

Preliminary Draft -- Comments Welcome

Anatomy of a Meltdown:
The Risk Neutral Density for the S&P 500 in the Fall of 2008

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 her tireless efforts in extracting it. We also thank OptionMetrics, LLC for providing interest rates and index dividend yields.

Anatomy of a Meltdown:

The Risk Neutral Density for the S&P 500 in the Fall of 2008

ABSTRACT

We examine the intraday behavior of the risk neutral probability density (RND) for the Standard and Poor's 500 Index extracted from a continuous real-time data feed of bid and ask quotes for index options. This allows an exceptionally detailed view of how investors' expectations about returns and attitudes towards risk fluctuated during the financial crisis in the fall of 2008. The increase in risk measures was extraordinary, such as a fivefold increase in minute-to-minute volatility from October 2006 to October 2008. In contrast to moderate positive autocorrelation in the S&P index, the analysis reveals unusually large negative autocorrelation in the mean and standard deviation of the RND, which actually moderated considerably during the crisis. Using quantile regressions, we find a strong pattern in how much different portions of the RND move when the level of the stock index changes, with the middle portion of the RND amplifying the change in the index by a factor of as much as 1.5 or more in some cases. This phenomenon increased in size during the crisis and, surprisingly, was stronger for up moves than for down moves in the market.

I. Introduction

The financial crisis that struck with fury in the fall of 2008 began in the credit market and particularly the market for mortgage-backed collateralized debt obligations in the summer of 2007. It did not affect the stock market right away. In fact, U.S. stock prices hit their all-time high in October 2007. Although it has since been determined that the economy entered a recession in December 2007, the S&P 500 index was still around 1300 at the end of August 2008.

Over the next couple of months, it would fall more than 500 points, and trading below 800 by mid-November. The "meltdown" of fall 2008 ushered in a period of extreme price volatility, and general uncertainty, such as had not been seen in the U.S. since the Great Depression of the 1930s. Not only were expectations about the future of the U.S. and the world economy both highly uncertain and also highly volatile, the enormous financial losses sustained by investors sharply reduced their willingness, and their ability, to bear risk.

These factors, both uncertainty and price volatility, are reflected in the prices of options and influence the market's "risk neutral" probability distribution. The risk neutral density (RND) combines both investors' objective estimate of the probability distribution for the level of the underlying asset on the option's expiration date and the effective deformation of those probabilities induced by their attitudes towards risk. This paper will study how the way investor's valued the stock market portfolio was altered during this period, as reflected in the behavior of the RND

Thirty years ago, Breeden and Litzenberger (1978) showed how the RND could be extracted from the prices of options with a continuum of strikes. Unfortunately, there are a number of significant difficulties in adapting their theoretical result to use with option prices observed in the market. 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 next section offers a brief review of the literature on extracting risk neutral densities from option prices. Section III describes our methodology, which combines procedures used by earlier researchers with some innovations introduced in Figlewski (2009), notably using the quoted bids and offers in the market rather than transactions prices, and using the Generalized Extreme Value distribution to construct the tails of the RND that can not be extracted from options prices. Section IV describes the real-time S&P 500 index options data used in the analysis. In Section V, we present summary statistics that illustrate along several dimensions how sharply the behavior of the stock market changed in the fall of 2008, as reflected in the risk neutral density. Section VI looks more closely at how the minute-to-minute changes in the different quantiles of the RND are related to fluctuations in the level of the stock market (the forward index). Section VII concludes.

II. Literature

There exists a rather wide and continuously evolving literature on the extraction and analysis of option-implied risk-neutral distributions. To date most of the literature has focused on identifying the best methodologies for estimating the option-implied RND. We abstain from analyzing this particular strand of the literature in depth, as both Jackwerth (2004) and Figlewski (2009) give excellent reviews of the prior literature on extracting option-implied distributions.

Less work has been done in utilizing the RND as a tool to infer the market's probability estimates, although a few studies have analyzed option-implied RNDs from stock index options to derive market expectations. Bates (1991) was one of the first. He utilized S&P 500 futures options in order to analyze market forecasts in the period leading up to the 1987 market crash, as a means to determine if the market predicted the impending crash. Bates (2000) examines the options market subsequent to the 1987 crash, and finds that the option-implied RND of the S&P 500 consistently over-estimated left tail events. Jackwerth and Rubinstein (1996) arrive at a similar conclusion in their analysis of S&P options, determining that there is a much higher probability of significant index decline inferred from option-implied distributions in the post-crash period relative to the pre-1987 data period. A number of stylized facts and summary statistics for the RND of the S&P index are presented by Lynch and Panigirtzoglou (2008) for the 1985-2001 data period.

Outside the US, Gemmill and Saflekos (2000) used FTSE options to study the market's expectations ahead of British elections, while Liu et al. (2007) obtain real-world distributions from option-implied RNDs and assess their explanatory power for observed index levels relative to historical densities. The forecasting ability of index options is tested in the Spanish market by Alonso, Blanco, and Rubio (2005), in the Japanese market by Shiratsuka (2001)

A number of papers that explicitly analyze the ability of index options to predict financial crises. As mentioned above, Bates (1991) finds that S&P 500 futures options are unable to predict the October 1987 market crash. Bhabra et al. (2001) examine whether index option implied volatilities were able to predict the 1997 Korean financial crisis. Their results suggest that the options market reacted to, rather than predicted the crisis. Malz (2000) examines a number of markets and provides evidence that option implied volatilities contain information on future large magnitude returns. Like Bhabra et al. (2001), Fung (2007) studies whether option implied volatility gives an early warning sign in predicting a crisis. He finds it performs favorably compared to other measures in predicting future volatility during the 1997 Hong Kong stock market crash. Finally, a number of studies have analyzed the relationship between option implied volatility and market returns. A negative asymmetric relationship between returns and implied volatility has been well-documented in the literature. Whaley (2000) uses the implied volatility index (VIX)-as an investor fear gauge and documents a negative asymmetric relation between returns and volatility, with larger responses of the VIX to negative movements in return.

Furthermore, the perception of the VIX as a measure of the “investor fear gauge” has led to the association of negative returns with increasing investor fear. Like Whaley (2000), Malz (2000), Giot (2002), and Low (2004) use the VIX as a measure of investor fear and again find a negative return volatility relationship. Skiadopoulos (2002) undertakes a similar study in the context of emerging markets. He uses the implied volatility index from the Greek derivatives market (GVIX) and documents a negative relationship between Greek stock market returns and the GVIX. These studies all provide evidence of a negative correlation between investor fear and returns. As increases in volatility can be attributable to either an increased probability of large negative or positive returns however, no attempt is made in these prior studies to disentangle these competing effects.

III. Fitting RNDs

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.

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

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

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

Solving for the risk neutral distribution $F(X)$ yields

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

Taking the derivative with respect to X a second time gives the Risk Neutral Density function

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

In practice, we approximate the solution to (3) using finite differences of option prices observed at discrete exercise prices in the market. Let there be 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 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 probability in the right tail from

$$X_{N-1} \text{ to infinity is approximated by } 1 - \left(e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}} + 1 \right) = -e^{rT} \frac{C_N - C_{N-2}}{X_N - X_{N-2}} .$$

The density $f(X_n)$ is approximated as

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

Equations (1) - (4) show how 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:

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

and

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

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

Obtaining a well-behaved RND from market option prices is a nontrivial exercise. There are several key problems that need to be dealt with, and numerous alternative approaches have been explored in the literature. Figlewski (2009) reviews the methodological issues and develops a consistent approach that works well. The interested reader can refer to that article for full details; we will summarize the steps here.

1. Use bid and ask quotes, eliminating options too far in or out of the money: The RND is a snapshot of the risk-neutralized probability density that is embedded in option prices at a moment in time. It must be extracted from a set of simultaneously observed option prices. Given that trading is sporadic for many strike prices even in active equity options markets, one can not use transactions data to obtain a plausible density. However,

marketmakers quote firm bids and offers continuously throughout the trading day, so it is possible, and much better, to get simultaneously recorded option prices from those quotes. In this exercise we use daily closing bid and ask quotes for S&P 500 index options and eliminate strike prices that are too far in or out of the money, for which the optionality value is small relative to the bid-ask spread.

2. Construct a smooth curve in strike-implied volatility space: While the theory envisions a continuum of strike prices, 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 those strikes by interpolation. Interpolating the option prices directly does not work well, so the standard approach, originally proposed by Shimko (1993), is to convert the option prices into Black-Scholes implied volatilities (IVs), interpolate the curve in Strike-IV space and then convert the IV curve back into a dense set of option prices.¹

3. Interpolate the IVs using a 4th degree smoothing spline: The most common tool for interpolation in finance is a cubic spline, but this gives rise to two problems, that have not been fully appreciated in the literature. An "interpolating spline" fits a continuous curve that goes through every observation exactly. This essentially forces every bit of market noise and pricing inaccuracy in the recorded option prices to be incorporated into the RND. Better results are obtained with a "smoothing spline," which is not required to go through every data point and applies a penalty function to lack of smoothness in the fitted curve. The second issue is that the curve generated by a cubic spline is not smooth enough.² Interpolating with a 4th order spline solves the problem. The results are insensitive to the number of knots used, so we use a single knot placed on the at the money exercise price.

4. Fit the spline to the bid-ask spread: Typically, the spline is fitted by least squares to the midpoint of the bid and ask IVs from the market. This applies equal weight to a squared deviation regardless of whether the spline would fall inside or outside the quoted spread. But because the spreads are quite wide, we are more concerned when the spline falls outside the quoted spread than if it stays within it. We therefore increase the weighting of deviations falling outside the quoted spread relative to those that remain within it. To do this efficiently, we adapt the cumulative normal distribution function to construct a weighting function that allows weights between 0 to 1 as a function of a single parameter σ .

¹ It is important to understand that this procedure does not assume that the Black-Scholes model holds for these option prices. It simply uses the Black-Scholes equation as a computational device to transform the data into a space which is more conducive to the kind of smoothing we wish to do, not unlike taking logarithms. The reason to do this is that we want to obtain a good estimate of the risk neutral density throughout its range, but the translation from probabilities to option prices is highly nonlinear. Deep in the money and deep out of the money option prices are much less sensitive to the RND than at the money options. Converting to IVs permits a more balanced fit across the whole range of strikes.

² 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.

$$(7) \quad w(IV_s) = \begin{cases} N[IV_s - IV_{Ask}, \sigma] & \text{if } IV_{Midpoint} \leq IV_s \\ N[IV_{Bid} - IV_s, \sigma] & \text{if } IV_s \leq IV_{Midpoint} \end{cases}$$

The dependence on the exercise price X in (7) 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 $\sigma = 100$ effectively weights all deviations equally. In the results reported below, we set $\sigma = .001$, thus placing very little weight on the distance of the spline from the midpoint of the bid and ask IVs, so long as it stays within the quoted spread.

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 their intrinsic values (which give no information about probabilities). It is generally felt that better information about the market's risk neutral probability estimates is obtained from out of the money and at the money contracts.³ Because puts and calls at the same strike price regularly trade on slightly different implied volatilities, switching 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.⁴ 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. In the analysis presented below, we have chosen a range of 20 points on either side of the current index value S_0 .⁵

Specifically, let X_{low} be the lowest traded strike such that $(S_0 - 20) \leq X_{low}$ and X_{high} be the highest traded strike such that $X_{high} \leq (S_0 + 20)$. For traded strikes between X_{low} and X_{high} we use a blended value between $IV_{put}(X)$ and $IV_{call}(X)$, computed as

$$(8) \quad IV_{blend}(X) = w IV_{put}(X) + (1 - w) IV_{call}(X)$$

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

³ For example, the current methodology for constructing the well-known VIX volatility index uses only out of the money puts and calls. See Chicago Board Options Exchange (2003).

⁴ How far these two implied volatilities can deviate from one another is limited by arbitrage, which in turn depends on transactions costs of putting on the trade. In our S&P 500 index option data, even though they are European options, put and call IVs can easily be 1 to 2 percentage points apart at the money.

⁵ In the data sample analyzed below, the average value of the index was 1141, so that 20 points was on average less than 2% of the current level of the index. The width of the range over which to blend put and call IVs is arbitrary. A small amount of experimentation suggested that the specific choice has little impact on performance of the methodology for this data set.

This is done for the bid and ask IVs separately to preserve the bid-ask spread for use in the spline calculation.

6. Convert the interpolated IV curve back to option prices and extract the middle portion of the risk neutral density: Taking numerical second derivatives as described above 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.

7. Add tails to the Risk Neutral Density: We are trying to approximate the market's aggregation of the individual risk neutralized subjective probability beliefs in the investor population. The resulting density 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. Many investigators impose a known distribution on the data, either explicitly or implicitly, which then fixes the behavior of the tails by assumption. For example, assuming the Black-Scholes implied volatilities are constant outside the range spanned by the data constrains the tails to be lognormal. However, this can easily produce anomalous densities that either deviate systematically from the market's RND in the observable portion of its tail, 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 added on.⁶ We adopt a more general approach and extend the empirical RND by grafting onto it tails drawn from Generalized Extreme Value (GEV) distributions fitted to match the shape of the RND estimated from the market data over the portions of the left and right tail regions for which it is available.

Similar to the way the Central Limit Theorem makes the Normal a natural model for the sample average from an unknown distribution, the Generalized Extreme Value distribution is a natural candidate for modeling the tails of an unknown density. 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.⁷ We therefore use the GEV distribution to construct tails for the RND. The following is an overview of the tail-fitting procedure. Complete details can be found in Figlewski (2009).

The GEV distribution has three parameters, which are set so that the tail satisfies three constraints. Let $X(\alpha)$ denote the exercise price corresponding to the α -quantile of the

⁶ The observable portion of the RND determines both the total probability in the tail and the density at the point the new tail must begin. The problem is that for the lognormal (or whatever density is chosen), matching both the total tail probability and the density at the connection point generally produces either a sharp kink or even a discontinuous jump in the fitted RND at that point.

⁷ Specifically, let x_1, x_2, \dots be an i.i.d. sequence of draws from some distribution F and let M_n denote the maximum of the first n observations. If we can find sequences of real numbers a_n and b_n such that the sequence of normalized maxima $(M_n - b_n)/a_n$ converges in distribution to some non-degenerate distribution $H(x)$, i.e., $P((M_n - b_n)/a_n \leq x) \rightarrow H(x)$ as $n \rightarrow \infty$ then H is a GEV distribution. The class of distribution functions that satisfy this condition is very broad, including all of those commonly used in finance. See Embrechts, et al (1997) or McNeil, et al (2005) for further detail.

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

$$(9a) \quad F_{EV}(X(\alpha_0)) = \alpha_0$$

$$(9b) \quad f_{EV}(X(\alpha_0)) = f(X(\alpha_0))$$

$$(9c) \quad f_{EV}(X(\alpha_1)) = f(X(\alpha_1))$$

where F_{EV} and f_{EV} denote the GEV distribution function and density, respectively. (9a) requires the fitted tail to contain the same total probability as the missing empirical tail; (9b) and (9c) require the density functions for the empirical RND constructed in steps 1-6 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 was possible with our S&P 500 option data for the right tail on nearly all dates, but after the market sold off sharply during the crisis, available option prices often did not extend as far into the left tail, especially given the increase in volatility, which widened the range of the distribution. Where possible, we used 5% and 95% as the connection points, and otherwise we set α_1 equal to the furthest connection point into the tail that was available from the data and $\alpha_0 = \alpha_1 - 0.03$.

Figure 1 provides an illustration of how this procedure works.

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 distributes to the public firm bid and offer quotes, trade prices and related information in real-time. The NBBO represents the 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.⁸

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

⁸ In the present case, S&P 500 index options are only traded on the Chicago Board Options Exchange, due to a licensing agreement.

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. Indeed, 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. New York time. Options trading begins shortly after that, 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, leading to about 366 observations of the RND per day.⁹ 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, which in this study all are December expirations. 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.¹⁰ Thus, the data consist of continuously updated quotes on options with a fixed maturity that telescopes downward over time. Among other things, this means that if the annualized volatility of the underlying does not change, the RND standard deviation will shrink over time. Where relevant, we adjust for this effect by converting values to a common time horizon, either annual or daily.

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

⁹ The market closes at 4:00 P.M. but there are often prices that come in for a few minutes after that. We start the day at 10:00 A.M. (observation 31) and stop at 4:05 (observation 396) at the latest, giving us up to 366 trading minutes per day.

¹⁰ Following the convention adopted by OptionMetrics, 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.

by OptionMetrics. OptionMetrics interpolates U.S. dollar LIBOR to match option maturity and converts it into a continuously compounded rate. The projected dividends on the index are also converted to a continuous annual rate.

We are going to use the RND as a tool to examine the behavior of the S&P index options market during the financial crisis of Fall 2008. Given the sheer volume of data, we do not try to include every day in the sample. Our full sample consists of 32 days over a 3 year period. To provide a base for comparison, we examine the RND on 4 Wednesdays from October 2006 and 3 from October 2007.¹¹ The Fall 2008 sample is in two parts. We have RNDs for every day in the period from September 8 through September 30, when the crisis first broke open and then for the next 8 Wednesdays, from October 1 through November 19, 2008. Each RND is fitted at index levels from 0 to 2000 in increments of 0.50, making 4001 values in each of over 11,000 minutes.

Figure 2 provides a graphic illustration of the differences in the RNDs across the three years included in the sample. To make the comparison easier, we have chosen three dates in early October, each with exactly 71 days to expiration. Under the same conditions these curves should have the same shape, with means equal to the forward index values for their respective expiration dates. Clearly that is not the case here. In October 2006, index volatilities and many other risk measures, like credit default swap spreads, were close to their all-time record lows. On October 4, the VIX index closed at 11.86 and a one step ahead forecast of annualized volatility from the GARCH model we will be using below was only 8.32%. October 10, 2007 featured substantially higher volatility (16.67 on the VIX and 12.61% from the GARCH model), but the key fact about the rightmost curve in Figure 2 is that the S&P 500 closed at 1565.15 on the previous day, its highest level in history. But by the time the meltdown was in full swing one year later on October 8, 2008, the S&P was down to 984.08, the VIX was 57.53 and the GARCH model forecast was 65.88%.

Table 1 presents some summary data for the data sample. The full sample contains 11,712 observations, drawn from 5 different months in 2006-2008. As Figure 2 illustrates, there was considerable variability in the location and shape of the risk neutral density prior to the crisis period in the fall of 2008. Below, we present some results for each of these subperiods separately, and some that combine October 2006 and 2007, and October and November 2008.

The highest and lowest levels of the expiration day forward for the S&P 500 index in each subsample show how extreme the changes in some of these periods were. We compute RND standard deviations in terms of the index value on expiration day. As described above, this shrinks over time as maturity approaches. To allow comparisons against volatility estimates from a GARCH model and from the VIX volatility index, that are expressed as annualized percents, we annualize the RND standard deviation as a percent of the forward index level at each point in time.

¹¹ There were 4 Wednesdays in October 2006 and 5 in October 2007, but data problems forced us to eliminate two from 2007.

The GARCH model estimates come from the standard T-GARCH ("threshold GARCH") specification, which allows the asymmetric response to positive and negative returns that has been found to be strongly supported in the data. The numbers reported here are for one day ahead forecasts, with the model parameters refitted at each point using S&P daily index returns over the previous ### days. The rightmost columns report the high and low closing values of the VIX volatility index over the days in each subsample. The VIX is actually computed in a somewhat similar way to our method of extracting the RND, so it is expected to be closely related to the RND standard deviation we calculate, except that it is designed to focus on (synthetic) options with exactly 30 days to maturity, which is short than the time horizons for all but the last two of our sample dates.¹²

Figure 3 plots the values of these three volatility measures for the 32 dates in our sample. It is interesting to note that in 2006-07, the RND standard deviation was above both the GARCH volatility and the VIX. This may indicate that the market expected that volatilities would rise over time. Because both the VIX and the RND standard deviation are risk neutral values while the GARCH model estimates the empirical volatility, it would be normal to find them both above the GARCH volatility. But in the fall of 2008, extremely volatile returns caused the GARCH forecasts to rise sharply, to reach values between 4 and 5 percent per day. The pattern seen here is consistent with the market anticipating that these high volatilities would diminish fairly rapidly over time.

V. The Properties of the Risk Neutral Density in Before and During the Meltdown

Table 2 presents a variety of summary statistics regarding the behavior of the S&P 500 index and the RND during several periods, both before and during the financial meltdown. The first two columns cover the 7 days in October 2006 and 2007 that we use for comparison. The next three columns summarize the results for September, October, and November 2008, respectively.

September was when the crisis hit in force. Our data begin on September 8, the day after Fannie Mae and Freddie Mac were taken over by the Federal government. The following Monday, September 15, Lehman Brothers declared bankruptcy and Merrill Lynch sold itself unexpectedly to Bank of America. Two weeks later, on Monday, Sept. 29, the U.S. House of Representatives shocked the financial markets by voting down the first comprehensive "TARP" bailout plan. When the news hit the market shortly after 1:30 P.M., the S&P 500 index responded by falling more than 4.8% in two hours.

By October, the crisis had spread to many other markets throughout the world. The stock index continued to fall and trading was marked by extreme volatility, such as had not been seen since the 1930s. On 55% of the trading days in October and November 2008, the index moved up or down by more than 3% (which corresponds to annualized

¹² For full details of how the VIX is now calculated, see Chicago Board Options Exchange (2003).

volatility in excess of 47%). Interestingly, while it well-known that the market tends to fall faster and further on the downside, in this extraordinary period sharp moves to the upside were just as common. On the two days with the largest price changes, October 14 and 29, the market rose more than 10%.

The first two lines in Table x show the average level of the S&P cash index and its forward value during these subperiods. The forward is defined as

$$(10) \quad F_t = S_t e^{(r_t - d_t)(T-t)}$$

where F_t is the forward level of the index, S_t is the current spot index, r_t is the riskless interest rate for the period from date t to expiration at date T and d_t is the annual dividend yield on the index. The expected value of the level of the index at option expiration under the risk neutral distribution should be equal to the forward price.

The second section of Table 2 provides some statistics on the intraday variability of the forward price. The day's trading range contains quite a lot of information about the volatility of the index. When the market gyrates over a wide trading range within a day, investors experience volatility in real-time, which may also have an effect on investor confidence, that is, how strongly they stick to their prior expectations about index returns. Moreover, options marketmakers find risk control via delta hedging substantially more difficult when the prices of the underlying assets cover a broad range in a short period of time.

Remarkably, in October 2006, the forward traded over a range of less than 1% during the day. This rose to a more typical range of about 1 1/4 percent in October 2007. But in the fall of 2008, the daily range widened out to the point that on an average day the index fluctuated around 5%.

Volatility of returns is the most common way to measure price fluctuation for an option's underlying asset. Estimating realized volatility from intraday data presents several important issues related to correcting for the effects of market microstructure and non-diffusive price jumps.¹³ We do not make any effort to deal with those issues here, and simply report the volatility computed as the standard deviation of log returns over 1-minute intervals. To provide a reasonable scale for the results, average volatilities are reported in terms of basis points per hour,. The differences in intraday volatility across the subperiods are striking.

Finally, we report the autocorrelation in the minute-to-minute returns. Consistent with the common observation that the index contains some stale prices, it does show a moderate amount of positive serial correlation. Interestingly, autocorrelation appears to have gone down during the crisis. However, we would not put too much faith in the value of -0.009 for November, since there were only three days in the sample for that month.

¹³ See, for example #####.

The next section of the table presents statistics on the average moments of the RND. As expected, the RND mean is very close to the forward on average. Here, it is slightly below the forward in each of the subperiods. We report two measures of the spread of the RND around its mean, the interquartile range and the risk neutral standard deviation. These measures are expressed in terms of their ratio to the RND mean (the 10:00 A.M. value for the interquartile range and the contemporaneous RND mean for the standard deviation). This is appropriate if the volatility in index points is expected to be roughly proportional to the level, *ceteris paribus*, as would be expected if returns volatility was independent of the level of the index.

To allow proper comparisons, a second adjustment to these measures of the spread of the RND is needed. As explained above, as expiration approaches, the RND collapses around the forward index level, which is itself converging to the final level of the index on expiration day. The RND standard deviation and interquartile range are also functions of the time to expiration. In the two lines labeled as "annualized," we adjust the raw figures using the "square root of T rule" by multiplying them by $\sqrt{365/\text{days to option maturity}}$. Expressed in this way, we see that both measures were 4 to 5 times bigger by the end of the sample than there were in 2006.

The final two lines in this section show the risk neutral skewness and kurtosis. The negative skewness is expected, as are the kurtosis values above 3.0, indicating fatter tails than the normal. Interestingly, though, both of these higher moments of the RND go down in size during the crisis, meaning the risk neutral density looks considerably more like a normal distribution at the end than at the beginning.¹⁴

Next the table reports average values of the tail shape parameters for the two GEV tails that we appended to the density. In the early periods, and as we have found with other daily and intraday datasets, the left tail generally has a positive estimated ξ , meaning a fatter tail than the normal, and the right tail has a negative ξ , implying the density does not extend out to positive infinity. Both of these properties are a little counterintuitive. Since the index can not go below 0, the left tail of the distribution must be bounded from below, so it can not extend to negative infinity. There is no such bound on the right tail, however. The typical parameter values give an indication of the effect of the market's risk neutralization of the expected empirical density.

Here we see that in the crisis, (our approximation of) the left tail became truncated, to an even greater extent than the right tail. This may be an indication that once the market had fallen so far, investors may have become more aware that there has to be a lower limit. Another possible explanation for the change in tail shape has to do with the way puts were priced at the very lowest part of the range of available strikes. If those deep

¹⁴ Note that the RND, as we have calculated it here, is in price space rather than returns. Under the Black-Scholes assumption of lognormality, the density in price space should be lognormal, i.e., positively skewed. Finding it close to normal here does not indicate that the density became consistent with Black-Scholes.

out of the money puts are bid unusually high, the density in that range has to be high, as well. But given the known total probability that must lie in the lowest part of the RND, having a very high density at the start of the appended tail means that a smooth tail can not extend very far without using up the total probability it covers.

The next section of the table provides a little more information about the relationship between the mean of the RND and the forward price. The first line shows that the RND mean is lower than the forward on average, but the difference is tiny, only a few basis points. The second line shows that these two values are rarely very far apart. The standard deviation of the difference between them is on the order of 10 basis points. Neither of these measures appears to have been much altered in the meltdown.

The final section of Table 2 examines the minute to minute variability of the RND mean and volatility. We compute the variance of the 1-minute changes, multiply by 60, take the square root, and express the result in basis points per hour. Both of these risk neutral moments are found to be highly variable, especially the standard deviation. This is in sharp contrast to what was found for the volatility of the forward index value.

The autocorrelation statistics provide the explanation for the difference. Both the RND mean and standard deviation are highly negatively autocorrelated. This would be easy to understand as an artifact of the marketmaking process. For example, it is common for transactions prices to bounce rapidly between bid and ask prices that evolve much more slowly. In that case, volatility of transactions prices measured at very short intervals is much higher than at longer intervals. But that explanation does not work here, first because our data comes from the slowly moving quotes, not from trades, and second, because the negative serial correlation does not disappear when the differencing interval is increased. Autocorrelation at 5 and 10 minute intervals is about the same as for 1-minute changes. It appears that what we are seeing is more likely to be due to overshooting the movements in the forward index, followed by corrections. It is as if investors are skittish, overreacting to short run market fluctuations and then reversing themselves. This phenomenon bears fuller investigation.

VI. The Response of the Risk Neutral Density to Fluctuations in the Stock Market

The RND is a deformation, induced by risk preferences, of the market's aggregate subjective estimate of the true probability distribution for the S&P index on option expiration day. While the true density may reasonably be assumed to obey some common probability law, lognormal perhaps, or a standard fat-tailed alternative, there is no reason to expect that of the risk neutral density. We would like to analyze how the RND changes over time in response to fluctuations in the underlying spot market, with greater precision than is possible by looking at its moments, as in Table 2.

To study how the shape of the RND changes when the stock market moves, in this section we break the RND down and look at the behavior of different quantiles. Under Black-Scholes assumptions, the RND is lognormal. Changes in its mean value cause it to shift back and forth along the price axis, but its shape does not change. If the level of the forward index goes down 1 dollar, each point on the RND shifts 1 dollar to the left. If some quantile does not move by the same amount as the index does, the shape of the RND will change.

To examine this relationship, we regress the minute to minute changes of each quantile against the forward index, with the following regression,

$$(11) \quad \Delta Q_{jt} = a_j + b_j \Delta F_t$$

where Q_{jt} refers to the j th quantile at date t and F_t is the contemporaneous change in the forward index level. As above, we exclude the first half hour of trading and begin the option trading "day" at 10:00 A.M.

The first three lines in Table 3 report the estimated values, standard errors, and t-statistics of the b_j coefficients over the whole sample for 15 percentiles of the RND, from 1% to 99%. The pattern across quantiles, which corresponds to a similar pattern across option exercise prices and moneyness, is very striking. While all but the estimates for the far rightmost tail are positive and highly statistically significant, they vary widely in size. In the middle of the distribution, from around the 20th to the 60th percentiles, the RND moves further than the index forward, nearly 50% more for the 30th percentile.

The RND remains tied to the forward price, so its quantiles can not continue to move further than the forward indefinitely. Instead, the negative autocorrelation seen above for the mean is plainly operating for these quantiles, as well. A move in the stock market is amplified in the change of the RND, some of which is then reversed over time.

The effect of this pattern of coefficients across quantiles is that when the market falls, the quantiles in the middle and left side of the density mover downward further than the right tail does. Visually, it appears as if the density is flattening out and the left side is shifting downwards while the right end remains relatively fixed. When the market rises, again the left end moves up more than the right, and the height of the mode increases, giving the impression that the density is stacking up against its right end.

The remote tails at both ends are distinctly less sensitive to the change in the index. A good reason to expect this result is that both tails are extracted from deep out of the money contracts, puts on the left and calls on the right. These options have very small deltas and wide bid-ask spreads, so their fair values are quite insensitive to changes in the forward index that might occur over a period as short as a minute.

The next three sets of lines show how the pattern of RND quantile sensitivity varied across our sample periods. The same general pattern is seen in all three, but as the crisis unfolded, it became much stronger in the middle quantiles, from 30% to 60% (where

most of the options trading occurs), and distinctly less sensitive in the wings, even reversing sign in the furthest tails. Since these are synthetic tails that have been appended to the market-determined middle portion of the RND, we hesitate to draw strong conclusions from their minute-to-minute fluctuation.

For comparison, the last three lines in Table 3 provide results for this quantile regression from a different dataset, that has been studied more extensively in Figlewski (2009). The RND was constructed on a daily basis from the bids and asks at the market close, as reported by OptionMetrics. The sample covers options maturing on the quarterly March-June-September-December cycle, with maturities ranging from about 90 down to 14 days. The sample period ran from January 4, 1996 through February 20, 2008, yielding a total of 2761 observations.

Here, every coefficient is highly significant, the coefficients are still quite different from 1.0, but the pattern across percentiles is different. The left tail moves more than the forward, but the coefficients become monotonically larger as one looks further out in the tail. By contrast, from the median (50%) up, the RND quantiles do not move as much as the forward and the sensitivity goes down uniformly for higher quantiles.

The conventional wisdom, reflecting the results from many studies, is that the stock market responds more to negative than to positive returns. Table 4 examines this proposition with our quantile regressions. The top half of the table displays the coefficients and their standard errors for minutes when the forward index fell, and the bottom half for minutes when it rose.

The standard result does appear to hold in the wings of the density, but it is not obviously true for the middle portion. The estimates of b_j were larger for negative returns than for positive returns up to about the 30th percentile, but those for positive returns were greater than for negative returns from the 40th to the 80th percentiles.

Comparing these results across subperiods, we see that the first two, covering the days from 2006 and 2007, and September 2008 are not too different. For negative returns, the September 2008 sensitivity to the stock market was stronger than before the crisis in the left and right wings of the density, up to 30% and from about 80% up, but weaker in the middle. For positive returns, the September 2008 b_j values were smaller up to about the 60% and larger above that.

In the final "full meltdown" period, the sensitivity of the middle portion of the RND became even stronger, with values above 1.5 for some of the quantiles. The largest coefficients were found for positive price changes. One is tempted to venture the hypothesis that under normal circumstances, the market is expected to go up on average, so a down day is more of a surprise relative to expectations than an up day. By November 2008, the mood in the stock market may have been to expect prices to fall further, so that it was a shock to see the market go up.

VII. Concluding Comments

The risk neutral probability density that can be extracted from market prices for options contains a large amount of information about price expectations and about risk preferences. The challenge is to extract it. We began by describing a procedure that allowed us to construct well-behaved estimates of the RND from quotes on S&P 500 index options.

Applying the methodology to a new data set from the OPRA real-time data feed, we were able to analyze the intraday fluctuations in the RND on a continuous basis. This allowed an extremely detailed picture of the sharp changes in the behavior of the U.S. stock during the period of financial crisis, September - November 2008.

Most obvious was the extraordinary increase in risk measures in conjunction with an extraordinary fall in prices. For example, between October 2006 and October 2008 the average daily trading range for the S&P index expanded from 0.83 percent of the index to about 5%. Intraday volatility rose by a factor of more than 6, from 18.6 to 133.5 basis points per hour.

Our analysis also revealed an important contrast between the intraday dynamics of the forward value of the stock index and the mean of the RND. The levels of the two are very closely connected, with the standard deviation of the difference being only around 10 b.p., but fluctuations of the RND mean are much more volatile in b.p. per hour. This is possible because both the mean and the standard deviation of the RND exhibit very high negative autocorrelation. This is consistent with a view that when new information hits both the stock and options markets simultaneously, the options market is "skittish" and its strong initial reaction overshoots the new equilibrium at first, then corrects afterwards. These results are just suggestive, of course, and further evidence would be required to support such a provocative hypothesis. Interestingly, the autocorrelation appears to have diminished during the crisis of 2008, with the result that the ratio of volatilities for the RND mean relative to the S&P forward fell from 2.85 in Oct. 2006 to 1.09 in October 2008.

We used quantile regressions to examine how different portions of RND respond when the stock market moves and found a striking pattern. In the region at and somewhat below the current index, where most option trading takes place, we found that the RND quantile moves substantially more than the change in index. This effect became noticeably stronger during the crisis, with the 40th percentile of the density moving 1.37 times as much as the forward index on a minute-to-minute basis. In Oct-Nov 2008, this factor increased to 1.49 when the market return was negative and a full 1.72 on rises.

This is the first study to examine how the options market and the stock market are connected using real-time data on the risk neutral density drawn from the OPRA data feed. The RND is an extremely sensitive tool that reveals the fluctuations in market return expectations and risk tolerance. It has great potential for advancing our

understanding of how prices are formed in our financial markets. Plainly, "more research is called for."

References

Bates, D. (1996). "Jumps and Stochastic Volatility: Exchange Rate Process Implicit in Deutsche Mark Options." *Review of Financial Studies* 9, 69-107.

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

Chicago Board Options Exchange (2003). VIX CBOE Volatility Index.
<http://www.cboe.com/micro/vix/vixwhite.pdf>.

Embrechts, Paul, C. Klüppelberg, and T. Mikosch (1997). Modelling Extremal Values for Insurance and Finance. Springer.

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.

Jackwerth, J.C. (2000). "Recovering Risk Aversion from Option Prices and Realized Returns." *Review of Financial Studies* 13, 433-451.

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

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

Jiang, G.J. and Y.S. Tian (2005). "The Model-Free Implied Volatility and Its Information Content." *Review of Financial Studies* 18, 1305-1342.

McNeil, Alexander .J., Rudiger Frei, and Paul Embrechts (2005). Quantitative Risk Management. Princeton University Press.

Optionmetrics (2003). "Ivy DB File and Data Reference Manual, Version 2.0." Downloadable .pdf file, available on WRDS.

Shimko, David (1993). "The Bounds of Probability." *RISK* 6, 33-37.