions that the Federal Reserve changed its target for the fed funds rate, the fed funds futures price for that month changed by less than half a basis point. In other words, over the last 3 years, the market has known with virtual certainty what the Fed was going to do well before the Fed actually changed the rate.

4 The possible bias in fed funds futures prices.

Several studies including Sack (2004) and Piazzesi and Swanson (2006) have noted a systematic tendency of the fed funds futures to overestimate the value of r_m , with the bias

increasing with the forecast horizon. The first three columns of Table 2 reproduce this result, finding an average value for f_{id} between negative one- and two-tenths of a basis point per day over the full sample. The usual t-test suggests that this bias is highly statistically significant.

Figure 1 displays the sample histogram for f_{1d} , drawn for comparison with the Normal distribution. Forty-six percent of the observations are exactly zero, while 25 observations exceed 5 standard deviations. If f_{id} were an i.i.d. Gaussian time series, one would not expect to see even one 5-standard-deviation outlier. Often these outliers occur on days that Gürkaynak, Sack, and Swanson (2005) associated with significant monetary policy announcements.

Figure 2 plots the actual time series for f_{1d} . In addition to the extreme outliers, one sees in this graph that the variance has a very clear declining trend over time, and that there is also serial correlation in the variance suggestive of strong GARCH effects. One can also see a trend in the variance within each month: given the nature of the discovery process, r_m is largely known by the last day of the month, and the first few days of each month are often characterized by a bigger variance for f_{id} .

The theory sketched in Section 2 suggests that f_{id} should follow a martingale difference sequence, but this does not mean the data must be Normal or even i.i.d. I was interested to see how the inference about the mean might change if one allows for predictability of the variance and departures from the Normal distribution. I modeled the distribution as a mixture in which some fraction p of the values are exactly zero, and the remaining $(1 - p)$

are drawn from a Student t distribution with scale parameter h_t and ν degrees of freedom.⁵ Specifically, for y_t corresponding to f_{id} and $\mathbf{Y}_{t-1} = (y_{t-1}, y_{t-2}, \dots, y_1)'$ the set of observations through date $t - 1$, the conditional likelihood of the t th observation is taken to be

$$f(y_t | \mathbf{Y}_{t-1}; \boldsymbol{\theta}) = p\delta_{y_t=0} + (1-p)(1-\delta_{y_t=0})g(y_t; h_t, \boldsymbol{\theta})$$

$$\delta_{y_t=0} = \begin{cases} 1 & \text{if } y_t = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$g(y_t; h_t, \boldsymbol{\theta}) = \left(k_{1\nu} / \sqrt{h_t} \right) [1 + (u_t^2/\nu)]^{-(\nu+1)/2} \quad (4)$$

$$k_{1\nu} = \Gamma[(\nu + 1)/2] / [\Gamma(\nu/2)\sqrt{\nu\pi}]$$

$$u_t = (y_t - \mu) / \sqrt{h_t}. \quad (5)$$

The maximum likelihood estimate of p will then be equal to the fraction of observations that are equal to zero, while a low value for ν , the degrees of freedom for the Student t distribution, could allow for the tendency for big outliers.

The conditional scale factor h_t is modeled as following Nelson's (1991) EGARCH(1,1) specification,

$$\log h_t - \boldsymbol{\gamma}'\mathbf{z}_t = \delta(\log h_{t-1} - \boldsymbol{\gamma}'\mathbf{z}_{t-1}) + \alpha(|u_{t-1}| - k_{2\nu}),$$

where $k_{2\nu}$ denotes the expected absolute value for a standard Student t variable with ν degrees of freedom, which can be shown to be given by

$$k_{2\nu} = E|u_t| = \frac{2\sqrt{\nu}\Gamma[(\nu + 1)/2]}{(\nu - 1)\sqrt{\pi}\Gamma(\nu/2)}.$$

The vector \mathbf{z}_t contains deterministic calendar variables that influence the variance. These consisted of a constant, a scaled time trend ($d/1000$ for daily data), and a dummy that

is equal to unity if d is one of the first two days of the month. The scaling on the time trend was used to make the numerical maximization a little better behaved (so that values of different parameters were not of starkly different orders of magnitudes)⁶.

The parameter vector $\boldsymbol{\theta} = (\mu, \gamma_1, \gamma_2, \gamma_3, \delta, \alpha, p, \nu)'$ was then found by numerically maximizing the sample log likelihood,

$$\sum_{t=1}^T \log f(y_t | \mathbf{Y}_{t-1}; \boldsymbol{\theta}). \quad (6)$$

Maximum likelihood estimates for y_t corresponding to f_{1d} , f_{2d} , and f_{3d} are given in columns 4-6 of Table 2, respectively. Asymptotic standard errors (based on second derivatives of (6) as in Hamilton 1994, equation [5.8.3]) are given in parentheses.

The addition of 6 parameters $(\gamma_2, \gamma_3, \delta, \alpha, p, \nu)'$ relative to an i.i.d. Gaussian specification increases the log likelihood by about 1,500; (a difference of only 12 would be enough to establish statistical significance at the 0.001 level). The t statistic on each of the individual new parameters is typically around 10, indicating that there is overwhelming evidence in the data for all of these effects.

These estimates imply that, other things equal, an observation f_{1d} at the beginning of the sample would be 10 times more variable than an observation at the end of the sample ($\exp[(0.52568)(4454)/1000] = 10.4$, and the first two days of each month nearly twice that of others ($\exp[(0.63570)] = 1.89$). The EGARCH effects are highly persistent, with autoregressive parameter of 0.83, and the innovations are quite fat-tailed, with the estimated degrees of freedom for the Student t distribution around 4.

Of particular interest is the fact that, with this more involved specification of the distrib-

ution and variance, the maximum likelihood estimate of the mean of each f_{id} becomes much smaller in absolute value or even positive, and would be judged in every case to be far from statistically significant. In other words, the negative sample means for f_{id} are dominated by the fact that the outliers are more likely to be negative than positive, and an estimation method such as MLE that downweights the outliers would no longer see evidence of any negative bias. The sample medians of f_{1d} , f_{2d} , and f_{3d} are all zero as well.

The maximum likelihood estimates in Table 2 downweight the outliers because of the low degrees of freedom on the Student t distribution. One can also try to model some of these outliers directly. The most important single factor in my data set appear to be monetary policy announcements.⁷ Gürkaynak, Sack, and Swanson (2005) identified 139 days between 1990 and 2005 on which the Federal Reserve made a significant monetary policy announcement. Let $z_{4d} = 1$ if a monetary policy announcement was reported by the these authors for day d and is zero otherwise. Note that this variable is allowed to influence the variance h_t but is not presumed to have any effect on the mean μ . The first three columns of Table 3 report the result when monetary policy announcements are included as a factor shifting the variance. The contribution to the variance is highly statistically significant, and increases the estimated degrees of freedom for the Student t distribution. But recognizing the influence of outliers of this form does not change the finding that the estimated means μ are still statistically insignificant.

The last three columns of Table 3 add to the description of the variance the variable z_{5d} , which is unity if an employment report was released on day d and zero otherwise. Again the

estimated degrees of freedom increase when one allows for this factor shifting the variance as well. If we added enough such determinants of the conditional variance, presumably we would eventually be left with a residual that is closer to a Normal distribution. However, the inference about the mean μ remains the same. Hence, whether outliers are taken into account through a Student t distribution with low degrees of freedom or as calendar factors that change the conditional variance, the conclusion is that, when outliers are allowed for in one way or the other, one does not find statistically significant evidence of a nonzero mean for f_{id} .

Some might argue that we shouldn't want to downweight outliers at all, since monetary policy announcements do not represent contaminated data but are instead the heart of what we are interested in investigating here. But once one recognizes that we are then talking about the mean not of a few thousand observations but rather of a much smaller sample of those days on which the markets received a major monetary policy announcement, then even if one were to maintain that this smaller sample did indeed have a negative population mean, this might still be consistent with perfectly rational, risk-neutral pricing of F_{dm} insofar as there was a learning process throughout the sample, in which market participants could not foresee perfectly in 1988 the average change in policy stance that turned out to characterize the sample.

Of course, much of the inference in the previous literature about bias was based not on the series f_{id} used here, but instead was based on monthly data, where the detailed corrections for announcement days, day-of-the-month effects, and observations that are exactly zero

are all unnecessary. The first column of Table 4 looks at the sample mean of monthly observations on u_{1m} , the difference between what the monthly fed funds rate actually turned out to be (r_m) and the value for that month's fed funds contract as of the last day of the preceding month. The sample mean of u_{1m} is of course based on the identical summary statistic as the first column of Table 2, namely $\sum_{d=1}^{T_D} f_{1d}$, except that whereas the first column of Table 2 estimate of the mean divides this sum by the number of days T_D , the first column of Table 4 divides by the number of months T_M . The standard errors for the first row of these respective columns involve different calculations, but the statistical significance with monthly data remains.

Letting now y_t in equation (5) correspond to the monthly series u_{1m} , I repeated the above EGARCH estimation using monthly data with $\mathbf{z}_m = (1, m/1000)'$ and $p \equiv 0$. Maximum likelihood estimates are reported in the fourth column of Table 4. Once again a model that allows for EGARCH with Student t innovations and a time trend for the variance is a vastly better description of the data, improving the log likelihood by 80. And once again with these corrections one arrives at a statistically insignificant, and positive rather than negative, estimate for the mean.

I found similar results for u_{2m} the monthly accumulations of f_{2d} , and u_{3m} , the monthly accumulations of f_{3d} :

$$u_{im} = F_{d^*(m),m+i-1} - F_{d^*(m-1),m+i-1}.$$

These results are also reported in Table 4, and confirm that the finding of a statistically significant mean for u_{im} is not reproduced in a model that allows for GARCH effects and

outliers.

5 Possible time-varying predictability of futures prices.

Consider next the serial correlation of f_{id} , as measured by OLS regression of f_{id} on a constant and five of its own lags. The estimated OLS coefficients and 95% confidence intervals are plotted in Figure 3. Coefficients on the first lag range from 0.14 to 0.16 and are highly statistically significant. All other lagged coefficients are less than 0.03 in absolute value, with only the coefficient relating f_{3d} to $f_{3,d-5}$ statistically significant at the 0.05 level ($p = 0.017$). Unlike the findings for the constant term, the estimated coefficient on $f_{i,d-1}$ and its statistical significance do not change much if one relies on maximum likelihood estimation of an EGARCH specification.⁸

Although this serial correlation is statistically significant, it is hard to claim that it has much economic significance. The predictability for a one-day-ahead forecast of f_{id} is quite limited, with the R^2 for all 3 regressions below 0.03. Moreover, the predictability two days ahead implied by these coefficients is essentially zero. Such very limited, very short-run serial correlation seems more likely to be attributed to measurement problems such as resolving bid-ask effects into settlement prices rather than to some fundamental predictability of the risk premium h_{dm} . In the remainder of this section, however, I include a single lag of f_{id} , partly to ensure correct calculation of standard errors.

Piazzesi and Swanson (2006) established using monthly data that e_{im} can be predicted using a number of macroeconomic and financial variables, particularly for longer horizons

i. Table 5 investigates several of the Piazzesi-Swanson indicators as possible predictors of the daily series f_{id} . I find no statistically significant contribution of the previous day's spread between Treasury yields of any maturity and the value of f_{id} . Noting that since the average forecast horizon of the farthest-forward daily series used here, f_{3d} , is 2.5 months, this finding is broadly consistent with that of Piazzesi and Swanson, who generally reported very little predictability of e_{im} for $i \leq 2$ months. The spread between Baa-corporate⁹ and 10-year-Treasury yields is only marginally statistically significant for f_{3d} , with a t-statistic of -1.958.

Piazzesi and Swanson also found that the 12-month change in nonfarm payrolls can be used to predict monthly e_{im} . Let n_m denote the seasonally unadjusted total quantity of non-farm payroll employment in month m . The next-to-last row of Table 5 replicates Piazzesi and Swanson's result with daily data, for which the regressor x_{d-1} used was $100 \log(n_{m^*(d-1)-1}/n_{m^*(d-1)-13})$. This makes a highly statistically significant contribution to predicting f_{2d} and f_{3d} , with faster employment growth over the preceding year signalling that the funds rate is likely to be higher than predicted by the futures market.

As did Piazzesi and Swanson, I also constructed a monthly data set for the annual growth rate of seasonally adjusted nonfarm employment $100 \log(\tilde{n}_{m-2}/\tilde{n}_{m-14})$ as it would actually have been reported and known to market participants as of the beginning of month m , using the real-time data archive described by Croushore and Stark (2001) and maintained by the Federal Reserve Bank of Philadelphia. Estimates are reported in the last row of Table 5, and are quite similar to the results using revised data. I also obtained essentially the identical

coefficient and standard error with EGARCH MLE (results not shown).

Again, although these coefficients are statistically significant, it would be a mistake to view them as of great economic importance. Even with both a lagged dependent variable and employment growth, the R^2 in these regressions is barely over 2%. Moreover, although nonfarm payrolls turn out to be a variable that helped to forecast the fed funds rate r_m over this sample period, that does not prove that the same variable will help predict it in the future. Viewing nonfarm payroll employment growth as a factor that turned out to be correlated with r_m within the observed sample but that markets overlooked, rather than as a factor that determines the risk premium h_{dm} , seems to me the most natural interpretation.

6 Conclusions.

While one can find some statistical evidence of predictability of price changes for near-term fed funds futures contracts, any daily fluctuations in the implicit risk premium h_{dm} account for at most a very small part of the variance of f_{id} for $i \leq 3$. Daily changes in the near-term fed funds futures contracts primarily reflect changes in market participants' assessments of where the federal funds rate is likely to be over the next few months.

Although these conclusions might appear to differ from those by Piazzesi and Swanson, I believe the results are broadly consistent. First, their strongest results came from contracts with horizons greater than or equal to 3 months; by contrast, the average duration of f_{3d} , the longest contract studied here, is 2-1/2 months. Second, Piazzesi and Swanson observe that

“risk premia seem to change primarily at business-cycle frequencies, which suggests that we may be able to “difference them out” by looking at one-day changes in near-dated federal funds futures on the day of a monetary policy announcement. Indeed, our results confirm that differencing improves these policy measures.”

The present study confirms that daily changes in near-term futures prices are indeed an excellent indicator of changes in market expectations of near-term Fed policy. Moreover, in recent years, these expectations have proven remarkably accurate. Daily changes in futures prices appear to offer us a useful tool for measuring the effects that anticipated near-term policy changes may have on the economy.

Footnotes

¹Piazzesi and Swanson (2006) studied changes on those days on which major monetary policy announcements occurred, and arrived at a similar conclusion to that uncovered here using all days.

²Data were purchased from the Chicago Board of Trade, with the settlement price at the end of day d used to calculate F_{dm} . Letting P_{dm} denote the quoted settlement price for day d , F_{dm} was constructed as $100 \times [100 - P_{dm}]$.

³For example, the interest rate implied by the May 2006 contract traded on Tuesday, May 30, 2006 was 4.945%. Because Monday was a holiday, the previous banking day was Friday, May 26, on which the implied May interest rate was 4.95%. Hence the value for f_{1d} for $d = \text{May 30, 2006}$ was $f_{1d} = -0.5$.

⁴The studies cited attributed this improvement in part to better communication by the Fed of its intentions. For example, Swanson noted that private sector forecasts of GDP and inflation do not exhibit the post-2000 reduction in mean squared error and dispersion across individual forecasters that is observed for forecasts of the fed funds rate.

⁵An earlier version of this paper treated the data as all discretely valued (due to number of significant digits reported) rather than a continuous-discrete mixture as here. The former requires specifying the bin width for which data are reported. Most observations are in multiples of 0.5 basis points, though there are 17 observations reported in multiples of 0.1 basis points, so an exact characterization of the nature of the discretization in the data would be rather involved. Results reported in the original paper are very similar to those

presented here.

⁶The numerical search was also kept robust with respect to wild numerical guesses (e.g., Hamilton, 1994, p. 146) by parameterizing $p = \lambda_p^2/(1 + \lambda_p^2)$, $\delta = \lambda_\delta^2/(1 + \lambda_\delta^2)$, and $\nu = 1 + \lambda_\nu^2$; none of the inequality constraints implicit in these parameterizations ended up being binding.

⁷Poole and Rasche (2000) examined days with large changes in the spot-month contract over 1989-1999 and concluded that economic news was more often the most important factor.

⁸That is, with μ in equation (5) replaced with $k + \phi y_{t-1}$. These MLE results are not reported separately.

⁹Piazzesi and Swanson use the BBB- rather than Baa-corporate yield, but I was unable to locate a daily series for the former. All daily interest rate data used in Table 5 were obtained from the FRED database of the Federal Reserve Bank of St. Louis.

References

- Carlson, John B., Ben Craig, Patrick Higgins, and William R. Melick. (2006) “FOMC Communication and the Predictability of Near-Term Policy Decisions.” *Federal Reserve Bank of Cleveland Economic Commentary*, June.
- Croushore, Dean, and Tom Stark. (2001) “A Real-Time Data Set for Macroeconomists.” *Journal of Econometrics*, 105, 2001, 111-130.
- Gürkaynak, Refet S. (2005) “Using Federal Funds Futures Contracts for Monetary Policy Analysis.” Working paper, Federal Reserve Board.
- Gürkaynak, Refet S., Brian Sack, and Eric Swanson. (2005) “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements.” *International Journal of Central Banking*, 1, 55-93.
- Gürkaynak, Refet S., Brian Sack, and Eric Swanson. (2007) “Market-Based Measures of Monetary Policy Expectations,” *Journal of Business and Economic Statistics*, 25, 201-212.
- Hamilton, James D. (1994) *Time Series Analysis*. Princeton: Princeton University Press.
- Hamilton, James D. (2006) “Daily Monetary Policy Shocks and the Delayed Response of New Home Sales,” working paper, UCSD.
- Kuttner, Kenneth N. (2001) “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market.” *Journal of Monetary Economics*, 47, 523-544.
- Krueger, Joel T., and Kenneth N. Kuttner. (1996) “The Fed Funds Futures Rate as a Predictor of Federal Reserve Policy.” *Journal of Futures Markets*, 16, 865-879.

Lang, Joe, Brian Sack and William Whitesell. (2003) "Anticipations of Monetary Policy in Financial Markets." *Journal of Money, Credit, and Banking*, 35(6, part 1), 889-909.

Nelson, Daniel B. (1991) "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica*, 59, 347-370.

Piazzesi, Monika, and Eric Swanson. (2006) "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy." Working paper, University of Chicago.

Poole, William, and Robert H. Rasche. (2000) "Perfecting the Market's Knowledge of Monetary Policy," *Journal of Financial Services Research*, 18, 255-298.

Poole, William, Robert H. Rasche, and Daniel L. Thornton. (2002) "Market Anticipations of Monetary Policy Actions," *Federal Reserve Bank of St. Louis Review*, 84(July/August), 65-94.

Rudebusch, Glenn D. (1998) "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review*, 39, 907-931.

Sack, Brian. (2004) "Extracting the Expected Path of Monetary Policy from Futures Rates." *Journal of Futures Markets*, 24, 733-754.

Swanson, Eric T. (2006) "Have Increases in Federal Reserve Transparency Improved Private Sector Interest Rate Forecasts?" *Journal of Money, Credit, and Banking*, 38, 791-819.

Table 1. Mean squared errors and mean absolute errors (in basis points) of forecasts from futures values compared with those for random walk.

Forecast horizon	Random walk MSE	Futures MSE	Percent MSE improvement	Futures MAE
Full data set (1988:12-2006:06)				
1 month ahead	389	128	67%	6.90
2 months ahead	1248	392	69%	12.76
3 months ahead	2522	914	64%	20.03
Recent data (2003:01-2006:06)				
1 month ahead	183	5	97%	1.50
2 months ahead	665	19	97%	3.18
3 months ahead	1484	48	97%	5.40

Table 2. Maximum likelihood estimates for i.i.d. Gaussian and EGARCH non-Gaussian descriptions of daily changes in fed funds futures (4454 observations, $d =$ Oct 5, 1988 to June 30, 2006, standard errors in parentheses).

	i.i.d. Normal			EGARCH		
	f_{1d}	f_{2d}	f_{3d}	f_{1d}	f_{2d}	f_{3d}
mean (μ)	-0.12 (0.03)	-0.15 (0.04)	-0.18 (0.05)	-0.03 (0.03)	0.00 (0.04)	0.02 (0.04)
log average variance (γ_1)	1.44 (0.02)	2.06 (0.02)	2.39 (0.03)	2.40 (0.14)	3.27 (0.17)	3.70 (0.19)
log h_{t-1} (δ)	----	----	----	0.83 (0.03)	0.91 (0.02)	0.94 (0.01)
$ u_{t-1} $ (α)	----	----	----	0.28 (0.03)	0.22 (0.03)	0.16 (0.02)
trend in variance (γ_2)	----	----	----	-0.53 (0.05)	-0.35 (0.06)	-0.38 (0.06)
first 2 days variance effect (γ_3)	----	----	----	0.64 (0.12)	0.56 (0.12)	0.37 (0.11)
probability of zero change (p)	----	----	----	0.462 (0.007)	0.426 (0.007)	0.317 (0.007)
Student t degrees of freedom (ν)	----	----	----	4.2 (0.3)	3.8 (0.3)	3.8 (0.2)
log likelihood	-9,528.35	-10,908.68	-11,638.62	-7,865.59	-9,160.47	-10,386.76

Table 3. Maximum likelihood estimates for EGARCH non-Gaussian descriptions of daily changes in fed funds futures with specific allowance for effect on variance of monetary policy announcements and employment releases (4454 observations, $d = \text{Oct 5, 1988 to June 30, 2006}$, standard errors in parentheses).

	Policy announcements			Policy and employment		
	f_{1d}	f_{2d}	f_{3d}	f_{1d}	f_{2d}	f_{3d}
mean (μ)	-0.04 (0.03)	0.01 (0.04)	-0.01 (0.04)	-0.04 (0.03)	0.01 (0.03)	0.03 (0.04)
log average variance (γ_1)	2.40 (0.13)	3.17 (0.17)	3.60 (0.18)	2.35 (0.13)	3.10 (0.17)	3.53 (0.17)
log h_{t-1} (δ)	0.82 (0.03)	0.89 (0.02)	0.94 (0.01)	0.83 (0.03)	0.90 (0.01)	0.94 (0.01)
$ u_{t-1} $ (α)	0.35 (0.04)	0.28 (0.03)	0.17 (0.02)	0.36 (0.04)	0.31 (0.03)	0.21 (0.02)
trend in variance (γ_2)	-0.54 (0.05)	-0.35 (0.06)	-0.37 (0.06)	-0.53 (0.05)	-0.33 (0.06)	-0.37 (0.06)
first 2 days variance effect (γ_3)	0.62 (0.11)	0.59 (0.11)	0.38 (0.11)	0.54 (0.11)	0.54 (0.11)	0.28 (0.10)
monetary policy announcement (γ_4)	2.01 (0.18)	2.04 (0.18)	1.57 (0.18)	1.94 (0.18)	1.97 (0.17)	1.51 (0.16)
employment data release (γ_5)	----	----	----	1.03 (0.14)	1.42 (0.13)	1.64 (0.13)
probability of zero change (p)	0.462 (0.007)	0.426 (0.007)	0.317 (0.007)	0.462 (0.007)	0.426 (0.007)	0.317 (0.007)
Student t degrees of freedom (ν)	5.6 (0.5)	4.7 (0.39)	4.3 (0.3)	6.5 (0.7)	6.0 (0.6)	5.8 (0.5)
log likelihood	-7,789.30	-9,077.39	-10,340.02	-7,760.16	-9,010.91	-10,241.41

Table 4. Maximum likelihood estimates for i.i.d. Gaussian and EGARCH non-Gaussian descriptions of monthly changes in fed funds futures (213 observations, m = Oct 1988 to June 2006, standard errors in parentheses).

	i.i.d. Normal			EGARCH		
	u_{1m}	u_{2m}	u_{3m}	u_{1m}	u_{2m}	u_{3m}
mean (μ)	-2.66 (0.75)	-3.17 (1.06)	-3.74 (1.27)	0.12 (0.24)	0.43 (0.34)	0.27 (0.67)
log average variance (γ_1)	4.79 (0.10)	5.47 (0.10)	5.83 (0.10)	5.73 (0.42)	6.47 (0.51)	7.01 (0.54)
log h_{t-1} (δ)	----	----	----	0.63 (0.16)	0.74 (0.22)	0.84 (0.11)
$ u_{t-1} $ (α)	----	----	----	0.18 (0.07)	0.15 (0.07)	0.30 (0.12)
trend in variance (γ_2)	----	----	----	-22.7 (3.1)	-23.6 (3.3)	-17.1 (3.8)
Student t degrees of freedom (ν)	----	----	----	2.1 (0.4)	2.2 (0.4)	4.1 (1.2)
log likelihood	-812.61	-884.70	-922.80	-731.08	-793.38	-860.16

Table 5. OLS coefficients on x_{d-1} in regression of f_{id} on constant, its own lag, and lagged value of indicated explanatory variable (standard errors in parentheses).

Explanatory variable	Dependent variable		
	f_{1d}	f_{2d}	f_{3d}
10-year minus 5-year Treasury spread	0.058 (0.086)	-0.036 (0.117)	-0.070 (0.138)
5-year minus 2-year Treasury spread	-0.009 (0.058)	-0.085 (0.079)	-0.126 (0.093)
2-year minus 1-year Treasury spread	-0.072 (0.112)	-0.136 (0.153)	-0.172 (0.181)
1-year minus 6-month Treasury spread	0.006 (0.173)	0.302 (0.236)	0.439 (0.279)
Baa minus 10-year Treasury spread	-0.035 (0.058)	-0.126 (0.079)	-0.184* (0.094)
12-month job growth as currently reported for period ending previous month	0.017 (0.023)	0.089** (0.031)	0.125** (0.036)
12-month job growth as reported at the time for most recent period that would have been known by end of previous month	0.016 (0.024)	0.093** (0.033)	0.121** (0.039)

* denotes statistically significant at 5% level, ** at 1% level.

Figure 1. Sample histogram (rectangles) of f_{1d} with bin-width of 0.5 basis points and Normal distribution (continuous curve). Height of rectangle is fraction of observations falling in that 0.5-basis-point interval, while height of curve is 0.5 times the $N(0,4)$ density.




Figure 2. Plot of f_{1d} .




Figure 3. OLS coefficients and 95% confidence intervals from regressions of f_{id} on a constant and five of its own lagged values.

