Research Department **Bank of Israel**

An Upgrade of the Models Used To Forecast the Distribution of the Exchange Rate

Daniel Nathan*

Discussion Paper No. 2010.01
February 2011

Research Department, Bank of Israel. http://www.boi.org.il

* Daniel Nathan – E-mail: Daniel.Nathan@boi.org.il; Phone: 972-2-655-2551

My thanks to Akiva Offenbacher, Roy Stein, Inon Gamrasni, Eliezer Borenstein, Polina Dovman, the participants in the seminar of the Bank of Israel Research Department, and to Dr. Ben Z. Schreiber for his helpful comments.

חטיבת המחקר, בנק ישראל ת״ד 780 ירושלים 91007
Research Department, Bank of Israel, POB 780, 91007 Jerusalem, Israel

שדרוג המודלים הגוזרים תחזית התפלגות בשער החליפין

דניאל נתן

תקציר

בבנק ישראל משתמשים בשני סוגים של מודלים (פרמטרי וא-פרמטרי) לבניית פונקצית התפלגות לשער העתידי של השקל/דולר. בניית ההתפלגויות מתבססת על פיתוחים תיאורטיים המוכרים בספרות, תוך התאמתם לשוק האופציות המט״חיות הנסחרות בבורסה. על מנת לכייל את המודלים הנ״ל, משתמשים בבנק ישראל במחירי אופציות הנדגמות מספר הפקודות במהלך היום.
מטרת נייר זה הינה להציע מתודולוגיה לסינון תצפיות חריגות הנובעות מהדלילות היחסית במסחר של חלק ניכר מהאופציות בשערי המימוש השונים על ידי שימוש בסינון המבוסס על כללי ארביטראז׳ המקובלים בספרות. סינון זה משפר את יציבות שני המודלים לעומת הסינונים הקיימים כעת בבנק ישראל.

An Upgrade of the Models Used To Forecast the Distribution of the Exchange Rate

Daniel Nathan

Abstract

The Bank of Israel uses two models, one parametric and one non-parametric, to construct a distribution function for the future shekel/dollar exchange rate. The construction of the distribution is based on theoretical developments as described in the literature, adjusted to the market for forex options traded on the Israeli stock exchange. To calibrate the models, the Bank of Israel uses option prices sampled from the order book during the day.

The purpose of this study is to propose a method for filtering out abnormal observations that arise due to the relatively low level of trade in a considerable share of options at the various strike prices using a filter based on the arbitrage rules found in the literature. The method improved the stability of both models in comparison to the filtering methods currently used at the Bank of Israel.

---

* My thanks to Akiva Offenbacher, Roy Stein, Inon Gamrasni, Eliezer Borenstein, Polina Dovman, the participants in the seminar of the Bank of Israel Research Department and to Dr. Ben Z. Schreiber for his helpful comments.

1. Introduction and Motivation

Many central banks,[1] including the Bank of Israel (Galai and Schreiber 2003, Hecht and Stein 2004, Hecht and Pompushko 2005), use the risk-neutral density function in order to derive the investors' expectations concerning future prices in the stock and foreign currency markets.

In order to derive the density function, it is necessary to sample a cross-section of prices of options traded simultaneously for the same underlying asset. These prices are subjected to errors stemming from the fact that trading is not frictionless, among other things. These errors (Bliss and Panigirtzoglou 2002) include:

A. Data errors – errors in recording the data or in reporting the prices

B. Lack of synchronicity – resulting from the need to use as large as possible a cross-section.

C. Liquidity premium.

Bliss and Panigirtzoglou (2002) compared the various methods of deriving the density function by testing their sensitivity to sampling errors. They tested high-order moments derived from the distribution. The article showed that:

1) The two methods are sensitive to sampling errors, while excluding extreme observations contributes to the stability of both methods.

2) The non-parametric method is the less sensitive of the two methods.

Since the Bank of Israel uses both of these methods, the goal of this paper is not to compare the methods, but to improve the stability of both of them. In the paper, we will present a method of identifying and excluding sampling errors that is consistent with the current theory in the field's literature, thereby improving the stability of the parameters estimated by the various methods. In addition, in order to solve the problem of lack of synchronicity, we will use the orders book, not the transactions book. At the same time, it should be noted that there is no way of identifying the source of the error (an error resulting from a liquidity premium, or in recording the data).

The paper is structured as follows: Part 2 presents a short theoretical background for pricing financial assets in general, and options in particular. Part 3 reviews the literature concerning the existing methods of deriving the density function in general, and the methods used at the Bank of Israel in particular, and the empirical limitations involved in these methods. Part 4 presents the proposed changes for improving the stability of the existing models. Part 5 presents various comparisons between the existing models at the Bank of Israel and the model with the proposed filter. Part 6 summarizes the paper.

---

[1] The following are a number of examples: At the Bank of England – Bahra (1997); at the Bank of Spain – Manzano and Sanches (1998); at the Bank of Canada – Macmanus (1999); at the Banque de France – Coutant (1999); and at the Bank of Japan – Nakamura and Shiratsuka (1999).

2. A Short Theoretical Background

2.1 Pricing Financial Assets

We assume that we are in an economy with N different financial instruments (shares, bonds, etc.). We further assume that there are only two points in time in this economy: $t = 0$ and $t = 1$.The price of the i-th instrument at time t is S_t^i. At time $t = 0$, S_1^i is a random variable dependent on the sample space, Ω. Therefore, if the economy reaches situation $\omega \in \Omega$, the value of the i-th financial instrument at time $t = 1$ will be $S_1^i(\omega)$.

A portfolio is an N-dimensional row vector $h = \{h_1, h_2, \ldots h_n\}$, such that h_i is the quantity of units that we buy from S_0^i and hold until time $t = 1$. We can therefore define the value of our portfolio at time t as:

$$(1)\ V_t^h = \sum_{i=1}^{N} h_i S_t^i, \quad t = 0,1$$

Portfolio h is called a portfolio with arbitrage if $V_0^h \le 0$ and $V_1^h \ge 0$, where one of the inequalities is strong. In simple language, in a portfolio that allows arbitrage, we do not pay money in advance, and we make a profit at the end of the period.

If there is no possibility of constructing a portfolio of this type, it can be seen (Harrison and Kreps 1979), that in a discrete sample space $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$, a vector of positive values $p_1, p_2, \ldots p_m$ exists such that:

$$(2)\ \sum_{j=1}^{m} p_j = 1$$

$$(3)\ S_0^i = \frac{1}{1+r} \sum_{j=1}^{m} p_j S_1^i(\omega_j)$$

The theoretical background described above describes an economy with two times and discrete global situations. In reality, time is continuous, and there are an infinite number of global situations. In this situation, it can be seen (Harrison and Pilska 1981) that when there is no possibility of constructing an arbitrage portfolio, the return on financial assets is equal to the risk-free interest rate:

$$(4)\ E\left(\frac{S_T}{S_t} \mid S_t\right) = \int q(S_T) \frac{S_T}{S_t} dS_T = e^{r(T-t)}$$

Corresponding to the vector of values in a discrete situation, there exists a density function in a continuous situation, $q(S_T)$, called the risk-neutral density (RND) function, that characterizes the future development of the asset.

2.2 Pricing Options

When some option $g(\cdot)$ written on underlying asset S_T is involved, pricing equation (4) becomes:

(5) $E\left(\frac{g(S_T)}{g(S_t)}|S_t\right) = \frac{1}{g(S_t)} E\big(g(S_T)|S_t\big) = e^{r(T-t)}$

It therefore follows that:

(6) $g(S_t) = e^{-r(T-t)} \int g(S_T)q(S_T)dS_T$

In other words, the distribution characterizing any underlying asset is implied in any option written on that underlying asset.

Furthermore, it can be seen (Breeden and Litzenberger 1978) that for a put option (the same result can be demonstrated with call options), the following equations hold:

(7) $\frac{\partial P(S_t, K)}{\partial K} = e^{-rt}\left[\int_0^K q(S_T)dS_T\right] = P\{S_T \le K\}e^{-rt}$

(8) $\frac{\partial^2 P(S_T, K)}{\partial K^2} = q(K)$

This means that the cumulative density function is implied in the first derivative of a put option according to the strike price, and the density function itself is implied in the second derivative.

2.3 The Connection between the Density Function and an Absence of Arbitrage in the Options Market

In Section 2.1, it was stated that the existence of the density function depended on an absence of arbitrage. In order to see an example of this in the options market, we examined one strategy in the options market:

A bear market spread – we assume that an underlying asset, S_t, is traded on the market, and that put options on that underlying asset exist at strike prices K_{i+1}, K_i, where $K_i < K_{i+1}$. A bear market spread strategy consists of buying a put option at strike price K_{i+1} and selling an additional put option at strike price K_i. Since the maximum profit on the expiry date cannot exceed K_{i+1} - K_i, it follows that:

(9) $P(K_{i+1}) - P(K_i) \le (K_{i+1} - K_i)e^{-rt} \Rightarrow \frac{P(K_{i+1}) - P(K_i)}{(K_{i+1} - K_i)} e^{rt} \le 1$

At the limit, the equation becomes:

(10) $\lim_{K_{i+1} \to K_i} \frac{P(K_{i+1}) - P(K_i)}{(K_{i+1} - K_i)} e^{rt} = \frac{\partial P(K)}{\partial K} e^{rt} = \int_0^K q(S_T)dS_T = P\{S_T \le K\} \le 1$

where the second equation is derived from Equation 7. In other words, were it possible to construct an arbitrage portfolio with the help of a bull spread strategy, the cumulative distribution function would be completely undefined, because its value could not rise to 1. The practical significance of this is that individuals will not price options according to the risk-neutral distribution if the possibility of arbitrage exists.

3. Deriving the Distributions – A Review of the Literature

3.1 The Existing Methods in the Literature

There are five principal methods[2] of deriving the risk-neutral density function: stochastic processes methods, binomial trees, parametric methods for approximating the density function, finite difference methods, and methods involving smoothing the implied standard deviation. This paper will focus on the two most popular methods among central banks: a parametric method of approximating a density function through a double lognormal distribution, and a non-parametric method involving smoothing the implied standard deviation from option prices.

In the parametric method, it can be assumed that the distribution is double lognormal;[3] the parameters characterizing the distribution (two expectations, two standard deviations, and the weight between the two situations) are derived by through the definition of a minimization problem (between the theoretical price of the option and the actual price). The non-parametric method, on the other hand, assumes nothing about the form of the probability. It derives the prices according to the strike prices, thereby deriving the form of the distribution. The two methods share common empirical problems related to the fact that a cross-section of simultaneously traded options is necessary in order to derive the distribution, because otherwise, they will not reflect the same information.

3.2 Filtering Methods Currently in Use at the Bank of Israel

Tradability in the TASE-listed options market is inadequate for sampling a large number of transactions that took place around a given point in time in order to construction the probability distribution. One possible solution for the lack of synchronicity is to use the TASE order book, while assuming that the "real" market price is based on an average of the best supply and demand prices. Use of the orders book solves the problem of a lack of synchronicity, but this approach generates empirical problems, since some options have no quote at a reasonable price. These problems are expressed in the spread between the supply price and the demand price, which reflects an inverse measure of the significance of the average as a measure of the "real" market price, and which is likely to reflect a market price that does not reflect the fair price.[4] As a result, it is unclear to what degree the parameters estimated in the model can be relied on. Furthermore, these average prices can cause long-term instability in the estimated parameters.

Filtering according to implied standard deviations – in filtering according to implied standard deviations, it was decided to filter out options with an implied standard deviation of less than 2% and greater than 20%. This method does not filter effectively when it is also exposed to Type I errors, filtering completely acceptable options as a result of periods in which the standard deviation is especially high, and Type II errors when every option with an implied standard deviation in this range is not filtered out, even though it is still likely to be an extreme observation. As a result, Hecht and Pompuskho (2005) decided to use weighting according to turnovers in order to deal with these problems (the new model is called RND in the article).

---

[2] For further discussion of the various methods, see Jackwerth (1999).

[3] Or even beyond this, such as triple lognormal – a less restrictive assumption, but one that reduces the degrees of freedom, as used as an example in Hecht and Pompushko (2005).

[4] A fair price is one that does not allow arbitrage profits.

Weights according to turnovers – weighting with the help of trading turnovers is calculated so that every option receives a weight according to its actual trading turnover on that day up until the time the sample was taken, in proportion to the total turnover of all the options on that day up until the time of the sample. According to this method, options traded on a significant turnover in comparison with the daily turnover are assigned a larger weight than options traded on small turnovers. This assumes that trading turnover is a significant characteristic of market efficiency, and the estimate for the price of an option with a relatively large turnover more faithfully reflects the real option price, and makes it possible to derive a more exact and more stable distribution of future exchange rates. While this method is not defined as a filter, it effectively filters out certain options whose turnovers are relatively negligible in comparison to other options. The weakness of this method is that deriving the distribution and standard deviation based on options data is suitable for a single common point in time, reflecting the state of the market and the various expectations in it at precisely that point in time. The trading turnovers used to determine the weights, on the other hand, are composed of a figure that accumulates during the entire trading day, and therefore do not reflect the state of the market at the point in time when the data sample was taken. In certain cases, the asymmetry in the timeliness of the data causes bias in the weights assigned to the observations. For example, an option that was close to the money at the start of the trading day, and was therefore traded on a large turnover during the day, can become further out of the money[5] as a result of a change in the exchange rate before the time when the sample was taken, thereby significantly reducing the trading turnovers. Because of light trading, at the time of sampling, the option price data will be less reliable with respect to the real option price, and with respect to information about the investors' expectations reflected in the option price. The weight the option will receive, however, will be as if the options prices sampled were very reliable.

4. The Proposed Method for Filtering Out Extreme Observations

The proposed filtering is based on defining a fair price set according the arbitrage rules.[6]

In a world in which options (call and put) are traded concurrently with trading in the underlying asset, arbitrage profits will not be generated only if the following conditions hold:[7]

a. The option price is greater than the internal value[8] of the option.

b. The option price increases/decreases monotonically for the various strike prices.

c. The gap in prices between two options of the same type (put or call) is smaller than the difference in their strike prices.

d. The cost of a butterfly strategy is positive.

---

[5] An option in the money is one that, were its expiry date today, would yield a profit to its buyer.

[6] The Bank of England uses a similar approach. See "Notes on the Bank of England – Implied Probability Density Functions."

[7] In practice, other rules can be found. For example, see Merton (1973). Carr and Madan(2005), however, provide the **minimum** number of conditions.

[8] The internal value is the amount that the option holder would receive today, were he able to exercise the option.

e. The implied standard deviation between any pair of options (put or call) with the same strike price should be equal (the put-call parity equation), with a deviation of no more than 2%.[9]

The above conditions constitute strategies that can be constructed in the market. If one of the conditions does not hold, this means that they yield a profit with no risk. Condition a. is the only one that checks each option separately from the others, and in effect filters out the options that are under-priced. The other conditions, which involve a number of options in the testing process (for example, in condition d., three options are checked at various strike prices), are designed to filter out options that are over-priced. It can therefore be stated that if these conditions do not hold, it will be as a result of over-pricing, not under-pricing, and the filtering will be conducted on the over-priced options (i.e. on the option with the highest monetary value that contributed to the profit from the strategy). Note also that condition e. is exceptional in not determining which option is "guilty" of arbitrage, because it involves two options of different types. For this reason, in this case, both options will be filtered out.[10]

The main advantage of filtering according to arbitrage rules lies in the fact that individuals price options according to the risk-neutral probability only on condition that a portfolio allowing arbitrage cannot be constructed (see Part 2). It is therefore possible that a theoretical problem exists with other filters (and consequently in the parameters obtained from the models) that do not filter out options that allow arbitrage.

4.1 Percentage of the Filtered Out Observations

The following figure displays the filtering percentages under the various methods,[11] where the division in the proposed filtering is according to conditions:

[9] Theoretically, they should be equal, if the depth of the market and the transaction costs are ignored.
[10] According to the professional literature, the first four conditions, which test the arbitrage profits in a combination of options of the same type (put or call) do not manage to eliminate arbitrage profits in this strategy, and it is therefore recommended to use this filter also.
[11] An option is catalogued as filtered out under weighting according to turnovers if its proportionate share of turnover is less than 0.1%.

Figure 2 – A Comparison of the Percentage of Observations Filtered Out Under the Various Methods

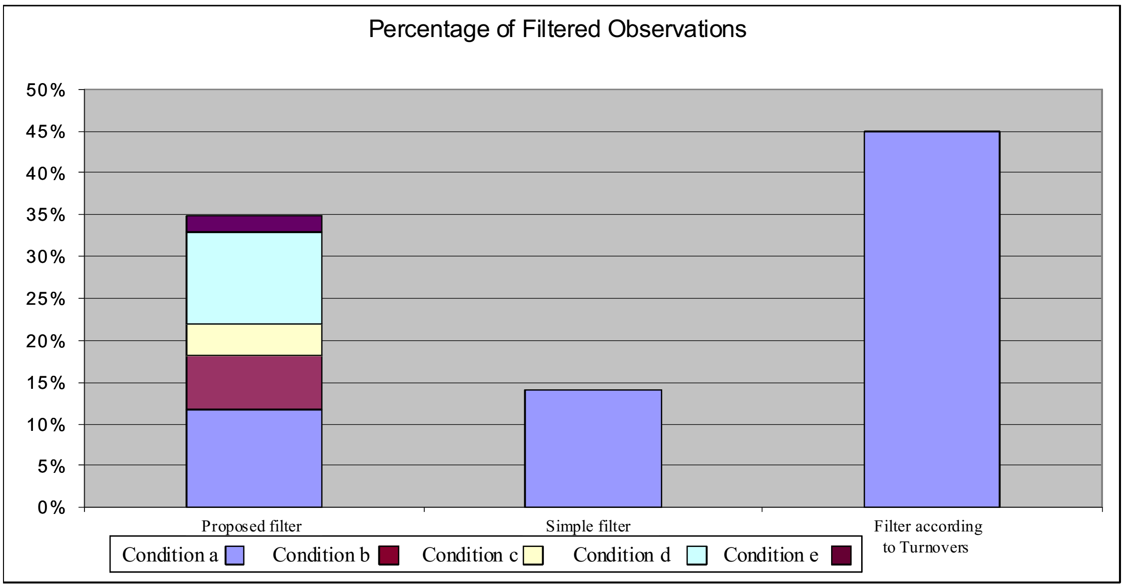

It can be seen in Figure 2 that under the proposed filter, 35% of the observations[12] are filtered out, compared with 14% under a simple filter (below 2% or over 20%) and 45% under weighting by turnovers. It is possible that the relatively small quantity of filtered out options under a simple filter is due to Type II errors – failure to filter out options that reflect unfair prices. Furthermore, the relatively large quantity of filtered out observations under filtering according to turnover is due to Type I errors. The reason for this is that most of the turnover involves options that are close to the money; options that are well in the money or far out of the money are therefore systematically filtered out, even though it is possible that they reflect fair prices. The following graph highlights this point, and displays the cumulative density function of the options passing through the filter under weighting according to turnover and under arbitrage rules, where the options are characterized by how far they are out of the money (moneyness).[13]

---

[12] Note that an average of about 31 observations were taken.
[13] The calculation is the natural logarithm of the forward price, divided by the strike price.

Figure 3 – Cumulative Empirical Density Functions

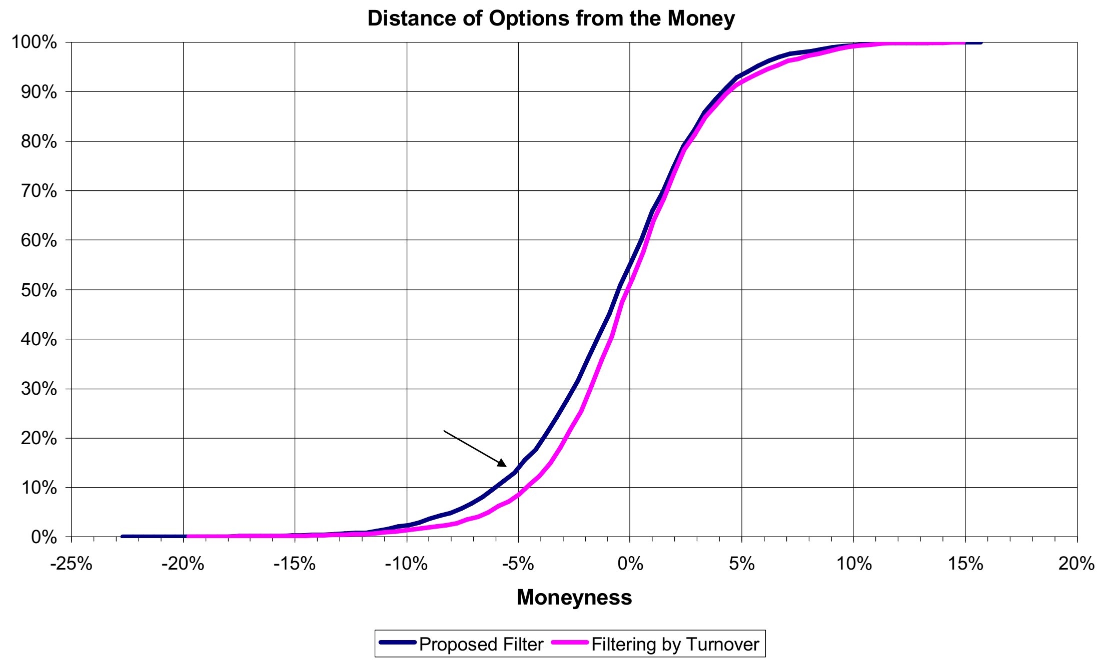

Figure 3 indicates that in filtering according to arbitrage rules, a higher percentage of options that are further out of the money are not filtered out than in filtering by turnover. For example, when the distance from the money is –5% or less, the percentage of options that are not filtered out is higher.

5. Testing and Comparing the Stability of Parameters Under the Various Methods

The market of TASE-listed options does not feature high tradability; sometimes, the option prices sampled for the purpose of estimating the probability distributions do not reflect market prices. For this reason, as in any empirical research, it is necessary to exclude extreme observations from the sample, therefore lowering the standard deviation of the estimated model and improving the results obtained from the estimation. Proper exclusion of extreme observations should therefore contribute to the stability of the results and reduce the white noise, and the autocorrelation coefficient will rise accordingly. Bliss and Panigirtzoglou (2002) compared various methods of estimating the probability distribution by testing the sensitivity of the methods to sample errors through a test of the variance of the high-order moments. Therefore, in the framework of evaluating the filter proposed in this study, we examined the time series of probability of a devaluation and the probability of an appreciation using both the parametric method and the non-parametric method.[14]

---

[14] Because the moments cannot be derived from the non-parametric model, I chose to conduct the comparison between all the methods for the probabilities of appreciation/deviation in the underlying asset.

The following models were tested:

A. A parametric (double lognormal) model:
   A.1. Proposed filter
   A.2. Simple filter (below 2% and above 20%)

B. A parametric model – The RND method with weighting by turnovers.

C. A non-parametric model:
   C.1. Proposed filter
   C.2. Simple filter

5.1 Testing the Variance of the Parameters

In this section, we statistically test the differences between the variances of the daily changes in the probability of an extreme appreciation and the probability of an extreme devaluation under the various models, where the sample period for the sake of the comparison is 2008.[15] This period was characterized by high volatility, both in Israel and overseas, given the global financial crisis. It is therefore particularly important to derive reliable and precise information from the database of shekel/dollar options trading. The assumption in the study underlying the comparison is that a more robust method is less volatile. The reason is twofold: (1) excluding extreme observations contributes to the stability of the estimated moments, as mentioned in the preceding section, and (2) the time series of fluctuations in prices has a long memory. In the empirical literature, this characteristic is called "volatility clustering",[16] meaning that the large changes in the price are accompanied by large changes in the next period. We therefore expect the time series to be both more stable and to see a slower decay. A statistical test of the variance of the estimates obtained from the probability distributions over the sample period under the various methods will show the most robust method, which will provide the most useful information (reducing the power of the white elephants). The following are the test results[17] for the various methods described above:

[15] There are 203 observations during this period.
[16] For example, see Ding, Granger, and Engle (1993). For a longer review, see Cont (2001).
[17] The statistical test that I used is the Brown-Forsythe Modified Levene Test. The hypothesis in the study is that the variance of the probability of an appreciation/devaluation of the parametric method using the proposed filter is lower than the variance using the simple filter and using the RND method. The same test was also conducted separately for the non-parametric method using the proposed filter, versus the non-parametric method using the simple filter. The hypotheses were tested in pairs against the proposed method.

Table 1 – Test of Stability of the Parameters Under the Various Methods

	Volatility in Change in Probability of an Appreciation of Over 3%	Volatility in Change in Probability of a Devaluation of Over 3%
A.1. Parametric distribution – proposed filter	0.021	0.022
A.2. Parametric distribution – simple filter	0.044 **(p<0.01)**	0.088 **(p<0.01)**
B. RND distribution – simple filter and weighting the observations by turnover	0.038 **(p<0.01)**	0.032 **(p<0.01)**
C.1. Non-parametric distribution – proposed filter	0.027	0.023
C.2. Non-parametric distribution – simple filter	0.044 **(p<0.01)**	0.033 **(p=0.37)**

It can be seen in Table 1 above that the lowest standard deviation is obtained under the proposed method for both the parametric and the non-parametric model. Note that these differences are significant, except in one case: the variance of the daily changes in the non-parametric model under the simple filter method is not greater than in applying the proposed filter. The robsustness of the model that applies the proposed filter is therefore greater, and the information obtained is cleaner.

5.2 Testing the Degree of Smoothing of the Data

When testing the development of the principal parameters of the distribution over the sample period (Figure 6 and Figure 10), it can be seen that the development trends are reflected more clearly in the models calculated with the proposed filter, for example, the rise in the standard deviation and in the risks of an extreme devaluation and appreciation in mid-2007.

In order to test this empirically, I will present the autocorrelation graph (correlogram) in this section. This graph displays the autocorrelation value at various lags, where the time series with long-term trends shows a slower decay and higher positive values.

Figure 4

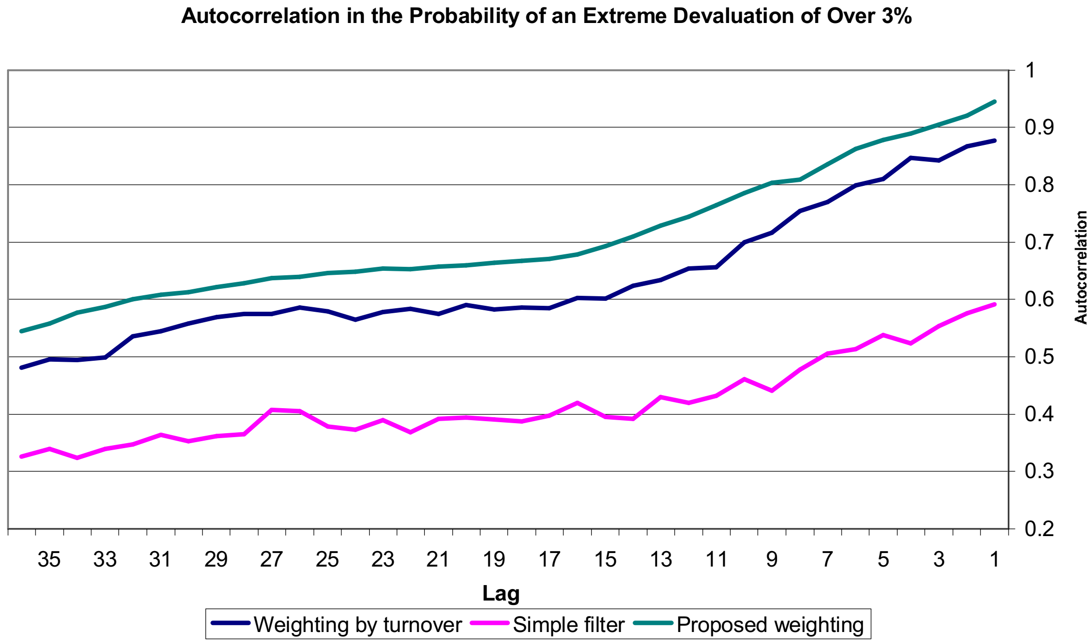

Figure 5

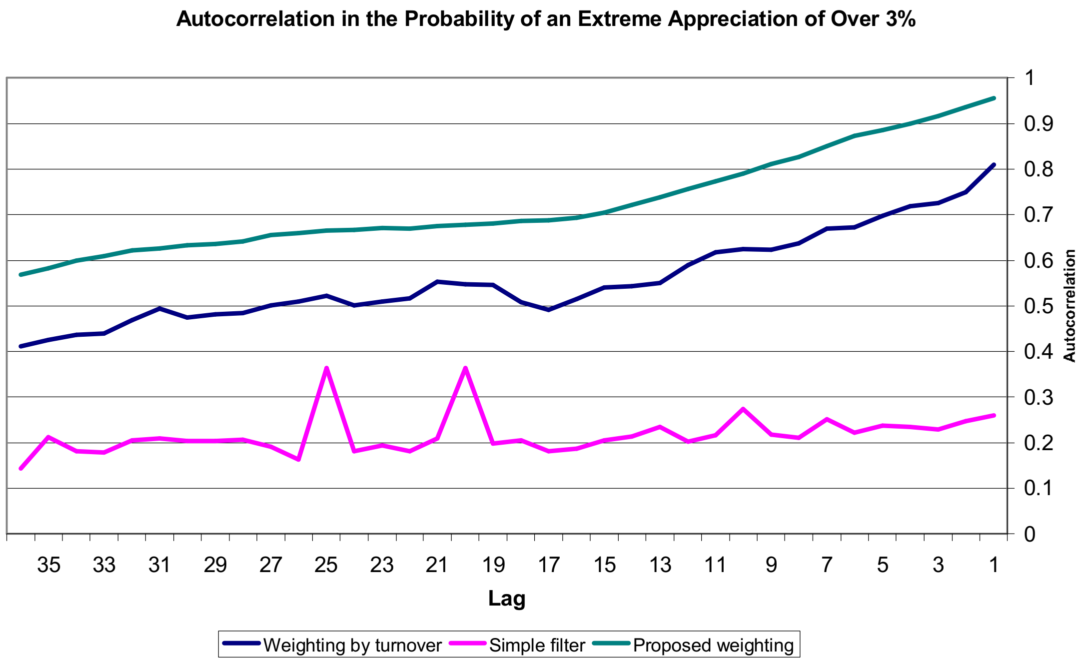

Figures 4 and 5 indicate that the slowest decay with the highest autocorrelation values for each lag occurs with the proposed filter. The following two pages display the time series of probabilities of an extreme devaluation/appreciation and the implied standard deviation under the various methods during the sample period. These give the impression that the development of the estimates can be evaluated more easily and at a higher level of confidence in Figures 7 and 10, which display the results of the models (parametric and non-parametric) under the filter proposed in this study.

Figure 6

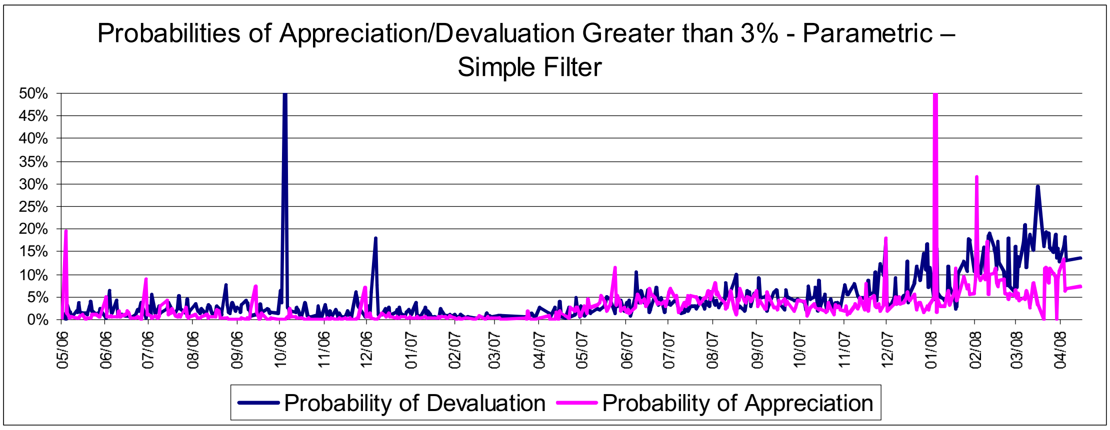

Figure 7

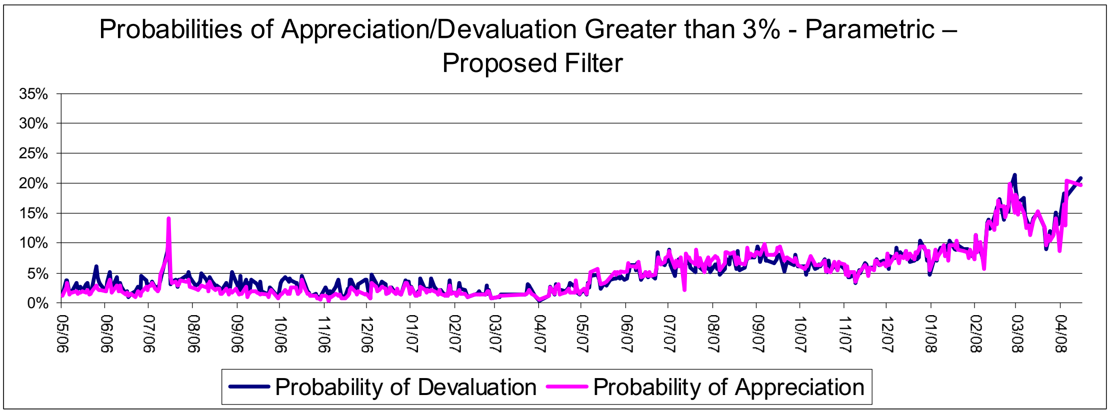

Figure 8

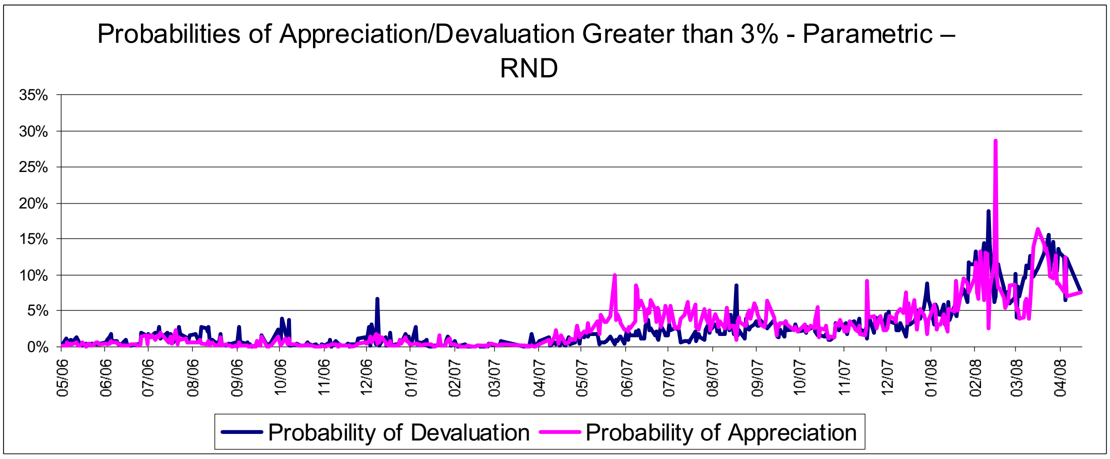

Figure 9

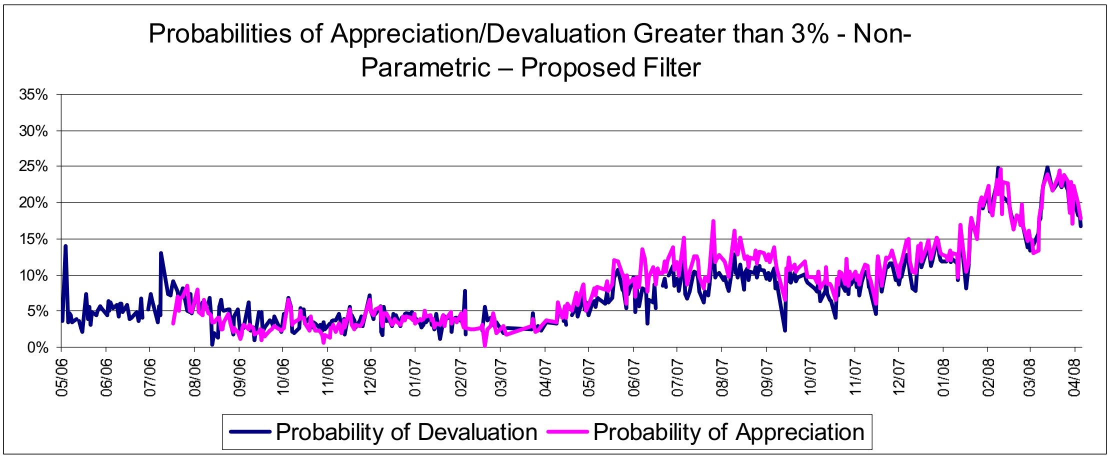

Figure 10

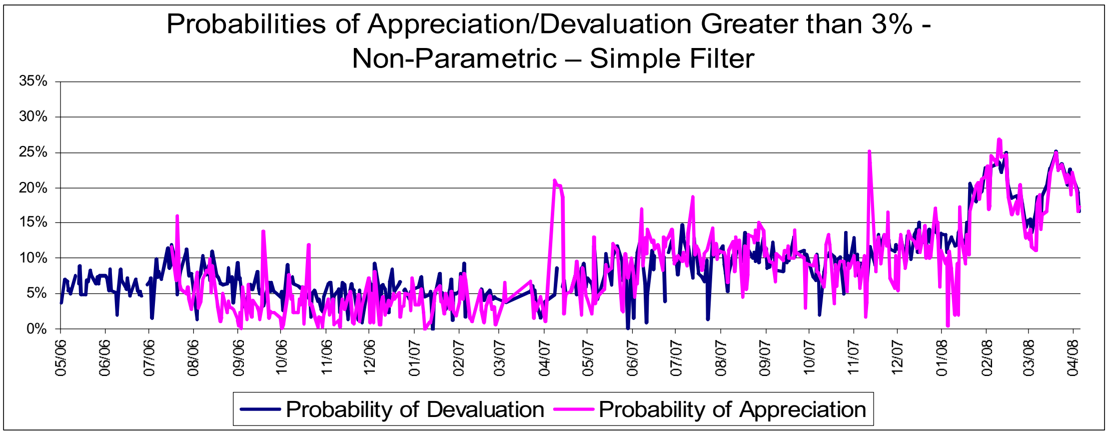

5.3 Testing and Comparing Probability Distributions on Three Selected Trading Days

In this section, we will examine to what extent the forms of the distributions obtained under the various methods differ on three different trading days. We will see graphically that the information derived from the two methods estimated using the filter proposed in this paper is more reliable.

The three trading days selected, which highlight this point, are June 14, 2006, August 12, 2007, and January 21, 2008.

Table 2 – The Parameters Estimated Under the Various Methods - June 14, 2006

	Standard Deviation	Probability of an Appreciation of Over 3%	Probability of a Devaluation of Over 3%
A.1. Parametric distribution – proposed filter	5.9%	2%	3%
A.2. Parametric distribution – simple filter	6.6%	4.5%	4.5%
B. RND distribution – simple filter and weighting the observations by turnover	6.5%	2%	3%
C.1. Non-parametric distribution – proposed filter	---	3%	3%
C.2. Non-parametric distribution – simple filter	---	---	8.4%

Table 3 – The Parameters Estimated Under the Various Methods - August 12, 2007

	Standard Deviation	Probability of an Appreciation of Over 3%	Probability of a Devaluation of Over 3%
A.1. Parametric distribution – proposed filter	8.4%	8%	7%
A.2. Parametric distribution – simple filter	10%	2%	7%
B. RND distribution – simple filter and weighting the observations by turnover	10%	3%	3%
C.1. Non-parametric distribution – proposed filter	---	10%	10%
C.2. Non-parametric distribution – simple filter	---	10%	12%

Table 4 – The Parameters Estimated Under the Various Methods – January 21, 2008

	Standard Deviation	Probability of an Appreciation of Over 3%	Probability of a Devaluation of Over 3%
A.1. Parametric distribution – proposed filter	9%	13%	12%
A.2. Parametric distribution – simple filter	10%	4%	2%
B. RND distribution – simple filter and weighting the observations by turnover	11%	8%	4%
C.1. Non-parametric distribution – proposed filter	---	13%	14%
C.2. Non-parametric distribution – simple filter	---	13%	15%

One of the most conspicuous things in these data is the great difference between the various methods in the probabilities derived. It can be seen that there is a large difference between the values of the parametric method using the simple filter, the RND method, and the non-parametric method using the simple filter. On the other hand, using the proposed filter, there is no great difference between the values obtained using the two methods: parametric and non-parametric. This fact strengthens the reliability of the proposed filter.

6. Conclusions

The goal of this paper is to propose a method based on arbitrage rules for filtering out extreme observations stemming from errors in sampling option prices that are due to the use of the order book, so that it will be possible to derive more reliable parameters for forecasting the distribution of the shekel-dollar exchange rate, as implied in the option prices on the TASE. It can be seen that applying the recommendation in this paper adds to the stability of the parameters in the long term, as reflected in lower volatility in the changes, which makes it possible to more easily follow trends in foreign currency risks. In addition, the advantage of the proposed filter over the filtering methods currently in use at the Bank of Israel is that the filter rests on theoretical foundations of conditions of arbitrage profits, and is not defined ad hoc.

Bibliography

Galai, Dan and Schreiber, B. (2003), "The Information Implied in Strategies and NIS/$ Options Traded in the OTC Market in Israel," Foreign Currency Issues, Discussion Paper Series, Bank of Israel

Hecht, Y. and Pompushko, H. (2006), "RND," Monetary Studies, Discussion Paper Series, Bank of Israel

Stein, R. and Hecht, Y. (2004), "Estimate of Expected Distribution of NIS/$ Exchange Rate Implicit in Option Prices," *The Economic Quarterly*

Stein, R., "Expected Distribution of NIS/$ Exchange Rate," Banking Issues 17

Bahra, B. (1997), "Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application", Bank of England Working Paper, No 66.

Bliss, R. and Panigirtzoglou, N. (2002), "Testing the Stability of Implied Probability Density Functions", *Journal of Banking and Finance* 26 (2-3), 381-422.

Carr, P. and Dilip, M. (2005), "A Note on Sufficient Conditions for No Arbitrage", *Finance Research Letters*, 2 (3), 125-130.

Cont, R. (2001), "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues", *Quantitative Finance* 1, 223-236.

Coutant, S. (1999), "Implied Risk Aversion in Option Prices", Banque de France Working Paper.

Harrison, J. M. and Kreps, D. (1979), "Martingales and Arbitrage in Multiperiod Securities Markets", *Journal of Economic Theory* 20, 381-408.

Harrison, J. M. and Pliska, S. (1981), "Martingales and Stochastic Integrals in the Theory of Continuous Trading", *Stochastic Processes and Their Applications* 11, 215-260.

Jackwerth, J. (1999), "Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review", *Journal of Derivatives,* 7 (2) 66-82.

Manzano, C. and Sanchez, I. (1998), "Indicators of Short-Term Interest Rate Expectations: The information Contained in the Options Market", Documento de Trabajo 9816, Servicio de Estudios, Banco de Espana.

Merton, R. (1973), "Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*, 4 (1), 141–183.

Nakamura, H. and Shiratsuka, S. (1999), "Extracting Market Expectations from Option Prices", Working paper, Bank of Japan.