coremont | Quantitative Analysis

Equity Implied Volatility Dynamics

September 2020

Mark Gibbs
mark.gibbs@coremont.com

Contents

1. Overview **3**

2. Conclusions **4**

3. Definitions & Dynamics **5**

4. Data Coverage **5**
4.1. Assets 5
4.2. Strikes & Maturities 5
4.3. Vol Data Source 8

5. Metrics **9**
5.1. Measure#1 : P&L Variance Reduction 9
5.2. Measure#2 : Vega P&L / Spot Independence 9

6. Results **10**
6.1. Example : Eurostoxx 50 EUR 10
6.2. Results : All Indices 15
6.3. Results : Filtered 18
6.3.1. Filter: By Absolute Return 18
6.3.2. Filter: By Direction 18

7. Analysis Limitations **18**

A. Dynamic Definitions **20**
A.1. Sticky Strike 20
A.2. Sticky Delta 20
A.3. Sticky Surface 20

Equity Implied Volatility Dynamics

1. Overview

The failure of the Black-Scholes model to fit markets in vanilla options has led to the specification of an implied volatility for each option. This, in turn, has led to the creation of implied volatility surfaces – interpolations and extrapolations of volatilities across all strikes and expiries.

However, even these complex data structures have not been enough to describe the variation of market prices and practitioners have been required to develop dynamics for these surfaces, i.e., a description of how the implied volatilities move as a function of the state variables of the market (see e.g. [1], [2]). Most prominent has been the dynamics of the implied volatilities as a function of the underlying asset price – the so-called vol/spot dynamics (see e.g. [3]).

In this empirical analysis, we will consider the realized variation in Equity implied volatilities between one trading day and the next and see how well the most common vol/spot dynamics describe these moves. We will consider these dynamics across a range of major Equity indices both in normal and in volatile markets. Specifically, we will focus on the daily variation in implied volatilities in two periods - from 1st Nov 2019 to 25th Feb 2020 and from 26th Feb 2020 to 6th July 2020. The first period represents a period of well balanced markets and relatively low realized volatility ("Low Vol"). The second period represents a period of high realized volatility ("High Vol").

Within Equity derivatives, there has been a long reliance on Sticky Strike dynamics. This analysis aims to test how well served we are by this old friend and whether two other dynamics, Sticky Delta and Sticky Surface, might be more appropriate. We will also examine the variation in our dynamics across the surface – by strike and by maturity. We will consider whether the performance of a dynamic is influenced by the direction of the market moves or their size. Performance will be measured in two ways:

- Measure#1 : The ability to minimize the variance of P&L due to movements in the implied vol surface. This approach could be viewed as most appropriate for traders who wish to minimize overall risk.
- Measure#2 : The ability to make our volatility P&L independent of movements in the underlying asset. Our dynamic may not then necessarily minimize the variance of P&L. This approach could be viewed as most appropriate for traders who wish to take on a specific risk (e.g. vega) but want to be insulated from others (e.g. spot).

2. Conclusions

In the following sections, we present our analysis on the three vol/spot dynamics. From the details given in Sections 3 to 7, we make the following conclusions:

- Sticky Surface gives the best performance across almost all indices for Short (1w to 3m) and Medium (6m to 2y) maturities both in the Low and High Vol periods. This is true using either measures of performance. Sticky Strike comes a close second - typically within the standard deviation of the sample means. This is independent of the size of the absolute spot move considered or its direction.
- For non-Asia/Pacific (APAC) indices, Sticky Delta proves a poor choice of dynamic even for Long (3y and beyond) maturities.
- For APAC indices, Sticky Delta is the best performer for Long maturities but Sticky Surface comes a close second - within the standard deviation of the sample means. However, even for the Low Vol period, P&L variance reduction is poor for all three dynamics.
- All three vol/spot dynamics do a poor job of minimizing P&L variation across the High Vol period. Performance on reducing P&L / spot correlation is reasonable on some assets (e.g. Nasdaq 100) but is generally worse than performance across the Low Vol period.
- Filtering by direction of movement does yield some interesting results on APAC indices but not on others. Filtering by size of move does not yield any obvious conclusions.

Overall, we see the need both for dynamics to be specific to a given underlying and to be a function of maturity. For non-APAC underlyings, we would recommend Sticky Surface vol/spot dynamics for all maturities. For APAC underlyings, we would recommend using Sticky Surface vol/spot dynamics for maturities up to 5y and Sticky Delta for longer. However there is clearly scope for a new vol/spot dynamic for APAC indices if the goal is to minimize P&L variance successfully. In periods of high volatility and high vol-of-vol, vol/spot dynamics are only one small part of hedging to minimize P&L variance and need to be supplemented with option hedges.

Our choice of vol/spot dynamic can be incorporated into our risk management in three main ways. Firstly, we can incorporate it into our scenarios to give more realistic estimations of our market data under perturbation of our state variables. Secondly, we can incorporate it into our Greeks. Doing this, we need to be comfortable with hedging volatility risk using stock/index futures. Thirdly, we can incorporate it into our P&L Explain to decide on how we partition our P&L between spot and vol factors. These decisions have a large subjective component to them but it is hoped that this analysis provides tools with which to assess existing and new vol/spot dynamics and show where maturity-based and regional flexibility is required.

3. Definitions & Dynamics

Let's begin by defining implied volatility as the volatility entered into the Black-Scholes model that matches the market price.

$$C_{Mkt}\left(K,T|S(t),t\right) = C_{BS}\left(K,T|t,F\left(S(t),t\right),\sigma\left(K,T|S(t),t\right),DF_{rf}(t,T)\right)$$

where

- $C_{Mkt}\left(K,T|S(t),t\right)$ is the market price of the vanilla option of strike K and maturity T as of time t and underlying asset price, $S(t)$.
- $C_{BS}\left(K,T|t,F,\sigma,DF\right)$ is the Black-Scholes price of a vanilla option of strike K and maturity T as of time t using forward F, volatility σ and risk-free discount factor DF.
- $\sigma\left(K,T|S(t),t\right)$ is the implied volatility for strike K and maturity T as of time t and underlying asset price, $S(t)$.

We can now define the three dynamics that we want to analyse. These are described in Table 3.1 and in Figure 3.2. More detailed definitions can be found in Appendix A.

4. Data Coverage

We now want to assess the reasonableness of these three dynamics. To do this, we will consider the daily variation in implied volatilities in two periods - from 1st Nov 2019 to 25th Feb 2020 (referred to as the "Low Vol" period) and from 26th Feb 2020 to 6th July 2020 (referred to as the "High Vol" period). The division between the two periods was decided by finding the point where the level of the VIX had deviated more than six standard deviations from its rolling one month average for three or more consecutive days.

4.1. Assets

We consider the volatilities of the indices listed in Table 4.1. These were picked to give a reasonable global coverage and to only include indices where sufficient liquidity exists within options markets to make it possible to discern daily changes in implied volatilities.

4.2. Strikes & Maturities

We consider a range of different tenors when attempting to observe the vol/spot dynamics : 1w, 1m, 2m, 3m, 6m, 9m 1y, 2y, 3y, 4y, 5y. Beyond 5y maturity, we do not believe there are enough sources of information to achieve much beyond

	StickyStrike
Description	Vol is constant by absolute strike
Formulation	$\sigma_{SSTR}\left(K,T\middle\vert S^{new}(t),t\right)=\sigma\left(K,T\middle\vert S(t),t\right)$
ATM Behavior	ATM vol moves along the skew
Justification	Consistent with Black-Scholes risks. Historically observed in Equity markets.
	StickyDelta
Description	Vol is constant by relative strike / Delta
Formulation	$\sigma_{SDEL}\left(K,T\middle\vert S^{new}(t),t\right)=\sigma\left(K\times\frac{F(S(t),t)}{F(S^{new}(t),t)},T\middle\vert S(t),t\right)$
ATM Behavior	ATM vol is constant
Justification	Over longer periods of time, we expect the implied volatility surface to be independent of market level. Historically observed in FX markets
	StickySurface
Description	Surface shape is constant by normalized strike, $z(.)$, while overall level follows the skew
Formulation	$\sigma_{SSUR}\left(K,T\middle\vert S^{new}(t),t\right) = \sigma\left(F\left(S^{new}(t),t\right),T\middle\vert S(t),t\right) + \Delta\sigma_t(z(K,S^{new}(t),t),T\middle\vert S(t),t)$
ATM Behavior	ATM Vol moves along the skew
Justification	Describing the shape of a surface by normalized strike is intuitive and matches many stochastic processes.

Table 3.1: Main characteristics of the three vol/spot dynamics under consideration in this analysis. *ATM Behavior* refers to the behavior of the strike for the prevailing level of the spot, i.e., the ATM strike moves as the spot moves. More detailed definitions of the dynamics can be found in Appendix A.

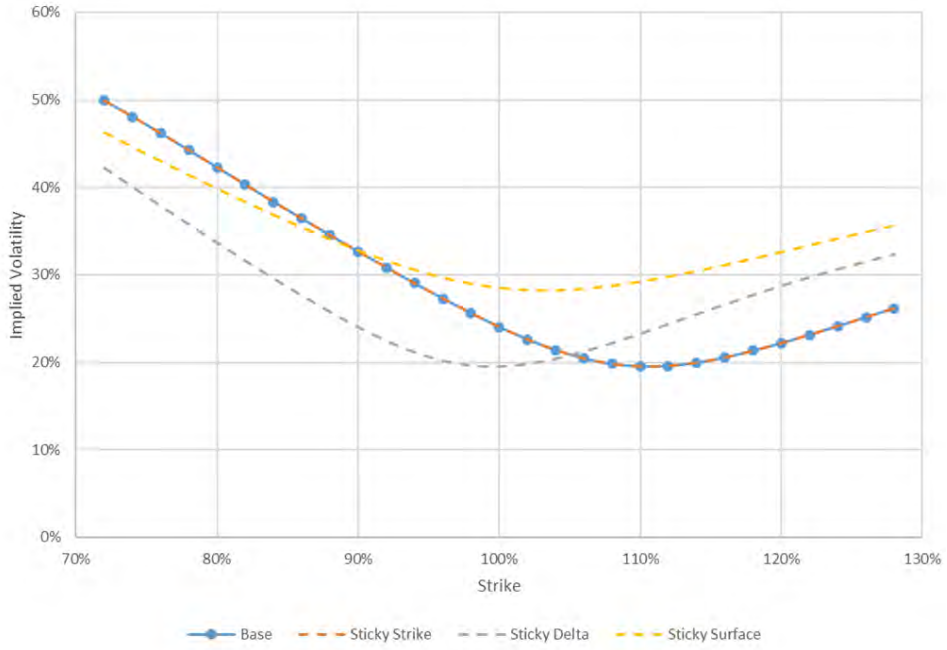

(a) 10% downward spot move

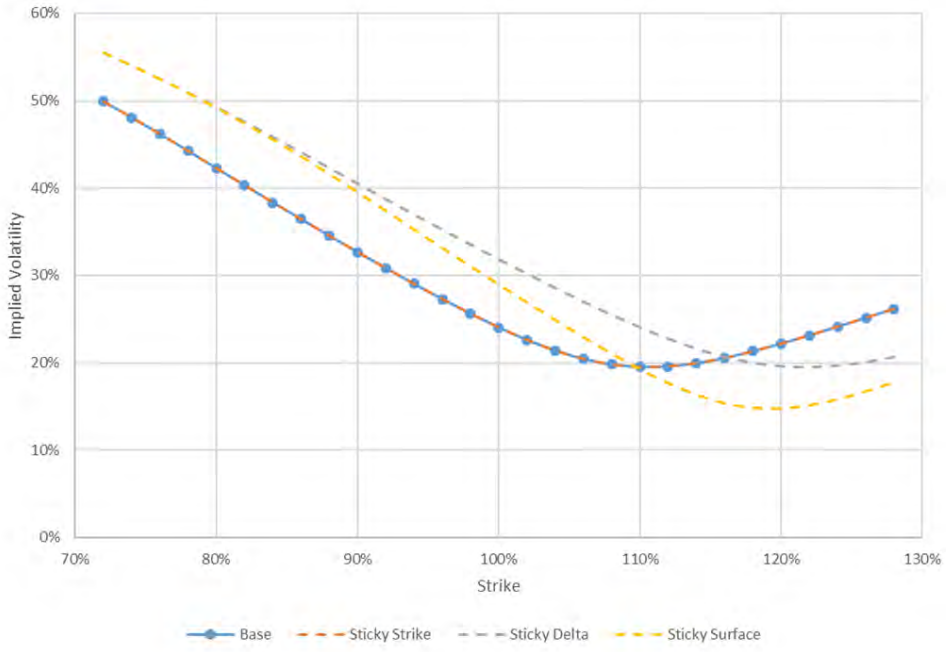

(b) 10% upward spot move

Table 3.2: Impact of a 10% down and up spot move on the 1 month implied volatility surface under three different vol/spot dynamics. Strikes are quoted as percentages of the unbumped forward to the expiry date. Based on data for SX5E as of 6th July 2020.

Index Name	Region
AEX	EMEA
ASX 200	APAC
CAC 40	EMEA
DAX 30	EMEA
EuroStoxx 50 EUR	EMEA
HSCEI	APAC
HSI	APAC
KOSPI 200	APAC
Nasdaq 100	NAM
Nikkei 225	APAC
OMX 30	EMEA
S&P500	NAM
SMI	EMEA
TOPIX	APAC

Table 4.1: List of all indices included within vol/spot dynamics analysis

reflecting the assumed dynamics of those extrapolating the more visible parts of the volatility market. We calculate results where possible for 7y and 10y but they should be viewed with caution.

We consider a range of strikes between -1.0 and +1.0 normalized strikes at each tenor. Hence the absolute spread of strikes is much smaller for shorter maturities than for longer maturities. One might comment that ± 1 normalized strike is not an especially deep in/out of the money strike.

4.3. Vol Data Source

The volatility information is extracted from the Coremont historical data store. The surfaces are calibrated from a range of different sources including listed option prices, broker option quotes, internal and external dividend projections and discount curves. All volatilities are quoted on a calendar Act/365 basis and, where necessary, are interpolated on the same basis. The volatilities and prices) are sampled at 4pm London local time (referred to as the LDN1600 mark). If the index in question has closed for the day (i.e., the APAC indices) then the market close and final calibration for that day will be used. As we are not examining covariances between assets in this analysis, we do not need contemporaneous data.

5. Metrics

So how are we going to measure the appropriateness of our vol/spot dynamics? On each day, t, we take the implied vol surface as of our LDN1600 mark. We then apply the dynamic to predict what surface we would expect on the following day, $t+1$, given that we know $S(t+1)$. For each strike and maturity, we then calculate the difference between the predicted implied volatility and the observed one from the $t+1$ implied vol surface,

$$X(K,T|t) = \sigma_X\left(K,T|S(t+1),t\right) - \sigma\left(K,T|S(t+1),t+1\right)$$

where $\sigma_X(.)$ is the implied vol surface evolved using the dynamic X. We then consider two different metrics to measure performance of the dynamics.

5.1. Measure#1 : P&L Variance Reduction

For each strike and maturity, we take the absolute value of $X(K,T|t)$. To prevent large absolute errors in prediction when we have large spot moves from dominating our data, we normalize the prediction error using the expected change in the volatility of the ATM strike, $K = F(S(t),t)$, from $S(t)$ to $S(t+1)$ using the day t implied vol surface. This gives us the relative error statistic

$$Err_X(K,T|t) = \left| \frac{X(K,T|t)}{\frac{\partial\sigma(K,T|S(t),t)}{\partial K}\Big|_{K=F(S(t),t)} \times \left(F\left(S(t+1),t\right) - F\left(S(t),t\right)\right)} \right|$$

We calculate these error statistics for every day in our two data periods, "Low Vol" and "High Vol" but filter out every day where the return was less than 25 basis points. This is to reduce the chance of the denominator of $Err_X(K,T|t)$ being very small and hence the statistics being dominated by noise.

We then consider the average error and its standard deviation across those days. We average further across strikes to create three statistics per maturity

- Put Error - the average of Err_X across all days for all normalized strikes, z_i, less than zero.
- ATM Error - the average of Err_X across all days for $z_i = 0$.
- Call Error - the average of Err_X across all days for all normalized strikes, z_i, greater than zero.

5.2. Measure#2 : Vega P&L / Spot Independence

For each strike and maturity, we calculate the Beta of the prediction error, $X(K,T|t)$, with respect to the log returns of the underlying index,

$$\beta_X(K,T) = \frac{Covar\left[X(K,T|t_i), \log\left(S(t_i)/S(t_{i-1})\right)\right]}{Var\left[\log\left(S(t_i)/S(t_{i-1})\right)\right]}$$

We calculate these betas using the prediction errors and index levels across all days in our two data periods, "Low Vol" and "High Vol". We then normalize by the gradient of the implied vol surface at the forward strike with respect to the forward-moneyness,

$$Err_X^{\beta}(K,T) = \frac{\beta_X(K,T)}{\left|\left.\frac{\partial\sigma(K,T|S(t),t)}{\partial K}\right|_{K=F(S(t),t)} \times F(S(t),t)\right|}$$

We average the beta error statistics across strikes to create three statistics per maturity

- Put Error - the average of Err_X^{β} across all normalized strikes, z_i, less than zero.
- ATM Error - the average of Err_X^{β} for $z_i = 0$.
- Call Error - the average of Err_X^{β} across all normalized strikes, z_i, greater than zero.

6. Results

6.1. Example : Eurostoxx 50 EUR

Let us first consider a specific index in order to demonstrate various features that are common across many underlyings before we single out certain indices for special attention.

The error statistics for the Eurostoxx50 Eur and their standard deviations can been see in Figures 6.1 and 6.2. Any relative error statistic under 100% suggests that the dynamic is making a passable job of predicting the movement in the surface and, hence, reducing the P&L volatility. A number above 100% suggests the opposite - that the error in the predicted vol is more than the predicted movement of the ATM vol.

The beta statistics for the Eurostoxx50 Eur and their standard deviations can been see in Figures 6.3 and 6.4. Any beta between $\pm 50\%$ suggests that the dynamic is making a passable job of making the P&L independent from spot movements. Any beta with a larger absolute value is performing poorly.

Various observations from the Eurostoxx 50 Eur results

- The performance of the dynamics using both measures is significantly better for the Low Vol period than the High Vol period.

	Strike			Moneyness			Surface		
Tenor	Put	ATM	Call	Put	ATM	Call	Put	ATM	Call
1w	99%	98%	99%	186%	176%	166%	111%	99%	88%
1m	78%	76%	74%	154%	149%	135%	83%	74%	64%
2m	69%	67%	63%	145%	141%	127%	74%	65%	54%
3m	64%	60%	55%	141%	137%	123%	68%	58%	49%
6m	55%	50%	46%	126%	123%	109%	58%	49%	42%
9m	56%	56%	54%	123%	123%	111%	60%	55%	51%
1y	64%	67%	63%	126%	128%	118%	66%	66%	62%
2y	85%	87%	83%	134%	141%	123%	86%	87%	84%
3y	77%	71%	71%	104%	102%	101%	70%	71%	75%
4y	105%	92%	89%	103%	101%	106%	86%	92%	98%
5y	119%	100%	99%	113%	111%	110%	94%	101%	109%
7y	136%	103%	101%	110%	108%	110%	100%	104%	114%
10y	108%	93%	87%	109%	101%	92%	101%	93%	103%

(a) Low Vol data period

	Strike			Moneyness			Surface		
Tenor	Put	ATM	Call	Put	ATM	Call	Put	ATM	Call
1w	199%	202%	202%	240%	242%	238%	198%	200%	197%
1m	160%	161%	152%	211%	214%	200%	160%	160%	148%
2m	151%	147%	139%	211%	209%	196%	153%	146%	133%
3m	151%	146%	142%	211%	207%	197%	154%	146%	135%
6m	140%	132%	128%	203%	197%	184%	144%	131%	120%
9m	136%	127%	122%	199%	193%	180%	140%	126%	114%
1y	135%	126%	120%	199%	192%	178%	140%	125%	112%
2y	122%	120%	117%	185%	180%	167%	128%	119%	108%
3y	132%	130%	137%	190%	186%	183%	137%	128%	128%
4y	143%	140%	146%	197%	191%	184%	147%	138%	134%
5y	156%	155%	165%	210%	205%	203%	159%	153%	150%
7y	186%	183%	187%	237%	235%	220%	188%	182%	174%
10y	258%	246%	237%	310%	299%	272%	257%	244%	233%

(b) High Vol data period

Table 6.1: Error statistics for the three vol/spot dynamics using the two data periods for EuroStoxx 50 EUR. The number quoted is the average of Err_X over the days in the period and further averaged across strikes for negative, zero and positive normalized strikes as described in Section 5.

	Strike			Moneyness			Surface		
Tenor	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**
1w	11%	10%	10%	11%	11%	10%	10%	10%	9%
1m	7%	7%	7%	9%	9%	8%	7%	7%	6%
2m	6%	6%	6%	8%	8%	8%	6%	6%	5%
3m	6%	6%	6%	8%	8%	7%	6%	6%	5%
6m	5%	5%	5%	7%	7%	6%	5%	5%	4%
9m	6%	7%	8%	7%	9%	8%	6%	7%	7%
1y	8%	13%	14%	10%	14%	15%	8%	13%	14%
2y	13%	16%	17%	16%	18%	19%	14%	16%	17%
3y	7%	7%	8%	10%	10%	10%	7%	7%	8%
4y	11%	10%	10%	9%	10%	11%	9%	10%	11%
5y	13%	13%	12%	12%	12%	12%	12%	13%	13%
7y	15%	13%	11%	13%	12%	12%	13%	13%	11%
10y	14%	12%	10%	13%	11%	10%	14%	12%	11%

(a) Low Vol data period

	Strike			Moneyness			Surface		
Tenor	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**
1w	32%	31%	30%	32%	32%	31%	32%	32%	30%
1m	25%	27%	24%	25%	27%	25%	25%	27%	24%
2m	24%	23%	21%	24%	23%	22%	24%	23%	21%
3m	23%	22%	21%	23%	22%	22%	23%	22%	21%
6m	22%	21%	19%	22%	21%	20%	22%	21%	19%
9m	19%	17%	16%	19%	18%	17%	19%	17%	16%
1y	19%	17%	16%	20%	19%	17%	19%	17%	16%
2y	17%	16%	13%	19%	18%	15%	17%	16%	13%
3y	17%	16%	16%	19%	18%	17%	17%	16%	15%
4y	18%	17%	16%	21%	19%	18%	19%	17%	16%
5y	22%	20%	20%	25%	23%	21%	23%	20%	18%
7y	33%	30%	28%	37%	33%	30%	33%	30%	27%
10y	69%	62%	56%	77%	69%	59%	69%	62%	52%

(b) High Vol data period

Table 6.2: Standard deviation of the relative error statistics, Err_X, for the three vol/spot dynamics using both data periods for EuroStoxx 50 EUR. The number quoted is the estimated standard deviation of the mean and is based on around 55 / 85 data points for the Low / High period respectively.

	Strike			Delta			Surface		
Tenor	Put	ATM	Call	Put	ATM	Call	Put	ATM	Call
1w	42%	46%	50%	137%	132%	124%	48%	39%	29%
1m	53%	52%	54%	154%	148%	132%	58%	43%	24%
2m	47%	44%	45%	146%	140%	124%	52%	37%	16%
3m	41%	38%	39%	140%	136%	119%	47%	32%	11%
6m	28%	25%	29%	126%	124%	110%	35%	21%	1%
9m	17%	14%	18%	116%	113%	100%	25%	11%	-8%
1y	11%	8%	13%	112%	107%	94%	21%	5%	-14%
2y	-10%	-9%	-1%	99%	90%	78%	7%	-12%	-29%
3y	-41%	-22%	4%	79%	75%	68%	-12%	-26%	-35%
4y	-61%	-32%	3%	63%	58%	52%	-22%	-37%	-44%
5y	-70%	-37%	0%	50%	42%	33%	-26%	-41%	-51%
7y	-51%	-29%	-4%	30%	24%	12%	-21%	-32%	-44%
10y	-16%	-19%	-10%	18%	11%	-3%	-12%	-21%	-34%

(a) Low Vol data period

	Strike			Delta			Surface		
Tenor	Put	ATM	Call	Put	ATM	Call	Put	ATM	Call
1w	36%	37%	45%	143%	141%	130%	42%	37%	26%
1m	40%	31%	36%	140%	139%	127%	46%	35%	20%
2m	54%	45%	46%	151%	151%	140%	59%	47%	31%
3m	59%	51%	50%	155%	153%	144%	64%	51%	34%
6m	61%	52%	56%	156%	153%	144%	66%	51%	34%
9m	63%	52%	54%	158%	153%	141%	69%	51%	31%
1y	67%	59%	62%	165%	160%	148%	74%	56%	37%
2y	74%	71%	78%	179%	171%	158%	87%	67%	48%
3y	80%	82%	94%	192%	181%	167%	101%	77%	57%
4y	83%	90%	107%	203%	190%	173%	107%	83%	62%
5y	90%	96%	114%	213%	200%	170%	113%	88%	56%
7y	106%	99%	124%	222%	206%	181%	116%	91%	63%
10y	102%	91%	87%	213%	202%	148%	99%	82%	33%

(b) High Vol data period

Table 6.3: Beta statistics for the three vol/spot dynamics using the two data periods for EuroStoxx 50 EUR. The number quoted is the average of Err_X^β over the days in the period and further averaged across strikes for negative, zero and positive normalized strikes as described in Section 5.

	Strike			Delta			Surface		
Tenor	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**
1w	9%	8%	8%	8%	8%	8%	8%	8%	8%
1m	7%	6%	6%	6%	6%	6%	6%	6%	5%
2m	6%	6%	5%	6%	6%	5%	5%	5%	5%
3m	5%	5%	5%	5%	5%	5%	5%	5%	4%
6m	4%	4%	4%	5%	4%	4%	4%	4%	4%
9m	5%	5%	5%	5%	5%	5%	5%	5%	5%
1y	6%	7%	8%	6%	7%	8%	6%	7%	8%
2y	8%	9%	11%	8%	9%	11%	8%	9%	11%
3y	6%	6%	7%	7%	6%	7%	6%	6%	7%
4y	8%	7%	8%	8%	8%	8%	8%	7%	7%
5y	9%	8%	8%	9%	9%	9%	9%	8%	8%
7y	10%	8%	8%	9%	8%	8%	8%	8%	9%
10y	8%	6%	7%	7%	6%	7%	7%	7%	8%

(a) Low Vol data period

	Strike			Delta			Surface		
Tenor	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**	**Put**	**ATM**	**Call**
1w	13%	14%	16%	13%	14%	16%	13%	14%	16%
1m	10%	10%	10%	10%	10%	10%	9%	10%	10%
2m	10%	10%	10%	11%	10%	10%	10%	10%	10%
3m	11%	10%	11%	11%	11%	11%	11%	10%	11%
6m	10%	10%	10%	10%	10%	10%	10%	10%	10%
9m	10%	9%	9%	11%	10%	10%	11%	10%	10%
1y	11%	10%	10%	11%	10%	11%	11%	10%	10%
2y	13%	12%	12%	13%	13%	13%	13%	12%	12%
3y	16%	15%	15%	17%	16%	15%	16%	15%	15%
4y	19%	18%	17%	20%	19%	18%	19%	18%	17%
5y	23%	20%	20%	23%	22%	21%	22%	21%	20%
7y	26%	24%	23%	28%	26%	24%	26%	25%	23%
10y	30%	28%	27%	32%	31%	29%	30%	29%	27%

(b) High Vol data period

Table 6.4: Standard deviation of the beta error statistics, Err_X^β, for the three vol/spot dynamics using both data periods for EuroStoxx 50 EUR. The number quoted is the estimated standard deviation of the mean and is based on around 55 / 85 data points for the Low / High period respectively.

- The quality of predictions is significantly better for Sticky Strike and Sticky Surface than for Sticky Delta for all maturities under 5y in both periods but especially so for the shorter dated ones.
- For maturities greater than 5y (in the Low Vol period), the absolute error statistics for all three dynamics are of similar limited quality. For these maturities, the beta error statistics for Sticky Delta are superior to those for Sticky Strike and Sticky Surface. As mentioned before, we should be cautious of the longer dated marks as they may reflect a very limited range of market participants.
- Comparing the beta error statistics in the Call wing, Sticky Surface outperforms for maturities 1y and below and Sticky Strike outperforms maturities greater than 1y.

6.2. Results : All Indices

We can now broaden our focus to the other indices. A top level summary of the dynamic with the best performance based on the average relative error statistics, Err_X, and average beta error statistics, Err_X^β, associated with Short (1w to 3m), Medium (6m to 2y) and Long (3y and beyond) maturities can be seen in Figure 6.5 and Figure 6.6.

We make various observations from these tables.

- Across all assets and both data periods, Sticky Surface is the best performing of all the dynamics. However, the difference in performance between Sticky Surface and Sticky Strike is often within the standard deviation of the sample mean.
- Sticky Delta is only the best performing dynamic for some APAC underlyings for long maturities. Even then, its performance is fairly poor and Sticky Surface and Sticky Strike are only marginally worse.
- For only a sub-set of indices do the even the best performing dynamics do a reasonable job of predicting movements in the surface for the Low Vol period. For APAC indices, none of the dynamics do a particularly convincing job.
- For the Low Vol period, the dynamics are fairly successful at removing the correlation between P&L and spot movement. The dynamics perform less well for APAC indices although it is only the Nikkei 225 where performance is poor.
- For the High Vol period, none of the dynamics do a good job of predicting movements in the surface. The dynamics also do not achieve good decorrelation between P&L and spot movement.

Index	Short		Medium		Long	
AEX	Surface	102%	Surface	74%	Surface	83%
ASX 200	Surface	94%	Surface	87%	Surface	151%
CAC 40	Surface	81%	Surface	59%	Surface	57%
DAX 30	Surface	70%	Surface	55%	Strike	57%
HSCEI	Surface	184%	Surface	151%	Delta	143%
HSI	Surface	168%	Surface	129%	Delta	192%
KOSPI 200	Surface	114%	Surface	90%	Delta	134%
Nasdaq 100	Surface	70%	Strike	60%	Strike	87%
Nikkei 225	Surface	144%	Surface	108%	Delta	251%
OMX 30	Surface	82%	Surface	65%	Strike	53%
SMI	Surface	94%	Surface	76%	Strike	56%
S&P 500	Surface	66%	Surface	53%	Strike	82%
Eurostoxx 50 EUR	Surface	74%	Surface	64%	Surface	94%
TOPIX	Surface	114%	Surface	108%	Surface	174%

(a) Low Vol data period

Index	Short		Medium		Long	
AEX	Surface	138%	Surface	131%	Surface	133%
ASX 200	Surface	117%	Strike	127%	Strike	171%
CAC 40	Surface	126%	Surface	105%	Surface	99%
DAX 30	Surface	146%	Surface	152%	Surface	142%
HSCEI	Surface	153%	Surface	175%	Surface	285%
HSI	Surface	154%	Surface	167%	Strike	219%
KOSPI 200	Surface	152%	Surface	149%	Surface	248%
Nasdaq 100	Surface	104%	Surface	139%	Surface	177%
Nikkei 225	Surface	161%	Surface	133%	Surface	247%
OMX 30	Surface	133%	Surface	122%	Surface	108%
SMI	Surface	178%	Surface	172%	Surface	146%
S&P 500	Surface	113%	Surface	117%	Surface	134%
Eurostoxx 50 EUR	Surface	127%	Surface	113%	Surface	145%
TOPIX	Surface	173%	Surface	183%	Surface	228%

(b) High Vol data period

Table 6.5: A summary of the dynamic with the lowest relative error statistics, Err_X, for each index bucketed by Short (1w to 3m), Medium (6m to 2y) and Long (3y and beyond) maturities and bucketed across all strikes for both data periods. The number quoted is the average error statistic for the dynamic in question.

Index	Short		Medium		Long	
AEX	Surface	36%	Strike	12%	Surface	8%
ASX 200	Surface	18%	Surface	13%	Surface	11%
CAC 40	Surface	40%	Surface	14%	Surface	12%
DAX 30	Surface	30%	Surface	12%	Surface	2%
HSCEI	Strike	1%	Strike	4%	Strike	5%
HSI	Surface	42%	Surface	17%	Surface	33%
KOSPI 200	Surface	35%	Surface	14%	Surface	17%
Nasdaq 100	Surface	21%	Surface	11%	Strike	20%
Nikkei 225	Strike	85%	Surface	25%	Delta	54%
OMX 30	Surface	26%	Surface	8%	Strike	27%
SMI	Surface	35%	Surface	10%	Strike	19%
S&P 500	Surface	45%	Surface	23%	Strike	5%
Eurostoxx 50 EUR	Surface	39%	Strike	12%	Strike	32%
TOPIX	Surface	63%	Surface	41%	Surface	25%

(a) Low Vol data period

Index	Short		Medium		Long	
AEX	Surface	44%	Surface	70%	Surface	99%
ASX 200	Strike	6%	Strike	11%	Strike	53%
CAC 40	Surface	36%	Surface	56%	Surface	39%
DAX 30	Surface	39%	Surface	71%	Surface	97%
HSCEI	Strike	17%	Strike	13%	Strike	9%
HSI	Surface	100%	Surface	96%	Surface	95%
KOSPI 200	Strike	66%	Surface	72%	Surface	104%
Nasdaq 100	Surface	24%	Surface	31%	Surface	13%
Nikkei 225	Strike	78%	Surface	95%	Surface	191%
OMX 30	Surface	27%	Surface	16%	Strike	26%
SMI	Surface	62%	Surface	70%	Surface	55%
S&P 500	Surface	31%	Surface	30%	Strike	10%
Eurostoxx 50 EUR	Surface	42%	Surface	58%	Surface	85%
TOPIX	Surface	65%	Surface	97%	Surface	139%

(b) High Vol data period

Table 6.6: A summary of the dynamic with the lowest beta error statistics, Err_X^{β}, for each index bucketed by Short (1w to 3m), Medium (6m to 2y) and Long (3y and beyond) maturities and bucketed across all strikes for both data periods. The number quoted is the average beta error statistic for the dynamic in question.

6.3. Results : Filtered

Are there patterns within this data beneath the simple averages that we have taken so far? Let's consider two possible ways to divide the data.

6.3.1. Filter: By Absolute Return

Previously, we have filtered our data for daily returns greater than 25 bps. What happens if we raise that threshold? Does the performance of the dynamics materially change? If we look at the relative error statistics for the Low Vol period where only absolute daily returns of greater than 2% were included in the sample, this suggests that the dynamics do a better job of predicting changes in the implied volatility surface for larger moves in the spot. However we must take care in interpreting this as the filter will remove the data points where a small move in the spot significantly increases the Err_X denominator and hence creates large contributions to the averages. When considering the beta error statistics for the Low Vol period, we see that filtering the data makes no significant impact on the performance.

6.3.2. Filter: By Direction

Previously, we have been averaging across positive and negative returns. What if the performance of the dynamics depends on the direction of the market move? If we divide the data into positive and negative returns, we find that both the relative error statistics are similar for the two different groups - well within the standard deviation of the sample mean. The beta error statistics however show a significant difference for APAC indices. For example, for the Nikkei 225, the beta error statistic for maturities less than 1y is significantly positive for downward movements but much closer to zero for upward movements. This suggests that, when the spot moves down, the surface rises relative to a StickyStrike dynamic.

7. Analysis Limitations

There are various limitations to the analysis presented here

- We consider a relatively small set of dates. This was driven by a desire to contrast two different vol regimes (Low vs High) rather than taking an average of many years of data. It is a focus of future work to extend the scope to cover a wider period.
- We are considering vol/spot dynamics but, by comparing implied volatilities from different days, our vol/time dynamic choice is also implicitly included. For maturities longer than one month, this is a minor source of noise. However, the one week maturity may be significantly influenced by this. In this analysis, we have not attempted to quantify this impact.

- We divide our data into Low and High vol regimes based on the level of the VIX. It could be argued that the impact of Covid-19 was already being incorporated into APAC surface dynamics before the 26th Feb 2020 cut off date. For the sake of simplicity, we do not attempt to use a different cut-off date for different regions.
- The two vol regimes do not have contain the same number of business days. This is, in part, due to differing holiday calendars across the indices considered.
- We only consider indices in this analysis. The study could be expanded to include single stocks although care would need to be taken on variation around ex-dividend dates.
- We only considered ± 1 normalized strikes in this analysis. It is a focus of future work to consider the dynamics of the wings although, given the results of this analysis, there is a concern that it will be very hard to draw a clear conclusion from such noisy data.
- As mentioned in previous sections, care must be taken to avoid drawing conclusions driven by statistical features of our metrics rather than our data. Across all the analysis, we considered both relative and absolute error statistics even when only relative results were presented.
- We assume at several points that the data that we are analysing is Normally distributed. This may not be the case and, for example, the estimates of sample mean uncertainty may be imprecise.

A. Dynamic Definitions

A.1. Sticky Strike

Sticky Strike is the most commonly used vol/spot dynamic within Equity derivatives. This dynamics simply states that the implied volatility for a fixed absolute strike and absolute expiry date is constant as a function of spot, $S(t)$.

$$\sigma_{SSTR}\left(K,T|S^{new}(t),t\right)=\sigma\left(K,T|S(t),t\right)$$

A.2. Sticky Delta

Sticky Delta is where the implied volatility for a fixed relative strike and absolute expiry date is constant as a function of spot, $S(t)$.

$$\sigma_{SDEL}\left(K,T|S^{new}(t),t\right)=\sigma\left(K\times\frac{F\left(S(t),t\right)}{F\left(S^{new}(t),t\right)},T\middle|S(t),t\right)$$

This is equivalent to saying that the implied volatility for a fixed Black-Scholes Delta and absolute expiry date is constant.

A.3. Sticky Surface

Sticky Surface attempts to describe both the level and the shape of the implied vol surface. Let's consider dividing our surface as follows

$$\sigma_t(K,T|S(t),t)=\sigma_{ATM}(T|S(t),t)+\Delta\sigma_t(z(K,S(t),t),T|S(t),t)$$

where $z(.)$ is the normalized strike

$$z(K,S(t),t)=\frac{\log\left(\frac{K}{F(S(t),t)}\right)}{\sigma_{ATM}(T|S(t),t)\sqrt{T-t}}$$

Here we have an ATM implied vol backbone, $\sigma_{ATM}(T)$, and a shape which is defined with respect to the normalized strike, $\Delta\sigma_t(z)$. We can define a dynamic where the ATM implied vol follows the skew but the shape with respect to normalized strike stays the same.

$$\sigma_{SSUR}\left(K,T|S^{new}(t),t\right)=\sigma\left(F\left(S^{new}(t),t\right),T|S(t),t\right)+\Delta\sigma_t(z(K,S^{new}(t),t),T|S(t),t)$$

References

[1] R. Cont and J. da Fonseca, "Dynamics of implied volatility surfaces," *Quantitative Finance*, vol. 2, p. 45, 2002.

[2] R. Cont, V. Durrelmann, and J. da Fonseca, "Stochastic Models of Implied Volatility Surfaces," *Economic Notes*, vol. 31, 2002.

[3] T. Daglish, J. Hull, and W. Suo, "Volatility surfaces: theory, rules of thumb and empirical evidence," *Quantitative Finance*, vol. 7, p. 507, 2007.

IMPORTANT INFORMATION

CONFIDENTIAL – DO NOT COPY OR DISTRIBUTE

This document, the information contained herein and any attached materials (collectively the “document") have been provided specifically for the use of the intended recipient only and must be treated as proprietary and confidential. If you have received this document in error please delete it and notify the sender immediately. This document remains the property of Coremont, which reserves the right to require its return (together with any copies or extracts thereof) at any time upon request. It may not be passed on, nor reproduced in any form, in whole or in part, under any circumstances without express prior written consent from Coremont. Any unauthorised review, retransmission, dissemination, forwarding, printing, copying or other use of this document is prohibited. Coremont may be legally required to produce the document to regulatory authorities and others with legal rights to the information. Internet communications cannot be guaranteed to be secure or error free as information could be intercepted, changed corrupted, lost, arrive late or contain viruses. Any information in this document should be regarded strictly as guidance and for informational purposes only. Coremont does not provide legal, accounting or tax advice. Coremont and the sender make no warranties regarding the accuracy or completeness of the information in this document and it should not be relied upon for any purpose and is subject to change without notice. Information contained in this document is unaudited and Coremont accepts no liability for any errors or omissions in this document. Coremont and its representatives, members, officers and employees accept no responsibility for any losses suffered as a result of reliance on the information in this document or the reliability, accuracy, or completeness thereof. Coremont assumes no obligation to update any information contained in this document.

In this document, "Coremont" means Coremont LLP and its affiliates. Coremont LLP is a limited liability partnership authorised and regulated by the Financial Conduct Authority of the United Kingdom and registered in England & Wales (reg. no. OC420985).