Dynamic Hedging of Foreign Exchange Risk using Stochastic Model Predictive Control

Farzad Noorian, Philip H. W. Leong
Computer Engineering Laboratory
School of Electrical and Information Engineering
The University of Sydney, NSW, 2006 Australia
Email: {farzad.noorian, philip.leong}@sydney.edu.au

Abstract—A risk management system for foreign exchange (FX) brokers is described. Stochastic model predictive control (SMPC) is used to reduce positions in foreign holdings over a receding horizon, while minimising a mean-variance cost function. Computation of the broker's position incorporates elements which model client flow, transaction costs, market impact, and exchange rate. Using both synthetic and historical data, the technique is shown to outperform two simple hedging strategies on a risk-cost Pareto frontier. Prediction of client and market behaviour are shown to further enhance the hedging outcome.

I. Introduction

The foreign exchange (FX) market plays an important role in international trade, most transactions being made by major banks and financial institutions. Unpredictable variations in exchange rates, combined with an accumulation of large transactions being made on behalf of clients, creates a significant risk for FX brokers. Such risk is routinely mitigated via financial derivatives, bid-ask spreads, and actions to reduce accumulated positions in foreign currencies over time.

While exposure to exchange rate volatility can be minimised by reducing positions, this conflicts with a desire to minimise transaction costs. In this paper, we formalize a methodology to manage FX risk by increasing or decreasing the broker's position subject to the current positions and market conditions.

To address the FX hedging problem, stochastic model predictive control (SMPC) is applied. We define a quadratic transaction cost model for the inter-bank trades and *risk* caused by exchange rate volatility. A stochastic formulation is used to describe exchange rate and client flow dynamics. A mean-variance optimization is then performed to minimise a combination of cost and risk over a receding horizon. The result is that, through a series of hedging actions, the broker's position can be reduced to almost zero at a given time, while minimising cost and risk. Results with both synthetic and historical data are used to demonstrate the effectiveness of this technique and it is shown that by improving the quality of prediction, the cost-risk frontier can be considerably improved.

The technique only requires controlling current holdings of the broker, and therefore does not impose additional infrastructure costs. To the best of our knowledge, this method has not been published publicly before, and implementing this method will lead to reduction of transaction costs for the broker's clients while increasing market competitivity.

The rest of this paper is as follows. Section II explores the background on hedging, FX risk management and SMPC. In section III, a formal definition of the problem is given and cost functions are defined. Our SMPC formulation is introduced in section IV, followed by numerical results is section V.

II. Related work

FX order flow risk management of market intermediaries are described in [1], explaining how they engage in *selective hedging*, either hedging their risk in a derivatives market (namely FX forwards) or holding their excessive inventories if they are compensated with a risk premium.

Ullrich [2] suggested an ad-hoc algorithm for changing hedging *intensity* based on predicting the high frequency FX rate. Evans and Lyons [3], Bates [4] and Dellacorte [5] investigated using additional information including customer order flow to predict high frequency FX rate changes. They claimed that order flow information can explain FX rate movements, although critical studies reveal that these methods are not successful in obtaining results better than a random walk benchmark [6], [7].

Model predictive control (MPC) and SMPC have recently become of interest in finance, especially in portfolio selection [8] and option pricing [9]. MPC is also known as receding horizon control (RHC). At each time-step a finite horizon optimization is performed to obtain the optimum control sequence, but only the first action of this sequence is applied. As time moves forward, the model is updated based on the observations, and the finite horizon optimization repeated. Feedback allows better control of the system in presence of external disturbance and model deficiencies, at the expense of computational power required for optimization at each time step. SMPC use stochastic optimization, which is more versatile in cases where a deterministic model of the underlying system is unfeasible. Stochastic optimisation has been successfully applied to hedging problems since it is a natural way to model the non-deterministic manner in which assets change over time [10], [11].

Primb's work on option hedging [12], [13] used SMPC with a linear quadratic cost function. Recent works by Bemporad et al. [14]–[16], have shown that a stochastic model predictive control can perform extremely well, with performance approaching that of prescient MPC models for European options

hedging.

To the best of our knowledge, MPC has not been applied to the specific area of hedging foreign exchange risk, the subject of this paper. As will be demonstrated, this problem has its unique challenges.

III. Problem formulation

In this section, we define a FX broker acting as an intermediary between different clients and the interbank market.

A. Broker position

At any t, the broker holds position $x_k(t)$ for currency k, starting with $x_k(0)$. This is affected by client initiated trades $f_k(t)$ and hedging actions $h_k(t)$ which are initiated by the broker and usually occur in the inter-bank market. For simplicity, we only allow trades at discrete timeslots $t \in [0, 1, ..., N]$ and use a discrete notation for all functions. Therefore the dynamics are formulated as

$$x_k(t+1) = x_k(t) + h_k(t) + f_k(t) \tag{1}$$

The broker must clear all positions at the end of trading session (e.g. daily or weekly), therefore $x_k(N+1) = 0, \forall k$ where N is the end-of-trading session time. This is common practice to avoid carrying an open position's risk during closing hours [17]. Assuming no client trades at the last time-step, i.e. $f_k(N) = 0$, the last hedging action becomes:

$$h_k(N) = -x_k(N) \tag{2}$$

Profit and loss of the broker is made by three major sources:
1) Transaction cost received from clients.
2) Hedging action transaction cost paid to other banks.
3) Market volatility.

As the profit from client-initiated trades is not affected by hedging, it is not used in this study.

At time t, the cumulative transaction cost $C_k(t)$ paid by broker for hedging is:

$$C_k(t) = \sum_{i=0}^{t} f_{cost}(h_k(i), i, k) \tag{3}$$

where f_{cost} is a time dependent cost function based on the bid-ask spread quoted for $h_k(t)$. This spread is known to be highly volatile, and is affected by market impact (increasing with the size of $h_k(t)$), market conditions (e.g. the spread decreases in liquid markets), and the history of trades between counterparties.

In our model, we use a quadratic cost function based on a fixed spread:

$$f_{cost}(h_k(t), t) = (\delta_k h_k(t))^2 \tag{4}$$

where δ_k is the bid-ask half spread for currency k. The quadratic has the advantage of taking the market into account, as larger trades generally cost more under real market conditions.

To avoid the market impact and risk of larger trades, we add an additional constraint on the size of hedging actions:

$$|h_k(t)| \leq h_{k,max} \tag{5}$$

Exchange rate volatility changes the value of positions, causing profit and loss:

$$L_k(t) = \sum_{i=1}^{t} x_k(i) r_k(i) \tag{6}$$

where $L_k(t)$ is the profit and loss due to changes in exchange rate and $r_k(t)$ is the exchange rate returns for currency k.

B. Control problem

A broker wishes to reduce his cost and risk at the end of trading period, subject to the dynamics and constraints defined in section III-A. This is formulated as a mean-variance optimization $E[C(N)] + \lambda \text{Var}[L(N)]$ where λ is the risk aversion factor. The following equation, considering the probabilistic nature of client flow and exchange rate, formalizes the broker's target:

$$\underset{h_k(i); \, \forall i,k}{\operatorname{argmin}} \quad E\left[\sum_{k=1}^{m} \sum_{i=0}^{N} \delta_k^2 h_k^2(i)\right]$$
$$+ \lambda \text{Var}\left[\sum_{k=1}^{m} \sum_{i=1}^{N} x_k(i) r_k(i)\right] \tag{7}$$
$$\text{subject to} \quad x_k(t+1) = x_k(t) + h_k(t) + f_k(t)$$
$$x_k(N+1) = 0$$
$$|h_k(i)| \leq h_{k,max}$$

where m is the number of currencies in the portfolio.

C. Exchange Rate Model

In this paper, the following stochastic differential equation is used to model the exchange rate:

$$dp(\tau) = j(\tau) + \mu p(\tau) d\tau + v(\tau) dW(\tau) \tag{8}$$

where $p(\tau)$ is the exchange rate for the continues time τ, $j(\tau)$ is the jump component influenced by events and announcements, μ is risk free interest, $W(\tau)$ is the Wiener process and $v(\tau)$ is the drift (volatility).

Here, we assume $\mu \approx 0$ as the optimization horizon is too short for the interest rate to be effective. The jump component is defined as:

$$j(\tau) = \begin{cases} I(\tau) & \tau \in \mathbf{A} \\ 0 & else \end{cases} \tag{9}$$

where $\mathbf{A}$ is the set of announced time of events known in advance and I is a random variable determining the intensity of the jump.

We model the returns $r(t)$ as the following discrete time equation:

$$r(t) = p(t\Delta\tau) - p((t-1)\Delta\tau) \tag{10}$$

where $\Delta\tau$ is the time-steps used in MPC optimization.

This model can be extended to a multivariate form, accounting for correlation between different exchange rates. For a given correlation matrix $\mathbf{C}$, the following equation will generate correlated random variables:

$$\mathbf{W} = \mathbf{W}_{i.i.d.}\mathbf{L} \tag{11}$$

where $\mathbf{L}$ is obtained from a Cholesky decomposition $\mathbf{C} = \mathbf{L}\mathbf{L}^*$, $\mathbf{W}_{i.i.d.}$ is a matrix of independent and identically distributed random variables, and $\mathbf{W}$ is the resulting matrix of correlated random variables.

Using (11), correlated Wiener processes can be generated which are then replaced in (8). Addition of events or drift are performed in a manner analogous to the univariate case.

IV. Dynamic Hedging of FX Risk

Based on the model defined in section III-B, the problem can be solved using stochastic control theory. Stochastic model predictive control, assumes that a state-space model is defined, which can be used to predict the system's behaviour.

Based on (7), the cost function is defined as:

$$
\begin{aligned}
J =&\mathrm{E}\Bigg[\sum_{k=1}^{m} \sum_{i=0}^{N-1} \delta_k^2 h_k(i)^2 + \\
&\sum_{k=1}^{m} \delta_k^2 \left(\sum_{i=0}^{N-1} h_k(i) + \sum_{i=0}^{N-1} f_k(i) + x_k(0) \right)^2 \Bigg] + \\
&\lambda \mathrm{Var}\left[\sum_{k=1}^{m} \sum_{i=1}^{N} r_k(i) \sum_{j=0}^{i-1} (h_k(j) + f_k(j) + x_k(0)) \right]
\end{aligned} \tag{12}
$$

To simplify notation, we consider each vector to be a concatenation of multiple currencies, e.g. for 3 currencies USD, EUR and GBP:

$$F(t) = [f_{USD}(t)\ f_{EUR}(t)\ f_{GBP}(t)] \tag{13a}$$
$$H(t) = [h_{USD}(t)\ h_{EUR}(t)\ h_{GBP}(t)] \tag{13b}$$
$$R(t) = [r_{USD}(t)\ r_{EUR}(t)\ r_{GBP}(t)] \tag{13c}$$
$$\delta = [\delta_{USD}\ \delta_{EUR}\ \delta_{GBP}], \tag{13d}$$

and define the vector representing them as:

$$\mathbf{F}(t) = [F(0)\ F(1)\ ...\ F(N-1)] \tag{14a}$$
$$\mathbf{H}(t) = [H(0)\ H(1)\ ...\ H(N-1)] \tag{14b}$$
$$\mathbf{R}(t) = [R(1)\ R(2)\ ...\ R(N)] \tag{14c}$$

Therefore (12) can be rewritten as:

$$
\begin{aligned}
J =&\mathrm{E}\Big[\mathbf{H}^T\boldsymbol{\Delta}\mathbf{H}^T + \\
&(\boldsymbol{\Upsilon}^T\mathbf{H} + \boldsymbol{\Upsilon}^T\mathbf{F} + X_0)^T \boldsymbol{\Delta} (\boldsymbol{\Upsilon}^T\mathbf{H} + \boldsymbol{\Upsilon}^T\mathbf{F} + X_0)^T\Big] + \\
&\lambda \mathrm{Var}\left[(\boldsymbol{\Upsilon}X_0 + \boldsymbol{\Psi}\mathbf{H} + \boldsymbol{\Psi}\mathbf{F})^T \mathbf{R} \right] \\
=&\mathrm{E}\Big[\mathbf{H}^T(\boldsymbol{\Delta} + \mathbf{Q})\mathbf{H} + 2\mathbf{H}^T(\mathbf{Q}\mathbf{F} + \boldsymbol{\Upsilon}\mathbf{D}X_0) + \\
&\left(\mathbf{F}^T\mathbf{Q}\mathbf{F} + 2\mathbf{F}^T\boldsymbol{\Upsilon}\mathbf{D}X_0 + X_0^T\mathbf{D}X_0 \right)\Big] + \\
&\lambda \mathrm{Var}\left[\mathbf{H}^T\boldsymbol{\Psi}^T\mathbf{R} + \mathbf{U}^T\mathbf{R}\right]
\end{aligned} \tag{15}
$$

where $X_0 = X(0)$ is the starting position and

$$\boldsymbol{\Psi} = \mathbf{I}_{m\times m} \otimes \mathbf{S}_{N\times N} = \begin{bmatrix} \mathbf{S} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{S} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{S} \end{bmatrix} \tag{16}$$

$$\boldsymbol{\Delta}_{mN\times mN} = diag(\boldsymbol{\delta}_{m\times 1} \otimes \vec{\mathbf{1}}_{N\times 1})^T \\ diag(\boldsymbol{\delta}_{m\times 1} \otimes \vec{\mathbf{1}}_{N\times 1}) \tag{17}$$

$$\boldsymbol{\Upsilon}_{mN\times m} = \mathbf{I}_{m\times m} \otimes \vec{\mathbf{1}}_{N\times 1} \tag{18}$$
$$\mathbf{D}_{m\times m} = diag(\boldsymbol{\delta})^T diag(\boldsymbol{\delta}) \tag{19}$$
$$\mathbf{Q} = \boldsymbol{\Upsilon}\mathbf{D}\boldsymbol{\Upsilon}^T \tag{20}$$
$$\mathbf{U} = \boldsymbol{\Psi}\mathbf{F} + \boldsymbol{\Upsilon}X_0 \tag{21}$$

and $\mathbf{S}$ is lower triangular matrix of 1s, $\vec{\mathbf{1}}_{N\times 1}$ is $[1\ 1\ \ldots\ 1]_{N\times 1}$, $\mathbf{I}$ is the identity matrix, m is the number of different currencies, $diag(x)$ creates a diagonal matrix from vector x and $\otimes$ is the Kronecker product.

The problem of minimising (15) is a stochastic optimisation problem which can be reduced to quadratic optimization. Here, E[ ] and Var[ ] operators are replaced by their definitions $\mathrm{E}[x] = \bar{x} = \frac{1}{n}(\sum_{i=1}^{n} x^{(i)})$ and $\mathrm{Var}[x] = \frac{1}{n}(\sum_{i=1}^{n}(x^{(i)} - \bar{x})^2)$. Also, random variables $\mathbf{F}$ and $\mathbf{R}$ are replaced with η scenarios, giving $\mathbf{F}^{(1)}, \mathbf{F}^{(2)}, \ldots, \mathbf{F}^{(\eta)}$ and $\mathbf{R}^{(1)}, \mathbf{R}^{(2)}, \ldots, \mathbf{R}^{(\eta)}$, generated by Monte-Carlo methods and/or prediction. The resulting cost function is:

$$
\begin{aligned}
J =&\mathbf{H}^T\mathbf{A}\mathbf{H} + 2\mathbf{H}^T\mathbf{B} + \mathbf{C} \\
\mathbf{A} =&\boldsymbol{\Delta} + \mathbf{Q} + \lambda\boldsymbol{\Psi}^T\left(\frac{1}{\eta}\sum_{i=1}^{\eta} \hat{\mathbf{R}}^{(i)}\hat{\mathbf{R}}^{(i)T}\right)\boldsymbol{\Psi} \\
\mathbf{B} =&\mathbf{Q}\bar{\mathbf{F}} + \boldsymbol{\Upsilon}\mathbf{D}X_0 + \lambda\boldsymbol{\Psi}^T\left(\frac{1}{\eta}\sum_{i=1}^{\eta} \hat{\mathbf{R}}^{(i)}\hat{\mathbf{Y}}^{(i)}\right) \\
\mathbf{C} =&\left(\frac{1}{\eta}\sum_{i=1}^{\eta} \mathbf{F}^{(i)T}\mathbf{Q}\mathbf{F}^{(i)}\right) + 2\bar{\mathbf{F}}^T\boldsymbol{\Upsilon}\mathbf{D}X_0 + \\
&X_0^T\mathbf{D}X_0 + \lambda\left(\frac{1}{\eta}\sum_{i=1}^{\eta} \hat{\mathbf{Y}}^{(i)T}\hat{\mathbf{Y}}^{(i)}\right)
\end{aligned} \tag{22}
$$

where

$$\bar{\mathbf{F}} = \frac{1}{\eta}\sum_{i=1}^{\eta} \mathbf{F}^{(i)} \tag{23}$$

$$\hat{\mathbf{R}}^{(i)} = \mathbf{R}^{(i)} - \frac{1}{\eta}\sum_{j=1}^{\eta} \mathbf{R}^{(j)} \tag{24}$$

$$\hat{\mathbf{Y}}^{(i)} = \mathbf{R}^{(i)T}\mathbf{U}^{(i)} - \frac{1}{\eta}\sum_{j=1}^{\eta} \mathbf{R}^{(j)T}\mathbf{U}^{(j)} \tag{25}$$

Numerical quadratic programming algorithms can efficiently minimize (22) with the constraint $-h_{k,max} \le h_k(i) \le h_{k,max}$. One can use a receding horizon scheme to control the hedging actions, with client flow and FX rate models
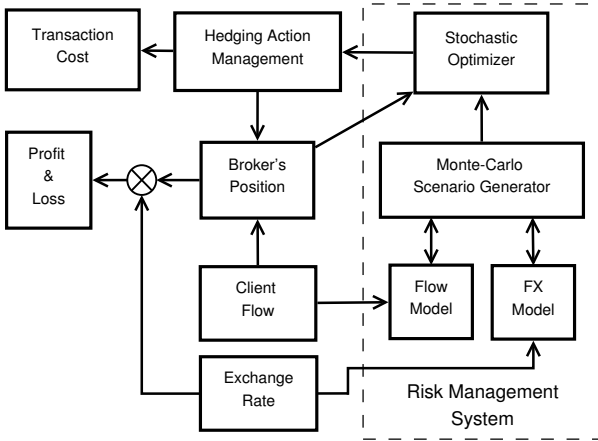
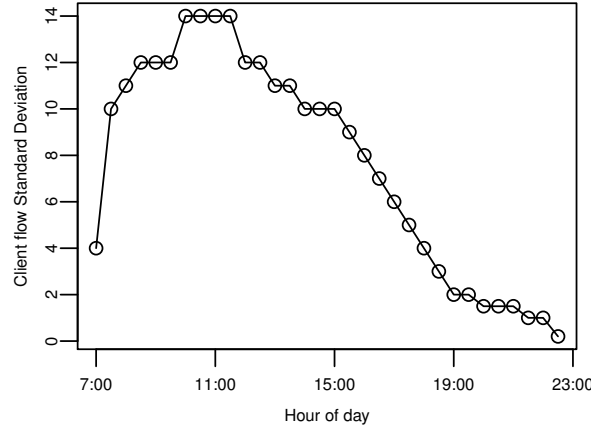Fig. 1. The proposed FX risk management system.

Fig. 2. The synthetic client flow standard deviation vs. time.

being updated at each time step based on observed market conditions. The Monte-Carlo simulator generates scenarios based on the models, which are then given to the stochastic optimizer to obtain the optimal hedging action $\mathbf{H}$. The first value of $\mathbf{H}$ is hedged, which adds to the transaction cost. At the next time-step, FX rate changes cause changes in the value of broker's open positions, generating profit or loss. The process continues until the end-of-trading time is reached.

Fig. 1 illustrates the relationship of the proposed risk management system to the rest of model. The hedging algorithm is summarised in algorithm 1.

Algorithm 1 SMPC Hedging Algorithm

$i \leftarrow 1$
while $i \leq N$ **do**
   Update client flow model from historic data
   Update FX rate model from market conditions
   Generate η scenarios for $\mathbf{F}$ and $\mathbf{R}$ for time $t \in [i, ..., N]$
   Compute the optimal $\mathbf{H}(i : n)$
   Hedge by $\mathbf{H}(i)$
   $i \leftarrow i + 1$
end while

V. RESULTS

In this section, two numerical examples are provided to compare the SMPC hedging with two benchmark strategies, both of which consider the exchange rate and client flow as being unpredictable and therefore do not use this information.

The first benchmark strategy sets a hard limit x_{max} for the broker's position:

$$h(t) = \begin{cases} x_{max} - (x(t) + f(t)) & x(t) + f(t) > x_{max} \\ x(t) + f(t) - x_{max} & x(t) + f(t) < -x_{max} \\ 0 & \text{otherwise} \end{cases} \quad (26)$$

By changing this limit, the strategy tries to move from minimum risk ($x_{max} = 0$) to minimum cost ($x_{max} = \infty$); however, as the limit passes a certain threshold, the transaction

cost caused by large end-of-trading session clearance imposes an additional cost to the broker.

The second strategy divides hedging instalments over all remaining trading hours to avoid market impact effects and additional transaction costs; this is further parametrized by λ as a risk aversion parameter:

$$h(t) = -(x(t) + f(t)) \frac{(N - t)\lambda + 1}{N - t + 1} \quad (27)$$

where $\lambda = 1$ gives the minimum risk and $\lambda = 0$ results in the minimum cost.

It must be noted that the *minimum cost* is only for variations of $\lambda \in [0, 1]$ and not the global minima for transaction cost. Given full knowledge of client flow, one can use the following equation to obtain the minimum theoretical transaction cost:

$$C_{min} = \sum_{i=1}^{N} (\delta \bar{f})^2 \quad (28)$$

where $\bar{f} = \frac{1}{N} \sum_{i=1}^{N} f(i)$. Minimum risk, in any case, is 0 when $h(t) = -f(t)$.

The experiments were implemented in the R programming language [18] and used quadprog package [19] for quadratic programming.

A. Synthetic test and results

It has been observed that the clients exhibit a certain seasonality, e.g. more trades happen mid-day rather than at the end-of-day hours. Therefore the client flow is modelled as a zero-mean heteroskedastic process, with a time-dependent variance σ^2:

$$f(t) \sim \mathcal{N}(0, \sigma^2(t)) \quad (29)$$

We assume that $f(i)$ and $f(j)$ are independent for $\forall i \neq j$. The client flow standard deviation, $\sigma(t)$, for synthetic data is shown in Fig. 2.
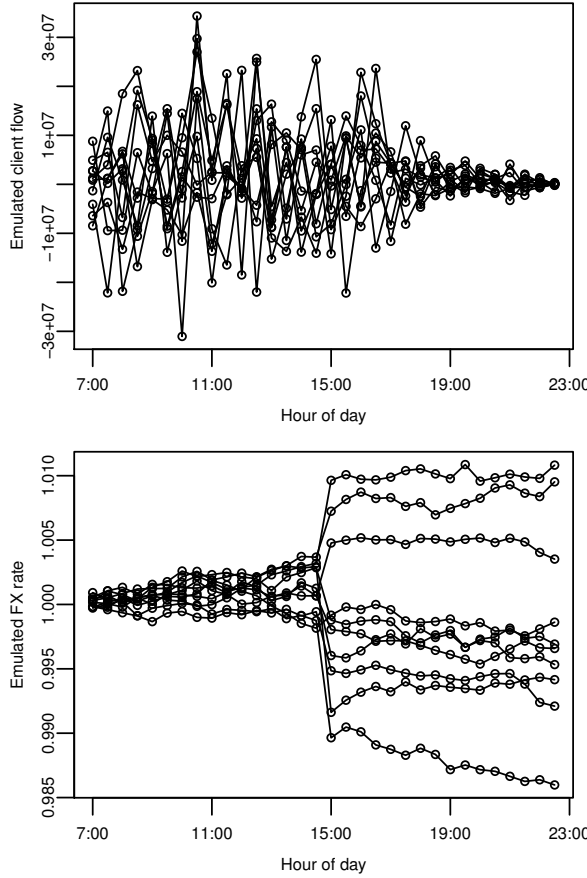
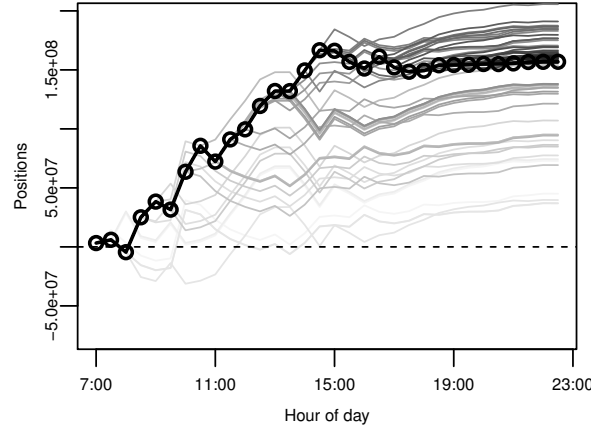Fig. 3. Example of 10 FX rate and client flow stochastic scenarios.

Fig. 4. Example of broker's position $x(t)$ changing through the trading session. Different shades of grey show stochastic paths generated at various times of trading session, with black being the observed position and brightness increasing proportional to uncertainty.
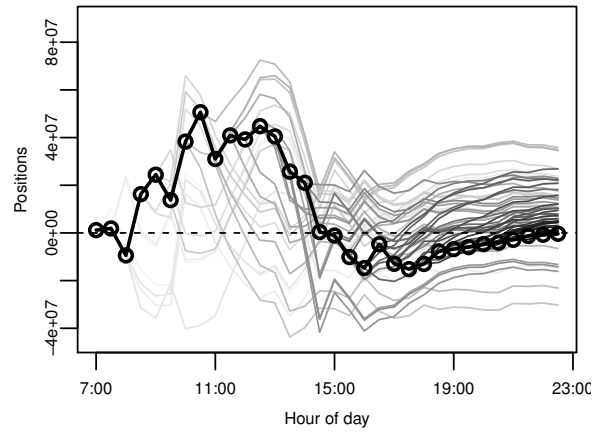

Fig. 5. Effect of hedging on broker's position. Different shades of grey show stochastic paths generated at various times of trading session.

The exchange rate is simulated using the model developed in section III-C. For each day, the number of events and their associated times are selected; this data is marked as publicly available, and is re-used in all simulations. A time-constant value of $v = 0.0005$ and $I = 0.0050$ were selected for volatility and event intensity of (8).

Fig. 3 shows 10 generated exemplar scenarios. An event at 15:00 causes a step change in the FX rate.

First, a simple test with one foreign asset for one trading session was performed. Settings demonstrated in Fig. 3 were used to generate FX rate and client flow. Fig. 4 shows the accumulation of position against different stochastic paths. Fig. 5 and Fig. 6 show hedged positions and hedging actions respectively using SMPC algorithm with $\lambda = 5 \times 10^{-4}$. A comparison between Fig. 4 and Fig. 6 shows the algorithm has reduced the position to 0 at the end of trading session. Furthermore, by using the knowledge of the announced event at 15:00, the open positions were minimized before the event to reduce the possibility of losses in case of sudden FX rate movement.

The main hedging simulation was performed with five foreign assets. The half-spreads δ_k were selected as 0.5, 1,

1, 1.5 and 2 per 10000 for the individual assets respectively. The asset correlation matrix $\mathbf{C}$ was regenerated randomly at the start of each session.

The simulation was ran for 50 trading sessions and end-of-trading costs and profits for different scenarios were measured. Average costs and risks (standard deviation of profits) are reported in Fig. 7. Each cost-risk frontier is generated by changing x_{max} in (26) for the first strategy, while for other strategies, different values for λ were used. The prescient strategy assumes perfect prediction, and uses (22) with future values of client flow and standard deviation of FX rate returns.

Since the first strategy uses hard limits, it only performs well up to a certain x_{max}. The second strategy performs better, but does not reach the minimum cost as calculated by (28). SMPC hedging, however, outperforms both approaches, and is very close to hedging with prescient flow and volatility.
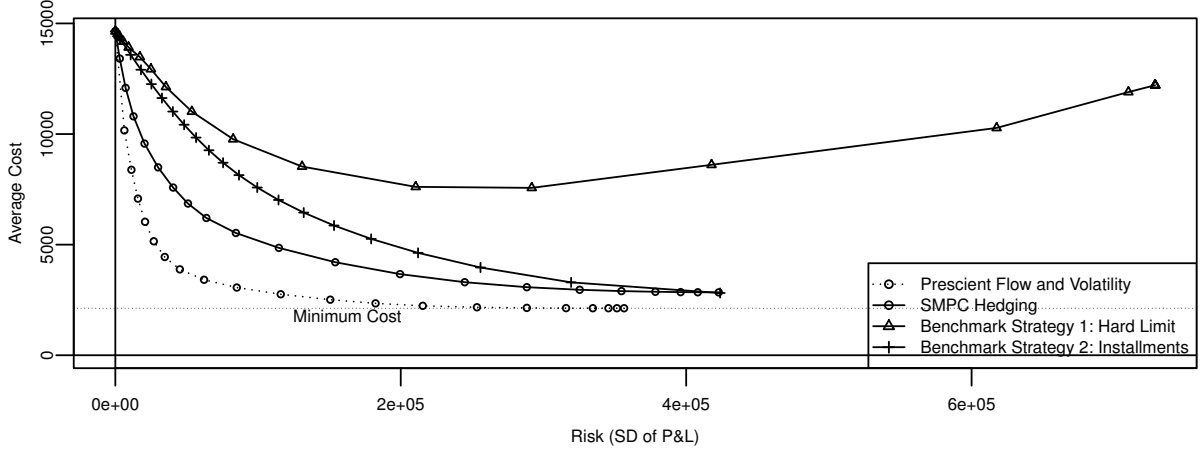
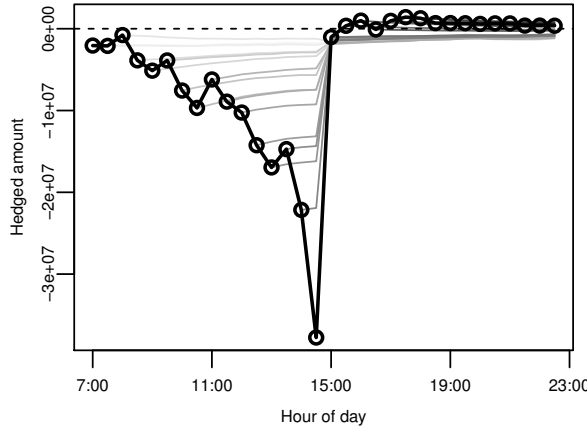Fig. 7. Risk-cost frontiers for the synthetic example.

Fig. 6. Hedging actions $h(t)$ in the single asset example. Different shades of grey show optimization results performed at various time-steps.
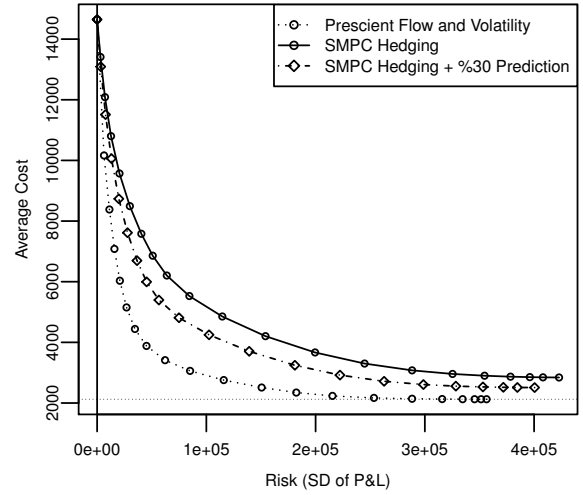

Fig. 8. Risk-cost frontiers for synthetic example, with and without prediction.

TABLE I
COST REDUCTION FOR DIFFERENT PREDICTION ACCURACIES FOR
SYNTHETIC EXAMPLE.

ρ	Prediction RMSE $(\times 10^6)$	Cost Reduction $(\times 10^6)$	Relative Cost Reduction %
0.00	8.69	6.92	0.00
0.10	7.81	6.62	9.50
0.20	6.93	6.33	18.20
0.30	6.09	5.71	37.70
0.40	5.19	5.38	48.10
0.50	4.33	5.18	54.30
0.60	3.45	4.56	73.70
0.70	2.60	4.43	77.80
0.80	1.73	4.38	79.60
0.90	0.87	4.11	87.80
1.00	0.00	3.73	100.00

The client flow model defined in (29) assumes that it is independent and identically distributed (i.i.d.) random variable. However, some of the external factors that influence client flow are measurable and can be used to forecast future flow. As prediction of client flow is beyond the scope of this paper, we study the effect of prediction by corrupting the observed flow with noise:

$$f^{(p)}(t) = \rho f(t) + (1 - \rho) N(t) \qquad (30)$$

where $f^{(p)}(t)$ is the predicted flow used in hedging, $f(t)$ is the observed flow, $N(t) \sim \mathcal{N}(0, \sigma^2(t))$ and ρ is the forecast accuracy ($\rho = 1$ corresponds to perfect hindsight).

Fig. 8 shows the effect of prediction with $\rho = 0.3$ versus prescient ($\rho = 1$) and stochastic hedging. Table I compares different values of ρ against accuracy of prediction in (30) and the reduction of cost relative to C_{min} for risk = 10^5. It is observed that increasing certainty of client flow considerably improves the risk-cost frontier.

B. Historical data

We used 6 months of retail FX client transaction flow data supplied by Westpac banking corporation. AUD was chosen

446

as the home currency and USD, EUR, NZD and JPY were selected as foreign assets.

The client flow was filtered and aggregated to create 32 half-hourly values per day. The first 2 month of data was used to obtain model parameters of (29), and the next 4 months were used for outsample testing.

The exchange rate simulation was performed by fitting historical FX rates to the model of section III-C. Individual variances v and correlation matrix C were computed from covariances of the previous day, and jump component timings were extracted from the publically available DailyFX event calendar [20]. Jump intensity was set to $I = 10v$. Only events classified as high impact were considered in this simulation and the rest were discarded. The half-spread δ for each currency was selected as 0.5, 1, 2 and 1 per 10000 for USD, EUR, NZD and JPY respectively.

Two hundred ($\eta = 200$) scenarios were generated per day for SMPC simulation. Cost and profit and loss were computed according to (3) and (6) respectively and were normalized by the daily total value of transactions $\sum |f(i)|$. The average normalized cost is reported on the vertical axis and the standard deviation of profit and loss is reported on horizontal axis as risk. Run-time of the algorithm for each time-step was less than a second on a 3.4 GHz Intel Core-i7 machine, making it suitable for deployment in real-time systems.

The average risk-cost curves for a simulation over 50 days are shown in Fig. 9. Although SMPC hedging outperforms both benchmark hedging algorithms, there is a significant gap between SMPC results and the prescient curve. This problem can be traced to client modelling, where the simple model of (29) is not able to capture certain features such as the fat-tailed distribution of flow, and dependencies between adjacent time periods.

To test this hypothesis, similar to the test in section V-A for synthetic example, (30) was used to determine the impact of better client flow models and forecasting on hedging. Fig. 10 shows the effect of prediction with $\rho = 0.3$ and $\rho = 0.6$ versus prescient ($\rho = 1$) and stochastic hedging, while table II compares different values of ρ against accuracy of prediction in (30) and the reduction of cost relative to C_{min} for normalized risk = 2.5. As with the synthetic test, increasing ρ improves hedging results.

VI. CONCLUSION

In this paper, a strategy for hedging FX trading risk was presented. The approach utilises stochastic model predictive control which has a solid theoretical background, and handles unknowns in exchange rate and client flow using stochastic models. Feedback control is used to ensure that constraints can be satisfied, and prediction utilised to enhance its performance. The cost function is constructed so that it can be efficiently minimised via quadratic programming, making the scheme suitable for real-time implementations.

The hedging strategy was verified using both synthetic and real-world data. SMPC was shown to significantly outperform
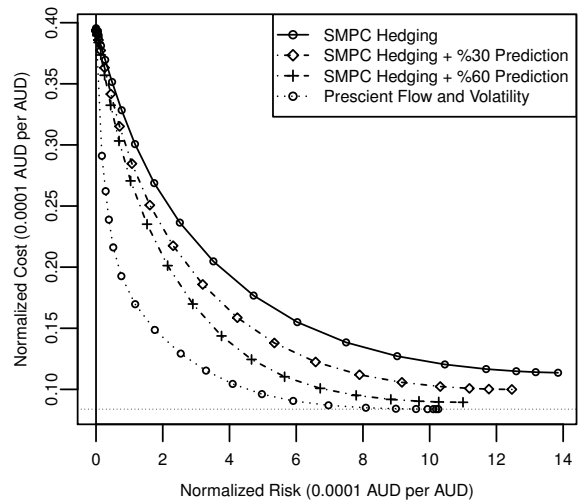

Fig. 10. Risk-cost frontiers for historical example, with and without prediction.

TABLE II
COST REDUCTION FOR DIFFERENT PREDICTION ACCURACIES FOR HISTORICAL EXAMPLE.

ρ	Normalized Prediction Error	Normalized Cost	Relative Cost Reduction %
0.0	5.29	0.237	0.0
0.1	4.92	0.228	8.0
0.2	4.36	0.222	14.2
0.3	3.81	0.210	25.0
0.4	3.27	0.202	32.4
0.5	2.71	0.192	42.2
0.6	2.16	0.185	48.3
0.7	1.63	0.178	54.9
0.8	1.09	0.173	59.9
0.9	0.54	0.167	65.1
1.0	0.00	0.130	100.0

simple trading strategies, and the performance of the SMPC with different levels of prediction ability was quantified.

In future work, we intend to pursue the following ideas to further improve hedging performance:

- The addition of a linear term to the cost function would be more realistic; however (22) will no longer be quadratic and other techniques are required for optimization.
- An alternative model for risk can be constructed using a covariance matrix of FX rate returns:

$$R(t) = \sum_{i=1}^{t} X^T(i)V(i)X(i) \tag{31}$$

where $R(t)$ is the risk until time t and $V(i)$ is the volatility covariance matrix for the duration of ith hedging time-step. Replacing (6) with (31) in (12) will remove the cost function's dependence on the unpredictable changes in FX rate *direction*, reducing the number of scenarios required to simulate future market trends.
- Our current client flow model is overly simplified since we only consider hourly seasonality modelled via a
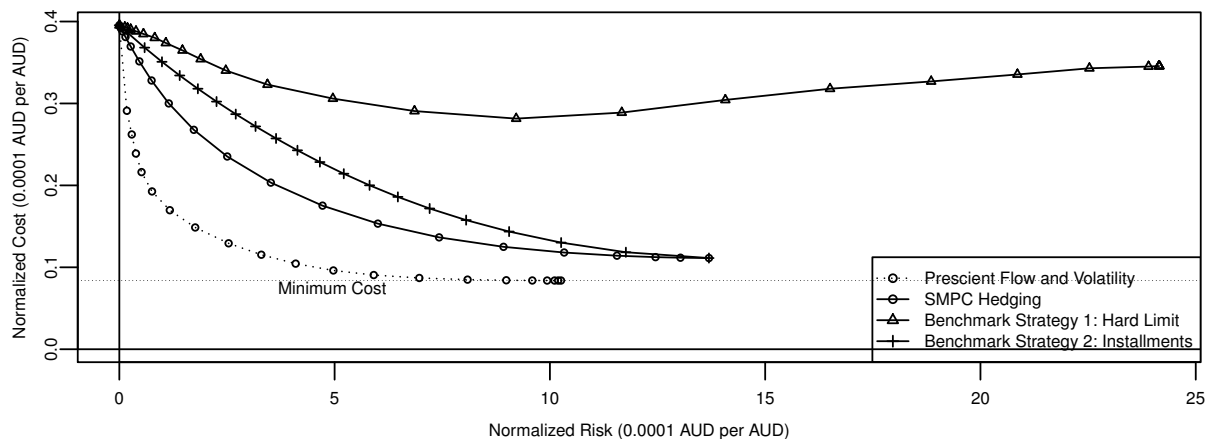
Fig. 9. Risk-cost frontiers for an example using historical data.

Gaussian distribution. Using better stochastic models or machine learning techniques to predict client flow will improve the hedging, as demonstrated in the paper. Similarly, improved prediction of volatility can further reduce risk.

REFERENCES

[1] C. DŚouza, *How Do Canadian Banks that Deal in Foreign Exchange Hedge Their Exposure to Risk?*, ser. Working papers. Bank of Canada, 2002.

[2] C. Ullrich, *Forecasting and Hedging in the Foreign Exchange Markets*, ser. Lecture Notes in Economics and Mathematical Systems. Springer, 2009, vol. 623.

[3] M. D. D. Evans and R. K. Lyons, "Forecasting exchange rate fundamentals with order flow," National Bureau of Economic Research, Working Paper, July 2009.

[4] R. Bates, M. A. H. Dempster, and Y. S. Romahi, "Evolutionary reinforcement learning in FX order book and order flow analysis," in *Computational Intelligence for Financial Engineering, 2003. Proceedings. 2003 IEEE International Conference on*. IEEE, 2003, pp. 355–362.

[5] P. D. Corte, D. Rime, L. Sarno, and I. Tsiakas, "(why) does order flow forecast exchange rates," *Preliminary*, 2011. [Online]. Available: www3.imperial.ac.uk/fileupload/download/7.dellacorte.pdf

[6] Aron Gereben and Gyrgy Gyomai and Norbert Kiss M., "Customer order flow, information and liquidity on the hungarian foreign exchange market," Magyar Nemzeti Bank (the central bank of Hungary), MNB Working Papers 2006/8, 2006. [Online]. Available: http://ideas.repec.org/p/mnb/wpaper/2006-8.html

[7] M. Sager and M. P. Taylor, "Commercially available order flow data and exchange rate movements: "caveat emptor"," *Journal of Money, Credit and Banking*, vol. 40, no. 4, pp. 583–625, 2008.

[8] F. Herzog, S. Keel, G. Dondi, L. Schumann, and H. Geering, "Model predictive control for portfolio selection," in *American Control Conference, 2006*, 2006, pp. 1252–1259.

[9] D. Voesenek, "Option Pricing and Hedging under Realistic Market Conditions An Optimal Control Approach," Swiss Federal Institute Of Technology Zurich, Master of Science Thesis, March 2013.

[10] J. Gondzio, R. Kouwenberg, and T. Vorst, "Hedging options under transaction costs and stochastic volatility," *Journal of Economic Dynamics and Control*, vol. 27, no. 6, pp. 1045 – 1068, 2003, High-Performance Computing for Financial Planning.

[11] P. J. Meindl and J. A. Primbs, "Dynamic hedging of single and multi-dimensional options with transaction costs: a generalized utility maximization approach," *Quantitative Finance*, vol. 8, no. 3, pp. 299–312, 2008.

[12] J. A. Primbs, "Dynamic hedging of basket options under proportional transaction costs using receding horizon control," *International Journal of Control*, vol. 82, no. 10, pp. 1841–1855, 2009.

[13] ——, "LQR and receding horizon approaches to multi-dimensional option hedging under transaction costs," in *American Control Conference (ACC), 2010*. IEEE, 2010, pp. 6891–6896.

[14] A. Bemporad, T. Gabbriellini, L. Puglia, and L. Bellucci, "Scenario-based stochastic model predictive control for dynamic option hedging," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 6089–6094.

[15] A. Bemporad, L. Puglia, and T. Gabbriellini, "A stochastic model predictive control approach to dynamic option hedging with transaction costs," in *American Control Conference (ACC), 2011*, 2011, pp. 3862–3867.

[16] A. Bemporad, L. Bellucci, and T. Gabbriellini, "Dynamic option hedging via stochastic model predictive control based on scenario simulation," *Quantitative Finance*, pp. 1–13, 2012.

[17] R. K. Lyons, "Profits and position control: A week of FX dealing," *Journal of International Money and Finance*, vol. 17, no. 1, pp. 97–115, 1998.

[18] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2013, ISBN 3-900051-07-0. [Online]. Available: http://www.R-project.org/

[19] A. Weingessel, *quadprog: Functions to solve Quadratic Programming Problems.*, 2013, R package version 1.5-5. [Online]. Available: http://CRAN.R-project.org/package=quadprog

[20] Forex Economic Calendar. DailyFX Market News and Analysis. [Online]. Available: http://www.dailyfx.com/calendar