



Unbiasing and robustifying implied volatility calibration in a cryptocurrency market with large bid-ask spreads and missing quotes

Mnacho ECHENIM* Emmanuel GOBET† Anne-Claire MAURICE‡

October 10, 2022

Abstract

We design a novel calibration procedure that is designed to handle the specific characteristics of options on cryptocurrency markets, namely large bid-ask spreads and the possibility of missing or incoherent prices in the considered data sets. We show that this calibration procedure is significantly more robust and accurate than the standard one based on trade and mid-prices.

Keywords: implied volatility, calibration, bid-ask spread, missing data, data augmentation

MSC2020: 91G20, 62G35, 62G09, 62D10, 62P05

JEL: G13, C80

1 Introduction

Context. European option prices are traditionally quoted in terms of the corresponding implied volatility, in order to have an adimensional measure of how expensive the option is, independently of the price magnitude of the underlying asset. Over the past two decades, a large number of papers (both from practitioners and academics) have focused on:

*Laboratoire d'Informatique de Grenoble (LIG), CNRS, Grenoble INP, UGA. 700 avenue Centrale, Domaine Universitaire, 38401 Saint Martin d'Hères, France. Email: MNACHO.ECHENIM@UNIV-GRENOBLE-ALPES.FR

†Centre de Mathématiques Appliquées (CMAP), CNRS, Ecole Polytechnique, Institut Polytechnique de Paris. Route de Saclay, 91128 Palaiseau Cedex, France. Email: EMMANUEL.GOBET@POLYTECHNIQUE.EDU

‡Kaiko - Quantitative Data. 2 rue de Choiseul 75002 Paris, France. Email: ANNE-CLAIRE.MAURICE@KAIKO.COM

- deriving some closed-form approximations of this price when the underlying asset is modeled using a specific volatility model (local, stochastic, local-stochastic, with jumps, etc...), see, e.g., [24], [32], [1], [14], [2], [19], [4], [33] among many others. For an overview, see the book of Gatheral [18]; other references include [3] or [15];
- understanding how to calibrate simplified parametric forms (such as SVI [17] or its extensions [7]) with the goal of designing robust optimization techniques, with good initialization guesses [8, 31, 7].

This is now a mature research area when applied to usual markets like equity/index, fixed income, fiat foreign exchange, or commodities. The aim of this work is to adapt these techniques to cryptocurrency markets, which have grown at a huge pace over the past years.

Crypto options history. Cryptocurrency markets have gone through an unprecedented period of growth, and the demand for sophisticated trading products has surged. Previously, options trading was limited to traditional financial markets, but there is an increasing adoption of options throughout cryptocurrency markets. Options trading for retail investors seems to align well with the speculative nature of retail crypto trading and as a speculative trading product, crypto options have grown at a similar pace as the broader crypto market. Speculation is not the only use case however, particularly for institutions, as options contracts also offer the ability to hedge against the risk of sudden price changes in crypto markets.

The exchange with the most market share of options trading is Deribit; the company claims an 80 to 90% market share of all options trading. Deribit was launched in 2016 and as a result certainly benefited from first-mover advantage in the market. Deribit option trading products have continued to become more and more complex, moving in line with the increased liquidity in futures markets as well as the increased demand for complexity in crypto trading products. This added complexity and continued stream of new products has allowed Deribit to maintain their market share, fending off competition from other exchanges.

The Chicago Mercantile Exchange launched regulated crypto options in 2020, and this marked an important milestone in the industry, paving the way for institutional adoption of crypto options. In response, Deribit launched daily traded options shortly after in order to satisfy the increased market demand for more complexity. 2022 was the first time an over-the-counter options trade was made by a major U.S bank, as Goldman Sachs executed a trade involving a non-deliverable (cash settled) bitcoin option. This transaction was evidence of a growing institutional presence and gives an insight into how the biggest players in traditional finance are increasingly utilizing crypto options today.




Figure 1: Implied volatility in the Raw SVI model (4.1) with the parameters $(a, b, \rho, m, \sigma) = (0.02, 0.2, -0.3, 0.05, 0.6)$.

On calibration. Traditionally, fitting a smile volatility (with time-to-expiry τ) model boils down to finding the optimal parameter χ that minimizes the Mean Square Error (MSE)

$$\sum_i (w(\tau, k_i; \chi) - w^{\text{Mkt}}(\tau, k_i))^2,$$

where:

- each k_i is a log-forward-moneyness at which market data for time-to-expiry τ is available;
- $k \mapsto w(\tau, k; \chi)$ is a parametrization of the total variance (the squared implied volatility multiplied by the time-to-expiry), such as, e.g., the arbitrage-free SVI parametrization presented in (4.1) and depicted in Figure 1 for a given set of parameters;
- $w^{\text{Mkt}}(\tau, k_i)$ is the market total variance at k_i .

The MSE is not the only loss criterion that practitioners tend to consider: the L_2 -norm can be replaced by the Huber loss [26], or simply the L_1 norm; the errors may be evaluated on the implied volatilities or the prices instead of total variances; weighted sums can also be considered, with weights that depend on the volume.

Statement of the problem. Our paper sheds light on an issue that is often ignored on traditional markets (where it is negligible), but has a significant impact in crypto markets. Usually the quantity $w^{\text{Mkt}}(\tau, k_i)$ is computed based on trade and mid-prices, the latter being an excellent approximation of the so-called *efficient price* when the bid-ask spread is small. However, it appears that the relative tick size (and thus the relative spread) for options on cryptos is roughly 10 to 50 times larger than for options on traditional markets (see Table 3). Hence the mid-price appears to be a very noisy

estimate of the efficient price, and as a consequence it may result in a higher error of the calibrated implied volatility. Besides, it is not uncommon for quotes to be incomplete (only the ask or the bid is available), and this differs with usual markets where assigned market-makers play the role of liquidity providers.

On the other hand, instead of using the mid-price information, one could restrict the considered data to trades. Unfortunately, at the time of writing this article, the option market liquidity is limited, which is why it is necessary to use the bid-ask information in the calibration scheme in order to have more stable results. As an illustration, over one week (from June 27 to July 3, 2022), we counted the number of BTC-USD option trades on the Deribit platform: there was an average of 30 trades per hour, all strikes combined, for the most liquid expiries (1 week and 1 month), around 5 trades per hour for the quarterly 3-month and 6-month expiries and that only 1 or 2 trades for the 9-month and 1-year expiries. The need for all trade and bid-ask prices is also reinforced because the time scale of the crypto market is shorter than that of traditional markets: indeed this is a market that is open 24 hours a day, 7 days a week. In addition, the volatility (resp. total variance) is around three (resp. nine) times higher than on equity/index markets, highlighting the need to update market prices more often, typically every 2 to 3 hours instead of once a day. The data that is used throughout the paper was provided by Kaiko¹, whose clients have shown a strong interest in obtaining calibrated implied volatilities at an hourly scale.

Our contributions. For such situations where tick sizes are relatively large (and thus spreads are also large), one may try to use the bid and ask prices, instead of only the mid-price. In a nutshell, how to adapt the calibration criterion written for mid-prices to account for bid and ask prices that may differ a lot from this mid-price? How to handle cases where one side of the quote is missing, discarding any attempt to use the naive mid-price? To the best of our knowledge, these questions are overlooked in the literature, which can be easily understandable in view of the usually small spreads and tick sizes of standard markets. Our contributions are the following.

- The market implied volatility is much more sensitive to large spreads when computed from ITM puts or from ITM calls, see Figure 4 for a precise statement. As a result, we filter call and put prices taking only OTM options.
- We prove that calibrating on the mid-price yields a systematic underestimation of the implied volatility parametrization in the wings (for small or large log-forward-moneynesses), see Theorem 3.1 for a precise statement.
- In order to remove the bias in the calibration process, we define a so-called *anchor price*, which is built from the bid and ask prices, and has a negligible bias (of order three w.r.t. the tick size) with respect to the efficient price. See Theorem 3.2 for a precise statement.

¹<https://www.kaiko.com/>

- In the case where there are missing or incoherent quotes, we design a data augmentation procedure based on a Beta distribution, the parameters of which are tuned adaptively depending on market data. This procedure can be used even when the considered quote is not corrupted, for the sake of better stability.
- We perform several tests on synthetic data, using the SVI model, and on real data on BTC-USD options; these experiments confirm the gain in accuracy and stability of the procedure. We observe quite systematically an improvement of calibration error by a factor between 25% and 30%.

Organization of the paper. In Section 2, for the sake of completeness and since this is a new emerging market, we provide a short description of the market for options on cryptocurrencies (cryptos for short), along with some comparisons with traditional markets. Notations are also defined in this section. In Section 3 we develop the analysis related to large spreads or missing quotes, and develop the proposed methodology. Numerical tests are presented in Section 4. Some technical results are postponed to the Appendix.

2 Market data

We describe the main exchanges for option-type derivatives, as well as some operating characteristics of the contracts. This is a description as of the time of writing, keeping in mind that this is a quickly-evolving market.

2.1 Market exchanges

As with decentralized finance, there are numerous option trading platforms, among which we can mention ² Binance, Bybit, Delta Exchange, Deribit, FTX, Huobi or OKX. The first platforms appeared in 2016, some of them are more recent. Most proposed options are European-style options, but Huobi gives access to American-style options. Almost all options are vanilla types, i.e. calls and puts; however, Huobi provides *Touch Options Double One-Touch* and *Touch Options Double No-Touch*, that are barrier options.

The underling asset is usually a currency pair crypto-fiat. The fiat currency (fiat for short) is often USD, although it can be replaced by USDC³ (or any another stablecoin), this is done for instance on Bybit. The usual cryptocurrencies used as underlyings are BTC or ETH (or more recently SOL). However, we anticipate that other cryptos will also be used in the future, as well as other fiats, such as Japanese Yen (JPY) or South Korean Won (KRW). Delta Exchange also offers

²<https://www.binance.com>, <https://www.bybit.com>, <https://www.delta.exchange>, <https://www.deribit.com>, <https://ftx.com>, <https://www.huobi.com>, <https://www.okx.com>

³USDC is a stablecoin pegged to USD. USDC has been developed by two companies, Circle et Coinbase, and it is integrated into Visa payments.

options on other cryptos, such as BNB, XRP, LINK, AVAX or MATIC. The payoff of the options can be written in units of crypto or units of fiat, which changes the pricing rules from a mathematical finance point of view [34, Chap. 4]. The settlement usually occurs in fiat. The global volume of option trading in billions USD is depicted in Figure 2.




Figure 2: Call and put weekly volumes, in billions USD, from December 2020 to April 2022. On the left, the underlying crypto is BTC; on the right, ETH. Source: <https://www.kaiko.com>.

2.2 Some comparative statistics

Despite the numerous option exchanges, Deribit claims to provide 80% to 90% of the global Open Interest. In what follows, we focus the discussion to Deribit contracts, which contain the most common features. We assume the reader is familiar with traditional markets, and we provide a few comparisons to highlight similarities and main differences with crypto-markets. As a benchmark, we consider options on SP500 on May 24th 2022, and we compare them with options on BTC-USD the same day.

Maturities/Expiries. For options on SP500, 43 maturities are listed: 31 in 2022, 9 in 2023, 1 in 2024, 1 in 2025 and 1 in 2026. On BTC-USD, only 8 expiries⁴ are available, from May 27 2022 until March 31 2023. The scarcity of expiries for cryptocurrency options is certainly a factor in increasing the liquidity of available contracts by concentrating trading on these contracts. On the other hand, this scarcity makes it difficult to construct a full implied volatility surface, because of the lack of available expiries in the term-structure of IVs.

Log-forward-moneyness. In Tables 1 and 2, we list the strikes and the related log-forward-moneynesses for the available quotes for options. Observe that the width of available log-forward-

⁴It appears that practitioners dealing with options on cryptos are more likely to refer to expiries rather than maturities.

moneynesses for both underlyings is quite similar, regardless of the expiry. The main difference comes from the center of the log-forward-moneyness range: it is biased to the left for SP500, showing the interest for options with small strikes (to help asset manager hedge themselves against market crashes), and it is biased to the right for BTC-USD, probably because the market has often been bullish.

Maturity	Strike (quotes) (min, max)	Log-forward-moneyness range (quotes)	Length range (quotes)
1 D (2022-05-25)	[1000, 5200]	[-0.5991, 0.1168]	0.7159
1 W (2022-06-01)	[1000, 5000]	[-0.5991, 0.0998]	0.6989
1 M (2022-06-24)	[1000, 5600]	[-0.5991, 0.1490]	0.7481
3.5 M (2022-09-16)	[200, 7400]	[-1.2980, 0.2701]	1.5681
7 M (2022-12-16)	[100, 7300]	[-1.5991, 0.2642]	1.8633
1 Y (2023-06-16)	[200, 7400]	[-1.2980, 0.2701]	1.5681
2 Y (2024-12-20)	[200, 9200]	[-1.2980, 0.3646]	1.6626

Table 1: SP500: available log-moneyness on quotes, at 2022 May 24th, for a few maturities

Expiry	Strike (quotes) (min, max)	Log-moneyness range (quotes)	Length range (quotes)
1 M (2022-06-24)	[10000, 75000]	[-0.4655, 0.4095]	0.875
3.5 M (2022-09-30)	[10000, 200000]	[-0.4655, 0.8355]	1.301
7 M (2022-12-30)	[10000, 300000]	[-0.4655, 1.0116]	1.4771
11M (2023-03-31)	[10000, 300000]	[-0.4655, 1.0116]	1.4771

Table 2: BTC-USD: available log-moneyness on quotes, at 2022 May 24th, for a few expiries.

Tick size. As a final comparison, we consider the tick size on option contracts, and compare it with different underlying assets, ranging from SP500 to EUR-USD, see Table 3. Compared to traditional markets, the relative tick size (defined as the ratio between the option tick size and the underlying asset value) is much larger for options on cryptos, by a factor ranging between 10 and 50. As a result, the relative spread must be significantly higher and hence, inferring the efficient price as the mid-price entails larger errors, that undeniably impact the accuracy and stability of the calibration.

For the sake of completeness, we show in Figure 3 the boxplots of spreads over a period of 1 week, for different expiries. Observe that extreme spreads can occur (over 100 ticks). For extreme

Option \ Exchange	CBOE	EUREX	EURONEXT	CME	Deribit
Underlying Name	SP500	Eurostoxx	AEX	EUR-USD	BTC-USD
Option tick size	0.05	0.10	0.01	0.00005	0.0005
Underlying Value (close) on 2022, May 24th	3 941.48	3 647.56	679.88	1.0737	1
Ratio in basis points (bps, i.e. 1E-4)	0.127	0.274	0.147	0.466	5

Table 3: Comparison of tick sizes for option prices written on various underlyings on different exchanges. On Deribit, the tick size is expressed in BTC.

log-moneynesses, the spread is usually 1 or 2 ticks for the shortest expiry, and up to 10 ticks for longer expiries; at the money, the spreads are often of order 5 to 15 ticks, depending on the expiries. This highlights the needs for calibration methodologies that are accurate and robust even when the spreads are (extremely) large. The calibration process we describe in the upcoming sections improves the one that simply takes mid-prices as inputs.

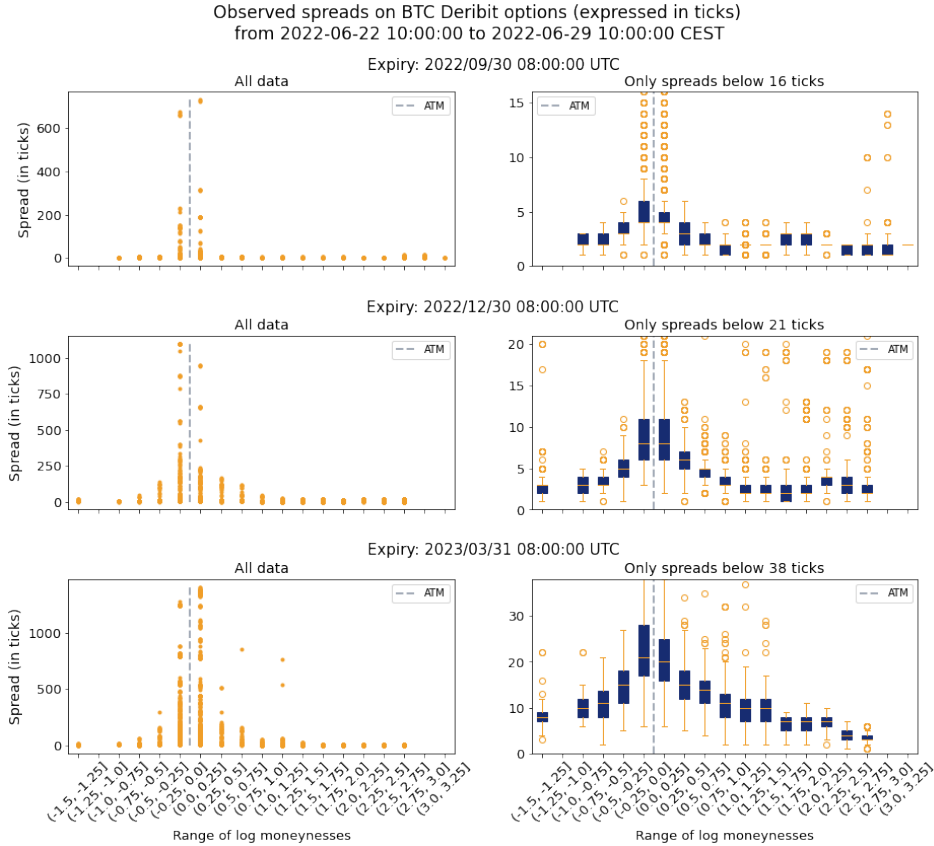


Figure 3: Boxplot of spreads (expressed in ticks) for data collected over a period of a week (2022-06-22 10:00 to 2022-06-28, 10:00 CEST). Option expiries are: 2022 Sept 30th (top), 2022 Dec 30th (middle), 2023 March 31th (bottom). The x -axis corresponds to observed log-moneynesses, the data points are those from BTC-USD options quoted on Deribit. On the left: all data; on the right: data with spreads below a threshold (respectively 16 ticks, 21 ticks, 39 ticks).

In addition, we note that on average, 10% of the quotes have no bid price and 5% of the bid prices are at 1 tick. This type of missing or out-of-arbitrage-bounds quotes prevent us from obtaining accurate and stable results using the naive mid-price in the calibration process. This problem is addressed with a data augmentation procedure in Section 3.4.

2.3 Quantitative modeling

We fix a few notations, used in all the sequel for the quantitative analysis.

- We consider two currencies: one for the "Fiat economy" with prices in units of F (typically USD), and one for the "Crypto economy" with prices in units of C (typically BTC, ETH).

- X stands for the crypto-fiat rate: 1 unit of \mathbf{C} is worth X units of \mathbf{F} .
- We denote the prices of zero-coupon bonds by $B_{t,T}^{\mathbf{F}}$ and $B_{t,T}^{\mathbf{C}}$, respectively in the fiat and crypto currencies.
- We denote by $C_t^{\mathbf{F}}(\Phi_T^{\mathbf{F}}, T)$ and $C_t^{\mathbf{C}}(\Phi_T^{\mathbf{C}}, T)$ the cash-prices in \mathbf{F} and \mathbf{C} currencies at time t , for some cashflows $\Phi_T^{\mathbf{F}}$ or $\Phi_T^{\mathbf{C}}$ delivered at time T .

The no-arbitrage condition between currencies \mathbf{C} and \mathbf{F} yields the relation

$$\frac{1}{X_t} C_t^{\mathbf{F}}(\Phi_T^{\mathbf{F}}, T) = C_t^{\mathbf{C}}\left(\frac{\Phi_T^{\mathbf{F}}}{X_T}, T\right), \quad (2.1)$$

meaning that it must be equivalent to compute the \mathbf{F} -price at time t of payoff $\Phi_T^{\mathbf{F}}$ at expiry T , then convert it in \mathbf{C} ; and to compute directly the \mathbf{C} -price at time t of payoff $\Phi_T^{\mathbf{F}}/X_T$ at expiry T .

2.3.1 Options

According to *Deribit - Options Documentation*⁵, call and put options are European-style options, written on the crypto-fiat rate X . The options are priced in \mathbf{C} .

Example 2.1 (from Deribit). *A trader buys a call option with a strike price of 10,000 USD for 0.05 BTC. Now they have the right to buy 1 BTC for 10,000 USD.*

- *At expiry, the BTC Index is at 12,500 USD, and the delivery price is 12,500 USD.*
In this case, the option is settled for 2,500 USD per 1 BTC. At expiry, the trader's account is credited with 0.2 BTC (2,500/12,500), and the seller's account is debited with 0.2 BTC. The initial purchase price was 0.05 BTC; therefore, the trader's profit is 0.15 BTC.
- *Any call option with an exercise price (strike price) above 12,500 USD will expire worthless.*

This example shows that the payoff (in \mathbf{C}) of a call option on the exchange rate \mathbf{C} - \mathbf{F} with strike K is therefore

$$\frac{(X_T - K)_+}{X_T}.$$

In view of (2.1), the Deribit call price is

$$\text{Call}_t^{\text{Der.}}(T, K) = C_t^{\mathbf{C}}\left(\frac{(X_T - K)_+}{X_T}, T\right) = \frac{1}{X_t} C_t^{\mathbf{F}}((X_T - K)_+, T). \quad (2.2)$$

⁵<https://legacy.deribit.com/pages/docs/options>

2.3.2 Futures

We refer to the *Deribit - Futures Documentation*⁶ to describe its so-called Future contract. Actually, Deribit is not a usual central clearing counterparty (CCP) because there are no daily margin calls. Hence, a Future in the Deribit terminology is not a Future in the traditional sense and the pricing rule must be clarified.

Expirations take place on Fridays at 08:00 UTC for weekly Futures, and on the last Friday of the month for monthly and quarterly Futures. In all cases there is a Cash settlement in \mathbb{C} . The value of the Deribit Future $\text{Fut}_t^{\text{Der.},T}$ at time t for expiry T is 1 USD per Index Point, with contract size USD 10 (on BTC) and USD 1 (on ETH); let us denote generically this size by \mathbb{S} .

Example 2.2 (adapted from Deribit). *A trader buys 100 futures contracts on the BTC index, at $\text{Fut}_t^{\text{Der.},T} = 10,000$ USD per BTC. The trader is now long 1,000 USD worth of BTC with a price of 10,000 USD (100 contracts \times 10 USD = 1,000 USD).*

- *At inception, the trader does not pay anything, but is committed to paying 1,000 USD at expiry; this amount is worth $1,000 B_{t,T}^{\mathbb{F}}$ USD at inception.*
- *At expiry, the trader receives 1,000/10,000 BTC, because of the futures contract.*
- *They can sell this amount in BTC at the expiry rate X_T ; for instance if $X_T = 12,000$ USD, then this amount is worth $\frac{1,000}{10,000} \times 12,000 = 1200$ USD.*

This example (Reverse Cash-and-Carry Arbitrage) shows that the payoff of this strategy at expiry is

$$\frac{100 \mathbb{S}}{\text{Fut}_t^{\text{Der.}}} X_T \text{ USD.}$$

Equating the \mathbb{F} -values, we get

$$100 \mathbb{S} B_{t,T}^{\mathbb{F}} = C_t^{\mathbb{F}} \left(\frac{100 \mathbb{S}}{\text{Fut}_t^{\text{Der.},T}} X_T, T \right) \iff \text{Fut}_t^{\text{Der.},T} = \frac{C_t^{\mathbb{F}}(X_T, T)}{B_{t,T}^{\mathbb{F}}}.$$

This shows that the Deribit Future contract is like an *OTC-forward contract*. Invoking again the no-arbitrage relation (2.1), we get

$$\text{Fut}_t^{\text{Der.},T} = \frac{X_t C_t^{\mathbb{C}}(1, T)}{B_{t,T}^{\mathbb{F}}} = X_t \frac{B_{t,T}^{\mathbb{C}}}{B_{t,T}^{\mathbb{F}}}.$$

This relationship between forward and spot exchange rates on the one hand, and interest rate products in \mathbb{C} and \mathbb{F} economies on the other hand, is known as Interest Rate Parity (IRP) relationship [34, p.153].

⁶<https://legacy.deribit.com/pages/docs/futures>

2.3.3 Implied volatility

We are now ready to define the implied volatility of the Deribit call price with expiry T and strike K . We first recall the application of the Black formula [21], under the assumption that the *forward* $C_t^F(X_T, T)/B_{t,T}^F = \text{Fut}_t^{\text{Der.}}$ admits a deterministic volatility σ_u^{Fut} : the call price in F is then

$$\text{Call}_t^F(T, K) = B_{t,T}^F \left(\text{Fut}_t^{\text{Der.}} \mathcal{N}(d_+) - K \mathcal{N}(d_-) \right), \quad (2.3)$$

$$\text{with} \quad d_{\pm} = \frac{1}{\sqrt{\int_t^T |\sigma_u^{\text{Fut}}|^2 du}} \log \left(\frac{\text{Fut}_t^{\text{Der.}}}{K} \right) \pm \frac{1}{2} \sqrt{\int_t^T |\sigma_u^{\text{Fut}}|^2 du}.$$

This is simply an extension of the Garman-Kohlhagen formula [16] with possibly stochastic interest rates. Combining the above and (2.2), we define the implied volatility IV of the Deribit call price as the constant parameter in (2.3) that matches market prices.

Definition 2.3. *The implied volatility IV at time t for expiry T and strike K is given implicitly by*

$$\frac{X_t \text{Call}_t^{\text{Der.}}(T, K)}{B_{t,T}^F} \stackrel{\text{def}}{=} \text{Fut}_t^{\text{Der.,}T} \mathcal{N}(d_+) - K \mathcal{N}(d_-) \quad (2.4)$$

$$\text{with} \quad d_{\pm} = \frac{1}{\text{IV} \sqrt{\tau}} \log \left(\frac{\text{Fut}_t^{\text{Der.,}T}}{K} \right) \pm \frac{1}{2} \text{IV} \sqrt{\tau}, \quad \tau = T - t.$$

The implied volatility can be retrieved similarly from put prices, with obvious adjustments to the formulas.

3 Improving calibration

In the following, we assume that calibration is performed at the current time, i.e., the time t is fixed at 0, and we omit to indicate this time dependency in the upcoming notations. To emphasize the dependency of the implied volatility (2.4) w.r.t. expiry and strike, we specialize notations:

- $w(\tau, k) \stackrel{\text{def}}{=} \tau \cdot \text{IV}^2(\tau, k)$ for the implied total variance at time 0;
- $k \stackrel{\text{def}}{=} \log\left(\frac{K}{\text{Fut}_0^{\text{Der.,}T}}\right)$ for the log-forward-moneyness;
- $\tau \stackrel{\text{def}}{=} T - t = T$ for the time-to-expiry;
- whenever the dependency w.r.t. the price P is important, we will write $\text{IV}(P, \tau, k)$ instead of $\text{IV}(\tau, k)$.

The bid and ask prices for log-forward-moneyness k and time-to-expiry τ are respectively denoted by $\text{Bid}(\tau, k)$ and $\text{Ask}(\tau, k)$. We write

$$\text{md}(\tau, k) \stackrel{\text{def}}{=} \frac{1}{2} \left(\text{Bid}(\tau, k) + \text{Ask}(\tau, k) \right) \quad \text{and} \quad \text{sp}(\tau, k) \stackrel{\text{def}}{=} \text{Ask}(\tau, k) - \text{Bid}(\tau, k),$$

for the mid-price and the spread respectively. The notation \mathbf{ts} stands for the tick size value on options.

3.1 Filtering OTM options in the data set

Thanks to the call-put parity relationship, if we assume a zero spread on the call and put prices, then the implied volatility is the same for both products, provided they have the same expiry and strike. However, when there is a non-zero spread, the implied volatilities will differ and we investigate below the way the choice of the option type (put or call) impacts the computation of implied volatilities. The larger the tick size, the larger the spread, and the larger the impact on the computed implied volatility.

The impact analysis can be visualized using a sensitivity indicator, when the price is bumped by the relative error ε , i.e.

$$\text{Initial price} \cdot (1 + \varepsilon) = \text{Black-Scholes price}(\text{IV}^\varepsilon(\tau, k)),$$

for some bumped implied volatility $\text{IV}^\varepsilon(\tau, k)$. The parameter ε encodes the fact that the efficient price is not observed exactly but with an error. Alternatively to a relative sensitivity, one could have considered an absolute sensitivity but we believe that reasoning with relative changes makes the argumentation more intuitive.

By differentiating w.r.t. ε , we get

$$\text{relative IV sensitivity} \stackrel{\text{def}}{=} \left. \partial_\varepsilon \text{IV}^\varepsilon \right|_{\varepsilon=0} = \frac{\left. \text{Initial price} \right|_{\varepsilon=0}}{\left. \text{Vega}^{\text{BS}}(\text{IV}^\varepsilon) \right|_{\varepsilon=0}}.$$

The above sensitivity depends in particular on the log-forward-moneyness. We depict it on Figure 4 under two working assumptions on volatilities: one is computed with an SVI model with parameters $(a, b, \rho, m, \sigma) = (0.02, 0.2, -0.3, 0.05, 0.6)$ (see Equation (4.1) below), and the other with a constant volatility equal to 50%. The other parameters are $\tau = 0.25$ and $\text{Fut} = 1$. The values in the plots depend on the considered IV but the shape globally remains the same.




Figure 4: Sensitivity $\partial_{\varepsilon} \text{IV}^{\varepsilon} \big|_{\varepsilon=0}$ as a function of the log-forward-moneyness k , under different IV models: SVI and Constant volatility. Relative sensitivities for calls are depicted on the left, those for puts on the right.

Figure 4 reads as follows: a relative error on market call price of 10% ($\varepsilon = 0.1$) entails an absolute IV error that ranges from $0.1 \times 10^2 \times 10^4 = 100\,000$ bps to $0.1 \times 10^4 \times 10^4 = 10\,000\,000$ bps for ITM options ($k \ll 0$), and that is approximately $0.1 \times 10^{-1} \times 10^4 = 100$ bps for OTM options ($k \gg 0$). The situation is symmetric for puts. All in all, the relative sensitivity is much smaller for OTM options than for ITM options.

This is the reason why in our calibration procedures, we filter bid-ask prices by keeping only OTM options, i.e. we only consider call prices (resp. put prices) for $k \geq 0$ (resp. $k \leq 0$).

3.2 Bias on the calibration

In this section we evidence that the calibration of the implied total variance from noisy observed prices systematically induces an underestimation of the efficient total variance in the left and right wings. The discussion is based on a statistical learning point-of-view. The removal of the bias is discussed in the next section.

To formalize the bias result, we assume that the mid-price is a centered-noisy estimate of the efficient price, which is a standard and reasonable assumption from a microstructural point of view:

(H-1) For any (τ, k) , the mid-price is an unbiased estimate of the efficient price Eff:

$$\text{md}(\tau, k) = \text{Eff}(\tau, k) + \varepsilon(\tau, k),$$

where $\varepsilon(\tau, k)$ is not identically zero and is centered.

This assumption is fulfilled for instance by the model in [35], where the author assumes the efficient price data is observed with a round-off error, which leads to an observation noise $\varepsilon(\tau, k)$ satisfying **(H-1)**. Further details on this modeling point-of-view are provided in Section 3.3. The fact that $\varepsilon(\tau, k)$ is not identically zero is a technical condition for Theorem 3.1, it is obviously satisfied

in practice. Here the underlying probability space $(\Omega, \mathbb{P}, \mathcal{F})$ accounts for the randomness of the noise $\varepsilon(\tau, k)$ for all τ, k , and the efficient prices $\{\text{Eff}(\tau, k) \mid \tau, k\}$ are considered as fixed (calibration is performed at time $t = 0$).

Under the sole Assumption **(H-1)**, we now justify that the calibration is strongly biased. As a calibration problem, we aim at fitting a parametric model on the efficient IV (i.e., the one related to efficient prices), by solving a calibration problem on n mid-price data points:

$$\inf_{\chi} \frac{1}{n} \sum_{i=1}^n \left(\tau \cdot \text{IV}^2(\text{md}(\tau, k_i), \tau, k_i) - \omega(\tau, k_i; \chi) \right)^2, \quad (3.1)$$

where $\omega(\tau, k_i; \chi)$ is a given total variance parametrization (for some parameters χ) at the log-forward-moneyness k_i for a given observed data point. The goal is to obtain the best parameters χ^* for which

$$\omega(\tau, k; \chi^*) \approx \tau \cdot \text{IV}^2(\text{Eff}(\tau, k), \tau, k), \quad \text{for all } k.$$

In the subsequent experiments of Section 4, we consider the SVI parametrization (4.1) for the total variance parametrization ω .

The above problem is similar to a regression problem [23] that consists in finding a function \mathcal{M} such that $Y = \mathcal{M}(X) + \text{noise}$, given the observation of an n -sample of (Y, X) called "Response, Design". The main difference is that in (3.1), we do not minimize the squared difference between Y and functions of X , but between a nonlinear function of Y (here $\tau \cdot \text{IV}^2(\cdot)$) and functions of X . This composition with a non-linear function changes the theoretical analysis.

To summarize, problem (3.1) is an optimization problem over a finite-sample data, which can be compared to the ideal optimization problem

$$\inf_{\chi} \int \mathbb{E} \left[\left(\tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k) - \omega(\tau, k; \chi) \right)^2 \right] \pi(\text{d}k), \quad (3.2)$$

where π is the distribution of the observed log-forward-moneyness. Denote by φ the squared function within the above expectation. Then the difference between (3.1) and (3.2) is bounded by the so-called supremum of empirical process

$$\sup_{\chi} \left(\frac{1}{n} \sum_{i=1}^n \varphi(\text{md}(\tau, k_i), \tau, k_i; \chi) - \int \mathbb{E} [\varphi(\text{md}(\tau, k), \tau, k; \chi)] \pi(\text{d}k) \right). \quad (3.3)$$

The above statistical error can be quantified depending on the assumptions on the tails of the quantity $\varphi(\text{md}(\tau, k), \tau, k; \chi)$: under boundedness conditions, see [28]; or under light and heavy tail conditions, see [6] and [5]. See also [29] for a broad overview of error bounds in statistical learning. Essentially, these results state that the probability of (3.3) being greater than $C\delta/\sqrt{n}$ is exponentially small in $\delta > 0$: in other words, the minimum in (3.1) is close to (3.2) with high probability.




Figure 5: Visualization of the concavity (in the wings) or convexity (ATM) of the function $P \mapsto \text{IV}^2(P)$. The underlying contract is a call (resp. put) for positive (resp. negative) log-forward-moneyness. From the left to the right: $k = -1.40, k = -0.87, k = 0.02, k = 0.90$. The dashed line corresponds to the volatility in an SVI model as in Figure 1, and with $\tau = 0.25$.

Furthermore, a bias-variance decomposition of (3.2) allows to write

$$\inf_{\chi} \int \left[\text{Var} [\tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k)] + \left(\mathbb{E} [\tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k) - \omega(\tau, k; \chi)] \right)^2 \right] \pi(\text{d}k).$$

Since the first term does not depend on the parameter χ , the calibration algorithm will take mid-prices as inputs and attempt to minimize the integral of the squared bias error, i.e., to compute

$$\inf_{\chi} \int \left(\mathbb{E} [\tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k) - \omega(\tau, k; \chi)] \right)^2 \pi(\text{d}k). \quad (3.4)$$

Theorem 3.1. *Assume (H-1) holds. In the wings (i.e., when $|k| \gg 0$), the calibration algorithm with mid-prices as inputs calibrates a negatively-biased total implied variance, while at the money ($k \approx 0$), the calibration algorithm calibrates a positively-biased total implied variance.*

The above under-estimation in the wings is similar if the mid-price error

$$(\tau \cdot \text{IV}^2(\text{md}(\tau, k_i), \tau, k_i) - \omega(\tau, k_i; \chi))^2$$

in (3.1) is replaced by the average over bid and ask prices

$$\frac{1}{2} (\tau \cdot \text{IV}^2(\text{Bid}(\tau, k_i), \tau, k_i) - \omega(\tau, k_i; \chi))^2 + \frac{1}{2} (\tau \cdot \text{IV}^2(\text{Ask}(\tau, k_i), \tau, k_i) - \omega(\tau, k_i; \chi))^2.$$

Proof. From Lemma A.1 in Appendix A, the function $P \mapsto \tau \cdot \text{IV}^2(P)$ is strictly concave when $|k|$ is large enough, and strictly convex for small values of $|k|$, see Figure 5.

Because $\varepsilon(\tau, k)$ is not identically zero, the Jensen inequality becomes a strict inequality and writes for log-forward-moneyness in the wings ($|k| \gg 0$):

$$\tau \cdot \text{IV}^2(\text{Eff}(\tau, k), \tau, k) \stackrel{\text{(H-1)}}{=} \tau \cdot \text{IV}^2(\mathbb{E} [\text{md}(\tau, k)], \tau, k) > \mathbb{E} [\tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k)]. \quad (3.5)$$




Figure 6: Distribution over 10 000 samples of the mid-price IV in the model **(H-2)**, when the effective price is given by an SVI smile, with the same parameters as in Figure 1. The spread is set to 4 ticks in the wings, and 20 ticks at the money. Colors: the exact IV is in red; the IV mean over the samples is in black. Values of the log-forward-moneyness k from the left to the right: $k = -1.31, k = -0.73, k = 0.04, k = 0.62, k = 1$.

The reverse inequality holds when $k \approx 0$. This shows that in the minimization problem (3.4), the parametrization $\omega(\tau, k; \chi)$ attempts to fit a value below (resp. above) the expected one (which is $\tau \cdot \text{IV}^2(\text{Eff}(\tau, k), \tau, k)$) in the wings (resp. at the money). This proves the first theorem statement.

When the mid-price is replaced by the average of bid and ask prices, the analysis is similar and shows that the bias is even higher in the wings (when $|k| \gg 0$). Indeed,

$$\begin{aligned} \tau \cdot \text{IV}^2(\text{md}(\tau, k), \tau, k) &= \tau \cdot \text{IV}^2\left(\frac{1}{2}\text{Bid}(\tau, k) + \frac{1}{2}\text{Ask}(\tau, k), \tau, k\right) \\ &> \frac{1}{2}\tau \cdot \text{IV}^2(\text{Bid}(\tau, k), \tau, k) + \frac{1}{2}\tau \cdot \text{IV}^2(\text{Ask}(\tau, k), \tau, k), \end{aligned}$$

which is a larger inequality than that of (3.5). \square

Theorem 3.1 shows that any parametric form $\omega(\tau, k; \chi)$ within the calibration procedure will lead to an underestimation of IV in the wings and an overestimation at the money. Presumably, this phenomenon might be mitigated by the additional approximation error due to the parametric form, but nonetheless, the overall behavior should remain the same.

On traditional markets where the spread is small, the above over/under-estimation is definitely negligible but for options on cryptos for which the relative tick size is larger (see Table 3 and Figure 3), the bias becomes a more significant concern. A visualization of the calibration bias from Theorem 3.1 is depicted in Figures 6 and 7, where we use synthetic data generated by an SVI model that gives the so-called efficient price, to which we add a noise according to model **(H-2)** to get the mid-price on which we calibrate. Consistently with Lemma A.1 which shows that $\partial_P(\text{IV}^2)$ goes to $+\infty$ as $|k|$ gets larger all other things being equal, the bid and ask prices result in an increasing spread on IVs as $|k|$ increases (see the x -ranges in Figure 6). Figure 6 evidences that as $|k|$ gets larger, the mid IV (in black) increasingly underestimates the true IV (in red) – the IV bias can be larger than 100 bps – while at the money, the bias remains small, and even negligible. In Figure 7, an SVI model is calibrated on the mean of the mid total variance (according to model **(H-2)** below): in other words, this implements criterion (3.4) which is the so-called square bias term in the calibration. Since the

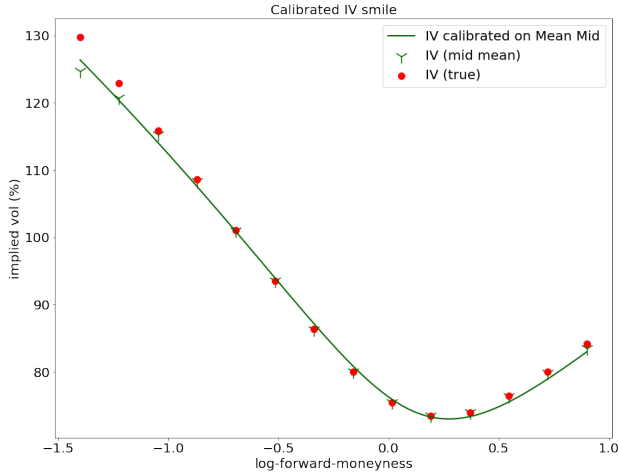


Figure 7: Recalibration of a SVI smile with the same parameters as in Figure 1. The spread is set to 10 ticks for any log-forward-moneyness. The calibration data comes from the mean of mid total variance (the related mean mid IV is depicted) where the mid-price is generated according to **(H-2)**.

mid IV underestimates in mean the true IV, the calibrated IV curve underestimates the true one in the wings, as predicted by Theorem 3.1.

3.3 Getting rid of the mid-price bias

The previous analysis shows that taking mid-prices as inputs of the calibration procedure induces a systematic bias, which is not corrected by taking the bid and the ask prices with equal probabilities. To compensate for this, we develop a novel approach that uses all the information from the bid and ask interval. This approach is based on determining an anchor price that depends on the bid and ask prices, and ideally leads to an unbiased estimator of the IV corresponding to efficient prices (called the *efficient IV* in what follows).

Anchor price. Our analysis is based on a probabilistic model of the efficient price, the mid-price and the spread.

(H-2) The noise $\varepsilon(\tau, k)$ in **(H-1)** has a symmetric distribution, uniformly distributed on an interval with a size equal to the spread; more formally,

$$\text{md}(\tau, k) = \text{Eff}(\tau, k) + \text{sp}(\tau, k) \cdot U,$$

where U is uniformly distributed on $[-1/2, 1/2]$, and independent from $\text{sp}(\tau, k)$.

Here again, we assume that $\text{Eff}(\tau, k)$ is fixed; the spread $\text{sp}(\tau, k)$ on the other hand can be random.

This model is inspired by [35], which assumes the efficient price data is observed with a round-off error. The model derivation is as follows, where we may omit the index τ, k for the sake of clarity. From Lemma A.2, and assuming for the time being that the spread is one tick (in other words, that $\delta = 1$), we can assert that the efficient price is related to the bid by the relation

$$\text{Eff} - \text{Bid} \stackrel{\text{def}}{=} V \cdot \delta$$

where V is asymptotically (as $\delta \rightarrow 0$) uniformly distributed on $[0, 1]$. Using again that the spread is taken to be one tick, we easily get

$$\text{Ask} - \text{Eff} = (1 - V) \cdot \delta \quad \text{and} \quad \text{md} = \text{Eff} + \delta \cdot \left(\frac{1}{2} - V\right).$$

Assumption **(H-2)** is obtained by letting $U \stackrel{\text{def}}{=} \frac{1}{2} - V$, which is uniformly distributed on $[-1/2, 1/2]$, and assuming that the above formula remains valid if the spread differs from one tick.

Theorem 3.2. *Under assumption **(H-2)**, define*

$$\rho \stackrel{\text{def}}{=} -\frac{\partial_P \text{IV}^2(\text{md}(\tau, k), \tau, k)}{\partial_P^2 \text{IV}^2(\text{md}(\tau, k), \tau, k) \cdot \text{sp}(\tau, k)},$$

and assume that $|\rho| \geq \frac{1}{\sqrt{12}}$. Then set

$$\nu \stackrel{\text{def}}{=} \rho - \text{Sgn}(\rho) \sqrt{\rho^2 - \frac{1}{12}} = \frac{\text{Sgn}(\rho)}{12 \left(|\rho| + \sqrt{\rho^2 - \frac{1}{12}} \right)}, \quad (3.6)$$

and define the anchor price by $\text{Anc}(\tau, k) \stackrel{\text{def}}{=} \text{md}(\tau, k) + \text{sp}(\tau, k) \cdot \nu$, with the convention $|\rho| = +\infty$ and $\nu = 0$ if $\partial_P^2 \text{IV}^2(\text{md}(\tau, k), \tau, k) = 0$.

Then the anchor price $\text{Anc}(\tau, k)$ depends only on $\text{Bid}(\tau, k)$ and $\text{Ask}(\tau, k)$, lies in the interval $(\text{Bid}(\tau, k), \text{Ask}(\tau, k))$, and its bias regarding the effective price is negligible at order 3 w.r.t. the tick size:

$$\mathbb{E} [\text{IV}^2(\text{Anc}(\tau, k), \tau, k)] = \text{IV}^2(\text{Eff}(\tau, k), \tau, k) + O(\text{ts}^3). \quad (3.7)$$

The proof is postponed to Appendix A.3. Figure 8 depicts the boxplots of implied volatilities based on the mid-price and on the anchor price. We observe that, as expected, the anchor price has an implied volatility that is much better aligned (in mean) with the efficient IV. Interestingly, we also note that the fluctuations of the anchor IV is slightly smaller than those of the mid-price IV, suggesting that the use of the anchor price yields better variance properties compared to the mid-price.




Figure 8: Boxplots, for different log-forward-moneynesses on x -axis, of implied volatilities based on mid-price or anchor price. Size of the sample: 100 000. The solid lines in the boxplots correspond to the empirical mean over the sample. The parameters for the efficient price and the spread are the same as in Figure 1.

Remark. The most natural choice for the distribution of U is uniform on $[-\frac{1}{2}, \frac{1}{2}]$ conditionally to sp , in view of Lemma A.2. We explore what results alternative modeling choices would have given regarding the bias correction from Theorem 3.2. Had we considered another distribution for U , we would have obtained

$$\nu \stackrel{\text{def}}{=} \frac{\sigma^2 \text{Sgn}(\rho)}{|\rho| + \sqrt{\rho^2 - \sigma^2}} \quad \text{where} \quad \sigma \stackrel{\text{def}}{=} \sqrt{\text{Var}(U)},$$

with the constraint $|\rho| \geq \sigma$. If $\sigma > 1/\sqrt{12}$, then the constraint on ρ is stronger than the one in Theorem 3.2. The extremal ν (corresponding to largest unbiasing parameter) is obtained for $\rho = \pm\sigma$, for which $\nu = \pm\sigma$. Besides, the largest σ corresponds to the Bernoulli distribution ($\sigma = 1/2$).

In Appendix A.4, we study under what conditions the constraint $|\rho| \geq 1/\sqrt{12}$ from Theorem 3.2 holds. Although at first sight this constraint seems quite restrictive w.r.t. k , it is not so if we take care to restrict the considered market data to prices greater than 1 tick, which anyone would do within a calibration framework.

3.4 Data augmentation to handle missing data

About missing data. As previously observed (see Section 2.2), on 10% to 15% of the data, the bid price is outside the arbitrage bounds. With the filter on the OTM options, this corresponds to very low prices at 1 or even zero ticks (when the data is missing). In this case of fully incoherent or missing data, a first approach is to project the observed bid price on the lower arbitrage bound rounded up to the tick. Once this is done, it is then possible to compute a mid-price that can be used as an input for the calibration process. However, with a large tick size and/or large spreads, the mid-price is a quite noisy representative of the efficient price and we propose a method that better takes into account all the uncertainty around the efficient price.

Generally speaking, the strategy of replacing missing or incoherent data by a reasonable guess (here, replacing the wrong bid by the arbitrage bound) is one of the possible methodologies that can be used when such an event occurs. There is an abundant literature on this topic in other fields (clinical trials, sample surveys, agricultural experiments, . . .), see for instance [11] for a broad overview, and to the best of our knowledge, there is no such related work in the context of volatility calibration. It is known that replacing missing volatility data by a single value is simple but not without risk. Among the possible methods, one can take advantage of past implied volatility data, in order to naively replace the missing or incoherent bid (over the period considered for the calibration) by a recent value. This method, which is called *Last Observation Carried Forward*, does not fit well with the calibration framework because of the speed at which data changes (in finance, and especially on crypto-markets) and the use of past data does not give the best up-to-date estimation on the current implied volatility. For instance, it could be the case that in the recent past the spread was small but, due to new market conditions, the spread gets much higher with no available bid at the next calibration step, which would result in a likely large error if the recent bid was used.

Other methods using parametric predictions based on past data may suffer from the same issue of quickly changing data. Alternatively, in the 70's Rubin developed a statistical methodology based on multiple imputations, i.e. each missing value is replaced by two or more imputed values in order to represent the uncertainty about which value to impute. See [36] for a monograph by this author, and [37] for a more recent reference where an asymptotic analysis is made in parametric model. This strategy allows for better robustness when missing data and improved variance due to effect of data augmentation.

Data augmentation. Our approach follows the same idea of multiple imputation, i.e. data augmentation, but we use this principle not to represent the uncertainty around the bid but directly the uncertainty around the efficient price. To allow for a general treatment, we denote by π_k^{aug} the distribution of augmented data at the log-forward-moneyness k ; this distribution has a support on the interval⁷ $(\text{Bid}(\tau, k_i), \text{Ask}(\tau, k_i))$ or equivalently $(0, 1)$ after a change of variable. Thus, the initial

⁷Here, $\text{Bid}(\tau, k_i)$ is the observed bid, or the lower arbitrage bound rounded up to the tick in case of a missing or incoherent bid

calibration problem (3.1) becomes

$$\inf_x \frac{1}{n} \sum_{i=1}^n \int_0^1 \left(\tau \cdot \text{IV}^2(\text{Bid}(\tau, k_i) + x \cdot \text{sp}(\tau, k_i), \tau, k_i) - \omega(\tau, k_i; \chi) \right)^2 \pi_{k_i}^{\text{aug.}}(dx).$$

For a given k , if we were using only the mid-price, then $\pi_k^{\text{aug.}}$ would be a Dirac measure at $1/2$; if we were equally using the bid and the ask prices, then $\pi_k^{\text{aug.}}$ would be the average of two Dirac measures (one at 0, the other at 1). One could take $\pi_k^{\text{aug.}}$ as the uniform distribution on $(0, 1)$, but this distribution would have the same mean as the mid-price and thus, in view of Theorem 3.1, it would yield the same bias, i.e. a strong underestimation of the effective IV in the wings (see Figure 6).

Data augmentation with the anchor price. To account for the bias, we impose the distribution $\pi_k^{\text{aug.}}$ to be centered on the anchor price $\text{Anc}(\tau, k)$, so that

$$\int_0^1 \text{IV}^2(\text{Bid}(\tau, k_i) + x \cdot \text{sp}(\tau, k_i), \tau, k_i) \pi_{k_i}^{\text{aug.}}(dx) = \text{IV}^2(\text{Anc}(\tau, k_i), \tau, k_i), \quad (3.8)$$

and the final calibration is expected to have a reduced bias.

All in all, the method that we propose, called the *Data Augmentation* method, combines two effects: multiple imputations of the possible price values in the bid-ask interval for a better robustness when there is uncertain and missing data, and alignment in mean to the anchor price to reduce the calibration bias. Although this method is motivated by missing or incoherent bids, we recommend using this method in all cases, as confirmed by our numerical tests in Section 4.

For practical considerations, there remains to define a distribution with support in $(0, 1)$ with a prescribed statistics like (3.8). We propose the use of a Beta distribution $\text{Beta}(a, b)$, the density of which is

$$f_{a,b}(x) := \frac{x^{a-1}(1-x)^{b-1}}{B(a,b)} \mathbf{1}_{0 < x < 1} \quad \text{with} \quad B(a,b) := \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)},$$

with some parameters $a > 0, b > 0$; it has the advantage of including the uniform distribution ($a = b = 1$), and being easily tunable because it is defined by two parameters only. Since the target in (3.8) depends on the anchor price, the calibrated parameters a, b will depend on the log-forward-moneyness k_i , and on the bid and ask prices. Finding parameters satisfying (3.8) is made possible thanks to the following result (proved in Appendix A.5).

Proposition 3.3. *Denote by $\beta_{a,b}$ the random variable having the $\text{Beta}(a, b)$ distribution. For any strictly increasing continuous function $\varphi : [0, 1] \mapsto \mathbb{R}$, the mapping*

$$a \in (0, +\infty) \mapsto \bar{\varphi}(a) := \mathbb{E}[\varphi(\beta_{a,1})]$$

is continuous, strictly increasing from $\varphi(0)$ to $\varphi(1)$.

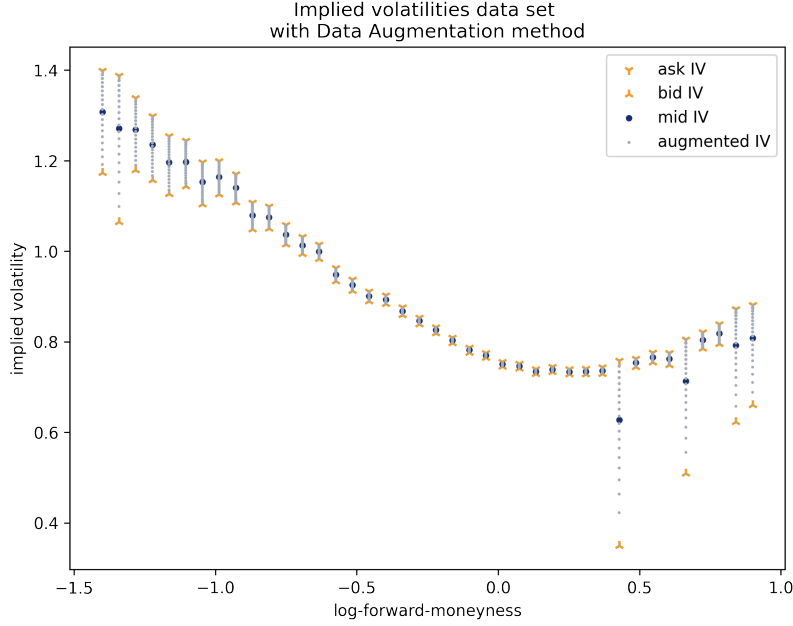


Figure 9: Illustration of the data augmentation method

As a consequence, by setting $\varphi(x) \stackrel{\text{def}}{=} \text{IV}^2(\text{Bid}(\tau, k_i) + x \cdot \text{sp}(\tau, k_i), k_i, \tau)$ (which is strictly increasing in x), there is a unique $a_i \in (0, +\infty)$ such that (3.8) holds: indeed the range of $\varphi(\cdot)$ is the interval $[\text{IV}^2(\text{Bid}(\tau, k_i), k_i, \tau), \text{IV}^2(\text{Ask}(\tau, k_i), k_i, \tau)]$ and the value $\text{IV}^2(\text{Anc}(\tau, k_i), k_i, \tau)$ lies strictly in this interval. Because

$$\mathbb{E} [\text{IV}^2(\text{Bid}(\tau, k_i) + \beta_{a,1} \cdot \text{sp}(\tau, k_i), k_i, \tau)] = \int_0^1 \text{IV}^2(\text{Bid}(\tau, k_i) + x \cdot \text{sp}(\tau, k_i), \tau, k_i) \pi_{k_i}^{\text{aug}}(dx)$$

is strictly increasing in a , one can use a bisection method to find the (unique) root of (3.8).

The principle of the data augmentation step is illustrated in Figure 9, in the case where there are large spreads and corrupted bids. Instead of considering a continuous distribution of points in the bid-ask interval at each log-forward-moneyness, we impute a finite number of points (here 20), with calibration weights associated to the Beta distribution, tuned as above.

4 Numerical tests

4.1 Results on synthetic data

SVI parametrization. SVI is a quite popular parametrization of implied variance which is key in describing call and put option prices at a given expiry τ . This is a five-parameters model,

Parameters	# of log-forward-moneyness	spread level	% of spurious bids
<i>Use case 1</i>	30	2 ticks	0%
<i>Use case 2</i>	20	4 ticks	15%
<i>Use case 3</i>	10	10 ticks	30%

Table 4: Parameters used for each use case.

$\chi^R = (a, b, \rho, m, \sigma)$, on the implied total variance, this writes

$$w_\tau(k; \chi^R) := a + b\{\rho(k - m) + \sqrt{(k - m)^2 + \sigma^2}\}. \quad (4.1)$$

Here $k = \log(F_\tau/K)$ is the log-forward-moneyness, which depends on the forward with expiry τ .

The SVI parametrization was first introduced in [17] at Merrill Lynch in 1999. The consistency of the SVI parametrization with large log-moneyness interpolations is studied in [27]. In [19], it is shown that the implied Heston volatility curve in large maturities is of SVI-type (hence the name SVI for *Stochastic Volatility Inspired*). Issues related to numerical calibration of an SVI parameterization are studied in [8]. Works about interpolation without arbitrage are available in [12, 22]. There are 3 known SVI parameterizations (Raw, Natural, Jump-Wing) that we can switch between and they are associated to different parameter interpretations; in (4.1), we choose the Raw SVI parameterization. See [18] for a broad overview of quantitative analysis on volatility.

In the SVI calibration, we must have $a \in \mathbb{R}$, $b \geq 0$, $|\rho| < 1$, $m \in \mathbb{R}$, $\sigma > 0$ and the condition $a + b\sigma\sqrt{1 - \rho^2} \geq 0$ ensures that $w_\tau(k; \chi^R) \geq 0$, $\forall k \in \mathbb{R}$. In addition, a certain function (depending on the calibrated parameters) has to be positive to ensure no butterfly arbitrage, see [20, Section 2.2] for details.

For the upcoming tests on synthetic data, we use the SVI IV of Figure 1.

Results on missing data. To show the efficiency of the *Data Augmentation* method, we compared it to the *Mid* method on three different use cases:

- (1) The first one, hereafter referred to as *Use Case 1*, represents a standard market case with a high level of liquidity.
- (2) The second use case (*Use Case 2*), represents a market with poor liquidity.
- (3) The last use case (*Use Case 3*) finally represents a chaotic market and a liquidity crisis.

For each use case, we initialize test parameters to describe the market they should represent. These parameters are:

- The number of log-forward-moneynesses simulated, i.e. the number of available strikes.
- The bid/ask spread level.

Use Case	Statistics	L1 error	L2 error
<i>Use Case 1</i>	Mid method	17.21 (± 0.60)	26.68 (± 1.04)
	Data Augmentation method	13.06 (± 0.41)	19.62 (± 0.71)
<i>Use Case 2</i>	Mid method	17.21 (± 0.60)	26.68 (± 1.04)
	Data Augmentation method	13.06 (± 0.41)	19.62 (± 0.71)
<i>Use Case 3</i>	Mid method	69.14 (± 2.76)	95.78 (± 3.88)
	Data Augmentation method	49.41 (± 1.55)	67.69 (± 2.34)

Table 5: Mean of the L1 and L2 errors and their 95%-confidence intervals

- The percentage of spurious bid prices (representing the bids at 0 or 1 tick observed in Section 2.2).

Note that the last parameter acts more as an aggravating factor for the liquidity. The missing or abnormal bid prices are not removed but projected onto the arbitrage boundary (rounded up to the next tick). Table 4 provides the values of the parameters for each use case.

Apart from the first step of simulating the 1000 samples, the rest of the test is identical for each case:

- Computing the mid-prices and their associated implied volatilities.
- Augmenting the data in the bid-ask interval by 100 equally spaced points and weighted with the Beta distribution, see Section 3.3, for the *Data Augmentation* method. Note that we have also tested adding only 10 data points and we do not observe any significant difference.
- Calibrating two SVI smiles using BFGS for the optimization algorithm [13, Chapter 3] and the Huber loss function [26], on both data sets (the one obtained using the *Mid* method and the other using the *Data Augmentation* method).
- Computing calibration errors over all log-forward-moneynesses between the calibrated implied volatilities and the *efficient* ones for each method:
 - The bias: mean of the differences between the volatilities.
 - The L1 error: mean of the absolute values of the differences between the volatilities.
 - The L2 error: quadratic mean of the differences between the volatilities.

The following tables and figures show the calibration errors and some statistics. All results are expressed in bps.

- Table 5 contains the error averages over the 1000 scenarii, along with their confidence intervals.

- In Figure 10, we plot the Q-Q plots of the bias and the L1 and L2 errors over the 1000 scenarii.
- Figure 11 presents the Box plots of the L1 and L2 errors corresponding to the 1000 scenarii.

Table 5 shows that the errors are significantly reduced. The L1 error is down by 24% (resp. 29% and 27%) for *Use Case 1* (resp. *Use Case 2* and *Use Case 3*) and the L2 error is reduced by 26% (resp. 29% and 30%). This reduction by on average 27.5% (regardless of the use case, i.e., the number of log-forward-moneynesses, the spread or the percentage of missing bids) shows the efficiency of the *Data Augmentation* method and also confirms its stability. We have conducted other tests with different use-cases, results on the performance are not reported here but they are similar.

In the Q-Q plots we note that the L1 and L2 errors are much lower with the *Data Augmentation* method than with the *Mid* method, and this applies to the mean (already reported in Table 5) but also to the whole distribution which is closer to zero. As for the bias, whether positive or negative, it is more concentrated around zero with the *Data Augmentation* method.

Finally the Box plots show that the problematic cases (those with the largest errors) are more limited with the *Data Augmentation* method, since the dispersion of the errors is smaller and the extreme values are lower. This shows the better robustness of this method over the *Mid* method.

4.2 Results on real data

The main goal of this work is the design of an implied volatility calibrator for cryptocurrency options that does not require any human intervention to generate consistent, robust and stable smiles over time. In most cases, enriching trades with bid-ask quotes and running the calibration process with trades and mid-prices is sufficient to construct reasonable smiles; for these situations such as the one in Figure 12, it is not clear how to quantify the gain that is obtained using the *Data Augmentation* method. This method appears qualitatively similar on this example but the theoretical analysis allows us to expect an overall benefit on all smiles that are calibrated.

In other cases, such as those involving liquidity issues (large bid-ask spreads and missing or incoherent bid prices), the *Data Augmentation* method corrects the bias due to these liquidity problems. Figure 13 depicts one of these cases. One can note that for all log-forward-moneynesses above 1.6, bid prices are missing and are replaced by the minimum price level, i.e. 1 tick. As a consequence, mid-prices are skewed downward and implied volatilities are underestimated. The *Data Augmentation* method allows to correct this.

5 Conclusion

Cryptocurrency options are growing increasingly popular, both for retail and institutional investors. This has led to a higher demand for tools that are standard in traditional markets, but less so in cryptocurrency markets, such as calibrators for implied volatilities. Although derivative products – and in particular options – that can be traded on cryptocurrency markets are similar to those that are

available on traditional ones, standard calibration procedures cannot be applied in a straightforward way to crypto options. Indeed, there are major differences between both markets, such as relative tick sizes, bid-ask spread sizes and amount of missing data that entail a problematic loss of accuracy in the calibrated implied volatilities.

In this paper we have designed a new calibration procedure specifically designed to overcome the shortcomings of standard procedures based on trade or mid-prices. Our calibration procedure is based on two features:

- the construction of an anchor price that permits to get rid of the bias that occurs in the wings for calibration procedures based on mid-prices;
- a data augmentation step that permits to handle in a robust way the case where some of the prices to be used in the calibration are missing or incoherent.

The experimental results from Section 4 show that the combination of both features results in a calibration procedure that provides more accurate results than standard ones.

As far as future work is concerned, we intend to investigate how our methods can be extended to handle the calibration of implied volatility surfaces taking the calendar arbitrage into account. In addition, we will study the evolution over time of calibrated parameters and provide some prediction methodologies of implied volatility surfaces.

Acknowledgments. The authors wish to thank Clara Medalie and Conor Ryder from Kaiko for their valuable feedback, as well as the Kaiko teams for their inputs and the data they provided.

A Technical results

A.1 Sensitivity of IV

Lemma A.1. *Let τ, k be fixed. Then the implied volatility function $P \mapsto \text{IV}(P)$ as a function of the input price P is such*

$$\begin{aligned}\partial_P(\text{IV}^2) &= 2 \text{IV}/\text{Vega}, \\ \partial_P^2(\text{IV}^2) &= \frac{2}{\text{Vega}^2} \left(-\frac{k^2}{\text{IV}^2\tau} + \frac{\text{IV}^2\tau}{4} + 1 \right),\end{aligned}$$

where Vega is the Black-Scholes Vega with parameters τ, k, IV .

Proof. Writing $P = \text{Black-Scholes price}(\text{IV}(P))$ and omitting the indices τ, k, BS for the sake of simplicity, we have $1 = \text{Vega} \cdot \partial_P \text{IV}(P)$ and $0 = \text{Vomma} \cdot (\partial_P \text{IV}(P))^2 + \text{Vega} \cdot \partial_P^2 \text{IV}(P)$. Using $\text{Vomma} = \text{Vega} \frac{d_+ d_-}{\text{IV}}$ [25, p.23] with the usual definitions for d_{\pm} , we get

$$\partial_P \text{IV}(P) = 1/\text{Vega}, \quad \partial_P^2 \text{IV}(P) = -\text{Vomma}/\text{Vega}^3 = -d_+ d_- / (\text{IV} \cdot \text{Vega}^2),$$

$$\partial_P^2(\text{IV}^2) = \frac{2}{\text{Vega}^2}(-d_+d_- + 1) = \frac{2}{\text{Vega}^2} \left(-\frac{k^2}{\text{IV}^2\tau} + \frac{\text{IV}^2\tau}{4} + 1 \right).$$

□

A.2 Distribution of round-off error

The following result states a nice and explicit approximation of the joint distribution of a random variable and its round-off error, in the asymptotics of small round-off level. Such results were initiated by [30] and [38], and stated in a refined form later in [9]. Define the fractional part of a scalar x as follows:

$$\{x\} = x - \sup\{n \in \mathbb{N} : n \leq x\}.$$

Lemma A.2. *Let δ be a round-off level and let Y be a scalar random variable, which distribution is assumed to have a density which is N -times differentiable ($N \geq 1$) with integrable derivatives. Then, for any function f assumed to be N times differentiable with bounded derivatives, we have*

$$\mathbb{E}[f(Y, \{Y/\delta\})] - \mathbb{E}[f(Y, V)] = \mathcal{O}(\delta^N),$$

where V is a random variable, uniformly distributed on $[0, 1]$ and independent from Y .

In words, the round-error of Y defined by $\delta\{Y/\delta\}$ is asymptotically independent from Y and uniformly distributed on $[0, \delta]$.

A.3 Proof of Theorem 3.2

We omit the index (τ, k) in the notations. A second-order Taylor expansion of $\text{IV}^2(\cdot)$ gives, keeping in mind that $|\text{Eff} - \text{md}| \leq \frac{1}{2}\text{sp}$, $|\text{Anc} - \text{md}| \leq \frac{1}{2}\text{sp}$ and $\text{sp} = \mathcal{O}(\text{ts})$:

$$\text{IV}^2(\text{Anc}) = \text{IV}^2(\text{Eff}) + \partial_P \text{IV}^2(\text{Eff})(\text{Anc} - \text{Eff}) + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{Eff})(\text{Anc} - \text{Eff})^2 + \mathcal{O}(\text{ts}^3).$$

Because $\text{Anc} - \text{Eff} = \text{sp} \cdot (\nu + U)$, observe that (3.7) is equivalent to

$$\begin{aligned} 0 &= \mathbb{E} \left[\partial_P \text{IV}^2(\text{Eff}) \cdot \text{sp} \cdot (\nu + U) + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot (\nu + U)^2 \right] + \mathcal{O}(\text{ts}^3) \\ &= \mathbb{E} \left[\partial_P \text{IV}^2(\text{Eff}) \cdot \text{sp} \cdot \nu + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot (\nu^2 + 2\nu \cdot U + \frac{1}{12}) \right] + \mathcal{O}(\text{ts}^3), \end{aligned}$$

where we used the fact that U is uniformly distributed on $[-1/2, 1/2]$, conditionally to sp . Note that this conditioning and centering argument does not apply to the term $2\nu \cdot U$, since ν depends on (sp, md) and not on sp only. However, observe that

$$\mathbb{E} \left[\partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot U \cdot \nu(\text{md}, \text{sp}) \right] = \mathcal{O}(\text{ts}^3). \quad (\text{A.1})$$

Indeed the function ν is smooth in the variable md , so that $\nu(\text{md}, \text{sp}) = \nu(\text{Eff}, \text{sp}) + O(\text{ts})$, and since U is centered conditionally to sp we get

$$\mathbb{E} [\partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot U \cdot \nu(\text{md}, \text{sp})] = \mathbb{E} [\partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot U \cdot \nu(\text{Eff}, \text{sp})] + O(\text{ts}^3) = O(\text{ts}^3).$$

It follows that

$$0 = \mathbb{E} \left[\partial_P \text{IV}^2(\text{Eff}) \cdot \text{sp} \cdot \nu + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot \left(\nu^2 + \frac{1}{12} \right) \right] + O(\text{ts}^3). \quad (\text{A.2})$$

Moreover, writing

$$\begin{aligned} \partial_P^2 \text{IV}^2(\text{Eff}) &= \partial_P^2 \text{IV}^2(\text{md}) + O(\text{ts}), \\ \partial_P \text{IV}^2(\text{Eff}) &= \partial_P \text{IV}^2(\text{md}) + \partial_P^2 \text{IV}^2(\text{Eff})(\text{Eff} - \text{md}) + O(\text{ts}^2), \end{aligned}$$

(A.2) becomes

$$\begin{aligned} 0 &= \mathbb{E} \left[\partial_P \text{IV}^2(\text{md}) \cdot \text{sp} \cdot \nu - \partial_P^2 \text{IV}^2(\text{Eff}) \cdot \text{sp}^2 \cdot U \cdot \nu + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{md}) \cdot \text{sp}^2 \cdot \left(\nu^2 + \frac{1}{12} \right) \right] + O(\text{ts}^3) \\ &= \mathbb{E} \left[\partial_P \text{IV}^2(\text{md}) \cdot \text{sp} \cdot \nu + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{md}) \cdot \text{sp}^2 \cdot \left(\nu^2 + \frac{1}{12} \right) \right] + O(\text{ts}^3), \end{aligned}$$

using (A.1) once again.

Hence, there remains to check that ν in (3.6) solves

$$0 = \partial_P \text{IV}^2(\text{md}) \cdot \text{sp} \cdot \nu + \frac{1}{2} \partial_P^2 \text{IV}^2(\text{md}) \cdot \text{sp}^2 \cdot \left(\nu^2 + \frac{1}{12} \right).$$

If $\partial_P^2 \text{IV}^2(\text{Eff}) = 0$, then $\nu = 0$ solves the above and this coincides with (3.6) by taking $|\rho| \uparrow +\infty$. In the other cases, this writes $(\nu - \rho)^2 = \rho^2 - \frac{1}{12}$. This equation has one or two roots for $|\rho| \geq 1/12$ and the one in (3.6) corresponds to the closest to zero.

Last we easily check that from the right hand side of (3.6), that $|\nu| \leq 1/\sqrt{12} < 1/2$ so that the anchor price lies strictly in the bid-ask interval. This concludes the proof. \square

A.4 Verification of the condition on ρ in Theorem 3.2

Let us now study under which conditions the constraint $|\rho| \geq 1/\sqrt{12}$ from Theorem 3.2 is met. First, we can simplify the expression of ρ using Lemma A.1:

$$\rho = -\frac{\partial_P \text{IV}^2(\text{md}(\tau, k), \tau, k)}{\partial_P^2 \text{IV}^2(\text{md}(\tau, k), \tau, k) \text{sp}(\tau, k)} = -\frac{\text{IV} \cdot \text{Vega}}{\left(-\frac{k^2}{\text{IV}^2 \tau} + \frac{\text{IV}^2 \tau}{4} + 1 \right) \text{sp}}. \quad (\text{A.3})$$

Around the log-forward-money $k \approx 0$, the situation is quite clear: indeed, from [25, p.23], $\text{Vega} = F\sqrt{\tau}\mathcal{N}'(d_+)$ and neglecting $\frac{\text{IV}^2 \tau}{4}$ for usual expiries and volatilities, we get $\rho \approx -\text{IV} \frac{F}{\text{sp}} \sqrt{\frac{\tau}{2\pi}} \approx -\text{IV}\sqrt{\tau}$.

800, which is much larger (in absolute value) than $1/\sqrt{12}$ for any realistic market data. Thus, our concern is rather for OTM values of k . As $k \rightarrow \pm\infty$, we observe that $\rho \rightarrow 0$ all other things being equal, making it seemingly hopeless to satisfy the condition $|\rho| \geq 1/\sqrt{12}$. But in practice, the option prices that are used for calibration, are above one or two ticks; this significantly restricts the range of possible values of k and \mathbb{IV} –at least, this imposes a constraint on \mathbb{IV} w.r.t. k . More precisely, define $\underline{\mathbb{IV}}(k)$ as the Black-Scholes volatility for which the OTM Call or the OTM Put (with forward-log-moneyness k) is worth n ticks. In other words, any volatility larger than $\underline{\mathbb{IV}}(k)$ leads to an option price larger than n ticks (i.e. to a market data point that is considered for calibration) and conversely, any market data point in the calibration set (i.e. with a value above n ticks) with log-forward-moneyness k has an implied volatility larger than $\underline{\mathbb{IV}}(k)$. Figure 14 depicts this minimal Implied Volatility $\underline{\mathbb{IV}}$ as a function⁸ of k , in the case where the minimal option price is equal to $n = 1$ tick (this represents the most permissive case). We argue that these values give a worst-case volatility for checking the condition $|\rho| \geq 1/\sqrt{12}$. Indeed, on Figure 15, we can visually check that the function

$$\sigma \mapsto \rho(\sigma) = -\frac{\sigma \cdot \text{Vega}(\sigma, k)}{\left(-\frac{k^2}{\sigma^2\tau} + \frac{\sigma^2\tau}{4} + 1\right) \text{sp}}$$

is increasing for OTM log-forward-moneyness k . Hence, it is sufficient for such OTM options to check the condition $|\rho| \geq 1/\sqrt{12}$ when the input price is generated by the implied volatility $\underline{\mathbb{IV}}(k)$. Figure 16 gives the graphs of such extremal ρ -functions for different time-to-expiries (when the minimal price is 1 tick): in all cases, the required condition on ρ is largely met. Had we only considered option with prices greater than 2 ticks or more, we would have fulfilled even more easily the condition $|\rho| \geq 1/\sqrt{12}$.

A.5 Proof of Proposition 3.3

The mean and variance of Beta(a, b) distribution is

$$\mathbb{E}[\varepsilon_i] = \frac{a}{a+b}, \quad \text{Var}(\varepsilon_i) = \frac{ab}{(a+b)^2(a+b+1)}.$$

Therefore, for $b = 1$,

- as $a \downarrow 0$, $\mathbb{E}[\varepsilon_i] \rightarrow 0$ and $\text{Var}(\varepsilon_i) \rightarrow 0$, hence $\varepsilon_i \rightarrow 0$ in probability, thus $\lim_{a \downarrow 0} \mathbb{E}[\varphi(\beta_{a,1})] = \varphi(0)$.
- as $a \uparrow +\infty$, $\mathbb{E}[\varepsilon_i] \rightarrow 1$ and $\text{Var}(\varepsilon_i) \rightarrow 0$, $\varepsilon_i \rightarrow 1$ in probability and $\lim_{a \uparrow +\infty} \mathbb{E}[\varphi(\beta_{a,1})] = \varphi(1)$.

It remains to show that $\bar{\varphi}(a)$ is strictly increasing and continuous in a . The continuity is clear writing $\bar{\varphi}$ as an integral with respect to the density $f_{a,1}$. The increasingness can be shown using the Beta-Gamma algebra property [10]: $\varepsilon^{(a,b)} \stackrel{d}{=} \frac{\Gamma(a)}{\Gamma(a)+\Gamma(b)} < \frac{\Gamma(a)+\Gamma(\varepsilon)}{\Gamma(a)+\Gamma(\varepsilon)+\Gamma(b)} \stackrel{d}{=} \frac{\Gamma(a+\varepsilon)}{\Gamma(a+\varepsilon)+\Gamma(b)} \stackrel{d}{=} \varepsilon^{(a+\varepsilon,b)}$ where all Gamma distributed random variables are independent; then $\mathbb{E}[\varphi(\beta_{a,1})] < \mathbb{E}[\varphi(\beta_{a+\varepsilon,1})]$ using the strict monotony of φ , hence the announced property. \square

⁸The range of k is defined as $[-2\sqrt{\tau}, +2\sqrt{\tau}]$, which easily encompasses the observed ranges of k as discussed in Section 2.

References

- [1] E. Alòs, J. León, and J. Vives. On the short-time behavior of the implied volatility for jump-diffusion models with stochastic volatility. *Finance and Stochastics*, 11:571–589, 2007.
- [2] L. Bergomi and J. Guyon. The smile in stochastic volatility models. *Preprint available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1967470*, 2011.
- [3] R. Bompis and E. Gobet. Asymptotic and non asymptotic approximations for option valuation. In T. Gerstner and P. Kloeden, editors, *Recent Developments in Computational Finance: Foundations, Algorithms and Applications*, chapter 4, pages 159–241. World Scientific Publishing Company, 2012.
- [4] R. Bompis and E. Gobet. Analytical approximations of local-Heston volatility model and error analysis. *Mathematical Finance*, 28:920–961, July 2018.
- [5] L. Chamakh, E. Gobet, and W. Liu. Orlicz norms and concentration inequalities for β -heavy tailed random variables. *In minor revision for Bernoulli*, 2021.
- [6] L. Chamakh, E. Gobet, and Z. Szabo. Orlicz random Fourier feature. *Journal of Machine Learning Research*, 21:1–37, 2020.
- [7] J. Corbetta, P. Cohort, I. Laachir, and C. Martini. Robust calibration and arbitrage-free interpolation of SSVI slices. *Decisions in Economics and Finance*, 42(2):665–677, 2019.
- [8] S. De Marco and C. Martini. Quasi-explicit calibration of Gatheral’s SVI model. *Zeliade White Paper*, pages 1–15, 2009.
- [9] S. Delattre and J. Jacod. A central limit theorem for normalized functions of the increments of a diffusion process, in the presence of round-off errors. *Bernoulli*, 3(1):1–28, 1997.
- [10] D. Dufresne. Algebraic properties of beta and gamma distributions, and applications. *Advances in Applied Mathematics*, 20(3):285–299, 1998.
- [11] C. K. Enders. *Applied missing data analysis*. The Guilford Press, 2010.
- [12] M. Fengler. Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance*, 9(4):417–428, 2009.
- [13] R. Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, second edition, 2020.
- [14] J.-P. Fouque, G. Papanicolaou, R. Sircar, and S. Knut. *Multiscale stochastic volatility for equity, interest rate, and credit derivatives*. Cambridge University Press, Cambridge, 2011.

- [15] P. Friz, J. Gatheral, A. Gulisashvili, A. Jacquier, and J. Teichmann, editors. *Large Deviations and Asymptotic Methods in Finance*. Springer Proceedings in Mathematics and Statistics. Springer-Verlag, Berlin, 2015.
- [16] M. B. Garman and S. W. Kohlhagen. Foreign currency option values. *Journal of international Money and Finance*, 2(3):231–237, 1983.
- [17] J. Gatheral. A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives. *Presentation at Global Derivatives & Risk Management, Madrid*, www.math.nyu.edu/fellows_fin_math/gatheral/madrid2004.pdf, May 2004.
- [18] J. Gatheral. *The volatility surface: a practitioner’s guide*, volume 357. John Wiley & Sons, 2011.
- [19] J. Gatheral and A. Jacquier. Convergence of Heston to SVI. *Quantitative Finance*, 11(8):1129–1132, 2011.
- [20] J. Gatheral and A. Jacquier. Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1):59–71, 2014.
- [21] H. Geman, N. El Karoui, and J.-C. Rochet. Changes of numeraire, changes of probability measure and option pricing. *Journal of Applied probability*, 32(2):443–458, 1995.
- [22] J. Glaser and P. Heider. Arbitrage-free approximation of call price surfaces and input data risk. *Quantitative Finance*, 12(1):61–73, 2012.
- [23] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer Series in Statistics, 2002.
- [24] P. Hagan, D. Kumar, A. Lesniewski, and D. Woodward. Managing smile risk. *Willmott Magazine*, pages 84–108, 2002.
- [25] E. Haug. *The complete guide to option pricing formulas*. McGraw-Hill Professional, second edition, 2007.
- [26] P. J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73 – 101, 1964.
- [27] P. Jäckel and C. Kahl. Hyp hyp hooray. *Wilmott Magazine*, 34:70–81, 2008.
- [28] T. Klein and E. Rio. Concentration around the mean for maxima of empirical processes. *The Annals of Probability*, 33(3):1060–1077, 2005.
- [29] V. Koltchinskii. *Oracle inequalities in empirical risk minimization and sparse recovery problems: École D’Été de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.

- [30] P. Kosulajeff. Sur la répartition de la partie fractionnaire d'une variable. *Mat. Sb. (N.S.)*, 2:1017–1019, 1937.
- [31] F. Le Floch. Initial guesses for SVI calibration. *Available at SSRN 2501898*, 2014.
- [32] R. Lee. Implied volatility: Statics, dynamics, and probabilistic interpretation. In *Recent Advances in Applied Probability*. Springer, 2005.
- [33] M. Lorig, S. Pagliarani, and A. Pascucci. Explicit implied volatilities for multifactor local-stochastic volatility models. *Mathematical finance*, 27:926–960, 2017.
- [34] M. Musiela and M. Rutkowski. *Martingale methods in financial modelling*. Springer Verlag, second edition, 2005.
- [35] M. Rosenbaum. Integrated volatility and round-off error. *Bernoulli*, 15(3):687–720, 2009.
- [36] D. B. Rubin. *Multiple imputation for nonresponse in surveys*. NJ: John Wiley & Sons, 1987.
- [37] A. A. Tsiatis. *Semiparametric theory and missing data*. Springer, 2006.
- [38] J. W. Tukey. On the distribution of the fractional part of a statistical variable. *Mat. Sb. (N.S.)*, 4:561–562, 1939.

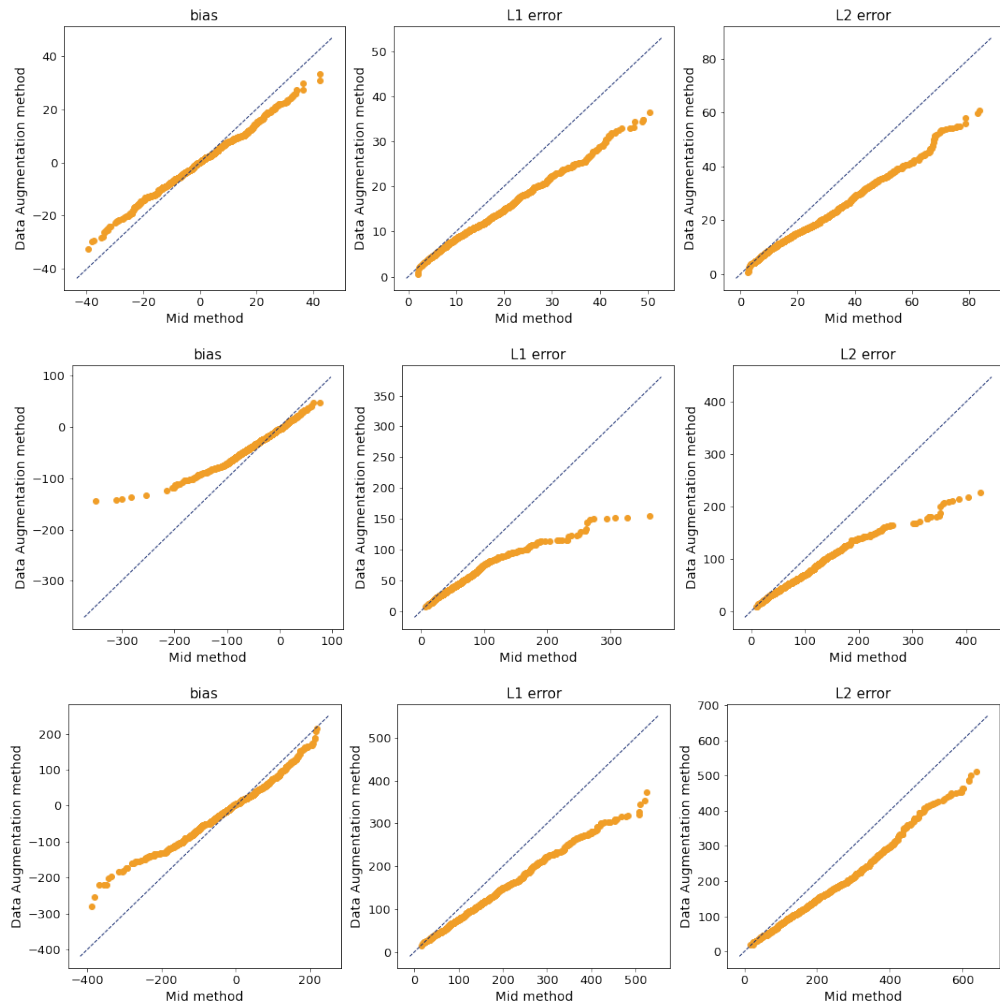


Figure 10: Q-Q plot of the calibration errors for all samples. Top row: *Use Case 1*, middle row: *Use Case 2*, bottom row: *Use Case 3*.

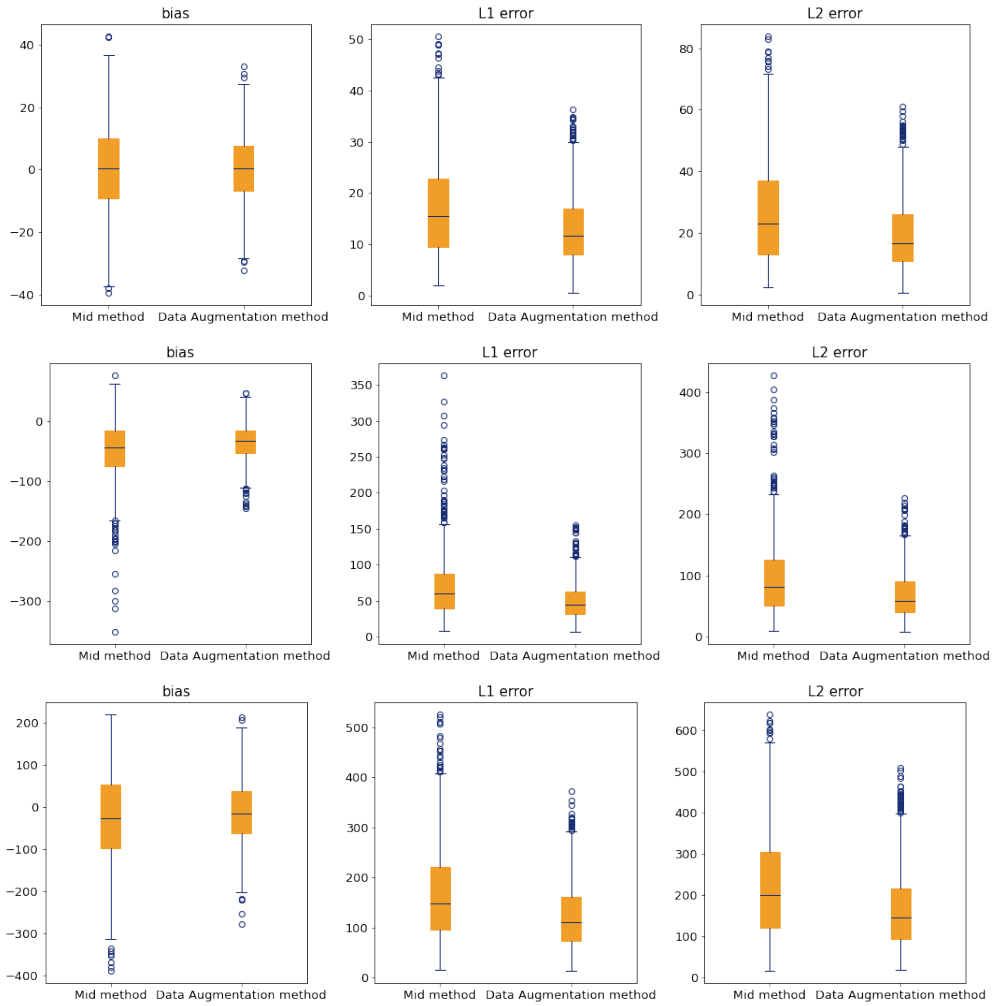


Figure 11: Box plots of the L1 and L2 errors for all samples. Top row: *Use Case 1*, middle row: *Use Case 2*, bottom row: *Use Case 3*.




Figure 12: Calibration of quotes for BTC-USD options with 2022/12/30 expiry. Quotes are extracted from the order book snapshots of June 22, 2022 at 17:07:00 CEST on the Deribit platform.




Figure 13: Calibration of quotes for BTC-USD options with 2022/09/30 expiry. Quotes are extracted from the order book snapshots of June 22, 2022 at 17:07:00 CEST on the Deribit platform.




Figure 14: Minimal implied volatility \underline{IV} as a function of k , when the minimal price is $n = 1$ tick on the left and $n = 2$ ticks on the right. Time-to-expiries τ are 2 Weeks, 1M, 2M, 3M, 6M, 1Y and are expressed in years.

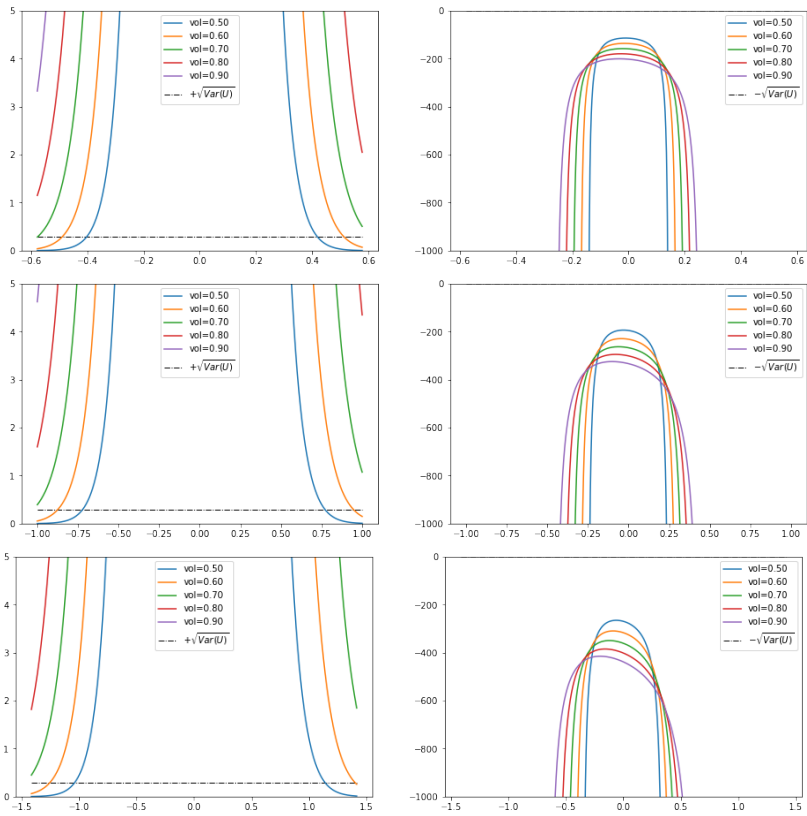


Figure 15: Behavior of $\sigma \mapsto \rho(\sigma)$, depending on the log-forward-moneyness k , when the spread in (A.3) is 1 tick. The time-to-expiry is 1 month at the top, 3 months in the middle, 6 months at the bottom.

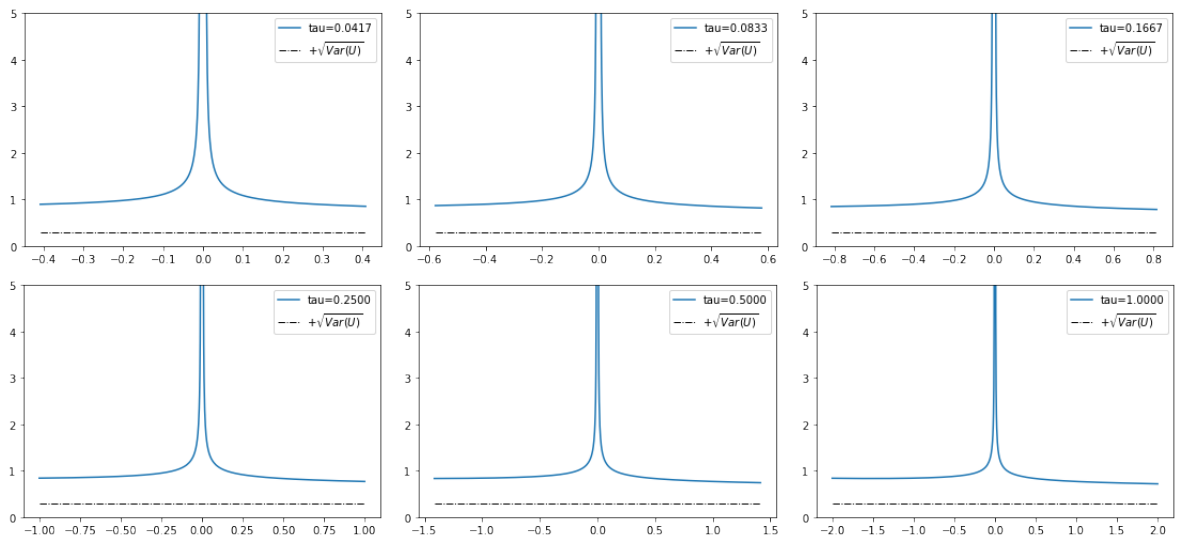


Figure 16: The function $k \mapsto \rho(k)$ when the input price is computed with the most extreme volatility $\underline{\text{IV}}(k)$ (for a minimal price of one tick). From top left to bottom right, different time-to-expiries τ in years (2W, 1M, 2M, 3M, 6M, 1Y).