

Price Limits Hide Volatility: Censored GARCH Evidence from the Pakistan Stock Exchange

Muhammad Uzair

September 2026

Abstract

Daily price limits censor the returns exchanges report, which can bias GARCH estimates of volatility. I extend the expected-shock censored-GARCH approach of Morgan and Trevor (1999) to Student- t innovations and validate it by Monte Carlo, showing that a common shortcut, updating the variance recursion with the clipped observed return rather than its conditional expectation, over-corrects, more so under fat tails. Applying the corrected model to 449 distinct companies on the Pakistan Stock Exchange (2016–2026), I find that ignoring censoring understates shock sensitivity α by a median of 7.6% under Student- t GARCH (17.3% under normal GARCH), rising with how often a stock’s price hits its limit; the effect is a highly significant central tendency (sign test $p < 10^{-9}$) rather than a bound that holds for every stock; naive α is not always understated. Consistent with PSX’s Re. 1 minimum band, cheap stocks censor less often than the percentage rule alone implies. The 2020 widening of the band is associated with lower censoring and a smaller bias in affected stocks relative to floor-bound stocks whose effective band did not change; a placebo test fails, so I report this result using an illustrative deterministic bound in the spirit of Rambachan and Roth (2023) rather than as clean causal evidence.

JEL classification: C22, C58, G14, G15

1 Introduction

Many exchanges cap how far a price can move in one day. On a day a stock wants to move further than the limit allows, the exchange only reports the limited move, not the true one: the return is censored. This matters for GARCH, which learns volatility from past squared returns. If large moves are censored, the model concludes the stock is calmer than it is.

PSX’s limit is $\max(\pi_t P_{t-1}, 1)$ rupee, where π_t moved from 5% to 7.5% in 2020 and to 10% in 2024. The Re. 1 floor makes the effective band for a cheap stock wider than π_t implies, giving a built-in source of variation to study.

Morgan and Trevor (1999) show how to estimate GARCH under censoring by replacing the censored squared return with its conditional expectation; Wei (2002) and Hsieh and Yang

(2009) extend this line of work. I contribute in three ways. First, I extend the expected-shock approach to Student- t innovations, deriving the truncated second moment in closed form; PSX returns are heavy-tailed (median $\hat{\nu} \approx 4.7$). Second, I show by Monte Carlo that a natural shortcut, a censored likelihood combined with an uncorrected variance recursion, over-corrects, and more so under fat tails. Third, I apply the corrected model to the full PSX equity market: 468 stocks over three price-limit regimes.

Ignoring censoring understates α by a median of about 8% under Student- t GARCH and about 17% under normal GARCH, growing with a stock’s limit-hit frequency, though the direction reverses for a meaningful minority of individual stocks. Cheap stocks censor less often than the percentage rule implies, consistent with the Re. 1 floor. The 2020 widening is associated with lower censoring and a smaller bias in affected stocks relative to floor-bound controls, though a failed placebo test means I treat this as suggestive rather than causal.

2 Data and institutional setting

PSX sets each day’s band from the previous close, $b_t = \max(\pi_t P_{t-1}, 1)$. The percentage π_t moved in five steps from 5% to 7.5% (Jan–Mar 2020) and from 7.5% to 10% (May–Jul 2024), and was otherwise flat. The Re. 1 floor binds below Rs. 20 at the 5% regime, Rs. 13.33 at 7.5%, and Rs. 10 at 10%: below this price every stock faces the same rupee band regardless of price, so the effective band, as a share of price, is wider the cheaper the stock.

I use daily OHLCV data for every PSX-listed equity (`psxdata`), September 2016 to September 2026. From 1,028 listed instruments I drop non-equity instruments (659 remain, 652 with usable data). The data vendor returns a company’s complete price history under each of its historical ticker codes after a rename or re-listing; I identify 27 such duplicate-alias codes (e.g. BYCO/CENERGY, ICI/LCI) by near-perfect correlation ($r > 0.999$, median absolute difference $< 2\%$) between two symbols’ closing prices on overlapping dates, and keep only the longer-history alias for each. After this and requiring at least 500 trading days and a zero-volume share below 20%, and excluding preference shares, 449 distinct companies remain (Table 1). A day is censored when the close lies within one tick of that day’s limit; a day is a breach when the close lies outside the band entirely (a corporate action or data error). Breaches are excluded from the likelihood and, in the variance recursion, are treated identically across all three models below: the recursion carries forward the current conditional variance rather than the (possibly spurious) observed squared return.

Table 1: Sample construction

	Value
Listed instruments $\rightarrow$ equities $\rightarrow$ final sample	1,028 $\rightarrow$ 659 $\rightarrow$ 449
Duplicate ticker-alias codes excluded	27
Stock-days (final sample)	$\approx 834,000$
Median trading days per stock	2,028
Median censoring rate / floor-bound share	4.8% / 3.4%

3 Models

Let $r_t^* = \sigma_t z_t$ be the latent return, with $\sigma_t^2 = \omega + \alpha e_{t-1} + \beta \sigma_{t-1}^2$ and z_t standard normal or Student- t_ν . I observe $r_t = \min(\max(r_t^*, \ell_t), u_t)$, where u_t, ℓ_t are the log-return limits implied by b_t . A day is censored when r_t sits on a limit.

Naive ignores the limits: $e_{t-1} = r_{t-1}^2$, standard likelihood.

Observed-shock censored uses the correct censored likelihood (an upper hit contributes $1 - F(u_t/\sigma_t)$, a lower hit $F(\ell_t/\sigma_t)$) but still sets $e_{t-1} = r_{t-1}^2$ in the recursion.

Expected-shock censored (MT) follows Morgan and Trevor (1999) and replaces the unobserved shock with $e_t = s_t^2 \mathbb{E}[z^2 \mid z \geq c]$, $c = u_t/s_t$. Under normality this is $1 + c\phi(c)/(1 - \Phi(c))$. For $X \sim t_\nu$ I derive, by integration by parts,

$$\mathbb{E}[X^2 \mathbf{1}\{X > c\}] = \frac{\nu}{\nu-2} S_{\nu-2}(c \sqrt{\frac{\nu-2}{\nu}}) + \frac{cf_\nu(c)(\nu+c^2)}{\nu-1},$$

dividing by $S_\nu(c)$ for the conditional moment; this matches numerical integration to 10^{-6} .

All models are fitted by maximum likelihood (Nelder–Mead; $\omega > 0$, $\alpha, \beta \geq 0$, $\alpha + \beta < 1$, $2 < \nu < 200$). Breach days are excluded from the likelihood and set to $e_t = \sigma_t^2$ in the recursion. I validate the models by Monte Carlo: GARCH(1,1), $\alpha = 0.15$, $\beta = 0.80$, 2,500 days, 50 replications at three censoring levels, under normal and t_5 innovations (Table 2, Figure 1). The naive model’s bias in $\hat{\alpha}$ grows with censoring (to -0.040); the observed-shock model overshoots (to $+0.057$ under fat tails); the MT model stays close to the truth throughout ($+0.003$ to $+0.009$), though it has a small upward bias under t innovations and slightly worse RMSE than naive at low censoring.

Table 2: Monte Carlo bias in $\hat{\alpha}$ (true $\alpha = 0.15$), 50 reps

Innovations	Censoring	Naive	Observed-shock	MT
Normal	2.0%–11.6%	−0.012 to −0.040	+0.011 to +0.042	+0.003 to +0.006
t_5	2.6%–9.2%	−0.001 to −0.036	+0.037 to +0.057	+0.007 to +0.009

4 Results

The Re. 1 floor and censoring frequency. Among the 233 stocks with both floor- and percentage-bound days, 135 (57.9%) censor less often on floor-bound days (sign test $p = 0.018$; median within-stock gap -0.007 ; pooled censoring rate 4.0% on floor days vs. 5.7% on percentage days), consistent with the floor mechanically loosening the effective band for cheap stocks. Pooled rates are descriptive only, since stock-days are not independent.

Understatement of shock sensitivity. Table 3 reports the gap $\hat{\alpha}_{\text{MT}} - \hat{\alpha}_{\text{naive}}$, both as a share of the naive estimate (how much larger the corrected estimate is) and as a share of the corrected estimate (the share of the true sensitivity that naive misses – the number that supports an “understatement” claim). Under Student- t GARCH the gap is positive in 65–69% of stocks (sign test $p < 10^{-9}$ across specifications), naive missing a median 7.6–10.4% of the corrected estimate. Under normal GARCH the effect is both more prevalent (97.0%) and




Figure 1: Monte Carlo bias in $\hat{\alpha}$ vs. censoring rate, by innovation distribution. Naive attenuates and worsens with censoring; the observed-shock correction overshoots, more so under fat tails; MT stays close to zero throughout.

larger (17.3%), consistent with the Monte Carlo finding that part of the apparent bias under normality is fat tails rather than censoring alone. The gap rises with a stock’s censoring rate (Spearman $\rho = 0.42$ – 0.74). The lower quartile of the Student- t (main) distribution is negative (IQR -4.3% to 24.8%): for a meaningful minority of stocks naive α is not understated, so the median effect should be read as a strong central tendency, not a bound that holds stock by stock.

Table 3: Understatement of α : naive vs. MT GARCH

Specification	n	Gap > 0	Share missed by naive	IQR	Spearman ρ
Student- t (main)	416	65.4%	7.6%	-4.3% to 24.8%	0.42
Student- t (strict)	277	69.3%	10.4%	-3.5% to 27.6%	0.62
Normal	436	97.0%	17.3%	7.7% to 26.7%	0.74

Evidence from band changes. Floor-bound stocks give a natural control group unaffected by the 2020 and 2024 rule changes. For 2020, censoring fell 4.6 percentage points more, and the α gap fell 0.044 more, in treated stocks than controls ($p < 10^{-4}$ for both). A placebo test on a matched 2018 window with no rule change shows a smaller but non-zero divergence (censoring +4.2pp, $p < 10^{-6}$; gap +0.021, $p = 0.02$ – 0.08 depending on test), violating parallel trends. I therefore report an illustrative deterministic bound in the spirit of Rambachan and Roth (2023) – $\hat{\theta} \pm M \cdot |\text{placebo}|$ – rather than their full honest-DiD confidence-set procedure, which would require a constrained optimization over the joint pre/post covariance structure that this sample does not attempt. On this bound, the censoring result is fragile (a placebo-sized violation only 1.10 times as large would erase it) while the α -gap result, the one of substantive interest, is more robust (breakdown at 2.06 times the placebo violation). The 2024 change replicates the censoring result ($p < 10^{-10}$) but not the α -gap result ($p = 0.85$; control group $n = 15$), which I treat as inconclusive rather than a null result given the small




Figure 2: Absolute α gap (latent – naive) vs. censoring rate, Student- t main specification. Six of 416 points fall outside the plotted range and are noted rather than silently dropped.

control group.

Table 4: Band changes: treated – control, first differences

Event	$n_{\text{treat}}/n_{\text{ctrl}}$	Δ censoring	Δ α gap
Placebo (2018)	91/42	+0.042 ($p < 10^{-6}$)	+0.021 ($p = 0.02\text{--}0.08$)
2020 widening	174/51	−0.046 ($p < 10^{-14}$)	−0.043 ($p < 10^{-4}$)
2024 widening	208/15	−0.031 ($p < 10^{-10}$)	+0.003 ($p = 0.85$)

Robustness. The direction and significance of the main result are stable whether boundary-adjacent fits (near-zero α , near-integrated persistence) are included or excluded (Table 3, main vs. strict): the Spearman correlation with censoring strengthens (0.42 to 0.62) once those fits are dropped, though the median effect size also rises, so I report both specifications rather than only the more favorable one.

5 Conclusion

Price limits are common in emerging and several developed markets, and they censor the returns standard volatility models take as given. I show, by Monte Carlo and on 449 distinct PSX companies, that replacing the censored shock with its conditional expectation recovers volatility parameters that a naive model typically understates and a cruder correction overstates, by a median of 8–17% depending on the assumed return distribution. This is a strong central tendency, not a universal one: for a meaningful minority of individual stocks

the naive estimate is not understated at all, so researchers estimating GARCH volatility in limit-order markets without accounting for censoring should treat naive α as typically, not uniformly, biased downward – the direction is well established, the size is heterogeneous across stocks, and it should not be read as a per-stock lower bound.

Data and code availability. Data pipeline, estimation, simulation code and all results are available at <https://github.com/uzair1509/psx-censored-volatility>.

AI-use statement. AI tools (Claude, Anthropic) assisted with code implementation, debugging and drafting. The research question, hypotheses, model choices, interpretation and conclusions are the author’s own.

References

- [1] Morgan, I.G. and Trevor, R.G. (1999). Limit moves as censored observations of equilibrium futures prices in GARCH processes. *Journal of Business & Economic Statistics*, 17(4), 397–408.
- [2] Wei, S.X. (2002). A censored-GARCH model of asset returns with price limits. *Journal of Empirical Finance*, 9(2), 197–223.
- [3] Hsieh, P.-H. and Yang, J.J. (2009). A censored stochastic volatility approach to the estimation of price limit moves. *Journal of Empirical Finance*, 16(2), 337–351.
- [4] Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5), 2555–2591.