

Explicit Rational Formulae for Bachelier (Normal) Implied Volatility

Fabien Le Floc’h 

Independent researcher; fabien@2ipi.com

Abstract

We present two explicit rational formulae for Bachelier, or normal, implied volatility. The formulae take the option price, forward, strike, and expiry as inputs and return the implied normal volatility without iteration. They follow the branch structure of LFK-4, but use the simpler near-the-money variable given by the absolute forward-strike difference divided by the tail time value, avoiding a logarithm and a small-argument Taylor branch in that region. LFK-2026 is the accuracy-oriented formula and approximates reciprocal absolute standardized moneyness directly in the far tail. LFK-2026C keeps the same shifted out-of-the-money rational tail approximation, but splits the near-the-money branch into a very small low- u rational and a mid-range rational. In double precision tests both remain close to machine accuracy, while LFK-2026C is the faster scalar implementation on the current benchmark mix.

Keywords: Bachelier model; normal implied volatility; bp vol; implied volatility; rational approximation; computational finance

1. Introduction

The Bachelier model, introduced in Bachelier’s 1900 thesis [1], is widely used when prices are naturally described by a normal volatility rather than a lognormal volatility. Its call price is explicit once the volatility is known, but the inverse problem—recovering the normal volatility from an observed option price—has no elementary closed form. A root finder is robust, but it is often too expensive for calibration, risk, and large grid calculations where the inverse must be evaluated millions of times.

This is not only a numerical issue. In interest-rate derivatives, normal or Bachelier volatilities are the usual convention when forwards, strikes, or rates can be close to zero or negative. Empirical work on low-rate swaption markets therefore considers Black, Bachelier, and shifted-lognormal volatility quotes across currencies including EUR, USD, GBP, and JPY [3]; in particular, the convention has long been familiar in Japanese-yen interest-rate markets. The same modelling issue briefly appeared outside rates in April 2020, when WTI crude oil futures traded below zero. CME Clearing then announced a switch of its options pricing and valuation model to Bachelier to accommodate negative underlying futures prices and negative strikes [4]; Choi, Kwak, Tee, and Wang [5] review this temporary commodity-exchange transition from Black–Scholes to Bachelier.

In interest-rate markets the same implied quantity is often quoted as *basis-point volatility*, or bp vol. This is not a separate volatility model: it is the normal, equivalently Bachelier, volatility expressed in the natural absolute units of the underlying. For a rate forward quoted as a decimal, a normal volatility σ corresponds to $10^4\sigma$ basis points per square-root year. This contrasts with Black or lognormal volatility, which is dimensionless and measures proportional rather than absolute moves.

We use the standard option terminology throughout: at the money (ATM) means $F = K$; for a call, in the money (ITM) means $F > K$, and out of the money (OTM) means $F < K$. The formulae are written with the OTM time value, so the same expressions also cover the reflected ITM side through put–call symmetry.

Explicit approximations fill this gap. The LFK-4 approximation [2] is a standard reference: it transforms the inverse problem into a small number of rational approximations, using one branch near the money and

three branches in the OTM tail. Related work includes the direct arithmetic Brownian-motion implied-volatility approximation of Choi, Kim, and Kwak [6]. That work is useful as a compact closed-form reference, but it is not aimed at the near-machine-epsilon accuracy sought by LFK-4 and by LFK-2026 here. For context, Jäckel's rational inversion of the Black–Scholes formula [7] shows the same high-accuracy philosophy in a different model; his normal implied-volatility approximation is also a popular production reference. QuantLib [8] includes an implementation in this family, but it is not a fully faithful transcription of Jäckel's original code and is less accurate. The timings below therefore compare the local Java implementations used in this study rather than the QuantLib port. The goal here is to keep the practical LFK-4 structure while improving the accuracy–cost trade-off. The two formulae below are designed to be portable: they use only arithmetic operations and a logarithm in the OTM routing variable. Both LFK-2026 and LFK-2026C avoid the final OTM square root by fitting $1/|d|$ directly.

Accurate validation of these formulae requires an accurate forward Bachelier price. Direct formulas for the OTM time value suffer cancellation in the tails, which can make an inverse approximation look worse than it is. For benchmarking we therefore compute the OTM time value with a symmetric erfcx expression; this makes the comparison primarily about the inverse formula rather than about loss of precision in the price used as input.

2. Setup and Notation

Let

$$x = F - K, \quad v = \sigma\sqrt{T}, \quad d = \frac{x}{v}.$$

The undiscounted Bachelier call price is

$$C = x\Phi(d) + v\phi(d).$$

Thus ATM corresponds to $m = |x| = 0$. Away from ATM, the call is ITM for $x > 0$ and OTM for $x < 0$; the reflected put side gives the same time value. As in LFK-4, the inverse is expressed in terms of the out-of-the-money time value and the absolute moneyness,

$$c_{\text{otm}} = \begin{cases} C - x, & x > 0, \\ C, & x \leq 0, \end{cases} \quad m = |x|.$$

At the money, where $m = 0$, the inverse is exact:

$$\sigma = \frac{C\sqrt{2\pi}}{\sqrt{T}}.$$

For $m > 0$, define

$$g = \frac{c_{\text{otm}}}{m}, \quad \alpha = 0.15, \quad \beta_s = -\ln \alpha, \quad \beta_e = 300 \ln(10).$$

The ITM/near-ATM branch is used when $g > \alpha$. The OTM branch uses

$$\tilde{\eta} = -\frac{\ln g + \beta_s}{\beta_e - \beta_s}.$$

The three OTM zones are

j	$\tilde{\eta}$ range
1	$[0, 0.011)$
2	$[0.011, 0.105)$
3	$[0.105, 1]$.

For all coefficient tables, rational functions are written in ascending powers:

$$R_{p,q}(y; a, b) = \frac{P_p(y)}{Q_q(y)}, \quad P_p(y) = \sum_{i=0}^p a_i y^i, \quad Q_q(y) = \sum_{i=0}^q b_i y^i, \quad b_0 = 1.$$

For numerical evaluation, the same polynomials may be evaluated by ordinary Horner steps in descending powers.

3. Reference: LFK-4

LFK-4 [2] uses

$$z = \frac{m}{m + c_{\text{otm}}}, \quad \eta = \frac{-z}{\ln(1 - z)}.$$

For small z , the original algorithm evaluates η by a Taylor expansion; otherwise it uses $\log 1p(-z)$. This avoids cancellation in the near-ATM limit. For $g > \alpha$, LFK-4 returns

$$\sigma = \frac{m + c_{\text{otm}}}{\sqrt{T}} h_{\text{LFK}}(\eta),$$

with a P7/Q5 rational. For $g \leq \alpha$, it fits d^2 with three P9/Q7 rationals in $\tilde{\eta}$, using the zones listed above.

4. LFK-2026

LFK-2026 uses the no-log ITM variable

$$u = \frac{z}{1 - z} = \frac{m}{c_{\text{otm}}}.$$

For $g > \alpha$, it uses a P10/Q9 rational:

$$\sigma = \frac{m + c_{\text{otm}}}{\sqrt{T}} R_{10,9}(u; a_I^{26}, b_I^{26}).$$

In the OTM branch, LFK-2026 approximates $1/|d|$ directly:

$$W_j^{26}(\tilde{\eta}) = R_{10,9}(\tilde{\eta}; a_j^{26}, b_j^{26}), \quad \sigma = \frac{m}{\sqrt{T}} W_j^{26}(\tilde{\eta}).$$

This removes the final square root in the OTM branch. The method uses more coefficients in the ITM branch than LFK-2026C, but the two methods share the same OTM rational approximation and the same direct OTM accuracy envelope.

5. LFK-2026C

LFK-2026C uses the same no-log ITM variable $u = m/c_{\text{otm}}$, but splits the ITM branch at $u = 0.20$. For small u it uses a compact P4/Q4 rational, and otherwise it uses a P9/Q8 rational:

$$\sigma = \frac{m + c_{\text{otm}}}{\sqrt{T}} \begin{cases} R_{4,4}(u; a_L^{26C}, b_L^{26C}), & u < 0.20, \\ R_{9,8}(u; a_I^{26C}, b_I^{26C}), & u \geq 0.20. \end{cases}$$

In the OTM branch it uses the same shifted-zone $1/|d|$ approximation as LFK-2026:

$$W_j^{26C}(\tilde{\eta}) = R_{10,9}(\tilde{\eta}; a_j^{26}, b_j^{26}), \quad \sigma = \frac{m}{\sqrt{T}} W_j^{26C}(\tilde{\eta}).$$

This is the fast variant. It keeps the OTM accuracy of LFK-2026 while reducing the average cost of the ITM/near-ATM branch.

6. Accuracy and Timing

The validation uses a symmetric erfcx expression for the OTM time value. With $a = m/v$,

$$c_{\text{otm}} = e^{-a^2/2} \left(\frac{v}{\sqrt{2\pi}} - \frac{m}{2} \operatorname{erfcx} \left(\frac{a}{\sqrt{2}} \right) \right),$$

which avoids cancellation in both deep ITM and deep OTM cases.

The accuracy table reports maximum relative implied-volatility error, $|\hat{\sigma}/\sigma - 1|$. Unless stated otherwise, the double precision tests use $F = 1$, $\sigma = 1$, and $T = 1$, with input prices generated by the stable erfcx formula above. The “double precision strike grid” row uses the strikes

$$K \in \{1.001, 1.01, 1.1, 1.5, 2, 3, 5, 7, 8, \\ 0.999, 0.99, 0.9, 0.5, 0, -0.5, -1\}.$$

The “deep ITM reflection” row starts from the deep ITM call strikes

$$K \in \{-1, -2, -3, -5, -7, -8, \\ -9, -10, -15, -20, -25, -30\},$$

computes their OTM time values, and feeds the reflected OTM calls with strike $2F - K$. The “direct OTM algorithm grid” row uses standardized moneyness values

$$d \in \{-8, -5, -3, -2, -1, -0.5, -0.1, 0, \\ 0.1, 0.5, 1, 2, 3, 5, 8\}.$$

The BigFloat rows use 1024-bit arithmetic and normalized OTM Bachelier time values on uniform grids in d : $|d| \leq 8$ and $|d| \leq 30$, both with step 0.025. In those rows the decimal coefficient literals are first rounded to binary64, then promoted to BigFloat, matching a portable double precision implementation. The adversarial row combines normalized OTM and reflected ITM cases on core values from $|d| = 0$ to $|d| = 37$, points immediately around the route and OTM zone boundaries, tail-cutoff checks down to subnormal prices, and scaled cases with $F \in \{-10, -0.1, 0, 1, 100\}$, $\sigma \in \{10^{-6}, 0.003, 1, 40, 10^6\}$, $T \in \{10^{-8}, 0.01, 1, 100, 10^8\}$, and $d \in \{-5, -2, -0.1, -10^{-8}, 0, 10^{-8}, 0.1, 2, 5\}$.

Table 1. Double precision test results based on 256 consecutive floating point numbers.

Test Configuration	Algorithm	Max Error	P99
Strike grid (Samples = 4096)	LFK-4	4.99×10^{-15}	3.89×10^{-15}
	LFK2026	4.22×10^{-15}	3.55×10^{-15}
	LFK2026C	4.22×10^{-15}	3.55×10^{-15}
	Jäckel 2017	4.22×10^{-15}	3.55×10^{-15}
Deep ITM reflection (Samples = 3072)	LFK-4	3.49×10^{-14}	2.82×10^{-14}
	LFK2026	3.40×10^{-14}	2.80×10^{-14}
	LFK2026C	3.40×10^{-14}	2.80×10^{-14}
	Jäckel 2017	3.44×10^{-14}	2.81×10^{-14}
Direct OTM grid (Samples = 3840)	LFK-4	3.53×10^{-14}	2.47×10^{-14}
	LFK2026	8.88×10^{-16}	6.66×10^{-16}
	LFK2026C	8.88×10^{-16}	6.66×10^{-16}
	Jäckel 2017	4.44×10^{-16}	2.22×10^{-16}

On direct high-precision inputs, Jäckel’s procedural formula has the smallest maximum error, while LFK-2026 and LFK-2026C are the most accurate of the explicit LFK-style rational approximations because they share the same shifted OTM approximation. For Jäckel, the BigFloat rows use a high-precision port of the original formula with the decimal constants first rounded to 64-bit floating point format. The few iterations

Table 2. Multiple precision test results based on 256 consecutive floating point numbers.

Test Configuration	Algorithm	Max Error	P99	Worst d
Dense grid $ d \leq 8$ Step 0.025 (Cases = 641)	LFK-4	1.18×10^{-15}	1.12×10^{-15}	-2.450
	LFK2026	1.14×10^{-16}	1.09×10^{-16}	-3.225
	LFK2026C	1.14×10^{-16}	1.09×10^{-16}	-3.225
	Jäckel	8.36×10^{-18}	7.53×10^{-18}	-1.125
Dense grid $ d \leq 30$ Step 0.025 (Cases = 2401)	LFK-4	1.18×10^{-15}	9.77×10^{-16}	-2.450
	LFK2026	1.21×10^{-16}	1.14×10^{-16}	-10.700
	LFK2026C	1.21×10^{-16}	1.14×10^{-16}	-10.700
	Jäckel	1.06×10^{-17}	9.36×10^{-18}	-30.000
Fixed tail/ATM points including ± 35 (Cases = 29)	LFK-4	9.77×10^{-16}	9.77×10^{-16}	-2.000
	LFK2026	1.03×10^{-16}	1.03×10^{-16}	-8.000
	LFK2026C	1.03×10^{-16}	1.03×10^{-16}	-8.000
	Jäckel	1.06×10^{-17}	1.06×10^{-17}	-30.000

used for polishing run on multiple precision arithmetic, hence the increased accuracy. The adversarial harness is dominated by ITM cases where the input call price is close to intrinsic; all four methods remain in the same practical accuracy envelope there.

Figure 1 shows the most diagnostic direct-accuracy view. The input is the normalized OTM Bachelier time value generated in 1024-bit arithmetic on the same $|d| \leq 30$, step-0.025 grid as the BigFloat table row, and the coefficient literals are first rounded to binary64 before being promoted for the validation. This removes call-price rounding and intrinsic-subtraction effects, so the plot shows the rational approximation error itself.




Figure 1. Direct BigFloat relative implied-volatility error for LFK-4, LFK-2026, and LFK-2026C. The left panel focuses on the central region $|d| \leq 8$, while the right panel shows the full direct grid $|d| \leq 30$. The dotted vertical lines mark the route boundary $g = \alpha$ and the two OTM zone boundaries.

The timing table below uses a rotated-order scalar benchmark and reports medians over fifty rounds. This is more stable than the fixed-order main report, whose single sequential timing slices can move when CPU frequency or scheduling changes during the run. Absolute timings are machine- and power-state-dependent; the relative ordering is the more reproducible part of the table.

Approximation	median ns/call	Relative to LFK-4
LFK-4	10.25	1.00
LFK-2026	8.05	0.79
LFK-2026C	7.42	0.72
Jäckel	25.22	2.46

LFK-2026 is the single-ITM-rational accuracy-first candidate. LFK-2026C keeps the same OTM accuracy and is the faster variant on this scalar benchmark.

Appendix G LFK-2026 Coefficients

Appendix G.1 ITM/near-ATM branch

i	$a_{l,i}^{26}$	$b_{l,i}^{26}$
0	2.50662827463100069e+00	1.0000000000000000e+00
1	8.26914966237441540e+00	3.79891342328838499e+00
2	1.10083895644891214e+01	5.87074621959482457e+00
3	7.62695942399991100e+00	4.76157470081986212e+00
4	2.96457253391146036e+00	2.18581925252758946e+00
5	6.51812439800918964e-01	5.72905243689549204e-01
6	7.81257091685455263e-02	8.27294623127229206e-02
7	4.67776338554095131e-03	6.05262799121958784e-03
8	1.17427118047196583e-04	1.90358881940080979e-04
9	8.06959950534549627e-07	1.76070970713414280e-06
10	-3.39451638335349906e-10	

Appendix G.2 OTM branch, zone 1

i	$a_{l,i}^{26}$	$b_{l,i}^{26}$
0	1.24910559446641112e+00	1.0000000000000000e+00
1	8.30013684602045146e+02	9.50178311644246946e+02
2	2.89118707775471907e+05	4.30671117130989209e+05
3	6.20616261157058701e+07	1.19319078802154481e+08
4	8.81948242946073532e+09	2.18139815883858109e+10
5	8.14953597568851318e+11	2.66113331187980811e+12
6	4.61473754119971406e+13	2.08194332615256938e+14
7	1.32385751716178975e+15	9.31612904921204000e+15
8	1.35683441709766740e+16	1.78457095251951296e+17
9	2.23602102096865640e+16	9.01018167981477888e+17
10	-1.15188057059215700e+16	

Appendix G.3 OTM branch, zone 2

i	$a_{2,i}^{26}$	$b_{2,i}^{26}$
0	1.25087969033276813e+00	1.0000000000000000e+00
1	2.09344504116440078e+02	4.55619235679538008e+02
2	-1.74963499510044603e+04	4.19868074702521844e+04
3	-1.28046460022976995e+07	-1.48653150107890852e+07
4	-1.43888265763526964e+09	-4.63359666129231930e+09
5	-5.70761754440745773e+10	-3.39232384615881042e+11
6	-8.98414557472468994e+11	-8.93831987833861328e+12
7	-5.46607883881420703e+12	-9.10822517732905625e+13
8	-1.07736581582367109e+13	-3.32157260465782750e+14
9	-3.54669640321780615e+12	-3.16204789033010312e+14
10	4.53512764886659485e+11	

Appendix G.4 OTM branch, zone 3

i	$a_{3,i}^{26}$	$b_{3,i}^{26}$
0	1.61713874667576762e+00	1.0000000000000000e+00
1	1.63839367097254751e+02	6.08432004257079598e+02
2	4.52092618390189637e+03	3.37041720211591382e+04
3	4.73452995425928821e+04	5.83977840443553636e+05
4	2.11536138072810922e+05	4.03973802769196732e+06
5	4.14000232322855853e+05	1.20766363495907933e+07
6	3.38490556724923430e+05	1.55296990844987314e+07
7	9.82243936325080576e+04	7.83292053799574263e+06
8	5.74901275345015438e+03	1.16810498733679834e+06
9	-1.25257952774401772e+02	5.54442814599942540e+03
10	4.62144628906322019e+00	

Appendix H LFK-2026C Coefficients

Appendix H.1 Low- u ITM/near-ATM branch, $u < 0.20$

i	$a_{L,i}^{26C}$	$b_{L,i}^{26C}$
0	2.50662827463100069e+00	1.0000000000000000e+00
1	4.33803460770355898e+00	2.23062541885758314e+00
2	2.36504153278287266e+00	1.63840524330855919e+00
3	4.35189695439266722e-01	4.35646805136993498e-01
4	1.77075775698172025e-02	2.96543644408861981e-02

Appendix H.2 High- u ITM/near-ATM branch, $u \geq 0.20$

i	$a_{I,i}^{26C}$	$b_{I,i}^{26C}$
0	2.50662827462874782e+00	1.0000000000000000e+00
1	7.07601239545194982e+00	3.32292052116718439e+00
2	7.79838748842165330e+00	4.35214422048696914e+00
3	4.28317532357895026e+00	2.86841498456426436e+00
4	1.24534893861795548e+00	1.01064393092251747e+00
5	1.88932466067008031e-01	1.87748433761637329e-01
6	1.38619741425921022e-02	1.70940757504186268e-02
7	4.15460128562968034e-04	6.51089901103861003e-04
8	3.33224453410284549e-06	7.12577601891527465e-06
9	-1.65051672752885574e-09	

LFK-2026C uses the same OTM coefficients as LFK-2026 in all three shifted $\tilde{\eta}$ zones.

Acknowledgments: This paper would not exist without the initial idea from Gary Kennedy.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Bachelier, L. Théorie de la spéculation. *Annales Scientifiques de l'École Normale Supérieure* **1900**, *17*, 21–86.
2. Le Floch, F. Fast and Accurate Analytic Basis Point Volatility. Available online: <https://ssrn.com/abstract=2420757> (accessed on 17 May 2026).
3. Patel, J.; Russo, V.; Fabozzi, F.J. Using the Right Implied Volatility Quotes in Times of Low Interest Rates: An Empirical Analysis across Different Currencies. *Finance Research Letters* **2018**, *25*, 196–201. <https://doi.org/10.1016/j.frl.2017.10.013>.
4. CME Clearing. Switch to Bachelier Options Pricing Model—Effective April 22, 2020. CME Clearing Advisory 20-171, 21 April 2020. Available online: <https://www.cmegroup.com/notices/clearing/2020/04/Chadv20-171.html> (accessed on 18 May 2026).
5. Choi, J.; Kwak, M.; Tee, C.W.; Wang, Y. A Black–Scholes User’s Guide to the Bachelier Model. *Journal of Futures Markets* **2022**, *42*, 959–980. <https://doi.org/10.1002/fut.22315>.
6. Choi, J.; Kim, K.; Kwak, M. Numerical Approximation of the Implied Volatility under Arithmetic Brownian Motion. *Applied Mathematical Finance* **2009**, *16*, 261–268.
7. Jäckel, P. Let’s Be Rational. *Wilmott* **2017**, *2015*, 40–53. <https://doi.org/10.1002/wilm.10395>.
8. QuantLib Project. QuantLib: A Free/Open-Source Library for Quantitative Finance. Available online: <https://www.quantlib.org/> (accessed on 18 May 2026).