Article from

Risk and Rewards

February 2018

Issue 71

ISSUE 71 FEBRUARY 2018

Risks & Rewards

SOCIETY OF ACTUARIES
INVESTMENT SECTION

Hedging Variable Annuities: How Often Should the Hedging Portfolio be Rebalanced?

By Maciej Augustyniak and Mathieu Boudreault

In the last decade, many insurers have implemented dynamic hedging programs to defend against market risks embedded in their variable annuity (VA) blocks of business. At the core of these programs are the so-called Greeks which correspond to price sensitivities with respect to various market risks such as movements in equity indices, interest rates and volatility. These Greeks indicate to the insurer how much to invest in equities, bonds and financial derivatives to offset market exposures in its VA contracts. Due to changes in market factors, Greeks vary in time and the insurer is therefore required to rebalance its hedging portfolio (i.e., adjust its hedging positions) periodically to ensure that the hedging strategy is achieving its objective.

When managing a VA hedging program, the choice of the rebalancing frequency is an important practical issue because of the high monitoring and trading costs that ensue when hedging positions are revised. It is well-known that in a Black-Scholes world hedging more frequently reduces the hedging error. In fact, groundbreaking work in financial theory showed that this error can theoretically be eliminated in a Black-Scholes setting with a continuously rebalanced delta hedge. However, in the real world perfect hedging is generally not feasible due to sudden price jumps, to market frictions, to the impossibility of trading in continuous time and to the limited availability of traded assets. Therefore, hedging

CONTINUED ON PAGE 3

in real market conditions entails a risk. It would be tempting to conclude based on Black-Scholes theory that this risk can be reduced with a more frequent rebalancing of the hedging portfolio. However, this is not necessarily the case because every hedging strategy carried out in the real world is exposed to model risk; that is, there is inevitably a discrepancy between the insurer's hedging model used to compute Greeks and the true (unknown) financial model or data-generating process. Consequently, adjusting hedging positions too often with the wrong model can lead to a larger accumulation of hedging errors than if less frequent revisions were made. This issue is especially important to investigate in the context of VAs because hedging is performed over long-term periods.

The objective of this article is to investigate how the choice of the rebalancing frequency in a VA hedging program impacts hedging effectiveness. More precisely, we examine the performance of daily, weekly, monthly and move-based delta hedging strategies for managing the underlying equity risk of a simple guaranteed minimum accumulation benefit (GMAB) VA indexed to historical S&P 500 returns. This allows us to conduct a back-testing exercise and determine what choice of rebalancing strategy would have been preferable to use in the past. Overall, we find that a monthly rebalanced delta hedging strategy consistently led to the smallest losses when dynamically hedging 10-year GMAB contracts maturing in the period 1990–2017. It must be emphasized that this conclusion is valid with and without transaction costs. Therefore, recent empirical evidence strongly favors a less frequent rebalancing of the hedging portfolio and we examine some explanations of this phenomenon.

The objective of this article is to investigate how the choice of the rebalancing frequency in a VA hedging program impacts hedging effectiveness.

GMAB CONTRACT AND ASSUMPTIONS

We assume that the insurer sells 10-year VA contracts with a GMAB rider. The value of the VA account in time is denoted by $\{A_t: t = 0,1, \ldots, T\}$, where t is measured in trading days from inception of the contract. Since there are approximately 252 trading days in each calendar year, the term-to-maturity of the contract is set to $T = 2520$ days. The VA account is invested in an investment fund, denoted by $\{S_t: t = 0,1, \ldots, T\}$ (in our hedging experiment, this investment fund will mimic historical returns on the S&P 500 price index). We assume an initial investment of $A_0 = S_0 = \$100$. The GMAB rider ensures that the policyholder will be able to recover the greater of the account value A_T and a guaranteed amount $G = 116$ at maturity (the guaranteed amount corresponds to the initial investment accumulated at an effective annual roll-up rate of 1.5 percent). The GMAB rider therefore creates a liability for the insurer in the form of a long-term put option guarantee; the insurer's liability at maturity is $\max(G - A_T, 0)$. This guarantee is financed via a fee withdrawn daily as a fraction of the account value at an annual nominal rate of $\alpha = 2\%$, that is, at the beginning of each trading day the insurer withdraws $A_t(\alpha/252)$ from the account value. As a result, fee cash flows are risky and should be hedged along with the guarantee. The relationship between the investment fund S_t and the VA account A_t at time t (right before the withdrawal of fees) is therefore given by:

$$A_t = S_t(1 - \alpha/252)^t.$$

Finally, we suppose that the VA contract is held to maturity (i.e., surrender and death are not possible) and assume a continuously compounded annual risk-free rate of r = 3%.

HEDGED LOSS

If the insurer does not use a hedging strategy, its **unhedged loss** on the VA contract at maturity, denoted by L_T, corresponds to the payoff on the GMAB rider less accumulated fees that were collected throughout the contract:

$$\begin{aligned} L_T &= \text{GMAB payoff} - \text{accumulated fees} \\ &= \max(G - A_T, 0) - \sum_{t=0}^{T-1} A_t(\alpha/252)\, e^{r(T-t)/252}. \end{aligned}$$

To manage the market risk embedded in the GMAB rider, we assume that the insurer establishes a dynamic delta hedging strategy under the Black-Scholes model. This strategy entails

CONTINUE ON PAGE 4

holding a position of Δ_t (the Greek *delta*) in the fund S_t at time t (the computation of Δ_t is detailed in the following section). This can be accomplished using futures or, equivalently, by taking a long position in Δ_t shares of the underlying fund and borrowing the costs or lending the proceeds.

The **hedged loss** on the VA contract at maturity, denoted by HL_T, corresponds to:

$$\begin{aligned} HL_T = &\text{ unhedged loss} - \\ &\text{cumulative mark-to-market} \\ &\text{gains on the hedge} \\ = &\ L_T - H_T. \end{aligned}$$

The mark-to-market gain at time $t+1$ associated with the delta hedge established at time t is:

$$\Delta_t(S_{t+1} - S_t e^{r/252}).$$

The **cumulative mark-to-market gains** on the hedge, denoted by H_T, correspond to the accumulated values of these gains to maturity:

$$H_T = \sum_{t=0}^{T-1} \Delta_t\left(S_{t+1} - S_t e^{r/252}\right)e^{r(T-t-1)/252}.$$

The objective of the delta hedging strategy is to generate cumulative mark-to-market gains at maturity that will allow the insurer to offset its loss on the VA contract.

COMPUTATION OF DELTA

To achieve its objective, the delta hedging strategy must protect the insurer against changes in the net value of the VA contract due to fluctuations in the underlying investment fund S_t. In a Black-Scholes setting, the net value of the VA contract is computed as an expected present value (PV) under the risk-neutral measure. The net value of the VA contract at time t (in the eyes of the insurer), denoted by V_t, corresponds to:

$$\begin{aligned} V_t = &\text{ Black-Scholes put price} - \text{expected PV of future fees} \\ &- \text{past fees accumulated to time t.} \end{aligned}$$

Note that the first two terms on the right-hand side of this equation are a function of S_t (or A_t) whereas the last term is not. The position Δ_t is then defined as the first-order sensitivity of V_t with respect to a change in S_t:

$$\Delta_t = \frac{\partial V_t}{\partial S_t} = \Delta_t^{\text{put}} - [(1-\alpha/252)^t - (1-\alpha/252)^T],$$

where

$$\Delta_t^{\text{put}} = -(1-\alpha/252)^T \Phi(-d_1),$$

$$d_1 = \frac{\log\left(\frac{A_t(1-\alpha/252)^{T-t}}{G}\right) + (r + \sigma_t^2/2)(T-t)}{\sigma_t\sqrt{T-t}}$$

is the formula for the *delta* of a put option (a document detailing the derivation of Δ_t is available on the author's website). Note that delta hedging the net value of the contract entails hedging both the guarantee offered **and** future fee cash flows.

We remark that the hedge ratio Δ_t is computed from the above formula only when the hedging position is revised. Otherwise, when the portfolio is not rebalanced at time t, we simply set $\Delta_t = \Delta_{t-1}$.

DATA AND VOLATILITY CALIBRATION

The variable of interest in our hedging experiment is the insurer's hedged loss at maturity denoted by HL_T. The goal of our back-test is to compute the realized values of HL_T assuming that the VA is exposed to 10-year rolling S&P 500 daily return data over the period 1960–2017. More precisely, the first VA contract is assumed to be issued on Dec. 31, 1959, and matures 2520 trading days later on Feb. 13, 1970. The second contract is issued on the following trading day, that is, on Jan. 4, 1960, and matures on Feb. 16, 1970. The process then continues until the final 10-year period which begins on Aug. 29, 2007, and ends on Aug. 31, 2017. In total, we obtain 11,998 10-year return paths. It must be noted that since these paths are based on series of overlapping returns, the hedging losses computed on these paths are not all independent. Nevertheless, the back-testing experiment allows us to assess the effectiveness of delta hedging strategies over time and determine what choice of rebalancing strategy would have been preferable to use in the past.

The computation of the Black-Scholes hedge ratio Δ_t requires a volatility assumption σ_t. Due to the long-term nature of the contract, we allow this parameter to be time-varying. In fact, it would be unrealistic to use a constant volatility assumption over a 10-year period. In our analysis, we assume that σ_t is calibrated at time t to the past annualized realized volatility computed from daily returns over a three-year period (756 trading days). We experimented alternative ways to set this parameter and found that our conclusions are robust to different calibration methods.

RESULTS

Figure 1 (page 5) illustrates the results of our back-testing exercise. We consider daily, weekly, monthly and move-based delta hedging strategies. The daily, weekly and monthly strategies are rebalanced every one, five and 21 trading days, respectively. The move-based strategy is rebalanced only when the value of the hedge ratio Δ_t changes by more than 0.05 in absolute value.

The chart in the upper part of Figure 1 shows the insurer's hedged loss at maturity (HL_T) for 10-year VA contracts maturing every trading day over the period 1970–2017 (the horizontal axis corresponds to the contract's maturity date). Note that whenever $HL_T < 0$, the hedging strategy results in a terminal gain for the insurer. We have not incorporated transaction costs into the variable HL_T because we first want to evaluate the performance of the rebalancing strategies without imposing a penalty for more frequent trading. In other words, the deck is somewhat stacked in favor of the daily strategy. The impact of transaction costs is discussed separately in a subsequent section.

The chart in the lower part of Figure 1 displays the terminal fund value (A_T) for every contract maturity. The shaded areas in the charts indicate maturities where the VA contract terminated in-the-money (i.e. $A_T < G$).

For contracts maturing in the period 1970–1990, the daily rebalanced delta hedge led to the smallest hedging losses among the strategies considered. However, for contracts maturing after 1990, the tide turned and the monthly rebalancing scheme generally resulted in the best performance. The outperformance of this strategy is particularly evident for contracts maturing in the last decade. The move-based strategy never surpassed all of its competitors and performed particularly poorly during 1990–2000. We experimented with alternative threshold levels, but the overall performance of these move-based strategies remained inferior.

EXPLANATION OF RESULTS

The fact that a monthly rebalanced delta hedge displays the best performance over an extended period may at first sight seem surprising. After all, in a Black-Scholes setting a more frequent rebalancing leads to a more effective hedge. However, this well-known result derived from financial theory assumes that the hedger uses the true data-generating model to construct his positions, that is, the hedging strategy is not exposed to model risk.

In the past 50 years, the financial econometrics literature has vastly documented a set of statistical properties which are common to a large number of financial series: these are known as **stylized facts**. They include fat tails of the return's distribution and volatility clustering, among others (see Cont., 2001), and strongly contradict the assumption underlying the Black-Scholes model that financial assets follow geometric Brownian motions (i.e., returns are independent and identically distributed according to a normal distribution). Therefore, a Black-Scholes

Figure 1
Results of the Back-Testing Experiment

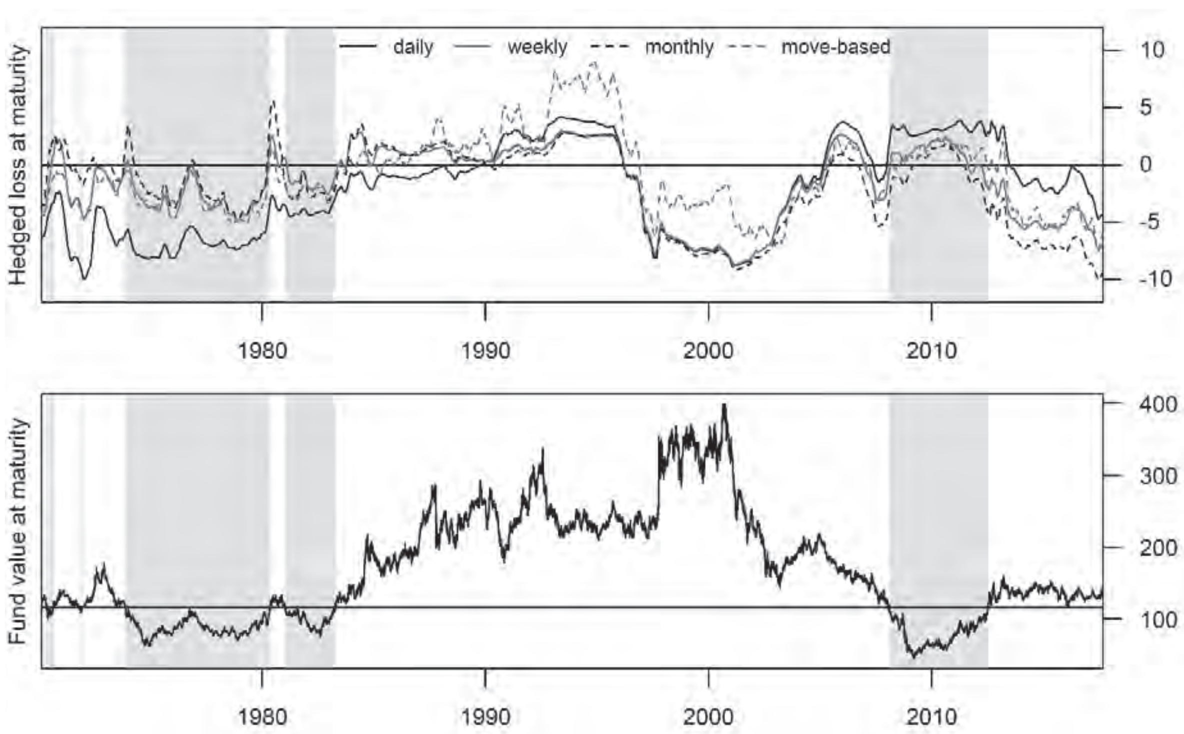

Figure 2
Annualized realized Volatilities

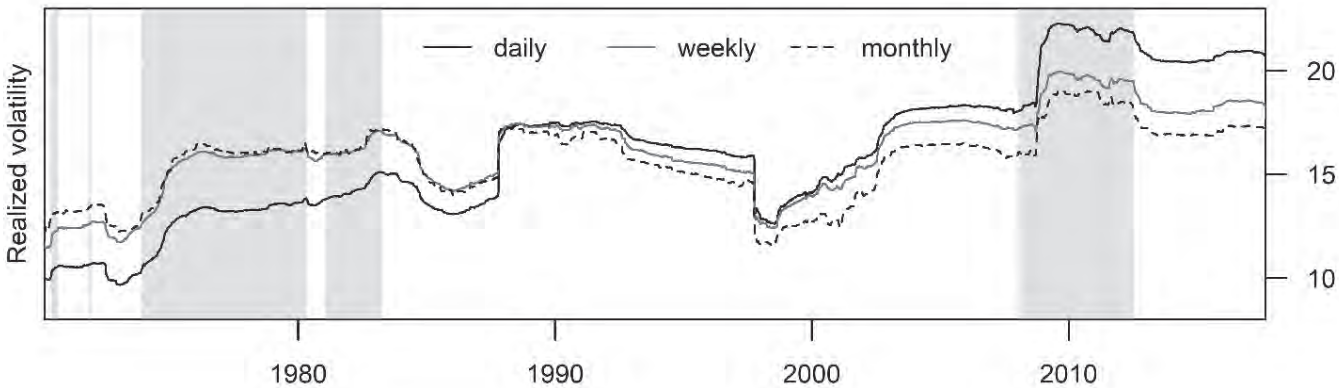

delta hedge in the real world is exposed to a large amount of model risk and there is no guarantee that conclusions derived in the idealized Black-Scholes setting will continue to hold in reality.

Aggregational Gaussianity (see Cont., 2001) is a stylized fact of financial data that stipulates that as one increases the time scale over which returns are calculated, their distribution looks more and more like a normal distribution. In particular, monthly returns tend to conform better to the Gaussian hypothesis than daily returns. One way to illustrate this statistically is to compare the kurtosis of daily and monthly returns (see also Table 1 of Boudreault, 2013). The kurtosis is a statistical measure of whether the data are heavy-tailed or light-tailed; data sets with high kurtosis tend to have heavy tails (data conforming to a Gaussian assumption have a kurtosis of three). Over the period 1995–2005, the kurtosis of S&P 500 daily returns is 6.1 versus 3.4 for monthly returns, whereas over the period 2007–2017, these numbers are 13.5 and 5.7, respectively. Consequently, a monthly Black-Scholes delta hedge is generally exposed to less model risk than a daily hedge.

A further reason that is perhaps more vital in explaining the better performance of the monthly hedge for contracts maturing after 1990 relates to the fact that S&P 500 daily returns exhibited from that time downward trending negative autocorrelations at short lags. For instance, during the 10-year period 2007–2017, the autocorrelations of S&P 500 daily returns at lags 1 and 2 were -10 percent and -6 percent, respectively. Such negative autocorrelations, although small, contribute to reducing the noise and volatility of aggregated returns.

Figure 2 illustrates the annualized realized volatilities of daily, weekly and monthly returns computed over rolling periods of 10 years (the horizontal axis indicates the date when the 10-year period ends). Note that daily volatilities are based on 2520 daily returns, whereas monthly volatilities are based on 120 returns constructed by aggregating daily returns over periods of 21 trading days. A monthly return therefore does not necessarily refer to the return in a calendar month.

We observe that for 10-year periods ending after 1990, the annualized volatility of monthly returns is below that of daily and weekly returns. This is a direct consequence of negative autocorrelations observed in daily returns. In fact, it can be shown (see Campbell et al., 1997, chapter 2) that the ratio of the annualized variance of h-period aggregated returns to one-period returns is theoretically equal to:

$$\frac{\text{annualized variance of } h\text{-period aggregated returns}}{\text{annualized variance of one-period returns}} = 1 + 2\sum_{k=1}^{h-1}(1 - k/h)\rho(k),$$

where $\rho(k)$ corresponds to the return autocorrelation at lag k. This result explains the pattern observed in Figure 2. For instance, returns prior to 1990 generally displayed positive autocorrelations at short lags (e.g., over the period 1970–1980, the return autocorrelation at lag 1 was 25 percent), which resulted in larger volatilities at a monthly frequency. Since the discrepancies between daily and monthly volatilities have been growing in recent 10-year periods, this indicates that autocorrelations have overall been trending downwards in recent years.

These observations offer an explanation as to why the monthly rebalanced Black-Scholes delta hedge performed better in our hedging experiment for contracts maturing after 1990: this

On average, the turnover for the daily rebalancing strategy was four times greater than the one for the monthly strategy. ...

strategy was exposed to returns exhibiting less noise and volatility. Moreover, the distribution of these returns was closer to the normal due to aggregational Gaussianity which implies a smaller degree of model risk in the hedging strategy. This also explains the underperformance of move-based strategies as they require more frequent rebalancing in periods of higher volatility/kurtosis (i.e., when returns further deviate from normality).

IMPACT OF TRANSACTION COSTS

The accumulated value of transaction costs to maturity can be taken as approximately proportional to the total turnover in the hedging position defined as:

$$\text{total turnover in the hedging position} = \sum_{t=1}^{T-1} S_t |\Delta_t - \Delta_{t-1}| e^{r(T-t)}.$$

On average, the turnover for the daily rebalancing strategy was four times greater than the one for the monthly strategy, which implies that transactions costs would be expected to be four times greater as well. Assuming that these costs are 0.25 percent times the turnover in the hedging position, the margin by which the daily strategy performed better than the monthly one for contracts maturing before 1990 is almost completely erased by trading frictions. Therefore, after accounting for transactions costs, there are essentially no 10-year periods in our hedging experiment where a daily rebalancing strategy performed significantly better than the others.

Finally, we note that the ratio of the turnover between the move-based and monthly strategies fluctuated between 0.5 and 1.5, which entails that the move-based method sometimes required less frequent trading than the monthly rebalancing scheme. However, whenever it involved less transaction costs, its performance still remained inferior to the monthly strategy.

CONCLUSION

Based on S&P 500 return data over the period 1960–2017, we have provided empirical evidence suggesting that hedging effectiveness may be improved by rebalancing the hedging portfolio less frequently than on a daily time scale. This conclusion emerges from three observations: (1) returns on larger time scales such as monthly are closer to being normally distributed than daily returns; this stylized fact known as aggregational Gaussianity implies that a Black-Scholes hedging strategy is exposed to less model risk at larger time scales, (2) negative autocorrelations in daily returns at short lags were observed in our data set; they imply some level of short-term mean reversion which contributes to reducing noise and volatility in aggregated returns, and (3) a more frequent rebalancing of the hedging portfolio entails larger transaction costs. ■

Maciej Augustyniak, FSA, ACIA, Ph.D., is an Assistant Professor at the Université de Montréal. He can be contacted at *augusty@dms.umontreal.ca.*

Mathieu Boudreault, FSA, ACIA, Ph.D., is an Associate Professor at the Université du Québec à Montréal. He can be contacted at *boudreault.mathieu@uqam.ca.*

REFERENCES

Boudreault, M. (2013). Pricing and hedging financial and insurance products Part 2: Black-Scholes' model and beyond. Risks & Rewards, February 2013, Society of Actuaries.

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). The Econometrics of Financial Markets. Princeton University Press, New Jersey.

Cont, R. (2001). Model uncertainty and its impact on the pricing of derivative instruments. Mathematical Finance, 16(3):519–547.

ACKNOWLEDGMENTS

This research was partly funded by *Autorité des Marchés Financiers*, the regulating body for insurance companies chartered in the province of Quebec, Canada.