

Stochastic Receding Horizon Control for Short-Term Risk Management in Foreign Exchange

Farzad Noorian^{a,*}, Barry Flower^b, Philip H.W. Leong^a

^a*Computer Engineering Laboratory, School of Electrical and Information Engineering,
The University of Sydney, NSW 2006, Australia*

^b*Westpac Institutional Bank, Westpac Place, 275 Kent Street,
Sydney, NSW 2000, Australia*

Abstract

Foreign exchange (FX) dealers are exposed to currency risk through both market and counterparty activities. Research in FX risk management has mainly focused on long-term risks, yet trading costs associated with long-term strategies make them undesirable for short-term risk hedging. In this paper, a short-term risk management system for FX dealers is described in which the optimal risk-cost profiles are obtained through dynamic control of the dealer's positions on the spot market. This approach is formulated as a stochastic receding horizon control (SRHC) problem, incorporating elements which model client flow, transaction cost, market impact, exchange rate volatility and fluctuations caused by macroeconomic announcements.

The proposed technique is backtested using both synthetic and historical client trade data. The results obtained outperform three benchmark hedging strategies on a risk-cost Pareto frontier, achieving up to 47.6% cost improvement over benchmark strategies. A flexible scenario generation oracle is also introduced and used to quantify the effects of predictive model quality on risk management.

Keywords: Foreign exchange hedging, short-term risk management, stochastic models, receding horizon control, scenario generation

*Corresponding author

Email addresses: farzad.noorian@sydney.edu.au (Farzad Noorian),
bflower@westpac.com.au (Barry Flower), philip.leong@sydney.edu.au (Philip H.W. Leong)

1. Introduction

Foreign exchange (FX) market dealers facilitate international trade by absorbing corporate, retail and institutional clients' cash flow and managing the resulting FX risk in the interbank market. Unexpected or even anticipated fluctuation in FX rates, combined with an accumulation of large transactions being made on behalf of their clients creates a significant risk for these dealers. While this risk can be mitigated by using financial derivatives or reducing positions, there are alternative strategies available that have potential to reduce contract and transaction costs.

In this paper, we formalize a methodology to manage short-term FX risk where the dealer *hedges* only by opening or closing new FX spot positions subject to market conditions and client flow. To find the optimal portfolio of positions in foreign currency, stochastic receding horizon control (SRHC) is used. First, a multi-period mean-variance cost function is formulated to minimize both the transaction costs due to trades and the risk caused by rate fluctuations. Several real-world constraints are accounted for and the resulting stochastic programming problem is approximated to a quadratic programming (QP) problem via Monte Carlo methods. Each stochastic scenario in the Monte Carlo simulation is generated by randomly sampling from custom distributions, modeling client flow, FX rate volatility and transaction cost. This data-driven approach enables a more versatile definition of data models compared to model-based analytic approaches. These result in an optimal cost-risk Pareto frontier which outperforms single-stage and rule-based hedging strategies.

We also offer statistical models for client flow, FX rate volatility and transaction cost. Specifically, time-varying models are used and special attention is paid to expected price jumps caused by macroeconomic announcements, as it has been shown that they heavily contribute to total daily price variation (Evans, 2011). Furthermore, a scenario generation *oracle* is introduced to analyze the effect of each model on the hedging strategy. By combining observations with synthetic data, the oracle generates scenarios with different predictive accuracies, allowing the dealer to identify the possible shortcomings of each component via backtesting.

The main contribution of this work is the formulation of a QP based FX risk management strategy which can accommodate stochastic FX rate and client flow models. While the idea of hedging risk, either by using derivatives or by employing a hedging portfolio (Joseph et al., 2013) has been

thoroughly explored in the literature, to the best of our knowledge the proposed method has been formally introduced only in previous work by the authors (Noorian and Leong, 2014). This paper extends that work by imposing limit constraints on positions, more realistic volatility models and time-varying bid-ask spreads, further analysis and simplification of the analytic cost function, and finally the formalization of a scenario generation oracle with exponential decay.

This paper is organized as follows. Section 2 explores the background of FX risk management and SRHC in finance. In Section 3, the problem is formally defined and the hedging objective is determined. Our optimization model is introduced in Section 4, followed by data models for scenario generation in Section 5. Finally, a series of numerical tests using synthetic and real-world data are used to validate the proposed method in Section 6.

2. Background

Receding horizon control (RHC), also known as model predictive control (MPC), is a process control technique where optimal policies are dynamically devised by solving a future looking multi-period optimization (Borrelli et al., 2014). This is achieved by embedding a predictive model of the system, which determines the system’s future performance with regards to different inputs, in the optimization’s cost function.

RHC is commonly implemented in discrete-time: at the start of each time-step, the predictive model is updated based on the latest observations and an optimization is performed to find the optimal control actions over a finite horizon. Only the optimization result for the current period is applied to the system, the rest being discarded. At the next time-step, the horizon is *shifted* one step forward (hence the name *receding horizon*) and the process is repeated. If the optimization objective has a strict deadline, the number of remaining time-steps shrinks as time advances and this technique is alternatively referred to as shrinking horizon control (Skaf et al., 2010) (Figure 1).

Updating the system model at each step, referred to as *feedback* in control theory, allows better regulation of the system in presence of external disturbances and model misspecification compared to an *open-loop* control system, at the expense of increased computational demand for the repeated optimization.




Figure 1: Flowchart of shrinking horizon control. The algorithm operates on time-steps $t \in [0, N]$.

Stochastic RHC (SRHC) introduces random variables and uncertainty into the optimization, and is more versatile in cases where a deterministic model of the underlying system is not available. Ordinarily, such a problem requires solving a multi-stage stochastic dynamic programming, and quickly becomes infeasible to solve as the horizon length grows. Recent literature have proposed Monte Carlo techniques (Glasserman, 2003) for solving this problem: by discretizing the stochastic variable evolutions to a finite number of *scenarios*, the problem is approximated by a sub-optimal random convex program (Calafiore, 2010). For linear-quadratic systems, this translates to a quadratic programming (QP) formulation, and is solved efficiently using numeric solvers. This technique has been shown to be stable (Bernardini and Bemporad, 2012), with analytic bounds on the number of scenarios required to obtain any level of robustness (Schildbach et al., 2013).

Additionally, scenarios allow a versatile range of models to be adopted compared to the limited nature of analytical formulations. Hence, scenario generation techniques have been widely studied (Römisch, 2011), and are applied to different aspects of quantitative finance including portfolio optimization (Jabbour et al., 2007) and risk minimization (Gilli et al., 2006).

There has been a growing interest in using SRHC in many financial applications, as this technique is a natural solution for multi-period stochastic systems with uncertain models. Portfolio optimization and asset allocation is one such application. Historically, techniques such as amortizing the expected costs and including the probability of profit realization over the expected holding period were used to extend the result of single-period optimization to multiple periods (diBartolomeo, 2012). By comparing dynamic asset allocation using stochastic dynamic programming against single-stage techniques, Infanger (2006) proved that dynamic techniques offer superior results. A receding horizon approach for portfolio selection, under the as-

sumption of zero transaction cost, has been formulated by Herzog et al. (2006). Topaloglou et al. (2008) included currency forwards in the portfolio to combine hedging and portfolio allocation using dynamic stochastic programming. Primbs (2007) considered both problems of wealth maximization and index tracking for a portfolio with different constraints using SRHC. Calafiore (2008) used an RHC strategy for portfolio allocation. By employing affine policies, he managed to convert the stochastic programming problem to a convex quadratic programming problem. In further publications, this method was enhanced with transaction cost (Calafiore, 2009) and asymmetric measures of risk (Calafiore and Kharaman, 2014).

Options hedging is another area where RHC approaches have been extensively used. Gondzio et al. (2003) extended the simple delta-vega hedging approach for hedging contingent claims using stochastic optimization. Hedging a basket of European call options using receding horizon techniques has been studied both with a mean-variance model approach (Primbs, 2009) and a scenario based approach (Primbs, 2010). In their other work, Meindl and Primbs (2008) incorporated a utility function as a part of their optimization function for option hedging. Recent works by Bemporad et al. (2010, 2011, 2014) have shown that hedging options using stochastic model predictive control can perform extremely well, with performance approaching that of prescient hedging models for European style options.

Empirical studies have been undertaken to quantify FX rate volatility, transaction costs and effect of news announcements in FX market. McGroarty et al. (2009) confirmed many previous studies regarding existence of intra-day patterns in FX spot market. While FX market is a true global market, liquidity and dealer participation depends on the geographical distribution of the currency being traded. As a result, volatility and bid-ask spreads exhibit M-shaped and U-shaped daily patterns respectively, as the markets go through the daily cycle of open, trade and close. Evans (2011) showed that jumps accompanying the announced events are prevalent and significant. Furthermore, approximately one third of jumps are caused by US macroeconomic news announcements, and the size of jumps are correlated with the informational surprise of the announcement. Scholtus et al. (2014) found that the bid-ask spread widens and volatility increases in the period around the announcements, which also coincides with a decrease in market depth.

Short-term FX risk management, specially from a dealer's inventory management view, has not been explored widely in the literature. Lyons (1998)

studied a week of an FX dealer’s trades, revealing informed decisions regarding hedging and speculation; nonetheless, no insights on the inventory management techniques were given. D’Souza (2002) described FX order flow risk management of market intermediaries, explaining how the Bank of Canada engages in selective hedging, by either hedging their risk in a derivatives market (namely FX forwards) or holding their excess inventories if they are compensated with a risk premium. A two-stage stochastic programming decision model has been suggested by Volosov et al. (2005) to hedge deterministic foreign cash flow, with spot and forward rates random behavior predicted by a vector error correction model. Other approaches have proposed ad-hoc rules for changing hedging intensity, mainly based on high frequency FX rate predictions (Ullrich, 2009).

We were unable to find any previous work on short-term FX inventory management or risk hedging using a stochastic multi-period optimization technique. As will be demonstrated, this problem has its own unique challenges.

3. Problem formulation

In this section, we define the specific problem of an FX dealer, who accepts flow from corporate, retail and institutional clients. The dealer wishes to *hedge* the risk of his position only by opening or closing new positions in the spot market. Here we assume:

- The dealer does not perform speculation and is only interested in hedging.
- The dealer makes decision in discrete time-steps.
- There is a limit on how much the dealer can trade at each time-step.
- There is a limit on how much open position the dealer is allowed or willing to keep at each time-step.
- The dealer must close all positions at the end of a trading session (eg, daily or weekly). This is common practice to avoid carrying an open position’s risk during non-business hours (Lyons, 1998).
- Market impact affects the interbank market’s bid-ask spread (Borkovec et al., 2013).

- Client flow is always initiated by the clients, and the dealer is unable to refuse them, assuming credit status and other conditions (eg, anti-money laundering laws) are satisfied.
- The dealer works with local clients (either corporate, retail or institutional), under their own geographic time-zone. These type of clients usually buy more foreign currency than sell (Ranaldo, 2009).
- FX rate jumps resulting from macroeconomic news announcements are a large contributor to the dealer’s risk (Evans, 2011). The timing of these events is known, but the direction and magnitude of the jumps are not.

The dealer’s profit and loss (P&L) arises from three major sources:

1. Transaction costs received from clients.
2. Transaction costs paid to interbank market counterparties for hedging trades.
3. Market volatility.

Here, we assume that the profit from transaction costs received from clients are not influenced by hedging. As a result, they are not considered and only volatility and transaction costs paid to interbank market counterparties are included in the optimization formulation.

3.1. Dealer dynamics

We formulate the problem in discrete-time notation, with the dealer trading at $t \in [0, 1, \dots, N]$. The trading session ends at $t = N + 1$. At any time t , the dealer holds a position of $x_k(t) \in \mathbb{R}$ for currency k . The dealer’s initial position is denoted by $x_k(0)$. Positions vary in time as a result of the accumulated client flow $f_k(t) \in \mathbb{R}$ and the dealer’s hedging actions $h_k(t) \in \mathbb{R}$ (Figure 2):

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

The dealer must close all positions at the end of a trading session (eg, daily or weekly), therefore $x_k(N + 1) = 0, \forall k$. As a result, the last hedging action has to be

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




Figure 2: Hedging state-space dynamics. The dealer's position at time $t + 1$, $x(t + 1)$, is a result of his position at time t being updated with accumulated client flow in this period $f(t)$ and hedged with $h(t)$.

At time t , the cumulative transaction costs paid by the dealer for hedging is

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

where $f_{cost}(h, t, k) = s_k(h, t)|h|$ is a time dependent transaction cost function based on the bid-ask spread $s_k(h, t)$ quoted for trading h at time t , and $|h|$ is the size (ie, absolute value) of h . The spread is known to be highly volatile, and is affected by market impact (ie, increases with the size of h), market conditions (eg, the spread decreases in liquid markets), and the history of trades between counterparties. The spreads are modeled as having a linear relationship with the size of trade (Borkovec et al., 2013),

$$s_k(h, t) = \delta_k(t)|h| \quad (4)$$

where $\delta_k(t)$ is the time-varying market impact coefficient. This results in a quadratic cost function:

$$f_{cost}(h, t, k) = \delta_k(t)h^2 \quad (5)$$

Additionally, we limit the size of hedging actions to $h_{k,max}$:

$$|h_k(t)| \leq h_{k,max} \quad (6)$$

This is realistic as the dealer will have trading limits with its counterparties in the interbank market and usually wishes to avoid impacting the market's liquidity beyond a certain degree.

FX rate volatility changes the value of positions, causing profit and loss. We denote the logarithmic market price of currency k with $p_k(t)$ and its returns $r_k(t)$ with

$$r_k(t) = p_k(t) - p_k(t-1).$$

The profit and loss (P&L), $L_k(t)$, is therefore given by

$$L_k(t) = \sum_{i=1}^t x_k(i)r_k(i). \quad (7)$$

The dealer usually wishes to avoid the risk of excessive loss, thus the maximum position of each currency is restricted to

$$|x_k(t)| \leq x_{k,max}. \quad (8)$$

4. Dynamic FX risk management

Dealers wish to reduce their transaction costs and risk exposure, subject to the discrete space representation and constraints defined in Section 3.1. Considering the probabilistic nature of client flow and FX rate, we formulate this as a mean-variance optimization, $\mathbb{E}[C(N)] + \lambda \text{Var}[L(N)]$, where $0 \leq \lambda < \infty$ is the risk aversion factor. The following equation, considering the probabilistic nature of client flow and FX rate, describes the minimization problem:

$$\begin{aligned} & \underset{h_k(t); \forall t,k}{\text{argmin}} \quad \mathbb{E} \left[\sum_{k=1}^m \sum_{t=0}^N f_{cost}(h_k(t)) \right] + \lambda \text{Var} \left[\sum_{k=1}^m \sum_{t=1}^N x_k(t)r_k(t) \right] \\ \text{subject to} \quad & x_k(t+1) = x_k(t) + h_k(t) + f_k(t) \\ & x_k(N+1) = 0 \\ & |h_k(t)| \leq h_{k,max} \\ & |x_k(t)| \leq x_{k,max} \end{aligned} \quad (9)$$

where m is the number of currencies in the dealer's portfolio.

Replacing the transaction cost function in (9) with (5), applying constraint (2), and expanding $x_k(t)$ recursively using (1) results in the following

cost function:

$$\begin{aligned}
J = & \mathbb{E} \left[\sum_{k=1}^m \sum_{t=0}^{N-1} \delta_k(t) h_k(t)^2 + \right. \\
& \left. \sum_{k=1}^m \delta_k(N) \left(\sum_{i=0}^{N-1} h_k(i) + \sum_{i=0}^N f_k(i) + x_k(0) \right)^2 \right] + \\
& \lambda \text{Var} \left[\sum_{k=1}^m \sum_{t=1}^N \left(\sum_{i=0}^{t-1} h_k(i) + \sum_{i=0}^{t-1} f_k(i) + x_k(0) \right) r_k(t) \right]
\end{aligned} \tag{10}$$

To simplify notation, we introduce variable $y_k(t)$, as the unhedged position of the dealer:

$$\begin{cases} y_k(0) = x_k(0) \\ y_k(t) = \sum_{i=0}^{t-1} f_k(i) + x_k(0) \end{cases} \tag{11}$$

This simplifies (10) to

$$\begin{aligned}
J = & \mathbb{E} \left[\sum_{k=1}^m \sum_{t=0}^{N-1} \delta_k(t) h_k(t)^2 + \right. \\
& \left. \sum_{k=1}^m \delta_k(N) \left(\sum_{i=0}^{N-1} h_k(i) + y_k(N+1) \right)^2 \right] + \\
& \lambda \text{Var} \left[\sum_{k=1}^m \sum_{t=1}^N \left(\sum_{i=0}^{t-1} h_k(i) + y_k(t) \right) r_k(t) \right].
\end{aligned} \tag{12}$$

Optimizing (12) results in a risk-cost Pareto frontier parametrized by λ . Using this information, the dealers choose the best hedging strategy considering contextual factors, such as their utility of risk.

4.1. Matrix Notation

We use the following matrix notation to simplify representation of the dealer dynamics. Vectors $\mathbf{x}_k, \mathbf{f}_k, \mathbf{h}_k, \boldsymbol{\delta}_k, \mathbf{r}_k, \mathbf{y}_k \in \mathbb{R}^N$ are the collection of

variables $x_k(t)$, $f_k(t)$, $h_k(t)$, $\delta_k(t)$, $r_k(t)$ and $y_k(t)$ in time:

$$\begin{aligned}\mathbf{x}_k &= [x_k(0) \ x_k(1) \ \cdots \ x_k(N-1)]^\top \\ \mathbf{f}_k &= [f_k(0) \ f_k(1) \ \cdots \ f_k(N-1)]^\top \\ \mathbf{h}_k &= [h_k(0) \ h_k(1) \ \cdots \ h_k(N-1)]^\top \\ \boldsymbol{\delta}_k &= [\delta_k(0) \ \delta_k(1) \ \cdots \ \delta_k(N-1)]^\top \\ \mathbf{r}_k &= [r_k(1) \ r_k(2) \ \cdots \ r_k(N)]^\top \\ \mathbf{y}_k &= [y_k(1) \ y_k(2) \ \cdots \ y_k(N)]^\top\end{aligned}$$

Notice that \mathbf{r}_k and \mathbf{y}_k are indexed differently.

Vectors $\mathbf{X}, \mathbf{H}, \mathbf{F}, \boldsymbol{\delta}, \mathbf{R}, \mathbf{Y} \in \mathbb{R}^{mN}$ are defined as concatenations of \mathbf{x}_k , \mathbf{f}_k , \mathbf{h}_k , $\boldsymbol{\delta}_k$, \mathbf{r}_k and \mathbf{y}_k over all currencies. For example, for 3 currencies USD, EUR and GBP,

$$\mathbf{X} = [\mathbf{x}_{USD} \ \mathbf{x}_{EUR} \ \mathbf{x}_{GBP}]^\top. \quad (14)$$

Similarly, vectors $\mathbf{x}(t)$, $\mathbf{f}(t)$, $\mathbf{h}(t)$, $\boldsymbol{\delta}(t)$, $\mathbf{r}(t)$, $\mathbf{y}(k) \in \mathbb{R}^m$ are the concatenation of their respective variables at time t for all k , eg,

$$\mathbf{x}(t) = [x_{USD}(t) \ x_{EUR}(t) \ x_{GBP}(t)]^\top. \quad (15)$$

Notice that all vectors except optimization objectives \mathbf{H} , $\mathbf{h}(t)$ and \mathbf{h}_k , are stochastic processes.

The matrix form of (12) is

$$\begin{aligned}J &= \mathbb{E} \left[\mathbf{H}^\top \boldsymbol{\Delta} \mathbf{H} + (\boldsymbol{\Upsilon}^\top \mathbf{H} + \mathbf{y}(N+1))^\top \mathbf{D} (\boldsymbol{\Upsilon}^\top \mathbf{H} + \mathbf{y}(N+1)) \right] + \\ &\quad \lambda \text{Var} \left[(\boldsymbol{\Psi} \mathbf{H} + \mathbf{Y})^\top \mathbf{R} \right] \quad (16)\end{aligned}$$

where

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

$$\begin{aligned}\boldsymbol{\Upsilon}_{mN \times m} &= \mathbf{I}_{m \times m} \otimes \vec{\mathbf{1}}_{N \times 1}, \\ \mathbf{D}_{m \times m} &= \text{diag}(\boldsymbol{\delta}(N)), \\ \boldsymbol{\Delta}_{mN \times mN} &= \text{diag}(\boldsymbol{\delta}),\end{aligned}$$

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

Additionally, matrix form of (11) is

$$\mathbf{Y} = \Psi \mathbf{F} + \Upsilon \mathbf{x}(0). \quad (17)$$

Stochastic operators in (16) can be expanded,

$$\begin{aligned} J = & \mathbf{H}^T \mathbb{E}[\Delta] \mathbf{H} + \mathbf{H}^T \Upsilon \mathbb{E}[\mathbf{D}] \Upsilon^T \mathbf{H} + \\ & 2\mathbf{H}^T \Upsilon \mathbb{E}[\mathbf{D} \mathbf{y}(N+1)] + \mathbb{E}[\mathbf{y}(N+1)^T \mathbf{D} \mathbf{y}(N+1)] + \\ & \lambda (\text{Var}[\mathbf{Y}^T \mathbf{R}] + \mathbf{H}^T \Psi^T \text{Var}[\mathbf{R}] \Psi \mathbf{H} + 2\mathbf{H}^T \Psi^T \text{Cov}[\mathbf{R}, \mathbf{Y}^T \mathbf{R}]) \end{aligned}$$

and then reordered by \mathbf{H} to

$$\begin{aligned} J = & \mathbf{H}^T (\mathbb{E}[\Delta] + \Upsilon \mathbb{E}[\mathbf{D}] \Upsilon^T + \lambda \Psi^T \text{Var}[\mathbf{R}] \Psi) \mathbf{H} + \\ & 2\mathbf{H}^T (\Upsilon \mathbb{E}[\mathbf{D} \mathbf{y}(N+1)] + \lambda \Psi^T \text{Cov}[\mathbf{R}, \mathbf{Y}^T \mathbf{R}]) + \\ & \mathbb{E}[\mathbf{y}(N+1)^T \mathbf{D} \mathbf{y}(N+1)] + \lambda \text{Var}[\mathbf{Y}^T \mathbf{R}], \end{aligned} \quad (18)$$

which is a quadratic function of hedging actions \mathbf{H} .

4.2. Hedging constraints

Unconstrained RHC problems with quadratic cost function and linear state-space are referred to as linear-quadratic (LQ) control problems and their solution can be found using the matrix algebraic Riccati equation (Borrelli et al. (2014, Chapter 9)). With deterministic constraints, (18) can be solved efficiently using numeric quadratic programming techniques.

The problem formulation in (9) includes four inequality constraints. The first two constraints were implicitly implemented as a part of the cost function in Section 4. The third constraint, $|h_k(i)| \leq h_{k,max}$, is deterministic and easily described by:

$$\begin{aligned} -\mathbf{H} & \leq \mathbf{H}_{max} \\ \mathbf{H} & \leq \mathbf{H}_{max} \end{aligned}$$

The fourth constraint requires special attention, as $x_k(i), i \in [1, \dots, N]$ are stochastic variables, and therefore are not deterministically constrainable before being observed. However, in a real-world multi-stage optimization, one

can use $\mathbf{f}(0)$ and $\mathbf{x}(0)$, which are known before solving for $\mathbf{h}(0)$ (as shown in Figure 2), to make the first inequality, $\mathbf{x}(1) < \mathbf{x}_{max}$, deterministic. The rest of inequalities can be replaced with their expected value (Calafiore, 2008), ie, $|\mathbb{E}[x_k(i)]| \leq x_{k,max}$. While this results in a less conservative control policy, it still guarantees constraints fulfillment (Bernardini and Bemporad, 2012).

Other approaches are also possible, including imposing constraints to hold with a given sufficiently high probability, or using the worst case of a scenario tree. Both of these approaches may not hold for some situations (eg, in case of scenario trees, when the outcome is not covered by the generated scenarios), and suffer from higher computational complexity.

In this paper, the expected value constraint is implemented. Rewritten in matrix form, and given $\mathbb{E}[\mathbf{X}] = \mathbb{E}[\Psi\mathbf{H} + \mathbf{Y}]$, the following constraints are imposed:

$$\begin{aligned} -\Psi\mathbf{H} &\leq \mathbf{X}_{max} + \mathbb{E}[\mathbf{Y}] \\ \Psi\mathbf{H} &\leq \mathbf{X}_{max} - \mathbb{E}[\mathbf{Y}] \end{aligned}$$

4.3. Simplification of the hedging cost function

If client flow is uncorrelated to the market conditions (ie, the FX rate returns, volatility and market impact coefficient), (18) can be further simplified to

$$\begin{aligned} J = &\mathbf{H}^T (\bar{\Delta} + \Upsilon\bar{\mathbf{D}}\Upsilon^T + \lambda\Psi^T\Sigma\Psi) \mathbf{H} + \\ &2\mathbf{H}^T (\Upsilon\bar{\mathbf{D}}\bar{\mathbf{y}}(N+1) + \lambda\Psi^T\Sigma\bar{\mathbf{Y}}), \end{aligned} \tag{20}$$

where $\Sigma = \text{Var}[\mathbf{R}]$ is the returns covariance matrix and for every variable the expected value $\mathbb{E}[x]$ is denoted by \bar{x} . Notice that optimization's constant terms are also dropped. Consequently the only requirements are

- expected position of the dealer in absence of hedging $\bar{\mathbf{Y}} = \mathbb{E}[\mathbf{Y}]$, or the expected client flow $\bar{\mathbf{F}} = \mathbb{E}[\mathbf{F}]$ as $\mathbb{E}[\mathbf{Y}] = \Psi\mathbb{E}[\mathbf{F}] + \mathbf{x}(0)$,
- expected market impact coefficient $\bar{\delta} = \mathbb{E}[\delta]$, which is then formed into diagonal matrices $\bar{\Delta}$ and $\bar{\mathbf{D}}$, and
- FX rate returns covariance matrix Σ .

One must note that this assumption is not generally true. It has been shown that order flow can exhibit a strong relationship with FX rates (Evans and Lyons, 2002), specially cumulative order flow being highly correlated

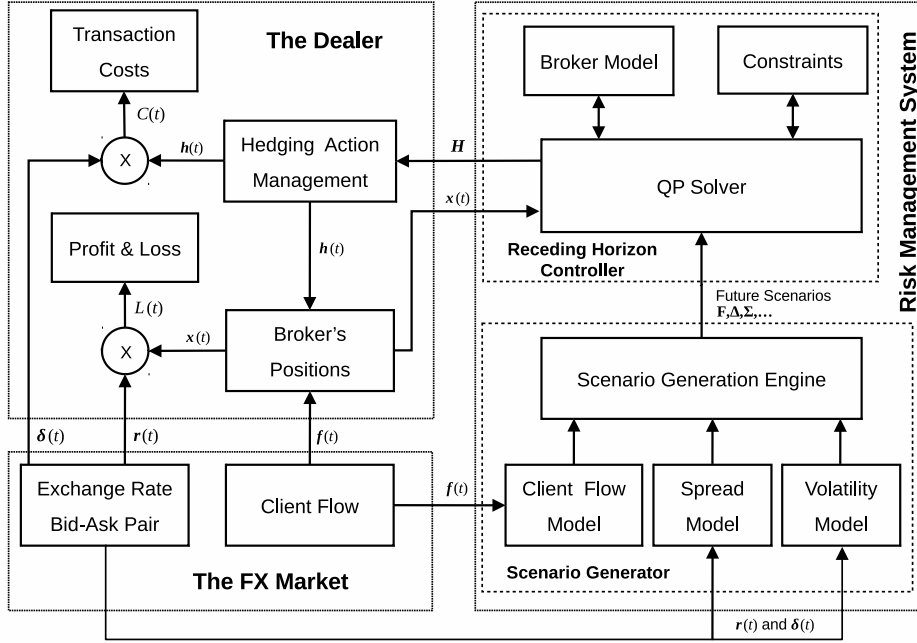


Figure 3: The proposed FX risk management system and its components.

with cumulative price change (McGroarty et al., 2009). Hence, while (20) makes the analysis more similar to the original Markovitz portfolio theory formulation, it may produce inferior results compared to the full model of (18).

4.4. The risk management system architecture

Figure 3 illustrates the relationship of the proposed risk management system to the dealer and the FX market. Notice that the system is separated from the dealer, and its control actions are passed through an *action management* subsystem which is arbitrated by the dealer. This subsystem acts as a circuit breaker, monitoring the behavior of the hedging system and alerting the dealer if any problems are found. Additionally, the design allows fault tolerance through N-module redundancy; multiple identical instance of the risk management systems can be employed and queried, and the correct results can be chosen on the basis of a quorum, thus making the system tolerant to hardware failure.

Algorithm 1: Stochastic receding horizon hedging algorithm.

```
1  $t \leftarrow 0$ .
2 while  $t \leq N$  do
3   Update client flow models using new trades data.
4   Update FX rate returns models from market conditions.
5   Update market impact coefficient models from market
   conditions.
6   Generate  $\eta$  scenarios for  $\mathbf{F}$ ,  $\mathbf{R}$  and  $\boldsymbol{\delta}$  for time  $t \in [t, \dots, N]$ .
7   Compute the optimal hedging action
    $\mathbf{H} = [\mathbf{h}(t), \mathbf{h}(t+1), \dots, \mathbf{h}(N)]$  by minimizing (18).
8   Hedge by  $\mathbf{h}(t)$ .
9   Update positions  $\mathbf{x}(t)$ .
10   $t \leftarrow t + 1$ .
11 end
12 Close all position using  $\mathbf{h}(N) = -(\mathbf{x}(N) + \mathbf{f}(N))$ .
```

4.5. Receding horizon optimization

In real-world systems, exact modeling of the client flow and the FX market is impossible due to incomplete data and approximations, and consequently the accuracy of hedging actions is lost over time. To improve performance, one can include new information in the optimizer using a receding horizon scheme, as summarized in Algorithm 1: at each time-step, client flow, FX rate returns and market impact coefficient models are updated based on observed market conditions. Monte Carlo methods sample scenarios based on the models, which are then given to the quadratic programming (QP) solver to obtain the optimal hedging action \mathbf{H} by minimizing (18) subject to the constraints of Section 4.2. The positions are hedged by $\mathbf{h}(0)$, which adds to the transaction costs. At the next time-step, FX volatility may change the rates and bid-ask spreads, and consequently the value of dealer’s open positions, thus generating profit or loss. The process is then repeated from its first step, and continues until the end-of-trading time is reached.

5. Scenario generation and data models

In this section, we describe several statistical models which are used to generate synthetic scenarios for in-depth testing of the proposed hedging system.

As we are also interested in scenario generating models with different degrees of predictive modeling accuracy, we define an oracle such that it generates scenarios by perturbing observations with alternative scenarios:

$$x^{(p)}(t) = \alpha e^{-\beta t} x(t) + (1 - \alpha e^{-\beta t}) x'(t) \quad (21)$$

Here, $x^{(p)}(t)$ is the oracle generated (ie, perturbed) scenario used in optimizations, $x(t)$ is the actual future data used for backtesting, $x'(t)$ is an alternative scenario sampled from the models' statistical distribution, and $0 \leq \alpha \leq 1$ and $\beta \geq 0$ are initial accuracy and oracle decay rate parameters. To generate extreme cases, $\alpha = 1$ and $\beta = 0$ are used for prescient hedging, while $\alpha = 0$ tests the quality of data models with ordinary scenario generation. Furthermore, $\alpha > 0$ and $\beta > 0$ create an exponentially decaying accuracy, simulating a more realistic signal decay model.

Each scenario consists of three correlated items: FX rate returns, client flow volume and market impact coefficients. Each of these components are parametrized differently for every currency in a multi-currency scenario.

5.1. FX rate and volatility model

The FX logarithmic price process can be expressed as a continuous-time jump-diffusion process (Evans, 2011),

$$dp(\tau) = \mu_p(\tau)d\tau + v(\tau)dW(\tau) + k(\tau)dq(\tau) \quad (22)$$

where $p(\tau)$ is the FX rate for the continuous time τ , $\mu_p(\tau)$ (the drift coefficient) is risk free interest, $W(\tau)$ is the Wiener process, $v(\tau)$ is the diffusion coefficient (the volatility), $k(\tau)$ measures the jumps' intensity and $q(\tau)$ is a counting process which is determined by the scheduled time of macroeconomic announcements. Here, we assume $\mu_p \approx 0$ as the optimization horizon is too short for the interest rate to be effective.

Figure 4 shows an example of different FX rate scenarios generated for an event occurring at 12:00.

We model the discrete-time returns $r(t)$ using

$$r(t) = p(t\Delta\tau) - p((t-1)\Delta\tau) \quad (23)$$

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

The discrete-time form of volatility, $v(t)$, is modeled using an M-shaped pattern to account for daily pattern of liquidity variation (McGroartya et al.,




Figure 4: Simulated FX rate scenarios with an announced event at 12:00 PM.

2009):

$$v(t) = \nu \left(1 + \frac{\rho - 1}{2} \left(1 + \cos \left(2\pi \frac{t + \frac{N-w}{2} - t_{min}}{N - w} \right) \right) \right) \quad (24)$$

Here, ν is the minimum daily volatility, ρ is the ratio of maximum daily volatility to its minimum, N is the number of time-steps in the day, t_{min} is the time of volatility minima (ie, maximum liquidity) and w determines width of the M-shape peaks. Figure 5 presents an example of daily volatility scenarios generated using (24) and their average.

This model can be extended to a multivariate form, accounting for correlation between different FX rates. For a given correlation matrix \mathbf{C} , the following equation will generate correlated stochastic variables (Glasserman (2003, Chapter 2)):

$$\mathbf{W} = \mathbf{W}_{i.i.d.} \mathbf{L} \quad (25)$$

Here, \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 stochastic variables, and \mathbf{W} is the resulting matrix of correlated stochastic variables.




Figure 5: Simulated M-shaped volatility during the day.

Correlated Wiener processes generated using (25) are then used in (22). Addition of events or drift are performed in a manner analogous to the univariate case.

5.2. Transaction cost model

The main component of transaction cost, as defined in (5), is the bid-ask spread, which in (4) was modeled as an affine function of trade size.

In addition to the size of trade, the bid-ask spread is also affected by liquidity. The liquidity increases during mid-trading hours, based on the geographic distribution of currency traders, and is reduced in non-trading hours (McGroarty et al., 2009).

We model this time-varying liquidity effect with a U-shaped market impact coefficient:

$$\delta(t) = \delta \left(1 + \frac{\rho - 1}{2} \left(1 + \cos \left(2\pi \frac{t + \frac{N}{2} - t_{min}}{N} \right) \right) \right) \quad (26)$$

Here, $\delta(t)$ is the time varying market impact coefficient, δ is the coefficient's daily minimum, ρ is the ratio of maximum daily market impact coefficient to its minimum, N is the number of time-steps in the day, and t_{min} determines the time of maximum liquidity.




Figure 6: Simulated U-shaped bid-ask spreads during the day.

Figure 6 shows an example of scenarios for daily bid-ask spread and their average generated by (26).

5.3. Client flow model

The choice of client flow model is a dealer dependent task. Cerrato et al. (2011) noted that in response to price moves, corporate and private clients act more as liquidity providers while profit-motivated traders (eg, hedge fund investors and asset managers) display a directional behavior as a result of being more informed. Furthermore, Chen et al. (2012) found significant relationship between volatility, spreads and the order flow of speculative traders. This information can be used by dealers to better model their client flow.

In this paper, we model private domestic clients who mainly trade during their own country's working hours. These clients exhibit a certain periodicity, eg, more trades happen mid-day rather than at the end-of-day hours. Also the dealers expect a bias in buys versus sells in a certain direction relative to their home currency, which is not influenced by price movement. Therefore the client flow is modeled as a heteroskedastic process, with a time-dependent mean μ_f and variance σ_f^2 :

$$f(t) \sim \mathcal{N}(\mu_f(t), \sigma_f^2(t)) \quad (27)$$




Figure 7: The synthetic client flow standard deviation versus time.

Figure 7 shows an example of $\sigma_f(t)$ for a dealer trading from 7:00 to 23:00, with the assumptions presented earlier.

6. Results

Several tests using synthetic and real-world data were performed to validate and measure the performance of the proposed hedging algorithm.

The experiments were implemented in R programming language (R Core Team, 2014) and used the quadprog package (Weingessel, 2013) for quadratic programming. Run-time of the hedging algorithm for each trading session with 50 scenarios on a 3.4 GHz Intel Core i7-2600 processor was measured to be less than 250 ms, making it suitable not only for backtesting, but also for deployment in online hedging systems. The source codes to our program are available from <https://sydney.edu.au/engineering/electrical/cel/farzad/JoR2016>.

6.1. Test data

6.1.1. Synthetic data

The synthetic data was generated according to the models described in Section 5. The dealer is assumed to work from 7:00 to 23:00 and hedge at 30 minute steps.

Client flow was generated with standard deviation σ_f as depicted in Figure 7 and a positive mean $\mu_f = \frac{1}{4}\sigma_f$.

For the FX rate volatility and market impact coefficients, each currency was assigned different δ , ν , ρ , ω and t_{min} parameters. At the start of each trading session, the asset correlation matrix \mathbf{C} , and number, time and impact of announced events $k(t)$ were regenerated randomly.

6.1.2. Real-world data

For real-world tests, 16 weeks of actual FX client transaction data supplied by Westpac institutional bank was used¹. AUD was chosen as the home currency and USD, EUR, NZD and JPY were selected as foreign assets. This data was filtered and aggregated to create 32 half-hourly values per day, from 7:00 to 23:00 Australian Eastern Daylight Time (AEDT). The first six weeks (30 days) of data was used to obtain model parameters as described in Section 5.3 and the rest (ie, 50 days) were used for out-of-sample backtesting.

Accordingly, half-hourly values of historical FX rates were fit to the model defined in Section 5.1. Individual variances ν and correlation matrix \mathbf{C} were computed from covariances of the previous day, and jump component timings were extracted from the publicly available DailyFX.com event calendar. Jump intensity was set to $k = 5\nu$. Only events classified in DailyFX.com as high impact were considered in this simulation and the rest were discarded.

As the interbank market impact data was not available, U-shaped daily seasonality was generated synthetically. The market impact coefficient δ was selected as the square of average spread, 0.5, 1, 2 and 1 per 10000 for USD, EUR, NZD and JPY respectively. The time of maximum liquidity was chosen as the 12:00 PM according to each market's geographic distribution, ie, New York, Frankfurt, Auckland and Tokyo respectively.

6.2. Benchmark strategies

We used three benchmark hedging strategies to compare against the proposed hedging method. The first strategy limits the dealer's position to x_{max} :

$$h_k(t) = \begin{cases} x_{max} - (x_k(t) + f_k(t)) & x_k(t) + f_k(t) > x_{max} \\ -x_{max} - (x_k(t) + f_k(t)) & x_k(t) + f_k(t) < -x_{max} \\ 0 & \text{otherwise} \end{cases} \quad (28)$$

¹This data was anonymized to protect Westpac client confidentiality.

As this limit changes, the strategy sweeps the cost-risk profiles from minimum risk ($x_{max} = 0$, ie, holding no risk) to minimum cost ($x_{max} \rightarrow \infty$, ie, letting the incoming client buy and sell orders neutralize each other).

Despite its simplicity, this strategy has major shortcomings. It does not utilize client flow and volatility information, nor it models the market impact. As a result, not only the hedging is far from optimal, but also the dealer could be left with a large open position at the end-of-trading session, imposing additional transaction costs.

The second benchmark strategy distributes the hedging actions required to close the current positions over the remaining trading time-steps, thus *gradually closing* all positions in order to avoid market impact effects and additional transaction costs. This strategy is further parametrized by λ as a risk aversion parameter, resulting in actions determined by

$$h_k(t) = -(x_k(t) + f_k(t)) \frac{(N-t)\lambda + 1}{N-t+1} \quad (29)$$

where $\lambda = 1$ gives the minimum risk by instantly closing all positions, and $\lambda = 0$ results in equally sized hedging actions, seeking the minimum cost by avoiding market impact.

For the third benchmark, a receding horizon control methodology was applied on a single-period form of (20), ie, $\mathbf{H} = \mathbf{h}(0)$; consequently, $\bar{\mathbf{Y}} = \mathbb{E}[\mathbf{y}(0)] = \mathbf{y}(0)$ and Ψ becomes an identity matrix. It is assumed that the same $\mathbf{h}(t) = \mathbf{h}(0), t > 0$ is used in the next $N - 1$ time-steps to close the expected final position $\bar{\mathbf{y}}(N+1)$; as a result, the transaction cost coefficients are modeled by $\bar{\boldsymbol{\delta}} = \mathbb{E}[\boldsymbol{\delta}]$, $\bar{\mathbf{D}} = \text{diag}(\bar{\boldsymbol{\delta}})$, $\bar{\boldsymbol{\Delta}} = (N-1)\text{diag}(\bar{\boldsymbol{\delta}})$ and Υ is set to the scalar value of $N - 1$. The risk covariance matrix $\boldsymbol{\Sigma}$ is computed as the volatility of log price returns from $t = 0$ to $t = N$, ie, $p_k(N) - p_k(0)$. Therefore for each time-step, the optimization cost function (20) transforms to

$$J = \mathbf{H}^T ((N^2 - N)\bar{\mathbf{D}} + \lambda\boldsymbol{\Sigma}) \mathbf{H} + 2\mathbf{H}^T ((N-1)\bar{\mathbf{D}}\bar{\mathbf{y}}(N+1) + \lambda\boldsymbol{\Sigma}\mathbf{y}(0)). \quad (30)$$

For $\lambda = 0$, the solution of $\mathbf{H} = \frac{-\bar{\mathbf{y}}(N+1)}{N}$ results in the minimum transaction cost and for the $\lambda \rightarrow \infty$, $\mathbf{H} = -\mathbf{y}(0)$ gives the minimum risk, similar to the second benchmark strategy.

It must be noted that while minimum risk is universal for all strategies (ie, $h_k(t) = -f_k(t)$ results in $x_k(t) = 0$ and therefore $L_k(t) = 0$ according

to (7)), the actual minimum cost is only obtainable by minimizing (18) with $\lambda = 0$ and full knowledge of future client flow and spreads. Therefore results of minimizing (18) using future data for different parameterizations of $0 \leq \lambda < \infty$ are included in comparisons as the *prescient hedging* Pareto frontier.

6.3. Single-asset tests

To test the correctness of the proposed hedging scheme, first a simple test with one foreign currency asset for one trading session was performed. Data was synthetically generated, as described in Section 6.1.1. FX rate returns were created with only one announcement event occurring at 12:00, as shown in Figure 4.

Figure 8 compares the accumulation of positions in absence of hedging against different stochastic paths. Figure 9 and Figure 10 show the hedged positions and hedging actions respectively using the SRHC hedging with $\lambda = 0.01$. Figure 9 and 10 also show results of open-loop optimization in gray; each line represents the results if the hedging actions were not recalibrated again at each time-step according to Algorithm 1.

It can be seen from Figure 10 that the proposed method has gradually reduced the positions to zero at the end of trading session, therefore reducing market impact and costs. Furthermore, comparing hedged and unhedged positions in Figure 11 reveals that by using the knowledge of the announced event at 12:00, the open positions were gradually minimized before the event to reduce the possibility of losses in case of a major unfavorable FX rate movement.

6.4. Multi-asset tests

To measure the hedging performance in a multi-asset portfolio, back-testing was performed with four foreign currencies using both synthetic and real-world data.

A total of 50 different trading sessions were backtested for the three benchmark strategies, scenario based SRHC hedging and prescient hedging. To generate cost-risk profiles, each run included different parametrizations of x_{max} in (28) for the first benchmark strategy and λ for other strategies. The end-of-trading costs and profits were computed according to (3) and (7) respectively and normalized by the daily total value of transactions $\sum |f(i)|$ to basis points (bps). The normalized values were then converted to cost-risk profiles.

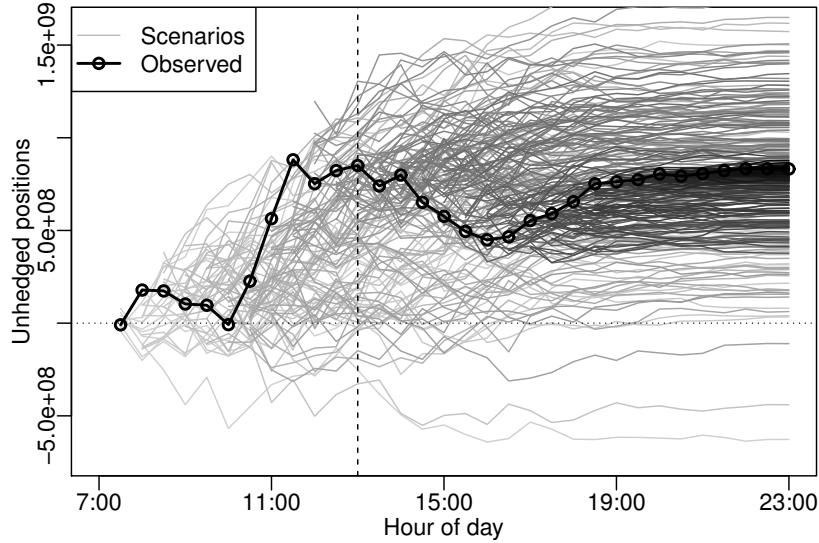


Figure 8: Dealer’s unhedged position $y(t)$ through the trading session in the single-asset test. Each gray line shows a stochastic scenario estimated at various times of trading session, with later scenarios colored darker. The vertical line marks the location of the announcement event.

Figure 12 shows the cost-risk profiles for synthetic data test, while real-world data’s profiles are reported in Figure 13. It is shown that SRHC hedging outperforms all benchmark hedging strategies by offering a lower cost for any given risk in the synthetic test. For selected levels of risk, Table 1 presents quantified values of the normalized cost and the percentage of cost savings scaled between the second benchmark and the prescient hedging. It is observed that the SRHC technique can improve costs up to 47.6%.

It was noted that in the real-world experiment, the cost saving improvements are noticeable only compared to the first two benchmarks which do not employ any information about the future. For example, for risk = 10 bps, the SRHC improvement is merely 2.4% better than single-stage hedging. Furthermore, there is considerable room for improvement compared to the prescient hedging frontier. This problem can be traced to scenario generation, where the simple models of Section 5 are unable to capture certain features such as the fat-tailed distribution of data, correlation between FX rate and client flow, and variation of transaction cost around events for real-world data test.

To confirm the superiority of the proposed algorithm over the single-




Figure 9: Dealer's hedged position $x(t)$ in the single-asset test. Each gray line shows result of open-loop hedging at various times of trading session, with later results colored darker. The vertical line marks the location of the announcement event.

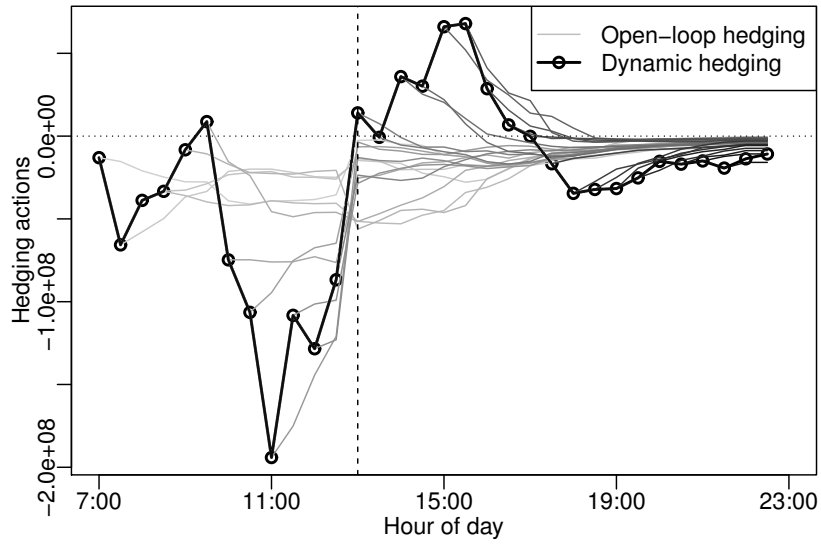


Figure 10: Hedging actions $h(t)$ in the single-asset test. Each gray line shows results of optimization performed at various times of trading session, with later optimizations colored darker. The vertical line marks the location of the announcement event.

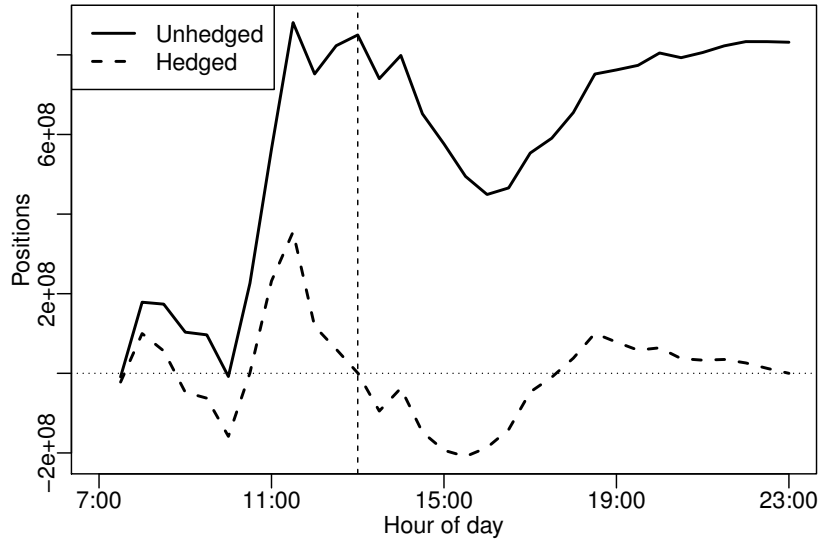


Figure 11: Comparison of the dealer's positions with and without hedging in the single-asset test.




Figure 12: Risk-cost profiles for the synthetic data experiment.

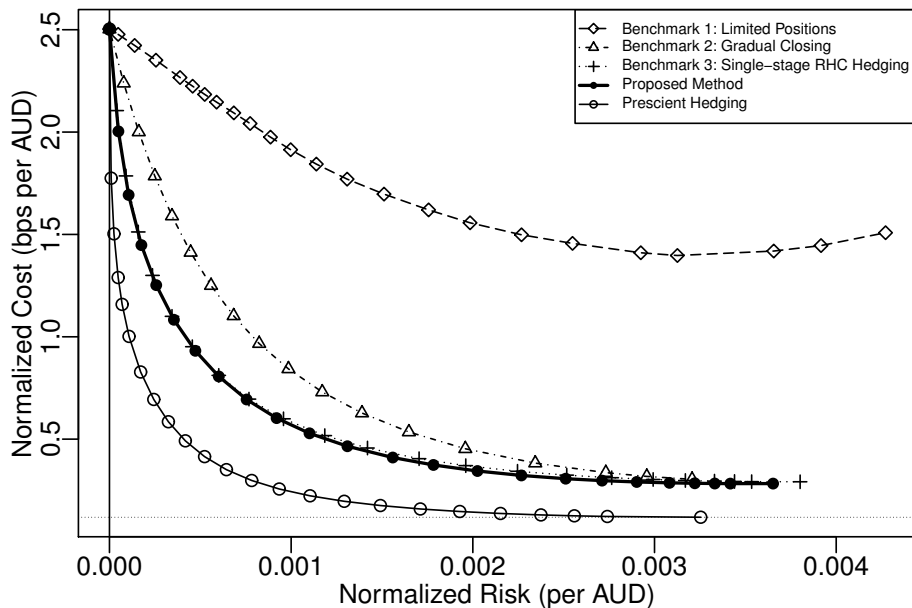


Figure 13: Risk-cost profiles for the real-world data experiment.

Table 1: Costs of different hedging strategies for selected levels of risk in the real-world data experiment. The costs and risk levels are normalized and presented in basis points. The values in percentage denote cost saving improvements, scaled between the gradual closing strategy and prescient hedging.

| Hedging strategy | Normalized cost for | | | | | |
|------------------|---------------------|-------|---------------|-------|---------------|-------|
| | Risk = 5 bps | | Risk = 10 bps | | Risk = 20 bps | |
| | bps | % | bps | % | bps | % |
| Gradual closing | 1.336 | 0% | 0.834 | 0% | 0.445 | 0% |
| Single-stage RHC | 0.909 | 47.2% | 0.584 | 42.3% | 0.367 | 26.1% |
| SRHC (proposed) | 0.905 | 47.6% | 0.570 | 44.7% | 0.347 | 32.4% |
| Prescient | 0.432 | 100% | 0.244 | 100% | 0.143 | 100% |

stage method regardless of the data model quality, the real-world test was repeated with prescient data in Figure 14. A comparison shows that the proposed method significantly outperforms other strategies when its input scenarios are accurate.




Figure 14: Risk-cost profiles for the real-world data experiment with prescient data.

6.5. Scenario quality and oracle tests

In the previous experiments, the impact of scenario generation accuracy on hedging was demonstrated. In this section, this impact is quantified using the oracle model of (21).

Figure 15 and Figure 16 compare hedging using the scenarios generated by the oracle with $\alpha = 0.5, \beta = 0.9$ and $\alpha = 0.9, \beta = 0.5$ for all scenario components, against ordinary scenario generation and prescient hedging for synthetic and real-world data respectively. As expected, improved scenario generation (ie, increasing initial accuracy α and reducing decay rate β) results in less cost for any chosen level of risk.

The effect of improving accuracy for individual components on risk management is compared in Figure 17 for synthetic and in Figure 18 for real-world data. Here, a high degree of accuracy ($\alpha = 1, \beta = 0.2$) is used to generate different scenarios for client flow, volatility and market impact coefficients separately.

In Table 2 and Table 3, this test is repeated with a wider range of α and β 's for synthetic and real-world data respectively. The results were averaged over the number of backtesting sessions and the normalized costs for risk




Figure 15: Effect of oracle accuracy on risk-cost profile improvement in the synthetic data experiment.

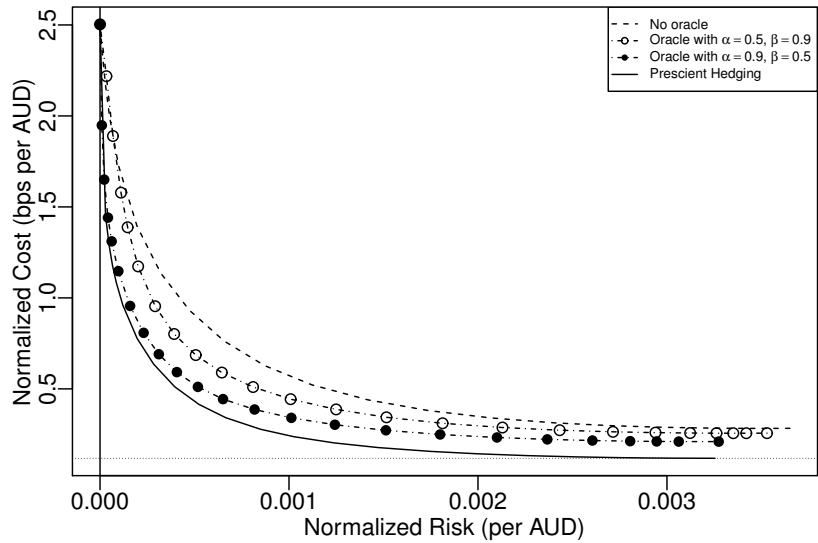


Figure 16: Effect of oracle accuracy on risk-cost profile improvement in the real-world data experiment.

= 10 bps were reported, scaled to denote 0% for the single-stage hedging strategy (ie, the third benchmark) and 100% for prescient hedging. For both tests, it is observed from the tables that improving scenario generation produces better hedging profiles. Yet, hedging enjoys the most improvements with better client flow modeling, while any further market impact coefficient model enhancements have minor effects on the final results. For example in the synthetic test, the ordinary scenario generation reduces costs by 26.7%, and using a perfect market impact model and a perfect volatility model improves this to 28.5% and 38.6% respectively. In comparison, a perfect client flow modeling boosts the cost savings up to 87.8%. A similar phenomenon is noticeable for real-world data, improving cost reduction by 1.2%, 27.3% and 73% with perfect transaction cost, volatility and client flow modeling respectively. A similar trend is also present in Figure 17 and Figure 18.

These differences are the result of using more definite models for market impact and FX rate, versus a less precise model for describing client flow. Most of the actual market impact and FX rate variations are captured by the ordinary scenario generation, with not much left to be improved by the oracle. On the other hand, the simple model of Section 5.3 does not produce accurate client flow scenarios, hence major improvements are observed using the oracle.

Furthermore, Figure 18 shows an interesting effect: improving volatility scenarios offers a better cost versus risk compared to improving client flow for the low risk region (ie, the left side) of the plot. Comparably, in the low cost region (ie, the right side), the cost-risk profile converges with the ordinary scenario generation curve while the client flow only oracle gives a considerably better cost versus risk, converging with the prescient hedging frontier. This is a direct effect of λ 's influence on minimizing the hedging cost function (18); when $\lambda \rightarrow 0$ (ie, only minimizing transaction cost), the effect of client flow modeling error is more evident, and when $\lambda \rightarrow \infty$ (ie, only minimizing risk), the client flow modeling errors become negligible compared to FX rate modeling error. In conclusion, the dealers can choose to invest only on better client flow or volatility modeling depending on their risk preference.

7. Discussion and Future Work

In this paper, a quadratic formulation was devised for minimizing FX hedging transaction costs, resulting in a numerically efficient quadratic programming problem. Here, the bid-ask spread approaches zero as the size of




Figure 17: Effect of using the oracle for individual scenario components on synthetic data experiment's risk-cost profiles.

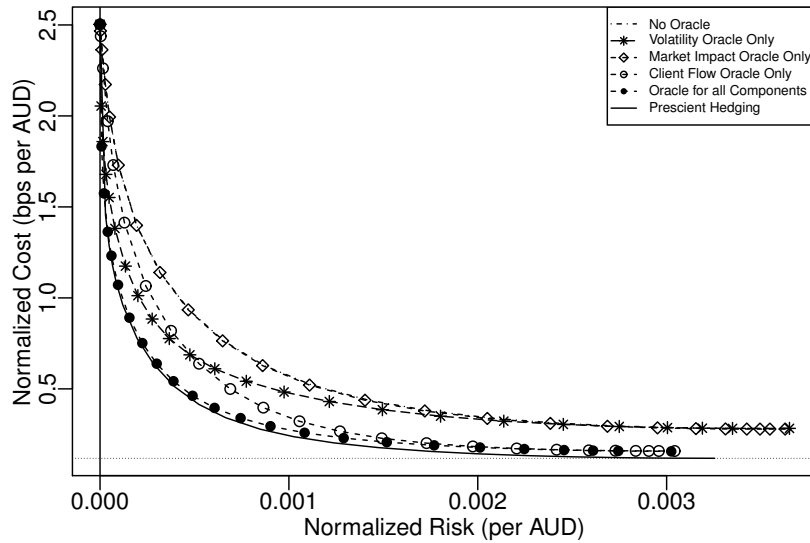


Figure 18: Effect of using the oracle for individual scenario components on real-world data experiment's risk-cost profiles.

Table 2: Cost savings for different oracle accuracies in the synthetic data experiment. The values show costs for risk = 10 bps, scaled as percentage between the single-stage strategy and prescient hedging results.

| Oracle | $\beta = 1$ | $\beta = 0.8$ | $\beta = 0.6$ | $\beta = 0.4$ | $\beta = 0.2$ | $\beta = 0$ |
|-----------------|-------------|---------------|---------------|---------------|---------------|-------------|
| $\alpha = 0$ | 26.7% | | | | | |
| Market impact: | | | | | | |
| $\alpha = 0.2$ | 27.4% | 26.8% | 26.6% | 27.2% | 27.4% | 27.6% |
| $\alpha = 0.4$ | 27.3% | 27.0% | 28.0% | 28.4% | 27.1% | 28.3% |
| $\alpha = 0.6$ | 27.3% | 28.0% | 28.1% | 26.9% | 26.8% | 27.7% |
| $\alpha = 0.8$ | 26.4% | 26.9% | 28.0% | 27.0% | 27.5% | 28.0% |
| $\alpha = 1$ | 28.6% | 27.6% | 28.0% | 29.0% | 27.5% | 28.5% |
| Volatility: | | | | | | |
| $\alpha = 0.2$ | 28.1% | 28.2% | 29.0% | 28.1% | 29.3% | 28.8% |
| $\alpha = 0.4$ | 32.5% | 32.6% | 33.0% | 34.5% | 34.4% | 35.2% |
| $\alpha = 0.6$ | 37.1% | 36.7% | 36.5% | 39.3% | 38.3% | 38.3% |
| $\alpha = 0.8$ | 37.2% | 38.1% | 38.0% | 39.5% | 38.9% | 38.4% |
| $\alpha = 1$ | 38.8% | 39.3% | 38.2% | 38.5% | 39.2% | 38.6% |
| Client flow: | | | | | | |
| $\alpha = 0.2$ | 33.6% | 34.7% | 38.0% | 41.2% | 44.2% | 49.1% |
| $\alpha = 0.4$ | 41.1% | 43.8% | 48.3% | 52.8% | 59.2% | 67.1% |
| $\alpha = 0.6$ | 47.2% | 51.0% | 56.3% | 62.9% | 70.7% | 78.8% |
| $\alpha = 0.8$ | 51.9% | 57.5% | 64.5% | 71.3% | 78.8% | 85.7% |
| $\alpha = 1$ | 57.4% | 63.1% | 70.3% | 76.3% | 84.1% | 87.8% |
| All components: | | | | | | |
| $\alpha = 0.2$ | 35.4% | 37.2% | 38.8% | 42.4% | 45.3% | 50.4% |
| $\alpha = 0.4$ | 46.7% | 50.8% | 53.7% | 58.3% | 64.6% | 71.8% |
| $\alpha = 0.6$ | 56.1% | 59.9% | 65.2% | 71.5% | 78.6% | 87.6% |
| $\alpha = 0.8$ | 62.5% | 67.7% | 72.8% | 80.3% | 88.2% | 96.7% |
| $\alpha = 1$ | 66.9% | 73.1% | 79.5% | 87.0% | 94.2% | 100.0% |

the trade is reduced. In practice, minimum bid-ask spread is limited by a constant term. A consequence of adding this constant to the bid-ask spread is the appearance of absolute value terms in the transaction cost and therefore the optimization cost function. Although no longer quadratic, the cost

Table 3: Cost savings for different oracle accuracies in the real-world data experiment. The values show costs for risk = 10 bps, scaled as percentage between the single-stage strategy and prescient hedging results.

| Oracle | $\beta = 1$ | $\beta = 0.8$ | $\beta = 0.6$ | $\beta = 0.4$ | $\beta = 0.2$ | $\beta = 0$ |
|-----------------|-------------|---------------|---------------|---------------|---------------|-------------|
| $\alpha = 0$ | 5.2% | | | | | |
| Market impact: | | | | | | |
| $\alpha = 0.2$ | 5.4% | 5.5% | 5.5% | 5.6% | 5.4% | 5.4% |
| $\alpha = 0.4$ | 5.8% | 5.8% | 5.9% | 5.8% | 5.6% | 5.8% |
| $\alpha = 0.6$ | 6.1% | 6.0% | 5.9% | 6.1% | 6.1% | 6.0% |
| $\alpha = 0.8$ | 6.3% | 6.2% | 6.2% | 6.3% | 6.3% | 6.3% |
| $\alpha = 1$ | 6.4% | 6.4% | 6.4% | 6.4% | 6.4% | 6.4% |
| Volatility: | | | | | | |
| $\alpha = 0.2$ | 6.0% | 5.5% | 5.2% | 5.2% | 5.4% | 5.3% |
| $\alpha = 0.4$ | 15.5% | 15.6% | 15.7% | 16.1% | 17.5% | 20.3% |
| $\alpha = 0.6$ | 24.1% | 24.6% | 25.1% | 26.3% | 28.3% | 30.8% |
| $\alpha = 0.8$ | 27.6% | 28.0% | 29.0% | 29.9% | 31.6% | 32.8% |
| $\alpha = 1$ | 28.2% | 28.7% | 29.6% | 30.5% | 31.4% | 32.5% |
| Client flow: | | | | | | |
| $\alpha = 0.2$ | 13.6% | 15.9% | 18.0% | 21.5% | 25.6% | 31.2% |
| $\alpha = 0.4$ | 20.9% | 24.7% | 29.0% | 34.7% | 42.0% | 50.6% |
| $\alpha = 0.6$ | 27.7% | 32.6% | 38.6% | 45.5% | 54.5% | 64.8% |
| $\alpha = 0.8$ | 33.7% | 39.5% | 46.3% | 54.9% | 64.4% | 73.8% |
| $\alpha = 1$ | 39.0% | 45.7% | 53.2% | 62.0% | 71.1% | 78.2% |
| All components: | | | | | | |
| $\alpha = 0.2$ | 11.0% | 12.8% | 14.9% | 17.7% | 21.5% | 26.6% |
| $\alpha = 0.4$ | 31.4% | 34.7% | 38.9% | 44.6% | 52.7% | 64.0% |
| $\alpha = 0.6$ | 44.1% | 48.5% | 54.4% | 62.2% | 72.8% | 86.4% |
| $\alpha = 0.8$ | 50.4% | 55.9% | 62.9% | 72.1% | 84.0% | 96.8% |
| $\alpha = 1$ | 54.3% | 60.5% | 68.5% | 78.6% | 90.7% | 100% |

function remains convex, and can be solved using other numerical optimization techniques. This requires higher computational power and reduces the system's overall performance; therefore its use has to be justified in terms of cost-risk improvements versus speed reduction in backtesting.

In future work, we also intend to improve hedging by using better scenario generation approaches. Machine learning techniques will be employed to capture the intrinsic but hidden relationships between client flow, FX rate and transaction costs, thereby increasing the quality of scenarios compared to model based scenario generation or bootstrapping from historical observations.

8. Conclusion

In this paper, a novel strategy for hedging FX dealer’s short-term risk was presented. The approach utilizes stochastic receding horizon control (SRHC), with stochastic models to describe client flow, FX rate and transaction costs, and dynamic updates to realign misspecified models with real-world observations. Using a Monte Carlo approach for approximating the stochastic programming problem to quadratic programming resulted in a fast algorithm, optimizing and backtesting a daily trading session under 250 ms.

The hedging strategy was verified using both synthetic and real-world data. SRHC hedging was shown to significantly outperform benchmark hedging strategies, with up to 47.6% improvement in reducing costs for the same level of risk using real-world data. Additionally, the performance of the SRHC with different predictive accuracy of models was quantified, and it was shown that the proposed volatility and transactions cost models captured most of the underlying stochastic processes properties. The shortcomings of the client flow model were noted and solutions for improving its precision were discussed.

Declaration of Interest

This research was supported under Australian Research Council’s Linkage Projects funding scheme (project number LP110200413) and Westpac Banking Corporation.

Acknowledgment

We thank Dr. Giulio Katis, Westpac Institutional Bank, for providing insights regarding this work.

References

- Bemporad, A., Bellucci, L., and Gabriellini, T. (2014). Dynamic option hedging via stochastic model predictive control based on scenario simulation. *Quantitative Finance*, 14(10):1739–1751.
- Bemporad, A., Gabriellini, T., Puglia, L., and Bellucci, L. (2010). Scenario-based stochastic model predictive control for dynamic option hedging. In *IEEE Conference on Decision and Control (CDC)*, pages 6089–6094.
- Bemporad, A., Puglia, L., and Gabriellini, T. (2011). A stochastic model predictive control approach to dynamic option hedging with transaction costs. In *American Control Conference (ACC)*, pages 3862–3867.
- Bernardini, D. and Bemporad, A. (2012). Stabilizing model predictive control of stochastic constrained linear systems. *IEEE Transactions on Automatic Control*, 57(6):1468–1480.
- Borkovec, M., Domowitz, I., and Escobar, C. (2013). How big is “Big” - some evidence from aggregate trading costs in the FX market. Report, Investment Technology Group, Inc.
- Borrelli, F., Bemporad, A., and Morari, M. (2014). *Predictive Control*. Cambridge University Press.
- Calafiore, G. C. (2008). Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473.
- Calafiore, G. C. (2009). An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM Journal on Control and Optimization*, 48(4):2254–2274.
- Calafiore, G. C. (2010). Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464.
- Calafiore, G. C. and Kharaman, F. (2014). Multi-period asset allocation with lower partial moments criteria and affine policies. In *IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFER)*, pages 100–106.

- Cerrato, M., Sarantis, N., and Saunders, A. (2011). An investigation of customer order flow in the foreign exchange market. *Journal of Banking & Finance*, 35(8):1892–1906.
- Chen, S., Chien, C., and Chang, M. (2012). Order flow, bid-ask spread and trading density in foreign exchange markets. *Journal of Banking & Finance*, 36(2):597–612.
- diBartolomeo, D. (2012). Smarter rebalancing: Using single period optimization in a multi-period world. Research publications, Northfield Information Services.
- D’Souza, C. (2002). How do canadian banks that deal in foreign exchange hedge their exposure to risk? Working paper, Bank of Canada.
- Evans, K. P. (2011). Intraday jumps and US macroeconomic news announcements. *Journal of Banking & Finance*, 35(10):2511–2527.
- Evans, M. D. and Lyons, R. K. (2002). Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1):170–180.
- Gilli, M., Këllezi, E., and Hysi, H. (2006). A data-driven optimization heuristic for downside risk minimization. *Journal of Risk*, 8(3):1–19.
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*, volume 53 of *Stochastic Modelling and Applied Probability*. Springer.
- Gondzio, J., Kouwenberg, R., and Vorst, T. (2003). Hedging options under transaction costs and stochastic volatility. *Journal of Economic Dynamics and Control*, 27(6):1045–1068.
- Herzog, F., Keel, S., Dondi, G., Schumann, L., and Geering, H. P. (2006). Model predictive control for portfolio selection. In *American Control Conference (ACC)*, pages 1252–1259.
- Infanger, G. (2006). Dynamic asset allocation strategies using a stochastic dynamic programming approach. *Handbook of asset and liability management*, 1:199–251.
- Jabbour, C., Peña, J., Vera, J., and Zuluaga, L. (2007). An estimation-free, robust CVaR portfolio allocation model. *Journal of Risk*, 11(1):1–22.

- Josephy, N., Kimball, L., and Steblovskaya, V. (2013). Alternative hedging in a discrete-time incomplete market. *Journal of Risk*, 16(1):85–117.
- Lyons, R. K. (1998). Profits and position control: A week of FX dealing. *Journal of International Money and Finance*, 17(1):97–115.
- McGroartya, F., ap Gwilymb, O., and Thomasc, S. (2009). The role of private information in return volatility, bidask spreads and price levels in the foreign exchange market. *Journal of International Financial Markets, Institutions and Money*, 19(2):387–401.
- Meindl, P. J. and Primbs, J. A. (2008). Dynamic hedging of single and multi-dimensional options with transaction costs: A generalized utility maximization approach. *Quantitative Finance*, 8(3):299–312.
- Noorian, F. and Leong, P. H. W. (2014). Dynamic hedging of foreign exchange risk using stochastic model predictive control. In *IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr)*, pages 441–448.
- Primbs, J. A. (2007). Portfolio optimization applications of stochastic receding horizon control. In *American Control Conference (ACC)*, pages 1811–1816.
- Primbs, J. A. (2009). Dynamic hedging of basket options under proportional transaction costs using receding horizon control. *International Journal of Control*, 82(10):1841–1855.
- Primbs, J. A. (2010). LQR and receding horizon approaches to multi-dimensional option hedging under transaction costs. In *American Control Conference (ACC)*, pages 6891–6896.
- R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ranaldo, A. (2009). Segmentation and time-of-day patterns in foreign exchange markets. *Journal of Banking & Finance*, 33(12):2199–2206.
- Römisch, W. (2011). Scenario generation. In *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Inc.

- Schildbach, G., Fagiano, L., and Morari, M. (2013). Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4):2479–2501.
- Scholtus, M., van Dijk, D., and Frijns, B. (2014). Speed, algorithmic trading, and market quality around macroeconomic news announcements. *Journal of Banking & Finance*, 38:89–105.
- Skaf, J., Boyd, S., and Zeevi, A. (2010). Shrinking-horizon dynamic programming. *International Journal of Robust and Nonlinear Control*, 20(17):1993–2002.
- Topaloglou, N., Vladimirov, H., and Zenios, S. A. (2008). A dynamic stochastic programming model for international portfolio management. *European Journal of Operational Research*, 185(3):1501–1524.
- Ullrich, C. (2009). *Forecasting and Hedging in the Foreign Exchange Markets*, volume 623 of *Lecture Notes in Economics and Mathematical Systems*. Springer.
- Volosov, K., Mitra, G., Spagnolo, F., and Lucas, C. (2005). Treasury management model with foreign exchange exposure. *Computational Optimization and Applications*, 32(1-2):179–207.
- Weingessel, A. (2013). *quadprog: Functions to Solve Quadratic Programming Problems*. R package version 1.5-5.