Statistical modeling of stochastic level shifts in time series

Thomas M. Trimbur
Statistical Research Division (F.O.B. 4 - 3228)
U.S. Census Bureau
Washington, D.C. 20023
Email : Thomas.M.Trimbur@census.gov

Abstract

Methods are developed for estimating trends in time series subject to level shifts. The approach is based on specifying stochastic models for breaks as part of the model structure, using heavy-tailed densities to allow for a positive probability of such a large change at any given time. Examining changes in trend movements, estimated from the dynamics of the dataset, provides more information than a yes/no criterion for making decisions on level shift events. Continuous-valued innovations in the trend are assessed using a statistical model; with the arrival of a data point that constitutes a break, timely warning is given with a smooth shift in the assessment. The empirical illustrations show how more robust trend estimates are obtained in practice.

KEY WORDS : Level shifts, Unobserved components, Heavy-tailed density, Non-Gaussian model, Robustness, Trend estimation

1. Introduction

Many economic series are subject to occasional sudden, large changes that have a lasting impact. Widespread recognition of the importance of level shifts or structural breaks in time series has motivated a great deal of econometric research. One approach is to assume that the breaks have occurred at fixed points in time. Bai (1997) develops a testing procedure for the case where there may be multiple breakpoints in the sample. In this paper, I investigate models where breaks in the trend occur stochastically. This is an appealing feature, as in practice, except in situations of planned intervention, the timing of large shifts in the level of a process is unknown in advance.

By using a continuous classification for large changes, the method also allows for stochastic variation in the magnitude and direction of shifts. In this article I implement unobserved components models where the trend innovation is assigned a continuous heavy-tailed probability density. This means that large shifts in the level, whose timing and size are stochastic, are directly specified as part of the model. Aston and Koopman (2003) provide an analogous treatment of series with outlying observations, showing how estimates of seasonal components may be made more robust, compared to binary procedures for outlier detection and removal.

The rest of this paper develops the framework for trend estimation with stochastic level shifts and provides empirical illustrations. The methods are straightforward to implement and general formulations are discussed in Durbin and Koopman (2001). The calculations are performed by a computer program written in the Ox language of Doornik (1999) and rely on the SsfPack library of routines of Koopman et. al (1999). The aim is to make efficient use of state space methods available for the wide range of linear Gaussian time series models, by embedding them in simulation routines designed to handle the non-Gaussian extensions.

Below, I discuss the details of the estimation method in the important special case of the local level model with t-disturbances for the trend. The computational strategy focusses on the use importance sampling in a multivariate setting, to handle the intractable and high-dimensional non-Gaussian density of the series of smoothed components. Explanation of the robust local level model case helps one understand the basic issues.

Section 2 first reviews the most basic model for trend analysis and sets the foundation for the application of the Student-t class of probability distributions. In Section 2.1 the estimation of the smoothed trend series and the likelihood function are discussed for non-Gaussian models. Section 2.2 introduces the use of the Student-t density for the trend error process within the local level model. The

---

* **Disclaimer.** This report is released to inform interested parties of ongoing research and to encourage discussion of work in progress. Any views expressed on statistical or methodological issues are those of the author and not necessarily those of the U.S. Census Bureau.

approach is illustrated in Section 3; computational issues and simulation results are discussed, and I investigate some real time series. Section 4 concludes.

2. Non-Gaussian local level model

Interest centers on the use of heavy-tailed distributions to model level shifts, and I start with the most basic framework. A simple model for a series subject to nonstationary trend and irregular movements is the local level model:

$$y_t = \mu_t + \varepsilon_t, \qquad t = 1, ..., T \tag{1}$$
$$\mu_{t+1} = \mu_t + \eta_t, \qquad t = 1, ..., T$$

where ε_t is serially independent with mean zero and constant (finite) variance σ^2_ε, and likewise for η_t. In the Gaussian case, this is equivalent to assuming white noise error terms, that is, a process whose values are uncorrelated at different time periods. However, in using more general distributions, the serial independence condition is stronger than the white noise assumption. In the case of uncorrelated non-Gaussian processes, there may be dependencies over time in the error process that are difficult to interpret. In assuming serially independent disturbances, it becomes clear that for each innovation, the values at different points in time are unrelated.

I further assume that ε_t and η_t are independent of each other contemporaneously and across all time periods. The vector of sample observations is written as $\mathbf{y} = (y_1, ..., y_T)'$, while the series of values for the level is arranged in a $T \times 1$ vector denoted by $\boldsymbol{\mu} = (\mu_1, ..., \mu_T)'$. Similarly, the disturbance series are denoted by $\boldsymbol{\eta} = (\eta_1, ..., \eta_T)'$, $\boldsymbol{\varepsilon} = (\varepsilon_1, ..., \varepsilon_T)'$. To begin, the irregular is assumed to be normally distributed, $\varepsilon_t \overset{iid}{\sim} NID(0, \sigma^2_\varepsilon)$, so that the focus lies on probabilistic description of breaks in the trend.

Level shifts are directly incorporated into the specification of model (1) by allowing for the random occurrence of very large $|\eta_t|$. By using a continuous probability density, consistency is maintained in defining the process. There is a significant possibility of unusual innovations in the level, that constitute shifts, but there is also a steady fall in probabilities at the extremes. This is a plausible framework where the size of breaks varies, and the chance of a break, exceeding a certain threshold, declines gradually with the magnitude of the threshold.

The Student-t class of distributions allows for a flexible degree of mass in the tails of the distribution. Its density function is

$$t_\upsilon(x; \mu, \sigma^2) = C(\upsilon, \sigma)\left(1 + \frac{(x-\mu)^2}{\upsilon\sigma^2}\right)^{-\frac{1}{2}(\upsilon+1)} \tag{2}$$
$$C(\upsilon, \sigma) = \frac{\Gamma((\upsilon+1)/2)}{\Gamma(\upsilon/2)\sqrt{\upsilon\pi}\sigma} \tag{3}$$

for $-\infty < x < \infty$, where $\upsilon > 0$ is the degrees of freedom parameter, μ gives the location, and σ^2 is the scale parameter of the density. The mean and mode are both equal to μ, as long as $\upsilon > 1$. The variance of the Student-t is given by $\sigma^2\upsilon/(\upsilon-2)$, $\upsilon > 2$. As $\upsilon \to \infty$ the shape becomes Gaussian with variance σ^2. For finite $\upsilon > 2$, the coefficient $\upsilon/(\upsilon-2) > 1$; for lower values, the total variance of the distribution is significantly higher than the scale, due to the contribution made by the increasing probability mass at the extremes. The second moment rises indefinitely as $\upsilon \to 2^+$.

Values in the range $2.5 < \upsilon < 12$ are sufficient to produce a good deal of variation in the tail thickness. This is illustrated in Figure 1, which shows the density functions for the upper and lower bound of the range, as well as an intermediate value, for the standardized ($\sigma^2 = 1$) Student-t with mean zero. Compared with the other two curves, the height of the tails is significantly greater for the distribution with $\upsilon = 2.5$ so the probability of a draw exceeding any given magnitude is considerably higher. Setting $\upsilon = 12$ already produces a rough approximation to the Gaussian case.

For the random disturbances that underpin the stochastic trend processes, I assume Student-t densities centered at zero as part of the time series model. Thus the independent errors have a density function with $\mu = 0$ denoted by $t_\upsilon(x; \sigma^2)$. Working with logarithmic expressions is often helpful for computational reasons. The logarithm of the t-density function is

$$\log t_\nu(x; \sigma^2) = -\tfrac{1}{2}(\nu+1)\log\left(\tfrac{1}{\nu}\left[\nu + \tfrac{x^2}{\sigma^2}\right]\right)$$
$$+ \log\left[\Gamma((\upsilon+1)/2)\right] - \log\left[\sqrt{\nu\pi}\sigma\Gamma(\nu/2)\right]$$

The use of such a heavy-tailed distribution for the trend innovation process enables one to anticipate, within the model structure, level shifts of varying magnitude.

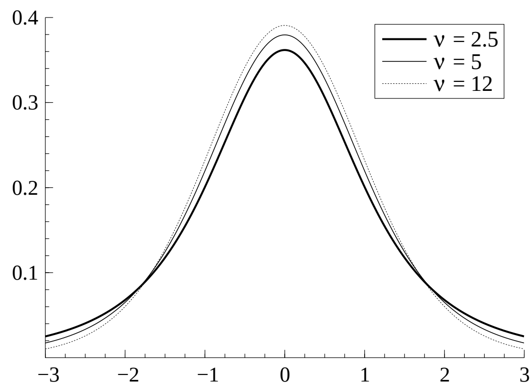

Figure 1: Probability density function for a standardized (scale factor equal to one) Student-t random variable with different degrees of freedom.

2.1 Trend estimation

This section describes how the non-Gaussian case is treated in applications. The expectation of the trend series, conditional on the information in the data, is of key interest. Estimation may be conducted efficiently by way of the importance sampling method, which has been widely applied in Bayesian applications in econometrics and statistics. Though the approach taken in this paper is classical, the link with Bayesian methodology arises since there is conditioning on the sample, just as in computation of a posterior density. State space methods, combined with the importance sampling technique, provide a convenient route to computing conditional expectations of interest. Below I describe how trend and parameter estimates are obtained for non-Gaussian models.

The expectation of the trend series, given the dataset $\mathbf{y}$ is

$$E[\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2] = \int \boldsymbol{\mu} p(\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2) d\boldsymbol{\mu} \quad (4)$$

The smoothed trend gives optimal estimates of how the underlying level has evolved over the sample period, using all the information in the data. However, direct computation of the quantity in (4) is not viable in many applications of interest, and in general it is difficult to draw from $p(\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2)$ in the non-Gaussian case. A viable strategy may be found by rewriting (4); first, consider the following equivalent formulation:

$$E[\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2] = \int g(\boldsymbol{\mu} \mid \mathbf{y}) \boldsymbol{\mu} \frac{p(\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2)}{g(\boldsymbol{\mu} \mid \mathbf{y})} d\boldsymbol{\mu} \quad (5)$$

This represents the expectation, of the quantity in brackets, taken with respect to $g(\boldsymbol{\mu} \mid \mathbf{y})$, which is referred to as the importance, or candidate, density. There is flexibility in the choice of importance density; in practice, the desirable properties are that the density is easy to sample from and that it provides a sufficiently good approximation for the application at hand. Here, $g(\boldsymbol{\mu} \mid \mathbf{y})$ refers to a Gaussian importance density. In particular, I consider the class of distributions derived from a local level model with normally distributed errors, whose variance may change over time.

Hence, the target density $p(\boldsymbol{\mu} \mid \mathbf{y}, \sigma_\eta^2, \sigma_\varepsilon^2)$ is approximated by $g(\boldsymbol{\mu} \mid \mathbf{y}, \{\widetilde{\sigma}_{\eta,t}^2, \widetilde{\sigma}_{\varepsilon,t}^2\})$, where $\{\widetilde{\sigma}_{\eta,t}^2, \widetilde{\sigma}_{\varepsilon,t}^2\}$ denotes the set of Gaussian variances for model (1) for $t = 1, ..., T$. Thus the candidate densities derive from linear Gaussian state space models with time-varying system matrices. Random draws from $g(\boldsymbol{\mu} \mid \mathbf{y}, \{\widetilde{\sigma}_{\eta,t}^2, \widetilde{\sigma}_{\varepsilon,t}^2\})$ are easily produced, for instance by the simulation smoother of Durbin and Koopman (2002). A simple criterion for selecting the parameters $\{\widetilde{\sigma}_{\eta,t}^2, \widetilde{\sigma}_{\varepsilon,t}^2\}$ is to compute those values that match the multivariate modes of the candidate and target density. Once the preferred approximating density $g(\boldsymbol{\mu} \mid \mathbf{y}, \{\widetilde{\sigma}_{\eta,t}^2, \widetilde{\sigma}_{\varepsilon,t}^2\})$ is obtained, then a number of key quantities are easily estimated. Further details on applying the importance sampling method may be found in Trimbur (2004).

2.2 Student-t trend innovations

In (1), to model the possibility of large innovations to the trend I specify a Student-t distribution for the level disturbances. That is, $\eta_t = \mu_{t+1} - \mu_t$, $t = 1, ..., T$, is independently and identically distributed with density function $p(\eta_t \mid \sigma_\eta^2) = t_\upsilon(\eta_t; \sigma_\eta^2)$. An example is shown in Figure 2, where $\upsilon = 5$ for a sample of 60 observations generated from the local level model, with $\sigma_\eta^2 = 1, \sigma_\varepsilon^2 = 4$. The signal-noise ratio $q = \sigma_\eta^2 / \sigma_\varepsilon^2$ is equal to 1/4 for the simulated series.

A shift in the level is apparent near the beginning of the series. Over the simulation period numerous moderate changes in the level are accompanied by occasional large ones, that occur more frequently and that vary more in

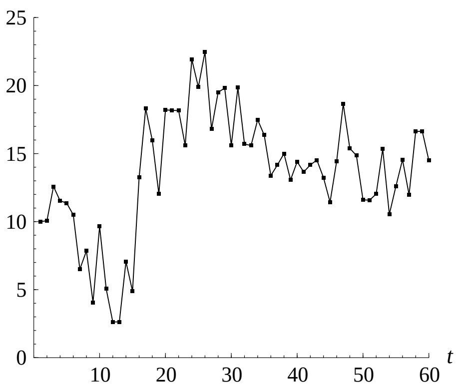

Figure 2: Simulated data ($T = 60$) generated from local level model in (1) where trend innovations η_t have a Student-t distribution with 5 degrees of freedom and scale parameter $\sigma_\eta^2 = 1$. The irregular component is Gaussian with variance $\sigma_\varepsilon^2 = 4$. The signal-noise ratio $q = \sigma_\eta^2/\sigma_\varepsilon^2 = 0.25$.

size than could be accounted for by a Gaussian structure. Many economic and financial series may be better characterized by models accounting for such additional variation. This would be reflected in the quality of smoothed trend estimates and forecasts. Later I consider examples where clear breaks in the level of the process are present, and illustrate robust trend estimation in some real time series.

Signal-noise ratios, defined as quotients of innovation variances, provide summary measures of time series dynamics. For the artificial dataset in Figure 2, q is less than one, reflecting the smaller relative variation in the trend innovation. Note that q has been defined in terms of the scale parameter of the Student-t density. In the limit as v becomes very large, the density becomes Gaussian and q is equal to the ratio of innovation variances. However, for lower degrees of freedom, the scale parameter σ^2 in (2) may be more representative of overall density shape than the variance $\nu\sigma^2/(\nu-2)$. The variance of the t-distribution becomes inflated by the increase in tail thickness as ν decreases toward two (even moderate increments, in the probability of draws with large magnitude, may have a significant impact on the variance), and in general the scale parameter may serve as a better basis for comparison.

The degrees of freedom was set to 5 in the simulated process in Figure 2. Lower values of ν would give rise to more frequent and larger shifts. In practice, the appropriate value will naturally depend on the particular dataset as the nature and importance of level breaks will differ across time series. As discussed in Trimbur (2004), the potential for precise estimation based on the properties of the likelihood may be limited, due to the finite length of the series, as the frequency of large movements, needed to characterize the density tail shape, is typically low. Thus, it may be desirable to fix ν to help ensure the effectiveness of the method for treating the stochastic level shift problem.

3. Empirical Illustrations

In this section, some examples are presented to illustrate the methodology. First, an annual time series that exhibits a definitive level shift, associated with a known historical event, is examined. This example illustrates how, once an extreme observation occurs, its classification adapts smoothly over time, and there is a gradual increase in the recognition of a level shift. The weighting function for estimating the trend is updated in an optimal fashion that depends on the information that becomes available after the occurrence of the event.

Next, a series is analyzed where the level shifts are more frequent and vary in their dynamics. This type of situation may be commonly encountered in practice, and it helps to illustrate the versatility of the model-based approach. The occurrence of the level shifts and their role in the evolution of the series is analyzed in a consistent framework. As the trend innovations may take on a continuous range, more information is available in characterizing changes in the level of the series, and the relative role of irregular fluctuations and level shifts is more precisely quantified.

Once an extreme observation occurs, its classification adapts smoothly over time, and there is a gradual increase in the recognition of a level shift. The weighting function for estimating the trend is updated in an optimal fashion, that depends on the information that becomes available after the occurrence of the event. This is illustrated for the annual time series of observations on the average level of the Nile over a hundred year period. In this case, a clear level shift is present, for which the underlying cause is known. A Gaussian model has some difficulty in pinpointing the change, while the method adopted in this paper provides a clear indication of the magnitude and timing of the shift in the level. Further, the implied weighting pattern, that is used in estimating the trend, is examined at different time points, which shows exactly how the model-based estimates adapt to the shift in the level of the series.

I also examine the quarterly average oil price over the last 25 years. For this economic series, large changes in the level of the series are apparent on a number of occasions. However, the occurrence and relative influence of shifts in the trend, in relation to the large irregular movements, is unclear from casual inspection. The model-based approach provides a consistent basis for making a quantitative assessment of how the trend in the market has evolved.

3.1 Computational issues

A key issue is to what extent the tail thickness may be assessed in practical situations, as reflected in the degrees of freedom for the Student-t trend innovation. As a general issue, one may expect the parameter υ to be difficult to discern, particularly for series of shorter length, as the frequency of extreme events is, by their nature, relatively low. Thus it would not be surprising if the available information about the extremes of the distribution was typically limited, and this would be reflected in a good deal of uncertainty in estimates of the degrees of freedom parameter. Another important consideration is the performance of the importance sampling method as a means of obtaining maximum likelihood estimates.

A simulated example is given in Trimbur (2004) where υ is set to 3. The results for this example produce estimates well above the true value, thus illustrating the challenges in estimating the crucial degrees of freedom parameter; in contrast the remaining parameters are estimated more effectively. The heavy-tailed characteristic enables one to emulate occasional large changes in the level of a series. However, given the computational difficulties that may be involved in unrestricted estimation, in applications where the level shift problem is of particular interest, it may be useful to assign a relatively low value to ν. This strategy guarantees a certain probability of such events. In a similar way, one may assume particular values for the signal-noise ratio to reflect the relationship between the disturbances driving the nonstationary and stationary parts of the model.

3.2 Shift in Nile level

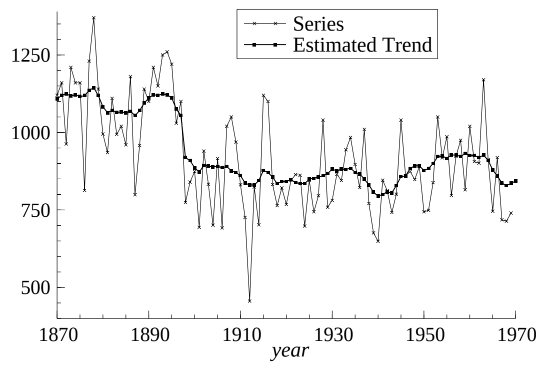

Figure 3: Annual recordings of the average downstream level of the Nile River over the hundred-year period 1870 to 1969, shown with estimated trend from model (1) with $\upsilon = 3$.

Figure 3 shows recordings of the yearly average downstream level of the Nile River in Egypt, originally studied in Cobb (1978). The series covers a hundred-year period during which the construction of a large dam led to a fall in the average water depth recorded; this shift is somewhat apparent in the graph but the dynamics of the transition are obscured by noise. When model (1) is fit with the restriction $\upsilon = 3$, so that the possibility of large trend movements is ensured, the smoothed trend in Figure 3 results. The change shows up clearly as a drop of about $\Delta = \widehat{\mu}_{t|T} - \widehat{\mu}_{t+1|T} = 136.3$ from year $t = 1897$, the year in which the construction of the dam took place.

Maximum likelihood estimation was conducted with $J = 1000$ draws in the importance sampler; convergence to the candidate density, computed for each set of parameters using the mode-matching algorithm explained earlier, occurred within 12 iterations or less. The computations were implemented in the Ox programming language of Doornik (1999) with a quasi-Newtonian method used to optimize over the free parameters. The implementation of the state space methods relied on the functions provided in the SsfPack package documented in Koopman et. al (1999). The initial values used in the maximization were $\sigma_\eta^2 = 10^3, \sigma_\varepsilon^2 = 10^4$, and the algorithm converged strongly (in the sense that numerical second derivatives gave clear indication of local optimality based on the curvature of the likelihood surface) after 16 steps, with resulting parameter estimates of $\widehat{\sigma}_\eta = 31.7, \widehat{\sigma}_\varepsilon = 120.1$. The signal noise ratio is then approximated as $\widehat{q} = \widehat{\sigma}^2_\eta / \widehat{\sigma}^2_\varepsilon = 0.07$. Note that the trend shown in Figure 3 was calculated using $J = 10,000$ draws from the importance sampler. As the summation need only be performed once when obtaining the final smoothed estimates of the component, it is computationally inexpensive to increase precision by sampling at greater frequency, compared to the parameter estimation procedure whereby each iteration uses an additional J draws from the simulation smoother.

The smoothed trend in Figure 3 is able to capture the relatively large changes in level in the late 1890's, as the choice of degrees of freedom reflects the assumption of heavy-tailed trend innovations that allow for stochastic level shifts. This result may be compared with those for a Gaussian model, as in Koopman et. al. (1999), where the change in the level of the series around the construction of the dam is spread over several years, in which case the characterization of the break is imprecise. It is clear from Figure 3 that a number of extreme observations in individual periods are also present, and an outlier process may be included in model (1) to reflect this. Hence, ε_t would be assigned a Student-t density, as in Aston and Koopman (2003), and the importance sampling method, explained previously for the heavy-tailed trend disturbance, naturally extends to this setup. Details may be found in Durbin and Koopman (2001).

The implied weighting pattern for estimating the trend immediately prior to the structural break, that is in 1897, is shown in Figure 4. The corresponding graph for the following year, where the break occurred, is displayed in Figure 5; now there is a striking transition in the weighting kernel for computing the level. Note that both patterns are naturally asymmetric as a consequence of the ability to accommodate level shifts. In Figure 5, the emphasis has flipped dramatically toward those observations that follow the time of the break in the sample. This shows what the models are doing in assessing a rapid shift in the level of the series.

The robust quality refers to how the trend estimator adapts quickly and makes a more transparent distinction of the break. The tail

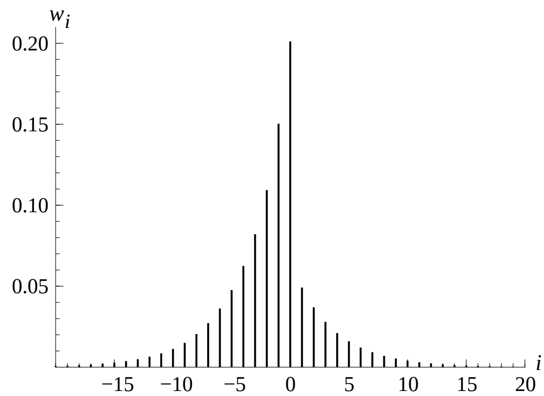

Figure 4: Weights for extracting the trend in the annual Nile time series immediately prior to the level shift ($t = 1897$).

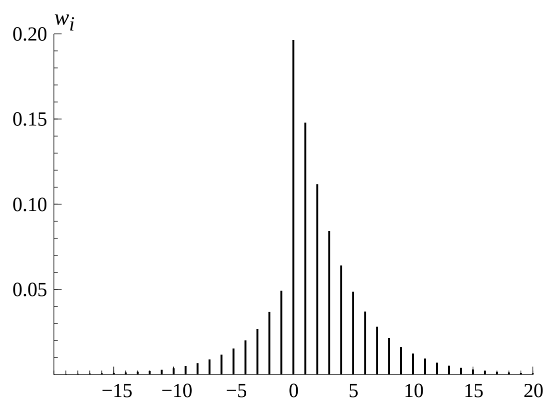

Figure 5: Weights for extracting the trend in the annual Nile time series immediately following the level shift ($t = 1898$).

thickness of the Student-t density depends crucially on the degrees of freedom parameter in formula (2), particularly for values of ν that approach the lower limit of 2. Hence, in the Nile example, choosing values slightly below 3 can make a difference in how well the event of the level shift is pinpointed. This is illustrated in Figure 6, which shows the results of trend estimation with ν fixed at 2.8. This represents a relatively minor reduction of only 0.2, or about 6.7% in ν, and yet the level shift is now significantly more discernible. The estimated magnitude of the underlying change around $t = 1897$ has risen by around 50%, to $\Delta = \widehat{\mu}_{t|T} - \widehat{\mu}_{t+1|T} = 179.6$.

On a real-time basis, one may look at how the detection of a level shift evolves as infor-

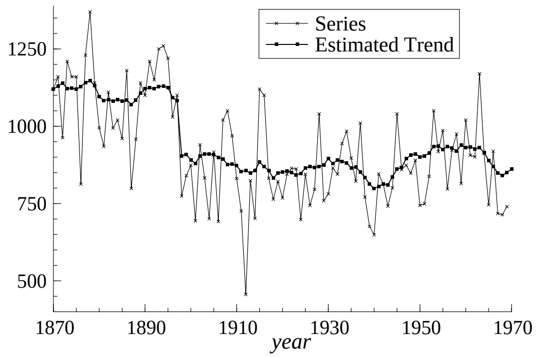

Figure 6: Annual recordings of the average downstream level of the Nile River over the hundred-year period 1870 to 1969, shown with estimated trend from model (1) with $\upsilon = 2.8$.

mation accrues. Naturally, it is a practical impossibility to immediately detect an underlying change in the trend, as initially, all that is known is that a sudden large innovation in the series has occurred. Thus, one aims for a method where the classification of each time point adjusts optimally, and where large innovation outliers do not have an adverse impact on parameter estimation. The use of robust models provides these advantages; furthermore, as there is a continuous-valued component, these models provides a great deal of information. One may expect outlier and level shift detection, in such a model-based framework, to be both more timely and more informative. An example of how trend estimates adapt over time for the Nile series is shown in Trimbur (2004).

3.3 Oil price changes

The occurrence of the level shift for the above example is relatively transparent as there is an associated historical event. Some apparent level shifts in time series data may stem from reclassifications or definitional changes in the quantity being measured. In other cases there may be more fundamental shifts in the dynamics of the economy, due to major financial or economic events, *e.g.*, oil price shocks. In both cases, non-Gaussian modeling strategies may have some advantages. In cases of known redefinitions in the composition of a data series, it is sensible to incorporate the relevant prior information. Some types of level shifts, however, may stem from extreme innovations that arise unexpectedly, representing surprise changes in the economic, political, or financial environment. Specifying the model to allow for the possibility of such events, before they occur, makes the method potentially valuable in forecasting applications. Statistical analysis in a setting with continuous-valued quantities may also provide additional insight beyond that of judgmental procedures based on historical assessment. Individual level shift events may vary a great deal in their dynamics, and in practice, a strategy based on time series modeling gives a good foundation in the general situation. The economic illustration examined next is characterized by a number of diverse, significant changes in the level of the series.

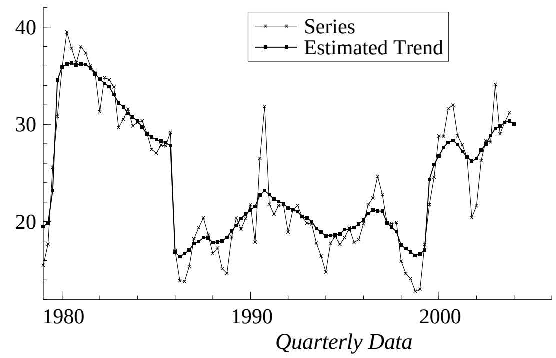

Figure 7: Quarterly average oil price from 1979:1 to 2003:4, shown with estimated trend where signal noise ratio is assumed equal to 1/10. The data are expressed as US Dollars per barrel for the West Texas Intermediate standard blend of crude oil.

Figure 7 shows the average price of crude oil over a recent sample of one-hundred quarterly observations, from 1979 to 2003 (USD per barrel for WTI blend, Source: U.S. Department of Energy). The series is rather volatile, and the fluctuations show substantial variation in magnitude and in degree of persistence. Though consumption of petroleum products may be linked to seasonal influences, e.g., the increased usage of heating oil in winter, seasonality is not apparent in the quarterly price series in Figure 7. If the primary determinants are related to expectations of future market conditions, then the average price may be expected to respond mostly to surprises in anticipated supply and demand factors. For instance, news of possible threats to shipments from the Middle East will induce positive movements, and if the potential supply disturbances are deemed long-lasting, one may expect the innovation outlier to be incorporated in the trend component. In contrast, temporary changes will not represent level shifts, but will instead be reflected in large values of ε_t. I estimate (1) with the signal-noise ratio constrained to be 0.1; this ensures a sensible degree of relative variation in the nonstationary and stationary parts of the model. The degrees of freedom parameter is estimated as $\nu = 2.92$, while the scale parameter is 0.73. Thus, the variance of the trend innovation is about one-third that of the irregular.

The resulting smoothed trend, shown in Figure 7, represents optimal estimates of the level over the sample period. Now $J = 100,000$ draws are used to better approximate the conditional mean series $\widehat{\mu}_t$, $t = 1, ..., T$; this precaution is taken as the initial trend movements are rather substantial (recall that a diffuse initialization is used). The sharp and lasting rise in 1979 shows up as an upward shock in the average level of the oil price. Similarly, the sudden drop in the mid-80's reflected a persistent change in the market balance as world oil demand fell and supply continued to expand steadily. On the other hand, the temporary price hike in the early 90's was followed, shortly thereafter, by an equally dramatic decline, as the previous concerns about possible interruptions in Mid-East regional oil supply rapidly dissipated. Recent movements suggest a high probability of an upward shift in the average level of prices, as reflected in the estimated trend series.

4. Conclusions

I have presented a statistical method for treating level shifts in time series. Such events are commonplace for many datasets in the social sciences and other fields and can have a substantial impact in analyses of dynamic properties. I have focussed on the issue of trend estimation. Some examples have shown how more robust trend estimates may be obtained whereby the model structure adapts more quickly to shifts in the underlying level. The continuous classification, implicit in our model-based framework, is more informative than binary labelling of structural changes. Since the unobserved components model with heavy-tailed innovations is designed to account for the probability of unexpectedly large changes before they occur, such models may prove useful in forecasting applications with real time series.

ACKNOWLEDGMENTS

I would like to thank Mark Asiala, John Aston, Bill Bell, David Findley, Andrew Harvey, Siem Koopman, and Brian Monsell for helpful comments.

REFERENCES

ASTON, J. and KOOPMAN, S. (2003) The Non-Gaussian Airline Model for Robust Seasonal Adjustment. *Mimeo.*

BAI, J. (1997) Estimating Multiple Breaks One at a Time. *Econometric Theory* 13, 315-52.

COBB, G. W. (1978) The problem of the Nile : conditional solution to a change point problem. *Biometrika* 65, 243-51.

DOORNIK, J. (1999) *Ox: An Object-Oriented Matrix Programming Language.* London: Timberlake Consultants Ltd.

DURBIN, J. and KOOPMAN, S. (2001) *Time Series Analysis by State Space Methods.* Oxford: Oxford University Press.

DURBIN, J. and KOOPMAN, S. (2002) A simple and efficient simulation smoother. *Biometrika* 89, 603-16.

KOOPMAN, S. J., SHEPHARD, N., and DOORNIK, J. A. (1999) Statistical algorithms for models in state space using SsfPack 2.2. *Econometrics Journal* 2**,** 113-66**.**

TRIMBUR, T. M. (2004) Analysis of stochastic level shifts in time series using heavy-tailed distributions : statistical and computational issues. *Research report in progress.*