

2002s-53

Multi-Task Learning For Option Pricing

Joumana Ghosn, Yoshua Bengio

Série Scientifique
Scientific Series



CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

Montréal
Mai 2002

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de la Recherche, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de la Recherche, de la Science et de la Technologie, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

- École des Hautes Études Commerciales
- École Polytechnique de Montréal
- Université Concordia
- Université de Montréal
- Université du Québec à Montréal
- Université Laval
- Université McGill
- Ministère des Finances du Québec
- MRST
- Alcan inc.
- AXA Canada
- Banque du Canada
- Banque Laurentienne du Canada
- Banque Nationale du Canada
- Banque Royale du Canada
- Bell Canada
- Bombardier
- Bourse de Montréal
- Développement des ressources humaines Canada (DRHC)
- Fédération des caisses Desjardins du Québec
- Hydro-Québec
- Industrie Canada
- Pratt & Whitney Canada Inc.
- Raymond Chabot Grant Thornton
- Ville de Montréal

© 2002 *Joumana Ghosn, Yoshua Bengio*. Tous droits réservés. *All rights reserved*. Reproduction partielle permise avec citation du document source, incluant la notice ©.

Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source.

Les cahiers de la série scientifique (CS) visent à rendre accessibles des résultats de recherche effectuée au CIRANO afin de susciter échanges et commentaires. Ces cahiers sont écrits dans le style des publications scientifiques. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

Multi-Task Learning For Option Pricing

Joumana Ghosn^{*} and Yoshua Bengio[†]

Résumé / Abstract

L'apprentissage multi-tâches est une manière d'apprendre des particularités d'un domaine (le biais) qui comprend plusieurs tâches possibles. On entraîne simultanément plusieurs modèles, un par tâche, en imposant des contraintes sur les paramètres de manière à capturer ce qui est en commun entre les tâches, afin d'obtenir une meilleure généralisation sur chaque tâche, et pour pouvoir rapidement généraliser (avec peu d'exemples) sur une nouvelle tâche provenant du même domaine. Ici cette commonalité est définie par une variété affine dans l'espace des paramètres. Dans cet article, nous appliquons ces méthodes à la prédiction du prix d'options d'achat de l'indice S&P 500 entre 1987 et 1993. Une analyse de la variance des résultats est présentée, démontrant des améliorations significatives de la prédiction hors-échantillon.

Multi-task learning is a process used to learn domain-specific bias. It consists in simultaneously training models on different tasks derived from the same domain and forcing them to exchange domain information. This transfer of knowledge is performed by imposing constraints on the parameters defining the models and can lead to improved generalization performance. In this paper, we explore a particular multi-task learning method that forces the parameters of the models to lie on an affine manifold defined in parameter space and embedding domain information. We apply this method to the prediction of the prices of call options on the S&P index for a period of time ranging from 1987 to 1993. An analysis of variance of the results is presented that shows significant improvements of the generalization performance.

Mots clés : valorisation d'options d'achat, apprentissage multi-tâches, réseau de neurones artificiels

Keywords: option call pricing, multi-task learning, artificial neural networks

^{*} Département d'informatique et recherche opérationnelle, Université de Montréal, Montréal, Québec, Canada, H3C 3J7. Email: ghosn@iro.umontreal.ca

[†] CIRANO and Département d'informatique et recherche opérationnelle, Université de Montréal, Montréal, Québec, Canada, H3C 3J7. Tel: +1 (514) 343-6804, email: bengioy@iro.umontreal.ca

1 Introduction

Biassing the hypothesis space of a learner, by embedding in it domain information or properties specific to the task being tackled, can greatly improve its generalization performance. Many techniques have been developed to incorporate domain knowledge in a learning model. Still, deriving an appropriate bias remains a very complex problem. Multi-task learning can be used to circumvent this difficulty by trying to “learn” a domain-specific bias. This technique can be used when several similar tasks derived from the same domain are available. All the tasks are learned simultaneously and the models being trained on these tasks are forced to exchange domain information. The transfer of information is performed by imposing constraints on the parameters of the models. Different types of constraints exist which correspond to different multi-task learning methods.

An important issue, related to multi-task learning, consists in determining the required number of tasks necessary to generate an adequate bias that can help improve the generalization performance of the corresponding learning models, as well as the relationship between the required number of tasks and the size of their training sets. A theoretical analysis [1] was performed which showed that the number of necessary tasks needs to be “large” if the training sets are “small” and vice versa. Experimental results [1, 4], obtained by applying several multi-task learning techniques to artificially generated tasks, confirm this analysis. They also show that generalization performance improve with increasing sizes of the training sets for a fixed number of tasks and/or with increasing numbers of tasks for a fixed size of the training sets. Tests on a real-world problem presented in this paper also confirm those results.

2 Multi-Task Learning

Multi-task learning methods were developed to solve three types of problems: reducing the training time of learning models [8], limiting the overfitting tendencies related to the use of tasks with very low signal-to-noise ratio [5] and improving the generalization performance.

Several methods exist that deal with the last problem. They include the “Task Clustering” technique [10] which groups classification tasks in different clusters, each containing a set of similar tasks that are used to devise a common distance metric used for classification. Another method consists in learning internal representations [1, 2]. This is achieved by training several neural networks on a set of similar tasks and forcing them to share their first hidden layers. This process therefore requires that all the networks apply the same preprocessing to their examples and then separately analyze the preprocessed examples. An alternative to this method corresponds to forcing the neural networks to share their last hidden layers [3]. In that case, each network applies a different preprocessing to its examples and then uses a common analysis method to interpret the preprocessed information. Other variants of the internal representation learning technique include learning representations that do not have the same level of categorization [6] and learning a set of tasks containing only one task of interest and several complementary tasks which are used to bias the hypothesis space of the main task [9].

Another example of methods that improve the generalization performance is the “Family Discovery” method [7] that is used to build a parameterized family of models. This iterative method forces all the learning models to have parameters which are close to a mixture of manifolds defined in parameter space and that was derived by using the largest eigenvectors obtained in a principal component analysis of the parameters of the models. In [4], a related approach is introduced that forces the parameters of the learning models to lie on a surface defined in parameter space. The parameters representing the surface are learned along with the parameters defining the learning models. This approach, which is described in the next

section, was compared to the ‘‘Family Discovery’’ method on a set of classification tasks [4] and was showed to have significantly better generalization performance.

3 Domain-Specific Surface Learning

Domain-specific bias learning using models trained on similar tasks can be performed by re-parameterizing the parameters P_i of each learning model i as follows: $P_i = f(\theta, \alpha_i)$, where $P_i \in \mathbb{R}^{n_1}$, $\theta \in \mathbb{R}^{n_2}$, $\alpha_i \in \mathbb{R}^{n_3}$ with $n_1 \geq n_3$, and n_1 corresponds to the parameter space dimensionality [4]. Each learner i in this model is defined by a set of ‘‘private’’ parameters α_i which are transformed according to a chosen function f defined by a set of ‘‘shared’’ parameters θ . The ‘‘shared’’ parameters are updated using examples sampled from all the tasks while the ‘‘private’’ parameters are updated using examples only sampled from the corresponding task. The ‘‘learned’’ domain bias is coded in the parameters θ defining the function f . Thus both θ and α_i are simultaneously learned, using a gradient-based technique (the conjugate gradients method, in the experiments).

This framework can be used to restrict the size of the hypothesis space of the learning models by trying to identify in this space a surface embedding domain knowledge, and by forcing the models to only explore hypotheses that lie on this surface. In this context, f represents the shape of the surface whose position in the parameter space is defined by the learned parameters θ , and the parameters P_i of each model i represent a point lying on this surface (the position of P_i on the surface being represented by α_i).

The simplest surface form is an affine manifold. More complex choices include mixtures of affine manifolds and non-linear surfaces. In the absence of a priori knowledge that could guide the choice of the surface type, it is important to be aware of the trade-off between the complexity of the surface on the one hand and on the other hand, the number of available similar tasks, the size of the training sets defining the tasks as well as the complexity of the tasks to be learned. In [4], experimental results using an affine manifold are presented that show improved generalization performance for a set of classification tasks.

A bias learning system using an affine manifold can be described as follows: let $P_i = \{P_{ij}\}$, $j = 1, \dots, n_1$, represent the n_1 parameters of a model i and let θ , a $(n_1 - d) \times d$ matrix, represent the parameters defining the direction of an affine manifold whose dimensionality is equal to d ($d \leq n_1$), and β , a $n_1 \times 1$ vector, represent the offset of the manifold with respect to the origin of the parameter space. If the position of each parameter vector P_i on the affine manifold is represented by $\alpha_i = \{\alpha_{ik}\}$, $k = 1, \dots, d$, the relationship between the P_i parameters and their corresponding α_i position can be expressed as follows:

$$P_i = \tilde{\theta}\alpha_i + \beta \quad (1)$$

where $\tilde{\theta}$ is an $n_1 \times d$ matrix whose first d rows correspond to a $d \times d$ identity matrix and whose last $n_1 - d$ rows correspond to the θ matrix.

Besides being used to represent the offset of the manifold, β can also be used to enforce a proximity constraint between the P_i points. Indeed, β can serve as a reference point on the manifold towards which all the P_i points should move (see [4] for more details). In this case, the cost function associated with such a system can be expressed as follows:

$$\begin{aligned} E &= \frac{1}{M} \sum_{i=1}^M \frac{1}{K} \sum_{k=1}^K (g_i(P_i, X_{ik}) - Y_{ik})^2 + \lambda \sum_{i=1}^M \sum_{j=1}^{n_1} \left(\sum_{t=1}^d \tilde{\theta}_{jt} \alpha_{it} \right)^2 \\ &= \frac{1}{M} \sum_{i=1}^M \frac{1}{K} \sum_{k=1}^K \left(g_i(\tilde{\theta}\alpha_i + \beta, X_{ik}) - Y_{ik} \right)^2 + \lambda \sum_{i=1}^M \sum_{j=1}^{n_1} \left(\sum_{t=1}^d \tilde{\theta}_{jt} \alpha_{it} \right)^2 \end{aligned} \quad (2)$$

where M corresponds to the number of tasks trained simultaneously, K is the size of each training set, $\{X_{ik}, Y_{ik}\}, k = 1, \dots, K$ are the input and desired values for task i , $g_i(P_i, X_{ik})$ is the output produced by the learning model i for the input example X_{ik} using the parameters P_i and λ is a constant used to weight the proximity constraint.

4 Application to Option Pricing

4.1 European Options Data

The model described in the previous section was used for the prediction of the prices of European call options on the S&P 500 index. A European call option entitles the buyer of the option the right to buy a specific stock at a specific moment in the future at a smaller price equal to the *strike price* K agreed upon when the option was bought. In other terms, a person can buy at time t at the price p_t the right, but not the obligation, to buy a stock at time $t + \tau$ (where τ is called the *maturity*) at a price K . If at time $t + \tau$ the actual stock price $s_{t+\tau} > K$, the option buyer can *exercise* the option to pay K rather than $s_{t+\tau}$, thus making a profit of $s_{t+\tau} - K$. On the other hand, if $s_{t+\tau} < K$, the buyer would not exercise the option, and would make no profit (but in both cases the buyer would have initially made a “loss” of p_t for buying the option). The option price p_t depends on several factors including the current stock price s_t , the maturity τ , the strike price K , the underlying volatility (variability) of the price sequence, etc.

We worked on data ranging from 1987 to 1993 containing around 51 000 examples, each corresponding to the tuple (date t , s_t/K , τ , stock price variance estimated over 5, 22, 66, 250 and 1250 days, p_t/K), and we trained models to predict the ratio of the option price and the strike (p_t/K) using all or some of the remaining variables. Preliminary experiments were performed using neural networks trained on different periods of time in the range 1987/1993 and using different subsets of input variables. All the subsets contained s_t/K and τ . The difference between the subsets consisted in using all, one or no volatility variables. The results showed that any subset including the volatility computed over a period of 1250 days led to very poor prediction performance. Similar performance were sometimes obtained when using subsets of inputs including the volatility computed over 250 days. As for the remaining volatilities, we obtained similar results when using subsets of inputs including them, and when using the simple subset of inputs only containing the variables s_t/K and τ . Given those results, all the experiments presented in this paper were performed using only these 2 input variables.

4.2 Experimental Setup

The data were used to compare two training approaches differing in the way the training data were used to generate a prediction model: the first approach consists in training a single neural network using all the training data, while the second approach consists in dividing the training data in different subsets spanning different (possibly overlapping) periods of time and using the multi-task learning technique described in the previous section to simultaneously train several neural networks, each using a different subset of the training data. In that case, multi-task learning is used to analyze different periods of time differently and perform domain-knowledge transfer through time. Given that the data can be non-stationary, we argue that allowing a model to apply different analyses to different periods of time and exchange information about these periods can lead to a model with improved prediction performance. In what follows, we explain in detail the training procedure for each approach and then present results to compare them.

The single-task training method was tested according to the following procedure: 20 different feed-forward neural network architectures were selected, each having two input

units and one output unit. The architectures had different numbers of hidden layers and/or different numbers of units per hidden layer. Each time an experiment using a single neural network was performed, all 20 architectures were trained and for each architecture, the experiment was repeated twice using different initial parameters. For each experiment, model selection among the architectures was performed using the prediction performance of a validation set.

The experiments performed using the single-task learning method were chosen in such a way as to test its generalization performance on different periods of time, and for each period of time to test the effect of the size of the training set on the generalization performance. The training sets had sizes equal to 3, 6, 12 and 24 months. All validation sets were formed of the 3 months following the training set, and the test sets (on which generalization performance were evaluated) contained the 12 months following the validation set. The test sets were chosen to include a long period of time in order to compare the stability of the single-task method with that of the multi-task learning technique. Surprisingly, both methods had a relatively stable and consistent behavior over the 12 months period included in each test set. Five test years were chosen to evaluate the single-task method: 1989 to 1993. For each test year, single neural networks were trained using the data corresponding to October, November and December of the previous year as a validation set and the 3, 6, 12 or 24 months preceding October as a training set (for the test period corresponding to 1989, we didn't perform experiments using 24 months training sets because the data we were provided starts in 1987).

The multi-task learning method was tested as follows: 2 different feed-forward neural network architectures were chosen (among the set of 20 architectures used in the single-task method). One architecture had 25 parameters and the other one had 59 parameters. The number of parameters in an architecture corresponds to n_1 , the parameter space dimensionality defined in the previous section. When using the affine manifold learning method, decisions had to be made concerning the dimensionality d of the manifold and the "weight" λ of the proximity constraint. For a specific value of n_1 , several values of d were chosen, ranging between 1 and n_1 , in steps of 5. And for each value of d , 5 different values of λ were chosen including $\lambda = 0$ which corresponds to removing the proximity constraint. Each time an experiment was performed, the two chosen neural network architectures were used. For each architecture, all the above combinations of d and λ were tried. Model selection among all those combined choices was performed using the prediction performance of the validation set.

The experiments, using multi-tasks learning, were designed in such a way as to evaluate the effect of the number of tasks that are trained simultaneously and the effect of the sizes of the training set of each task on the generalization performance. Experiments were performed by simultaneously training 2, 3 or 4 tasks and for each such choice, using training sets of 3, 6 or 12 months for each task. As with the single-task method, the validation sets of the different tasks that were trained simultaneously correspond to 3 months of data and the test set is formed of 12 months. The same five test years used for the single-task method were used when applying the multi-task learning technique. Following is an example based on test year 1990 to explain the method used to choose the training and validation sets. In this example, 3 tasks are trained simultaneously. When using training sets of 3 months, the training and validation sets of task 1 respectively correspond to the time period 1/1989-3/1989 and 4/1989-6/1989, while the training and validation sets of task 2 respectively correspond to the period 4/1989-6/1989 and 7/1989-9/1989, and the training and validation sets of task 3 correspond to the period 7/1989-9/1989 and 10/1989-12/1989. For training sets of 6 months, the validation sets are kept fixed and the training sets are modified only by adding to them the 3 months that precede the current training dates. The corresponding training and validation sets therefore correspond to 10/1988-3/1989 and 4/1989-6/1989 for task 1, 1/1989-6/1989 and 7/1989-9/1989 for task 2 and 4/1989-9/1989 and 10/1989-

		N		
		2	3	4
M	3	-	0	+
	6	-	+	+
	12	0	+	+

Table 1: Analysis of variance results obtained when comparing the generalization performance of single models trained without bias learning with models obtained when using affine manifolds and simultaneously training N tasks, each having a training set corresponding to M months. The “-”, “0” and “+” signs respectively mean that the single model is significantly better, equivalent or worse than the bias learning method. The significance level was set to 5%.

12/1989 for task 3. When using training sets of 12 months, the training and validation test sets correspond respectively to 4/1988-3/1989 and 4/1989-6/1989 for task 1, 7/1988-6/1989 and 7/1989-9/1989 for task 2 and 10/1988-9/1989 and 10/1989-12/1989 for task 3. The validation sets were kept fixed for different sizes of the training sets to control the experiment setting. For training sets of 6 and 12 months, there is an overlap between the training sets associated to different tasks. We decided to explore this idea to verify if improved generalization performance could be obtained without having to increase the total amount of data used to train all the tasks. In the example presented for test year 1990, when training sets of 12 months are used, the total number of **different** training months is equal to 18 months (and not 36 months).

With the multi-task learning technique, how should we perform predictions on new data (i.e. on the test year)? Two options were possible: we could either combine the outputs of all the models trained simultaneously or we could use the model trained on the most recent data to make new predictions. We chose the second option and are exploring ways to apply the first one.

4.3 Results

In the first comparison tests of the two methods, a model selection based on the validation set performance was performed on the neural networks obtained using the single-task method with different sizes of the training sets. For each year, one network was chosen among the four networks trained on 3, 6, 12 and 24 months. The generalization performance of the chosen networks (for the test years 1989 to 1993) were compared to those obtained when using the multi-task learning method. An analysis of variance was performed whose results are presented in Table 4.3. Those results show that the generalization performance of the bias learning method improve when the number of tasks learned simultaneously increases and/or when the size of the training sets increases. In Figure 1, the generalization mean squared errors of the first method and those obtained with the bias learning method using 3 tasks and 6 months training sets are plotted. **Apart from year 1991, the multi-task learning method nearly always outperforms the single-task method, in particular in 1989 when the performance of the single-task method deteriorate drastically.**

In a second set of comparison tests, we decided to “favor” the single-task method by applying an analysis of variance to its generalization performance and choosing the training set size that leads to the best predictions. We found that using 24 months training sets leads to significantly better performance. So we compared models trained with 24 months using the single-task method with the bias learning method. The results obtained by performing an analysis of variance show that the bias learning technique trained with 3 or 4 tasks and 12 months training sets still manages to achieve significantly better performance, while the remaining models that led to significant improvements in Table 4.3 have now performance

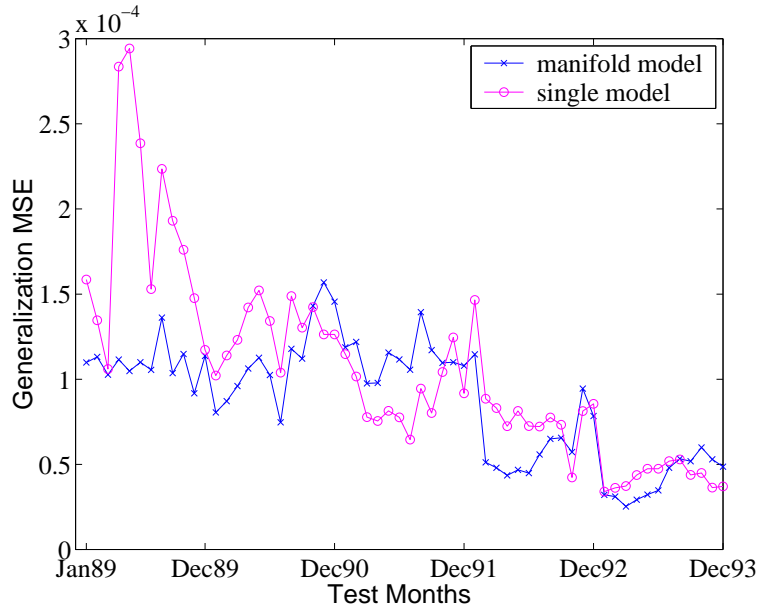


Figure 1: Generalization mean squared errors measured from January 1989 to December 1993 using on the one hand a model trained with no bias learning and on the other hand the affine manifold learning method trained on 3 tasks each using 6 months training sets.

similar to those of the single-task method. These results show that when using multi-task learning, it is possible to use a smaller amount of data (which is “recycled” and used several times) and manage to obtain results that are similar or better than those obtained when no bias is learned.

5 Future Work And Conclusion

When using the affine manifold learning method, only the model trained on the most recent data was used to estimate the ratio of the option price and the strike price. An alternative to this choice would be to combine (using a weighted average) the predictions of all the models that were trained simultaneously. Different combinations could be used that would consider the past generalization performance of each model, the length of time that separates the data that was used to train the model and the time at which predictions must be made.

We showed in this paper how multi-task learning can be applied to transfer knowledge through time. We also presented a framework that allows to recycle the data and use it in different contexts leading to different domain-related knowledge. Results presented on a real-world financial task show that multi-task learning leads to significantly improved performance.

References

- [1] J. Baxter. Learning model bias. In M. Mozer, D. Touretzky, and M. Perrone, editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1996.

- [2] R. Caruana. Algorithms and applications for multitask learning. In *The 13th International Conference on Machine Learning*, Bari, Italy, 1996.
- [3] J. Ghosh and Y. Bengio. Multi-task learning for stock selection. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9, pages 946–952. MIT Press, 1997.
- [4] J. Ghosh and Y. Bengio. Bias leaning, knowledge sharing. In *International Joint Conference on Neural Networks*, 2000. to appear.
- [5] T. Heskes. Solving a huge number of similar tasks: a combination of multi-task learning and a hierarchical bayesian approach. In *International Conference On Machine Learning*, 1998.
- [6] N. Intrator and S. Edelman. How to make a low-dimensional representation suitable for diverse tasks. *Connection Science, Special issue on Transfer in Neural Networks*, 1996.
- [7] S. M. Omohundro. Family discovery. In M. Mozer, D. Touretzky, and M. Perrone, editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1996.
- [8] L. Y. Pratt. Non-literal transfer among neural network learners. In R. J. Mammone, editor, *Artificial Neural Networks for Speech and Vision*, pages 143–169. Chapman and Hall, 1993.
- [9] D. L. Silver and R. E. Mercer. The parallel transfer of task knowledge using dynamic learning rates based on a measure of relatedness. *Connection Science, Special issue on Transfer in Inductive Systems*, 8(2):277–294, 1996.
- [10] S. Thrun and J. O’Sullivan. Discovering structure in multiple learning tasks: The tc algorithm. In *International Conference on Machine Learning*. Morgan Kaufmann, 1996.

Liste des publications au CIRANO*

Série Scientifique / *Scientific Series* (ISSN 1198-8177)

- 2002s-53 Multi-Task Learning For Option Pricing / J. Ghosn et Y. Bengio
- 2002s-52 Input Decay: Simple and Effective Soft Variable Selection / N. Chapados et Y. Bengio
- 2002s-51 On Out-of-Sample Statistics for Time-Series / F. Gingras, Y. Bengio et C. Nadeau
- 2002s-50 Forecasting Non-Stationary Volatility with Hyper-Parameters / Y. Bengio et C. Dugas
- 2002s-49 Cost Functions and Model Combination for VaR-based Asset Allocation using Neural Networks / N. Chapados et Y. Bengio
- 2002s-48 Experiments on the Application of IOHMMs to Model Financial Returns Series / Y. Bengio, V.-P. Lauzon et R. Ducharme
- 2002s-47 Valorisation d'Options par Optimisation du Sharpe Ratio / Y. Bengio, R. Ducharme, O. Bardou et N. Chapados
- 2002s-46 Incorporating Second-Order Functional Knowledge for Better Option Pricing / C. Dugas, Y. Bengio, F. Bélisle, C. Nadeau et R. Garcia
- 2002s-45 Étude du Biais dans le Prix des Options / C. Dugas et Y. Bengio
- 2002s-44 Régularisation du Prix des Options : Stacking / O. Bardou et Y. Bengio
- 2002s-43 Monotonicity and Bounds for Cost Shares under the Path Serial Rule / Michel Truchon et Cyril Tétédo
- 2002s-42 Maximal Decompositions of Cost Games into Specific and Joint Costs / Michel Moreaux et Michel Truchon
- 2002s-41 Maximum Likelihood and the Bootstrap for Nonlinear Dynamic Models / Sílvia Gonçalves, Halbert White
- 2002s-40 Selective Penalization Of Polluters: An Inf-Convolution Approach / Ngo Van Long et Antoine Soubeyran
- 2002s-39 On the Mediational Role of Feelings of Self-Determination in the Workplace: Further Evidence and Generalization / Marc R. Blais et Nathalie M. Brière
- 2002s-38 The Interaction Between Global Task Motivation and the Motivational Function of Events on Self-Regulation: Is Sauce for the Goose, Sauce for the Gander? / Marc R. Blais et Ursula Hess
- 2002s-37 Static Versus Dynamic Structural Models of Depression: The Case of the CES-D / Andrea S. Riddle, Marc R. Blais et Ursula Hess
- 2002s-36 A Multi-Group Investigation of the CES-D's Measurement Structure Across Adolescents, Young Adults and Middle-Aged Adults / Andrea S. Riddle, Marc R. Blais et Ursula Hess

* Consultez la liste complète des publications du CIRANO et les publications elles-mêmes sur notre site Internet :

- 2002s-35 Comparative Advantage, Learning, and Sectoral Wage Determination / Robert Gibbons, Lawrence F. Katz, Thomas Lemieux et Daniel Parent
- 2002s-34 European Economic Integration and the Labour Compact, 1850-1913 / Michael Huberman et Wayne Lewchuk
- 2002s-33 Which Volatility Model for Option Valuation? / Peter Christoffersen et Kris Jacobs
- 2002s-32 Production Technology, Information Technology, and Vertical Integration under Asymmetric Information / Gamal Atallah
- 2002s-31 Dynamique Motivationnelle de l'Épuisement et du Bien-être chez des Enseignants Africains / Manon Levesque, Marc R. Blais, Ursula Hess
- 2002s-30 Motivation, Comportements Organisationnels Discrétionnaires et Bien-être en Milieu Africain : Quand le Devoir Oblige / Manon Levesque, Marc R. Blais et Ursula Hess
- 2002s-29 Tax Incentives and Fertility in Canada: Permanent vs. Transitory Effects / Daniel Parent et Ling Wang
- 2002s-28 The Causal Effect of High School Employment on Educational Attainment in Canada / Daniel Parent
- 2002s-27 Employer-Supported Training in Canada and Its Impact on Mobility and Wages / Daniel Parent
- 2002s-26 Restructuring and Economic Performance: The Experience of the Tunisian Economy / Sofiane Ghali and Pierre Mohnen
- 2002s-25 What Type of Enterprise Forges Close Links With Universities and Government Labs? Evidence From CIS 2 / Pierre Mohnen et Cathy Hoareau
- 2002s-24 Environmental Performance of Canadian Pulp and Paper Plants : Why Some Do Well and Others Do Not ? / Julie Doonan, Paul Lanoie et Benoit Laplante
- 2002s-23 A Rule-driven Approach for Defining the Behavior of Negotiating Software Agents / Morad Benyoucef, Hakim Alj, Kim Levy et Rudolf K. Keller
- 2002s-22 Occupational Gender Segregation and Women's Wages in Canada: An Historical Perspective / Nicole M. Fortin et Michael Huberman
- 2002s-21 Information Content of Volatility Forecasts at Medium-term Horizons / John W. Galbraith et Turgut Kisinbay
- 2002s-20 Earnings Dispersion, Risk Aversion and Education / Christian Belzil et Jörgen Hansen
- 2002s-19 Unobserved Ability and the Return to Schooling / Christian Belzil et Jörgen Hansen
- 2002s-18 Auditing Policies and Information Systems in Principal-Agent Analysis / Marie-Cécile Fagart et Bernard Sinclair-Desgagné
- 2002s-17 The Choice of Instruments for Environmental Policy: Liability or Regulation? / Marcel Boyer, Donatella Porrini
- 2002s-16 Asymmetric Information and Product Differentiation / Marcel Boyer, Philippe Mahenc et Michel Moreaux
- 2002s-15 Entry Preventing Locations Under Incomplete Information / Marcel Boyer, Philippe Mahenc et Michel Moreaux
- 2002s-14 On the Relationship Between Financial Status and Investment in Technological Flexibility / Marcel Boyer, Arnel Jacques et Michel Moreaux

- 2002s-13 Modeling the Choice Between Regulation and Liability in Terms of Social Welfare / Marcel Boyer et Donatella Porrini
- 2002s-12 Observation, Flexibilité et Structures Technologiques des Industries / Marcel Boyer, Armel Jacques et Michel Moreaux
- 2002s-11 Idiosyncratic Consumption Risk and the Cross-Section of Asset Returns / Kris Jacobs et Kevin Q. Wang
- 2002s-10 The Demand for the Arts / Louis Lévy-Garboua et Claude Montmarquette
- 2002s-09 Relative Wealth, Status Seeking, and Catching Up / Ngo Van Long, Koji Shimomura
- 2002s-08 The Rate of Risk Aversion May Be Lower Than You Think / Kris Jacobs
- 2002s-07 A Structural Analysis of the Correlated Random Coefficient Wage Regression Model / Christian Belzil et Jörgen Hansen
- 2002s-06 Information Asymmetry, Insurance, and the Decision to Hospitalize / Åke Blomqvist et Pierre Thomas Léger
- 2002s-05 Coping with Stressful Decisions: Individual Differences, Appraisals and Choice / Ann-Renée Blais
- 2002s-04 A New Proof Of The Maximum Principle / Ngo Van Long et Koji Shimomura
- 2002s-03 Macro Surprises And Short-Term Behaviour In Bond Futures / Eugene Durenard et David Veredas
- 2002s-02 Financial Asset Returns, Market Timing, and Volatility Dynamics / Peter F. Christoffersen et Francis X. Diebold
- 2002s-01 An Empirical Analysis of Water Supply Contracts / Serge Garcia et Alban Thomas
- 2001s-71 A Theoretical Comparison Between Integrated and Realized Volatilities Modeling / Nour Meddahi
- 2001s-70 An Eigenfunction Approach for Volatility Modeling / Nour Meddahi
- 2001s-69 Dynamic Prevention in Short Term Insurance Contracts / M. Martin Boyer et Karine Gobert
- 2001s-68 Serial Cost Sharing in Multidimensional Contexts / Cyril Tétédo et Michel Truchon
- 2001s-67 Learning from Strike / Fabienne Tournadre et Marie-Claire Villeval
- 2001s-66 Incentives in Common Agency / Bernard Sinclair-Desagné
- 2001s-65 Detecting Multiple Breaks in Financial Market Volatility Dynamics / Elena Andreou et Eric Ghysels
- 2001s-64 Real Options, Preemption, and the Dynamics of Industry Investments / Marcel Boyer, Pierre Lasserre, Thomas Mariotti et Michel Moreaux
- 2001s-63 Dropout, School Performance and Working while in School: An Econometric Model with Heterogeneous Groups / Marcel Dagenais, Claude Montmarquette et Nathalie Viennot-Briot
- 2001s-62 Derivatives Do Affect Mutual Funds Returns : How and When? / Charles Cao, Eric Ghysels et Frank Hatheway
- 2001s-61 Conditional Quantiles of Volatility in Equity Index and Foreign Exchange Data / John W. Galbraith, Serguei Zernov and Victoria Zinde-Walsh