

Target Volatility Strategies as Mean-Preserving Contractions

Kais Dachraoui

Independent Financial Advisor, Ottawa, Canada – kdachraoui1@gmail.com

Abstract

This note establishes a distributional result for target-volatility strategies. Under a multiplicative scale representation and conditional mean-scale stability, with the target set to the average scale, the TVS return equals the conditional mean of the underlying return. The underlying return is therefore a mean-preserving spread of the TVS return, implying dominance in concave order and, in particular, second-order stochastic dominance. TVS is therefore not merely a variance-reduction device; it is a mean-preserving contraction of the underlying return distribution. This is a risk-reduction result, not a return-premium one.

Keywords: *target volatility strategies; second-order stochastic dominance; mean-preserving spread; volatility timing; risk management*

JEL classification: G11; G17; D81; C02

1. Introduction

Target volatility strategies (TVS) adjust exposure to an underlying index inversely with an estimate or forecast of its volatility, formed using information available before the return is realized, so as to hold total portfolio volatility near a fixed target σ_T . Empirically, Dopfel and Ramkumar (2013) document that high-volatility regimes coincide with lower risk-adjusted equity returns, and Hocquard, Ng, and Papageorgiou (2013) show that a constant-volatility overlay normalizes negatively skewed return distributions and curbs maximum drawdown, an early empirical link between volatility targeting and the shape of the return distribution. Existing theory (Dachraoui, 2018; Moreira and Muir, 2017), however, evaluates TVS primarily through its effect on expected return or the Sharpe ratio. Such results primarily characterize moments of the return distribution rather than establishing preference over the full distribution for risk-averse investors.

This note establishes such a distributional result. Under a multiplicative scale representation of returns and a conditional mean-scale stability condition, and provided the volatility target equals the average scale, Y_t and X_t have identical expected value, and the TVS return second-order stochastically dominates the underlying index return for every risk-averse investor; the dominance

is a pure risk-reduction result, not a return-premium one. It in fact holds in the stronger concave order: the underlying return is a mean-preserving spread of the TVS return in the sense of Rothschild and Stiglitz (1970), so preference for the TVS return holds for every concave utility function, not only non-decreasing ones.

The empirical literature on volatility-managed portfolios documents favorable distributional effects consistent with this mechanism: Moreira and Muir (2017) show that scaling exposure inversely with recent variance raises Sharpe ratios and reduces the frequency of extreme drawdowns across equity and factor portfolios, while Cirelli, Vitali, Ortobelli Lozza, and Moriggia (2017) test a specific discontinuous target-volatility rule directly against stochastic dominance criteria and find it preferred by risk-averse investors empirically. This note instead derives the ordering analytically, as a structural consequence of the return's scale representation rather than a property of one rule verified in one sample. The result is therefore distinct from the question of whether volatility-managed strategies generate positive realized performance: Cederburg, O'Doherty, Wang, and Yan (2020) find that many volatility-managed factor portfolios fail to outperform out of sample, and Barroso and Detzel (2021) find that transaction costs erode most such gains outside the market portfolio, yet Proposition 3 requires no positive alpha, since $E[Y_t] = E[X_t]$ exactly.

2. Setup

Let $X_t := r_t - r^f$ denote the excess return of an Index portfolio over the risk-free rate, and let $S_t > 0$ denote the conditional volatility signal for X_t , formed using information available at $t - 1$, used to set exposure $L_t = \sigma_T / S_t$, as in the target-volatility signal rule of Dachraoui (2018). The TVS excess return is $Y_t := L_t X_t$.

Definition (Predetermined scale decomposition).

$$Z_t := X_t / S_t,$$

so that $X_t = S_t Z_t$ holds automatically, with Z_t interpreted as the scale-adjusted shock. This is not a claim that S_t equals some latent “true” volatility: for *any* predetermined positive process S_t ; a trailing realized-volatility estimate, an EWMA, a GARCH forecast, or any other signal observable at $t - 1$; this decomposition applies without alteration.

Under Definition alone, the TVS return takes the simple form:

$$Y_t = L_t X_t = (\sigma_T / S_t) S_t Z_t = \sigma_T Z_t.$$

TVS replaces the random scale S_t with the fixed target σ_T while leaving the scale-adjusted shock Z_t untouched: $X_t = S_t Z_t$ becomes $Y_t = \sigma_T Z_t$.

Assumption (Conditional mean-scale stability).

$$E[S_t | Z_t] = E[S_t] \text{ a.s.}$$

This condition is strictly weaker than full independence of Z_t and S_t : it permits, for instance, the conditional dispersion of S_t given Z_t to vary with Z_t , so long as the conditional *mean* scale does not drift with the scale-adjusted shock. It is a same-period condition, distinct from the cross-period leverage effect of Black (1976), which concerns how *past* shocks predict *future* volatility.

In plain terms, the assumption requires that the average level of the predetermined volatility signal does not systematically depend on the size or sign of the subsequent residual. The name parses accordingly: “scale” refers to the signal S_t , “mean” to its conditional average $E[S_t | Z_t]$, and “stability” to that average remaining unchanged across realizations of Z_t .

In standard conditional-volatility models, such as GARCH and EWMA, the volatility signal S_t is measurable with respect to information available at $t - 1$. When S_t coincides with the model's conditional scale and the corresponding innovation is independent of the information used to construct S_t , then S_t and Z_t are independent outright, delivering Assumption immediately, since independence is stronger than mean-independence. The assumption itself only requires the weaker property: it allows the scale and the realized scale-adjusted shock to be dependent, but rules out dependence in the conditional mean of the scale given it.

We also set the volatility target equal to the mean scale:

$$\sigma_T = E[S_t].$$

The choice $\sigma_T = E[S_t]$ is deliberate: it isolates the distributional effect of volatility scaling by holding the mean volatility scale fixed. For targets different from $E[S_t]$, TVS combines this distributional effect with a change in the average scale, and the mean-preserving-spread result no longer follows in general.

Proposition 1 (Central identity). For $\sigma_T = E[S_t]$, under Assumption,

$$Y_t = E[X_t | Z_t] \text{ a.s.}$$

Proof. From above, $Y_t = \sigma_T Z_t = E[S_t] Z_t$. Separately, $E[X_t | Z_t] = E[S_t Z_t | Z_t] = Z_t E[S_t | Z_t] = Z_t E[S_t]$, using that Z_t factors out of its own conditioning σ -field regardless of any dependence structure, and Assumption for the last step. The two expressions coincide. ■

3. Main result

Proposition 2 (Mean-preserving-spread decomposition). Under Proposition 1,

$$X_t = Y_t + \varepsilon_t, \quad \varepsilon_t := (S_t - E[S_t]) Z_t, \quad E[\varepsilon_t | Z_t] = 0 \text{ a.s.}$$

Proof. $E[\varepsilon_t | Z_t] = E[X_t | Z_t] - Y_t = Y_t - Y_t = 0$ by Proposition 1. ■

This provides the Rothschild and Stiglitz (1970) mean-preserving-spread representation of X_t as Y_t plus a conditionally mean-zero random scale component: X_t is a mean-preserving spread of Y_t . Since $Y_t = \sigma_T Z_t$ with $\sigma_T > 0$, conditioning on Y_t is equivalent to conditioning on Z_t , so $E[\varepsilon_t | Z_t] = 0$ is the same statement as $E[\varepsilon_t | Y_t] = 0$. Equivalently, Y_t is a mean-preserving *contraction* of X_t : the same mean, with the dispersion reduction made precise in Proposition 4 and the preference ordering in Proposition 3.

Proposition 3 (Concave-order dominance). Under Proposition 1, for every concave function u (not necessarily non-decreasing),

$$E[u(X_t)] \leq E[u(Y_t)].$$

That is, Y_t dominates X_t in the concave order. Restricting to non-decreasing concave u gives, as a corollary, that Y_t second-order stochastically dominates X_t .

Proof. By the conditional Jensen inequality, which requires only concavity of u , not monotonicity,

$$E[u(X_t) | Z_t] \leq u(E[X_t | Z_t]) = u(Y_t),$$

using Proposition 1 in the equality. Taking expectations over Z_t gives $E[u(X_t)] \leq E[u(Y_t)]$. ■

Strictness. Since $E[\varepsilon_t | Z_t] = 0$, the law of total variance gives $\text{Var}(\varepsilon_t) = E[\text{Var}(X_t | Z_t)]$ (used again in Proposition 4 below). Whenever

$$\text{Var}[(S_t - E[S_t]) Z_t] > 0,$$

the concave-order dominance is strict for every strictly concave utility function for which the relevant expectations exist.

Proposition 4 (Equal mean, lower variance). Under Proposition 1,

$$E[Y_t] = E[X_t]$$

(since $E[Y_t] = \sigma_T E[Z_t] = E[S_t] E[Z_t] = E[S_t Z_t] = E[X_t]$), and, applying the law of total variance to $Y_t = E[X_t | Z_t]$,

$$\text{Var}(X_t) = \text{Var}(Y_t) + E[\text{Var}(X_t | Z_t)],$$

so $\text{Var}(Y_t) \leq \text{Var}(X_t)$, strictly whenever $E[\text{Var}(X_t | Z_t)] > 0$. In this setting, Proposition 3's dominance implies the corresponding equal-mean variance reduction, but is strictly stronger than that mean-variance comparison, since it holds for every concave utility, not only quadratic ones.

4. Discussion

The central implication of the analysis is that volatility scaling can improve the distribution of returns without requiring an increase in expected return. The TVS return is the conditional mean of the underlying return, with the difference between the two representing a component of return variation associated with the random scale. The strategy therefore does not create a new source of return; it removes variation around the conditional mean.

This distinction is important for interpreting the economic value of target-volatility strategies. A higher Sharpe ratio or lower volatility can suggest an improvement in risk-adjusted performance, but neither result by itself establishes that investors with general concave preferences would prefer the transformed return. The concave-order result provides a stronger statement. It shows that the reduction is not specific to variance or to a particular choice of utility function: the entire return distribution is improved in the sense relevant for risk-averse preferences. The familiar reduction in variance is consequently only one implication of a more general distributional contraction.

The decomposition also clarifies the role of volatility scaling. When the underlying scale is above its average, the strategy reduces exposure; when it is below its average, the strategy increases exposure. These adjustments do not systematically favor positive or negative realizations. Rather, they replace the variable scale of the underlying return with a fixed scale while preserving the

conditional mean. The resulting gain comes from eliminating scale variation that is conditionally mean-zero, rather than from forecasting the direction of returns.

This perspective also connects the result to the empirical literature. Previous studies document improvements in Sharpe ratios, drawdowns, and skewness for volatility-managed portfolios, while others question the persistence of their realized performance; the analysis here identifies a structural condition under which such distributional improvements follow from the scale transformation itself, independently of whether a particular implementation generates positive alpha in a given sample. In this sense, it complements rather than competes with the empirical evidence, offering a distributional interpretation of the patterns reported by Dachraoui (2018) and Hocquard, Ng, and Papageorgiou (2013) without requiring those particular patterns to hold in every implementation.

5. Conclusion

This note provides a distributional characterization of target-volatility strategies, showing that their benefit extends beyond lower variance to a broader improvement in the distribution of investment outcomes. This shifts the interpretation of volatility targeting from a strategy that may improve risk-adjusted performance to one that can be understood as a mean-preserving contraction of return risk, providing a theoretical benchmark for interpreting existing empirical evidence and for studying extensions involving more realistic volatility signals and implementation constraints.

References

- Barroso, P., Detzel, A., 2021. Do limits to arbitrage explain the benefits of volatility-managed portfolios? *Journal of Financial Economics* 140 (3), 744–767.
- Black, F., 1976. Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association*, 177–181.
- Cederburg, S., O'Doherty, M.S., Wang, F., Yan, X., 2020. On the performance of volatility-managed portfolios. *Journal of Financial Economics* 138 (1), 95–117.
- Cirelli, S., Vitali, S., Ortobelli Lozza, S., Moriggia, V., 2017. A conservative discontinuous target volatility strategy. *Investment Management and Financial Innovations* 14 (2-1), 176–190.
- Dachraoui, K., 2018. On the optimality of target volatility strategies. *Journal of Portfolio Management, FBI Special Issue*, 1–10.

Dopfel, F.E., Ramkumar, S.R., 2013. Managed volatility strategies: applications to investment policy. *Journal of Portfolio Management* 40 (1), 27–39.

Hocquard, A., Ng, S., Papageorgiou, N., 2013. A constant-volatility framework for managing tail risk. *Journal of Portfolio Management* 39 (2), 28–40.

Moreira, A., Muir, T., 2017. Volatility-managed portfolios. *Journal of Finance* 72 (4), 1611–1644.

Rothschild, M., Stiglitz, J.E., 1970. Increasing risk: I. A definition. *Journal of Economic Theory* 2 (3), 225–243.