WORKING PAPER

Options Under a Defended Currency Floor

Stickiness, Break Risk and the EUR/CHF Episode, 2011–2015

Supervised by Dr. T. Zamrik

21-MAY-2024

Phase portrait of the FitzHugh–Nagumo model: $\dot{x} = x - \frac{x^3}{3} - y + I,\ \dot{y} = \varepsilon\,(x + a - b\,y),\quad a = 0.68,\ b = 0.79,\ I = 0.37,\ \varepsilon = 0.09$



WP-2024-33393745
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2024 T. Zamrik. All rights reserved. Published by Zamrik Quant Research. No part of this document may be reproduced, stored, distributed or adapted, in whole or in part, by any means, without the prior written permission of the publisher, except for brief quotation with full citation. Requests for permission: info@zamrik.com.

Suggested Citation

Derivatives & Volatility Desk, Zamrik Quant Research (2024). 'Options Under a Defended Currency Floor: Stickiness, Break Risk and the EUR/CHF Episode, 2011–2015'. *Zamrik Quant Research Working Paper* **WP-2024-33393745**. Available at https://zamrik.com/research-items/wp-2024-33393745/.

About Zamrik Quant Research

Zamrik Quant Research is the research arm of Zamrik Quant Lab. It works through five research desks, from market microstructure and derivatives to optimal control, stochastic analysis and applied mathematics, and publishes everything it finds in the open: research papers, notebooks that run on Kaggle, free tools that answer one practical question, and essays.

The five desks are Market Microstructure & Trading; Derivatives & Volatility; Optimal Control & Decision; Stochastic Analysis; and Applied Mathematics. Each paper, notebook, tool and essay belongs to exactly one desk. Zamrik Quant Research is run by a small, independent team, headed by Dr. T. Zamrik. We are self-funded: no sponsors, no advertisers, no outside capital.

Every publication passes through the same five functions before anyone outside sees it: research, data, software, verification and production. We never retrofit the mathematics to a conclusion chosen in advance. Every number is checked two independent ways. Every figure is re-run from its code. When we find an error, or replace a piece with a better one, we record it in the public corrections log at https://zamrik.com/corrections/.

Legal Notice

This paper is research and education. Nothing in it is investment, financial, legal or engineering advice; nothing is a recommendation to buy or sell shares of stock, derivative contracts, bonds, or any other instrument or property; and nothing is a forecast. You are responsible for your own decisions. Figures and code are provided as they are. Care and independent checks reduce errors but cannot rule them out, so no warranty is given. The views expressed are those of the Derivatives & Volatility Desk.

DERIVATIVES & VOLATILITY DESK ZAMRIK QUANT RESEARCH



WORKING PAPER

Options Under a Defended Currency Floor

Stickiness, Break Risk and the EUR/CHF Episode, 2011–2015

Supervised by Dr. T. Zamrik[1]

21-MAY-2024 WP-2024-33393745

Abstract. When a central bank defends a currency floor, the exchange rate does not bounce off it: it sits on it. We model the log distance to the floor as a diffusion that is sticky at zero, pushed up at rate θ while it sits there, and let the bank abandon the floor at hazard κ per unit of floor time, the rate jumping by e^J. The discounted rate is a martingale exactly when the drift above the floor is the carry less half the variance and $\kappa(1-e^J)=\theta-q$, with q the interest differential. Because θ is a path functional, local time over twice the floor time, it is the same under every equivalent measure, and the integrated break hazard equals the floor's push net of carry divided by the relative jump. Calls, puts and digitals are explicit Laplace transforms; they agree with finite differences to 1.2×10^{-5} and with Monte Carlo within 2.83 standard errors, and the floor puts a steep downside skew on options. On 858 ECB fixings of the EUR/CHF floor of 2011 to 2015 the stickiness is weakly identified, 0.20 to 1.71 per year across bands, yet the risk-neutral twelve-month probability of a rate below 1.20 from the floor stays between 16.9 and 18.7 per cent with the realized jump to 1.028. With the realignment to 1.15 that Hertrich and Zimmermann describe, the same path gives 35.6 per cent, inside their option-implied range: the path prices the product of break probability and jump size, not either factor alone.

WP-2024-33393745 · Derivatives & Volatility Desk · Working Paper · Published 21-MAY-2024 · Reading time: about 42 min

1. Introduction

On 6 September 2011 the Swiss National Bank set a minimum exchange rate of 1.20 Swiss francs per euro, and on 15 January 2015 it discontinued it. The European Central Bank's reference rate for the franc printed 1.2010 on 14 January 2015 and 1.0280 the next day, 14.3 per cent below the floor. In between, the rate spent long stretches just above 1.20 (Figure 3): over 858 daily fixings it never printed 1.2000, its lowest fixing was 1.2008, and 63 fixings sat at or below 1.2010. The floor was not a wall the rate touched and left. It was a place where the rate stayed.

Option markets priced the floor throughout. Hertrich and Zimmermann [5] extract risk-neutral probabilities that the rate falls below 1.20 from EUR/CHF option prices, and Jermann [6] reads the market's view of the floor's survival from the same market. The episode is the cleanest natural experiment we know for a defended floor: a known level, a

[1]Correspondence: Derivatives & Volatility Desk, Zamrik Quant Research, info@zamrik.com.

known date of abandonment, a measured jump, and a liquid option market on both sides of the break.

This paper asks three questions. When the rate sticks to a defended floor that can break, what does the absence of arbitrage say about the risk of the break? How are options priced? And what did the EUR/CHF floor of 2011 to 2015 imply, read from its path alone?

Two models of the EUR/CHF floor are priced on option data. Hertrich and Zimmermann [5] value options with a reflecting lower barrier, using the formulas of Veestraeten [8] with the correction of Hertrich and Veestraeten [9], and imply a barrier that sits below 1.20; they write that the rate "only prevails for an infinitely short time period on the floor in order to avoid arbitrage". Ingersoll [10] values foreign-exchange derivatives when the exchange process is bounded. Jermann [6] works in discrete time, with the observed rate the larger of a fundamental rate and the floor, and a constant probability that the floor continues. In the first model the rate leaves the floor at once; in the second it rests on the floor whenever the fundamental rate is below it, but in discrete time, with no diffusion and no boundary condition.

The target-zone literature begins with Krugman [11]. Lera and Sornette [7] model EUR/CHF near the floor with target-zone dynamics in which volatility falls towards the floor; their model has no break risk and prices no options.

A diffusion that spends positive time at a boundary is sticky. Feller [2] classified the boundary conditions of one-dimensional diffusions, stickiness among them; Itô and McKean [3] built sticky Brownian motion on a half line by a time change of reflected Brownian motion; Engelbert and Peskir [1] showed that its stochastic differential equation has a weak solution unique in law and no strong solution. In finance, Nie and Linetsky [4] study a Vasicek short rate with a sticky zero lower bound. Anagnostakis [17] studies a market whose risky asset is a sticky geometric Brownian motion and shows that it is free of arbitrage if and only if the interest rate is zero. None of these models has a boundary that can be abandoned. In the work we have read, no model prices options under a sticky floor that can break, or links the stickiness to the break risk through the absence of arbitrage.

We model the log distance to the floor as a diffusion that is sticky at zero, with a push rate θ while it sits there, and let the floor break at hazard κ per unit of floor time with a deterministic log jump J. The model is set out in the next section, with its construction by a time change and its generator. The section after it proves the no-arbitrage relation and its consequences, and the one after that prices calls, puts and digitals three ways. We then show what the floor does to the implied-volatility smile and to the term structure of break probabilities, calibrate the model to the ECB fixings of the EUR/CHF floor, set the result beside the option-implied numbers of Hertrich and Zimmermann, and close with what the calibration does not settle. Nothing here is advice on trading; it is pricing mathematics and a historical calibration.

Main findings

- No arbitrage holds exactly when $\kappa_Q(1-e^J) = \theta - q$, where the push rate $\theta = L_t/(2A_t)$

is a path functional and so the same under every equivalent measure, which makes the integrated break hazard equal to the floor's push net of carry divided by the relative jump (Theorem 1, Corollary 1).

- Calls, puts and digitals are explicit Laplace transforms that agree with an independent finite-difference solver to 1.2×10^{-5} and with Monte Carlo within 2.83 standard errors over three seeds, while a wrong hazard fails the martingale test at $|z| \geq 214$ (Proposition 2).
- The floor puts a steep downside skew on options: one month from a spot of 1.21, implied volatility is 18.9 per cent at a strike of 1.10 against 4.1 per cent at the money, and puts struck at or below the floor are pure break-risk claims (Example 2, Proposition 4).
- On the ECB fixings of 2011 to 2015 the stickiness is weakly identified, 0.20 to 1.71 per year across bands, but the risk-neutral twelve-month probability of a rate below 1.20 from the floor is not: it stays between 16.9 and 18.7 per cent (block bootstrap 14.6 to 19.3 per cent).
- The path prices the product of break probability and jump size: with the realized jump to 1.028 it implies 17.5 per cent at twelve months, well below Hertrich and Zimmermann's twelve-month estimate, which does not fall much below 30 per cent, and with their realignment to 1.15 it implies 35.6 per cent, inside their range.

2. The Model

Definition 1 (The sticky floor and its abandonment)**.** Let S be the EUR/CHF rate, S_t Swiss francs per euro at time t in years, let $F = 1.20$ be the floor and $X_t = \log(S_t/F)$ the log distance to it. Let r_d and r_f be the Swiss franc and euro interest rates and $q = r_d - r_f$ the interest differential. Before the break,

$$dX_t = \mathbf{1}\{X_t > 0\}\,(\mu\,dt + \sigma\,dW_t) + \theta\,\mathbf{1}\{X_t = 0\}\,dt, \qquad X_t \geq 0,$$

with W a Brownian motion, drift μ and volatility $\sigma > 0$ above the floor, and push rate $\theta \in (0, \infty)$ on it. The floor time is $A_t = \int_0^t \mathbf{1}\{X_s = 0\}\,ds$ and the floor's push is $\ell_t = \theta A_t$. The floor is abandoned at $\tau = \inf\{t : A_t > E\}$, where E is exponential with rate $\kappa \geq 0$ and independent of W ($E = \infty$ when $\kappa = 0$), so that κ is the hazard of a break per unit of floor time. At τ the rate jumps to $S_\tau = Fe^J$ with $J < 0$; afterwards $\log S_t = \log F + J + \nu\,(t - \tau) + \sigma_p\,(B_t - B_\tau)$, with drift ν, volatility $\sigma_p > 0$ and a Brownian motion B independent of W and E.

The convention is the one used throughout: one sticky equation, with the push written as a drift on the floor. The limit $\theta \to \infty$ is instant reflection and $\theta \to 0$ is absorption. Since A grows only while $X = 0$, the break always comes from the floor: $S_{\tau-} = F$.

Assumption 1 (Standing assumptions)**.** Trading is frictionless and continuous in the Swiss franc and euro money markets, at constant rates r_d and r_f. The Swiss franc account

is the numeraire, so the traded asset is one euro on deposit, worth $S_t e^{r_f t}$ francs, and its discounted value is $e^{-qt}S_t$. The parameters $\mu, \sigma, \theta, \kappa, J, \nu, \sigma_p$ of Definition 1 are constants, the jump J is deterministic, and prices are expectations under a measure Q equivalent to P on $\mathcal{F}_T$ for a finite horizon T.

Constant rates and a deterministic jump are the simplest setting that keeps the mechanism visible. Corollary 2 lets the jump be random, and the calibration reports how the results move with the rates and with the jump.

Lemma 1 (Construction by a time change, and the push as a local time)**.** Let β be a Brownian motion, $x_0 \geq 0$, $Y_s = x_0 + \mu s + \sigma\beta_s$, and let $R_s = Y_s + \ell^R_s$ be its reflection at zero, with regulator $\ell^R_s = \max\{0, \max_{u\leq s}(-Y_u)\}$. Put $T(s) = s + \ell^R_s/\theta$, let $s(t)$ be its inverse and $X_t = R_{s(t)}$. Then: (a) X solves the equation of Definition 1 for a Brownian motion W, with floor time $A_t = \ell^R_{s(t)}/\theta$ and $s(t) = t - A_t$; (b) for every solution of that equation, $\ell_t = \theta A_t = \frac{1}{2}L_t$, where L_t is the right local time of X at zero, and $\langle X\rangle_t = \sigma^2(t - A_t)$; hence

$$\theta = \frac{L_t}{2A_t} \quad \text{whenever } A_t > 0;$$

(c) with E independent of β, the break time is $\tau = T(\varsigma)$ with $\varsigma = \inf\{s : \ell^R_s > \theta E\}$: the floor breaks when the regulator first exceeds θE.

Proof. T is continuous and strictly increasing with $T(0) = 0$ and $T(s) \geq s$, so its inverse $s(t)$ is continuous and strictly increasing, with $s(t) \leq t$. The regulator grows only on $\{R = 0\}$, and that set has Lebesgue measure zero, so

$$A_t = \int_0^t \mathbf{1}\{X_r = 0\}\, dr = \int_0^{s(t)} \mathbf{1}\{R_u = 0\}\, (du + \theta^{-1} d\ell^R_u) = \frac{\ell^R_{s(t)}}{\theta},$$

and $t = T(s(t))$ gives $s(t) = t - A_t$. Hence $X_t = x_0 + \mu\, s(t) + \sigma\beta_{s(t)} + \ell^R_{s(t)} = x_0 + \mu\,(t - A_t) + \sigma\beta_{s(t)} + \theta A_t$, with $t - A_t = \int_0^t \mathbf{1}\{X_r > 0\}\, dr$. Each $s(t)$ is a stopping time of β, since $\{s(t) \leq u\} = \{T(u) \geq t\}$, so $M_t = \beta_{s(t)}$ is a continuous local martingale with $\langle M\rangle_t = s(t) = \int_0^t \mathbf{1}\{X_r > 0\}\, dr$. With an independent Brownian motion γ, put $W_t = M_t + \int_0^t \mathbf{1}\{X_r = 0\}\, d\gamma_r$. Then $\langle W\rangle_t = t$, so W is a Brownian motion by Lévy's characterization, and $\int_0^t \mathbf{1}\{X_r > 0\}\, dW_r = M_t$ because $\int \mathbf{1}\{X > 0\}\, dM = M$. This is the equation of Definition 1. (b) Let X be any solution. It is a nonnegative continuous semimartingale, and Tanaka's formula for the right local time gives $X_t = X^+_t = X_0 + \int_0^t \mathbf{1}\{X_r > 0\}\, dX_r + \frac{1}{2}L_t$. The push term carries the factor $\mathbf{1}\{X = 0\}$, so $\int_0^t \mathbf{1}\{X_r > 0\}\, dX_r = \int_0^t \mathbf{1}\{X_r > 0\}(\mu\, dr + \sigma\, dW_r)$. Comparing with the equation, $\theta A_t = \frac{1}{2}L_t$. The quadratic variation is $\sigma^2 \int_0^t \mathbf{1}\{X_r > 0\}\, dr = \sigma^2(t - A_t)$. (c) $A_t > E$ if and only if $\ell^R_{s(t)} > \theta E$, and s is an increasing bijection of $[0, \infty)$. □

Itô and McKean [3] built sticky Brownian motion on a half line by this time change. Engelbert and Peskir [1] establish existence, uniqueness in law and the absence of a strong solution for the stochastic differential equations of sticky Brownian motion. We need less: the verification argument of Proposition 2 uses only the equation, so the prices below are

the same for every solution, and the Monte Carlo scheme of the prices section is built on part (c).

Proposition 1 (The generator and the boundary condition at the floor)**.** Let $f \in C^2([0,\infty))$ with f, f' and f'' bounded, and let $\kappa \geq 0$. Under Definition 1,

$$e^{-\kappa A_t} f(X_t) - f(X_0) - \int_0^t e^{-\kappa A_r}\, \mathcal{G}_\kappa f(X_r)\, dr$$

is a martingale, where $\mathcal{G}_\kappa f(x) = \mu f'(x) + \frac{1}{2}\sigma^2 f''(x)$ for $x > 0$ and $\mathcal{G}_\kappa f(0) = \theta f'(0) - \kappa f(0)$; and $\mathbb{E}_x[f(X_t)\,\mathbf{1}\{t < \tau\}] = \mathbb{E}_x[e^{-\kappa A_t} f(X_t)]$. The function $\mathcal{G}_\kappa f$ is continuous at the floor exactly when

$$\mu f'(0) + \tfrac{1}{2}\sigma^2 f''(0) = \theta f'(0) - \kappa f(0),$$

the Feller–Wentzell condition. In Feller's form $p_1 f(0) - p_2 f'(0) + p_3\, \mathcal{G} f(0) = 0$ [2], with $\mathcal{G} f(0) = \mu f'(0) + \frac{1}{2}\sigma^2 f''(0)$, the weights are proportional to $(\kappa, \theta, 1)$. In natural scale, $s'(x) = e^{-2\mu x/\sigma^2}$, the speed measure has density $2/(\sigma^2 s'(x))$ on $(0,\infty)$ and an atom $m(\{0\}) = 1/\theta$ at the floor.

Proof. By Itô's formula and $d\langle X\rangle_t = \sigma^2 \mathbf{1}\{X_t > 0\}\, dt$ (Lemma 1),

$$df(X_t) = \mathbf{1}\{X_t > 0\}(\mu f' + \tfrac{1}{2}\sigma^2 f'')(X_t)\, dt + \theta f'(0)\, \mathbf{1}\{X_t = 0\}\, dt + \sigma f'(X_t)\, \mathbf{1}\{X_t > 0\}\, dW_t,$$

and $de^{-\kappa A_t} = -\kappa\, e^{-\kappa A_t}\, \mathbf{1}\{X_t = 0\}\, dt$. The product rule gives the drift $e^{-\kappa A_t}\mathcal{G}_\kappa f(X_t)$; the stochastic integral has a bounded integrand and is a martingale. Since E is independent of X, $P(\tau > t \mid X) = P(E \geq A_t \mid X) = e^{-\kappa A_t}$, which gives the second identity. The limit of $\mathcal{G}_\kappa f(x)$ as $x \downarrow 0$ is $\mu f'(0) + \frac{1}{2}\sigma^2 f''(0)$, so continuity at the floor is the stated condition, which reads $\kappa f(0) - \theta f'(0) + \mathcal{G} f(0) = 0$. For the speed measure, write $D_s f = f'/s'$ and $dm = 2\, dx/(\sigma^2 s')$ on $(0,\infty)$. As $s''/s' = -2\mu/\sigma^2$,

$$D_m D_s f = \frac{\sigma^2 s'}{2}\Big(\frac{f'}{s'}\Big)' = \tfrac{1}{2}\sigma^2 f'' + \mu f'$$

on $(0,\infty)$. At an atom, $D_m D_s f(0) = D_s f(0+)/m(\{0\}) = f'(0)/m(\{0\})$, with nothing below the floor; equating with $\theta f'(0)$, the generator at the floor when $\kappa = 0$, gives $m(\{0\}) = 1/\theta$. □

For the free diffusion, zero is a regular boundary and infinity is natural. The condition says what happens at the floor: the rate is pushed up at rate θ, and is killed, to restart after the jump, at rate κ. The finite-difference solver of the prices section discretizes exactly this generator, with this condition as its boundary row.

3. No Arbitrage on the Floor

Theorem 1 (The floor's push prices the break)**.** Under Assumption 1, let Q be a measure equivalent to P on $\mathcal{F}_T$ under which the rate has the form of Definition 1 with drift μ_Q, hazard κ_Q and post-break drift ν_Q, all constant. Then $e^{-qt} S_t$, $0 \leq t \leq T$, is a Q-martingale

if and only if

$$\mu_Q = q - \tfrac{1}{2}\sigma^2, \qquad \kappa_Q\,(1 - e^J) = \theta - q, \qquad \nu_Q = q - \tfrac{1}{2}\sigma_p^2,$$

the last only when $\kappa_Q > 0$. The push rate θ and the volatilities σ and σ_p are the same under P and under every such Q, because $\theta = L_t/(2A_t)$ and $d\langle X\rangle_t = \sigma^2 \mathbf{1}\{X_t > 0\}\,dt$ are path functionals. The hazard κ_Q is therefore fixed by θ, q and J, while the real-world hazard κ_P is free; their ratio κ_Q/κ_P is the market price of break risk.

Proof. Before the break $S_t = Fe^{X_t}$, and Itô's formula with Lemma 1 gives

$$\frac{dS_t}{S_t} = \mathbf{1}\{X_t > 0\}((\mu_Q + \tfrac{1}{2}\sigma^2)\,dt + \sigma\,dW_t^Q) + \theta\,\mathbf{1}\{X_t = 0\}\,dt,$$

with W^Q a Q-Brownian motion. The break comes from the floor, where $S_{\tau-} = F$, and moves the rate by $S_{\tau-}(e^J - 1)$. Its counting process $N_t = \mathbf{1}\{\tau \le t\}$ has Q-compensator $\int_0^t \kappa_Q\,\mathbf{1}\{X_r = 0,\ r < \tau\}\,dr$. After the break, $dS_t/S_t = (\nu_Q + \tfrac{1}{2}\sigma_p^2)\,dt + \sigma_p\,dB_t^Q$. Collecting terms,

$$d(e^{-qt}S_t) = e^{-qt}S_{t-}\,b_t\,dt + dM_t,$$

$$b_t = \mathbf{1}\{t < \tau\}\Big[\mathbf{1}\{X_t > 0\}(\mu_Q + \tfrac{1}{2}\sigma^2 - q) + \mathbf{1}\{X_t = 0\}(\theta + \kappa_Q(e^J - 1) - q)\Big] + \mathbf{1}\{t \ge \tau\}(\nu_Q + \tfrac{1}{2}\sigma_p^2 - q),$$

where M is a Q-local martingale: the two Brownian integrals and the compensated jump $e^{-qt}S_{t-}(e^J - 1)(dN_t - \kappa_Q\mathbf{1}\{X_t = 0, t < \tau\}\,dt)$. If $e^{-qt}S_t$ is a Q-martingale, then $\int_0^t e^{-qr}S_{r-}\,b_r\,dr$ is a continuous local martingale of finite variation, hence zero, and $b = 0$ almost everywhere for $dt \times dQ$. The set $\{X > 0, t < \tau\}$ has positive measure, which gives the first equality. So does $\{X = 0, t < \tau\}$: by Lemma 1, $A_t = L_t/(2\theta) > 0$ as soon as the floor is reached, and the floor is reached before τ and T with positive Q-probability, since Q is equivalent to P. This gives the second equality. When $\kappa_Q > 0$ the break comes before T with positive probability, and the last bracket gives the third. Conversely, under the three equalities $b = 0$ and $e^{-qt}S_t$ is a local martingale. It is a true martingale because $\mathbb{E}_Q \sup_{t\le T} S_t < \infty$. Before the break, $X_t \le x_0 + (|\mu_Q| + \theta)\,T + \sigma \sup_{u\le T} |\int_0^u \mathbf{1}\{X_r > 0\}\,dW_r^Q|$, and the stochastic integral is a Brownian motion run on the clock $t - A_t \le t$, so its supremum has Gaussian tails; after it, S is Fe^J times a geometric Brownian motion, whose supremum over $[0, T]$ is integrable. Finally, invariance. By Girsanov's theorem a change to an equivalent measure turns W into W^Q plus a drift $\int \phi\,dr$, which changes the drift of X by $\sigma\phi\,\mathbf{1}\{X > 0\}$ only, because $d\langle X, W\rangle = \sigma\,\mathbf{1}\{X > 0\}\,dt$ vanishes on the floor. The push term θA_t is a pathwise process of finite variation and is not touched. The functionals $L_t/(2A_t)$ and $\langle X\rangle_t$ are almost sure limits of functionals of the path, and P and Q have the same null sets. The jump J is deterministic. Hence θ, σ, σ_p and J are common to P and Q, and κ_Q is fixed. By the fundamental theorem of asset pricing [12], an equivalent martingale measure of this form rules out a free lunch with vanishing risk. □

A floor's push cannot be removed by a change of measure; it can only be offset by the compensator of a downward jump. A positive hazard solves the relation only if $\theta >$

q, which holds whenever $q \le 0$. The theorem fixes the hazard per unit of floor time, not the probability of a break in calendar time, which depends on how much floor time accumulates; Corollary 1 makes this precise.

Corollary 1 (The integrated hazard is the push, and the reflecting limit)**.** Under the conditions of Theorem 1, for $t \le \tau$,

$$\kappa_Q A_t = \frac{\ell_t - qA_t}{1 - e^J}.$$

Let $\theta \to \infty$ with $\kappa_Q = (\theta - q)/(1 - e^J)$ and $\mu_Q = q - \frac{1}{2}\sigma^2$. Then for $x_0 \ge 0$ and $T > 0$,

$$Q(\tau \le T) \to 1 - \mathbb{E}\exp\Big(-\frac{(M_T - x_0)^+}{1 - e^J}\Big), \qquad M_T = \max_{s \le T}(-\mu_Q s + \sigma\beta_s),$$

with β a Brownian motion: under instant reflection, the floor breaks when the reflection regulator exceeds $1 - e^J$ times a unit exponential variable. With the drift set to zero and $x_0 = 0$, the limit is

$$1 - 2e^{c^2/2}\,\Phi(-c), \qquad c = \frac{\sigma\sqrt{T}}{1 - e^J}.$$

Proof. The identity is Theorem 1 multiplied by A_t, with $\ell_t = \theta A_t$. For the limit, the law of τ is the same for every solution of the equation (its Laplace transform is given in Proposition 2), so we may compute it on the construction of Lemma 1, with one Brownian motion β for every θ and $E = E_1/\kappa_Q$, E_1 a unit exponential variable independent of β. Then $\tau \le T$ if and only if $\kappa_Q A_T \ge E_1$, up to a null event. On the construction, $A_T = \ell^R_{s(T)}/\theta$ and $T - s(T) = A_T \le \ell^R_T/\theta$, so $s(T) \to T$ as $\theta \to \infty$, and

$$\kappa_Q A_T = \frac{\theta - q}{\theta}\,\frac{\ell^R_{s(T)}}{1 - e^J} \;\to\; \frac{\ell^R_T}{1 - e^J}$$

almost surely, by continuity of ℓ^R. Since E_1 has no atom, dominated convergence gives

$$Q(\tau \le T) \to Q(\ell^R_T \ge (1 - e^J)E_1) = 1 - \mathbb{E}\exp(-\ell^R_T/(1 - e^J)).$$

Finally $\ell^R_T = \max\{0, \max_{s \le T}(-x_0 - \mu_Q s - \sigma\beta_s)\}$, which is $(M_T - x_0)^+$, with $-\beta$ in place of β. With zero drift and $x_0 = 0$, M_T has the law of $\sigma\sqrt{T}\,|N|$, N standard normal, and

$$\mathbb{E}\,e^{-c|N|} = 2\int_0^\infty e^{-cz}\,\varphi(z)\,dz = 2e^{c^2/2}\int_0^\infty \varphi(z + c)\,dz = 2e^{c^2/2}\,\Phi(-c). \qquad \square$$

$\square$

The identity is the mechanism of the paper. A large push rate means a large hazard, but also a short floor time, and the two cancel in $\kappa_Q A_t$: the break is driven by the push ℓ_t, half the local time, net of a carry term that is small when q is small. In the limit of instant reflection a reflecting barrier on a traded rate becomes free of arbitrage once the break is priced per unit of push rather than per unit of time.

Corollary 2 (The relation in other conventions)**.** Under the conditions of Theorem 1: (i)

with $Y = X/\sigma$ and the stickiness ν written as in the sticky Brownian equations of [1], $\mathbf{1}\{Y_t = 0\}\,dt = \nu^{-1}\,dL_t(Y)$ with $L(Y)$ the right local time of Y at zero, one has $\nu = 2\theta/\sigma$ and

$$\kappa_Q\,(1 - e^J) = \tfrac{1}{2}\,\nu\sigma - q;$$

(ii) with the speed-measure atom of Proposition 1, $\kappa_Q\,(1 - e^J) = 1/m(\{0\}) - q$; (iii) if J is random, independent of W, B and E, with $\mathbb{E}_Q e^J < \infty$, the relation becomes $\kappa_Q\,(1 - \mathbb{E}_Q e^J) = \theta - q$.

Proof. Tanaka's formula for $(cX)^+ = cX^+$, $c > 0$, shows that the right local time scales, $L(cX) = c\,L(X)$, so $L(Y) = L(X)/\sigma$. By Lemma 1, $\mathbf{1}\{X_t = 0\}\,dt = dA_t = dL_t(X)/(2\theta) = \sigma\,dL_t(Y)/(2\theta)$, hence $\nu = 2\theta/\sigma$ and $\theta = \frac{1}{2}\nu\sigma$. (ii) Proposition 1 gives $m(\{0\}) = 1/\theta$. (iii) In the proof of Theorem 1 the compensator of the jump term becomes $\kappa_Q\,\mathbb{E}_Q(e^J - 1)\,S_{t-}\,\mathbf{1}\{X_t = 0, t < \tau\}\,dt$, and the rest is unchanged. □

Form (i) is the one with the volatility in it: in standardized units the same stickiness means a larger push when the rate is more volatile. Form (iii) is the one a random post-break level needs; only $\mathbb{E}_Q e^J$ enters.

Example 1 (The reference parameter set)**.** Take $\sigma = 0.04$ per root year, $\theta = 0.25$ per year, $J = -0.15$, $\sigma_p = 0.10$, $r_d = 0$ and $r_f = 0.003$, so $q = -0.003$. Theorem 1 gives the drift $\mu_Q = q - \frac{1}{2}\sigma^2 = -0.0038$ above the floor and, with $1 - e^J = 0.1393$, the hazard

$$\kappa_Q = \frac{0.25 + 0.003}{0.1393} = 1.8163$$

per year of floor time. The break comes after an exponential floor time with mean $1/\kappa_Q = 0.55$ years. From the floor the one-year break probability is 18.9 per cent (Table 2); the driftless reflecting limit of Corollary 1, with $c = 0.2872$, gives 19.3 per cent.

This set is fixed before any computation and is used for the prices, the smile and the term structure. Its volatility is close to the 3.5 per cent measured on the fixings, its push rate is inside the calibrated range, and its jump is close to the realized one.

Remark 1 (A sticky floor without break risk). With $\kappa_Q = 0$ the floor never breaks, and Theorem 1 reads $\theta = q$: a sticky floor on a traded rate is free of arbitrage only if the carry equals the push. With $q \le 0$, as for the Swiss franc against the euro in our calibration, no equivalent martingale measure exists, and by the fundamental theorem of asset pricing [12] the market admits a free lunch with vanishing risk. This is consistent with the result of Anagnostakis [17], whose main theorem, in the section of [17] on arbitrage, shows that a market whose risky asset is a sticky geometric Brownian motion, sticky at an interior level, satisfies the no-arbitrage condition if and only if the interest rate is zero; a remark that follows it notes that the stickiness persists under the risk-neutral measure, as θ does here. At his interior sticky level the price does not move while it sits there, so the carry must vanish; at a one-sided floor the price is pushed up at rate θ while it sits there, so the carry must equal θ. The case $\kappa_Q = 0$ of Theorem 1 thus extends his condition to a floor and a rate differential. With break risk the condition is no longer a restriction on the market: any push rate above q is matched by a hazard.

Remark 2 (The reflecting-barrier assumption). Hertrich and Zimmermann [5] price EUR/CHF options with the reflecting lower barrier of Veestraeten [8], implied below 1.20, and write that the rate "only prevails for an infinitely short time period on the floor in order to avoid arbitrage". Instant reflection does not avoid it on a traded rate without break risk. If $X = Y + \ell^R$ is a reflected diffusion, the discounted rate has the finite-variation part $e^{-qt}S_t((\mu + \frac{1}{2}\sigma^2 - q)\,dt + d\ell^R_t)$, and $d\ell^R$ is singular with respect to dt, because it grows on the set $\{X = 0\}$ of Lebesgue measure zero. A change to an equivalent measure moves the drift only by a term absolutely continuous in dt, so no equivalent measure makes the discounted rate a local martingale once the barrier is reached with positive probability. Corollary 1 says what repairs it: a break priced per unit of push, at intensity $d\ell^R_t/(1-e^J)$. Ingersoll [10] values foreign-exchange derivatives when the exchange process is bounded.

Method. Theorem 1 says that the forward is a martingale at the hazard $\kappa^*_Q = (\theta - q)/(1 - e^J)$ and at no other. In transform space this is exact: with $g(x) = Fe^x$, Proposition 2 must return $U(\lambda) = S_0/(\lambda - q)$, and we evaluate $|U(\lambda)(\lambda - q)/S_0 - 1|$ at every node of the Euler inversion. By Monte Carlo (the scheme of the prices section) we test $z = (\overline{S_T}\,e^{-qT} - S_0)/\text{s.e.}$. The test must also be able to fail, so it is run at the wrong hazards $\kappa = 0$ and $\kappa = 2\kappa^*_Q$.

Parameters. The reference set of Example 1, $\kappa^*_Q = 1.8163$; spots 1.20, 1.21 and 1.23; transform at $T = 1/12$, $1/4$, 1 and 2 years with 33 Euler nodes each; Monte Carlo at $T = 1/4$ and 1 year with 2^{20} paths per run and seeds 20220614, 11 and 2015. The tolerances were fixed before the computation: 10^{-10} in transform space, $|z| \le 4$ for every run at κ^*_Q, and $|z| \ge 10$ at each wrong hazard.

Result. In transform space the largest deviation is 5.7×10^{-16} over 3 spots, 4 maturities and 33 nodes. By Monte Carlo the largest $|z|$ over 3 seeds, 3 spots and 2 maturities is 2.62, with per-seed maxima 1.16, 1.45 and 2.62, and the largest relative error of the discounted mean is 1.8×10^{-4}.

Error. At the wrong hazards the test fails as it should: $z = +570$ at $\kappa = 0$ and $T = 1/4$, $z = +965$ at $\kappa = 0$ and $T = 1$, and $z = -214$ at $\kappa = 2\kappa^*_Q$ and $T = 1$. In transform space the one-year forward from 1.21 is then off by +2.3 and −1.8 per cent. Every run at κ^*_Q is within the tolerance.

Method. Corollary 1 gives the break probability in the limit of instant reflection as a one-dimensional integral against the law of the maximum of a drifted Brownian motion. We evaluate it by quadrature and compare it with the transform break probability $L^{-1}[\varphi/\lambda](T)$ of Proposition 2 at a large push rate, with the hazard $\kappa_Q = (\theta - q)/(1 - e^J)$ moving with θ. We also report the driftless closed form, which has no pass or fail.

Parameters. The reference set of Example 1, $T = 1$ year, spots 1.20 and 1.21; push rates $\theta = 1$, 10, 100, 1000 and 10^4; tolerance 10^{-4} at $\theta = 10^4$, fixed before the computation.

Result. At $\theta = 10^4$ the transform gives 0.203830 from 1.20 and 0.159449 from 1.21, and the closed form gives the same to six decimals. From the floor the transform moves from 0.20013 at $\theta = 1$ to 0.20346, 0.20379, 0.203826 and 0.203830 as θ grows. The driftless formula, with $c = 0.2872$, gives 0.1934, against 0.1895 for the transform at the reference $\theta = 0.25$.

Error. The largest absolute difference between transform and closed form is 3.7×10^{-7}, within the tolerance. The driftless formula is off by 0.010 from the limit because it drops the drift $\mu_Q = -0.0038$. On the calibrated parameters it gives 17.0 per cent at one year against 17.6 per cent for the transform, and 9.1 against 8.9 per cent at three months: the limit already carries most of the answer at finite θ.

4. Prices

Proposition 2 (Prices as Laplace transforms)**.** Under Assumption 1, let μ_Q, κ_Q and ν_Q be as in Theorem 1, and let the payoff be $g(x) = a_i + c_i e^x$ on finitely many intervals $[x_i, x_{i+1})$ of the log distance $x = \log(S/F) \in \mathbb{R}$, which covers calls, puts and digitals. For $\operatorname{Re}\lambda > \max(0, q)$ put $D = \mu_Q^2 + 2\lambda\sigma^2$, $r_\pm = (-\mu_Q \pm \sqrt{D})/\sigma^2$ with the principal root, and

$$K(x \to y) = \frac{e^{-r_+(y-x)}}{\sqrt{D}} \quad (y \geq x), \qquad K(x \to y) = \frac{e^{r_-(x-y)}}{\sqrt{D}} \quad (y < x).$$

Let $u_p(x) = \int_0^\infty K(x \to y)\, g(y)\, dy$ and

$$u(x) = u_p(x) + c\, e^{r_- x}, \quad c = \frac{g(0) - u_p(0)\,(\lambda + \kappa_Q - \theta r_+)}{\lambda + \kappa_Q - \theta r_-}, \qquad \varphi(x) = \frac{\kappa_Q\, e^{r_- x}}{\lambda + \kappa_Q - \theta r_-},$$

and let $h(\lambda) = \int_{\mathbb{R}} K_p(J \to y)\, g(y)\, dy$, with K_p the same kernel for drift ν_Q and volatility σ_p. Then, with $x_0 = \log(S_0/F)$,

$$\int_0^\infty e^{-\lambda T}\, \mathbb{E}_Q\big[g(\log(S_T/F))\big]\, dT = u(x_0) + \varphi(x_0)\, h(\lambda), \qquad \mathbb{E}_Q\, e^{-\lambda\tau} = \varphi(x_0).$$

The price at T is $e^{-r_d T}\, L^{-1}[u(x_0) + \varphi(x_0)h](T)$, the break probability is $Q(\tau \leq T) = L^{-1}[\varphi(x_0)/\lambda](T)$, and every integral is elementary.

Proof. Both sides are analytic in $\operatorname{Re}\lambda > \max(0, q)$*: on the left* $|\mathbb{E}_Q g| \leq C(1 + S_0 e^{qT})$*; on the right* $\operatorname{Re} r_- < 0$*, so the denominators have positive real part, and* $\operatorname{Re} r_+ > 1$*, so the call pieces converge, because* $\operatorname{Re}\sqrt{D} \geq \sqrt{\mu_Q^2 + 2\sigma^2 \operatorname{Re}\lambda}$ *and* $r_+ > 1$ *for real* $\lambda > q$ *when* $\mu_Q = q - \frac{1}{2}\sigma^2$*. It is enough to prove the identity for large real* λ*. The kernel.* $K(x \to y)$ *is the Laplace transform in time of the Gaussian transition density of* $\mu_Q t + \sigma\beta_t$ *from* x *to* y*. As a function of* x *it solves* $\lambda K - \mu_Q K' - \frac{1}{2}\sigma^2 K'' = 0$ *for* $x \neq y$*, is continuous, decays on both sides, and its derivative jumps by* $(r_- - r_+)/\sqrt{D} = -2/\sigma^2$ *at* $x = y$*. Hence* $\lambda u_p - \mu_Q u_p' - \frac{1}{2}\sigma^2 u_p'' = g$ *on* $(0, \infty)$ *off the breakpoints of* g*. Writing* $u_p(x) = D^{-1/2}\big(\int_0^x e^{r_-(x-y)} g(y)\, dy + \int_x^\infty e^{-r_+(y-x)} g(y)\, dy\big)$ *and differentiating, the boundary terms cancel and* $u_p'(0) = r_+ u_p(0)$*. The killed resolvent. The function* $v = u_p + c\, e^{r_- x}$ *solves the same equation, grows at most like* e^x*, and satisfies* $(\lambda + \kappa_Q)v(0) - \theta v'(0) = g(0)$*: with* $v'(0) = r_+ u_p(0) + c\, r_-$ *this is the stated* c*. Let* $Z_t = e^{-\lambda t - \kappa_Q A_t} v(X_t) + \int_0^t e^{-\lambda r - \kappa_Q A_r} g(X_r)\, dr$*. The function* v *is* C^1 *with a piecewise continuous second derivative, and* X *spends no time*

at the breakpoints, so Itô's formula applies as in Proposition 1 and gives

$$dZ_t = e^{-\lambda t-\kappa_Q A_t}\Big[\mathbf{1}\{X_t>0\}(g-\lambda v+\mu_Q v'+\tfrac12\sigma^2 v'')(X_t) + \mathbf{1}\{X_t=0\}(g(0)-(\lambda+\kappa_Q)v(0)+\theta v'(0))\Big]dt + dN_t$$

with N a local martingale, so both brackets vanish and Z is a local martingale. For large real λ it is uniformly integrable, by the bound on $\sup_{t\le T} e^{X_t}$ in the proof of Theorem 1, and letting $t\to\infty$, $v(x_0)=\mathbb{E}\int_0^\infty e^{-\lambda r-\kappa_Q A_r}g(X_r)\,dr=\mathbb{E}\int_0^\tau e^{-\lambda r}g(X_r)\,dr$, since $Q(\tau>r\mid X)=e^{-\kappa_Q A_r}$. So $u(x_0)$ is the transform of $\mathbb{E}_Q[g(X_T)\,\mathbf{1}\{T<\tau\}]$. The break. The same argument with the running reward $\kappa_Q\mathbf{1}\{X=0\}$ in place of g and the condition $(\lambda+\kappa_Q)w(0)-\theta w'(0)=\kappa_Q$ gives $w=\varphi$ and $\varphi(x_0)=\mathbb{E}\int_0^\infty e^{-\lambda r}\kappa_Q e^{-\kappa_Q A_r}\,dA_r=\mathbb{E}\,e^{-\lambda\tau}$, because given X the break time has distribution function $1-e^{-\kappa_Q A_t}$. After the break. For $T\ge\tau$, $\log(S_T/F)=J+\nu_Q(T-\tau)+\sigma_p(B_T-B_\tau)$ with $B_{\tau+\cdot}-B_\tau$ independent of $\mathcal{F}_\tau$. By Fubini, $\int_0^\infty e^{-\lambda T}\,\mathbb{E}[g(\log(S_T/F))\,\mathbf{1}\{\tau\le T\}]\,dT=\mathbb{E}[e^{-\lambda\tau}]\,h(\lambda)=\varphi(x_0)\,h(\lambda)$. Adding the two parts gives the transform; the break probability has transform $\mathbb{E}\,e^{-\lambda\tau}/\lambda$. On each interval the integrand is $a+c\,e^y$ times an exponential in y, so every integral is elementary. □

The proof uses only the equation of Definition 1, so every solution of it carries these prices. The transform of the forward, $g(x)=Fe^x$, is $S_0/(\lambda-q)$ exactly when κ_Q satisfies Theorem 1, which is the transform-space martingale test.

Algorithm 1 — Pricing by the sticky, killed resolvent

Inputs: σ, θ, J, σ_p, r_d, r_f; the spot S_0 and maturity T; the payoff as pieces $(a_i,c_i,[x_i,x_{i+1}))$ in $x=\log(S/F)$; the Euler parameters $A=18.4$, $n=20$, $m=12$ of Abate and Whitt [15]. *Output:* the price at T and the break probability $Q(\tau\le T)$.

1. Set $q=r_d-r_f$, $\mu_Q=q-\frac12\sigma^2$, $\kappa_Q=(\theta-q)/(1-e^J)$, $\nu_Q=q-\frac12\sigma_p^2$ and $x_0=\log(S_0/F)$.
2. For $k=0,1,\dots,n+m$ set the node $\lambda_k=(A+2\pi ik)/(2T)$.
3. At each node compute D, r_+ and r_-, then $u_p(x_0)$ and $u_p(0)$ piece by piece: a piece below the starting point contributes $\int e^{r_-(x-y)}(a+c\,e^y)\,dy$, a piece above it $\int e^{-r_+(y-x)}(a+c\,e^y)\,dy$, both in closed form.
4. Compute c and $\varphi(x_0)$ from Proposition 2, and $h(\lambda_k)$ from the post-break kernel over the whole line, started at J; set $U_k=u_p(x_0)+c\,e^{r_-x_0}+\varphi(x_0)\,h(\lambda_k)$.
5. Form the partial sums $s_j=(e^{A/2}/T)(\frac12\operatorname{Re}U_0+\sum_{k=1}^{j}(-1)^k\operatorname{Re}U_k)$ for $j=n,\dots,n+m$, and average them binomially, $f(T)=\sum_{j=0}^{m}\binom{m}{j}2^{-m}s_{n+j}$.
6. Return the price $e^{-r_dT}f(T)$, and the break probability from the same two steps applied to $\varphi(x_0)/\lambda_k$.

Cost. Thirty-three evaluations of elementary functions per payoff, spot and maturity; no grid and no simulation. A smile of 41 strikes at four maturities and three spots is 492 prices.

Accuracy. The Euler sum's discretization error is of order e^{-A}, and the binomial average accelerates the alternating tail [15]. Against the finite-difference solver the transform

agrees to 1.2×10^{-5} over 156 prices, and the transform-space martingale identity holds to 5.7×10^{-16} at every node.

Proposition 3 (Parity and the limits of no break and of instant reflection)**.** The formulas of Proposition 2 hold for any hazard $\kappa \geq 0$ in place of κ_Q, since their proof does not use Theorem 1. Then: (a) at $\kappa = \kappa_Q$, a call and a put of strike K and maturity T satisfy

$$C - P = S_0\, e^{-r_f T} - K\, e^{-r_d T};$$

(b) the transform is continuous in κ at $\kappa = 0$, where $\varphi = 0$, and at $\kappa = 0$ every put with $K \leq F$ is worth exactly zero; (c) at $\kappa = 0$, as $\theta \to \infty$ the transform converges to the transform of the reflected diffusion, $u_p + c_R\, e^{r_- x}$ with $c_R = -r_+ u_p(0)/r_-$, and at each node the difference is $O(1/\theta)$.

Proof. $C-P = e^{-r_d T}\,\mathbb{E}_Q[S_T - K]$*, and Theorem 1 gives* $\mathbb{E}_Q S_T = S_0\, e^{qT}$*. (b) The transform depends on* κ *through* $\lambda + \kappa$ *in two denominators whose real parts stay positive, so it is analytic in* κ *near zero, and* φ *carries the factor* κ*. At* $\kappa = 0$ *the floor never breaks,* $S_T \geq F \geq K$*, and* $(K - S_T)^+ = 0$*. (c) The reflected diffusion's resolvent satisfies the Neumann condition* $v'(0) = 0$ *(the verification of Proposition 2 with the regulator in place of the push), and* $u_p'(0) = r_+ u_p(0)$ *gives* c_R*. At* $\kappa = 0$*, subtracting,*

$$c(\theta) - c_R = \frac{g(0) - \lambda u_p(0) + \lambda u_p(0)\, r_+/r_-}{\lambda - \theta r_-},$$

because the terms in $\theta r_+ u_p(0)$ *cancel. The numerator does not depend on* θ *and* $|\lambda - \theta r_-| \geq \theta\,|\operatorname{Re} r_-|$*, so the difference is* $O(1/\theta)$*.* □

The limits in (b) and (c) are checks on the machinery, not models of the market: with $\kappa = 0$ the relation of Theorem 1 fails unless $\theta = q$ (Remark 1). The arbitrage-free limit of instant reflection lets the hazard grow with θ, and it is Corollary 1.

Method. Three routes that share only the model. (i) The transform of Proposition 2 inverted by Algorithm 1. (ii) Finite differences on the generator of Proposition 1, with its Feller–Wentzell condition as the boundary row and the break as a source term on the floor row, never using the transform: Crank–Nicolson after four implicit half steps, a second-order one-sided derivative at the floor node, cell-averaged payoffs and the free solution as the far boundary. (iii) Monte Carlo on the time change of Lemma 1: the free path Y is simulated exactly on a grid of the reflected clock, with exact Brownian-bridge minima inside each step, which gives the regulator ℓ^R exactly at grid points; real time is $s + \ell^R_s/\theta$ and the break comes when $\ell^R \geq \theta E$. Only the location inside the final step is approximated: X_T is placed on the floor with probability equal to that step's floor time over its real time, and otherwise at a uniform time of the step. Markov chain approximations [16] are another way to simulate diffusions with sticky boundaries.

Parameters. The reference set of Example 1; spots 1.20, 1.21 and 1.23; 13 payoffs: puts struck at 1.10, 1.15, 1.20, 1.21, 1.23 and 1.25, calls at 1.20, 1.21, 1.23, 1.25 and 1.30, a digital put at 1.20 and a digital call at 1.22. Transform and finite differences at $T = 1/12$,

1/4, 1 and 2 years, 156 cells; grid step 10^{-4} on $[0, 0.6]$ with $\max(500, \lceil 2000T \rceil)$ time steps. Monte Carlo at $T = 1/4$ and 1 year, 78 cells, 2^{20} paths per run, reflected-clock step 5×10^{-5} years, seeds 20220614, 11 and 2015. Tolerances, fixed before the computation: $|\text{transform} - \text{FD}| \leq \max(2 \times 10^{-5}, 10^{-3}|\text{price}|)$ and $|z| \leq 4$ per cell, per seed and pooled.

Result. Table 1 and Figure 1. The largest absolute difference between transform and finite differences is 1.2×10^{-5} over 156 cells, 0.094 of the tolerance at worst. Against Monte Carlo the largest $|z|$ per seed is 1.70, 2.40 and 2.83, and 1.78 pooled over the three seeds. The break probability of the transform against the Monte Carlo break frequency has largest $|z| = 1.47$ over 18 runs.

Error. Every cell is within its tolerance. The finite-difference error is a discretization error of the solver; the Monte Carlo deviations are those of sampling noise, spread evenly across payoffs, spots and seeds.

Source. Proposition 2 by Algorithm 1, the finite-difference solver and the Monte Carlo scheme of the price comparison, reference set of Example 1, spot 1.21, $T = 1$ year. The Monte Carlo value is the mean over three seeds of 2^{20} paths each, with its standard error in units of the last digit in parentheses.

Columns. The payoff; its price in Swiss francs per euro of notional by the transform, by finite differences and by Monte Carlo; and z, the Monte Carlo value less the transform in standard errors.

payoff	**transform**	**finite differences**	**Monte Carlo**	z
put 1.10	0.011413	0.011413	0.011438(21)	1.17
put 1.15	0.017942	0.017942	0.017967(29)	0.86
put 1.20	0.024972	0.024972	0.024993(38)	0.57
put 1.21	0.027599	0.027599	0.027622(39)	0.57
put 1.23	0.036697	0.036697	0.036722(42)	0.60
put 1.25	0.050095	0.050095	0.050127(44)	0.74
call 1.20	0.031347	0.031347	0.031335(18)	−0.71
call 1.21	0.023975	0.023975	0.023963(16)	−0.69
call 1.23	0.013072	0.013072	0.013064(13)	−0.66
call 1.25	0.006470	0.006470	0.006469(9)	−0.12
call 1.30	0.000723	0.000723	0.000724(3)	0.24
digital put 1.20	0.143283	0.143283	0.143191(197)	−0.46
digital call 1.22	0.543117	0.543111	0.543096(281)	−0.08

Transform and finite differences agree to the sixth decimal except for the digital call, where they differ by 6.2×10^{-6}; no Monte Carlo value is more than 1.17 standard errors from the transform.

Script. `fig_prices.py`: reads `results.npz` and plots the stored prices and z-scores; nothing

is recomputed.

Shows. Left, puts from spot 1.21 at three months and one year by the three methods, with four-standard-error bars on the Monte Carlo means. Right, the z-score of every Monte Carlo cell against the transform, for each of the three seeds, with the tolerance $|z| = 4$ in red.

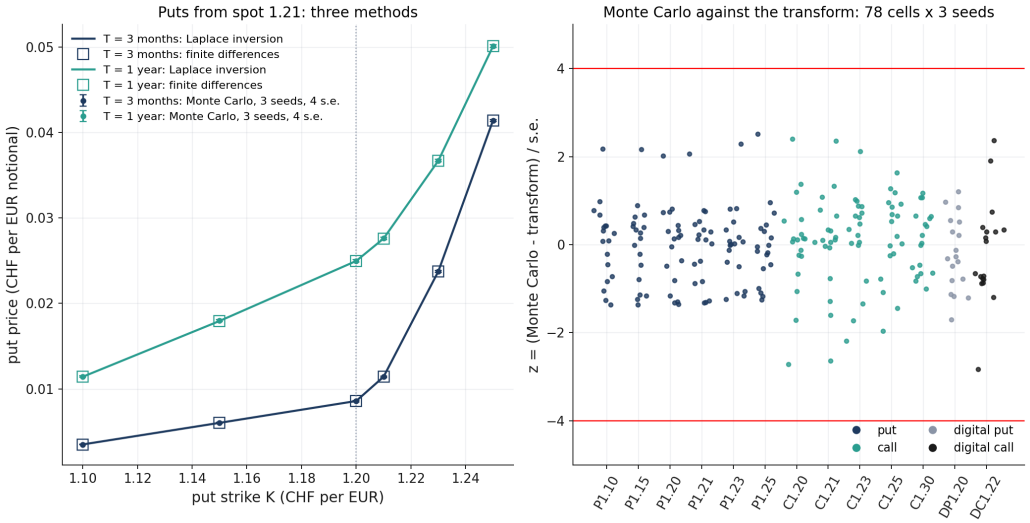


Figure 1: Put prices from spot 1.21 by Laplace inversion, finite differences and Monte Carlo at three months and one year, and the z-scores of all 78 Monte Carlo cells against the transform for each of three seeds, all inside four standard errors.

The put curve bends at the floor: puts struck below 1.20 pay only after a break and are cheap, puts struck above it pay on the ordinary diffusion as well. The 234 z-scores scatter without pattern by payoff, and none reaches 3.

Method. Proposition 3 checked on the transform: parity across strikes; the transform at $\kappa = 10^{-9}$ against $\kappa = 0$, the puts at or below the floor at $\kappa = 0$, and the transform at $\kappa = 0$ against Monte Carlo at $\kappa = 0$; and the transform at $\kappa = 0$ for growing θ against the closed-form law of the reflected diffusion of Abate and Whitt [14], integrated by quadrature.

Parameters. The reference set of Example 1. Parity over spots 1.20, 1.21 and 1.23, maturities 1/12, 1/4, 1 and 2 years and the seven strikes of the price grid; the $\kappa \to 0$ checks on the put struck at 1.23, at $T = 1$ and spot 1.21, and on all puts struck at 1.10, 1.15 and 1.20; the $\theta \to \infty$ check on the same put at $\theta = 1$, 10, 100, 1000 and 10^4. Tolerances fixed before the computation: 10^{-7} for parity, 10^{-8} for continuity in κ, 10^{-12} for the puts, and 10^{-4} at $\theta = 10^4$ with the error falling at least fivefold per decade.

Result. The largest parity violation is 1.3×10^{-9}. The transform moves by 1.7×10^{-11} between $\kappa = 10^{-9}$ and $\kappa = 0$, every put struck at or below 1.20 is exactly 0 at $\kappa = 0$, and Monte Carlo at $\kappa = 0$ agrees with the transform within $|z| = 2.17$ over the 18 cells with nonzero sample variance. The reflected closed form gives 0.007662 for the put; the transform's errors at the five push rates are 4.4×10^{-4}, 4.4×10^{-5}, 4.5×10^{-6}, 4.5×10^{-7}

and 4.5×10^{-8}.

Error. All within tolerance. The error times θ stays between 4.40×10^{-4} and 4.46×10^{-4}, the $O(1/\theta)$ rate of Proposition 3(c), and falls tenfold per decade.

5. What the Floor Does to Options

Example 2 (The implied-volatility smile under the reference set)**.** Price the out-of-the-money option at each strike from 1.00 to 1.40 in steps of 0.01 by Algorithm 1, and invert the Garman–Kohlhagen formula [13] for its implied volatility. From spot 1.21 at one month the implied volatility is 18.9 per cent at $K = 1.10$, 13.0 at 1.15, 4.5 at the floor, 4.1 at the money and 4.0 at 1.25; at one year the same strikes give 10.1, 8.5, 5.8, 5.3 and 4.4 per cent. From the floor at one month the left wing peaks at 24.8 per cent at $K = 1.06$; from 1.21 it peaks at 21.7 per cent at 1.05. Far above the floor the smile returns to the diffusion volatility: 4.1 per cent at $K = 1.35$ and one year.

The skew points steeply down, and it is made by the break, not by the diffusion. A put struck below the floor pays only if the floor breaks (Proposition 4), so its price is a break probability times an expected shortfall, and the Garman–Kohlhagen formula, which knows only one volatility, needs a large one to produce it. The wing peaks at strikes of 1.05 to 1.06, a little above the post-break level $Fe^{J} = 1.03$. As the maturity grows the diffusion's own spread grows and the wing flattens; as the spot moves away from the floor the break recedes and the wing falls (Figure 2). The 56 points of 492 whose out-of-the-money price is below 10^{-9}, all calls struck at 1.27 or above at one and three months, are dropped rather than inverted.

Script. `fig_smile.py`: reads the implied volatilities stored in `results.npz`; nothing is recomputed.

Shows. Left, the Garman–Kohlhagen implied volatility against the strike from spot 1.21 at one, three, six and twelve months. Right, the three-month smile from spots 1.20, 1.21 and 1.23. The dotted vertical line is the floor and the grey horizontal line the diffusion volatility of 4 per cent.

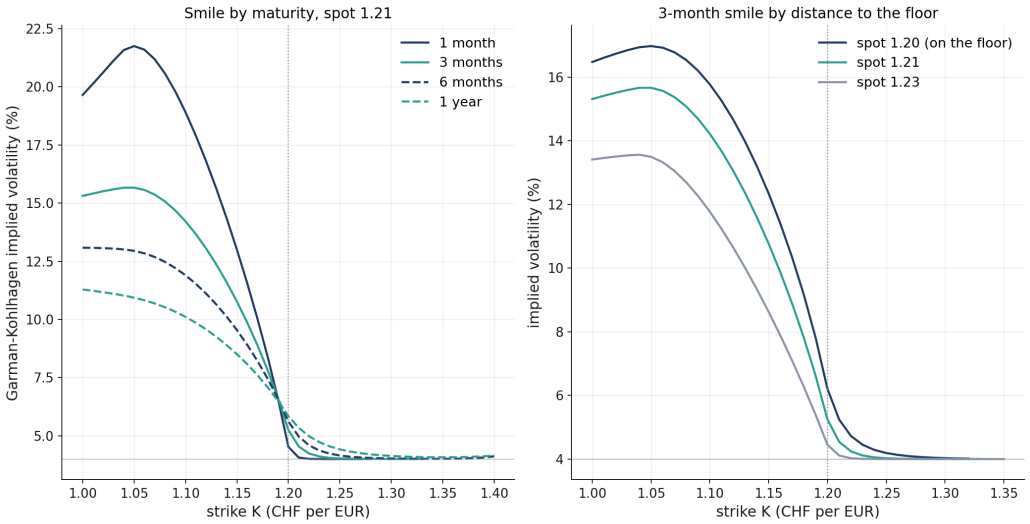


Figure 2: Implied volatility under the sticky floor against the strike, by maturity from spot 1.21 and by distance to the floor at three months, with a steep left wing below the floor and a flat right wing at the diffusion volatility.

At three months the left wing peaks at 17.0 per cent from the floor, 15.7 from 1.21 and 13.6 from 1.23, at strikes of 1.04 to 1.05; to the right of the floor every curve falls onto the diffusion volatility within a few strikes.

Proposition 4 (Options struck at or below the floor are break-risk claims)**.** Under Definition 1, let $K \le F$. A put of strike K and maturity T pays only after a break, and

$$P(K,T) = e^{-r_d T}\,\mathbb{E}_Q\big[(K - S_T)^+\,\mathbf{1}\{\tau \le T\}\big] \le K\,e^{-r_d T}\,Q(\tau \le T).$$

The probability of a rate below the floor at T is at most the probability of a break by T,

$$Q(S_T < F) = Q(\tau \le T) - Q(\tau \le T,\ S_T \ge F),$$

the difference being the probability that a broken rate is back at or above the floor at T.

Proof. Before the break $X_t \ge 0$, so $S_t \ge F \ge K$ on $\{T < \tau\}$, where $(K - S_T)^+ = 0$ and $\mathbf{1}\{S_T < F\} = 0$. Hence both payoffs vanish off $\{\tau \le T\}$, and $(K - S_T)^+ \le K$. □

The proposition separates the two quantities that are easy to confuse: the probability that the floor breaks by T, and the probability that the rate is below the floor at T. They coincide at short horizons, where a broken rate has no time to climb back, and drift apart when the jump is small or the horizon long.

Method. The break probability $Q(\tau \le T) = L^{-1}[\varphi/\lambda](T)$ and the probability of a rate below the floor, $Q(S_T < F)$, the undiscounted digital put of Proposition 2, are two separate transform inversions. Proposition 4 says the second is at most the first; we measure the gap on a term grid.

Parameters. The reference set of Example 1; spots 1.20, 1.21 and 1.23; 60 maturities

spaced geometrically from one week to two years. The values in Table 2 at one and three months are read off this grid by linear interpolation.

Result. Table 2. From the floor the break probability is 4.7 per cent at one month, 9.1 at three months, 18.9 at one year and 26.3 at two years; from 1.23 it is 0.1, 1.2, 8.2 and 15.3 per cent. The probability of a rate below the floor is 4.7, 9.1, 18.4 and 24.2 per cent from the floor at the same horizons.

Error. The gap $Q(\tau \le T) - Q(S_T < F)$ is never negative beyond 6×10^{-16}, rounding in the two inversions, as Proposition 4 requires. It is at most 6.0×10^{-10} at one month, where a broken rate has no time to climb back, and grows to 2.07 percentage points from the floor at two years.

Source. Proposition 2, $Q(\tau \le T) = L^{-1}[\varphi/\lambda](T)$ by Algorithm 1, reference set of Example 1, read off a 60-point term grid.

Columns. The spot, then the risk-neutral probability, in per cent, that the floor breaks within one month, three months, one year and two years.

spot	1 month	3 months	1 year	2 years
1.20	4.7	9.1	18.9	26.3
1.21	1.5	5.1	14.7	22.2
1.23	0.1	1.2	8.2	15.3

The term structure is concave and rises fastest at short horizons from the floor, where floor time accumulates at once; from 1.23 the rate must first return to the floor, and the curve starts flat.

6. The EUR/CHF Episode

The data are the European Central Bank's euro reference rates for the Swiss franc, series EXR.D.CHF.EUR.SP00.A, the daily fixing at 14:15 Central European Time to four decimals, from the ECB Data Portal for 2011 to 2015 (Source: ECB statistics). The file holds 1279 fixings from 3 January 2011 to 31 December 2015 and six empty rows with status H, holidays, which are skipped. We use the statistics unmodified and derive only the log distance to 1.20, counts and realized variances.

The floor sample runs from 6 September 2011 to 14 January 2015: 858 fixings, one for every TARGET working day, over 3.36 years, or 255.6 fixings a year, which sets the clock of one fixing to 1/255.6 of a year. No working day is missing. The fixing never printed 1.2000: its lowest value in the sample is 1.2008 and its highest 1.2599, and in 2012 it stayed just above 1.20 for months (Figure 3). In these data the floor is therefore a band, not a level. We call a fixing a floor day when it is at or below a band edge p_δ, with $\delta = \log(p_\delta/1.20)$ and p_δ one of 1.2005, 1.2010, 1.2015, 1.2020, 1.2030 and 1.2050; the primary edge is 1.2010, with 63 floor days, and the edge 1.2005 holds none.

The break is read from the fixing of 15 January 2015, 1.0280: $J = \log(1.0280/1.20) =$

-0.1547 and $1-e^{J}=0.1433$. The post-break volatility is the realized volatility of daily log changes from 16 January to 31 December 2015, $\sigma_p = 7.0$ per cent. The interest differential is not estimated: we take $r_d = 0$ and $r_f = 0.003$, so $q = -0.003$, and report the results at $q = 0$ and $q = -0.006$ as well.

Script. `fig_path.py`: reads the ECB file `ecb_exr_d_chf_eur_2011_2015.csv` beside it and plots the fixings as published.

Shows. Left, the EUR/CHF fixing from 2011 to 2015, the floor, the 63 floor-regime fixings at or below 1.2010, and the fixing of 15 January 2015. Right, April to September 2012, with the band edges 1.2010, 1.2020 and 1.2030.

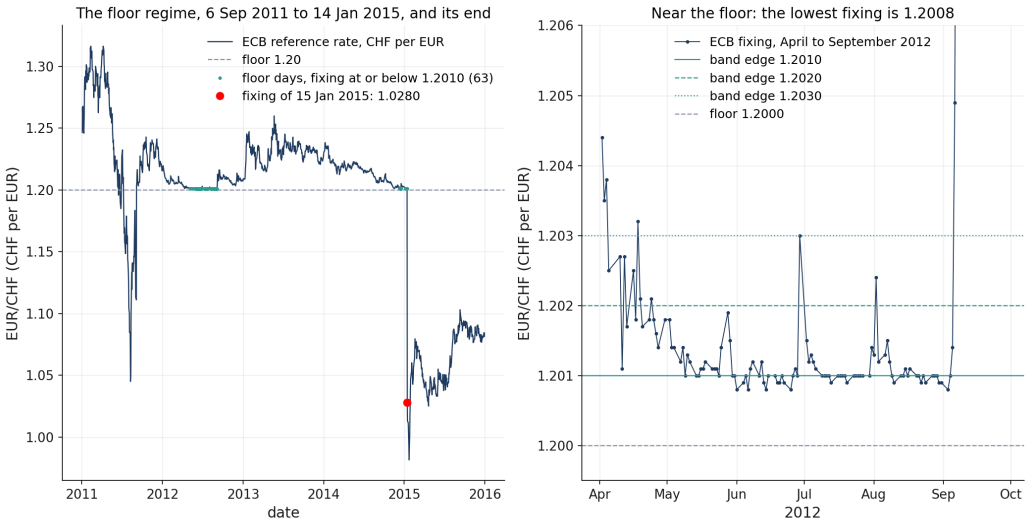


Figure 3: The EUR/CHF reference rate from 2011 to 2015 with the 1.20 floor, the floor days at or below 1.2010 and the 15 January 2015 fixing of 1.0280, and a zoom on 2012 showing the fixing resting just above the floor.

The rate does not touch the floor and leave: it rests about a thousandth above it for weeks at a time, and the lowest fixing of the regime is 1.2008. The jump of 15 January 2015 takes it to 1.0280 in one fixing.

Lemma 2 (The stationary atom of the floor without break)**.** Let $\kappa = 0$ and $\mu < 0$ in Definition 1. The probability measure

$$\pi(dx) = \pi_0\,\delta_0(dx) + (1-\pi_0)\,\frac{2|\mu|}{\sigma^2}\,e^{-2|\mu|x/\sigma^2}\,dx, \qquad \pi_0 = \frac{|\mu|}{|\mu|+\theta},$$

satisfies $\int \mathcal{G}f\,d\pi = 0$ for every $f \in C^2([0,\infty))$ with compact support, and it is the speed measure of Proposition 1 normalized to mass one. The long-run fraction of time on the floor is π_0.

Proof. Write $p(x) = (1-\pi_0)(2|\mu|/\sigma^2)e^{2\mu x/\sigma^2}$, *so that* $\frac{1}{2}\sigma^2 p' = \mu p$. *For such* f, *integration*

by parts gives

$$\int_0^\infty (\mu f' + \tfrac{1}{2}\sigma^2 f'')p\,dx = -\tfrac{1}{2}\sigma^2 f'(0)\,p(0) + \int_0^\infty (\mu p - \tfrac{1}{2}\sigma^2 p')f'\,dx = -\tfrac{1}{2}\sigma^2 f'(0)\,p(0),$$

and the atom contributes $\pi_0\,\theta f'(0)$*. The sum is* $f'(0)(\pi_0\theta - (1-\pi_0)|\mu|)$*, which vanishes for the stated* π_0*. The speed measure of Proposition 1 has density* $(2/\sigma^2)e^{2\mu x/\sigma^2}$*, of total mass* $1/|\mu|$*, and the atom* $1/\theta$*; normalized, its atom is* $(1/\theta)/(1/\theta + 1/|\mu|) = \pi_0$*.* □

For a positively recurrent diffusion the normalized speed measure is the stationary law, and the ergodic theorem makes π_0 the long-run fraction of time on the floor. Inverting, $\theta = |\mu|(1-\pi_0)/\pi_0$: a second estimator of the push rate from the fraction of floor days, usable only when the drift is negative and measured.

Start. The push rate is a ratio of path functionals, $\theta = L_t/(2A_t)$ (Lemma 1), and daily fixings see neither the local time nor the floor time directly. The occupation-density formula for the right local times L_t^a of X links both to counts of fixings: for bounded measurable h,

$$\int_0^t h(X_r)\,d\langle X\rangle_r = \int_{\mathbb{R}} h(a)\,L_t^a\,da, \qquad d\langle X\rangle_r = \sigma^2\,\mathbf{1}\{X_r > 0\}\,dr. \tag{6.1}$$

Steps. Take a band $(0,\delta]$ of fixings counted as floor days and a layer $(\delta,\delta+\varepsilon]$ above it. With h the indicator of the layer, σ^2 times the time spent in the layer is $\int_\delta^{\delta+\varepsilon} L_t^a\,da$. If the local-time profile is flat over $[0,\delta+\varepsilon]$, that is $L_t^a \approx L_t^0 = L_t$, the time in the layer is $\varepsilon L_t/\sigma^2$ and the time in $(0,\delta]$ is $\delta L_t/\sigma^2$. Let n_0 be the number of fixings with $X \le \delta$, n_1 the number in the layer and N the fixings per year. A floor day is floor time or time in $(0,\delta]$, so

$$\frac{n_0}{N} \approx A_t + \frac{\delta L_t}{\sigma^2}, \qquad \frac{n_1}{N} \approx \frac{\varepsilon L_t}{\sigma^2}. \tag{6.2}$$

Solving, $L_t \approx \sigma^2 n_1/(\varepsilon N)$ and $A_t \approx (n_0 - (\delta/\varepsilon)n_1)/N$.

End. The band-corrected and the raw estimators are

$$\hat\theta_c = \frac{\sigma^2\,n_1}{2\varepsilon\,(n_0 - (\delta/\varepsilon)\,n_1)}, \qquad \hat\theta_{\text{raw}} = \frac{\sigma^2\,n_1}{2\varepsilon\,n_0}, \tag{6.3}$$

the second ignoring the width of the band. The volatility is estimated from daily log changes with both ends above the band, $\hat\sigma^2 = N\,\overline{(\Delta x)^2}$, and the drift as $\hat\mu_P = N\,\overline{\Delta x}$; neither uses the floor days. The push, $\ell_t = L_t/2$, is carried by n_1 alone, while the floor time is a difference of two comparable counts and is much noisier, so $\hat\theta$ inherits the noise of the floor time. Lemma 2 gives a second estimator, $\hat\theta_{\text{stat}} = |\hat\mu_P|(1-\hat\pi_0)/\hat\pi_0$ with $\hat\pi_0 = n_0/n$ the fraction of floor days among the n fixings, defined when $\hat\mu_P < 0$.

Method. The estimators of the derivation above on the floor sample, for every band edge and layer; the hazard $\kappa_Q = (\hat\theta_c - q)/(1-e^J)$ of Theorem 1 with the realized jump; the risk-neutral probability of a rate below 1.20 by Proposition 2; a moving-block bootstrap of

the floor-sample fixings; and the fraction-of-days estimator of Lemma 2 as a cross-check. Corollary 1 predicts that the break probability depends on the push and only weakly on how it is split between push rate and floor time, so it should move far less than $\hat{\theta}$ across the grid.

Parameters. Band edges 1.2005 to 1.2050, layers $\varepsilon = 0.002$, 0.004 and 0.008, primary cell 1.2010 and 0.004; $q = -0.003$, with 0 and -0.006 as sensitivities; $J = -0.1547$, $\sigma_p = 7.0$ per cent; bootstrap blocks of 20 fixings, 2000 resamples, seed 20220614, 90 per cent percentile intervals.

Result. Above the primary band, 781 daily changes give $\hat{\sigma} = 3.53$ per cent, and 3.53 to 3.77 per cent across bands; the drift is $\hat{\mu}_P = -0.23 \pm 2.02$ per cent a year, -0.11 standard errors from zero. In the primary cell $n_0 = 63$ and $n_1 = 146$, so $\hat{\theta}_{\text{raw}} = 0.361$ and $\hat{\theta}_c = 0.699$ per year, with a bootstrap interval of 0.23 to 4.66; the hazard is $\kappa_Q = 4.89$ per year of floor time (bootstrap 1.61 to 32.5). The risk-neutral probability of a rate below 1.20 is 5.0, 8.9 and 17.5 per cent at one, three and twelve months from the floor, 1.5, 4.7 and 13.2 from 1.21, and under 0.1, 0.9 and 6.8 from 1.23. The interest differential moves the twelve-month value from the floor to 16.6, 17.5 and 18.4 per cent at $q = 0$, -0.003 and -0.006. The fraction-of-days estimator gives $\hat{\pi}_0 = 7.3$ per cent and $\hat{\theta}_{\text{stat}} = 0.029$, hence $\kappa_Q = 0.22$ and 11.0 per cent at twelve months from the floor; it rests on a drift indistinguishable from zero and carries no information.

Error. Across the 15 cells with floor days (Table 3), $\hat{\theta}_c$ runs from 0.20 to 1.71 per year, a factor of 8.4, and κ_Q from 1.45 to 11.94; the twelve-month probability from the floor runs only from 16.9 to 18.7 per cent, a factor of 1.11, and from 1.21 from 12.7 to 14.4 per cent. The bootstrap interval of the twelve-month probability is 14.6 to 19.3 per cent from the floor and 10.3 to 15.0 from 1.21, against 0.23 to 4.66 for $\hat{\theta}_c$; 1808 of the 2000 resamples have a positive corrected floor count and are used. The stickiness is weakly identified and the break probability is not, as Corollary 1 predicts.

Source. The derivation of the stickiness estimators and Theorem 1 on the 858 ECB fixings of the floor sample, with $q = -0.003$ and the realized jump; probabilities by Proposition 2. The band edge 1.2005 holds no fixings and is omitted.

Columns. The band edge p_δ and the layer width ε; the band-corrected push rate $\hat{\theta}_c$ and the hazard κ_Q, per year; the risk-neutral probability, in per cent, of a rate below 1.20 in twelve months, from 1.20 and from 1.21.

band edge, ε	$\hat{\theta}_c$	κ_Q	**from** 1.20	**from** 1.21
1.2010, 0.002	0.987	6.91	17.6	13.3
1.2010, 0.004	0.699	4.89	17.5	13.2
1.2010, 0.008	0.649	4.55	17.4	13.2
1.2015, 0.002	0.218	1.54	16.9	12.7
1.2015, 0.004	0.293	2.07	17.2	13.0
1.2015, 0.008	0.299	2.11	17.2	13.0
1.2020, 0.002	0.204	1.45	16.9	12.8
1.2020, 0.004	0.306	2.16	17.4	13.2
1.2020, 0.008	0.299	2.11	17.4	13.1
1.2030, 0.002	0.232	1.64	17.2	13.1
1.2030, 0.004	0.372	2.62	17.7	13.5
1.2030, 0.008	0.253	1.79	17.4	13.2
1.2050, 0.002	1.709	11.94	18.7	14.4
1.2050, 0.004	0.867	6.07	18.5	14.2
1.2050, 0.008	0.239	1.69	17.6	13.4

The floor-day counts are 63, 106, 118, 144 and 183 for the five edges. The push rate varies eightfold down the table and the probabilities by less than two percentage points.

Method. Hertrich and Zimmermann [5] report risk-neutral probabilities, implied from EUR/CHF options by a reflecting-barrier model, that the rate falls below 1.20 within one, three and twelve months. We take their numbers as they state them in the text of the working-paper version of [5], never digitized from a figure, and set beside them the calibrated model's probabilities from the primary cell: the break probability $Q(\tau \le T)$ and the probability $Q(S_T < 1.20)$ of a rate below the floor at T. Proposition 4 says the second is at most the first. To separate the probability of a break from its size, we recompute both with the post-break level at 1.10 and at 1.15, the realignment Hertrich and Zimmermann describe, in place of the fixing 1.028.

Parameters. The primary cell of the calibration, $\hat{\theta}_c = 0.699$, $\hat{\sigma} = 3.53$ per cent, $\sigma_p = 7.0$ per cent, $q = -0.003$; post-break levels 1.028, 1.10 and 1.15; spots 1.20 and 1.21; horizons of one, three and twelve months.

Result. Their three-month probability is about 50 per cent at the start of the floor, under 40 per cent by the end of 2011, 50 per cent again in the second quarter of 2012, about 35 per cent after September 2012, "between zero and 30%" from January 2013 to August 2014, 40 per cent in August 2014 and about 50 per cent from October and November 2014 until the abandonment; their twelve-month probability "does not fall much below 30%" even when the one-month probability is essentially zero. With the realized jump the path implies 8.9 per cent at three months and 17.5 per cent at twelve months from the floor: the first inside their range for January 2013 to August 2014, the second well below their

twelve-month level, and the three-month value well below the 40 to 50 per cent of the final months. With the jump to 1.15 the same path implies 24.4 and 35.6 per cent from the floor and 13.5 and 28.0 per cent from 1.21, inside their ranges; with the jump to 1.10, 14.6 and 25.8 per cent from the floor (Table 4 and Figure 4).

Error. Against Proposition 4, no break probability falls below the probability of a rate under the floor. With the realized jump the two differ by less than 0.1 percentage point at twelve months, 17.6 against 17.5 per cent from the floor, so the comparison with Hertrich and Zimmermann does not depend on which of the two their numbers are read as. With the jump to 1.15 they differ by 9.0 points, 44.6 against 35.6 per cent, because a rate that breaks to 1.15 often climbs back above 1.20 within the year.

Source. Theorem 1 and Proposition 2 on the primary calibration cell, $q = -0.003$, spot 1.20, horizon twelve months, with the post-break level set to the fixing of 15 January 2015 and to two alternatives.

Columns. The post-break level Fe^{J}; the relative jump $1 - e^{J}$; the hazard κ_Q per year of floor time; the risk-neutral break probability $Q(\tau \le T)$ and the probability $Q(S_T < 1.20)$ of a rate below the floor, in per cent.

post-break level	$1 - e^{J}$	κ_Q	$Q(\tau \le T)$	$Q(S_T < 1.20)$
1.028	0.1433	4.89	17.6	17.5
1.10	0.0833	8.42	27.4	25.8
1.15	0.0417	16.84	44.6	35.6

The product $\kappa_Q(1 - e^{J}) = 0.7016$ is the same in every row: it is $\hat{\theta}_c - q$, which the path fixes.

Script. `fig_break.py`: reads the stored term structures in `results.npz` and the jump in `numbers.json`; the estimates of Hertrich and Zimmermann are drawn from the numbers stated in their text.

Shows. The calibrated risk-neutral probability of a rate below 1.20 against the horizon, from spots 1.20, 1.21 and 1.23 with the realized jump to 1.028, the range over the band and layer grid from the floor (shaded), and the curves from 1.20 and 1.21 with the jump to 1.15 instead (dashed). The grey bar at three months is Hertrich and Zimmermann's range for January 2013 to August 2014, the black bar their level for August 2014 to January 2015, and the arrow at twelve months their statement that the twelve-month probability does not fall much below 30 per cent.

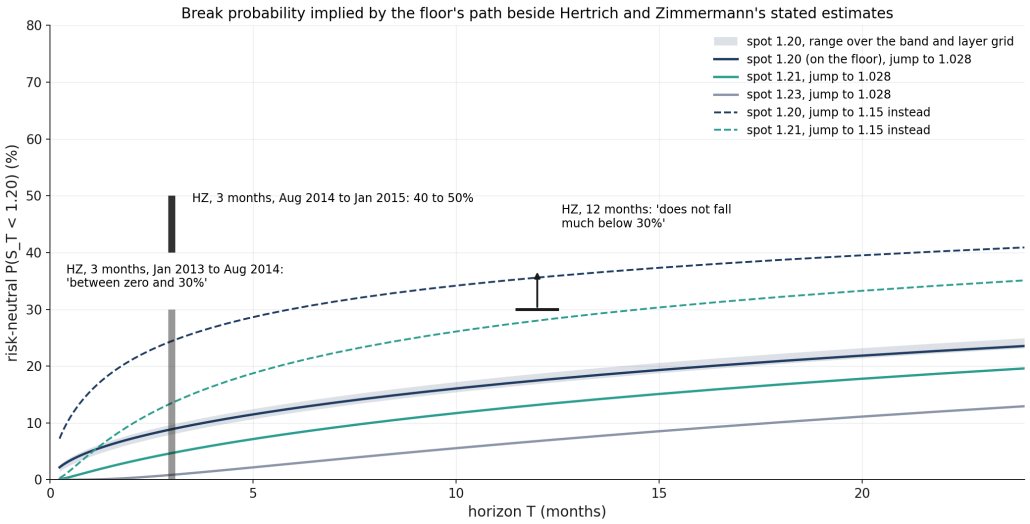


Figure 4: The calibrated risk-neutral probability of a EUR/CHF rate below 1.20 against the horizon, with the realized jump and with a jump to 1.15, beside the three-month and twelve-month estimates that Hertrich and Zimmermann state in words.

The grid band around the floor curve is narrow at every horizon, although the push rate behind it varies eightfold. The realized-jump curves sit below the stated twelve-month level; the dashed curves reach it.

Remark 3 (Why the product is pinned and the factors are not). Theorem 1 fixes the product $\kappa_Q(1-e^J) = \theta - q$, the risk-neutral rate of expected loss per unit of floor time, and the floor's path fixes θ, through the push. It does not fix κ_Q and J separately: a smaller jump needs a larger hazard, in exact proportion, and Table 4 shows the product at 0.7016 for post-break levels of 1.028, 1.10 and 1.15. The path therefore prices the product of break probability and jump size. Hertrich and Zimmermann [5] describe a realignment to about 1.15; with that jump the floor's path gives 35.6 per cent at twelve months, inside their option-implied range, and with the realized jump it gives about half as much. The two readings are consistent if the option market priced a small jump with a high probability, and the jump that came was large. Corollary 1 explains the other half of the result, that the probability is pinned although θ is not: the integrated hazard is the push over the relative jump, $\kappa_Q A_t \approx \ell_t/(1-e^J)$ when q is small, and the push is carried by the layer counts, which are stable, while θ divides it by the floor time, which is not.

7. Limits

Remark 4 (What the calibration does not settle). Daily fixings at 14:15 miss intraday visits to the floor, and the fixing never shows 1.2000; floor time is measured with a band δ and the push with a layer ε, and $\hat{\theta}$ moves from 0.20 to 1.71 per year with them while the twelve-month break probability moves from 16.9 to 18.7 per cent. The estimator assumes a flat local-time profile over the band and the layer; a drift or a volatility that depends on the distance to the floor breaks it, and Lera and Sornette [7] find volatility falling towards the floor, which a constant σ reads partly as stickiness. The hazard and

the parameters are constant, while Hertrich and Zimmermann [5] and Jermann [6] find the floor's credibility moving over time: ours is one calibration for the whole episode, set against their time-varying numbers. The post-break law rests on one fixing, 1.0280, and does not use the intraday low; σ_p is the realized volatility of 2015, known only after the event; and the break probabilities depend on the jump, 17.5 per cent at twelve months becoming 35.6 per cent with a jump to 1.15. A random jump enters only through $\mathbb{E}_Q e^J$ (Corollary 2). The interest differential is assumed, not estimated; $q = 0$ and $q = -0.006$ give 16.6 and 18.4 per cent. The push rate and the volatility are pathwise and carry over to Q, but the real-world hazard κ_P is not identified by Theorem 1, and every probability here is risk-neutral, like those of Hertrich and Zimmermann.

Two points of procedure. The Monte Carlo scheme approximates only the location inside the last step of the reflected clock. A first run interpolated that location deterministically, which smeared the floor's atom and failed the $\kappa = 0$ check with a largest $|z|$ of 4.97, a bias that shrank with the step; the step now places X_T on the floor with probability equal to its floor time over its real time, and every Monte Carlo line was rerun and passes. The parameters, grids, seeds and tolerances were fixed in a parameter file before any computation; four further items, the reflecting-limit check, the driftless formula, the jump sensitivity and the fraction-of-days estimator, were added to that file after the second run and before they were computed.

8. Conclusion

A defended floor makes the exchange rate sticky, and stickiness on a traded rate has a price. We have shown that the absence of arbitrage ties the push the floor gives the rate to the risk that it breaks: the hazard per unit of floor time times the relative jump equals the push rate net of carry, $\kappa_Q(1 - e^J) = \theta - q$, and since the push rate is a functional of the path, the same under every equivalent measure, the risk-neutral hazard is fixed by the path. Integrated, the hazard is the floor's push over the relative jump, so the break probability depends on the push and hardly on how it splits into push rate and floor time. Calls, puts and digitals follow as explicit Laplace transforms, confirmed by an independent finite-difference solver and by Monte Carlo on an exact time change, and the floor gives options a steep downside skew in which puts below the floor price only the break.

On the ECB fixings of the EUR/CHF floor the push rate is poorly measured, 0.20 to 1.71 per year, but the twelve-month probability of a rate below 1.20 from the floor is not: 16.9 to 18.7 per cent with the realized jump, and 35.6 per cent with the jump to 1.15 that the option-implied analysis of Hertrich and Zimmermann describes. A floor's path prices the product of break probability and jump size. How the option market split that product between probability and size, the path alone cannot say.

References

[1] Engelbert, H.-J., & Peskir, G. (2014). Stochastic differential equations for sticky Brownian motion. *Stochastics*. https://doi.org/10.1080/17442508.2014.899600

[2] Feller, W. (1952). The parabolic differential equations and the associated semi-groups of transformations. *Annals of Mathematics*. https://doi.org/10.2307/1969644

[3] Itô, K., & McKean, H. P. (1963). Brownian motions on a half line. *Illinois Journal of Mathematics*. https://doi.org/10.1215/ijm/1255644633

[4] Nie, Y., & Linetsky, V. (2020). Sticky reflecting Ornstein-Uhlenbeck diffusions and the Vasicek interest rate model with the sticky zero lower bound. *Stochastic Models*. https://doi.org/10.1080/15326349.2019.1630287

[5] Hertrich, M., & Zimmermann, H. (2017). On the credibility of the euro/Swiss franc floor: a financial market perspective. *Journal of Money, Credit and Banking*. https://doi.org/10.1111/jmcb.12390

[6] Jermann, U. J. (2017). Financial markets' views about the euro–Swiss franc floor. *Journal of Money, Credit and Banking*. https://doi.org/10.1111/jmcb.12389

[7] Lera, S. C., & Sornette, D. (2015). Currency target-zone modeling: an interplay between physics and economics. *Physical Review E*. https://doi.org/10.1103/PhysRevE.92.062828

[8] Veestraeten, D. (2008). Valuing stock options when prices are subject to a lower boundary. *Journal of Futures Markets*. https://doi.org/10.1002/fut.20299

[9] Hertrich, M., & Veestraeten, D. (2013). Valuing stock options when prices are subject to a lower boundary: a correction. *Journal of Futures Markets*. https://doi.org/10.1002/fut.21612

[10] Ingersoll, J. E. (1996). Valuing foreign exchange rate derivatives with a bounded exchange process. *Review of Derivatives Research*. https://doi.org/10.1007/BF01531597

[11] Krugman, P. R. (1991). Target zones and exchange rate dynamics. *The Quarterly Journal of Economics*. https://doi.org/10.2307/2937922

[12] Delbaen, F., & Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*. https://doi.org/10.1007/BF01450498

[13] Garman, M. B., & Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of International Money and Finance*. https://doi.org/10.1016/S0261-5606(83)80001-1

[14] Abate, J., & Whitt, W. (1987). Transient behavior of regulated Brownian motion, I: starting at the origin. *Advances in Applied Probability*.
https://doi.org/10.2307/1427408

[15] Abate, J., & Whitt, W. (1995). Numerical inversion of Laplace transforms of probability distributions. *ORSA Journal on Computing*.
https://doi.org/10.1287/ijoc.7.1.36

[16] Meier, C., Li, L., & Zhang, G. (2021). Simulation of multidimensional diffusions with sticky boundaries via Markov chain approximation. *arXiv*. arXiv:2107.04260

[17] Anagnostakis, A. (2023). Pricing and hedging for a sticky diffusion. *arXiv*. arXiv:2311.17011

Zamrik Quant Research
The research arm of Zamrik Quant Lab

About the Desk

Derivatives & Volatility Desk

Zamrik Quant Research

Options, volatility and rates: pricing, hedging and the free boundary.

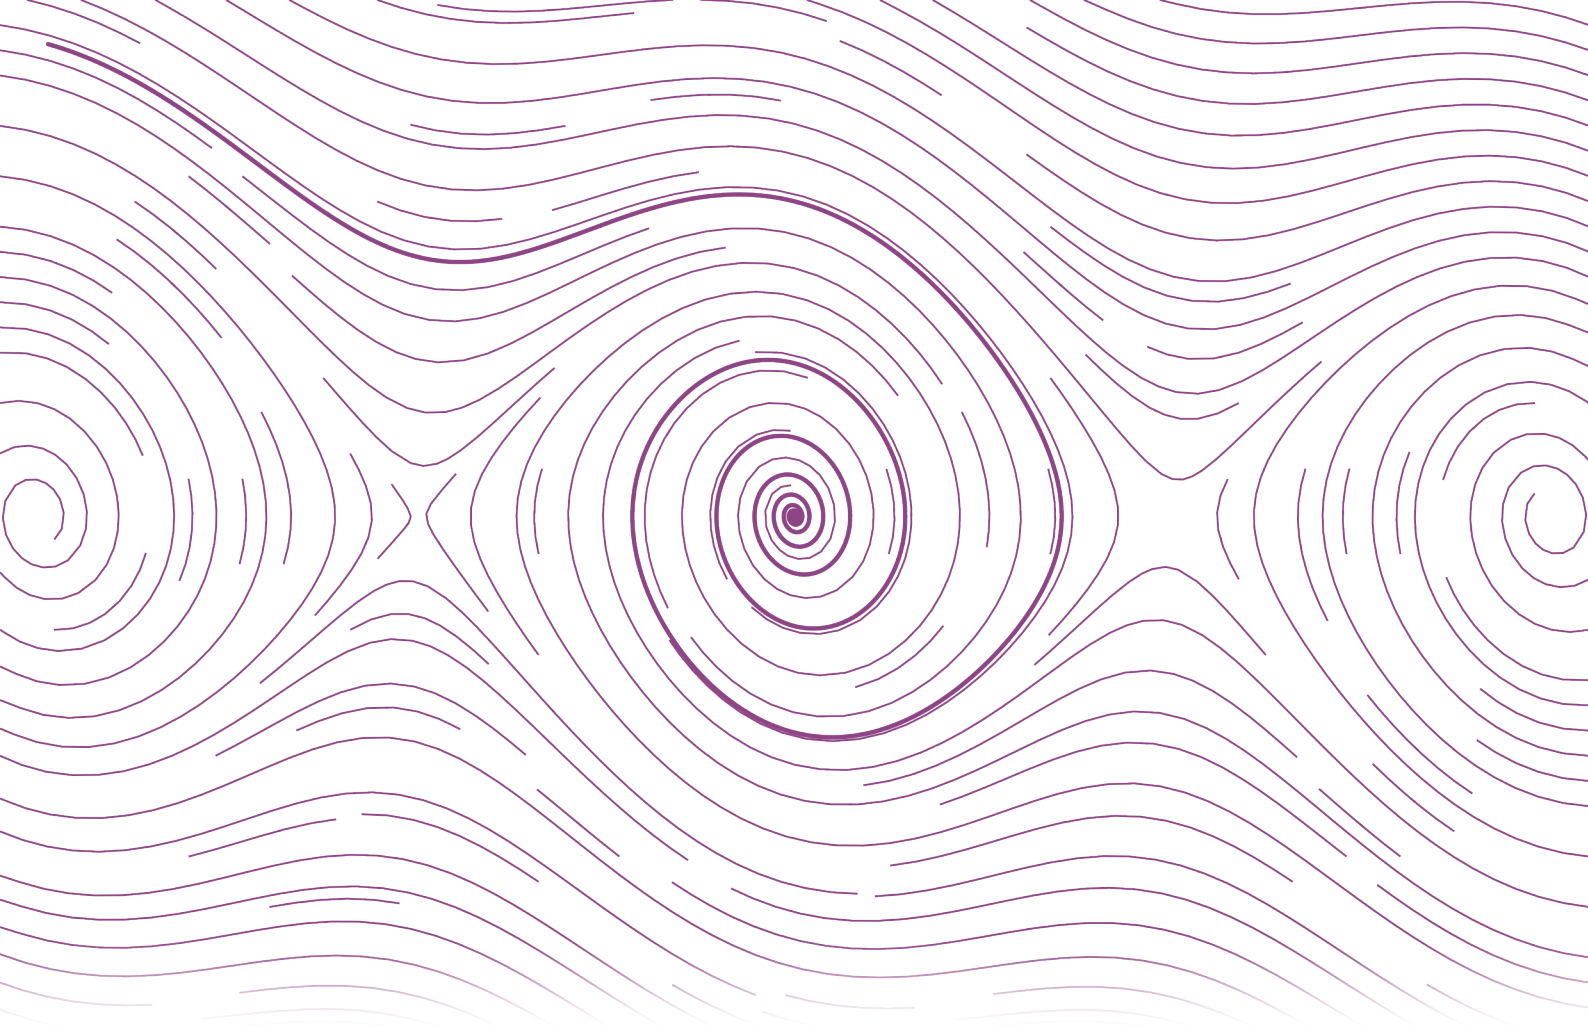
A derivative is a promise whose value depends on a path nobody controls. The Derivatives & Volatility desk works out what such promises are worth, how to hedge them, and when their holder should act.

The foundations come first: arbitrage, risk-neutral measures and the entropy view of them, and the benchmark approach, which prices with the growth-optimal portfolio instead. From there the desk takes on contracts whose exercise is itself a decision — the Russian option on a running maximum, and the Israeli option that either side may cancel — solved two independent ways, by Hamilton–Jacobi–Bellman equations and by reflected backward SDEs, checked against each other, and extended to ambiguity about the drift.

Volatility is modelled as a random process in its own right, including Wishart stochastic volatility for several assets at once, linked to the order flow that drives it. Rates are treated in the Heath–Jarrow–Morton framework, through bond prices and the derivatives written on forward rates. The desk also asks how much risk a position should carry at all, through the Kelly criterion and the growth-optimal portfolio, and its essays explain products that surprise their owners, such as leveraged funds whose long-run return falls short of their multiple.

The desk is supervised by **Dr. T. Zamrik**, Head of Zamrik Quant Research, PhD, Imperial College London (2011).

Notes



Phase portrait of the damped pendulum: $\dot{\theta} = \omega,\ \dot{\omega} = -\sin\theta - \delta\,\omega, \quad \delta = 0.21$

ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



WP-2024-33393745
zamrik.com