Journal of Business Valuation and Economic Loss Analysis

Volume 4, *Issue* 1 2009 *Article* 4

Forecasting the Periodic Net Discount Rate with Genetic Programming

Neal F. Wagner, *SolveIT Software Pty Ltd*
Mark A. Thompson, *Texas Tech University*

Recommended Citation:
Wagner, Neal F. and Thompson, Mark A. (2009) "Forecasting the Periodic Net Discount Rate with Genetic Programming," *Journal of Business Valuation and Economic Loss Analysis*: Vol. 4 : Iss. 1, Article 4.
Available at: http://www.bepress.com/jbvela/vol4/iss1/art4
DOI: 10.2202/1932-9156.1072

Forecasting the Periodic Net Discount Rate with Genetic Programming

Neal F. Wagner and Mark A. Thompson

Abstract

This paper examines the periodic net discount rate using genetic programming (GP) techniques to build better short-term forecasts. Standard GP techniques require human judgment as to which data window to use, which may be problematic due to structural breaks and persistence (or long memory) in the net discount rate. We use a recently developed extension of GP to overcome this problem. While our results show no significant out-of-sample forecast improvement relative to the linear alternative or random walk model over the full sample, they do provide evidence as to the stochastic nature of the net discount rate considering the AR(3) model yielded lower forecasting errors in the post-1982 sample.

KEYWORDS: periodic net discount rate, genetic programming, forecasts

Introduction

Forensic economists are called upon to estimate the present value of lost future earnings in personal injury, wrongful death, and discrimination cases as well as valuing future medical needs. A net discount rate approach is often used; however, there continues to be much debate in the forensic economics literature as to the time series properties of the net discount rate. The underlying issue is whether the net discount rate is stationary or not. A stationary time series implies that a random shock to the series is temporary as the series will revert back to the (long-run) mean.[1] Hence, a non-stationary series will not revert back to its long-run mean following a random shock. As such, the use of historical data to calculate net discount rates could be misleading depending on the time series properties of the series under investigation.

In this paper, we are interested in the overall predictive capability of different models. In particular, we utilize advanced statistical and modeling techniques such as Genetic Programming (GP) to assist in specifying the model. Since the end result is to have a better forecast of the net discount rate to use, we compare the out-of-sample forecasting performance of an autoregressive (AR) model and a Random Walk (RW) model. If the series is stationary, then the AR model may perform well; if the series is nonstationary, then the RW model may perform well as historical data may be of little value. The literature surrounding the stationarity of net discount rates is mixed. In addition, structural breaks and persistence in the time series add to the problems of whether historical data can or should be used in valuation studies.

The next section discusses the literature on testing for unit roots and the time series properties of net discount rates. Section 3 provides a brief discussion of GP techniques, while section 4 describes the data and forecasting experiment. The results are discussed in section 5 with concluding remarks in section 6.

Literature Review

Early studies on the net discount rate examine the time series properties of the components in constructing the rate (i.e., interest rates and earnings growth).[2]

[1] A time series that is stationary has a constant mean and variance over time and the covariance between the two time periods is a function of the distance of the time periods. A non-stationary time series is one that contains a unit root. A time series that is nonstationary often is first-differenced to render it stationary. See Enders (2004) for further information on unit roots.

[2] For a good survey of the empirical literature on testing unit roots in the net discount rate, see Payne (2007).

However, this research led to examining the stationarity of the net discount rate. Braun et al. (2005) highlights the importance of this issue by stating that, "The practice of calculating economic losses using a historical measure is implicitly based on the assumption that net discount rates are a stationary or mean reverting process." They go on to say the use of the historical average of the net discount rate is valid when the series is stationary. If the net discount rate is nonstationary, then the current value may be best to use in present value calculations.

Another problem is the role of structural breaks in the respective series. Such breaks can be problematic when testing for unit roots. Perron (1989) indicates that standard unit root testing procedures may suggest the existence of a unit root when the series may be stationary around a shift in the mean. Gamber and Sorensen (1994) apply the concept to the net discount rate and point out the implications of using a historical average over different time periods. After a shift in the net discount rate, the appropriate historical average would only include the data since the shift and not over the entire sample. The structural break issue led to several studies examining the net discount rate over different time periods. The consensus is that there was a major structural break around late 1979 (Bonham and LaCroix, 1992; Gamber and Sorensen, 1993, 1994; Haslag et al., 1994; Payne et al., 1999; Hays et al., 2000; Braun et al., 2005). Such a break must be taken into account when examining the properties of the net discount rate. In particular, Gamber and Sorensen find that the net discount rate follows a stationary process around a shifting mean using the Perron (1989) unit root test. The structural break corresponds to a change in the Federal Reserve's operating procedure. Gamber and Sorensen (1994) further support this finding using monthly data with a structural break in October 1979 (and January 1980 for the other net discount rates).

However, some of the findings question whether a permanent shift took place or whether the time series exhibits a higher degree of persistence. For example, Pelaez (1996) examines the persistence in the periodic net discount rate by measuring the extent to which a times series will deviate from its mean following a random shock using the Cochrane (1988) variance ratio test. Using a monthly net discount rate for manufacturing hourly earnings and the 1-year U.S. Treasury note, Pelaez (1996) finds evidence that the periodic net discount rate is stationary. Payne et al. (1999) extend the Pelaez (1996) study for eight employment sectors and also find evidence that the periodic net discount is stationary. Using the Campbell-Mankiw (1987) along with the Cochrane (1988) test, Hays et al. (2000) examine the degree of persistence across a variety of net discount rates. They conclude the use of historical real net discount ratios can be used for present value calculations for longer time spans as real net discount ratios are stationary over those longer time periods.

DOI: 10.2202/1932-9156.1072

Finally, the latest research examines the net discount rate using endogenous-break unit root tests. Sen et al. (2002) investigate the periodic net discount rate using high grade municipal bonds (as oppose to Treasury bond yields) with the Sen (2003) unit root test. The Sen (2003) unit root test allows for the break and the form of the break to be unknown. Sen et al. (2002) find the periodic net discount rate, when considering a structural (unknown) break in 1981, rejects the null hypothesis of a unit root. Braun et al. (2005), however, report evidence that the net discount rate contains a unit root even when allowing for two structural breaks (identified as 1979 and 1981).

While the results on the stationarity of the net discount rate are mixed, our approach uses GP techniques to examine the historical data for purposes of generating short-term forecasts. In particular, we allow the functional form of the model to be discovered within the historical data. This technique may better handle the problems that previous researchers have had in examining the time series properties of the net discount rate. That is, the unknown structural breaks and persistence (or long memory) issues are problematic for unit root tests. As a result, the AR model may be misspecified and lead to forecasting errors. The GP techniques may be better suited to deal with these econometric issues and improve out-of-sample forecasts of the net discount rate. The next section provides a brief overview of the GP technique applied to time series forecasting as well as a new extension, Dynamic Forecasting Genetic Programming (DFGP).

Forecasting with Genetic Programming

Evolutionary Computation (EC) refers to the use of the Darwinian theory of natural selection as a model for building computer programs that "evolve" solutions to problems rather than have the procedure for reaching a solution explicitly designated by the programmer.[3] In EC, a population of candidate solutions to a particular problem is generated and then ranked based on their proximity to the optimal solution. Higher-ranking solutions are used to produce new populations of candidates through genetic operators of "mating" or "mutation." This procedure is repeated until an optimal solution is found or some maximum number of iterations is reached.

Genetic Programming (GP) is a sub-category of EC invented by Koza (1992) in which solutions are represented as tree structures. When applied to time series forecasting, GP solution trees represent mathematical expressions where internal nodes are mathematical operators and leaf nodes are explanatory

[3] Please see Back et al. (1997) and Eiben and Smith (2003) for an overview of EC.

variables or constants. Figure 1 shows an example GP forecasting solution. The figure represents the expression $x_{t1} + sin(5.31x_{t2})$.

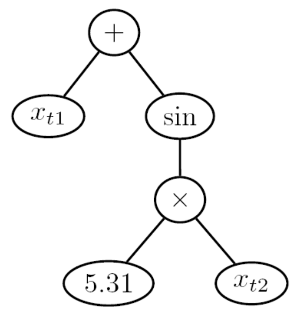

Figure 1: GP Forecasting Solution

A random population of candidate solution trees is initially created and the solutions are ranked based on their prediction error over a set of training data. Higher ranking solutions (i.e., those with smaller error) are selected for reproduction into the next generation via mating or mutation. Mating is accomplished by swapping sub-trees from two parent trees to produce two offspring trees. Figure 2 illustrates this operation. Parent trees *p1* and *p2* produce offspring trees *o1* and *o2*. In the figure, the dotted lines represent sub-trees to be swapped. Mutation is accomplished by replacing a sub-tree from a parent tree with a randomly-generated tree.

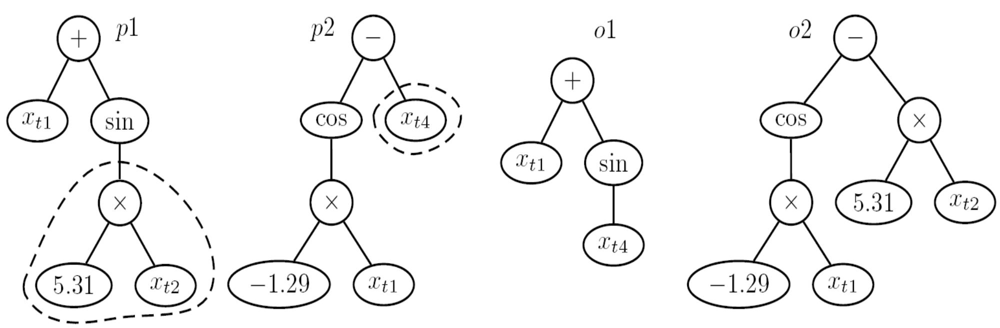

Figure 2: GP Mating Operation

Numerous studies apply GP to time series forecasting with favorable results. Some examples of these include Andrew and Prager (1994), Chen and Yeh (1996, 1998), Jonsson and Barklund (1996), Oakley (1997), Hiden et al

DOI: 10.2202/1932-9156.1072

(1998), Iba and Sasaki (1999), Iba and Nicholaev (2000), Koubadoun (2000, 2003), Neely and Weeler (2002), and Wagner et al (2007).

Because GP builds non-linear models, it is able to capture many aspects displayed by real-world data. GP also has the advantage that the forecasting model need not be prescribed, allowing for automatic discovery of a relevant functional form. This is especially important when an efficient forecasting model is not known. GP produces a model built empirically from data rather than a pre-determined model with empirically-fitted parameters.

One issue that a forecaster must decide on regardless of the method is selecting the training data. Typically, some set of historical data is chosen for training and is used in building a forecasting model. The problem with this practice can be seen when the underlying data-generating process of the time series is nonstatic. GP, as well as most statistical methods, assumes a static underlying process. If the underlying process shifts, the methods must be re-evaluated in order to accommodate the new process. Additionally, these methods require the number of historical time series data for analysis is designated *a priori*. This presents a problem in non-static environments because different segments of the time series may have different underlying data-generating processes. For example, a time series representing the daily stock value of a major U.S. airline is likely to have a different underlying process before September 11, 2001 than it does afterwards. If time series data span more than one underlying process, forecasting models built using that data may have unduly large errors.

$$\underbrace{22, 33, 30, 27, 24, 20, 21, 20, 20}_{\text{segment1}}, \underbrace{23, 26, 29, 30, 28, 29, 32, 30, 31}_{\text{segment2}} | \underbrace{\ldots}_{\text{future}}$$

Figure 3: Segments with Different Underlying Processes

Consider the subset of time series data shown in Figure 3. Suppose this represents the most recent historical data and has been chosen for training. Furthermore, the subset consists of two segments each with a different underlying process. The second segment represents the current environment and is valid for forecasting future data, while the first segment represents an older environment that no longer exists. Because both segments are utilized for training, the forecasting model is distorted. This situation highlights the need for forecasting methods that can automatically determine the correct training "window" (i.e., the correct number of historical data to be used for training). In Wagner et al. (2007), a new extension to GP, called Dynamic Forecasting GP (DFGP), addresses this issue.

In DFGP, two novel heuristics are employed that seek to automatically adjust to shifts in the underlying process. The first heuristic is a sliding window of time. Training starts at the beginning of the available historical data. Some initial window size (number of data observations to analyze) is set and several generations of GP are run to evolve a population of solutions. Then, the data window slides to include the next time series observation. Several generations are run with the new data window and then the window slides again. This process is repeated until all available data have been trained on up to and including the most recent historical data. Figure 4 illustrates this process and the vertical bar marks the end of available historical data. The set of several generations run on a single analysis (training) window is referred to as a "dynamic generation." Thus, a single run of the DFGP includes several dynamic generations (one for each window slide) on several different consecutive analysis windows. This sliding window feature allows the DFGP to analyze all existing data and potentially take advantage of previously observed patterns.

$$\underbrace{22,33,30,27,24,20,21,20,20}_{\text{window}-1},23,26,29,30,28,29,32,30,31|\underbrace{\ldots}_{\text{future}}$$

$$22,\underbrace{33,30,27,24,20,21,20,20,23}_{\text{window}-2},26,29,30,28,29,32,30,31|\underbrace{\ldots}_{\text{future}}$$

$$\vdots$$

$$22,33,30,27,24,20,21,20,20,\underbrace{23,26,29,30,28,29,32,30,31}_{\text{window}-\text{i}}|\underbrace{\ldots}_{\text{future}}$$

Figure 4: A Sliding Data Analysis

The second heuristic DFGP employs is the adaptive windowing technique. In this technique, not one but two initial windows are specified (one larger and one smaller) at the beginning of available historical data. Both windows undergo GP training and then each is used to predict a small number of future (as yet untrained upon) data. The sizes of the two windows adjust in the direction determined by the window with the better prediction. That is, if the smaller window predicts better, then both windows are reduced in size. Alternatively, if the larger window predicts better, then both windows increase in size. This adjustment takes place at each slide of the window. Once DFGP analyzes a sufficient number of historical data, the training windows will have “honed in” on the currently active process. As shown in Wagner et al. (2007) and Wagner and Michalewicz (2008), this procedure can automatically adjust the training window correctly in the presence of shifting underlying processes.

Data and Forecasting Experiments

We define the periodic net discount rate as follows:

(1) $\text{NDR} = (i - \text{g})/(1 + \text{g})$

where g is the annual earnings growth rate and i denotes the 1-year constant maturity U.S. Treasury security. The use of the periodic net discount rate allows us to compare the time series behavior of the net discount rate in such previous studies as Pelaez (1996), Payne et al. (1999), and Sen et al. (2002). A plot of the period net discount rate over the full sample of January 1965 to September 2007 is shown in Figure 5.

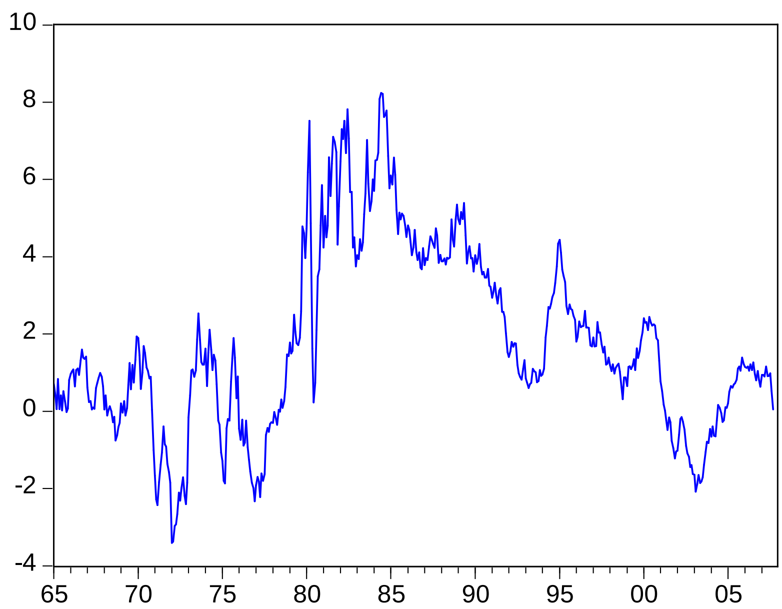

Figure 5. Periodic Net Discount Rate

The mean periodic net discount rate over the full sample is 1.75 with a standard deviation of 2.35. As can be seen in Figure 5 and discussed in section 2, the net discount rate is widely considered to contain structural breaks occurring around 1979. A structural break is synonymous with a shift in the underlying data-generating process and, thus, the use of DFGP (described in the previous section) is appropriate. We investigate both the use of standard GP and DFGP in forecasting the periodic net discount rate using historical data after the structural break of 1979-1982 as well as prior to the structural break. The idea is to see how GP and DFGP react to the structural break and whether using data before the break for training is beneficial or harmful for forecasting the series in periods after the break.

In the first set of forecasting experiments, we use historical data from January 1982 through September 2002 (i.e., post-1982 break sample) for training both the GP and DFGP models. Forecasts are generated for a 5-year period from October 2002 through September 2007. For the second set of experiments, we use historical data from January 1965 through September 2002 (i.e., full sample) to train the models and generate forecasts for the same 5-year period. For all GP experiments, models are built using 3 autoregressive lags of the periodic net discount rate and one-step-ahead forecasts are produced. Table 1 reports the out-of-sample forecasting results for several different models using the post-1982 break sample and full sample. Along with GP models, we estimate an AR(3) model and a Random Walk (RW) model for comparison. If the periodic net discount rate is stationary as supported by Pelaez (1996), Payne et al. (1999), and Sen et al. (2002), then the AR(3) model may perform well out-of-sample. However, if the periodic net discount rate is nonstationary as supported by Braun et al. (2005), then the RW model may perform well since historical data is of no value as the best forecast is last period's rate.

Table 1. Forecasting Errors

	Post-1982 Break Sample	*Full Sample*
AR(3)	4.90	5.01
RW	5.00	5.00
GP	5.19	5.00
DFGP	5.27	5.37

Notes: The forecast error measure is the mean squared error (denoted as 10^{-2}) of the 60 one-month-ahead observations to examine the forecast accuracy of the different models.

Since the GP algorithm is stochastic in nature and results vary from one run to another, we execute 20 runs and then average the 20 forecasts to generate a single out-of-sample forecast for each time period. An alternative approach by Kaboudan (2001, 2003) is to use the best run as the result. However, in a real

DOI: 10.2202/1932-9156.1072

forecasting application this practice is not useful since the actual value is unknown (i.e., one cannot know which run produces the best forecast for a given time period without first knowing the corresponding actual value of that time period).

Discussion of Results

There are several interesting results from the forecasting experiment. First, the AR(3) model out-performs the other models in the post-1982 sample. This finding tends to support the stochastic nature of the periodic net discount rate. However, we expected the model performance to be considerably better using just the post-1982 sample compared to the full sample (including the break) for the AR(3) model and the standard GP model. Second, the standard GP model actually improved its short-term forecasts having trained on the full sample. The results indicate the structural break did not harm the efficiency of the standard GP model (over the full sample). For the DFGP model, we also did not expect a difference in forecasting performance between the post-1982 break sample and the full sample as DFGP automatically adjusts its training window to focus on data corresponding to the currently active underlying process. To our surprise, the DFGP model is not appropriate for short-term forecasts as it has the highest mean square error in the post-1982 sample as well as the full sample.

The RW model did just as well over the full sample as the more complex GP and linear AR(3) model. While the GP model did not yield substantial better forecasts out-of-sample, there may be several advantages to their use in modeling net discount rates. The literature documents mixed results as to the stationarity of the net discount rate. Some researchers find issues of structural breaks and long memory leading to concerns as to the appropriate functional form. While our findings provide additional support of the stochastic nature of the periodic net discount rate, GP techniques do not require a specified functional form as it lets the data decide in an effort to provide the best forecast. The results indicate the use of the GP model was no worse than the linear alternative or RW model over the full sample (i.e., including the structural break).

Concluding Remarks

This paper sets out to improve our understanding of how the periodic net discount rate behaves over time in an effort to yield better short-term forecasts. In pursuit of this goal, we examine the time series behavior of the net discount rate using GP techniques. However, standard GP techniques require human judgment as to the data window to use. This has been shown to be a problem in the literature with the possibility of structural breaks and persistence (or long memory). To overcome this problem, we use a recently developed extension of GP, known as DFGP.

The DFGP model is conceptually pleasing as there is no clear direction to model the periodic net discount rate in the forensic economics literature. However, our results show that while GP models have several advantages in dealing with econometric issues of the net discount rate, there is no significant out-of-sample forecast improvement relative to the linear alternative or random walk model over the full sample. These findings also provide evidence of the stochastic nature of the periodic net discount rate considering the AR(3) model yielded lower forecasting errors in the post-1982 sample and similar to the other models over the full sample.

References

Andrew, M. and R. Prager (1994) Genetic programming for the acquisition of double auction market strategies. In K. Kinnear (ed.), *Advances in Genetic Programming*, MIT Press, 355-368.

Bäck, T., D. Fogel, Z. Michalewicz (1997) *Handbook of Evolutionary Computation*, Oxford University Press.

Bonhan, C. and L. Sumner (1992) Forecasting earnings growth and discount rates: New evidence from time series analysis. *Journal of Forensic Economics*, 5, 221-231.

Braun, B., J. Lee, and M. Strazicich (2005) Historical net discount rates and future economic losses: Refuting the common practice. In R. Kaufman, J. Rodgers, and G. Martin, *Economic Foundations of Injury and Death Damages*, Edward Elgar Publishing: United Kingdom, 468-491.

Campbell, J. and N.G. Mankiw (1987) Permanent and transitory component in macroeconomic fluctuations. *American Economic Review*, 77, 111-117.

Chen S. and C. Yeh (1996) Toward a computable approach to the efficient market hypothesis: An application of genetic programming. *Journal of Economics Dynamics and Control*, 21, 1043–1063.

DOI: 10.2202/1932-9156.1072

Chen S., C. Yeh, and W. Lee (1998) Option pricing with genetic programming. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, 32–37.

Cochrane, J. (1988) How big is the random walk in GDP? *Journal of Political Economy,* 96, 893-920.

Eiben, A. and J. Smith (2003) *Introduction to Evolutionary Computing*, Springer.

Enders, W. (2004) *Applied Econometric Time Series*, 2nd ed., John Wiley & Sons: New York.

Gamber, E. and R. Sorensen (1993) On testing for the stability of the net discount rate. *Journal of Forensic Economics*, 7, 69-79.

Gamber, E. and R. Sorensen (1994) Are net discount rates stationary? Implications for present value calculations: Comment. *Journal of Risk and Insurance*, 61, 503-512.

Haslag, J., M. Nieswiadomy, D. Slottje (1994) Are net discount rates stationary: Some further evidence. *Journal of Risk and Insurance*, 61, 513-518.

Hays, P., M. Schrieber, J. Payne, B. Ewing, and M. Piette (2000) Are net discount ratios stationary? Evidence of mean reversion and persistence. *Journal of Risk and Insurance,* 67, 439-449.

Hiden H., B. McKay, M. Willis, and M. Tham (1998) Non-linear partial least squares using gentic programming. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, 128–133.

Iba, H. and N. Nikolaev (2000) Genetic programming polynomial models of financial data series. In *Proceedings of the 2000 Congress of Evolutionary Computation*, 1459–1466.

Iba, H. and T. Sasaki (1999) Using genetic programming to predict financial data. In *Proceedings of the Congress of Evolutionary Computation*, 244–251.

Jonsson, P. and J. Barklund (1996) Characterizing signal behavior using genetic programming. *Evolutionary Computing: Lecture Notes in Computer Science*, 1143, 62–72.

Kaboudan, M. (2000) Genetic programming prediction of stock prices. *Computational Economics*, 6, 207–236.

Kaboudan, M. (2001) Genetically evolved models and normality of their residuals. *Journal of Economics Dynamics and Control*, 25, 1719-1749.

Kaboudan, M. (2003) Forecasting with computer-evolved model specifications: a genetic programming application. *Computers and Operations Research*, 30, 1661–1681.

Koza, J. (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press.

Neely, C. and P. Weller (2002) Predicting exchange rate volatility: genetic programming versus GARCH and RiskMetrics. The Federal Reserve Bank of St. Louis *Review*, 84, 43-54.

Oakley, H. (1997) Genetic programming for nonlinear equation fitting to chaotic data. In T. Back, D. Fogel, and Z. Michalewicz (eds.), *Handbook of Evolutionary Computation*, Oxford University Press.

Payne, J. (2007) Testing for a unit root in the net discount rate: A survey of the empirical results. *Journal of Business Valuation and Economic Loss Analysis*, 2, Article 4.

Payne, J., B. Ewing, and M. Piette (1999) Mean reversion in net discount rates. *Journal of Legal Economics*, 9, 69-80.

Pelaez, R. (1996) Mean-reversion in the net discount rate: Evidence from the manufacturing sector. *Journal of Legal Economics*, 6, 29-39.

Perron, P. (1989) The Great Crash, the oil price shock and the unit root hypothesis. *Econometrica*, 57, 1361-1401.

Sen, A., G. Gelles, and W. Johnson (2002) Structural instability in the net discount rate series based on high grade municipal bond yields. *Journal of Legal Economics*, 12, 87-100.

Sen, A. (2003) On unit root tests when the alternative is a trend-break stationary process. *Journal of Business and Economic Statistics*, 21, 174-184.

DOI: 10.2202/1932-9156.1072

Wagner N., Z. Michalewicz, M. Khouja, and R. McGregor (2007) Time series forecasting for dynamic environments: The DyFor genetic program model. *IEEE Transactions on Evolutionary Computation*, 11, 433-452.

Wagner N. and Z. Michalewicz (2008) An analysis of adaptive windowing for times series forecasting in dynamic environments: Further tests of the DyFor GP model. *Proceedings of the Genetic and Evolutionary Computation Conference*.