Open

Original Article

A simple scheme for allocating capital in a foreign exchange proprietary trading firm

Received (in revised form): 22nd October 2014

Antony Jackson

is Assistant Professor in Financial Economics in the School of Economics at the University of East Anglia in the United Kingdom. He holds an MA in Economics from the University of Cambridge, an MSc in Management Science & Operational Research from Warwick Business School, and a PhD in Economics from the University of Leicester. Jackson has an investment banking background, having been Vice President in Credit Risk for Credit Suisse First Boston in Tokyo. He holds PRIMIA's Professional Risk Manager (PRM) designation, and is a Regular Member of the CFA Institute.

Correspondence: Department of Economics, University of East Anglia, Norwich, NR4 7TJ, UK.
E-mail: antony.jackson@uea.ac.uk

ABSTRACT We present a model of capital allocation in a foreign exchange proprietary trading firm. The owner allocates capital to individual traders, who operate within strict risk limits. Traders specialize in individual currencies, but are given discretion over their choice of trading rule. The owner provides the simple formula that determines position sizes – a formula that does not require estimation of the firm-level covariance matrix. We provide supporting empirical evidence of excess risk-adjusted returns to the firm-level portfolio, and we discuss a modification of the model in which the owner dictates the choice of trading rule.
Journal of Asset Management advance online publication, 18 December 2014; doi:10.1057/jam.2014.40

Keywords: foreign exchange trading; technical trading rules; portfolio management

The online version of this article is available Open Access

INTRODUCTION

We present a model in which a foreign exchange trading firm owner shares capital among a group of traders. His objective is to earn excess risk-adjusted returns to the firm-level portfolio, but under the constraint that his employees trade as individuals. They specialize in individual currencies, concentrating solely on the exchange rate between their designated currency and the US dollar.

We contribute to the literature on risk-adjusted performance measurement (RAPM) in two distinct ways. First, we derive a simple plug-in formula for capital allocation that may be adapted to different risk measures[1] and, second, we highlight the impact organizational constraints may have on the allocation problem.

A full mean–variance optimization exercise (based on the firm-level covariance matrix) would likely generate a volatile capital allocation scheme. This is due in part to the high sensitivity of the optimal weights to changes in the problem set-up brought about by time-varying sample estimates of variances and covariances, but also due to estimation error. A literature has developed that deals with the problem – the *shrinkage* literature.[2] We propose that the owner adopts a simple risk budgeting scheme that is based on the conditional volatilities of individual currencies, but not on the correlations between them.

We discuss two versions of the model: the 'discretionary' and the 'automated' model. In the former, each trader is tasked with trading a single currency, but is given discretion over his choice of trading rule. In the latter, owners dictate both the risk limits and the trading rule choices. In both versions, the owner provides the formula that determines position sizes. The key distinction between the models lies in the way information is processed. The discretionary model is an 'as-if' model, in which optimal rules are chosen with the benefit of hindsight – accordingly, a higher statistical threshold is put in place if the results are to be deemed significant. In the automated model, the owner's choice of rule is adapted to the information actually available in real time. His choice of rule is based on an ever-expanding information set, and can be seen as an exercise in statistical learning.

The remainder of this article is set out as follows. In the section 'Position-sizing formula', we develop the position-sizing formula for the case of a single trader. In the case of multiple traders, we introduce a simple method of risk budgeting that is based on an equitable distribution of capital. We assume that risk budgets are binding, and that variations in position sizes are driven primarily by time variation in conditional volatilities. The section 'Expectations and conditional volatility' introduces the owner's forward-looking risk-adjusted performance target – the target Sharpe ratio acts as the transmission mechanism from conditional volatilities to position sizes. The 'Results' section provides supporting empirical evidence of the ability of both versions of our model to deliver statistically and economically significant excess returns to the firm-level portfolio. The final section concludes, and discusses concurrent research in which we allow traders to use their full risk allocations selectively.

POSITION-SIZING FORMULA

We begin by solving the owner's capital allocation problem for the case of one trader. In this case, the solution provides an optimal level of leverage. The model follows Campbell and Viceira (2002), modified to the choice between a foreign currency and a risk-free domestic asset.

The owner exhibits constant relative risk aversion (CRRA), with power utility defined over next-period wealth:

$$U(W_{t+1}) = \frac{W_{t+1}^{1-\gamma}}{1-\gamma}, \qquad (1)$$

where W is wealth and γ is the coefficient of relative risk aversion.

We assume portfolio returns are log-normally distributed, which implies that next-period wealth is also log-normally distributed. In conjunction with CRRA preferences, the assumption of log-normally distributed portfolio returns leads to a closed-form solution for capital allocation to the risky asset, equation (18), which is increasing in expected log returns and decreasing in the variance of log returns. Our specification has the convenient property that the percentage of capital allocated to the risky asset is independent of the owner's current level of wealth.

The owner's objective is to maximize the expected utility of next-period wealth:

$$\max \mathbb{E}\left(\frac{W_{t+1}^{1-\gamma}}{1-\gamma}\right)$$

subject to the budget constraint

$$W_{t+1} = \left(1+R_{p,t+1}\right) W_t, \qquad (2)$$

where $R_{p,t+1}$ is the simple portfolio return. We detail the algebraic steps necessary to restate this problem equivalently in terms of $\log(W_t)$, rather than W_t. The idea is to modify the problem set-up in order to take advantage of the log-normality of returns assumption.

Observe first that, because $(1-\gamma)$ is a constant, maximizing $\mathbb{E}[W_{t+1}^{1-\gamma}/(1-\gamma)]$ is the same as maximizing $\mathbb{E}[W_{t+1}^{1-\gamma}]$. Furthermore, applying the natural log function results in a monotonic transformation, thus maximizing $\mathbb{E}[W_{t+1}^{1-\gamma}]$ is the same as maximizing $\log(\mathbb{E}[W_{t+1}^{1-\gamma}])$.

Our assumption of log-normally distributed wealth now enables convenient manipulation of the awkward-looking expression $\log(\mathbb{E}[W_{t+1}^{1-\gamma}])$. We make use of the standard result of the log-normal density function that, if $\ln(W_{t+1})$ is normally distributed, then

$$\mathbb{E}[W_{t+1}] = \exp\left(\mathbb{E}[\log(W_{t+1})] + \frac{1}{2}\mathbb{V}[\log(W_{t+1})]\right), \quad (3)$$

where $\mathbb{E}$ and $\mathbb{V}$ are the expectation and variance operators, respectively. Taking natural logs of both sides of equation (3) yields

$$\log(\mathbb{E}[W_{t+1}]) = \mathbb{E}[\log(W_{t+1})] + \frac{1}{2}\mathbb{V}[\log(W_{t+1})], \quad (4)$$

from which it follows that

$$\log\left(\mathbb{E}\left[W_{t+1}^{1-\gamma}\right]\right) = \mathbb{E}\left[\log\left(W_{t+1}^{1-\gamma}\right)\right] + \frac{1}{2}\mathbb{V}\left[\log\left(W_{t+1}^{1-\gamma}\right)\right]. \quad (5)$$

By using standard rules for manipulating expectations, variances and natural logs, the expression in equation (5) simplifies to

$$(1-\gamma)\log(\mathbb{E}[W_{t+1}]) = (1-\gamma)\mathbb{E}[\log(W_{t+1})] + \frac{1}{2}(1-\gamma)^2\mathbb{V}[\log(W_{t+1})], \quad (6)$$

and after dividing equation (6) throughout by $(1-\gamma)$, we are left with the useful expression

$$\log(\mathbb{E}[W_{t+1}]) = \mathbb{E}[\log(W_{t+1})] + \frac{1}{2}(1-\gamma)\mathbb{V}[\log(W_{t+1})]. \quad (7)$$

Applying natural logs to the budget equation, that is equation (2), yields

$$\log(W_{t+1}) = \log\left(1+R_{p,t+1}\right) + \log(W_t). \quad (8)$$

Substituting for $\log(W_{t+1})$ from equation (8) in equation (7) yields the objective function

$$\log(\mathbb{E}[W_{t+1}]) = \mathbb{E}\left[\log\left(1+R_{p,t+1}\right)\right] + \mathbb{E}[\log(W_t)] + \frac{1}{2}(1-\gamma)\mathbb{V}\left[\log\left(R_{p,t+1}\right)\right]. \quad (9)$$

Of course, $\mathbb{E}[\log(W_t)]$ is the same as log (W_t), because W_t is known at time t. Constants do not affect the maximization exercise, thus maximizing $\log(\mathbb{E}[W_{t+1}])$ is equivalent to the following maximization problem:

$$\max\ \mathbb{E}\left[\log\left(1+R_{p,t+1}\right)\right] + \frac{1}{2}(1-\gamma)\mathbb{V}\left[\log\left(R_{p,t+1}\right)\right]. \quad (10)$$

After introducing the lower-case notation $r_{p,t+1} \equiv \log(1+R_{p,t+1})$, we arrive at the much simpler-looking maximization objective:

$$\max \mathbb{E}\left[r_{p,t+1}\right] + \frac{1}{2}(1-\gamma)\mathbb{V}\left[r_{p,t+1}\right]. \quad (11)$$

The maximization problem is now stated in terms of log portfolio returns, which are a non-linear combination of the individual assets in the portfolio. We follow the Campbell and Viceira (2002) linear approximation of these returns, adapted to the special case of foreign exchange. The equation of the simple return R_{t+1} to the foreign currency is

$$1+R_{t+1} = \left(1+R^{\star}_{0,t+1}\right)\left(\frac{S_{t+1}}{S_t}\right), \quad (12)$$

where $R^{\star}_{0,t+1}$ is the foreign risk-free interest rate and S_t is the exchange rate expressed as domestic currency units per unit of foreign currency. Taking natural logarithms throughout equation (12) gives

$$\log(1+R_{t+1}) = \log\left(1+R^{\star}_{0,t+1}\right) + \log(S_{t+1}) - \log(S_t). \quad (13)$$

The trader allocates α_t of his wealth to the foreign currency and $1-\alpha_t$ to the domestic risk-free asset, giving an equation in terms of the portfolio simple return $R_{p,t+1}$ of

$$\begin{aligned}1+R_{p,t+1} &= 1+\alpha_t R_{t+1}+(1-\alpha_t)R_{0,t+1}\\ &= 1+R_{0,t+1}+\alpha_t\left(R_{t+1}-R_{0,t+1}\right),\end{aligned}$$

where $R_{0,t+1}$ denotes the rate of interest on the domestic risk-free asset. This equation can be rearranged to give

$$\frac{1+R_{p,t+1}}{1+R_{0,t+1}} = 1+\alpha_t\left(\frac{1+R_{t+1}}{1-R_{0,t+1}}-1\right). \quad (14)$$

Substituting for $1+R_{t+1}$ from equation (12) in equation (14), and taking natural logarithms throughout, yields

$$\begin{aligned}r_{p,t+1}-r_{0,t+1} &= \log\{1+\alpha_t[\exp(r^{\star}_{0,t+1}-r_{0,t+1}\\ &\quad +s_{t+1}-s_t)-1]\}. \quad (15)\end{aligned}$$

Here we have introduced the lower-case notation $r_{p,t+1}\equiv\log(1+R_{p,t+1})$, $r_{0,t+1}\equiv\log(1+R_{0,t+1})$, $r^{\star}_{p,t+1}\equiv\log(1+R^{\star}_{0,t+1})$ and $s_t\equiv\log(S_t)$. Equation (15) shows that portfolio excess returns are a function of the log interest rate differential, $r_{0,t+1}-r^{\star}_{0,t+1}$, and the log return on the exchange rate, $s_{t+1}-s_t$. Let

$$\begin{aligned}&f\left(r^{\star}_{0,t+1}-r_{0,t+1},s_{t+1}-s_t\right)\\ &\quad=\log\left\{1+\alpha_t\left[\exp\left(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t\right)-1\right]\right\}. \quad (16)\end{aligned}$$

The second-order Taylor expansion of the function f around the point $r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t=0$ yields

$$\begin{aligned}&f\left(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t\right)\\ &\quad\approx f(0)+f'(0)\left(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t\right)\\ &\qquad+\frac{1}{2}f''(0)\left(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t\right)^2,\end{aligned}$$

with $f'(0)=\alpha_t$ and $f''(0)=\alpha_t(1-\alpha_t)$. After replacing $(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t)^2$ with its conditional expectation σ^2_t, the linear approximation for excess log returns can be written as

$$\begin{aligned}r_{p,t+1}-r_{0,t+1} &= \alpha_t\left(r^{\star}_{0,t+1}+-r_{0,t+1}+s_{t+1}-s_t\right)\\ &\quad+\frac{1}{2}\alpha_t(1-\alpha_t)\sigma^2_t. \quad (17)\end{aligned}$$

The final stage is to substitute equation (17) and the corresponding variance of portfolio log returns, $\alpha^2_t\sigma^2_t$, into the objective function, that is equation (11):

$$\begin{aligned}&\max\alpha_t\mathbb{E}_t\left(r^{\star}_{0,t+1}-r_{0,t+1}+s_{t+1}-s_t\right)\\ &\quad+\frac{1}{2}\alpha_t(1-\alpha_t)\sigma^2_t+\frac{1}{2}(1-\gamma)\alpha^2_t\sigma^2_t.\end{aligned}$$

The first-order condition yields an optimal allocation of wealth to the foreign currency of

$$\alpha_t=\frac{\mathbb{E}_t(s_{t+1}-s_t)+r^{\star}_{0,t+1}-r_{0,t+1}+\sigma^2_t/2}{\gamma\sigma^2_t}. \quad (18)$$

The proportion of wealth allocated to the foreign currency is increasing in the expected appreciation of the foreign currency and increasing in the differential between the foreign interest rate and the domestic risk-free rate. The proportion is decreasing in conditional volatility and decreasing in the owner's coefficient of relative risk aversion.

Portfolio construction

Of course, the owner's allocation problem is complicated by there being multiple currencies and by each currency being assigned to an individual trader. Although we do not discuss the principal-agent problem in this article, employees in proprietary firms commonly have individual contracts that detail percentage profit splits with the owner, who in return offers substantially reduced commissions that reflect the large trading volumes the firm places with third parties. That traders act as individuals – rather than as part of a team – restricts the ability of the owner to allocate capital in accordance with a full mean-variance optimization scheme. The allocations would likely be too volatile, with certain traders being allocated large shares of capital purely on the basis of pairwise correlations between currencies. We suggest that the owner is more likely to adopt an equitable scheme, with departures reflecting the discipline of traders in staying within their risk limits, or by their ability to choose 'good' trading rules. As an abstraction, we adopt the purely equitable approach and leave for future research the interesting questions of how to measure and reward individual performance.

We suggest that the owner achieves an equitable allocation simply by multiplying his coefficient of relative risk aversion by the number of traders in the firm. Diversification benefits to the firm-level portfolio ensure that portfolio risk lies within the owner's original risk tolerance. An equitable allocation can be thought of as a risk-adjusted naive diversification[3] strategy without short-sale constraints. A justification for this approach could be to avoid the problem of distinguishing current profitability due to chance from that of trading skill. Logistically, the owner then presents a simple position-sizing formula to each trader, with variations in position size reflecting changes in conditional volatility and the level of the firm's capital:

$$\alpha_t = \frac{1}{N}\frac{\mathbb{E}_t(s_{t+1} - s_t) + r^{\star}_{0,t+1} - r_{0,t+1} + \sigma_t^2/2}{\gamma\sigma_t^2}, \tag{19}$$

where N is the number of traders. It now remains to discuss how the owner calculates conditional volatility and how traders use trading rules in forming expectations.

EXPECTATIONS AND CONDITIONAL VOLATILITY

We now introduce the notion of a 'target Sharpe Ratio'. The idea offers a practical solution to the problem of mapping binary signals into expectations, but also offers a practical insight into the way traders use simple rules-of-thumb in their decision-making. The method is grounded in several references in the literature to firms' use of threshold levels of risk-adjusted profitability. Lyons (2001) provides anecdotal evidence that foreign exchange trading firms only allocate capital to those strategies expected to yield annualized Sharpe ratios in the range 0.5–1.0; Grinold and Kahn (2000) and Menkhoff and Taylor (2007) suggest that 0.5 is a common benchmark used for identifying 'good' trading rules.

We propose, uncontroversially, that the owner expects to earn a risk premium as compensation for being exposed to exchange rate risk. Owners adopt a target Sharpe ratio, which when rearranged and augmented by a trading rule signal gives an expression for the expected appreciation of each foreign currency:

$$\mathbb{E}_t(s_{t+1} - s_t) = I_t \frac{\mathrm{SR}_{\mathrm{target}} \times \sigma_t}{\sqrt{250}}. \tag{20}$$

Here $\mathrm{SR}_{\mathrm{target}}$ is an annualized measure of the Sharpe ratio, $I_t \in \{-1, 1\}$ is a binary signal, and we assume that there are 250 trading days in a year. Owners use an exponentially weighted moving average (EWMA) estimate of volatility. The advantage of the method – which is well established in the risk management industry – is that it can be used to produce estimates extremely quickly. The EWMA estimator is defined by

$$\hat{\sigma}_t^2 = (1-\lambda)\mu_{t-1}^2 + \lambda\hat{\sigma}_{t-1}^2, \tag{21}$$

where λ is a smoothing parameter and μ_{t-1} is last period's return. A smoothing parameter of 0.94 is generally regarded as appropriate for daily observations (Alexander and Sheedy, 2010).

Substituting equations (20) and (21) into the capital allocation equation, that is equation (19), yields the bottom-line position-sizing formula:

$$\alpha_t = \frac{1}{N\gamma}\left(\frac{I_t \mathrm{SR}_{\mathrm{target}}}{\sqrt{250}\hat{\sigma}_t} + \frac{r^{\star}_{0,t+1} - r_{0,t+1}}{\hat{\sigma}_t^2} + \frac{1}{2}\right). \tag{22}$$

Let us examine the sensitivity of the position size to each variable in equation (22).

First, we use a simple risk-budgeting scheme (rather than full mean-variance optimization). To allocate $1/N$ of optimal position sizes to each currency is to be conservative; likely there would be benefits to diversification at the portfolio level that would allow larger position sizes in each currency. Second, position sizes are inversely related to the risk aversion, γ, of the owner. Now, examining the terms in the parentheses, a higher target Sharpe ratio implies higher expected returns, and hence larger position sizes. The target Sharpe ratio – or the expected market price of risk – is the mechanism that maps volatility into expected returns. On balance, however, higher estimates of conditional volatility are accompanied by smaller position sizes. Even though the target Sharpe ratio heuristic maps higher volatilities into greater absolute expected exchange rate movements, smaller position sizes result, as the risk term dominates. Larger position sizes are taken when trading rule signals act in the same direction as the interest rate differential – the 'carry trade' effect.

Trading rules

In the 'discretionary' version of our model, traders are free to choose the trading rule that generates signals in their particular currency. In the 'automated' version, the owner dictates the choice of trading rule. We now describe the choice set of trading rules.

We include four types of rule – designed to broadly follow those of Qi and Wu (2006), who in turn apply the stock index rules of Sullivan *et al* (1999). With the exception of the 'filter' rule, the variable of interest is the number of days of sample data. There are four types of trading rule, each having 250 possible parameter values, giving a total trading rule universe of 1000 rules. The 'momentum', 'moving average' and 'trading range break' rules are trend-following rules, whereas the 'filter' is a contrarian rule. Each rule is described below.

Momentum

The momentum indicator signals whether the rate of change of the exchange rate has been positive or negative over a historical time period $n \in \{1, 2, 3, \ldots, 250\}$:

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > S_{t-n} \\ 0 & \text{if } S_t = S_{t-n} \\ -1 & \text{if } S_t < S_{t-n} \end{cases}.$$

Moving average

The moving average indicator offers a slightly more complicated version of the momentum rule; it includes all sample points in its calculation:

$$\mathrm{SMA}_t = \sum_{t=t-n}^{t} \frac{S_t}{n+1}.$$

The indicator function for the simple moving average rule is

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > \mathrm{SMA}_t \\ 0 & \text{if } S_t = \mathrm{SMA}_t \\ -1 & \text{if } S_t < \mathrm{SMA}_t \end{cases}.$$

Trading range break

The n-day breakout indicator generates a positive signal if today's close is greater than the highest high of the previous n prices:

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > \max(S_{t-1},\ S_{t-2}, \ldots, S_{t-n}) \\ -1 & \text{if } S_t < \min(S_{t-1},\ S_{t-2}, \ldots, S_{t-n}) \\ I_{t-1} & \text{otherwise} \end{cases}.$$

The indicator is zero until the first breakout, and maintains this value until a breakout occurs in the opposite direction.

Filter

The filter rule focuses on recent highs and lows. Consider a falling market. The low is reset at subsequent lower lows until an n-per cent rise generates a buy signal. We consider $n \in \{1.000$ per cent, 1.006 per cent, 1.012 per cent, $\ldots$, 2.494 per cent$\}$, designed to capture 250 rules in a range consistent with the previous literature. Sell signals in a rising market are generated similarly.

The trading rule signals provide the final piece of information required by the position-sizing formula, that is equation (22). We provide empirical evidence that the owner is able to allocate capital equitably, while still generating firm-level excess risk-adjusted returns. Our exchange rates are drawn from the Federal Reserve Board's H.10 series, and interest rates are British Bankers Association 3-month LIBOR rates. They cover the period from 4 January 1999 to 20 January 2012, and comprise six major foreign currencies: the Australian dollar, British pound, Canadian dollar, Euro, Japanese yen and Swiss franc. We now compare the returns to the 'discretionary' and 'automated' versions of our model.

RESULTS

In the discretionary model, the owner dictates the position-sizing formula, but allows traders discretion in their choice of trading rule. This version offers the owner diversification across methods, as well as diversification across currencies. The discretionary model raises the question of how traders choose their trading rule. Although this is an interesting question in itself, we sidestep the modelling problem by allowing traders to choose the optimal in-sample rule for their particular currency. Clearly this is not achievable in reality, and thus any statistical inference drawn from the exercise must take into account the so-called 'data snooping' problem (Lo and MacKinlay, 1990).

The Reality Check (White, 2000) tests whether the best rule beats the null hypothesis of zero excess profitability. The idea is that the researcher can search aggressively across a wide variety of rules, safe in the knowledge that the distribution of the test statistic under the null hypothesis adjusts to compensate for the increased chance of achieving 'lucky' results across many searches. The time-series bootstrap (Politis and Romano, 1994) generates pseudo-time series of returns by sampling blocks of observations from the empirical series, where the size of each block is drawn from a geometric distribution with mean size q. The size of the block is an increasing function of the dependency evident in the empirical data – we use a conservative block size of $q=10$, as in Sullivan *et al* (1999).

The following iterative procedure (White, 2000) obtains the P-value for the best model. Starting with the first model, and $B=1000$ bootstrap replicate series, the test statistic $\overline{V}_1$ is defined as

$$\overline{V}_1 = n^{\frac{1}{2}}\overline{R}_1,$$

where $\overline{R}_1$ denotes the mean excess return of the first model and n is the number of returns. For each of the $B=1000$ bootstrap replicate series, one calculates the statistic

$$\overline{V}^{\star}_{1,i} = n^{\frac{1}{2}}\left(\overline{R}^{\star}_{1,i} - \overline{R}_1\right), \quad i = 1, \ldots, 1,000 \tag{23}$$

where the superscript '$\star$' identifies simulated series. The P-value of the first rule is obtained by comparing $\overline{V}_1$ to the percentiles of $\overline{V}^{\star}_{1,i}$. One then proceeds to examine the second trading rule. Compute

$$\overline{V}_2 = \max\left\{n^{\frac{1}{2}}\overline{R}_2, \overline{V}_1\right\} \tag{24}$$

and

$$\overline{V}^{\star}_{2,i} = \max\left\{n^{\frac{1}{2}}\left(\overline{R}^{\star}_{2,i} - \overline{R}_2\right), \overline{V}^{\star}_{1,i}\right\}, \quad i = 1, \ldots, 1000 \tag{25}$$

noting that equation (25) uses the same replicate series as in equation (23).

One proceeds recursively through the remaining k models, obtaining

$$\overline{V}_k = \max\left\{n^{\frac{1}{2}}\overline{R}_k, \overline{V}_{k-1}\right\}$$

and

$$\overline{V}^{\star}_{k,i} = \max\left\{n^{\frac{1}{2}}\left(\overline{R}^{\star}_{2,i} - \overline{R}_2\right), \overline{V}^{\star}_{k-1,i}\right\}, \quad i = 1, \ldots, 1000.$$

The P-value of the optimal rule is obtained by comparing $\overline{V}_k$ with the percentiles of $\overline{V}^{\star}_{k,i}$.

Table 1 presents the results for discretionary traders. Panel A presents, for comparison, the results for the time series of returns conditioned on the trading rule signals of the optimal rule. To enable meaningful comparisons between absolute levels of excess return, we calibrate the coefficient of relative risk aversion, γ, to a restricted version of equation (22), in which the conditional volatility estimate is fixed and the trader ignores the interest rate differential:

$$\alpha_t = \frac{1}{\gamma}\left(\frac{I_t \mathrm{SR}_{\mathrm{target}}}{\sqrt{250}\tilde{\sigma}_t}\right) \equiv k,$$

where k is a constant. Now γ merely acts as a leverage parameter – the same percentage of capital is invested or borrowed for all signals. Reality Check P-values and Sharpe ratios are unaffected by the particular value of γ chosen, but the calibrated value of $\gamma = 3.7$ ensures that, on average, position sizes are equal with and without the owner's position-sizing formula. An equivalent exercise – and the one followed by most studies of technical trading rules – is to condition the time series of returns by a sequence of 1s and −1s (corresponding to long and short positions), and to then analyse the conditioned time series of returns. Our restricted version of equation (22) merely changes the sequence of conditioning variables to $+/-k$, where k is chosen to generate position sizes that are, on average, equal to those generated by the owner's formula.

Panel B presents the results for discretionary traders using the optimal rule in combination with the owner's position-sizing formula. This is an 'as-if' analysis, where we study the situation in which each trader chooses the single in-sample rule that is optimal for their currency. The improvement in Reality Check P-values – as traders actively manage their position sizes – is evident across all currencies. The economic significance of the results – the Sharpe ratio – increases markedly once the traders actively manage their position sizes. However, the results in Panel B present an interesting dilemma to the owner. After adjusting for data snooping bias, half of the traders appear to generate excess returns. In practice, the owner may have to balance the competing claims of the star

Table 1: Table presenting the excess returns, Reality Check P-values and Sharpe ratios of individual traders in the 'discretionary' model

	Best trading rule	*Round-trip transaction costs*					
		Excess return (annualized)		*Reality check (P-value)*		*Sharpe (annualized)*	
		0.00%	0.05%	0.00%	0.05%	0.00%	0.05%
Panel A: Without Position Sizing							
Australian dollar	103-day momentum	9.4	9.0	0.224	0.271	0.66	0.63
Canadian dollar	113-day momentum	6.2	5.9	0.252	0.327	0.63	0.60
Swiss franc	22-day breakout	5.4	5.1	0.619	0.687	0.47	0.45
Euro	26-day momentum	9.8	9.0	0.046*	0.077	0.93	0.86
British pound	101-day breakout	6.7	6.7	0.217	0.235	0.68	0.67
Japanese yen	176-day breakout	5.5	5.5	0.514	0.531	0.54	0.54
Panel B: With Position Sizing							
Australian dollar	1.7% filter	21.0	19.6	0.009**	0.023*	1.08	1.00
Canadian dollar	113-day momentum	12.0	11.0	0.134	0.226	0.73	0.67
Swiss franc	108-day momentum	8.7	7.9	0.452	0.558	0.52	0.46
Euro	26-day momentum	17.2	15.6	0.021*	0.044*	1.04	0.95
British pound	101-day breakout	15.9	15.3	0.046*	0.069	0.87	0.84
Japanese yen	176-day breakout	15.2	14.7	0.100	0.131	0.73	0.70

All simulations use a target Sharpe ratio of 0.5 and deduct 0.05 per cent proportional round-trip transaction costs, as in Qi and Wu (2006).
Levels of significance are * = 5 per cent and ** = 1 per cent.

Table 2: Table presenting excess returns, Reality Check *P*-values and Sharpe ratios for the firm-level portfolio in the 'automated' model

	Best trading rule	*Excess return (annualized)*	*Reality check (P-value)*	*Sharpe (annualized)*
Panel A: Without Position Sizing				
0.00% round-trip costs	108-day momentum	4.4	0.193	0.60
0.05% round-trip costs	108-day momentum	4.0	0.287	0.54
Panel B: With Position Sizing				
0.00% round-trip costs	108-day momentum	11.1	0.013*	0.99
0.05% round-trip costs	108-day momentum	10.0	0.024*	0.89

All simulations use a target Sharpe ratio of 0.5 and deduct 0.05% proportional round-trip transaction costs, as in Qi and Wu (2006).
Levels of significance are * = 5 per cent and ** = 1 per cent.

traders (those delegated with the responsibility of trading the Australian dollar, Euro and British pound) with those of the under-performers – traders who still offer value to the owner in terms of method and currency diversification. This principal/agent problem is one we are actively researching, but from which we abstract in the current article.

The corresponding performance of the firm-level portfolio is described in Table 2. The value to the owner in imposing the position-sizing formula is striking. Reality Check *P*-values for constant position sizing are 0.193 with transaction costs, and 0.287 without transaction costs. The Sharpe ratio of the best rule – the 108-day momentum rule – is 0.60, reducing to 0.54 when transaction costs are properly taken into account. In stark contrast, the returns to the firm-level portfolio when traders are charged with using the owner's position-sizing formula are both statistically and economically significant. The 108-day momentum rule is still the optimal in-sample rule, but the Reality Check *P*-value is now 0.024, with a corresponding Sharpe ratio of 0.89. That the Sharpe ratio has also increased provides evidence that the results have economic, as well as statistical, merit.

Table 3 examines the impact on individual currency traders of being forced to move away from the currency-specific optimal trading rule to the portfolio optimal trading rule, the 108-day momentum strategy. It can be seen that, with the exception of the Japanese yen, which has been forced into losses, individual currency trading remains profitable. In Panel B, one observes consistent improvement in returns once the position-sizing formula is applied. Interestingly, position sizing turns the Japanese Yen returns into profit, suggesting that economic benefit is being derived from estimating time-varying volatility in accordance with the EWMA estimator.

In the 'automated' model, the owner assumes control of the trading rule choice. We assume the owner instructs every trader to use the same trading rule – the best-performing rule on the available historical data. At the beginning of the sample – where there is little historical information – the trading rule choice is volatile. But eventually the owner's choice converges to the 108-day momentum rule that was found to be the best in-sample rule in Table 2.

Table 4 presents the results for the firm-level portfolio in the 'automated' model. Naturally, mean returns are lower than for the in-sample 'discretionary' model: the best net returns in Table 2 were 10.0 per cent, whereas they have decreased to 7.5 per cent in Table 4. Nevertheless, the *P*-value has decreased further to 0.021. The reduction in mean returns has been more than offset by the elimination of data-snooping bias – the owner only uses historical information in the 'automated' model. The qualitative conclusion of the

Table 3: Table presenting the excess returns and Sharpe ratios of individual traders when all traders use the 108-day momentum rule

	Best trading rule	Excess return (annualized)	Sharpe (annualized)	108-day momentum Excess return (annualized)	108-day momentum Sharpe (annualized)
Panel A: Without Position Sizing					
Australian dollar	103-day momentum	9.0	0.63	8.3	0.58
Canadian dollar	113-day momentum	5.9	0.60	4.1	0.42
Swiss franc	22-day breakout	5.1	0.45	4.4	0.38
Euro	26-day momentum	9.0	0.86	4.2	0.40
British pound	101-day breakout	6.7	0.67	4.6	0.46
Japanese yen	176-day breakout	5.5	0.54	–1.7	–0.17
Panel B: With Position Sizing					
Australian dollar	1.7% filter	19.6	1.00	17.8	0.87
Canadian dollar	113-day momentum	11.0	0.67	8.9	0.54
Swiss franc	108-day momentum	7.9	0.46	7.9	0.46
Euro	26-day momentum	15.6	0.95	9.3	0.55
British pound	101-day breakout	15.3	0.84	11.4	0.64
Japanese yen	176-day breakout	14.7	0.70	4.6	0.21

All simulations use a target Sharpe ratio of 0.5 and deduct 0.05% proportional round-trip transaction costs, as in Qi and Wu (2006).

Table 4: Table presenting excess returns and Sharpe ratios for the firm-level portfolio in the 'automated' model

	Excess return (annualized)	Time-series bootstrap (P-value)	Sharpe (annualized)
Panel A: Without Position Sizing			
0.00% round-trip costs	1.5	0.240	0.21
0.05% round-trip costs	0.9	0.344	0.12
Panel B: With Position Sizing			
0.00% round-trip costs	8.0	0.020*	0.70
0.05% round-trip costs	7.5	0.021*	0.66

Levels of significance are * = 5 per cent and ** = 1 per cent.

previous section remains – the position-sizing formula is crucial in generating excess returns to the firm-level portfolio.

Sensitivity analysis

The owner's coefficient of relative risk aversion and his choice of target Sharpe ratio are leverage parameters that do not affect the results, other than by scaling the mean returns. Variations in capital allocation are driven by time-varying estimates of volatility. The degree to which the EWMA estimator (equation (21)) discounts past information is determined by the smoothing parameter λ. Setting $\lambda = 1.00$ is equivalent to using a constant volatility estimate, which in combination with a target Sharpe ratio, implies expectations of a constant risk premium. It is evident from the first row of Table 5 that the worst results follow from an assumption of constant volatility. If, however, the trader uses a value of λ in the region of 0.94 – as suggested by J.P. Morgan/Reuters (1996) – then the results are robust to variations around this level. The results are more forgiving of a trader who errs on the side of placing greater weight on recent returns, than of one who errs on the side of treating volatility as constant. The implication is that the owner should be confident using the J.P. Morgan/Reuters (1996) parameter.

CONCLUSION

We consider an organizational structure in which a trading firm owner shares capital among a group of traders, each of whom trades in isolation from his colleagues. We propose an equitable risk-budgeting scheme, a key strength of which is its computational simplicity. The allocation scheme offers

Table 5: Table demonstrating the importance of the EWMA estimator

	EWMA parameter (λ)	Excess return (annualized)	P-value	Sharpe (annualized)
Panel A: 'Discretionary' Model				
Constant Volatility	1.00	4.8	0.147	0.65
	0.98	3.7	0.037*	0.82
	0.96	9.7	0.022*	0.87
RiskMetrics™	0.94	10.0	0.024*	0.89
	0.92	10.1	0.026*	0.89
	0.90	10.1	0.028*	0.88
	0.88	10.0	0.030*	0.86
	0.86	9.9	0.035*	0.84
Panel B: 'Automated' Model				
Constant Volatility	1.00	3.2	0.069	0.41
	0.98	6.4	0.030*	0.58
	0.96	7.2	0.021*	0.64
RiskMetrics™	0.94	7.5	0.021*	0.66
	0.92	7.6	0.021*	0.66
	0.90	7.7	0.021*	0.66
	0.88	7.7	0.022*	0.65
	0.86	7.7	0.023*	0.64

The smoothing parameter $\lambda = 1.00$ is equivalent to a constant volatility assumption; $\lambda = 0.94$ is the level recommended by J.P. Morgan/Reuters (1996). All simulations use a target Sharpe ratio of 0.5 and deduct 0.05% proportional round-trip transaction costs, as in Qi and Wu (2006).
Levels of significance are * = 5 per cent and ** = 1 per cent.

some stability, with variations in position sizes resulting from changes in conditional volatility, rather than from traders' past performance. We show that this abstraction of the owner's capital allocation problem generates statistically and economically significant excess returns to the firm-level portfolio in both the 'discretionary' and 'automated' versions of the model.

We are actively researching extensions to the model that reflect the finer organizational details of actual proprietary trading firms. One departure from our abstract model is in the way owners compensate traders – traders tend to be self-employed, with formal profit-sharing contracts. Successful traders see their capital accounts grow, whereas unsuccessful traders see their accounts shrink, creating interesting intra-firm capital dynamics.

A second extension lies in the way traders use their risk limits. In this article, we assume risk limits are binding, as the owner dictates the position-sizing formula. A relaxation of this assumption creates an incentive for traders exhibiting the traits associated with Prospect Theory (Kahneman and Tversky, 1979) to trade within their limits when in the domain of profits, and only at their limits when in the domain of losses. Should the owner wish to adopt our equitable structure, his contracts will need to be structured accordingly.

NOTES

1. For a survey of the problem approached from a Value-at-Risk perspective, see Aziz and Rosen (2010).
2. See, for example, Ledoit and Wolf (2003, 2004).
3. Naive diversification is a strategy that allocates an equal share of capital to the constituents of a portfolio constructed with a no-sales constraint. See, for example, DeMiguel *et al* (2009).

REFERENCES

Alexander, C. and Sheedy, E. (eds.) (2010) Advanced value at risk models. In: *The Professional Risk Managers' Handbook*. Wilmington, DE: PRMIA Publications, pp. 114–155.
Aziz, A. and Rosen, D. (2010) Capital allocation and RAPM. In: C. Alexander and E. Sheedy (eds.). *The Professional Risk Managers' Handbook*. Wilmington, DE: PRMIA Publications, pp. 12–40.
Campbell, J.Y. and Viceira, L.M. (2002) *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. New York: Oxford University Press.
DeMiguel, V., Garlappi, L. and Uppal, R. (2009) Optimal versus naive diversification: How inefficient is the 1/N

portfolio strategy? *Review of Financial Studies* 22(5): 1915–1953.
Grinold, R.C. and Kahn, R.N. (2000) *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Selecting Superior Money Managers*. 2nd edn. New York: McGraw-Hill.
Kahneman, D. and Tversky, A. (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2): 263–291.
Ledoit, O. and Wolf, M. (2003) Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10(5): 603–621.
Ledoit, O. and Wolf, M. (2004) Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management* 30(4): 110–119.
Lo, A.W. and MacKinlay, A.C. (1990) Data-snooping biases in tests of financial asset pricing models. *Review of Financial Studies* 3(3): 431–467.
Lyons, R.K. (2001) *The Microstructure Approach to Exchange Rates*. Cambridge: MIT Press.
Menkhoff, L. and Taylor, M.P. (2007) The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature* 45(4): 936–972.
J.P. Morgan/Reuters (1996) *RiskMetrics-Technical Document*. 4th edn. New York: Morgan Guaranty Trust Company.
Politis, D. and Romano, J. (1994) The stationary bootstrap. *Journal of the American Statistical Association* 89(428): 1303–1313.
Qi, M. and Wu, Y. (2006) Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market. *Journal of Money, Credit, and Banking* 38(8): 2135–2158.
Sullivan, R., Timmermann, A. and White, H. (1999) Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance* 54(5): 1647–1691.
White, H. (2000) A reality check for data snooping. *Econometrica* 68(5): 1097–1126.