American Binary FX Options:
From Theoretical Value to Market Price

Kurt Smith

Master of Philosophy (Cantab)

Bachelor of Business (Hons)

This thesis is presented for the degree of Doctor of Philosophy of The University of Western Australia

UWA Business School

2010

ABSTRACT

There is no universally accepted benchmark model for pricing exotic FX options to market, such as that for European vanilla FX options. The use of not only different models but also of different methodologies, results in widely dispersed model-dependent exotic option prices for any given set of market and contract inputs. The severity of the resultant model price dispersion is strong evidence of model risk. Model risk is especially acute for price-makers in sell-side banks who, owing to the heterogeneity of over-the-counter (OTC) exotic options, do not have actual traded daily closing mark-to-market prices for all exotic options in their book, and so must mark-to-model instead. If the model does not perform well, it will not reflect market reality, and neither will the reported daily profit and loss. Given that sell-side banks use models throughout the price-making process, from pricing market risk to identifying hedging strategies, defining risk limits, reporting to key stakeholders internally and externally, as well as determining trader bonuses and Basel II capital retention levels, model risk is an important consideration in the OTC exotic option space.

The orthodox response to model risk is price centric. Orthodoxy develops models that rely on complex and esoteric mathematics in order to improve pricing accuracy, even if it results in models that are opaque and inaccessible to most price-makers and risk managers who use them. In contrast, this research focuses on hedging strategies. This is because price-makers, irrespective of the model used to price, hedge unwanted imbalances in exotic FX option risk with liquid, traded European vanilla FX option strategies like butterflies and risk reversals. Since price-maker hedging activity is relatively model-independent, whereas price is highly model-dependent, it follows that the actual hedging behaviour of price-makers should dictate the form of the pricing model if model risk is to be minimised. In this context, the traded vanilla volatility surface is only relevant to exotic option prices insofar as it prices the cost of a traded hedge. Since it is not possible, let alone practical to trade a whole-of-volatility-surface hedge, there is no economic substance to bind orthodoxy's whole-of-volatility-surface calibrations to traded exotic option prices.

This research presents a variant of the vanna-volga model which, in addition to pricing smile risk and skew risk, also prices term risk, which is the risk that an exotic option terminates prior to expiry owing to the spot price trading at or through a barrier price. The model presented here uniquely prices to market the cost of a European vanilla FX option hedge portfolio that matures at the exotic FX option's expected stopping time, instead of at the expiration date like other vanna-volga models. By expiring at the exotic option's expected stopping time, the hedge uniquely prices the non-trivial term structure that is present not only in the level of volatility, but also in its smile and skew. Why price the cost of expensive long-dated butterflies if the exotic option is expected to terminate sooner, leaving a residual butterfly open position that has to be unwound? It is eminently logical to price instead butterflies that expire at the same time as the exotic option being hedged. As well as making redundant the use of arbitrary constants, empirics, calibrations, simulations, etc. that introduce model risk into other models, the model in this research also identifies the traded market hedge upon which the model price is based. As a result, price-makers can trade the calendar effect in exotic options in an intuitive manner that is analogous to, and internally consistent with trading the calendar effect in vanilla options.

The empirical research in this thesis involves the following steps, which together makes it larger in scale and scope than existing papers in the published literature. Model prices for American binary FX options are tested against actual traded market prices, and against a competitor model that is widely used in the exotic FX option market. For the first time, model prices are not just tested to establish their proximity to traded market prices, but also, by using additional information about the known micro-structure of the exotic FX option market, whether model prices reflect actual traded market behaviour. It was found that the model exceeded challenging performance benchmarks, in that its prices were both extremely accurate and reflected known market behaviour, for a large sample of actual traded market prices.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

I am indebted to my supervisors, Professor I. H. Y. Izan and Professor I. Watson, not only for their insightful comments, but also for their encouragement and patience which kept me motivated over the years.

I am also grateful to Mr P. Donovan and Ms J. Lee at GFI Group Inc., and to Ms O. Yap at Reuters LLC, who collectively supplied the financial market data that allowed the model presented in this thesis to be tested empirically.

I wish also to thank my wife Mary, who has traversed the peaks and troughs of my research with understanding and sympathy, and who by suffering through many drafts of my work helped to improve it. To my children, Maddy and Gabe, thank you for showing me what is important in life.

ABBREVIATIONS

AUD/USD.	AUD
Black-Scholes-Merton model.	BSM
Calendar Time.	t
Call (Put).	$\phi = 1$ ($\phi = -1$)
Delta.	Δ
Deposit Rate (Domestic).	r_d
Deposit Rate (Foreign).	r_f
EUR/USD.	EUR
Expiry Time.	T
Gamma.	Γ
GBP/USD.	GBP
(Implied) Volatility.	σ
(Implied) Volatility Surface.	IVS, $\sigma(\Delta, T)$
Market Traded Price.	Mkt
Over-the-counter.	OTC
Profit and Loss.	P&L
Spot FX Rate.	S_t
Strike Price.	K
SuperDerivatives Ask Price.	SD_{Ask}
SuperDerivatives Bid Price.	SD_{Bid}
SuperDerivatives Mid Price.	SD_{Mid}
Strike Structure of Volatility.	$\sigma(\Delta)$
Term Structure of Volatility.	$\sigma(T)$
Theoretical Value.	TV
Trader Model Ask Price.	TM_{Ask}
Trader Model Bid Price.	TM_{Bid}
Trader Model Mid Price.	TM_{Mid}
Trader Model Ask Price to Match SuperDerivatives Bid-Ask Spread Width.	TM_{Ask}^{SD}
Trader Model Bid Price to Match SuperDerivatives Bid-Ask Spread Width.	TM_{Bid}^{SD}

USD/CAD.	CAD
USD/JPY.	JPY
Vanna.	Δ_{σ} , ∂delta/∂vol, ∂vega/∂spot
Vega.	Φ
Volatility.	σ
Volatility Surface.	$\sigma(\Delta, \mathrm{T})$
Volga.	Φ_{σ} , ∂vega/∂vol

CHAPTER 1

INTRODUCTION

1.1. Motivation for the research

This thesis is motivated by the scale and scope of unresolved problems caused by the absence of a universally accepted benchmark model for pricing exotic options to market. The lack of consensus on market pricing within both academia and industry exists not only at model level, but even at the more rudimentary methodological level. The large number of competing methodologies for pricing exotic options to market is proof of both the lack of consensus in the solution and the economic significance of the problem.

Given the presence of competing models, it is not only plausible but likely that two sell-side bank counterparties to an exotic option will calculate different model prices, despite using identical inputs. The potential for both banks to show an immediate revaluation profit, owing to heterogeneous over-the-counter exotic options having to be marked-to-model rather than to market, is untenable, especially for internal and external stakeholders responsible for supervision.

The overall aim of this research is to develop a model for pricing exotic options to market which: (i) predicts actual traded market prices with sufficient accuracy to be a useful decision-making tool for price-makers in practice; (ii) identifies and quantifies market risk in a manner which provides unique insights into risk management of exotic options for price-makers in practice; and (iii) achieves real savings in computational efficiency relative to best practice quantitative models exemplified by the universal volatility models of Jex, Henderson and Wang (1999), and Lipton and McGhee (2002). A model which satisfies (i), (ii) and (iii) will potentially be a leading contender for becoming a universally accepted benchmark model for pricing exotic options to market, as there are no other published models that meet all of these criteria.

1.2. Economic significance of the research

1.2.1. Model risk

Model risk is the risk that model prices do not reflect financial economic substance. Model risk is one of the most significant challenges confronting exotic option trading businesses in the banking industry, firstly, because sell-side bank dependency upon exotic option models is high; and secondly, the absence of a universally accepted benchmark model for pricing exotic options to market increases the likelihood that model predictions will diverge from market reality.

Sell-side banks rely on exotic option models for a broad range of critical tasks from pricing market risk to identifying hedging strategies, calculating daily profit and loss, defining risk limits, reporting to key stakeholders both internally and externally, as well as determining trader bonuses and Basel II capital retention levels. It is because of their widespread use throughout the sell-side process that models have a key role not only in retaining capital, but, paradoxically, in depleting it. For example:

- in the absence of a benchmark model there is more potential for the economic cost of exotic option market risk to be mis-priced at the point of transfer from end-user to bank, which undermines bank profitability;[1]
- the lack of articulation between exotic option model prices and the market price of the vanilla options traded to hedge them, exposes banks to the risk of failing to recover the cost of intermediating market risk;
- it is because book risk limits are defined in terms of model sensitivities to factor inputs, and daily profit and loss revaluations are by necessity marked-to-model, that unrealised model profits can diverge appreciably from profits that would be realised if the options were liquidated in the market; and
- price-makers can exploit their intimate knowledge of models and their unique access to the market to game model revaluations, i.e., to maximise bonuses based on unrealised model profits at the expense of profits realised in a later reporting period.[2]

Therefore, model risk is economically significant for sell-side banks active in exotic options. Furthermore, given that the bulk of exotic option trading volume is concentrated in a relatively small subset of sell-side banks, the likelihood of model risk in one bank compounding into systemic risk across the banking industry is increased.

1.2.2. Market efficiency

Bid-ask (trading) spreads and internal (sales) distribution margins are two examples of market inefficiencies that increase the cost of optionality for buy-side end-users. While non-zero spreads and margins are necessary to encourage liquidity provision, excessive spreads and margins are a friction on the transfer of risk in the real economy. In this context, the European vanilla FX option market is relatively efficient. The market universally accepts the Garman and Kohlhagen (1983) extension of the Black and Scholes (1973) and Merton (1973) model, with its accompanying exogenous volatility surface, to price European vanilla FX options to market. As a result, European vanilla FX option market prices are so transparent they are commoditised. Competition among banks for vanilla franchise flows is so intense that market inefficiencies like bid-ask spreads and internal distribution margins are minimised to the point of profitability being crucially dependent upon high transaction volume.[3]

In contrast, the exotic FX option market is relatively inefficient. The absence of a universally accepted benchmark model for pricing exotic FX options to market makes price discovery opaque. Opacity results in wider bid-ask spreads as the risk of mis-pricing is greater. Opacity also allows excessive internal distribution margins to be concealed from end-users. In both cases, end-users pay more for optionality than if the exotic FX option market was more efficient.

[1] If a bank underprices optionality, it does not receive sufficient premium to intermediate the market risk. Conversely, if a bank overprices optionality, its franchise flow is undermined as end-users either leave for competitors or have their credit quality deteriorate by overpaying to transfer risk.

[2] A practice known as system-arbitrage in sell-side banks.

[3] Since profit is the product of volume and spread / margin, a decrease in spread / margin requires an increase in volume to maintain, let alone grow, profits.

1.2.3. Market completeness

Exotic options improve market completeness by increasing the range of payoffs available to end-users to hedge away, or speculate on, market risk. Exotic options, individually or in combination with other exotic and / or vanilla options, make it possible for market risk to be divided and subdivided into its constituent elements. As a result, hedging and speculation can be implemented with precision, according to the end-user's unique needs. However, the market is not as complete as it could be, because liquidity provision for first generation exotic options is sustained by the current size of bid-ask spreads and internal distribution margins. If there was consensus on a market pricing benchmark model, then competition would force a decline in spreads and margins for first generation exotic options that would require liquidity providers to innovate in order to maintain and grow profits sourced from franchise flows.[4]

1.3. Contribution of the research

The major contribution of this research is the discovery of a key variable crucial for pricing exotic options to market, which results in: (i) extremely strong pricing performance both absolutely and relative to a best practice competitor; (ii) unique insights into risk management for price-makers in practice; and (iii) significant computational efficiencies relative to best practice quantitative models. The key variable is the expected stopping time of exotic options, which is an analytic function dependent upon market and contract inputs only.

The discovery of the crucial role of expected stopping time in pricing exotic options to market occurred because this research is focused on the financial economics of the market, and not the financial mathematics that has come to define contemporary exotic option modelling orthodoxy. If one recognises the pivotal role of price-makers in price discovery, i.e., that price-makers make prices, not models, then it follows that it must be price-makers' hedging activity which binds the traded European vanilla volatility surface to exotic option market prices, and not arbitrary mathematical assumptions about the functional form of volatility dynamics imposed by financial engineers. That is why price-makers originally used the market price of high-order greeks defined by the traded European vanilla volatility surface, to price to market the contribution to hedging costs of high-order greeks in exotic options, as per the heuristic models of Savery (2000), Famery and Cornu (2000), Lipton and McGhee (2002), and Wystup (2003). Despite their logic and pragmatism, heuristic models had disappointing pricing performance that resulted in their becoming discredited and ultimately rejected in favour of increasingly complex mathematical models with a much weaker connection to the actual traded market mechanism.

This research makes a significant contribution to the literature by demonstrating that the disappointing pricing performance of heuristic models in the past was not because of a failure of heuristic models in general, but rather a failure of specific applications of the heuristic model. This finding is profound, not just in abstract, theoretical terms, but, as shown in sections 1.3.1-1.3.4 inclusive, in terms of the real economic contribution to participants active in exotic options.

4 There are second and third generation exotic FX options, however, these options trade in much smaller volumes than first generation options. In fact, second and third generation options are considered highly illiquid, and are rarely brokered through interbank brokers.

1.3.1. Reduction in model risk

The strength of the pricing results in this research is prima facie evidence that the model presented herein captures the essence of the actual traded market mechanism. The model's accuracy is compelling, given that the empirical research in this thesis is much larger in scale and scope than exotic option pricing research published in the literature.

But model risk is not just about accuracy, it is also about the process. Model risk increased because orthodox quantitative modelling diverged from actual traded market practice. Their arbitrary price-sensitive assumptions, while plausible, are also "unverified, indeed unverifiable" (Derman, 2003, p. 13). In contrast, price-maker's in the market, irrespective of the model they use to price, hedge unwanted imbalances in book high-order greeks with liquid, commoditised European vanilla option trading strategies like (delta neutral) risk reversals and (vega neutral) butterflies. Hence, price is highly model-dependent, whereas hedging is relatively model-independent. As a result, to reduce model risk exotic option model prices must articulate with the cost of the European vanilla options that hedge them, otherwise mark-to-model exotic option profits and mark-to-market vanilla option profits will be asymmetric and will permit arbitrage. The model presented in this thesis preserves articulation and symmetry. Orthodox quantitative models, on the other hand, take specific, traded Garman and Kohlhagen volatilities and, via calibration, turn them into generic, non-traded whole-of-volatility-surface parameters. Since it is not possible to trade a whole-of-volatility-surface hedge, there is no economic substance to bind arbitrary model parameters to traded exotic option market prices, as "there are no obvious relationships between market and model parameters" (Hakala and Kirch, 2002, p. 249). Unlike the model in this research, orthodox quantitative models dis-articulate exotic option prices from traded vanilla hedging costs, thereby undermining the economics of wholesale financial intermediation, and increasing model risk.

1.3.2. Improvement in market efficiency

In essence, the model in this research is a simple and transparent mechanism for mapping the universally accepted market supplement for European vanilla options - the traded volatility surface - to the market supplement[5] for American and European exotic options. The model in this research effectively reduces complex exotic option risk to simple combinations of traded European vanilla options, analogous to Black and Scholes' (1973) and Merton's (1973) reduction of European vanilla option risk to simple combinations of traded linear instruments. It is by working within the universally accepted vanilla paradigm that the full extent of prior knowledge in vanilla options can be directly and conveniently applied to gaining unique insights into exotic options. Conversely, orthodox quantitative models make vanilla option pricing more complex in order to obtain exotic option prices. That is, they replace the universally accepted exogenous traded volatility surface with arbitrary endogenous volatility dynamics, and, in the process, dilute and distort the information content in the European vanilla market supplement, and sever the internally consistent link with the market supplement for exotic options.[6]

---

[5] The market supplement for exotic FX options is the difference between the the market traded price and the theoretical value. This difference can be positive, negative or zero.

[6] It is incorrect to retain the exogenous traded volatility surface for vanilla options, and then to use arbitrary endogenous volatility dynamics for exotic options, as the calibration that gives effect to the transition to endogeneity for exotic options is only accurate up to a non-zero error metric. A non-zero error for in-sample closeness-of-fit represents arbitrage from using exogeneity for vanillas and endogeneity for exotics. Nevertheless, in commercial systems that use orthodox quantitative models, endogeneity

An extremely useful corollary of discovering that expected stopping time can be used to map the traded vanilla market supplement directly to the exotic market supplement, is the identification of the traded market hedge upon which the price is based. At the moment, price-makers struggle with the contradiction of hedging exotic options with variable duration with European vanilla options with fixed duration, but they do it anyway, as there are few practical alternatives.[7] Orthodox quantitative models provide no risk management guidance for price-makers, as these models price independently of a traded market hedge. The model presented in this research, however, uniquely identifies and quantifies the market risk resulting from the lengthening and shortening of exotic option risk vis-à-vis the fixed duration of traded vanilla option hedges. Expected stopping time seamlessly reconciles the temporal dimension of smile risk and skew risk sourced from both vanilla and exotic options, and hence, allows price-makers to trade the temporal dimension of exotic option risk explicitly and scientifically, in contrast to the implicit and ad hoc methods of contemporary practice. Therefore, in this research, a simple and transparent model price is underwritten by a simple and transparent traded market hedge. Both contribute to lowering entry barriers and increasing competition, which improves market efficiency.

1.3.3. Improvement in market completeness

Simplicity and transparency are also vital for making markets more complete. As the market becomes more efficient, earnings sourced from first generation exotic options decline, requiring a compensating: (i) increase in first generation exotic option volume; and / or (ii) development beyond niche markets of second, third and higher generations of exotic options, to increase spreads and margins. The former is part of the commoditisation process, and the latter makes niche payoffs widespread and / or develops new innovative payoffs. In each case, markets are more complete because end-users gain access to payoffs that were hitherto limited, either by illiquidity or unavailability. By reducing complex exotic option risk to simple combinations of traded European vanilla options, the model in this research promotes the evolution of exotics from commoditisation (of the ‘old’) to (‘new’) innovation. In contrast, by using methodologies that are complex, opaque and inaccessible to most price-makers and end-users, orthodoxy sustains the status quo and stifles innovation.

1.3.4. Improvement in computational efficiency

The model in this research is not dependent on calibration, simulation, optimisation or estimation. Orthodox quantitative models, on the other hand, are dependent on at least one of these computationally expensive numerical procedures. By having no need for computationally expensive numerical procedures, the model in this research makes one of the most significant contributions to improving computational efficiency in the exotic option literature. This research makes redundant the protracted search for marginal computational efficiencies in intermediate calculations, with the resulting diminished explanatory power in final calculations, that Ayache et

---

for exotics coexists with exogeneity for vanillas. The coexistence of inconsistent dynamics is tolerated because European vanilla option price-makers will not accept any error in vanilla pricing, and thus, reject outright orthodox quantitative models.

7 In practice, exotic option risk is not only hedged with European vanilla options, but with exotic options too. However, as shown in Chapter 3, franchise flows typically get price-makers long-the-barrier, and so initial hedging of unwanted imbalances in book high-order greeks is usually completed with liquid European vanilla options to avoid paying away to the interbank exotic option market spreads just earned from franchise flows. That is, “although every [exotic] option is relatively illiquid, the market as a whole for the greeks is very liquid” (Taleb, 1997, p. 53). Once severe imbalances are reduced, price-makers can then time their entry into the exotic option market to finesse the shape of their book when market conditions are more favourable.

el. (2004, p. 33) lament as “quantitative finance . . . wasting itself in sophisticated mathematical exercise”.

In this thesis, computational efficiencies are material not marginal because complex, numerically expensive procedures like calibration are not required. Calibration is a process of reverse-engineering European vanilla option prices under orthodox dynamics from European vanilla option prices traded in the market.[8] All calibrations have a non-zero error metric because fitting models to the market is not perfect. However, a non-zero error metric means that orthodox models mis-price European vanilla options, even though the market they calibrate to does not. Given that small errors in vanillas turn into large errors in exotics, it is surprising that exotic option price-makers tolerate errors in vanilla option prices. European vanilla option price-makers do not tolerate errors. That is why they still price under Garman-Kohlhagen dynamics and the exogenous volatility surface, not orthodox quantitative models. Therefore, by not numerically modifying the traded volatility surface, the model in this thesis not only achieves significant computational efficiency savings, but also ensures that no mis-pricing is introduced.

Since calibration is not necessary, orthodoxy’s problems with under-fitting, over-fitting, non-stationarity of parameters, etc., are neutralised. Even traditional model verification tests like in-sample closeness-of-fit are redundant. In this thesis, the error metric is zero, and hence, in-sample closeness-of-fit is perfect, because the traded volatility surface is not modified in any way.

In essence, the model in this research achieves much greater computational efficiency savings than orthodox quantitative models because it follows Derman’s (2003, p. 13) advice that “one good strategy in attempting to value exotic options that are sensitive to the smile is to try to avoid modeling the dynamics of volatility as much as possible”. Computational efficiency savings reduce the computation time required to calculate model prices and model sensitivities to factor inputs, especially for a large book of exotic options. This is an obvious economic benefit for price-makers in practice.

1.4. Structure of the research

This thesis is structured as follows. Chapter 2 reviews the theoretical and empirical literature on exotic options in general, and exotic FX options in particular. Special attention is given to the contrast between the evolutionary progress of orthodox, quantitative exotic option models published in the literature, relative to the revolutionary attributes of the model described in Chapter 3. The aims of the research and the research methodology, which includes the testing framework, price tests, performance criteria and data description, are also described in Chapter 3. Chapter 4 presents the results of empirical testing, both absolutely and relative to a best practice competitor model. Chapter 4 reports on coarse grade pricing performance, which is routinely performed in the published literature on exotic option pricing, as well as fine grade pricing performance, which is unique to this thesis, and which takes into account the known market microstructure of the interbank exotic FX option market. The conclusion is in Chapter 5.

[8] The European vanilla FX option market prices under Garman-Kohlhagen dynamics and the traded volatility surface.

CHAPTER 2

LITERATURE REVIEW

2.1. Introduction

The pronounced trend in contemporary exotic option pricing research is to develop ever more intricate models in an attempt to make real world complexity endogenous. Extensive and intensive numerical and / or empirical regimes are the norm, not the exception. In this context, developing a model that achieves both pricing accuracy and computational efficiency is challenging. There is a natural tension between accuracy, defined by the extremely fine pricing tolerances that are a prominent feature of the fiercely competitive and liquid foreign exchange option market,[9] and efficiency. Since the minimum pricing tolerance is defined by the market, most researchers sacrifice computational efficiency and rely on technological brute-force to solve exotic option pricing problems mathematically. In contrast, this research is guided by the fundamental financial economics of the problem, thereby avoiding the protracted search for marginal efficiencies in numerical routines which is currently popular in academia and industry, and concentrating instead on eliminating such routines altogether, making the prospect of substantial computational efficiency savings possible.

Most published research in exotic option pricing also introduces greater model risk than is necessary, making the identification, measurement and communication of market risk to market practitioners more opaque as a result. All credible alternatives to the Black and Scholes (1973) and Merton (1973) theoretical paradigm require, in some form, calibration to the traded volatility surface. The traded volatility surface is an exogenous correction to the BSM model to convert theoretical values for European vanilla options into market prices.[10] These exogenous corrections are made in proportion to the ∂vega/∂vol and ∂delta/∂vol of European vanilla options. However, when pricing exotic options to market it is popular to convert, through calibration, these highly specific exogenous corrections into arbitrary endogenous estimates. For example, Jex, Henderson and Wang (1999) and Lipton and McGhee (2002) calibrate to the traded volatility surface to estimate non-traded diffusion parameters such as volatility of variance, long-run variance, mean reversion speed, and relative jump heights. These parameters are weakly related at best with the original correction, and have little to no practical relevance for price-makers hedging market risk with market-traded instruments. The loss of correspondence between highly specific, BSM static market corrections and estimated dynamic model parameters increases model risk, and is a major reason why alternative models calibrated to the same traded volatility surface produce markedly different exotic option prices. Inferring model dynamics from market statics in the manner described above, is a routine practice in contemporary exotic option research which has dangerous repercussions in practice.

The model developed in this research, hereafter referred to as the Trader Model, does not require calibration to the traded volatility surface. The direct correspondence between the original highly specialised vanilla market correction and the exotic market correction is

[9] The market demands at least five digit (0.00001) accuracy for American binary FX options.

[10] The volatility surface is a matrix of volatilities with dimension tenor by delta (that is, $\Sigma = \sigma(T, \Delta)$). BSM has a single point estimate for all tenors and deltas for a specific currency pair (σ).

maintained. The information contained in the traded volatility surface is not distorted or diluted by forcing it to estimate phenomena for which it is ill-equipped, namely arbitrary, model-specific diffusion parameters. As a result, calibration empirics and numerics are redundant, significantly improving computational efficiency.[11] Furthermore, since there is a direct, one-to-one mapping between the original vanilla market correction, the exotic market correction and liquid market trading strategies, not only is model risk minimised but unique insights for hedging exotic options in practice are also obtained.

2.2. Theoretical research

2.2.1. Context

It is now conventional in academia and industry for models of the market price of exotic FX options to be based on volatility dynamics which are local, stochastic or jump, or some arbitrary combination thereof. Consequently, most recent research focuses on developing new specifications of the functional form of volatility dynamics and / or improving computational efficiency. The Trader Model takes a diametrically opposed perspective. Rather than developing incomplete endogenous models crucially dependent upon arbitrary volatility dynamics, the Trader Model instead proposes a simple exogenous procedure for converting theoretical values of exotic options into market prices which is analogous to, and internally consistent with, the universally accepted benchmark correction methodology practiced in the European vanilla option market.

This general approach has been tried before but it has produced unsatisfactory results. Savery (2000), Famery and Cornu (2000), the heuristic model of Lipton and McGhee (2002) and Wystup (2003) all exogenously correct exotic option theoretical values for crucial convexities not priced by BSM.[12] BSM price delta convexity to underlying spot asset prices, but do not price vega or delta convexity to volatility. Since ∂vega/∂vol and ∂delta/∂vol are measures of these non-BSM convexities, and (vega neutral) butterflies[13] and (delta neutral) risk reversals[14] are their respective market prices, it appeared that a practical solution to a theoretical conundrum had been found. However, the market prices of ∂vega/∂vol and ∂delta/∂vol are for European vanilla options and the adjustment to theoretical value is for American exotic options. To date, the 'solution' has been to scale convexity adjustments by empirics (Lipton and McGhee, 2002) or by touch probabilities (Wystup, 2003) to account for the possible early termination of American options.[15] Both methods are unsatisfactory. Empirical scaling suffers from the same problems as other more complex methods, such as arbitrary choices of parameters, sample period and sample frequency; and instability of estimates.

[11] Calibration is an inverse problem where cross-sectional fitting of time-dependent parameters to the strike structure is commonplace in the literature. It is redundant in the proposed research.

[12] This approach is also a feature of interbank option pricing software, such as Fenics' ∂Vega.

[13] A vega neutral butterfly is a liquid, commoditised European vanilla option strategy consisting of a long (short) strangle and a short (long) straddle, weighted such that net vega and delta is zero.

[14] A delta neutral risk reversal is: 1. A liquid European vanilla option strategy consisting of a long (short) OTM Call and a short (long) OTM Put, where the Call and Put have different strikes and identical delta. Trades in the interbank FX option market with a delta hedge to make it delta-neutral. 2. Any option strategy where the slope of the risk changes sign. This is consistent with Taleb's (1997, p. 275) definition that "a risk reversal for a book manager is the switch in risk across one point", such as "where the gamma and / or vegas flip from positive to negative across one point".

[15] Savery (2000) and Famery and Cornu (2000) are silent on calculating the quantum of the convexity adjustments for American options.

Scaling by touch probability is more theoretically appealing, but it does not work for all American binary FX options.[16]

The Trader Model specifies an exogenous procedure for calculating the market price of all American binary FX options which is internally consistent with the underlying European vanilla FX option market. It is widely known in relation to models of this class that the requisite adjustment to theoretical value is equivalent to the additional hedging costs, positive or negative, faced by a price-maker when volatility is variable.[17] However, it is identified for the first time in this research that these hedging costs are dependent upon *when* the touch level of the American binary FX option is expected to trade, not the *probability* of the touch level trading. It is this significant and unique departure from conventional wisdom which allows market prices to be obtained simply, efficiently and transparently.

For example, using EUR data from 7 March 2002, a double-no-touch (DNT) option with a theoretical value of 10% has a touch probability of 90%, and an expected stopping time of 45%.[18] That is, whilst the DNT has an extremely high probability of touching, it is still expected to last nearly half of its nominal duration. The Trader Model uniquely calculates the price of the DNT option as the cost of the hedge portfolio expiring at the expected stopping time of the DNT option. As a result, the traded volatility and interest rate data at the expected stopping time is important, not expiry data.[19] Since the DNT option is highly likely to terminate prior to expiry, constructing a hedge portfolio consisting of European vanilla options which cannot terminate until expiry as per Lipton and McGhee's heuristic model, and Wystup, creates a residual unhedged risk once a touch level trades. This thesis shows that Lipton and McGhee and Wystup's use of expiry dates in conjunction with arbitrary scaling factors is only required because the risk of early termination has not been correctly valued.

The volatility surface is an exogenous correction which compromises the strict no-arbitrage replication theory of BSM, but it is universally accepted as the market benchmark for European vanilla options owing to its simplicity, efficiency and transparency. The market accepts sub-optimal volatility dynamics because the exogenous correction delivers both price discovery and hedging insights easily understood by price-makers in practice.[20] It is conspicuous then that advances in exotic option pricing are complex, inefficient and opaque. The recent preoccupation with volatility dynamics in the literature has resulted in interpretations of market risk which are unintuitive or counter-intuitive for hedging. For example, most incomplete models define market risk as an arbitrary combination of non-traded parameters, which does not provide any insight to price-makers hedging market risk with market traded instruments. In contrast, the Trader Model defines market risk as a simple portfolio of directly quoted, high-volume market trading strategies such as butterflies and risk reversals. Therefore,

[16] Wystup's (2003) method is specific to American One Touch options (single touch level). It fails for even the simple extension to two touch levels (American Double No Touch options). Therefore, Wystup's (2003) model prices a close proxy of, but not the 'true' risk inherent in American One Touch options.

[17] Option prices have long been interpreted as accumulated hedging costs. For example, dynamic delta hedging, which is synonymous with BSM, is a hedging strategy with an accumulated cost (value).

[18] Spot of 0.8750, ATM vol of 0.0865, domestic (foreign) interest rates of 0.0195 (0.0341). A 3 month EUR DNT option with a theoretical value of 10%, has a touch probability of 90.02% and expected stopping time of 44.6%.

[19] Volatility data is for the period from today to the expected stopping time. Interest rate data is for the period from the spot delivery date to the delivery date for the expected stopping time.

[20] Pricing and hedging are inextricably linked, as hedging is concerned with how option prices change given a change in one or more of the models' arguments.

it is possible for a price-maker to understand market risk and to transact in the market to reduce those risks he or she is unwilling to take.

It is not possible to find a single specification of volatility dynamics which will explain all past, present and future market prices of exotic options. Instead there is an abundance of plausible volatility dynamics which produce different, model-specific prices. Even if 'true' volatility dynamics were found, the plethora of non-traded parameters endemic to this class of models means that market risk is defined in a manner which cannot be translated into hedging strategies in the market. If market risk is to be successfully intermediated by sell-side banks and transferred at economic value by their corporate and institutional clients, then accurate pricing and hedging insights cannot continue to be mutually exclusive. Model mathematics and market economics must be in synch, not unconnected.

Ayache et al. (2004, p. 36) conclude with the "disappearance of the model" because it is impossible to find "the absolutely true process and the absolute pricing algorithm". While being sympathetic with their conclusion, this thesis refutes their 'solution' of introducing even more calibration and parameterisation to infer market processes. Like Ayache et al. (2004), the Trader Model proposes less structure not more. Unlike Ayache et al. (2004), pre-eminence is given to traded volatility surface statics, not smile dynamics. Market prices of American binary FX options are dependent upon the shape and level of the volatility surface because price-makers use European vanilla options to hedge unwanted net book risks. This is the real direct economic relationship between the vanilla volatility surface and pricing exotic options to market. The artificially imposed mathematical relationship - calibrating arbitrary volatility dynamics to the surface - is not a close representation of this market activity. Volatility surfaces are a static construct, and exotic option pricing problems exist because dynamics cannot be inferred from statics without introducing model specificity. Ayache et al. (2004, p. 36) rely on numerical brute-force to obtain smile dynamics from which a "hedging strategy should more or less impose itself naturally". The Trader Model, on the other hand, directly identifies a unique model-independent optimal hedging strategy consisting of a portfolio of liquid market-traded instruments, thereby rendering the unverifiable functional form of smile dynamics an unnecessary distraction.[21] Using a common approach to value both sides of the book, that is, by valuing American exotic options and their European vanilla option hedges consistently, reduces model risk as one gains "insulation from the risks of the formula" (Taleb, 1997, p. 259).

As attributed to Avellaneda in Ayache et al. (2004, p. 11), guessing the volatility process is the same as guessing the price. For some reason, the former is acceptable in academia and industry, even though the latter is not. In contrast, the approach taken in this research is to tackle the problem from the perspective of the price-maker who must ultimately hedge their book. In the Trader Model, hedging is a real world concept, not a theoretical construct. It is not restricted to completeness, where self-financed dynamic delta hedging with the underlying asset eliminates all option risk. Instead, it recognises that hedging is a book, not individual option, phenomenon. Price-makers do not hedge each option perfectly, they hedge books optimally, as "it is always preferable to be roughly hedged against a broad set of eventualities

[21] In BSM, self-financed dynamic hedging with the underlying asset completely eliminates all risk. In Ayache et. al (2004), 'optimal' hedging does not eliminate all risk, but rather, minimises the variance of the profit and loss of the hedge portfolio.

than exactly hedged against a narrow parameter" (O'Connell cited in Taleb, 1997; p. 115). Therefore, in the Trader Model framework, the market price of an American binary FX option, just like the proposed optimal hedging strategy, is not perfect. Rather, the Trader Model has a more modest, pragmatic ambition similar to Ayache et al.'s (2004, p. 36) search for the "right tool" to guide decision-making in practice. However, instead of the hedging strategy being a byproduct which "very often corresponded to the trader's, model-independent intuition", the Trader Model explicitly constructs a hedging strategy based on traders' model-independent intuition, which subsequently results in a unique market price (Fig. 2.1).[22]

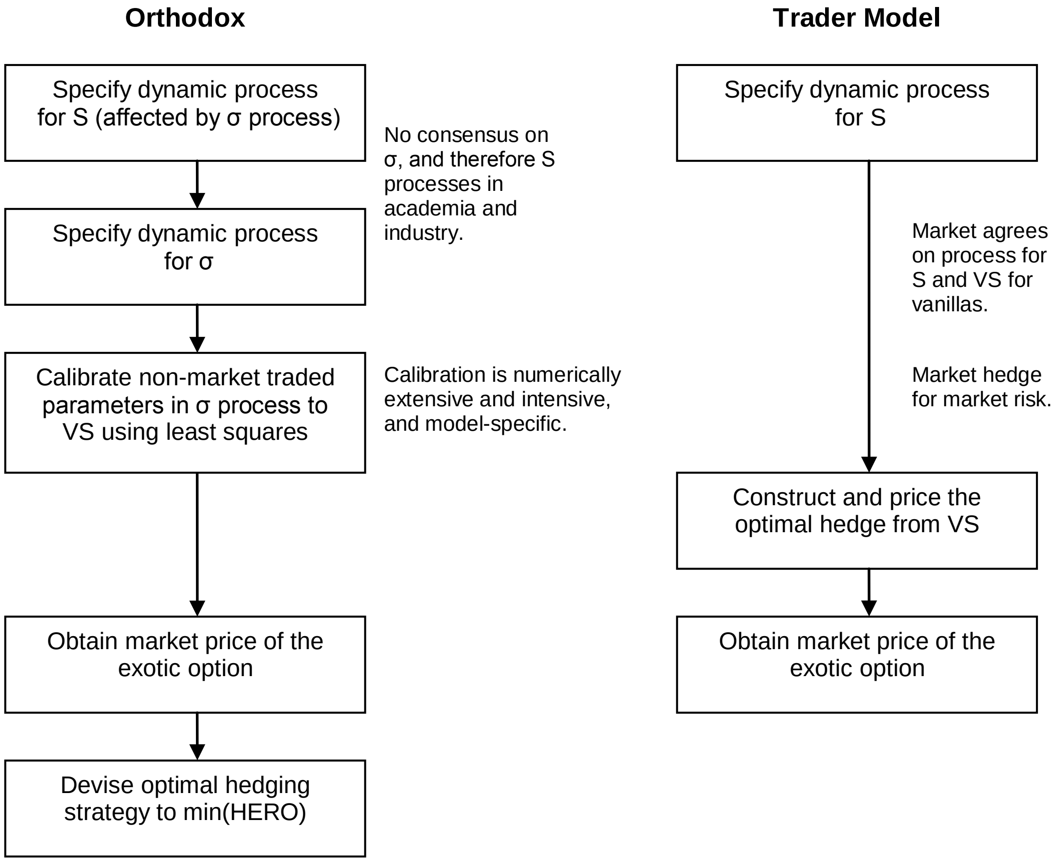

S is the underlying asset price, σ is volatility, VS is the volatility surface, and HERO is hedging error at replicating optimum.

Fig. 2.1. Schematic summary of orthodox and Trader Model approaches.

2.2.2. Significance

American binary options, both directly and indirectly through reverse barrier options, are amongst the most popular and risky of all exotic financial derivative products in the foreign exchange market. Their riskiness poses significant problems for the sell-side financial institutions who offer exotic derivative products to corporate and institutional clients, as well as

[22] All price-makers, irrespective of the modelling methodology they use to price, hedge net book risks with liquid, commoditised European vanilla option strategies like zero delta straddles, vega neutral butterflies and delta neutral risk reversals. Hence, even though pricing is highly model-dependent, hedging is relatively model-independent. Whilst it is preferable to trade on both sides of the exotic option market to minimise model risk, price-makers in the first instance at least hedge an unwanted accumulation of risk with European vanilla options, to avoid paying away as a price-taker in the interbank market, exotic option bid-ask spreads just earned as a price-maker from their franchise flows. Once primary risks are smoothed, price-makers can then time their entry to the exotic option market to finesse their book hedge, when conditions are more favourable.

the financial regulators responsible for their prudential supervision. The problems stem from the lack of consensus in, and yet widespread application of, exotic option valuation models in the foreign exchange market. For example, valuation models are used by sell-side financial institutions:

- to price hedging and other strategies for corporate and institutional clients;
- to price bank competitors directly or indirectly via brokers in the interbank market;
- to manage market risk at book and bank levels;
- to calculate regulatory capital under Basel II;
- to report profits internally and externally; and
- to review trader performance and award performance bonuses.

Whereas the European vanilla option market has BSM and the exogenous traded volatility surface, there is no comparable market price benchmark for the most popular exotic options, by volume, in the foreign exchange market. Therefore, it is not just possible but likely that two interbank counterparties to the same transaction report different profits, risk profiles and capital adequacy requirements.

The BSM methodology was first applied to European vanilla FX options by Garman and Kohlhagen (1983) and Grabbe (1983), and was extended to American binary options by Rubinstein and Reiner (1991a), Kunitomo and Ikeda (1992) and Hui (1996). The market refers to these extensions as the theoretical value of the American binary option. Theoretical values are a key reference point for exotic option price-makers and interbank brokers.[23] The widespread acceptance of the BSM inspired theoretical value allows for a significant and convenient reduction in the dimension of the market price problem. Theoretical values are derived under spot FX rate dynamics[24] described by Eq. (2.1), where S_t is the spot FX rate, r_d (r_f) is the domestic (foreign) interest rate, σ is the volatility, and dW_t is a Wiener process. Theoretical values for European vanilla FX options, and American binary FX options with (i) a single continuously monitored barrier (B), and (ii) two continuously monitored barriers (U, L), are shown in Eq's (2.2), (2.3) and (2.4), respectively. In each case, T-t is the annualised term to maturity and N(·) is the cumulative normal distribution function.

$$dS_t = \left(r_d - r_f\right) S_t dt + \sigma S_t dW_t \qquad (2.1)$$

$$Vanilla^{TV}\left(\bullet\right) = \phi S_t e^{-r_f(T-t)} N\left(\phi d_+\right) - \phi K e^{-r_d(T-t)} N\left(\phi d_-\right) \qquad (2.2)$$

where $d_\pm \triangleq \dfrac{\ln\left(\dfrac{S_t}{K}\right) + \left[r_d - r_f \pm \dfrac{\sigma^2}{2}\right](T-t)}{\sigma\sqrt{(T-t)}}$;

and $\phi = 1$ for a Call option; $\phi = -1$ for a Put option

[23] The theoretical value of an exotic option must first be agreed between banks, or between banks and interbank brokers, before an exotic option can be priced to market.

[24] These dynamics are specified in domestic risk neutral measure terms.

$$OT^{TV}(\bullet) = e^{-\omega r_d(T-t)}\left[\left(\frac{B}{S_t}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} N(-\eta e_+) + \left(\frac{B}{S_t}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} N(\eta e_-)\right] \tag{2.3}$$

where $\vartheta_- \triangleq \sqrt{\theta_-^2 + 2(1-\omega) r_d}$;

$$e_\pm \triangleq \frac{\pm \ln \frac{S_t}{B} - \sigma \vartheta_- (T-t)}{\sigma\sqrt{T-t}};$$

$$\theta_\pm \triangleq \frac{r_d - r_f}{\sigma} \pm \frac{\sigma}{2};$$

$\eta = 1$ if $H < S_t$; $\eta = -1$ if $S_t < H$;

and $\omega = 1$ for payout at end; $\omega = 0$ for payout at hit

$$DNT^{TV}(\bullet) = e^{-r_d(T-t)} \sum_{n=-\infty}^{\infty}\Big[e^{-2n\theta(h-l)}\left\{N(h+y_n) - N(l+y_n)\right\} - e^{-2n\theta(h-l)+2\theta h}\left\{N(h-2h+y_n) - N(l-2h+y_n)\right\}\Big] \tag{2.4}$$

where $\theta = \left(\dfrac{r_d - r_f - \frac{1}{2}\sigma^2}{\sigma}\right)\sqrt{T-t}$;

$$h = \frac{\frac{1}{\sigma}\ln\frac{U}{S_t}}{\sqrt{T-t}};$$

$$l = \frac{\frac{1}{\sigma}\ln\frac{L}{S_t}}{\sqrt{T-t}};$$

and $y_n = 2n(h-l) - \theta$

It is well known that there is often a substantial difference between theoretical values for American binary FX options obtained under the BSM paradigm, and actual traded market prices. The difference is called the market supplement and it can be positive or negative depending upon financial market conditions and exotic option contract specifications. Table 2.1 shows, by way of illustration, the variation in the size of the market supplement for American binary FX options in published research.

Table 2.1
Size of the market supplement for American binary FX options.

Option (FX)	Min. (Pct)	Max. (Pct)	Author
One Touch (JPY)	-5.0	+1.5	Jex, Henderson and Wang (1999)
One Touch (EUR)	-3.2	+2.0	Hakala and Wystup (2002)
Double No Touch (EUR)	+1.7	+5.7	Lipton and McGhee (2002)

In all instances, option maturity was three months. As an example, in Jex et al. (1999), a OT with a theoretical value of 30% had a market value of 25%, with bid-ask prices of 23.75% and 26.25%, respectively.

The implications of a non-zero market supplement are profound for sell-side financial institutions who use orthodox BSM methodology in the foreign exchange market:

- price-making in the FX option market is a low-margin, high-volume business, so pricing errors undermine the viability of the business as a going concern;
- market risk is incorrectly measured and inefficiently transferred, leading to extra costs in the real economy; and
- profitability, dealer bonuses and regulatory capital bear little relation to the actual market risk intermediated by the financial institution.

The market supplement is non-zero because BSM methodology does not price vega convexity to volatility and delta convexity to volatility, which are crucial factors affecting American binary FX option prices in practice. BSM only prices delta convexity to the underlying spot exchange rate. There have been several attempts to model market prices of exotic options, with mixed success. The following is a review of representative models. The review starts from the simplest extension of BSM and ends with the most popular contemporary models used in academia and industry.

2.2.3. Contemporary models

Term structure of volatility models

BSM's assumption of constant volatility (σ) in the dynamics of the underlying asset resulted in a term structure of volatility equal to zero. Merton (1973) generalised the BSM model to incorporate the non-zero term structure of volatility that was observed in the market by specifying volatility as $\sigma = \sigma_t(t)$ instead. However, Merton's generalisation did not improve the pricing of out-of-the-money (OTM) Calls and OTM Puts to market.

Local volatility models

Local volatility models are one of the simplest extensions to BSM and were pioneered by Dupire (1994), Derman and Kani (1994) and Rubinstein (1994). Instead of modelling the underlying asset price dynamics as a geometric Brownian motion with constant volatility (σ), it is generalised by specifying volatility as a deterministic function of underlying asset prices and time, $\sigma = \sigma_L(S_t, t)$, per Eq. (2.5).

$$dS_t = (r_d - r_f)S_t dt + \sigma_L(S_t, t)S_t dW_t \qquad (2.5)$$

The simplicity and transparency of the method are its major strengths. However, there are several key weaknesses. In a modelling context, the dynamics of volatility under this process are unrealistic in that smiles dissipate over time, whereas they persist in the market. Furthermore, (local) transition probabilities can be negative, which is counter-intuitive, and one must also interpolate and extrapolate sparse data to define the deterministic functional form of $\sigma_L(S_t,t)$, which is known as an ill-posed inverse problem (Ayache et al., 2004).[25] Even though market prices of European vanilla options are matched by construction, ∂vega/∂vol is typically underpriced by local volatility models, leading to significant discrepancies between model and market prices of American binary FX options (Jex, Henderson and Wang, 1999; Lipton and

[25] Interpolations and extrapolations also require the financial engineer to make further assumptions as to the functional form of the smoothing equation, e.g. quadratic, cubic, polynomial, etc. Method choice is price sensitive for both European vanilla and American exotic options (e.g. in one simple ad hoc test conducted by the author on a EUR $DNT_{TV} = 0.10$, the variation in the market supplement was 0.0034 between polynomial and cubic spline methods, ceteris paribus).

McGhee, 2002). Fig. 2.2 shows the magnitude of the problem for EUR DNT options in Lipton and McGhee (2002), where local volatility model prices are consistently well below market bid prices. Model prices falling outside market bid-ask spreads is economically unsustainable for a low-margin, high-volume business such as exotic FX option price-making.[26]

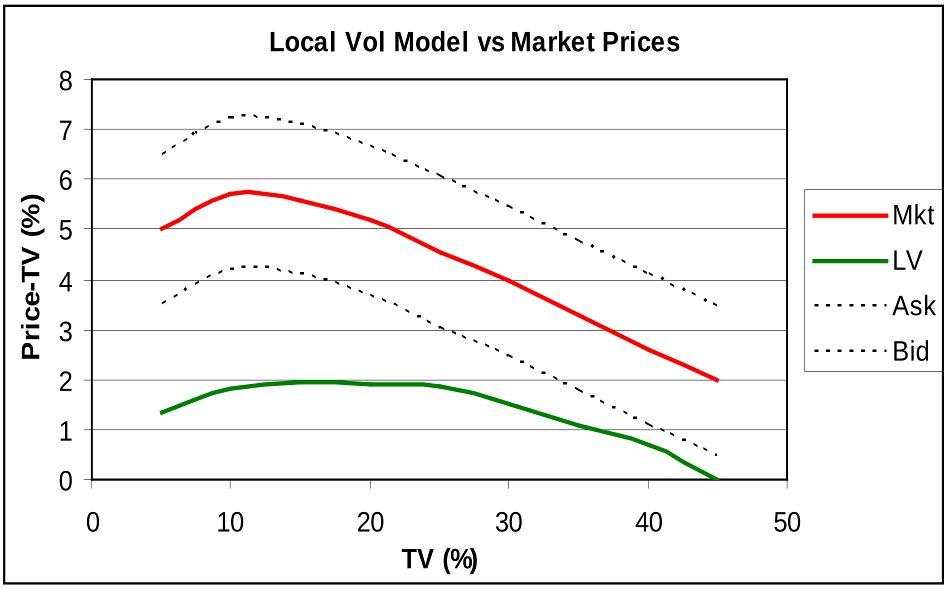

Fig. 2.2. Local volatility model prices versus market prices for DNT options reported in Lipton and McGhee (2002). LV is the local volatility model price subtract the theoretical value (TV), and Mkt is the actual traded market price subtract the theoretical value. If DNT options actually traded at theoretical value (i.e. the market supplement is zero), Mkt would plot on the x-axis.

Static hedging models

Static hedging models were pioneered by Derman, Ergener and Kani (1994), Bowie and Carr (1994), Chriss and Ong (1995), and Carr, Ellis and Gupta (1998). The modus operandi of static hedging is to create a portfolio of European vanilla options whose value is identical to the payoffs of the exotic option along its temporal (expiry) and spacial (barrier price level(s)) boundaries. Payoffs are the focus, not probabilities. Unfortunately, one has to choose arbitrary discrete points in time to replicate the payoff at the barrier level(s), as theory provides no assistance in choosing replication points for continuously monitored American binary FX options. Furthermore, the choice is price sensitive. While a 'set and forget' static hedge is attractive, as a concept, because of the extreme instability of greek risks for American binary options, one "should be warned against the static replication of instruments that have a stopping time (i.e. an unstable duration) with instruments that have a constant duration" (Taleb, 1997, p. 256). Also, in the presence of even minor transaction costs static hedging portfolios cost a lot to establish and unwind, such that "in most cases . . . the [static] replication will be impractical" in financial markets (Taleb, 1997, p. 256). Static hedging with options is a significant departure from BSM's dynamic hedging with the underlying asset. Whereas dynamic replication under BSM is riskless and strictly no-arbitrage, static replication is not. Static

[26] Using the same bid-ask spreads as Lipton and McGhee (2002), for $DNT_{TV} = 0.10$, a price-maker can lift a local volatility model ask price, and simultaneously give a market bid price, and earn a riskless immediate profit of €10,000 per €1,000,000 payout.

replication is more pragmatic, aiming for a portfolio loss distribution peaked at zero while minimising the transaction costs of replication.[27]

Heuristic models

The theory underpinning heuristic models is that the traded volatility surface is an exogenous market price correction for ∂vega/∂vol and ∂delta/∂vol in European vanilla options which is not priced by BSM, and butterflies and risk reversals are the respective market prices of these convexities. Therefore, it should be possible to obtain the market price of an American binary FX option by making an exogenous correction for its ∂vega/∂vol and ∂delta/∂vol. A hedge portfolio is constructed with OTM Put and OTM Call options with the same expiry as the American binary FX option and arbitrary delta (15 delta in Lipton and McGhee, 25 delta in Wystup).[28] If:

W is a 2x1 column vector of European vanilla option weights;

M is a 2x2 matrix of European vanilla option ∂vega/∂vol and ∂delta/∂vol; and

E is 2x1 column vector of American binary option ∂vega/∂vol and ∂delta/∂vol,

then the hedge portfolio that eliminates (local) market risk is $W=M^{-1}E$, and the exotic market supplement is calculated by multiplying the weights (W) of the OTM Put and OTM Call options by their respective market supplements (Lipton and McGhee, 2002, p. 82 has a complete description and worked example of this process).

This intuitive approach was pioneered by market practitioners (Savery, 2000; Famery and Cornu, 2000) because it is internally consistent with the market price benchmark for European vanilla options which is universally accepted in practice, and it clearly expresses the hedge portfolio in terms of market-traded instruments. However, to date, heuristic models have not lived up to their promise. Reconciling the price of volatility convexities defined by European vanilla options with corrections required for American exotic options has not been possible without resorting to weighting schemes such as empirics (Lipton and McGhee, 2002), touch probabilities (Wystup, 2003), and arbitrary constants (Wystup, 2006). To date, these published weighting schemes do not work for all American binary options. ∂Vega is a commercial vendor system which weights the pure convexity correction by a reverse-engineered scaling factor, which is a curious response borne out of need rather than logic, especially for products which are multi-dimensionally unstable like exotic options. In effect, ∂Vega needs a price to calculate the price.[29] Unlike its heuristic predecessors, the Trader Model solves this conundrum by defining market risk and the optimal hedge portfolio in terms of low- and high-order greeks at the expected stopping time of the American binary option, not at its expiry. Furthermore, the Trader Model does not eliminate local risk, but instead focuses on the net contribution to book risk of the American binary option, and how price-makers smooth these global risks in practice.

[27] It can be argued that transaction costs are transferred, rather than minimised, owing to the non-trivial establishment costs of the static hedge portfolio.

[28] 15 delta means a Call option with a delta of 0.15 and a Put option with a delta of -0.15. For FX options, the delta is usually the spot delta. Longer-dated instruments use forward deltas by convention. 10 delta and 25 delta pillars are always quoted directly in the market. 15 delta is quoted on request, but is more usually obtained by interpolation, which makes that delta pillar subject to the interpolation constraints imposed by the financial engineer. This is a source of model risk.

[29] ∂Vega is also undermined by the fact that it violates no-arbitrage boundaries, such that it is possible to extract riskless profits by trading exotic options with its users. Therefore, it is not used as a pricing performance benchmark in this thesis.

Implied probability models

It is possible to reverse-engineer market, as distinct from theoretic probability densities from option prices (Breeden and Litzenberger, 1979). While it is possible to re-specify the underlying asset price diffusion to factor in the market inspired implied probability, it suffers from the same problem, namely, information from European option prices being used to make inferences about American option prices. As noted in Jex et al. (1999, p. 5):

> "European option [market] prices depend only upon the expected distribution for the asset value at the maturity of the option, and as such provides information about how this distribution differs from the [BSM] lognormal distribution. The smile does not directly provide information about the process that leads to this non-lognormal distribution . . . a number of different processes could be postulated which would match the observed volatility smile and yet give different values for the same path dependent [exotic] option".

Implied probability models use the probability structure for the nominal duration to re-specify the underlying asset price dynamics. Since smiles imply distributions, and smiles are different for each maturity, time dependent parameters are required to incorporate all the information in the volatility surface. Brigo and Mercurio (2000, 2002) propose explicit asset price dynamics that are consistent with a given parametric risk-neutral (forward) distribution. While they offer some flexibility, it comes at the non-trivial cost of an arbitrarily large number of non-market traded parameters, potentially causing slow computation times and overfitting problems as well as providing no financial economic insights into hedging in the market.

Stochastic volatility models

Stochastic volatility models are the most popular in the literature for modelling the market price of exotic options (e.g. Hull and White, 1987; Melino and Turnbull, 1990; Stein and Stein, 1991; Heston, 1993; Derman, 1998; Papanicolaou and Sircar, 1999; Britten-Jones and Neuberger, 2000). The principle of volatility varying stochastically is plausible, and the dynamics of stochastic volatility models are more realistic in that smiles persist. But very little is known about volatility dynamics. Hull and White's (1987) process for the instantaneous variance of spot FX returns[30] follows a geometric Brownian motion to ensure that variance is strictly positive. Stein and Stein's (1991) mean-reverting Ornstein-Uhlenbeck process captures the important real-life characteristic that variance mean-reverts, but at the considerable cost of allowing variance to be negative, which is counter-intuitive. Heston's (1993) mean-reverting Feller (square-root) process ensures that variance not only mean-reverts but is also strictly positive (Eq. (2.6)).

$$\begin{aligned} dS_t &= (r_d - r_f)S_t dt + \sqrt{\upsilon}S_t dW_t \\ d\upsilon &= \kappa(\theta - \upsilon)dt + \varepsilon\sqrt{\upsilon}dZ \end{aligned} \qquad (2.6)$$

However, Heston (1993) is not unique or even optimal but one of a family of plausible stochastic volatility models. This is problematic because:

[30] Instantaneous variance equals the instantaneous volatility squared.

- Plausibility is measured by reference to time-series of historic spot FX rates and / or current European vanilla option prices, and both sources are unstable. For example, Derman (1999) shows that volatility regimes can be highly temporally unstable, such that the actual regime is unknown at present and can only be established with hindsight. Since discovery only occurs when it is no longer useful, there is a real risk that prices obtained via a plausible process calibrated under one regime will become less plausible and possibly even implausible under another regime.
- Different stochastic volatility dynamics return identical European vanilla option prices by construction, but different exotic option prices. Theory cannot a priori offer guidance as to which dynamics are better for pricing an exotic option at present. Given that prices calculated by stochastic volatility models are highly sensitive to the arbitrary choice of volatility dynamics, this problem has great economic significance, not just mathematical inconvenience.

Non-traded parameters of the volatility process such as long-run variance (θ), mean reversion speed (κ) and volatility of variance (ε) are fitted to the traded volatility surface via numerically cumbersome and intensive calibration routines. The volatility surface does not define these factors in the European vanilla option market, so it is tenuous at best to use the surface to define these parameters in the American exotic option market.[31] While fitting is possible, Andersen and Andreasen (2000, p. 232) note that it often requires "unrealistically high negative correlation between the [underlying asset] and [its] volatility". This is because "stochastic volatility is modelled as a diffusion and hence only allowed to follow a continuous sample path, [therefore] its ability to internalize enough short-term kurtosis and thus to price short-term options properly is limited" (Bakshi, Cao and Chen, 1997, p. 2005). Using statics in the present (volatility surface) to define dynamics in the future (arbitrary volatility diffusion) is questionable, and it is only really done because it is preferred to the alternative of using historic time series, even though "there are no obvious relationships between market and model parameters which makes estimation of model parameters difficult to verify" (Hakala and Kirch, 2002, p. 249). Stochastic volatility models typically over-price ∂vega/∂vol, causing exotic option model prices to diverge significantly from market prices (Jex, Henderson and Wang, 1999; Lipton and McGhee, 2002). Fig. 2.3 shows the magnitude of this over-pricing for EUR DNT options. In Fig. 2.3, the model price consistently exceeds the market's ask price, which is unsustainable for a high volume, low margin exotic option price-making business.[32]

The price-maker's risk is also not described in terms of market-based hedging strategies using market-traded instruments. In fact, hedging is an afterthought, considered only after the price has already been calculated. Given that there is no whole-of-volatility-surface hedge traded in the market, price-makers are by necessity forced to trade a subset of the surface to hedge risk. Stochastic volatility models are silent on how to reconcile prices derived from a non-traded whole-of-volatility-surface, to the cost of hedging with a subset of traded

[31] The volatility surface is defined by the term and strike structures of volatility, which are based on zero delta straddles, butterflies and risk reversals. Price-makers do not consider θ, κ, ε or any other arbitrary non-traded parameter prior to making prices in straddles, butterflies or risk reversals.

[32] Using the same bid-ask spreads as Lipton and McGhee (2002), for $DNT_{TV} = 0.20$, a price-maker can give a stochastic volatility model bid price, and simultaneously lift a market ask price, and earn a riskless immediate profit of $5,000 per $1,000,000 payout.

options.[33] Because volatility is stochastic and not perfectly correlated with the underlying asset price's geometric Brownian motion, it is not possible to hedge perfectly all risk with the underlying asset alone, another option must be introduced.

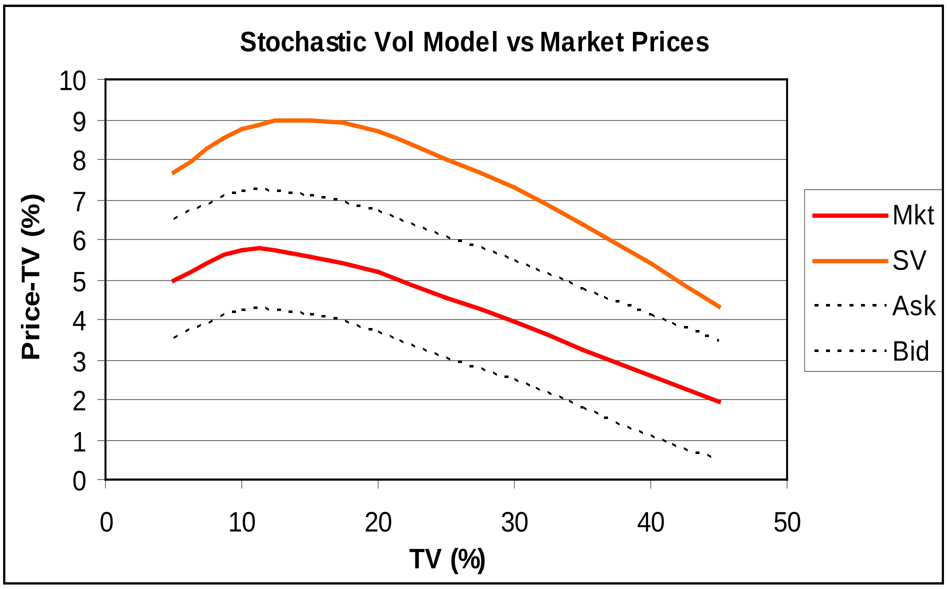

Fig. 2.3. Stochastic volatility model prices versus market prices for DNT options reported in Lipton and McGhee (2002). SV is the stochastic volatility model price subtract the theoretical value (TV), and Mkt is the actual traded market price subtract the theoretical value. If DNT options actually traded at theoretical value (i.e. the market supplement is zero), Mkt would plot on the x-axis.

Jump diffusion models

Models based on jump diffusion processes for the underlying asset have been explored by Merton (1976), Andersen and Andreasen (2000), Martinez and Senge (2002), among others. Martinez and Senge's (2002) jump diffusion process is shown in Eq. (2.7). Discontinuous jump processes have some theoretical appeal and also some empirical support, for example, they model steep skews in short-dated European vanilla options particularly well. However, "like stochastic volatility models, jump diffusion models are challenging to handle numerically" (Andersen and Andreasen, 2000, p. 232), since the calibration of non-traded jump diffusion parameters such as jump intensities (λ) to the volatility surface is notoriously difficult. Jump models are incomplete, in that all risk cannot be eliminated by dynamic delta hedging with the underlying asset alone. Like stochastic volatility models, jump models also do not give intuitive insights into how to hedge the resulting market risk.

$$dS_t = (r_d - r_f - \lambda\kappa)S_t dt + \sigma S_t dW_t + \gamma_t S_t dN_t \qquad (2.7)$$

Other authors, such as Matytsin (1999) and Duffie, Pan and Singleton (2000) propose dynamics which include jumps in volatility as well as jumps in asset price. The main reason for adding this non-trivial layer of complexity is to achieve a closer degree of fit to the volatility surface through the introduction of additional free parameters. In addition to the risk of over-fitting, Duffie et al. (2000, p. 1365) shows that "adding jumps in volatility may attenuate . . . the overpricing for [European vanilla] options that are not too far out-of-the-money", but "actually exacerbates the overpricing for far out-of-the-money [European vanilla options]". Unfortunately,

[33] Whilst each option underlying the volatility surface is traded, it is not possible to trade them all at once in a way that locks-in the model calibration.

errors in the pricing of European vanilla options compound in the pricing of exotic options. Matytsin (1999) shows that adding volatility jumps leads to very different exotic option prices compared to other models calibrated to the same volatility surface.

Universal volatility models

Bates (1996), Jex et al. (1999), Blacher (2001) and Lipton and McGhee (2002) are examples of what are now called universal volatility models. The Lipton and McGhee (2002) model is shown in Eq. (2.8).

$$\begin{aligned} dS_t &= (r_d - r_f) S_t dt + \sqrt{\upsilon}\sigma_L(S_t, t) S_t dW_t + \left(e^j - 1\right) S_t dN_t \\ d\upsilon &= \kappa(\theta - \upsilon)dt + \varepsilon\sqrt{\upsilon} dZ \end{aligned} \tag{2.8}$$

Universal volatility models are an arbitrary mix of local, stochastic and / or jump models. Mixing processes is theoretically questionable but appears to have some merit on pragmatic grounds, in that if local volatility underprices ∂vega/∂vol and stochastic volatility overprices ∂vega/∂vol, then a combination of the two ought to generate more realistic exotic option prices. Fig. 2.4 shows that this is the case in Lipton and McGhee (2002). Nevertheless, calibration and hedging difficulties remain, especially where jump diffusion is included, as there are now even more non-traded parameters in the diffusion, and "such a model would . . . be difficult to handle numerically and slow to calibrate accurately to traded prices" (Andersen and Andreasen, 2000, p. 233). Taleb's (1997, p. 109) counsel that "it is better to use a model with the smallest number of parameters to estimate" is ignored. No matter how complex the guess of the volatility dynamics, it still remains a guess of the exotic option price because "smile models (e.g. local volatility, jump-diffusion, stochastic volatility, etc.) may agree on the vanilla prices and totally disagree on the exotic prices and the hedging strategies" (Ayache et al., 2004, p. 1). Theory gives us little to no guidance, especially given Derman's (1999) unstable volatility regimes, as to how to differentiate from among the many competing volatility models which model will be best for pricing an exotic option at present.

Universal volatility models, like their constituent local, stochastic and jump elements, are silent on the issue of hedging. In fact, model prices are obtained independently of hedge construction. Hedging is an afterthought in these models because hedging is by necessity conducted with traded instruments in the market, whereas universal volatility model risk is defined in terms of non-traded parameters. This fundamental disconnect between model price and market hedging cost is a significant source of model risk which undermines both the quality of reported profits and their sustainability over the medium-term.[34] Calibration to the volatility surface does not overcome this problem, because price-makers hedge with a small subset of European vanilla options, not the entire surface.[35]

Furthermore, calibration has a non-zero error metric. That is, even though the traded volatility surface is arbitrage-free, the universal volatility model will mis-price European vanilla

[34] Sell-side OTC exotic option trading desks report asymmetric profits. Exotic options are marked-to-model, and the vanilla options used to hedge them are marked-to-market. If exotic options are marked to the universal volatility model, then the profit and loss so calculated bears no relation to the cost of hedging them. Price-makers hedge high-order greek risks with (vega neutral) butterflies and (delta neutral) risk reversals, irrespective of which model they use to price. Jump intensity, mean-reversion speed, etc. do not quantify the costs of these hedges. This is a significant problem as exotic option trading desks are a low-margin, high volume business.

[35] There is no whole-of-volatility-surface hedge traded in the market. Therefore, traders by necessity must use a subset of traded European vanilla options to hedge exotic options.

options because the calibration is not perfect. This is one reason why the universal volatility model is not used by price-makers to price European vanilla options to market. Given that errors in the vanilla surface compound into large errors in exotic option prices, it is curious that exotic option price-makers are prepared to use a model that cannot price European vanilla options without error.

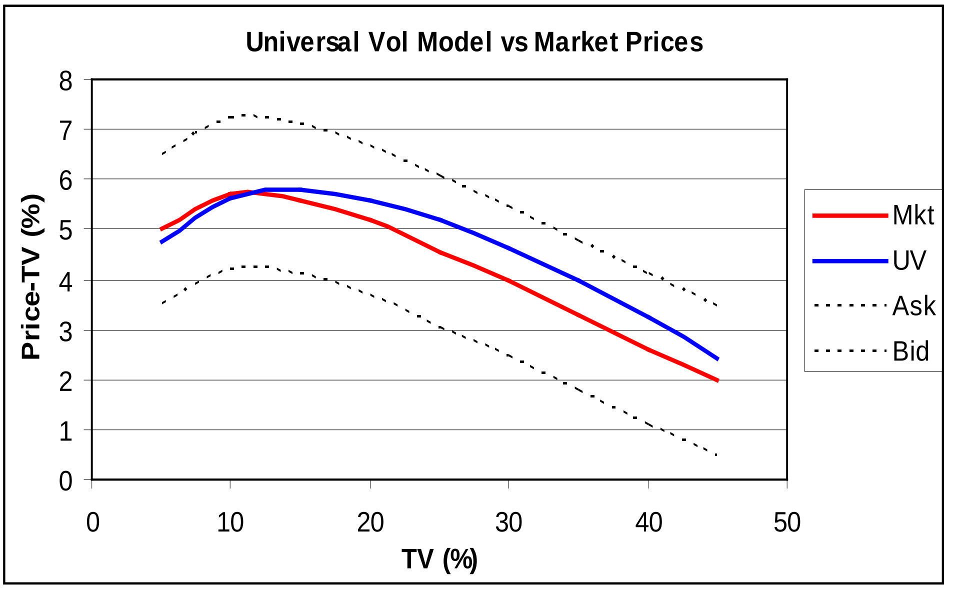

Fig. 2.4. Universal volatility model prices versus market prices reported in Lipton and McGhee (2002). UV is the universal volatility model price subtract the theoretical value (TV), and Mkt is the actual traded market price subtract the theoretical value. If DNT options actually traded at theoretical value (i.e. the market supplement is zero), Mkt would plot on the x-axis.

Lipton and McGhee (2002, p. 85) conclude that "while both local and stochastic volatility models produce price corrections in qualitative agreement with the market, only a universal volatility model is capable of matching the market properly". Jex et al. (1999) and Ayache et al. (2004) arrive at similar conclusions. Table 2.2 is a summary of contemporary exotic option modelling methods.

Table 2.2
Summary of contemporary exotic option modelling methods.

Local Volatility Models
Good Easy to implement. Returns the market price of European vanilla options.
Bad Ill-posed inverse problem. Interpolate and extrapolate sparse data to define the deterministic functional form of σ_L. Conditional probabilities can be negative.
Severe Smile dynamics are incorrect. Underprices the smile effect in exotic options.
Stochastic Volatility, Jump Diffusion and Universal Volatility Models
Good Plausible smile dynamics. Returns the market price of European vanilla options up to a closeness-of-fit error metric.
Bad Arbitrary guess of the functional form of the volatility process (there are entire families of stochastic vol, jump diffusion and universal vol models). Calibrating to the volatility surface is theoretically tenuous and computationally inefficient. Fitting is poor for short-maturity options (stochastic vol) and extremely cumbersome computationally (jump, universal).
Severe Guessing the volatility process is the same as guessing the price. The model price is highly sensitive to the arbitrary choice of the functional form of volatility process. The model price is highly sensitive to the arbitrary choice of calibration method. Stochastic vol overprices the smile effect in exotic options.

2.3. Empirical research

2.3.1. Introduction

Empirical research in exotic options is constrained by the limited availability of traded market input and output data. As a result of this binding constraint, there are few examples of market-grounded, empirical research in the literature. In this thesis, published empirical research is classified according to the two central themes of (i) model pricing accuracy and (ii) the relative efficiency of computational methods. Only pricing accuracy is directly relevant to this thesis. The Trader Model does not employ computational methods synonymous with contemporary competitor approaches, namely parameter estimation and calibration, thereby making the conventional search for marginal computational efficiencies redundant.

2.3.2. Pricing

Jex, Henderson and Wang (1999) and Lipton and McGhee (2002) are two examples of published exotic option pricing research that are particularly relevant to this thesis, owing to their exclusive focus on American binary FX options and their empirical testing is grounded in market traded data. Jex et al.'s and Lipton and McGhee's empirical research have many similarities, in that model pricing performance is analysed against actual traded market prices as well as competitor model benchmarks for:

- one option maturity;
- one currency pair;
- one binary option type; and
- one valuation date.

Both of these papers are handicapped by small data sets, which means that they can conduct only a very simple analysis of model pricing performance. Both papers graphically plot model price outputs against traded market prices and competitor model benchmark prices, they do not report any descriptive statistics or other performance measures. The inference is that because their model prices plot closer to market prices than competitor model benchmarks, then their pricing performance is better.

This inference may be misleading. It is bad practice to generalise from very specific market scenarios. If only one option maturity, one currency pair, one binary option type and one valuation date is analysed, one cannot ascertain whether pricing performance is robust to a change in maturity, currency, option type and / or valuation date. In addition, Jex et al. and Lipton and McGhee both chose competitor model benchmarks which were known in advance to be deficient. Setting a low performance standard undermines the rigour of their analysis and, therefore, makes one less confident in the conclusions drawn.

In contrast, the scale and scope of the data in this thesis is much more comprehensive than both Jex et al. and Lipton and McGhee (refer to section 3.5 'Data Description'), such that the analysis is more rigorous and the conclusions more robust.[36] Furthermore, the competitor model benchmark chosen in this thesis is widely used in the financial markets, has won numerous industry based awards and, according to a prominent, high-profile financial engineer, has "become a standard reference for pricing exotic FX options up to the market" (Wystup, 2003).

2.3.3. Computational methods

Schoutens, Simons and Tistaert (2004) empirically tested, among other things, the impact of process choice on exotic option prices. This is research with considerable economic significance given that, as noted by Avellaneda, "guessing the volatility process is the same as guessing the price" (cited in Ayache et al., 2004, p. 11). Schoutens et al. found that the initial 'guess' of the volatility process resulted in extraordinarily large price variation for exotic options relevant to this thesis (Table 2.3). They additionally found that price variation owing to the arbitrary choice of calibration method was as much as 13%. Given that American binary FX options are quoted to five decimal point accuracy in the interbank market, price variation of this magnitude is unacceptably large.

[36] The quality of the data in this thesis also appears to be of higher quality than Jex et al. and Lipton and McGhee. In this thesis, every exotic option market price actually traded in the interbank exotic FX option market. In contrast, Jex et al. report the market prices for eight JPY up OT options, and eight JPY down OT options. All 16 OT options had three month maturities. It is extremely unusual for 16 JPY OT options with identical maturities to trade in the interbank market in the same day, let alone at the same time. Hence, at best there is a non-synchronous data problem, and at worst, the 'market' prices did not actually trade but were merely quoted by different counterparties, which means they are only model prices not market traded prices. Lipton and McGhee have the same issue, except for EUR DNT options. It is not known how many data points were used in Lipton and McGhee, because output data were disclosed as continuous lines, not discrete points.

Table 2.3
Process variation reported in Schoutens et al. (2004).

Exotic Option	Percent Change in Price Min. (%)	Max. (%)
One Touch	4.0	16.2
Reverse Knockout	34.2	512.5
Reverse Knockin	0.6	14.7

Percent change is defined as (Max-Min)/Min. Reverse knockout and knockin options are barrier options which contain an implied OT option, and hence, are directly relevant for this thesis. Schoutens et al. reported results for reverse barrier Call options only. Source: Schoutens et al. (2004).

Detlefsen and Hardle (2006, p. 25) focused on calibration risk, and concluded that:

> “We have shown that different ways to measure the error between the model and the market in the calibration routine lead to significant price differences of exotic options in the sense that these differences often exceed the profit margins of the products”.

These model price variations are already too great for the market, in that they exceed by a significant margin the fine pricing tolerances established by interbank convention. Even worse, there are non-trivial computational risks which Schoutens et al. and Detlefsen et al. do not analyse. For example, they do not quantify the risk of price variation from arbitrary choices of parameter constraints or discretisation methods. It is unlikely that either of these sources of risk exceed process or calibration risk, however, they contribute to rather than offset model price variation. Table 2.4 shows some of the computational methods underlying contemporary models. This table is not meant to be exhaustive, but rather indicative of the arbitrary choices routinely made by financial engineers implementing stochastic volatility, universal volatility and jump diffusion models in practice. Each choice contributes to the variation in model prices by artificially constraining the size of the solution space. If the price variation were small, it would not matter that there are several arbitrary choices for each step of the model implementation. However, the first column alone is where Schoutens et al. (2004) found 200% variation in model prices, making prices highly model-dependent. Theory cannot differentiate ex ante which model variant is ‘best’, and, given Derman’s (1999) unstable volatility regimes, what is ‘best’ in the past is unlikely to be ‘best’ in the present or remain ‘best’ in the future.

Table 2.4
Summary of computational methods underlying contemporary models.

Process	Calibration Method	Error Functional	Minimisation Algorithm	Parameter Constraints	Discretisation Method
Heston	Historic Time Series	Absolute Price	Differential Evolution	'Hard' Range	Finite Element
Blacher	Vol Surface	Relative Price	Gradient Descent	'Soft' Convex Penalty	Finite Difference
Hull	Vol Surface & Exotics	Absolute Volatility	Levenberg-Marquardt	Time-Dependent	2D-Trinomial
Lipton		Relative Volatility	Simulated Annealing	Time-Independent	Monte Carlo
Bates			Method of Moments	Piecewise Constant	Quasi-Monte Carlo
Stein and Stein			Maximum Likelihood	Equal Weights	Wavelets
Merton			Downhill Simplex	Vega Weights	

Each column contributes to the variation in model prices. The first column is where Schoutens et al. (2004) found 200% variation in model prices. The table should be read as independent columns, the rows do not articulate. This table is not meant to be exhaustive, but rather indicative of the arbitrary choices made by financial engineers implementing stochastic volatility, universal volatility and jump diffusion models in practice.

2.4. Conclusion

This research subscribes to the view that an imperfect but simple and transparent model is more useful for price-makers in practice than over-engineered opaque models, especially when markets are unstable and / or undergoing structural change. This is consistent with the position taken by Taleb (1997) and Derman (2000) which favours pragmatism, flexibility and responsiveness to market conditions, rather than over-reliance on the mathematical methods of the physical sciences. It is through the identification of unique optimal hedging portfolios consisting of market-traded instruments, that this research renders redundant both the traditional approach of 'guessing' plausible volatility dynamics and solving for exotic option prices numerically, with little to no regard for hedging; and the more modern approach of using hedging considerations to 'guess' the specification of the dynamic processes for the underlying asset and its volatility (Ayache et al., 2004; Blacher, 2001). This research focuses on the financial economics of the problem because "financially relevant questions can only be answered by relevant financial theory", and because "the need to go back to 'basics' is a very welcome conclusion, to say the least, at a time when quantitative finance seems to be wasting itself in sophisticated mathematical exercise" (Ayache et al., 2004, p. 33).

CHAPTER 3

METHODOLOGY

3.1. Introduction

This chapter presents the aims of the research, the philosophical and methodological rationale underpinning the Trader Model, a description of and justification for the research methods used to test the Trader Model, and a description of the data obtained for empirical testing.

3.2. Research aims

The overall aim of the research is to develop a model for pricing exotic options to market which: (i) predicts actual traded market prices with sufficient accuracy to be a useful decision-making tool for price-makers in practice; (ii) identifies and quantifies market risk in a manner which provides unique insights into risk management of exotic options for price-makers in practice; and (iii) achieves real savings in computational efficiency relative to best practice quantitative models exemplified by the universal volatility models of Jex et al. (1999) and Lipton and McGhee (2002).[37] The model will specifically:

- be accurate and independent of arbitrary specifications of volatility dynamics;
- define the market risk of exotic options in a manner useful for hedging in the market in practice and which is also independent of arbitrary specifications of volatility dynamics; and
- identify unique option-level hedging strategies which are simple portfolios of common, market-traded instruments.

This research will be a significant and unique contribution to knowledge in exotic option pricing and hedging. Just as Cox, Ross and Rubinstein (1979) provided an insightful interpretation of BSM which, in its day, demystified complex mathematical concepts for financial economic theorists and practitioners, it is expected that this research will likewise demystify the pricing of exotic options to market for a broader audience beyond financial engineers. This research is positioned squarely at the nexus between theory and financial market practice and, like Bakshi, Cao and Chen (1997, p. 2004), it asks:

> "While the search for that perfect option pricing model can be endless . . . what do we gain from each generalized feature? Is the gain, if any, from a more realistic feature worth the additional complexity or implementational costs"?

The Trader Model which is developed and discussed in greater detail in Section 3.3., is not intended to be "that perfectly specified option pricing model [which] is bound to be too complex for applications" (Bakshi et al., 1997, p. 2004), but instead, to be a simple model specifically designed for application in the market by price-makers, where pragmatism and flexibility are preferred to the idealistic pursuit of perfection in imperfect markets. The Trader Model will be tested to see how well it can explain market traded exotic option prices, and the

[37] Ayache et al. (2004, p. 15) states that "only the universal volatility model . . . manages to fit the smile dynamics and at the same time to fit the barrier option prices".

extent to which it can identify unique market traded optimal hedging strategies which reduce the market risk of running a book of exotic options.

Contemporary exotic option pricing research in academia and industry is crucially dependent upon arbitrary specifications of volatility dynamics, non-market traded parameters and / or theoretically baseless constants, and, as a result, it introduces unnecessary complexity and inefficiency into solutions, as well as significant model risk. In the Trader Model, market prices and optimal hedging strategies are dependent upon nothing other than market spot and deposit rates, European vanilla option volatilities and exotic option contract specifications. Therefore, intermediate calculations, such as numerically extensive and intensive calibration schemes that are a prominent feature of competitor methodologies, are redundant in the Trader Model. As the theoretical basis of the Trader Model is internally consistent with the universally accepted benchmark methodology for pricing European vanilla options to market, model risk is also minimised.

3.3. Trader Model

3.3.1. Philosophy

The Trader Model is a radical philosophical departure from contemporary exotic option pricing orthodoxy. Contemporary methods, such as universal volatility models (Jex et al., 1999; Blacher, 2001; Lipton and McGhee, 2002), rely on and make a virtue of layer upon layer of advanced mathematics and technological brute-force to compute prices. In contrast, the Trader Model builds on fundamental financial economic intuition and has modest computational requirements.

The central tenet guiding the development of the Trader Model was Derman's (2003, p. 13) claim that "all financial models are wrong". Bearing this in mind, along with the fact that price-makers in practice use model outputs as inputs in the price-making process, it follows that a simple approach with less structure and fewer binding and unverifiable assumptions would be intuitively appealing. Price transparency is just as important as price discovery for a model to be the "right tool" for decision-making for price-makers in practice (Ayache et al., 2004, p. 36).

The Black-Scholes-Merton (BSM) model and its exogenous volatility surface is a classic example of simplicity and pragmatism. It is well known that it is wrong to use the BSM model outside of its original theoretical context. Nevertheless, BSM is used in practice because price-makers understand how to modify the BSM model with the volatility surface to price factors crucial to the market, and that is why BSM and its volatility surface together constitute the universally accepted benchmark for pricing European vanilla options to market, dominating stochastic, jump and universal volatility models. "Traders . . . are comfortable with [the BSM model] because they have learned the necessary tricks to make it work" (Taleb, 1997, p. 109). The Trader Model is dependent upon the same structure as BSM, and it has a unique mechanism for applying the same trader-inspired 'necessary tricks' for European vanilla options to the computation of market prices for exotic options in general, and, in this thesis, for American binary FX options in particular. In this sense, the Trader Model fills the void between heuristic financial economic approaches of the past and the mathematically laden style of orthodox methods.

The contrarian philosophy of focusing on the fundamental financial economics of exotic option pricing by explicitly taking into account price-makers' pivotal role in price discovery in traded markets is gaining support, even among leading financial engineers. For example, Ayache (2004b, p 29) advocates a new theoretical representation based on "the pair composed of the trader and his knowledge of the [model]".[38] Bates (2003, p. 400) expresses a desire for "a renewed focus on the explicit financial intermediation of the underlying risks by option market makers" and for "plausible models of market maker behaviour". And Derman (2002, p. 82) observes that "'return' and 'volatility' lie in the realm of quantitative finance, but 'expected' and 'implied' lie in the domain of behaviour. A future theory that married the quantitative to the behavioural would be a worthy goal". The Trader Model is not a theory, but it marries quantitative methods to actual trader behaviour more closely than any other exotic option pricing model.

In addition to this growing theoretical support, there is also compelling practical support. If market prices are obtained with one basis (arbitrary non-market traded parameters) and hedging is by necessity conducted with a different basis (a subset of traded instruments), then the fundamental economic link between market price and hedging cost is diluted, if not lost, crucially affecting business line profitability given the low-margin, high-volume nature of bank exotic FX option trading businesses.

3.3.2. Method

The Trader Model is dependent upon the same structure as implementing BSM in the market. This thesis focuses on the FX market and the assumed dynamics for the underlying spot FX rate take the traditional functional form of Eq. (3.1), which is used in conjunction with the exogenous traded volatility surface for European vanilla options $(\sigma = \sigma[\Delta, T])$.

$$dS_t = (r_d - r_f)S_t dt + \sigma S_t dW_t \tag{3.1}$$

In academia and industry these dynamics are used to obtain theoretical values and market prices for European vanilla FX options. Although they are also used for deriving the theoretical value of exotic FX options, they are considered by financial engineers to be inadequate for pricing exotic FX options to market. This inconsistency is not explained by theory. It is simply because:

(i) to date, no one has been able to price exotic options to market with these parsimonious dynamics; and

(ii) the existence of non-zero calibration error means that if an orthodox model is calibrated to BSM theoretical conditions (a flat volatility surface and flat yield curves), it will not return the BSM theoretical value, and hence, the market must use BSM dynamics to calculate the theoretical value.[39]

[38] The original quote uses the term "weapon", not model, which is a reference to Haug's two part series in Wilmott Magazine called "Know Your Weapon", which was all about the Black-Scholes-Merton model.

[39] Price-makers will universally reject an exotic option pricing model that cannot return the theoretical value of an exotic option, because the theoretical value must be agreed between price-makers for all interbank exotic FX option trades, before market price negotiation between counterparties even commences. The coexistence of two different processes to describe the dynamics of the same FX pair is theoretically questionable.

Therefore, to price exotic options to market, orthodox models introduce additional structure via increasingly more esoteric and complex assumptions about volatility dynamics which may improve in-sample quality of fit, but usually at the expense of divorcing the model from traded market reality. In contrast, in rejecting the need for any additional structure, this thesis is consistent with Derman's recommendation that "one good strategy in attempting to value exotic options that are sensitive to the smile is to try to avoid modeling the dynamics of volatility as much as possible" (2003, p. 13).[40]

The existence and uniqueness of exotic option theoretical values and their widespread use by interbank price-makers as a convenient and essential reference point for quoting market prices, reduces the dimension of the exotic option market pricing problem. One only needs to price the difference between market price and theoretical value, known as the market supplement. For FX options with expiries of twelve months or less, which is the focus of this thesis, volatility effects dominate interest rate effects (Taleb, 1997; Hakala and Wystup, 2002). Since the market supplement for European vanilla options is the set of smiles and skews for all expiries, the valuation of the market supplement for exotic options reduces to the problem of quantifying smile and skew effects.

Unlike contemporary approaches, the Trader Model preserves the direct link between smiles and skews and the size of the market supplement for exotic options. Smiles and skews are not used to define non-traded parameters like jump intensity and mean reversion speed in arbitrarily specified volatility dynamics. The information contained in the volatility surface is static and specific, and it is theoretically tenuous to infer dynamics from statics and to generalise from specifics. It is common in academia and industry to infer and generalise in this way because:

- contemporary exotic option pricing research is evolutionary not revolutionary, because existing models dependent on volatility dynamics are refined rather than redefined;
- the alternative of calibrating volatility diffusion parameters to historic time-series is discredited, even though the current practice of calibrating the dynamic evolution of future volatility in a model by fitting to static smiles and skews in the present is not any better (refer to Subection 2.2.1.);
- Ayache et al.'s recommendation (2004, p. 11) to calibrate to the European vanilla volatility surface and exotic option prices is not practical for first generation exotic FX options,[41] because one effectively needs exotic option prices to find an exotic option price, something which is difficult to achieve in the highly competitive interbank market where heterogeneous OTC exotic products are traded;[42] and

[40] Derman's reference to the smile also includes the skew. A skew is simply an asymmetric smile.
[41] First generation exotic FX options include binary and barrier options, and they represent approximately 95% of traded volume.
[42] A model that is dependent on prices for exotic options traded between other banks is problematic for a price-maker because: (i) the trader is reduced to being a price-taker not a price-maker; and (ii) traded prices reflect the shape of the exotic option books of the banks that are counterparties to the trade, not the trader's own book.

- modern computing power means it is possible to solve significant computational problems with advanced numerical techniques within a reasonable time frame, irrespective of whether it is optimal or even desirable to do so.

In the European vanilla FX option market, smiles are priced by (vega neutral) butterflies and skews are priced by (delta neutral) risk reversals. Table 3.1 shows that butterflies are essentially neutral to all exposure other than ∂vega/∂vol, and risk reversals are essentially neutral to all exposure other than ∂delta/∂vol. Therefore, smiles price ∂vega/∂vol and skews price ∂delta/∂vol, and to use them to price anything else introduces model risk.[43] The Trader Model preserves the one-to-one mapping by valuing exotic option ∂vega/∂vol and ∂delta/∂vol with smiles and skews, respectively. Hence, what is already done and universally accepted as the market pricing benchmark for European vanilla FX options is simply extended to the American binary FX option market. Therefore, the Trader Model achieves what others' have not, a unification of the market pricing and risk management framework for European vanilla and American exotic options. This unification is crucial because "it's safe to say that there is no area where model risk is more of an issue than in the modeling of the volatility smile" (Derman, 2003, p. 4).

Table 3.1
The market risk of zero delta straddles, vega neutral butterflies and delta neutral risk reversals in the FX option market.

Risk	ZD Straddle	VN Fly	DN RR
Delta	0	0	0
Vega	>> 0	0	≈ 0
Volga	0	>> 0	≈ 0
Vanna	0	≈ 0	>> 0

Zero delta straddles, vega neutral butterflies and delta neutral risk reversals are constituent elements of the traded volatility surface. These commoditised trading strategies have (almost) mutually exclusive market risks. A vega neutral butterfly is a source of volga, and in the market the vega neutral butterfly is quoted as the volatility spread between the OTM options and the ATM options, which is the smile. A delta neutral risk reversal is a source of vanna, and in the market, the delta neutral risk reversal is quoted as the volatility spread between the OTM Call option and the OTM Put option, which is the skew. Therefore, smiles price volga, and skews price vanna. Zero delta straddles, which have zero volga and zero vanna, also have zero smile and skew adjustment. Therefore, zero delta straddles define the level of volatility. The data is for illustrative purposes only, and it is for a long straddle, long butterfly and a risk reversal with a long Call / short Put.

Extending the BSM and exogenous volatility surface market pricing benchmark to exotic options has been tried before by practitioners and has failed (Famery and Cornu, 2000; Savery, 2000; Lipton and McGhee, 2002; Wystup, 2003 and 2006; and commercial vendor system ∂Vega). It failed because practitioners were unable to reconcile European volatility surfaces with American optionality. Famery and Cornu, and Savery describe micro-mechanics but are silent on this crucial issue. Lipton and McGhee attempted reconciliation by empirically scaling the convexity correction, Wystup scaled by touch probabilities and / or constants, and ∂Vega scaled by a reverse-engineered 'theoretical value adjustment' matrix supplied by the GFI broker desk. All of these scaling methods are unsatisfactory:

[43] Since OTC exotic FX options are contractually heterogeneous, they are marked-to-model not marked-to-market for profit and loss, bonus, limits and reporting purposes. Model risk is a big focus of Basel II capital adequacy requirements.

- *Empirical scaling*. Empirically derived scaling factors are dependent on the arbitrary choice of sample period and sample frequency. A different sample period and / or sample frequency will produce a different scaling factor and hence, a different exotic option price. Since theory cannot distinguish ex ante which sample period and frequency is best, and Derman (1999) has shown that volatility regimes change over time, empirical scaling is a poor choice for price-makers in multi-dimensionally unstable markets like exotic FX options.
- *Scaling by touch probabilities*. Touch probabilities are theoretically more appealing as they are dependent only upon market and option contract inputs, and solutions for American binary touch probabilities are available analytically. While scaling by touch probabilities is presented by Wystup as a reasonable proxy for obtaining the market price of a binary option with a single barrier (OT), the method fails even for the modest extension of pricing to market binary options with two barriers (DNT). Clearly, touch probabilities do not capture the essence of the market mechanism.
- *Reverse-engineering a scaling factor*. The ∂Vega commercial vendor system is dependent upon a matrix of 'theoretical value adjustments'.[44] Theoretical value adjustments are reverse-engineered from exotic option market prices traded through the GFI broker desk. In other words, an adjustment factor is artificially changed until ∂Vega reproduces the actual traded market price for the exotic option. This method is a poor alternative. Firstly, a price must be quoted in the interbank broker market before a price-maker can obtain the adjustment coefficient and the price. This is a luxury not available to price-makers in practice.[45] Secondly, the adjustment coefficient is multi-dimensionally unstable, rendering it meaningless for options which differ contractually (a feature of heterogeneous OTC exotic FX options), and for markets which change over time (all markets). Casual observation of Fig. 3.1 shows that for EUR DNT options, theoretical value adjustments are unstable. For a TV of about 0.035, actual theoretical value adjustments ranged between 0.22-0.72, resulting in significantly different ∂Vega model prices. This range is especially large when one considers that in the interbank market American binary FX options are priced to five decimal point accuracy. Ex ante, a price-maker cannot know whether 0.22 or 0.72 or some other adjustment coefficient is accurate, as reverse-engineering is by construction ex post. The theoretical value adjustment is unstable by theoretical value and by expiry days, for all exotic FX options. There is no discernible pattern that can be exploited to produce reliable accurate market prices, which is why ∂Vega is unsuccessful and discredited in academia and industry. The Trader Model is not tested

[44] Refer to www.gfigroup.com for details.

[45] Quoting speed and aggressive spreads are a measure of a price-maker's professionalism in the interbank FX option market. Delay of even a few minutes to obtain a recently reverse-engineered adjustment coefficient would be considered unprofessional and undermine the credibility of the business. Furthermore, it is considered unprofessional to ask for reference prices without intending to deal.(Taleb, 1997, p. 57). A model that requires reference exotic option traded prices before a price can be made essentially reduces the price-maker to a price-taker.

empirically against ∂Vega because ∂Vega also has many violations of no-arbitrage boundaries, and produces prices which are mathematically impossible. It is not known whether the reverse-engineered scaling factors cause or contribute to arbitrage violations, or whether there is some other flaw in ∂Vega.

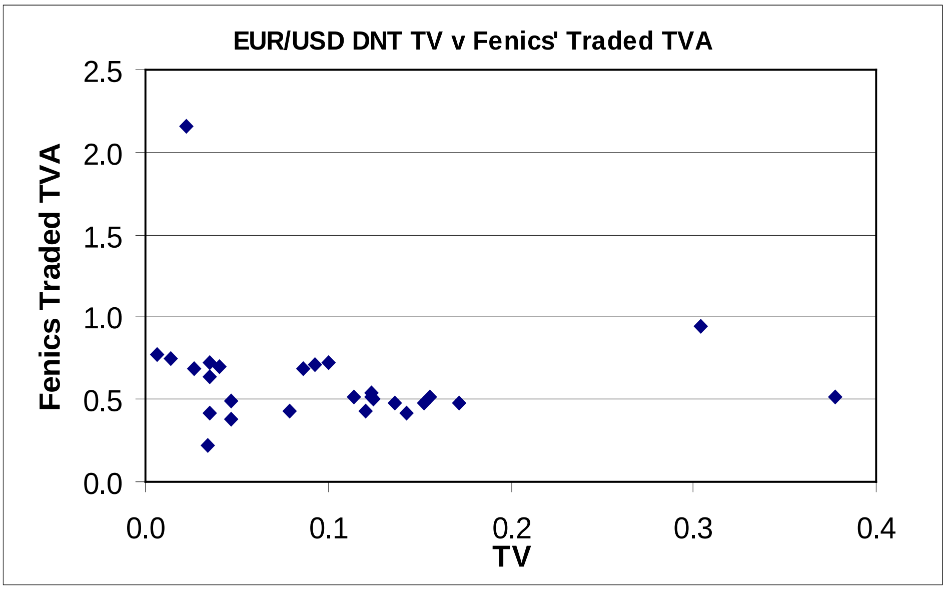

Fig. 3.1. ∂Vega theoretical value adjustments (TVA) for EUR DNT options. TVA are reverse-engineered scaling factors which ensure that the heuristic ∂Vega model price equals the actual traded market price. TVA's are derived ex post, but are required by price-makers ex ante. The figure is based on actual traded exotic FX option prices for the period 7 January to 29 September 2004. TV is the acronym for theoretical value. Source: GFI Group Inc.

- *Scaling by constants*. Scaling by arbitrary and theoretically baseless constants is an obvious weakness and a source of model risk as exotic option prices are multi-dimensionally unstable (e.g., American binary FX option prices are non-linear functions of, and hyper-sensitive to, changes in spatial, temporal and volatility dimensions). It is reasonable to conclude that the mere existence of arbitrary and theoretically baseless constants is evidence of the failure of that model to capture adequately prices traded in the market. An example of this overly simplistic approach is in Wystup (2006), where the recommended weighting for a DNT option is 0.5. Wystup notes that this weighting has no theoretical justification. As BSM contains no theoretically baseless constants, there is an immediate misalignment between exotic option model prices and their European vanilla option hedging costs, resulting in model risk.

The Trader Model, in contrast, develops a unique method for reconciling European vanilla volatility surfaces and American exotic optionality. Whilst American binary options have a known nominal duration, in most practical circumstances they are expected to terminate sometime prior. It is possible to calculate this expected stopping time analytically, and the Trader Model uniquely uses expected stopping time to scale the size of the exotic option's vega and delta convexities to volatility, to establish the expiry date of the relevant butterfly and risk reversal, and, as a result, to price American vega and delta convexities to volatility. Scaling by the expected stopping time is intuitive. A six month DNT option with a theoretical value of 0.10, has a touch probability of approximately 0.90.[46] This DNT option is extremely unlikely to survive until expiry and so any hedging strategy ought to have a duration which matches the expected termination date of the exotic option, otherwise the hedge portfolio will turn into an open position once the exotic option terminates. Expected stopping time is independent of empirics, reference exotic option prices from brokers and ill-conceived constants, and it works for all binary options. Expected stopping time is a function of option contract and market inputs only, and does not suffer from the rampant arbitrariness endemic to other scaling schemes. For example, Eq. (3.2) shows the expected stopping time for a single, continuously monitored option barrier derived in Taleb (1997, p. 476):

$$
\begin{aligned}
E\left(\tau_H^T\right) &= \frac{h}{\lambda} + \left(T - \frac{h}{\lambda}\right) N\left(\frac{h}{\sqrt{T}} - \lambda\sqrt{T}\right) \\
&\quad - e^{(2\lambda h)}\left(T + \frac{h}{\lambda}\right) N\left(-\frac{h}{\sqrt{T}} - \lambda\sqrt{T}\right)
\end{aligned}
\tag{3.2}
$$

$$\text{where } \lambda = \frac{\left(r_d - r_f\right)}{\sigma} - \frac{\sigma}{2};$$

$$\text{and } h = \frac{1}{\sigma}\ln\frac{H}{S_t}$$

and S_t is the spot FX rate, r_d (r_f) is the domestic (foreign) interest rate, σ is the volatility, H is a continuously monitored barrier $H > S_t$; and T is the annualised term to maturity of the option.

Appendix A contains schematic diagrams outlining how the Trader Model converts theoretical values into market prices for American binary FX options. Fig. 3.2-3.5 inclusive, show vega and delta convexities to volatility for representative DNT options and OT options, for illustrative purposes only. When an option has positive (negative) ∂vega/∂vol, it must trade at a premium (discount) to theoretical value, ceteris paribus, as its vega changes beneficially (detrimentally) with respect to volatility. When an option has positive (negative) ∂delta/∂vol, it must trade at a premium (discount) when the skew is positive, ceteris paribus, as its delta changes beneficially (detrimentally) with respect to volatility. This is common to Famery and Cornu (2000), Savery (2000), Lipton and McGhee's (2002) heuristic model, Wystup (2003), and ∂Vega. However, the Trader Model uniquely:

- Scales the size of the vega and delta convexities by the expected stopping time of the exotic option. For example, if a DNT option has an expected stopping time which is 0.4 of its nominal duration, then the vega and delta convexities are scaled by 0.4. If DNT option vega convexity to volatility is 200 points, then the vega convexity to volatility used in the Trader Model is 80 points. In an intuitive way, the Trader Model accounts for the American optionality in the exotic option.

[46] Assuming interest rates are zero, to illustrate the principle simply.

- Chooses the relevant unique butterfly and risk reversal to price the vega and delta convexities, respectively. If the exotic option has an expected stopping time which is 0.4 of its nominal duration, then the Trader Model uses a butterfly and a risk reversal with expiries equal to the expected stopping time to price the convexity. That is, if the DNT option has a nominal duration of 90 days, then the Trader Model will price vega and delta convexities with a butterfly and a risk reversal with expiries of 36 days. As a result, the exotic option and its hedge terminate at the same time, there is no unwanted residual hedge as per other competitor models which use a butterfly and a risk reversal with expiries of 90 days, even though the exotic option in most practical circumstances will terminate much earlier. Again, in an intuitive way, the Trader Model reconciles European vanilla volatility surface information with American exotic option pricing. This is essential, as all major markets exhibit pronounced term structures of vega and delta convexity to volatility.[47] All other models of this class implicitly assume that such term structures are equal to zero. Savery (2000) and Wystup (2003) show the mechanics of calculating the per unit price of vega and delta convexities to volatility.

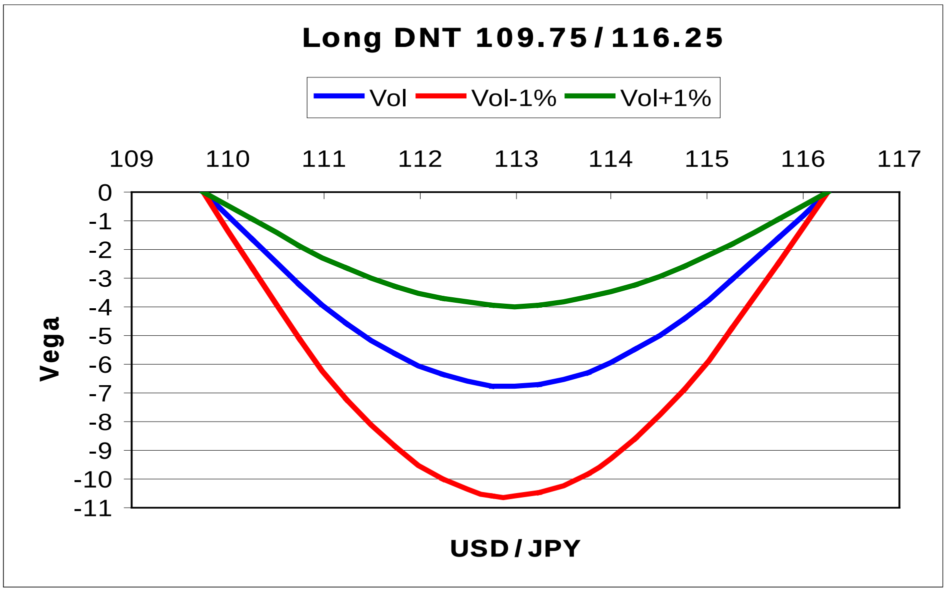

Fig. 3.2. ∂vega/∂vol for a DNT option. This figure shows how BSM-inspired vega for a DNT option with L=109.75 and U=116.25 changes when volatility changes. For this DNT option, vega convexity to volatility is positive. When volatility increases (σ+1%), vega gets longer. When volatility decreases (σ-1%), vega gets shorter. Since vega changes in a beneficial manner whenever volatility changes, the DNT option must trade at a premium to theoretical value, ceteris paribus. This premium will be greatest when the spot FX rate is around 113.

[47] Which is why contemporary orthodox quantitative modelling methods introduce even more degrees of freedom, namely time dependent parameters, to try to incorporate this phenomenon.

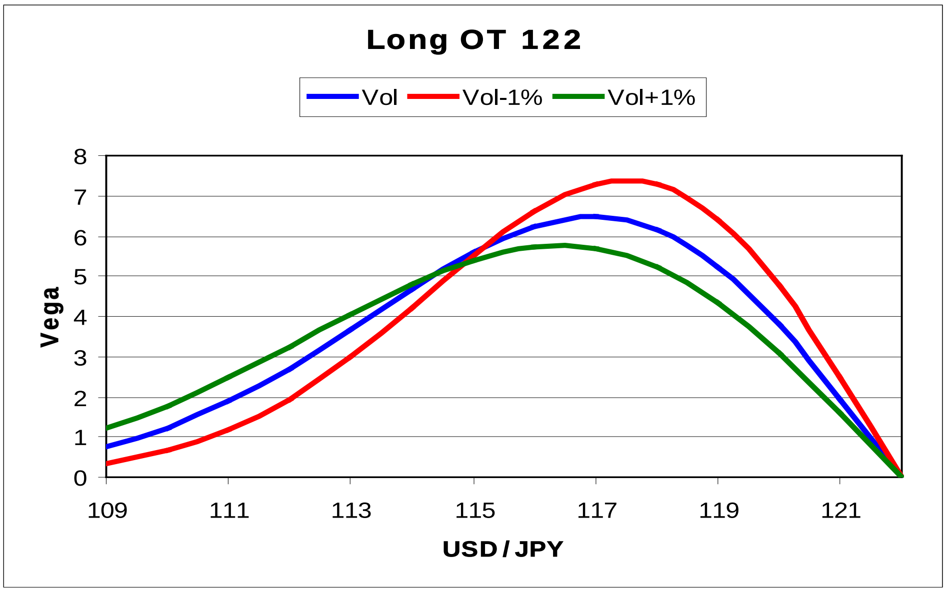

Fig. 3.3. ∂vega/∂vol for a OT option. This figure shows how BSM-inspired vega for a OT option with U=122.00 changes when volatility changes. For this OT option, vega convexity to volatility is positive (negative) when the spot FX rate is less (greater) than approx. 115.00. When the spot FX rate is less (greater) than 115, if volatility increases to σ+1%, vega gets longer (shorter); and if volatility decreases to σ-1%, vega gets shorter (longer). Since vega changes in a beneficial (detrimental) manner whenever volatility changes, the OT option must trade at a premium (discount) to theoretical value, ceteris paribus.

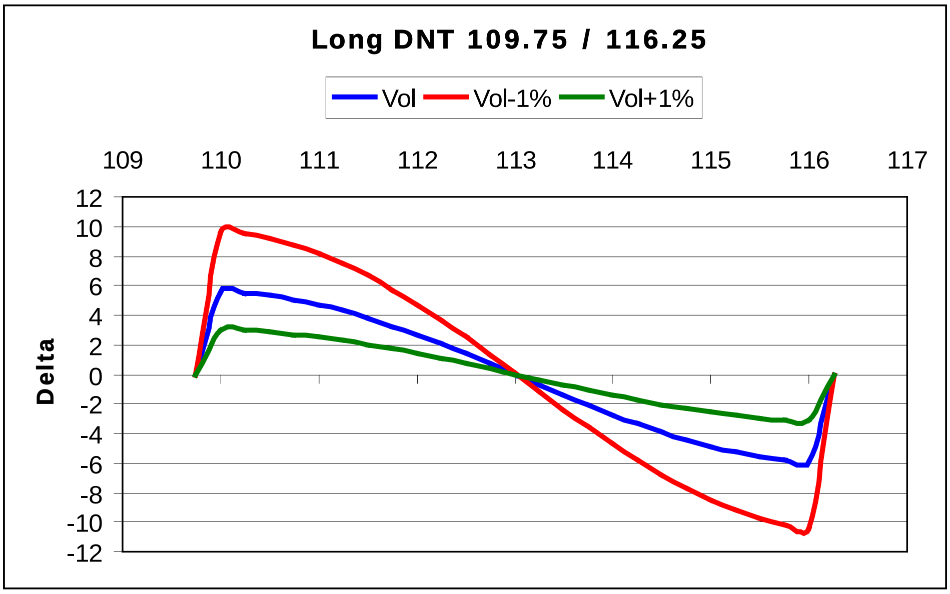

Fig. 3.4. ∂delta/∂vol for a DNT option. This figure shows how BSM-inspired delta for a DNT option with L=109.75 and U=116.25 changes when volatility changes. For this DNT option, delta convexity to volatility is negative (positive) when the spot FX rate is less (greater) than 113.00. When volatility increases to σ+1%, delta gets shorter (longer). When volatility decreases to σ-1%, delta gets longer (shorter). If the skew is negative, delta changes in a beneficial (detrimental) manner when volatility changes, so the DNT option must trade at a premium (discount) to theoretical value, ceteris paribus. If the skew is positive, delta changes in a detrimental (beneficial) manner when volatility changes, so the DNT option must trade at a discount (premium) to theoretical value, ceteris paribus. DNT option's with symmetric barriers have negligible local delta convexity to volatility.

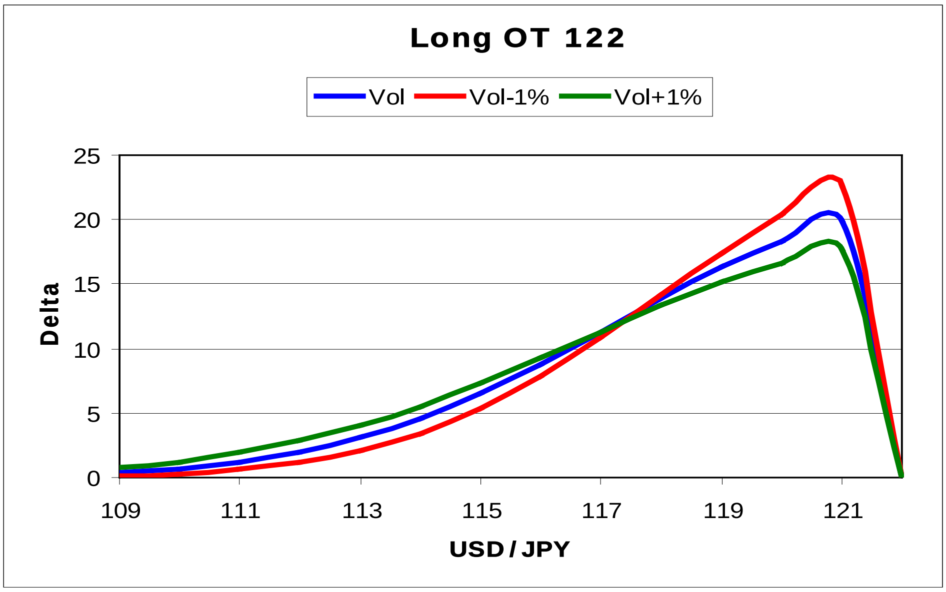

Fig. 3.5. ∂delta/∂vol for a OT option. This figure shows how BSM-inspired delta for a OT option with U=122.00 changes when volatility changes. For this OT option, delta convexity to volatility is positive (negative) when the spot FX rate is less (greater) than approx. 117.00. When volatility increases to σ+1%, delta gets longer (shorter). When volatility decreases to σ-1%, delta gets shorter (longer). If the skew is negative, delta changes in a detrimental (beneficial) manner when volatility changes, so the OT option must trade at a discount (premium) to theoretical value, ceteris paribus. If the skew is positive, delta changes in a beneficial (detrimental) manner when volatility changes, so the OT option must trade at a premium (discount) to theoretical value, ceteris paribus.

Table 3.2 shows a worked example of the process outlined in Appendix A.

Table: 3.2
Trader Model numerical example for 1yr EUR/JPY DNT 120/140.

Theoretical Value		0.1389
Volga	0.02077	
Vanna	0.80412	
EST/Maturity	0.46435	
Fly Zeta p.u.	6.42834	
RR Zeta p.u.	-0.04736	
Smile Adj	0.0620	
Skew Adj	-0.0177	
Mkt Supplement		0.0443
Market Value		0.1832
Vega Spread	0.008137	
Volga Spread	0.026604	
Vanna Spread	0.004702	
Slippage	0.004816	
Model Bid		0.1635
Model Ask		0.2077
Mkt Trade		0.1875

This table refers to a EUR/JPY DNT option with barrier FX rates of 120 and 140, with a maturity of one year. EST is expected stopping time. Zeta per unit is the per unit price of volga (vanna) calculated from the butterfly (risk reversal). The smile (skew) adjustment is the product of volga (vanna), EST/Maturity and fly (RR) zeta p.u. The market supplement is the sum of smile and skew adjustments. Market value is the sum of theoretical value and the market supplement. Spreads are calculated by finding how many straddles are required to hedge the DNT option vega, and determining the cost of crossing the vanilla bid-ask spread to implement that trade. Ditto for volga and vanna. Slippage accounts for the discontinuity at the barrier, and for a DNT option, the asymmetric slippage applies to the Ask side.

The issue of early termination is pivotal, because American binary options are "options on time rather than options on the asset" (Taleb, 1997, p. 305). All models except the Trader Model use smile and skew volatilities at the expiry date of the American binary option to value the market supplement to theoretical value. For example, heuristic models like Wystup (2003) and ∂Vega use expiry smiles and skews exclusively; and even the most complex universal volatility models define diffusion parameters by calibrating to expiry smiles and skews, either exclusively (time-independent) or inclusively (time-dependent). It is clear from Merton (1973, p. 256) that market-traded volatilities are volatilities of the forward, not volatilities of the spot. In FX the volatility surface defines volatilities of the forward temporally (expiry) and spatially (delta) specific to the Garman and Kohlhagen (1983) extension of BSM. However, the early termination condition for American binary options is spot dependent. Early termination occurs if and only if the spot FX rate trades at or beyond the barrier FX rate. The forward FX rate is irrelevant. By using volatilities at the exotic option expiry, strong American optionality is valued with volatilities exhibiting the strongest European characteristics.

The Trader Model reconciles American and European optionality uniquely and intuitively. As the strength of the binary's American optionality intensifies, such as low (high) theoretical value DNT (OT) options, the expected stopping time of the binary shortens. The Trader Model uses volatilities at the expected stopping time to value the market supplement, not the much longer expiry date. Volatilities of the forward at the expected stopping time closely resemble volatilities of the spot, as short-dated European vanilla FX options have low levels of interest rate risk (rho and phi are small). As American optionality weakens, such as high (low) theoretical value DNT (OT) options, the expected stopping time lengthens. Since in these cases American binary options more closely resemble European options, the Trader Model uses volatilities approaching the expiry date to value the market supplement to theoretical value.

Tables 3.3-3.5 inclusive demonstrate the significance of this issue, using JPY as an example. Table 3.3 shows a base case scenario, where JPY and USD interest rates both remain constant when the spot rate changes. This is the same scenario as the BSM paradigm and, as expected, the change in the forward is symmetric to and of the same order of magnitude as the change in spot. In contrast, if the change in spot is accompanied by a decline in JPY interest rates (Table 3.4) or by a decline in USD interest rates (Table 3.5), then the change in the forward is highly asymmetric and different in magnitude to the spot.

Table 3.3
Changes in the JPY forward FX rate when the spot FX rate changes, and JPY and USD interest rates are constant.

Spot Change	S-4.00	S	S+4.00
Spot (S)	100.50	104.50	108.50
r_d (JPY)	0.01	0.01	0.01
r_f (USD)	0.05	0.05	0.05
1Yr Fwd (F)	96.56	100.40	104.25
Fwd Change	-3.84		3.84

If JPY and USD interest rates remain constant when spot changes, the 1yr forward changes symmetrically by approximately the same magnitude as the spot. i.e. spot declines (rises) by 4.00, the 1yr forward declines (rises) by 3.84, where $F=S\exp[(r_d-r_f)(T-t)]$.

Table 3.4
Changes in the JPY forward FX rate when the spot FX rate changes, JPY interest rates decline by 0.01, and USD interest rates are constant.

Spot Change	S-4.00	S	S+4.00
Spot (S)	100.50	104.50	108.50
r_d (JPY)	0.00	0.01	0.00
r_f (USD)	0.05	0.05	0.05
1Yr Fwd (F)	95.60	100.40	103.21
Fwd Change	-4.80		2.81

If JPY interest rates decline when spot changes, the 1yr forward changes asymmetrically to the spot. i.e. spot declines (rises) by 4.00, the 1yr forward declines (rises) by 4.80 (2.81) where $F=S\exp[(r_d-r_f)(T-t)]$.

Table 3.5
Changes in the JPY forward FX rate when the spot FX rate changes, USD interest rates decline by 0.01, and JPY interest rates are constant.

Spot Change	S-4.00	S	S+4.00
Spot (S)	100.50	104.50	108.50
r_d (JPY)	0.01	0.01	0.01
r_f (USD)	0.04	0.05	0.04
1Yr Fwd (F)	97.53	100.40	105.29
Fwd Change	-2.87		4.89

If USD interest rates decline when spot changes, the 1yr forward changes asymmetrically to the spot. i.e. spot declines (rises) by 4.00, the 1yr forward declines (rises) by 2.87 (4.89) where $F=S\exp[(r_d-r_f)(T-t)]$.

Tables 3.3-3.5 inclusive demonstrate that a policy of pricing American binary options with expiry volatility is flawed. The volatility surface is a set of directly quoted European vanilla option price adjustments defined in terms of the volatility of the forward. Tables 3.4 and 3.5 are plausible market scenarios where the volatility of the forward differs significantly from the volatility of the spot. For example, the Bank of Japan's quantitative easing policy[48] reduced JPY interest rates as JPY was appreciating, consistent with Table 3.4.[49] In contrast, in the nine months to May 2008 the Federal Reserve implemented an aggressive monetary easing policy reducing USD interest rates[50] by 3.25% as JPY appreciated (Warsh, 2008), consistent with Table 3.5.[51] If long-dated American binary options are priced under a market scenario like Tables 3.4 or 3.5, then the volatilities of the forward are markedly different to the volatilities of the spot, and it is the latter which is relevant for the termination of American binary options. Volatilities of the forward and volatilities of the spot diverge for all maturities, though the effect is more pronounced as exotic option maturity lengthens. By way of example, Table 3.6 shows JPY forward sensitivity to a decline in JPY interest rates, analogous to Table 3.4.

[48] Commenced March 2001 and lifted in March 2006. This stimulatory monetary policy was principally designed to reduce long-term JPY interest rates (Spiegel, 2001 and Spiegel, 2006).
[49] JPY Call (Put) options increase (decrease) in value when spot JPY appreciates and JPY interest rates fall. Therefore, JPY skew is strongly negative for options with significant exposure to the forward, that is, long-dated options.
[50] Fed funds rate.
[51] In this case, as JPY appreciates the volatility of the spot is greater than the volatility of the forward.

Table 3.6
JPY forward FX rate sensitivity to a JPY interest rate decline of 0.01.

Spot	1mo Fwd	3mo Fwd	1yr Fwd	3yr Fwd	5yr Fwd	10yr Fwd
S – 4.00	-4.07	-4.21	-4.80	-6.18	-7.29	-9.09
S + 4.00	3.90	3.69	2.81	0.70	-1.06	-4.24

If the change in JPY spot FX rates is accompanied by a decline in JPY interest rates of 0.01, then the volatilities of the forward and the volatilities of the spot diverge in a more pronounced manner as option maturity lengthens. Row one shows that a spot decline of 4.00 accompanied by a decline of 0.01 in JPY interest rates, causes the forward rate to fall by only 4.07 for 1mo, and 6.18 for 3yr maturity. This table matches the scenario in Table 3.4. Mo is the acronym for 'month' and Yr is the acronym for 'year'. The results in zero-delta straddle terms, quoted as USD (FOR percent) strike is approximately the same, as expected. The variation is from zero-delta straddle strike dependency upon the level of volatility, whereas the forward rate is independent of volatility. If one sets volatility to zero, the zero delta straddle terms are identical to the forward FX rate terms. Nevertheless, the conclusions are the same under both methods.

Table 3.7 summarises the significant inconsistencies between instantaneous volatilities that are defined by arbitrary, quant-imposed dynamics like stochastic volatility, jump diffusion and universal volatility; and the actual market-traded volatilities comprising the European vanilla volatility surface.

Table 3.7
Inconsistencies between quant-imposed instantaneous volatilities and actual market-traded European vanilla volatility surfaces.

Instantaneous Volatility	Volatility Surface
Dynamics.	Statics.
Arbitrary.	BSM.
Non-traded.	Traded.
Instantaneous.	Forward.
Path dependent.	Path independent.
Quant defined.	Market defined.
Continuous.	Discrete.

BSM is the acronym for Black-Scholes-Merton model. Quant is a colloquial financial market term for financial engineer.

It is clear from Table 3.7 that financial engineers, either explicitly or implicitly, make a substantial leap of faith when they calibrate instantaneous volatility diffusions to the volatility surface. In effect, they rely on all the inconsistencies conveniently cancelling each other out so that only useful information remains. However, empirical research shows that this is not the case, as there are significant pricing departures both within and between classes of exotic option pricing models calibrating to the same volatility surface (e.g. Schoutens et al., 2004; and Detlefsen et al., 2006). Financial engineers calibrate their exotic option pricing models to the European vanilla volatility surface because it is the only credible alternative to get their models to price non-trivial economic factors crucial to the market. However, Table 3.7 shows that this credibility is superficial, not theoretically sound.

Fig. 3.6 and Fig. 3.7 are schematic representations of the main classes of existing American exotic option pricing models and the Trader Model, respectively. The figures show how the Trader Model returns to the financial economic roots of American exotic option pricing by not imposing an artificial structure for volatility dynamics, and, thereby, simplifying the market pricing task both theoretically and computationally, without constraining the solution space like orthodox methods.

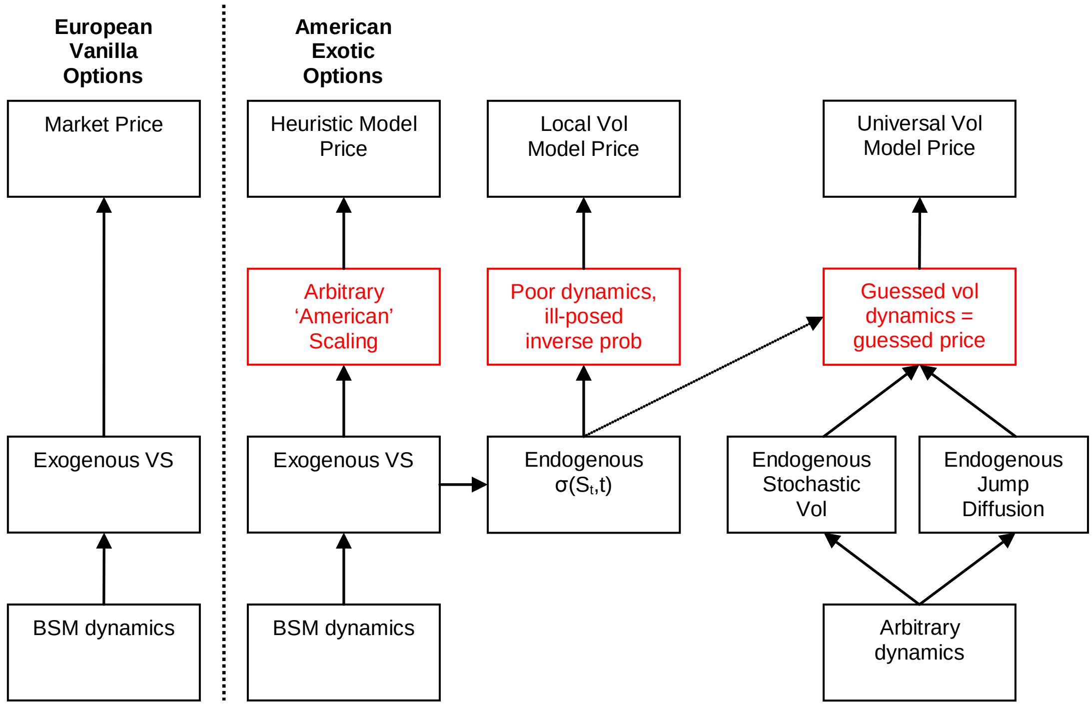

Fig. 3.6. Schematic of existing American exotic option pricing models. The red text boxes succinctly describe the major problems with the existing American exotic option pricing models. Model complexity, both theoretically and computationally, increases significantly from left to right. VS is the European vanilla option volatility surface.

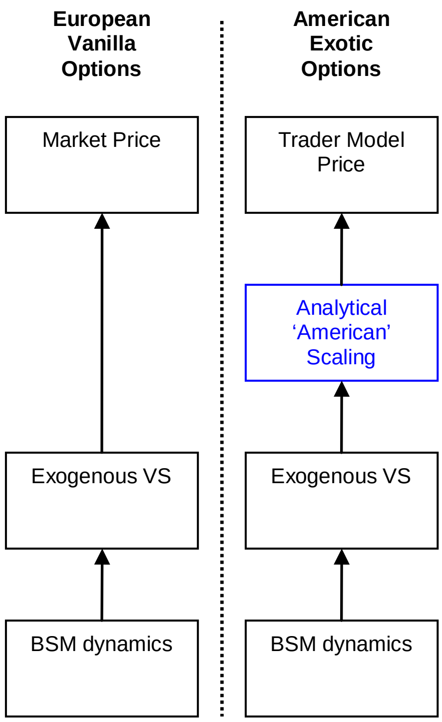

Fig. 3.7. Schematic of the Trader Model. The blue text box shows the pricing and hedging step unique to the Trader Model. Crucially, the step is analytical, as the scaling is performed with the expected stopping time, which is dependent on market and option contract inputs only. VS is the European vanilla option volatility surface.

A key benefit of using this approach to price the market supplement is that the additional market risk is quantified directly and a unique portfolio consisting of liquid market-traded options is identified to hedge that market risk. Since the additional market risk is priced with butterflies and risk reversals that expire at the expected stopping time, it is logical for the hedge portfolio to consist of those European vanilla options.[52] When an American binary FX option is first dealt, its main risk is the net contribution to the price-maker's book of high-order greeks, such as vega and delta convexity to implied volatility:

> "In mathematical terms, the book, neutral in its lower moments, can easily lose its stability in the higher moments. An option book, we will see, is not as 'compact' as mathematicians believe. It will generally be neutral in the lower moments and exposed to various risks in the higher moments" (Taleb, 1997, p. 149).[53]

So butterflies and risk reversals are the obvious immediate hedging choice even for price-makers with models based on arbitrary volatility dynamics. This is a crucial point not only theoretically, but also from a practical economic perspective. It is the price-makers' hedging activity which binds the volatility surface to exotic option prices, not the dynamical behaviour of non-traded parameters imposed by a financial engineer. Since price-makers must hedge in practice with market traded instruments, the fact that the Trader Model specifically preserves the direct link between price and hedging cost, and orthodox methods do not, has significant implications for the assessment of bank model risk under Basel II as well as trading desk profitability and sustainability.

Fig. 3.8 and Table 3.8 show the errors which occur when using a European vanilla option hedge portfolio with expiry matching the expiry date of the American binary option. As shown in Fig. 3.8, a hedge portfolio is only a hedge when it offsets the risk of something else. If the American binary option terminates, the European vanilla option hedge must terminate too, to be considered a hedge. The Trader Model hedge terminates at the expected stopping time (EST) of the American binary option it is hedging. In contrast, Wystup's (2003) hedge turns into an open position once the American binary option terminates. For low (high) theoretical value DNT (OT) options, the length of the residual open position EST→T is significant.

[52] In reality, price-makers will first attempt to warehouse market risk to exploit natural hedges within their book. If that is insufficient because of their risk appetite or the bank's (via trading limits), they will then try to avoid crossing interbank bid-ask spreads by aggressively pricing structured deals to lure franchise flows into flattening the exposures of concern. If franchise flows are insufficient, price-makers will clear excess risk through the interbank market. Usually via vanillas in the first instance, at least, to exploit their liquidity and tighter bid-ask spreads. Since hedging trades are usually traded as price-takers, hedging with exotic options is timed to coincide with favourable market conditions.

[53] Higher moments are the skew (3rd), smile (4th) and their instabilities (5th, 6th, 7th) (Taleb, 1997, p. 202-3).

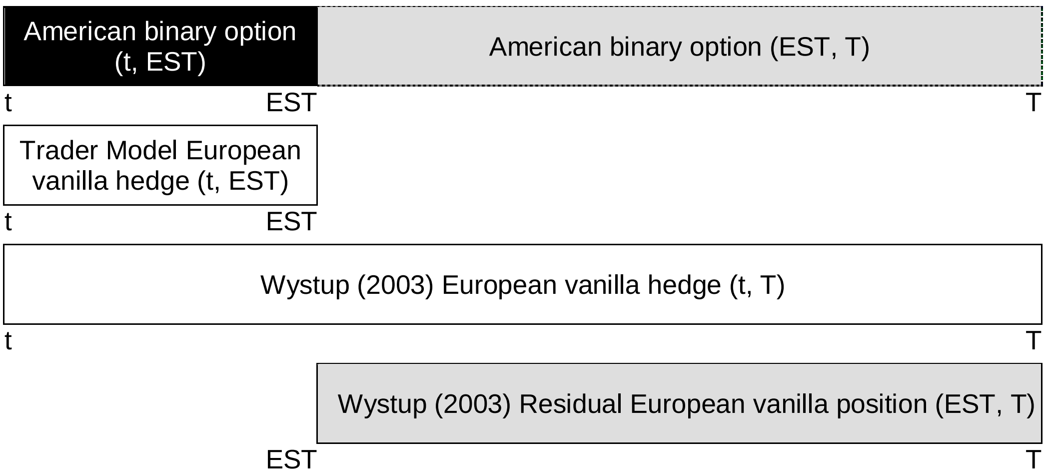

Fig. 3.8. The residual (European vanilla option) open position remaining once an American binary option terminates. The American binary option trades at t and expires at T. It has an expected stopping time equal to EST. The Trader Model hedges the American binary with European vanilla options that expire at EST, leaving no residual position. Wystup (2003) and ∂Vega use European vanilla options that expire at T, and therefore, leave a residual open position of length (EST, T).

Table 3.8 shows that the term structures of convexity can be significantly different between expected stopping time and expiry. Using the methodology in Savery (2000), for 13 December 2007, JPY per unit price of vega convexity to volatility was $10.69 for three months, and $16.99 for twelve months. Similarly, the per unit price of delta convexity to volatility was $0.1002 and $0.2068, for three and twelve month expiries, respectively. These prices of convexity will result in significantly different prices for American binary options.

Table 3.8
Term Structures of Convexity for JPY.

	Term	0d Straddle	10d Fly	10d RR
est	3m	10.25	1.75	-5.75
T	12m	9.10	3.00	-7.50

JPY traded volatility data from 13 December 2007 in percent. For an American binary option with EST = 0.25(T), using expiry data to price convexity will grossly overstate the impact of both the (vega neutral) butterfly and the (delta neutral) risk reversal.

Furthermore, price-makers do not construct option risk eliminating hedges for vega, and vega and delta convexities to volatility as per Lipton and McGhee (2002, p. 82) and Ayache et al. (2004, p. 12), because the excess risk is net book, not gross option, in nature. Price-makers use butterflies and risk reversals to smooth book smile and skew risk optimally, they do not construct hedge portfolios to eliminate a single exotic option's risk.[54] They explicitly take into account the term risk caused by the lengthening and shortening of American exotic option expected stopping time versus the constant duration of the European vanilla option hedge portfolio. Since price-makers do not use the entire set of all European vanilla options spanning the volatility surface to hedge their net book risk, it is illogical to calibrate exotic option pricing

[54] The economics which underpin sell-side price-making desks is to charge end-user clients the gross cost of hedging options via their price, but then only have to incur the much smaller cost of hedging net book risk, by exploiting warehousing. For example, a price-maker charges an exporter and an importer the gross cost of hedging implicitly in their option prices, but the exporter and importer options offset each other to some extent, significantly reducing the price-maker's net hedging cost. There are some wrinkles in the market. E.g. if market conditions make exporter hedging attractive, importer hedging will be unattractive. Hence, risk typically accrues asymmetrically, requiring price-makers in the first instance to smooth unwanted book imbalances in the liquid, European vanilla option market using commoditised strategies with tight bid-ask spreads.

models to the entire volatility surface. As time elapses, the main risk transfers to discontinuity risk, where the smoothing of gap deltas becomes paramount. This is achieved best with opposing binary or barrier trades, but one can also use European vanilla option spreads and spot delta unwinds.

3.3.3. Attributes of the Trader Model

The Trader Model exhibits the following characteristic qualities:

1. Internal consistency with the universally accepted benchmark BSM / exogenous volatility surface model for pricing European vanilla options. Instead of unnecessarily and sub-optimally complicating the European vanilla pricing environment in order to price exotics to market, as per contemporary research (e.g. Jex et al., 1999; Lipton and McGhee, 2002), the Trader Model instead retains the universally accepted vanilla benchmark paradigm and simplifies the exotic option pricing environment. Since no additional structure is required, model risk is significantly reduced. Model risk is reduced because the Trader Model does not suffer from a crucial weakness that undermines all local, stochastic, jump and universal volatility models: where all "agree on the vanilla prices [by construction] and totally disagree on the exotic prices and the hedging strategies" (Ayache et al. (2004, p. 1). The coexistence of a large number of models of this type with radically divergent prices, is a clear indication that no single model dominates all others. The Trader Model instead prices the cost of hedging smile (skew) risk with a butterfly (risk reversal) with expiry equal to the expected stopping time of the exotic option.

2. Does not require arbitrary specifications of volatility dynamics. Derman (1999) has shown that the same asset can exhibit different volatility regimes in different time periods, thus making the choice of 'correct' volatility dynamics, when that approach is used, both critical and impossible ex ante. This contradiction is at the core of Taleb's (1997, p. 383) observation that financial engineers prefer "a flawless model based on imperfect assumptions", whereas traders prefer "an imperfect model based on flawless assumptions". The Trader Model is imperfect, but it is based upon the fewest and universally accepted, if not flawless, assumptions. Since the Trader Model is independent of arbitrary specifications of volatility dynamics, the significant model risk associated with defining arbitrary volatility diffusion parameters with the volatility surface is avoided. Statics in the present (volatility surface) are not used to define dynamics in the future (arbitrary volatility diffusion).

3. Quantifies the time effect of American optionality. The impact of time is so critical that Taleb (1997, p. 305) describes American binary options as "options on time rather than options on the asset". The Trader Model is the only model which prices vega and delta convexities to volatility at the expected stopping time of the exotic option. Models of the same class (Famery and Cornu, 2000; the heuristic model of Lipton and McGhee, 2002; Wystup, 2003; and commercial vendor system ∂Vega) price these convexities at the nominal duration of the option, and contemporary approaches such as universal volatility models try a brute-force approach of introducing time-dependency for arbitrary volatility diffusion parameters. Both methods are sub-optimal on theoretical and computational grounds because they introduce term risk. American binary options are, in most practical circumstances, expected to terminate much sooner than the nominal duration, and so convexity prices at the nominal duration are irrelevant.[55] Time-dependent parameters are equivalent to using a computational sledgehammer to crack a financial economic walnut. Not only are parameters selected arbitrarily, but they are also allowed to float freely over time with the only condition being to minimise a mathematical optimisation algorithm. Attempting to improve computational efficiency by calibrating with few free parameters fails because:

> "In practice the smile is not well matched over all option maturities by a single volatility process. For markets where the bid-offer spreads for European options are at the basis point level any best-fit approach with a limited number of free parameters is likely to be outside the spread for most options" (Jex et al., 1999, p. 5).

It is precisely because contemporary methods ignore the fundamental economic relationships underpinning the volatility surface, that they are left with the computational problem of manipulating large matrices and optimisation algorithms to approximate prices. In contrast, the Trader Model uniquely focuses on a vector of information at the expected stopping time only, and this information is not corrupted by passing it through an algorithm to force it to define something for which it is ill-suited. It is used to price effects common to European vanilla and American exotic options.

4. Market risk is defined in terms of unique, liquid, market-traded hedging strategies. In the Trader Model, the market price of the exotic option is actually the cost of its unique market hedge, which corresponds to price-makers' model-independent financial economic intuition. In contrast to the pre-eminent role given to the hedging of market risk in the Trader Model, local, stochastic, jump and universal volatility models neglect hedging considerations until after the price has been obtained, and find that "having accomplished a good fit of a proposed model to market data does not tell anything about how to hedge" (Hakala and Wystup, 2002, p. 276). Even if price-makers use local, stochastic, jump or universal volatility models in practice, they can still only use hedging strategies which are available in the market. Except in this case, the cost of the hedging strategy is not overtly taken into account, rather it is the opaque interaction of arbitrarily chosen dynamics and non-traded parameters with little to no concession to real-life financial market practicalities. Pricing with one philosophy and hedging with another introduces model risk, since, to-date, there is no way to reconcile their differences. These differences strike at the core of price-maker desk sustainability and profitability over the medium-term.

5. Mathematically simple and computationally efficient. The Trader Model is dependent on the simplest most parsimonious dynamics of all exotic option pricing models. As a result, it renders intermediate calculations like extensive and intensive numerical calibration redundant, thereby achieving significant computational efficiency savings. This contrasts markedly with the popular pursuit of marginal efficiency gains by tinkering with ever more elaborate computational routines, a prime example of "quantitative finance . . . wasting itself in sophisticated mathematical exercise" (Ayache et al., 2004, p. 33). The Trader Model is not only a whole-of-[volatility]-surface model, it is also a term structure of interest rates model. However, instead of trying to identify dynamics which cover stochastic spot, volatility and interest rates for all assets,

[55] The economics underpinning the heuristic method is to value the European vanilla option portfolio used to hedge the risks of the exotic option. It does not make sense to choose a hedge portfolio which remains after the exotic option is expected to terminate.

time periods and regime changes, the Trader Model uniquely focuses on the relevant vector of volatility, and the relevant elements of interest rates, to make significant computational savings. The relevant information is at the expected stopping time of the American exotic option.

The Trader Model does not attempt to forecast the real distribution of future volatility, it simply takes as given the benchmark conditions universally accepted by the European vanilla FX option market, and prices exotic options in an internally consistent manner. Contemporary orthodox methods like universal volatility models, cannot justifiably claim the high ground on distributional issues. The distributional assumptions of Jex et al. (1999), Blacher (2001) and Lipton and McGhee (2002) do not dominate others in terms of pricing performance, otherwise there would be a single benchmark model and not a family of models. Furthermore, trading applications of universal volatility models perform intra-day re-calibration of numerous time-dependent free parameters[56] in order to obtain the requisite in-sample closeness of fit. Intra-day re-calibration is tacit recognition that advocates of universal volatility models do not have any insight as to the real distribution of future instantaneous volatility.[57] It remains a fact that price-makers using universal volatility models still hedge market risk by analysing book sensitivity to high-order greeks such as ∂vega/∂vol and ∂delta/∂vol. They do not hedge by analysing exposure to model-specific non-traded parameter sensitivities.

6. Price Transparency. Price transparency is just as important as price discovery. Transparency has three real economic benefits. Firstly, it allows exotic option price-makers to understand their true risk and modify prices and hedges consistently, even when markets are unstable or undergoing structural change. Secondly, it promotes the commoditisation of price-maker spreads and the reduction of internal distribution margins for the benefit of end-users in the real economy. Thirdly, commoditisation also forces sell-side financial institutions to innovate in order to maintain and grow profits, which helps to make markets more complete.

Despite the obvious benefits of price transparency, other models are deliberately opaque. For example, Gershon, CEO of SuperDerivatives, a high-profile exotic option pricing model used in the financial markets claims that:

> "The typical options user / hedger doesn't worry about the mathematical detail of the price if the outcome is reliable. We drive safely with our cars but never bother to learn how their engine works. We use electric machines and don't know the mechanism behind them" (cited in Wystup, 2006, p. 304).

Contrary to Gershon's claim, price-makers do need to understand how model prices are obtained if they are to add value in an environment where daily profit and loss is calculated as mark-to-model for exotic options and mark-to-market for the vanilla options used to hedge them.[58] Similarly, bank internal middle-office risk managers and external regulators must

[56] Parameters also usually have additional artificial constraints imposed by financial engineers to avoid modelling complications, like an error functional minimisation calibration which yields a negative 'volatility of volatility' parameter.

[57] If model dynamics closely represented the market, then only an initial calibration would be required, as thereafter, model dynamics would follow market changes closely. If frequent re-calibration is required, it is required because model dynamics do not closely follow market changes.

[58] The FX market is dominated by OTC markets. OTC exotic option markets are heterogeneous owing to franchise flows structuring unique hedging or speculative positions (contrast with a futures exchange with standardised products). Therefore, there is no end of day closing price for every exotic option traded OTC. As a result, banks calculate profits asymmetrically (exotics mark-to-model and vanillas mark-to-market), which can lead to reversals in profitability once unrealised exotic option P&L turns to realised P&L.

understand pricing mechanics in order to establish the correct amount of regulatory capital to retain under Basel II economic capital. Furthermore, given the requirement by new accounting standards such as IAS 39, which provides rules on how to include derivatives in the balance sheet and the profit and loss of companies, even price-takers must understand how prices are formed. The analogy with cars and electric machines is nonsensical, because both cars and electric machines must be manufactured according to minimum safety standards imposed and enforced by regulatory authorities, effectively relieving end-users of the need to know detail. Since SuperDerivatives models have not been verified and validated by external regulatory authorities, end-users cannot take comfort in opacity.

By virtue of its financial economic simplicity, pricing and hedging with the Trader Model is much easier to understand and more inclusive of the broader finance community than contemporary models whose high-level financial mathematics restricts their application to the exclusive domain of financial engineering specialists. Table 3.9 is a summary of comments from high-profile financial engineering specialists outlining (i) major problems with the prevailing mathematical orthodoxy for pricing and hedging exotic options; and (ii) recommended solutions. The final column shows that the Trader Model not only does not suffer from the problems, but also already incorporates the recommendations.

Table 3.9
Major financial engineering issues and Trader Model solutions.

Quant	Model Issue	Trader Model
Quant Imposed Model Processes and Dynamics		
Avellaneda	"Guessing the volatility process is the same as guessing the price". Cited in Ayache et al., 2004, p. 11.	TM does not constrain the solution space by imposing arbitrary volatility dynamics.
Jex et al.	"A number of different [volatility] processes could be postulated which would match the observed volatility smile and yet give different values for the same path dependent [exotic] option". 1999, p. 5. "In practice the smile is not well matched over all option maturities by a single volatility process". 1999, p. 5.	TM does not constrain the solution space by imposing an arbitrary volatility process.
Ayache et al.	"Nobody should be in a position to decide which particular smile dynamics will prevail". 2004, p. 11. "Your wrong guess about the smile dynamics can generate an immediate arbitrage opportunity against you". 2004, p. 11.	TM solutions are arbitrage-free and independent of smile guesses.
Quant Imposed Model Estimation, Numerics and Calibration Techniques		
Andersen & Andreasen	"Stochastic volatility and jump diffusion models . . . [are] difficult to handle numerically and slow to calibrate accurately to traded prices". 2000, p. 233.	TM does not calibrate.
Hakala & Kirsch	"There are no obvious relationships between market and model parameters which makes estimation of model parameters difficult to verify". 2002, p 249.	TM does not estimate parameters.
Andersen & Andreasen	While fitting of stochastic volatility models is possible it often requires "unrealistically high negative correlation between the [underlying asset] and [its] volatility". 2000, p. 232.	TM does not distort or dilute information by modifying input data via 'fitting'.
Lack of Hedging Insights		
Hakala & Wystup	"Having accomplished a good fit of a proposed model to market data does not automatically tell anything about how to hedge". 2002, p. 276.	TM identifies a unique traded hedge and ensures that option price equals hedging cost.
Ayache et al.	"Local volatility, jump-diffusion, stochastic volatility etc. may agree on the vanilla prices and totally disagree on the exotic prices and the hedging strategies". 2004, p. 1.	TM hedge is consistent with how traders actually hedge in the market.
Quant Recommendations to Improve Exotic Option Pricing Models		
Derman	"One good strategy in attempting to value exotic options that are sensitive to the smile is to try to avoid modeling the dynamics of volatility as much as possible". 2003, p. 13.	TM does not model the dynamics of volatility.
Taleb	"It is better to use a model with the smallest number of parameters to estimate". 1997, p. 109.	TM has no parameters to estimate.
Ayache et al.	"The need to go back to 'basics' is a very welcome conclusion, . . . , at a time when quantitative finance seems to be wasting itself in sophisticated mathematical exercise". 2004, p. 33.	TM is about the 'basics'. i.e. economic pragmatism not mathematical idealism.
Taleb	"It is better to improve on a simple but seasoned model than operate with a more advanced but newer model". 1997, p. 109.	TM uses the simplest exotic option pricing methodology.
Taleb	"Traders . . . are comfortable with [the BSM model] because they have learned the necessary tricks to make it work". 1997, p. 109.	TM prices smile, skew and term risks consistently with BSM.
Ayache et al.	"Financially relevant questions can only be answered by relevant financial theory". 2004, p. 33.	TM is based on financial economics not mathematics.
Bates	Recommends "a renewed focus on the explicit financial intermediation of the underlying risks by option market makers" and "plausible models of market maker behaviour". 2003, p. 400.	TM explicitly prices how price-makers actually trade in the market in practice.
Derman	"A future theory that married the quantitative to the behavioural would be a worthy goal". (2002, p. 82)	TM marries quantitative methods to actual trader behaviour more closely than any other model.

'TM' is the acronym for Trader Model.

3.4. Testing the Trader Model

3.4.1. Scope and Scale

Even though the Trader Model can be applied in general to all instruments and markets, the focus of the empirical analysis in this thesis is first generation binary options in the foreign exchange (FX) market with maturities of one year or less. The rationale for this focus is explained below.

First generation binary options. Published papers of high-profile competitor models use binary options to explain their research. For example, Lipton and McGhee (2002) use double-no-touch (DNT) options, while Jex et al. (1999) and Wystup (2003) use one touch (OT) options. Authors choose these options because they are highly sensitive to the shape of the volatility surface and because they are "smallest decomposable fragments" (Taleb, 1997, p. vi), thereby avoiding redundancy. Other binary option prices are determined by simple parity relationships,[59] and the prices of other popular (by traded volume) first generation exotic options, such as reverse barrier options, are primarily explained by the discontinuity priced by OT options because "typically a barrier option will be dynamically hedged with a combination of the underlying, a vanilla option, and a one-touch" (Ayache et al., 2004, p. 34). As well as the aforementioned reasons, Ayache et al. (2004, p. 11) note that:

> "Your wrong guess about the smile dynamics can generate an immediate arbitrage opportunity against you, if somebody picks the right security to trade against you. As a matter of fact, all FX option traders are aware of the existence of such a security. It is the barrier option, the simplest instance of which is the one touch".

To be precise, the focus of the empirical analysis in this research is continuously monitored American binary FX options with one (OT) and two (DNT) barriers. In a foreign exchange context, a OT option obliges the seller to pay a fixed cash amount to the buyer if the spot FX rate trades in the market at or beyond the barrier FX rate, prior to expiration.[60] While the liability is crystallised immediately, physical payment occurs on the delivery date of the option.[61] A DNT option obliges the seller to pay a fixed cash amount to the buyer if the spot FX rate trades in the market without ever touching or exceeding either barrier FX rate prior to expiration. The liability can only be crystallised at expiration, and physical payment occurs on the delivery date of the option. For illustration, expiry payoffs of OT and DNT options are shown in Fig. 3.9. It is self-evident that, ceteris paribus, the price of a OT (DNT) option increases (decreases) when the distance between the barrier(s) and the current spot FX rate diminishes. In the interbank exotic FX option market, the convention is to price OT and DNT options as a percentage of the fixed cash payment (e.g. a price of 10% is equivalent to a premium of

[59] Letting NT be the acronym for No Touch options, and R the cash payout received on the delivery day of the option, then OT + NT = Rexp($-r_d$(T-t)). Similarly for two barriers, where D stands for double, DOT + DNT = Rexp($-r_d$(T-t)).

[60] For a FX binary option to be considered 'touched', transactions at or beyond the barrier FX rate must trade in commercial size (USD3 million) between 5:00am Sydney time on Monday, to 5:00pm New York time on Friday (Wystup, 2003, p. 1).

[61] Each exchange rate has its own settling instructions. However, for options on the major exchange rates the delivery date is usually two good business days after the expiry date. Public holidays and weekends cause delivery dates to extend beyond two calendar days.

$1 million if the payout is $10 million). OT and DNT options are the binary options with greatest traded volume in the FX market.[62]

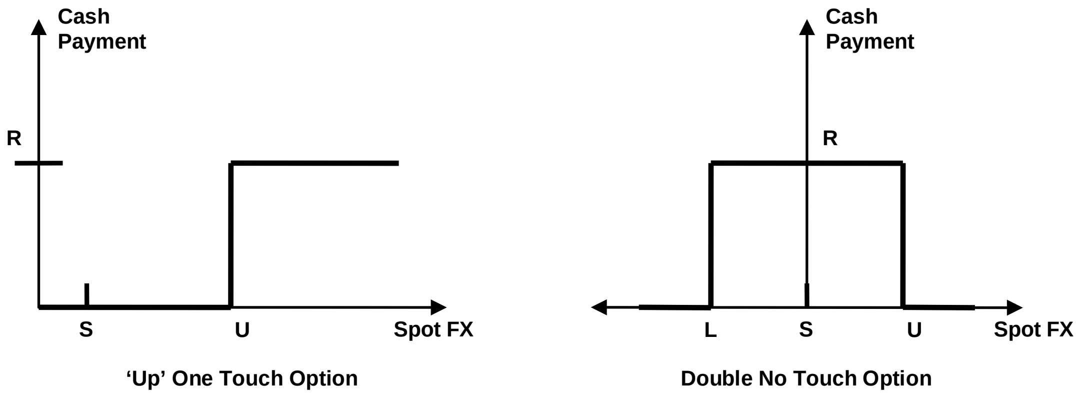

Fig. 3.9. Expiration payoff profiles for OT and DNT options. S is the spot FX rate. L(ower) and U(pper) are spot FX barrier rates. R is the fixed cash payment amount. An 'up' ('down') OT option approaches the barrier from below (above). For DNT options, L <S < U is a starting condition, and barrier prices L and U may be symmetric or asymmetric about S.

Foreign Exchange Market. Published papers of high-profile competitor models, such as Jex et al. (1999), Lipton and McGhee (2002) and Wystup(2003), all focus on the FX market. Having the same market focus for empirical analysis as competitor models facilitates ready comparison between this research and the research of others. The exotic FX option market is also priced to extremely fine tolerances[63] and has a liquid, actively quoted volatility surface,[64] making tests of pricing accuracy more rigorous and robust vis-à-vis equity and commodity markets, which are characterised by significant market frictions,[65] ill-posed inverse problems owing to sparse data, and non-synchronous data. Insofar as exotic options are concerned, FX markets are also much bigger and more sophisticated than their equity and commodity counterparts.

Maturities of one year or less. Most traded volume in the OTC FX option market is in relatively short maturities of one year or less. This short maturity structure also coincides with the perception and practice of FX option price-makers being "volatility managers" (Wystup, 2002, p. 16). While it is possible to extend the unique pricing and hedging method of the Trader Model to long maturities, that is not the focus of this thesis.[66]

[62] In the exotic option pricing database used in this thesis (Jan-Sep 2004 inclusive), NT and DOT options did not trade through the interbank GFI broker market. It is common in the interbank exotic FX option market for OT and DNT option trading to be liquid, and NT and DOT trading to be illiquid.
[63] Five decimal point accuracy for American binary FX options.
[64] The Trader Model methodology still works even in the absence of a liquid volatility surface. The entire volatility surface can be propagated from sparse data without the current methods of extrapolation or interpolation which undermines existing theory, via the method in Smith (2006).
[65] Market frictions include bid-ask spreads, non-continuous trading, stamp duty, brokerage etc. Equity and commodity market frictions are significantly larger than FX markets, resulting in much wider bid-ask spreads for exotic option prices.
[66] The author has not been able to obtain much traded data for long-dated maturities, owing to the relatively infrequent traded volume. In the exotic option pricing database in this thesis, only nine binary options (4.9%) had maturities greater than one year.

3.4.2. Testing Framework

Introduction

It is customary in contemporary published research to test the empirical performance of new exotic option pricing models against: (i) actual traded market prices and (ii) established best practice competitor models. For example, Jex et al. (1999) and Lipton and McGhee (2002) empirically tested the predictive power of their universal volatility model prices against both actual traded market prices and the prices of other models. The models chosen as comparative benchmarks were local volatility (Jex et al.), and local and stochastic volatility (Lipton and McGhee). In both cases, there were no statistics or diagnostics only visual representations of model prices relative to actual traded market prices. The inferences drawn from these observations were that the universal volatility model was *more* accurate because its prices were closest to the actual traded market prices; and *sufficiently* accurate because its prices traded within the proxy for the actual traded market bid-ask spread.

Hakala and Wystup (2002) took a less rigorous approach by assuming that the trader's rule-of-thumb model price (Wystup, 2003) was a de facto market price, and tested empirically an implementation of Heston's (1993) stochastic volatility model against it. Again, there were no statistics or diagnostics, just a visual representation with the inference that since the prices from Heston's model were "fairly close to the prices the trader's rule of thumb method yields" (p. 279), that a properly calibrated Heston model can "match prices quoted by traders" (p. 267). Testing model prices against model prices without reference to actual traded market prices is not only sub-optimal it is insufficient. Given its inability to explain market prices of other options in the same class, selecting the 'trader's-rule-of-thumb' model as a benchmark appears to be motivated more by convenience than relevance.

This thesis follows the same format as Jex et al. (1999) and Lipton and McGhee (2002), but with major modifications to improve the rigour and robustness of the empirical testing. Table 3.10 shows that the empirical research in this thesis is much more comprehensive as the Trader Model is tested with an extensive database of actual traded exotic FX option market prices which is significantly larger in scale and scope than any other published research. It is also tested against a competitor model which has gained considerable support among both academics and practitioners as a de facto benchmark for exotic FX option pricing. Therefore, to be considered a success, the Trader Model must achieve a much higher performance standard, under actual market conditions, than any other model proposed in published research.

Table 3.10
Summary of the scale and scope of empirical testing of models pricing American binary FX options to market.

	Trader Model	Jex et al.	Lipton & McGhee
Currency Pairs	8	1	1
Expiry Days	90	1	1
Spot FX Rates	159	1	1
Number of Deals	183	16[67]	Not reported.[68]
Exotic Option Class	OT & DNT	OT	DNT
Volatility Skews	Positive & Negative	Negative	Positive
Vol Term Structures	Normal & Inverse	Normal	Normal

For example, in this thesis, the database of traded market prices consisted of 90 unique exotic option expiry days and 159 unique spot FX rates. Collectively, this represents a diverse range of different exotic options being priced in an absolute sense, as well as relatively vis-à-vis Jex et al. (1999) and Lipton and McGhee (2002).

In this research, if the Trader Model performs strongly against the market, it is prima facie evidence that the model has explanatory power and is accurate, as it identifies and quantifies factors crucial to the formation of exotic option prices. However, whilst it is necessary to test against the market, it is not sufficient. Pricing models are the start of the price-making process not the end, and so any price formed in the market is a complex combination of mathematical and behavioural elements, some of which simply cannot be modelled. For example, price-makers with a trading view on the market, an asymmetric inventory or an opinion on the counterparty's trading intention, will aggressively lean model prices to the right or left to obtain the desired exposure and to optimise bid-ask spread income.[69] Price-makers lean model prices to the right (left) when they have a preference to buy (sell) the option. That is, if the model price is 0.06-0.08, a price-maker may show a preference to buy the option by making the counterparty a price of 0.07-0.09. If the counterparty lifts the price-maker's ask, the actual traded price in the market is 0.09, or a full 0.01 outside the pure model price without trader intervention. This does not mean that the model's price performance is poor, just that there are real economic factors affecting market prices other than market risk at present. Testing the model's pricing performance against competing models helps to 'level the playing field' by removing the effect of price-makers' unique trading intervention(s) on the price.

Fig. 3.10 is a schematic diagram outlining the price-making process in the interbank exotic FX option market. The Trader Model, or any other model, is located at step 3. Models return the price of risk given market[70] and contract inputs[71] at present. Traders use model bid and ask prices as the starting point of their value-add in the price-making process. To compare

[67] Jex et al. (1999) do not disclose model inputs.

[68] Lipton and McGhee (2002) present results as solid continuous lines, not discrete points. As a result, it is not possible to identify the underlying discrete deals at source.

[69] Wystup (2003, p. 3) also notes that "other prices for one-touch options can be caused by different vega profiles in the trader's portfolio, a marketing campaign or a hidden additional sales margin". References to 'marketing campaign' and 'sales margin' only apply to bank clients, not interbank counterparties. However, it is clear that there are many reasons why market prices can trade away from model prices.

[70] Market inputs are unique to each bank. For example, volatilities used in the model are those quoted by the vanilla option desk of the exotic option price-maker's bank, not the broker's amalgamated average of many banks' bid-ask volatilities. Similarly, deposit rates and forward swap points are quoted by the bank's interest rate and FX forward desks, respectively, not amalgamated broker averages. Different 'market' inputs is another source for variation between bank exotic option prices. This is one reason why interbank exotic FX option markets agree the theoretical value before pricing to market. Any material disagreement in market input will be captured in the theoretical value.

[71] Contract inputs are face value, barrier FX rate(s), expiry and delivery dates etc. Face value is usually only disclosed for pricing purposes as a professional courtesy, if it is unusually large or small relative to interbank exotic FX option pricing convention. Obviously, face value is crucial to the settlement of a dealt option, and its hedging.

model bid and ask prices (step 3) directly against broker collated prices (step 9) or actual traded market prices (step 10) is not sufficient by itself, as one is actually comparing inputs (model prices) to outputs (collated or traded prices), which is both internally inconsistent and illogical. Price-maker trading interventions like steps 4, 5 and 6 undermine the effectiveness of the analysis because they drive a wedge between the model and the market. The need for these trading interventions is a key reason why banks remunerate price-makers extremely well, otherwise, if models were enough, they would employ unskilled operators at much lower cost. In the interbank FX option market trading interventions are both common and crucial, and they are also behavioural not mathematical, and as such, cannot be modelled.

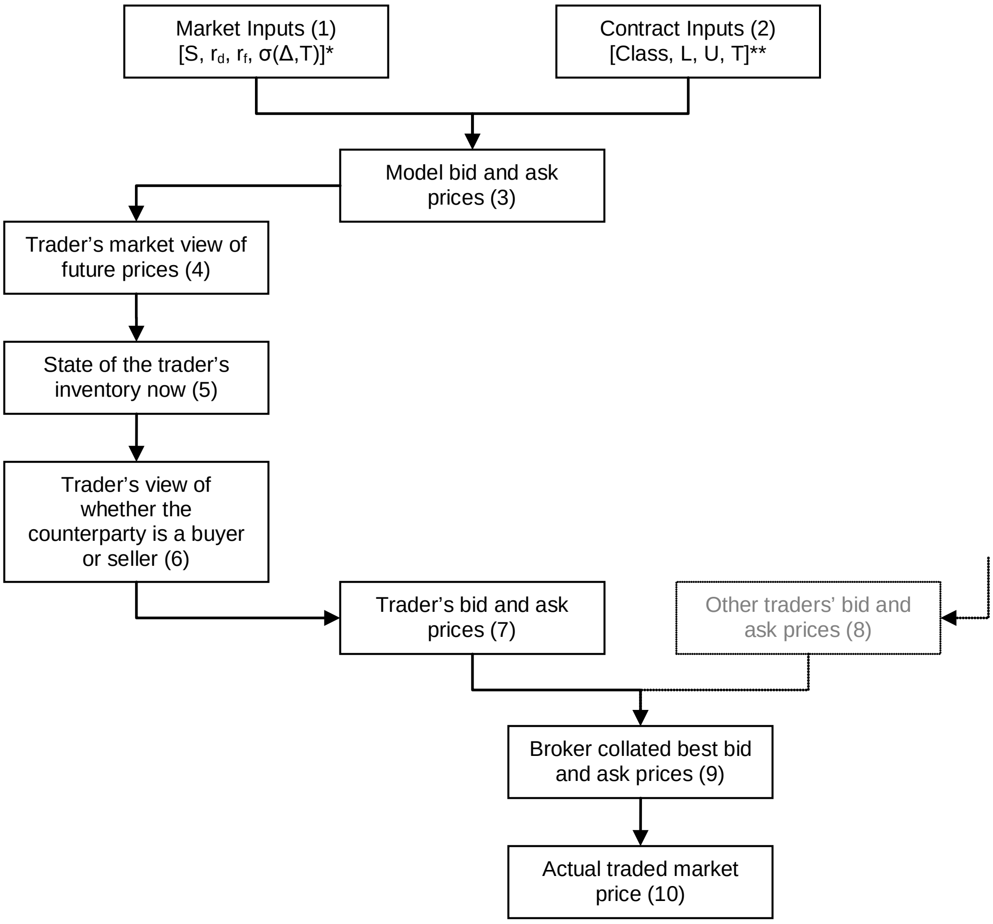

* S is the spot FX rate, r_d (r_f) is the domestic (foreign) deposit rate, and σ(Δ, T) is the European vanilla volatility surface.

** Class is OT or DNT, L (U) is the lower (upper) barrier rate and T is the expiry date.

Fig. 3.10. The interbank price-making process for exotic FX options. Model bid and ask prices are inputs in the price-making process, not outputs. The trader's market view (step 4), current inventory (step 5) and view of the counterparty's trading intention (step 6) have a significant impact on the price shown to the broker (step 7). Steps 4, 5, and 6 drive a wedge between model and market prices for options executed in the OTC interbank exotic FX option market. Sometimes banks deal direct without a broker.

Price Tests and Performance Criteria

There have been some references to 'explanatory power' and 'accuracy' in relation to the measurement of the Trader Model's pricing performance. Price tests and their performance criteria are designed to remove ambiguity from these generic terms and put empirical results on a firm objective footing.

Price tests and their performance criteria will be of two main types:

- Trader Model prices versus actual traded market prices; and
- Trader Model prices versus benchmark competitor model prices.

<u>Trader Model Versus Actual Traded Market Prices</u>

Given the prevalence of trading interventions in the interbank exotic FX option market, to analyse pricing performance effectively model prices must be tested against: (i) actual traded market prices, explicitly taking into account the impact of trading interventions; and (ii) a competitor model, to neutralise the impact of trading interventions.[72]

To take into account the impact of trading interventions we focus not just on model prices but the market within which the model is being used, in order to replicate the price-maker's role as closely as possible. It is possible to develop reasonable proxies for the current state of price-makers' inventory (step 5) and for establishing the price-maker's view of the counterparty's trading intention (step 6), which make the analysis of model and market prices more realistic. It is not possible to develop a reasonable proxy for price-makers' trading views (step 4) without more information. The empirical testing process for analysing Trader Model prices against actual traded market prices is shown in Fig. 3.11.

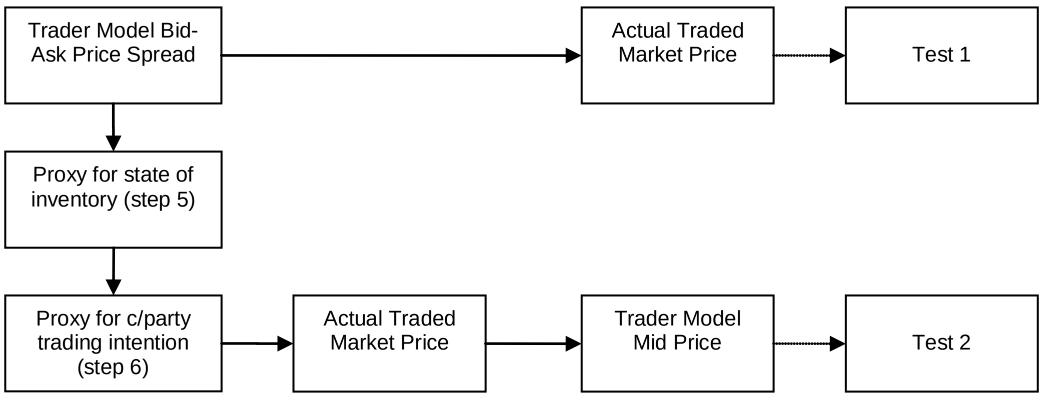

Fig. 3.11. The process for testing Trader Model prices versus actual traded market prices.

<u>Test 1</u>

Test 1 compares the Trader Model bid price and the Trader Model ask price to the actual traded market price for each OT option and DNT option in the sample. If test 1 shows that actual traded market prices consistently fall within Trader Model bid-ask spreads, then that is regarded as prima facie evidence of good pricing performance. However, test 1 has limited effectiveness on its own because of the non-trivial impact of trading interventions in the interbank exotic FX option market. If actual traded market prices do not trade within [TM_{Bid},

[72] If behavioural elements like steps 4, 5 and 6 cannot be modelled, then testing a model against another model has the effect of 'levelling the playing field'.

TM_{Ask}], it would be premature to conclude that the Trader Model is not a good decision-making tool. Further analysis is required.

It is common in the exotic option literature to reverse test 1 (Jex et al., 1999; Lipton and McGhee, 2002). That is, to place an arbitrary constant bid-ask spread symmetrically about the actual traded market price, and test whether the model mid-price falls within that range. Jex et al. (Lipton and McGhee) used an arbitrary constant spread of 0.025 (0.03) to test JPY OT (EUR DNT) options empirically. This approach is sub-optimal. Firstly, bid-ask spreads for American binary options are not symmetric. Market risk is asymmetric and so market bid-ask spreads, as distinct from constant spreads, are also asymmetric. Secondly, the database in this thesis consists of American binary option market prices traded in the interbank exotic FX option market. Pricing the interbank market (competitors) is very different from pricing corporate and institutional franchise flows (customers). In the interbank market, it is much more likely that options actually trade close to or even outside model bid and ask prices, because price-makers routinely try to extract the maximum initial revaluation profit via steps 4, 5 and 6 of Fig. 3.11.[73] This is much less likely when pricing franchise flows, as price-makers are pressured by senior management to make keen prices free of trading interventions to support the sales and distribution desk in growing the franchise business. Therefore, it is much better theoretically to develop model bid-ask spreads based on market risk, and test whether actual traded prices fall within that range. This approach also has a worthwhile corollary of ascertaining whether the model is a useful decision-making tool for price-makers in the market in practice, that is, whether the market trades in accordance with a priori expectations (e.g. whether the market trades less than or greater than model mid prices [test 2]).

Since test 1 is dependent upon Trader Model bid-ask spreads, it is important that the width of those spreads is realistic, otherwise, at the extreme, model bid-ask spreads could be made arbitrarily wide to ensure all market prices traded within the model generated spread. SuperDerivatives is a commercial vendor system which has gained a reputation for being "the standard reference for pricing exotic options up to market" (Wystup, 2003), winning academic accolades and numerous industry awards in the process. Trader Model bid-ask spreads will be analysed against those generated by SuperDerivatives as a test of their market credibility (test 4).

<u>Test 2</u>

Test 2 also analyses Trader Model prices against actual traded market prices, but this time by explicitly taking into account the impact of trading interventions. Since the database of traded exotic option prices used in this research was obtained from the interbank market, one would expect the frequency and size of trading interventions to be high. To implement test 2 in practice, one must develop reasonable proxies for price-maker trading interventions and then hypothesize ex ante whether the actual market price should trade below or above the Trader Model mid price (i.e. closer to model bid or ask prices, respectively). The proxies and hypotheses developed for this thesis are presented below.

[73] The only obligation on a price-maker making a market to a competitor bank in the interbank market, is to quote a professional bid-ask spread width shortly after being asked for a price. The price-maker has no obligation with respect to the level of the bid-ask spread, and so tries to 'squeeze' bank counterparty's for more re-valuation profit.

Proxy for the state of traders' inventory. Franchise flows, either directly or indirectly through reverse barrier options,[74] typically get price-makers net long OT options and net short DNT options (Taleb, 1997, p. 347; UBS). This is because of the prevalence of American binary options in exotic structured products, primarily to make the structured product 'zero cost'[75] and thereby cost competitive with forward contracts (Topper, 2002, p. 97). Table 3.11 shows structured products that are popular with franchise flows, and the resulting American binary position warehoused by the sell-side bank. Therefore, price-makers approaching the interbank market are more likely, ceteris paribus, to sell OT options and buy DNT options to alleviate asymmetric inventories caused by franchise flows. This is expected to result in actual interbank market prices for OT (DNT) options trading below (above) Trader Model mid prices.

Table 3.11
Franchise flow impact on sell-side bank exotic option books.

Franchise Product	Client Position	Bank Binary Position
KO Forward	RKO(Φ) - KO(-Φ)	Long OT Option
Shark Forward	Fwd + RKO	Long OT Option
Range Forward	Fwd + DNT	Short DNT Option
Range Accrual Forward	Fwd + DNTs	Short DNT Options
Boomerang Forward	KO(K_1,Φ) – RKI(K_2,-Φ) + KI(K_2,Φ)	Long OT Option
Double Shark Forward	Vanilla(Φ) - RKI(-Φ)	Long OT Option
Range Deposit	Depo(r<mkt) + DNT	Short DNT Option
Tower Deposit	Depo(r<mkt) + DNTs	Short DNT Options
Hanseatic Cross-Currency Swap	Swap – RKI(H_1,-Φ) + RKO(H_2,Φ)	Long OT Options
DNT-Linked Swap	Swap + 2 x DNT	Short DNT Options

Source: 'Franchise product' and 'client position' columns are from Wystup (2006). Φ represents whether the option is a Call (Φ = 1) or a Put (Φ = -1). K, H and r represent strike prices, barrier FX rates and deposit rates, respectively. Fwd is a FX forward contract, RKO (RKI) is a reverse knockout (knockin) barrier option; and KO (KI) is a regular knockout (knockin) barrier option.

Proxy for the trader's view of the counterparty's trading intention. This proxy operates on two levels: (i) the risk-reward ratio reflected in the theoretical value (TV) of the exotic option; and (ii) the physical location of the American binary's barrier FX rate(s). The exotic option TV is the starting reference point when quoting prices to banks and interbank brokers. The TV of the binary option can proxy for the trading intent of the counterparty because if the TV is low, such as 0.1, the buyer is risking only 0.1 for the maximum reward of 1.0, making it an attractive (unattractive) risk-reward ratio for the buyer (seller). In contrast, if the TV is 0.8, the buyer is risking 0.8 for a maximum reward of 1.0. Therefore, it is reasonable to conclude that, ceteris paribus, low TV American binary options are more likely to be bid than offered, such that low (high) TV American binary options are expected to trade in the market above (below) Trader Model mid prices.

Taleb (1997, p. 375) refers to barrier FX rates located outside the one year traded high and low in spot FX rates as a "densely mined market", owing to the extreme "buildup of gap

[74] Ayache et al. (2004, p. 34) supports the 'indirect' claim, by noting that "typically a barrier option will be dynamically hedged with a combination of the underlying, a vanilla option, and a one-touch".

[75] 'Zero cost' refers to a strategy with a net premium outlay of zero. This is achieved by a customer selling an option to finance the purchase of an option.

delta orders below the low and above the high".[76] In the event of a breach in the one year traded spot FX rate high or low, one would expect volatile movements in the spot FX rate as significant stop orders are filled in illiquid spot markets.[77] He goes on to conclude that "the expected slippage costs of gap deltas need to be increased when the risk manager has adequate information to conclude that their level is located in a densely mined market range". For OT (DNT) options, slippage cost is subtracted from (added to) the bid (ask) price of the option, as slippage costs and hence bid-ask spreads, are asymmetric.

To illustrate the application of Taleb's "densely mined market range" in this research, consider JPY binary options. For JPY binary options in this thesis, barrier FX rates located outside the range [100, 125] will be considered to be in a "densely mined market" and thus relatively unattractive to buyers (sellers) of OT (DNT) options. If the spot FX rate trades at or beyond the barrier FX rate before expiry, interbank buyers of OT options with barrier FX rates outside the range [100, 125] will have to unwind gap deltas in the same direction as the underlying spot market by buying (selling) spot FX in rising (falling) spot FX markets, thereby exacerbating losses.[78] In contrast, interbank sellers of OT options unwind gap deltas against the spot market direction, making the problem less acute. Therefore, a JPY OT option with a barrier FX rate less (more) than 100 (125) will be expected to trade below the Trader Model mid price.[79] The same argument applies for DNT options, except that the asymmetric risk is reversed. That is, the gap delta unwind is now pronounced for sellers of DNT options and less acute for buyers of DNT options. Hence, for JPY DNT options with barrier(s) outside the range [100, 125], one would expect market prices to trade above Trader Model mid prices. Whilst gap delta unwinds do affect DNT options, their impact is less than for OT options. This is because it is obvious that if a customer wants to trade a OT option with a barrier below (above) the one year spot minimum (maximum) as part of a structured product, they signal a sell interest, as they are hoping to sell an option that never pays off. Given that DNT options have another barrier in the opposite direction, the placement of a single barrier is less critical to whether the payoff occurs, and the strength of the signal is diluted accordingly.

In this thesis, Taleb's "densely mined market" is defined as a barrier FX rate located outside of the range [Low Barrier, High Barrier] in Table 3.12. For example, a JPY binary option with barrier FX rate less than 100 is considered to be in a "densely mined market", thereby

---

[76] A wider range in spot FX rates than the one year traded high and low is even more conservative, in that the gap deltas will be even more significant for those levels.

[77] A OT (DNT) option dynamically delta hedged by a price-maker will become an open spot delta position if the spot FX rate trades at or beyond a barrier FX rate.

[78] There is the facility in the interbank market for counterparties to agree to clear gap deltas with each other if the barrier FX rate trades in the spot FX market before expiry. This facility does not undermine Taleb's "densely mined market" theory or the proxy for counterparty trading intention, because in the vast majority of cases, price-makers are trading OT options interbank with the express intention of clearing excessive discontinuity risk from gap deltas caused by franchise flows trading in reverse barrier options. Franchise flows do not delta hedge as they are price-takers with natural market exposure. Banks are price-makers with neutral market exposure. Therefore, price-makers trade interbank to clear asymmetric gap delta exposure, and this cannot occur if the interbank trade neutralises gap delta.

[79] The effect of large gap deltas is to dilute the impact of the European vanilla skew on the price of a JPY up OT option with barrier level sub-100 (market convention is to refer to these options as JPY up rather than USD down). This is because the gamma of a European vanilla OTM JPY Call option with strike rate sub-100 benefits from the expected volatility of the spot FX rate until expiry. In contrast, the gamma of a JPY up OT option with barrier FX rate sub-100 benefits from the expected volatility of the spot FX rate only until the spot FX rate trades at or beyond the barrier FX rate, which could be much sooner than the OT option's expiry.

leading to expectations that OT (DNT) option prices will trade closer to model bid (ask) prices, ceteris paribus.[80]

Table 3.12
The level of barrier FX rates as a proxy of counterparty trading intention.

Spot FX Rate	Low Barrier	High Barrier
EUR	1.0400	1.3000
JPY	100.00	125.00
EUR/JPY	110.00	140.00
GBP	1.5000	1.9100
AUD	0.6500	0.8000
CAD	1.2620	1.5635
EUR/GBP	0.6485	0.7250
EUR/CHF	1.4500	1.6000

In this thesis, Taleb's (1997, p. 375) "densely mined market" is defined as a barrier FX rate located outside of the range [Low Barrier, High Barrier]. For example, a JPY binary option with barrier FX rate less than 100 is considered to be in a "densely mined market", thereby leading to expectations that OT (DNT) option prices will trade closer to model bid (ask) prices, ceteris paribus.

Table 3.13 shows that the impact of proxies can conflict for OT (DNT) options with low (high) theoretical values. Since banks only fund price-making businesses to support and grow franchise flows, it is reasonable to assume that the franchise flow effect is always prevalent and dominant. Ranking the subordinate risk-reward and gap delta effects is subjective. To reduce the impact of subjectivity, testing will be conducted across strata defined as strong (dominant proxies align) and medium (dominant proxies conflict). For OT options, since there is only one barrier, it is much easier for flow sell interests to place the barrier outside of the one year spot range to reduce the probability that the liability will be crystallised. Therefore, gap deltas dominate risk / reward effects for OT options. For DNT options, since there are two barriers, it is much harder for flow buy interests to place barriers outside the one year spot range and still have attractive terms in the exotic structured product.[81] Therefore, risk / reward considerations dominate gap deltas when flows consider structured products containing DNT options.

Table 3.13
Summary of Expected Proxy Impact.

	Inventory Proxy	Counterparty Intention Proxies	
	Franchise Flows	TV Risk-Reward	Gap Deltas
OT	Bid	Low (High) = Ask (Bid)	Bid
DNT	Ask	Low (High) = Ask (Bid)	Ask

By way of illustration, a low theoretical value (TV) DNT option is expected to trade in the market closer to the Trader Model ask price. The results of this analysis will be reported in the 'strong' stratum. A low TV OT option is expected to trade in the market closer to the Trader Model bid price, because franchise flows dominate other proxies. The results of this analysis would be reported in the 'strong' ('medium') stratum if the barrier is outside (within) the range specified in Table 3.12.

Fig. 3.12 shows a practical application of the additional performance criterion in test 2. Using the extra information contained in bid-ask spreads, which is lost by using mid-points only, allows one to infer whether the model's pricing performance is good under actual traded market

[80] The UBS exotic option trading desk refers to a "densely mined market" as a "bad neighbourhood", and bad neighbourhood risk is accounted for by modifying model prices asymmetrically as described above.
[81] As barrier width increases, the DNT option theoretical value increases. Thus, the buyer of a DNT option with widely placed barriers will have to dilute their risk / reward ratio (risk more for the same reward).

conditions and thus, whether it is a useful tool for price-makers in practice. In Fig. 3.12, example 1 and example 2 bid-ask spreads (A-B) are identical, but bid (B) and ask (A) prices are different. The market actually trades (X) within both bid-ask spreads, and test 1 takes this as prima facie evidence of good pricing performance. Test 2's additional criterion for distinguishing the level of performance within these 'good' results is to use proxies for franchise flows and counterparty trading intention to hypothesize ex ante whether X should trade below or above the mid (M) price (i.e. closer to B or A, respectively). If the proxies indicate the market price should trade below the Trader Model mid price, and it does, then this strengthens the claim that the model's pricing performance is good. Using test 1 only ignores this additional crucial information about the interbank exotic FX option market microstructure, and may result in a model with systematic pricing bias rating well if the magnitude of the bias is within the bounds of the bid-ask spread.

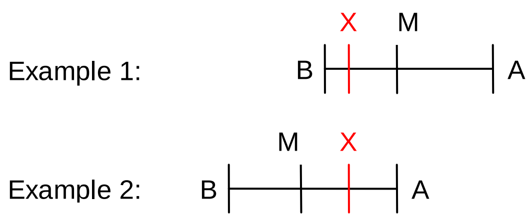

In examples 1 and 2, bid-ask spreads (A-B) are identical, but bid (B) and ask (A) prices are different. The market actually trades (X) within both bid-ask spreads, which is prima face evidence of 'good' performance in test 1. If the proxies imply that market prices should trade below model mid prices, then under test 2's additional criterion, example 1 performance is better than example 2.

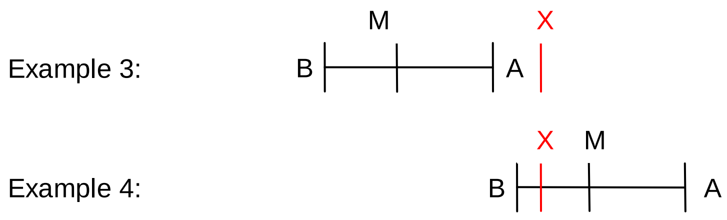

In examples 3 and 4, bid-ask spreads (A-B) are identical, but bid (B) and ask (A) prices are different. If the proxies imply that market prices should trade above model mid prices, then under test 2's additional criterion, example 3 performance is better than example 4, even though the market price trades outside the bid-ask spread of example 3 and within the bid-ask spread of example 4.

Fig. 3.12. Performance criterion for the analysis of Trader Model prices versus actual traded market prices, taking into account the impact of price-maker trading interventions. It is because of the impact of price-maker trading interventions that the position of the actual traded market price relative to the Trader Model mid-price (test 2) is as important as its position relative to the Trader Model bid-ask spread (test 1).

If the proxies imply that the market should trade above the Trader Model mid price, then under test 2, example 3's performance is better than example 4, even though the market price traded outside the bid-ask spread in the former, and within the latter. Whilst this may appear counter-intuitive, example 4 is more likely to indicate significant over-pricing which only by good fortune did not exceed the bid-ask spread. It is not uncommon for bid (offered) options to trade outside model ask (bid) prices in the interbank market. Price-makers usually know the counterparty's interest, as they typically share the same direction if not the same size of asymmetric inventories caused by franchise flows.[82] So if the price-maker prefers to buy (sell) an exotic option to clear some market risk, it is highly likely that the counterparty will as well. The price-maker's response will be to move model prices to the right (left) when quoting the counterparty, such that if their ask (bid) is paid (given) it will be to the right (left) of the model ask (bid) price. This is the very essence of interbank price-making and the wholesale intermediation of market risk: if you must increase rather than decrease market risk via an interbank trade, at least ensure that you extract a large initial revaluation profit to offset the cost of clearing risk by crossing spreads in another transaction.[83]

Trader Model vs Benchmark Competitor Model Prices

As well as testing the accuracy of the Trader Model against actual traded market prices, the Trader Model also will be tested against a competitor model in order to neutralise the impact of trading interventions. Inter-model analysis neutralises the impact of trading interventions because trading interventions are behavioural not mathematical, and, as such, all models are unable to price these effects from market and contract input data alone.

The choice of competitor model is paramount. In order to set a challenging 'best practice' pricing performance standard for the Trader Model, it must be tested empirically against a reputable, high-profile competitor model which represents the pinnacle of contemporary achievement in the exotic FX option market. However, as noted in the review of published literature, unlike the European vanilla FX option market, there is no universally accepted benchmark pricing methodology in the exotic FX option market. The plethora of published papers proposing alternate stochastic, jump and universal volatility dynamics is evidence of that. Nonetheless, there is a model that is widely used in industry called SuperDerivatives, which has clearly gained strong support among academics and financial market practitioners alike for its outstanding pricing performance:

- A leading academic financial engineer specialising in exotic FX option pricing stated that "SuperDerivatives has become a standard reference for pricing exotic FX options up to the market" (Wystup, 2003, 2006). Wystup has a strong publishing record in exotic FX option pricing, and he also has experience in investment banking, both as an employee and as a consultant, primarily in building exotic FX option pricing models. He is currently the Professor of Quantitative Finance at the Frankfurt School of Finance and Management and he is also the Managing Director of MathFinance Ag.
- Risk, the leading industry publication for derivative dealers and financial engineers claimed that "SuperDerivatives established its foreign exchange options pricing product SD-FX as something of a de facto standard globally";[84]
- SuperDerivatives was awarded "Best Pricing Technology" for two consecutive years by Risk;

[82] Sell-side price-making desks typically see the same type of franchise flows at the same time, as corporate and institutional customers try to exploit favourable market conditions to their advantage. If market conditions make an exotic structured product attractive, salesforces from different banks compete for flow business by pushing these products aggressively.

[83] Revaluation profits are calculated daily. On the trade date, revaluation profit for exotic options are equal to the difference between traded price and model mid price. Therefore, if a price-maker buys (sells) on their bid (ask) from (to) another price-maker, then they earn positive revaluation profits equal to mid-bid (ask-mid). If a price-maker wants to clear risk by acting as a price-taker in the interbank market, then they incur negative revaluation profits.

[84] Coping with complexity, 2005, 18, 12, p. 26-30.

- SuperDerivatives was awarded “Best Risk Management and Options Vendor” for three consecutive years by FX Week, an industry publication for the professional FX dealer market;
- SuperDerivatives was awarded “Best FX Option Pricing System” for two consecutive years by Euromoney, an industry publication for the professional dealer market;
- SuperDerivatives was also awarded “Best Overall System in FX” and “Most Used System in FX” in 2005 by Euromoney; and
- SuperDerivatives was commended for its sell-side innovation in the The Banker Technology Awards 2007 (a Financial Times Publication).

No other model has gained such strong support in academia and industry, and so SuperDerivatives is chosen as the benchmark competitor model for analysing the performance of the Trader Model in this research. The inter-model testing process for the Trader Model vis-à-vis SuperDerivatives is outlined schematically in Fig. 3.13. The Trader Model is not tested against SuperDerivatives for computational efficiency. The Trader Model calculates prices more quickly than SuperDerivatives, but it is not possible to separate whether SuperDerivatives' slower calculation speed is attributable to computational inefficiency or the relative inefficiency of their internet delivery mechanism.

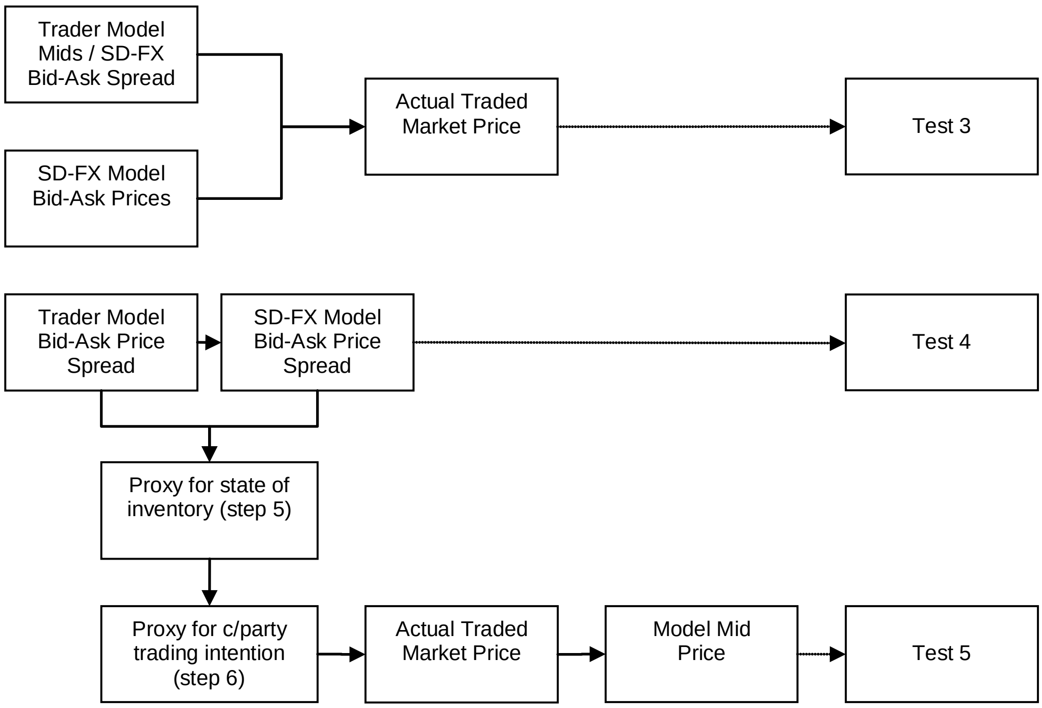

Fig. 3.13. The testing process for Trader Model prices versus SuperDerivatives model prices.

<u>Test 3</u>

Test 3 counts the number of occurrences where actual market prices traded within model bid-ask spreads. Therefore, it is essential that both models have identical bid-ask spread width, otherwise comparative pricing performance can be artificially improved by simply widening the spread of one model. Given that SuperDerivatives has an established reputation in the interbank exotic FX option market, its bid-ask spread width will be used as the benchmark for test 3. That is, Trader Model bid-ask spread width will be modified to match SuperDerivatives bid-ask spread width, as per Eq. (3.3) and Eq. (3.4):

$$TM_{Bid}^{SD} = TM_{Mid} - \left[\frac{TM_{Mid} - TM_{Bid}}{TM_{Ask} - TM_{Bid}} \times \left(SD_{Ask} - SD_{Bid} \right) \right] \tag{3.3}$$

$$TM_{Ask}^{SD} = TM_{Mid} + \left[\frac{TM_{Ask} - TM_{Mid}}{TM_{Ask} - TM_{Bid}} \times \left(SD_{Ask} - SD_{Bid} \right) \right] \tag{3.4}$$

Test 3 is analogous to test 1, with the additional step of inter-model analysis. Inter-model analysis helps to neutralise the impact of price-maker trading interventions on the analysis, because trading interventions are behavioural not mathematical, and, as such, all models struggle to price these effects from market and contract input data alone. The Trader Model's performance will be deemed better than SuperDerivatives if:

- actual market prices trade within $\left[TM_{Bid}^{SD}, TM_{Ask}^{SD} \right]$ with greater frequency than $[SD_{Bid}, SD_{Ask}]$; and
- if actual market prices trade outside $\left[TM_{Bid}^{SD}, TM_{Ask}^{SD} \right]$, then the departures are of smaller magnitude than the departures from $[SD_{Bid}, SD_{Ask}]$.

Test 4

Test 4 analyses the relative width of bid-ask spreads generated by the Trader Model and SuperDerivatives. It is essential that model bid-ask spreads are relevant and realistic in the context of the fiercely competitive interbank exotic FX option market, otherwise spreads could be made arbitrarily wide just to improve model pricing performance outcomes. Since SuperDerivatives is an industry system widely used in, and highly regarded by the interbank exotic FX option market, it is reasonable to conclude that their bid-ask spreads are relevant, realistic and competitive, and, as such, may serve as an appropriate performance benchmark.

In the extremely competitive interbank exotic FX option market, there is a natural tension between ensuring bid-ask spreads are sufficiently wide to reflect market risk accurately, but are narrow enough to compete and project professionalism. For the Trader Model to outperform under these constraints, its bid-ask spreads must be similar in magnitude to, but narrower on average than SuperDerivatives'. Since the frequency with which actual market prices trade within model-generated bid-ask spreads is a key performance criterion in test 1 and test 2, the results of test 4 will also provide additional information to interpret the results of these earlier tests.

Test 5

Test 5 is analogous to test 2, except that there is also comparative analysis between models to differentiate their relative pricing performance. Proxies of trading intervention are used to establish whether exotic options should trade below or above model mid prices. For example, if it is the former, the Trader Model outperforms SuperDerivatives if the market trades below TM_{Mid} more frequently than SD_{Mid}. Fig. 3.14 illustrates the application of test 5.

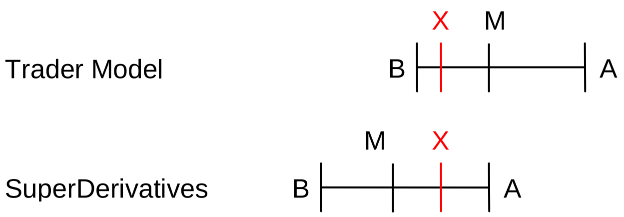

In this example, both models generate the same bid-ask spread (A-B), but have different bid (B) and ask (A) prices. The market actually trades (X) within both bid-ask spreads. A criterion for establishing better performance is to use the proxies for trading intervention to establish whether X should trade below or above the model mid price (M). If the proxies imply the market should trade below (above) the model mid price, then in this example the Trader Model (SuperDerivatives) performs best.

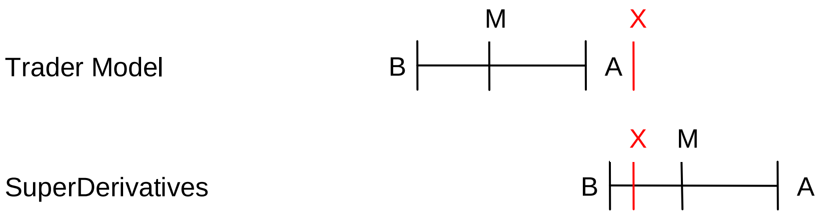

In this example, both models generate the same bid-ask spread (A-B), but have different bid (B) and ask (A) prices. The market actually trades (X) outside the Trader Model spread and within SuperDerivatives spread. If the proxies imply the market should trade above (below) the model mid price, then in this example the Trader Model (SuperDerivatives) performs best.

Fig. 3.14. Performance criterion for the analysis of Trader Model prices versus SuperDerivatives Model Prices, taking into account the impact of price-maker trading interventions. It is because of the impact of price-maker trading interventions that the position of the actual traded market price relative to model mid-prices (test 5) is as important as its position relative to the model bid-ask spreads (test 3).

<u>Traditional Tests Not Conducted in this Thesis</u>

Traditional empirical tests of in-sample closeness-of-fit and out-of-sample pricing and hedging performance, such as Belledin and Schlag (1999) are not relevant for this thesis. The Trader Model is dependent upon BSM dynamics and the exogenous volatility surface correction, and is independent of arbitrary specifications of volatility dynamics. Therefore, computationally inefficient in-sample fitting of arbitrary and possibly time-dependent free parameters to historic time series or present day cross-sectional exchange rate or option data is redundant. The non-trivial problems of over- or under-fitting, instability of estimates over time and lack of hedging insights are also redundant. As a result, we do not need to perform efficiency tests as is common in published literature, as we already know that the Trader Model is computationally much more efficient than contemporary orthodox models.

3.5. Data

3.5.1. Introduction

The data in this research was obtained from two global financial institutions that are an integral part of the interbank FX option market. It consists of time-series and cross-sectional data that is significantly larger in scale and scope than the data in any other published exotic option research. The comprehensiveness of the data provides a special opportunity to test not only the accuracy and usefulness of the Trader Model as a decision-making tool for price-makers, but also to test the robustness of the Trader Model's unique innovation under a diverse range of actual market traded scenarios.

In this research, sample input (section 3.5.2) and output (section 3.5.3) data spans an eight currency pair cross-section for the time-series 7 January 2004 to 29 September 2004, inclusive.[85] The dimensions of the cross-section and time-series were randomly chosen by the suppliers, and all data supplied was used in the empirical testing except for the exclusions described later in this chapter (specifically Table 3.24).[86]

3.5.2. Input Data

The Trader Model's structure is identical to that which is universally accepted as the benchmark for pricing European vanilla options to market. That is, in a foreign exchange context, Trader Model dynamics are identical to the Garman and Kohlhagen (1983) version of the BSM model, and the market supplement to theoretical value is priced exogenously by the volatility surface. Therefore, Trader Model prices for American binary FX options, $TM(S,r_d,r_f,t,T,\sigma[\Delta,T],L,U,\varphi)$, are dependent on actual traded market prices $(S,r_d,r_f,\sigma[\Delta,T])$ and option contract parameters (L,U,φ,t,T) only.[87] Since the Trader Model imposes no additional structure to obtain exotic option prices, such as arbitrary assumptions about the functional form of volatility dynamics, there is no requirement to process raw input data into modified input data via extensive and intensive intermediate calibrations as per contemporary orthodox models. In addition to achieving significant computational efficiency savings, making calibration redundant also ensures that information contained in raw input data is not distorted or diluted by arbitrary price-dependent model choices for calibration numerics, smoothness criteria, minimisation algorithms etc.[88]

Input data is dependent on the trade date of the exotic option. It is not uncommon for an option maturity pillar,[89] such as the three month maturity, to have a range of possible trade date dependent term lengths for volatilities and deposit rates. Volatilities are defined over expiry days (trade date to expiry date) and deposit rates are defined over delivery days (spot delivery date to option delivery date). As a result, a three month option maturity pillar may have variations like those depicted in Table 3.14 which undermine the accuracy of model calculations if not overtly taken into account. Model pricing errors from this source are eliminated in this thesis. The Trader Model was programmed into the commercial vendor system Fenics, which is owned by GFI Group Inc., so as to ensure historical day count conventions were accounted for

[85] The data is classified as input or output data. These classifications are referenced to the Trader Model, and should not be confused with the financial market practice of traders using model outputs as inputs in the price-making process (Section 3.4.2 and Fig. 3.11).

[86] The data was from the present to as far back as the providers were prepared to supply at the time the thesis first commenced (October 2004).

[87] S is the spot FX rate, r_d (r_f) is the domestic (foreign) deposit rate, t is calendar time, T is the option expiry date, $\sigma(\Delta,T)$ is the volatility surface, L (U) is the lower (upper) barrier rate, and φ is a binary variable where 1 (-1) represents a European vanilla Call (Put) option.

[88] Even if one assumes identical functional form for volatility dynamics and calibrates to the same volatility surface, using different calibration techniques will result in different exotic option model prices.

[89] OTC FX option markets have fixed maturity pillars that are routinely quoted to define the volatility surface in the maturity dimension. These pillars are overnight, one week, two week, one month, etc. A full specification of maturity pillars for EUR is shown in Table 3.19.

in empirical testing.[90] Furthermore, the data obtained for this thesis is already in Fenics format, eliminating compatibility conflicts and additional conversion processing.

Table 3.14
Example of the variation in the three month maturity pillar for EUR.

T = 3m / t	6 Jan 04	7 Jan 04	8 Jan 04	9 Jan 04
Expiry Days	91	91	90	89
Delivery Days	91	95	92	91

Volatilities are defined for expiry days, and deposit rates are defined for delivery days. These types of variations are common and must be explicitly taken into account in the pricing model, otherwise the accuracy of exotic option theoretical values will be undermined. By way of illustration, for the trade date 7 January 2004, a three month EUR option had expiry days of length 91 days, and delivery days of length 95 days.

Spot FX Rates (S)

It is convention in the interbank exotic FX option market to trade a spot delta hedge with the exotic option.[91] A delta hedge eliminates the tic-by-tic spot FX rate sensitivity of the price-maker's exotic option quote. It also conveniently establishes the reference spot FX rate specific to the actual traded market price of the exotic FX option. Therefore, the spot FX rates in this thesis are all market traded, deal-specific point estimates of which there are 183 in total, covering the currency pairs in Table 3.15. Table 3.15 shows the significant variation in spot FX rates over the sample period, making this sample an excellent vehicle for testing the robustness of the Trader Model's pricing performance. Deal-specific spot FX rate data was supplied by GFI Group Inc., one of the largest interbank OTC FX option brokers.

Table 3.15
Minimum and maximum values of spot FX rates for the period 7 January 2004 to 29 September 2004 inclusive.

Currency	Max	Min
EUR	1.2890	1.1940
JPY	114.45	105.45
EUR/JPY	136.40	131.85
GBP	1.8970	1.7670
AUD	0.7795	0.6865
EUR/CHF	1.5695	1.5080
EUR/GBP	0.6655	0.6655
CAD	1.3965	1.2920

EUR/GBP has zero variation because there was only one American binary FX option for this pair in the traded exotic option database. Source: GFI Group Inc.

Deposit Rates (r_d, r_f)

There are two sources for the deposit rates used in this thesis. GFI Group Inc. provided domestic and foreign deposit rate data specific to the delivery date of the traded exotic option as confirmed by the counterparties to the transaction. However, the Trader Model requires a term structure of deposit rates, not point estimates in order to price the impact of time on exotic

[90] Fenics stores historical day count conventions which take into account the impact of public holidays and weekends etc. on expiry and delivery days. Therefore, to value a historical transaction, one only needs to set the Fenics system date to the trade date of the transaction and set the maturity for the option, and Fenics ensures the correct expiry and delivery days are used according to interbank FX option market convention (which varies by currency pair). These system-generated dates were then cross-checked against the database of actual traded exotic option prices for accuracy.

[91] Sometimes, a vega hedge is transacted as well.

options. Vectors of deposit rates for conventional maturity pillars were obtained from Reuters LLC, one of the largest global financial service providers to the FX option markets, to supplement the deposit rate data from GFI Group Inc. Table 3.16 shows an example of the deposit rate data used in this thesis for EUR options traded on 20 January 2004. In total, there were over 9,200 deposit rate data points used in this research.

Table 3.16
Structure of deposit rate input data for EUR 20 January 2004.

	Deposit Rate Maturity Pillars					
Currency	**1w**	**1m**	**2m**	**3m**	**6m**	**1y**
EUR	1.975	2.04	2.05	2.06	2.09	2.11
USD	1.015	1.03	1.06	1.08	1.095	1.30

The one month EUR (USD) deposit rate is 2.04 (1.03) % p.a. There are two sources for deposit rates in this thesis, GFI Group Inc. and Reuters LLC. The former supplies deposit rates for the delivery date of the exotic option as agreed in the confirmation of the actual traded deal ticket. The latter supplies the vector of deposit rates for generic maturity pillars to complete the term structure. For exotic option contracts with maturity T > 1y, the deposit rates recorded in the GFI exotic option database for that maturity will be used to supplement the 1w – 1y dataset obtained from Reuters LLC.

Table 3.17 shows descriptive statistics for EUR deposit rates in this research. It is clear that there was variation in deposit rate levels and slopes during the period of empirical testing. Even though the focus in this thesis is exotic options with a maturity of one year or less, which means volatility risk is expected to dominate interest rate risk, the diversity in the range of actual traded deposit rate scenarios provides an excellent opportunity to measure the robustness of the Trader Model's pricing performance as well. Tables of descriptive statistics for the deposit rates of other currencies in this thesis are presented in Appendix B.

Table 3.17
EUR deposit rate input data description.

EUR	Max	Min	Diff
1w	2.080	1.960	0.120
1m	2.110	1.945	0.165
2m	2.110	1.895	0.215
3m	2.175	1.875	0.300
6m	2.285	1.805	0.480
12m	2.475	1.825	0.650
12m-1w	0.440	-0.245	0.685
12m-1m	0.410	-0.120	0.530

Minimum and maximum EUR-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. All rates are in percent per annum. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with slope of 44 basis points, to an inverse curve with slope of -24.5 basis points. Source: Reuters LLC.

Table 3.18 summarises the variation in level and slope for deposit rate term structures in this research. Again, there is evidence of a sufficiently wide range of values for both the level (-0.09% to 5.53%) and slope (-0.27% to 1.37%), such that it is reasonable to conclude that the actual traded market deposit rate input data supports a test of the Trader Model's robustness.

Table 3.18
Level and slope of deposit rate term structures.

	Level		Slope	
Currency	**Max**	**Min**	**Max**	**Min**
EUR	2.175	1.875	0.440	-0.245
USD	2.000	1.040	1.370	0.190
JPY	0.000	-0.090	0.155	0.005
AUD	5.530	5.300	0.620	0.115
GBP	4.970	3.940	1.015	-0.010
CHF	0.645	0.160	0.925	0.160
CAD	2.570	1.975	0.830	-0.270

'Level' is the size of the three month deposit rate in percent per annum, and 'slope' is the difference between the twelve month and one week deposit rates (12m-1w) in percent per annum. Three month deposit rates were chosen for illustration only.

European Vanilla FX Option Volatilities (σ[Δ,T])

In the interbank European vanilla FX option market volatility is traded directly as a price, it is not implied via an inversion of the BSM model.[92] Volatility is the input which is entered into the Garman and Kohlhagen (1983) version of the BSM model to find the amount of premium the buyer pays to the seller. Trading volatility directly is a convenient 'shorthand' because: (i) volatility is the only non-observable input, so it is the only variable which requires negotiation between counterparties; and (ii) interbank trading strategies are delta neutral by construction, hence spot direction is subordinated and volatility is elevated by this market practice.

The European vanilla FX option volatility data in this research is supplied by GFI Group Inc. The data consists of high frequency intraday volatilities amounting to 2,592,021 traded data points. The procedure to obtain the relevant volatility surface information per exotic option transaction is as follows:

- the GFI Group Inc. exotic FX option database contains the at-the-money (ATM) volatility used to establish the reference theoretical value (TV) of the exotic option at the time of the actual market trade;
- the GFI Group Inc. European vanilla FX option volatility database is scanned to match the trade date, currency pair, term and ATM volatility for the TV of the exotic option, so that the volatility surface data is synchronised with the actual market trade; and
- all of the European vanilla FX option volatilities which match the trade date, currency pair and time of the particular exotic option trade are collected and recorded in the format shown in Table 3.19, and illustrated in Fig. 3.15.

[92] Nevertheless, it is still commonly referred to as the implied volatility surface by market practitioners and academics alike.

Table 3.19
Structure of implied volatility input data for EUR 20 January 2004.

	0Δ		25Δ		10Δ	
Term	**Bid**	**Ask**	**Fly**	**RR**	**Fly**	**RR**
ONT	13.50	15.00	0.225	-0.400	0.500	-0.700
01W	11.00	11.80	0.225	-0.400	0.500	-0.700
02W	11.00	11.75	0.215	0.000	0.600	0.000
01M	11.00	11.20	0.250	0.100	0.625	0.175
02M	10.95	11.05	0.250	0.250	0.700	0.400
03M	10.75	10.95	0.250	0.450	0.800	0.850
06M	10.70	10.90	0.250	0.550	0.825	0.950
09M	10.55	10.70	0.250	0.600	0.850	1.000
01Y	10.65	10.75	0.250	0.700	0.875	1.100
02Y	10.55	10.85	0.275	0.700	0.900	1.200
03Y	10.60	10.90	0.300	0.650	0.950	1.100
04Y	10.50	10.85	0.325	0.625	1.050	1.050
05Y	10.40	10.80	0.350	0.600	1.150	1.000

ONT is overnight. 0Δ is zero delta straddle, which defines the at-the-money volatility. Fly and RR are (vega neutral) butterfly and (delta neutral) risk reversal, respectively. Fly and RR are available for 25 delta (i.e. European vanilla Call [Put] options with a delta of 0.25 [-0.25]) and 10 delta pillars. Together, term and delta pillars define the volatility surface. Source: GFI Group Inc.

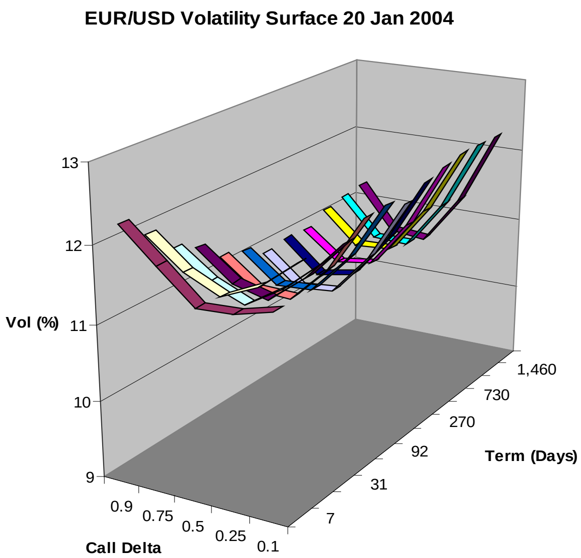

Fig. 3.15. EUR volatility surface input data 20 January 2004. A Call delta of 0.9 is equivalent to an OTM Put option. A Call delta of 0.1 is equivalent to an OTM Call option. This figure illustrates the volatility surface data in Table 3.19 from one week to five year maturities, inclusive. In contrast, the volatility surface implicit in BSM is a horizontal plane parallel to the x- and y-axes.

The depth of the volatility surface data varies by currency pair. Major units like JPY and EUR, and the cross EUR/JPY have term structure of volatility information out to 10 years, 5 years and 5 years, respectively. All other currencies have term structure of volatility information out to two years only. This was not a constraint for the exotic FX options in this thesis, as only JPY, EUR and EUR/JPY had option maturities greater than two years (the maximum option maturity was three years in length).

Table 3.20 shows maxima and minima for actual market traded volatility surface data by currency pair. It is clear insofar as the key defining characteristics of the surface are concerned, namely level, smile and skew, that volatility input data in this research was unstable and thereby offers an excellent test of robustness as well as accuracy of the Trader Model under extremely diverse, actual traded market scenarios. Appendix C shows the level, smile and skew of the volatility surface for DNT options and OT options by currency pair.

Table 3.20
Level, smile and skew of the volatility surface input data.

	Level		Smile		Skew	
FX	Max	Min	Max	Min	Max	Min
EUR	19.50	6.00	0.950	0.500	1.70	-1.75
JPY	18.00	5.00	2.300	0.750	1.75	-6.30
EUR/JPY	14.50	8.80	1.525	0.800	0.40	-2.90
GBP	17.50	9.00	0.725	0.500	2.50	-1.60
AUD	19.50	9.90	0.900	0.725	0.48	-3.20
EUR/CHF	5.90	3.40	0.770	0.385	0.45	-1.055
EUR/GBP	7.55	6.50	0.625	0.550	-0.05	-0.550
CAD	10.60	8.50	0.750	0.500	1.00	-0.600

'Level' is the size of ATM volatility in percent per annum. 'Smile' is the volatility premium for a 10 delta vega neutral butterfly in percent per annum. 'Skew' is the volatility premium for a 10 delta delta neutral risk reversal in percent per annum. 10 delta refers to OTM European vanilla options with spot deltas of 0.1 for the Call and -0.1 for the Put.

Exotic Option Contract Parameters

Exotic option contract parameters L, U, t, T are recorded in the GFI Group Inc. exotic FX option database. For example, Table 3.21 shows contract parameter inputs for American binary FX options that traded on 8 January 2004. Delivery dates automatically defined by FENICS by setting the system date to the trade date, were verified against those recorded in the exotic option database as a cross-check for accuracy.

Table 3.21
Example of American binary FX option contract parameter inputs.

FX 1	FX 2	Term	Style	Strategy	U	L	Trade	Expiry
USD	JPY	2 Yrs	DNT	Range	116.5	94	1/8/2004	1/11/2006
EUR	USD	9 Mo	OT	Up	1.4		1/8/2004	10/8/2004

Strategy refers to a range because the traded spot FX rate for the JPY DNT option was 106.15, which was within the range established by the lower (L) and upper (U) barrier rate levels. Similarly, the strategy was referred to as up because the traded spot FX rate for the EUR OT option was 1.2605, which was less than the upper barrier rate level. Source: GFI Group Inc.

3.5.3 Output Data

The output data consists of the set of all actual traded market prices for each American binary FX option in the database, expressed as a percent of payout. Tables 3.22 and 3.23 summarise by currency pair and by maturity the structure of output data for American binary FX options tested in this thesis. This data is representative of the interbank exotic FX option market in general, in that most trades are in EUR (37.7%) and JPY (38.8%); and in relatively short maturities (95% of maturities are one year or less). Nevertheless, there is a mix of FX pairs and maturities which ensures that the empirical testing in this thesis is much larger in breadth and depth than any other published research. For example, Jex et al. (1999) only tested their model with sixteen JPY OT options with three month maturities. Similarly, Lipton and McGhee (2002) only tested their model with EUR DNT options with three month maturities.[93]

Table 3.22
American binary FX option contract by currency.

FX	OT	DNT	Total
EUR	39	30	69
JPY	37	34	71
EUR/JPY	1	10	11
GBP	6	9	15
AUD	3	2	5
EUR/CHF	3	1	4
EUR/GBP	0	1	1
CAD	1	5	6
NZD	1	0	1
Total	91	92	183

OT is One Touch option, and DNT is Double-No-Touch option. Source: GFI Group Inc.

Table 3.23
American binary FX option contract by maturity (days).

Term	OT	DNT	Total
T ≤ 35	14	16	30
35 < T ≤ 95	28	27	55
95 < T ≤ 185	21	14	35
185 < T ≤ 275	7	10	17
275 < T ≤ 370	17	20	37
T > 370	4	5	9
Total	91	92	183

OT is One Touch option, and DNT is Double-No-Touch option. In the second row, there are 28 OT options with a maturity (T) 35 days < T ≤ 95 days. Rows approximate the maturity pillars of 1m, 3m, 6m, 9m, 1y and $1y^+$. Source: GFI Group Inc.

The output data of nine American binary FX options were excluded from the empirical testing in this research. Table 3.24 identifies the excluded options and the reason for their exclusion. Essentially, two exotic options did not have the corresponding complete input data, and seven exotic options had database recording errors. Fig. 3.16 and Fig. 3.17 provide a visual justification for the exclusion of the six outliers.

[93] The number of EUR DNT options tested by Lipton and McGhee is unknown, as outputs were recorded diagrammatically as solid, continuous lines, not discrete points.

Table 3.24
American binary FX option deals excluded from empirical testing.

FX	Style	Term	Barrier(s)	Reason for Exclusion
EUR	DNT*	35d	1.2490/ 1.1910	Duplicated deal with an incorrectly recorded traded market price of 0.0165 (other is 0.165).
EUR	DNT*	6m	1.3500/ 1.1950	Option length is unknown. Term is recorded as 6m but the expiry date is consistent with 12m.
JPY	DNT*	15m	126.25/ 106.25	Option length is unknown. Term is recorded as 15m but the expiry date is consistent with1m.
CAD	DNT	11m	1.365/1.215	No CAD volatility data for this deal.
EUR	OT*	2m	1.3500	Incorrectly recorded traded market price (a TV of 0.2987 cannot trade at 0.0047).
NZD	OT	6m	0.6575	No NZD volatility data for this deal.
JPY	OT	9m	97.00	Recording error in option expiry / maturity database field caused error in benchmark competitor model pricing.
EUR	OT*	6m	1.2450	Incorrectly recorded traded market price (a TV of 0.8428 cannot trade at 0.464).
JPY	OT*	95d	117.00	Incorrectly recorded traded market price (a TV of 0.1669 cannot trade at 0.0105).

* denotes options regarded as outliers in this research. Other exclusions are owing to a lack of volatility input data. d is days, and m is months. OT is One Touch option, and DNT is Double-No-Touch option.

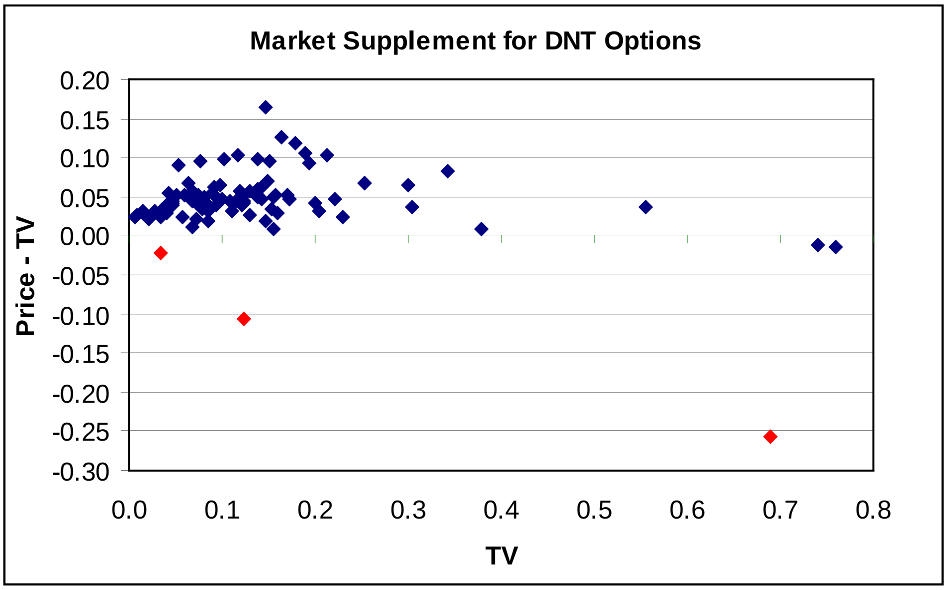

Fig. 3.16. DNT option outliers excluded from empirical testing. The green x-axis represents the set of all outputs if the market price actually traded at theoretical value (TV). The red dot-points are the three excluded DNT outliers recorded with an asterix in the option column in Table 3.24. The two data points at a theoretical value of 0.7415 and 0.7589 are not outliers. It is reasonable that high TV American binary FX options trade at a small discount to TV owing to the reversal of the risk-reward effect, as well as their high-order greeks. The data-point (0.1473, 0.1652) is not an outlier. This data-point is for a two year JPY DNT option, so it is reasonable that its market supplement can be larger than other options with much shorter maturities.

Fig. 3.16 and Fig. 3.17 show the market supplement for all American binary FX options in the exotic option database supplied by GFI Group Inc. It is clear that the sign and / or quantum of the market supplement is too extreme in the excluded cases, which are coloured red. Taking DNT options first, low theoretical value DNT options typically trade at a premium to theoretical value, resulting in a positive market supplement. Not only do the DNT data-points (0.0345, -0.0229) and (0.1232, -0.1067) have negative supplements, but, relative to their theoretical value, the negative supplements are significant in size. This is counter-intuitive. The market supplement for the data-point at (0.6891, -0.2554) is more likely to be negative because its high theoretical value reverses risk-reward effects, however, the quantum is clearly incorrect, and can only be explained by recording error. The next largest market supplement (0.1473, 0.1652) is not an outlier. This data-point is for a two year JPY DNT option, and it is reasonable that the market supplement for this option can be significantly larger than the market supplement for options with much shorter maturities.

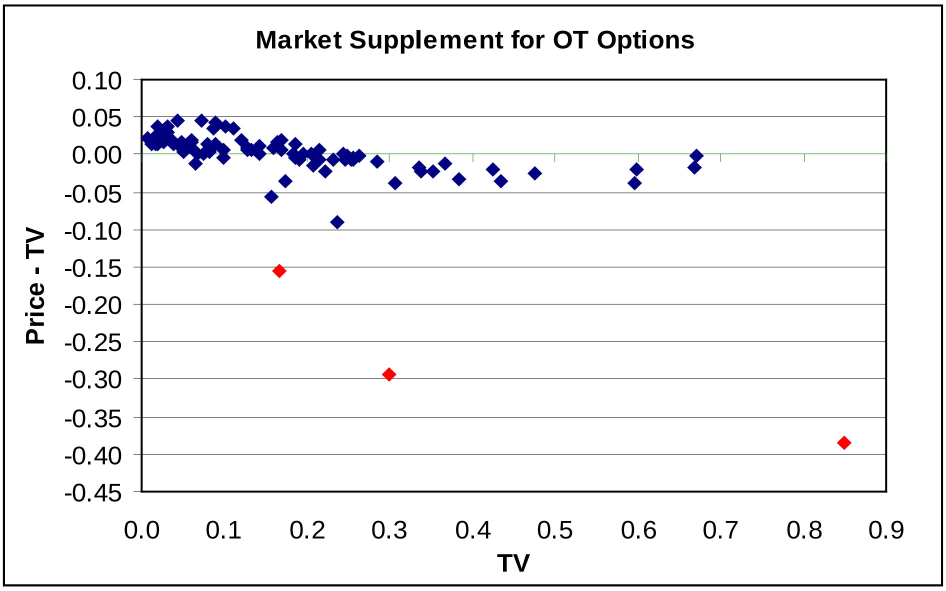

Fig. 3.17. OT option outliers excluded from the empirical testing. The green x-axis represents the set of all outputs if the market price actually traded at theoretical value (TV). The red dot-points are the three excluded OT outliers recorded in the option column in Table 3.24 with an asterix. Whilst its market supplement is larger than others', the data-point (0.2376, -0.0914) is not an outlier. This data-point is for a three year JPY option, which is the longest maturity in the database. It is reasonable that its market supplement can be significantly larger than other options with much shorter maturities.

The market supplements for the three excluded outliers for OT options have plausible sign, but implausible quantums. For comparison, the data-point (0.2376, -0.0914) is not an outlier. This data-point is a three year JPY OT option, which is the longest maturity in the exotic option database. Therefore, it is reasonable that the market supplement for this option can be significantly larger than the market supplement for options with much shorter maturities.

Table 3.25 summarises the composition of the American binary FX option database for empirical testing, by reconciling information provided separately in Tables 3.22, 3.23 and 3.24.

Table 3.25
Composition of the American binary FX option database for empirical testing.

Option	Inclusions	Exclusions	Total
OT	86	5	91
DNT	88	4	92
Total	174	9	183

This table reconciles information provided separately in Tables 3.22, 3.23 and 3.24.

3.6. Conclusion

This chapter presented the aims of the research, the philosophical and methodological rationale underpinning the model, a description of and justification for the research methods used to test the model, and a description of the data obtained for empirical testing.

The attributes of the Trader Model outlined in Section 3.3.3 add real economic value in the exotic option market. Unlike orthodox models, the Trader Model gives hedging insights, reduces model risk and improves computational efficiency, all of which contribute to making the Trader Model attractive as a decision-making tool for price-makers in practice. Chapter 4 will establish whether the Trader Model's performance substantiates this early promise.

CHAPTER 4

RESULTS

4.1. Introduction

The primary purpose of the empirical research in this chapter is to test the performance of the Trader Model. Since the data in this thesis is much more extensive than any other published research on pricing exotic options to market, the empirical results will be more rigorous and robust, and hence, a lot less susceptible to the spatial and temporal limitations of financial market data that undermine other research.

Section 4.2 analyses the performance of the Trader Model by comparing its prices against actual traded market prices, and Section 4.3 compares Trader Model prices against a best practice competitor model. Section 4.4 concludes the analysis.

4.2. Trader Model prices versus actual traded market prices

Two tests are applied in the analysis of Trader Model prices versus actual traded market prices. The first examines the number of times the actual market price traded within the Trader Model bid-ask spread [TM_{Bid}, TM_{Ask}].[94] It is a 'raw' test in the sense that price-maker trading interventions, which are common in the interbank exotic FX option market, are ignored, just like Jex et al. (1999), Lipton and McGhee (2002), and Wystup (2003). It is also implicitly assumed for now (to be verified or refuted in Section 4.3), that Trader Model bid-ask spreads are competitive and professional in an interbank exotic FX option context.

The second test incorporates the essential market microstructure of the price-making process by taking into account the non-trivial impact of franchise flows and other sources of trading interventions on interbank market prices, consistent with how prices are actually formed in wholesale financial markets in practice. This test is an empirical analogue of Bates' (2003, p. 400) and Derman's (2002, p. 82) argument that pricing models should reflect actual price-making behaviour. That is, to establish model usefulness to a price-maker, one must carry out the analysis of model pricing performance under the same realistic price-making conditions Bates and Derman advocate for price formation.

Table 4.1 shows the number of times the actual market price traded within the Trader Model bid-ask spread for OT options and DNT options. Appendix D (Appendix E) shows in figure-form the pricing results for all OT (DNT) options in this research. These appendices provide a visual representation of the overall pricing performance of the Trader Model by currency pair against both theoretical value and expiry days.

[94] Appendix A shows a schematic diagram of the bid-ask generation process for the Trader Model.

Table 4.1
Number of American binary options trading within the Trader Model bid-ask spread.

Option	$TM_{Bid} \le Mkt \le TM_{Ask}$	Total	Percent
OT	79	86	91.9
DNT	79	88	89.8
Total	158	174	90.8

For OT options 79 actual market prices traded within the Trader Model bid-ask spread, representing 91.9% of the total database of OT options traded in the interbank market. These are raw results, i.e. not including considerations for price-maker trading interventions. Mkt is the acronym for market price, and TM_{Bid} and TM_{Ask} represent Trader Model bid price and Trader Model ask price, respectively.

One must take care in interpreting the results of Table 4.1. On face value, the Trader Model appears to have merit as a decision-making tool in practice, as most American binary option market prices traded within [TM_{Bid}, TM_{Ask}]. Table 4.2, Fig. 4.1 and Fig. 4.2 give additional support to these raw results, as the size of the exceptions (where the market price trades outside [TM_{Bid}, TM_{Ask}]), are extremely small in most cases. The size of the mean exceptions shown in Table 4.2 are calculated as follows:

$$\forall \text{Mkt} \notin \left[\text{TM}_{\text{Bid}}, \text{TM}_{\text{Ask}}\right]$$

$$\text{Mkt} < \text{TM}_{\text{Bid}} \xrightarrow{\text{Defined}} \text{Lower Exception}\left(\text{LE}\right) = \text{TM}_{\text{Bid}} - \text{Mkt} \tag{4.1}$$

$$\text{Mkt} > \text{TM}_{\text{Ask}} \xrightarrow{\text{Defined}} \text{Upper Exception}\left(\text{UE}\right) = \text{Mkt} - \text{TM}_{\text{Ask}} \tag{4.2}$$

$$\text{Mean Exception} = \frac{\sum \text{LE} + \sum \text{UE}}{\text{No. of Exceptions}} \tag{4.3}$$

Table 4.2
American binary options trading outside the Trader Model bid-ask spread.

Option	Mean Exception	Percent of Spread
OT	0.0060	16.79
DNT	0.0066	17.46
Total	0.0063	17.17

For OT options where actual market prices traded outside of the Trader Model bid-ask spread, the mean size of the exception was 0.0060, or 16.79% of the bid-ask spread.

Table 4.2 shows that for American binary options that traded outside [TM_{Bid}, TM_{Ask}], the mean exception size was only a small proportion of the Trader Model bid-ask spread. To put these results in context, price-maker trading interventions can be as much as 50% of model bid-ask spreads before initial revaluation losses occur.[95] It is clear from Fig. 4.1 and Fig. 4.2 that most exceptions traded very close to TM_{Bid} or TM_{Ask}, hence most of the exception size is actually concentrated in a few isolated cases. For example, if the worst exception (out of a total of 9) is eliminated from the DNT option analysis, the mean exception size falls significantly to only 0.0038 or 9.43% of the bid-ask spread. Since price-maker trading interventions are common in the interbank exotic FX option market, the fact that most trades fall within or very close to the Trader Model bid-ask spread is prima facie evidence of strong pricing performance.

[95] For example, if a price-maker believes a bank counterparty wants to buy a DNT option with model bid-ask prices equal to 0.08-0.12, then the price-maker can show a price of 0.10-0.14 without initial revaluation loss. That is, if the bank counterparty does not behave as expected and gives the price-maker's bid (0.10), then the initial revaluation profit equals zero (trade price [0.10] minus mid [0.10] price). However, the price-maker acquires considerable benefit (volga) for zero revaluation cost (they do not pay away any spread to acquire the benefit).

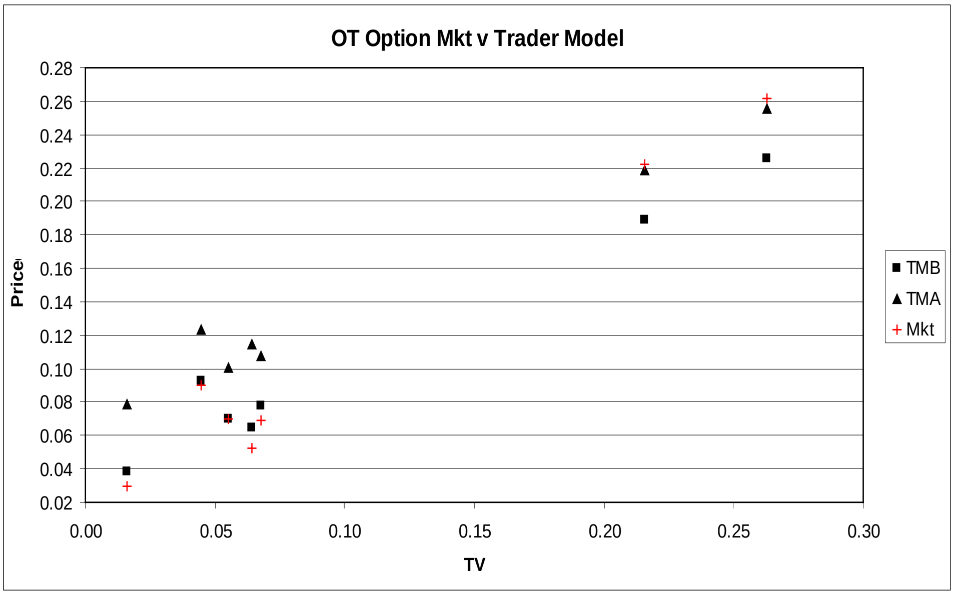

Fig. 4.1. OT option exceptions by theoretical value. OT option exceptions are when Mkt < TM_{Bid} or Mkt > TM_{Ask}, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively. Mkt is the acronym for market price, and TV is the acronym for theoretical value. All OT option exceptions are small, in that actual market prices trade very close to TM_{Bid} or TM_{Ask}.

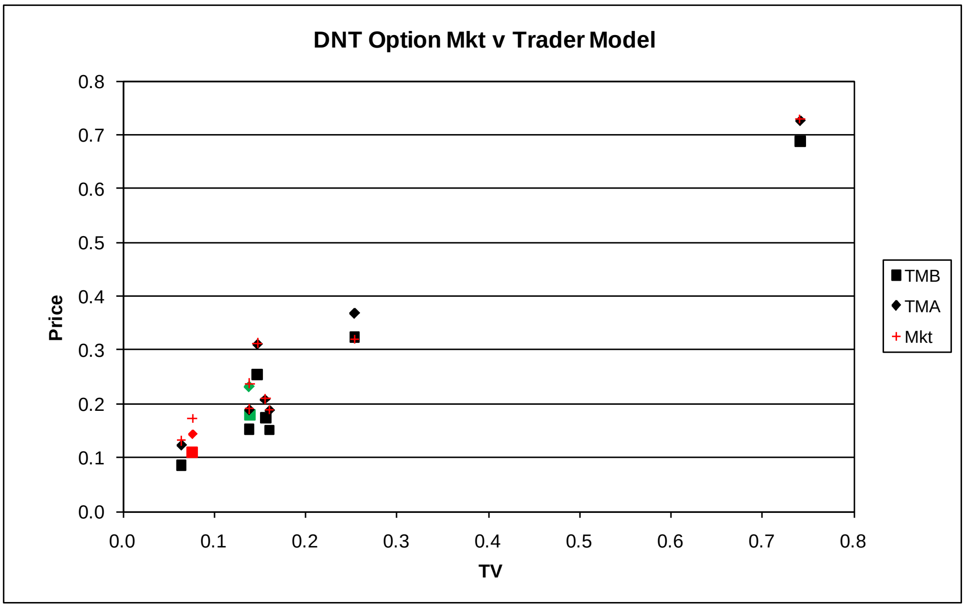

Fig. 4.2. DNT option exceptions by theoretical value. Exceptions are when Mkt < TM_{Bid} or Mkt > TM_{Ask}, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively. Mkt is the acronym for market price, and TV is the acronym for theoretical value. Most DNT option exceptions are small, in that actual market prices trade very close to TM_{Bid} or TM_{Ask}. The bid and ask price pair shown in red is the largest exception. The bid-ask prices denoted in green are to assist the reader in identifying matching groups of bid, ask and market prices, only.

Tables 4.3 and 4.4 show the breakdown of exceptions across currency pairs for OT options and DNT options, respectively. Given the much larger smiles and skews in JPY, and the larger variation in JPY smiles and skews across option maturity, it is not surprising that there are more exceptions for JPY denominated American binary options.[96]

[96] This claim is supported by the results from testing SuperDerivatives' model prices in Section 4.3. For example, Superderivatives had seven JPY DNT option exceptions, which was 21.21% of the total JPY DNT option sample. Of SuperDerivatives three OT Option exceptions, two were JPY OT options.

Table 4.3
Breakdown of OT option exceptions by FX pair.

	Up			Down				
OT	**Exceptions**	**Total**	**Percent**	**Exceptions**	**Total**	**Percent**	**Total**	**Pct**
JPY	5	23	21.74	1	12	8.33	35	17.14
EUR	1	24	4.17	0	13	0.00	37	2.70
Total	6	47	12.77	1	25	4.00	72	9.72

OT option exceptions are when $Mkt < TM_{Bid}$ or $Mkt > TM_{Ask}$, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively; and Mkt is the acronym for market price. There were five JPY up OT options out of 23 (21.74%) in the database whose market price traded outside [TM_{Bid}, TM_{Ask}]. Overall, there were six JPY OT option exceptions representing 17.14% of JPY OT options in the database.

Table 4.4
Breakdown of DNT option exceptions by FX pair.

DNT	Exceptions	Total	Percent
JPY	8	33	24.24
EUR	0	28	0.00
EUR/JPY	0	10	0.00
GBP	0	9	0.00
AUD	0	2	0.00
CAD	1	4	25.00
EUR/CHF	0	1	0.00
EUR/GBP	0	1	0.00
Total	9	88	10.23

DNT option exceptions are when $Mkt < TM_{Bid}$ or $Mkt > TM_{Ask}$, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively; and Mkt is the acronym for market price. There were eight JPY DNT options whose market price traded outside [TM_{Bid}, TM_{Ask}]. This amounted to 24.24% of the total number of JPY DNT options in the database.

Fig. 4.3 and Fig. 4.4 show that exceptions are spread across the spectrum of possible expiry days. In fact, both figures show that some of the exceptions are in very short-dated (less than one month expiry) and very long-dated (greater than one year expiry) options, where prices are affected by factors outside of volatility hedging.[97] Very short-dated options are dominated by the issue of whether the spot rate will trade at or beyond a barrier price, owing to the severity of the discontinuity risk associated with unwinding gap deltas; and long-dated options are, to a much greater extent than otherwise, influenced by interest rate risks. For OT options, two of the seven exceptions (28.6%) had expiry days less than one month. For DNT options, three of the nine exceptions (33.3%) had expiry days less than one month or greater than one year. Therefore, it is reasonable to conclude that 28.6% (33.3%) of the OT (DNT) option exceptions could be classed as extreme cases. Fig. 4.4 also shows in red the bid-ask prices for the largest DNT option exception.

[97] Whilst the scope of this thesis is restricted to American binary FX options with maturity less than or equal to one year, it is interesting to note the pricing performance of the Trader Model for all options in the database. Pricing performance was strong for long-dated options. Of the seven DNT options with maturity greater than one year, five traded within the Trader Model's bid-ask spread; and neither of the exceptions was large. Of the four OT options with maturity greater than one year, all four traded within the Trader Model's bid-ask spread.

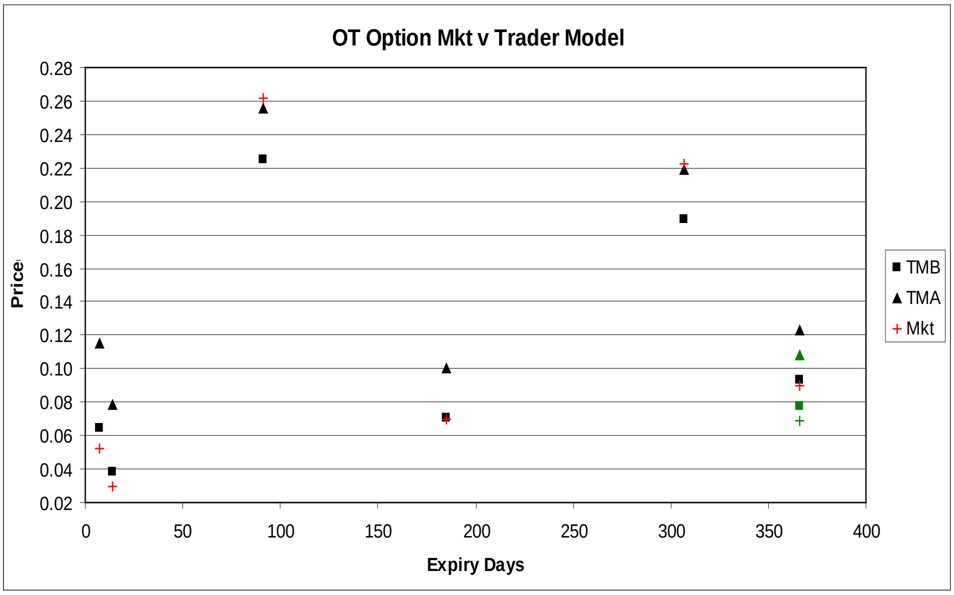

Fig. 4.3. OT option exceptions by expiry days. OT option exceptions are when Mkt < TM_{Bid} or Mkt > TM_{Ask}, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively; and Mkt is the acronym for market price. The prices denoted in green are to assist the reader in identifying matching groups of bid, ask and market prices, only.

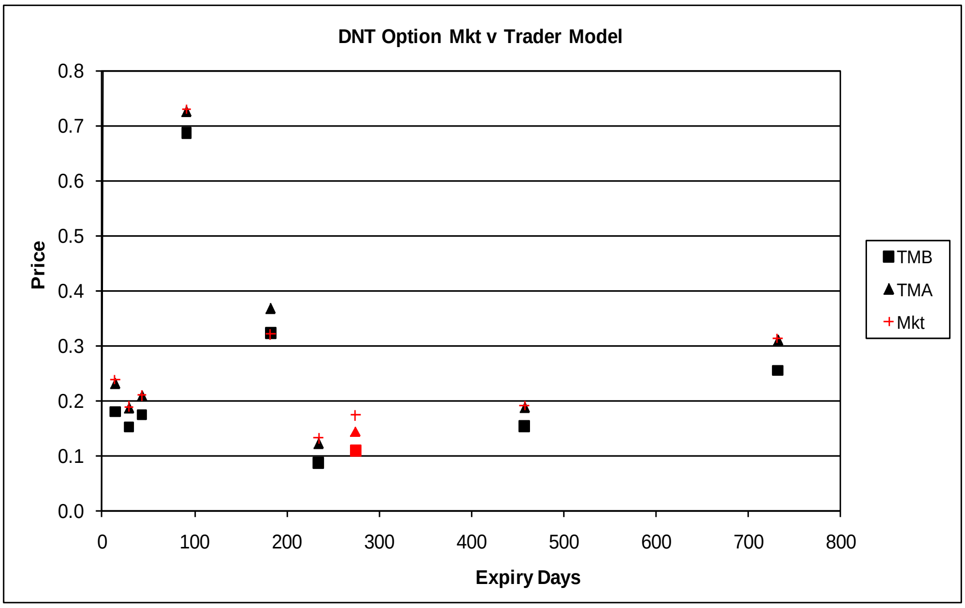

Fig. 4.4. DNT option exceptions by expiry days. DNT option exceptions are when Mkt < TM_{Bid} or Mkt > TM_{Ask}, where TM_{Bid} and TM_{Ask} are Trader Model bid price and Trader Model ask price, respectively; and Mkt is the acronym for market price. The red bid-ask price series is the largest exception.

Whilst these initial results are positive, further testing under actual market conditions is required to gain more confidence in the pricing performance of the Trader Model as a useful decision-making tool in practice. Table 4.5 introduces actual market conditions into the analysis of pricing performance by stratifying empirical results according to the relative impact of trading interventions on OT option prices. Table 4.5 shows that the relationships between Trader Model prices and actual traded market prices for OT options are consistent with hypotheses presented in Section 3.3.2. Market prices traded below Trader Model mid prices for 65.1% of OT options, which is consistent with the hypothesis that price-makers typically approach the interbank market with a sell interest to clear net long OT option positions from franchise flows.

Table 4.5 also shows that for the 40 OT options where trading interventions were expected to have strong to very strong intensity, actual market prices traded below Trader Model mid prices in 77.5% of cases. For the 46 OT options where trading interventions were expected to have only medium intensity, the result was a much more balanced 54.3%.

Table 4.5
Results for the impact of trading interventions on OT option prices.

Franchise	Risk / Reward		Gap Delta		Impact	
Flow	Low TV	High TV	$U,L \notin \left[S_{min}^{1yr}, S_{max}^{1yr}\right]$	$U,L \in \left[S_{min}^{1yr}, S_{max}^{1yr}\right]$	Intensity	No. < (≥) Mid
< Mid		< Mid	< Mid (BN)		VS	1 (1)
< Mid	> Mid		< Mid (BN)		St	30 (8)
< Mid		< Mid		< Mid	M	5 (6)
< Mid	> Mid			< Mid	M	20 (15)

U is an upper barrier price, and L a lower barrier price. S is the minimum or maximum spot rate over a one year horizon. VS is the acronym for 'very strong'; St is the acronym for 'strong'; and M stands for 'medium', impact intensity. Low TV is defined as $OT_{TV} < 0.30$, high TV is $OT_{TV} \geq 0.30$. BN is the acronym for 'bad neighbourhood', a term coined for when losses from gap delta unwinds are likely to be significant. The results reported in the last column refer to the number of OT options trading in the market at a price less (more) than the Trader Model mid price. Franchise flow effects are dominant (refer to the Methodology chapter for details). Most flows are short the barrier, and so prefer to place barriers in a bad neighbourhood to reduce the probability that they will be touched. Therefore, gap deltas dominate risk / reward for OT options. Accordingly, where flow and 'bad neighbourhood' align, this will have stronger intensity than where they do not. As expected, most OT options trade in the market at a price which is less than the Trader Model mid price. Also consistent with expectations, this effect is more pronounced as trading interventions intensify.

The empirical results for DNT options are similar to, but even more pronounced than those for OT options. This relative strength is consistent with expectations, as DNT option franchise flow and risk / reward effects compound in the strong to very strong intensities, whereas the same effects can offset for OT options. For example, low-TV DNT options are nearly always bid by franchise flows, not just corporates but also institutions like hedge funds. In contrast, low-TV OT options are usually offered by corporates, but hedge funds will sometimes take bid positions as a low risk / high reward speculative strategy. Fig. 4.5 and Table 4.6 show that the impact of franchise flows for DNT options is opposite to OT options. This is consistent with the hypothesis that price-makers typically approach the interbank market with a buy interest to clear net short DNT option positions from franchise flows, as 78.4% of DNT options traded at a market price above the Trader Model mid price. Table 4.6 confirms that for the 82 DNT options where trading interventions were expected to have strong to very strong intensity, 78.0% of the market prices for DNT options traded above the Trader Model mid price. Unlike OT options, there were not enough DNT options that traded with medium intensity to draw robust conclusions. However, in the very small sample of medium intensity DNT options, five of the six (83.3%) options traded above TM_{Mid}.

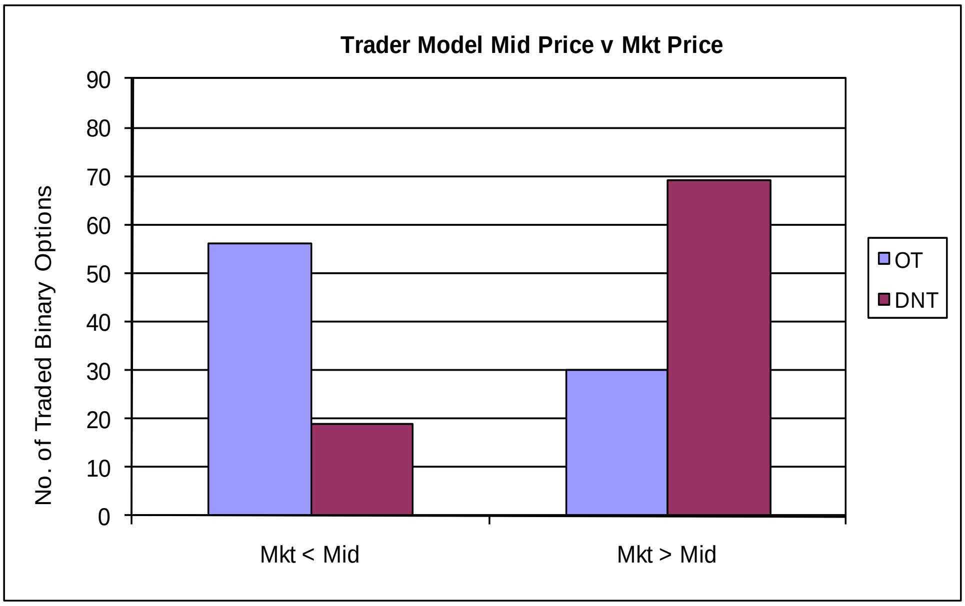

Fig. 4.5. American binary option traded prices versus Trader Model mid prices. The blue (burgundy) bars show the number of OT (DNT) options in each category. For OT options, actual market prices traded below (above) Trader Model mid-prices 56 (30) times. Fig. 4.5 shows that the Trader Model indicates that OT options are generally offered, and DNT options are generally bid, in the interbank exotic FX option market, consistent with expectations outlined in hypotheses in Section 3.4.2. The effect is stronger for DNT options, also as expected.

Table 4.6
Results for the impact of trading interventions on DNT option prices.

Franchise	Risk / Reward		Gap Delta		Impact	
Flow	Low TV	High TV	$U,L \notin \left[S^{1yr}_{min}, S^{1yr}_{max}\right]$	$U,L \in \left[S^{1yr}_{min}, S^{1yr}_{max}\right]$	Intensity	No. < (≥) Mid
> Mid	> Mid		> Mid (BN)		VS	7 (18)
> Mid	> Mid			> Mid	St	11 (46)
> Mid		< Mid	> Mid (BN)		M	0 (2)
> Mid		< Mid		> Mid	M	1 (3)

U is an upper barrier price, and L a lower barrier price. S is the minimum or maximum spot rate over a one year horizon. VS is the acronym for 'very strong'; St is the acronym for 'strong'; and M stands for 'medium', impact intensity. Low TV is defined as $DNT_{TV} < 0.30$, high TV is $DNT_{TV} \geq 0.30$. BN is the acronym for 'bad neighbourhood', a term coined for when losses from gap delta unwinds are likely to be significant. The results reported in the last column refer to the number of DNT options trading in the market at a price less (more) than the Trader Model mid price. Franchise flow effects are dominant (refer to the Methodology chapter for details). Most flows are short the barrier, but because there are two barriers, it is often not possible to place barriers in a bad neighbourhood to reduce the probability that they will be touched. Therefore, risk / reward dominates gap delta for DNT options. Accordingly, where flow and risk / reward align, this will have stronger intensity than where they do not. As expected, most DNT options trade in the market at a price which is greater than the Trader Model mid price. Too few DNT options trade in the medium intensity categories to draw robust conclusions. However, prima facie, DNT options are less affected by the intensity of trading interventions than OT options.

Therefore, the empirical results of the second test reinforce the claim that the Trader Model is a useful decision-making tool for price-makers in practice. Where market prices for OT (DNT) options were expected to trade below (above) the Trader Model mid price, they did in most cases, and especially in those cases with strong to very strong impact intensity. As a result, one can conclude that Trader Model not only generates an accurate bid-ask spread vis-à-vis the traded market price, it also positions that spread, via the mid price, in a manner which is free from bias and consistent with market expectations. To put that another way, the Trader Model does not exhibit systematic over- or under-pricing which is otherwise obscured by a wide bid-ask spread. Price-makers can be confident that the Trader Model is pricing the underlying economics of price-making, because "the fact that the model agrees to such a degree with the market prices provides confirmation that the assumptions behind the model are close to the actual market mechanism" (Jex et al., 1999, p. 12).

4.3. Trader Model prices versus benchmark competitor model prices

Since the market price data in this thesis is from the interbank exotic FX option market, it will be affected by price-maker trading interventions. Whereas Section 4.2 tested Trader Model pricing performance against actual traded market prices, explicitly taking into account the impact of trading interventions; this section tests the Trader Model against a benchmark competitor model to neutralise the impact of trading interventions. That is, as no model endogenously prices price-maker trading interventions, inter-model comparisons can be conducted 'on a level playing field'.

In this section, the Trader Model is tested against SuperDerivatives, a model widely used in the interbank exotic FX option market. Given its reputation for excellence in both academia and industry, benchmarking against SuperDerivatives is a challenging, best practice pricing performance standard for testing Trader Model's usefulness to price-makers in practice.

There are three tests for comparing the Trader Model's pricing performance against the competitor benchmark SuperDerivatives. Analogous to Section 4.2, there is a 'raw' test which examines the number of times market prices for American binary options traded within the normalised bid-ask spread of the Trader Model compared to the number of times for SuperDerivatives, where normalisation is defined by Eq. (3.3) and Eq. (3.4). The second test in this section compares the bid-ask spread width of the Trader Model against SuperDerivatives' bid-ask spread width; and the final test compares Trader Model prices against SuperDerivatives prices once price-maker trading interventions are taken into account.

In the first test, the Trader Model's 'raw' pricing performance will be regarded as better than SuperDerivatives if:

(i) actual market prices trade within the Trader Model's normalised bid-ask spread with greater frequency than they trade within SuperDerivatives' bid-ask spread, i.e. $\text{Mkt} \in \left[\text{TM}_{\text{Bid}}^{\text{SD}}, \text{TM}_{\text{Ask}}^{\text{SD}}\right] > \text{Mkt} \in \left[\text{SD}_{\text{Bid}}, \text{SD}_{\text{Ask}}\right]$; and

(ii) where actual market prices trade outside the bid-ask spread, then the mean exceptions from $\left[\text{TM}_{\text{Bid}}^{\text{SD}}, \text{TM}_{\text{Ask}}^{\text{SD}}\right]$ are of smaller magnitude than the mean exceptions from $\left[\text{SD}_{\text{Bid}}, \text{SD}_{\text{Ask}}\right]$.

As discussed in detail in the Methodology chapter, normalisation ensures that pricing performance is not artificially enhanced by making model bid-ask spreads unrealistically wide. In this test, the Trader Model's normalised bid-ask spread width is equal to SuperDerivatives' bid-ask spread width by construction.[98] However, for all American binary options in this thesis, bid and ask prices for both models differ. That is, even though $\text{TM}_{\text{Ask}}^{\text{SD}} - \text{TM}_{\text{Bid}}^{\text{SD}} \equiv \text{SD}_{\text{Ask}} - \text{SD}_{\text{Bid}}$, in all cases $\text{TM}_{\text{Bid}}^{\text{SD}} \neq \text{SD}_{\text{Bid}}$ and $\text{TM}_{\text{Ask}}^{\text{SD}} \neq \text{SD}_{\text{Ask}}$. As per Section 4.1, exceptions are defined as:

[98] Normalisation is defined by Eq. (3.3) and Eq. (3.4). Since SuperDerivatives has an established reputation in the interbank exotic FX option market, it is logical to make the Trader Model bid-ask spread width match SuperDerivatives' bid-ask spread width, rather than the other way round.

$$\forall \text{Mkt} \notin \left[\text{Model}_{\text{Bid}}, \text{Model}_{\text{Ask}}\right]$$

$$\text{Mkt} < \text{Model}_{\text{Bid}} \xrightarrow{\text{Defined}} \text{Lower Exception}(\text{LE}) = \text{Model}_{\text{Bid}} - \text{Mkt} \tag{4.4}$$

$$\text{Mkt} > \text{Model}_{\text{Ask}} \xrightarrow{\text{Defined}} \text{Upper Exception}(\text{UE}) = \text{Mkt} - \text{Model}_{\text{Ask}} \tag{4.5}$$

$$\text{Mean Exception} = \frac{\sum \text{LE} + \sum \text{UE}}{\text{No. of Exceptions}} \tag{4.6}$$

Tables 4.7 and 4.8 show the results for criteria (i) and (ii), respectively. Table 4.7 shows that actual market prices traded within the Trader Model's normalised bid-ask spread with marginally less frequency than SuperDerivatives (91% versus 92%). Table 4.8 shows that when actual market prices traded outside of the bid-ask spread, the mean exception size for the Trader Model was much smaller in magnitude for both OT options and DNT options than the exceptions of SuperDerivatives.

Table 4.7
Number of American binary options trading within SuperDerivatives' bid-ask spread.

Option	TM^{SD}	SD	Total
OT	81	83	86
DNT	77	77	88
Total	158	160	174

For OT options, 81 actual market prices traded within the normalised Trader Model bid-ask spread, and 83 actual market prices traded within the SuperDerivatives bid-ask spread. Normalisation is where Trader Model bid-ask spread width is equated to SuperDerivatives bid-ask spread width as per Eq. (3.3) and Eq. (3.4), and it is denoted TM^{SD}. SD is the acronym for SuperDerivatives. These are raw results, i.e. not including considerations for price-maker trading interventions.

Table 4.8
American binary options trading outside SuperDerivatives' bid-ask spread.

	Trader Model^{SD}		SuperDerivatives	
Option	Mean Exception	Pct of SD Spread	Mean Exception	Pct of SD Spread
OT	0.00297	7.92	0.01208	31.79
DNT	0.00823	20.59	0.01023	25.53

For OT options where the market price trades outside of the normalised Trader Model bid-ask spread, the mean exception size was 0.00297, representing 7.92% of SuperDerivatives' spread. For SuperDerivatives, the mean exception size for OT options was 0.01208, representing 31.79% of SuperDerivatives spread. Normalisation is where Trader Model bid-ask spread width is equated to SuperDerivatives spread width as per Eq. (3.3) and Eq. (3.4). These are raw results, i.e. not including considerations for price-maker trading interventions.

Therefore, without taking into account price-maker trading interventions, it is difficult to distinguish the pricing performance of the Trader Model from SuperDerivatives. The number of American binary options trading within the bid-ask spread marginally favours SuperDerivatives; whereas the mean exception size for OT and DNT options strongly favours the Trader Model. However, these results are not completely free from bias. The widespread use of SuperDerivatives by price-makers in the interbank exotic FX option market makes it an excellent benchmark, but it also means that actual traded market prices are not independent from SuperDerivatives' model prices. When price-makers use SuperDerivatives for price discovery, this model price dependency is directly transferred to actual traded market prices. In this context, it is remarkable that the empirical results of ‘raw’ pricing performance are indistinguishable. It is remarkable that SuperDerivatives’ significant advantage of being a widely used, incumbent model in the interbank market has not translated, prima facie, into significantly better raw pricing performance than the Trader Model.

The second test in this section analyses the relative width of bid-ask spreads generated by the Trader Model and SuperDerivatives. In the extremely competitive interbank exotic FX option market, there is a natural tension between ensuring bid-ask spreads are sufficiently wide to reflect market risk accurately, but are narrow enough to compete and project professionalism. For the Trader Model to outperform under these constraints, its bid-ask spreads must be similar in magnitude to, but narrower on average than SuperDerivatives’. Table 4.9 shows that Trader Model spreads are of similar magnitude to, but narrower than SuperDerivatives, for both OT options and DNT options.

Table 4.9
Mean model bid-ask spread width.

Option	Trader Model	SuperDerivatives
OT	0.0314	0.0386
DNT	0.0403	0.0410

For OT options, the mean bid-ask spread width for the Trader Model is 0.0314, and for SuperDerivatives, it is 0.0386.

Since SuperDerivatives is widely used in the interbank market in practice, the fact that Trader Model bid-ask spreads are similar to, but tighter on average than SuperDerivatives, means that on this measure the Trader Model outperforms SuperDerivatives. The results in Table 4.9 also support the results of the first test in Section 4.2, which relied on the Trader Model’s bid-ask spreads being competitive and professional in an interbank exotic FX option market setting to provide a challenging pricing performance standard. As a consequence, the results in Table 4.9 support the claim that the Trader Model’s absolute ‘raw’ pricing performance is strong.

The final test in this section compares the relative pricing performance of the Trader Model against SuperDerivatives, after taking into account the impact of price-maker trading interventions on model prices. As hypothesised in Section 3.4.2, it is expected that OT (DNT) options trading through the interbank exotic FX option market typically reflect a sell (buy) interest, such that actual market prices are expected to trade below (above) model mid-prices. That is, for this test, the Trader Model will outperform SuperDerivatives if:

- $\text{Mkt} < \text{TM}^{\text{OT}}_{\text{Mid}}$ with greater frequency than $\text{Mkt} < \text{SD}^{\text{OT}}_{\text{Mid}}$; and
- $\text{Mkt} > \text{TM}^{\text{DNT}}_{\text{Mid}}$ with greater frequency than $\text{Mkt} > \text{SD}^{\text{DNT}}_{\text{Mid}}$.

Tables 4.10 and 4.11 show the results for OT options and DNT options, respectively. For both OT options and DNT options, Trader Model prices are consistent with expectations outlined in Section 3.4.2, whereas SuperDerivatives prices are contrary to expectations. That is, the market traded below Trader Model mid-prices for 65.1% of OT options, compared to only 23.3% for SuperDerivatives. This effect is even more pronounced for those OT options where price-maker trading interventions are expected to be more frequent and of greater magnitude.

That is, for OT options with strong to very strong impact intensity, the market price traded below Trader Model mid-prices 77.5% of the time, compared to only 20% for SuperDerivatives.

Table 4.10
Inter-model comparison of pricing performance for OT options taking into account the impact of price-maker trading interventions.

	Trader Model		SuperDerivatives	
Intensity	Mkt < TM_{Mid}	Mkt > TM_{Mid}	Mkt < SD_{Mid}	Mkt > SD_{Mid}
VS	1	1	0	2
St	30	8	8	30
M	25	21	12	34
Total	56	30	20	66

Mkt is the actual traded market price of the OT option. TM_{Mid} and SD_{Mid} are the mid-prices from the Trader Model and SuperDerivatives, respectively. VS is the acronym for 'very strong'; St is the acronym for 'strong'; and M is the acronym for 'medium', impact intensity. For the 40 OT options with strong to very strong impact intensity, the market price was less than the Trader Model (SuperDerivatives) mid-price 77.5% (20%) of the time, consistent with (contrary to) hypotheses in Section 3.4.2.

As expected, DNT option results are the reverse of OT options. The market price traded above Trader Model mid-prices for 78.4% of DNT options, compared to only 29.5% for SuperDerivatives. Where price-maker trading interventions are expected to be more frequent and of greater magnitude, that is, for DNT options with strong to very strong impact intensity, the market price traded above Trader Model mid-prices 78% of the time, compared to only 28% for SuperDerivatives.

Table 4.11
Inter-model comparison of pricing performance for DNT options taking into account the impact of price-maker trading interventions.

	Trader Model		SuperDerivatives	
Intensity	Mkt < TM_{Mid}	Mkt > TM_{Mid}	Mkt < SD_{Mid}	Mkt > SD_{Mid}
VS	7	18	17	8
St	11	46	42	15
M	1	5	3	3
Total	19	69	62	26

Mkt is the actual traded market price of the DNT option. TM_{Mid} and SD_{Mid} are the mid-prices from the Trader Model and SuperDerivatives, respectively. VS is the acronym for 'very strong'; St is the acronym for 'strong'; and M is the acronym for 'medium', impact intensity. For the 82 DNT options with strong to very strong impact intensity, the market price was greater than the Trader Model (SuperDerivatives) mid-price 78% (28%) of the time, consistent with (contrary to) hypotheses in Section 3.4.2.

Therefore, not only were Trader Model prices consistent with expectations of price-maker trading interventions, but where price-maker trading interventions were expected to have greater frequency and greater magnitude, the results were even better. In contrast, SuperDerivatives' results were contrary to expectations, and where price-maker trading interventions were expected to have greater frequency and greater magnitude, the results were even worse. Fig. 4.6 and Fig. 4.7 show this effect diagrammatically. Fig. 4.6 shows that the results for the Trader Model are consistent with expectations, both within each American binary option class as well as across option class (e.g. results are more pronounced for DNT options). Whilst the Trader Model's absolute pricing performance is strong, the results are even stronger when pricing performance relative to SuperDerivatives is considered.

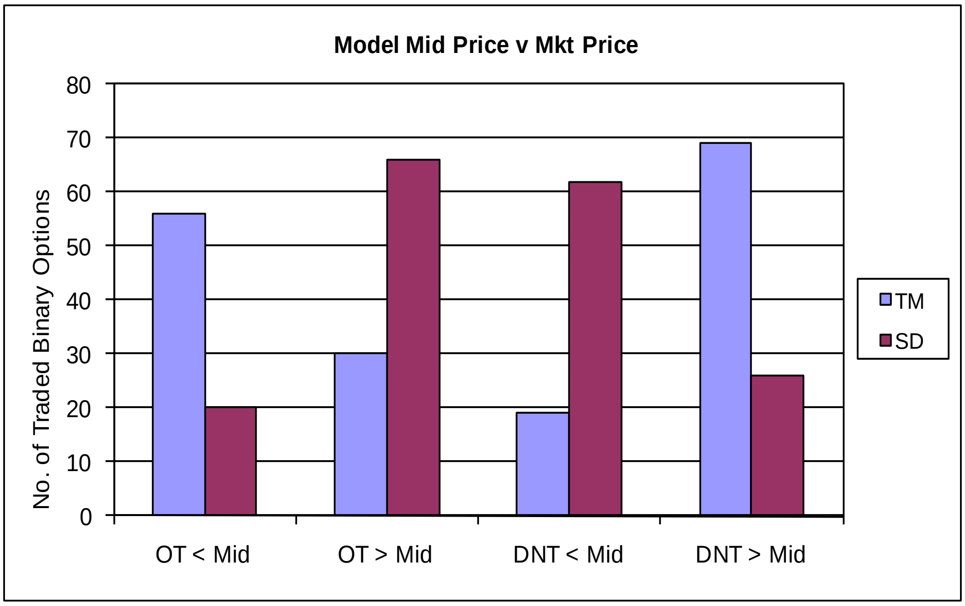

Fig. 4.6. American binary option traded market prices versus model mid prices. The blue (burgundy) bars show the number of American binary options in each category for the Trader Model (SuperDerivatives). For OT options, actual market prices traded below (above) Trader Model mid-prices 56 (30) times. Fig. 4.6 shows that for the Trader Model OT options are generally offered, and DNT options are generally bid in the interbank exotic FX option market, consistent with expectations outlined in hypotheses in Section 3.4.2. The effect is stronger for DNT options, also as expected. Fig. 4.6 shows that the Trader Model's relative pricing performance vis-à-vis SuperDerivatives was even stronger than its absolute pricing performance.

Fig. 4.7 shows that the pricing performance of the Trader Model improved, both absolutely and relative to SuperDerivatives, for American binary options in the strong to very strong impact intensity. American binary options within these categories of impact intensity should exhibit more pronounced trader interventions. Fig. 4.7 shows that the results are more pronounced for the Trader Model, both absolutely and (especially) relative to SuperDerivatives.

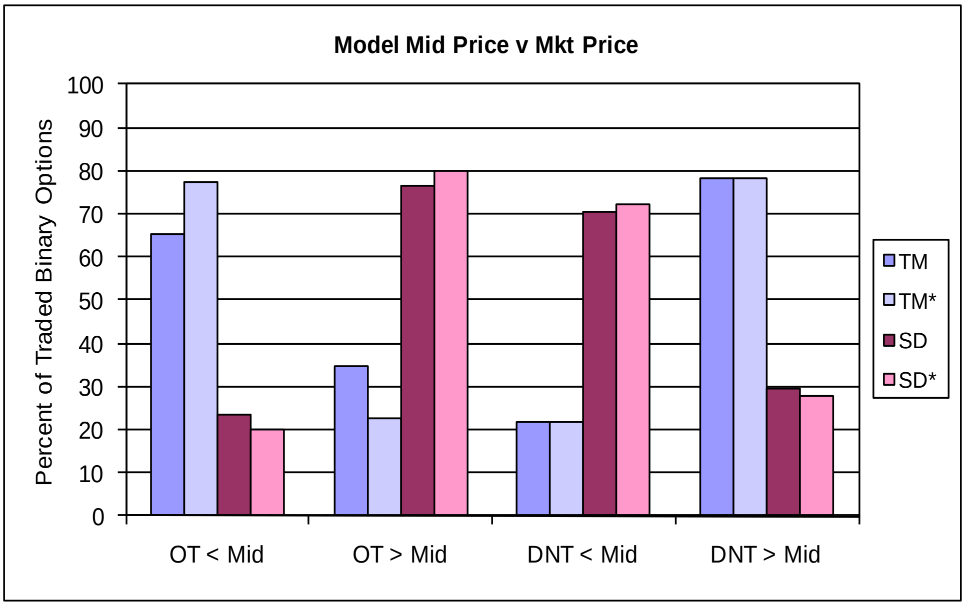

Fig. 4.7. American binary option traded market prices versus model mid prices, taking into account the impact of price-maker trading interventions. The blue (burgundy) bars show the percentage of American binary options in each category for the Trader Model (SuperDerivatives). The lighter shades of blue and burgundy, i.e., TM* and SD* show the results for the Trader Model and SuperDerivatives for American binary options in the strong to very strong impact intensity. Fig. 4.7 shows that for the Trader Model OT options are even more likely to be offered in the interbank exotic FX option market as impact intensity strengthens, consistent with expectations outlined in hypotheses in Section 3.4.2. Trader Model DNT option prices are insensitive to impact intensity, but that could reflect the small sample size for medium intensity options (6). SuperDerivatives OT and DNT option prices behave contrary to expectations. Fig. 4.7 shows that the already strong absolute pricing performance is improved further when pricing performance relative to SuperDerivatives is considered.

Therefore, the results in this section strongly support the claim that the Trader Model is a useful decision-making tool for price-makers in practice. The Trader Model is accurate relative to the competitor benchmark SuperDerivatives. Furthermore, empirical results strongly support the notion that the Trader Model captures the essential financial economics underpinning the wholesale intermediation of market risk in the price-making process, both in an absolute sense, but especially relative to SuperDerivatives. Unfortunately, it is not possible to reconcile or to explain price differences between the Trader Model and SuperDerivatives, as SuperDerivatives has a policy of not releasing details of their model, citing the need to keep proprietary information confidential for commercial reasons.

4.4. Conclusion

The empirical testing of the Trader Model in this thesis is much larger in both scale and scope than any other published research on the pricing of exotic options to market. The results of the empirical testing show that the Trader Model's pricing performance is robust to a broad range of diverse market conditions and option contract specifications, both absolutely and relative to a commercially successful competitor benchmark. The accuracy and robustness of the Trader Model, given the rigour of the empirical testing, provides scope to draw inferences not only about the Trader Model itself, but also about the key factors driving the market mechanism for pricing exotic options to market.

Lipton and McGhee (2002, p. 85) claim that "only a universal volatility model is capable of matching the market [price] properly. In our experience, this conclusion is valid for almost all path-dependent options". The pricing performance of the Trader Model, both absolutely against actual traded market prices, and relatively against SuperDerivatives, is sufficiently strong to question their claim. Furthermore, unlike Lipton and McGhee (and SuperDerivatives), the Trader Model identifies a unique hedge portfolio that is priced, and easily implemented in highly liquid, actual traded markets. This is a significant economic benefit to a price-maker managing an exotic option book in the market.

CHAPTER 5

CONCLUSION

5.1. Findings

The overall aim of this thesis was to develop and test empirically a model for pricing American binary FX options to market which: (i) predicts actual traded market prices with sufficient accuracy to be a useful decision-making tool for price-makers in practice; (ii) identifies and quantifies market risk in a manner which provides unique insights into risk management for price-makers in practice; and (iii) achieves real savings in computational efficiency relative to best practice quantitative models exemplified by the universal volatility models of Jex et al. (1999) and Lipton and McGhee (2002). This research was motivated by the extreme dependency of sell-side banks on exotic option pricing models. Sell-side banks rely on these models for a broad range of critical tasks from pricing market risk to identifying hedging strategies, calculating daily profit and loss, defining risk limits, reporting to key stakeholders both internally and externally, as well as determining trader bonuses and Basel II capital retention levels. It is through this extreme model dependency that banks are exposed to model risk, defined as the risk that reported daily mark-to-model profits do not reflect financial economic substance.

The model developed in this research, referred to as the Trader Model, was tested empirically in a manner similar to Jex et al. (1999) and Lipton and McGhee (2002). That is, Trader Model prices for American binary FX options were compared with actual traded market prices and the prices of a competitor model. American binary FX options were chosen not only to assist with comparisons to published literature, but also because:

- binary option prices are extremely sensitive to the traded volatility surface;
- as a “smallest decomposable fragment” (Taleb, 1997, p. vi) binary optionality is elemental to, and embedded in, other exotic options with significant traded volume; and
- their extreme sensitivity to the traded volatility surface means that binary options are “the right security” to exploit arbitrage opportunities emanating from arbitrary specifications of smile dynamics that are an essential component of contemporary orthodox pricing methodologies (Ayache et al., 2004, p. 11).

However, in contrast to Jex et al. and Lipton and McGhee, a large database of actual traded market prices was used, and the competitor model is widely regarded as best practice both within academia and industry. Therefore, the performance benchmarks for the Trader Model are much more challenging and rigorous than the performance benchmarks in published literature, which, in turn, underpins the robustness of inferences drawn in this thesis.

Furthermore, in stark contrast to published literature, the empirical research in this thesis explicitly takes into account that traded price data is formed in the interbank market. As a result, the empirical research in this thesis uniquely incorporates the price-maker / model nexus into model testing, which is internally consistent with recommendations by prominent financial engineers to incorporate this nexus into model pricing.[99] Therefore, testing is not restricted to the frequency with which the actual market price traded within the model bid-ask spread, but also whether the actual market price should, given our knowledge of sell-side bank franchise flows and its effect on price-maker behaviour, trade closer to the bid price or the ask price. The tests incorporate this additional information which is critical to the wholesale intermediation of market risk. Doing so enables not only the coarse grade pricing structure of the model bid-ask spread to be tested, but also the fine grade pricing structure of the position of the market price relative to the model bid-ask spread. Both are crucial to the economic fundamentals underpinning price-making in practice. Again, this additional performance criterion is more challenging than criteria used in the existing literature.

It was found in this thesis that the Trader Model predicts actual traded market prices with a high degree of accuracy, not only absolutely but relative to the best practice competitor model benchmark. Actual market prices traded within the Trader Model bid-ask spread for 92% of OT options, and 90% of DNT options. This is a remarkable outcome given that (i) price-makers routinely move model prices to the right (left) to reflect bid (offered) interest to bank counterparties; and (ii) Trader Model prices were completely independent of actual traded market prices, because the Trader Model was not used by interbank price-makers during the testing period.[100] The Trader Model also performed strongly against the competitor model. After neutralising differences in bid-ask spread width, the Trader Model's coarse grade pricing performance was almost identical to the competitor model. While the Trader Model had two fewer OT options trading within the normalised spread than the competitor model, the Trader Model's exceptions were of much smaller magnitude than the exceptions of the competitor model. This strong relative pricing performance was achieved despite the competitor model having the significant advantage of being a commercially successful model used by price-makers, and, therefore, having its model prices influence traded prices themselves.

Whilst the Trader Model's coarse grade pricing performance was excellent, its fine grade pricing performance was even better. This is crucial, as fine grade pricing performance is a key determinant of the usefulness of the Trader Model as a decision-making tool in practice. It is well known that price-makers are usually net long binary discontinuity from franchise flows, i.e., net long OT options[101] and net short DNT options (Wystup, 2006; UBS). Therefore, one would expect price-makers who clear excess discontinuity risk via the interbank exotic FX option market to show, in general, offered interest for OT options and bid interest for DNT options, ceteris paribus. The Trader Model reflected offered interest for OT options, as traded market prices were less than Trader Model mid-prices for 65% of OT options in the database. The Trader Model also reflected bid interest for DNT options, as traded market prices were greater than Trader Model mid-prices for 78% of DNT options. In contrast, the benchmark competitor model reflected the opposite, namely bid interest for OT options (77%) and offered interest for DNT options (70%).

99 Derman (1996, p. 5); Bates (2003, p. 400); Derman (2002, p. 82); and Ayache et al. (2004, p. 33).
100 The Trader Model was not licensed for use by sell-side banks until the last quarter of 2008.
101 Either directly or indirectly via reverse barrier options.

The strength of the Trader Model's pricing performance, not only in respect of its accuracy but also in respect of its robustness to inputs that vary on both time series and cross-sectional bases, demonstrates that the Trader Model's unique innovation closely reflects the actual market pricing mechanism for American binary FX options. That finding alone is significant. In addition, because the Trader Model directly prices the cost of liquid, commoditised, market traded hedging strategies, it is also able to identify and quantify market risk in a manner which provides unique insights into risk management for price-makers in practice. That is, price-makers can implement a liquid market traded hedge that is internally consistent with the model price. In contrast, Jex et al. (1999) and Lipton and McGhee (2002), for example, derive model prices independently of a market traded hedge. They calibrate to the entire volatility surface, but cannot identify nor quantify an implementable market traded hedge that is internally consistent with their model price. The benchmark competitor model also does not identify a market traded hedge.

In addition to making price discovery more opaque, and economic inference more difficult, numerical calibration to the volatility surface has also made contemporary orthodox pricing methodologies computationally inefficient. In fact, the protracted search for marginal gains in computational efficiency in numerical calibration is described well by Ayache et al.'s (2004, p. 33) observation that "quantitative finance seems to be wasting itself in sophisticated mathematical exercise". By making numerical calibration redundant, the Trader Model achieves significant real savings in computational efficiency relative to best practice quantitative models.

5.2. Trader Model attributes

The path of evolution of exotic option pricing models has increased bank exposure to model risk. Orthodox models for pricing exotic options to market are so mathematically complex that price discovery is opaque and only tenuously linked to the underlying financial economics of the market. To quote Rebonato (2002, p. 7), "very often . . . the adoption of new models has been driven not by their superior explanatory power, but by . . . technological and numerical advances". In contrast, the Trader Model is driven by financial economics because it is price-makers who make prices, not models. By explicitly taking into account the prominent role of price-makers in price discovery the Trader Model effectively redefines the modelling problem from pricing to hedging centric. This is consistent with Bates' (2003, p. 400) call for "a renewed focus on the explicit financial intermediation of the underlying risks by option market makers". Financial intermediation is primarily a hedging issue. Price, on the other hand, is a secondary consideration to establish the premium required by a price-maker to cover the cost of intermediation. Since the sell-side FX option business is a low margin / high volume proposition, it is imperative that price articulates closely with hedging costs, otherwise "selling optionality too cheaply is likely to cause an uneven but steady haemorrhaging of money out of the book, and can ultimately cause the demise of the trader" (Rebonato, 2002, p. 1).

The renewed focus on hedging is pivotal. In contrast to the universally accepted benchmark model for pricing European vanilla FX options to market,[102] there is no universally

[102] Garman and Kohlhagen (1983) model derived under the BSM paradigm, in conjunction with the exogenous volatility surface.

accepted benchmark model for pricing exotic FX options to market. However, there is general agreement on hedging, albeit with disagreement on modelling philosophy and therefore on the (model-dependent) price. This general agreement on hedging is not strictly observed, for example price-makers will not trade the same hedge at the same time, as they have different inventory and different views on the market. Nevertheless, a price-maker's "option book . . . will generally be neutral in the lower moments and exposed to various risks in the higher moments" (Taleb, 1997, p. 149).[103] Because price-makers must use traded instruments to hedge high-order risk, the obvious immediate hedging choices are liquid, highly commoditised trading strategies like butterflies and risk reversals, even for price-makers with orthodox models based on arbitrary volatility dynamics, such as universal volatility models. As a result, hedging by price-makers is relatively model-independent, whereas pricing is extremely model-dependent. This is a crucial point not only theoretically, but also from a practical financial economic perspective. It is the price-maker's hedging activity which binds the traded volatility surface to exotic option prices, not the arbitrary dynamical behaviour of non-traded parameters imposed by a financial engineer.

Price-makers must hedge in practice with market traded instruments. As the Trader Model specifically preserves the direct one-to-one mapping between price and hedging cost, trading desk profits calculated with the Trader Model will be internally consistent between exotic options and the vanilla options that hedge them. The key outcome for sell-side banks is significantly reduced model risk. Orthodox models calculate a price independently of a traded market hedge, thereby disarticulating price and hedging costs which can result in highly asymmetric profits from exotic option mark-to-model and vanilla option mark-to-market valuations. Price-makers' actual hedging activity cannot bind arbitrary whole-of-volatility-surface numerical calibrations to the market price of an exotic option, because there is no traded whole-of-volatility-surface hedge.

The Trader Model is a unique application of what is now known as the vanna-volga method. The vanna-volga method is intuitive, simple, transparent, and computationally efficient, and it reflects actual price-maker hedging behaviour. However, prices derived from the vanna-volga method have not achieved a satisfactory performance standard when compared to actual traded market prices. To overcome this unsatisfactory pricing performance, financial engineers introduced arbitrary constants to scale the raw solution; and developed other mathematically-laden and opaque methodologies that price independently of the cost of a traded market hedge. Both responses significantly increase model risk. The Trader Model's innovation is to recognise and exploit a key economic insight to develop a unique application of the vanna-volga method that achieves strong pricing performance without requiring arbitrary constants, thereby maximising pricing and hedging intuition, and minimising model risk.

The Trader Model's unique innovation is the discovery of an intuitive and simple concept for pricing exotic options to market that ensures there is articulation between the market price of exotic options and the market price of vanilla options. Articulation is essential for price-making to be sustainable in practice. In so doing, the Trader Model does not require

[103] Part of the reason for this phenomenon is that trading book limits are usually described in lower moment terms, i.e. first-order greeks.

arbitrary scaling constants to price the risk of early termination. The Trader Model also ensures that the divergence between volatilities of the forward and volatilities of the spot is taken into account in an internally consistent manner for all option contracts, irrespective of whether they are American or European.

The Trader Model prices exotic options by calculating the cost of the European vanilla option hedge, where the maturity of the hedge is equal to the expected stopping time of the exotic option. This is eminently logical, as a hedge is only a hedge if it offsets the risk of something else. In those cases where the exotic option terminates early, it is optimal to have the hedge terminate at the same time, otherwise a residual open position exists. Equating the maturity of the hedge to the expected stopping time of the exotic option means that only the cost of the hedge needs to be priced. This is computationally efficient and minimises hedging costs. In contrast, other applications of the vanna-volga method equate the maturity of the hedge to the maturity of the exotic option, thereby making it necessary to price not only the cost of the hedge, but also the cost of unwinding the hedge when the exotic option terminates early. All these applications of the vanna-volga method are computationally inefficient and cause excessive hedging costs.

By equating the maturity of the hedge to the expected stopping time of the exotic option, the Trader Model also successfully resolves the dilemma caused by the divergence between the volatility of the forward and the volatility of the spot. As the strength of American optionality intensifies, such as high (low) theoretical value OT (DNT) options, the expected stopping time of the exotic option shortens. The Trader Model uses volatilities at the expected stopping time to value the market supplement to theoretical value, not the much longer maturity. In these cases, volatilities of the forward at the expected stopping time closely resemble volatilities of the spot, as short-dated European vanilla options have low levels of interest rate risk. As American optionality weakens, such as low (high) theoretical value OT (DNT) options, the expected stopping time lengthens. Since in these cases American exotic options more closely resemble European options, the Trader Model uses volatilities approaching the maturity to value the market supplement to theoretical value. Given that all FX option markets exhibit term structures of volatility, and term structures are also present in smiles and skews, the Trader Model is the only model that prices the cost of a liquid market-traded hedge which incorporates term risk.

When pricing exotic options to market, Derman's (2003, p. 13) advice is germane:

> "As in most social sciences, the big and interesting battle in options theory and the smile is to avoid being Utopian, and instead to try to pick methods and models whose results depend as little as possible on unverified, indeed unverifiable, assumptions".

The Trader Model does not depend on "unverified, indeed unverifiable, assumptions", instead it relies on the same assumptions as the universally accepted benchmark pricing model for European vanilla options. In the FX market, the volatility surface is an exogenous market pricing correction to the Garman and Kohlhagen (1983) model for European vanilla options that reflects the additional costs of hedging volga and vanna. Therefore, volatilities are specific to market risk and model dynamics. The Trader Model is a pragmatic rather than Utopian mechanism for translating the European vanilla option market supplement - the volatility surface - to an internally consistent market supplement for American exotic options. The empirical results in this thesis, which are much larger in scale and scope than any other published research on pricing exotic options to market, show that the pricing performance of the Trader Model is very strong, not only vis-à-vis traded market prices, but also relative to a competitor benchmark with an established reputation in academia and industry as "a standard reference for FX options prices" (Wystup, 2006, p. 303). Like Jex et al. (1999, p. 12), we conclude that because of "the fact that the model agrees to such a degree with the market prices provides confirmation that the assumptions behind the model are close to the actual market mechanism". Unlike Jex et al., the Trader Model's very strong pricing performance was achieved without having to abandon the parsimonious Garman and Kohlhagen dynamics universally accepted for pricing European vanilla FX options, and without having to dilute or distort the information content in the traded volatility surface by applying it out of context. Jex et al., like all universal volatility, stochastic volatility and jump diffusion models, rely on elaborate, unverifiable dynamics that are not universally accepted, and where the volatility surface is no longer a set of model price adjustments for specific market risks under Garman and Kohlhagen dynamics, but is instead a convenient mechanism for quantifying arbitrary, non-traded parameters that have no connection to the price-maker's price adjustment.

5.3. Extensions of the research

This research has established that the Trader Model is an effective decision-making tool for price-makers who trade American binary FX options in practice. This research used the largest data sample of actual traded exotic option prices that could be obtained by the author, which, was larger than those used in published literature. Whilst the scale and scope of the sample of actual traded data supports the robustness of the Trader Model's pricing performance to a wide range of market conditions and contract specifications, extensions of this research would be to test empirically the pricing performance of the Trader Model for:

- OT options and DNT options traded during a time period different to the data sample in this research (7 January 2004 to 29 September 2004, inclusive);
- other exotic FX options, such as regular and reverse barrier options; and
- other markets, such as equity and commodity markets.

Each of these extensions test the robustness of the Trader Model's pricing performance, and hence its unique philosophy and methodology, to an even broader range of market conditions and contract specifications. However, each of these extensions must overcome the challenge of obtaining actual traded market data from over-the-counter financial markets that are extremely difficult for non-participants to access. Extensions to other contract specifications, such as second generation exotic FX options, will require the derivation of expected stopping time consistent with those specifications, for the approach in this thesis to be implemented.

BIBLIOGRAPHY

Andersen, L., and Andreasen, J. (2000). Jump-diffusion processes: volatility smile fitting and numerical methods for pricing. Review of Derivatives Research, 4, 231–262.

Andersen, L., and Brotherton-Ratcliffe (1998). The equity option volatility smile – an implicit finite-difference approach. Journal of Computational Finance, 1, 2.

Apel, T., Winkler, G., and Wystup, U. (2002). Valuation of options in Heston's stochastic volatility model using finite element methods. Chapter 24 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

Ayache, E., Henrotte, P., Nassar, S., and Wang, X. (2004). Can anyone solve the smile problem? Ito33 working paper.

Ayache, E. (2004a). www.fooledbyrandomness.com/ayache.pdf

Ayache, E. (2004b) www.fooledbyrandomness.com/ayache4.pdf. The back of beyond. Wilmott Magazine, 26-31.

Ayache, E. www.wilmott.com. One step beyond.

Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52, 5, 2003–2049.

Bates, D. (1996). Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. The Review of Financial Studies, 9, 69–107.

Bates, D. (2003). Empirical option pricing: a retrospection. Journal of Econometrics, 116, 1, 2, 387–404.

Belledin, M., and Schlag, C. (1999). An empirical comparison of alternative stochastic volatility models. Working paper, Department of Economics and Business Administration, Goethe-University of Frankfurt.

Blacher, G. (2001). A new approach for designing and calibrating stochastic volatility models for optimal delta-vega hedging of exotics. Global Derivatives Conference presentation at Juan-les-Pins.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81, 637–657.

Black, F. and Scholes, M. (1972). The valuation of option contracts and a test of market efficiency. Journal of Finance, 27, 5, 399-417.

Bowie, J., and Carr, P. (1994). Static simplicity. Risk, 7, 45 - 49.

Breeden, D., and Litzenberger, R. (1978). Prices of state-contingent claims implicit in options prices. Journal of Business, 5, 621–652.

Brigo, D., and Mercurio, F. (2000). A mixed-up smile. Risk, September, 123–126.

Brigo, D., and Mercurio, F. (2002). Lognormal-mixture dynamics and calibration to market volatility smiles. International Journal of Theoretical and Applied Finance, 5, 427–446.

Britten-Jones, M., and Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. Journal of Finance, 55, 2, 838–866.

Carr, P., Ellis, K, and Gupta, V. (1998). Static hedging of exotic options. Journal of Finance, 53, 3, 1165–1190.

Chriss, N., and Ong, M. (1995). Digitals diffused. Risk, 8, 12, 56–59.

Cox, J., Ingersoll, J., and Ross, S. (1985). A theory of the term structure of interest rates. Econometrica, 53, 385–407.

Cox, J., Ross, S., and Rubinstein, M. (1979). Option pricing: a simplified approach. Journal of Financial Economics, 7, 3, 229–263.

Das, S., and Sundarem, R. (1999). Of smiles and smirks: a term structure perspective. Journal of Financial and Quantitative Analysis, 34, 2, 211–239.

Derman, E. (1995). Static options replication. The Journal of Derivatives, 2, 4, 78–95.

Derman, E. (1997). Trading and hedging local volatility. The Journal of Financial Engineering, 6, 3, 1233–1268.

Derman, E. (1998). Stochastic implied trees: arbitrage pricing with stochastic term and strike structure of volatility. Journal of Theoretical and Applied Finance, 1, 1, 61–110.

Derman, E. (1999). Regimes of volatility. Risk, 12, 4, 55-59.

Derman, E. (2000). A guide for the perplexed quant. Journal of Derivatives, 8, 2, 58–68.

Derman, E. (2002). Nine billion ways to be snowed. www.risk.net, June.

Derman, E. (2003). Laughter in the dark: the problem of the volatility smile. www.ederman.com/new/docs/euronext-volatility_smile.pdf.

Derman, E., Ergener, D., and Kani, I. (1994). Forever hedged. Risk, 7, 39-45.

Derman, E., and Kani, I. (1994). Riding on a smile. Risk, February, 32–39.

Derman, E., Kani, I., and Chriss, N. (1996). Implied trinomial trees of the volatility smile. Journal of Derivatives, 3, 4, 7–22.

Derman, E., Kani, I., Ergener, D., and Bardhan, I. (1995). Enhanced numerical methods for options with barriers. Goldman Sachs Quantitative Strategies Research Notes.

Detlefsen, K., and Härdle, W. (2006). Calibration risk for exotic options. SFB 649 Discussion Paper, 2006-001.

Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. Econometrica, 68, 6, 1343-1376.

Dupire, B. (1994). Pricing with a smile. Risk, January, 18–20.

Dupire, B. (1997). A unified theory of volatility. Banque Paribas Working Paper.

Ebenfeld, S., Mayr, M., and Topper, J. (2002). An analysis of onion options and double no touch digitals. Wilmott Magazine, September, 68–77.

Famery, F. and Cornu, S. (2000). Exotic options: from theoretical value to dealing price. Euroderivatives Financial Products Guide, Risk Supplement, September 2000.

Fouque, J., Papanicolaou, G., and Sircar, R. (2000). Calibrating random volatility. Risk, February, 89–92.

Garman, M., and Kohlhagen, S. (1983). Foreign currency option values. Journal of International Money and Finance, 2, 231–237.

Grabbe, O. (1983). The pricing of call and put options on foreign exchange. Journal of International Money and Finance, 2, 239–253.

Hakala, J., and Kirch, M. (2002). Local volatility surfaces – tackling the smile. Chapter 22 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

Hakala, J., and Wystup, U. (2002). Heston's stochastic volatility model applied to foreign exchange options. Chapter 23 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

Hakala, J., Perissé, G. and Wystup, U. (2002). The pricing of first generation exotics. Chapter 6 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

Haug, E. (2003a). Know your weapon: part 1. Wilmott Magazine, May, 49–57.

Haug, E. (2003b). Know your weapon: part 2. Wilmott Magazine, July, 50–56.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies, 6, 2, 327–343.

Hui, C. (1996). One touch double barrier binary option values. Applied Financial Economics, 6, 343–346.

Hull, J., and Suo, W. (2002). A methodology for assessing model risk and its application to the implied volatility function model. Journal of Financial and Quantitative Analysis, 37, 2, 297–318.

Hull, J., and White, A. (1987). The pricing of options on assets with stochastic volatility. Journal of Finance, 42, 281–300.

Jex, M., Henderson, R., and Wang, D. (1999). Pricing exotics under the smile. Risk, November, 72–75.

Johnson, S., and Lee, H. (2003). Capturing the smile. Risk, March, 89–93

Jondeau, E., and Rockinger, M. (2000). Reading the smile: the message conveyed by methods which infer risk neutral densities. Journal of International Money and Finance, 19, 6, 885–915.

Kunitomo, N. and Ikeda, M. (1992). Pricing options with curved boundaries. Mathematical Finance, 2, 275–298.

Lipton, A. (2002). The vol smile problem. Risk, February, 61-65.

Lipton, A. (2001). Mathematical methods for foreign exchange: a financial engineer's approach. World Scientific, Singapore.

Lipton, A., and McGhee, W. (2002). Universal barriers. Risk, May, 81–85.

Martinez, M., and Senge, T. (2002). A jump-diffusion model applied to foreign exchange markets. Chapter 25 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

Matytsin, A. (1999). Modelling volatility and volatility derivatives. www.math.columbia.edu/~ smirnov/Matytsin.pdf.

Melino, A., and Turnbull, S. (1990). Pricing foreign currency options with stochastic volatility. Journal of Econometrics, 45, 239–265.

Melino, A., and Turnbull, S. (1995). Misspecification and the pricing and hedging of long-term foreign currency options. Journal of International Money and Finance, 14, 3 373–393.

Merton, R. (1973). Theory of rational option pricing. Bell Journal of Economics and Management Science, 4, 141–183.

Merton, R. (1976). Option pricing when the underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125-144.

Mikhailov, S., and Nögel, U. (2005). Heston's stochastic volatility model implementation, calibration and some extensions. Wilmott Magazine, 74-79.

Papanicolaou, G., and Sircar, K. (1999). Stochastic volatility, smile and asymptotics. Applied Mathematical Finance, 6, 107–145.

Rebonato, R. (2002). Theory and practice of model risk management. www.quarchome.org.

Rebonato, R. (1999). Volatility and correlation: in the pricing of equity, FX and interest rate options. John Wiley and Sons, New York.

Rubinstein, M. (1994). Implied binomial trees. Journal of Finance, 69, 3, 771–818.

Rubinstein, M. (1998). Edgeworth binomial trees. Journal of Derivatives, Spring, 20–27.

Rubinstein, M. and Reiner, E. (1991a). Unscrambling the binary code. Risk, 4, 10, 75–83.

Rubinstein, M. and Reiner, E. (1991b). Breaking down the barriers. Risk, 4, 9, 28–35.

Savery, H. (2000). Quantifying volatility convexity. DerivativesStrategy.com, February 2000.

Schoutens, W., Simons, E. and Tistaert, J. (2004). A Perfect Calibration! Now What? Wilmott Magazine, March, 66-78.

Shevchenko, P. (2003). Exact formulae for some barrier options written on one or two underlying assets. CSIRO, unpublished research paper.

Shimko, D. (1992). Finance in continuous time: a primer. Kolb Publishing Company. Florida.

Smith, K. (2006). No-arbitrage volatility surfaces. Commercial-in-Confidence Working Paper.

Smith, K. (2008). Expected stopping time for continuously monitored exotic FX options with two barriers. Commercial-in-Confidence Working paper.

Spiegel, M. (2001). Quantitative easing by the Bank of Japan, FRBSF Economic Letter, 31.

Spiegel, M. (2006). Did quantitative easing by the Bank of Japan "work"? FRBSF Economic Letter, 28.

Stein, E., and Stein, J. (1991). Stock price distributions with stochastic volatility: an analytic approach. The Review of Financial Studies, 4, 727-752.

Taleb, N. (1997). Dynamic hedging. Wiley, New York.

Technology Rankings 2005. A Risk Magazine Special Report, 18, 12, 32-38.

Tompkins, R. (1997). Static versus dynamic hedging of exotic options: an evaluation of hedge performance via simulation. Appeared on CD supplement to Wilmott. Journal of Risk Finance, Summer/2002 (2002), S. 6-34.

Topper, J. (2002). Taking a corporate view: zero-cost structures. Chapter 10 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

UBS. Effects of volatility smile on exotic option pricing. Presentation document.

Warsh, K. (2008). The Federal funds rate in extraordinary times. Speech, Board of Governors of the Federal Reserve System, http://www.federalreserve.gov/newsevents/speech/warsh20080521a.htm.

Wilmott, P. (1998). Derivatives – the theory and practice of financial engineering. Chichester, John Wiley and Sons.

Wystup, U. (2006). FX options and structured products. Chichester, John Wiley and Sons.

Wystup, U. (2003). The market price of one touch options in foreign exchange markets. *Derivatives Week*, 12 (13), p. 1-4.

Wystup, U. (2002). Volatility management. Chapter 2 of "Foreign exchange risk: models, instruments and strategies". Edited by Hakala, J. and Wystup, U. Risk Books, London.

APPENDIX A

A.1 Trader Model overview

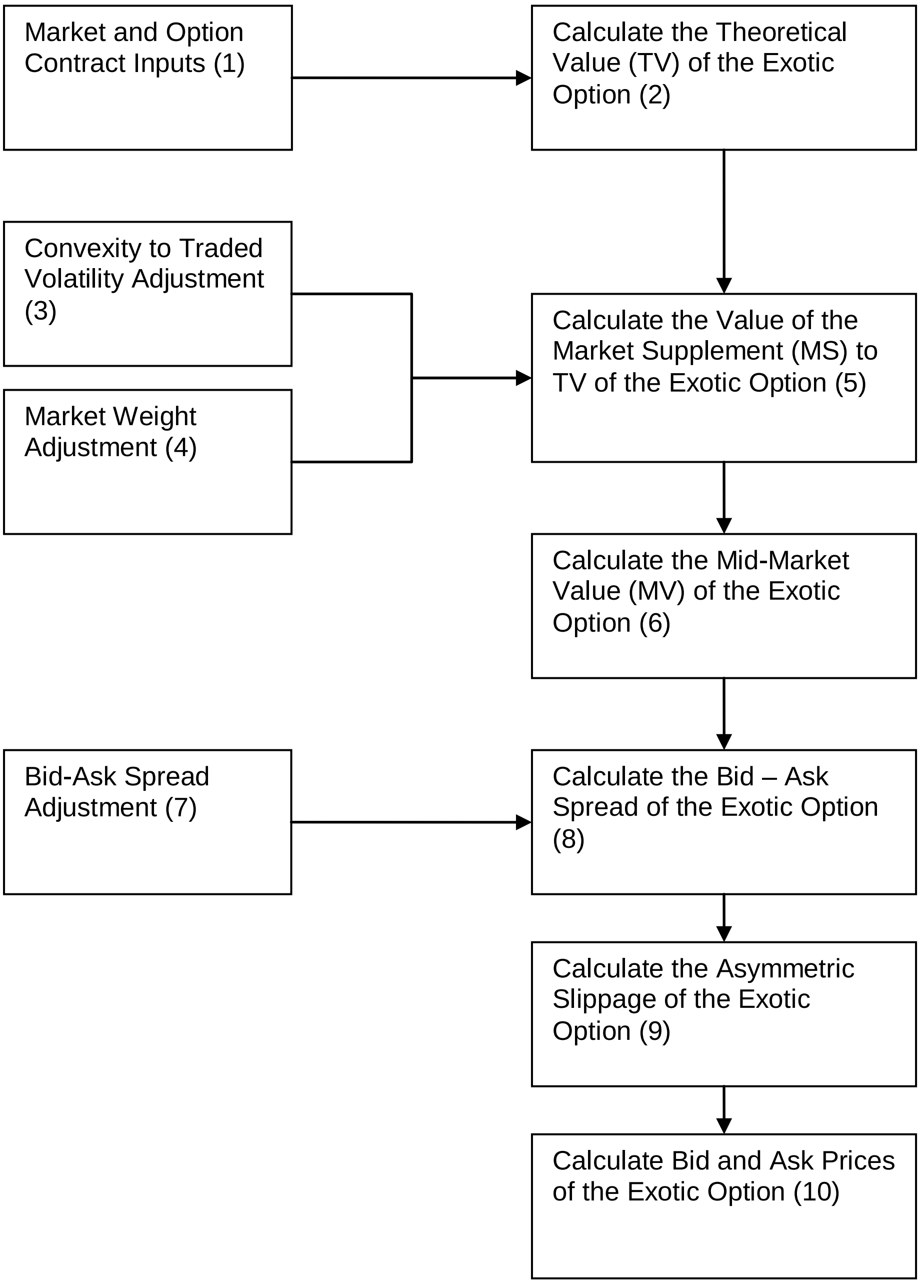

A.2 Trader Model vega convexity to traded volatility adjustment

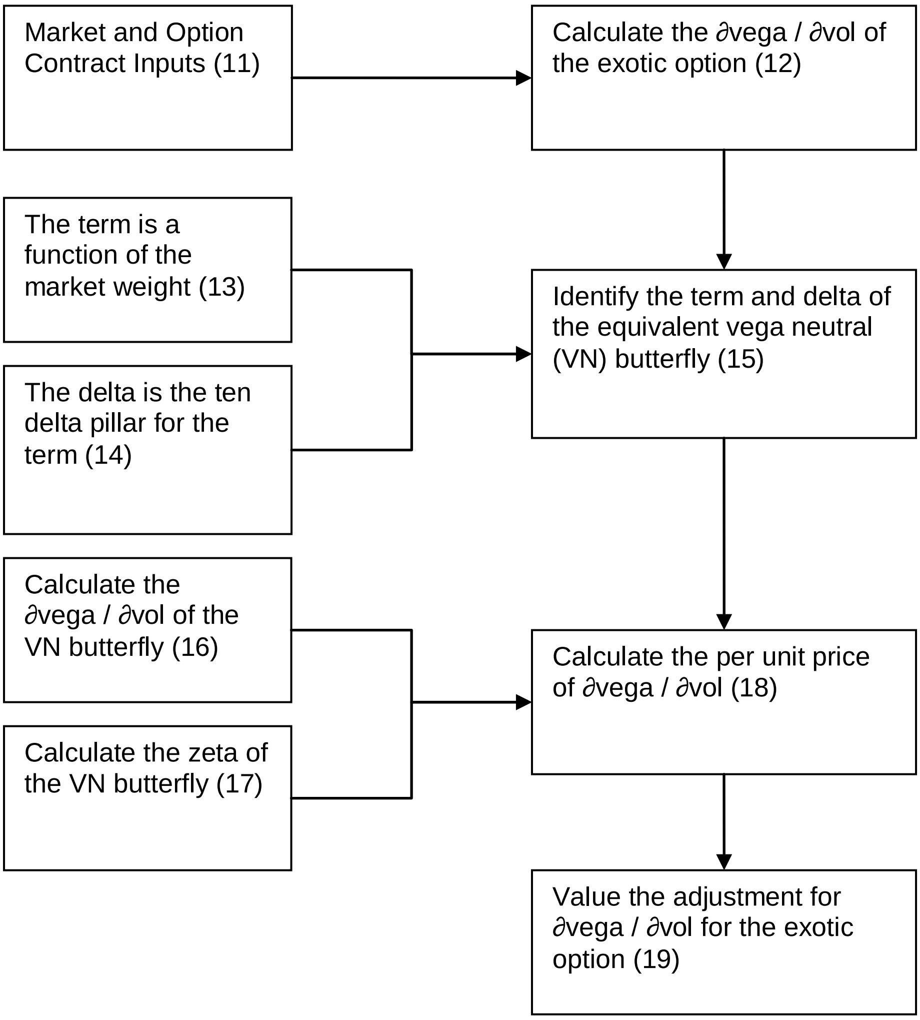

A.3 Trader Model delta convexity to traded volatility adjustment

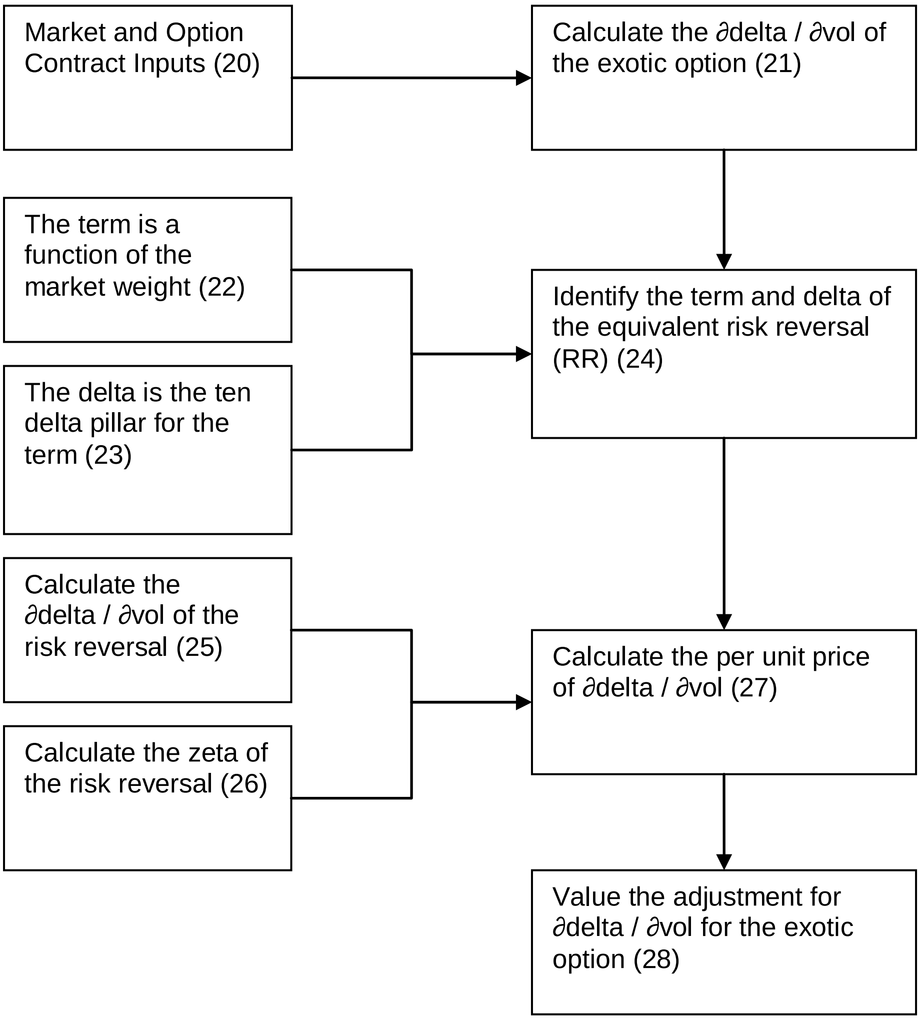

A.4 Trader Model bid-ask spread calculation overview

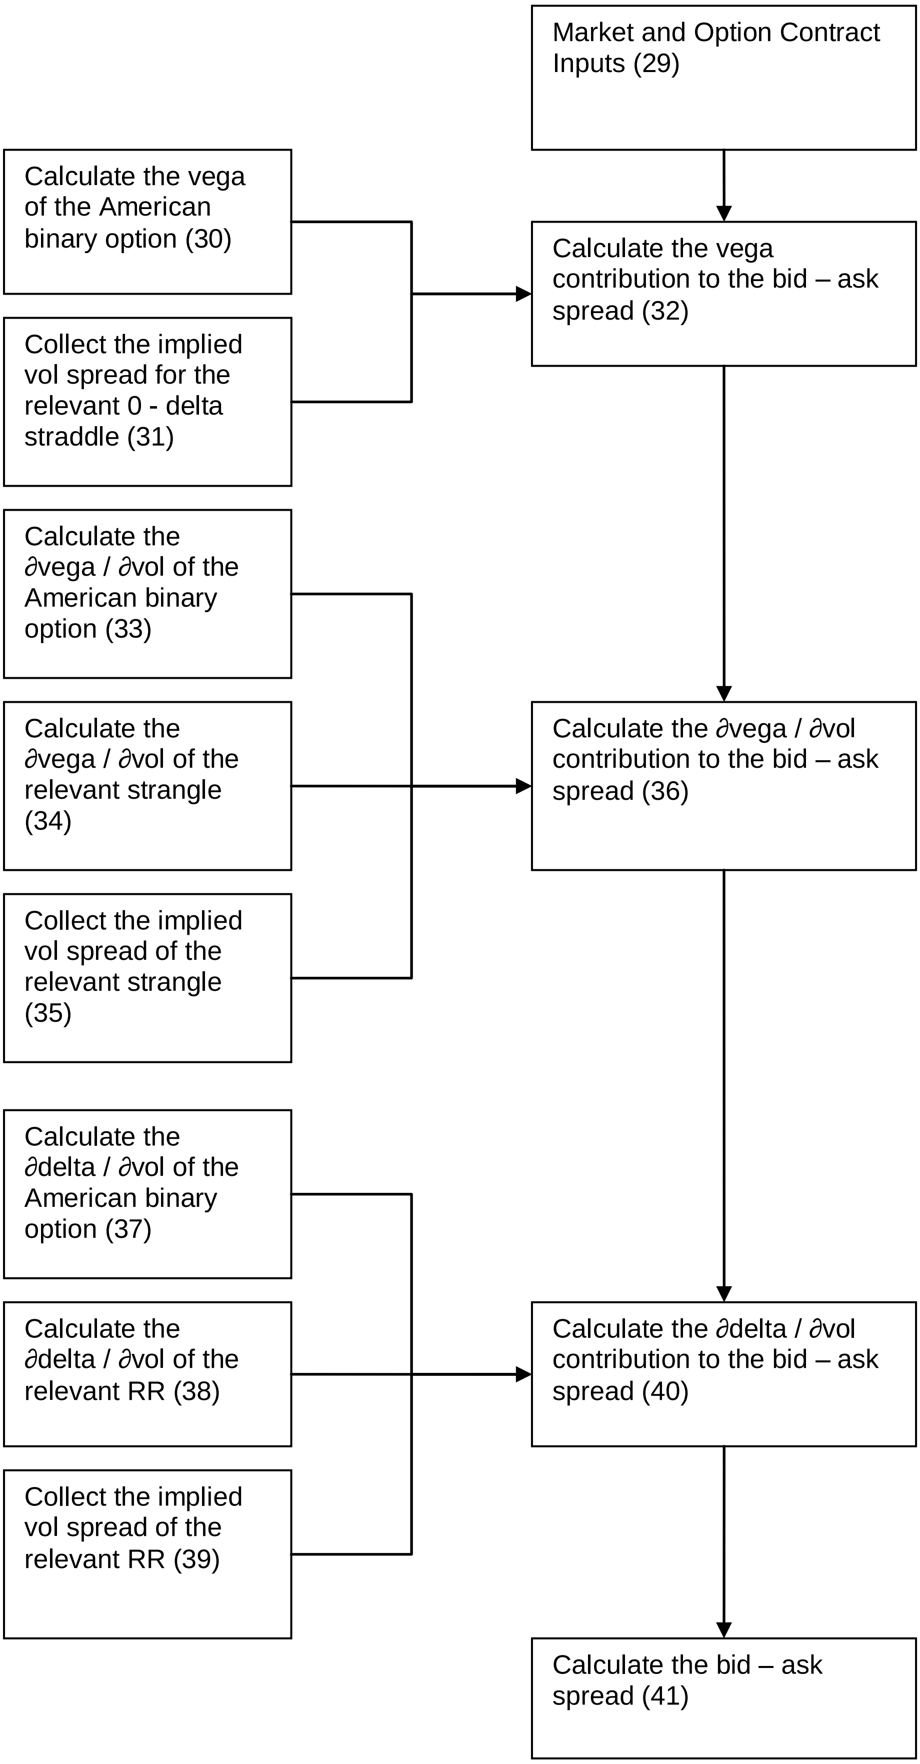

APPENDIX B

Descriptive statistics for deposit rates

Table B.1: EUR deposit rate input data description

EUR	Max	Min	Diff
1w	2.08	1.96	0.12
1m	2.11	1.945	0.165
2m	2.11	1.895	0.215
3m	2.175	1.875	0.3
6m	2.285	1.805	0.48
12m	2.475	1.825	0.65
12m-1w	0.44	-0.245	0.685
12m-1m	0.41	-0.12	0.53

Minimum and maximum EUR-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with slope of 44 basis points, to an inverse curve with slope of -24.5 basis points. Source: Reuters LLC.

Table B.2: USD deposit rate input data description

USD	Max	Min	Diff
1w	1.8	0.99	0.81
1m	1.81	1.02	0.79
2m	1.87	1.03	0.84
3m	2.0	1.04	0.96
6m	2.14	1.09	1.05
12m	2.435	1.21	1.225
12m-1w	1.37	0.19	1.18
12m-1m	1.2	0.17	1.03

Minimum and maximum USD-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with a slope of 19 to 137 basis points. Source: Reuters LLC.

Table B.3: JPY deposit rate input data description

JPY	Max	Min	Diff
1w	-0.005	-0.125	0.12
1m	0.03	-0.105	0.135
2m	0.05	-0.09	0.14
3m	0	-0.09	0.09
6m	0.05	-0.07	0.12
12m	0.06	-0.02	0.08
12m-1w	0.155	0.005	0.15
12m-1m	0.12	-0.095	0.215

Minimum and maximum JPY-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with a slope of 0.5 to 15.5 basis points (bp). The 12m to 1w term structure exhibited normal (12bp) and inverse (-9.5bp) deposit rate curves. Source: Reuters LLC.

Table B.4: AUD deposit rate input data description

AUD	Max	Min	Diff
1w	5.455	5.115	0.34
1m	5.455	5.235	0.22
2m	5.51	5.275	0.235
3m	5.53	5.3	0.23
6m	5.62	5.345	0.275
12m	5.805	5.4	0.405
12m-1w	0.62	0.115	0.505
12m-1m	0.505	0.05	0.455

Minimum and maximum AUD-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with a slope of 11.5 to 62 basis points. Source: Reuters LLC.

Table B.5: GBP deposit rate input data description

GBP	Max	Min	Diff
1w	5.01	3.55	1.46
1m	4.825	3.77	1.055
2m	4.87	3.885	0.985
3m	4.97	3.94	1.03
6m	5.15	4.09	1.06
12m	5.365	4.215	1.15
12m-1w	1.015	-0.01	1.025
12m-1m	0.98	0.28	0.7

Minimum and maximum GBP-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with slope of 101.5 basis points, to an inverse curve with slope of -1 basis point. Source: Reuters LLC.

Table B.6: CHF deposit rate input data description

CHF	Max	Min	Diff
1w	0.52	0.11	0.41
1m	0.555	0.12	0.435
2m	0.585	0.14	0.445
3m	0.645	0.16	0.485
6m	0.8	0.205	0.595
12m	1.085	0.355	0.73
12m-1w	0.925	0.16	0.765
12m-1m	0.8	0.17	0.63

Minimum and maximum CHF-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with a slope of 16 to 92.5 basis points. Source: Reuters LLC.

Table B.7: CAD deposit rate input data description

CAD	Max	Min	Diff
1w	2.77	1.1975	0.795
1m	2.68	1.975	0.705
2m	2.61	1.975	0.635
3m	2.57	1.975	0.595
6m	2.66	2.005	0.655
12m	2.97	2.04	0.93
12m-1w	0.83	-0.27	1.1
12m-1m	0.79	-0.24	1.03

Minimum and maximum CAD-denominated deposit rates for the period 7 January 2004 to 29 September 2004, inclusive. The last two rows show the slope of the term structure of deposit rates. For example, the term structure between twelve month (12m) and one week (1w) deposit rates ranged from a normal curve with slope of 83 basis points, to an inverse curve with slope of -27 basis poinst. Source: Reuters LLC.

APPENDIX C

Descriptive statistics for European vanilla traded volatility surfaces

Table C.1: Level, smile and skew of the volatility surface data

	Level		Smile		Skew	
FX Pair	**Max**	**Min**	**Max**	**Min**	**Max**	**Min**
EUR/USD	19.50	6.00	0.95	0.50	1.70	-1.75
USD/JPY	18.00	5.00	2.30	0.75	1.75	-6.30
EUR/JPY	14.50	8.80	1.525	0.80	0.40	-2.90
GBP/USD	17.50	9.00	0.725	0.50	2.50	-1.60
AUD/USD	19.50	9.90	0.90	0.725	0.48	-3.20
EUR/CHF	5.90	3.40	0.77	0.385	0.45	-1.055
EUR/GBP	7.55	6.50	0.625	0.55	-0.05	-0.55
USD/CAD	10.60	8.50	0.75	0.50	1.00	-0.60

Level is the size of ATM volatility in percent per annum. *Smile* is the volatility premium for a 10 delta vega neutral butterfly in percent per annum. *Skew* is the volatility premium for a 10 delta delta neutral risk reversal in percent per annum. 10 delta refers to OTM European vanilla options with spot deltas of 0.1 for the Call and -0.1 for the Put.

Table C.2: Level, smile and skew of the volatility surface for DNT options

DNT	Level		Smile		Skew	
FX Pair	**Max**	**Min**	**Max**	**Min**	**Max**	**Min**
EUR/USD	18.50	6.00	0.95	0.50	1.40	-1.75
USD/JPY	15.50	5.50	2.30	0.75	1.75	-6.00
EUR/JPY	14.50	8.80	1.50	0.80	0.40	-2.90
GBP/USD	17.50	9.15	0.725	0.525	1.30	-1.60
AUD/USD	14.50	10.25	0.90	0.75	0.48	-2.20
EUR/CHF	4.45	3.60	0.77	0.385	0.45	0.25
EUR/GBP	7.55	6.50	0.625	0.55	-0.05	-0.55
USD/CAD	10.60	8.50	0.75	0.50	1.00	-0.60

Level is the size of ATM volatility in percent per annum. *Smile* is the volatility premium for a 10 delta vega neutral butterfly in percent per annum. *Skew* is the volatility premium for a 10 delta delta neutral risk reversal in percent per annum. 10 delta refers to OTM European vanilla options with spot deltas of 0.1 for the Call and -0.1 for the Put. Volatility surface data is truncated at the three year maturity to match the length of the longest exotic option deal.

Table C.3: Level, smile and skew of the volatility surface for OT options

OT	Level		Smile		Skew	
FX Pair	**Max**	**Min**	**Max**	**Min**	**Max**	**Min**
EUR/USD	19.50	6.00	0.95	0.50	1.70	-1.60
USD/JPY	18.00	5.00	2.30	0.90	1.75	-6.30
EUR/JPY	10.75	9.30	1.525	0.80	0.00	-2.20
GBP/USD	16.50	9.00	0.725	0.50	2.50	-0.80
AUD/USD	19.50	9.90	0.90	0.725	0.175	-3.20
EUR/CHF	5.90	3.40	0.77	0.385	0.45	-1.055
EUR/GBP	N/A	N/A	N/A	N/A	N/A	N/A
USD/CAD	10.40	8.55	0.70	0.50	0.50	0.00

Level is the size of ATM volatility in percent per annum. *Smile* is the volatility premium for a 10 delta vega neutral butterfly in percent per annum. *Skew* is the volatility premium for a 10 delta delta neutral risk reversal in percent per annum. 10 delta refers to OTM European vanilla options with spot deltas of 0.1 for the Call and -0.1 for the Put. *N/A* is 'Not Applicable', as there were no OT option trades in the database for this FX pair.

APPENDIX D

Trader Model pricing results for OT options

EUR

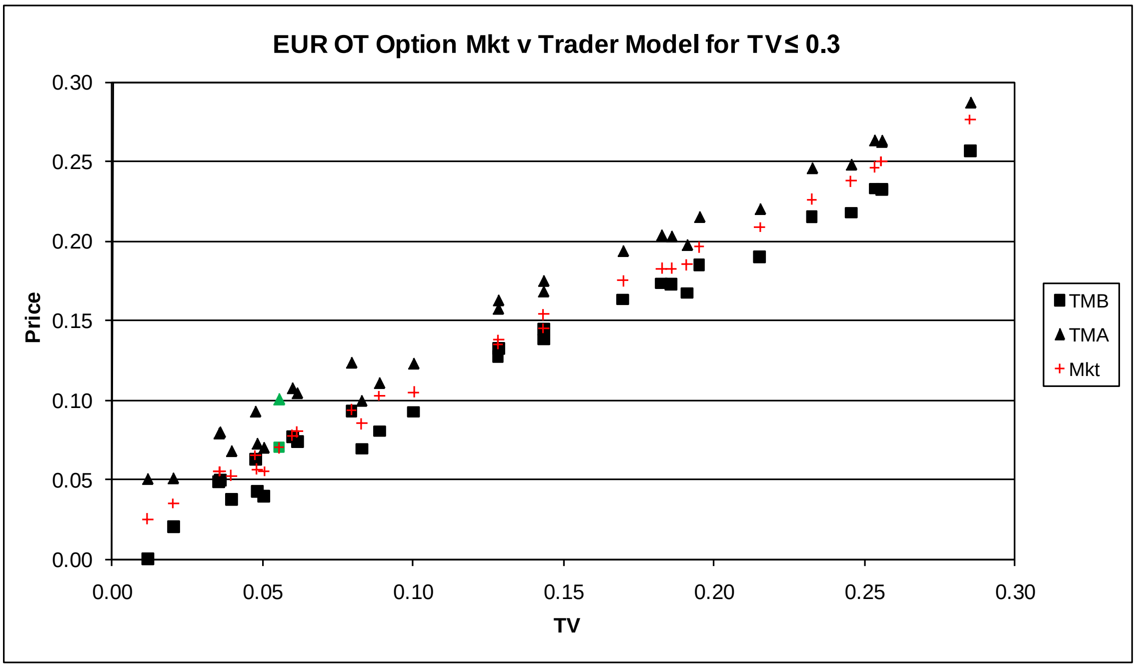

Fig. D1. EUR OT option traded market prices versus Trader Model bid-ask prices, for theoretical values less than or equal to 0.3. Prices in green show EUR OT options where market prices traded outside of the Trader Model bid-ask spread.

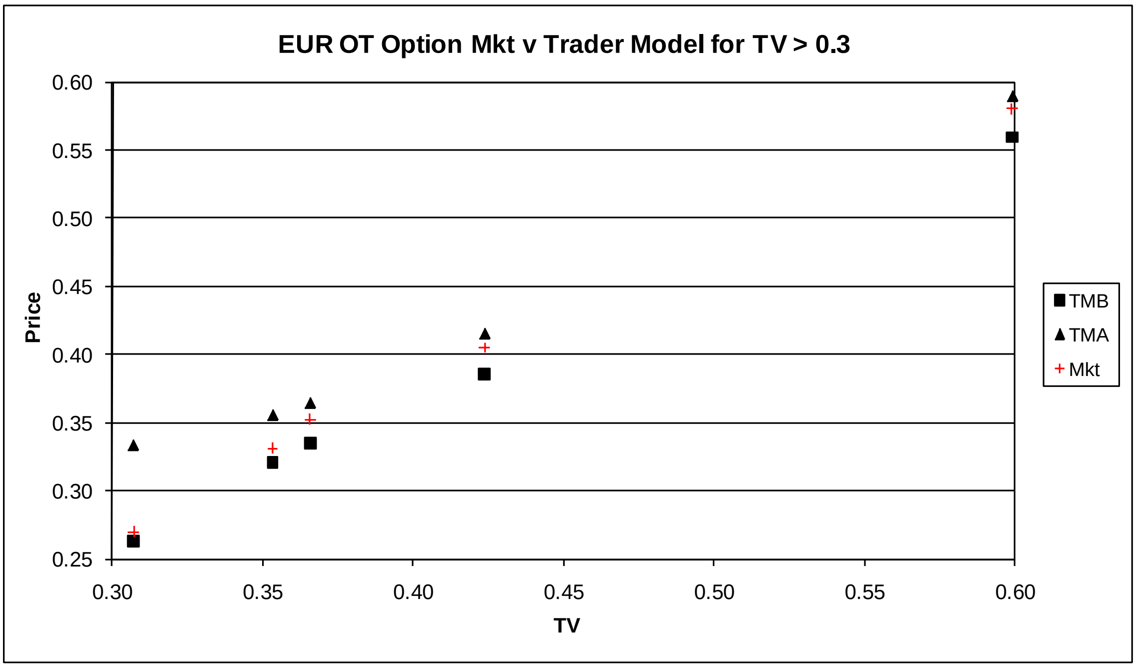

Fig. D2. EUR OT option traded market prices versus Trader Model bid-ask prices, for theoretical values greater than 0.3.

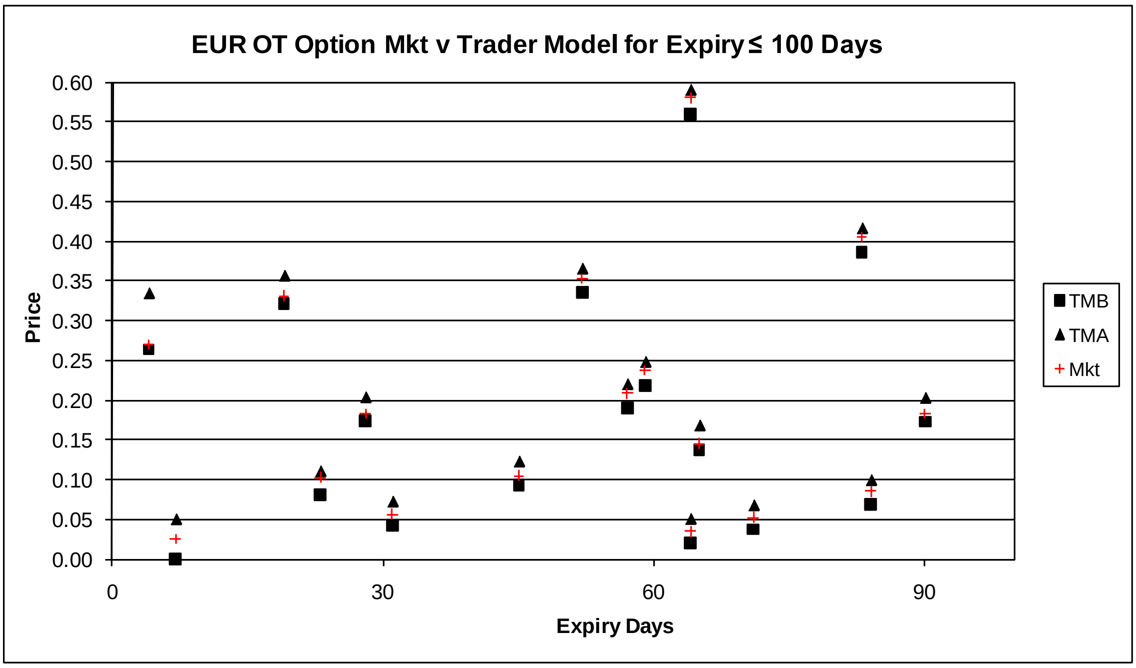

Fig. D3. EUR OT option traded market prices versus Trader Model bid-ask prices, for expiries less than or equal to 100 days.

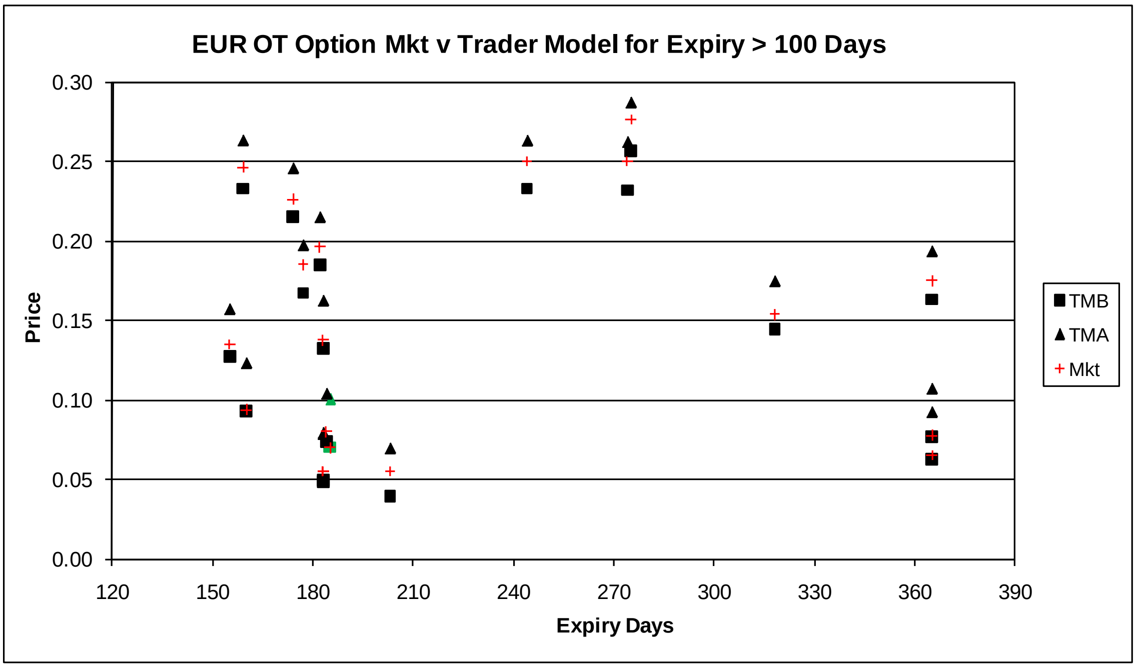

Fig. D4. EUR OT option traded market prices versus Trader Model bid-ask prices, for expiries greater than 100 days. Prices in green show EUR OT options where market prices traded outside of the Trader Model bid-ask spread.

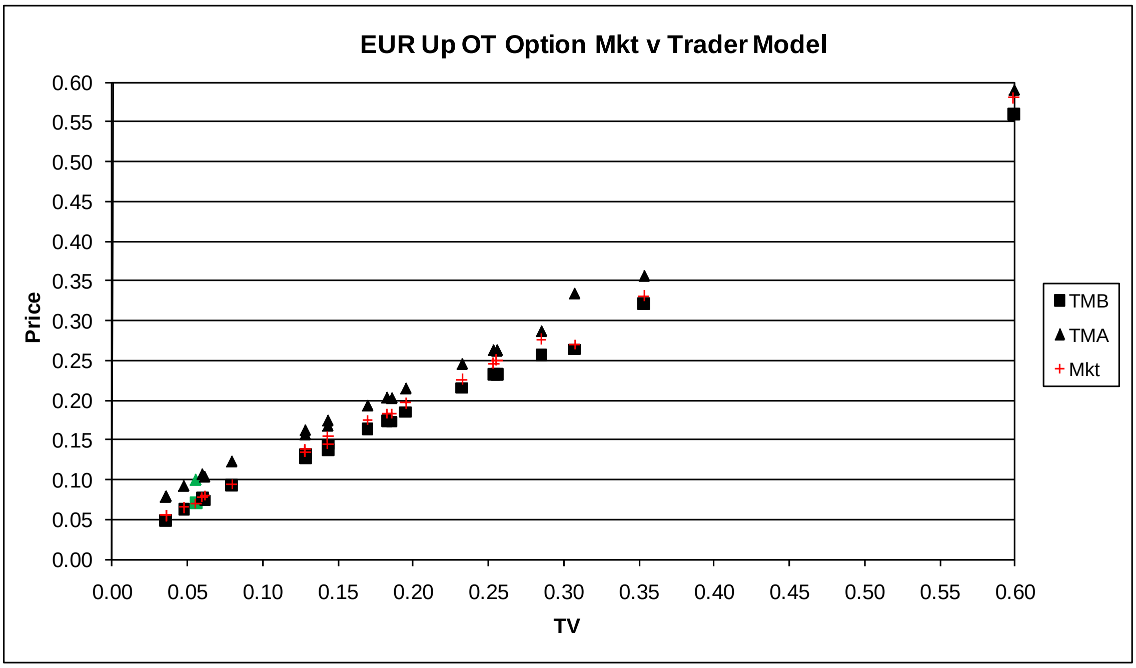

Fig. D5. EUR Up OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. Prices in green show EUR OT options where market prices traded outside of the Trader Model bid-ask spread.

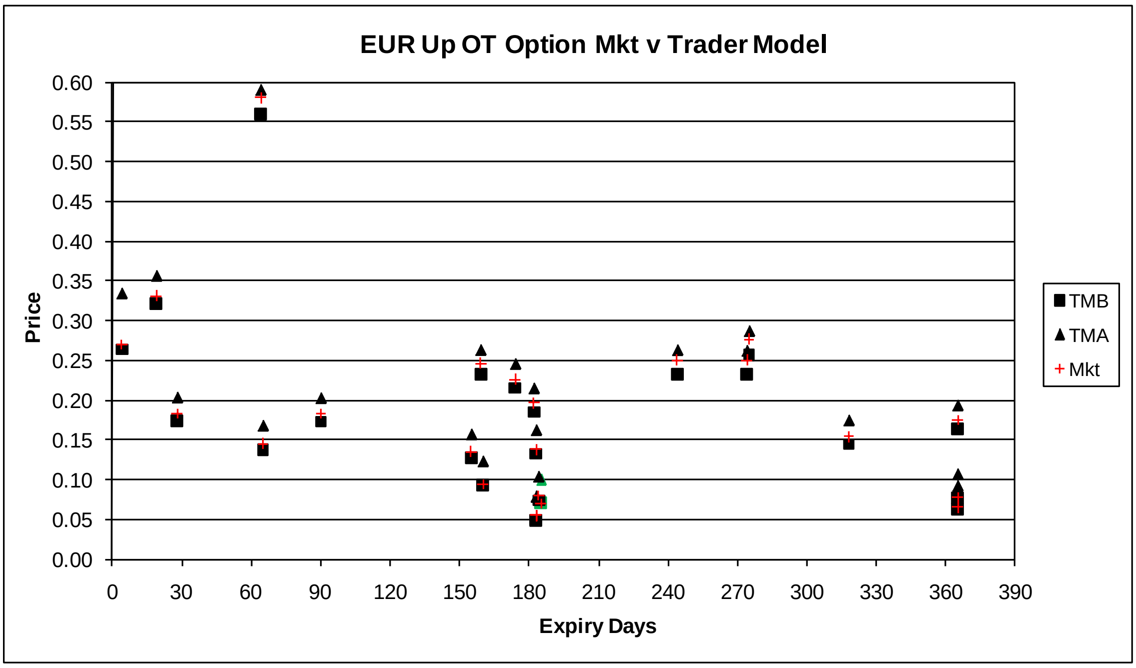

Fig. D6. EUR Up OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. Prices in green show EUR OT options where market prices traded outside of the Trader Model bid-ask spread.

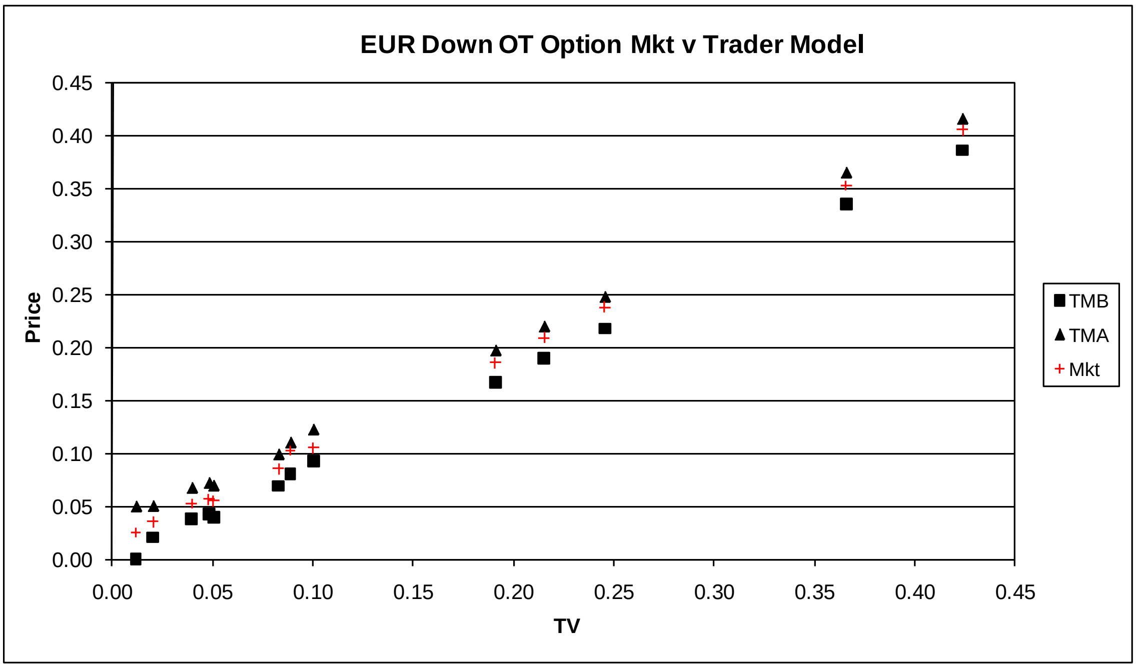

Fig. D7. EUR Down OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

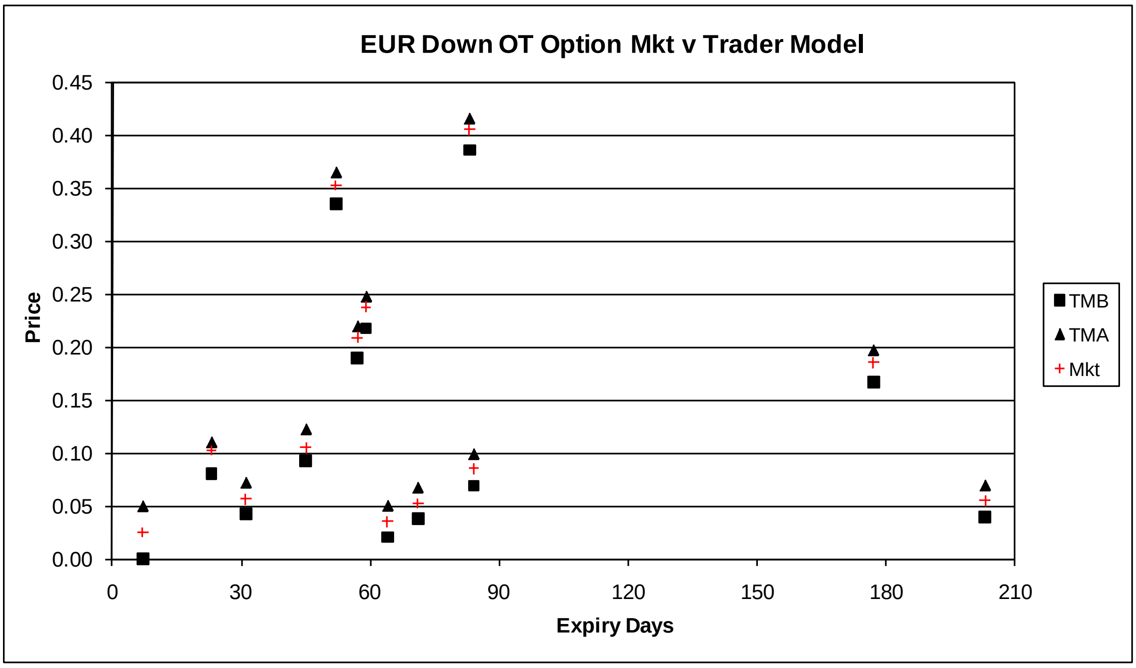

Fig. D8. EUR Down OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

JPY

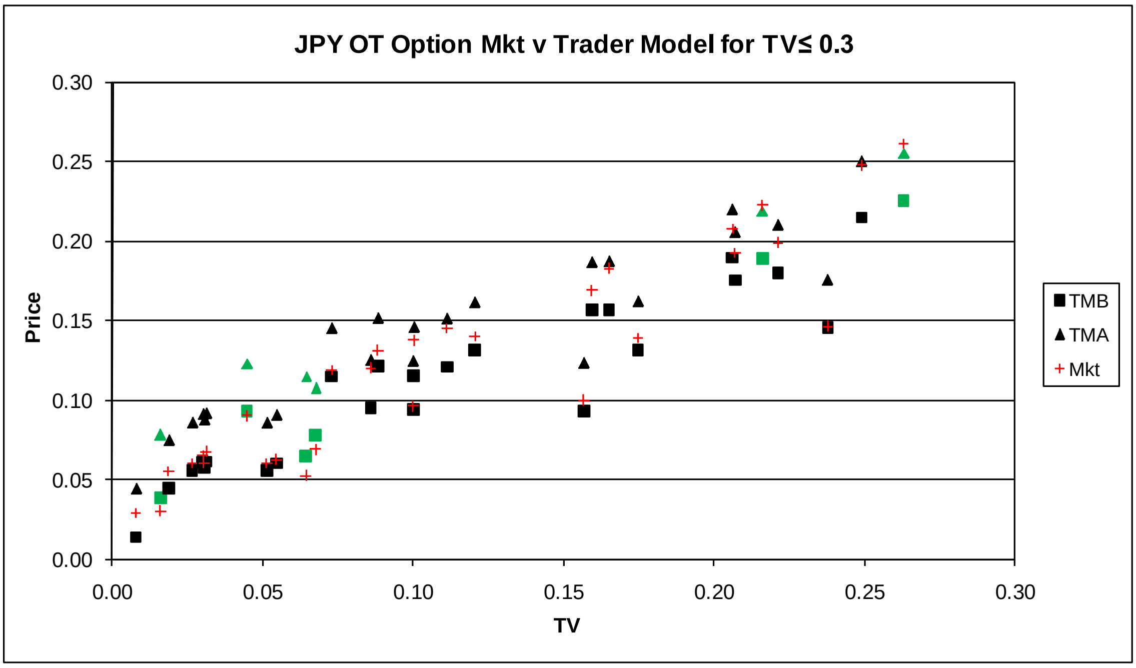

Fig. D9. JPY OT option traded market prices versus Trader Model bid-ask prices, for theoretical values less than or equal to 0.3. Prices in green show JPY OT options where market prices traded outside of the Trader Model bid-ask spread.

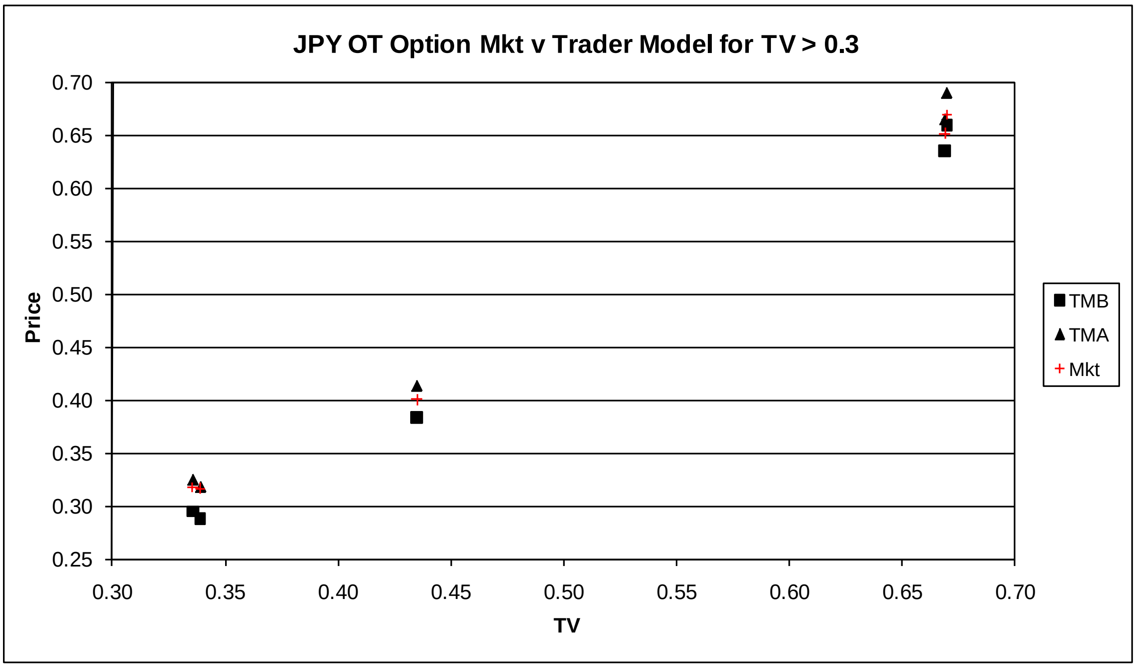

Fig. D10. JPY OT option traded market prices versus Trader Model bid-ask prices, for theoretical values greater than 0.3.

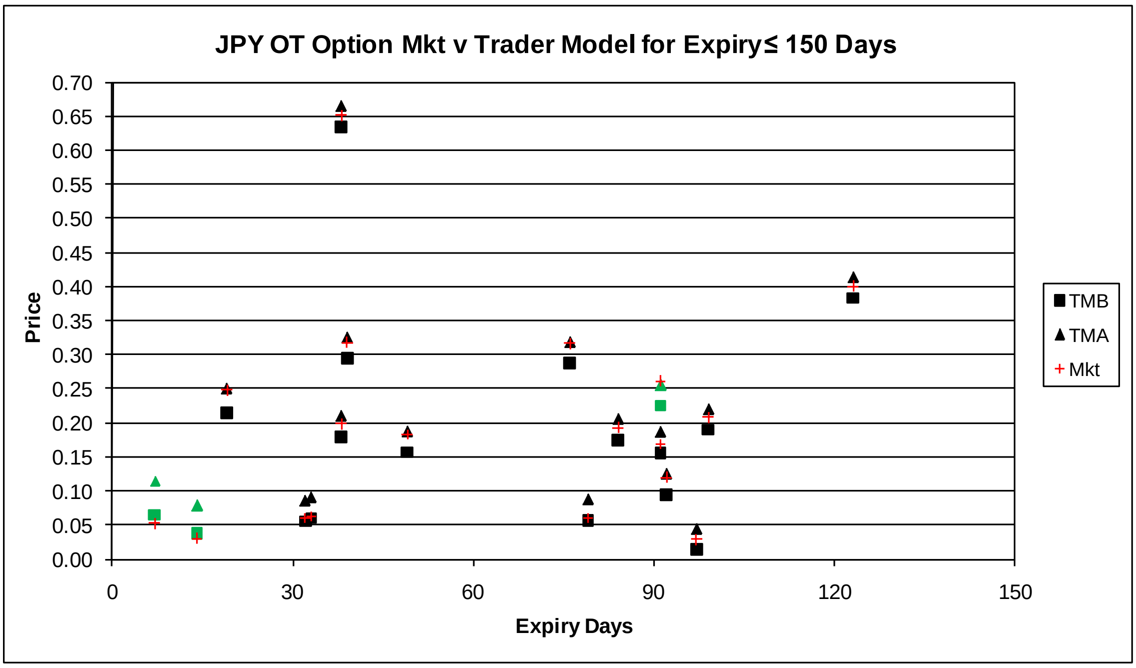

Fig. D11. JPY OT option traded market prices versus Trader Model bid-ask prices, for expiries less than or equal to 150 days. Prices in green show JPY OT options where market prices traded outside of the Trader Model bid-ask spread.

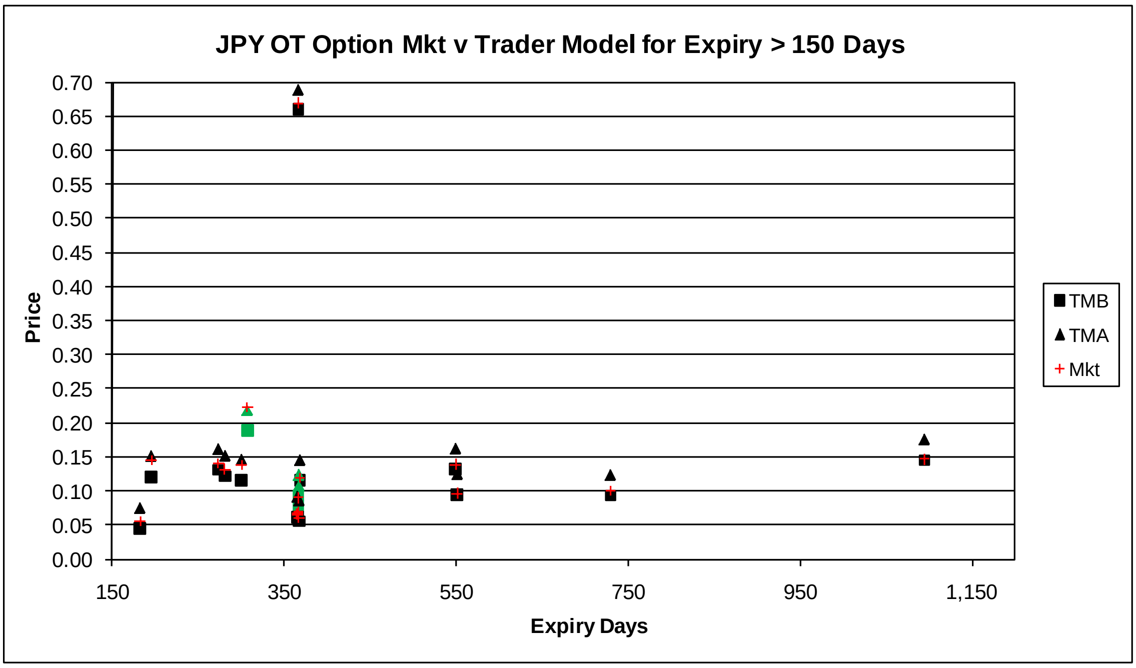

Fig. D12. JPY OT option traded market prices versus Trader Model bid-ask prices, for expiries greater than 150 days. Prices in green show JPY OT options where market prices traded outside of the Trader Model bid-ask spread.

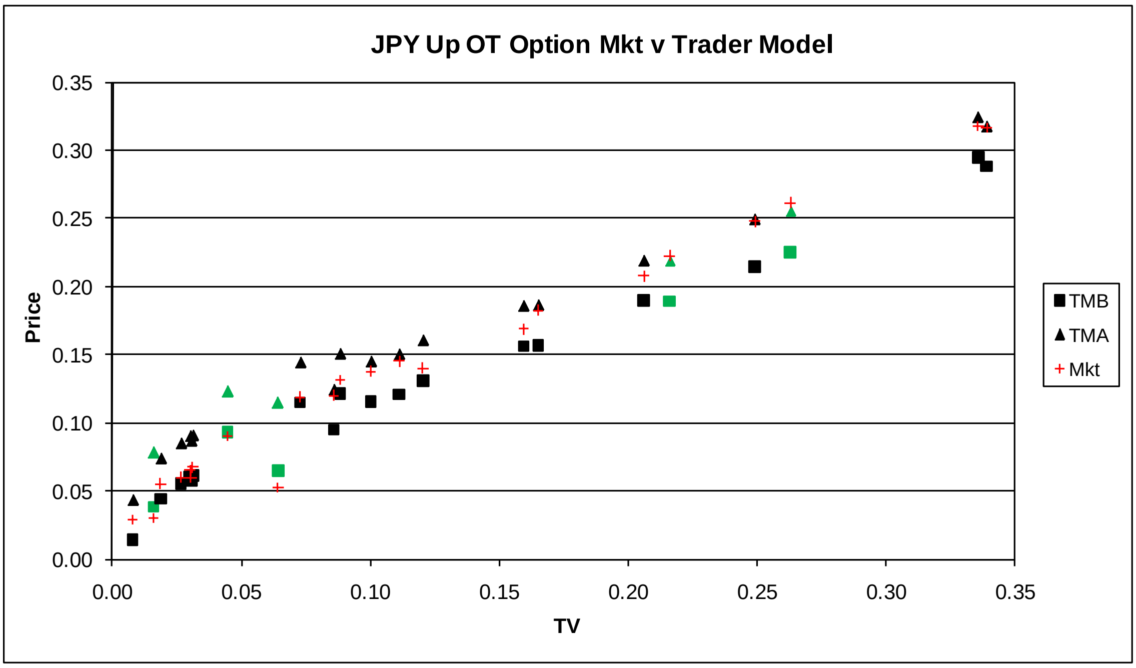

Fig. D13. JPY Up OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. Prices in green show JPY Up OT options where market prices traded outside of the Trader Model bid-ask spread.

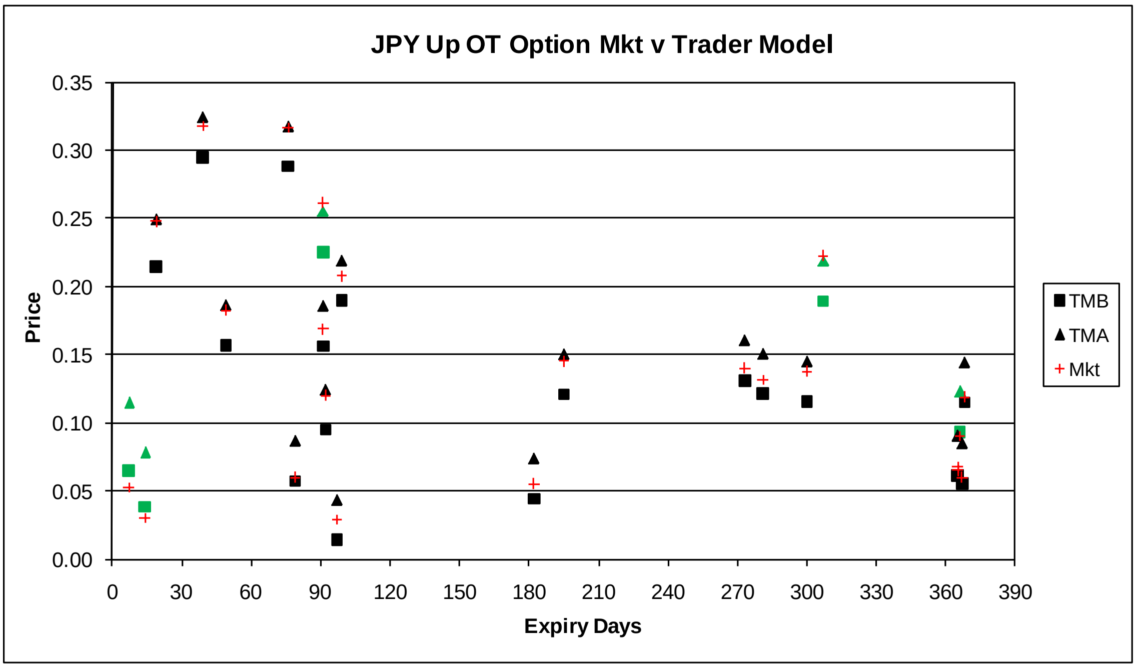

Fig. D14. JPY Up OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. Prices in green show JPY Up OT options where market prices traded outside of the Trader Model bid-ask spread.

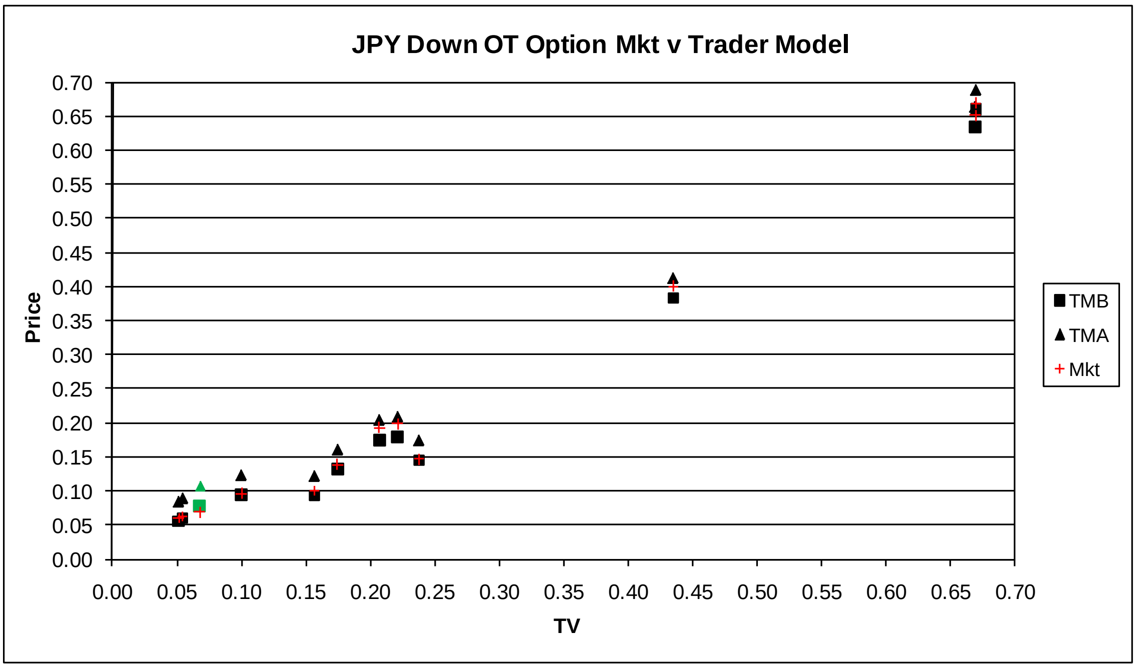

Fig. D15. JPY Down OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. Prices in green show JPY Down OT options where market prices traded outside of the Trader Model bid-ask spread.

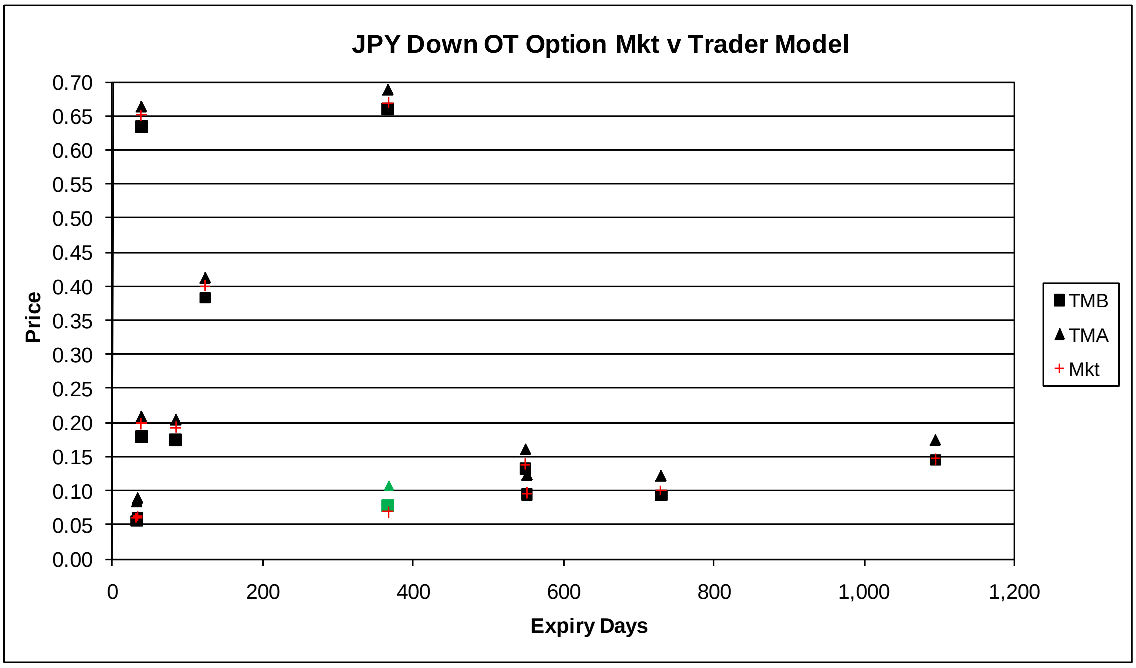

Fig. D16. JPY Down OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. Prices in green show JPY Down OT options where market prices traded outside of the Trader Model bid-ask spread.

EUR/JPY

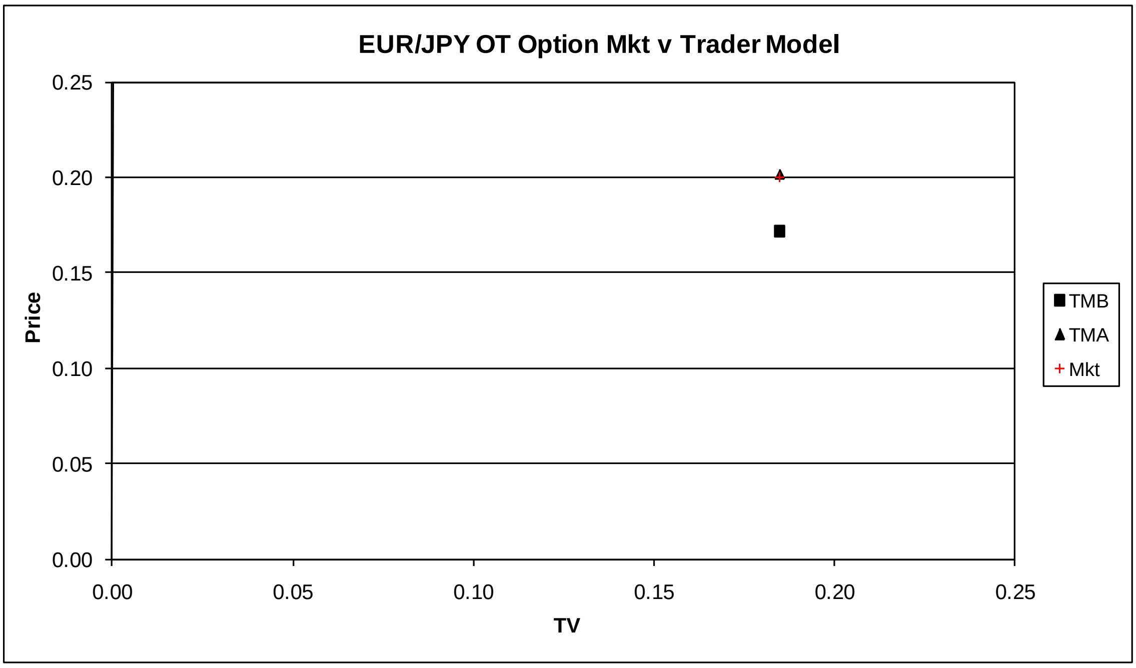

Fig. D17. EUR/JPY OT option traded market price versus Trader Model bid-ask prices, as a function of theoretical value. The option is a EUR down OT option.

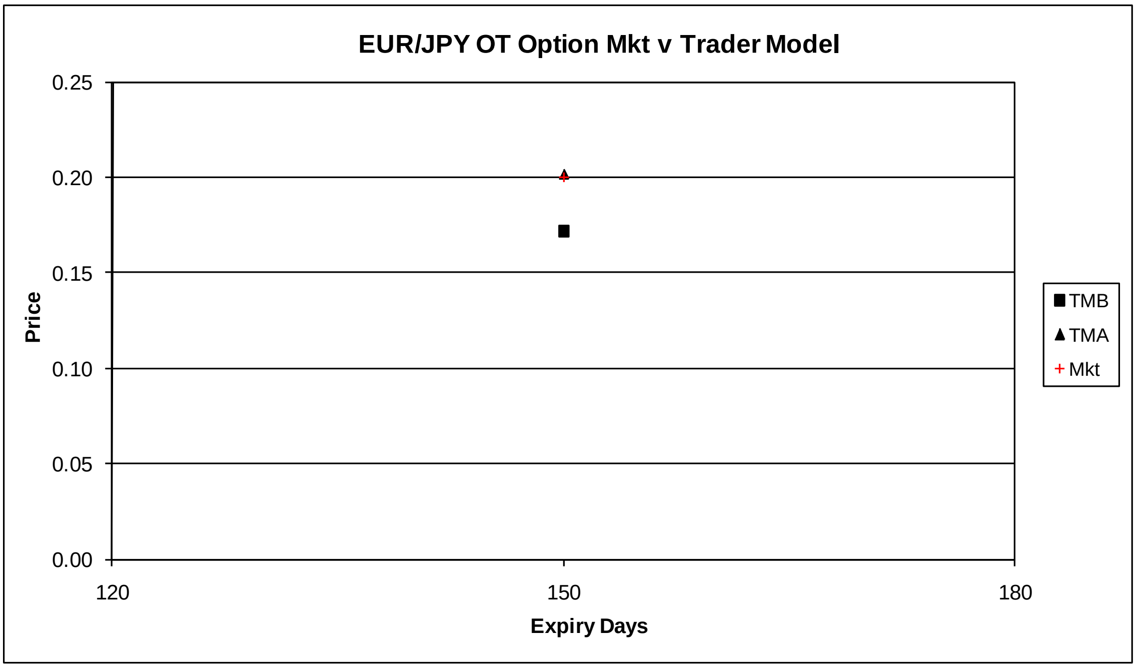

Fig. D18. EUR/JPY OT option traded market price versus Trader Model bid-ask prices, as a function of expiry days. The option is a EUR down OT option.

GBP

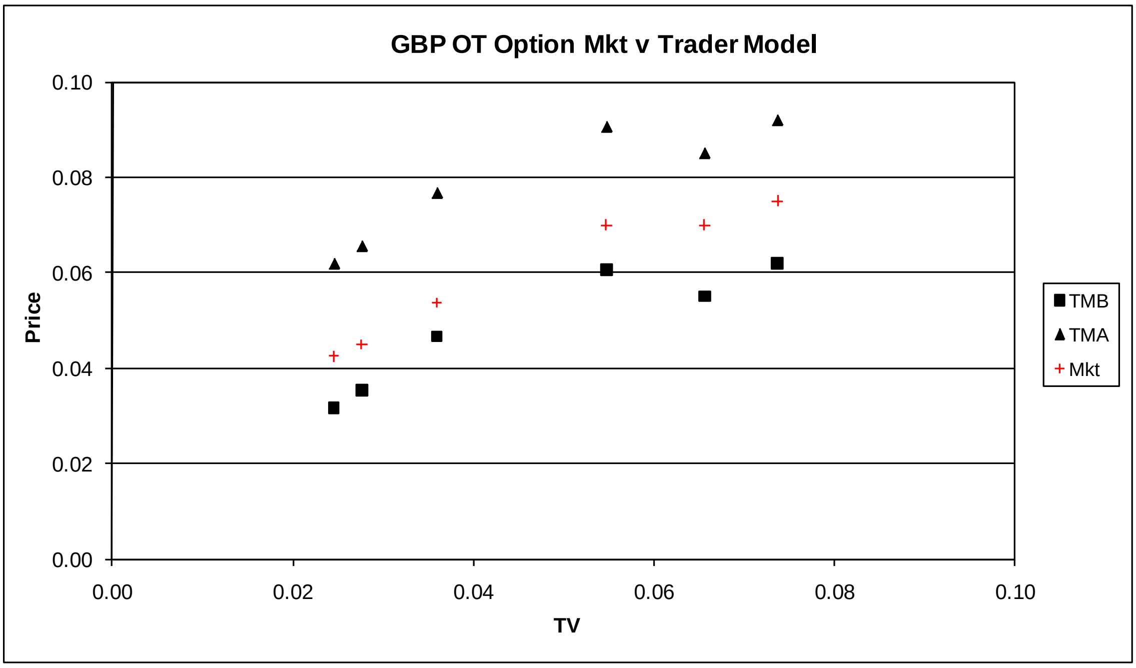

Fig. D19. GBP OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

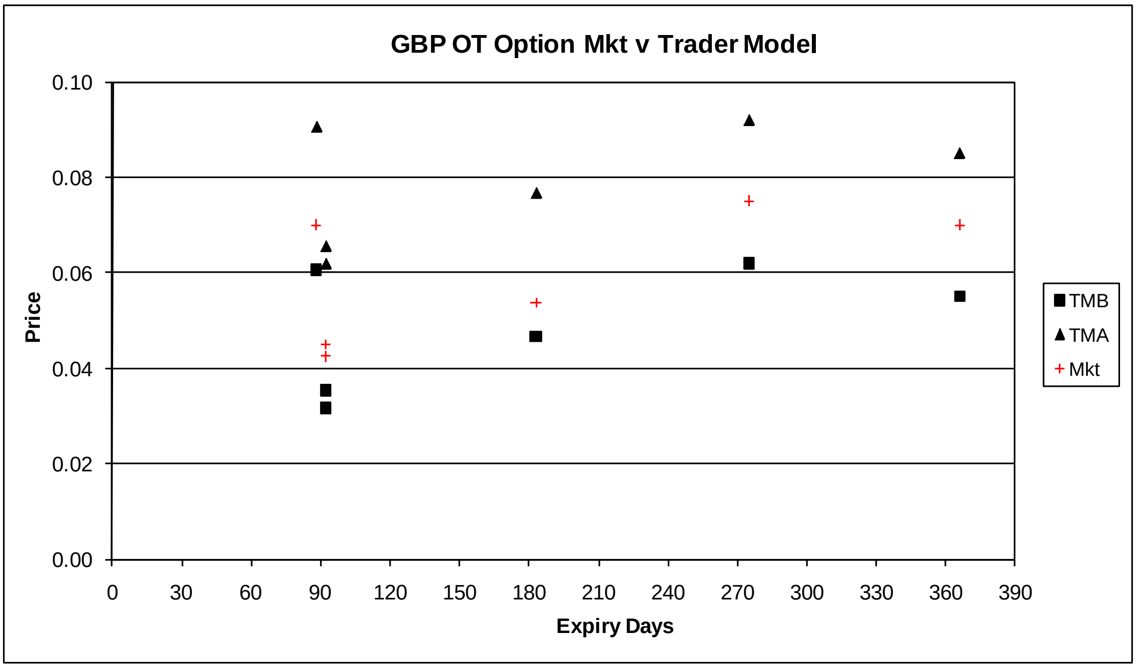

Fig. D20. GBP OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

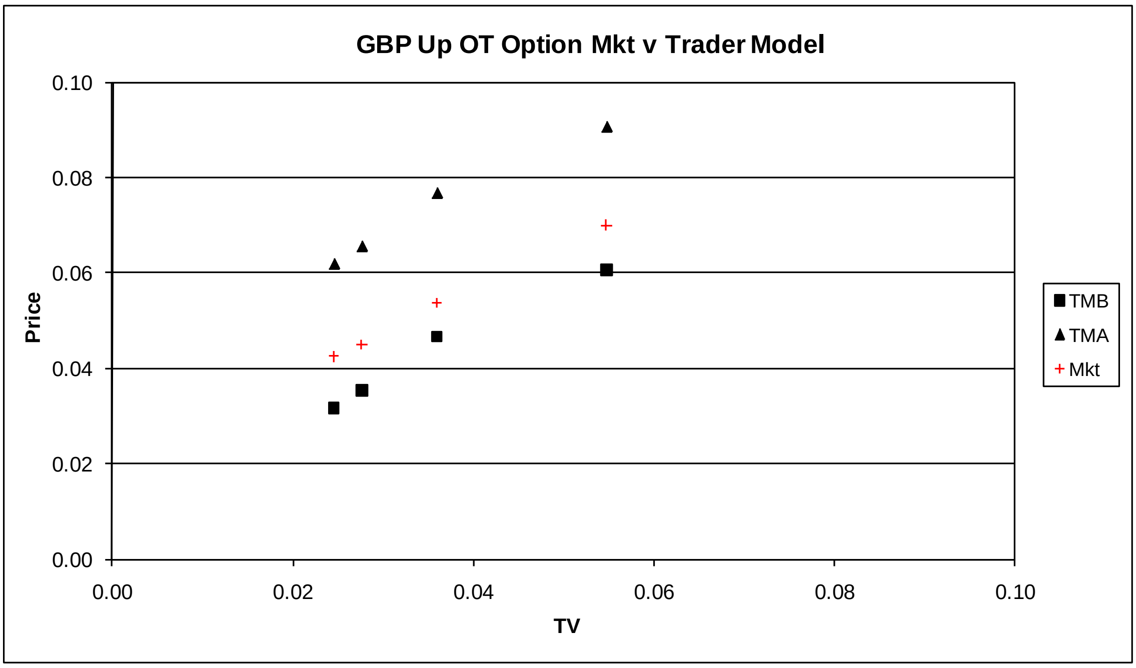

Fig. D21. GBP Up OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

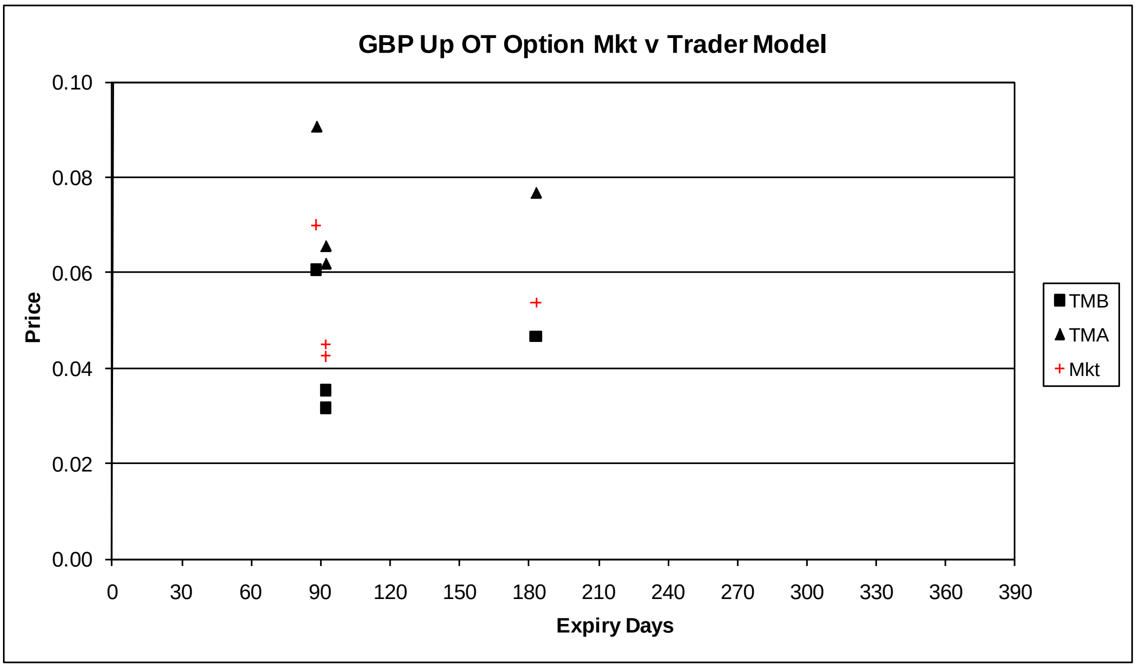

Fig. D22. GBP Up OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

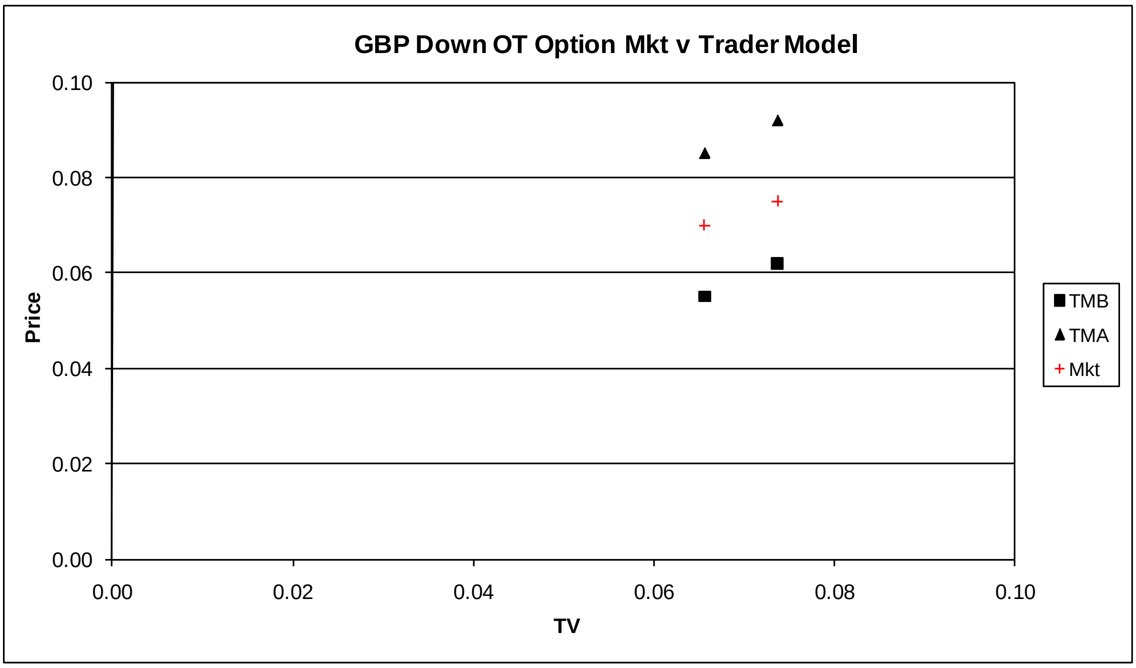

Fig. D23. GBP Down OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

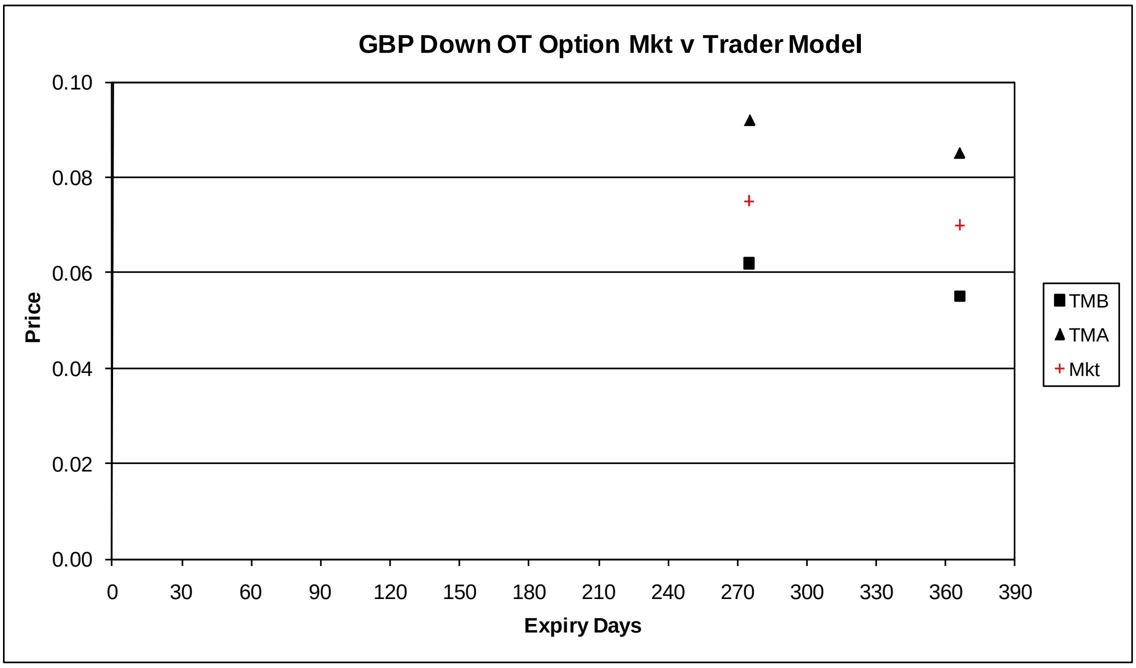

Fig. D24. GBP Down OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

AUD

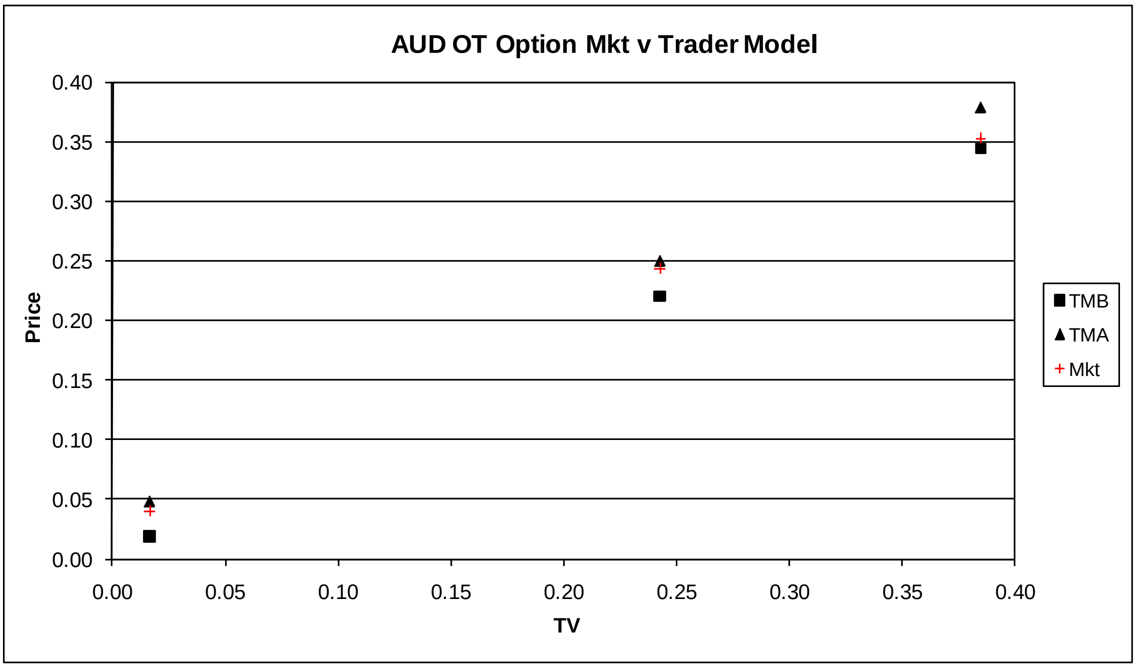

Fig. D25. AUD OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. The option with TV = 0.3846 is an AUD Up OT option, the others are AUD down.

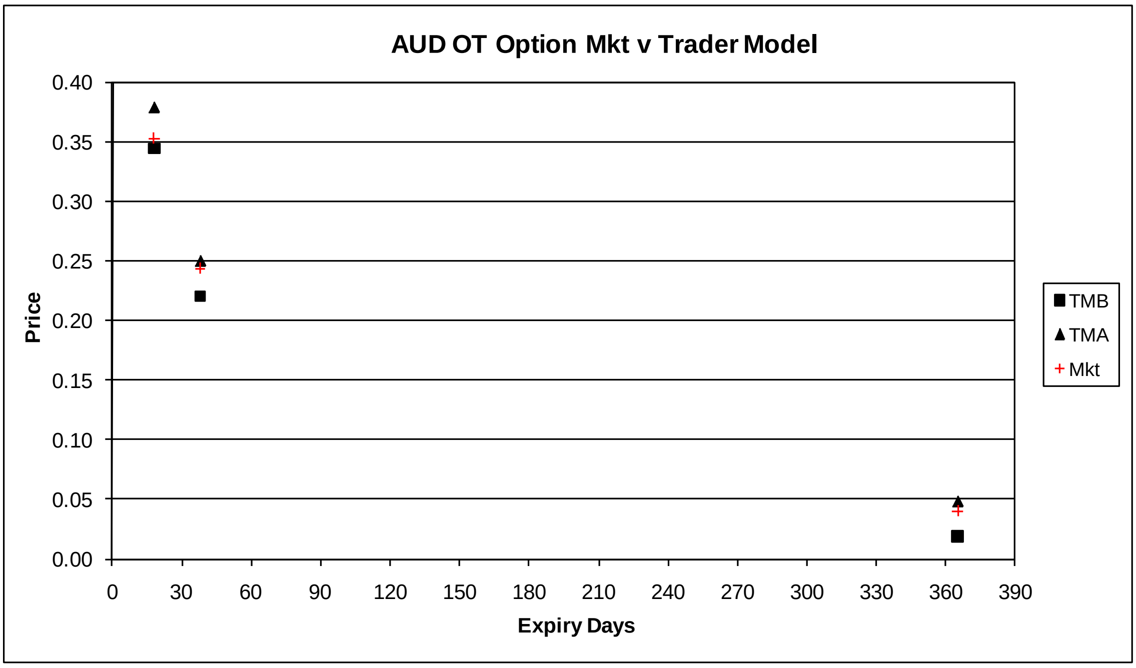

Fig. D26. AUD OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. The option with Expiry = 18 days is an AUD Up OT option, the others are AUD down.

CAD

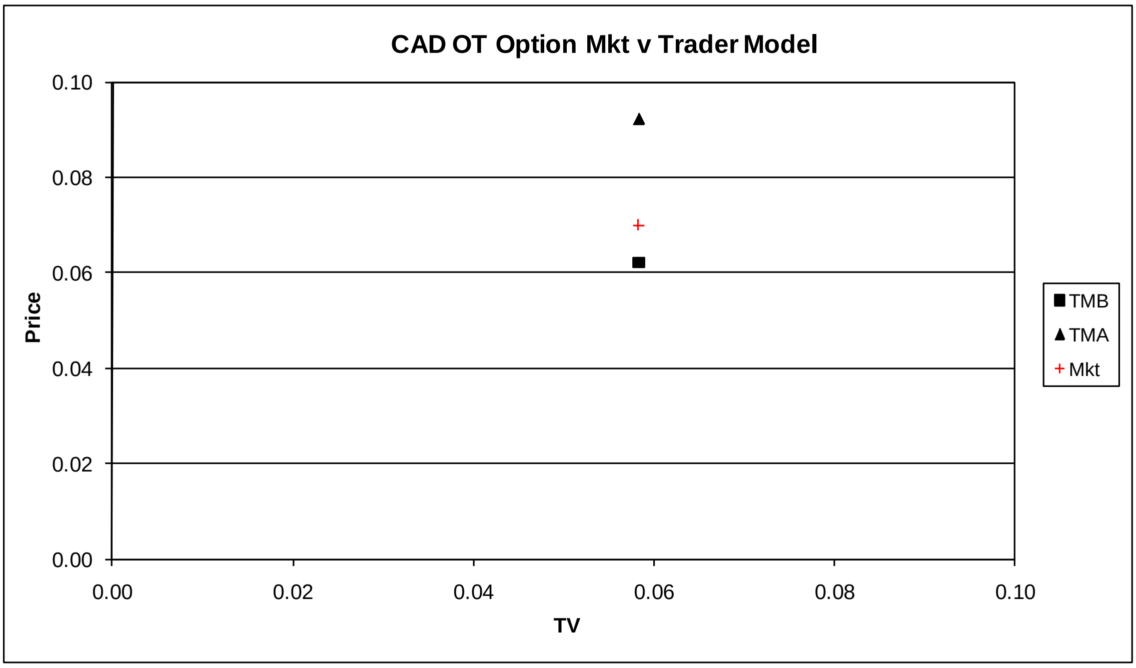

Fig. D27. CAD OT option traded market price versus Trader Model bid-ask prices, as a function of theoretical value. The option is a CAD Up OT option.

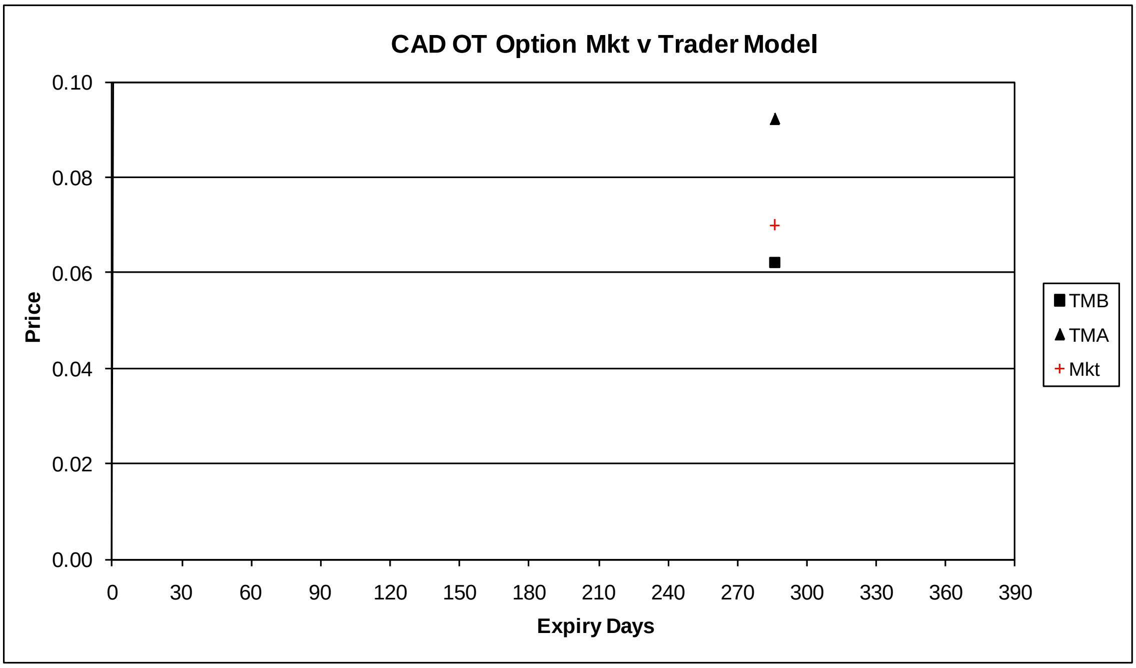

Fig. D28. CAD OT option traded market price versus Trader Model bid-ask prices, as a function of expiry days. The option is a CAD Up OT option.

EUR/CHF

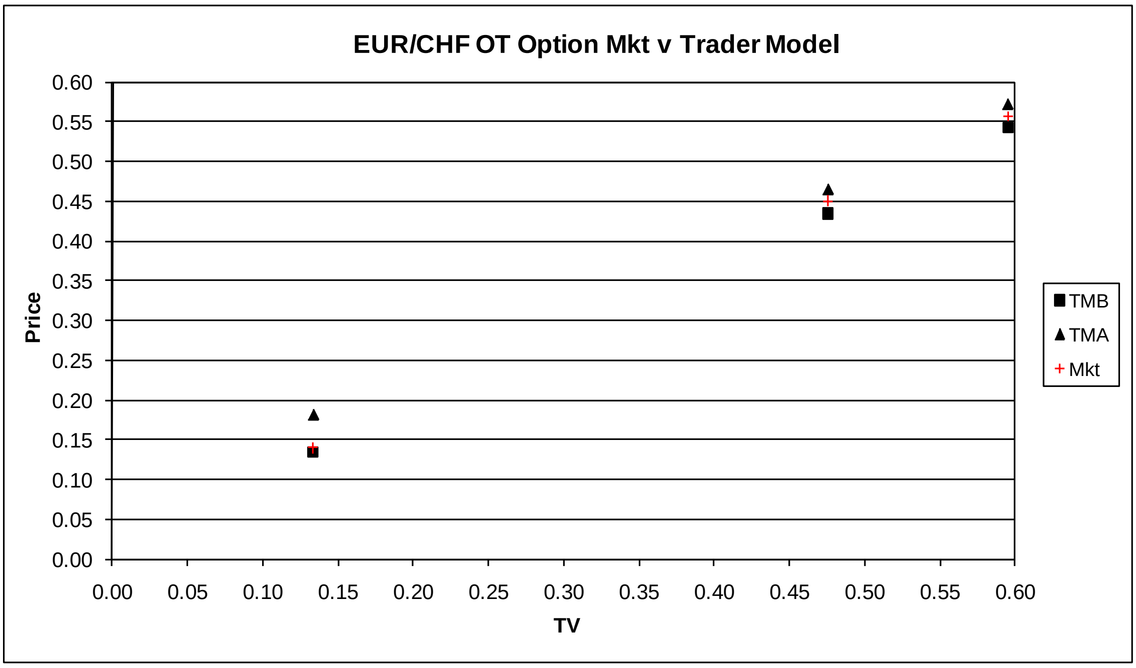

Fig. D29. EUR/CHF OT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. The option with TV = 0.595 is a EUR Up OT option, the others are EUR Down.

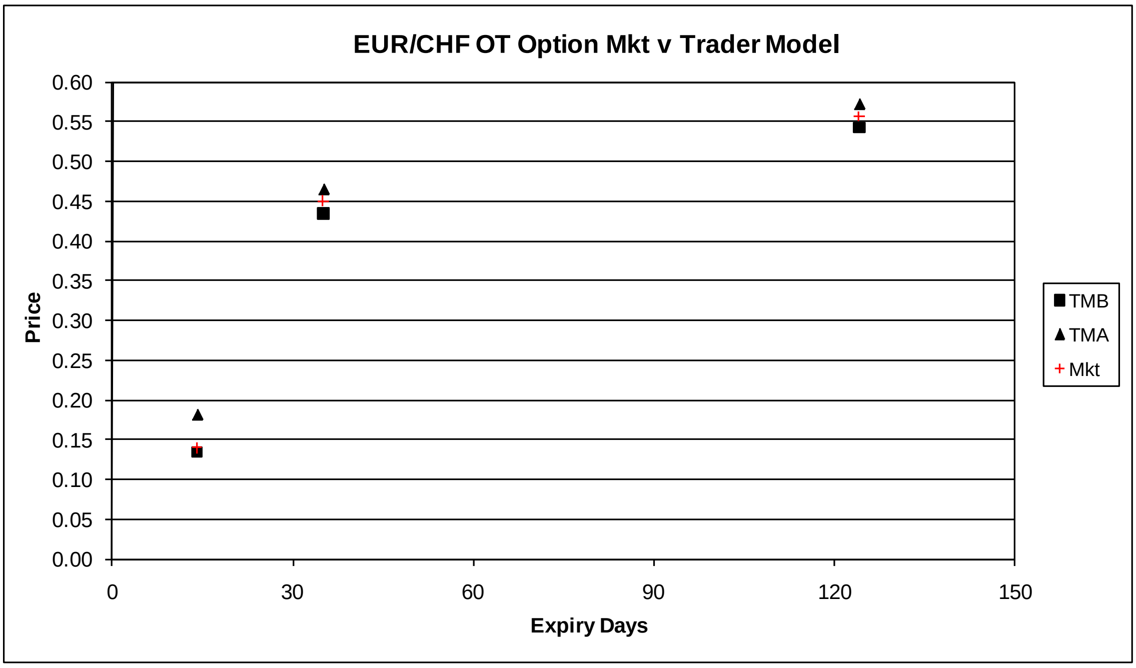

Fig. D30. EUR/CHF OT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. The option with expiry = 124 days is a EUR Up OT option, the others are EUR Down.

APPENDIX E

Trader Model pricing results for DNT options

EUR

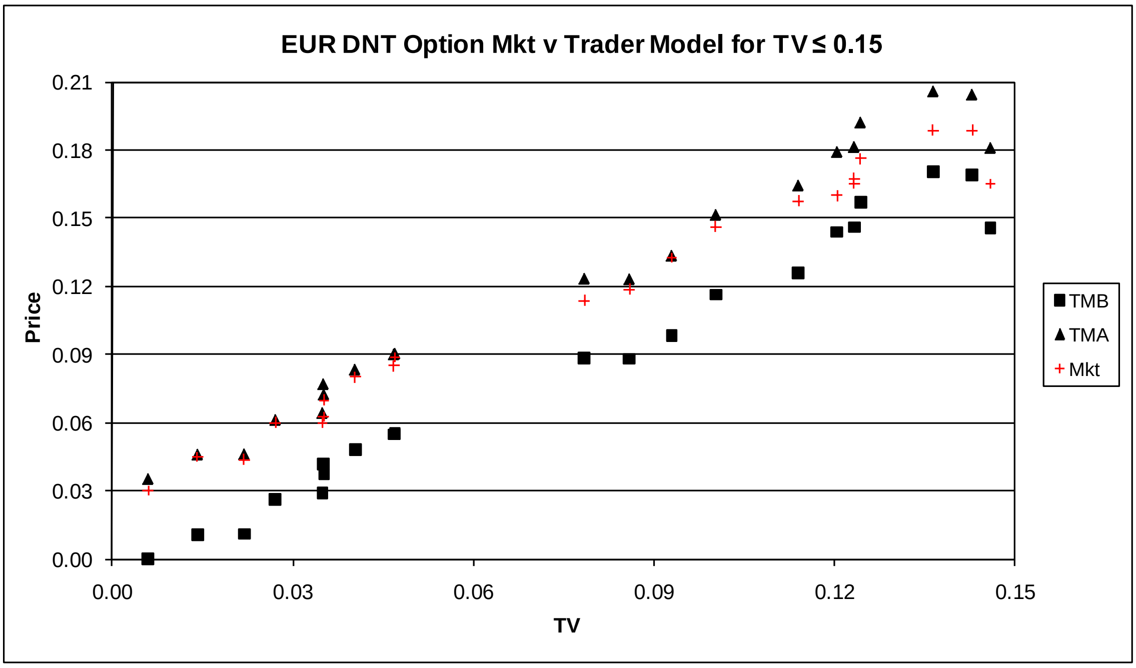

Fig. E1. EUR DNT option traded market prices versus Trader Model bid-ask prices, for theoretical values less than or equal to 0.15.

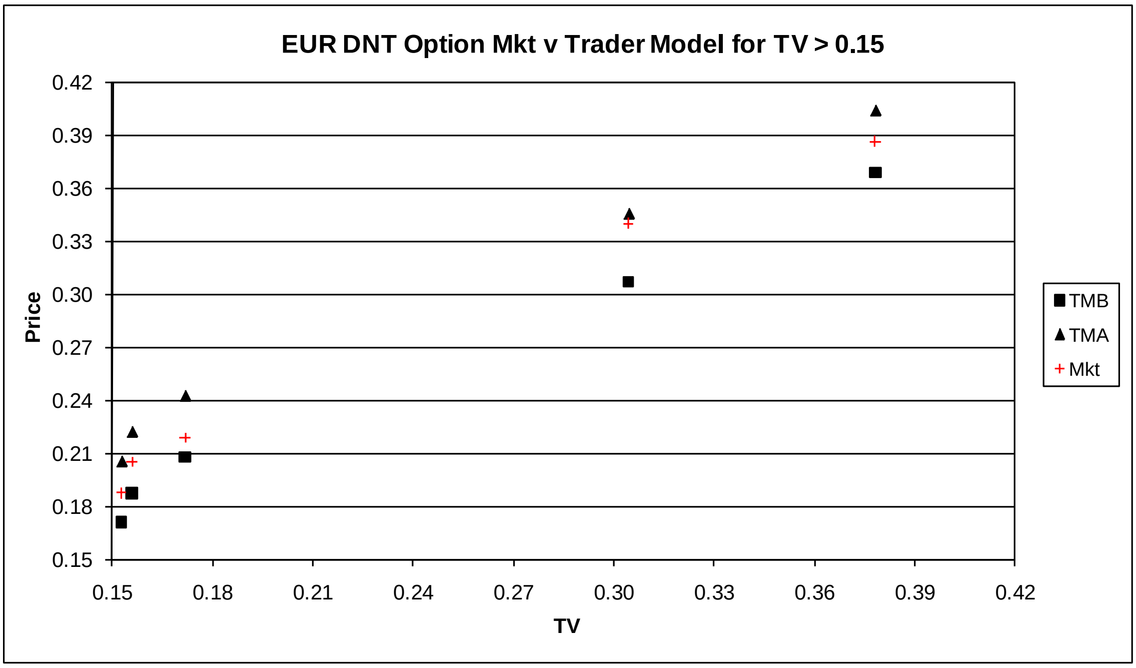

Fig. E2. EUR DNT option traded market prices versus Trader Model bid-ask prices, for theoretical values greater than 0.15.

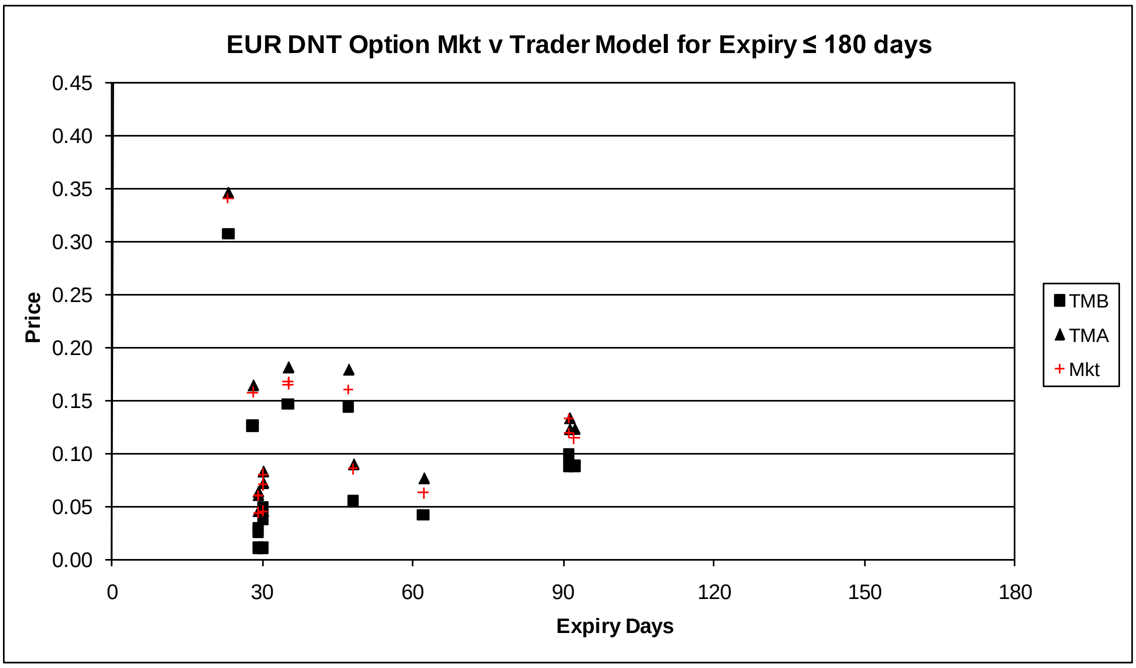

Fig. E3. EUR DNT option traded market prices versus Trader Model bid-ask prices, for expiries less than or equal to 180 days.

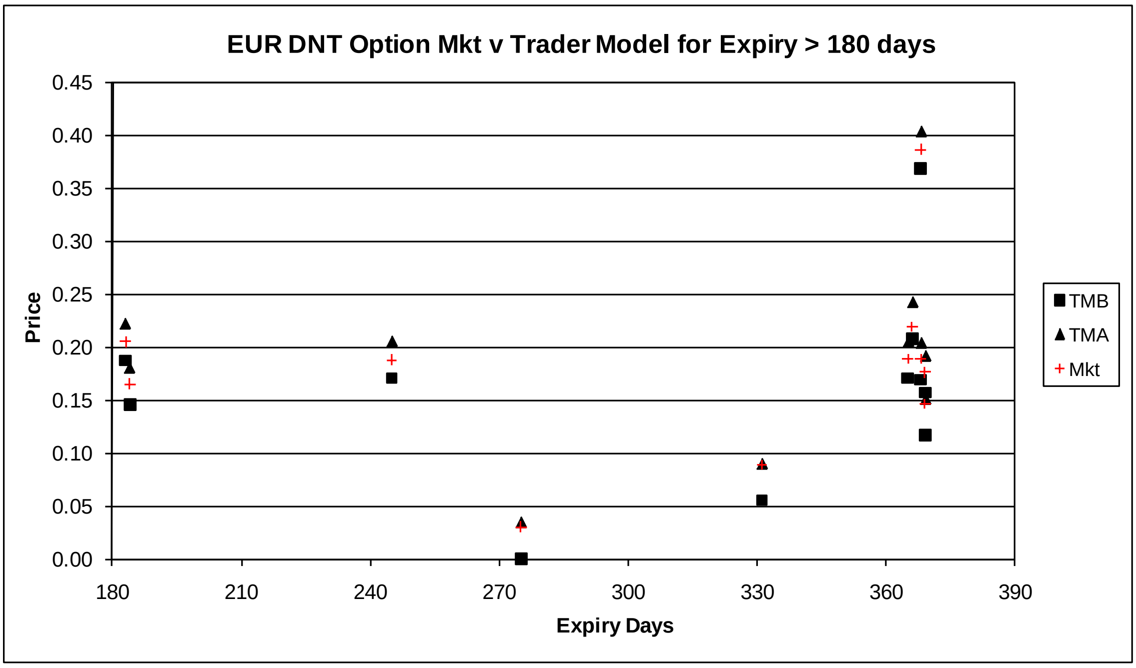

Fig. E4. EUR DNT option traded market prices versus Trader Model bid-ask prices, for expiries greater than 180 days.

JPY

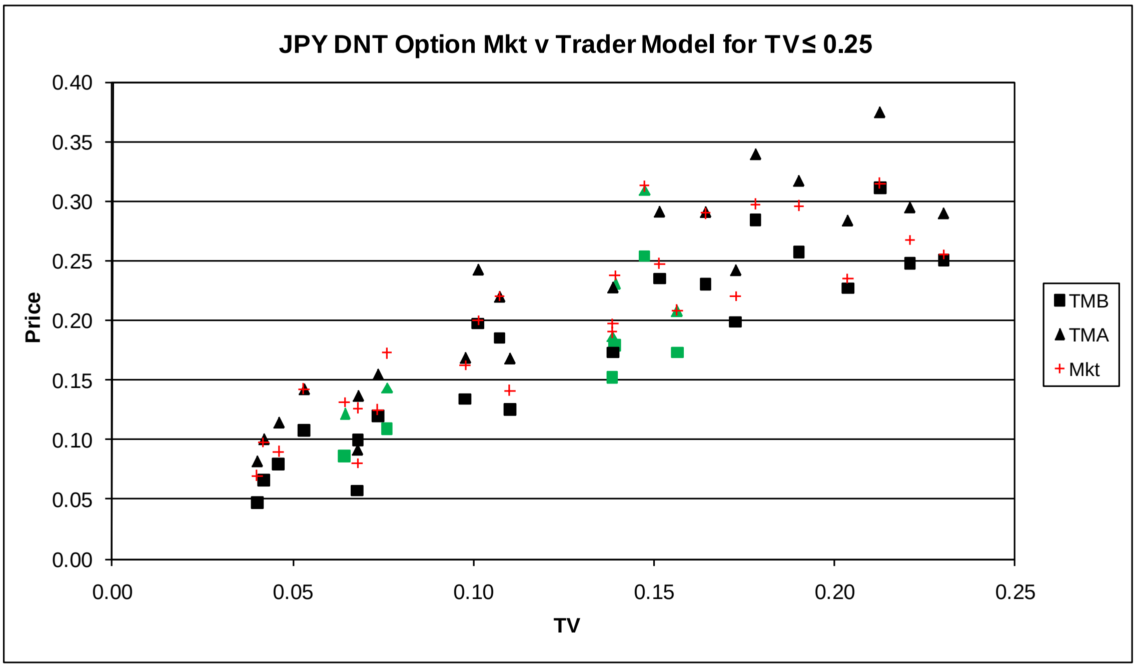

Fig. E5. JPY DNT option traded market prices versus Trader Model bid-ask prices, for theoretical values less than or equal to 0.25. Prices in green show JPY DNT options where market prices traded outside of the Trader Model bid-ask spread.

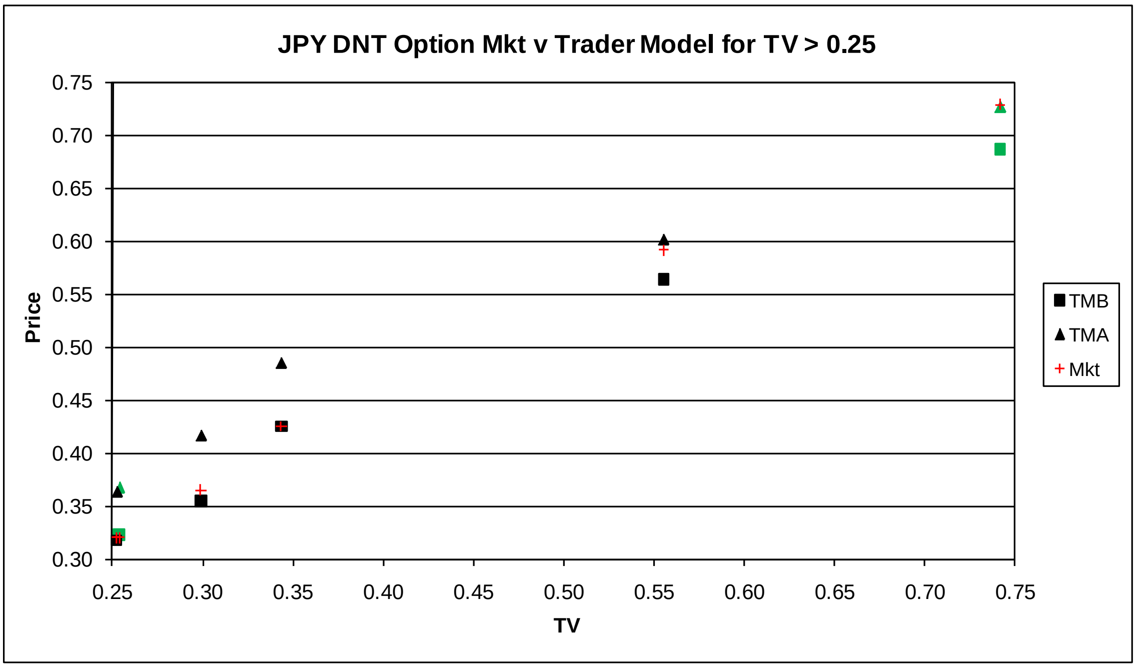

Fig. E6. JPY DNT option traded market prices versus Trader Model bid-ask prices, for theoretical values greater than 0.25. Prices in green show JPY DNT options where market prices traded outside of the Trader Model bid-ask spread.

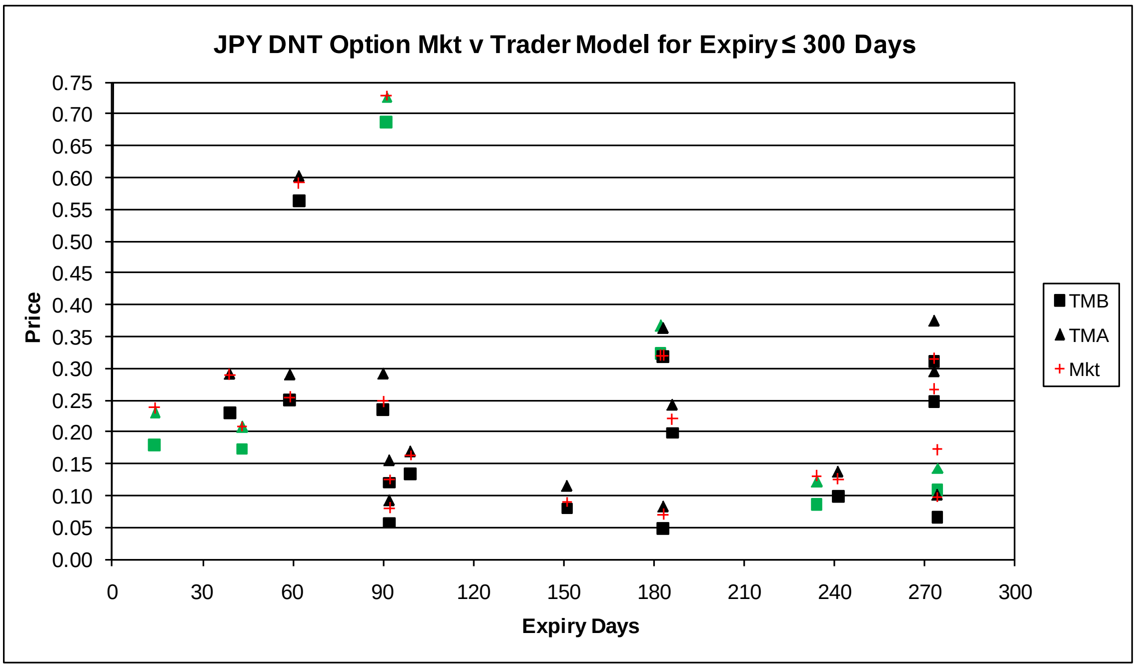

Fig. E7. JPY DNT option traded market prices versus Trader Model bid-ask prices, for expiries less than or equal to 300 days. Prices in green show JPY DNT options where market prices traded outside of the Trader Model bid-ask spread.

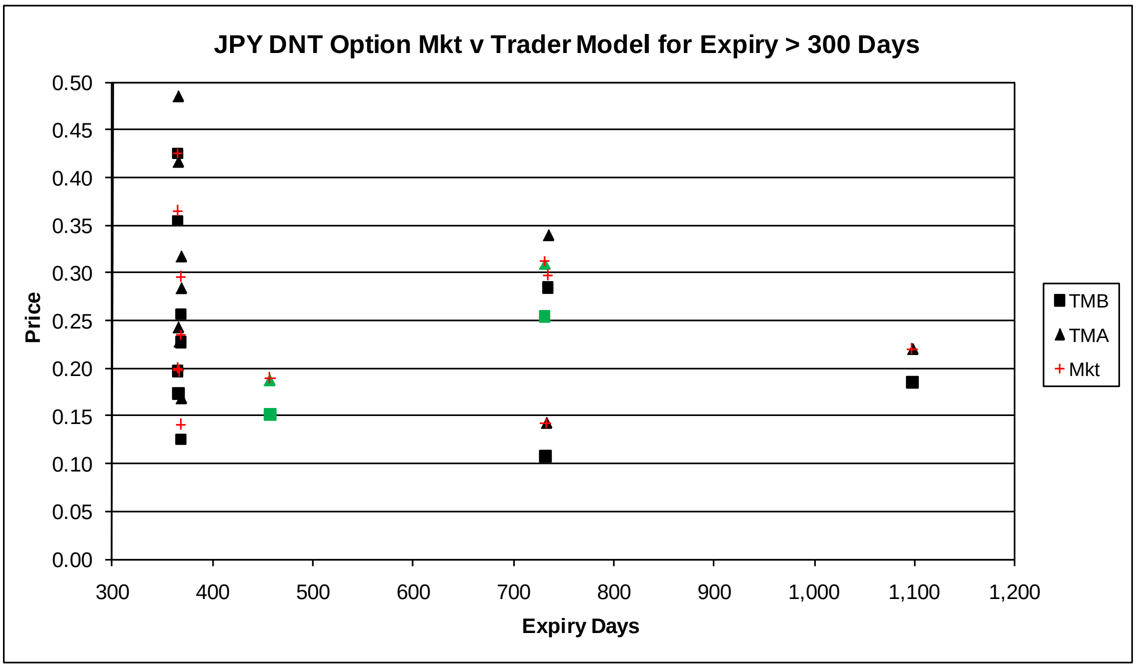

Fig. E8. JPY DNT option traded market prices versus Trader Model bid-ask prices, for expiries greater than 300 days. Prices in green show JPY DNT options where market prices traded outside of the Trader Model bid-ask spread.

EUR/JPY

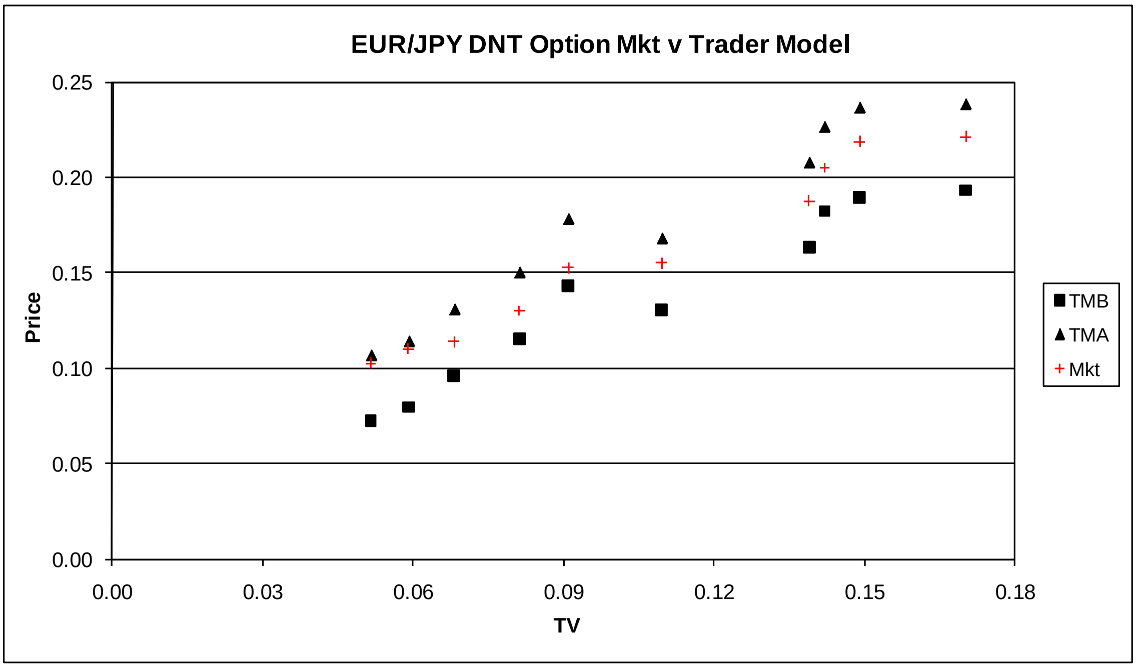

Fig. E9. EUR/JPY DNT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

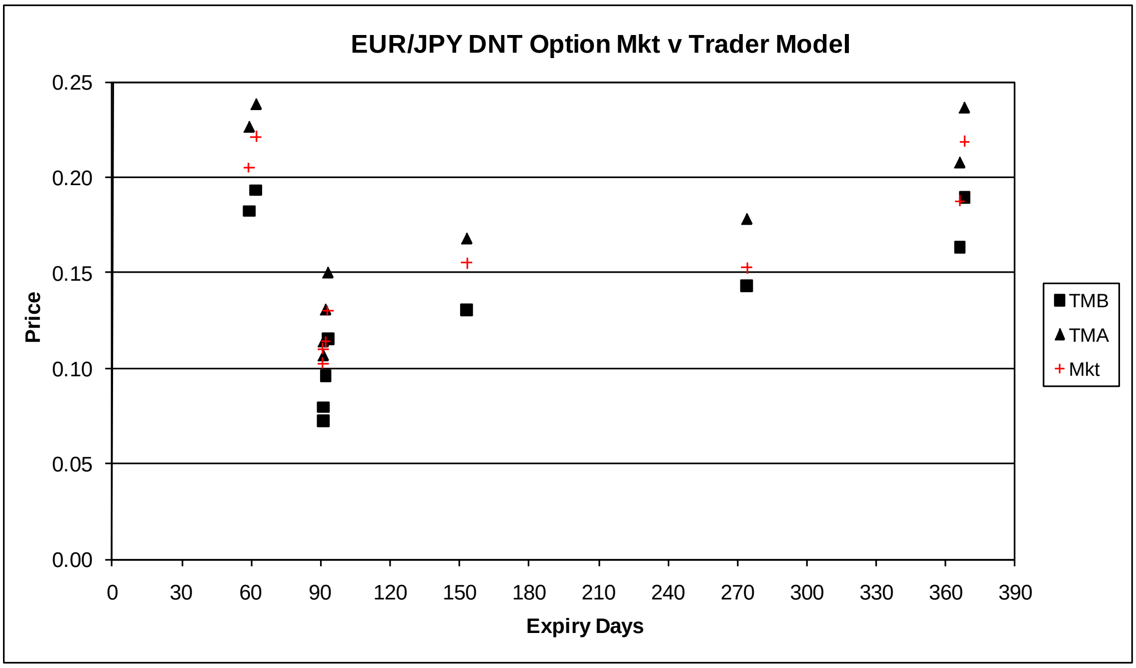

Fig. E10. EUR/JPY DNT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

GBP

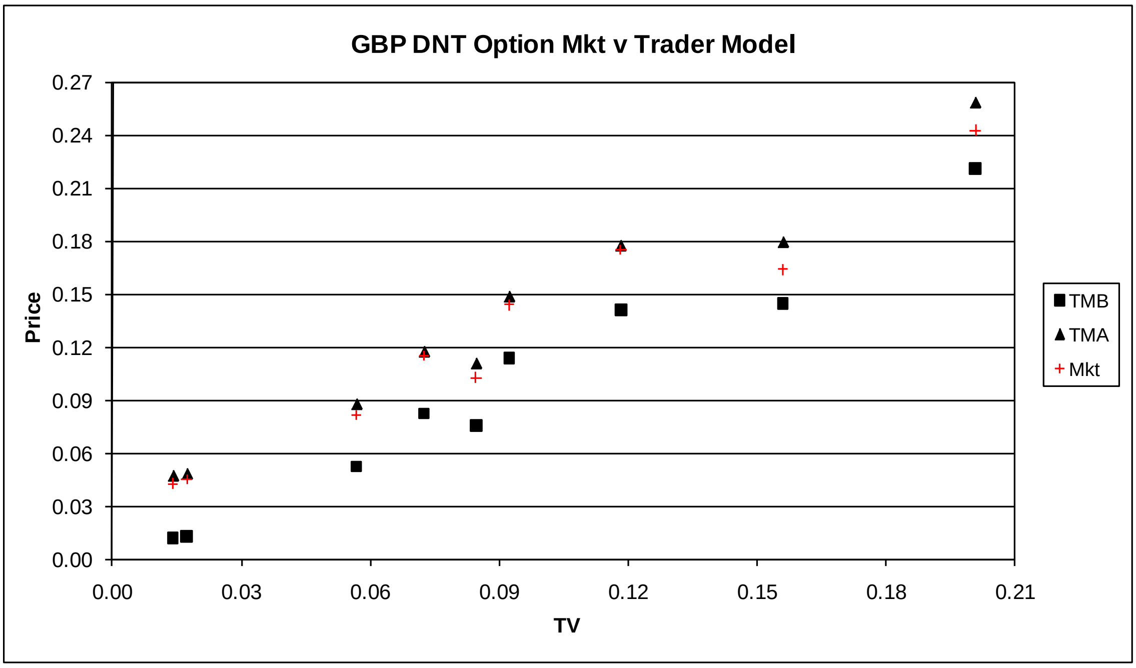

Fig. E11. GBP DNT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

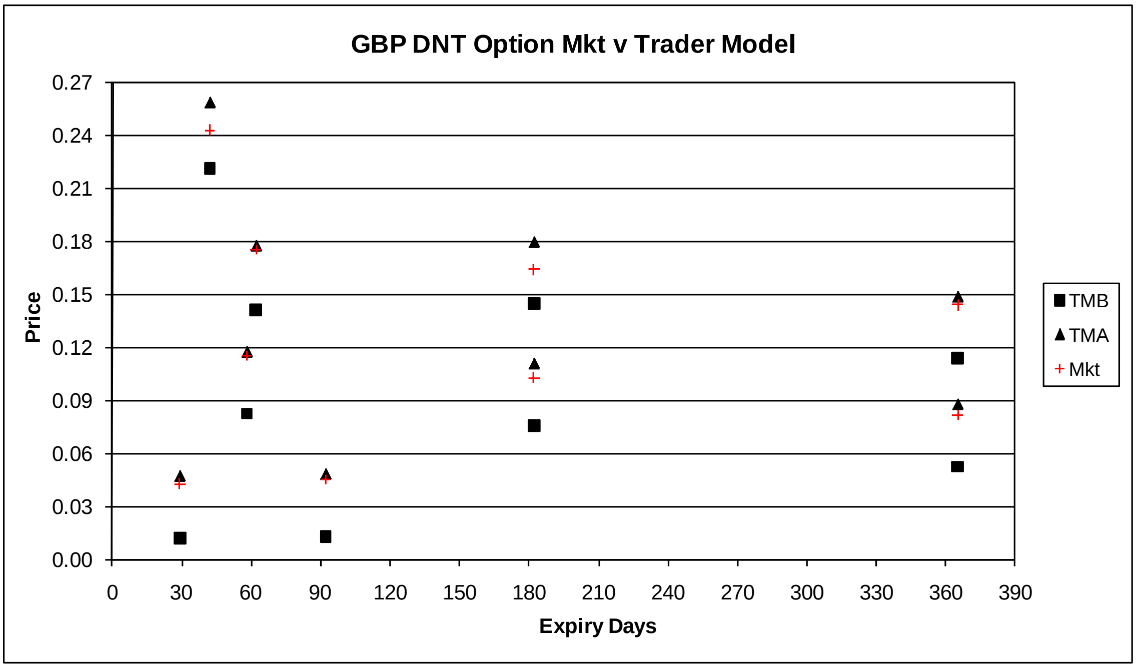

Fig. E12. GBP DNT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

AUD

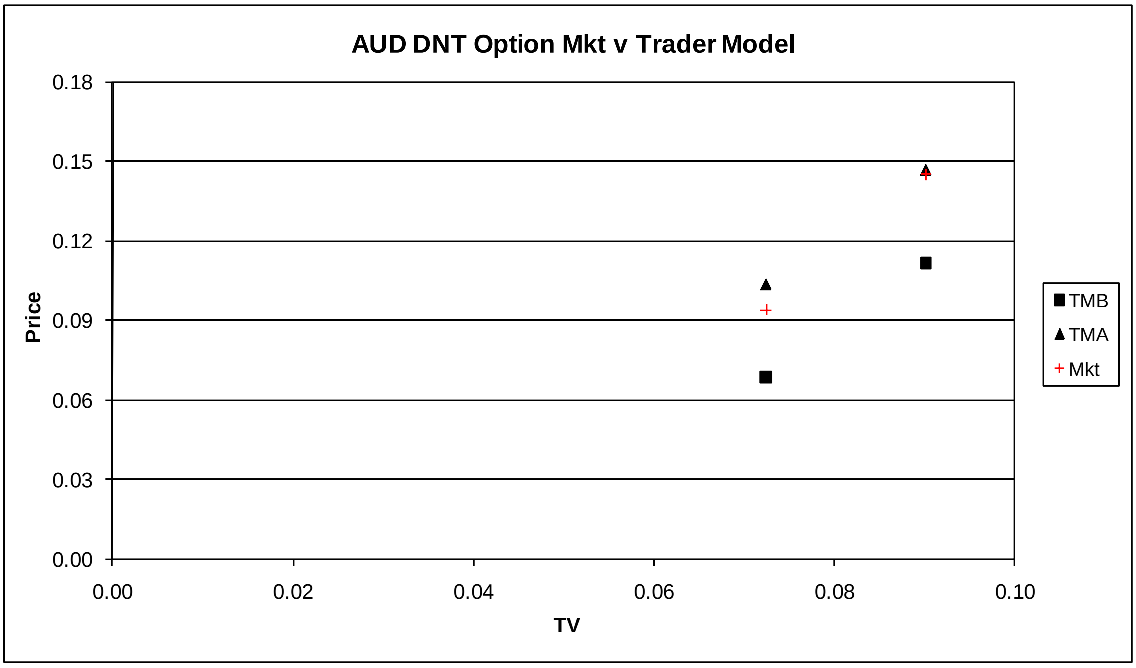

Fig. E13. AUD DNT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value.

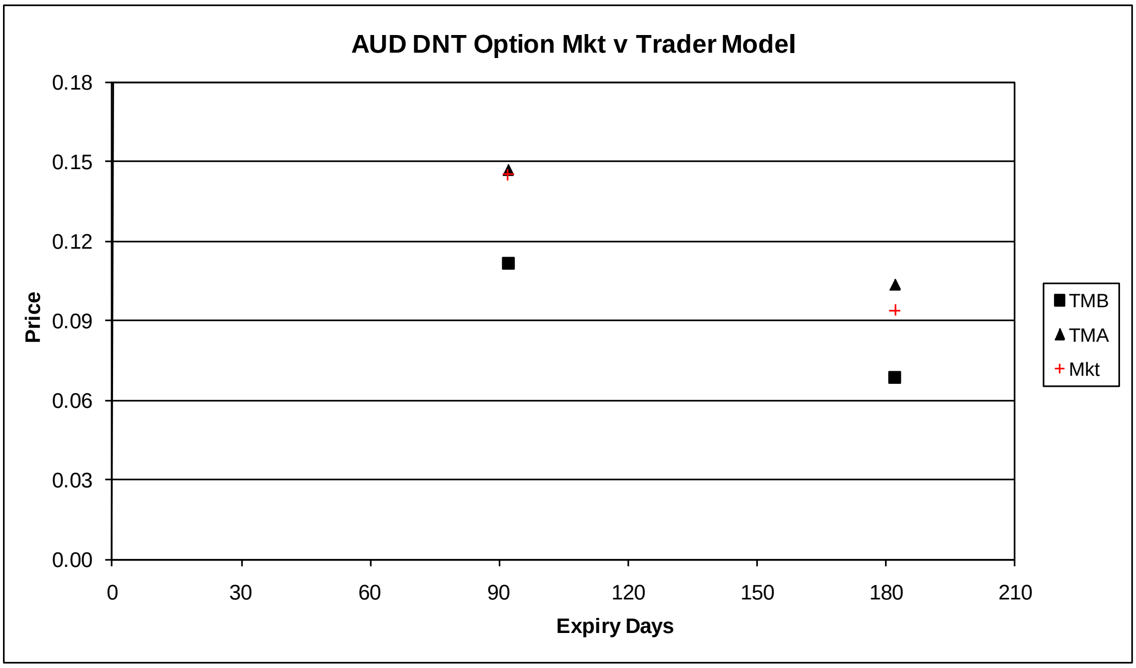

Fig. E14. AUD DNT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days.

CAD

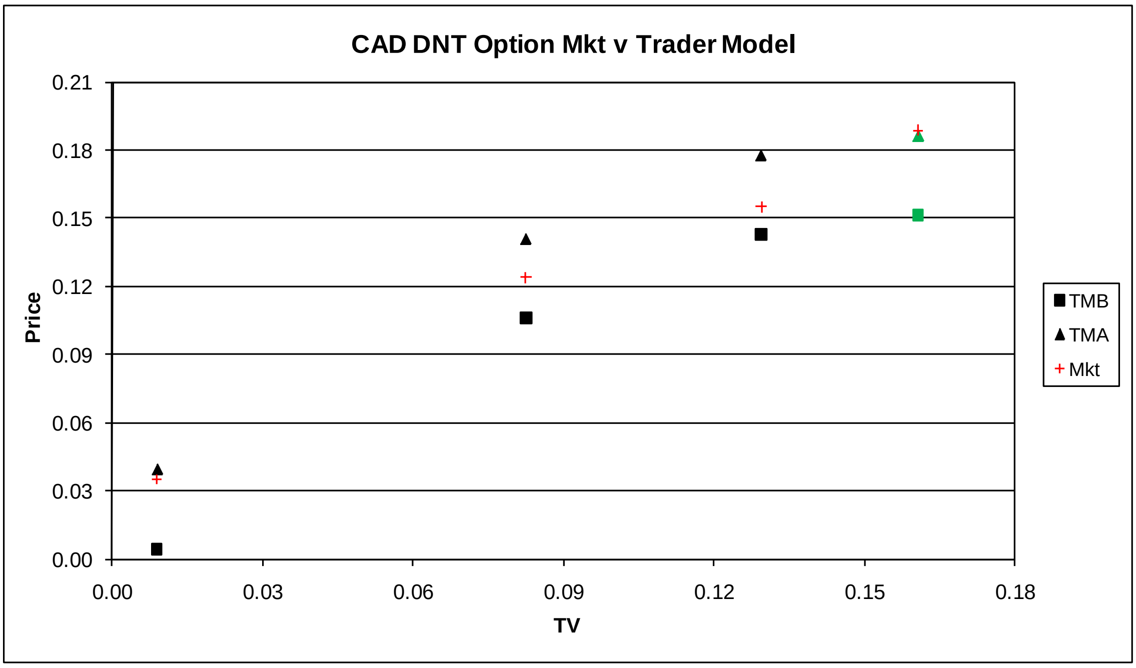

Fig. E15. CAD DNT option traded market prices versus Trader Model bid-ask prices, as a function of theoretical value. Prices in green show CAD DNT options where market prices traded outside of the Trader Model bid-ask spread.

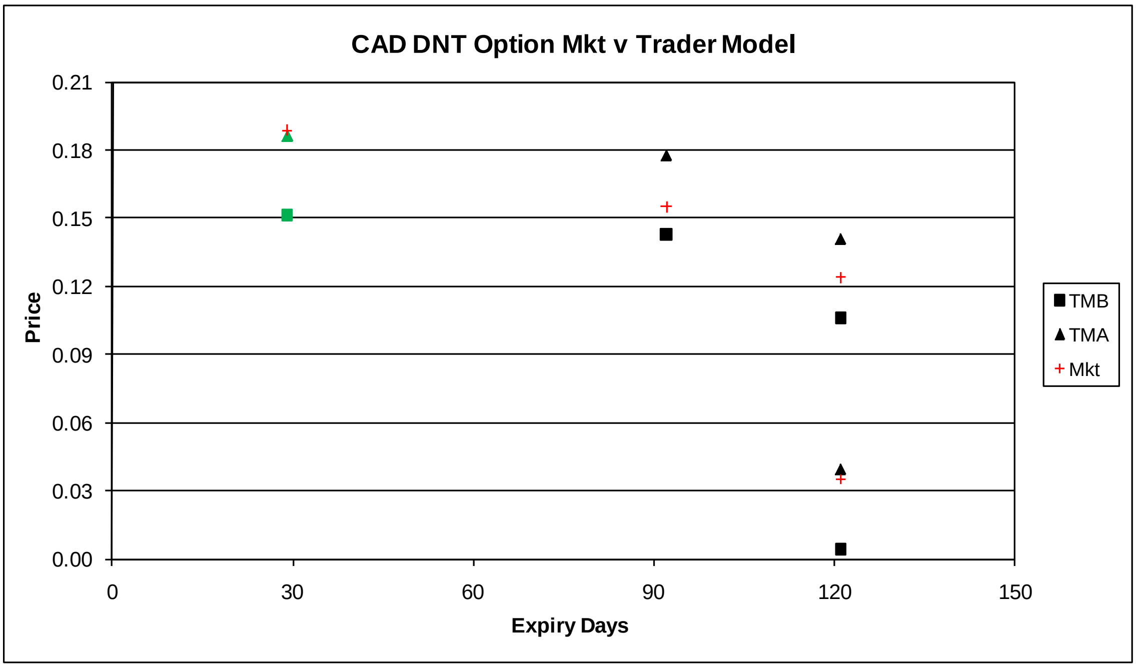

Fig. E16. CAD DNT option traded market prices versus Trader Model bid-ask prices, as a function of expiry days. Prices in green show CAD DNT options where market prices traded outside of the Trader Model bid-ask spread.

EUR/CHF

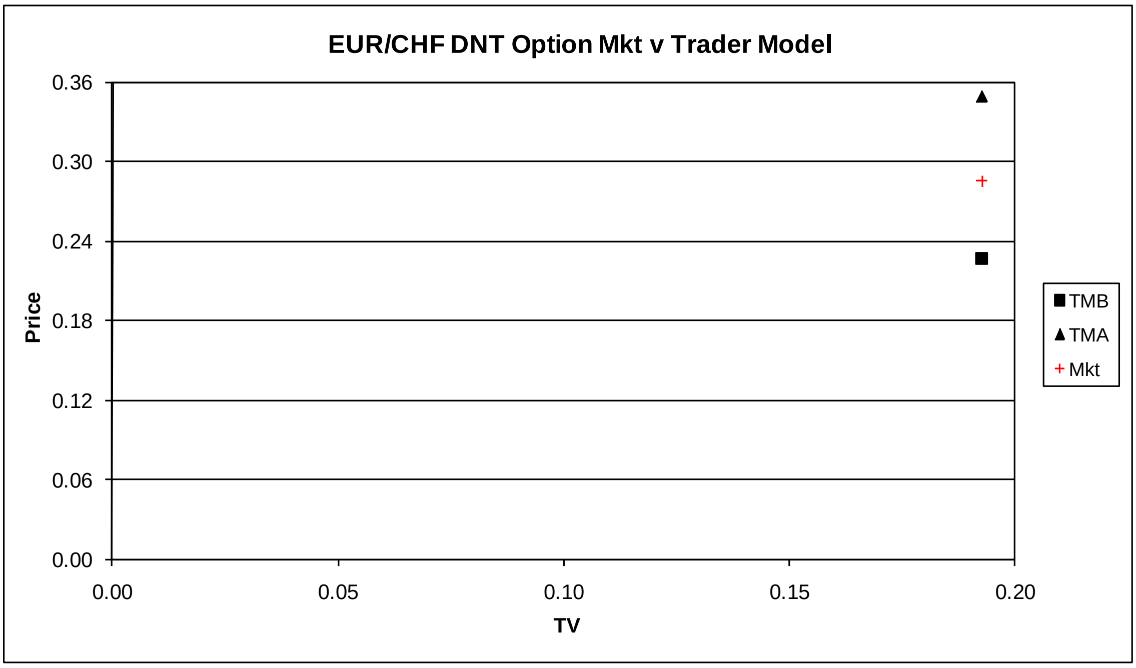

Fig. E17. EUR/CHF DNT option traded market price versus Trader Model bid-ask prices, as a function of theoretical value.

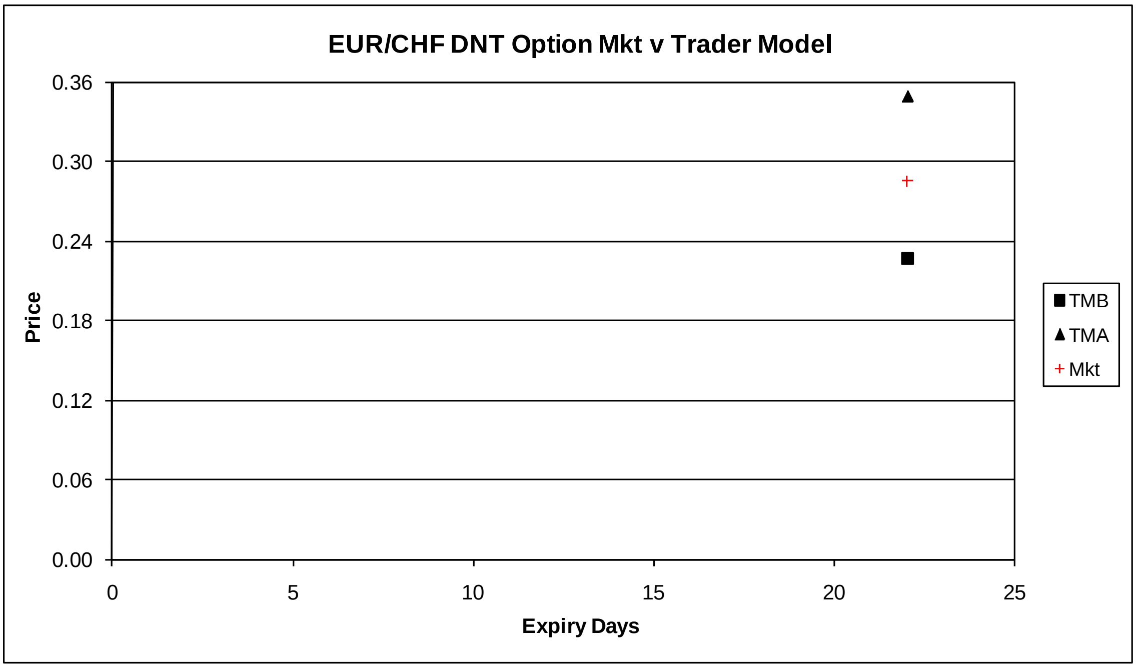

Fig. E18. EUR/CHF DNT option traded market price versus Trader Model bid-ask prices, as a function of expiry days.

EUR/GBP

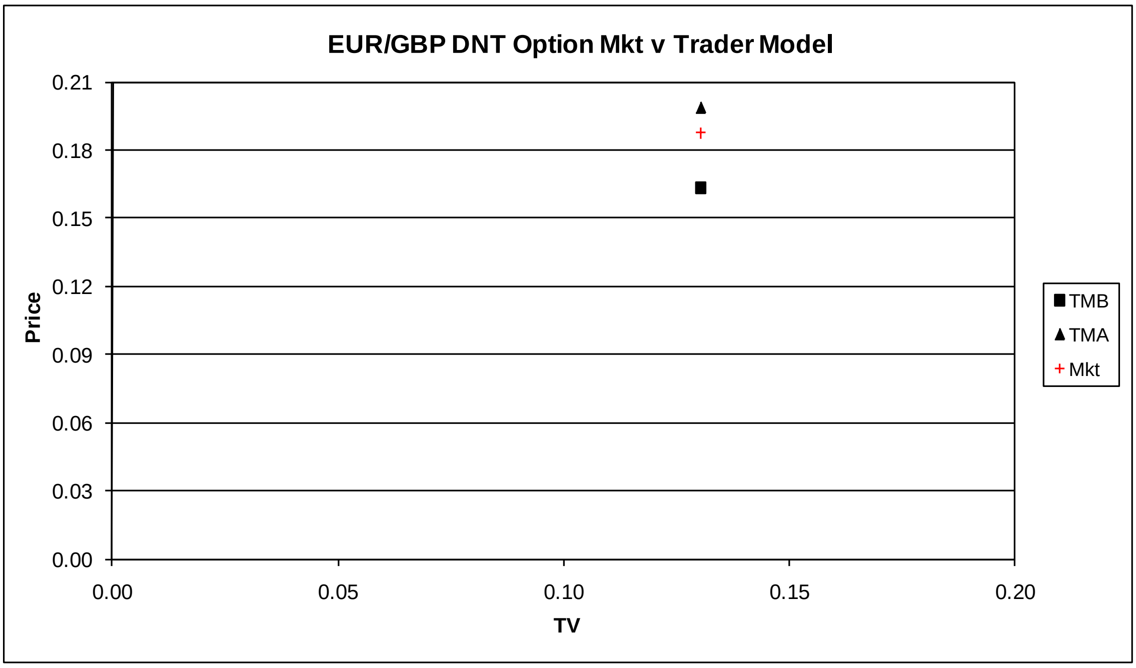

Fig. E19. EUR/GBP DNT option traded market price versus Trader Model bid-ask prices, as a function of theoretical value.

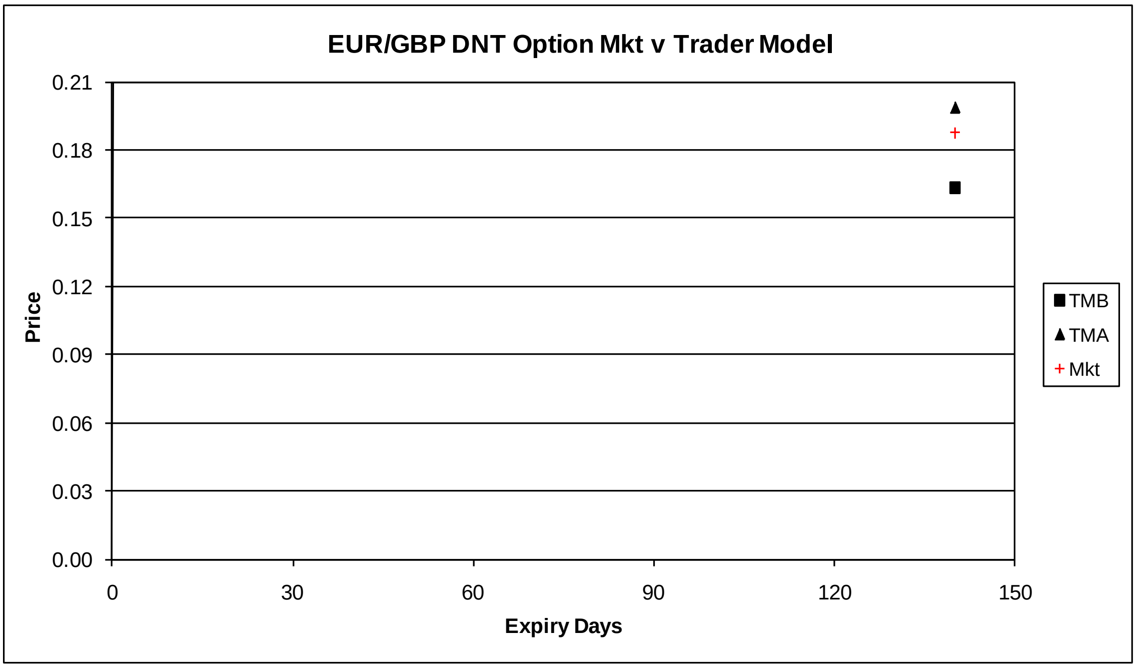

Fig. E20. EUR/GBP DNT option traded market price versus Trader Model bid-ask prices, as a function of expiry days.

GLOSSARY

If any of the following terms have multiple meanings, only the meaning which is relevant for the purpose of this thesis is shown.

American Option. A binary option which can terminate prior to expiry if the spot rate trades in the market at or beyond a barrier rate.

Ask. The lowest price at which a price-maker is prepared to sell the option.

Bid. The highest price at which a price-maker is prepared to buy the option.

Binary Option. An American One Touch (OT) or American Double-No-Touch (DNT) FX option, with continuously monitored barrier rates.

Book. Price-maker's portfolio of options.

BSM. The seminal option pricing papers of Black and Scholes (1973) and Merton (1973). In a FX option context, it has become associated with spot FX rate dynamics of the form $dS_t = (r_d - r_f)S_t dt + \sigma S_t dW_t$, where r_d, r_f and σ are constant, as per Garman and Kohlhagen (1983) and Grabbe (1983).

Buy-Side. A buy-side financial institution does not provide liquidity on demand to other market participants. It is instead a price-taker; a user of liquidity. Examples include hedge funds, portfolio managers, etc.

Calibration. Popular practice among financial engineers of using statics from the present (traded volatility surface) to define dynamics in the future (arbitrary volatility diffusions) by fitting arbitrary parameters to data via an optimisation algorithm.

Delta. Delta can be defined in terms of the spot or the forward. For American binary options, spot delta is important, because the termination condition is satisfied by the spot rate not the forward rate. That is, it is the spot rate which trades at or beyond a barrier rate to bring about early termination. The spot delta measures the change in the value of an option given a small change in the underlying spot FX rate. The spot delta is defined in the interbank FX option market in foreign percent or domestic points terms, depending on the FX pair. Delta is defined as $\left[\Delta = \frac{\partial G}{\partial S}\right]$.

Derivatives. Derivatives are products with payoffs and prices dependent upon the stochastic evolution of associated underlying financial variables (Bates, 2003, p. 388).

∂delta/∂vol. Aka vanna. ∂delta/∂vol is the change in the option delta given a small change in the option's volatility. It is equal to ∂vega/∂spot. i.e. $\left[\frac{\partial\Delta}{\partial\sigma}=\frac{\partial^2 G}{\partial\sigma\partial S}=\frac{\partial\Phi}{\partial S}\right]$.

Domestic Currency. FX rates are quoted as the number of units of domestic currency required to purchase one unit of foreign currency. For example, EUR is quoted as USD domestic, EUR foreign; and JPY is quoted as JPY domestic, USD foreign. The domestic currency is the numeraire.

Double-No-Touch Option. A DNT option obliges the seller to pay a fixed cash amount to the buyer if the spot FX rate trades in the market without ever touching or exceeding either barrier price prior to expiration. The liability can only be crystallised at expiration, and physical payment occurs on the delivery date of the option.

∂vega/∂vol. Also known as volga, vomma and vol of vol. ∂vega/∂vol is the change in the option vega given a small change in the option's volatility. i.e. $\Phi_\sigma=\frac{\partial\Phi}{\partial\sigma}=\frac{\partial^2 G}{\partial\sigma^2}$.

European Option. A vanilla option which either expires worthless or is exercised by the buyer of the option, on expiration.

Financial Engineer. Professional with qualifications principally in the mathematics of the physical sciences. Also known as a quant.

Foreign Currency. FX rates are quoted as the number of units of domestic currency required to purchase one unit of foreign currency. For example, EUR is quoted as USD domestic, EUR foreign; and JPY is quoted as JPY domestic, USD foreign.

Forward FX Rate. The foreign currency exchange rate for delivery later than spot delivery. In the interbank FX market, forward swap points are traded directly, not the outright forward rate. Where F_t is the outright forward rate, $F_t = S_t e^{(r_d - r_f)(T-t)}$. Swap points are equal to $F_t - S_t$.

Franchise Flows. From the perspective of the price-maker, an option traded with corporate and institutional customers of the counterparty. To be distinguished from other bank counterparties (interbank flows).

Gamma. Delta convexity to underlying spot FX rates (∂delta/∂spot). i.e. $\Gamma=\frac{\partial\Delta}{\partial S}=\frac{\partial^2 G}{\partial S^2}$.

Gap Delta. The size of the delta position which needs to be unwound in the spot FX market if an American binary FX option terminates. For example, if a long 'down' OT option with a delta of -$400 million is hedged in the spot dimension by $400 million of spot FX, if the barrier FX rate trades in the market the delta of the OT option goes to zero, and the price-maker is left with an open spot FX position of +$400 million in a falling spot FX market.

Geometric Brownian Motion. Diffusion of the form $dX = \alpha X dt + \sigma X dW$ with drift α and volatility σ. The conditional mean of $\ln(X_u)$ for $u > t$ is $\ln(X_t) + \alpha(u-t) - \frac{1}{2}\sigma^2(u-t)$ and the conditional standard deviation of $\ln(X_u)$ is $\sigma\sqrt{(u-t)}$. $\ln(X_u)$ is normally distributed. The conditional expected value of X_u is $X_t e^{\alpha(u-t)}$ (Shimko, 1992).

Give the bid. Price-taker sells at the price-maker's bid price.

Greeks. Option greeks are usually described as low-order and high-order. Low-order greeks include delta, gamma, vega, rho, phi and theta. High-order greeks include volga and vanna.

High TV. An American binary option with a theoretical value greater than 0.30. Whilst the cut-off point is somewhat arbitrary, the essential characteristic is that a buyer of the option risks a relatively large amount to earn potentially a modest reward.

Implied Volatility Surface. The set of all European vanilla option smiles and skews for each expiry. In FX option markets, implied volatility (σ) is a matrix of dimension ΔxT. The matrix is only relevant for models of the BSM paradigm. In the FX option market, 'implied volatility surface' and 'volatility surface' are used interchangeably. However, in FX option markets, the surface is not implied but directly traded, and hence, traded volatility surface is more accurate terminology.

Internal Distribution Margins. Trading floor sales practitioners widen price-maker bid-ask spreads when quoting franchise flows, in order to generate income as compensation for developing trading ideas, hedging structures etc. for business that is executed.

Interbank Market. The interbank market for FX options comprises banks and inter-dealer brokers only. In this market, each option trade ultimately has as counterparties two banks. The trade may be negotiated directly between banks or indirectly via inter-dealer brokers.

Inventory. The options in price-makers' books are known as inventory. Inventory usually consists of long and short option positions from a diverse range of interbank, corporate and institutional counterparties.

Lean prices to the 'right' or 'left'. Price-makers lean model prices to the right when they are aggressive bidders, and to the left when aggressive sellers. e.g. model bid-ask prices of 0.10/0.12, can be leaned to the right and made 0.1025/0.1225, or leaned to the left and made 0.0975/0.1175, to show mildly aggressive bid and ask interest, respectively.

Lift the ask. Price-taker buys at the price-maker's ask price.

Long the barrier. One is 'long the barrier' if one wants the spot rate to trade at or beyond a barrier rate. For example, if a price-maker buys (sells) a OT (DNT) option they are long the barrier. If the spot rate trades at or beyond a barrier rate before expiry, then, in the OT option case, their payoff is crystallised and they receive the full face value from the counterparty on the delivery date of the option. In the DNT option case, their contingent liability is extinguished.

Low TV. An American binary option with a theoretical value less than 0.30. Whilst the cut-off point is somewhat arbitrary, the essential characteristic is that a buyer of the option risks a relatively small amount to earn potentially a much larger reward.

Market-Maker. Trader in a sell-side financial institution responsible for preparing bid and ask prices on demand for corporate, institutional and interbank counterparties, either direct or through brokers. Provider of liquidity. Also known as a price-maker.

Market Supplement. The difference between an option's actual traded market price and its theoretical value. The market supplement can be positive, negative or zero.

Model Risk. The risk that the model being used to price a binary option does not reflect actual financial economic substance.

Moneyness. At-the-money (ATM) options in the FX option market are zero-delta straddles. Out-of-the-money (OTM) options have deltas of smaller magnitude than the options comprising the straddle. In-the-money (ITM) options have deltas of greater magnitude than the options comprising the straddle. e.g. if the Call option in the straddle has $\Delta_{ATM}^{Call} = 0.48$, then $\Delta_{OTM}^{Call} < 0.48$ and $\Delta_{ITM}^{Call} > 0.48$.

One Touch Option. A OT option obliges the seller to pay a fixed cash amount to the buyer if the spot FX rate trades in the market at or beyond the barrier price, prior to expiration. While the liability is crystallised immediately, physical payment occurs on the delivery date of the option.

Over-the-counter. Markets are either over-the-counter or exchange traded. For FX and FX options, over-the-counter traded volume dominates exchange traded volume by a significant margin.

Phi. The change in the value of the option given a small change in the foreign deposit rate. For example, the phi of a USD/JPY option is the sensitivity of the option price to a small change in the USD deposit rate. i.e. $\phi = \frac{\partial G}{\partial r_f}$.

Price-Maker. Trader in a sell-side financial institution responsible for preparing bid and ask prices on demand for corporate, institutional and interbank counterparties, either direct or through brokers. Provider of liquidity. Also known as a market-maker.

Price-Taker. Trader in a buy- or sell-side institution who demands bid and ask prices from a price-maker, either direct or through brokers. User of liquidity.

Quant. Financial engineer with qualifications principally in the mathematics of the physical sciences.

Rho. The change in the value of the option given a small change in the domestic deposit rate. For example, the rho of a USD/JPY option is the sensitivity of the option price to a small change in the JPY deposit rate. i.e. $\rho = \frac{\partial G}{\partial r_d}$.

Risk Reversal. **1.** A liquid European vanilla option strategy consisting of a long (short) OTM Call and a short (long) OTM Put, where the Call and Put have different strikes and identical delta. Trades in the interbank FX option market with a delta hedge to make it delta-neutral. It prices the skew in the European vanilla option market. **2.** Any option strategy where the slope of the risk changes sign. This is consistent with Taleb's (1997, p. 275) definition that "a risk reversal for a book manager is the switch in risk across one point", such as "where the gamma and / or vegas flip from positive to negative across one point". In this thesis, the broader definition (2) is used. In this thesis, OTM legs of the risk reversal have a delta equal to ±0.10.

Sell-Side. A sell-side financial institution provides liquidity on demand to other market participants. A sell-side financial institution is a price-maker; a provider of liquidity. Examples include commercial and investment banks.

Short the barrier. One is 'short the barrier' if one does not want the spot rate to trade at or beyond a barrier rate. For example, if a price-maker sells (buys) a OT (DNT) option they are short the barrier. If the spot rate trades at or beyond a barrier rate before expiry, then, in the OT option case, their contingent liability is crystallised and they have to pay the full face value to the counterparty at the delivery date of the option. In the DNT option case, their asset becomes worthless.

Skew. The set of asymmetric volatilities at expiry T, which represents the 'price' of (delta neutral) risk reversals at different delta pillars Δ_i. Skews can be positive, negative or zero.

Slippage. The risk that the price-maker unwinds the delta hedge for a binary option at a worse spot rate than the barrier rate that has just traded. This is a significant risk when the spot rate approaches the barrier rate close to expiry.

Smile. The set of symmetric volatilities at expiry T, which represents the 'price' of vega neutral butterflies at different delta pillars Δ_i. Smiles are usually positive.

Spot (FX) Rate. The foreign currency exchange rate for immediate delivery in theory, and for delivery in two business days in practice (except USD/CAD delivery which is one business day). Quoted as the amount of domestic currency required to buy one unit of foreign currency. E.g. USD/JPY is the amount of JPY required to buy one US dollar.

Spread. The difference between the price-maker's ask price and bid price. e.g. A bid price of 0.30 and an ask price of 0.33 has a spread equal to 0.03.

Straddle. A European vanilla option trading strategy with a long or short position in both a Put and a Call, with identical strikes, such that $\text{Delta}_{\text{Put}} + \text{Delta}_{\text{Call}} = 0$ when dealt. Strike prices (K) are $K_{0\Delta}^{\text{pct}} = S \big/ \exp\left(-\left(r_d - r_f - 0.5\sigma^2\right)T\right)$ or $K_{0\Delta}^{\text{pts}} = S \big/ \exp\left(-\left(r_d - r_f + 0.5\sigma^2\right)T\right)$, for delta expressed in percent foreign and domestic points, respectively.

Strangle. A European vanilla option trading strategy with a long or short position in both a Put and a Call, with different strikes, such that $\text{Delta}_{\text{Put}} + \text{Delta}_{\text{Call}} = 0$. In the interbank FX option market, only OTM options are used.

Strike Structure of Volatility. The vector of OTM volatilities for all option expiries T_i. i.e. the smile and / or skew for each maturity. In the FX option market, it is more accurate to describe it as the delta structure of volatility, as it is quoted by delta.

System Arbitrage. A practice whereby traders exploit model mis-pricing to book unrealised mark-to-model revaluation 'profits'. For example, a trader buys (sells) an exotic option in the market that is over- (under-) priced by their model. Buying (selling) in the market at a price below (above) the model price produces an immediate unrealised daily revaluation system 'profit' when the option is marked-to-model. It does not represent a real (realised) profit, only a real discrepancy between the model price and the market price. Taleb notes that "it is often easier to arbitrage one's accounting system than the market" (1997, p. 85).

Term Structure of Volatility. The vector of straddle volatilities for all option expiries T_i.

Theoretical Value. The value of vanilla and exotic options derived under the BSM paradigm.

Theta. The change in value of an option given a small reduction in the option's time to maturity. i.e. $\theta = \frac{\partial G}{\partial \tau}$.

Vanna. Also known as ∂delta/∂vol and ∂vega/∂spot. Vanna is the change in option delta given a small change in the option's volatility. i.e. $\Delta_\sigma = \frac{\partial \Delta}{\partial \sigma} = \frac{\partial \Phi}{\partial S} = \frac{\partial^2 G}{\partial \sigma \partial S}$.

Vega. The change in the value of the option given a small change in the option's volatility. For binary options, the relevant volatility is the zero-delta straddle volatility for the same expiry as the binary option. i.e. $\Phi = \frac{\partial G}{\partial \sigma}$.

Vega Neutral Butterfly. A liquid, commoditised European vanilla option strategy consisting of a long (short) strangle and a short (long) straddle, weighted such that net vega and delta is zero, that is, $FV_{strangle} = \frac{Vega_{straddle}}{Vega_{strangle}} \times FV_{straddle}$. In this thesis, the strangle consists of two OTM options with delta equal to ±0.10.

Volatility. The most common reference to volatility is as a constituent element of the volatility surface. Whilst colloquially referred to as volatility or vol for short, its full name in an FX option context is 'Garman-Kohlhagen volatility'. Garman-Kohlhagen volatilities are volatilities of the forward, not volatilities of the spot. In the theoretical value framework, volatilities of the forward and volatilities of the spot are the same, because domestic and foreign interest rates are constant. However, in a market price framework where interest rates are not constant, volatilities of the forward and volatilities of the spot are different. In the interbank FX option market, volatilities of the forward are directly traded between price-makers as pseudo prices, not implied from dollar prices. Nevertheless, consistent with the interbank FX option market, volatility and implied volatility is used interchangeably in this thesis.

Volga. Also known as ∂vega/∂vol, vomma and vol of vol. Volga is the change in the option vega given a small change in the option's volatility. i.e. $\Phi_\sigma = \frac{\partial \Phi}{\partial \sigma} = \frac{\partial^2 G}{\partial \sigma^2}$.