Bank of Canada Banque du Canada

Working Paper 2002-2 / Document de travail 2002-2

Asset Allocation Using Extreme Value Theory

by

Younes Bensalah

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-2

January 2002

Asset Allocation Using Extreme Value Theory

by

Younes Bensalah

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
ybensalah@bankofcanada.ca

Contents

Acknowledgements

I am grateful to John Kiff, Shafiq Ebrahim, Toni Gravelle, and Peter Thurlow for helpful comments, suggestions, and documentation. Also, I would like to thank André Bernier for providing me with data.

Abstract

This paper examines asset allocation strategies in an extreme value at risk (VaR) framework in which the risk measure is the *p*-quantile from the extreme value distribution. The main focus is on the allocation problem faced by an extremely risk-averse institution, such as a central bank. The optimal portfolio in terms of excess return over the risk-free rate per unit of risk is also described.

An example of asset allocation is presented using a 1-year treasury bill and a 5-year zero-coupon bond. The allocation is conducted using different risk measures: duration, standard VaR, the quantile of the empirical distribution, and the quantile of the extreme distribution. An approximation procedure is described for the allocation of *N*-assets. An example of allocating eight Canadian treasuries and bonds is given (covering the whole Canadian term structure).

The implications of the results on optimal allocation of capital under stressed market conditions are discussed. Some practical issues concerning the use of the results are described, such as who should allocate capital based on extreme values.

JEL classification: C0, C4, C5, G1
Bank classification: Financial markets

Résumé

L'auteur examine diverses stratégies de répartition de l'actif à l'aide d'un cadre d'analyse de la valeur exposée au risque dans lequel la mesure du risque est le quantile d'ordre *p* de la distribution des valeurs extrêmes. Il se penche principalement sur le problème de répartition auquel est confrontée une institution extrêmement réfractaire au risque, telle une banque centrale. L'auteur présente également le portefeuille « optimal », soit celui qui maximise l'excédent de rendement par rapport au taux sûr, par unité de risque.

L'étude propose l'exemple d'un portefeuille de bons du Trésor à 1 an et d'obligations coupon zéro à 5 ans réparti selon différentes mesures du risque : duration, quantile de la distribution normale, quantile de la distribution empirique et quantile de la distribution des valeurs extrêmes. L'auteur décrit une procédure d'approximation pour la répartition de *N* actifs. Un exemple fondé sur la répartition de huit obligations et bons du Trésor du Canada (couvrant l'ensemble de la structure des taux d'intérêt en place au Canada) illustre l'application de cette procédure.

L'auteur examine les implications des résultats pour la répartition optimale du capital lorsque des tensions s'exercent sur le marché. Il aborde également certaines questions pratiques soulevées par les résultats (p. ex. celle de savoir quels investisseurs devraient se fonder sur la théorie des valeurs extrêmes pour déterminer la composition de leur portefeuille).

Classification JEL : C0, C4, C5, G1
Classification de la Banque : Marchés financiers

1. Introduction

Asset allocation was first addressed in a return-variance framework (Markowitz 1952). The allocation of capital among a set of assets is an optimization problem: an investor could consider maximizing the return per unit of risk or solving for the minimum risk portfolio. Using variance as a measure of risk and assuming normality of returns leads to optimal allocation under normal market conditions. The assumption of normality gives the same weight to positive and negative returns and underweights extreme shocks relative to central observations. Many financial returns do not follow a normal distribution, and are better represented by fat-tailed distributions. In fact, ignoring the fat-tail phenomenon might lead to an optimistic allocation of assets, since the allocation underweights high-risk tail events. On the other hand, the assumption of normality makes the calculus much easier, because the optimization problem has a closed form; a mixture of N normal distributions is a normal distribution. Also, using fat-tailed distributions to represent the returns adds complexity to the problem (no closed form exists), and the allocation can be very time-consuming and may not be possible (in terms of the time calculation) for a large set of assets. Using fat-tailed distributions solves part of the problem, but how would one achieve a conservative allocation of assets? In this case, where an investor is mainly interested in extreme shocks, we need a model that fits the tail of the distribution without making any assumptions about the nature of the underlying distribution. Extreme value theory (EVT) can be used for that purpose (see Bensalah 2000).

In this paper, in addressing the problem of asset allocation we will focus on stressed market conditions, using the quantile from an extreme value distribution (EVD) as a measure of risk. Section 2 allocates two fixed-income securities using different measures of risk and using different assumptions about the underlying distribution of returns. Surprisingly, the allocation using the quantile from the EVD allocates some weight to the 5-year zero-coupon bond. This result goes against the principle that associates more risk with longer-duration assets, which is widely accepted in the world of fixed-income securities. We also present the minimum risk portfolio[1] and the optimal portfolio in terms of maximizing return per unit of risk. Section 3 presents an approximation procedure to solve the problem in the case of a large set of assets and a lack of a closed form. The results for allocating two and three fixed-income securities, using an intuitive and straightforward (but time-consuming) procedure, as well as an approximation procedure, are presented and compared. Section 4 offers some conclusions and introduces some ideas that might improve the accuracy of the solution.

1. The minimum risk portfolio is selected on the basis of minimizing a given measure of risk (for example, variance).

2. Asset Allocation using Different Measures of Risk

The asset-allocation problem was first addressed by Markowitz (1952) in a framework that considers the risk-return trade-off and using variance of returns as a measure of risk. The use of the variance and the assumption of normality of returns makes this problem easy to solve, because of the nature of the normal distribution: a mixture of N normal is a normal distribution.

An investor using a variance as a measure of risk to allocate assets should understand the following: (i) the same weight is given to negative and positive returns, and (ii) the optimal return is achieved using this measure only if the market remains under normal conditions, because more importance is given to the central observations (assumption of normality). Hence, the importance of extreme shocks is underestimated.

Since many financial distributions of returns are fat-tailed, an approximation by a normal distribution leads the importance of extreme shocks to be underestimated. Huisman, Koedijk, and Pownall (1999) compare the allocation of an index of bonds and an index of stocks using a normal value at risk (VaR) and a VaR from the empirical distribution as measures of risk. They conclude that, assuming normality, one obtains the same portfolio allocation for any choice of confidence level. The assumption of normality renders the investor's attitude towards risk (in terms of increasing the confidence level) unimportant in the optimization process (the normal VaR is a linear function of the variance). On the other hand, the non-parametric nature of the empirical distribution leads to a change in the optimal portfolio for different confidence levels. The problem with using the empirical distribution is that it does not allow for an allocation based on a very high confidence level (no out-of-sample projection[2]). Huisman, Koedijk, and Pownall also use a fat-tailed distribution to capture the effect of the tail. Although allocating two assets is not time-consuming under any distribution assumption, for non-normal distribution—since no closed form exists for the optimization problem—allocating a large set of assets could be impossible (see section 3). They conclude that the use of a measure of risk that can capture higher moments of the distribution is more adequate.

The use of fat-tailed distribution improves the accuracy of measuring risk and reduces model risk. Using the p-quantile of a fat-tailed parametric distribution as a measure of risk could be beneficial for investors who are mainly concerned with central observations but also want to capture more accurately the downside risk. For those investors, an allocation under the assumption of a fat-tailed

2. For a very high quantile (say, 99.99 per cent), one should choose the 0.01 per cent observation from the ranked sample. Even with very large samples of data, this quantile may not exist when empirical distributions are used. The lack of a parametric form in the distribution does not permit a projection to the high quantile measures of the distribution.

distribution (such as the *t*-student distribution) and based on a reasonable confidence level (not far into the tail) could be useful.

Some investors, however, are concerned with extreme shocks. They want a model that accurately fits the tail of the distribution. EVT provides such a tool. It has four major advantages over the classic methods: (i) it focuses on extreme shocks and ignores the central observations; (ii) it makes no assumption regarding the underlying distribution of returns, so the model risk is reduced; (iii) it is a parametric approach, so out-of-sample projection is possible; and, (iv) it is better able to capture event risk (i.e., crashes, currency devaluations).

Thus, using the quantile from an EVD can be interesting for some risk-averse investors, who want to achieve a very conservative allocation of assets—in other words, a portfolio allocation that is robust under stressed market conditions.

2.1 A two-asset allocation example using different measures of risk

In this section, two fixed-income securities are allocated based on four measures of risk. The results for each measure are described. Series of daily returns over a 5-year period (1995–2000) are used with mean returns of 10 July 2000.

2.1.1 Duration

Duration, which is probably the most popular fixed-income risk measure, is the first derivative of the price with respect to the interest rate, and so it represents the sensitivity of the price to a parallel change in interest rates. Duration is simple to use. For example, to choose between a 1-year treasury-bill and a 5-year zero-coupon bond, we choose the instrument that has a short duration to minimize the risk, which in our case means that we put all of our capital in the 1-year treasury bill. Duration, however, does not capture the curvature effect (convexity) and higher interest rate moments, and it assumes parallel changes for the whole term structure.

2.1.2 Normal VaR

If we assume that returns are normal, the VaR or the *p*-quantile of the normal distribution is a linear function of the variance. This implies that portfolio allocations are the same for any confidence level chosen by the investor. In Table 1, the optimal allocation using a normal VaR corresponds to an optimal allocation using duration.

Table 1: Allocation using Normal VaR (per cent)

	Minimum risk portfolio	**Optimal portfolio**
VaR	0.10410	0.10410
Return	6.2983	6.2983
Weight 1-year treasury bill	100	100
5-year zero-coupon bond	0	0

2.1.3 Historical VaR

For allocations based on a historical VaR (Tables 2 and 3), no assumption is made about the underlying distribution of returns, and we use the quantile of the empirical distribution as a measure of risk. We still give the same weight, however, to all the observations, and we cannot perform an allocation based on a very high confidence level (an out-of-sample projection is not possible; see Table 2, column 0.9999). These two facts show the weakness of using non-parametric distributions to model the distribution of returns.

Table 2: Allocation using Historical VaR (minimum risk portfolio) (per cent)

Confidence level	0.99	0.999	0.9999
VaR	0.3243	0.60411	NA
Return	6.2958	6.2983	NA
Weight 1-year treasury bill	99	100	NA
5-year zero-coupon bond	1	0	NA

Table 3: Allocation using Historical VaR (optimal portfolio) (per cent)

Confidence level	0.99	0.999	0.9999
VaR	0.32437	0.60411	NA
Return	6.2983	6.2983	NA
Weight 1-year treasury bill	100	100	NA
5-year zero-coupon bond	0	0	NA

2.1.4 Extreme VaR

In this subsection, we use the *p*-quantile from the EVD to allocate the 1-year treasury bill and the 5-year zero-coupon bond (Tables 4 and 5). The focus is on extreme shocks. Table 4 shows how the dependency between the two assets differs in the tail from what the normal distribution may suggest. If we take into account the extreme shocks, we should allocate part of our capital to the long-duration bond, and the same applies for the optimal portfolio in terms of return per unit of risk. Figure 1 shows this allocation graphically. The 1-year treasury bill seems to be more volatile in an extreme condition, so investing 100 per cent of the capital in this asset can be very risky under extreme conditions.

Table 4: Allocation using Extreme VaR (Minimum Risk Portfolio) (per cent)

Confidence level	0.99	0.999	0.9999
VaR	1.2120	1.8637	2.4127
Return	6.2703	6.2397	6.2065
Weight 1-year treasury bill	89	77	64
5-year zero-coupon bond	11	23	36

Table 5: Allocation using Extreme VaR (Optimal Portfolio) (per cent)

Confidence level	0.99	0.999	0.9999
VaR	1.2259	1.9145	2.4910
Return	6.2881	6.2626	6.2371
Weight 1-year treasury bill	94	86	76
5-year zero-coupon bond	6	14	24

In the end, we can say that there is no final answer or optimal measure of risk. Any allocation is optimal if its underlying assumptions correspond to the risk preference of the investor (i.e., the sensitivity of the investor to tail events). Understanding the risk preference is key before using any methodology. An institution that cares especially about extreme events should allocate assets conditioned on extreme events. To achieve this very conservative allocation of assets, however, some basis points of the portfolio return are conceded.

3. Allocating a Large Set of Assets using an Approximation Procedure

Asset allocation using the *p*-quantile of the normal distribution (a linear function of the variance, in this case) as a measure of risk is straightforward. The assumption of normality simplifies the calculus, because a mixture of *N* marginal normal distributions is a normal distribution, and the second moment of the resulting distribution can be written in terms of the marginal distribution's variances and covariances. Solving for the optimal portfolio in terms of minimum risk is well understood in the literature and maximizing the return per unit of risk is equivalent to maximizing the Sharpe ratio. The asset-allocation problem under the assumption of normality is reduced to a simple optimization problem, which can be solved for any number of assets and under any kind of constraints. Relaxing the assumption of normality adds a great deal of complexity to the problem. If we assume that the distribution of returns is not normal, no closed form for the optimization problem exists (i.e., we cannot write the resulting *p*-quantile of the portfolio distribution in terms of the parameters of the individual distributions of assets).

In our case, we chose the *p*-quantile of an EVD as a measure of risk. A mixture of *N* EVD is not necessarily an EVD, this being one of the limitations of EVT (see Bensalah 2000).

One can, nevertheless, compute the risk-return value of individual portfolios. As such, an intuitive procedure used to solve this problem is to draw the surface of all possible portfolios in a return-risk framework and then choose the optimal portfolio in terms of minimum risk or maximum return per unit of risk. Unfortunately, this procedure works well only for a limited number of assets, as the computational burden is large. To illustrate this fact, Algorithm I describes the procedure for allocating two assets.

Algorithm I: Allocating two assets

Step 1 - Define a step of 1 per cent.

Step 2 - Allocate 99 per cent to the first asset.
Allocate 1 per cent to the second asset.
Estimate the EVD of the portfolio, the *p*-quantile, and return.

. . . .

Step n - Allocate (100-n) per cent to the first asset.
Allocate n per cent to the second asset.
Estimate the EVD of the portfolio, the *p*-quantile, and return.

. . . .

Step 100 - Allocate 1 per cent to the first asset.
Allocate 99 per cent to the second asset.
Estimate the EVD of the portfolio, the *p*-quantile, and return.

End - Select the portfolio of minimum risk or the portfolio of maximum return per unit of risk.

We can see that we need 100 calculations if we choose a precision of 1 per cent; for three assets we need 100*100 calculations; and for *N* assets we need 100^{n-1} calculations.

This procedure quickly becomes very cumbersome in terms of computation time for a large set of assets. The problem is that we have no control over the number of portfolios to draw to properly map the surface of all possible portfolios. We can choose a smaller number of portfolios to draw combined with an interpolation between every two portfolios. However, for a large set of assets, the computation problem remains unsolved.

Algorithm II allows us to control the number of portfolios, gives an approximation of the solution, and can take into account the available computer resources.

Algorithm II: Allocating two assets

Step 1 - Define a number of portfolios as m.

Step 2 - Randomly generate the weight of the first asset, $W_1 \sim U[0,1]$, from a uniform distribution [0,1], and set the weight of the second asset to:

$W_2 = 1 - W_1.$

Estimate the EVD of the portfolio, the *p*-quantile, and return.

Repeat step 2 m times.

End - Select the portfolio of minimum risk or the portfolio of maximum return per unit of risk.

We can see that we have control over the number of points (portfolios) we want to draw and thus we can determine how accurate we want our approximation of the surface to be, taking into account the available computer resources. Parallel programming suits this kind of problem, because each estimation at each point is independent of the others. We just have to be careful not to generate the same random weights in all the parallel programs (by manipulating the seeds of the random generator).

Tables 6 and 7 apply this procedure to two and three assets. Also, the procedure is compared to the solution based on Algorithm I for giving *p*-quantile as a measure of risk. Figures 1 to 8 plot the results presented in Tables 6 and 7.

Table 6: Allocating Two Assets (minimum risk portfolio)[a] (per cent)

Confidence level (0.9999)	Algorithm I (100 portfolios)	Algorithm II (25 random portfolios)	Algorithm II (50 random portfolios)	Algorithm II (70 random portfolios)
VaR	2.4127	2.4132	2.4132	2.4132
Return	6.2065	6.2032	6.2032	6.2032
Weight 1-year treasury bill	64	62.6886	62.6886	62.6886
Weight 5-year zero-coupon bond	36	37.3114	37.3114	37.3114

a. See Figures 2, 3, and 4 for graphs.

Table 7: Allocating Three Assets (minimum risk portfolio)[a] (per cent)

Confidence level (0.9999)	Algorithm I (10,000 portfolios)	Algorithm II (80 random portfolios)	Algorithm II (100 random portfolios)	Algorithm II (200 random portfolios)	Algorithm II (1000 random portfolios)	Algorithm II (2000 random portfolios)	Algorithm II (5000 random portfolios)
VaR	2.0301	2.0884	2.0884	2.0550	2.0385	2.0291	2.0291
Return	6.2356	6.2430	6.2430	6.2448	6.2390	6.2351	6.2351
Weight 1-year treasury bill	91	90.8006	90.8006	92.2364	91.3157	90.9912	90.9912
Weight 5-year zero-coupon bond	0.9	2.0636	2.0636	0.2101	0.3707	0.0063	0.0063
Weight 30-year zero-coupon bond	8.91	7.1356	7.1356	7.5534	8.3135	9.0025	9.0025

a. See Figure 5, 6, 7, and 8 for graphs.

Now, we extend the procedure to allocate N number of assets and we apply it to eight Canadian treasury bills and bonds (covering the whole Canadian term structure, as regards maturity). Series of daily returns over a 5-year period (1995–2000) are used to estimate the EVD. The EVD is estimated using the maxima sampled over 90-day non-overlapping intervals[3] (see Bensalah 2000 for a discussion on sampling the maxima).

Algorithm III: Allocating N assets

Step 1 - Define a number of portfolios as m.

Step 2 - Randomly generate the weight of the first asset, $W_1 \sim U[0, 1]$, from a uniform distribution [0,1].

$W_2 \sim U[0, 1 - W_1$

$W_3 \sim U[0, 1 - W_1 - W_2]$

$W_4 \sim U[0, 1 - W_1 - W_2 - W_3]$

…………………………………………

$W_n = 1 - W_1 - W_2 - W_3 \ldots\ldots\ldots\ldots\ldots\ldots W_{n-1}$

Estimate the EVD of the portfolio, the p-quantile, and return.

Repeat step 2 m times.

End - Select the portfolio of minimum risk or the portfolio of maximum return per unit of risk.

3. One way to build the sample of extremes is to divide the original sample into non-overlapping equal intervals of time and take the maximum within each interval. The EVD is estimated to fit the samples of these maxima.

Tables 8 and 9 show the results of allocating the eight assets.

Table 8: Minimum Risk Portfolio Simulation (per cent)

Confidence level (0.9999)	100 random portfolios	200 random portfolios	400 random portfolios	1000 random portfolios	5000 random portfolios	10,000 random portfolios	20,000 random portfolios
VaR	0.45	1.64	0.54	0.52	0.19	0.19	0.19
Return	5.53	5.81	5.53	5.53	5.51	5.51	5.51
1-month treasury bill	70.98	16.28	92.08	91.30	94.25	94.25	95.39
2-month treasury bill	22.41	13.63	1.29	1.42	2.31	2.31	1.52
3-month treasury bill	4.36	29.77	2.82	3.09	2.46	2.46	2.47
6-month treasury bill	0.89	12.76	1.21	1.33	0.61	0.61	0.06
1-year treasury bill	0.05	17.81	1.68	1.85	0.17	0.17	0.06
5-year zero-coupon bond	0.05	0.82	0.08	0.08	0.09	0.09	0.45
10-year zero-coupon bond	0.12	0.74	0.07	0.08	0.04	0.04	0.05
30-year zero-coupon bond	1.14	8.19	0.77	0.85	0.07	0.07	0.02

Table 9: Optimal Portfolio Simulation (per cent)

Confidence level (0.9999)	100 random portfolios	200 random portfolios	400 random portfolios	1000 random portfolios	5000 random portfolios	10,000 random portfolios	20,000 random portfolios
VaR	2.82	2.82	2.52	1.02	1.60	1.60	1.98
Return	6.03	6.03	5.66	5.74	6.03	6.03	6.15
1-month treasury bill	16.99	16.99	75.68	57.56	11.64	11.64	0.29
2-month treasury bill	2.08	2.08	0.61	0.78	17.32	17.32	4.35
3-month treasury bill	13.12	13.12	3.84	6.14	1.73	1.73	4.93
6-month treasury bill	20.08	20.08	5.88	12.49	11.25	11.25	14.32
1-year treasury bill	43.68	43.68	12.79	10.48	51.71	51.71	52.62
5-year zero-coupon bond	3.70	3.70	1.08	10.61	3.22	3.22	21.56
10-year zero-coupon bond	0.24	0.24	0.07	0.44	0.28	0.28	0.83
30-year zero-coupon bond	0.12	0.12	0.04	1.50	2.85	2.85	1.10

We can summarize the results of the approximation procedure; see Figures 9 and 10 for a graphical example using 10,000 and 20,000 random portfolios. For an automatic allocation using a precision of 1 per cent (using Algorithm I), 100^7portfolios would have to be estimated.

The quality of the approximation depends on the number of points randomly generated and on the number of assets to allocate. Figures 11 and 12 illustrate how the random-mapping solution evolves with the number of random portfolios. We can set up a convergence rule based on the minimum fluctuation between the estimated portfolios. That is, we can increase the number of random portfolios calculated (m) until the optimal portfolio varies little between each increment in the number m. This is likely to lead to a sampling size that is significantly smaller than the one required by Algorithm I (100^7).

Using a more sophisticated measure of risk is difficult and can be very time-consuming. Mapping the original set of assets into a small set could be a first step before solving for the optimal benchmark portfolio.

4. Conclusion

Asset allocation is a complex problem that should take into account the risk profile, different kinds of constraints (minimum or maximum holdings; shortfall constraints, such as the confidence level), and the universe of assets. The emergence of VaR as an ex-post risk measure was a revolution in the world of risk management. Using VaR with more complex assumptions than normality in a proactive way (i.e., to allocate assets) is a new challenge.

A better understanding of the risk profile will lead to the use of appropriate assumptions and measures of risk. Investors who equally weight positive and negative returns (and focus on the central observations) can continue to use the variance as a measure of risk and the assumption of normal distribution of returns. This approach is compatible with their perception and understanding of their risk profile. Even investors who focus on central observations, however, might think that the downside risk should play a bigger role than the normal assumption might suggest. The use of fat-tailed distributions (such as *t*-student) and the *p*-quantile as a measure of risk can partially capture the effect of the tail. A third category of risk-averse investors needs to minimize the risk of extreme shocks. EVT provides a model that focuses on extreme shocks without making any assumptions about the underlying distribution of returns, reducing the model's risk. Central banks likely fall into this category and EVT captures the event risks that they would tend to focus on.

Also, in times of market stress, adjusting asset allocations on an EVT basis could be an alternative strategy to the classic flight to quality. Having different reaction strategies reduces the behaviour risk (the risk that all market participants react the same way, increasing the magnitude of a shock).

By relaxing the assumption of normality, however, the computation time can become explosive. The approximation procedure presented in this paper gives a reasonable approximation of the real solution based on a random mapping of the surface of all possible portfolios. This procedure was used to allocate eight Canadian treasury bills and Canada bonds (covering the whole Canadian term structure, as regards maturity). The resulting portfolio can be seen as a Canadian fixed-income benchmark for an extremely risk-averse institution. Other algorithms to generate the random weights could be considered. We might see some improvement to the procedure if, as a second step, we could sample the portfolios around the approximation to improve the solution's accuracy.

Bibliography

Bensalah, Y. 2000. “Steps in Applying Extreme Value Theory to Finance: A Review.” Bank of Canada Working Paper No. 2000–20 (http://www.bankofcanada.ca/en/res/wp00-20.htm).

Huisman, R., G.K. Koedijk, and A.J.R Pownall. 1999. “Asset Allocation in a Value at Risk Framework.” Erasmus University, Rotterdam, Faculty of Business Administration.

Lucas, A. and P. Klaassen. 1998. “Extreme Returns, Downside Risk, and Optimal Asset Allocation.” *Journal of Portfolio Management* (Fall): V25-1.

Markowitz, H. 1952. “Portfolio Selection.” *Journal of Finance* 7: 77–91.

Wanj, J. 2000. “Mean-Variance-VaR Based Portfolio Optimization.” Department of Mathematics and Computer Science, Valdosta State University.

Figure 1: The allocation of a 1-year treasury bill and a 5-year zero-coupon bond using a precision step of 1 per cent (100 portfolios). "*" is the minimum risk portfolio, and "O" is the optimal portfolio in terms of return/risk.

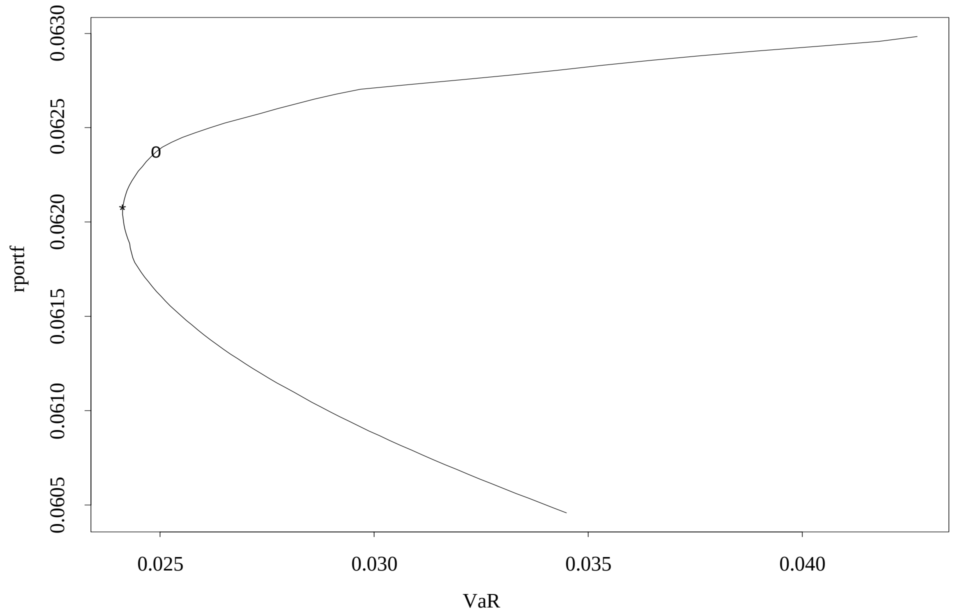

Figure 2: The allocation of a 1-year treasury bill and a 5-year zero-coupon bond using 25 random portfolios. "O" is the minimum risk portfolio.

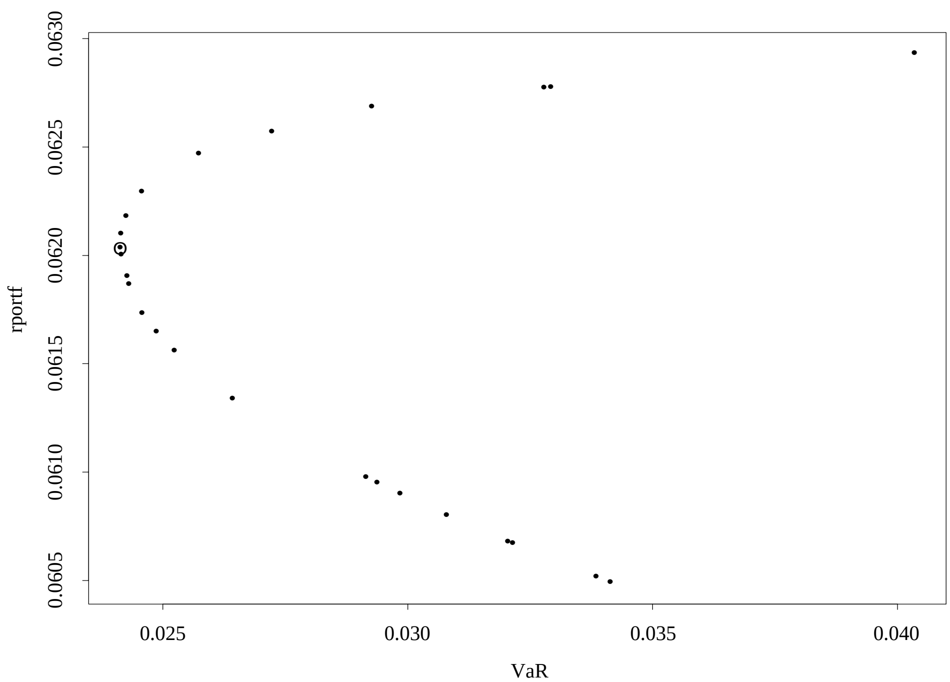

Figure 3: The allocation of a 1-year treasury bill and a 5-year zero-coupon bond using 50 random portfolios. "O" is the minimum risk portfolio.

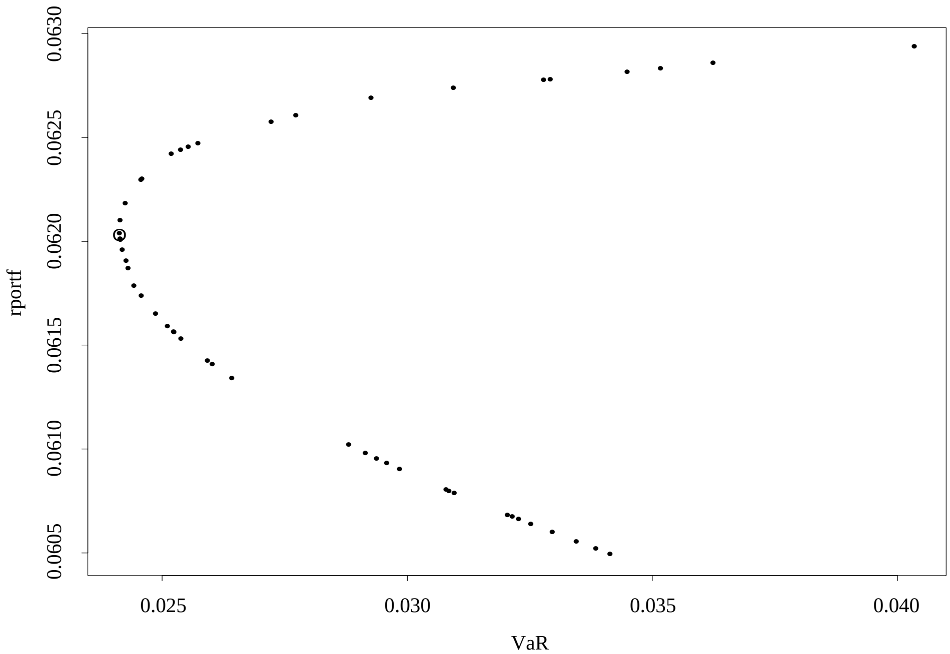

Figure 4: The allocation of a 1-year treasury bill and a 5-year zero-coupon bond using 70 random portfolios. "O" is the minimum risk portfolio.

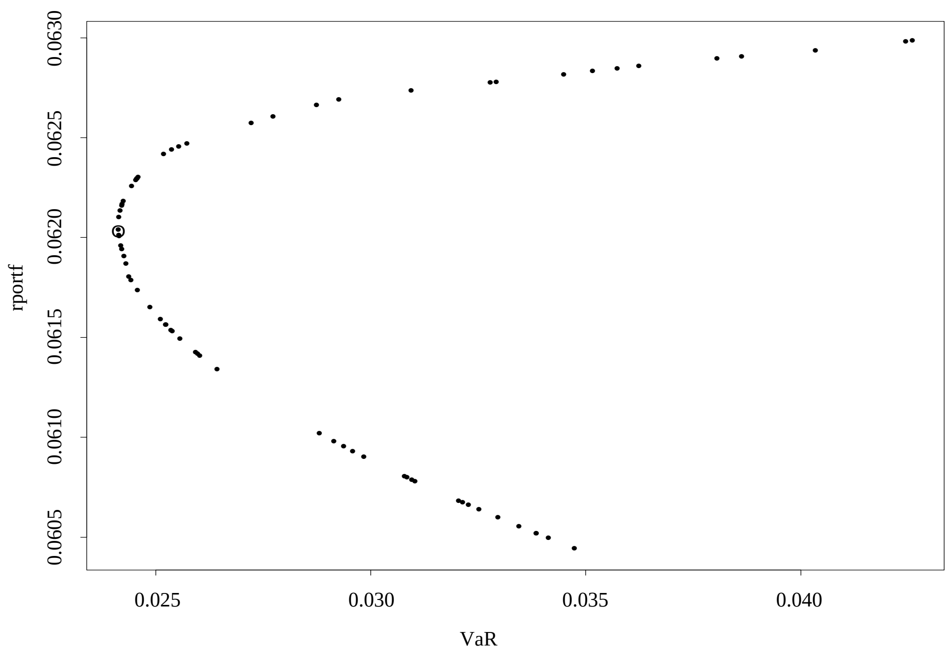

Figure 5: The allocation of a 1-year treasury bill, a 5-year zero-coupon bond, and a 30-year zero-coupon bond using a precision step of 1 per cent (10,000 portfolios). "*" is the minimum risk portfolio, and "O" is the optimal portfolio in terms of return/risk.

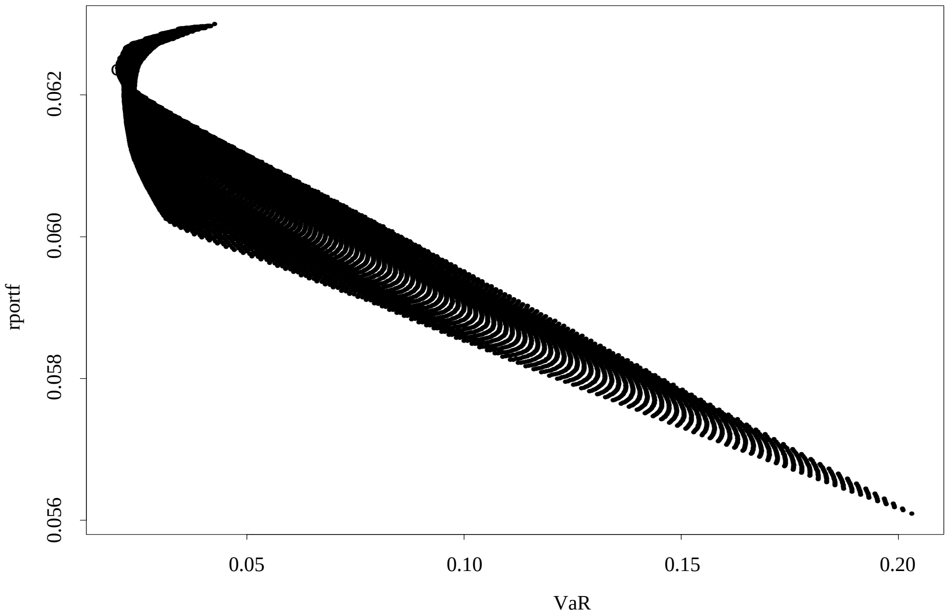

Figure 6: The allocation of a 1-year treasury bill, a 5-year zero-coupon bond, and a 30-year zero-coupon bond using 80 random portfolios. "O" is the minimum risk portfolio.

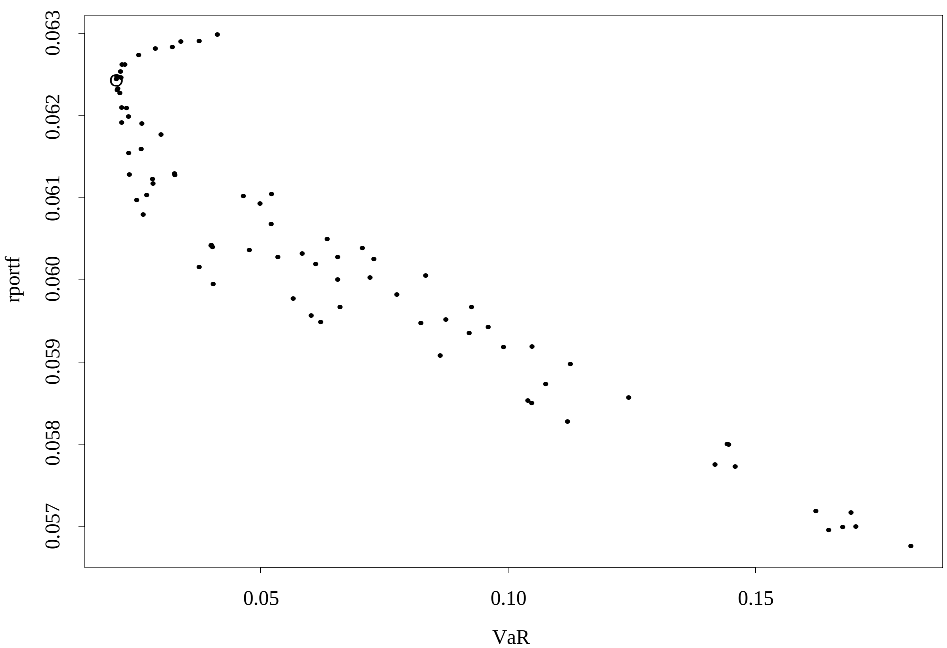

Figure 7: The allocation of a 1-year treasury bill, a 5-year zero-coupon bond, and a 30-year zero-coupon bond using 1000 random portfolios. “O” is the minimum risk portfolio.

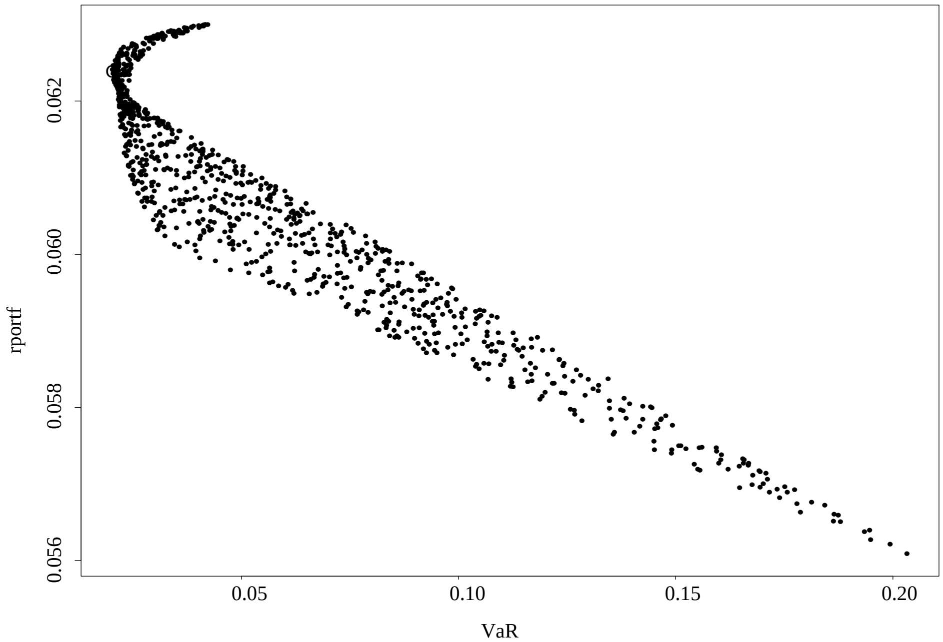

Figure 8: The allocation of a 1-year treasury bill, a 5-year zero-coupon bond, and a 30-year zero-coupon bond using 2000 random portfolios. “O” is the minimum risk portfolio.

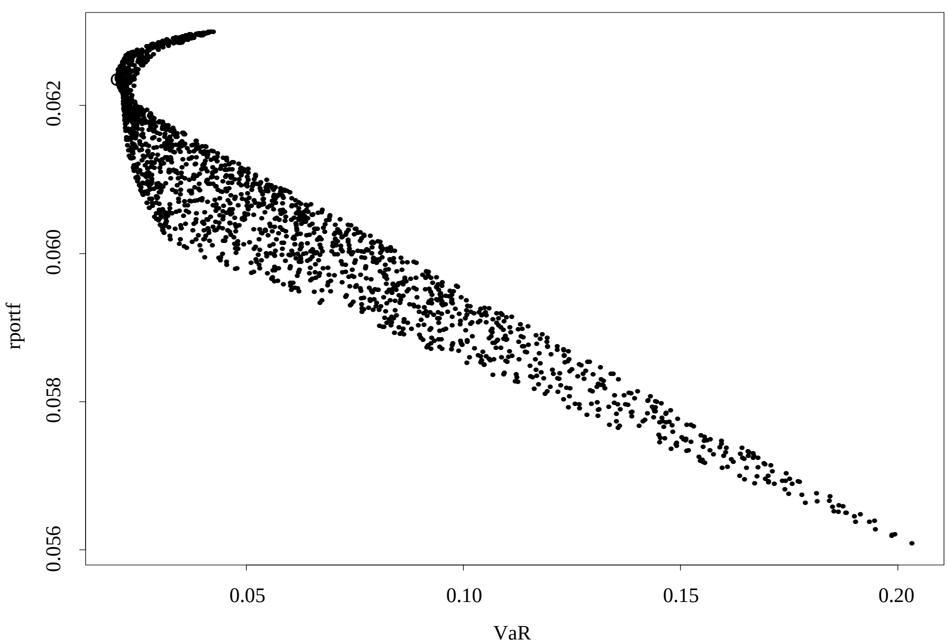

Figure 9: The allocation of eight Canadian fixed-income securities using an approximation of 10,000 random portfolios. “M” is the minimum risk portfolio, and “O” is the optimal portfolio in terms of return/risk.

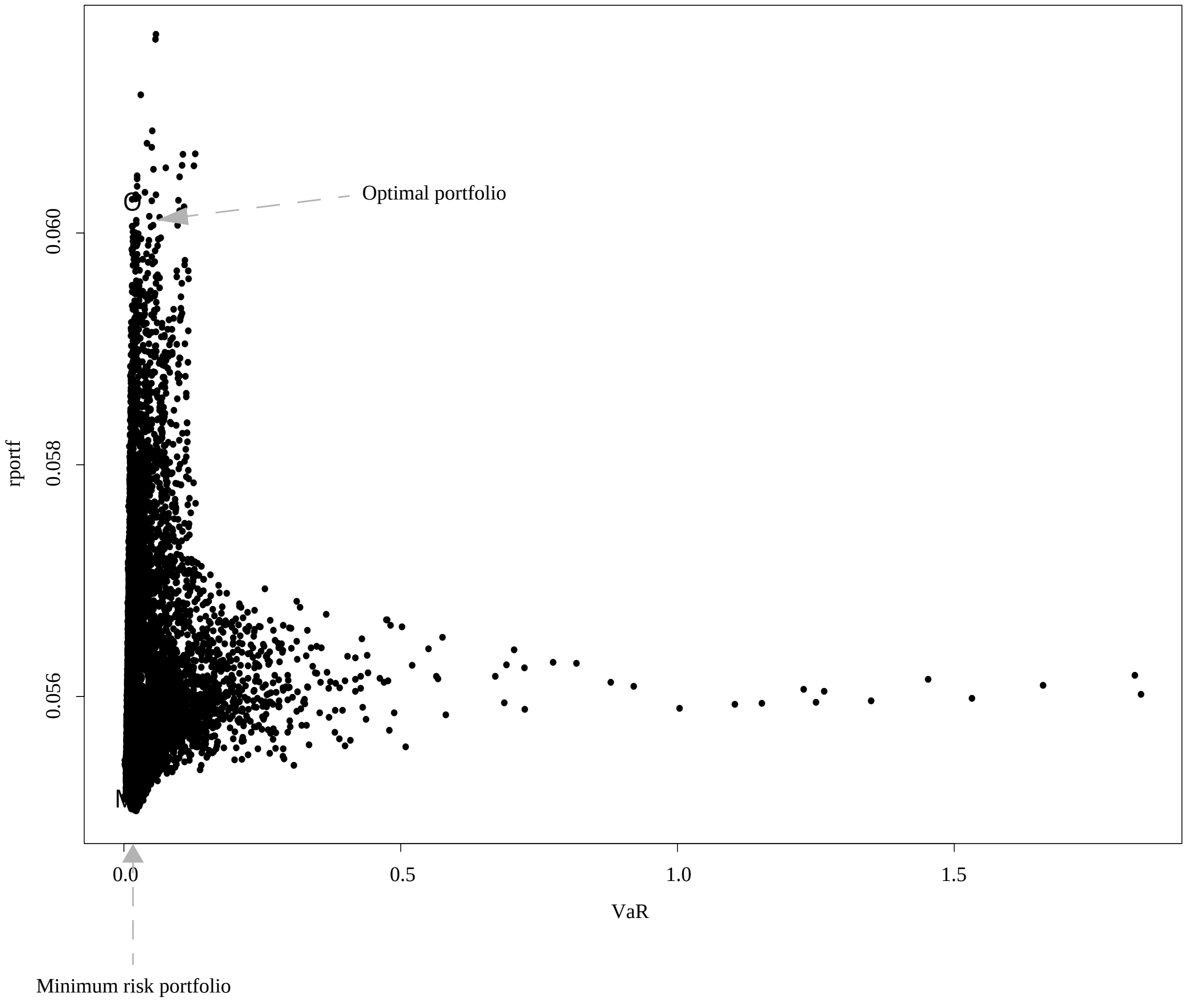

Figure 10: The allocation of eight Canadian fixed-income securities using an approximation of 20,000 random portfolios. “M” is the minimum risk portfolio, and “O” is the optimal portfolio in terms of return/risk.

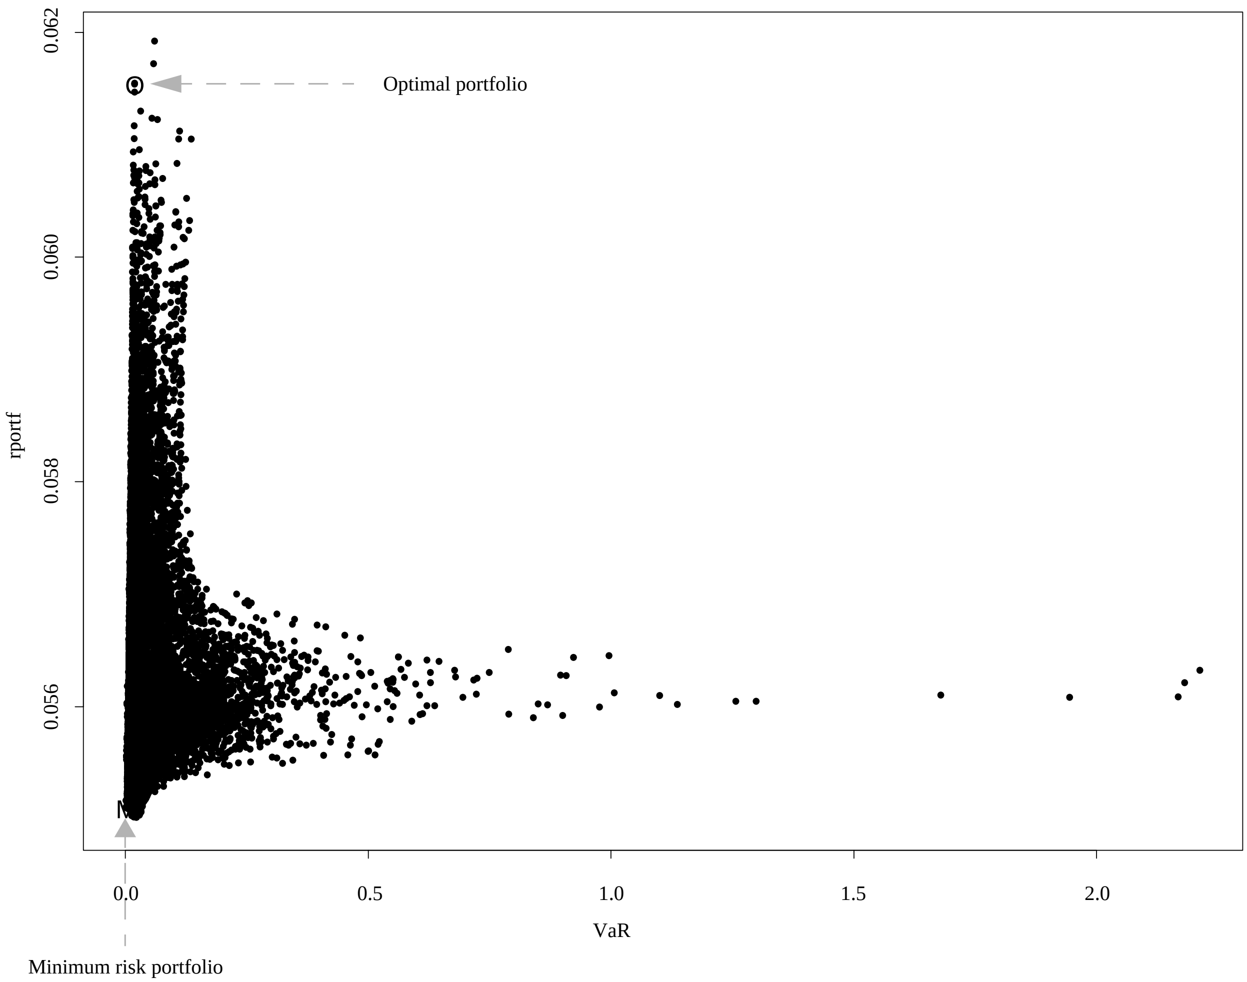

Figure 11: Random simulation for the minimum risk portfolio

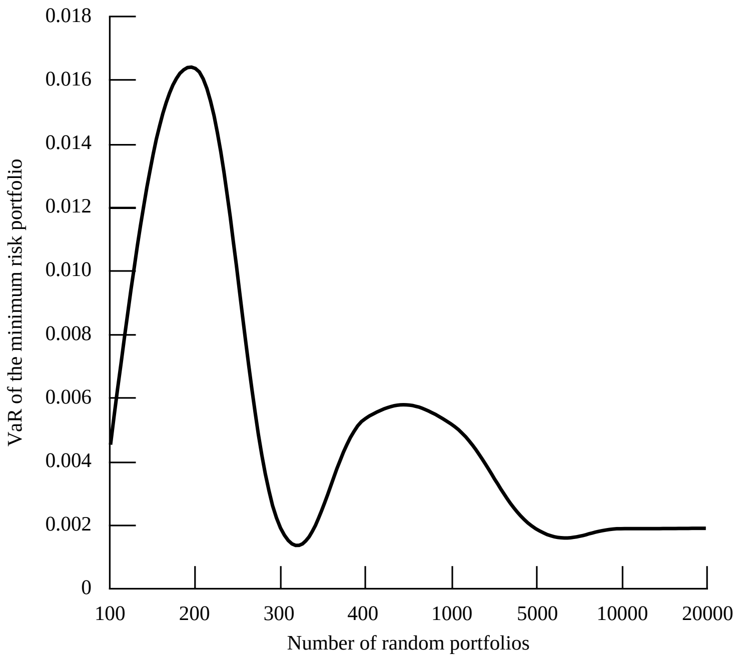

Figure 12: Random simulation for the minimum risk portfolio

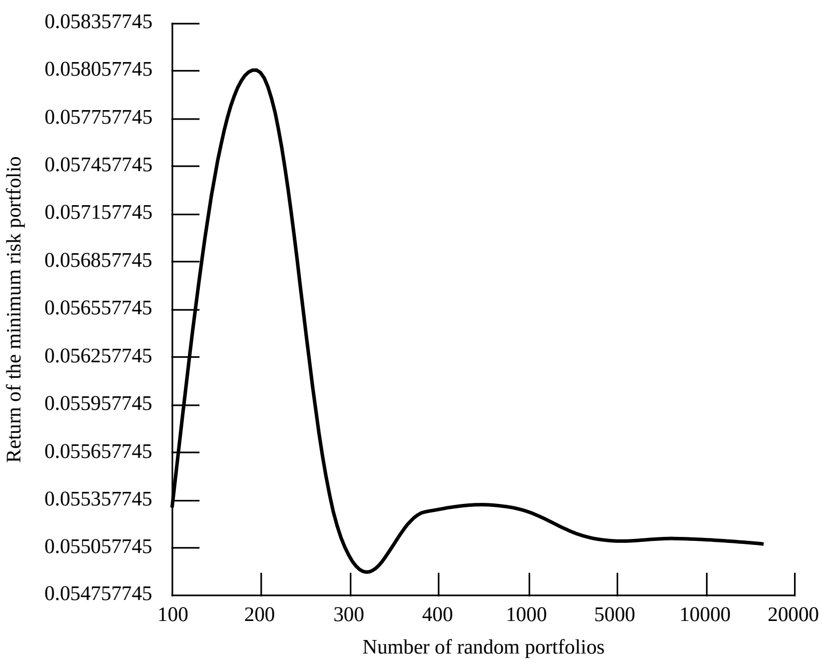

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue*.**

2002

2002-1	Taylor Rules in the Quarterly Projection Model	J. Armour, B. Fung, and D. Maclean

2001

2001-27	The Monetary Transmission Mechanism at the Sectoral Level	J. Farès and G. Srour
2001-26	An Estimated Canadian DSGE Model with Nominal and Real Rigidities	A. Dib
2001-25	New Phillips Curve with Alternative Marginal Cost Measures for Canada, the United States, and the Euro Area	E. Gagnon and H. Khan
2001-24	Price-Level versus Inflation Targeting in a Small Open Economy	G. Srour
2001-23	Modelling Mortgage Rate Changes with a Smooth Transition Error-Correction Model	Y. Liu
2001-22	On Inflation and the Persistence of Shocks to Output	M. Kichian and R. Luger
2001-21	A Consistent Bootstrap Test for Conditional Density Functions with Time-Dependent Data	F. Li and G. Tkacz
2001-20	The Resolution of International Financial Crises: Private Finance and Public Funds	A. Haldane and M. Kruger
2001-19	Employment Effects of Restructuring in the Public Sector in North America	P. Fenton, I. Ip, and G. Wright
2001-18	Evaluating Factor Models: An Application to Forecasting Inflation in Canada	M.-A. Gosselin and G. Tkacz
2001-17	Why Do Central Banks Smooth Interest Rates?	G. Srour
2001-16	Implications of Uncertainty about Long-Run Inflation and the Price Level	G. Stuber
2001-15	Affine Term-Structure Models: Theory and Implementation	D.J. Bolder
2001-14	L'effet de la richesse sur la consommation aux États-Unis	Y. Desnoyers
2001-13	Predetermined Prices and the Persistent Effects of Money on Output	M.B. Devereux and J. Yetman

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca