Insurance: Mathematics and Economics 70 (2016) 364–372

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

journal homepage: www.elsevier.com/locate/ime

Hedging insurance books

Peter Carr [a], Dilip B. Madan [b,*], Michael Melamed [c], Wim Schoutens [d]

[a] Finance and Risk Engineering, New York University Tandon School of Engineering, 12 Metrotech Center, Brooklyn, NY 11201, United States
[b] Robert H. Smith School of Business, University of Maryland, United States
[c] JMSOLVE LLC, United States
[d] Department of Mathematics, K. U. Leuven, Belgium

ARTICLE INFO

Article history:
Received April 2015
Received in revised form
March 2016
Accepted 9 May 2016
Available online 12 July 2016

Keywords:
Acceptable risks
Bid and ask prices
Concave distortions
Variance gamma model
Risk premiums
Arrival rates

ABSTRACT

Complex insurance risks typically have multiple exposures. If available, options on multiple underliers with a short maturity can be employed to hedge this exposure. More precisely, the present value of aggregate payouts is hedged using least squares, ask price minimization, and ask price minimization constrained to long only option positions. The proposed hedges are illustrated for hypothetical Variable Annuity contracts invested in the nine sector ETF's of the US economy. We simulate the insurance accounts by simulating risk-neutrally the underliers by writing them as transformed correlated normals; the physical and risk-neutral evolution is taken in the variance gamma class as a simple example of a non-Gaussian limit law. The hedges arising from ask price minimization constrained to long only option positions delivers a least cost and most stable result.

1. Introduction

Insurance books of Variable Annuity contracts combine exposures to financial and non-financial risks. The financial component can be quite complex, involving path-dependent aspects of account values that are themselves portfolio values of investments in numerous financial securities. One anticipates that movements in security prices over the near term can have considerable effects on the present value of payouts associated with the insurance liabilities. As a result, taking positions in the options markets of related securities may constitute a good hedge for the insurance book.

This paper illustrates a variety of ways of designing such multi-underlier option hedges and comments on the issues involved in choosing the one to implement. The hedging strategy requires one to build an expectation of the market value of remaining liabilities, conditional on the price levels of underlying assets whose options are being entertained as hedges. Issues in designing the Monte Carlo simulation of such valuations are addressed. Given that replication is not likely in such contexts, one has to also address the assessment of post hedge residual risks.

The methods presented can be applied in a variety of contexts and will be illustrated here in the context of an example provided by a hypothetical Guaranteed Minimum Withdrawal Benefit Variable Annuity rider (GMWBVA). The objective we seek to accomplish is to enhance the technologies available for hedging complex multidimensional risks using options on a number of underlying assets. The use of the GMWBVA contract is just an example of a potentially complex risk. Any entity capable of simulating the present value of its liabilities conditional on the multi-dimensional levels of suitable underlying assets may approach the hedging question using the methods of this paper. In this regard, we note that the general purpose of the hedge is not that of risk elimination, as this is generally not likely to be possible, but is instead one of improving the overall quality of the hedged position. This requires a quantification of attitudes towards residual risk and the paper proposes conservative market-based procedures along these lines. These are important questions that need to be addressed when we generalize away from the static perfect hedging context of Carr and Wu (2014) towards the wider context of imperfect hedging.

Any implementation of the methodology will require the joint simulation of time paths at a regular frequency, of relevant underlier prices, and will also require the construction of an associated present value of the related insurance payouts or other liabilities. As already noted, we illustrate with a sample of GMWBVA's written on account values invested in the nine sector ETF's for the US economy. We simulate paths for the underliers, the related account values, and we also build the present value of

* Corresponding author.
E-mail address: dbm@rhsmith.umd.edu (D.B. Madan).

http://dx.doi.org/10.1016/j.insmatheco.2016.05.002
0167-6687/

all associated insurance payouts. We use a time horizon of about 40 years, which is when the risks arising from the underliers in the initial book are assumed to have been dissipated.

One may then go forward in the path space, by say six months, to form a matrix of the level of the underlying asset price levels at this date on different paths, along with the associated aggregate present value of all payouts on the same paths. We then introduce as hedging assets, positions in the underlying stocks and out-of-the-money options on these stocks. The residual cash flow is given by earnings on the hedge positions, less the target cash flow represented by the present value of payouts. The hedge positions are sought to enhance the value of the residual cash flow. This is an optimization problem whose controls are positions in all options on all underliers as well as positions in the underlying stocks. The hedge is designed on numerically solving this optimization problem. The proposed process is applied over a lifetime of the GMWBVA contract and allows solving for optimal positions in the underlying stocks and options on these stocks at each time hedge rebalancing needs to occur.

The issues to be addressed are (i) the formulation of the financial path space on which the optimization is to be conducted, (ii) the modeling of the non-financial risks and their interaction with the financial path space generated, (iii) the selection of hedge instruments, (iv) the choice of an optimization criterion to be applied to the residual cash flows, and (v) the design of any constraints to be imposed on positions. After addressing these matters in separate sections, we present the details of our case study on GMWBVA accounts written on the sector ETF's. We close with the results of our case study.

The outline of the rest of the paper is as follows. Section 2 models the interaction of the non-financial risks with a financial path-space. Section 3 takes up the issue of financial path-space design. Section 4 presents the hedge instruments and addresses the choice of optimization criteria and constraints. Section 5 takes up the details for the GMWBVA case study. The results of this case study are presented in Section 6. Section 7 concludes.

2. Modeling non-financial risk interactions with the financial path space

The interactions of the non-financial risks with a financial path-space are particular to the structured insurance product. Here, we focus on the interactions relevant to our case study of *GMWBVA* contracts written on account values invested in the financial assets. A single account value process $A(t)$ and its associated base $B(t)$ evolve until the random time τ. This random time is the minimum of either the time of death of the account holder or the time at which the account value first hits zero. Positive account values at death are returned to the account holder's estate and all premium payments and withdrawals stop. If the account value hits zero before death, then the account holder receives a percentage of the base to death and all premiums stop. Before time τ, the risk-neutral account evolution process is given by

$$dA(t) = (r - \kappa)A(t)dt - (z + y)B(t)dt + A(t)\theta' dN(t), \quad (1)$$

where r is the interest rate, κ is a management fee based on the account value, y is a fee relative to the level of the base, and z is the withdrawal rate. The final term is the return to investment in the equity exposure, where θ is the vector of account proportions invested in the equity assets and $dN(t)$ is the martingale component of equity returns. For a vector of equity prices $S(t)$, the risk-neutral equity evolution is

$$\frac{dS(t)}{S(t)} = rdt + dN(t),$$

where the division is component-wise. Furthermore, there are annual resets on the base defined as

$$B(n) = \max(B(n-1)(1+\eta_{n-1}), A(n))$$

where η_n is a positive constant, η, for n less than or equal to $\tilde{T}$ and zero thereafter.

Let the complementary distribution function for the time of death of the contract holder be

$$G(t) = P(\text{contract holder lives longer than age } t).$$

If the account value hits zero and death has not occurred, then for an account holder who had an age α when entering the contract at time zero (conditioning of course on being alive), the present value at time t of the withdrawal payouts (with withdrawal rate z) on the base until death is

$$V(t) = zB(t) \int_0^\infty \exp(-rh) G(t + \alpha + h) \mathrm{d}h / G(\alpha). \quad (2)$$

3. Financial path-space simulation

The selection of a path-space to simulate the dynamics of the underlying equity (and hence the above martingale component in Eq. (1)) is both a critical and a difficult decision, especially when it comes to multiple underliers over long periods of time, that are the typical context for insurance products.

3.1. The physical process versus the risk-neutral process

There is a strong temptation to simulate the physical process as far as one may learn it from time series data, for this is the reality that was experienced and could be the future that one has to actually live through. Yet, we also learn from the market valuation of financial products that rare events with low probability have a much higher price relative to the probability of their occurrence (Bollerslev and Todorov, 2011). When simulating a finite sample space, the events driving prices may not even arise in a physical simulation. Furthermore, it is of little comfort that theoretically the two probabilities are equivalent, for on the finite sample space on which decisions are to be based, the relevant events are lost. They do not occur and measure changes cannot compensate for non-existence.

From a more heuristic perspective, one recognizes that time series data teaches us partially about movements over small periods of time like days, while our concerns are over longer periods of time like those embedded in option maturities. Option markets reflect market concerns about risk assessments over such longer periods. Using the physical measure may bias hedges towards linear instruments like the delta hedge, thereby underplaying the exposure to non-linearities built into option positioning. The resulting mishedge may be too costly to correct subsequently.

Such deficiencies in physical simulation have led researchers to endogenize measure changes into the simulation by directly simulating the risk-neutral process (see for example Korn et al., 2010). By raising the probabilities of the rarer events of greater import to valuation and simultaneously lowering the probability of smaller moves, more of the relevant events enter the decision sample. However, the risk-neutral process may be estimated well for a single underlier from data on option prices for this underlier. The risk-neutral process for the joint law on multiple underliers is not so readily available. For example, for even two underliers, one cannot synthesize the joint density from traded option prices. To extract such information one would require information on the prices of the product of two calls at different strike pairs and the same maturity. Such securities do not trade and the required price information is not available. In the absence of joint information, one may proceed further by making additional modeling assumptions.

In this direction, Madan (2011) supposes a set of multiple underliers to be a linear mixture of independent factors. Madan (2011) then employs time series data and applies independent components analysis (Fast ICA Hyvärinen and Oja, 2000) to estimate the mixing matrix. The mixing matrix is further assumed to be risk-neutrally relevant. Risk-neutral laws for the independent components are sought with a view to repricing all options on all underliers across their strikes and maturities at market close on a single day.

3.2. The market price of risk

The discomfort with the physical law is matched by an equal discomfort with the risk-neutral law. For the question arises as to which risk-neutral law is to be employed for hedging very long maturity liabilities using relatively short maturity claims at option maturities. From daily calibrations to option surfaces, one would extract different risk-neutral laws each day with considerable variations reflecting movements in both the actual probabilities of events and their associated market risk premia. If one just observes the risk-neutral law, it is not clear whether market prices of risk involved are exaggerated or not. Hence, we seek to pay attention to both the physical and the risk-neutral process making some assessment that the latter is not exaggerated and suggestive of unduly expensive hedges that are not realistically called for.

For an assessment about the levels of implicit risk premia, one has to be explicit about the risks that are being priced and the market prices that are being implicitly assumed. In the implementation that we enact, we take an explicit decomposition of risks and their prices. Hence, we estimate a physical joint law for the underliers and then explicitly risk-neutralize using average levels of observed market prices of risk. In this regard, we follow Madan (2016a,b) in the description of risks and their prices as seen in option markets.

It is interesting to contrast our approach with the situation arising from assuming that price processes in continuous time have continuous sample paths. When risk is modeled by these continuous processes, the market price of a particular risk at a particular point of time is a number given by the covariation between the risk and the change of measure density process between the risk-neutral and physical measures (see for example Skiadas, 2009). Covariations are proportional to time, as are excess returns, and the risks embedded in the former are compensated for in the latter. Both the risk and its compensation occur over time. Return distributions over small intervals are Gaussian and locally, risk is adequately characterized by variance. Such a paradigm has served us well as is effectively documented by the large literature based on such assumptions.

However, positions in risky assets also expose the holder to the risk of sudden and substantial moves in prices. The modeling of risk by purely continuous processes unrealistically assumes away such exposures. In fact, under continuous processes there are no instantaneous risk exposures and all risks occur over time by covariation. A good feature of the Gaussian model has always been the observation that it is a limit law: the sum of a large number of independent shocks tends to a Gaussian distribution. However, the Gaussian distribution is not the only limit law. Lévy (1937) and Khintchine (1938) characterized *all* of the limit laws. These turn out to be the self-decomposable laws that are a subclass of the infinitely divisible laws associated with Lévy processes.

3.3. Jump models

A significant part of the literature has accommodated large sudden moves by adding to the continuous component a separate jump process with a finite arrival rate of jumps and often an exponential or normally distributed jump size distribution. The resulting process is called a jump diffusion. A survey of these popular stochastic processes is given in Kou (2008). The jump diffusion process is, however, not a limit law, as it is not a self-decomposable law. In fact it is the sum of two orthogonal processes that have nothing to do with each other. We prefer to continue to use limit laws that have sufficient activity at small jump sizes to even permit dispensing with the continuous component altogether. Hence, we work with pure jump processes whose laws at all time intervals are self decomposable, i.e. limit laws.

For such processes, there is a separate market price for each jump size. The market price of risk is no longer a single number, but rather is a whole function of the jump size. A particularly tractable, empirically flexible, self-decomposable law is the variance gamma (VG) law developed in Madan and Seneta (1990) and extended in Madan et al. (1998). For the physical and risk-neutral VG laws, we observe that the market price of risk function can be simply described by a parametric function with three parameters. The parameters price the overall jump size, the arrival rate of down moves of a fixed magnitude, and the arrival rate of a comparable up move. The market price of risk is given by the ratio of risk-neutral to physical arrival rates of jumps. The parameters may be estimated from a combination of time series and option data. Our risk-neutralization will make use of typical market price of risk functions for our nine underliers.

Pure jump processes of independent and identically distributed increments are Lévy processes defined by the specification of the arrival rate $k(x)$ of jumps for all jump sizes, $x \in \mathbb{R} - \{0\}$. This arrival rate function $k(x)$ specifies a Poisson arrival rate for each jump size and is called the Lévy density. It must satisfy certain technical conditions specified in Sato (1999) and Schoutens (2003) for example. Suppose one demands that a Lévy process has self-decomposable laws for the marginal distribution of the evolution over any time interval. In this case, Sato (1999) shows that the function $h(x) = |x|k(x)$ must be decreasing for positive x and increasing for negative x. In the variance gamma case, the function $h(x) = C\exp\left((G-M)x/2 - (G+M)|x|/2\right)$, for $C, G, M > 0$. This function is clearly decreasing for positive x as it is $\exp(-Mx)$ and increasing for negative x as it is $\exp(-G|x|)$. Hence, the variance gamma process produces marginals which are self-decomposable.

The variance gamma process $X(t;\sigma,\nu,\theta)$ was originally constructed (Madan et al. 1998) as Brownian motion with drift θ and volatility σ, time-changed by a gamma process $g(t;\nu)$ of unit mean rate and variance rate ν. As a result, we have that

$$X(t;\sigma,\nu,\theta) = \theta g(t;\nu) + \sigma W(g(t;\nu)),$$

where $(W(t), t > 0)$ is a standard Brownian motion. The transformations between σ, ν, θ and C, G, M may be found in Schoutens (2003).

The monthly distributions for log price relatives of our nine underliers are in the variance gamma class with parameters that are risk-neutralized. However, each non-Gaussian distribution can be individually transformed into a standard normal distribution by using the VG distribution function composed with the inverse of the standard normal distribution function. The resulting normally distributed random variables can then be correlated using the empirical correlation matrix for the nine underliers. Hence, the single step simulation is that of correlated normals transformed to risk-neutralized variance gamma marginal distributions. Non-Gaussian simulations of correlated random variables along these lines are also presented in Bouchaud and Potters (2003). Further details of our specific simulation are provided in Section 5.

4. Hedging insurance contracts

4.1. Target cash-flows

Let $S_{k,t}^{(j)}$ be a simulation of the path-space of equity underlier k, at time t, and on path j, where $k = 1, \ldots, K$, $t = 1, \ldots, T$, and $j = 1, \ldots, J$. From such a simulation, one may simulate the account values of numerous account holders $A_i(t)$ and the associated payouts $C_{i,t}^{(j)}$ (which are negative for the premiums paid), both as a percentage of the account value and as a percentage of the base. The payouts are positive for the insurance component of withdrawals after the account value has gone to zero and prior to the account holder's death. We record this insurance benefit as a lump sum paid out at the time of the account value going to zero before death, as computed by Eq. (2). We then form the aggregate present value of all payouts to each account holder on each path by

$$\widetilde{C}_i^{(j)} = \sum_{t=1}^{T} e^{-rt} C_{i,t}^{(j)},$$

and the aggregate payout to all account holders as

$$\mathcal{C}^{(j)} = \sum_i \widetilde{C}_i^{(j)}.$$

The target cash flow is then the vector $\mathcal{C} = (\mathcal{C}^{(j)}, \ j = 1, \ldots, J)$.

4.2. Hedging assets and least squares hedging

The hedging assets consist of financed out-of-the-money options on all underliers for a fixed maturity H and investment in the underliers themselves, held to the same maturity H. The cash flows to the options and the investment in the underliers with payoffs occurring at time H are given by the $D \times J$ matrix $\mathcal{H}$, where $\mathcal{H}_d^{(j)}$ is the financed payoff to hedging asset d on path j in row d and column j of $\mathcal{H}$ for $d = 1, \ldots, D$ and $j = 1, \ldots, J$. We ensure that the hedging assets are zero cost on the sample space.

Let $x = (x_d, d = 1, \ldots, D)$ denote the positions x_d in hedging asset d. The first hedge reported on is the least squares hedge, where we just regress the target cash flow $\mathcal{C}'$ onto the matrix $\mathcal{H}'$, including a constant term representing a bond position. We thus write

$$\mathcal{C}' = a + \mathcal{H}' x_{LS} + u$$

with bond position a, least squares hedge x_{LS} and error term u. One may remove constants from the sample space to define

$$\widetilde{\mathcal{C}} = \mathcal{C} - \frac{1}{J} \sum_{j=1}^{J} \widetilde{\mathcal{C}}^{(j)}$$

$$\widetilde{\mathcal{H}} = \mathcal{H} - \frac{1}{J} \sum_{j=1}^{J} \widetilde{\mathcal{H}}^{(j)}$$

and to define the centered residual cash flow

$$\widetilde{\mathcal{R}}(x) = x\widetilde{\mathcal{H}} - \widetilde{\mathcal{C}}.$$

We may think of $\widetilde{\mathcal{C}}$ as a random variable and of

$$\widetilde{\mathcal{W}}(x) = \widetilde{\mathcal{C}} - x\widetilde{\mathcal{H}}$$

as a hedged position.

4.3. The market valuation of residual risks in partially hedged positions

We recognize that for many multi-dimensional risks, options trading on the set of underliers will not be able to replicate the present value of liabilities that constitute a fairly general multi-dimensional function, no matter what the underlying price processes. Hence, one has to come to terms with evaluating the post hedge risk exposure. As constants do not constitute risk, the residual risk evaluation deals directly with centered random variables thereby removing any incentives for positions in hedge instruments to access their favorable means when uncentered. Such positioning is not hedging, but rather amounts to speculating.

For centered random variables, one wishes to evaluate the improvement being made to the post hedge risk by the hedge. A simple criterion would just be to evaluate the variance for example. But in dealing with non-linear exposures engineered by a variety of optionalities, there would be a reluctance to significantly reduce a large positive skewness for a minimal lowering of the variance. One may then consider a variety of ways of forming tradeoffs between the set of higher moments. This is a difficult set of tradeoffs to get comfortable with. A solution that comes to mind is to maximize some expected utility function parametrizing a level of risk aversion. For a centered random variable, one has to be tolerant of losses. This suggests the use of an exponential utility function. However, the risk aversion parameter in such a utility function is scale dependent and difficult to choose. Furthermore, from the perspectives introduced below, all expectations are suspect as they treat all probabilities equally even when we know that we have little experience with the probabilities of tail events on both sides of the gain/loss spectrum. A conservative approach should reduce the probabilities of large gains, while magnifying the probabilities of large losses. This is accomplished below by taking expectations with respect to a non-additive probability or constructing a suitable (Choquet, 1953) expectation.

4.4. Two-price economies and acceptable risks

The interest here is to take a more market-oriented view of risk assessment, as opposed to some personal preference approach. This is because risks have to be ultimately off loaded in markets. However, the difficulty with using market valuation operators is that in classical financial economics with arbitrage-free markets satisfying the law of one price, the market valuation operators are linear and hence their maximization will never lead to interior solutions. One has to start by imposing a structure of constraints that essentially do the work of delivering a solution. Hence, we turn our attention to recent developments for arbitrage-free two price economies that distinguish buying prices from selling prices. In such economies, Cherny and Madan (2010) and Madan (2015) show that the bid and ask price operators are non-linear. The bid operator is concave and hence suitable for maximization, while the ask operator is convex, and hence suited to minimization. Furthermore, Madan et al. (forthcoming) address the general hedging problem and show that quite generally, the hedging problem has an interior solution. These developments lead us to apply such market-based criteria for the assessment of residual risk. Madan (2016b) describes how these operators may be calibrated to market data and how they build in tradeoffs with respect to the higher moments of the risk distribution.

A fundamental concept of equilibria in two price economies (Madan, 2012) is the definition of acceptable risks. A delineation of the set of acceptable risks is required whenever one recognizes realistically that one cannot eliminate all risks via hedging. Arbitrage opportunities are obviously acceptable, but one needs a more inclusive mathematical definition of acceptability. Graphically, arbitrages constitute the positive orthant i.e.non-negative random payoffs obtained at zero cost. One can generalize these arbitrages to a convex cone, $\mathcal{A}$, containing the positive orthant. This approach was originally proposed in Artzner et al. (1999). It is shown in

Artzner et al. (1999) that there exists a set $\mathcal{M}$ of probability measures such that a risk represented by a random variable X is acceptable, i.e. $X \in \mathcal{A}$, just if $E^Q[X] \geq 0$ for all $Q \in \mathcal{M}$. The bid and ask prices, b, a respectively, for X then require that

$$X - b(1+r) \in \mathcal{A}, \qquad a(1+r) - X \in \mathcal{A},$$

(see for example Carr et al., 2011). Equivalently, for all $Q \in \mathcal{M}$

$$E^Q[X] - b(1+r) \geq 0$$

$$a(1+r) - E^Q[X] \geq 0.$$

The best bid and ask prices for X provided by the market, $\widetilde{b}(X)$, $\widetilde{a}(X)$ respectively, are then given by

$$\widetilde{b}(X) = \frac{1}{1+r} \inf_{Q \in \mathcal{M}} E^Q[X] \tag{3}$$

$$\widetilde{a}(X) = \frac{1}{1+r} \sup_{Q \in \mathcal{M}} E^Q[X]. \tag{4}$$

By virtue of the infimum, the bid price operator is concave, while the use of the supremum makes the ask price operator convex.

Two further hypotheses introduced in Cherny and Madan (2009) lead to simple and easy to implement bid and ask price operators. Specifically, when risk acceptability is modeled solely in terms of the distribution function and we demand bid additivity for comonotone risks, the computation of the bid price may be reduced to an expectation under a concave distortion function. More precisely, there then exists a concave distribution function $\Psi(u)$, $0 \leq u \leq 1$ from the unit interval to itself, such that

$$\widetilde{b}(X) = \frac{1}{1+r} \int_{-\infty}^{\infty} x d\Psi\left(F_X(x)\right) \tag{5}$$

where $F_X(x)$ is the distribution of X. Integration by parts in the expression (5) leads to a Choquet (1953) expectation

$$\widetilde{b}(X) = \frac{1}{1+r}\left(-\int_{-\infty}^{0} \Psi\left(F_X(x)\right) dx + \int_{0}^{\infty} \left(1 - \Psi\left(F_X(x)\right)\right) dx\right) \tag{6}$$

where by the concavity of Ψ, one may observe the reduction of probabilities in the upper gain tail and the inflation of probabilities in the loss tail. Such an expectation is also called a non-linear expectation or expectation with respect to a non-additive probability. It is related to the dual theory of choice of Yaari (1987). Given the lack of experience with extreme tail events on both sides of the gain/loss spectrum, non-linear expectation can be interpreted as an imminently reasonable way to build in the required conservatism embedded in two price economies.

The distortion of a risk-neutral distribution function accomplishes the task of ensuring that the bid and ask prices being generated are conservative relative to market values. The particular distortion function that we work with is *minmaxvar*, introduced in Cherny and Madan (2009) and defined by

$$\Psi^{\gamma}(u) = 1 - \left(1 - u^{\frac{1}{1+\gamma}}\right)^{1+\gamma}.$$

It is also clear from the relationship of the sup to the inf that the ask price is the negative of the bid price for the negative risk. On a finite sample space, one may order outcomes in increasing order with $X_{(l)}$ being the lth largest outcome for $l = 1, \ldots, L$, and employ the empirical distribution function to evaluate the bid price as

$$\widetilde{b}(X) = \sum_l X_{(l)} \left(\Psi^{\gamma}\left(\frac{l}{L}\right) - \Psi^{\gamma}\left(\frac{l-1}{L}\right)\right).$$

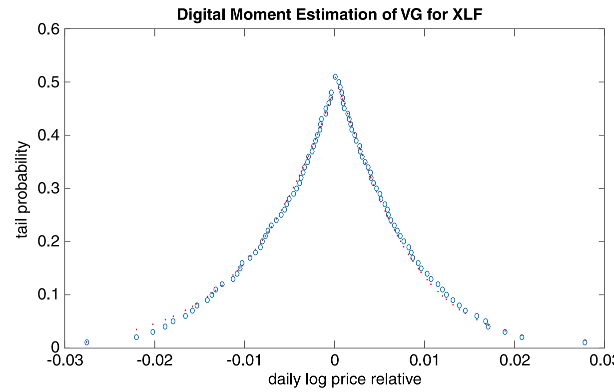

Fig. 1. The observed tail probabilities are shown as circles with dots representing the corresponding estimated model probabilities.

4.5. Conic hedging of insurance risks

For our insurance liability, we may seek a hedge position x with a view to minimizing the ask price of $\widetilde{\mathcal{W}}$. In our case, this is given by

$$z(x) = -\sum_{j=1}^{J} \widetilde{R}_{(j)}(x) \left(\Psi^{\gamma}\left(\frac{j}{J}\right) - \Psi^{\gamma}\left(\frac{j-1}{J}\right)\right).$$

Our second hedge position is chosen to minimize this ask price for the substantial stress level of 0.75. We refer to Madan (2016b) for further remarks on the calibration and choice of stress levels. Our third hedge is chosen to minimize the ask price for the hedged liability under the further constraint that the option positions are kept non-negative, so as to refrain from becoming an option seller. This constraint is imposed with the view of lowering ask prices computed on a finite sample space.

5. Details for the GMWBVA case study

The details are presented in subsections devoted to the estimation of the physical law for the underliers, their risk-neutralization, the risk-neutral simulation, the construction of the target cash flow, and the space of the hedging assets.

5.1. Physical estimation of variance gamma law for the underliers

The estimation was conducted on 501 observations for two years of data on the nine underliers *xlb*, *xle*, *xlf*, *xli*, *xlk*, *xlp*, *xlu*, *xlv*, and *xly*. We followed Madan (2014) in employing digital moment matching for the estimation of parameters. For each underlier, 99 strikes were chosen for the return defined as the log price relative corresponding to the percentiles ranging from 1 to 99 in steps of unity. For positive strike prices, the model probability of digital calls was matched to the probability observed for being above the selected strike. Similarly, for negative strike prices, the model probability of digital puts was matched to the observed probability of being below the corresponding strike. We thus match out-of-the-money digital put and call probabilities to observed tail probabilities on both sides. The variance gamma model was fit by least squares minimization of the percentage squared error between the observed probability of being in the tails, and the model probability for the same event.

As an example, Fig. 1 presents a graph of the observed tail probabilities and the corresponding model probabilities from the estimated *VG* model. Table 1 presents the estimated parameters in the σ, ν, θ format for the nine underliers.

The annualized parameters in *CGM* format are presented in Table 2.

Table 1
Digital moment estimation of VG parameters.

ETF	Sigma	Nu	Theta
XLB	0.010608	0.769945	−0.002265
XLE	0.010006	0.362770	−0.005162
XLF	0.010797	0.634793	−0.001855
XLI	0.009021	0.420204	−0.003797
XLK	0.007949	0.239938	−0.005311
XLP	0.006385	0.393545	−0.001448
XLU	0.007222	0.344631	−0.001208
XLV	0.006965	0.165835	−0.003199
XLY	0.008464	0.301066	−0.003845

Table 2
Estimated VG parameters in CGM format.

ETF	C	G	M
XLB	327.30	133.13	173.39
XLE	694.65	188.70	291.81
XLF	396.98	149.25	181.08
XLI	599.71	199.65	292.95
XLK	1050.27	288.75	456.85
XLP	640.33	319.34	390.37
XLU	731.22	311.21	357.53
XLV	1519.58	437.00	568.87
XLY	837.03	255.53	362.85

Table 3
Market prices of risk.

ETF	a	b	c
XLB	0.004111	17.7897	146.4575
XLE	0.001861	48.2356	234.1455
XLF	0.002849	14.6195	161.0551
XLI	0.001883	44.4990	240.7047
XLK	0.000924	80.9255	366.7470
XLP	0.001276	36.2157	354.6252
XLU	0.001006	22.2935	331.4028
XLV	0.000435	62.3980	495.7091
XLY	0.001019	54.0275	309.2121

5.2. Risk neutralization of physical laws

It is shown in Madan (2016a) that the function $Y(x)$ which relates the market price of jump risk Y to jump size x, has the form

$$Y(x) = a\exp\left(bx + c|x|\right)$$

when both the risk-neutral and physical processes are taken in the variance gamma class. Furthermore, the risk-neutral Lévy measure $\widetilde{k}(x)$ is related to its physical counterpart $k(x)$ by

$$\widetilde{k}(x) = Y(x)k(x). \tag{7}$$

Let C, G, and M denote the physical VG parameters, while their risk-neutral counterparts are denoted by $\widetilde{C}$, $\widetilde{G}$, and $\widetilde{M}$. For all nine underliers, we take the average market prices of risk to be given by the parameters a, b, and c where

$$a = \frac{\widetilde{C}}{C}; \qquad b = \frac{\widetilde{G}}{G}; \qquad c = \frac{\widetilde{M}}{M}.$$

Scaling the physical VG parameters in CGM format as estimated by least squares matching of digital moments by the average market prices of risk, we obtain annualized risk-neutral parameters $\widetilde{C}$, $\widetilde{G}$ and $\widetilde{M}$ for the nine underliers. Table 3 presents the average market prices of risk employed for the risk-neutralization.

For the purpose of simulating the VG processes, it is easier to consider each VG process as a time-changed Brownian motion. Accordingly, we transform back to the σ, ν, θ parametrization and report in Table 4 the risk-neutralized VG parameters for the nine underliers.

Table 4
Risk neutralized VG parameters.

ETF	Sigma	Nu	Theta
XLB	0.2568	0.7432	−0.1544
XLE	0.3135	0.7733	−0.3262
XLF	0.3859	0.8841	−0.1930
XLI	0.2911	0.8857	−0.1824
XLK	0.2688	1.0303	−0.2260
XLP	0.7172	1.2240	−0.2667
XLU	0.4274	1.3591	−0.1589
XLV	0.1824	1.5111	−0.1177
XLY	0.3127	1.1722	0.0330

5.3. Risk neutral simulation of underliers

The nine underliers are simulated as VG processes using these annualized parameters using a time step of one month. The one month marginal distributions are those of a VG process with these annualized parameters. However, the simulated underliers are generated as correlated normals that are transformed non-linearly so as to have the appropriate marginal VG distribution. For a standard normal variate Z, the transformation employed is $F_{VG}^{-1}(N(z))$ where $F_{VG}(x)$ is the appropriate VG distribution function. The path-space simulation uses mean-corrected exponentials of these variates to build the martingale component with the risk-neutral drift taken to be the riskfree interest rate. The simulation goes out for 40 years in steps of one month on 10,000 paths for the nine underliers. The simulated path-space is stored in a three dimensional matrix of size $480 \times 10\,000 \times 9$. The correlation matrix employed is shown in Table 5.

5.4. Construction of target cash flows

Ten account values were simulated out to 40 years with the accounts invested in the nine underliers using randomly generated long-only portfolio weights. The monthly return on the account was a portfolio weighted average of the returns on the underliers, adjusted for fees and withdrawals. The withdrawals were at the rate of five percent of the base, while the fees were at the rate of one percent of both the account value and the level of the base. The increment of the base was at the larger of five percent per year and the improvement in the account value, for the first ten years. Thereafter, it went up only due to improvement of the account value.

The account holders had ages ranging from 60 to 69 in steps of one year. The life-time distribution was taken as Weibull, with mean life times ranging from 80 to 89 in steps of one year. The standard deviation of the life-time distribution was taken at five years. The interest rate was 2% *peryear*.

For the ten accounts, 10,000 paths of length 480 were simulated for the account value, the level of the base, and the cash flows paid to the account holders. The account holders received as a lump sum the expected present value of five percent of the base from the time the account value goes to zero, or the time of death, when the former was earlier.

The target cash flow for the hedge is the present value of all payments through time on all accounts. Fig. 2 presents the probabilities for different aggregate present value payouts.

5.5. The space of hedging assets

For a maturity close to six months on December 22, 2014, we obtained data on the prices of options on the nine underliers with strikes within 50% of the spot. The number of options and the strike ranges on the nine underliers are shown in Table 6.

There are a total of 276 options which along with the nine underliers give us a hedging asset space of 285 hedging assets. The

Table 5
Correlation matrix.

	XLB	XLE	XLF	XLI	XLK	XLP	XLU	XLV	XLY
XLB	1.0000	0.8076	0.8046	0.8505	0.7638	0.5846	0.5150	0.6595	0.7858
XLE	0.8076	1.0000	0.7897	0.8006	0.7244	0.5909	0.4965	0.6383	0.7466
XLF	0.8046	0.7897	1.0000	0.8412	0.7633	0.6375	0.5071	0.7072	0.8073
XLI	0.8505	0.8006	0.8412	1.0000	0.8054	0.6541	0.5330	0.7380	0.8450
XLK	0.7638	0.7244	0.7633	0.8054	1.0000	0.6059	0.4484	0.6706	0.8061
XLP	0.5846	0.5909	0.6375	0.6541	0.6059	1.0000	0.6628	0.7286	0.7160
XLU	0.5150	0.4965	0.5071	0.5330	0.4484	0.6628	1.0000	0.5404	0.5321
XLV	0.6595	0.6383	0.7072	0.7380	0.6706	0.7286	0.5404	1.0000	0.7759
XLY	0.7858	0.7466	0.8073	0.8450	0.8061	0.7160	0.5321	0.7759	1.0000

Table 6
Number of options and strike ranges on underliers.

Variable	XLB	XLE	XLF	XLI	XLK	XLP	XLU	XLV	XLY
Number	30	79	14	26	18	20	22	34	33
Lowest strike	30	45	16	38	30	35	33	40	45
Highest strike	59	105	29	64	47	54	54	80	80

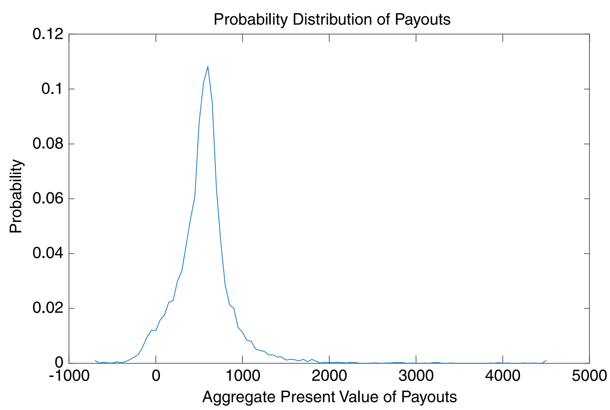

Fig. 2. Probability distribution of aggregate present value of payouts on the ten accounts.

matrix $\mathscr{H}$ of cash flows to the hedging assets six months out is $285 \times 10{,}000$.

The hedge is constructed on a censored sample space where we removed paths with the six month outcome below 50 or above 200 for an initial value reset to 100. The number of paths in the hedging exercise was 7024.

6. Hedging results for the case study

Three hedges were performed on this path-space, viz least squares, unconstrained ask price minimization, and ask price minimization with options constrained to be long-only. For each hedge, we determined the cost of the hedge. The three costs were 559.63, 6337.15 and 196.88, respectively. The bond components were 618.64 for the least squares hedge and 543.42 for the two ask price minimizing hedges. The stock components were respectively, −133.75, 5528.53, and −354.61. The three option components were 74.74, 265.19, and 8.07 respectively. It appears that the ask price minimizing hedge when unconstrained takes some large long positions in the stocks. It also takes bigger option positions. From a cost perspective, the constrained ask price minimization appears to be the most favorable alternative.

Figs. 3–5 present graphs of functions of the nine underliers accessed by the three hedges. We observe that the least squares hedge and the ask price minimization hedge are quite erratic, with many swings that may be smoothed out using a kernel estimator. This smoothing is done for these two optimization criteria in Figs. 6 and 7. The constrained ask price minimization already provides a smooth hedge and does not require a kernel smoothing exercise. Of course, the smoothed functions would have to be rehedged to determine the required option positions, but this is a simple exercise not conducted here (see Fig. 8).

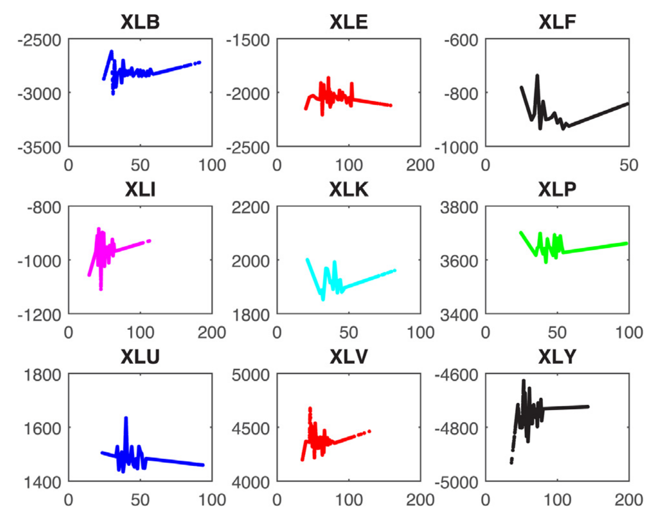

Fig. 3. Least squares hedge positions in the nine underliers.

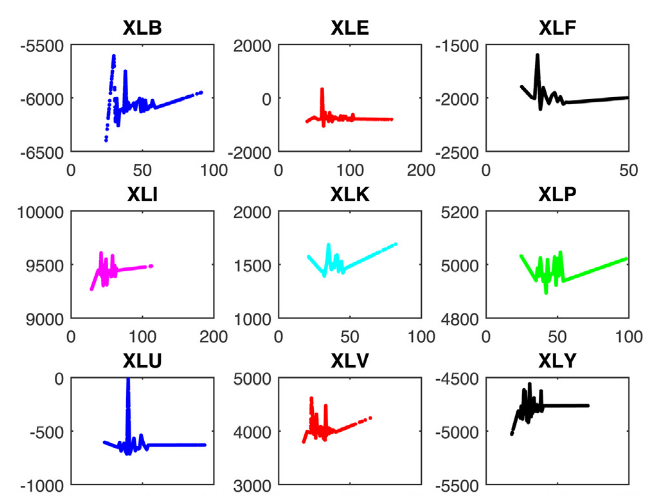

Fig. 4. Ask price minimizing hedge positions in the nine underliers.

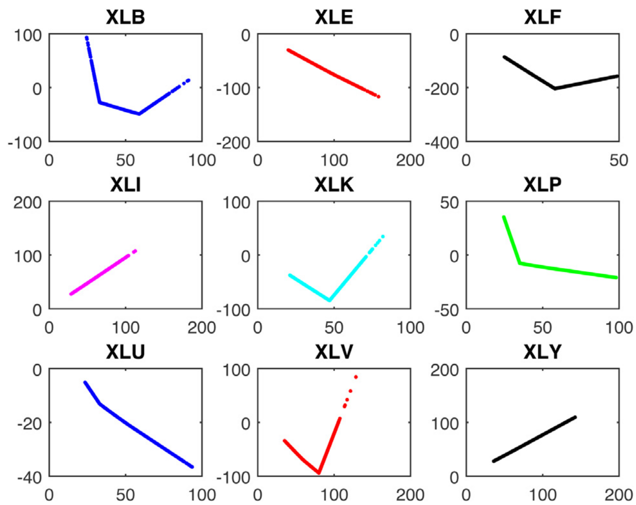

Fig. 5. Ask price minimizing hedge positions in the nine underliers constrained to long only option positions.

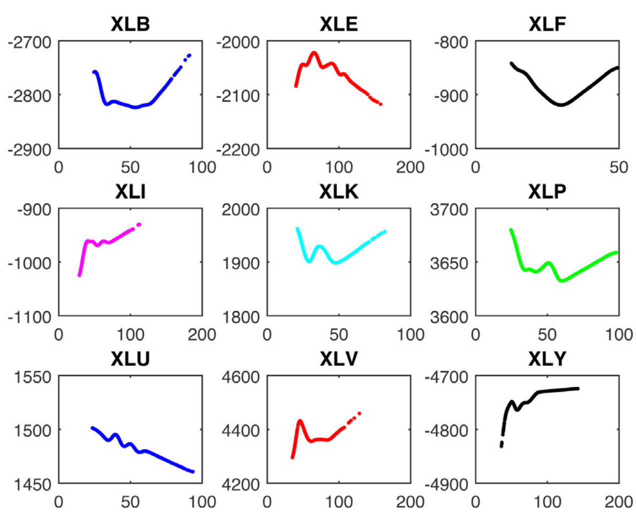

Fig. 6. Least squares positions after kernel smoothing.

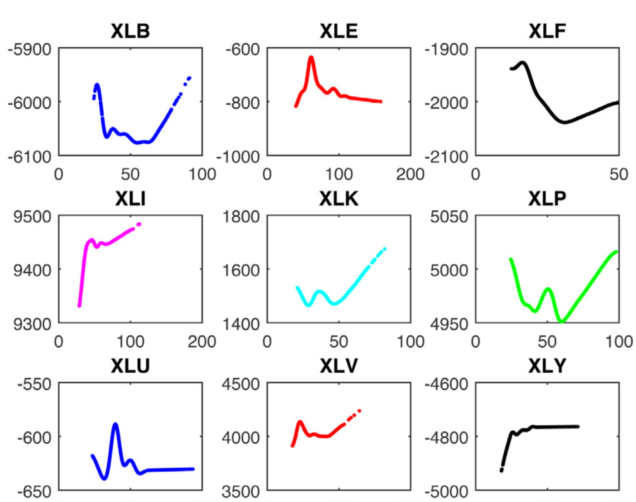

Fig. 7. Ask price minimizing positions after kernel smoothing.

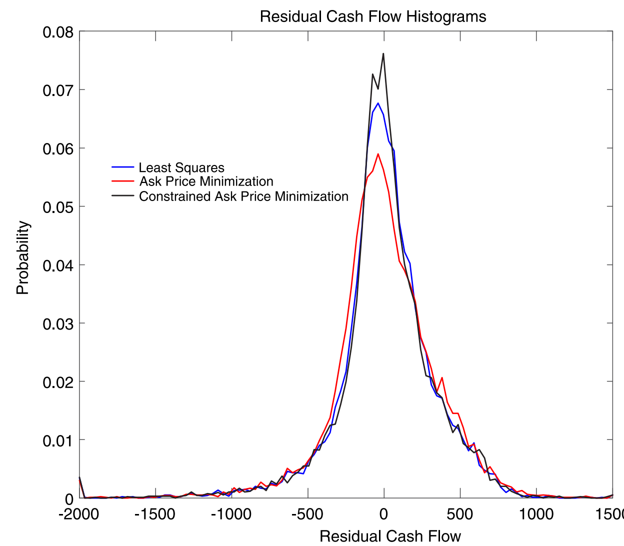

Fig. 8. Histograms of residual cash flows for the three hedges.

Table 7
Residual statistics.

Variable	Least squares	Ask price	Ask price constrained
std	371.5509	391.4379	378.8404
skew	−3.9751	−2.2213	−4.2935
kurtosis	56.0941	28.2784	63.1922
peakedness	0.8161	0.7977	0.8155
tailweightedness	0.0108	0.0108	0.0111
1%	−1061.3622	−1033.5319	−1106.1660
5%	−501.0562	−520.2784	−501.4194
10%	−314.0071	−347.6303	−316.1633
25%	−136.0459	−180.1364	−128.0253
50%	−0.5341	−14.1656	−3.0193
75%	173.8528	193.7751	169.0452
90%	385.0372	434.0794	389.3928
95%	517.9756	585.9301	528.4485
99%	724.4330	904.5527	715.0523

Finally, we present a table of statistics on the residual cash flows for the three hedges and a graph of the three histograms (see Table 7).

We observe that ask price minimization unconstrained at a higher cost does deliver a slightly higher upside with a comparable downside to the other two. A conservative strategy may well be to opt for the constrained ask price minimizing solution. The first two hedges if smoothed would have to be re-evaluated as well.

7. Summary

Insurance risks are complicated, involving exposures to numerous underlying assets. A strategy is developed to use short maturity options written on multiple underliers to hedge this exposure. The strategy is illustrated by hedging GMWBVA accounts invested in the nine sector ETF's of the US economy. The implementation requires a risk-neutral simulation of the underliers. This simulation is accomplished by expressing the underliers as transformed correlated normals with the transformation respecting risk-neutrality at the simulation horizon. The underlying physical and risk-neutral evolution is taken in the variance gamma class as a simple example of a non-Gaussian limit law. Insurance accounts for GMWBVA's are simulated as adapted to the path-space of the ETF's. The present value of aggregate payouts is then hedged using least squares, ask price minimization, and ask price minimization constrained to long only option positions. We conclude that the last of these three alternatives delivers the least costly and most stable result.

References

Artzner, P., Delbaen, F., Eber, M., Heath, D., 1999. Coherent measures of risk. Math. Finance 9, 203–228.

Bollerslev, T., Todorov, V., 2011. Tails, fears and risk premia. J. Finance 66, 2165–2211.

Bouchaud, J.P., Potters, M., 2003. Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management. Cambridge University Press, Cambridge, UK.

Carr, P., Madan, D.B., Vicente Alvarez, J.J., 2011. Markets, profits, capital, leverage and returns. J. Risk 14, 95–122.

Carr, P., Wu, L., 2014. Static hedging of standard options. J. Financ. Econ. 12, 3–46.

Cherny, A., Madan, D., 2009. New measures for performance evaluation. Rev. Financ. Stud. 22, 2571–2606.

Cherny, A., Madan, D.B., 2010. Markets as a counterparty: An introduction to conic finance. Int. J. Theor. Appl. Finance 13, 1149–1177.

Choquet, G., 1953. Theory of capacities. Ann. Inst. Fourier 5, 131–295.

Hyvärinen, A., Oja, E., 2000. Independent components analysis: Algorithms and applications. Neural Netw. 13, 411–430.

Khintchine, A. Ya., 1938. Limit Laws of Sums of Independent Random Variables. ONTI, Moscow, (in Russian).

Korn, R., Korn, E., Kroisandt, G., 2010. Monte Carlo Methods and Models in Finance and Insurance. CRC Press, Boca Raton, FL.

Kou, S.G., 2008. Jump diffusion models for asset pricing in financial engineering. In: Birge, J.R., Linetsky, V. (Eds.), Handbooks in OR & MS Volume 15. Elsevier, http://dx.doi.org/10.1016/S0927-0507(07)15002-7.

Lévy, P., 1937. Théorie de l'Addition des Variables Aléatoires. Gauthier-Villars, Paris.

Madan, D.B., 2011. Joint risk neutral laws and hedging. Trans. Inst. Ind. Eng. 43, 840–850.

Madan, D.B., 2012. A two price theory of financial equilibrium with risk management implications. Ann. Finance 8 (4), 489–505.

Madan, D.B., 2014. Estimating parametric models of probability distributions. Methodol. Comput. Appl. Probab. http://dx.doi.org/10.1007/s11009-014-9409-4.

Madan, D.B., 2015. Asset pricing theory for two price economies. Ann. Finance 11, 1–35.

Madan, D.B., 2016a. Risk premia in option markets. Ann. Finance 12, 71–94.

Madan, D.B., 2016b. Conic portfolio theory. Int. J. Theor. Appl. Probab. 19, http://dx.doi.org/10.1142/S0219024916500199.

Madan, D.B., Seneta, E., 1990. The variance gamma (VG) model for share market returns. J. Bus. 63, 511–524.

Madan, D.B., Pistorius, M., Schoutens, W., 2016. Dynamic conic hedging for competitiveness. Math. Financ. Econ. (forthcoming).

Sato, K., 1999. Lévy Processes and Infinitely Divisible Distributions. Cambridge University Press, Cambridge.

Schoutens, W., 2003. Lévy Processes in Finance. Wiley, New Jersey.

Skiadas, C., 2009. Asset Pricing Theory. Princeton University Press, Princeton, NJ.

Yaari, M., 1987. The dual theory of choice under risk. Econometrica 55, 95–115.