Version of 10/22/2010

Comments Welcome

The Impact of the Federal Reserve's Interest Rate Target Announcement on Stock Prices: A Closer Look at How the Market Impounds New Information

Justin Birru*

and

Stephen Figlewski**

* Ph.D. Student New York University Stern School of Business 44 West 4th Street, Suite 9-160 New York, NY 10012-1126	** Professor of Finance New York University Stern School of Business 44 West 4th Street, Suite 9-160 New York, NY 10012-1126
212-998-0576 jbirru@stern.nyu.edu	212-998-0712 sfiglews@stern.nyu.edu

The authors are grateful to the International Securities Exchange for providing the options data used in this study, and to Robin Wurl for extracting it. We also thank Rob Engle and other participants in the NYU Finance seminar for valuable comments.

The Impact of the Federal Reserve's Interest Rate Target Announcement on Stock Prices: A Closer Look at How the Market Impounds New Information

ABSTRACT

The Federal Reserve announces its new interest rate target while the stock market is open, at precisely 2:15 P.M. eight times a year. In the Efficient Markets model, an information shock is impounded in prices immediately and accurately as soon as it becomes public knowledge, and only the unanticipated portion moves prices. Responding accurately to news requires investors to judge how much other investors have been surprised and how their investment decisions will be affected, so how the market responds to the news generates additional information to be digested and acted upon. This suggests that the full process of returning to equilibrium after an information event can not be instantaneous. In this paper, we explore the "informational microstructure" of the stock market around Fed funds target announcements by examining the market's risk neutral probability density for future stock prices, that we extract from real-time option prices using a non-model dependent procedure. Our results show that the market's adjustment to the news continues well beyond the initial information release.

Keywords: risk neutral density, Fed funds target announcement, efficient markets

JEL Classification: G14, G13, D84

I. Introduction

In the 1960s, pathbreaking work on the concept of informationally efficient markets by Fama (1965), Samuelson (1965) and others, summarized in Fama's (1970) classic review article, laid one of the major foundations of modern financial theory. Before then it was widely believed that stock prices moved in trends or followed more complicated patterns that could be uncovered and exploited through "technical analysis" of past prices. Fama's powerful counterargument was that for any such regular pattern to exist, investors would essentially have to be throwing money away by ignoring, or consistently misjudging, important information that was public knowledge.

By contrast, rational investors will, at least, evaluate all publicly available information correctly. When market expectations are rational, every security will be priced to earn a fair expected return commensurate with its risk. Only surprises will move prices, and by definition, surprises can not follow a predictable time pattern, so return shocks must be serially uncorrelated. Today this principle is built into the specification of virtually every theoretical asset pricing model, by the assumption that the stochastic information process driving returns is Brownian motion or some other Markov process.

The Efficient Markets model has strong implications about how newly released information is incorporated into prices. Anything that can not be predicted ahead of time, such as a fire that destroys a major factory, will not affect the market until it occurs, but once a new piece of information is publicly released, it should be fully impounded in prices immediately. After the initial impact, there should be no opportunity for an investor to make an excess return by trading on what is now public knowledge. In the case of an event that is partially predictable beforehand, like a stock split or a takeover bid, as the information seeps into the market, prices will adjust appropriately at each point in time. There will be one final price impact on announcement day, after which no further predictable movement should occur.

Innumerable research studies over the years have documented that market responses to a broad range of events and information releases do indeed follow these patterns closely. The first "event study," by Fama, Fisher, Jensen and Roll (1969) (FFJR), looked at how stock prices responded to an announcement of a stock split, an event that could be partially but not fully anticipated ahead of time. FFJR produced what has since become the classic graph of how stock prices in an efficient market adjust to new information. On average, before the event the stock price rose consistently as the future split came to look increasingly probable. There was a final jump at the time of the announcement, and no further increase thereafter. FFJR has been on the reading list in Finance Ph.D. core theory courses since it was published, and the picture of price adjustment just described is ingrained in the collective psyche of the finance profession as the image of how new information is incorporated into stock prices.

As finance research techniques have advanced and as more extensive market data have become available to researchers, it has become possible to examine market price formation at close range. This revealed that while Brownian motion may be a useful approximation for market behavior over periods of days and months, things are different over very short time intervals in which price movements are dominated by market microstructure, the mechanics of the marketmaking process. In this paper, we wish to examine what might be called the market's informational microstructure: what happens when an important new piece of information is released to the market and how it is incorporated into stock prices.

The specific information we examine is the Federal Reserve's announcement of its future target for the federal funds rate. In addition to its direct impact on bank borrowing costs, this announcement is viewed more broadly by the market as an extremely important signal of how the Fed views the current state of the economy and how it will manage monetary policy going forward. The target rate is announced publicly at precisely 2:15 P.M. New York time, at the conclusion of the meeting of the Federal Open Market Committee (FOMC) eight times a year. It is an excellent information event for studying the impact of a news release on the stock market because the market views it as very important, so it has a large impact on prices, and it is announced at a pre-specified time during trading hours, so the timing of the release is well-known ahead of time and it happens while trading is in progress.

Figure 1 plots the behavior of the forward value of the S&P 500 index during the day on December 11, 2007 when the Fed announced a 25 basis point cut in the Fed funds target.[1] That morning the market opened well above the previous night's close and essentially proceeded to mark time, with low volatility, while waiting for the Fed announcement at 2:15 P.M. When the announcement came, although the target rate was cut, the market's response was strongly negative. In ten minutes between 2:14 P.M. and 2:24 P.M. the index fell over 26 points, more than 1.7%. Plainly the market had been expecting a bigger rate cut and was disappointed that it was only 25 b.p. But this picture does not quite look like the canonical FFJR plot. After the initial impact, the market continued trading with much greater volatility than before, and the index continued moving sharply downward for the rest of the day, ending finally at about 1489.

This is only a single day, but it exemplifies the nature of price behavior that is typical for Federal Reserve target rate announcements. There is an immediate significant impact when the news is released, as expected. But there is also a definite period afterwards during which the market is very volatile as it seeks a new equilibrium that takes into account all of the information generated by the new information.

The public announcement of the new target reveals an important fact to the market. But the raw information only becomes incorporated in prices by altering investors' expectations about the future and/or their assessment of risk and their attitudes about bearing that risk. A critical factor in that process, especially in building the news into

[1] We will be analyzing the risk neutral probability density extracted from S&P 500 index options prices. The mean of that density is the risk neutral expected value of the index on option expiration date, so the appropriate comparison to explore differences between the options market and the stock market is between the risk neutral mean and the expiration day forward value of the index, computed from the current spot S&P index.

return expectations, is how other investors are responding to the news. The behavior of the market itself following the information release produces new information that investors must also incorporate in their revised beliefs.

We are particularly interested in exploring what are two different and distinct phases of information assimilation. In an efficient market, as soon as information is released, investors immediately adjust their positions based on how they expect the news to affect prices. We expect the initial impact period for an important and highly visible public announcement like this one to be very short. Some of the market's instantaneous response may occur automatically, when the market's price change touches off standing price sensitive limit and stop-loss orders in the order book. The second phase, however, involves feedback from the market's response to investors' trading decisions, which then produce further market fluctuations for investors to consider, and so on. Reaching the new full informational equilibrium may take some time.

Empirical tests of models involving expectations are difficult because prices can be easily seen but the market expectations and risk attitudes that determine those prices are not directly observable. Harrison and Kreps (1979) proved that the expected objective probability distribution for returns could always be combined with risk preferences into a single risk neutral density (RND). This construct has the property that in equilibrium, the price of every asset will be its expected payoff under the risk neutral distribution discounted back to the present at the riskless rate of interest.

At about the same time, Breeden and Litzenberger (1978) showed how the RND could be extracted from a set of option prices. There are a number of significant difficulties in adapting their theoretical result for use with actual option prices observed in the market, but Figlewski (2009) develops a methodology that performs well. We will apply it to an extraordinarily detailed dataset of real-time best bid and offer quotes in the consolidated national options market, which allows a very close look at the behavior of the RND, essentially in real-time.

The idea of extracting information from market option prices is related to the calculation of implied volatilities, done extensively by option traders and researchers. But that procedure is heavily model-dependent. For it to deliver the market's true risk neutral expectations, the model used in the extraction must be exactly the same as the model the market is using to price options.[2] We also obtain market expectations by analyzing option prices, but one of the great advantages of our procedure is that it is not model dependent. The risk neutral density can be extracted without assuming market prices come from any particular pricing model.

---

[2] Despite the fact that Black-Scholes implied volatilities are extensively used by options traders, and researchers, the ubiquitous volatility smile they exhibit contradicts this assumption. A very large amount of research has explored the connection between implied volatility and subsequent realized volatility from many perspectives. Only in rare cases does statistical analysis fail to indicate that implied volatility is a biased forecast. See Jackwerth (2004) or Poon and Granger (2003) for extensive reviews of this literature.

We will use the behavior of the risk neutral density over the future value of the Standard and Poors 500 stock index extracted at one-minute intervals throughout the day when the Fed funds target rate is announced to explore how the market assimilates the new information, in both phases of the process. Very little previous work has been done to investigate the impact of the Fed's announcement on stock prices with intraday data, and virtually none using risk neutral densities from index options in real time. Our first investigation in this area sheds light on a number of important issues, and also uncovers a few provocatively anomalous results.

The next section gives a brief and incomplete review of the related literature, although as just mentioned, there is little that is directly comparable to what is done in this study. Section III describes how the RND can be extracted from option prices, as set forth by Breeden and Litzenberger. Additional detail is provided in the Appendix. The data is described in Section IV.

Section V looks at how the forward value of the stock index and the mean and variance of the risk neutral density behave on announcement days, before and after the news release. In addition to the mean, the risk neutral density reveals the variance of the market's expectation for the index at option expiration date, which allows us to measure directly the resolution of uncertainty produced by the Fed's announcement. In Section VI, we explore how the changes in the forward index and the RND mean over different time periods on announcement days are related to variables that might be expected to contain relevant information. Section VII looks at the behavior of price volatility within the announcement day. Volatility is related to the rate of information flow, and we investigate the process of re-equilibration after the initial impact of the announcement. Section VIII concludes.

II. Review of the Literature

Academics have long studied the effects of monetary policy and FOMC announcements of federal funds target rate decisions on various financial markets. The majority of these studies investigate the response of equity or bond markets in the hours or days around the announcement. Bernanke and Kuttner (2005) undertake an event-study approach to assess the effects of monetary policy on equity returns. They identify a negative relationship between unanticipated changes in the federal funds target rate and equity prices. Chulia-Soler, Martens, and van Dijk (2007) and Zebedee, Bentzen, and Lunde (2008) examine intraday data and also document a negative relationship between equity returns and the direction of the target rate surprise. Both studies also identify an increase in equity volatility arising from target rate surprises. Other studies using intraday data to analyze either return or volatility responses to monetary policy include Davig and Gerlach (2006), Hausman and Wongswan (2006), Andersson (2007), Basistha and Kurov (2008), Brenner, Pasquariello and Subrahmanyam (2009), and Farka (2009).

Far fewer studies have examined the reaction of the options market to FOMC announcements. The studies that do examine the link between monetary policy and option markets often focus on the effect of policy on option implied volatilities. Both Nikkinen and Sahlstrom (2004) and Chen and Clements (2007) study the effect of different monetary policy announcements on the VIX index. Nofsinger and Prucyk (2003) also examine S&P index option implied volatilities, as well as index option volume reactions to a number of different macro announcements. They find that bad news leads to increases in volatility and volume, while good news results in lower volume and is not associated with subsequent higher volatility. Finally, Beber and Brandt (2006) utilize options on bond futures to extract the option-implied state-price densities of bond prices around various macroeconomic announcements. They specifically examine the changes in higher order moments of the state-price density around the announcement conditional on the information content of the announcement.

A third strand of literature investigates RNDs. The option-implied RND studies most closely related to the current paper are those utilizing RNDs to analyze market forecasts. Bates (1991), Bates (2000) and Jackwerth and Rubinstein (1996), among others, examine the option-implied probabilities of left-tail events. Other papers, such as Gemmil and Saflekos (2000) and Alonso, Blanco and Rubio (2005) look at the forecasting ability of RNDs, while Lynch and Panigirtzoglou (2008) study the time-series properties of the RND for the S&P index from 1985-2001. Figlewski (2009) develops a new method for extracting RNDs from option prices, and utilizes this methodology to examine the option-implied RND for the S&P index from 1996-2008. Finally, Birru and Figlewski (2010) examine intraday option-implied S&P index RNDs during the financial meltdown of 2008.

The existing literature utilizing option-implied risk-neutral distributions to gauge market expectations has seen marked growth in the past decade. As discussed above, there also exists a large literature examining the effects of monetary policy on financial markets. However, the interactions of monetary policy and option-implied risk-neutral distributions have yet to be analyzed.

III. Extracting Risk Neutral Densities from Option Prices

In the following, the symbols C, S, X, r, and T all have the standard meanings of option valuation: C = call price; S = time 0 price of the underlying asset; X = exercise price; r = riskless interest rate; T = option expiration date, which is also the time to expiration. P will be the price of a put option. We will also use f(x) = risk neutral probability density function, also denoted RND, and $F(x) = \int_{-\infty}^{x} f(z)dz$ = risk neutral distribution function.

The value of a call option is the expected value of its payoff on the expiration date T, discounted back to the present. Under risk neutrality, the expectation is taken with respect to the risk neutral probabilities and discounting is at the risk free interest rate.

(1) $$C = e^{-rT} \int_X^\infty (S_T - X) f(S_T) dS_T$$

Taking the partial derivative in (1) with respect to the strike price X gives

$$\frac{\partial C}{\partial X} = -e^{-rT} \int_X^\infty f(S_T) dS_T = -e^{-rT} \left[1 - F(X)\right]$$

Solving for the risk neutral distribution F(X) yields

(2) $$F(X) = e^{rT} \frac{\partial C}{\partial X} + 1$$

Taking the derivative with respect to X a second time gives the risk neutral density function

(3) $$f(X) = e^{rT} \frac{\partial^2 C}{\partial X^2}$$

In the market, option prices are only available for a discrete set of exercise prices, so we approximate the solution to (3) using finite differences.

Let there be call options available for maturity T at N different exercise prices, with X_1 representing the lowest exercise price and X_N being the highest. In this procedure, the X's are structured to be equally spaced for convenience, that is, $X_n - X_{n-1}$ is a constant for all n.

To estimate the total probability in the left tail of the risk neutral distribution up to X_2, we approximate $\frac{\partial C}{\partial X}$ at X_2 by $e^{rT} \frac{C_3 - C_1}{X_3 - X_1} + 1$, and the total probability in the right tail from X_{N-1} to infinity is approximated by $1 - \left(e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}} + 1 \right) = -e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}}$.

The density $f(X_n)$ is approximated by the numerical second partial derivative at X_n,

(4) $$f(X_n) \approx e^{rT} \frac{C_{n+1} - 2C_n + C_{n-1}}{(\Delta X)^2}.$$

Equations (1) - (4) show how an approximation to the portion of the RND lying between X_2 and X_{N-1} can be extracted from a set of call option prices. A similar derivation can be done to yield a procedure for obtaining the RND from put prices. The equivalent expressions to (2) and (3) for puts are:

(5) $$F(X) = e^{rT} \frac{\partial P}{\partial X}$$

and

(6) $$f(X) = e^{rT} \frac{\partial^2 P}{\partial X^2}$$

These relationships are applied to market data by replacing the partial derivatives with numerical approximations, as with calls.

Implementing this strategy with market option prices to obtain a well-behaved RND is a nontrivial exercise. Several key problems need to be dealt with, including how to obtain or generate option prices at enough exercise prices for the numerical derivatives in (4) to be accurate and how to extend the density into the tails, and numerous alternative approaches have been explored in the literature. We have developed a consistent approach that works well. A brief summary of the key steps is given in the Appendix; for full details, the interested reader should refer to Figlewski (2009).

IV. Data

The intraday options data are the national best bid and offer (NBBO) extracted from the Option Price Reporting Authority (OPRA) data feed for all equity and equity index options. OPRA gathers pricing data from all exchanges, physical and electronic, and disseminates firm bid and offer quotes, trade prices and related information in real-time. The NBBO represents the current inside spread in the consolidated national market for options. Exchanges typically designate one or more "primary" or "lead" marketmakers, who are required to quote continuous two-sided markets in reasonable size for the options they cover, and trades can always be executed against these posted bids and offers.[3]

The quoted NBBO bid and ask prices are a much better reflection of current option pricing than trades are. Because each underlying stock or index has puts and calls with many different exercise prices and expiration dates, option trading for even an extremely active index like the S&P 500 is relatively sparse, especially for contracts that are away from the money. However the NBBO is available and continuously updated at all points in time for all contracts that are currently being traded.[4]

---

[3] In the present case, due to a licensing agreement S&P 500 index options are only traded on the Chicago Board Options Exchange.

[4] A large proportion of the data flow consists of quotes for deep in the money contracts simply to keep them current as the underlying index fluctuates. Even though there is little or no trading in those options, the marketmakers must update their quotes constantly to avoid being "picked off," which means that the posted quotes on OPRA reflect their best judgment at every point in time as to the correct values for all options, regardless of trading volume.

The stock market opens at 9:30 A.M. S&P index options are traded in Chicago, but to avoid confusion we will always state time of day in terms of New York time. Options trading begins shortly after 9:30, but it can take several minutes before all contracts have opened and the market has settled into its normal mode of operation. To avoid introducing potentially anomalous prices at the beginning of the day from contracts that have not yet begun trading freely, we start the options "day" for our analysis at 10:00 A.M. We extract the NBBO's for all S&P 500 options of the chosen maturity from the OPRA feed and record them in a pricing tableau. The full set of current bids and offers for all strikes is maintained and updated whenever a new quote is posted. Every quote is assumed to remain a current firm price until it is updated. Our data set for analysis consists of snapshots of this real-time price tableau taken once every minute, from 10:00 A.M. until 4:00 P.M., leading to 361 observations of the RND per day.[5] The current index level is also reported in the OPRA feed, which provides a price series for the underlying that is synchronous with the options data.

S&P options are traded on a quarterly March-June-September-December cycle for more than a year into the future. There are also monthly expirations for the next few nearby "off" months. We concentrate on the quarterlies. These are European options, so no uncertainty is introduced by the possibility of early exercise. The risk neutral density that can be extracted from December options is therefore the market's risk-neutralized probability distribution for the index level at the market open on the third Friday of December.[6] Thus, the data consist of continuously updated quotes on options with a fixed maturity that telescopes downward over time. There are 8 Federal Open Market Committee (FOMC) meetings per year, so there are typically two Fed days for each option maturity, about 45 days apart.

Bid and ask implied volatilities needed for fitting the RND are computed using Merton's continuous dividend version of the Black-Scholes model. The riskless rate and dividend yield data required for this and for computing forward values for the index were provided by OptionMetrics.[7] U.S. dollar LIBOR is interpolated to match option maturity and converted into a continuously compounded rate. The projected dividends on the index are also converted to a continuous annual rate.

Table 1 shows the Federal Reserve's announced interest rate target for each of the 28 announcement dates in the sample. The time period covered is from May 3, 2005 through December 16, 2008. There were actually 30 announcement dates during this interval, but because of technical problems with the data feed, we are missing data for two of them. Those days are grayed out in the table.

[5] The market closes at 4:00 P.M. but there are often prices that come in for a few minutes after that. For consistency across sample days, we start the day at 10:00 A.M. (observation 31) and stop at 4:00 (observation 391), giving us 360 trading minutes per day.

[6] Following the convention adopted by OptionMetrics, the standard source for daily option prices, we treat this as if the contracts actually expired at the close on the Thursday before expiration Friday. For example, on Wednesday of expiration week, we would treat the contracts as having one day to expiration.

[7] See the OptionMetrics (2003) manual for a full description of how the interest rate and dividend yields are calculated.

The time period began with a series of 10 consecutive target increases of 25 basis points each. This was the tail end of a long series of such target rate increases, and from the commentary in the news media, by this time they were widely anticipated. There followed 9 FOMC meetings at which the target was unchanged, with the last one on August 7, 2007, just at the beginning of the current financial crisis. In fact, financial conditions deteriorated so rapidly after that, that the Fed felt it necessary to cut the target rate 25 basis points less than two weeks later, outside of a scheduled FOMC meeting. As the crisis deepened, the Fed cut the target aggressively, from 5.25% at the beginning of August 2007 to 2.00% at the end of April 2008. During January 2008, they reduced the target by 75 b.p. on Jan. 22, in advance of that month's FOMC meeting and then another 50 b.p. at the meeting itself. Because we want to examine the impact on the market of an announcement that was expected by traders and was released during trading hours, we do not include August 17, 2007 or January 22, 2008 in the sample.

The target was held at 2.00% for several successive meetings, but the full force of the crisis hit in September 2008 and the target was reduced again, all the way to the extraordinarily low level of "0 to 25 basis points," by the end of the sample.

The conventional wisdom is that a cut in the target rate signals looser monetary policy and easier credit conditions, which is good for business. The stock market is expected to respond favorably. A rise in the target rate indicates that the Fed is trying to moderate economic activity to hold down inflation, and the stock market is expected to fall. The changes in the Standard and Poor's 500 stock index on announcement days displayed in Table 1 show that this pattern is mostly, but not always, confirmed. There were 10 rate increases over this period. The stock market fell on 7 of those days and rose on 3. There were 8 rate cuts (including the two days that are not in our data), of which 4 produced stock price increases and 4 produced decreases. The target was left unchanged at 12 FOMC meetings, with 10 stock market advances and only 2 declines.

In theory, one should not focus on the full change in the target rate but only on the portion of the change that was unanticipated by the market: the "surprise." To the extent that the market has already incorporated the FOMC's predicted action into prices, an announcement that the Fed did what was expected provides little news and should have little impact on the stock market. By contrast, when the Fed fails to do what was expected, whether that involves holding the rate steady when the market expects a change, or changing the target but by a different amount than was expected, the market response when the information is released should be strong and immediate.

A good illustration of the interplay between expectations and the market's response occurred with the December 11, 2007 announcement. The Fed cut the target by 25 basis points at that meeting, but it was widely reported that the market had been expecting a cut of 50 b.p. and was disappointed with the smaller reduction The index had closed at 1515.96 the night before, and during the day prior to the announcement it rose 6.22 points to 1522.18 at 2:14 P.M. When the 25 b.p. reduction was announced, the index fell by 26.45 within the next 10 minutes. Then, as the market digested the information over the remainder of the day, the index kept falling, reaching 1477.65 by the close, down 38.31

on the day, including an extraordinary loss of -44.53 points after the 2:15 P.M. announcement (2.93% in less than two hours).

Bernanke and Kuttner (2005) attempted to measure the market's surprise at the announcement of a new target rate using the behavior of Fed funds futures prices, concluding that an unanticipated 25 basis point cut in the target rate was associated with about a 1% increase in stock prices. They found essentially no response to the anticipated part of the rate target. We will explore how well surprises calculated by their method can explain the stock market reactions in our sample.

The Chicago Mercantile Exchange trades a futures contract based on the daily average of the funds rate over each calendar month. On any given date, the futures price for the current month should embody the average rate the market expects for that month. If the Fed's new target is correctly anticipated in the Fed funds futures market, that rate will already be impounded in the futures price before the FOMC meeting and there should be no change on the announcement date. On the other hand, if a rate change is completely unexpected, the futures price should change by the amount by which the average changes when computed with the new rate for the remainder of the month.

For example, suppose that in a month with 30 days, the announcement date is the 16th, and the target rate prior to the FOMC meeting is 5.00 percent. If the market expects no change in the rate, Fed funds futures for that month should be based on an expected 30 days of 5.00 percent, so the futures price (defined as 100 - rate) would be 95.00. If the market expected a 25 b.p. cut, the expected average would be 15 days of 5.00 and 15 days of 4.75, and the futures price would be 100 - 4.875 = 95.125. In this case, if the market expects no rate cut but the Fed cuts the target by -25 b.p., the futures price should rise from 95.00 to 95.125 over the announcement day. From that price change, one can compute that the market's surprise was -25 b.p.

The last column in Table 1 shows target rate surprises computed in this fashion from the change in the Fed funds futures prices over the announcement dates.[8] Unfortunately, it is apparent that there is considerable noise in this calculation. In the early part of the sample, when Fed policy was easily predicted, there were many dates with no surprise. But when market volatility increased in the later portion of the sample, the connection seems much weaker between the surprise in the Fed funds market and the changes in the target rate and the stock market. For example, the surprise calculated for December 11, 2007 from Fed funds futures was only 0.7 basis point, but as just described, the stock market's response was much stronger. In March 2008, when the target was reduced by a remarkable 75 b.p., the computed surprise was +15.5 b.p., i.e., the market was "surprised"

[8] The actual calculation provided in Bernanke and Kuttner (2005) is a little more complicated than described here, to take better account of rate changes that occur at the very end of a month, such that there is almost no time remaining for the new rate to change the average for that month. Also, it was necessary to modify their formula slightly: the published version is based on the assumption that the new target is only available after the futures market has closed on the announcement day. Currently the announcement is at 2:15 Eastern time, and Fed funds futures remain open until 3 PM, so that the new information is incorporated in the market closing price on the same day.

that the cut was not larger. Yet the stock market rose 54.14 points, consistent with the sharp drop in rates being largely unanticipated.

In the end, a news announcement should move market prices only to the extent that the information was not already anticipated by the market, so it is important to explore that aspect of the impact of Fed funds target changes. However, in this case, the potential measure of the unexpected rate change extracted from the behavior of Fed funds futures appears to be a noisy and inadequate proxy for the surprise in the stock market. We will see below that the target rate surprises shown in Table 1 provide virtually no explanatory power for the stock market's response to the announcement.

V. Behavior of the Stock Market and the Risk Neutral Density on Fed Funds Target Announcement Days

In this section we examine how the interest rate target information released by the Fed at precisely 2:15 P.M. on date t affects price behavior over several time intervals. The periods we focus on here are:

- Full Day: from the close on date t-1 to the close on date t;
- Overnight: from the close on date t-1 to the beginning of the trading day, assumed to be 10:00 A.M., on date t;
- Pre-Announcement period: the trading hours from 10:00 A.M. through 2:14 P.M. on date t;
- Announcement Impact: a 10-minute period from 2:14 to 2:24 on date t, which we assume captures the market's initial response to the news;
- Re-Equilibration: the remainder of the trading day after the Impact period on date t, during which traders weigh the news and also the responses of the other investors as the market digests the information and incorporates it into final prices. To gain a more detailed picture of this process, in some of our tests we subdivide the re-equilibration period into 12-minute subintervals.

Because the RND is the density over the final value of the index at option expiration, the appropriate measure of stock prices for this study is the current forward index value for that date, defined as

$$F_t = S_t \, e^{(r_t - d_t)(T-t)} \tag{7}$$

where F_t is the date t forward value for the level of the index on date T, S_t is the current spot index, r_t is the riskless interest rate for the period from date t to date T and d_t is the annual dividend yield on the index.

The expected value of the index level at a given future date follows a martingale. At any point in time, if the market is in informational equilibrium the forward index should equal the mean of the risk neutral density, and both should be the market's risk neutralized forecast of the value of the S&P index on date T. And by the martingale property, they are also the market's expectation of the level of the forward index and the RND mean at the close on date t, and at every other instant up to option expiration, as well.

We will use the Federal Reserve's interest rate target announcements to examine how these two measures behave when the arrival of a major piece of news knocks the market out of equilibrium and it then re-equilibrates to a level that fully incorporates the new information.

Let F_{period} denote the level of the stock market at the end of some specified period on an announcement date t. For example, F_{impact} is the forward index at 2:24 P.M., the end of the announcement impact period. To keep notation simple, the dependence on the date t will be left implicit.

Consider the nature of the information available to investors at different points in time around the Fed's announcement. At 2:14 P.M., the end of the pre-announcement period

$$(8) \qquad F_{pre-announcement} = E\ [\ F_T \mid \Phi_{2:14}\ ] = E\ [\ F_{close} \mid \Phi_{2:14}\ ]$$

where $\Phi_{2:14}$ denotes the information available to the market just before the announcement. Assuming information entering the market over the next 10 minutes from news other than the Fed's announcement is negligible, the forward index after the initial impact is

$$(9) \qquad F_{impact} = E\ [\ F_{close} \mid \Phi_{2:14}\ ,\ \Phi_{News}\ ]$$

where Φ_{News} denotes the new information contained in the target rate announcement. Thus the change from $F_{pre-announcement}$ to F_{impact} measures how much the initial impact of the news alters the market's risk neutral expectation for the level of the S&P forward at the close, and at all future times up to expiration at date T. In an informationally efficient market, the size of the immediate price change should reflect how accurately the announcement was anticipated by the market. These changes may be large on any given date, but they should average out to zero over the long run.

We assume that this immediate news impact largely takes place in the first 10 minutes after the announcement. But the process of assimilating the new information fully into prices is not over at this point. The market's initial reaction gives an investor important additional information about how other investors are responding to the news. This will cause each individual investor to reevaluate his or her own first assessment of the announcement in light of the market's reaction.

As trading continues after the initial impact, the market will seek its new equilibrium. But this may take some time, since the market's behavior during the process continues to generate additional information that causes investors to keep revising their expectations. The market can not be expected to settle into a new equilibrium immediately.

For convenience in this discussion, we will continue to ignore the possible arrival of independent new information unrelated to the Fed announcement, under the assumption that it is likely to be of much lesser importance than the Fed news and uncorrelated with it, so that even with unrelated news, the equations below continue to hold in expected value.

During re-equilibration, at first the additional information flowing to investors is mainly from price movements. For some time τ_1 , not too long after the end of the impact period, we can write

$$(10) \qquad F_{\tau_1} = E\ [\ F_{close} \mid \Phi_{2:14}\ ,\ \Phi_{News},\ \{F_\tau\}_{\tau=impact}^{\tau=\tau_1}\ ]$$

to indicate that the market's price behavior following the announcement adds a significant element to investors' information set.

As re-equilibration progresses, new information continues to enter the process in the form of commentary and analysis of the announcement and the market's reaction to it. Traders will discuss the event with each other on the trading desk and with counterparts at other firms. Commentary will begin to appear on Bloomberg, CNBC, financial media websites, and other business news outlets.

To the extent that investors feel this "softer" information is useful, it will lead to further revision of expectations. By the time the market has fully impounded the Fed's news and achieved its new equilibrium, which we assume happens by the close of trading on date t, it will have reached a level that reflects the market's revised risk neutral expectation of the S&P index at option expiration, given the announcement, the market's reaction to the announcement, and the further public commentary and analysis in the media, denoted $\Phi_{Analysis}$:

$$(11) \qquad F_{close} = E\ [\ S_T \mid \Phi_{2:14}\ ,\ \Phi_{News},\ \{F_\tau\}_{\tau=impact}^{\tau=\tau_1},\ \Phi_{Analysis}\ ]$$

It is this process that we will examine using the behavior of the risk neutral density.

Figure 1 showed how the stock market behaved on December 11, 2007. To illustrate how the Fed announcement affected the RND on that day, Figures 2 and 3 plot it at several points in time. The left panel in Figure 2 shows the RND just before the Fed announced a 25 basis point cut in the funds rate target, which we have noted above turned out to be viewed as a disappointment by the market. The S&P index at this time was

1522.18 and the forward level was 1533.95, shown by the vertical gray line in the figure. The RND at this time had a very typical shape, with a distinct left-skew and no apparent bulges.[9]

The right panel shows how the RND changed over the next 10 minutes. The S&P index dropped sharply, so that by 2:24 P.M., the forward was 1507.30, 26.65 points lower than before the announcement. To highlight the change in the RND, the dashed lines replicate the left plot, before the announcement and the solid lines indicate the new RND and S&P forward as of 2:24 P.M. The density moved left throughout most of its range, with a hint of bulging out in the left tail. Earlier research has found that when the market drops sharply, the left side of the RND typically moves by more than the right side or the middle.[10] Here the median (the 50th percentile) fell by -24.72 points, from 1564.78 to 1540.06, while the 5th percentile dropped -57.01 points, from1213.29 to 1156.28 but the 95th percentile only went down by -16.60, from 1744.89 to 1728.29. Also, it is not uncommon for a temporary bulge to develop on the left side of the RND during times of severe market stress on the downside.

Figure 3 displays RNDs during the re-equilibration period. In the left plot, we see that between 2:24 and 3:00 P.M., the forward index only dropped from 1507.30 to 1502.90. The two RNDs are quite close to one another over most of the range, with the exception that the bulge in the left tail has grown larger. This deformation indicates that at this moment, the market was willing to pay higher prices for option contracts that would only pay off if the market experienced a huge drop over the next three months, down to below about 1100. Such contracts might be thought of as "disaster insurance." To make the same observation from a different perspective that is likely to have the stronger influence on option quotes, at this time marketmakers required markedly higher prices to write such contracts.

Finally, in the right panel of Figure 4, we see the evolution of the RND from 3:00 P.M. to the close at 4:00 P.M. During this interval the market continued its downward trajectory, with the forward falling another 13.96 points to close at 1488.94. This time, the right side of the RND showed little movement, while the left side changed considerably. Essentially, probability mass was transferred from the middle of the density to the less extreme part of its left side. One might say that by the close the left portion of the density had caught up with the bulge seen at 3:00 P.M.

In Table 2, we examine the intraday behavior of the S&P index forward and the RND on the 28 announcement days in our sample. On average, the forward index rose 5.70 points on these dates. The standard deviation of this price change across days was 17.93, making the t-statistic on the mean 1.68, barely significant at the 5% level for a one-tailed test. Much of the average increase occurred overnight, while the market was closed and

[9] Note that left skewness is not consistent with a lognormal density, the form the RND takes under Black-Scholes assumptions, or with a mixture of lognormals, as others have assumed in order to obtain a density with fatter tails. Figlewski (2009) finds that in 12 years of daily closing prices for S&P 500 options, the fitted risk neutral density was never skewed to the right as it would be under lognormality.

[10] See Figlewski (2009) or Birru and Figlewski (2010).

during the Pre-announcement period up to 2:14 P.M. on date t. As we expected for an efficient market, the response to the announcement during the 10-minute Impact period was close to zero on average, while the standard deviation was a much larger 7.61.

The standard deviation shows that the announcement had a large impact, but on average the market's expectation was unbiased. Following the announcement, the standard deviation of the change over the remainder of the day as the market returned to equilibrium was 12.03, more than twice what it was during the more than four hours prior to the announcement. This indicates that much of the market's adjustment to the news occurred after the initial impact period.

The next two lines provide the same statistics for the change in the RND mean. Because they are tied together by arbitrage, the RND mean should equal the forward level of the index, yet a certain amount of discrepancy between them is commonly observed, due to the transactions costs of executing the arbitrage trade. This raises the key issue: which one is the better measure of the "true" level of the stock market? Here we see that the average change over the day is nearly 30% smaller for the RND mean than the forward, which might suggest that the options market more correctly predicted the impact of the Fed announcements. However, that hypothesis is not supported when we look at behavior in the subperiods. During the trading day, including the impact period, the mean and standard deviation are virtually identical for changes in the forward and the RND mean. The only real difference between them occurs in the overnight period. At best, we can say that at the close on date t-1, the RND mean predicts its own value the next morning at 10:00 A.M. better than the forward does.

The next four lines take a closer look at the difference between the RND mean and the forward, which we have called the premium. The first pair of lines relates to the level of the premium and the second pair relates to its change. The table shows the premium at the end of each subperiod, except for the Full day column, which reports the closing value of the premium on date t-1.

On average, the RND mean was 0.96 above the forward at the close on the day before a Fed announcement. If the premium has predictive power, that is, if the RND mean is better than the concurrent forward as a forecast of their common future value, then a positive premium should lead to convergence in which the forward rises relative to the RND mean and the premium goes to 0. That is not what we see here. The average premium is positive at the close on date t-1 but it becomes negative overnight and remains so during the trading day. This behavior suggests overshooting by the RND mean, not convergence to equality with the forward index. Does this indicate significant information inefficiency in the risk neutral density? We will examine this puzzling phenomenon in greater detail below.

The RND provides information not just about the market's expected value for the index at option expiration, but also about the market's degree of certainty over that expectation, as manifested by the standard deviation of the density. This allows us to look at the resolution of uncertainty when the new interest target is announced. A good measure of how much information is provided by the Fed's announcement is the change in the RND variance, which we can compare to the information flow on a typical non-announcement day.

Let $\Delta F_t = F_t - F_{t-1}$ denote the one-day change in the S&P forward. Looking ahead from the close at date t-1, we can write

$$F_T = F_{t-1} + \sum_{\tau=t}^{T} \Delta F_\tau \tag{12}$$

Under the risk neutral density, the expected value of F_T given F_{t-1} is

$$E[\,F_T \mid F_{t-1}\,] = F_{t-1}$$

and its variance is

$$\begin{aligned} VAR\,[\,F_T \mid F_{t-1}\,] &= VAR\,[\,\sum_{\tau=t}^{T} \Delta F_\tau\,] \\ &= \sum_{\tau=t}^{T} VAR[\,\Delta F_\tau\,] + 2 \sum_{\tau_1=t}^{T-1} \sum_{\tau_2>\tau_1}^{T} COV\,[\,\Delta F_{\tau_1},\ \Delta F_{\tau_2}\,] \end{aligned} \tag{13}$$

Return shocks in an efficient market should all be serially uncorrelated, so the covariance terms in equation (13) should be zero in expectation and small in a finite sample, and the variance of the price change from date t-1 to T is the sum of the daily variances over that interval. This can be expressed as (T-t) times the average variance per day:

$$\begin{aligned} VAR\,[\,F_T \mid F_{t-1}\,] &\cong \sum_{\tau=t}^{T} VAR[\,\Delta F_\tau\,] \\ &= (T-t+1)\left(\sum_{\tau=t}^{T} VAR[\,\Delta F_\tau\,] \Big/ (T-t+1) \right) \end{aligned} \tag{14}$$

Dividing the change in the variance of the risk neutral distribution over the announcement day by the average risk neutral variance per day as of date t-1 gives a measure of how much the market's uncertainty was changed by the information release, expressed in terms of the number of average days' worth of variance that is impounded in the RND on date t-1. The next two lines in Table 2 report those values for the different time intervals and we see that a large amount of uncertainty was typically resolved by the target rate announcement: On average, over an announcement day the RND variance fell about 5 times as much as would be expected for an average day.

Surprisingly, however, nearly all of that reduction in RND variance occurred prior to the announcement. This could indicate leakage of the information into the market before its official release. The Fed takes considerable pains to prevent its rate decisions from leaking out before the official announcement. Nevertheless, some FOMC meetings are held on two consecutive days, and high-intensity media scrutiny after the first day's session might produce some information about the tone of the meeting and the general orientation of the committee members even without any significant security breach. Moreover, information about how the market might respond to a given announcement may well enter the market prior to the news release, in the form of orders placed with marketmakers, as well as softer information about what key players are expecting.

A key factor in how the market assimilates the Fed's announcement is that the market responds very differently to positive information and price increases than to negative news and price drops. It is well-established that both option-implied volatilities and the realized volatility of the stock index increase when stock prices fall. To examine positive and negative market reactions to announcements separately, the next two sets of lines in Table 2 report the average changes in the S&P forward, the RND mean, and the RND variance for the 14 announcement days when the market closed higher than on the previous day and the 14 "Market Down Days" when it closed lower.

On "up" days, the forward index rose an average of over 18 points and the reduction in risk neutral variance represents more resolution of uncertainty than would be expected over a normal week. On these days, a large fraction of the day's return occurred after the announcement, with about 11% in the initial impact and another 49% during the re-equilibration period. There was considerable reduction in the RND variance prior to the announcement, too, but on these days, the announcement itself caused RND variance to drop by the equivalent of 1.39 average days over the next 10 minutes, and a further 0.82 days' worth before the market closed.

By contrast, on days when the market went down after the announcement, over the full day RND variance only fell about a third as much as on up days and all of this reduction occurred before the announcement. The initial impact of the unfavorable news increased RND variance by about half a day's worth in the first 10 minutes and by another 1.08 days during re-equilibration. The index forward and the RND mean both rose on average during the pre-announcement period, as they had on market up days, but this was more than reversed after the announcement. Finally, we note that for down days the RND mean anticipated the drop by falling 3.08 index points overnight while the S&P forward was essentially flat. But with only 14 observations in this subsample, it is not clear how much to make of this result.

VI. The Announcement Day Response of the Stock Market and the Risk Neutral Density to Information Variables

In this section, we look at regressions of the market variables on the informational variables to explore more closely the connection between the news and the behavior of

the S&P forward and the RND before and after the announcement. The dependent variables are the changes in the S&P forward and the RND mean over the interval, the level of the premium at the beginning and its change over the interval.
We consider three information-based explanatory variables: the change in the target rate announced at 2:15 P.M. on date t, the "surprise" calculated from the behavior of the Fed funds rate on that date, as described above, and the premium of the RND mean over the S&P forward at the beginning of a given time interval. The first treats the entire change in the target rate as being unanticipated, while the second attempts to measure the unexpected portion directly. The premium variable should reveal any differences in the informational efficiency of the options market versus the stock market.

In an efficient market in continuous equilibrium, if there is no information leakage prior to a news release, regressing the change in the forward price or the risk neutral mean on the information should show no explanatory power over any interval before the announcement. When news is released, both variables react instantly and the market jumps to its new equilibrium with the information fully impounded in prices. Since neither variable would contain information about the news, any difference between them would be unrelated to the actual news release, so a regression of the premium on the news would show no effect, either. In all cases, once the news was released, there should be no further effect after the initial impact.

One would expect a different pattern if there were some news leakage before the official announcement. In the extreme case where all of the "news" was actually known to the market and impounded in prices before the close on date t-1, none of our regressions should have any explanatory power for price changes on date t, including at the time of the public announcement. If there is some, but not complete, information leakage the market would have partially adjusted to it ahead of time, but we would still expect a reaction at the time of the announcement and an immediate jump to the new equilibrium as in the case with no leakage. There would be correlation between the news to be released and market price movements in the periods when the information is leaking into the market, but none after the announcement.

If there is information leakage and options traders have different access to the leaks than traders in the stock market, or if they have different ability to predict the news release from publicly available information, we might also see that prices in one market contain more information about the impending announcement than prices in the other. In that case, the coefficient on the news in the better-informed market prior to the announcement would be larger and more significant, and the premium would be a valid predictor of the news that would be released later. For example, suppose options traders had better information than the rest of the market that the Fed was about to announce an increase in the rate target. Before the announcement, the RND mean would be below the S&P forward. When the announcement was made, the forward would adjust downwards towards the RND mean, and the negative premium would go to zero.

Table 3 reports results from four regression specifications run over the five time intervals in Table 2. The first section gives results for the full day from the close on date t-1 to the close on date t. The first line shows that an increase in the interest rate target had a large and significant negative effect on stock prices over the day, indicating that the Fed announcement was not fully anticipated by the market as of date t-1. A quarter percent increase in the target was associated with about a 6.6 point drop in the index.

The "surprise" extracted from Fed funds futures did not have useful explanatory power. The coefficient is not statistically significant and is of the wrong sign. (The positive coefficient would mean that an unexpected rate increase raises stock prices). The coefficient on the level of the premium at the close on date t-1 is also anomalous. If the options market was better informed than the stock market, the coefficient should be positive. If the stock market was better informed than the options market, the premium should have no power to explain changes in the forward and the coefficient should be 0. But a significantly negative coefficient means that if the RND mean was above the forward at the close on day t-1, the forward was likely to <u>fall</u> the next day.

The next regression, using the same variables to explain the change in the RND mean, shows results very similar to the regression on the forward. The coefficient on the change in the target is of the same size and statistical significance and the two anomalous coefficients on the other variables are also quite similar, except that the coefficient on the premium is even more negative and significant for the RND mean. These regressions indicate that the announced change in the rate target does affect both the forward and the RND mean, because the information is not fully impounded in market prices before the announcement date, and the R^2 statistics of 0.367 and 0.514 show they have fairly good explanatory power.

The next regression is of the level of the premium at the close on date t -1 on the change in the target. The coefficient has the correct negative sign, but is not close to statistically significant, indicating that the premium had little to no value as a forecast of the news release. The last regression in this section simply relates the change in the premium over the interval to its level at the start. The coefficient is highly significant and very close to -1.0. Whatever the premium was at the close on date t-1, on average it almost entirely disappears by the close on date t.

The next section of Table 3 covers the changes that occurred overnight, between the close on date t-1 and 10:00 A.M. on date t. Again the regressions on the forward and the RND mean are quite similar, except that the reversion of the premium toward 0 during this period provides additional explanatory power for the change in the RND mean, so the equation R^2 is considerably higher. The not-yet-announced target rate change does have a negative coefficient in both regressions, that is about one third as large as the coefficient for the full day and is borderline significant at the 5% level. This could be considered weak evidence that some information about the impending announcement entered the market overnight. The regressions on the level and change in the premium show the same thing as for the full day.

The forward and RND mean regressions for the pre-announcement period, from 10:00 A.M. to 2:14 P.M. on date t, again have negative coefficients on the rate change, but these are not significant. But the coefficient on the premium has now become positive and even significant at 5% in the forward equation, consistent with the premium containing information about which way the forward will move over the period. If this result is information-related, the information in question is entering the market before the announcement. This possibility is bolstered by the fact that the coefficient on the target change in the premium equation is now negative and highly significant. It appears that the premium at 10:00 A.M. does contain significant information about that afternoon's announcement and the fourth regression shows that, unlike the first two periods, only about a third of that premium disappears by the end of the pre-announcement interval.

The next set of results shows what happened during the 10-minute impact period immediately following the Fed announcement. It is very interesting that the time when an informationally efficient market should respond most strongly to whatever portion of the announcement was unanticipated, neither the stock market nor the options market registered a significant impact on average. The R^2 statistics on those equations are virtually 0.

This result is consistent with rational forecasting. Only the unanticipated part of a news release should move prices. If the market has formed an unbiased expectation of the information by 2:14 P.M. on announcement day, the surprise should be uncorrelated with the announcement itself. A drop in the target rate may be more or less than was expected, so a smaller than expected reduction on a day like Dec. 11, 2007 is treated by the market in the same way as an unexpected increase when no change was expected. The only small problem with this hypothesis is that the variable specifically designed to measure surprises is far from statistically significant, even though it does at least have the right sign here.

We favor the market efficiency hypothesis and feel that the surprise variable in the regression probably does not capture the actual portion of the rate change that was unanticipated by the stock market. While it seems quite possible that in a sample of this size, one or a few very big responses like Dec. 11, 2007 could cause this result, Table 1 showed many cases in which the stock market moved in the direction opposite to what the surprise variable indicated, not just one or two large ones.

One final intriguing result for the impact period is that the premium does have a negative and significant coefficient. On days when the target rate is raised (lowered), the premium just before the announcement is significantly negative (positive). Also, in this period, the mean-reverting tendency of the premium has disappeared.

Finally, in the re-equilibration period following the initial impact, the S&P forward and the RND mean regressions have large negative coefficients on the change in the target rate, but they are not significant. The surprise variable again has the wrong sign, but is not quite significant and the premium gets large negative insignificant coefficients.

To sum up the results in this section, we have seen that the stock market and the options market respond very strongly to the Federal Reserve's announcement, but much of the effect occurs prior to the actual news release. The surprise measure computed from the behavior of Fed funds futures and the premium of the RND mean over the forward index, two variables that were thought ex ante might carry information, proved to have little explanatory power. There did seem to be some correlation between the premium and the change in the target, but this did not carry over in a reasonable way to help explain the changes in the forward or the RND mean. Finally, although we hesitate to draw conclusions from coefficients that were not statistically significant, these results are not inconsistent with the hypothesis that a substantial portion of the market's adjustment to the information in the 2:15 P.M. announcement occurs in the re-equilibration period following the initial impact.

VII. Volatility of the Stock Market and the Risk Neutral Mean on Fed Funds Target Announcement Days

Table 2 compared price changes across 5 time intervals on Fed announcement days and showed that much of the day's activity occurred after the 2:15 P.M. announcement. When new information is released into an informationally efficient market, the price change should be zero on average and the results are consistent with that principle. The amount and importance of new information contained in a news release determines the size of the price impact, which can be measured by the standard deviation of the price change. When new information enters the market not just as a single event, but as a flow over time, an increased rate of information arrival should produce higher volatility but zero autocorrelation in price changes. Under full efficiency, the price is always equal to the true expected value conditional on current information, so even large information flows are completely impounded immediately, and price innovations will be serially uncorrelated even at very short time intervals. By contrast, if the market adjusts sluggishly to news, price changes will show positive autocorrelation, while if the market is too jumpy and regularly overshoots the correct valuation, autocorrelation will be negative.

We now examine the standard deviation and serial correlation for changes in the S&P index forward and the RND mean on Fed announcement days. For comparison we also examine a set of days when there was no announcement. The 11 Non-Announcement days, comprising a total of 3890 trading minutes, are listed in Table 4.

In Table 5 we consider two measures of volatility, the standard deviation across sample days of the price change in specified time intervals, and the average minute-to-minute price volatility within those intervals. Autocorrelation is also computed from one-minute price changes. The breakdown of the day into time periods is the same as above, and we add one column for the full Trading Day, in order to relate price behavior over subintervals to the average during the time the market was open. The top panel shows results for the 28 announcement days and the bottom panel is for the 11 non-announcement days.

The first two lines in each panel of Table 5 indicate little difference between the index forward and the RND mean on either announcement or non-announcement days. But the sizes of the price changes over different periods within the day was quite different when there was an announcement than on an ordinary day. The price change for the whole trading day was about 40% larger for the announcement days, but the change over the pre-announcement period was more than twice as big on non-announcement than announcement days. By contrast, the price change in the 10 minute impact period from 2:14 P.M. to 2:24 P.M. was more than 50% bigger than during the preceding 4 hours and 14 minutes, and three times bigger than during the same 10 minute period on non-announcement days. After the announcement, the price change over the remainder of the day, the re-equilibration period, was about 70% higher than on non-announcement days.

Within both groups, these comparisons combine results from days with potentially quite different levels of baseline volatility. To adjust for this in measuring the intraday pattern of price changes, the next two lines of each subsection report the ratio of the standard deviation of the change in each subinterval relative to the price change over that entire trading day. For example, relative to the whole day, the price change in the period before 2:14 P.M. is about three times bigger on non-announcement than announcement days.

The next two sets of lines in each subsection look at the volatility of 1-minute price changes within each time interval, in index points and as a ratio relative to volatility over the whole day. There were large volatility differences between subperiods on announcement days, but virtually none on non-announcement days. Interestingly, the strong time pattern in the forward price volatility across subperiods was also present for the RND means, but it was somewhat attenuated. The post-announcement period was still more volatile than the pre-announcement period, but the differences were smaller.

A particularly interesting difference between the forward index level and the RND mean is that, while they change by about the same amounts on the day, for both announcement and non-announcement days the RND mean is much more volatile than the index forward, and more so on non-announcement days. The large differences in autocorrelation account for this discrepancy. On announcement days autocorrelation for the index forward is 0.08 over the whole day, but it varies considerably across subintervals, going as high as 0.17 in the impact period. The non-announcement days also exhibit mild positive autocorrelation for the forward index. The RND mean, however, shows very strong negative autocorrelation in all periods on non-announcement days, and during the pre-announcement period on announcement days. We will postpone discussion of this seemingly anomalous result briefly while we turn our attention to a finer breakdown of the re-equilibration period in Table 6.

To get a more precise look at what is happening during re-equilibration following a Fed announcement, Table 6 breaks the period after 2:24 P.M. into 12-minute subintervals. It presents the same volatility measures and contrasts announcement and non-announcement days, as in Table 5. Table 5 showed that on announcement days the standard deviation of the price change over the whole period before 2:14 P.M., was 4.96

index points. Here we see that once the impact occurred, 3 of the 8 subsequent 12-minute re-equilibration intervals also had larger average price changes than in the full pre-announcement period and a fourth one was essentially equal. By contrast, on non-announcement days, none of these intervals after 2:14 P.M. had a price change standard deviation even half as large as before 2:14.

The pattern of minute-to-minute volatility on announcement days was even more striking: Prices in the impact period were 4 1/2 times more volatile than before the announcement, and every subsequent subinterval exhibited much greater volatility than in the pre-announcement period. This indicates that the increase in information flow associated with the announcement persists well after the actual news release. Moreover, there is a pattern of high volatility within the re-equilibration period that dies off over time.

To explore the time decay of volatility a little further, Equation (15) reports results of a regression in logs of the absolute value of price change in each minute relative to the volatility within the impact period, as a function of the number of minutes since the end of the impact period. Here τ indicates a minute within date t, with τ_0 representing 2:24 P.M., the end of the impact period. $\sigma_{Ft,impact}$ is the standard deviation of 1-minute changes in the index forward during the impact period on date t. t-statistics are shown in parentheses.

$$(15)\quad \log(\,|\,\Delta F_{t,\tau}\,|\,/\,\sigma_{Ft,impact}) \;=\; \underset{(-4.84)}{-0.562} \;+\; \underset{(-8.82)}{-0.276}\,\log(\tau - \tau_0)$$

$R^2 = 0.027$
NOBS = 2649

The regression confirms that the size of price changes declines over time during the re-equilibration period. Clearly the amount of unexplained variance of 1-minute price changes is quite high, as indicated by the low regression R^2, but the estimate of the coefficient of decay is highly significant.

Equation (16) reports the comparable regression to eq. (15) for the absolute changes in the RND mean over the re-equilibration period, where $\sigma_{RNDt,impact}$ is the standard deviation of 1-minute changes in the RND mean during the impact period.

$$(16)\quad \log(\,|\,\Delta RNDmean_{t,\tau}\,|\,/\,\sigma_{RNDt,impact}) \;=\; \underset{(-7.99)}{-0.749} \;+\; \underset{(-7.23)}{-0.183}\,\log(\tau - \tau_0)$$

$R^2 = 0.019$
NOBS = 2649

The coefficient of decay is again highly significant, but the regression indicates that while the size of the RND mean changes starts at a higher level, it decays more slowly over the re-equilibration period. The subinterval standard deviations in these subintervals on non-announcement days display no time pattern.

Returning to the autocorrelation results, Table 5 showed that for the forward index on announcement days it is 0.08 over the whole day, but it varies considerably across subintervals. Autocorrelation is essentially zero during the pre-announcement period; it becomes strongly positive during the impact and first re-equilibration periods, and it also becomes strongly positive at the end of the trading day. By contrast, for the non-announcement sample, the forward index exhibits virtually zero autocorrelation over the full day and, while there are large values both positive and negative in the re-equilibration period subintervals, there is no apparent pattern. The very last period of the day does show strong positive autocorrelation even greater than that in the comparable period on announcement days, but the subperiods immediately preceding the last one are quite different.

The RND mean behaves very differently from the forward index. On non-announcement days the minute-to-minute changes are highly negatively autocorrelated in every single subperiod, averaging -0.26 over the whole trading day. Autocorrelation on announcement days is actually less extreme, averaging -0.09 for the trading day, and even showing positive values during the impact and the subperiod immediately after.

This surprising result was found with a different sample of dates in Birru and Figlewski (2010). That paper focused on intraday behavior of the stock market during the financial "meltdown" in Fall 2008.[11] While strong nonzero serial correlation in returns could be evidence of a severely inefficient market, we feel that this is almost certainly an incorrect interpretation.[12] Rather, we believe it is a reflection of the marketmaking process.

Negative serial correlation in quote revisions could simply reflect risk management by marketmakers following large option trades by outside investors. Marketmakers try to hedge all of their "Greek letter" risk exposures, such as delta, gamma, and so on. When a large trade is taken into inventory, they adjust their bids and offers to try get back to market neutrality. The changes in the risk neutral density will be largely reversed and disappear as the marketmakers trade back to their preferred inventory positions, resulting in higher short term volatility for the RND mean than for the forward index and strong negative serial correlation in the series of quote revisions.

We examined whether this effect could account for the negative autocorrelation in the RND, but large transactions are simply not frequent enough for this to be the main

[11] The control sample analyzed here is drawn from the data sample for that paper, but only Sept. 23, 2008 comes from the "meltdown" period after the precipitating events involving Merrill Lynch, AIG, and Lehman Brothers.

[12] Errors in the price data would be another possible explanation for negative autocorrelation. A bad price quote will produce a deformation in the fitted RND that will be reversed when a correct quote for that option is posted, causing negative autocorrelation. Birru and Figlewski (2010) looked at this by examining autocorrelation measured over different time intervals. Serial correlation should disappear once the time step is long enough that a bad point is corrected within the same period. But strong negative autocorrelation was found to persist even when changes over 30 minute periods were considered.

explanation. Instead, the quotes themselves are posted and revised many times in between trades, with strong negative autocorrelation in their changes. This behavior is present in the quotes for both calls and puts at all strike levels throughout the day.

Table 7 reports the average quote to quote serial correlation for calls and puts of all degrees of "moneyness" during the intraday periods we have been examining The option data is drawn from days with quite different volatility of the underlying index and time to option expiration. For comparability, we put the options into moneyness buckets defined relative to the risk neutral density at 10:00 A.M. For example, on a day when the RND median at 10:00 A.M. is 1100, both the 1100-strike call and put would go into the 45-55% bucket. If the 30th percentile of the RND was 980, the 980-strike options would be placed in the 25-35% bucket. The risk neutral probability of ending up in the money at expiration is 70% for the call and 30% for the put, independent of volatility and option maturity.

The table shows clearly the strong and pervasive negative serial correlation for the quote revisions in this market. Although a full analysis of the phenomenon is beyond the scope of this paper, we believe that in this case it is a result of market convention in the way in which quotes are revised.

Options are priced relative to the underlying using a model, so their theoretical values vary continuously with the index. But bid-ask spreads are relatively wide, and even though it is essentially costless to post fresh quotes, by convention they are only changed by relatively large increments. This can lead to price bands for the index in which it moves freely without inducing changes in the option quotes, separated by break points where the quotes repeatedly flip back and forth as the index fluctuates above and below the critical level, generating a pattern of extended periods with no quote revisions alternating with periods of strongly negatively autocorrelated changes.[13]

It is not clear why option marketmakers post quotes in this way, but the data clearly show that revisions are consistently in much larger increments than the minimum tick size allowed in the market. How much autocorrelation this quoting practice induces should depend on the volatility of the underlying, the bid-ask spread, the size of the conventional quote increment, the number of competing dealers who might revise their quotes at different index points, option trading volume, and possibly other factors. In the case of the S&P 500 index, many options are very deep in the money, with almost no trading activity and wide bid-ask spreads, and many are deep out of the money, with narrow

[13] A simple numerical example will show how this behavior can lead to strong negative autocorrelation. Suppose the current index is 1100, the 980-strike call is 136 bid, 138 ask. Focusing just on the bid, if the index goes up, there will be a point, say when the index hits 1101, at which the bid will be raised to 136.50. At a higher index level, say 1103, the bid is raised to 137. So, while the index fluctuates between 1100.00 and 1100.99 the option quotes are constant (even though the option's theoretical value moves with every change in the index). But when the index gets to 1101, every small tick up or down that causes it to go through 1101, will produce a 0.50 change in the option quotes, leading to a string of plus and minus 0.50 quote revisions. This will continue until the index moves into the middle of the range, away from the break points.

spreads in terms of the option price, but with very wide break points in terms of the index due to their very small deltas.

Table 7 shows considerable variability in option quote autocorrelation across moneyness and in different periods during the announcement day. It is strongest for deep out of the money options (high percentiles for the calls and low percentiles for the puts), and before the announcement. The latter is consistent with the low rate of information flow and index volatility during this period. By contrast, quote autocorrelation, while still negative, is greatly reduced during the impact period. Recall that the RND mean showed positive minute-to-minute serial correlation during the impact and first re-equilibration periods, suggesting that the much greater flow of information into the market dominated the effect on bid-ask quotes just discussed. As the re-equilibration progressed, autocorrelation in the different exercise price buckets returned toward its average level.

For comparison, Table 8 presents mid-quote autocorrelations for the same time periods on non-announcement days. The pattern across strikes is similar, with the deep out of the money contracts showing the most autocorrelation. Average values for the whole day are roughly comparable, but the time pattern around 2:15 P.M. and immediately after that is strikingly different. In contrast to Table 7, no particular pattern is visible for the non-announcement days. This lends support to the idea that the market's distinctly lower level of minute-to-minute autocorrelation in the RND mean following the Fed's interest rate target announcement reflects real revisions of the market's expectations as it digests the news.

VII. Concluding Comments

This is the first study using the risk neutral density extracted from real-time prices of S&P 500 index options to explore in fine detail how the public release of important information is incorporated into stock prices. The RND provides a direct measure of both the market's risk neutral expected value for the future level of the stock market and also the market's degree of uncertainty around that expected value. We found that on a day when a new target for the Fed funds rate is announced, resolution of uncertainty is about five times greater than on an ordinary day. We also found that uncertainty went down substantially on average, whether the market's response to the announcement was positive or negative, but the reduction was roughly three times larger when the market went up after the announcement than when it went down.

Our results strongly confirmed that the Fed announcement is important information to the market, in that it moves prices substantially (about 7 1/2 S&P points on average during the 10 minutes following the announcement, and about 18 points for the whole day). There was some evidence of information leakage into the market before the announcement. This could come about if there were security breaches around the Fed Open Market Committee meeting, or it could be that information about what rate the market was expecting and how it would respond to the new target became more available through order flow and other clues as the announcement approached.

The evidence largely supported the hypothesis that the market's expectation about this event is rational. Although the announcement does have a large impact on stock prices, the average change was approximately zero, meaning that the S&P index before the announcement was an unbiased forecast of its level afterwards. But we also saw strong evidence that fully adjusting to the information contained in the news release and in the market's path back to equilibrium afterwards is not instantaneous. Much of the total price change on the date occurs well after the initial impact of the announcement. This conclusion is supported by the fact that minute to minute volatility is much higher after the announcement than before, and it gradually tapers off during the re-equilibration period.

As is common when a new data set is analyzed with a relatively new technology, our investigation brought out a few puzzles that suggest directions for future research. One surprise was the degree to which adjustment of the market to the new target rate occurs before the announcement. Another was the anomalous behavior of the risk neutral mean minus the forward index, which we have called the premium. In some periods, the results indicated that if the RND mean was higher than the forward, both were likely to go down. Yet, the premium did appear to contain relevant information in the pre-announcement period. One minor surprise was that the "Surprise" variable extracted from the behavior of Fed funds futures prices on announcement day did not seem to contain useful information when examined at close range with intraday data, even though Bernanke and Kuttner (2005) had found it to be useful in an earlier time period with daily data. Finally, the strong autocorrelation in the RND mean on both announcement and non-announcement days is an important oddity that remains to be explained. We offered a hypothesis and some suggestive evidence that negative autocorrelation could be understood as a result of the options market making process, but rigorous examination of this idea is left for future research.

References

Alonso, Francisco, Roberto Blanco, and Gonzalo Rubio (2005). "Testing the Forecasting Performance of Ibex 35 Option-Implied Risk-Neutral Densities." Working Paper Series, No. 0505, Banco de Espana.

Andersson, Magnus (2007). "Using Intraday Data to Gauge Financial Market Responses to Fed and ECB Monetary Policy Decisions," European Central Bank Working Paper, No. 726.

Basistha, Arabinda and Alexander Kurov (2008). "Macroeconomic cycles and the Stock Market's Reaction to Monetary Policy." Journal of Banking and Finance 32, 2606-2616.

Bates, D. (1991). "The Crash of '87: Was it Expected? The Evidence from Options Markets." Journal of Finance 46, 1009-1044.

Bates, D. (2000). "Post-'87 Crash Fears in the S&P 500 Futures Option Market." Journal of Econometrics 94, 181-238.

Beber, Alessandro and Michael W. Brandt (2006). "The Effect of Macroeconomic News on Beliefs and Preferences: Evidence from the Options Market." Journal of Monetary Economics 53, 1997-2039.

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

Brenner, Menachem, Pasquariello, Paolo and Marti Subrahmanyam (2009). "On the Volatility and Comovement of U.S. Financial Markets Around Macroeconomic News Announcements." Journal of Financial and Quantitative Analysis 44, 1265-1289.

Birru, Justin and Stephen Figlewski (2010). "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008." Working paper, New York University.

Breeden, Douglas and Robert Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Option Prices." Journal of Business 51, 621-652.

Chen, En-Te and Adam Clements (2007). S&P 500 Implied Volatility and Monetary Policy Announcements." Finance Research Letters 4, 227-232.

Chulia-Soler, H, M. Martens, and D. van Dijk (2007). "The Effects of Fed Funds Target Rate Changes on S&P 100 Stock Returns, Volatilities and Correlations." Erasmus Research Institute of Management Report Series.

Davig, Troy and Jeffrey R. Gerlach (2006). "State-Dependent Stock Market Reactions to Monetary Policy." International Journal of Central Banking 2, 65-84.

Fama, Eugene (1965). "The Behavior of Stock Market Prices." Journal of Business 38, 34-105.

Fama, Eugene, Lawrence Fisher, Michael Jensen, and Richard Roll (1969). "The Adjustment of Stock Prices to New Information." International Economic Review 10, 1-21.

Farka, Mira (2009). “The Effect of Monetary Policy Shocks on Stock Prices Accounting for Endogeneity and Omitted Variable Biases.” Review of Financial Economics 18, 47-55.

Figlewski, Stephen (2009). "Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio," in Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle (eds. Tim Bollerslev, Jeffrey R. Russell and Mark Watson). Oxford, UK: Oxford University Press.

Gemmill, G. and A. Saflekos (2000). "How Useful are Implied Distributions? Evidence from Stock-Index Options." Journal of Derivatives 7, 83-98.

Harrison, J. Michael and Kreps, David M. (1979). "Martingales and arbitrage in multiperiod securities markets," Journal of Economic Theory, 20(3), 381-408.

Hausman, Joshua and Jon Wongswan (2006). “Global Asset Prices and FOMC Announcements.”FRB International Finance Discussion Paper No. 886.

Jackwerth, J.C. (2004). Option-Implied Risk-Neutral Distributions and Risk Aversion. Charlotteville: Research Foundation of AIMR.

Jackwerth, J.C. and Mark Rubinstein (1996). "Recovering Probability Distributions from Option Prices." Journal of Finance 51, 1611-1631.

Lynch, Damien and Nikolaos Panigirtzoglou (2008). "Summary Statistics of Option-Implied Probability Density Functions and their Properties." Bank of England Working Paper No. 345.

Nikkinen, Jussi and Petri Sahlstrom (2004). “Impact of the Federal Open Market Committee’s Meetings and Scheduled Macroeconomic News on Stock Market Uncertainty.” International Review of Financial Analysis 13, 1-12.

Nofsinger, J. and B. Prucyk (2003). "Option Volume and Volatility Response to Scheduled Economic News Releases." Journal of Futures Markets 23, 315-345.

Poon, Ser-Huang and Clive W. J. Granger (2003). " Forecasting Volatility in Financial Markets: A Review" Journal of Economic Literature 41, pp. 478-539.

Samuelson, Paul (1965). "Proof the Properly Anticipated Prices Fluctuate Randomly." Industrial Management Review, Spring, 41-49.

Zebedee, A.A., E. Bentzen, P.R. Hansen and A. Lunde (2008). “The Greenspan Effect on Equity Markets: An Intraday Examination of US Monetary Policy Announcements.” Financial Markets and Portfolio Management 22, 3-20.

Table 1
Federal Reserve Interest Rate Target Announcements

Note: The table reports the federal funds rate target announced at 2:15 P.M. on the specified date, the change from the previous rate target, the change of the S&P 500 stock index from the close on date (t-1) to date t, and the surprise component of the rate change computed from the change in the current Fed Funds futures prices, as described in the text. Grayed out dates are not included in the data set.

Date	Target Rate	Change	S&P 500 Index Change	Fed Funds Futures "Surprise"
5/3/2005	3	0.25	-0.99	0
6/30/2005	3.25	0.25	-8.52	0
8/9/2005	3.5	0.25	8.25	0
9/20/2005	3.75	0.25	-9.68	0.014
11/1/2005	4	0.25	-4.25	0.225
12/13/2005	4.25	0.25	7.00	0
1/31/2006	4.5	0.25	-5.11	0
3/28/2006	4.75	0.25	-8.38	0
5/10/2006	5	0.25	-2.29	-0.007
6/29/2006	5.25	0.25	26.87	-0.015
8/8/2006	5.25	0	-4.29	-0.039
9/20/2006	5.25	0	7.54	0
10/25/2006	5.25	0	4.84	0
12/12/2006	5.25	0	-1.48	0
1/31/2007	5.25	0	9.42	0
3/21/2007	5.25	0	24.10	0
5/9/2007	5.25	0	4.86	0
6/28/2007	5.25	0	-0.63	0
8/7/2007	5.25	0	9.04	0.025
9/18/2007	4.75	-0.5	43.13	-0.138
10/31/2007	4.5	-0.25	18.36	-0.020
12/11/2007	4.25	-0.25	-38.31	0.007
1/30/2008	3	-0.5	-6.49	-0.095
3/18/2008	2.25	-0.75	54.14	0.155
4/30/2008	2	-0.25	-5.35	-0.055
6/25/2008	2	0	7.68	-0.025
8/5/2008	2	0	35.87	-0.006
9/16/2008	2	0	20.90	0.056
10/29/2008	1	-0.5	-10.42	-0.060
12/16/2008	0.125*	-0.875	44.61	-0.110
Target Rate Changes Outside of Scheduled FOMC Meetings				
8/17/2007	5	-0.25	34.67	0.145
1/22/2008	3.5	-0.75	-14.69	-0.667

Table 2
Levels and Changes of Key Variables on Fed Announcement Days

Note: The table reports the means and standard deviations of changes in key variables from the close on (t-1) to the close on date t and over subperiods. The Forward is computed from the spot S&P 500 index increased by the riskless interest rate less the dividend yield to the option expiration date. The RND mean and variance are computed from the estimated risk neutral density. The Premium is the RND mean minus the Forward. The relative change in the RND variance is the change over the interval divided by the RND variance per day over the remaining life of the options.

		Full day: Close date t-1 to Close date t	Overnight: Close date t-1 to 10:00 AM date t	Pre-announcement: 10:00 AM to 2:14 PM	Announcement impact: 2:14 PM to 2:24 PM	Re-equilibration: 2:24 PM to Close date t
Change in Forward S&P	**mean**	5.70	2.12	2.79	-0.63	1.42
	std dev	17.93	8.46	4.96	7.61	12.03
Change in RND mean	**mean**	4.02	0.52	2.70	-0.66	1.45
	std dev	20.07	9.31	4.95	7.47	12.23
Premium	**mean**	0.96	-0.64	-0.73	-0.76	-0.72
	std dev	5.25	0.94	0.95	1.13	1.20
Change in Premium	**mean**	-1.68	-1.60	-0.08	-0.03	0.03
	std dev	5.33	5.29	0.46	0.49	0.94
Relative change in RND variance	**mean**	-4.97	-3.51	-1.16	-0.43	0.13
	std dev	9.10	5.50	2.18	3.04	4.65
Market Up after Announcement						
	change in S&P forward	18.37	3.94	3.36	1.99	9.07
	change in RND mean	18.36	4.13	3.12	1.93	9.17
	change in RND variance	-7.50	-4.00	-1.29	-1.39	-0.82
Market Down after Announcement						
	change in S&P forward	-6.97	0.30	2.21	-3.25	-6.23
	change in RND mean	-10.32	-3.08	2.28	-3.25	-6.27
	change in RND variance	-2.45	-3.03	-1.02	0.52	1.08

Table 3
Regressions of Changes in Market Variables on Information Variables by Interval

Notes: The table reports results of regressions on changes of information-related variables over the full announcement days and subperiods. The Premium is the RND mean minus the Forward, and the Surprise is computed from the change in the current Fed Funds futures prices, as described in the text. Each regression has 28 observations and t-statistics are shown in parentheses.

Dependent Variable	Constant	Change in Target Rate	Surprise	Premium	R-squared
FULL DAY					
Change in S&P Forward	6.258	-26.42	64.89	-1.373	0.367
	(2.15)	(-2.71)	(1.38)	(-2.46)	
Change in RND Mean	5.482	-27.32	65.66	-2.341	0.514
	(1.86)	(-2.76)	(1.37)	(-4.14)	
Premium	0.907	-2.18			0.016
	(0.88)	(-0.64)			
Change in Premium	-0.759			-0.962	0.965
	(-3.94)			(-26.23)	
OVERNIGHT					
Change in S&P Forward	1.762	-8.92	25.45	0.100	0.123
	(1.07)	(-1.61)	(0.95)	(0.32)	
Change in RND Mean	1.077	-9.92	25.04	-0.876	0.319
	(0.65)	(-1.79)	(0.94)	(-2.77)	
Premium	0.907	-2.18			0.016
	(0.88)	(-0.64)			
Change in Premium	-0.671			-0.968	0.982
	(-4.89)			(-37.07)	
PRE-ANNOUNCEMENT					
Change in S&P Forward	4.010	-3.70	0.93	2.041	0.277
	(3.67)	(-1.32)	(0.08)	(1.68)	
Change in RND Mean	3.699	-3.79	0.58	1.691	0.227
	(3.25)	(-1.29)	(0.05)	(1.34)	
Premium	-0.665	-1.07			0.223
	(-5.42)	(-2.66)			
Change in Premium	-0.294			-0.328	0.244
	(-2.69)			(-2.82)	

Table 3, continued

Dependent Variable	Constant	Change in Target Rate	Surprise	Premium	R-squared
IMPACT					
Change in S&P Forward	-1.484 (-0.56)	-0.85 (-0.15)	-10.59 (-0.40)	-1.191 (-0.42)	0.014
Change in RND Mean	-1.392 (-0.54)	-0.78 (-0.14)	-6.88 (-0.27)	-1.009 (-0.36)	0.008
Premium	-0.743 (-6.57)	-0.80 (-2.15)			0.158
Change in Premium	0.045 (0.30)			0.107 (0.66)	0.018
RE-EQUILIBRATION					
Change in S&P Forward	-0.121 (-0.04)	-10.19 (-1.21)	61.66 (1.54)	-1.530 (-0.51)	0.118
Change in RND Mean	-0.486 (-0.14)	-10.75 (-1.25)	60.86 (1.50)	-2.043 (-0.67)	0.118
Premium	-0.773 (-4.74)	-0.70 (-1.30)			0.064
Change in Premium	-0.321 (-1.38)			-0.470 (-2.28)	0.175

Table 4

Sample of Non-Announcement Days

Note: The table shows the dates and the S&P Index level and change for the 11 non-announcement date analyzed as a control sample.

	S&P Index Level	S&P Index Change
9/13/2006	1318.07	5.07
9/18/2006	1321.18	1.52
9/21/2006	1318.03	-7.15
10/4/2006	1350.2	16.09
10/18/2006	1365.80	1.75
9/20/2007	1518.75	-10.28
9/25/2007	1517.21	-0.52
10/10/2007	1562.47	-2.68
10/24/2007	1515.88	-3.71
9/10/2008	1232.04	7.53
9/23/2008	1188.22	-18.87

Table 5
Intraday Volatility of the Forward S&P Index and the RND Mean on Announcement and Non-Announcement Days

Notes: The table reports several volatility-related variables over the full announcement day and subperiods. The standard deviation of the change over the full interval is the standard deviation across dates of the change from the beginning to the end of each period, in index points. Interval std dev relative to the full trading day expresses the standard deviation in index points as a fraction of the full change during trading hours. The standard deviation of 1-minute changes is the minute-to-minute volatility within the period, and Interval 1-minute std dev relative to full day is the minute-to-minute volatility in the interval divided by that volatility over the whole trading day.

		Full Day	Overnight	Trading Day	Pre-Announcement	Impact	Re-Equilibration
	From	Close(t-1)	Close(t-1)	10:00 AM	10:00 AM	2:14 P.M.	2:24 P.M.
	To	Close(t)	10:00AM	4:00 PM	2:14 PM	2:24 P.M.	4:00 P.M.
Announcement Days							
Std dev of change over full interval	**S&P forward**	17.93	8.46	16.05	4.96	7.61	12.03
	RND mean	20.07	9.31	16.14	4.95	7.47	12.23
Interval std dev relative to full trading day	**S&P forward**	1.12	0.53	1.00	0.31	0.47	0.75
	RND mean	1.24	0.58	1.00	0.31	0.46	0.76
Std dev of 1-minute changes in interval	**S&P forward**	-	-	0.67	0.39	1.84	0.91
	RND mean	-	-	0.95	0.74	2.01	1.14
Interval 1-minute std dev relative to full day	**S&P forward**	-	-	1.00	0.58	2.76	1.37
	RND mean	-	-	1.00	0.78	2.11	1.20
Autocorrelation of 1-minute changes	**S&P forward**	-	-	0.08	0.00	0.17	0.06
	RND mean	-	-	-0.09	-0.21	0.10	-0.06
Non-Announcement Days							
Std dev of change over full interval	**S&P forward**	9.01	6.24	11.19	10.23	2.25	7.21
	RND mean	9.19	6.64	11.20	10.17	2.31	7.13
Interval std dev relative to full trading day	**S&P forward**	0.81	0.56	1.00	0.91	0.20	0.64
	RND mean	0.82	0.59	1.00	0.91	0.21	0.64
Std dev of 1-minute changes in interval	**S&P forward**	-	-	0.52	0.50	0.55	0.56
	RND mean	-	-	1.16	1.15	1.13	1.15
Interval 1-minute std dev relative to full day	**S&P forward**	-	-	1.00	0.96	1.07	1.08
	RND mean	-	-	1.00	1.00	0.97	1.00
Autocorrelation of 1-minute changes	**S&P forward**	-	-	0.02	0.01	0.08	0.04
	RND mean	-	-	-0.26	-0.27	-0.22	-0.24

Table 6
Volatility of the Forward S&P Index and the RND Mean during the Re-Equilibration Period, by Sub-Intervals

Notes: See the notes to Table 5. In this table, the Re-Equilibration period is broken down into 12-minute subperiods, within which the calculations are done as described in Table 5 for longer periods. The Full Re-Equil. column reports the values for the whole interval from 2:24 P.M. through 4:00 P.M.

		Impact	Re-Equilibration								Full Re-Equil
	From	2:14 P.M.	2:24 P.M.	2:36 P.M.	2:48 P.M.	3:00 P.M.	3:12 P.M.	3:24 P.M.	3:36 P.M.	3:48 P.M.	2:24 P.M.
	To	2:24 P.M.	2:36 P.M.	2:48 P.M.	3:00 P.M.	3:12 P.M.	3:24 P.M.	3:36 P.M.	3:48 P.M.	4:00 P.M.	4:00 P.M.
Announcement Days											
Std dev of change over full interval	**S&P forward**	7.61	3.72	5.69	4.89	2.84	5.51	3.01	3.94	8.61	12.03
	RND mean	7.47	3.83	5.88	5.21	2.89	5.57	3.04	4.22	8.44	12.23
Interval std dev relative to full trading day	**S&P forward**	0.47	0.23	0.35	0.30	0.18	0.34	0.19	0.25	0.54	0.75
	RND mean	0.46	0.24	0.36	0.32	0.18	0.34	0.19	0.26	0.52	0.76
Std dev of 1-minute changes in interval	**S&P forward**	1.84	1.07	1.10	0.88	0.77	0.77	0.72	0.72	0.75	0.91
	RND mean	2.01	1.18	1.31	1.07	1.06	0.96	1.01	0.99	1.03	1.14
Interval 1-minute std dev relative to full day	**S&P forward**	2.76	1.61	1.65	1.32	1.15	1.15	1.09	1.07	1.12	1.37
	RND mean	2.11	1.24	1.38	1.13	1.12	1.01	1.07	1.04	1.09	1.20
Autocorrelation of 1-minute changes	**S&P forward**	0.17	0.14	0.03	-0.02	0.03	0.05	-0.02	0.12	0.12	0.06
	RND mean	0.10	0.08	-0.01	-0.07	-0.03	-0.06	-0.13	-0.05	-0.01	-0.06
Non-Announcement Days											
Std dev of change over full interval	**S&P forward**	2.25	2.54	1.86	1.96	2.39	1.22	3.69	2.55	2.35	7.21
	RND mean	2.31	2.42	2.14	2.12	2.62	1.07	4.53	2.29	2.96	7.13
Interval std dev relative to full trading day	**S&P forward**	0.20	0.23	0.17	0.18	0.21	0.11	0.33	0.23	0.21	0.64
	RND mean	0.21	0.22	0.19	0.19	0.23	0.10	0.40	0.20	0.26	0.64
Std dev of 1-minute changes in interval	**S&P forward**	0.55	0.58	0.52	0.48	0.55	0.54	0.57	0.59	0.52	0.56
	RND mean	1.13	1.24	0.90	0.91	1.29	1.05	1.30	1.03	1.06	1.15
Interval 1-minute std dev relative to full day	**S&P forward**	1.07	1.13	1.00	0.94	1.06	1.05	1.10	1.14	1.00	1.08
	RND mean	0.97	1.07	0.78	0.79	1.12	0.91	1.12	0.89	0.91	1.00
Autocorrelation of 1-minute changes	**S&P forward**	0.08	0.00	0.06	0.03	0.12	0.03	0.06	-0.05	0.20	0.04
	RND mean	-0.22	-0.18	-0.12	-0.19	-0.26	-0.25	-0.16	-0.14	-0.06	-0.24

Table 7
Autocorrelation in Intraday Index Option Mid-Quote Changes on Federal Reserve Interest Target Announcement Days

Notes: The table shows serial correlation coefficients for changes in the bid-ask midpoint during subintervals of the trading day. Weighted averages are reported in percentile buckets based on where the option strike lies relative to the 10:00 A.M. Risk Neutral Density. Weights are proportional to the number of quotes reported for each option strike within the specified interval.

	Trading Day	Pre-Announcement	Impact	Re-equilibration subintervals								Full Re-equilibration
	10:00 AM 4:00 PM	10:00 AM 2:14 PM	2:14 PM 2:24 PM	2:24 PM 2:36 PM	2:36 PM 2:48 PM	2:48 PM 3:00 PM	3:00 PM 3:12 PM	3:12 PM 3:24 PM	3:24 PM 3:36 PM	3:36 PM 3:48 PM	3:48 PM 4:00 PM	2:24 PM 4:00 PM
Calls												
All	-0.363	-0.488	-0.167	-0.239	-0.253	-0.272	-0.343	-0.311	-0.339	-0.375	-0.382	-0.303
0-5%	-0.446	-0.562	-0.204	-0.307	-0.354	-0.373	-0.441	-0.406	-0.443	-0.502	-0.520	-0.405
5-15%	-0.299	-0.430	-0.122	-0.192	-0.194	-0.224	-0.276	-0.203	-0.262	-0.323	-0.313	-0.241
15-25%	-0.260	-0.403	-0.106	-0.133	-0.161	-0.169	-0.214	-0.202	-0.227	-0.267	-0.247	-0.186
25-35%	-0.302	-0.445	-0.119	-0.182	-0.141	-0.159	-0.248	-0.241	-0.265	-0.316	-0.320	-0.219
35-45%	-0.308	-0.429	-0.152	-0.195	-0.152	-0.176	-0.283	-0.259	-0.291	-0.305	-0.328	-0.234
45-55%	-0.292	-0.422	-0.134	-0.162	-0.151	-0.206	-0.269	-0.252	-0.294	-0.293	-0.295	-0.217
55-65%	-0.312	-0.418	-0.174	-0.204	-0.177	-0.223	-0.313	-0.316	-0.301	-0.298	-0.309	-0.247
65-75%	-0.383	-0.498	-0.219	-0.270	-0.245	-0.317	-0.371	-0.368	-0.341	-0.237	-0.286	-0.304
75-85%	-0.405	-0.539	-0.216	-0.303	-0.251	-0.347	-0.397	-0.397	-0.300	-0.243	-0.311	-0.314
85-95%	-0.470	-0.574	-0.266	-0.351	-0.386	-0.393	-0.470	-0.480	-0.410	-0.356	-0.347	-0.408
95-100%	-0.588	-0.656	-0.570	-0.651	-0.508	-0.545	-0.693	-0.696	-0.574	-0.547	-0.614	-0.523
Puts												
All	-0.272	-0.378	-0.135	-0.198	-0.172	-0.206	-0.254	-0.226	-0.248	-0.293	-0.285	-0.222
0-5%	-0.712	-0.798	-0.504	-0.640	-0.664	-0.706	-0.723	-0.716	-0.743	-0.765	-0.763	-0.665
5-15%	-0.536	-0.651	-0.278	-0.409	-0.443	-0.493	-0.501	-0.463	-0.493	-0.540	-0.465	-0.471
15-25%	-0.414	-0.542	-0.177	-0.299	-0.326	-0.338	-0.413	-0.351	-0.359	-0.400	-0.362	-0.347
25-35%	-0.333	-0.464	-0.150	-0.216	-0.213	-0.208	-0.303	-0.274	-0.285	-0.298	-0.306	-0.250
35-45%	-0.311	-0.439	-0.151	-0.186	-0.158	-0.196	-0.267	-0.253	-0.267	-0.272	-0.313	-0.222
45-55%	-0.271	-0.397	-0.118	-0.161	-0.140	-0.176	-0.234	-0.231	-0.269	-0.287	-0.274	-0.201
55-65%	-0.256	-0.378	-0.101	-0.152	-0.120	-0.180	-0.228	-0.219	-0.249	-0.278	-0.260	-0.185
65-75%	-0.234	-0.351	-0.094	-0.140	-0.102	-0.159	-0.207	-0.213	-0.243	-0.271	-0.257	-0.170
75-85%	-0.233	-0.364	-0.077	-0.135	-0.117	-0.153	-0.198	-0.181	-0.215	-0.260	-0.250	-0.168
85-95%	-0.222	-0.337	-0.095	-0.181	-0.122	-0.170	-0.225	-0.172	-0.207	-0.252	-0.254	-0.192
95-100%	-0.229	-0.329	-0.115	-0.165	-0.138	-0.179	-0.204	-0.171	-0.194	-0.265	-0.262	-0.186

Table 8
Autocorrelation in Intraday Index Option Mid-Quote Changes on Non-Announcement Days

Notes: The table shows serial correlation coefficients for changes in the bid-ask midpoint during subintervals of the trading day. Weighted averages are reported in percentile buckets based on where the option strike lies relative to the 10:00 A.M. Risk Neutral Density. Weights are proportional to the number of quotes reported for each option strike within the specified interval.

	Trading Day	Pre-Announce-ment	Impact	Re-equilibration subintervals								Full Re-equili-bration
	10:00 AM 4:00 PM	10:00 AM 2:14 PM	2:14 PM 2:24 PM	2:24 PM 2:36 PM	2:36 PM 2:48 PM	2:48 PM 3:00 PM	3:00 PM 3:12 PM	3:12 PM 3:24 PM	3:24 PM 3:36 PM	3:36 PM 3:48 PM	3:48 PM 4:00 PM	2:24 PM 4:00 PM
Calls												
All	-0.347	-0.362	-0.386	-0.323	-0.394	-0.298	-0.313	-0.248	-0.316	-0.300	-0.314	-0.307
0-5%	-0.419	-0.436	-0.443	-0.353	-0.465	-0.365	-0.386	-0.290	-0.372	-0.408	-0.370	-0.372
5-15%	-0.354	-0.375	-0.372	-0.295	-0.380	-0.286	-0.333	-0.214	-0.311	-0.303	-0.305	-0.297
15-25%	-0.299	-0.309	-0.357	-0.293	-0.367	-0.267	-0.259	-0.227	-0.290	-0.232	-0.261	-0.268
25-35%	-0.280	-0.292	-0.331	-0.279	-0.314	-0.225	-0.225	-0.221	-0.257	-0.212	-0.258	-0.242
35-45%	-0.296	-0.312	-0.343	-0.305	-0.339	-0.248	-0.242	-0.227	-0.282	-0.219	-0.252	-0.260
45-55%	-0.271	-0.279	-0.327	-0.298	-0.311	-0.232	-0.242	-0.198	-0.254	-0.199	-0.281	-0.247
55-65%	-0.257	-0.270	-0.314	-0.292	-0.309	-0.211	-0.243	-0.198	-0.251	-0.204	-0.281	-0.244
65-75%	-0.331	-0.349	-0.390	-0.382	-0.421	-0.285	-0.275	-0.209	-0.304	-0.263	-0.373	-0.303
75-85%	-0.431	-0.445	-0.485	-0.495	-0.495	-0.408	-0.417	-0.365	-0.408	-0.428	-0.519	-0.419
85-95%	-0.544	-0.548	-0.597	-0.612	-0.580	-0.613	-0.591	-0.570	-0.596	-0.536	-0.570	-0.544
95-100%	-0.675	-0.664	-0.731	-0.735	-0.767	-0.807	-0.762	-0.773	-0.714	-0.768	-0.802	-0.717
Puts												
All	-0.287	-0.302	-0.330	-0.286	-0.318	-0.231	-0.246	-0.198	-0.274	-0.211	-0.266	-0.249
0-5%	-0.519	-0.525	-0.570	-0.530	-0.466	-0.553	-0.594	-0.473	-0.535	-0.568	-0.593	-0.545
5-15%	-0.428	-0.447	-0.391	-0.420	-0.421	-0.332	-0.420	-0.297	-0.404	-0.352	-0.405	-0.379
15-25%	-0.335	-0.361	-0.332	-0.290	-0.354	-0.273	-0.257	-0.226	-0.267	-0.226	-0.302	-0.272
25-35%	-0.279	-0.292	-0.308	-0.289	-0.300	-0.230	-0.232	-0.203	-0.253	-0.174	-0.239	-0.234
35-45%	-0.286	-0.303	-0.340	-0.288	-0.318	-0.218	-0.227	-0.231	-0.276	-0.199	-0.249	-0.246
45-55%	-0.264	-0.278	-0.335	-0.274	-0.290	-0.214	-0.211	-0.190	-0.257	-0.196	-0.266	-0.230
55-65%	-0.229	-0.246	-0.286	-0.243	-0.248	-0.158	-0.208	-0.127	-0.228	-0.162	-0.221	-0.192
65-75%	-0.236	-0.246	-0.303	-0.246	-0.287	-0.176	-0.213	-0.149	-0.250	-0.164	-0.236	-0.209
75-85%	-0.248	-0.259	-0.335	-0.261	-0.312	-0.213	-0.223	-0.146	-0.251	-0.200	-0.232	-0.224
85-95%	-0.251	-0.268	-0.306	-0.238	-0.288	-0.190	-0.220	-0.112	-0.252	-0.183	-0.240	-0.209
95-100%	-0.267	-0.278	-0.325	-0.281	-0.314	-0.217	-0.216	-0.180	-0.268	-0.192	-0.239	-0.232

Figure 1: Intraday Behavior of the Forward Value of the S&P Index on Dec. 11, 2007

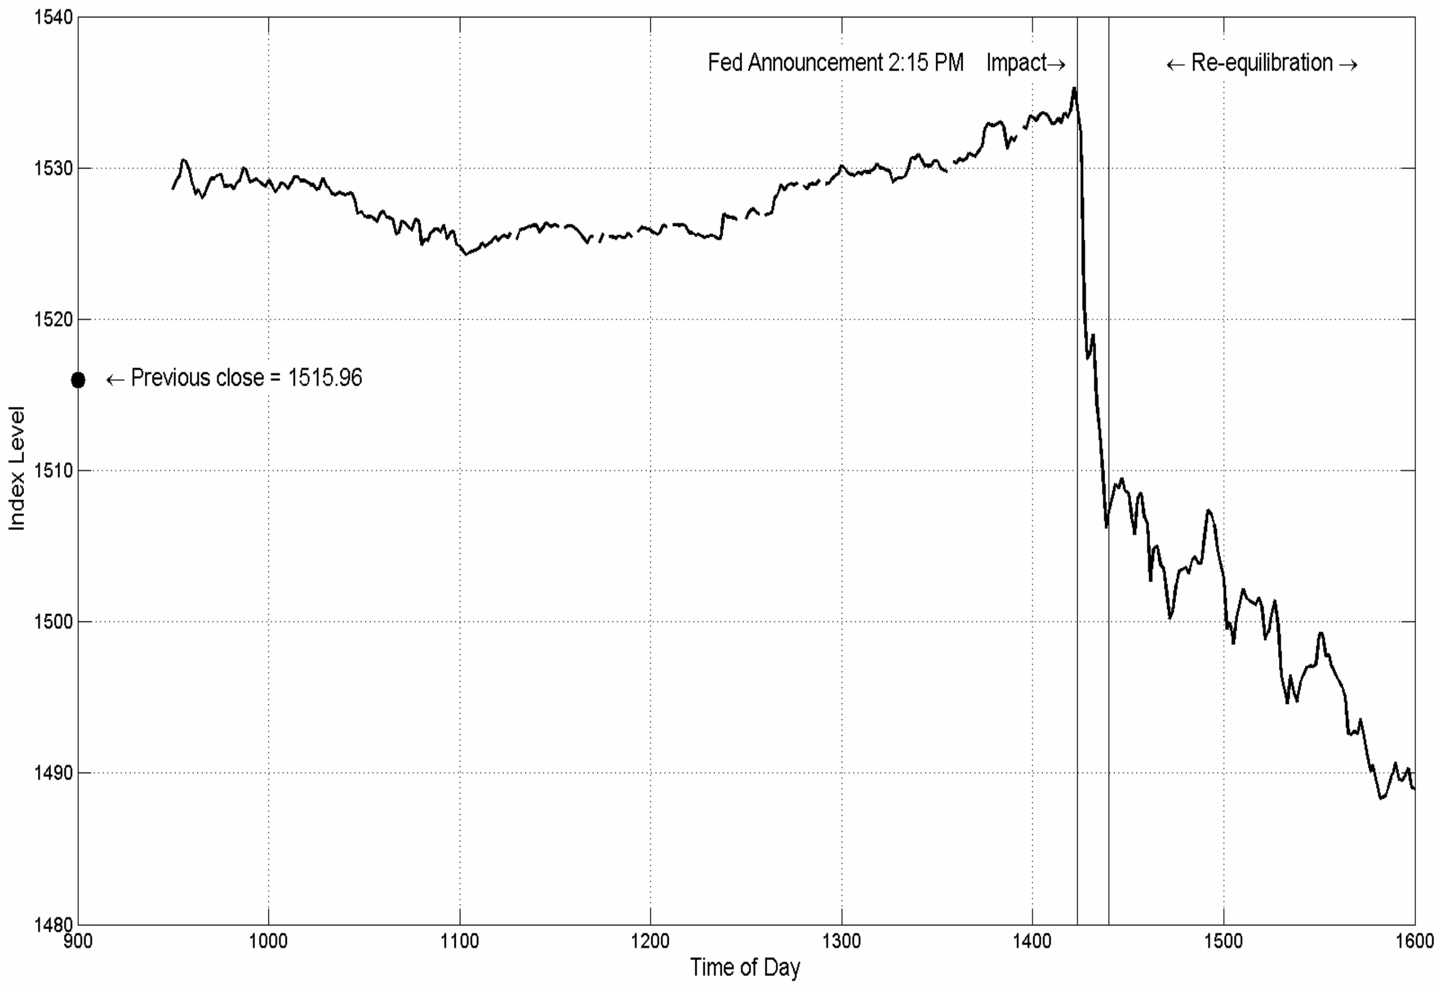

Figure 2: Initial Impact of the Fed Announcement on the Risk Neutral Density, Dec. 11, 2007

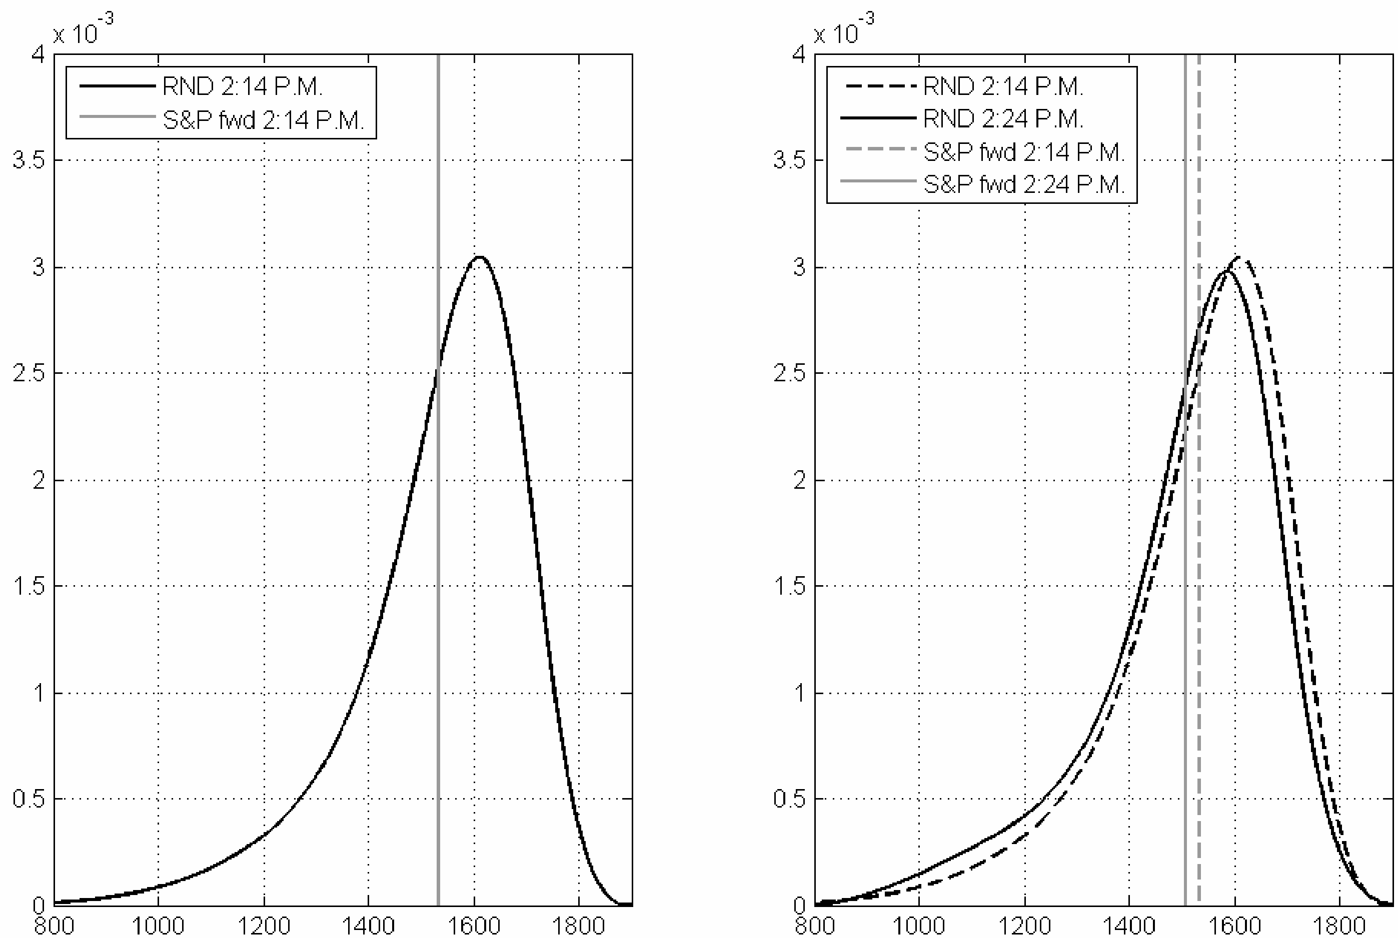

Figure 3: Re-equilibration of the Stock Market after the Fed Announcement, Dec. 11, 2007

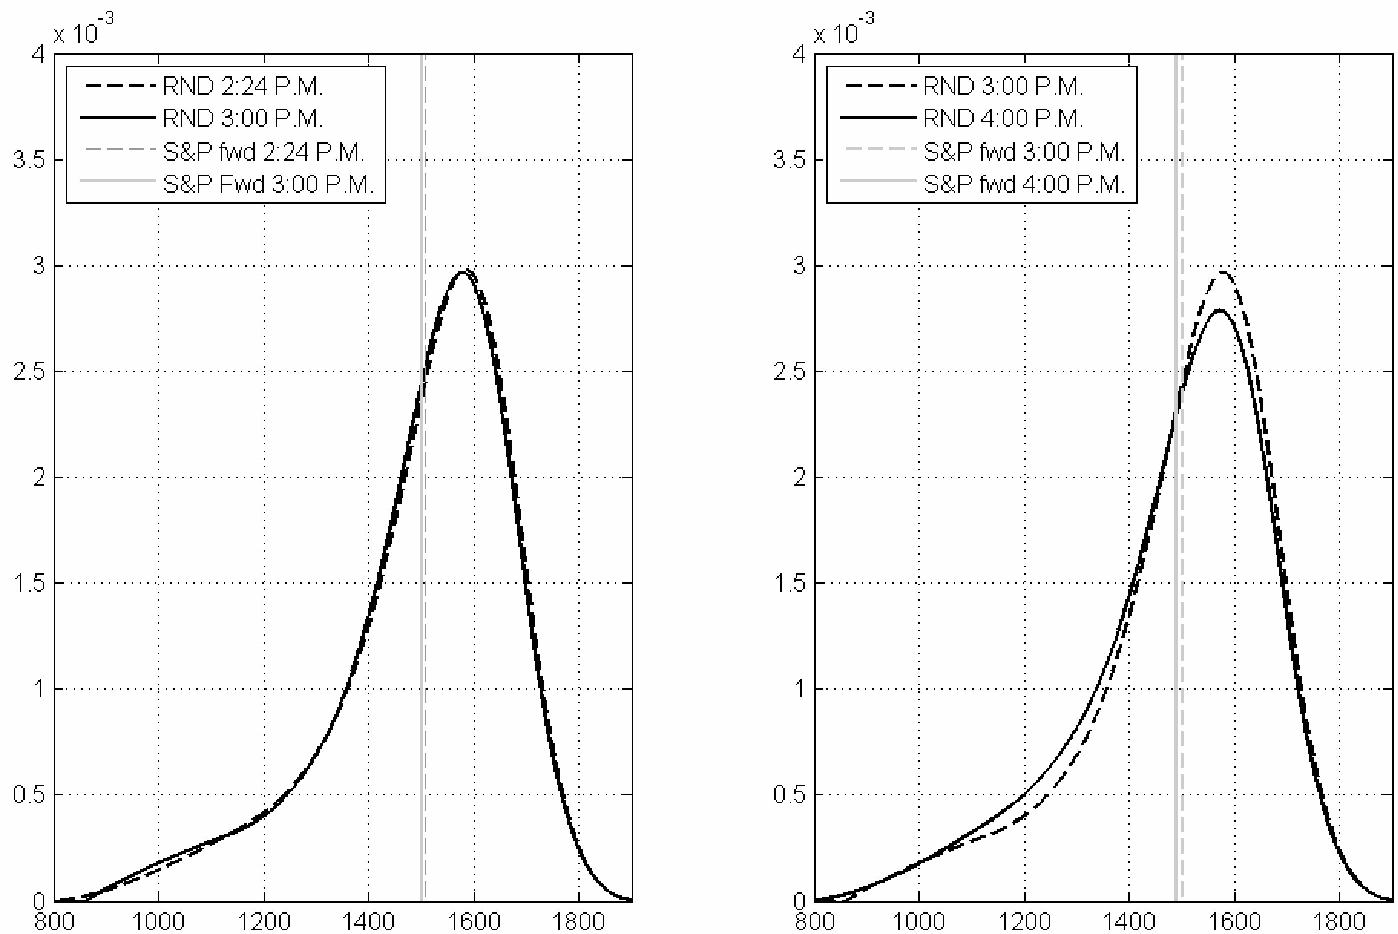

Appendix

Extracting the Risk Neutral Density from Market Options Prices

This appendix sketches out the steps we use to extract a well-behaved estimate of the risk neutral density (RND) from a set of market options data. A full exposition and discussion of alternative approaches can be found in Figlewski (2009).

1. Use bid and ask quotes, eliminating options that are too far in or out of the money: The RND is a snapshot of the risk-neutralized probability density embedded in option prices at a moment in time, so the prices must all be observed simultaneously. Transactions are sporadic at best for most exercise prices, but market makers quote firm bids and offers continuously throughout the trading day, so it is much better to take option prices from quotes than from trades. We use bid and ask quotes for S&P 500 index options obtained from the Options Price Reporting Authority (OPRA) real-time data feed of the national best bid and offer. Options that are too far out of the money are eliminated by requiring a minimum bid price of $0.50.

2. Construct a smooth curve by interpolation using a 4th degree smoothing spline in strike-implied volatility space: The theory outlined above envisions a continuum of strike prices, but in practice even very active options markets only trade in a relatively sparse set of strikes. To get an RND that is reasonably smooth, it is necessary to fill in option prices between the traded strikes by interpolation. A standard approach is to convert the option prices into Black-Scholes implied volatilities, interpolate the curve in Strike-IV space and then convert the IV curve back into a dense set of option prices. It is important to understand that doing this does not assume that the Black-Scholes model holds in the data. It simply uses the BS equation as a computational device that exploits the fact that while the market does not believe the basic BS model is correct, it does believe that the volatility "smile" should be fairly smooth. Other researchers have adopted a closely related smoothing approach of translating into strike-delta space, where delta is the first derivative of the option value with respect to the underlying.

The most common tool for interpolation in finance is a cubic spline, but the curve generated by a cubic spline is not smooth enough.[14] Interpolating with a 4th order spline solves the problem. An "interpolating spline" fits a continuous curve that goes through every observation exactly, which generally gives quite unsatisfactory densities, because it essentially forces every bit of market noise and pricing inaccuracy in the recorded option prices to be incorporated into the RND. We use a "smoothing spline" that does not force the fitted curve to go through every data point exactly. The results are insensitive to the number of knots used, so we use a single knot placed on the at the money exercise price.

---

[14] A cubic spline consists of a set of curve segments joined together at their endpoints, called "knot points," such that the resulting curve is continuous up to its second derivative, but the third derivative changes at the knots. Since the RND is obtained as the second derivative of the option value with respect to the strike price, cubic spline interpolation forces it to be continuous but allows sharp spikes to occur at the knots.

4. Fit the spline to the bid-ask spread: Typically, the average of the bid and ask quotes is assumed to represent the option price in the market, and the spline is fitted by least squares to the midpoint of the bid and ask IVs. This applies equal weight to the squared deviation between each data point and the spline approximation, regardless of whether the spline would fall inside or outside the quoted spread. But the spreads are quite wide, so we should be more concerned when the spline falls outside the quoted spread than if it stays within it. Therefore, in the optimization we use a weighting function to increase the penalty on deviations falling outside the quoted spread relative to those that remain within it.[15]

5. Use out of the money calls, out of the money puts, and a blend of the two at the money: Deep In the Money Options have wide bid-ask spreads, very little trading volume, and high prices that are almost entirely due to intrinsic value not optionality. Better information about risk neutral probabilities is obtained from out of the money and at the money contracts. But puts and calls at the same strike price regularly trade on slightly different implied volatilities, with puts often 1-2 percentage points above calls for S&P index options, so switching abruptly from one to the other at a single strike price would create an artificial jump in the IV curve, and a badly behaved density function. To avoid this, we blend the put and call bid and ask IVs to produce a smooth transition in the region around the current stock price.[16] This is done for the bid and ask IVs separately to preserve the bid-ask spread for use in the spline calculation.

---

[15] To do this efficiently, we adapt the cumulative normal distribution function to construct a weighting function that allows weights between 0 and 1 as a function of a single parameter σ.

(A1)
$$w(IV_s) = \begin{cases} N[\, IV_s - IV_{Ask} \;,\; \sigma] \;\text{ if }\; IV_{Midpoint} \le IV_s \\ N[\, IV_{Bid} - IV_s \;,\; \sigma] \;\text{ if }\; IV_s \le IV_{Midpoint} \end{cases}$$

The dependence on the exercise price X in (A1) is implicit. For the option with strike price X, IV_s is the fitted spline IV, IV_{Ask} ,IV_{Bid} and $IV_{Midpoint}$ are, respectively, the implied volatilities at the market's Ask and Bid prices, and the average of the two. N[.] denotes the cumulative normal distribution function with mean 0 and standard deviation σ and $w(IV_s)$ is the weight applied to the squared deviation $(IV_s - IV_{Midpoint})^2$. The value of σ is set by the user. A high value, such as σ = 100 in our case, effectively weights all deviations equally. Here we set σ = .001, thus placing very little weight on the difference between the spline and the midpoint of the bid and ask IVs, so long as the spline stays within the quoted spread.

[16] Here we have chosen a range of 20 points on either side of the current forward index value F_0. Specifically, let X_{low} be the lowest traded strike such that $(F_0 - 20) \le X_{low}$ and X_{high} be the highest traded strike such that $X_{high} \le (F_0 + 20)$. For traded strikes between X_{low} and X_{high} we use a blended value between $IV_{put}(X)$ and $IV_{call}(X)$, computed as

(A2)
$$IV_{blend}(X) \;=\; w\, IV_{put}(X) \;+\; (1-w)\, IV_{call}(X)$$

where
$$w \;=\; \frac{X_{high} - X}{X_{high} - X_{low}}$$

The width of the range over which to blend put and call IVs is arbitrary. In the data sample analyzed here, the forward value of the index was greater than 1100 in all cases, so that 20 points was less than 2% of the current level.

6. Add tails to the risk neutral density: Taking numerical second derivatives produces the portion of the RND that lies between the lowest and the highest strikes used in the calculations (not including the endpoints). To complete the density, it is necessary to extend it into the left and right tails.

The market's aggregation of individual investors' risk neutralized subjective probability beliefs need not obey any particular probability law, nor is it even a transformation of the true (but unobservable) distribution of realized returns on the underlying asset. Imposing a specific distribution on the data, either explicitly or implicitly, can easily produce anomalous densities that either deviate systematically from the market's RND in the region where it is observable, or that match the empirical RND out to the lowest and highest strikes, but then sharply change shape at the point where the new tail is appended.

The observable portion of the RND determines both the total probability in the tail and the density at the point where the new tail must begin. We use the empirical RND and extend it by grafting on tails drawn from Generalized Extreme Value (GEV) distributions. The tails are fitted to match the shape of the market RND at the left and right ends where it is observable.

The Generalized Extreme Value distribution is a natural candidate for modeling the tails of an unknown density, because the Fisher-Tippett Theorem proves that under weak regularity conditions the largest value in a sample drawn from an unknown distribution will converge in distribution to one of three types of probability laws, all of which belong to the generalized extreme value (GEV) family.

The GEV distribution has three parameters, which we set so that the tail satisfies three constraints. Let $X(\alpha)$ denote the exercise price corresponding to the α-quantile of the risk neutral distribution. That is, $F(X(\alpha)) = \alpha$. For simplicity, consider fitting the right tail. We first choose a value α_0 where the GEV tail is to begin, and then a second, more extreme point α_1, that will be used in matching the GEV tail shape to that of the empirical RND. The three conditions are

(A3) $$F_{EV}(X(\alpha_0)) = \alpha_0$$

(A4) $$f_{EV}(X(\alpha_0)) = f(X(\alpha_0))$$

(A5) $$f_{EV}(X(\alpha_1)) = f(X(\alpha_1))$$

where F_{EV} and f_{EV} denote the GEV distribution function and density, respectively. (A3) requires the fitted tail to contain the same total probability as the missing empirical tail; (A4) and (A5) require the density functions for the empirical RND constructed in steps 1-5 above and the GEV tail to be equal at both α_0 and α_1.

The choice of values for α_0 and α_1 is arbitrary. Our initial preference is to connect the left and right tails at α_0 values of 5% and 95%, with α_1 set at 2% and 98%, respectively. This is not always possible with our S&P 500 option data, because the price range spanned by available option strikes may not extend far enough into the tail to reach these values. In that case, we set the α_1s equal to the lowest (left tail) and highest (right tail) values available from the empirical RND and the α_0 values 3% closer to the mean.

Figure A1 provides an illustration of how this procedure works.

Figure A1: Appending GEV Tails to the RND, 2:14 P.M., Dec. 11, 2007