Journal of Monetary Economics 55 (2008) 677–691

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Futures prices as risk-adjusted forecasts of monetary policy ☆

Monika Piazzesi [a,*], Eric T. Swanson [b]

[a] University of Chicago, NBER, Chicago, USA
[b] Federal Reserve Bank of San Francisco, USA

ARTICLE INFO

Article history:
Received 7 August 2006
Received in revised form
4 April 2008
Accepted 4 April 2008
Available online 2 May 2008

Keywords:
Federal funds futures
Monetary policy
Risk premia

ABSTRACT

Many researchers have used federal funds futures rates as measures of financial markets' expectations of future monetary policy. However, to the extent that federal funds futures reflect risk premia, these measures require some adjustment. In this paper, we document that excess returns on federal funds futures have been positive on average and strongly countercyclical. In particular, excess returns are surprisingly well predicted by macroeconomic indicators such as employment growth and financial business-cycle indicators such as Treasury yield spreads and corporate bond spreads. Excess returns on eurodollar futures display similar patterns. We document that simply ignoring these risk premia significantly biases forecasts of the future path of monetary policy. We also show that risk premia matter for some futures-based measures of monetary policy shocks used in the literature.

1. Introduction

Predicting the future course of monetary policy is of tremendous importance to financial market participants. The current state of the art in this area is to use futures contracts on the short-term interest rate that is targeted by the central bank and to interpret the futures rate on, say, the December federal funds futures contract as the market expectation of what the federal funds rate will be in December. This procedure is widely used in the financial press (e.g., The Wall Street Journal, 2005; Financial Times, 2005), by Fed watchers (e.g., Altig, 2005; Hamilton, 2006), by central banks (e.g., European Central Bank Monthly Bulletin, 2005, p. 24; Federal Reserve Monetary Policy Report to Congress, 2005, p. 22), and in the academic literature (e.g., Krueger and Kuttner, 1996; Rudebusch, 1998, 2002; Bernanke and Kuttner, 2005).[1]

The standard practice is appealing for many reasons. First, producing the forecasts is simple—the rates on various contracts can be obtained directly from futures exchanges at any time during the day. Second, the forecasts work well—federal funds futures outperform forecasts based on alternative methods, such as sophisticated time series specifications, monetary policy rules, and forecasts derived from Treasury bills or other financial market instruments

☆ We are particularly indebted to Charlie Evans and Frank Schorfheide (our discussants at the Stanford/FRBSF conference), and Donald Kohn for helpful suggestions. We thank Claire Hausman for excellent research assistance, and John Cochrane, Lou Crandall, Darrell Duffie, Bob Hall, David Laibson, Andy Levin, Brian Sack, Martin Schneider, Jiang Wang, Jonathan Wright, and seminar participants at the Bundesbank, European Central Bank, Federal Reserve Board, Universities at Columbia, Heidelberg, Mannheim, Stanford, Rochester, Harvard, and the Wharton School, the IMA Workshop at the University of Minnesota, the NBER Summer Institute, and the Society of Economic Dynamics Meetings for comments. The views expressed in this paper, and any errors and omissions, are those of the authors, and are not necessarily those of the individuals listed above, the management of the Federal Reserve Bank of San Francisco, or any other individual within the Federal Reserve System.

* Corresponding author.
E-mail address: piazzesi@uchicago.edu (M. Piazzesi).

[1] Some of these studies allow for constant risk premia.

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2008.04.003

(e.g., Evans, 1998; Gürkaynak et al., 2007). Third, previous studies did not find any large time variation in risk premia in fed funds futures (e.g., Krueger and Kuttner, 1996; Sack, 2004; Durham, 2003).[2]

However, there is by now a large and well-accepted body of evidence in the finance literature against the expectations hypothesis for Treasury yields (e.g., Fama and Bliss, 1987; Stambaugh, 1988; Campbell and Shiller, 1991; Cochrane and Piazzesi, 2005). Over a very wide range of sample periods and bond maturities, excess returns on Treasury securities have been positive on average, time-varying, and significantly predictable. Time-varying risk premia in these markets may well carry over to related markets and therefore lead to systematic deviations of fed funds futures rates from expectations of the subsequently realized fed funds rate.

In this paper, we show that the expectations hypothesis also fails for federal funds futures. In particular, excess returns on fed funds futures contracts at even short horizons have been positive on average and significantly predictable. The R^2's depend on the forecast horizon and range from 10% at a two-month horizon up to 39% at a six-month horizon. We find that macroeconomic business-cycle indicators such as employment growth capture this predictability surprisingly well. We also find that financial business-cycle indicators such as corporate bond spreads and Treasury yield spreads do well at predicting excess returns. These findings stand up to a battery of robustness checks, including bootstrapped test statistics, real-time data, subsample stability pre- and post-1994, rolling-endpoint regressions, out-of-sample forecasts, and a comparison to excess returns on eurodollar futures, for which we have a somewhat longer sample.

We exploit the significant predictability of excess returns on futures to propose a risk adjustment to forecasts of monetary policy. We find that not implementing our risk adjustment can produce very misleading results. Specifically, forecasts based on the expectations hypothesis make large mean errors and large mean-squared errors. Moreover, errors from unadjusted forecasts vary systematically over the business cycle; futures rates tend to overpredict in recessions and underpredict in booms. Non-risk-adjusted forecasts also tend to perform very poorly around economic turning points, adapting too slowly to changes in the direction of monetary policy. For example, right before recessions, when the Fed has already started easing, fed funds futures keep forecasting high funds rates. As a consequence, forecast errors using unadjusted futures rates are more highly autocorrelated than are forecast errors using our risk-adjusted futures rates.

Our findings also suggest that monetary policy shocks may not be accurately measured by the difference between the fed funds rate target and an ex ante market expectation based on fed funds futures. Indeed, we document that the amount by which we need to adjust these shocks can be substantial, at least relative to the size of the shocks themselves. However, 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 fed funds futures on the day of a monetary policy announcement. Indeed, our results confirm that differencing improves these policy measures.

Our findings for federal funds futures complement those in the traditional finance literature on Treasuries in several ways. First, we find that the most important predictive variables for excess returns are macroeconomic variables, such as employment growth. This finding allows us to link the predictability in excess returns directly to the business cycle, while the existing literature on Treasuries has focused mainly on predictability using financial variables such as term spreads (e.g., Cochrane and Piazzesi, 2005).

Second, fed funds futures are actually traded securities, while the zero-coupon yield data used in Fama and Bliss (1987) and many other papers are data constructed by interpolation schemes. While the predictability patterns in this artificial data may not lead to profitable trading rules based on actual securities, investors can implement our results directly by trading in fed funds futures. Interestingly, we document evidence that suggests that futures market participants were aware of these excess returns in real time: traders that are classified as "not hedging" by the U.S. Commodity Futures Trading Commission (CFTC) went long in these contracts precisely when we estimate that expected excess returns on these contracts were high, and they went short precisely in times when we estimate expected excess returns were very low.

Finally, fed funds futures contracts have maturities of just a few months and may therefore be less risky than Treasury notes and bonds, which have durations of several years; moreover, the holding periods relevant for measuring excess returns on fed funds futures are less than one year, while the results for Treasuries typically assume that the investor holds the securities for an entire year (an exception is Stambaugh, 1988, who studies Treasury bills). Given the short maturities and required holding periods to realize excess returns in the fed funds futures market, one might think that risk premia in this market would be very small or nonexistent. We find that this is not the case.

Throughout this paper, we will often use the label "risk premia" to refer to "predictable returns in excess of the risk-free rate." This use of language should *not* be interpreted as taking a particular stance on the structural interpretation of our results. The existing literature has proposed several appealing explanations for why excess returns on these contracts might be predictable. Some of these explanations are based on preferences: for example, investors may exhibit risk aversion which varies over the business cycle, or care about the slow-moving, cyclical consumption of items like housing. Other explanations are based on beliefs that deviate from rational expectations, for example because of learning or for psychological reasons. It is not easy to make the case for just one of these explanations: beliefs and other preference

[2] These studies run regressions of one-month excess returns on fed funds futures on a variety of variables, including macroeconomic variables. While some of the regression coefficients are statistically significant, they are economically small. Our results are different: we show that for holding periods longer than one month, risk premia are large on average and vary over time substantially.

parameters can often not be identified separately. We therefore set aside these issues as beyond the scope of the present paper.

The remainder of the paper proceeds as follows. Section 2 shows that measures of excess returns on fed funds futures are identical to monetary-policy forecast errors and can be predicted using business-cycle indicators such as employment growth or corporate bond spreads. Section 3 performs a battery of robustness checks. Section 4 presents our risk-adjustment to policy forecasts and shows that failing to implement the adjustment can lead to substantial mistakes, so that the predictability of excess returns is economically as well as statistically significant. Section 5 investigates the implications of our results for measures of monetary policy "shocks" used in the literature. Section 6 concludes.

2. Excess returns on federal funds futures

Federal funds futures contracts have traded on the Chicago Board of Trade exchange since October 1988 and settle based on the average federal funds rate that prevails over a given calendar month.[3] Let $f_t^{(n)}$ denote the fed funds futures contract rate for month $t+n$ as quoted at the end of month t. We will refer to $n=1$ as the one-month-ahead futures contract, $n=2$ as the two-month-ahead contract, and so on. Let r_{t+n} denote the ex post realized value of the federal funds rate for month $t+n$, calculated as the average of the daily fed funds rates in month $t+n$ for comparability to the fed funds futures contracts.

The buyer of a fed funds futures contract locks in the contracted rate $f_t^{(n)}$ for the contract month $t+n$ on a \$5 million deposit (the \$5 million deposit is never actually made by the buyer; this is only the number that is used to compute the payoff of the contract at maturity). The contracts are cash-settled the day after expiration, with expiration occurring at the end of the contract month. At that time, the buyer receives \$5 million times the difference between $f_t^{(n)}$ and the realized fed funds rate r_{t+n} converted to a monthly rate.[4] As is standard for futures contracts, there are no up-front costs to either party of entering into the contract; both parties simply commit to the contract rate at time t and receive their payoffs at time $t+n$.

For the buyer of the futures contract, the amount $(f_t^{(n)} - r_{t+n}) \times \$5$ million represents the payoff of a zero-cost portfolio. Using common terminology (e.g., Cochrane, 2001, p. 11), we refer to this difference, $f_t^{(n)} - r_{t+n}$, as an "excess return." We denote the excess return for the buyer of the contract by

$$rx_{t+n}^{(n)} = f_t^{(n)} - r_{t+n}. \tag{1}$$

Since we will consider futures contracts with maturities n ranging from one to six months, the excess returns in (1) will correspond to different holding periods for different values of n. To make excess returns on these different contracts more directly comparable, we also report statistics for annualized excess returns, which are computed by multiplying the excess returns in (1) by $12/n$. Also, we measure returns in basis points. These conventions will apply throughout the paper.

Eq. (1) treats fed funds futures contracts as forward contracts, and thereby abstracts from the fact that futures contracts are "marked to market" every day, based on collateral posted in interest-bearing accounts. The appendix of the working paper version of this article (Piazzesi and Swanson, 2006) explains this procedure in more detail, computes excess returns on fed funds futures contracts taking into account the effects of marking to market, and shows that the difference between a precise definition of excess returns on futures contracts and the simplification in Eq. (1) is very small and does not matter for our results below. We therefore use (1) as the definition of excess returns in this paper for simplicity.

The simplification in Eq. (1) also has the advantage that excess returns are easily linked to forecasting. Under the expectations hypothesis, futures are expected future short rates: $f_t^{(n)} = \mathrm{E}_t(r_{t+n})$. Thus, Eq. (1) not only represents excess returns, but also minus the forecast error under the expectations hypothesis. This coincidence makes it easy to see how we can adjust futures-based forecasts for risk premia.

2.1. Average excess returns

To investigate whether average excess returns on federal funds futures contracts are zero, we run the regression:

$$rx_{t+n}^{(n)} = \alpha^{(n)} + \varepsilon_{t+n}^{(n)} \tag{2}$$

for different contract horizons n.

Panel A in Table 1 presents results from regression (2) for the forecast horizons $n = 1, \ldots, 6$ months over the entire period for which we have fed funds futures data, October 1988 through December 2005. This period will be the baseline for all of our regressions below. We run the regression at monthly frequency, sampling the futures data on the last day of each month t.[5] We compute standard errors using the heteroskedasticity- and autocorrelation-consistent procedure from

[3] The average federal funds rate is calculated as the simple mean of the daily averages published by the Federal Reserve Bank of New York, and the federal funds rate on a non-business day is defined to be the rate that prevailed on the preceding business day.

[4] This means that $f_t^n - r_{t+n}$ gets multiplied by $\frac{1}{12}$ according to the quoting convention in the fed funds futures market, which uses a 30-day month and 360-day year. See the CBOT web site for additional details.

[5] We restrict attention to monthly data in order to avoid variations in the maturity of the contracts that would arise over the course of each month: for example, with daily data, the one-month-ahead contract could have as few as 28 and as many as 61 days until maturity, which is a significant variation

Table 1
Unconditional and conditional expected excess returns

n	1	2	3	4	5	6
Panel A: average excess returns						
$\alpha^{(n)}$, unannualized	2.9	6.3	10.5	16.1	23.2	30.7
$\alpha^{(n)}$, annualized	35.0	38.1	42.2	48.3	55.6	61.4
(t-Stat)	(3.7)	(3.4)	(3.0)	(2.9)	(2.9)	(2.7)
[Bootstrap p-value]	[0.001]	[0.001]	[0.007]	[0.016]	[0.022]	[0.037]
Panel B: recession dummy						
Constant	28.2	26.6	28.8	33.3	40.0	49.4
(t-Stat)	(3.3)	(2.8)	(2.3)	(2.3)	(2.3)	(2.4)
[Bootstrap p-value]	[0.000]	[0.005]	[0.029]	[0.045]	[0.044]	[0.037]
Recession dummy	77.8	131.3	151.2	169.3	177.0	168.6
(t-Stat)	(1.4)	(2.1)	(3.1)	(5.1)	(6.6)	(4.7)
[Bootstrap p-value]	[0.326]	[0.047]	[0.042]	[0.030]	[0.018]	[0.041]
R^2	0.03	0.10	0.14	0.16	0.17	0.13
Panel C: nonfarm payrolls						
Constant	−0.3	1.7	−6.6	−16.8	−31.3	−45.7
$f_t^{(n)}$	0.13	0.19	0.25	0.32	0.38	0.44
(t-Stat)	(2.1)	(2.4)	(3.1)	(4.5)	(7.0)	(10.9)
[Bootstrap p-value]	[0.078]	[0.017]	[0.003]	[0.001]	[0.000]	[0.000]
ΔNFP_{t-1}	−0.20	−0.37	−0.51	−0.62	−0.70	−0.73
(t-Stat)	(−1.8)	(−3.0)	(−3.8)	(−5.1)	(−7.2)	(−11.1)
[Bootstrap p-value]	[0.082]	[0.006]	[0.003]	[0.000]	[0.000]	[0.000]
R^2	0.03	0.10	0.18	0.25	0.32	0.39
Panel D: corporate bond spread						
Constant	−43.7	−67.6	−89.1	−105.4	−125.3	−142.7
$f_t^{(n)}$	0.06	0.06	0.08	0.10	0.14	0.18
(t-Stat)	(1.3)	(1.0)	(1.4)	(2.1)	(2.6)	(3.0)
[Bootstrap p-value]	[0.260]	[0.266]	[0.133]	[0.072]	[0.009]	[0.013]
BBB spread	0.34	0.52	0.62	0.69	0.75	0.80
(t-Stat)	(1.8)	(2.6)	(2.9)	(2.9)	(3.0)	(2.8)
[Bootstrap p-value]	[0.093]	[0.031]	[0.026]	[0.027]	[0.003]	[0.036]
R^2	0.03	0.08	0.12	0.16	0.20	0.22

Note: n denotes the horizon of the fed funds futures contract in months. The sample for each regression is 1988:10–2005:12 at monthly frequency, sampled on the last day of each month. Excess returns are measured in annualized basis points, except for Panel A, which reports both unannualized and annualized results. HAC t-statistics are reported in parentheses, and bootstrapped p-values for the t-statistics are in brackets; t-statistics for constants in Panels C–D not reported. See text for details.

Hodrick (1992), which generalizes the Hansen and Hodrick (1988) procedure for overlapping contracts to the case of heteroskedasticity, allowing for $n-1$ lags of excess returns to be serially correlated due to contract overlap. Throughout this paper, we report HAC t-statistics based on these standard errors. Finally, to help account for the small-sample distributional properties of these t-statistics, Table 1 reports the bootstrapped p-value for each t-statistic using 50,000 bootstrap draws following Horowitz (2004).[6]

As can be seen in Panel A, average excess returns on fed funds futures have been significantly positive over our sample, ranging from about 3 to 5 basis points per month (35–61 bp per year). For example, buying the six-month-ahead futures contract and holding it to maturity is a strategy that generated a return of 61.4 bp per year on average.[7] Longer-horizon contracts have had greater excess returns on average even on a per-month or per-year basis.

(*footnote continued*)
in the holding period required to realize the excess return on the contract. In turn, these differences in maturities and holding periods influence the size and time variation of risk premia, as we will show below. Also, these variations would translate into different forecasting horizons when we later use our results to forecast the fed funds rate. Nonetheless, our results are all similar when we sample the data at daily rather than monthly frequency.

[6] Following Horowitz (2004), observations are resampled from the data with replacement to generate 50,000 synthetic samples of the same size as the original. To account for possible serial dependence of the data generating process, we resample the data in blocks, with block sizes of 4, 5, 6, 8, 10, and 12 observations for the cases $n = 1, 2, 3, 4, 5, 6$, respectively. The bootstrap p-value is computed as the percentage of synthetically generated t-statistics that exceed the actual t-statistic in absolute value.

[7] The unannualized excess return on the six-month-ahead contract is, on average, \$5 million times 0.307% times $\frac{1}{12}$. The annualized excess return is just double this amount (multiplying by $\frac{12}{6}$). Bid–ask spreads on fed funds futures contracts are typically small, usually around 1 bp even for long-maturity contracts. On days of very high volatility, spreads for long-maturity contracts ($n = 5, 6$ months) can reach 20 bp, but these spikes are very short-lived and can easily be avoided by trading just one or two days later. All of our findings of excess returns are thus robust to the presence of bid–ask spreads on these contracts.

The averages for the post-1994 period are a little lower but still significantly positive at 24.9, 27.2, 29.0, 32.7, 37.0, and 42.9 bp per year.

2.2. Excess returns and recessions

Under the expectations hypothesis, excess returns on bonds and interest-rate futures are assumed to be constant over time. However, there is by now a large body of literature that finds excess returns on Treasury securities to be significantly time-varying and predictable (e.g., Fama and Bliss, 1987; Campbell and Shiller, 1991; Cochrane and Piazzesi, 2005). Fig. 1 plots the realized excess return on the four-month-ahead federal funds futures contract, $rx^{(4)}_{t+4}$, from October 1988 through December 2005, with each point in the graph depicting the realized excess return, $f^{(4)}_t - r_{t+4}$, at date t. Certainly, the time variation in this realized excess return series has been large, ranging from -315 to 413 bp at an annualized rate. The graph also suggests that there have been several periods during which fed funds futures generated particularly large excess returns: the years 1991–1992, early 1995, the fall of 1998, and the years 2001–2002 (these are also the periods during which the Federal Reserve lowered interest rates). Two of these periods, 1991–1992 and 2001–2002, coincided with the two recessions in our sample. The other two periods were not recessions, but were also periods with slower economic growth.

Panel B in Table 1 reports the results from regressing the excess returns on fed funds futures on a constant and a recession dummy D_t:

$$rx^{(n)}_{t+n} = \alpha^{(n)} + \beta^{(n)} D_t + \varepsilon^{(n)}_{t+n}. \tag{3}$$

The recession dummy is significant for all contracts with maturities longer than just one month. The estimated coefficient on the recession dummy suggests that expected excess returns are countercyclical; expected excess returns are about 3–5 times higher in recessions than they are on average during other periods. (Note that annualizing the excess returns is a normalization that does not affect the t-statistics or R^2 in any of our regressions.) The fitted values from regression (3) for the four-month-ahead contract are depicted by the gray line in Fig. 1.

Of course, recession dummies are not useful as predictive variables in real time, since the NBER's business-cycle dating committee declares recession peaks and troughs as long as two years after they have actually occurred. Fig. 1 suggests, however, that any macroeconomic or financial business-cycle indicator may be a good candidate for forecasting excess returns on fed funds futures contracts, an exercise to which we now turn.

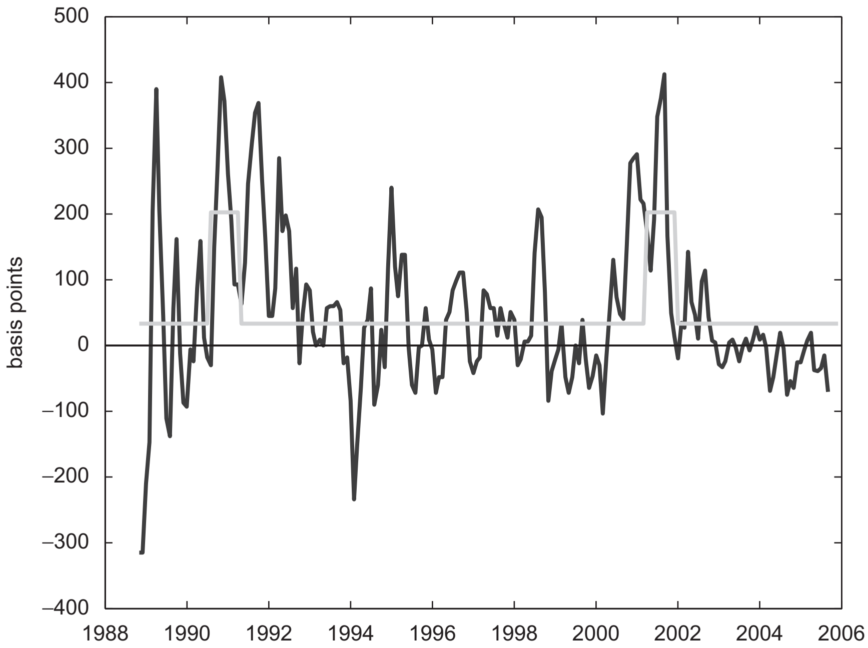

Fig. 1. Annualized excess returns on the federal funds futures contract four months ahead. The gray step function represents the fitted values from a regression of $rx^{(4)}_{t+4}$ on a constant and a recession dummy.

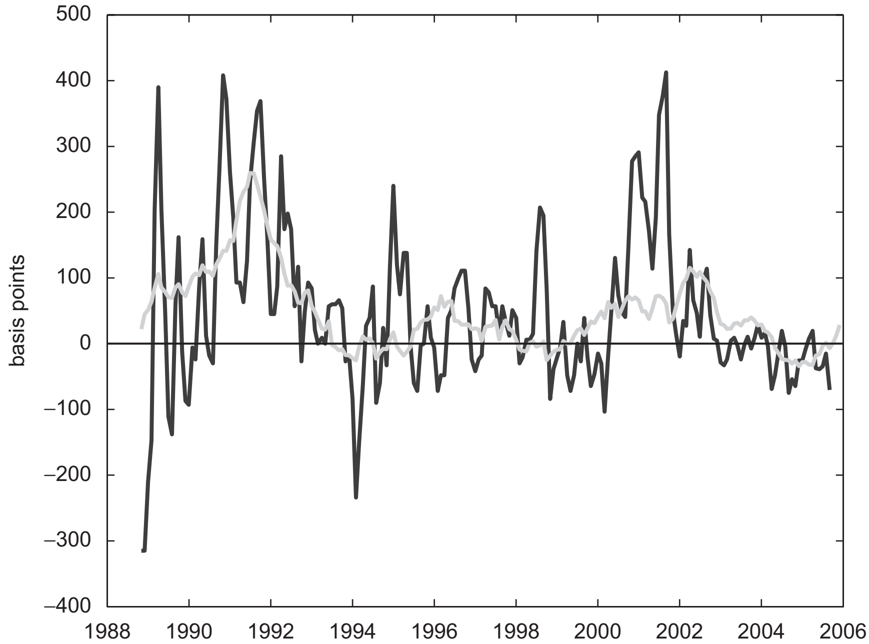

Fig. 2. Annualized excess returns on the federal funds futures contract four months ahead. The gray function represents the fitted values from a regression of $rx_{t+4}^{(4)}$ on a constant, employment growth and $f_t^{(4)}$ itself.

2.3. *Employment growth as a predictor of excess returns*

To investigate whether a variable or set of variables forecasts excess returns on federal funds futures, we run predictive regressions of the form:

$$rx_{t+n}^{(n)} = \alpha^{(n)} + \beta^{(n)} X_t + \varepsilon_{t+n}^{(n)}, \tag{4}$$

where X_t is a vector of variables known to financial markets in month t. Since GDP is only available at quarterly frequency, it is not a useful variable for forecasting monthly excess returns. We therefore turn to a closely related indicator of real activity: employment. More precisely, we use the year-on-year change in the logarithm of U.S. nonfarm payroll employment, being careful to use data that were available to financial market participants in real time (that is, as of the last day of month t).[8]

Panel C of Table 1 reports the forecasting results from regression equation (4) based on these real-time nonfarm payroll numbers. The regression also includes the futures rate itself on the right-hand side, as is common practice. The results show that employment growth is a significant predictor of excess returns for contracts with two months or more to maturity. Fitted values from regression (4) for the four-month-ahead contract are depicted by the gray line in Fig. 2. As we would expect from our results using the recession dummy, expected excess returns are countercyclical in Fig. 2: expected excess returns and employment growth are inversely related. The estimated slope coefficients in Panel C increase with the maturity of the contract and lie between -0.20 and -0.73 for annualized returns.

To understand the magnitude of these coefficients, note that employment growth is measured in basis points, which means that a 1 percentage point drop in employment growth increases expected excess returns on federal funds futures by about 20–73 bp per year. Over our sample, the mean and standard deviation of employment growth were 135 and 132 bp, respectively, which means that a one-standard deviation shock to employment makes us expect around 95 bp more in annualized excess returns on the six-month-ahead futures contract. The own futures contract rate $f_t^{(n)}$ is also a significant predictor of excess returns for these contracts, and the positive coefficient implies that, all else equal, excess returns are lower when the level of interest rates is lower.[9]

[8] Two data issues arise if we wish to run the predictive regressions (4) with data that were available to financial market participants in real time. First, nonfarm payroll numbers for a given month are not released by the Bureau of Labor Statistics until the first Friday of the following month. Thus, to perform the predictive regressions (4) with data that were available at the end of month t, we must lag the employment numbers by an entire month. Second, nonfarm payroll numbers are revised twice after their initial release and undergo an annual benchmark revision every June, so the final vintage numbers are not available for forecasting in real time. We therefore collected the real-time nonfarm payroll numbers, and use the first release of nonfarm payrolls for month $t-1$ and the revised value for nonfarm payrolls for month $t-13$ to compute the year-on-year change. Even the revised value for month $t-13$ is not quite equal to the final vintage of data for that month, because the BLS performs occasional benchmark revisions.

[9] As a robustness check, we re-ran the predictability regressions in Table 1 using the current-month fed funds rate r_t instead of the own futures contract rate $f_t^{(n)}$. The results (not reported) are similar to those based on the futures rate. For example, the R^2 are identical to those in Panel C. We also ran

The R^2 in Panel C suggest that we can predict up to 39% of the variation in excess returns on federal funds futures with employment growth and the futures rate itself. This result is remarkable, since these R^2 are comparable in size to those reported in Cochrane and Piazzesi (2005), who study excess returns on Treasuries over much longer holding periods (one year, as compared to just one to six months for our fed funds futures regressions above).

2.4. Financial market variables as predictors of excess returns

We can also try to forecast excess returns on fed funds futures using financial indicators of the business cycle, which have the advantage of being available at higher frequencies than macroeconomic indicators. Panel D in Table 1 reports the results for one such indicator, the spread between BBB-rated 10-year corporate bonds and the 10-year Treasury yield. (In the NBER working version of this paper, we report similar results for Treasury yield spreads, another widely used financial indicator.) The results lend additional support to the hypothesis that business-cycle indicators are useful predictors of excess returns in the fed funds futures market. The estimated coefficients on the corporate bond spread in these regressions are significant for fed funds futures contracts with two months or more to maturity, with R^2 that range from 8% to 22%. Although we do not graph the corporate bond spread here in the interest of space, it is most successful at capturing the runups in excess returns in 1990–1993 and 2001–2002, further suggesting that the predictive power of these financial indicators may be largely due to the relationship between excess returns and the business cycle.

By contrast, a time-varying measure of the risk premium proposed by Sack (2004), which is based on the slope of the eurodollar futures curve from four to five years ahead and is *not* closely tied to the business cycle, performs poorly at forecasting excess returns in the futures markets. Sack's long-horizon eurodollar slope is never statistically significant for any of our fed funds future contracts (t-statistics never exceed 0.7 in absolute value), with R^2 below 5%.[10] Moreover, the coefficient estimates on this slope have the "wrong" sign: Sack's analysis suggests that they should be positive—a high slope represents a higher risk premium and higher excess return—but in fact our coefficient estimates find that this relationship is negative.

3. Robustness and real-time predictability of excess returns

Our results above provide substantial evidence for time variation of risk premia in federal funds futures. We now perform a variety of robustness checks and sensitivity analysis of this basic result.[11] We devote particular attention to assessing whether these excess returns were predictable in real time, looking at rolling-endpoint regressions and also some intriguing historical evidence on the actual positions taken by non-hedging traders in the federal funds futures and eurodollar futures markets.

3.1. One-month holding period excess returns

Our sample period spans only about 17 years, which results in as few as 34 independent windows for our longest-horizon (six-month-ahead) fed funds futures contracts. A way to increase the number of independent observations in our regressions and check the robustness of our results is to consider the excess returns an investor would realize from holding an n-month-ahead fed funds futures contract for just one month—by purchasing the contract and then selling it back as an $(n-1)$-month-ahead contract in one month's time—rather than holding the contract all the way through to maturity. By considering one-month holding period returns on fed funds futures, we reduce potential problems of serial correlation and sample size for the longer-horizon contracts, and give ourselves 206 independent windows of data (under the null hypothesis of no predictability of excess returns) for all contracts.

We thus consider regressions of the form

$$f_t^{(n)} - f_{t+1}^{(n-1)} = \alpha^{(n)} + \beta^{(n)} X_t + \varepsilon_{t+1}^{(n)}, \tag{5}$$

where $f_t^{(n)}$ denotes the n-month-ahead contract rate on the last day of month t, $f_{t+1}^{(n-1)}$ denotes the $(n-1)$-month-ahead contract rate on the last day of month $t+1$, and the difference between these two rates is the ex post realized one-month holding period return on the n-month-ahead contract. Using specification (5), the residuals are serially uncorrelated under

(*footnote continued*)
predictability regressions including both r_t and $f_t^{(n)}$. The results (also not reported) did not significantly improve upon those in Panel C, except for the one-month-ahead contract. For $n=1$, we obtain significant slope coefficients for all RHS variables with an R^2 of 5%. The estimated slope coefficients on $f_t^{(n)}$ and r_t almost offset each other, the former being somewhat bigger. We also ran regressions with the most recent revised vintage of the employment data, and our results are similar: in particular, we find that employment growth predicted excess returns with R^2 values of 1%, 7%, 14%, 20%, 28%, and 37%. The results are also similar if we use contemporaneous rather than lagged nonfarm payrolls growth as a regressor.

[10] The estimated coefficients on Sack's long-horizon eurodollar slope are 0.94, −0.51, −3.42, −11.3, −17.5, and −19.3 at the $n = 1, 2, 3, 4, 5, 6$ horizons, with t-statistics of 0.1, −0.0, −0.2, −0.5, −0.7, and −0.7, respectively.

[11] In Piazzesi and Swanson (2006), we also extended the results of Section 2 to eurodollar futures, which have traded on the Chicago Mercantile Exchange since 1981 and have quarterly expirations which settle based on the spot eurodollar (LIBOR) rate in effect at expiration. Our findings for eurodollar futures contracts were very much in line with those for fed funds futures above.

Table 2
One-month holding period excess returns and nonfarm payrolls

n	1	2	3	4	5	6
Constant	0.3	6.3	−5.8	−5.4	−14.2	−69.1
(t-Stat)	(0.0)	(0.2)	(−0.2)	(−0.1)	(−0.3)	(−1.2)
$f_t^{(n)}$	0.13	0.23	0.34	0.41	0.47	0.69
(t-Stat)	(2.4)	(3.1)	(3.7)	(3.8)	(3.8)	(4.7)
ΔNFP_{t-1}	−0.19	−0.55	−0.76	−0.92	−1.02	−1.15
(t-Stat)	(−2.0)	(−4.1)	(−4.8)	(−5.0)	(−4.9)	(−5.0)
R^2	0.03	0.08	0.10	0.11	0.11	0.14

Note: n denotes the horizon of the fed funds futures contract in months. The sample for each regression is 1988:10–2005:12 at monthly frequency, sampled on the last day of each month. Excess returns are measured in annualized basis points (i.e., one-month excess returns are multiplied by 12). HAC t-statistics are reported in parentheses. The regression equation is (5), where X_t contains $f_t^{(n)}$ and nonfarm payroll employment growth ΔNFP_{t-1} from $t-13$ to $t-1$. ΔNFP_{t-1} is measured in basis points.

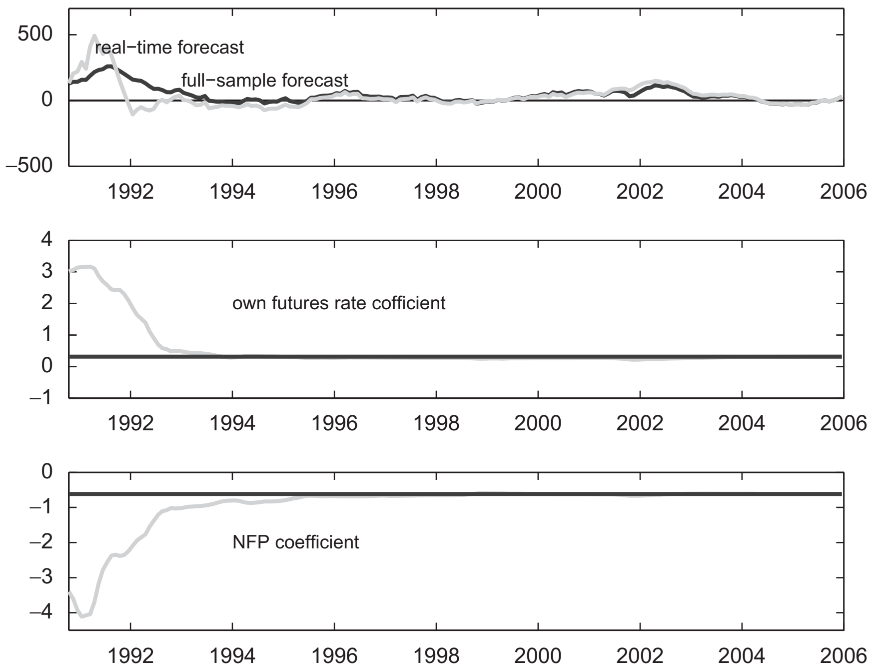

Fig. 3. The top panel shows real-time and full-sample forecasts of $rx_{t+4}^{(4)}$. The middle panel shows the gray rolling estimates of the coefficient on the own futures rate $f_t^{(4)}$. The flat black line is the full-sample coefficient from Table 1. The lower panel shows the gray rolling estimates of the coefficient on employment growth. Again, the flat black line is the full-sample coefficient from Table 1.

the null hypothesis of no predictability of excess returns, because all variables in Eq. (5) are in financial markets' information set by the start of the next period.

Table 2 presents the results of our previous analysis applied to this alternative specification, where the regressors are the own contract rate and employment growth. Although the R^2 values are uniformly lower, as is to be expected from quasi-first-differencing the left-hand side variable, our previous results are robust to this alternative specification. Results for term spreads and corporate bond spreads are similarly robust across specifications. (These results are not presented to conserve space.)

3.2. Rolling-endpoint regressions

We have documented that excess returns on fed funds futures were predictable using business-cycle indicators such as employment growth, corporate bond spreads, and Treasury yield spreads. To what extent could an investor have predicted these returns in real time, using only data that was available up to that point in time? To answer this question, we perform a set of rolling-endpoint regressions.

Fig. 3 shows real-time forecasts of excess returns on the four-month-ahead fed funds futures contract together with the full-sample forecasts from Panel C in Table 1 based on employment growth and the own futures contract rate.

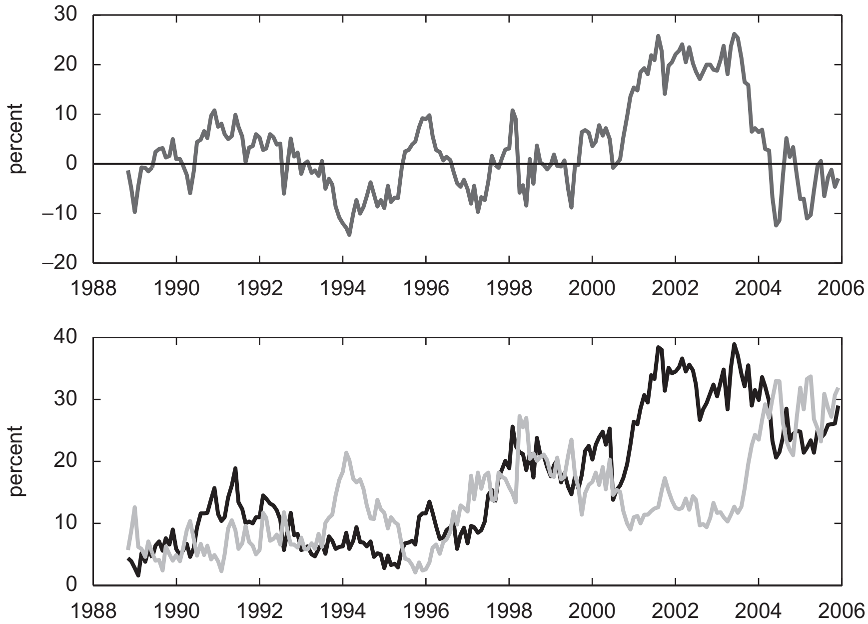

Fig. 4. The upper panel shows net positions in eurodollar futures. The lower panel shows long (black) and short (gray) positions separately.

The real-time forecasts for each month t are constructed by estimating the slope coefficients with data from October 1988 up through what was available at the end of the previous month $t-1$. Fig. 3 graphs these forecasts starting in October 1990, when we have only 24 months of data to estimate the three parameters of the model. The graph suggests that the real-time fitted values are quite close to the full-sample fitted values over most of the sample—indeed, the two series are essentially identical from the beginning of 1994 onward. The middle and lower panels in Fig. 3 show the rolling estimates of the slope coefficients together with their full-sample counterparts (the horizontal black line), and again suggest that the rolling point estimates have largely converged to their full-sample values by 1994.

3.3. Data on non-hedging market participants' positions

The previous section shows that excess returns on fed funds futures were potentially predictable to investors in real time using rolling regressions. In this section, we present some evidence indicating that informed investors at the time actually *did* correctly forecast the excess returns that were subsequently realized.

The U.S. CFTC requires all individuals or institutions with positions above a certain size to report their positions to the CFTC each week, and the extent to which each position is hedged. In the eurodollar (federal funds) futures markets, about 90% (95%) of open interest is held by individuals or institutions that must report to the CFTC as a result of this requirement. The CFTC reports the aggregates of these data with a three-day lag, broken down into hedging and non-hedging categories and into long and short positions, in the weekly Commitments of Traders report, available online.

The lower panel in Fig. 4 plots the percentage of long and short open interest in eurodollar futures held by noncommercial market participants—those market participants that are classified by the CFTC as *not* hedging offsetting positions that arise out of their normal (non-futures related) business operations.[12] The number of open long positions in these contracts held by noncommercial market participants (as a percentage of total reportable open interest) is depicted in the bottom panel of the figure by the black line, and the number of open short positions (as a percentage of reportable open interest) by the gray line.[13] The upper panel of Fig. 4 plots the difference between the noncommercial percentage long and short series as the "net long position" of noncommercial market participants.

The net long position graphed in Fig. 4 displays a clear positive correlation with the excess returns in federal funds futures contracts depicted in previous figures: for example, noncommercial market participants began taking on a huge net long position in late 2000, just a few months before excess returns in these contracts began to soar, and they took on

[12] The primary example of a *commercial* participant in the federal funds or eurodollar futures market would be a financial institution seeking to hedge its commercial and industrial loan portfolio. For more details on the institutional features of these markets, see Stigum (1990).

[13] Analogous data are available for fed funds futures positions as well, but we focus on eurodollar futures positions here as this market is thicker and contracts run off less frequently—only once per quarter rather than every month—which reduces some high-frequency variation in the percentage long and short series. Open interest is almost always highest in the front-month or front-quarter contract, so the running off of these contracts can create jumps. The patterns in fed funds futures noncommercial market participant holdings are similar to those in eurodollar futures, albeit noisier for this reason.

substantial long positions in mid-1990 through mid-1991 and late 1995 as well, again correctly forecasting excess returns over these periods; noncommercial market participants also took on a very substantial *short* position in late 1993 through mid-1994, correctly anticipating the low or even negative excess returns that were subsequently realized when the Federal Reserve began tightening policy in 1994.

These casual observations are confirmed in regression results (reported in Piazzesi and Swanson, 2006). The net long position variable is a significant predictor of excess returns in the fed funds futures market at horizons of two months or more, with R^2 values from 5% to 25%. Interestingly, the statistical significance of the net long position variable disappears if we also include any of employment growth, term spreads, or corporate risk spreads in the regression, suggesting that the information content of noncommercial market participants' net position is spanned by the business-cycle indicators that we have already considered in Section 2.

These results suggest that noncommercial market participants were aware of the upcoming excess returns in the market and positioned themselves accordingly, at the expense of those engaged in hedging other financial activities. The hedgers—primarily banks—essentially paid an insurance premium to noncommercial participants for providing hedging services. There are two primary explanations for why these premia were not "competed away" by the market. First, the futures market may not be perfectly competitive, with barriers to entry and noncommercial market participants facing limits on the size of the positions that they may take; commercial participants with hedging demand thus may not face a perfectly elastic supply curve for either the long or short side of these futures contracts. Second, noncommercial market participants may themselves be risk averse; for example, futures traders in these markets may be most averse to taking on large bets or risky positions precisely when their own jobs are most in jeopardy, around the times of recessions. The hypothesis that excess returns in these markets would be competed away requires both an assumption of perfectly competitive futures markets and of risk-neutral market participants, and both of these assumptions may not apply.

4. Risk-adjusted measures of monetary policy expectations

How misleading would it be to ignore risk premia on federal funds and eurodollar futures and treat the unadjusted prices of these securities as measures of monetary policy expectations? Using futures, the forecast errors are just minus the excess returns on the fed funds futures contract:

$$r_{t+n} - f_t^{(n)} = -rx_{t+n}^{(n)}. \tag{6}$$

To the extent that excess returns on federal funds and eurodollar futures are forecastable, one would be making systematic forecast errors if one used unadjusted futures rates as measures of monetary policy expectations.

However, we can risk-adjust these futures rates using our previous results. To do that, we take expectations of both sides of Eq. (4) and solve for the expected n-month-ahead federal funds rate:

$$\mathrm{E}_t[r_{t+n}] = f_t^{(n)} - (\alpha^{(n)} + \beta^{(n)} X_t). \tag{7}$$

From Table 1, we know that the expected excess return, $\alpha^{(n)} + \beta^{(n)} X_t$, is on average positive. This suggests that risk-adjusted forecasts lie on average below the futures rate. Moreover, Table 1 shows that expected excess returns are countercyclical. This suggests that risk-adjusted forecasts subtract a countercyclical term from the futures rate or, equivalently, add a procyclical term to the futures rate: risk-adjusted forecasts will tend to lie above the unadjusted futures rate in booms and below the futures rate in recessions.

These features of our risk-adjustment are illustrated in Fig. 5, which plots forecasts of the federal funds rate out to a horizon of 12 months on two different dates: December 1993 and December 2000. We plot a number of alternative forecasts based on federal funds and eurodollar futures rates:

1. Unadjusted futures: $\alpha^{(n)} = 0$ and $\beta^{(n)} = 0$.
2. Rule-of-thumb-adjusted futures: a constant risk adjustment of 1 bp per month, which is a rule of thumb that has been used by staff at the Federal Reserve Board for these interest-rate futures,[14] so $\alpha^{(n)} = n$ and $\beta^{(n)} = 0$. For eurodollar futures, the rule of thumb is 13 bp plus 3 bp per quarter: $\alpha^{(n)} = 3n + 13$ and $\beta^{(n)} = 0$.
3. Risk-adjusted futures: rolling OLS estimates of $\alpha^{(n)}$ and $\beta^{(n)}$, where X_t includes the own futures rate $f_t^{(n)}$ and NFP growth $\Delta\mathrm{NFP}_{t-1}$, as in Panel C of Table 1 for fed funds futures and eurodollar futures.

Fig. 5 illustrates that unadjusted futures forecasts are always higher than rule-of-thumb-adjusted forecasts. The two panels in Fig. 5 suggest that in times when the funds rate is expected to rise—such as December 1993—the higher, unadjusted futures therefore do better than rule-of-thumb-adjusted futures. However, the lower, rule-of-thumb-adjusted futures do better in times when the funds rate is expected to fall, such as December 2000. This is exactly the mechanism

[14] In private communication dated May 2004, Donald Kohn mentioned that staff at the Federal Reserve Board came up with this adjustment factor informed by their reading of the historical data on ex post errors in the federal funds and eurodollar futures markets and in interest-rate surveys. Although this adjustment factor was 1 bp per month at the short end of the futures curve at the time of that communication, he noted it has not always been that and would change as events warrant.

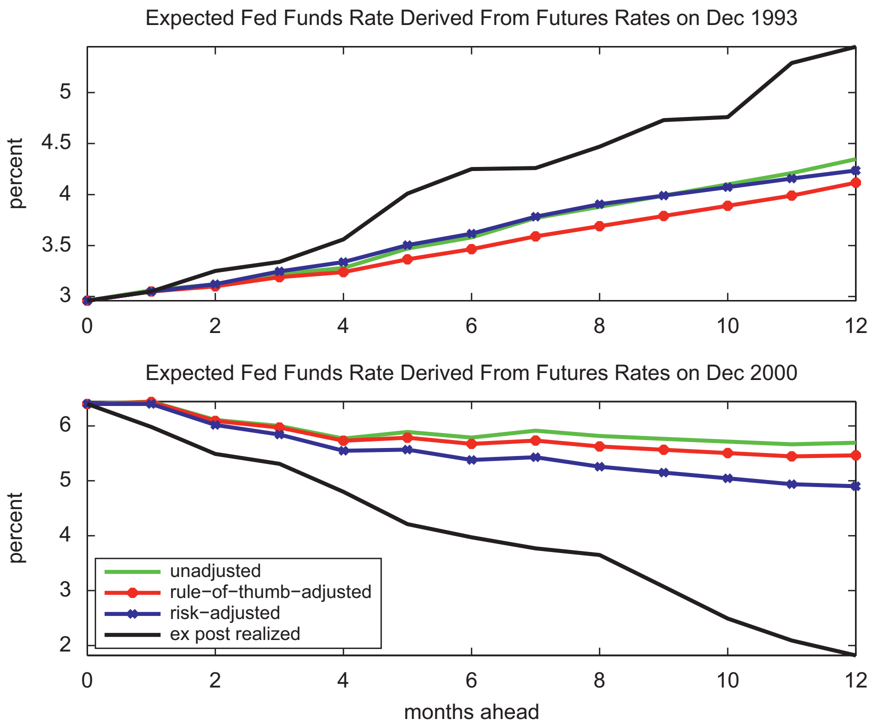

Fig. 5. Federal funds rate forecasts on two illustrative dates, and subsequent realized funds rate. Funds rate forecasts are constructed from unadjusted and risk-adjusted futures rates, and using three different risk adjustments: an estimated constant adjustment, a rule-of-thumb-constant adjustment, and a time-varying risk adjustment based on employment growth.

exploited by our time-varying risk adjustment: in December 1993, our risk-adjusted futures forecast (the x-line) is closer to the unadjusted futures forecast, while in December 2000, the x-line is closer to the rule-of-thumb-adjusted futures forecast.

Fig. 5 suggests that unadjusted futures rates, or futures adjusted by a constant, can be wrong over long periods of time. The forecast errors tend to be negative during periods of falling rates and positive during periods of rate hikes. The forecast errors are largest when the funds rate changes direction, and keep being large for substantial amounts of time. The reason is that unadjusted or constant-adjusted futures rates only slowly adapt to changes in direction. As a result, these forecasts tend to *lag behind actual market expectations* around economic turning points; they generate forecast errors that are more autocorrelated than forecast errors from risk-adjusted futures.

To see this point more clearly, Table 3 reports some summary statistics for forecast errors, for each of a number of different forecasts. We compute forecast errors from futures-based forecasts and also from a simple vector autoregression (VAR) as a benchmark.[15] The VAR is computed at monthly frequency using four lags of each of the fed funds rate, the year-on-year percentage change in the core CPI, and the year-on-year percentage change in nonfarm payroll employment.[16] We compute forecasts for the n-month-ahead fed funds rate as it would have been made at each time t, using real-time data and rolling-endpoint regressions. For example, when we compute forecasts for r_{t+n} using the VAR benchmark, we estimate the parameters of the VAR using only data up through time $t-1$ and then use the values of the fed funds rate, core CPI inflation, and nonfarm payrolls growth at time t as the conditioning variables for the forecast. Similarly, we use our rolling "out-of-sample" forecasts for risk premia based on nonfarm payrolls growth to make our risk adjustments. The forecast errors are computed over the October 1990–December 2005 period, so that we have two years of data to estimate the parameters for the October 1990 forecast.

Table 3 reports mean forecast errors (ME), root-mean-squared errors (RMSE), and the nth autocorrelation (ρ_n) for the n-month-ahead forecast. (Note that even for an efficient n-month-ahead forecasts, the forecast errors would have an MA($n-1$) autocorrelation because of forecast overlap; Table 3 therefore reports the nth autocorrelation which, ideally, should be zero.) The last column of Table 3 shows that risk-adjusted futures still made autocorrelated forecast errors over

[15] We also considered forecasts from an AR(1) and a random walk. The resulting forecasts, however, were outperformed by the VAR, so we did not include them in Table 3 to conserve space. Another alterative are Taylor-rule forecasts. For forecasts up to three months ahead, Evans (1998) documents that they tend to be dominated by forecasts based on (unadjusted) futures.

[16] The lag length was selected to yield good empirical forecast performance: more lags than 4 tended to lead to overfitting and poor forecast performance, while fewer lags tended to lead to overly simple dynamics and poor forecast performance.

Table 3
Forecasts of the federal funds rate

	Benchmark			Federal funds futures-based forecasts								
	VAR(4)			Unadjusted future			Rule-of-thumb-adjusted futures			Risk-adjusted futures		
n	ME	RMSE	ρ_n	ME	RMSE	ρ_n	ME	RMSE	ρ_n	ME	RMSE	ρ_n
(A) Federal funds rate forecasts												
1	−1.4	27	0.67	−3.2	11	0.10	−2.2	11	0.06	−0.5	12	0.08
2	−1.1	36	0.57	−7.6	19	0.28	−5.6	18	0.23	−2.1	19	0.19
3	−1.9	47	0.39	−12.4	29	0.37	−9.4	28	0.33	−4.3	29	0.20
4	−2.6	57	0.24	−18.2	41	0.39	−14.2	40	0.34	−7.0	38	0.16
5	−3.6	66	0.13	−24.7	54	0.40	−19.8	52	0.40	−8.7	46	0.14
6	−5.2	75	0.08	−30.8	66	0.49	−24.8	64	0.45	−8.9	50	0.16
	Benchmark			Eurodollar futures-based forecasts								
	VAR(3)			Unadjusted future			Rule-of-thumb-adjusted futures			Risk-adjusted futures		
n	ME	RMSE	ρ_n	ME	RMSE	ρ_n	ME	RMSE	ρ_n	ME	RMSE	ρ_n
(B) Eurodollar rate forecasts												
1	15	120	0.80	−17	44	0.40	−1	40	0.29	−7	38	0.07
2	27	135	0.45	−43	89	0.56	−24	81	0.48	−21	64	0.28
3	31	144	0.10	−73	134	0.57	−51	123	0.51	−37	90	0.17
4	25	161	−0.22	−105	182	0.47	−80	169	0.39	−54	116	0.04
5	11	180	−0.34	−135	223	0.38	−107	207	0.30	−68	138	0.05
6	−17	198	−0.26	−163	256	0.31	−132	238	0.22	−83	159	0.16
7	−45	209	−0.13	−181	281	0.26	−147	260	0.16	−86	179	0.16
8	−72	206	0.17	−200	297	0.25	−163	274	0.15	−92	195	0.06

Note: n is the forecasting horizon in months (quarters for eurodollar rate). ME denotes the mean forecast error (in basis points), RMSE the root-mean-squared error (in bp), and ρ_n is the nth autocorrelation of the forecast error, all over the period 1990:10–2005:12 (1990Q3–2005Q4 for eurodollars). VAR(4) is a monthly VAR forecast based on four lags of each of the federal funds rate, year-on-year percentage change in the core CPI, and year-on-year percentage change in nonfarm payrolls. VAR(3) is a quarterly VAR based on three quarterly lags of the 90-day eurodollar rate, year-on-year percentage change in the core CPI, and year-on-year percentage change in nonfarm payrolls. The risk-adjusted futures-based forecast adjusts the fed funds futures rate for risk premia using the own futures rate and the year-on-year percentage change in nonfarm payrolls. Coefficients of the VAR and the risk-adjustment regression are recomputed at each date t using data only up through month $t-1$, so all forecasts are pseudo-out-of-sample.

our sample, but the autocorrelation is much smaller than for any other forecast in the table. This is especially true for longer forecasting horizons. Moreover, risk-adjusted futures generate smaller average errors and lower root-mean-square errors.

Interestingly, Panel A makes a strong case for fed funds futures in general, even on a risk-unadjusted basis. The futures-based forecasts produce lower root-mean-square errors than a VAR. However, unadjusted futures made large, negative errors on average, ranging from 3 to 31 bp. The rule-of-thumb-adjusted futures improve upon this: average forecast errors are lower by exactly the amount of the adjustment, and the adjustment also lowers mean-square errors. However, this adjustment only represents a small improvement over unadjusted forecasts. The risk-adjusted forecasts we estimate in this paper generate forecast errors that are always smaller on average and almost always smaller in root-mean-square terms, especially for longer forecasting horizons. Panel B confirms these findings for longer-horizon forecasts using eurodollar futures.[17] Again, risk-adjusted futures do much better than unadjusted futures or the rule-of-thumb-adjusted futures.

5. Monetary policy shocks

Fed funds futures have also been used by a number of recent authors to separate systematic changes in monetary policy from monetary policy "shocks".[18] The idea is to use fed funds futures market forecast errors as measures of exogenous, unforecastable changes in the stance of monetary policy.[19] The fed funds futures market expectation is measured assuming the expectations hypothesis. Since we have shown in the previous section that futures rates should be adjusted for time-

[17] The VAR for the 90-day eurodollar rate uses the same variables as the monthly VAR but sampled at quarterly frequency beginning in 1990Q3 with data going back to 1985Q1. We chose a lag length of 3 since this seemed to give good empirical forecast performance: a lag length of 4 performed worse and a lag length of 2 performed better at the shortest horizons but worse at longer horizons.

[18] See, e.g., Rudebusch (1998), Cochrane and Piazzesi (2002), and Faust et al. (2004) for different approaches. All of these studies treat the federal funds rate as the monetary policy instrument, as in Bernanke and Blinder (1992), and attempt to improve upon the earlier VAR-based identification of monetary policy shocks surveyed in Christiano et al. (1999).

[19] Faust et al. (2004) describe the procedure in detail and test many of the required assumptions. Alternatively, Piazzesi (2005) and Cochrane and Piazzesi (2002) measure market expectations from high-frequency data on short-term interest rates instead of fed funds futures. Piazzesi (2005)

Table 4
Risk-adjusting measures of monetary policy shocks

	Actual-futures		Change in FF futures	
	Original	Adjustment	Original	Adjustment
Panel A: summary statistics of policy shocks				
Mean	−3.0	−3.0	−1.2	−1.2
Std. dev.	11.0	3.2	8.2	2.0
Min	−43.8	−12.6	−42.5	−6.4
Max	17.1	4.1	14.5	15.2

	Constant	1 year–6 months	2–1 years	5–2 years	10–5 years	R^2	p-value
Panel B: t-statistics from regressions on Treasury spreads							
Actual-futures	−2.0	−1.7	2.6	−2.5	2.2	0.09	0.054
Change in FF futures	−1.5	−0.6	1.8	−1.6	1.3	0.06	0.431

Note: Daily observations on days of FOMC meetings and intermeeting policy moves, 1994–2005.

varying risk premia, we now investigate to what extent these risk premia might affect futures-based measures of monetary policy shocks.

Computing the futures market's forecast error of the next policy move is less straightforward than it may seem, because of some institutional features of the fed funds and fed funds futures markets. For example, the futures contract settles based on the average fed funds rate that is realized during the contract month, and not on the value of the funds rate on a particular date, such as the day following an FOMC meeting. Moreover, the Federal Reserve sets a target for the funds rate, but does not completely control the funds rate itself, and the difference between the actual funds rate and the target can be nonnegligible, even for monthly averages. In the literature, these complications have led to alternative approaches on how policy shocks are computed from futures rates.

Here, we consider the two primary approaches to measuring monetary policy shocks that have been used in the literature. First, Rudebusch (1998) suggests defining the monetary policy shock as the difference between the realized fed funds rate target and the expected fed funds rate derived from fed funds futures. While this might seem to be the most natural definition of the market's forecast error, it can suffer from the technical issues described above that cause the market expectation of the future realized funds rate to differ from the market expectation of the future target rate. More importantly, monetary policy shocks measured in this way will be contaminated by risk premia in the futures market even if those risk premia are constant. Our second measure of monetary policy shocks, suggested by Kuttner (2001), differences out both the technical factors in the fed funds market and any constant risk premia by using the *change* in the current-month or one-month-ahead fed funds futures rate on the day of an FOMC announcement. This approach uses daily fed funds futures data to make the interval $[t, t+1]$ around the FOMC announcement small and assumes that risk premia do not change over this small interval.[20]

We compute these two measures of monetary policy shocks over the sample period 1994–2005, when the Federal Reserve was explicitly announcing changes in its target for the fed funds rate. We include every FOMC meeting and every intermeeting policy move by the FOMC over this sample. Table 4 reports summary statistics for both measures of policy shocks. From Panel A it is apparent that the first measure of monetary policy shocks, labeled "actual-futures", is larger and more volatile than the second measure: the mean, standard deviation, and extremes of the shocks are all larger in the first shock series than in the second. The two shocks series do generally agree on the days of large monetary policy shocks, however—for example, the min and max of the two series both occur on the same days.

We investigate the robustness of the two monetary policy shock series to risk premia by regressing them on a set of conditioning variables that were known to financial markets right before the FOMC announcement—for this exercise, we pick Treasury yields as the regressors, because we have high-frequency data on these yields.[21] Under the expectations hypothesis, each monetary policy shock measure should be unpredictable on the basis of these conditioning variables.

Results for these regressions are summarized in Panel B of Table 4, which also reports the p-values from a zero-slopes F-test that all of the coefficients (excluding the constant term) in each regression are jointly equal to zero. As can be seen in

(*footnote continued*)
computes $E_t[r_{t+1}]$ from an arbitrage-free model of the term structure of interest rates. Cochrane and Piazzesi (2002) use the change in the one-month eurodollar rate and unrestricted regressions of r_{t+1} on a set of interest rates.

[20] This assumption is consistent with the finding by Evans and Marshall (1998) that risk premia in Treasuries show only a small response to monetary policy shocks in a VAR. It is also potentially consistent with our findings above and those of Cochrane and Piazzesi (2005) that risk premia vary substantially at business-cycle frequencies but perhaps do not vary as much at higher frequencies.

[21] Results using corporate bond spreads are very similar. In both cases, we reestimate the predictability regression coefficients for the monetary policy shock series for two reasons: first, risk premia depend on the maturity of the contract and FOMC meetings are typically not scheduled for the end of the month, and second, the risks associated with monetary policy shocks (and therefore the risk premia associated with these shocks) may have changed after 1994, when the Fed started announcing its policy moves at FOMC meetings, as argued in Piazzesi (2005).

the table, our first measure of monetary policy shocks (the realized target rate minus the futures market expectation) appears to be significantly contaminated both by a constant risk premium (as evidenced by the *t*-statistic for the constant term) and possibly also by time variation in risk premia (as evidenced by the zero-slopes *F*-test, which is borderline statistically significant). In contrast, our second policy shock measure (the one-day change in the federal funds futures rate) seems to do much better: no coefficient is statistically significant, and we do not reject the null of no contamination by time-varying risk premia.

Panel A of Table 4 also reports basic statistics for our estimated risk adjustments to each of the three monetary policy shock series. The risk adjustments for our first shock measure ("actual-futures") has a standard deviation of 3.2 bp, which seems substantial relative to the standard deviation of the shock series itself of 11 bp. In contrast, the estimated risk adjustments to our second measure of monetary policy shocks are much smaller as well as statistically insignificant.

We infer from this analysis that monetary policy shocks based on the difference between the realized federal funds rate target and the unadjusted federal funds futures-based forecast, as suggested by Rudebusch (1998), are significantly contaminated by risk premia, both on average and by an amount that varies over time. The primary alternative—measuring monetary policy shocks based on the one-day change in the federal funds futures rate around FOMC announcements, as suggested by Kuttner (2001)—seems to be much more robust to the presence of risk premia in these contracts. The difference-based measure may largely "difference out" risk premia that are moving primarily at lower, business-cycle frequencies, consistent with our analysis in Section 2 and the findings in Cochrane and Piazzesi (2005).

6. Conclusions

We document substantial and predictable time variation in excess returns on federal funds futures. We show that excess returns on these contracts are strongly countercyclical and can be predicted with R^2 of up to 39% using contemporaneous macroeconomic and financial business-cycle indicators such as employment growth, Treasury yield spreads, and corporate bond spreads. We also present evidence that suggests that market participants could have been and in fact were aware of these excess returns in real time, as evidence by real-time rolling-endpoint regressions and the observation that non-hedging futures market participants were large buyers of these contracts in times of high expected excess returns and large sellers in times of low or negative expected excess returns.

Our findings of predictable excess returns in federal funds futures contracts have important consequences for computing market expectations from these futures rates. We find that ignoring these risk premia substantially increases forecast errors, both on average and in terms of root-mean-squared error. Moreover, unadjusted futures make forecast errors that are more autocorrelated, because unadjusted futures-based forecasts lag behind our risk-adjusted forecasts around economic turning points. Finally, we show that measures of monetary policy shocks based on the realized funds rate target minus the ex ante unadjusted fed funds futures rate are significantly contaminated by risk premia. Instead, a measure of monetary policy shocks based on the one-day change in federal funds futures around FOMC announcements seems to be more robust—for instance, it may largely "difference out" risk premia that move primarily at lower, business-cycle frequencies.

References

Altig, D., 2005. Funds rate probabilities: September 25, November a toss-up. Macroblog, September 12, downloaded from ⟨http://macroblog.typepad.com/macroblog⟩.
Bernanke, B.S., Blinder, A.S., 1992. The federal funds rate and the channels of monetary transmission. American Economic Review 82, 901–921.
Bernanke, B.S., Kuttner, K.N., 2005. What explains the stock market's reaction to federal reserve policy? Journal of Finance 60, 1221–1257.
Board of Governors of the Federal Reserve System, 2005. Monetary Policy Report to the Congress: July 20, 2005. Washington, DC.
Campbell, J.Y., Shiller, R.J., 1991. Yield spreads and interest rate movements: a bird's eye view. Review of Economic Studies 58, 495–514.
Christiano, L., Eichenbaum, M., Evans, C., 1999. Monetary policy shocks: What have we learned and to what end? In: Woodford, M., Taylor, J. (Eds.), Handbook of Macroeconomics. North-Holland, Amsterdam.
Cochrane, J., 2001. Asset Pricing. Princeton University Press, Princeton, New Jersey.
Cochrane, J.Y., Piazzesi, M., 2002. The fed and interest rates: a high-frequency identification. American Economic Review 92 (2), 90–95.
Cochrane, J.Y., Piazzesi, M., 2005. Bond risk premia. American Economic Review 95 (1), 138–160.
Crane, A.T., Mackenzie, M., 2005. Investors reappraise rate outlook—more traders see a pause in fed's cycle of increases amid fallout from Katrina. The Wall Street Journal September 2, C4.
Durham, B., 2003. Estimates of the term premium on near-dated federal funds futures contracts. FEDS Working Paper 2003-19, Federal Reserve Board.
European Central Bank, 2005. Monthly Bulletin: September, Frankfurt.
Evans, C., 1998. Real-time Taylor rules and the federal funds futures market. Federal Reserve Bank of Chicago, Economic Perspectives 2, 44–55.
Evans, C., Marshall, D., 1998. Monetary policy and the term structure of nominal interest rates: evidence and theory. Carnegie-Rochester Conference Series on Public Policy 49, 53–111.
Fama, E.F., Bliss, R.R., 1987. The information in long-maturity forward rates. American Economic Review 77, 680–692.
Faust, J., Swanson, E., Wright, J., 2004. Identifying VARs based on high-frequency futures data. Journal of Monetary Economics 6, 1107–1131.
Gürkaynak, R., Sack, B., Swanson, E., 2007. Market-based measures of monetary policy expectations. Journal of Business and Economic Statistics 25 (2), 201–212.
Hamilton, J., 2006. A pause it shall be. Econbrowser, August 4 ⟨http://www.econbrowser.com⟩.
Hansen, L.P., Hodrick, R.J., 1988. Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. Journal of Political Economy 88, 829–853.
Hodrick, R.J., 1992. Dividend yields and expected stock returns: alternative procedures for influence and measurement. Review of Financial Studies 5, 357–386.
Horowitz, J.L., 2004. The bootstrap. In: Heckman, J.J., Leamer, E. (Eds.), Handbook of Econometrics, vol. 5. Elsevier, New York, pp. 3160–3228.

Hughes, J., Beales, R., 2005. Katrina could alter course of fed. Financial Times September 2, 21.
Krueger, J.T., Kuttner, K.N., 1996. The fed funds futures rate as a predictor of federal reserve policy. Journal of Futures Markets 16, 865–879.
Kuttner, K.N., 2001. Monetary policy surprises and interest rates: evidence from the federal funds futures market. Journal of Monetary Economics 47, 523–544.
Piazzesi, M., 2005. Bond yields and the federal reserve. Journal of Political Economy 113 (2), 311–344.
Piazzesi, M., Swanson, E., 2006. Futures prices as risk-adjusted forecasts of monetary policy. Working Paper 2006-23, Federal Reserve Bank of San Francisco.
Rudebusch, G., 1998. Do measures of monetary policy in a VAR make sense? International Economic Review 39, 907–931.
Rudebusch, G., 2002. Term structure evidence on interest rate smoothing and monetary policy inertia. Journal of Monetary Economics 49 (6), 1161–1187.
Sack, B., 2004. Extracting the expected path of monetary policy from futures rates. Journal of Futures Markets 24 (8), 733–754.
Stambaugh, R.F., 1988. The information in forward rates: implications for models of the term structure. Journal of Financial Economics 21, 41–70.
Stigum, M., 1990. The Money Market. McGraw-Hill, New York.