Joint Probability Fluvial-Tidal Analyses: Structure Functions and Historical Emulation

D. A. Jones

This report is an official document prepared under contract between the Ministry of Agriculture, Fisheries and Food and the Natural Environment Research Council, for MAFF Project FD0417 "Confluence flood, joint probability".

Institute of Hydrology
Crowmarsh Gifford
Wallingford
Oxfordshire
OX10 8BB
UK

Tel: 01491 838800
Fax: 01491 692424
Telex: 444293 ENVRE G

Project T04066S1

March 1998

Summary

A summary is given of a number of techniques for estimating extreme levels arising from disjoint and/or disparate causes – the "joint probability problem". Nearly all of the practically feasible methods of analysis are based on "structure function simulation" which employ simplified models of the underlying cause-and-effect mechanisms and which, in turn, make use of simplified versions (ie. statistics of) the "cause" and "effect" variables. A second group of methods based on "continuous simulation" make use of physically realistic numerical models and are applied directly to create time-series of "effect" variables from corresponding "cause" time-series. Thus the latter group of methods would represent the underlying physical processes at a much finer time-scale, and with more detail and less simplification, compared with the former. This initial dichotomy can be further developed into a two-way table of methods by distinguishing between those methods which make use of statistical (or stochastic) models of the causal variables and those methods which apply statistical methods to the outcome of the cause-and-effect model. The first group of methods are "joint probability modelling" methods which involve modelling the joint statistical dependence between the "cause" variables, either via explicit probability descriptions or via random simulation methods. In contrast the second group, "historical emulation" methods, avoid having to model the joint dependence of the "causes" by making use of observed data for these, which implicitly encompasses the underlying dependence. At first sight the main value of historical emulation approaches would lie where some man-made effect has changed the physics of a situation, but equally important are those cases where no observations of the "effects" are available.

Historical data for river-flow and river- and sea-levels for the Tidal Thames are re-evaluated to assess the appropriateness of the structure-function approach which has previously been adopted in this case. The existing approach involves estimating the tidal-peak river-levels for tidally-influenced river-sites on the basis of river-flows and the peak levels at the estuary mouth (Southend). It is found this can be improved upon if the underlying tidal harmonics, or "astronomical predictions" of tides, are brought into account. This suggests that it would be better to base a structure-function approach on three components: (i) river-flow; (ii) the astronomical prediction of sea-level; (iii) the residual or "surge" element of sea-level at the estuary mouth. Further analyses indicate that, even for a site near the estuary mouth such as Southend, the tidal-peak levels are statistically related to the river-flows. In addition, the size of the surge component is related, not only to the season of the year as might be expected, but also to the astronomical prediction of the tidal peak and to the river-flow. The historical data are also used to indicate the extent to which interaction between the flow, astronomical and surge components are important in determining the tidal-peak-levels reached upstream.

Contents

1. Introduction

1.1 Background

This report is a contribution to research on the topic of joint probability methods applied to estimating the required river defence levels in situations where flooding can arise from more than one main "cause". Such situations are to be found at confluences and estuaries:

where rivers meet, when floods may be due to high flows on individual rivers or because the combination of flows is high;

where a river meets the sea, when floods may be due to high fluvial flows, high tides, large sea-surges, or a combination of all three.

While the main part of this report is principally concerned with an application on an estuary, similar methods of analysis might be applied to river-confluence problems after careful consideration.

IH has previously been involved in calculations for the Tidal Thames Defence Levels problem, which provides an example of a joint probability study involving the interaction of fluvial, tidal and surge effects. The main study was undertaken with Sir William Halcrow & Partners Ltd. for Thames Water Authority in 1987-88 and there have been some re-evaluations, and studies of tributaries of the tidal Thames, for the NRA (now EA). Clearly a dominant feature of these studies was the Thames Barrier: the purpose of the studies was to define appropriate flood defence levels both upstream and downstream of the barrier, taking into account its existence and operation.

Outlines of the use of various approaches employed for the Tidal Thames Defence Levels studies are provided in Section 1.3, but Section 1.2 begins by giving a general introduction to these approaches. A method involving joint probability models (outlined here in Sections 1.2.1.3 and 1.3.1) was used for calculations relating to the main Thames, while a simpler technique called "historical emulation" (outlined in Sections 1.2.1.4 and 1.3.2) was developed and used for the tributaries. For this application, both the joint probability modelling and historical emulation approaches are versions of a "structure function simulation" approach which is discussed in a general context in Section 1.2.1. Another fundamentally different type of approach, that of "continuous simulation" (Section 1.2.2), is at present not practicable. It seems reasonable to divide the continuous simulation type of approach into two classes. Both would make use of a fine-detailed time-scale for modelling, but the first, "event-based" class of approach would work with relatively short periods of time over which detailed modelling is carried out, whereas the "long-period" class of approach would apply the same detailed modelling, but over a single long period of record, usually over 10 years in length. A summary of the various combinations of approaches is given in Table 1.1. As indicated in this table, the historical emulation and joint probability modelling methods of analysis can each be applied within any of the conceptual approaches.

The difference between the different conceptual approaches essentially concerns the level of simplification adopted within the overall modelling. In particular, structure function simulation involves rather extensive simplification. For example, the quantities used for modelling are typically summary statistics such as the peak level within an event, in contrast to the use of time-series of data in the continuous simulation approach. The "*structure functions*", from which the name of approach is adopted, can themselves be considered as simplifications: they are a simple way of relating the summary statistics of the various data-series, in contrast to complicated numerical models applied to complete data-series within continuous simulation approaches. In the case of the Tidal Thames Defence Levels study the structure functions were simple look-up tables from which the peak water-level at a river-site within a given tidal cycle could be estimated from the peak sea-level at the estuary mouth and the mean flow entering at the top of the tidal reach.

Later sections of this report (Chapter 3) re-examine the structure functions used within the studies of the Tidal Thames Defence Levels, based upon evaluations using the set of historical data described in Chapter 2. The conclusion is drawn that there is considerable room for improvement in the way that these structure functions are defined. Since similar definitions of structure functions have been used in many other studies of estuaries, this conclusion is potentially of wide importance. Chapter 3 presents strong evidence of the need to take into account, within a structure-function approach, not only the actual sea-level at the estuary mouth but also the underlying tidal, or astronomical, component of such levels: ie. the "predicted" tidal peak available from Admiralty tide-tables or harmonic analysis. One variant of such an approach would use as primary variables the astronomical prediction of tidal peak, the "surge", defined as the difference of the observed peak sea-level from its predicted value, together with the river-flow. Chapter 4 describes some analyses of simple structure-functions of this type and, in particular, examines the statistical properties of such models, given that a joint probability modelling approach might be founded on this type of three-variable structure function.

1.2 General Conceptual Approaches and Methods of Analysis

1.2.1 Structure Function Simulation

1.2.1.1 Introduction

The methods of analysis which were adopted for the Tidal Thames Defence Levels study, namely the joint probability modelling and historical emulation methods, can both be considered to have been carried out within the conceptual framework of "structure function simulation" which has been mentioned in Section 1.1. The following subsections given some further information about the general foundations of these methods within the structure function simulation approach. Section 1.1 noted that essentially similar methods of analysis can also be applied within the rather different conceptual approach of "continuous simulation", and a limited discussion of this given later in Section 1.2.2. Section 1.3 contains more details of the specific application of the methods of analysis adopted for the Tidal Thames Defence Levels study.

1.2.1.2 Simplifications for Modelling

When used in the context of structure function simulation, joint probability modelling and historical emulation involve considerable simplification of both the physical and statistical aspects of the real-world: some of these simplifications might be avoided by use of continuous simulation, but this would very much depend upon the level of detail retained within the modelling. The simplifications for structure function modelling can be described as follows.

Time-Blocking

The first stage of simplification concerns the basic division of the ordinary time-scale into blocks, or time-units, which are then treated as individual items within the analysis. Time-blocking might involve division of the time-scale into steps of equal length, but might also involve the identification of "events", with data in portions of time outside such blocks not being used explicitly within the analysis.

Summary Quantities

A second stage of simplification concerns the use of statistical summaries of the various important time-varying quantities within each block of time: for example, values of the maximum water-levels within each time-block might be used.

Structure Functions

The third, and arguably most important, stage of simplification is the use of one or more *structure functions*: here it is assumed that the quantity of main interest, and in particular the summary statistic for this quantity, can be expressed as a simple function of the summary statistics for the other quantities which are available. In this sense, the analysis makes a clear distinction between quantities which are "causes" and those which are "effects": effects are predicted from their causes.

To these key elements of simplification may be added some concerns related to unobservable causes. An implicit assumption is that data are available for all of the major physical variables that "cause" the "effect", or at least that any unobserved "causes" will be relatively unimportant. There is a somewhat related assumption that the "structure function" idea is good enough, given that only summary statistics are used for both causes and effects: while there may be a good relationship between cause and effect when considered at a detailed time-scale, this may be obscured by the use of summary statistics. Thus the use of summary statistics can induce some extraneous variation in required relationship which can be classed as having an unobservable cause if only the chosen summary statistics are available. To a certain extent, the limitations imposed by these assumptions can be overcome by extending the concept of the structure function to encompass information about the error with which the effect can typically be predicted if the values of the causal variables are known. An extended structure function would provide both a central estimate of the effect variable and a typical size of estimation error, both of which would vary with the causal variables. More sophisticated versions of this extension would provide conditional probability distributions for the effect given known values for the causes.

1.2.1.3 The Joint Probability Modelling Method

Within the structure function simulation approach, the joint probability modelling method is based upon defining a detailed statistical model for the summary statistics of the "causes". This model is then used to integrate across the range of possible values of the cause-variables of all types in order to derive the probability distribution of the "effect". The statistical model here needs to represent not only the joint dependence of the cause-variables within a single time-block but also the between -block aspects of the problem. While it may often be found that events within adjacent time-blocks are effectively independent if the time-blocking is defined in an appropriate way, this needs to be checked. Where time-blocking is such that only significant events are being treated explicitly, the statistical model also needs to describe about the rate at which such events occur.

Joint probability modelling methods can be applied in a number of ways, of which only the two at the ends of a range of variants are mentioned here. One possibility is to accomplish the required integrations of the probability distributions by numerical integration of these distributions: here the structure functions are taken into account when defining the ranges over which the joint distributions are to be integrated. For the second possibility, the integration would be accomplished by random simulation. That is, a large number of sets of cause-variables would be generated from the joint distribution of these quantities. Each such set would be converted to a value for the "effect" by making use of the structure function and, where applicable, this value might include a random contribution for the error with which the structure function predicts the effect. Finally, a statistical analysis would be made of the overall set of generated values for the effects.

1.2.1.4 The Historical Emulation Method

In contrast to the method of the previous sub-section, analyses using the historical emulation approach are rather simple, since they avoid much of the need for detailed statistical modelling. The simplest version of this approach is to use long records of the "causes" to create, via the structure function, an equivalent long record of the "effect" which is then analyzed using standard techniques as if it were a long record of observations. When this is done within a structure-function approach, only the summary statistics of the causes are used, creating corresponding summary statistics of the effect. The terminology "emulation" rather than "simulation" is used here to emphasize the fact that, although not observations as such, the values created are in effect tied to particular time-instances through the use of observed data for the causal variables. "Simulation" would be used for non-time-specific values created in a similar way, for example in the integration methods using randomly generated cause-variables described in Section 1.2.1.3.

In cases where an event-based time-blocking is used, or in order to extract the maximum information from the data via the use of an event-based analysis, there may be a need for a statistical model for the rate at which events occur. However, it would often suffice to use a continuous sequence of emulated values to derive a set of "annual maxima" which are then treated as if they were statistically independent observations.

Various extensions of this basic idea are available for practical use. For example, if the structure function idea has been extended to include a statistical model for the error, then the "emulated" values can be created so as to include randomly generated versions of these errors and a final analysis might combine information from several such sets of partly-random emulated series. It is clearly possible also to mix a short record of observed values of the "effect" with emulated values so as to create a long record for analysis.

1.2.2 Continuous Simulation

Section 1.2.1 introduced the "joint probability modelling" and "historical emulation" methods in the context of "structure function simulation". There are corresponding versions of these methods under the general heading "continuous simulation". Continuous simulation methods have the advantage of avoiding many of the simplifications, and hence errors, inherent in the use of structure function simulation. In particular, both the use of structure functions and the use of summary quantities to drive the structure functions are avoided with continuous simulation.

Here, "continuous", as in "continuous simulation" and "continuous-time", indicates that the time-series used, while actually sampled in discrete-time, will nevertheless be sampled frequently enough to encompass all relevant time-variations. This is in contrast to the use of summary statistics within the relatively large time-blocks used for structure function simulation. Continuous-time versions of the time-series of the cause-variables would be used to drive a numerical model which provides a reasonably accurate representation of the physics of the situation: this would create a continuous-time version of the "effect" from which various summary statistics of extremes would be extracted for later analysis. An example would be a numerical model based on the St. Venant equations for one-dimensional hydrodynamic channel-flow. While the same model might be used to create a structure function, the essential difference is that here the "results" are determined by the actual time-variations of the series driving the model, rather than by a simplified version of these inherent in the use of summary quantities in specifying the structure function.

A simple version of the "historical emulation" approach to continuous simulation would use long records of continuous-time observations of the "causes" to create a corresponding modelled series for the "effect". The "joint probability modelling" version of continuous simulation would be based on driving the numerical model with pseudo-random continuous-time records for the "causes". Here the process used to generate these pseudo-random series would be designed to reflect the statistical properties of these data, including both intra- and inter-series properties; that is, both serial dependence in individual series and joint dependence between series. In this case, the "joint probability modelling" method would rely on summarising the properties of many sets of pseudo-randomly-driven "effects" series, since numerical integration would not be feasible here whereas it often is for structure function simulation.

Once again, the main distinction between the historical emulation and joint probability modelling approaches is that the former avoids the need for constructing a joint probability model, which can be a considerable task. In the context of continuous simulation, a joint probability model would usually be structured directly as a stochastic random-simulation model, rather than being expressed in terms of probability distributions as more usual for structure function simulation. For continuous simulation, it may well be appropriate to frame the stochastic model in terms of other background or underlying physical processes. In particular, for applications involving river-confluences, it may be easier to frame a stochastic model for the flows in adjacent tributaries based on random spatial rainfall fields applied as inputs to rainfall-runoff models, rather than to build a joint model for the flows alone.

As mentioned in Section 1.1, the continuous simulation type of approach can be divided into two classes. Both of these would make use of a fine-detailed, "continuous", time-scale for modelling, but the first, "event-based" class of approach would work with relatively short periods of time over which detailed modelling is carried out, whereas the "long-period" class of approach would apply the same detailed modelling, but over a single long period of record, usually over 10 years in length. Whether or not an event-based continuous simulation approach is applicable in a given case will depend upon the intrinsic properties of the processes being modelled and, in particular, upon the amount of process-memory to be taken into account. As an example, one might consider the problem of river-levels at a confluence. If a numerical hydrodynamic model is used to represent the combined effects of river-flows on several tributaries, a warm-up period of as little as one or two days would often be enough to ensure that start-up conditions did not affect the results for flood peaks: thus for this particular purpose there is little process memory. However if, for the same problem, the river-flows used as inputs to the hydrodynamic model are generated by some stochastic model, it is clear that the process-memory here is rather longer: minimal estimates would be several weeks for soil-moisture conditions and one or two years for groundwater effects. Thus an event-based continuous simulation approach may be precluded for problems involving river-flows, unless one is prepared to accept the difficulties associated with adopting a statistical model for the initial conditions pertaining at the start of each event.

1.3 Application of Methods for the Tidal Thames Defence Levels

1.3.1 The Joint Probability Modelling Method

For the Tidal Thames Defence Levels Study, estimates were required of events of high return period, of the order 100 to 10,000 years. This was the reason for using the joint probability method for sites along the main Thames. Such return-periods are beyond the range that can be reasonably estimated using the historical emulation approach, particularly since the operation of the Thames Barrier introduces shelves and sharp up-turns in the effective level-versus-return-period curve in precisely the region in which a curve fitted to either observed or emulated data would have to be extrapolated in order to provide an estimate. The data available for this study consisted of flows at the upstream end of the reach (Teddington) and sea-levels at the downstream end (Southend).

The statistical analysis was based on a data-set created by extracting, from a basic data-set consisting of tidal-cycle maxima (2 per day), the maximum levels reached at Southend within each neap-to-neap tidal cycle (approximately 14¾ days), together with the flows at Teddington concurrent with those maxima. One reason for this step was to try to avoid the need to incorporate serial dependence into the model. While this was successful in regard to the sea-levels taken on their own, it was found that not only was there dependence between sea-levels and flow, but also serial dependence in the flows concurrent with sea-level maxima in adjacent 14¾-day periods. Note that the statistical models involved in the joint probability modelling also contained elements to represent trend, seasonality and the remaining astronomical cycles in the sea-levels, and seasonality in the flows.

In the joint probability calculations, the modelled long-term trend in sea-level was not included but was instead replaced by a fixed value. Any required assumption about future trends in sea-level would be incorporated by adding an appropriate allowance to the answers produced. One essential additional component of the overall model was to include consideration of "secondary events" to reflect the fact that only neap-to-neap maximum tides are incorporated directly, whereas a higher level at a given point on the river may occur on an adjacent tide if the flow happens to be higher then. A simple procedure was developed for including not only the sea-level and flow from the underlying statistical model but also a simple additional model for what might happen on adjacent tides (ie. slightly lower sea-levels with slightly higher flows).

A basic ingredient of the joint probability method here is the "structure function". This specifies what the maximum river-level at a chosen point will be for a given combination of maximum sea-level and flow and for a given assumption about whether or not the barrier is operated for that tidal cycle. It is clear that the use of structure functions hides a number of difficulties, since the maximum level reached at an estuary site actually depends on how sea-levels and flows vary over the tidal cycle and not just on single values of these quantities. The application essentially ignored the potentially moderately large effects arising from wind acting on the estuary and tidal-river, but some allowance for this was made by developing a description of the errors with which peak-levels can be predicted via the structure functions.

The required structure functions were obtained from multiple runs of a hydrodynamic model, in this case ONDA, configured and calibrated by Halcrow to the main Thames from Molesey, somewhat above Teddington, to Southend. For later purposes, it is convenient to note here that the structure functions were calculated by passing several tidal cycles through the hydrodynamic model, with the flow at the upstream boundary being represented by a constant value. The use of a constant fluvial flow, while obviously an approximation, is unlikely to cause problems considering the fairly slow variation in flows at Teddington. The tidal cycles for sea-level at Southend were constructed to be sinusoidal with an amplitude to match the require maximum sea-level, except in cases where this amplitude would be greater than the largest amplitude experienced in the "astronomical" component of sea-levels for this site. In the exceptional cases, the sea-level cycle was constructed to be the sum of a sinusoid of the designated maximum amplitude and a "surge" component.

This surge component was of such a size as to make up the required maximum sea-level but, apart from this, details of the relative timing and duration of the surge were based on the observed behaviours of a few large surge events.

For this application, used two tiers of time-blocking (Section 1.2.1) were used: the first, concerned with the identification of every high tide, was the one at which the structure function was applied; the second tier for time-blocking was concerned with the 14¾-day periods used within the main probability calculations. Each of these tiers had a corresponding set of summary quantities. The structure function for a given location predicted high-tide levels at that site (for any high-tide or, equivalently, for any given time-block at the first tier), from the corresponding high-tide level at Southend and from the daily-average flow at Teddington. The probability calculations required a summary quantity at the second tier of time-blocking: in particular, the maximum high-tide level at the target location within each 14¾-day period was required as the final output of the structure function modelling, while the statistical modelling was based on the maximum sea-level at Southend within each second-tier time-block, together with the daily-average flow at Teddington on the day of occurrence of the maximum sea-level at Southend.

The joint probability method allows various barrier closure rules based on sea-level and flow to be compared, and it even allows account to be taken of the uncertainty about whether or not the barrier will be closed. This uncertainty arises because the actual closure decision is based on a forecast of the maximum sea-level for the oncoming tide. Additional "random" factors can be incorporated. For example, an analysis of data for intermediate sites along the estuary allowed quantification of the errors involved in predicting maximum levels via the structure function, and these errors, which were treated as a random noise, were allowed for within the joint probability calculations.

Once the underlying joint probability model has been estimated, it is possible to calculate the required exceedence probabilities as precisely as may be desired. Thus it is tempting to regard the method as providing much more accurate answers than the historical emulation approach. However, this would be to ignore entirely the question of how accurately the parameters of the probability model can be estimated. No matter what method is used, only a rather limited reliability can be expected for estimates of river-levels of moderate-to-high return periods when derived from what will inevitably be a comparatively short set of historical data. However, the estimates are of some importance and interest for their own sake, particularly where valuable property is at stake. There does seem to be scope for making use of the apparent precision of the answers from the joint probability modelling approach in the context of making relative comparisons between the consequences of different operating policies for the Thames Barrier, since the difference between two answers from the calculations will be rather less affected by errors arising in estimating the underlying statistical model than will a single answer taken on its own.

1.3.2 The Historical Emulation Method

The "historical emulation" technique is comparatively simple to apply, since it avoids formal statistical modelling of the kind required for joint probability modelling calculations. IH employed the method on a number of the Thames tributaries, where the alternative would have been to extend the statistical modelling required for the joint probability method to include the upstream tributary flow as a third variable in addition to Teddington flows and sea-levels at Southend.

The historical emulation technique proceeds by constructing, for each tidal cycle in the basic data-set (2 per day), an estimate for what the river-level would have been at a chosen point. Annual maxima are then extracted from these constructed values and are used to estimate return periods in the same way as if the record had been "real". Where the operation of the Thames Barrier is in question, the effects of this are allowed for by deciding for each tide whether or not the barrier would have been closed on the basis of the recorded values for Teddington flow and sea-level at Southend. It should be recalled that part of the historical record is for a period for which the Thames Barrier did not exist, that the closure rules for the Barrier actually applied in the past may not be the same as those under test and also that various test-closures of the Barrier take place. If, for a given tide and a given closure-rule, it is decided that the Barrier would have closed, or if the Barrier actually did close, then any observed river-levels at intermediate points along the main Thames are treated as missing. The estimated value at the chosen point for a given tide is constructed on the basis of the best information available. Thus, if an observation is available for that point, or for a nearby site, that is used; otherwise, a value is constructed by reference to a structure function for the target point based on upstream tributary flows and maximum level at the junction of the tributary and main river. The river-flow into the tidal tributary for a particular tide is obtained from a record of daily-mean flows using an allowance to produce an estimate of peak flow. The river-level at the junction is obtained either from a nearby site on the main river or by using a structure function for the main river based on Teddington flows and Southend sea-levels and depending on whether or not the Barrier is assumed to have closed for that tide.

The "historical emulation" method was first used on a tributary on which rapid urbanisation was occurring. To provide estimates for present-day conditions, the recorded values for the tributary flows were modified in such a way as to reflect a return period relationship for flows (ie. a fluvial flood frequency relationship) which had been derived separately by the Water Authority. In this method, only the sizes of peak flows on the tributary are changed, not their times of occurrence, and it is to be hoped that, even though the joint statistical dependence of sea-levels and the tributary and main-channel flows is not explicitly represented in the model, they will still be broadly correctly represented. However, it is known that, in general, urbanisation affects both the river-catchment's sensitivity to rainfall duration and the seasonality of occurrence of high flows and these effects would not be accommodated by the above simple method. In principle, the method can be extended to deal with possible long-term trends in sea-level by replacing the set of observed values for sea-level by a set which has been first detrended and then had an appropriate trend superimposed. Note that, although this has not yet been implemented in practice, it would be a relatively simple matter to take account of a variety of minor random effects by generating a number of alternative sets of the estimated data-values based on different random outcomes. Examples of such random effects would be to take account of the uncertainty in whether or not the barrier would operate for a given tide, or to take account of estimation errors if it is necessary to use flow at a different site to construct an estimate of the upstream tributary flow.

This application of historical emulation used only a single tier of time-blocking (Section 1.2.1), compared with the double-time-blocking used in the joint probability modelling approach. The summary quantities and the structure functions used here are essentially the same as employed in the first tier of time-blocking described in Section 1.3.1. However, the structure functions used to predict levels on tributaries of the Thames were based on summary quantities in which the high-tide level at Southend was replaced by the high-tide level at the tributary mouth.

A disadvantage of the historical emulation technique is that it can only provide estimates for low to moderate return periods, up to about the length of record available. It does have the major advantage of being easier to understand than the joint probability method, since it essentially involves only simple manipulations of the available data and, furthermore, it avoids relying on the assumptions implicit in working with a statistical model.

1.3.3 Discussion

The availability of historical data is crucial to analyses of extreme water levels and, to be conveniently useful, they need to be in computerised form. The use of tidal maxima in our analyses was just about adequate, but continuous (e.g. hourly or ¼ hourly) records of sea-level are to be preferred for two reasons. Firstly, in the context of quality-control of the data, a more thorough check of the data against values from nearby sites can be accomplished if data for the complete tidal cycle are available. Secondly, a continuous record of data allows a wider range of analyses to be contemplated, even for the simpler statistical procedures suggested here. The actual methods of analysis adopted will depend on three things: the problem itself, how much time and resources it is worth expending on the problem, and the data available for the solution.

The classification of conceptual approaches and methods of analysis adopted here, and which is summarised in Table 1.1, should not be allowed to limit the scope of what might be considered for a given application. As an example, it seems worth mentioning the following variant of historical emulation within an event-based version of continuous simulation. In particular, this example illustrates an event-based approach which would make use of an annual maximum type of approach to subsequent analyses as opposed to a peaks-over-threshold approach. Consider an estuary problem involving records of sea-level and river-flows. Suppose that, for whatever period of record is available, 3 or 4 events per year can be selected which one is fairly certain would have between them produced the annual maximum level at all intermediate points of interest. Then, if continuous records of flow and sea-level are available for a limited period around each event, perhaps digitised especially for the purpose, the events can be simulated individually and thus eventually a set of annual maxima for all intermediate sites can be constructed. Note that it is the availability of continuous records which would allow this analysis to be done and it is clear that this approach has the distinct advantage (over the structure-function version of historical emulation) in that it allows proper account to be taken of the relative timings of peaks of flow and sea-level. The approach here avoids the need for specifying what an "event" is in the formal sense required for a peaks-over-threshold approach, since it is unnecessary to estimate a rate for the occurrence of events. In principle, one would be applying an annual-maximum approach to the results from a long-period continuous simulation, but making use of a computational and data-gathering short-cut by restricting the data employed to relatively few likely annual-maxima events. Where there is a lack of continuous-time data-sets in computerised form, a more complete "peaks-over-threshold" analysis of levels would be rather more difficult and costly to achieve, since this would require data for rather more events in order that one can be sure to have considered all events which would have produced a level exceeding a given threshold at an unobserved intermediate site. This "annual-maximum" event-based continuous simulation version of historical emulation is potentially applicable to both estuary and river-confluence problems.

Table 1.1 *Summary of the approaches and methods available for joint probability problems.*

Conceptual Approach	Method of Analysis	
	Historical Emulation Accuracy and precision of results directly related to the length of the historical record.	Joint Probability Modelling Accuracy equivalent to length of historical record. Increased *apparent* precision in results is possible.
Structure-function simulation	Creation of quasi-observations via a simple mechanism.	Direct calculations using explicit probability distributions, otherwise random simulations of long records.
Event-based continuous simulation	Creation of short-periods of continuous record of quasi-observations via a detailed model.	Not applicable if processes have a long internal memory.
Long-period continuous simulation	Creation of a long, continuous record of quasi-observations via a detailed model.	Likely to be feasible for random simulations only. A single very long, entirely fictitious (random) continuous record could be created.

2. Data for the Tidal Thames

2.1 Information taken from the Tidal Thames Defence Levels study

A substantial data set was accumulated for the study of the Tidal Thames Defence Levels (Chapter 1), and this has been used as the basis of the present study. In the first instance, this data set consisted of values for the high-tide levels recorded at Southend, Gallions, Tower Pier and Richmond for the period 1939 to 1985, which were transcribed into computer-compatible form from hand-written records held by the Port of London Authority (PLA). Note that, while the manuscript records did contain values for low-water levels as well as for high-water levels, and also information about the times at which the high- and low-water levels were reached, none of this extra detail was included in the transcription process in order to save costs. Thus, none of the analyses discussed in this report have made use of low-water (low-tide) levels or the recorded times of tidal-peaks.

The basic records, and the initial computerised versions of these, were such that either one or two high tides were recorded each day at each site, depending on each tidal-peak's time of occurrence within the day at the site. This means that different numbers of tides are recorded for the same day at the different sites. As part of the data-preparation effort for the original study, the records for the individual sites were aligned so that what are effectively the same high tides are identified consistently across the sites. Thus each tide is now recorded against the day at which the peak occurred at Southend, even if the peak reached the subject site on the next calendar day. In addition, these records of high-tide values were extended by including, for each high-tide, the value of the daily-mean flow on the Thames at Teddington, for the day on which the tidal-peak at Southend occurred. The data-set was extensively quality-controlled by a cross-comparison of the tidal-peak values held for the four sites, with problem values being compared with the manuscript records and, in a limited number of cases, with records in the form of charts held at the PLA. Details of the closures of the Thames Barrier, since its first effective employment in November 1982, were obtained from the Barrier operators and these were taken into account in the quality-control checking. In summary, the records available from the Tidal Thames Defence Levels study consisted of data for the 47 calendar years, 1939 to 1985: peak-levels for the 33171 high-tides in that period, together with corresponding river-flows at Teddington and an indicator of whether or not the Thames Barrier was operated for that tide (either for test purposes or for a real potential flood event).

The original study also led to the creation of "structure functions" for a number of sites along the tidal reaches of the River Thames. These are essentially tables of values whereby the tidal-peak river- or estuary-level at a given site can be predicted from values of the tidal-peak value at Southend and the river-flow at Teddington. As discussed in Chapter 1, these tables were constructed by use of the ONDA hydrodynamic river model which had been configured on the basis of extensive channel and flood-plain survey data together with calibration against river-level records for a 9-day period in December 1978 and checking against records for a 3-day period in December 1985. In fact, subsequent to the original study, a revised set of structure functions was created by Halcrow as part of a re-evaluation of the problem, and it is this revised set of structure functions that is employed here. However, for the present study, the difference between these two sets of structure functions can be considered minor.

In the comparisons made here, it should be remembered that the structure functions were defined to represent a particular set of conditions: specifically, they assume the existence of flood-defences sufficient to retain the river within prescribed geographical limits. While the model included the tidal washlands in the upper reaches of the Tidal Thames (Syon Park, Old Deer Park, Ham and Petersham), elsewhere the model assumed that the river would be retained "in-bank", behind defences the size of which it was the study's task to determine. It is clear that the river-levels recorded in the historical data reflect conditions in which such defences either did not exist or were at levels rather lower than required to contain the river. This makes it rather difficult to evaluate in a formal way the results of predictions from the structure functions in comparison with the levels in the historical record. In addition to effects from flood-defences, it is known that there have been changes to the configuration of the river channel and flood plains over the historical period which would potentially affect river-levels in all parts of the regime, not just during periods of high levels.

2.2 Data for the present study

The data-set used in this report has been derived directly from that created for the Tidal Thames Defence Levels study (Section 2.1). Because the record available for Gallions was comparatively short at 12 years, it was decided to concentrate the present study on the sea- and river-level records for Southend, Tower Pier and Richmond, although the data for Gallions were retained to aid in cross-checking any apparently unusual data-points shown-up by the analyses. For both Richmond and Tower Pier, and to a certain extent for Gallions also, it seems most natural to use the term "river-level" rather than "estuary-level", and this is the convention that has been adopted for this report.

Given the types of analyses being undertaken here, it was necessary to be able to ascribe a reference time to each of the high tides in the record. Although the original manuscript records had included the time at which each high-tide was observed, this information had not been incorporated into the computerised records made for the original study. Given the nature of the present study, it would have been inordinately expensive to digitise this information. To overcome this problem, it was decided to treat the sequence of high-tides as if they were equally spaced in time at a time interval of 0.517525 days, which corresponds to half a lunar day. Thus the first high tide at Southend on 1 January 1939, and corresponding tides at Tower Pier and Richmond, are each treated as occurring at the time of "0.517525 days", with the next at "1.03505 days" or, equivalently, one mean lunar day. Note that the 17167 days in the 47 year record provided 33171 high-tides, and that $33171 \times 0.517525 = 17166.8$ so that this matches well with the observed rate at which high-tides are observed. It should be recalled that the computerised data set contains no information at all about low-tides and hence there is no requirement for the construction of times of low-tides for this study.

The treatment of the record of high-tides as if they were an equally-spaced time-series is clearly an approximation. It is not clear what effect this approximate treatment of timing might have on the analyses being done, but one can hope that the effects would be small. However, there are almost certainly some important effects being introduced by this approximation. An examination of the times of high tides recorded in the manuscript record in the first month of the 1939 revealed time-intervals between observed high-tides that varied between 11 hours 52 minutes and 12 hours 59 minutes, or from 0.4944 to 0.5410 days, with an apparent systematic variation in these interval-lengths related to the tidal cycle. The analyses being undertaken here (Section 3) contain elements which are essentially spectral analyses. Two possible effects of the equal-interval approximation on these analyses are that the astronomical periodicities contained within the underlying continuous-time sea-level records may change in relative importance, or that new apparent "non-astronomical" period-lengths might be introduced.

Each of the three sea- or river-level records contains a number of periods during which levels were not recorded: for example, the gauge at Southend was under repair from 31 January to 6 July 1948. In addition, further missing data were introduced during the quality-control investigations of the original study. For the present study, the values for high-tides during which the Thames Barrier operated are also treated as missing: this rule was applied to the Southend record as well as to those for Tower Pier and Richmond because there was some suspicion that an effect from operating the Barrier might be felt even as far downstream as Southend. Taking into account these missing data, the numbers of high-tides available for analysis at Southend, Tower Pier and Richmond are 32374, 32071 and 30933, respectively. The numbers of high-tides on which levels for the pair Southend and Tower Pier, and for the pair Southend and Richmond, are both available are 31446 and 30155 respectively.

Figures 2.1 to 2.4 show the data available for a selection of time-periods within the available record. Figure 2.2 shows a period including the highest values of flow at Teddington within the period (709 cumecs on 20 March 1947). As further examples of extremes, Figures 2.3 and 2.4 include the highest and lowest values of sea-level at Southend. In addition, Figure 2.4 contains within the period shown a case of a high-tide on which the Thames Barrier closed: this is the reason for the missing value shown for all three level-records. The way in which the values for the flow at Teddington corresponding to the high-tide levels were constructed means that these data typically consist of a series of pairs of equal values, with an occasional singleton: this pattern can be discerned in Figures 2.1 to 2.4 . Note that the later study by Halcrow, leading to the revised structure functions already mentioned, also cast doubt upon the values of flow contained in the historical record for the period 18-23 March 1947 (Figure 2.2). This later study suggested that, because of the existence of a coffer dam not taken into account in the original calculation of flows, the actual flows may have been 9% less than those in the standard Teddington record for this period (ie. a peak flow of 645 rather than 709 cumecs). The analyses here continue to use the standard record.

2.3 Structure functions

As noted in Section 2.1, the previous studies have led to there being available a set of structure functions for various points along the Tidal Thames, including both Tower Pier and Richmond. The versions of the structure functions used here are based on the revisions of the structure functions created by Halcrow in 1996. Because of the context for which they were created, these structure functions covered a range of levels at Southend from 2.0 to 4.5 m and a range of flows from 50 to 1000 cumecs. Specifically, the original tables for the structure function presented predicted river-levels for pairs of sea-levels and flows from the following sets.

Sea-levels at Southend	2.0, 2.5, 3.0, 3.5, 4.0, 4.5 m AOD;
Flow at Teddington	50, 100, 200, 300, 400, 600, 800, 1000 cumecs.

The present study required values for the structure function outside of these ranges and it has therefore been necessary to create extrapolation procedures which will hopefully produce reasonable results. Examples of the extrapolated structure functions are included in Figures 2.5 to 2.8. The major requirement for extrapolation was for the range of sea-levels at Southend below 2 m: here the procedure used a linear extrapolation from the values given in the original structure function at 2 m, with a slope chosen to match the scatter plot of river-level against sea-level obtained from the historical record. It is arguable that a more realistic behaviour for the structure function might have been found in the case of Richmond where the original structure function showed signs of exhibiting a "purely fluvial" behaviour for high river flows: thus, for low amplitude tides, the river level reached at Richmond during extremely high flows should not depend much on the sea-level at Southend. However, most of the analyses here are concerned with the combinations of sea-level and flow actually experienced during the period of the historical record and there are very few occurrences of high flows with low peak sea-levels. For this and other reasons, the conclusions found here in respect to the present structure functions should not be affected by doubts over the appropriate extrapolation of these functions.

Figures 2.5 and 2.6 show scatter plots obtained from the historical record by plotting the observed river-levels at Tower Pier and Richmond, respectively, against the sea-level at Southend. Also included in these plots are the values given by the structure functions, shown as lines representing the predicted river-levels for Teddington flows of 50, 100, 200, 300, 400, 600, 800 and 1000 cumecs. Figures 2.7 and 2.8 show the scatter plots obtained by plotting the observed river-levels at Tower Pier and Richmond, respectively, against the corresponding values for the flow at Teddington. Again the structure functions are shown in these plots in the form of lines representing the predicted river-levels for peak sea-levels at Southend of 2.0, 2.5, 3.0, 3.5, 4.0 and 4.5 m AOD. These figures show that, for both Tower Pier and Richmond, the variation of peak river-levels is very strongly matched by the variation in sea-levels at Southend and is not at all well matched by the variation in the flows at Teddington. In neither case does the scatter plot of levels against flows reveal the definite existence of a relationship although it might be possible to imagine that such a relationship is beginning to appear for high flows in the case of Richmond. The variations shown by the structure functions agree with these conclusions: the structure functions predict a strong change of river-level with sea-level at Southend, across the range of values actually experienced, and only minor changes with flow at Teddington.

As noted earlier, the structure functions used reflect what are to a certain extent "current conditions", given that the required flood defences are essentially in place. In contrast, the historical record reflects changing conditions, relating to changes in both channel geometry and flood defences. Each time-series of sea- and river-level contains a trend over time, part of which may be related to such changes. While the installation of flood defences might be expected to lead to an abrupt change, implementation of a large number of local schemes for defence improvement over a long period could lead to the appearance of there being a slow drift. Even if the trends do not actually arise from these causes, one would still wish to be able to use the structure functions to impute what would happen "now" if events similar to those of the past were to occur, where "now" version of the past events would be found by removing the effects of trend. Thus there is some interest in examining scatter plots of detrended data. Figures 2.9 to 2.12 are direct equivalents of Figures 2.5 to 2.8 but using detrended data: here the linear-in-time trend component of the models (to be presented in Section 3) has been used to create values for conditions at the end of 1985. For present purposes it is sufficient to note that the historical data show rather greater trends at Tower Pier and Richmond than at Southend, and that separate values for the trend rate have been used for the three sites: 15, 39 and 42 centimetres per century at Southend, Tower Pier and Richmond, respectively. In comparison to Figures 2.5 to 2.8, the "detrended" scatter plots of Figures 2.9 to 2.12 seem to reveal a very slighted improved predictability of river-level from Southend sea-levels compared with the original data, but no apparent improvement in predictability based on flows at Teddington.

2.4 Simple data analyses

Figure 2.13 shows histograms of the sea- and river-levels at Southend, Tower Pier and Richmond. For the period of observed data used here, the minimum and maximum values of the tidal-peak levels at the three sites are shown in Table 2.1. Note that the lowest high-tide levels for the sites here (and for Gallions) occurred on 15 October 15 1983, on which occasion there was a deep depression moving over Scotland and high winds of 60 knots were widespread, with 80 knots in some localities. This sub-period of the record is shown in Figure 2.4. The maximum levels for Southend and Tower Pier occurred for the tidal-surge flood event of 1953 (1 February 1953, for which the tidal-peak level at Richmond was 5.26 m AOD, shown in Figure 2.3). The maximum river-level at Richmond (5.47 m) occurred on 31 December 1978, when the maximum level reached at Tower Pier was 5.21 m AOD.

The histograms in Figure 2.13 seem to show some granularity or preferred-values in the water-level observations and this is also evident in some of the earlier scatter-plots, most notably Figures 2.7 and 2.8. Any attempt to deal explicitly with this problem would be complicated by the fact that the present records have been derived from manuscript records of levels in units and precisions which have changed over time. For example, a brief scan of the manuscript records led to Table 2.2 as a summary of the precision for the values recorded for Southend. Here, the classification into "occasional values" and "some values" indicates that, respectively, a lot less than one half, and somewhat less than one half, of the values are recorded at the half-interval, whereas a consistent recording to the suggested half-interval precision would lead to half of the values being recorded at the half-interval. In addition to the change from imperial to metric units, the original data were recorded to various local datum-levels in different periods and these have been converted to a standard scale. Taken together, these data-recording effects mean that any "preferred" values in different periods of time are unlikely to be "round" numbers when expressed in metres AOD.

Table 2.1 *Minimum and maximum values of the tidal-peak levels recorded during 1939 to 1985.*

Site	Minimum level (m AOD)	Maximum level (m AOD)
Southend	-0.06	4.61
Tower Pier	0.63	5.41
Richmond	0.70	5.47

Table 2.2 *Apparent recording precision for the record of sea-levels at Southend*

Period	Apparent Precision
January 1939 - July 1946	nearest inch (with occasional values to ½ inch)
August 1946 - September 1949	nearest ½ inch
October 1949 - February 1950	nearest inch
February 1950 - March 1950	nearest ½ inch
April 1950 - May 1950	nearest inch
May 1950 - March 1953	nearest ½ inch
April 1953 - December 1953	nearest inch (with occasional values to ½ inch)
January 1954 - August 1966	nearest 0.1 foot
September 1966 - December 1973	nearest 0.1 foot (with some values to 0.05 foot)
January 1974 - December 1985	nearest 0.01 metre

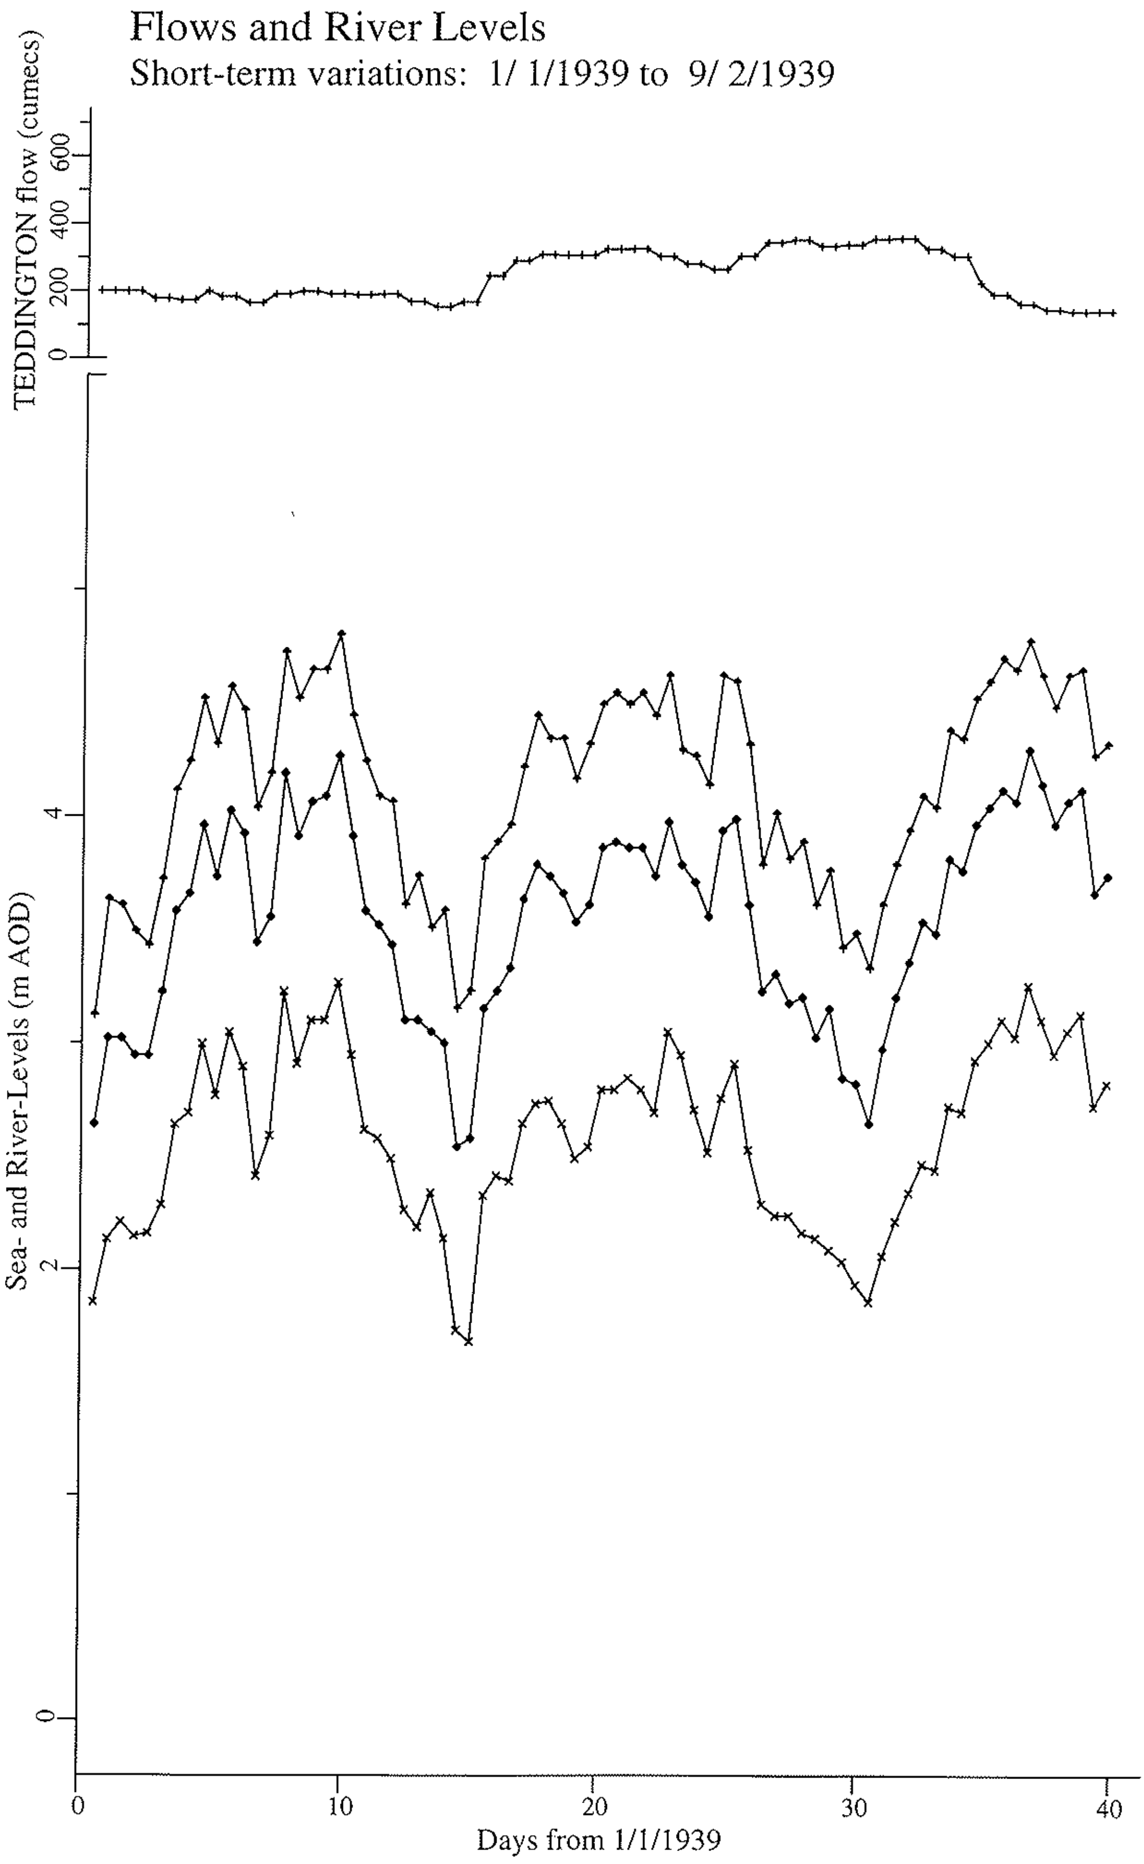

Figure 2.1 Example of short-term variations of observed levels and flows.

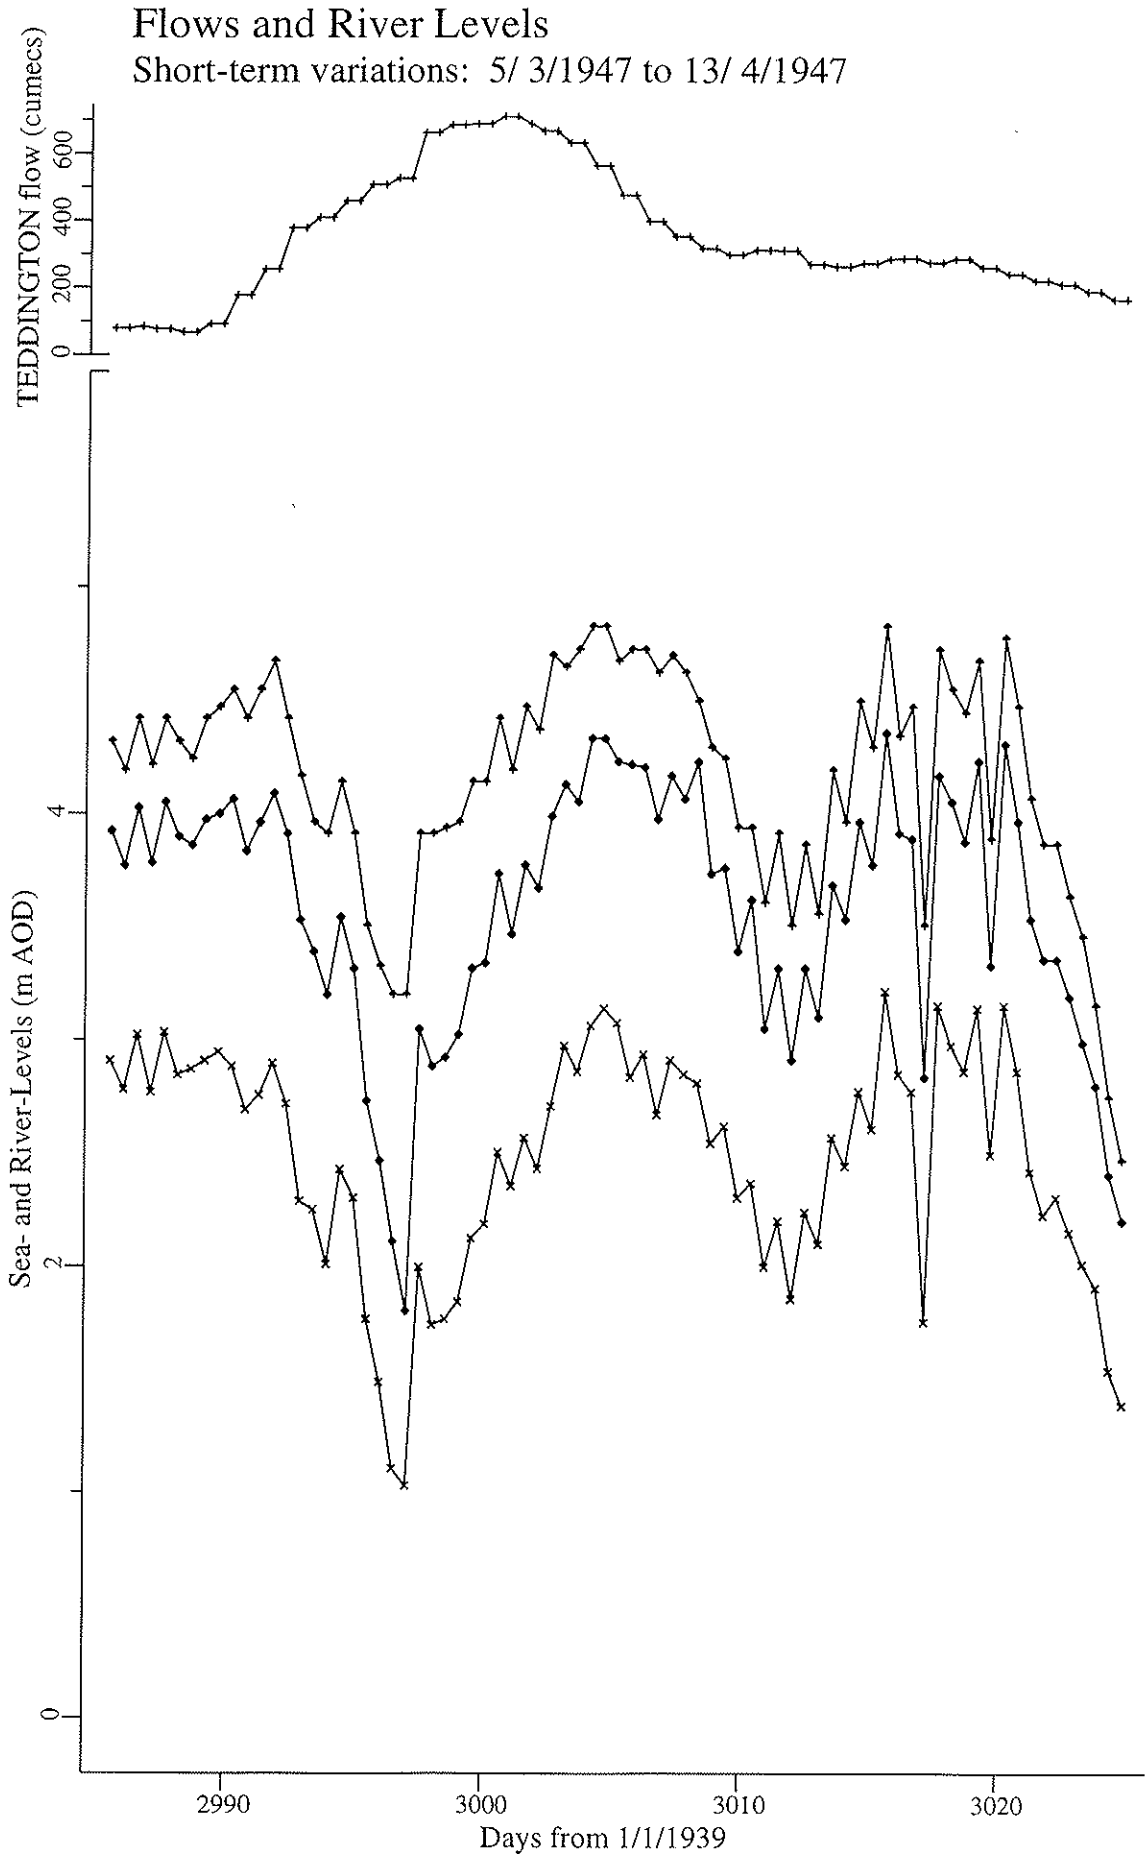

Figure 2.2 Example of short-term variations of observed levels and flows.

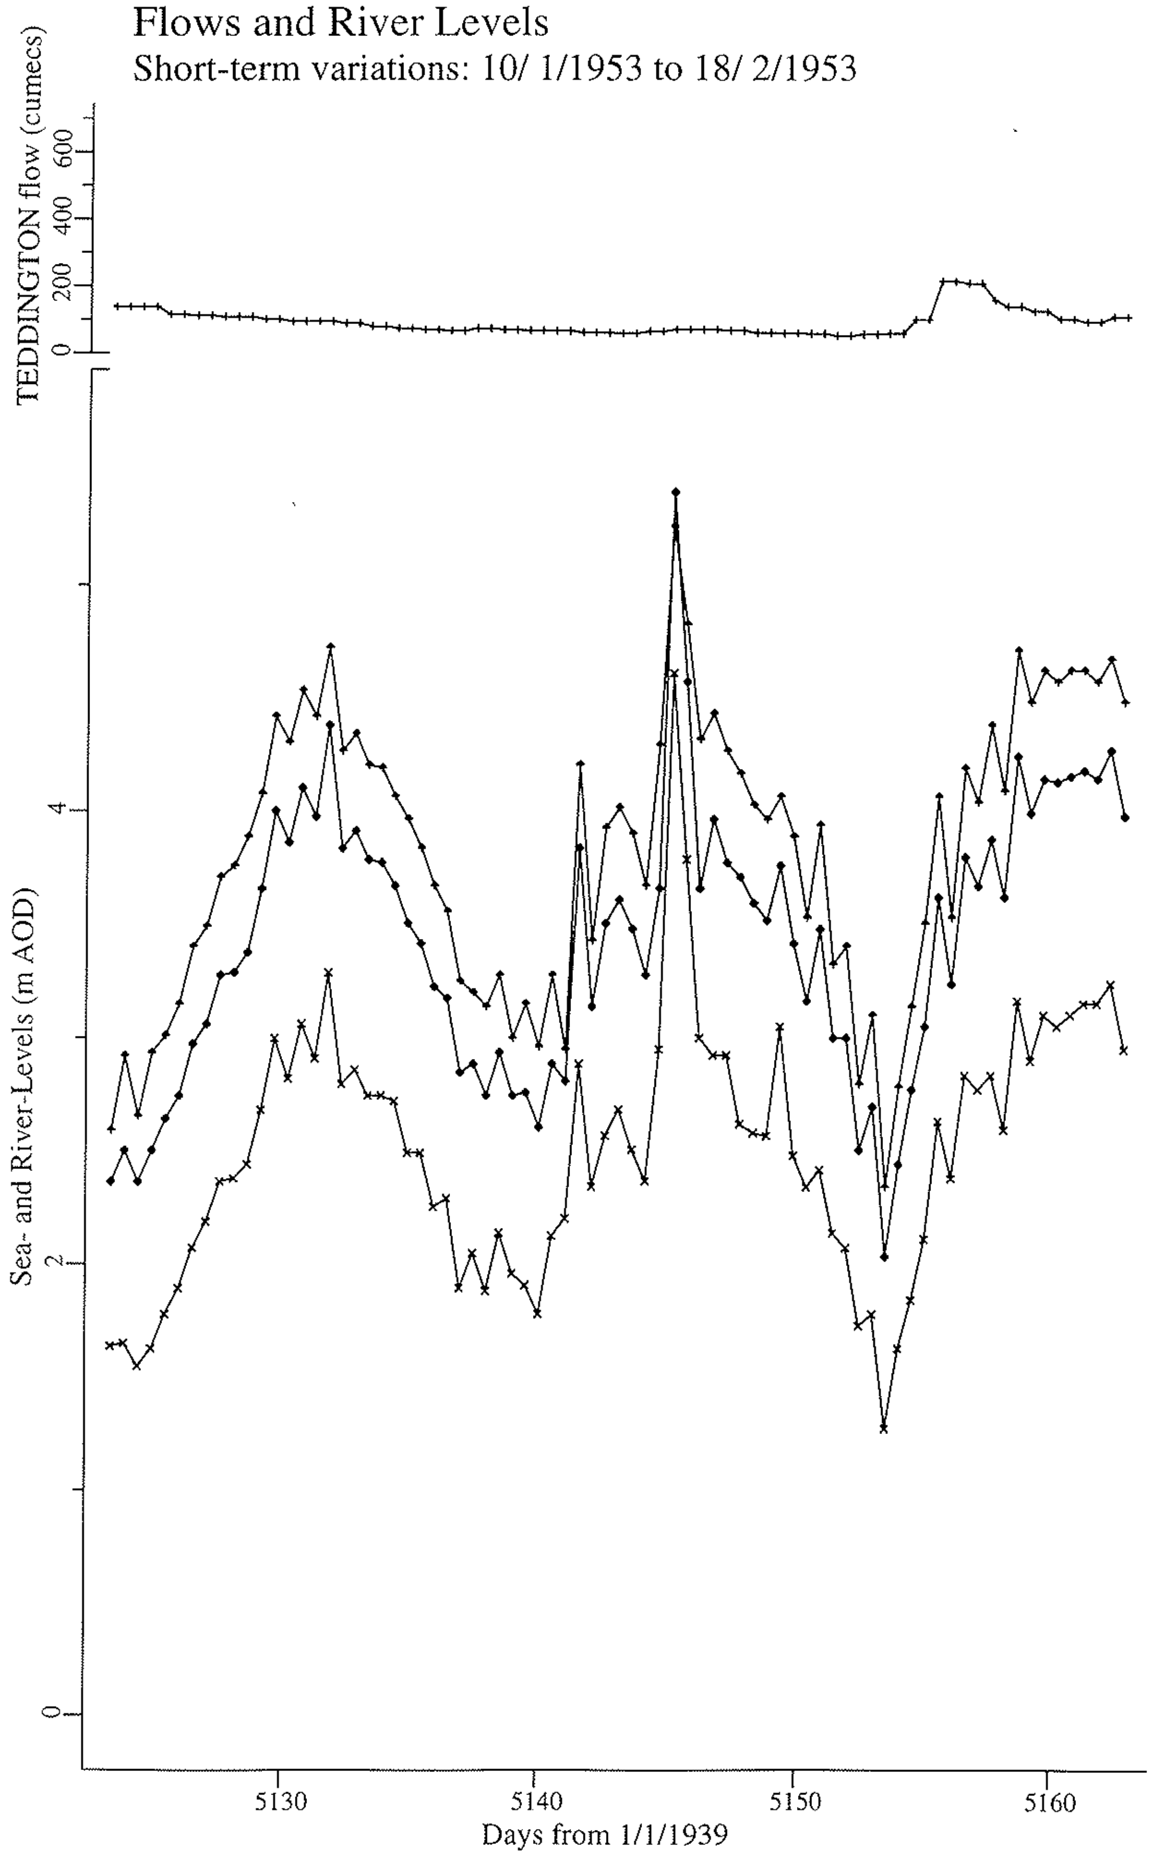

Figure 2.3 Example of short-term variations of observed levels and flows.

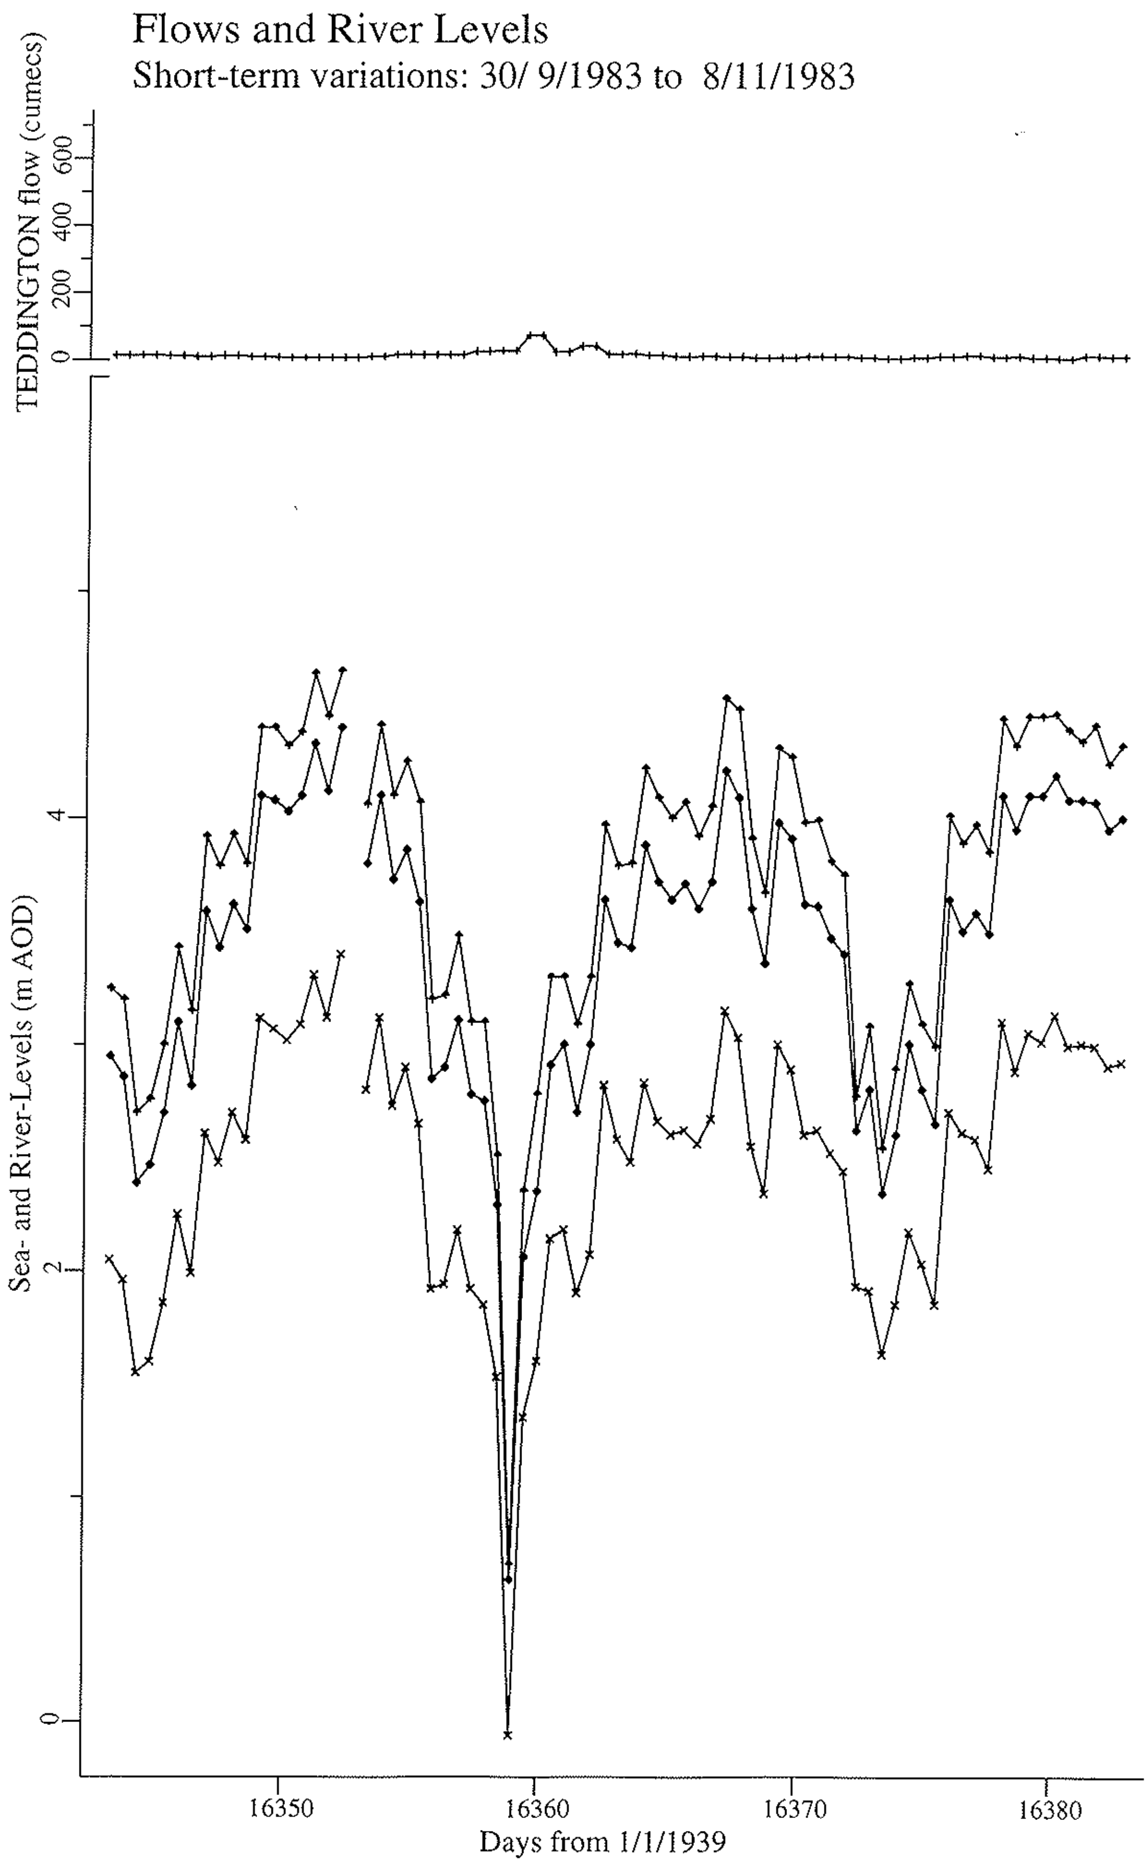

Figure 2.4 Example of short-term variations of observed levels and flows.

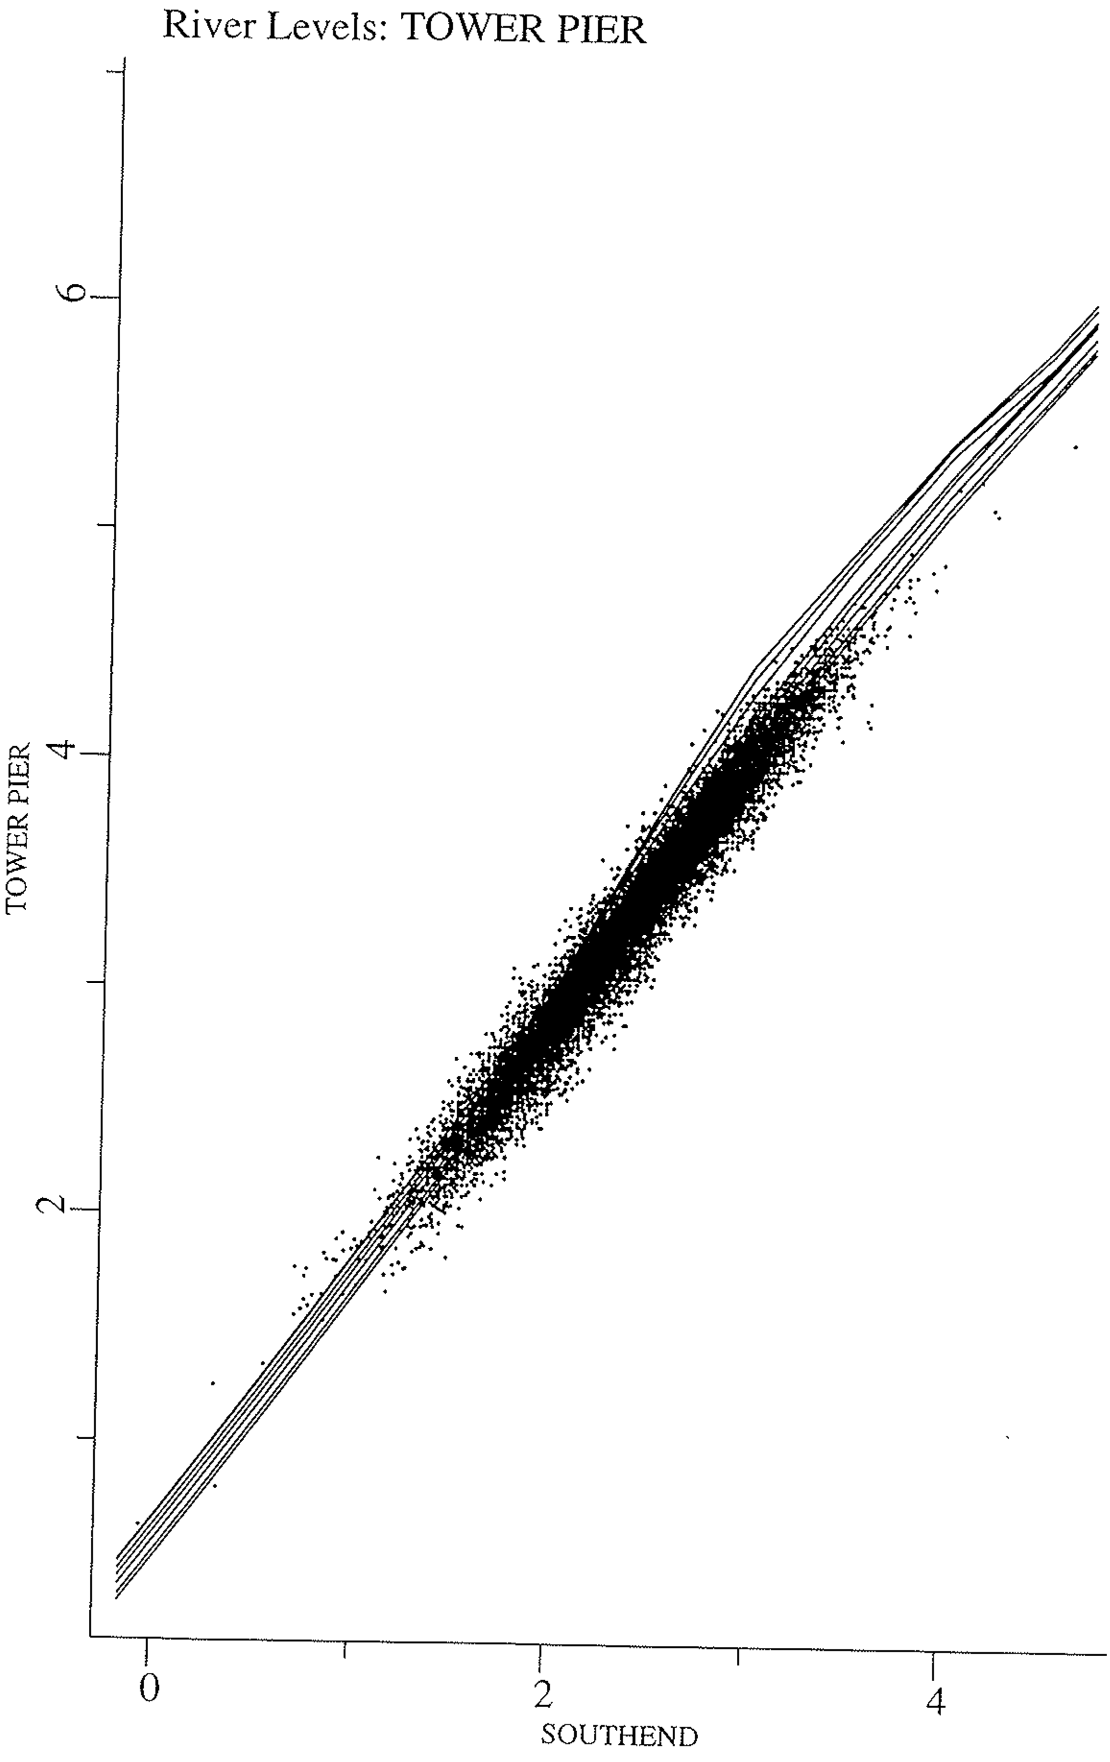

Figure 2.5 Scatter plot of levels at Tower Pier against levels at Southend, with the structure function for 50, 100, 200, 300, 400, 600, 800, 1000 cumecs.

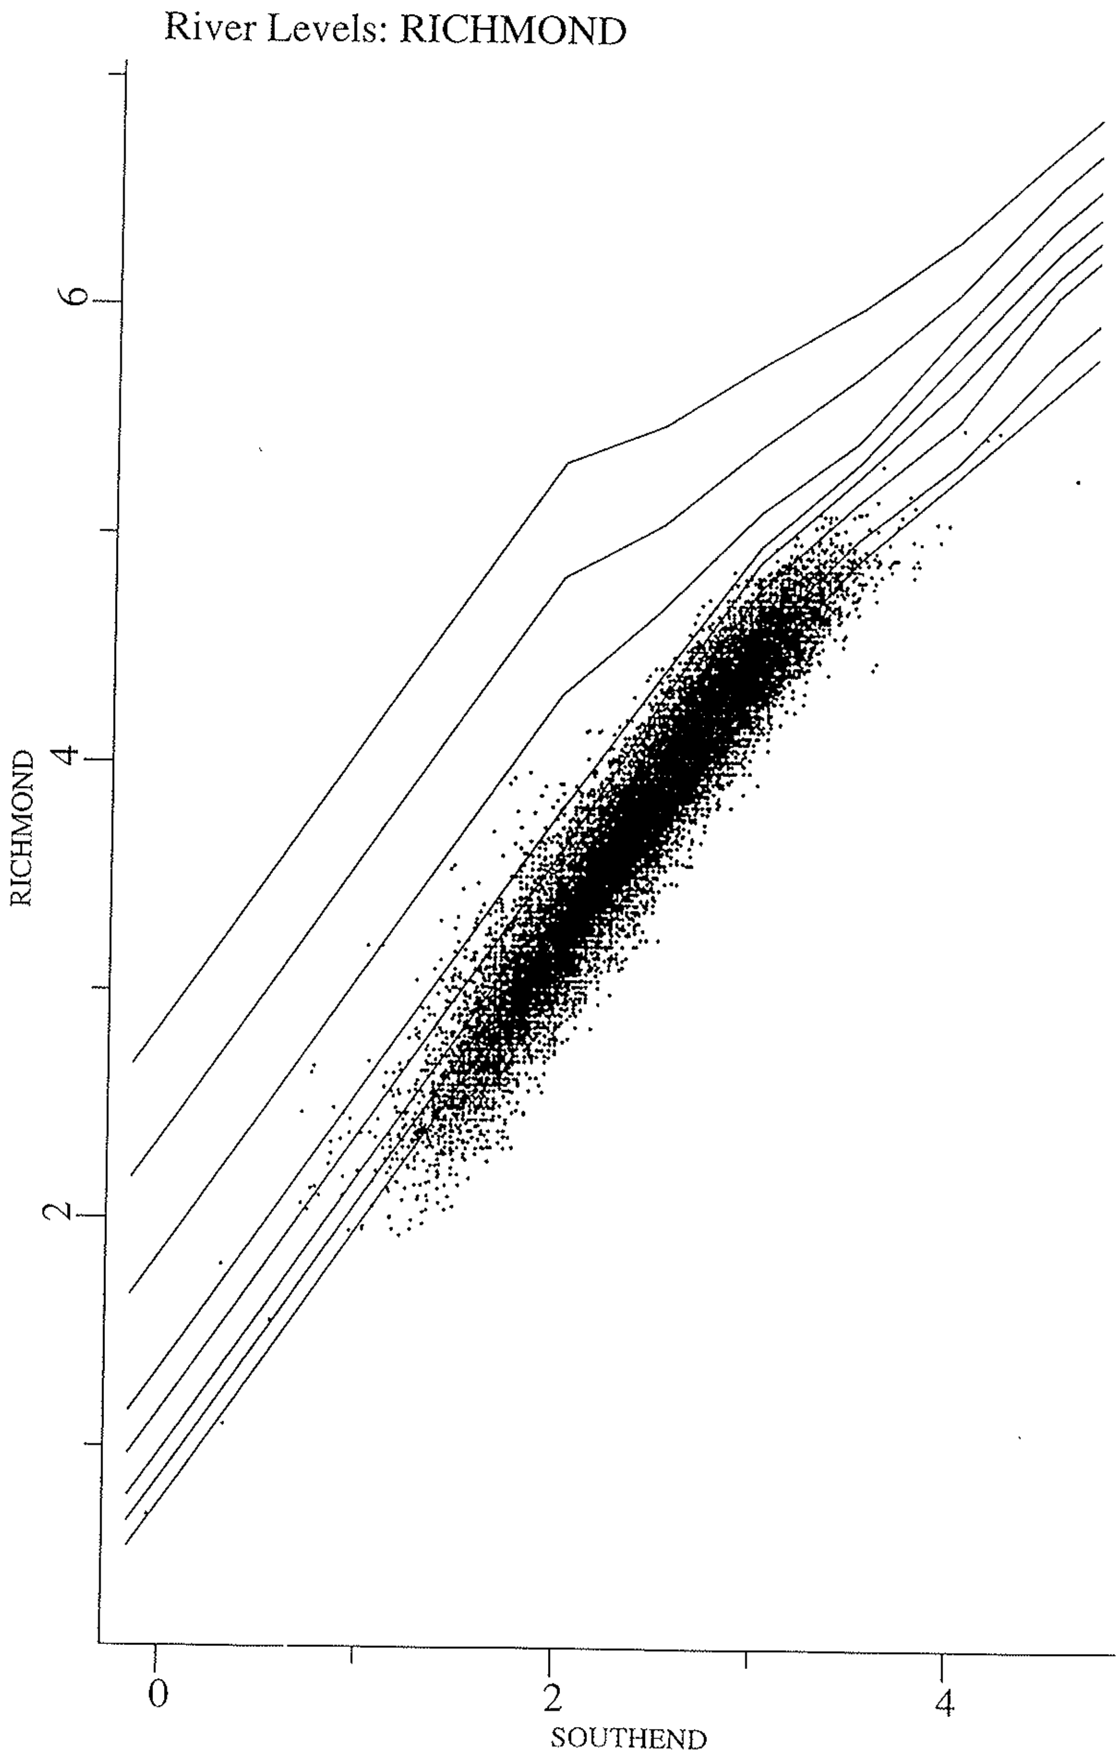

Figure 2.6 Scatter plot of levels at Richmond against levels at Southend, with the structure function for 50, 100, 200, 300, 400, 600, 800, 1000 cumecs.

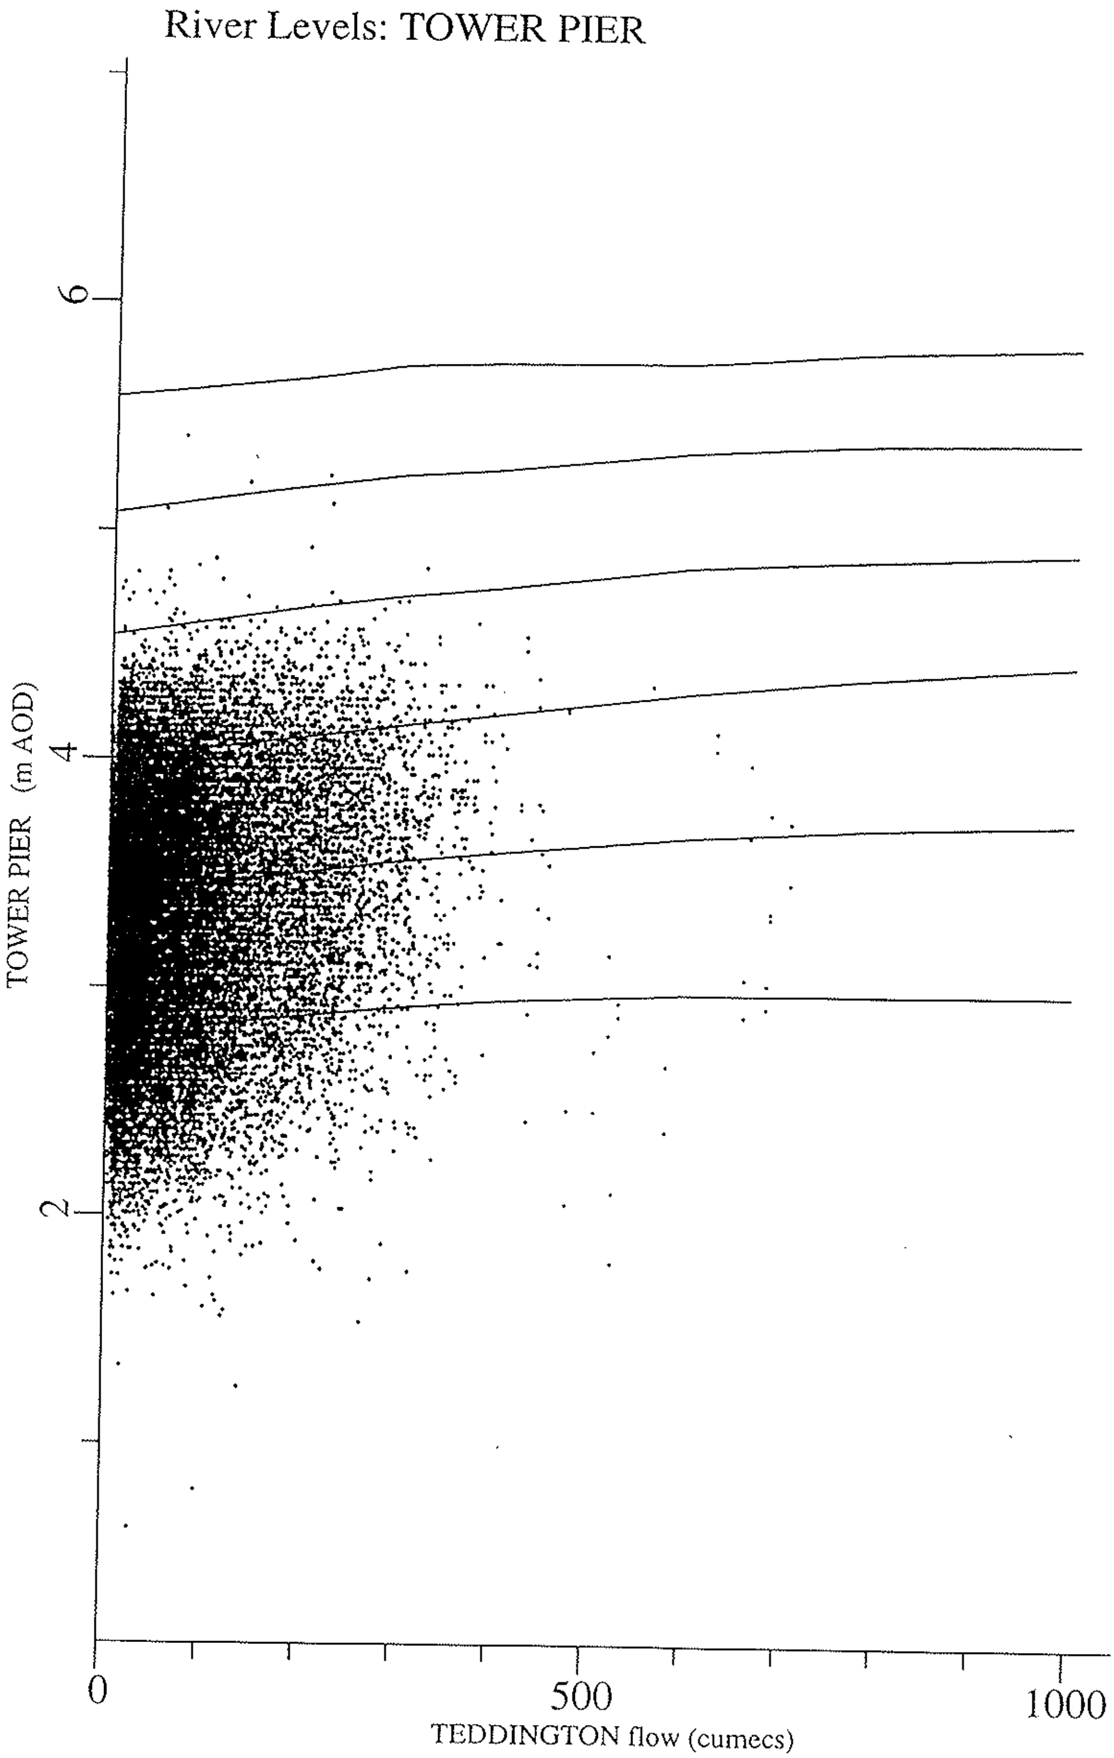

Figure 2.7 Scatter plot of levels at Tower Pier against flows at Teddington, with the structure function for Southend levels of 2, 2.5, 3, 3.5, 4, 4.5 m.

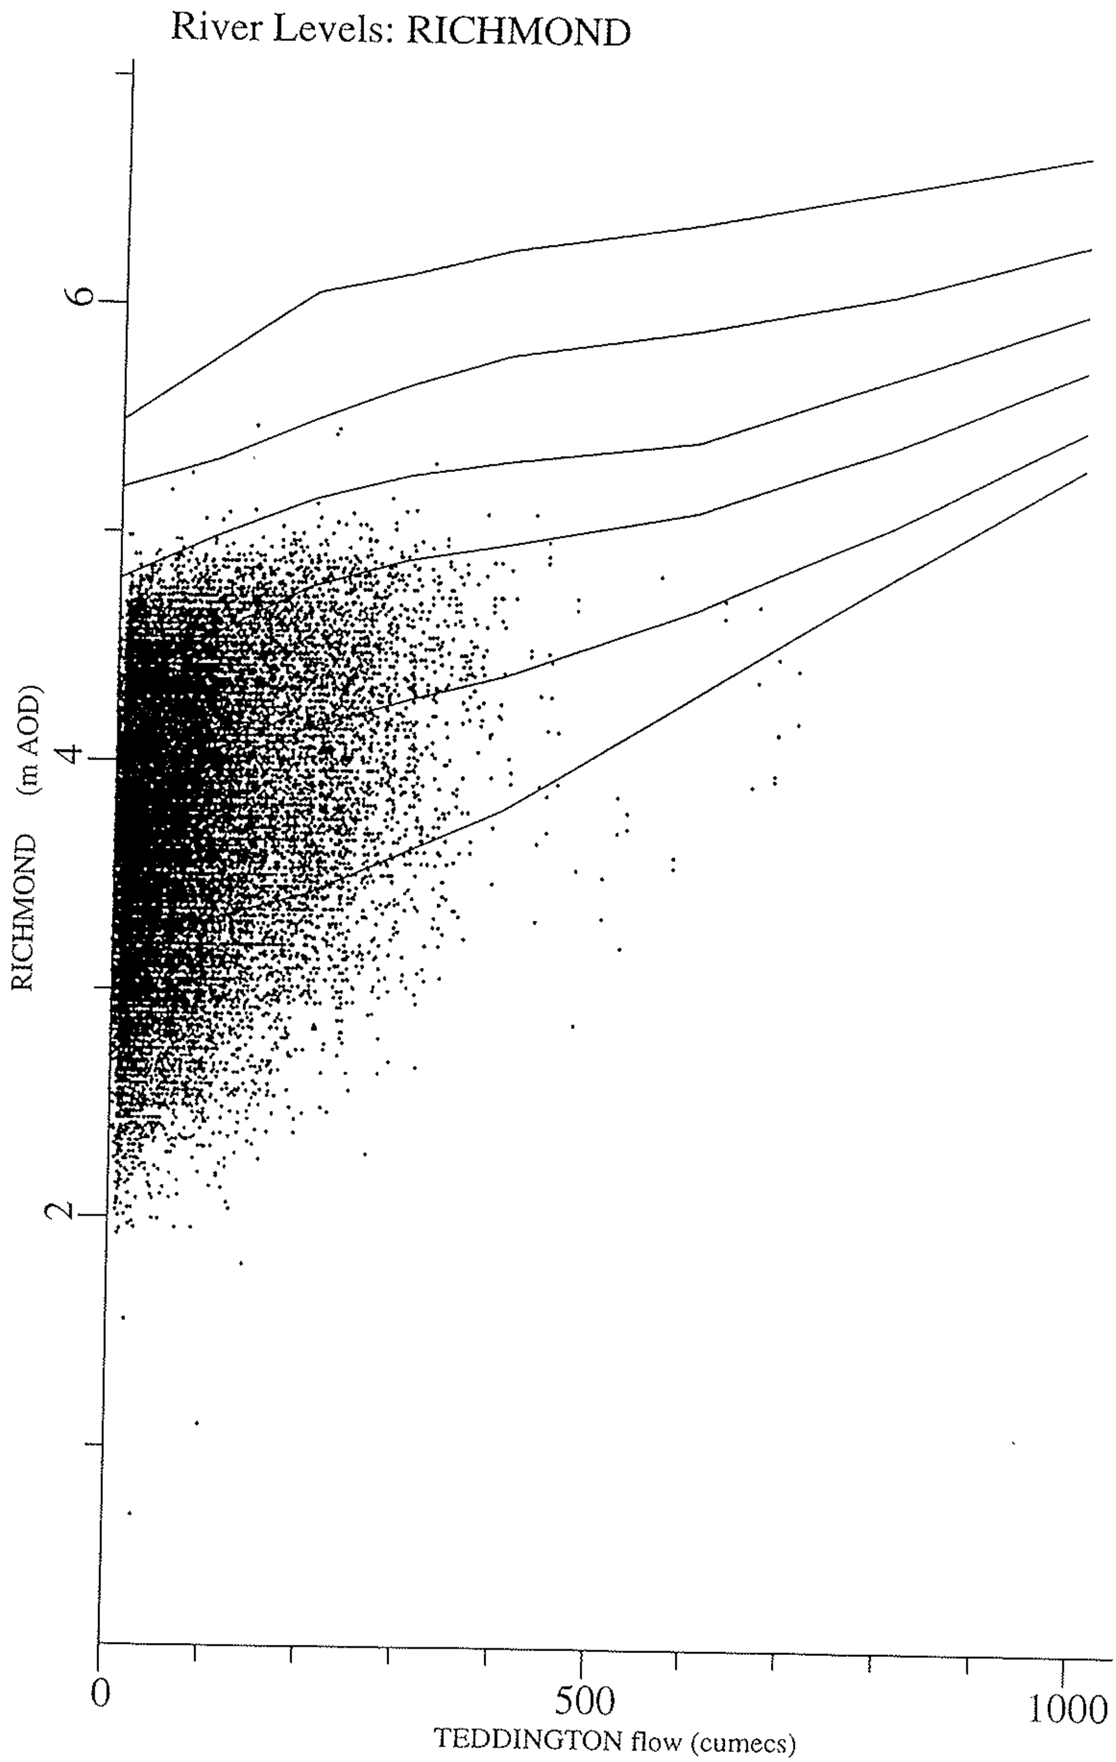

Figure 2.8 Scatter plot of levels at Richmond against flows at Teddington, with the structure function for Southend levels of 2, 2.5, 3, 3.5, 4, 4.5 m.

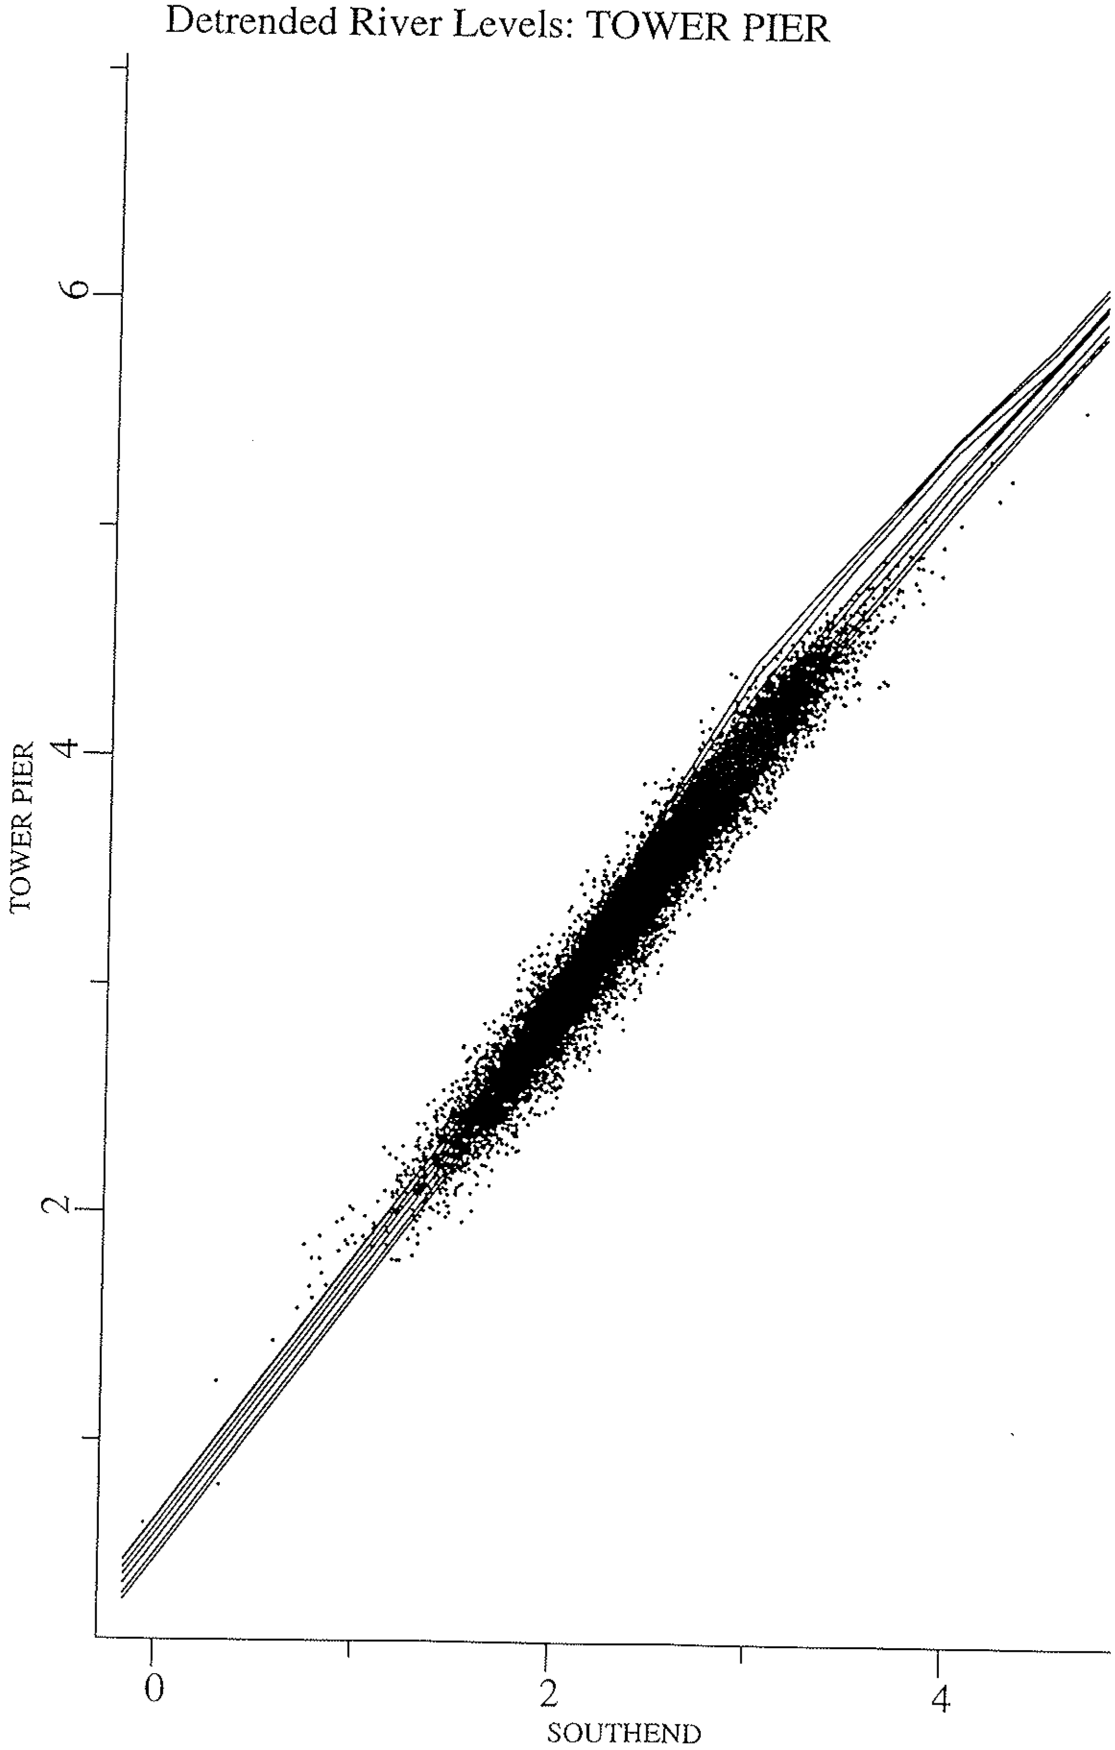

Figure 2.9 Scatter plot of levels at Tower Pier against levels at Southend. The levels have been detrended. The structure function is as in Figure 2.5.

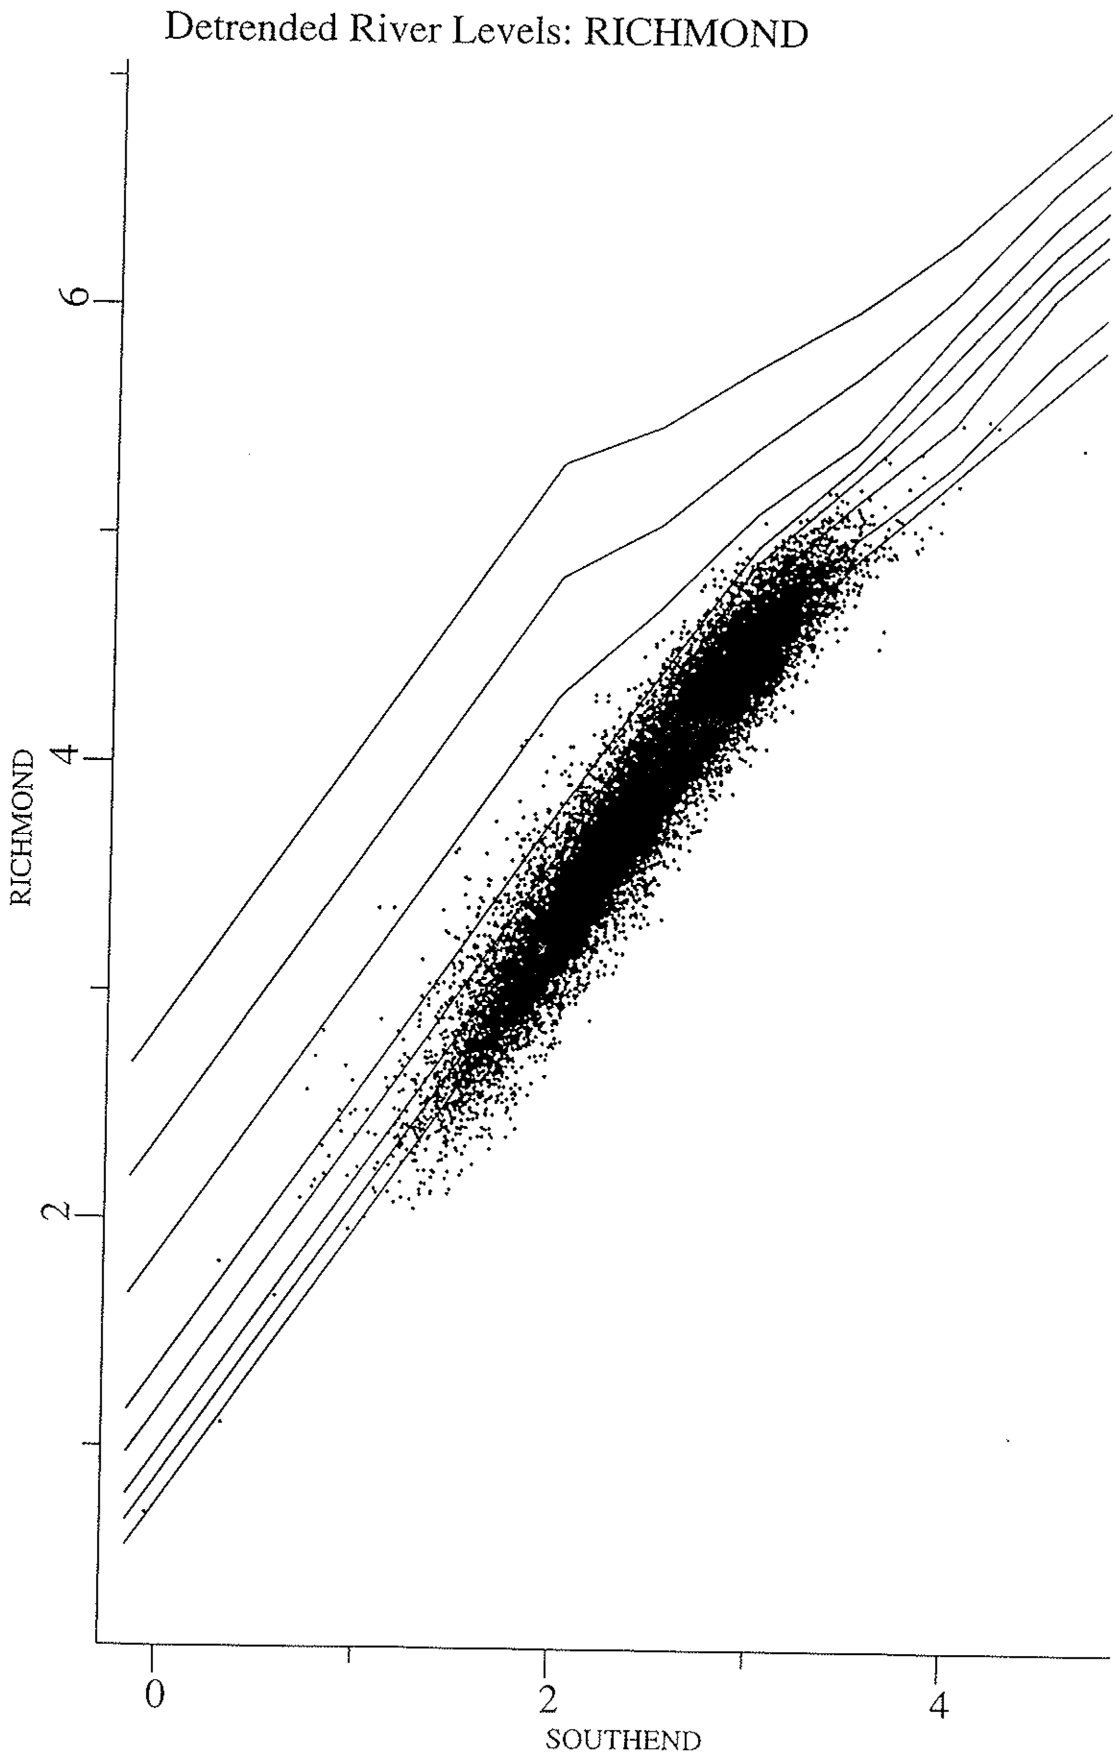

Figure 2.10 Scatter plot of levels at Richmond against levels at Southend. The levels have been detrended. The structure function is as in Figure 2.6.

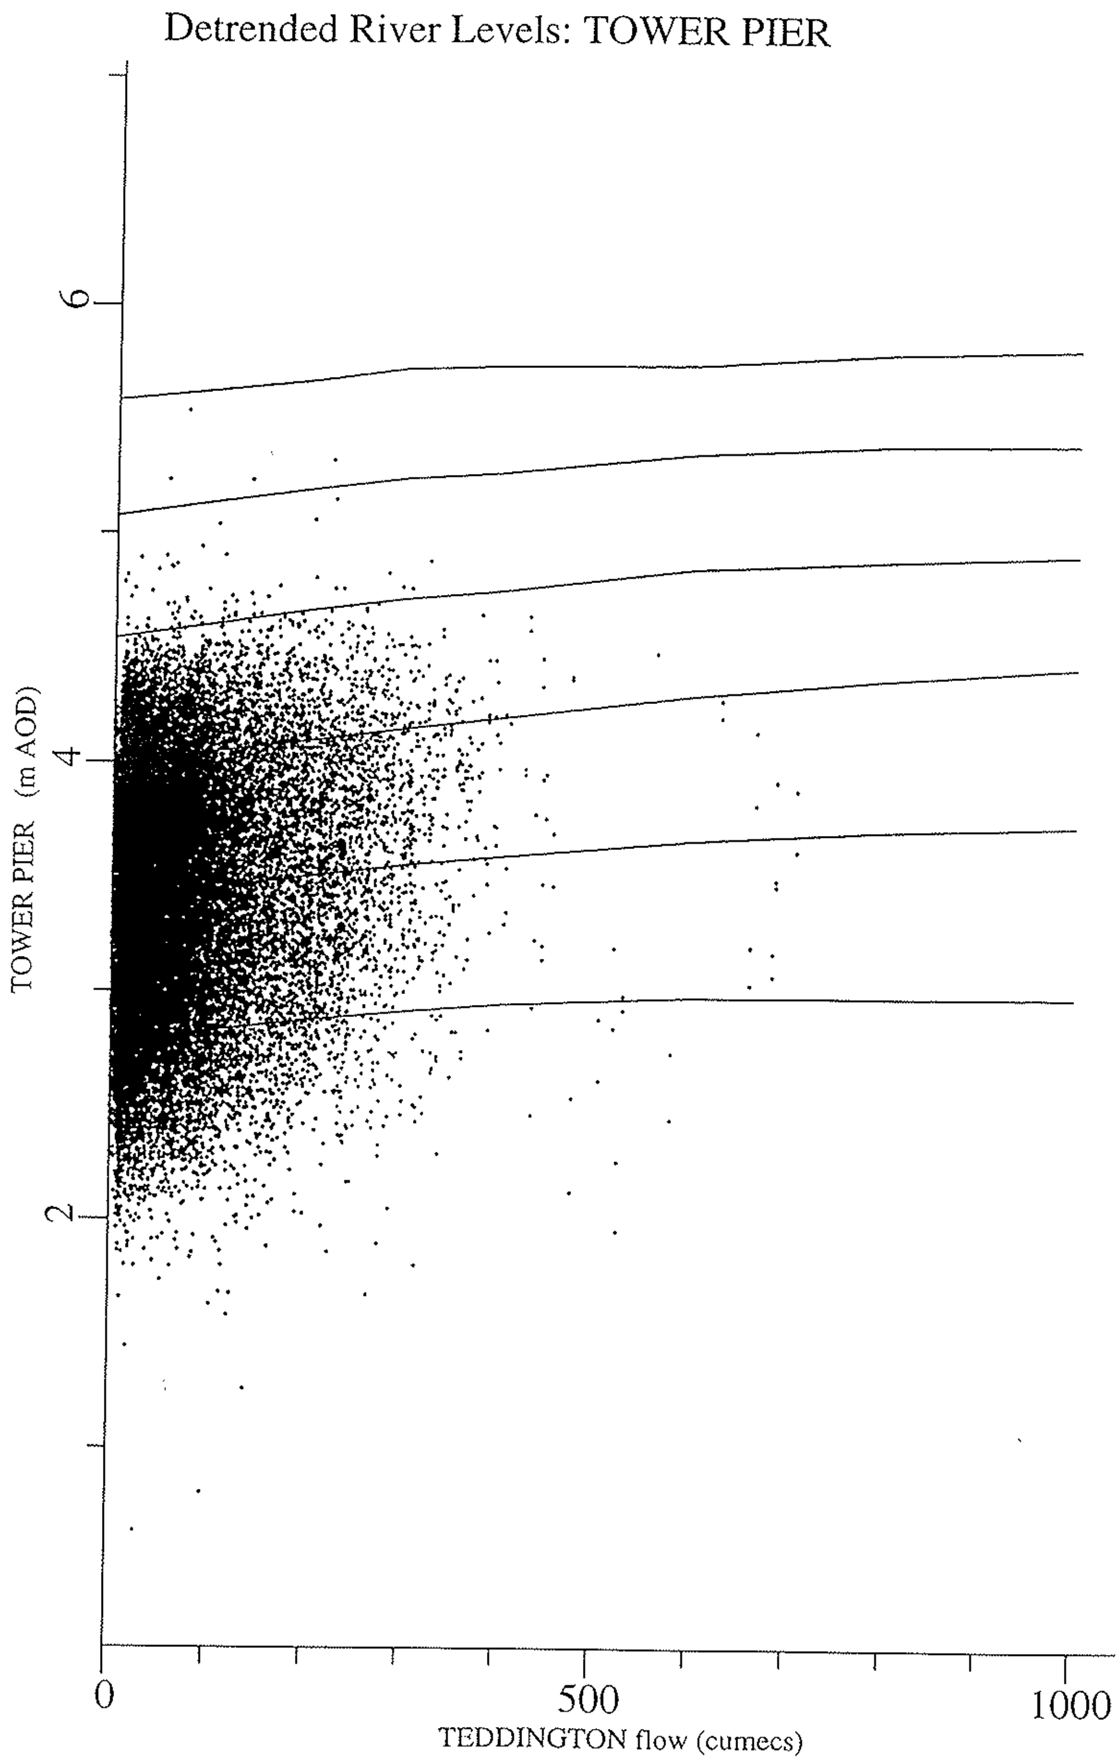

Figure 2.11 Scatter plot of levels at Tower Pier against flows at Teddington. The levels have been detrended. The structure function is as in Figure 2.7.

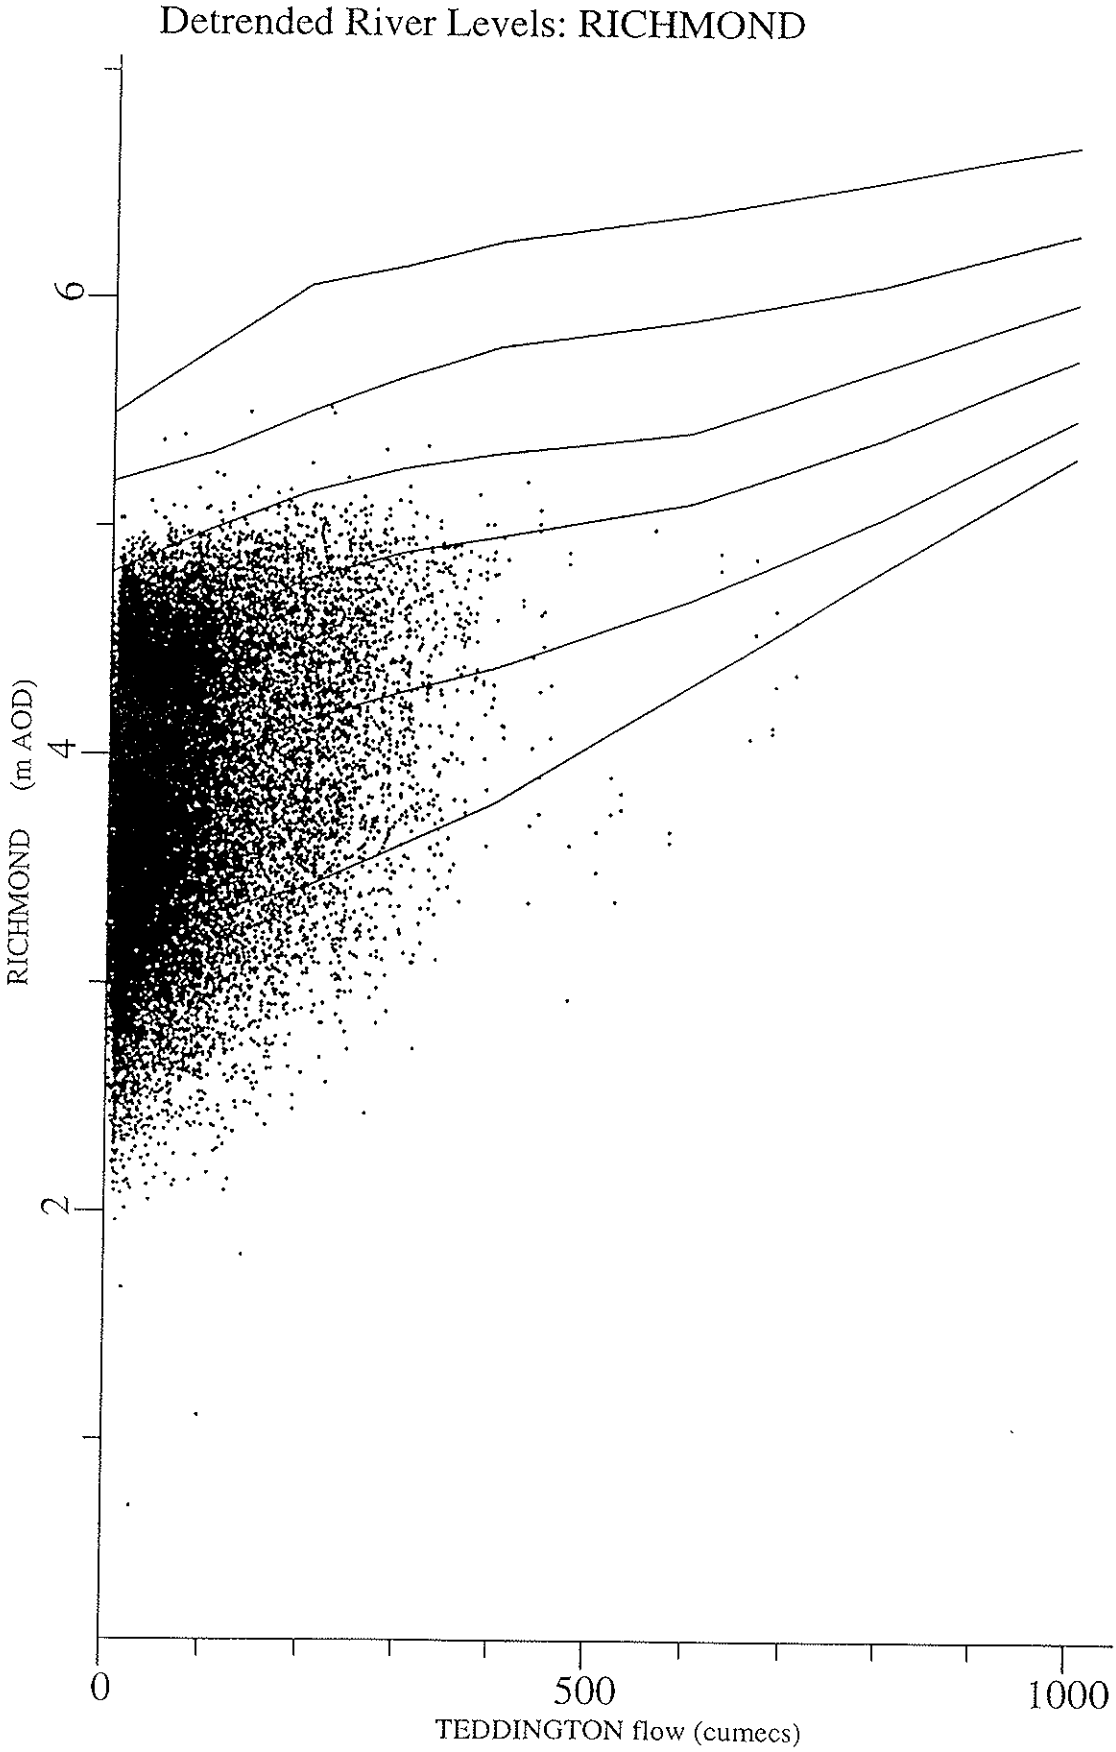

Figure 2.12 Scatter plot of levels at Richmond against flows at Teddington. The levels have been detrended. The structure function is as in Figure 2.8.

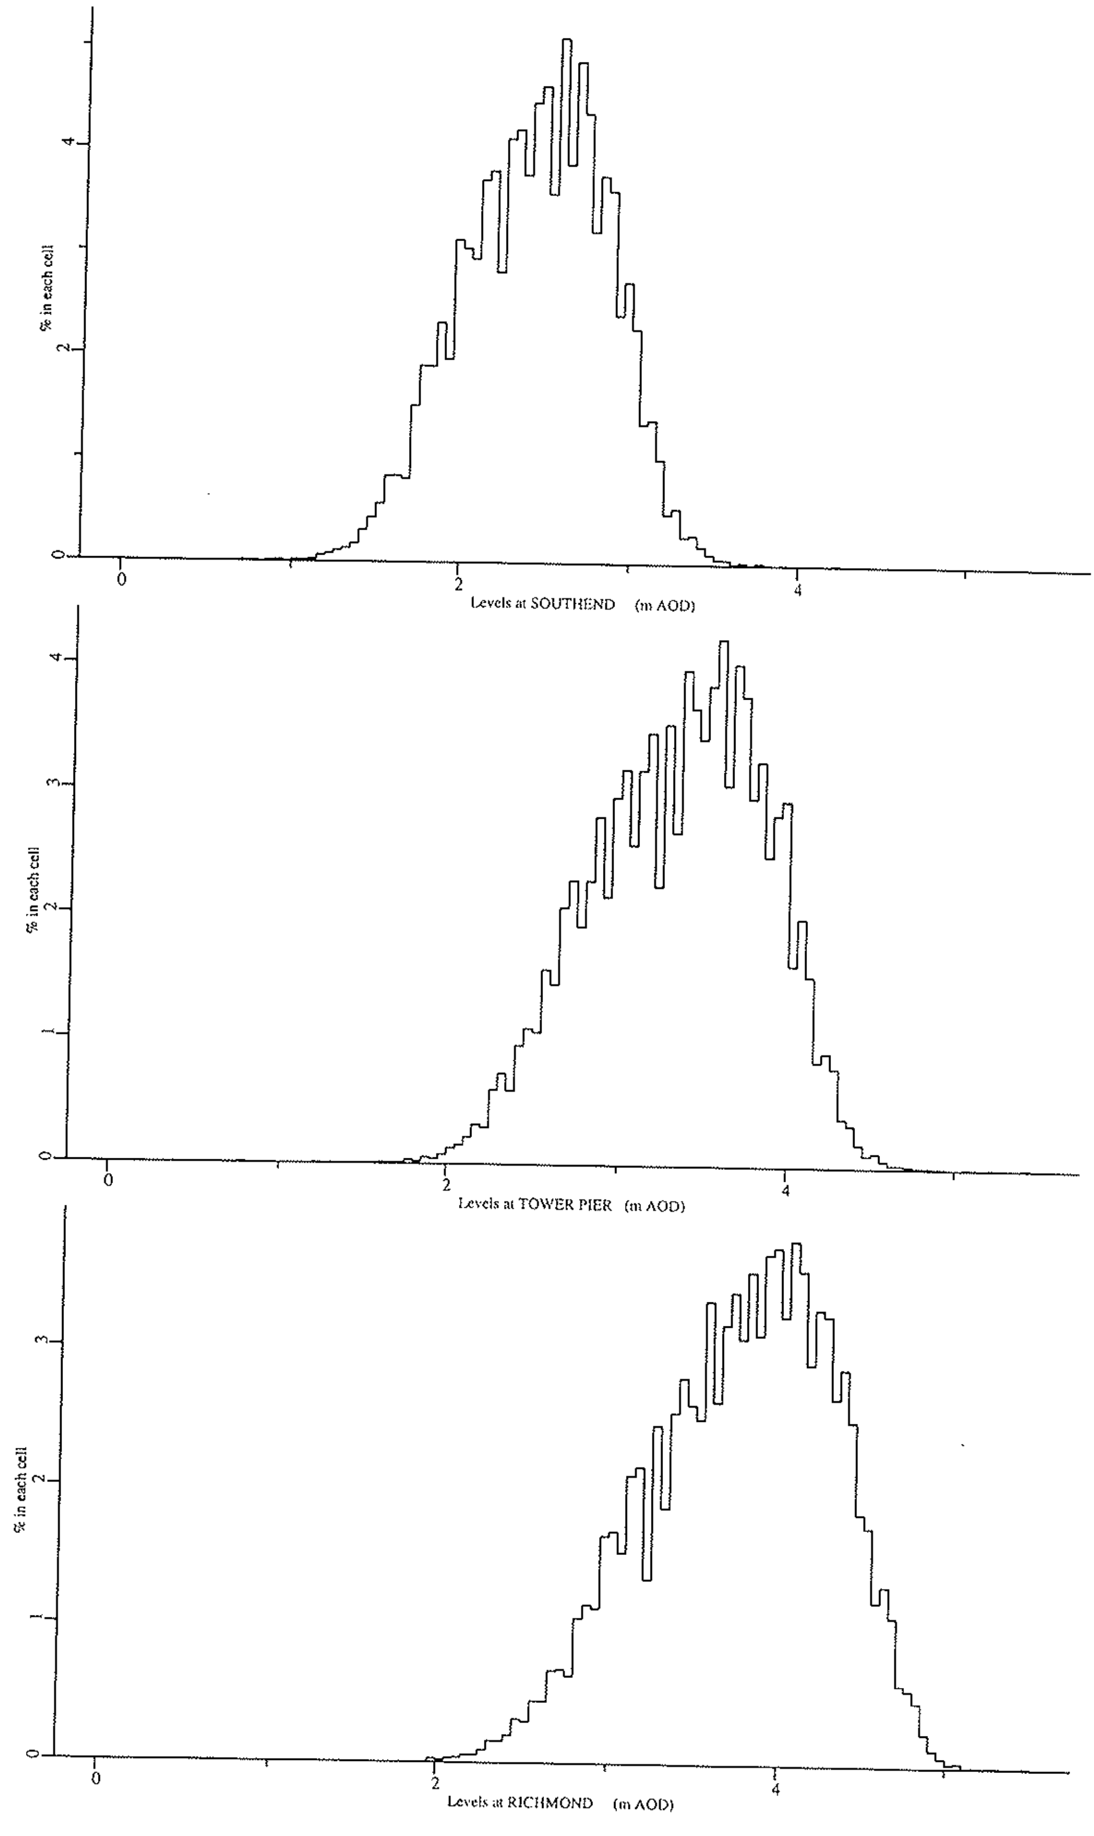

Figure 2.13 Histograms of observed levels at Southend, Tower Pier and Richmond.

3. Analysis of Structure Functions

3.1 Structure Functions based on Tidal-peak and River-flow

Given that a set of structure functions is available from the Tidal Thames Defence Levels study, it seems appropriate to start with these in an assessment of how well the peak river-levels reached along the tidal Thames can be predicted on the basis of the tidal-peak levels at Southend and the river-flows at Teddington. Earlier chapters have included discussion of the fact that the structure functions represent a fixed set of modelled-conditions, whereas the data-set represents an evolving set of real conditions. Furthermore, the series of sea- and river-levels exhibit trends of varying extent. Because of these trends and the doubts over the direct relevance of the structure functions to conditions pertaining during the historical record, it seems reasonable to base an initial assessment of the performance of the structure functions on modified versions of these in which extra allowances are made for bias and trend.

The following notation will be used here. Let the subscript j denote the sequence number for the tidal-peak in the set of high-tides described in Section 2.2., and let the following definitions be made for values referring to peak number j.

S_j Southend peak level (m AOD);
T_j Tower Pier peak level (m AOD);
R_j Richmond peak level (m AOD);
Q_j Daily mean flow (cumecs) at Teddington;

t_j Time reference (solar years from 1 January 1939).

Here, the time reference t_j is as described in Section 2.2, except that here it is convenient to describe the model in terms of solar years rather than solar days: later on, both units are used as part of the description.

The structure functions will be denoted by F_T and F_R, in the cases of Tower Pier and Richmond respectively. Thus, for generic tidal-peak values, T is predicted by $F_T(S,Q)$, and R is predicted by $F_R(S,Q)$: ie. for a peak sea-level at Southend of S, and a daily mean flow at Teddington of Q, the basic structure function predicts that the peak river-level at Tower Pier will be $F_T(S,Q)$.

In order to assess the performance of the structure functions, modified predictors of the following form are considered,

$$\text{est}\,(R_j) = a + b\,\{t_j/100\} + c\,F_R(S_j, Q_j), \qquad (3.1)$$

with a similar form for the predictions at Tower Pier. Here the coefficients a, b and c are either fixed at their default values of 0, 0 and 1, respectively, or are fitted by the method of least squares to the whole of the data-set. The scaling of the time-reference point is chosen so that the parameter b measures the trend in metres per century. Because of the potential for trend in all the sea- and river-level series, it is clear that the trend parameter "b" measures a relative trend in this context.

Since the purpose here is to provide an assessment of the performances of the structure functions in predicting peak river-levels, it seems reasonable to construct a simple alternative predictor, based on the same set of information. If such a simple alternative did substantially better than the structure functions, this would be strong evidence against continued use of the structure functions. Since the structure functions derive from hydrodynamic modelling of the reaches concerned, one would expect that these would have good properties beyond the range of conditions experienced within the historical data-set. Hence, if the simple predictor did turn out to perform comparably to the structure functions over the historical period, one would not be justified in applying Occam's Razor to justify discarding the structure functions. A convenient simple prediction model for peak river-levels is one of the following form:

$$\text{est}\ (R_j) = a + b\ \{t_j\ /100\} + s\ S_j + f_1\ \{Q_j\ /500\} + f_2\ \{Q_j\ /500\}^{1/2}. \qquad (3.2)$$

Here the particular combination of terms involving the flow, Q_j, was chosen to be similar to one found useful in the regression analyses reported later. Note that this expression for the predictor contains no interaction effects between sea-levels at Southend and the flows at Teddington: no investigation of the possibility of improving this empirical predictor by including interaction effects has been made, since the purpose here was to use the predictor in Equation (3.2) to check that the structure function has no important deficiencies in comparison with simple predictors.

The performance of the predictors is assessed in the usual way, that is in terms of the properties of the prediction errors, ε_j, defined by

$$\varepsilon_j = R_j - \text{est}\ (R_j).$$

As well as considering the usual root mean square error (rmse) measure of performance, results are given for the minimum and maximum errors. Table 3.1 summarises the results and these are discussed below.

For the reasons stated at the beginning of this Section, the results for Model 3 in Table 3.1 are be taken as the best assessment of the likely errors in using a structure function approach to predicting river-levels, once adjustments are made at the hydrodynamic-modelling stage so that a particular set of assumed flood-defence conditions are incorporated. The root mean square errors for this model are 8.6 and 12 cm for Tower Pier and Richmond respectively. It can be argued that these values are in fact too high because of the changing flood defence levels over the historical period. In particular, there was a partial failure of flood defences during the 1953 surge event which resulted in peak-levels at upstream sites being rather lower than they might have been. Since the analyses here essentially involve comparisons of predictors over identical time-periods, the conclusions reached are unlikely to be affected by considerations such as these. In addition, the number of tidal-peaks included in the analysis (approximately 30,000) mean that the average performance will not be strongly affected by poor performance on only a few instances.

A comparison of Model 3 with the simpler Models 1 and 2 in Table 3.1 does not indicate a very strong effect on the root mean square error from adjusting for bias and trend, although the relative trend coefficients of 15 and 11 cm per century show that moderately-sized adjustments for trend are being made over the historical period (nearly 50 years). The sizes of the scaling coefficients, c, for the structure function in Model 4 of Table 3.1 (which are slightly smaller than unity) suggest that reducing the variability arising from the original predictor can make only a minor improvement and this is confirmed by the results for the root mean square error. However, one can note that the reductions of the root mean square errors arising from fitting the scaling coefficient are of the same order as those involved in fitting the coefficient for trend (going from Model 2 to Model 3) and thus the importance of these effects is comparable. For Model 5, the scaling coefficients, s, on the tidal peak levels at Southend (which are somewhat larger than unity) are suggestive of an "amplification factor" for sea-level effects. Note that later models (described in Chapter 4) suggest that this is misleading: a better summary may be that, while tidal cycles are amplified, "surge residuals" are not. The results for Model 5, which is the simple alternative to the structure function predictor, indicate a slight improvement over model 3 in terms of root mean square error. However, since this simple model involves fitting two more parameters than are being fitted for Model 3 (and moreover, has involved a data-based choice of model-structure in reaching the form in Equation (3.2)), the difference does not seem unwontedly large.

Table 3.1 *Comparison of observed levels with the predictions of river-level produced by modified versions of the existing structure functions.*

Model	Term	Tower Pier	Richmond	Error Criterion	Tower Pier	Richmond
1. unadjusted	a	0	0	rmse	0.0904	0.1228
	b	0	0	min	-0.5019	-0.7877
	c	1	1	min	0.4842	0.6050
2. adjusted for bias	a	0.0185	0.0205	rmse	0.0885	0.1211
	b	0	0	min	-0.5204	-0.8082
	c	1	1	max	0.4657	0.5845
3. adjusted for trend and bias	a	-0.0177	-0.0062	rmse	0.0861	0.1201
	b	0.1513	0.1125	min	-0.5052	-0.8040
	c	1	1	max	0.4490	0.6078
4. adjusted for trend, bias and scaling	a	0.0922	0.1095	rmse	0.0843	0.1189
	b	0.1569	0.1182	min	-0.5093	-0.7728
	c	0.9667	0.9688	max	0.4359	0.5554
5. empirical model based directly on level and flow	a	0.4450	0.6662	rmse	0.0851	0.1151
	b	0.1542	0.1171	min	-0.5131	-1.1557
	s	1.1575	1.2018	max	0.4365	0.6625
	f_1	0.0657	0.1668			
	f_2	0.1667	0.4383			

With the scaling included in Equation (3.2), the size of the effect attributable to flow can be judged by noting that $(f_1 + f_2)$ is the difference between the levels predicted for flows of 500 and 0 cumecs. The results for Model 5 of Table 3.1, which relates to Equation (3.2), give these flow-effects as 23 cm for Tower Pier and 61 cm for Richmond. The corresponding values from the structure functions (Model 3) vary with the value taken for the sea-level at Southend: for Tower Pier, differences of 19, 28 and 25 cm are predicted for sea-levels of 2, 3 and 4 m; for Richmond the differences predicted are 96, 77 and 65 cm for the same set of sea-levels. Thus there is moderately good correspondence between these results, taking into account that the observed levels for Richmond at high flows are potentially less than they might have been had a full level of flood-defence been in place.

The main conclusion to be drawn from Table 3.1 is that the existing structure functions do a reasonably good job in predicting river-levels from Southend sea-level and Teddington flow. There is no strong suggestion that a simple modification of the structure function, in the ways explored here, would produce substantially better results. All of the comparisons need to be moderated by the fact that the observed set of sea- and river-levels have not arisen under the set of conditions assumed in the construction of the structure functions. Although there is some suggestion that results as good as those for the structure function could be achieved by fitting a simple empirical model of a type similar to Equation (3.2), thus avoiding the hydrodynamic modelling underlying the structure functions, this ignores two points:

(i) there is a need to be able to predict river- and estuary-levels for sites for which there are no historical records;

(ii) there is a need to predict river-levels in conditions well outside the range experienced during the historical data-period.

The use of the hydrodynamic model within the construction of the structure functions should provide a sound basis for solving of these problems.

3.2 Re-analysis of original structure functions

The analysis reported in Section 3.1 appears to show that structure functions based on tidal-peak levels at Southend and flow at Teddington appear to do reasonably well in predicting river-levels at points on the Tidal Thames. However, the further analysis which is reported in this section suggests that these predictions can nevertheless be improved.

Figures 3.1 to 3.4 show the sample autocorrelation functions of the residuals from the models labelled 2 to 5 in Table 3.1, respectively. These plots reveal the strong presence of structure in the residuals. For both Tower Pier and Richmond, all four models show a seasonal effect with a yearly period. Models 4 and 5 also show a strong periodic effect with a cycle length of around 15 days: the same effect is present to a lesser extent for Models 2 and 3. It seems that allowing a scaling factor for the effect derived from Southend reveals a predictable component related to the usual tidal harmonics, in this case the neap-to-neap cycle of 14.765 days. In principle, both the yearly and neap-to-neap cycles could be included in a revised form of structure function to create an improved predictor. When this is done, the presence of cycles related to other tidal harmonics is revealed by the autocorrelation of the residuals. Essentially the same conclusions are found from a spectral analysis of the residuals. For reasons to be discussed in Section 3.3, it seems reasonable to limit analysis of direct extensions of the structure function models to variants including only a moderate number of tidal harmonics. The immediate intention in this section is to assess how much improvement in predictive performance is achieved compared with predictors which omit the seasonal and tidal cycles.

Table 3.2 lists the period-lengths of the cycles incorporated into the extended forms of the predictors introduced in Section 3.1. It should be recalled that, because the data are being treated as if they were regularly-sampled in time at a spacing of half a lunar day, or 0.517525 solar days, the Nyquist point corresponds to a period of 1.03505 days and a cycle with period length equal to one solar day would give equivalent results to one with a period of 1.072646 days, due to aliasing. The basic forms of estimators in Equations (3.1) and (3.2) are extended by the addition of pairs of terms, where each pair consists of sine and cosine terms of the required period, with coefficients to be estimated for each. An exception occurs for cycles with a period equal to a lunar day (harmonic component number 14 in Table 3.2), in which case only the cosine term is included: the reason for this is that the sine term is always zero because it is evaluated at integer multiples of half a lunar day.

Table 3.2 *Seasonal and tidal harmonics included in extended predictors.*

Term Number	Period Length	Cycle
1	18.613 years	
2	9.306 years	
3	6.204 years	
4	1 year	tropical year
5	½ year	
6	31.812 days	
7	27.554 days	
8	14.765 days	neap-to-neap
9	13.661 days	
10	7.383 days	
11	1.120 days	
12	1.076 days	
13	1.070 days	
14	1.035 days	lunar day
15	1 day	solar day

The selection of periods chosen for inclusion in the extended predictors in this section was based on a simple spectral analysis of residuals from a succession of prediction models. While an empirical spectrum might show a large peak, the period corresponding to this peak would not be precisely defined by the information in the spectrum: the actual period used for a new term was determined by a process in which combinations of the underlying astronomical frequencies were examined to find a corresponding period. A similar process was used for the models described later (in Chapter 4), for which rather more harmonic components were identified. Several of the periods included in Table 3.2 are aliases of components having a sub-daily period.

Table 3.3 shows results equivalent to those in Table 3.1, but for the extended predictors: in order to reduce the number of items of information reported in this table, the seasonal and tidal harmonic terms are summarised by quoting only the amplitudes (half-ranges) of the yearly and neap-to-neap cycles, together with the total amplitude of all the cycles. These amplitudes are listed on lines marked "neap", "year" and "total". Here, the period of the "neap" cycle is 14.765 days.

Table 3.3 *Comparison of observed levels with the predictions of river-level produced by modified versions of the existing structure functions, with seasonal and tidal harmonics included.*

Model	Term	Tower Pier	Richmond	Error Criterion	Tower Pier	Richmond
2a. adjusted for bias	a	0.0202	0.0203	rmse	0.0841	0.1099
	b	0	0	min	-0.5070	-0.7075
	c	1	1	max	0.5079	0.6501
	neap	0.0078	0.0222			
	year	0.0218	0.0475			
	total	0.1261	0.2319			
3a. adjusted for trend and bias	a	-0.0154	-0.0059	rmse	0.0817	0.1089
	b	0.1488	0.1094	min	-0.4827	-0.7077
	c	1	1	max	0.4872	0.6496
	neap	0.0078	0.0222			
	year	0.0220	0.0477			
	total	0.1238	0.2297			
4a. adjusted for trend, bias and scaling	a	0.4982	0.7105	rmse	0.0730	0.0953
	b	0.1796	0.1530	min	-0.4815	-0.6556
	c	0.8437	0.8062	max	0.3786	0.4804
	neap	0.0906	0.1322			
	year	0.0187	0.0317			
	total	0.2680	0.4202			
5a. empirical model based directly on level flow	a	0.8435	1.1786	rmse	0.0720	0.0960
	b	0.1775	0.1517	min	-0.4922	-0.6648
	s	0.9866	0.9760	max	0.3661	0.5675
	f_1	0.1212	0.1800			
	f_2	0.1801	0.5140			
	neap	0.1033	0.1350			
	year	0.0234	0.0329			
	total	0.3001	0.4377			

It is interesting to note that, while the scaling coefficients, c and s, of Models 4a and 5a are smaller than they were in the corresponding models in Table 3.1, the coefficients f_1 and f_2 are now larger than before. Thus Model 5a suggests that the difference between the levels predicted for flows of 500 and 0 cumecs, for a given time of year, should be 30 cm for Tower Pier and 69 cm for Richmond. For a given tidal-peak level at Southend, the prediction of the scaled structure function model would be changed by as much as ±27 cm at Tower Pier or ±42 cm at Richmond, depending on the size of tidal-peak being assumed. Consideration of the autocorrelations of the residuals from these models, which are shown in Figures 3.5 to 3.8, suggests that further small improvements to the predictors can be expected if further tidal harmonics are added: in particular, further long-period tidal cycles are indicated for Tower Pier. The present analysis is intended as a first indication of the likely improvement in the predictors that might be achieved by taking tidal components into account, and these small further improvements have not been pursued in this context.

3.3 Alternative types of structure functions

On comparing the models in Table 3.3 with their equivalents in Table 3.1, it is seen that in each case at least a moderate improvement is gained from incorporating seasonal and tidal effects. There are large gains in predictive performance in those cases (Models 4a and 5a) where the model is free to scale-down the effects of the tidal-peaks at Southend, either directly or via the structure function, and replace part of this variation with information from the seasonal and tidal cycles. In the case of the seasonal cycles (periods of 1 year and ½ year) it might be suggested that, when included, these components of the model could represent the effects of other fluvial contributions to the Thames which might not be adequately represented by the flows at Teddington. However, there is no similar partial explanation for the important contribution made to the predictors by the other period lengths usually associated with tidal analyses. The only plausible conclusion is that the tidal-peak levels reached on the Thames depend to a considerable extent on the tidal dynamics in the river-channel and estuary and that this cannot be fully represented by a predictor based simply on the maximum level reached at Southend. Once again, the fact that similar conclusions arise from the empirical model (Models 5 and 5a), as well as from those based on the original structure functions (for example, Models 4 and 4a), suggests that the fault lies not with the particular functions derived by hydrodynamic modelling but with the premise on which these are based.

There are several different ways in which the structure-function type of approach might be revised in an attempt to provide better predictors of tidal-peak level. One needs to recall the reasons for making use of a structure function at all, which have been discussed in Section 1.2.1 . The requirements are for a simple predictor based on a few "known" quantities, but also for one derived in an integrated way from a sound representation of the underlying reality. Thus one would not count the extended forms of the structure-function models used in this section as meeting this criterion, because it is not physically realistic simply to add tidal harmonics onto the existing structure functions. Revised structure functions might attempt to predict tidal-peak levels from the following sets of three "cause" variables.

(a) Flow, the astronomical prediction for the peak-level and the observed peak-level, where both the latter would refer to Southend. Here the astronomical and observed peak-levels could be combined to give a "surge-residual" for the peak, so the three variables might be taken to be flow, astronomical prediction and surge.

(b) Flow, the observed peak-level and the observed low-water level at Southend immediately preceding the given tide.

In both cases here, the third variable would be being used to control for the size of the change in level at Southend over the given tidal ½ cycle: this can be expected to influence the levels reached upstream via the input of momentum that this represents. In case (a), the underlying tidal cycle might be assumed to be approximately sinusoidal with an amplitude defined by the peak-level, thus defining the low-water level: however, the true "astronomical cycle" departs from this behaviour to some extent and it might be necessary to include a fourth variable to encompass this. The evaluation of structure functions based on these variables would be done in a similar way to that employed previously (Section 1.3.1). One step in this procedure is to use the supplied information to construct a detailed time-series representing the sea-levels at Southend over the course of several tidal-cycles leading up to the desired peak. In case (a), the analysis of past observed surges would contribute, as before, to defining precisely the course of the rise in sea-level up to the final peak. For case (b), some reasonable way of defining the shape of the rise would need to be found: note that the three variables would not be enough to define the size of a "surge" component and hence the need for a fourth variable may again be indicated.

Sets of variables other than those suggested above are, of course, possible. For example, the analysis to be described in Section 4 is concerned with an approach which, when posed within the context of the present discussion, would require the use of "astronomical predictions" of peak level at both Southend and the target site. It is not easy to know how to select between even the two sets listed above. There may be a preference for set (a) on the grounds that using this within a stochastic-modelling situation would require only modelling the "surge" component since the astronomical component can be treated as known whereas, for set (b), there would be a pair of random quantities associated with each tide. Ideally, the choice might be made on the grounds of the accuracy with which peak river-levels can be predicted. Such an analysis might involve a similar comparison of the observed levels in an historical data-set against the model predictions, although there is the possibility of basing the choice purely on the results from a hydrodynamic model.

The results of the models presented so far give a first indication of how much improvement in predictive performance might eventually be obtainable from a model using both observed values and astronomical predictions for sea-level, compared with one using the observed tidal-peak only. This should give an approximate assessment of how worthwhile it would be to develop structure-functions based on set (a) above. Unfortunately, because the data-set presently available does not include information about tidal minima, it is not possible to provide a similar analysis for case (b). To summarise the results in Tables 3.1 and 3.3, it seems that the root mean square error of the prediction could be reduced from 8½ cm to 7.3 cm in the case of Tower Pier and from nearly 12 cm to 9½ cm for Richmond. Note that each of these error statistics contains contributions arising from the effects of changing defence levels over time, etc., so that they are likely to be over-estimated: thus the true proportional improvement is likely to be under-estimated. Table 3.4 is concerned with an alternative way of comparing the models of Table 3.3 with those of Table 3.1: this lists the maximum differences between the predictors with and without the seasonal and tidal harmonics. These results show that rather substantial changes can be expected to the estimated river-levels for individual tides if improved types of structure functions can be developed.

Table 3.4 *Maximum difference between the predicted peak river-levels from modified structure-function models excluding and including the seasonal and tidal contributions.*

Model		Maximum difference (m)	
		Tower Pier	Richmond
2, 2a	adjusted for bias	0.091	0.179
3, 3a	adjusted for trend and bias	0.088	0.176
4, 4a	adjusted for trend, bias and scaling	0.332	0.446
5, 5a	empirical model based directly on level and flow	0.403	0.553

3.4 Notes on the analyses

(a) The sample autocorrelations of the residuals, shown in Figures 3.1 to 3.8, indicate strongly that it is not reasonable to make the standard assumption of statistical independence of the errors in the least-squares fitting procedures employed here: it is for this reason that "standard errors" for the estimated coefficients in the models have not been quoted. For the same reason, there has been no use of a formal statistical hypothesis-testing approach to deciding whether or not to allow particular coefficients to move away from their default values of 0 or 1. It would certainly have been possible to develop statistically valid procedures for accomplishing both of these tasks. However, for present purposes, it seems sufficient to rely on a common-sense type of approach to judging whether the differences between models represents a meaningfully large difference. Thus there has been emphasis both on differences in the root mean square errors of the various estimators of river-level, and on the differences between the estimators in individual cases, where both of these comparisons are made on the easily understood centimetre scale, rather than being made relative to some internally generated scale.

(b) As a supplement to the autocorrelation functions of the residuals, Figures 3.9 to 3.12, show histograms of the residuals for a selection of the models described in this chapter. These graphs have been drawn to uniform scales and therefore omit residuals with absolute values greater than 50 cm: the ranges of the residuals are quoted in Tables 3.1 and 3.3 . The histograms reveal a moderate degree of negative skewness, this effect being larger for Richmond than for Tower Pier. Since the residuals are defined as "observed minus predicted", this indicates a tendency for larger extremes of over-estimation then under-estimation. As would be expected, a visual comparison of the widths of these histograms between models agrees with the conclusions already drawn from the root mean square errors in Tables 3.1 and 3.3 .

(c) For comparison with the plots of observed river-levels given earlier, Figures 3.13 to 3.16 show plots of short sections of these data, together with the estimated river-levels obtained by using the original structure functions together with an adjustment for trend. Thus these plots correspond to Model 3 in Table 3.1. Similar plots are given later for some of the models developed in Section 4.

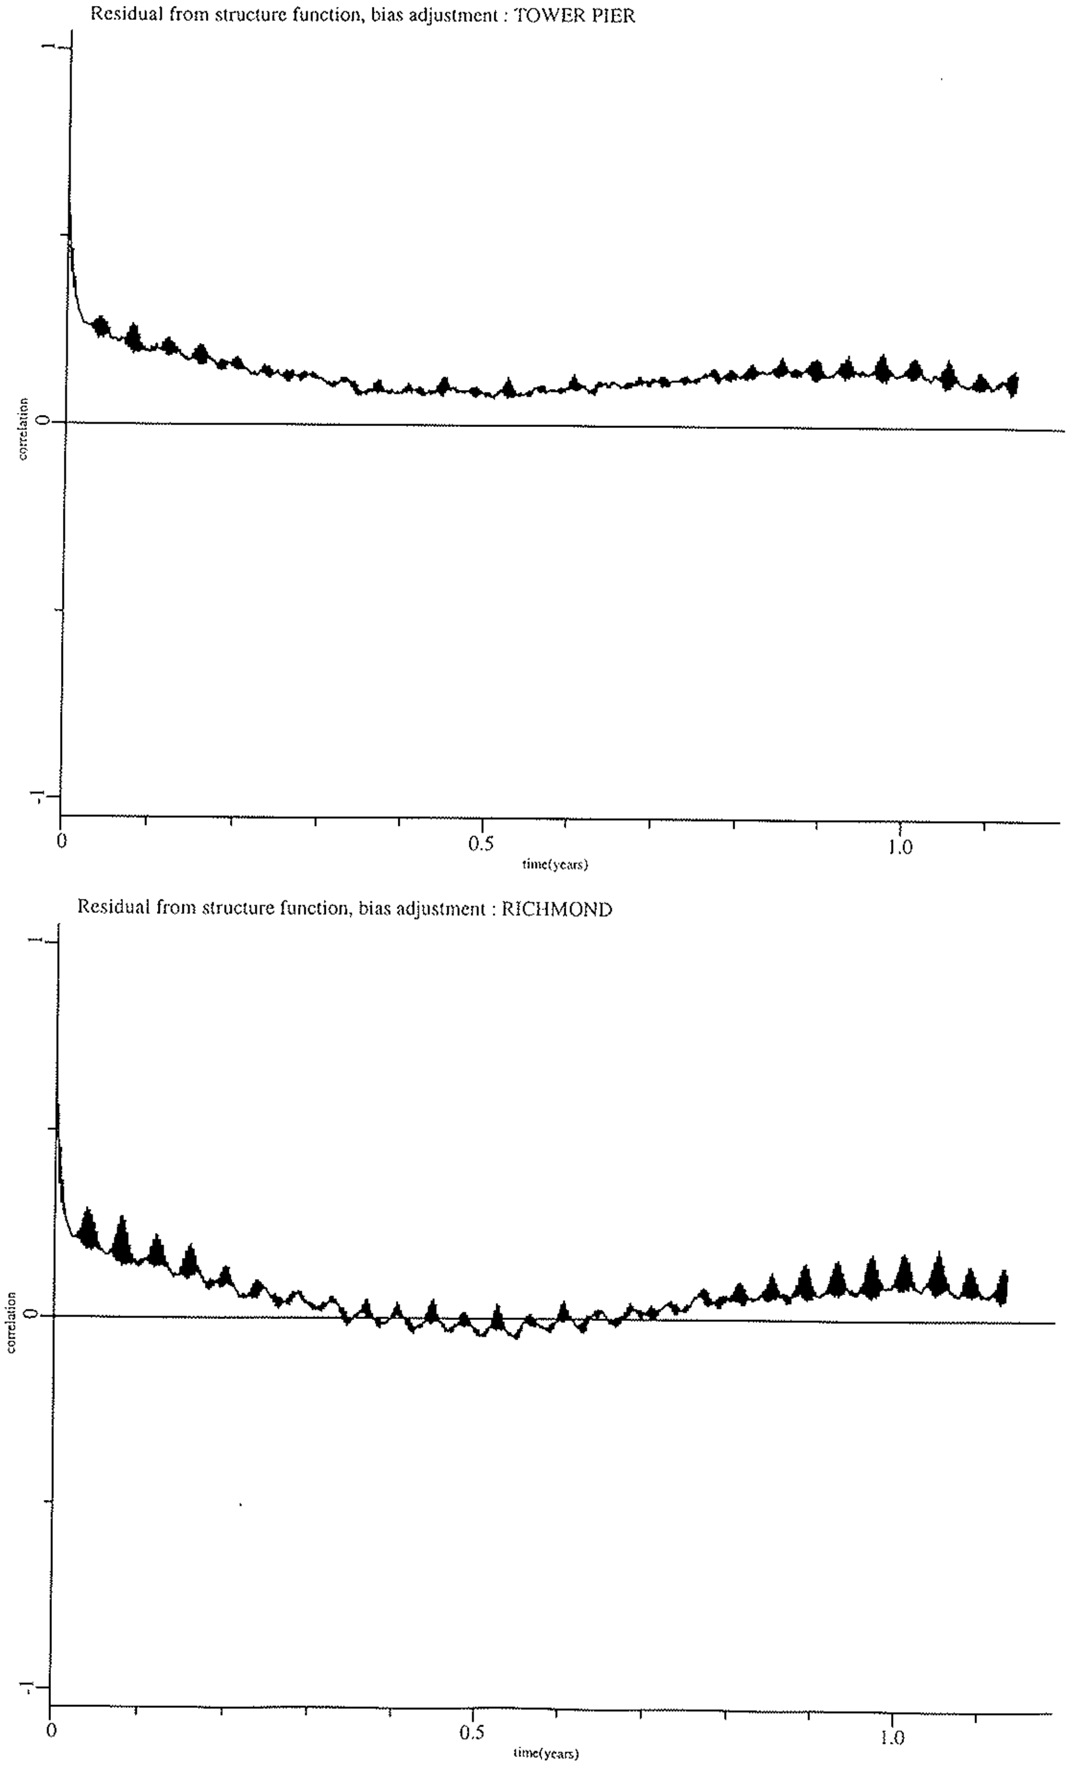

Figure 3.1 Autocorrelation functions of residuals from Model 2.

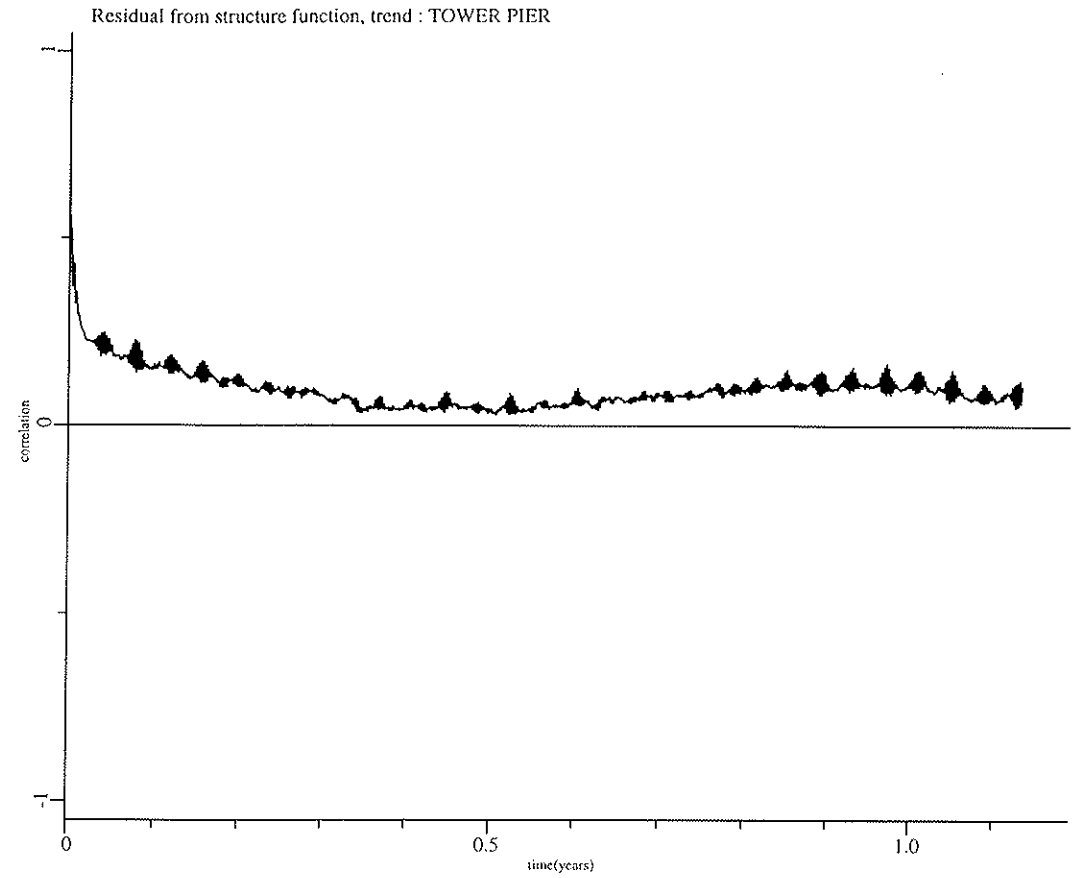

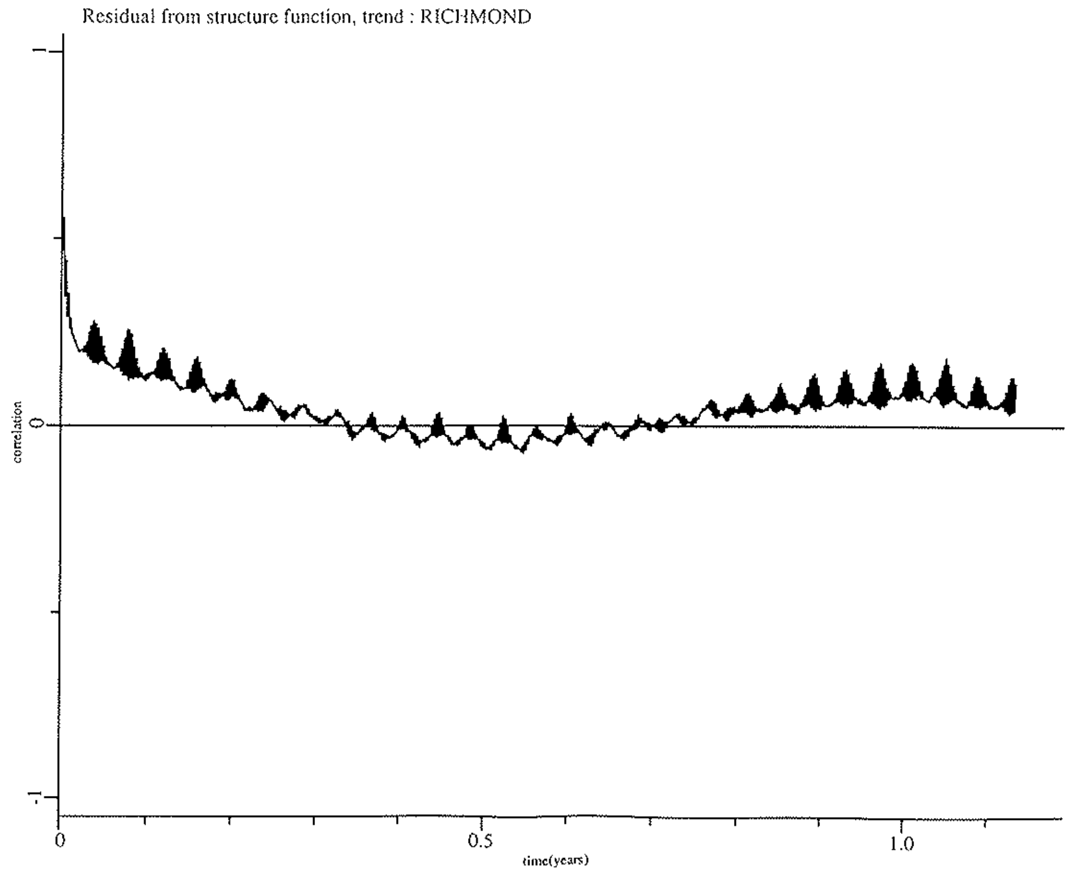

Figure 3.2 Autocorrelation functions of residuals from Model 3.

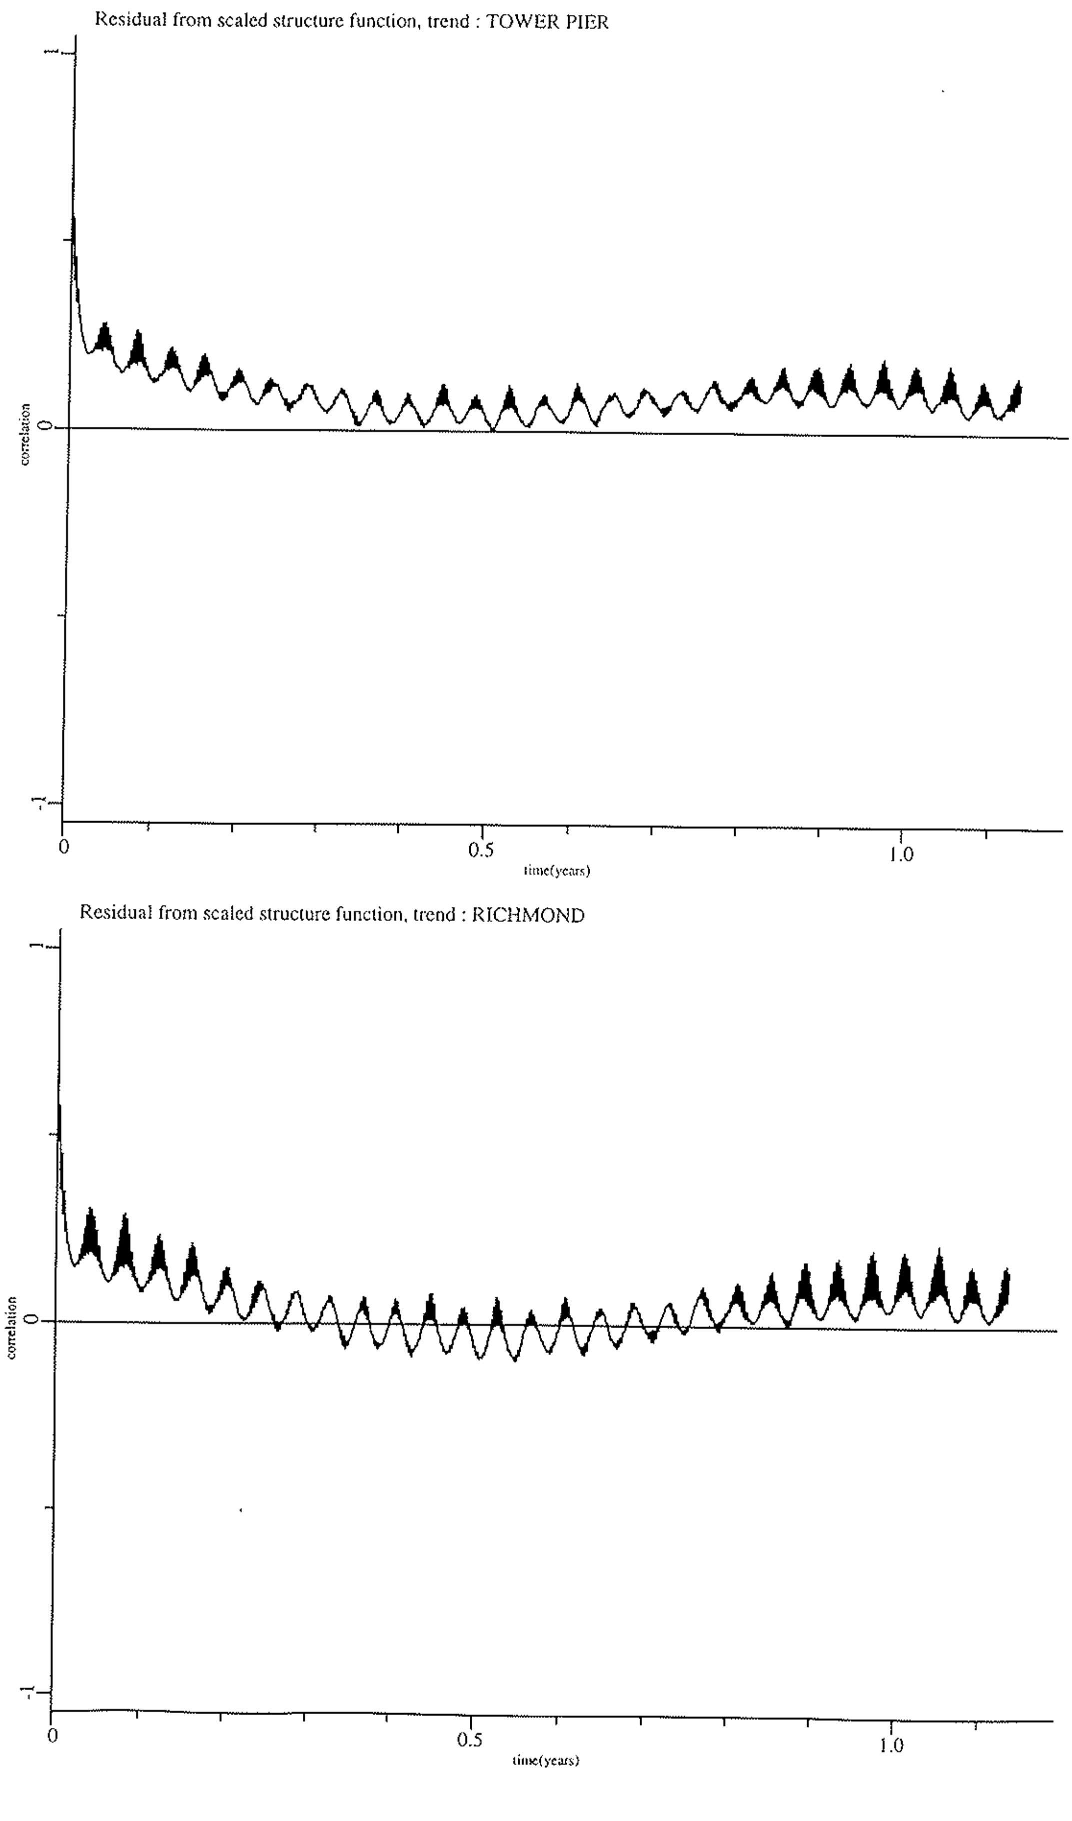

Figure 3.3 Autocorrelation functions of residuals from Model 4.

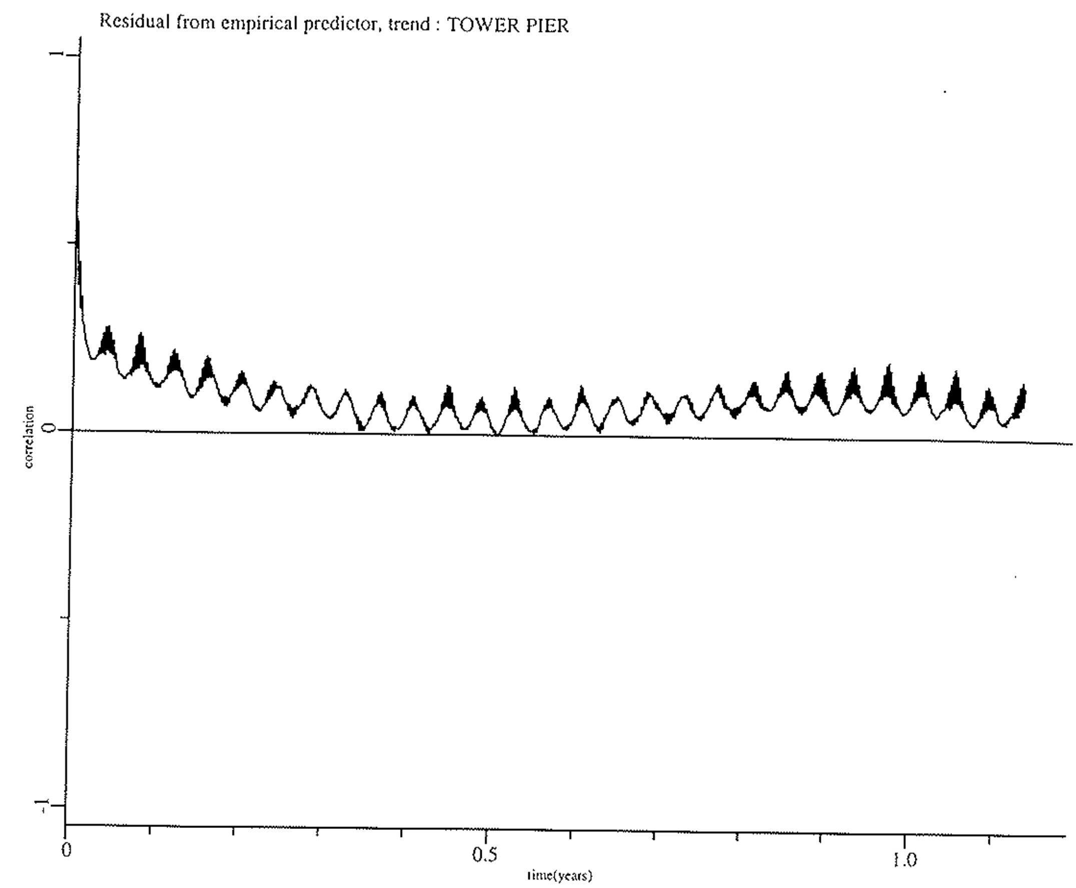

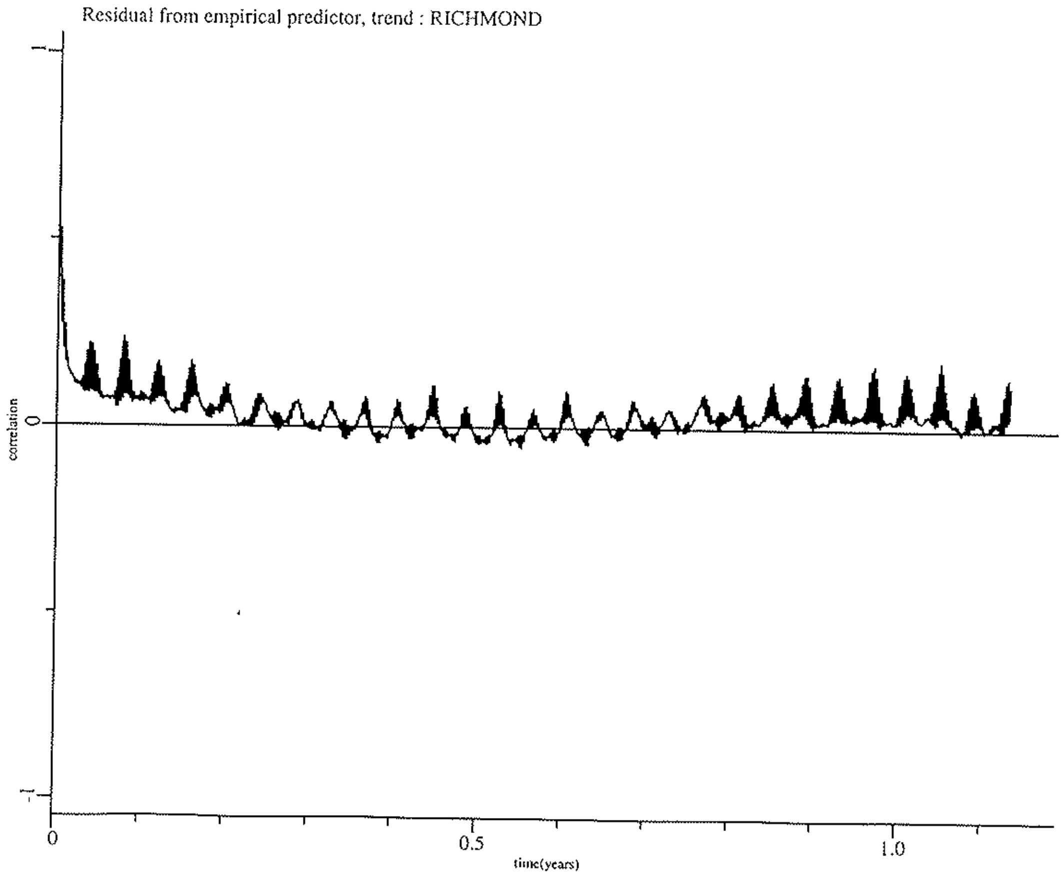

Figure 3.4 Autocorrelation functions of residuals from Model 5.

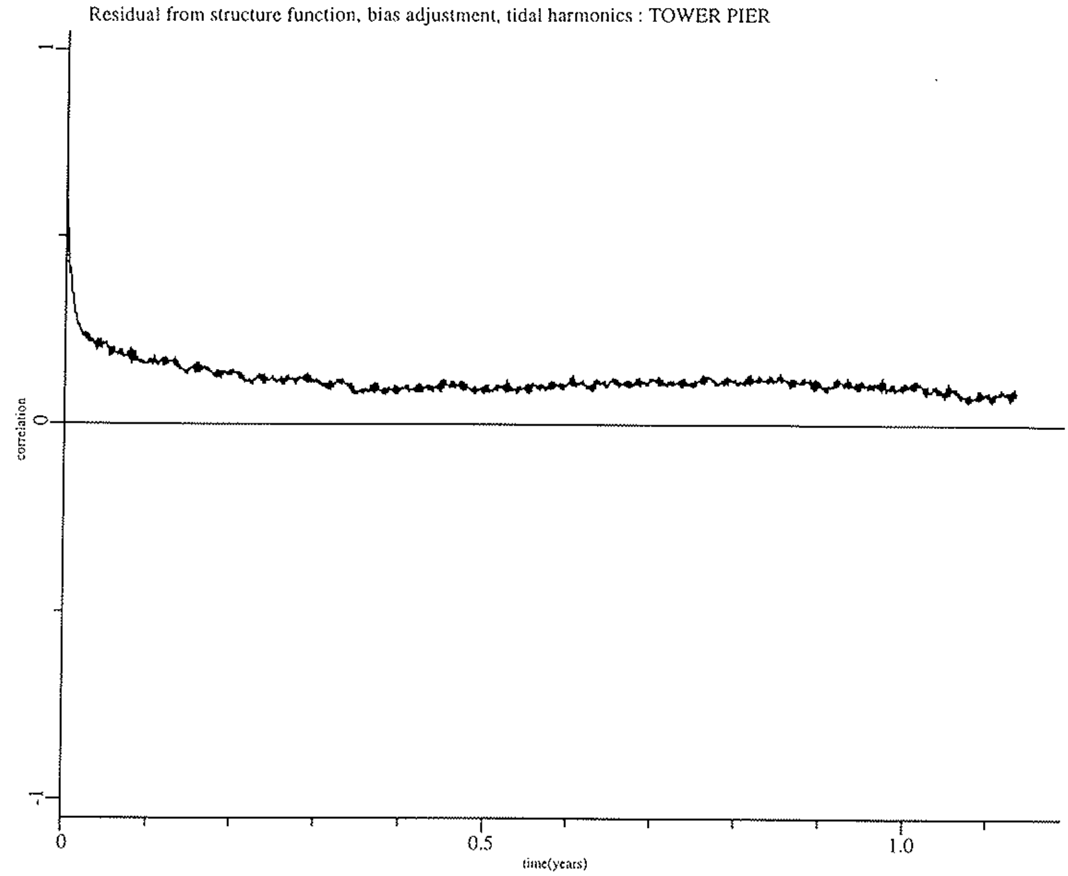

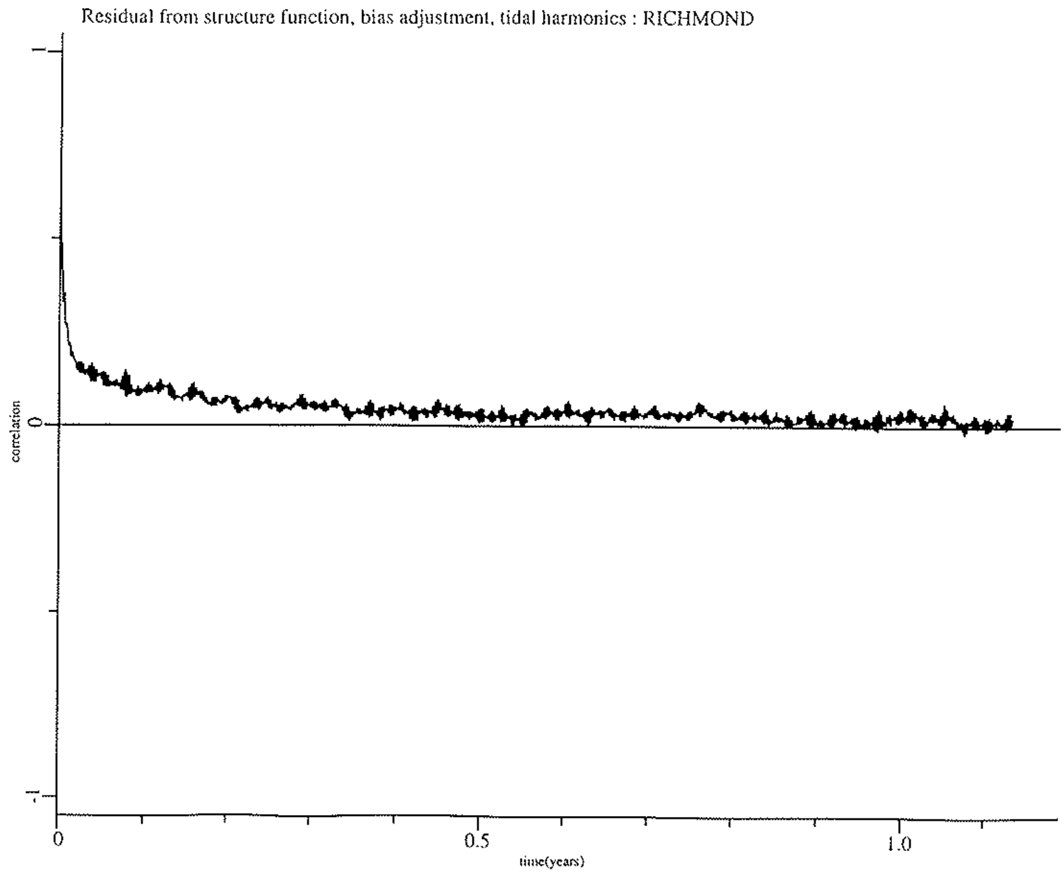

Figure 3.5 Autocorrelation functions of residuals from Model 2a.

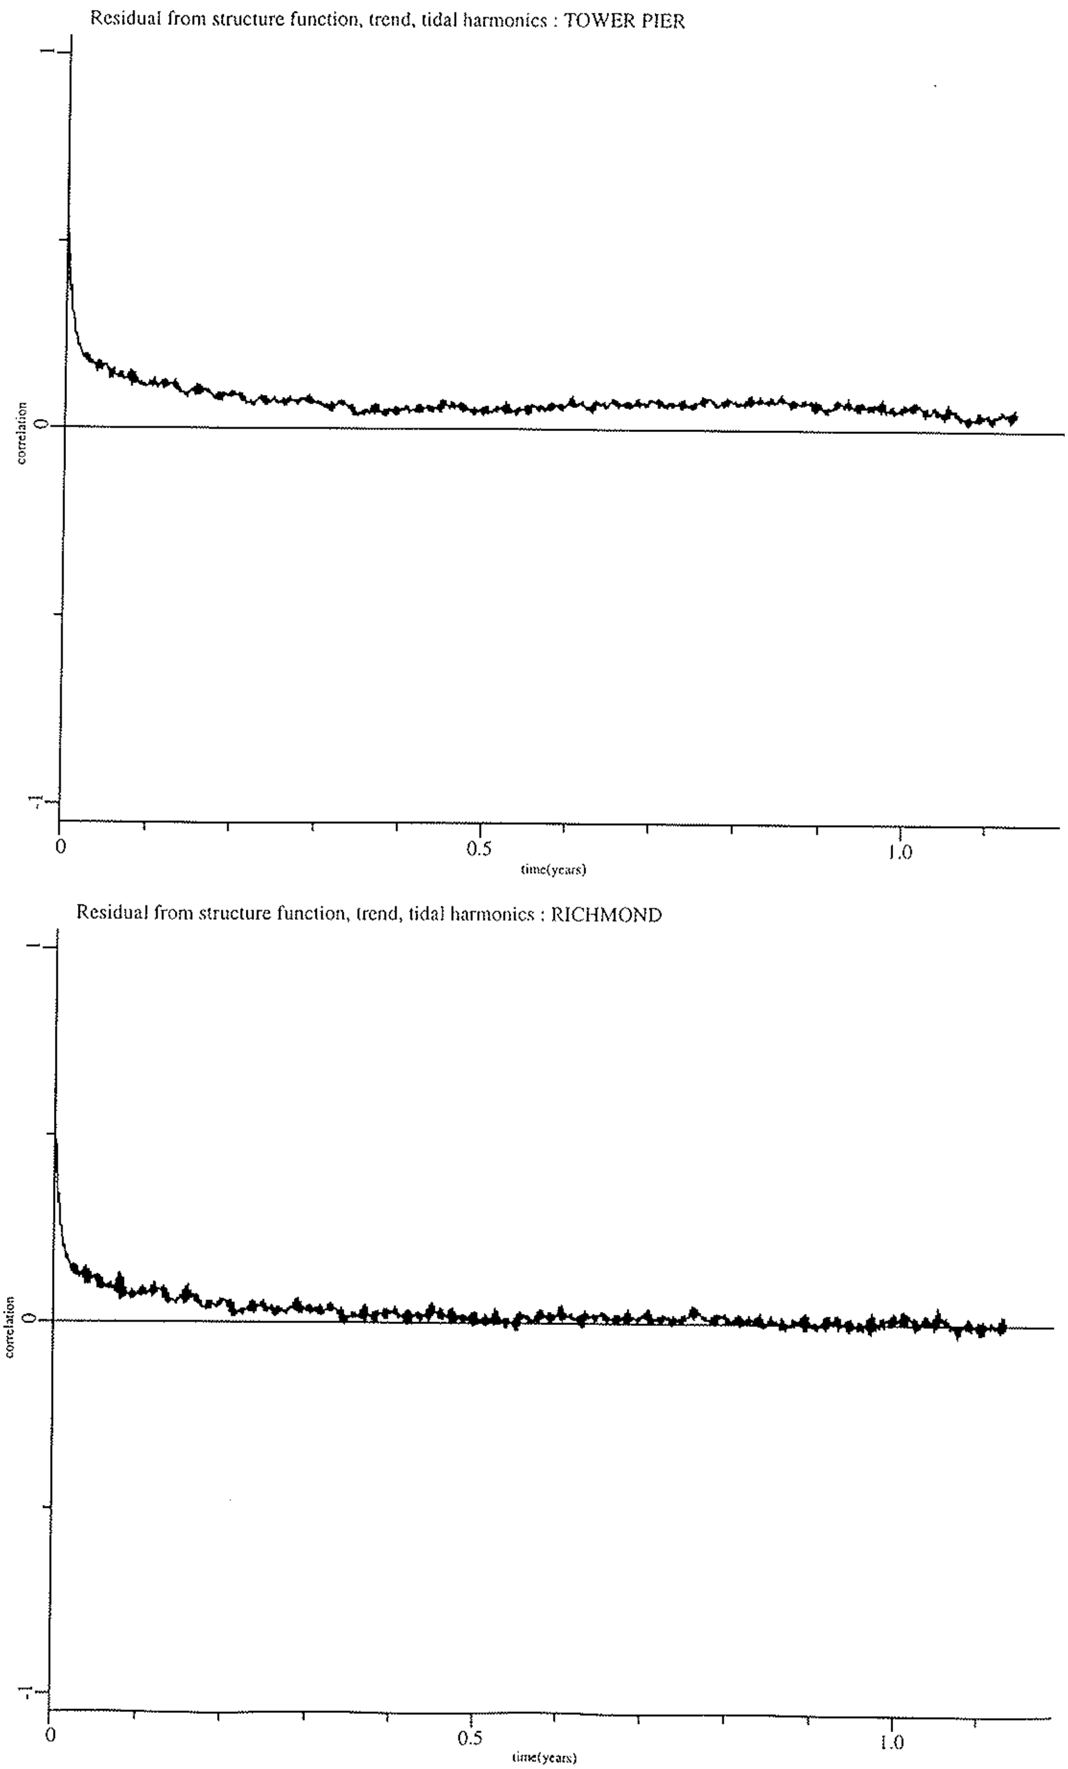

Figure 3.6 Autocorrelation functions of residuals from Model 3a.

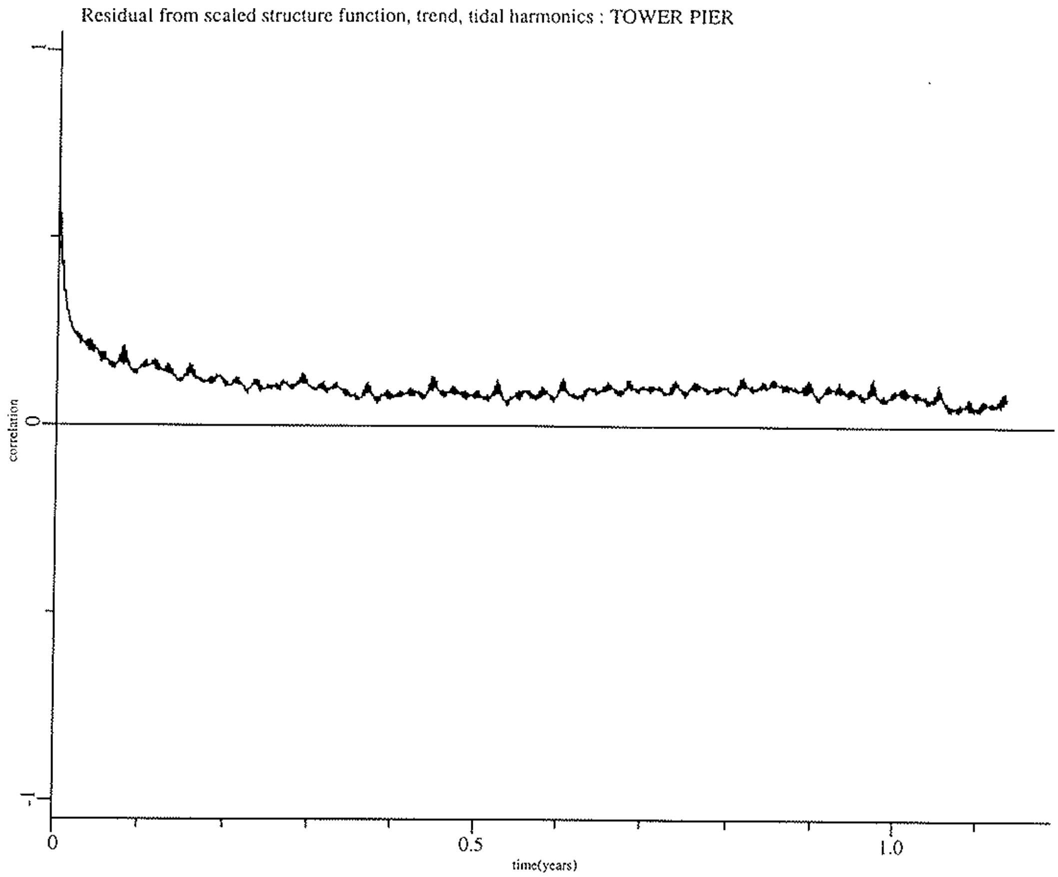

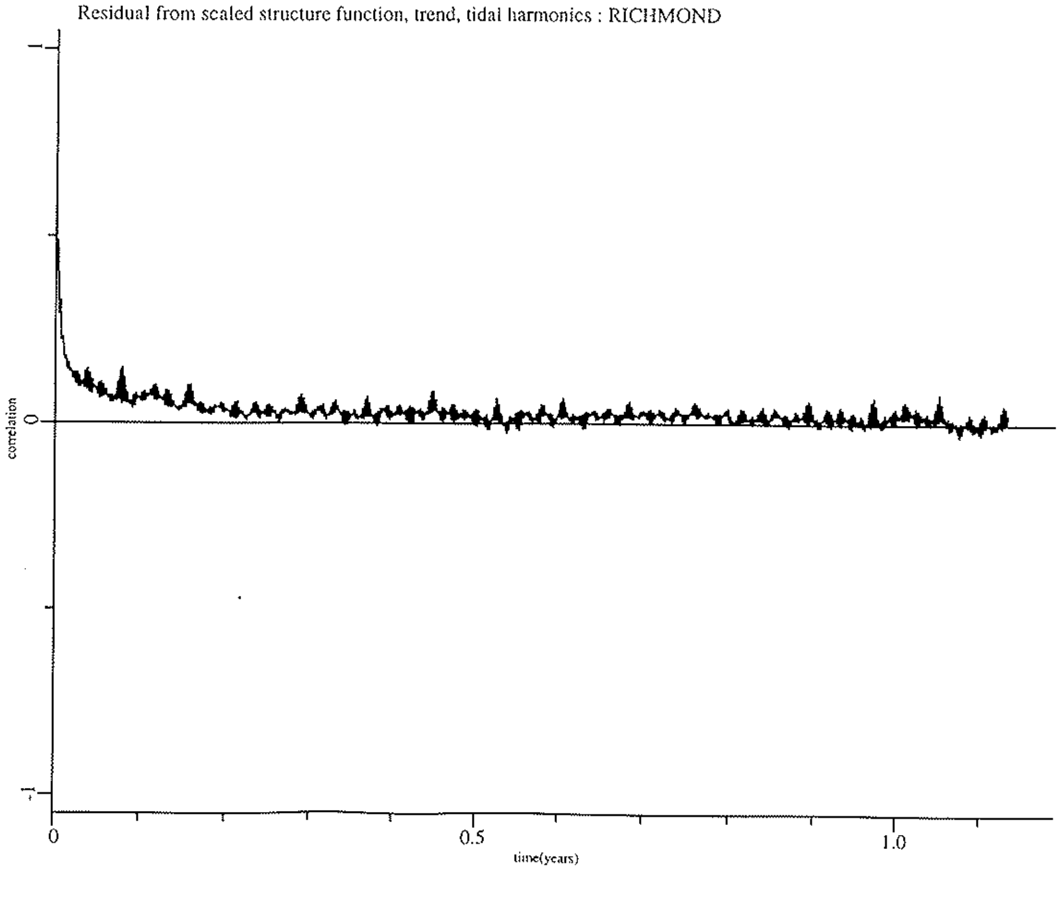

Figure 3.7 Autocorrelation functions of residuals from Model 4a.

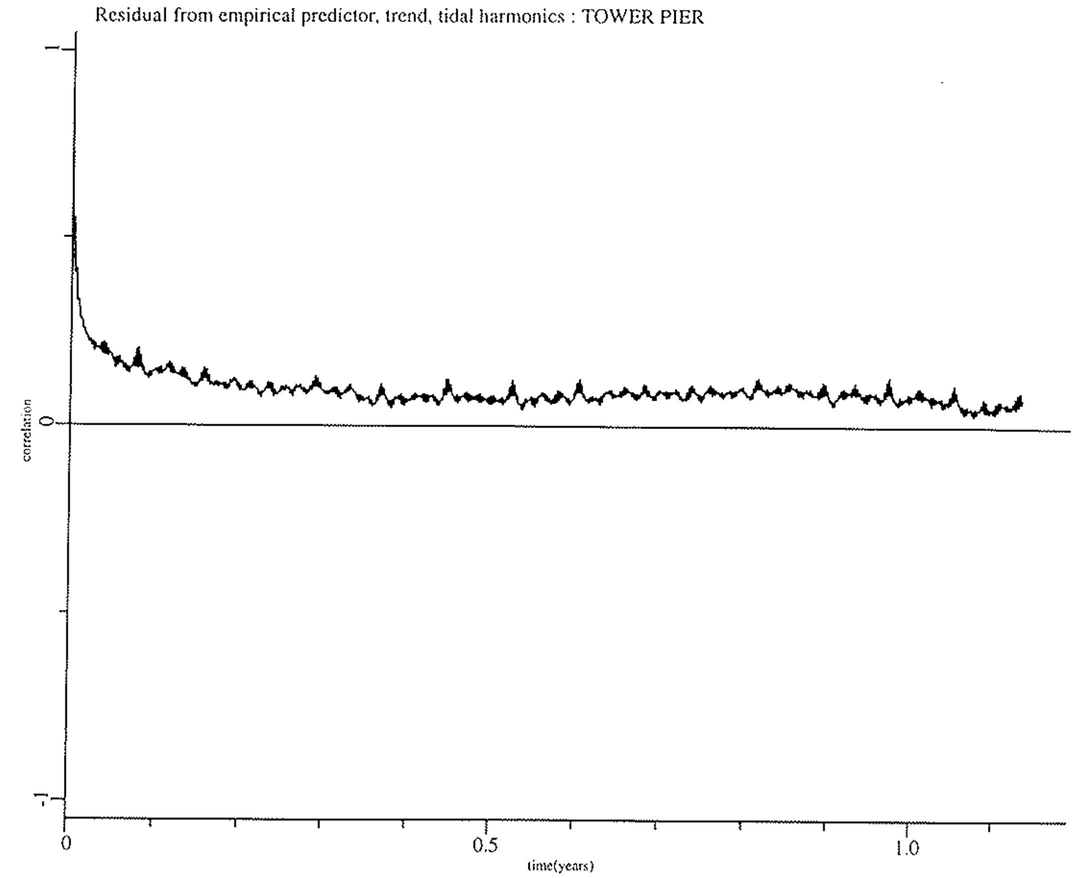

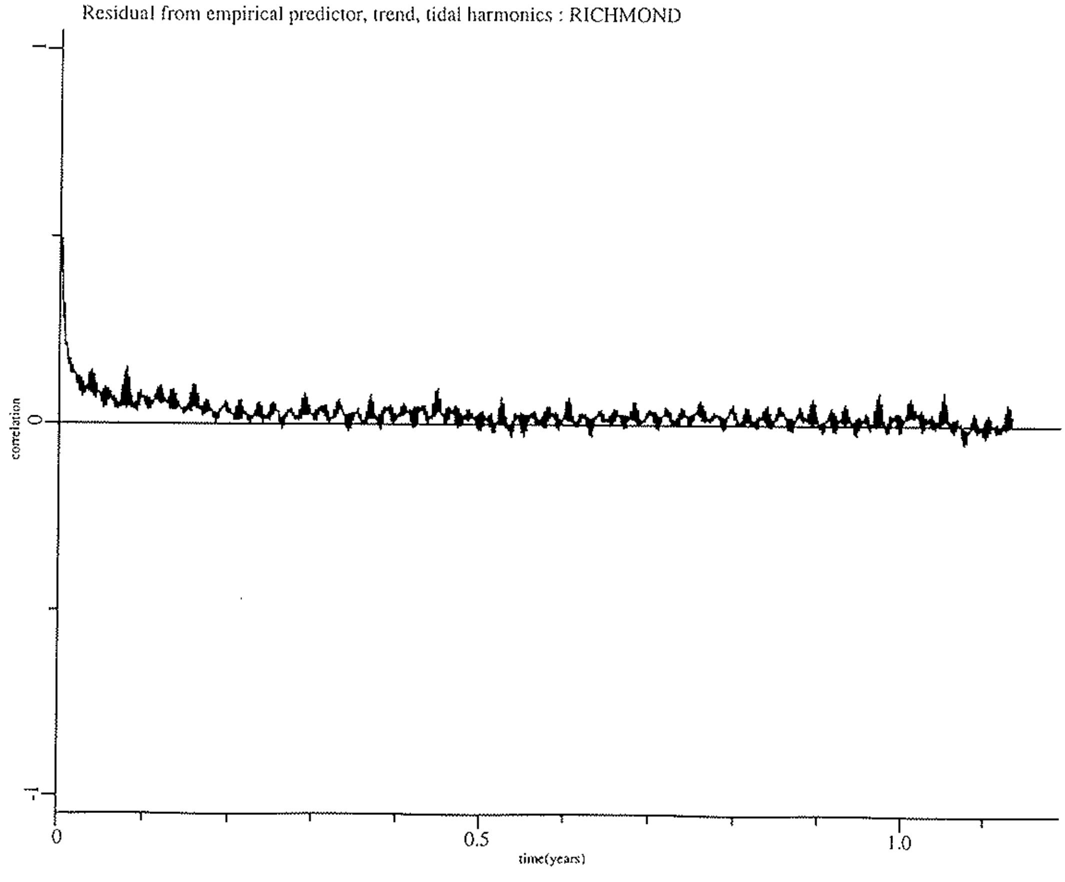

Figure 3.8 Autocorrelation functions of residuals from Model 5a.

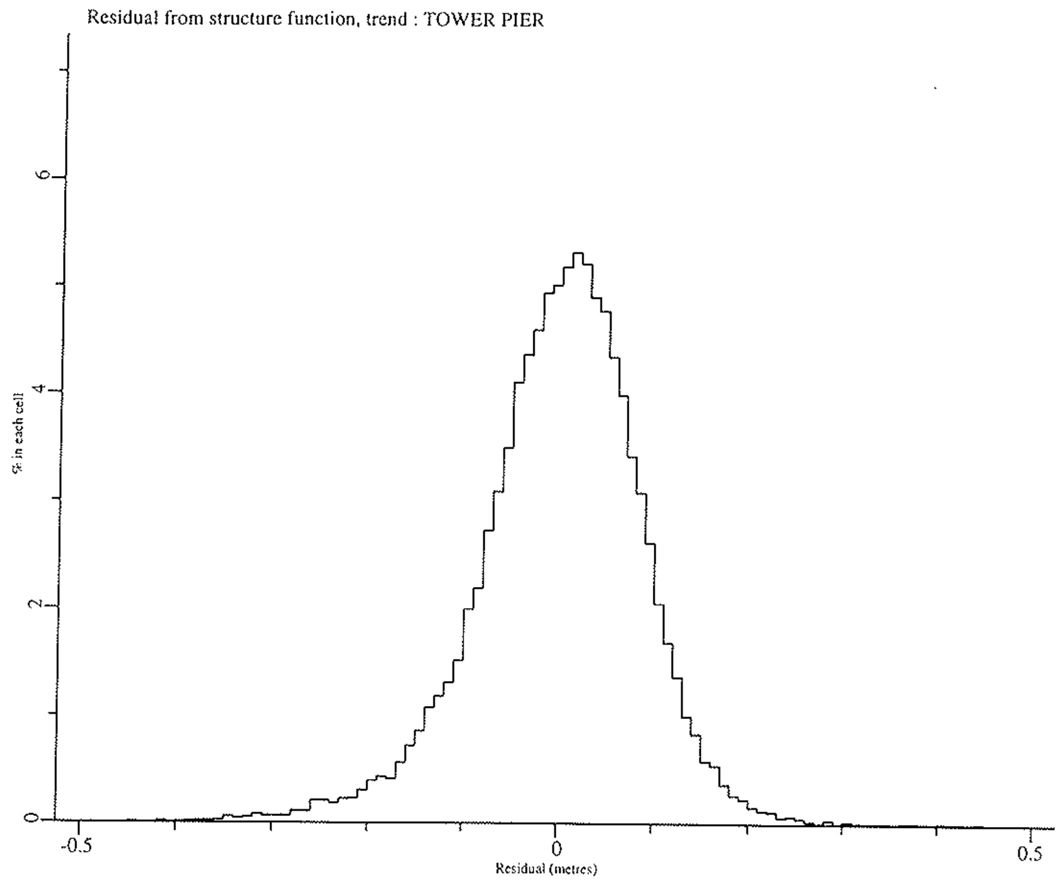

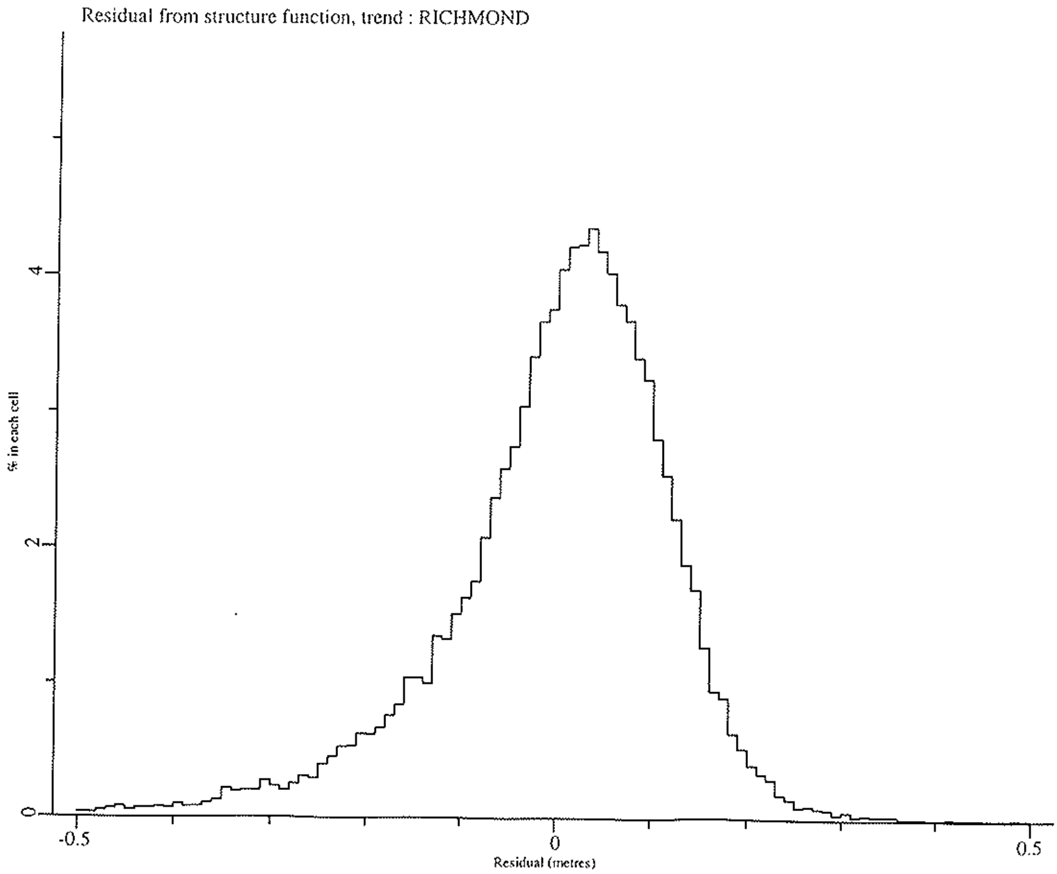

Figure 3.9 Histograms of the residuals from Model 3.

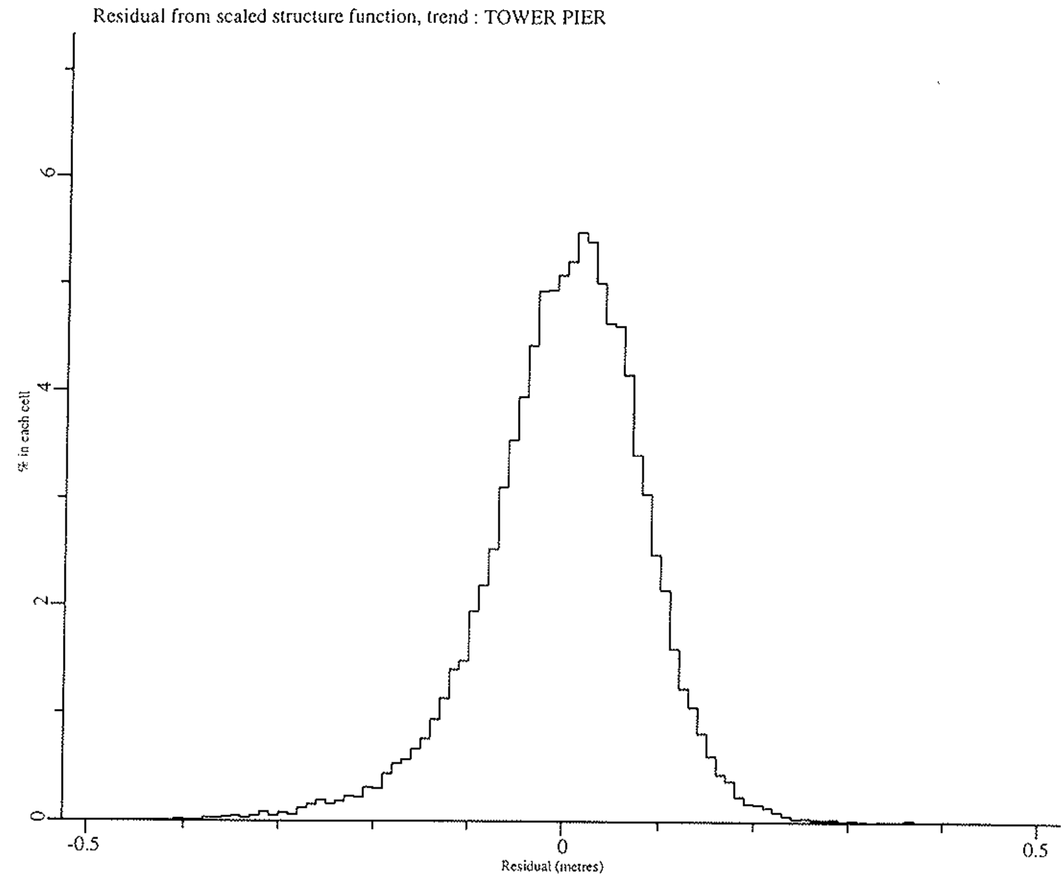

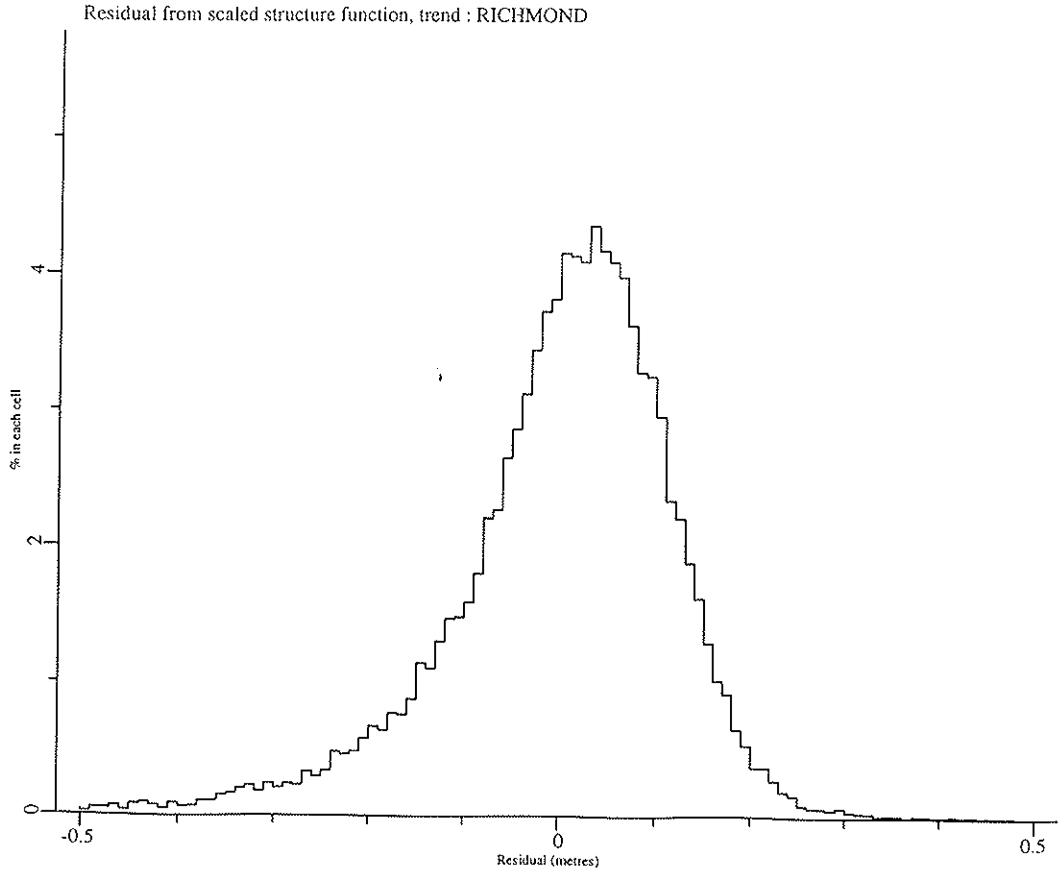

Figure 3.10 Histograms of the residuals from Model 4.

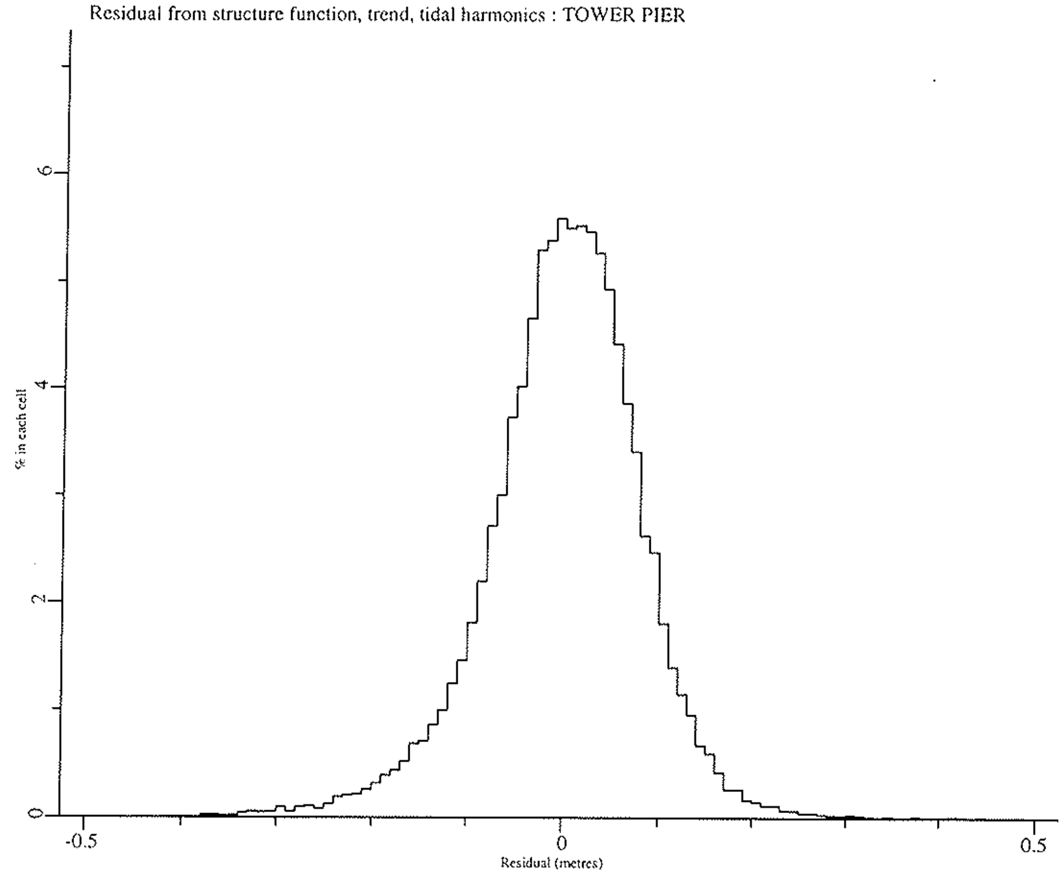

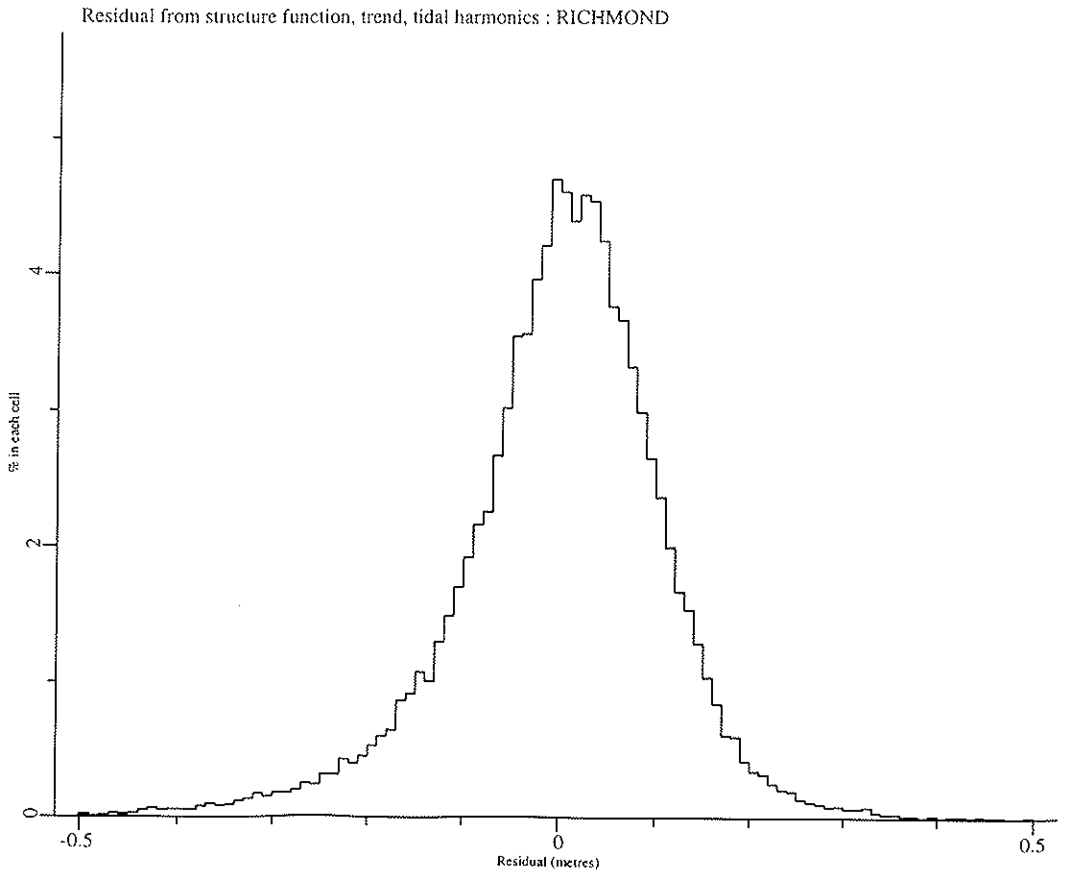

Figure 3.11 Histograms of the residuals from Model 3a.

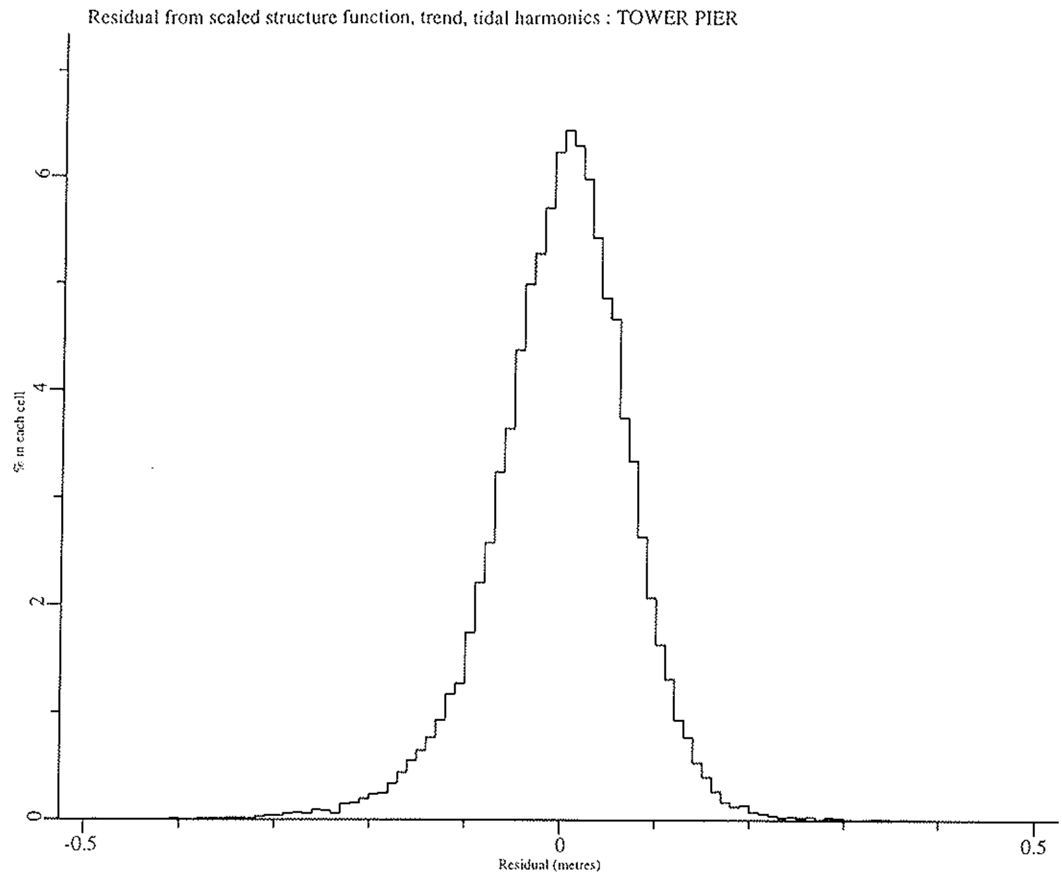

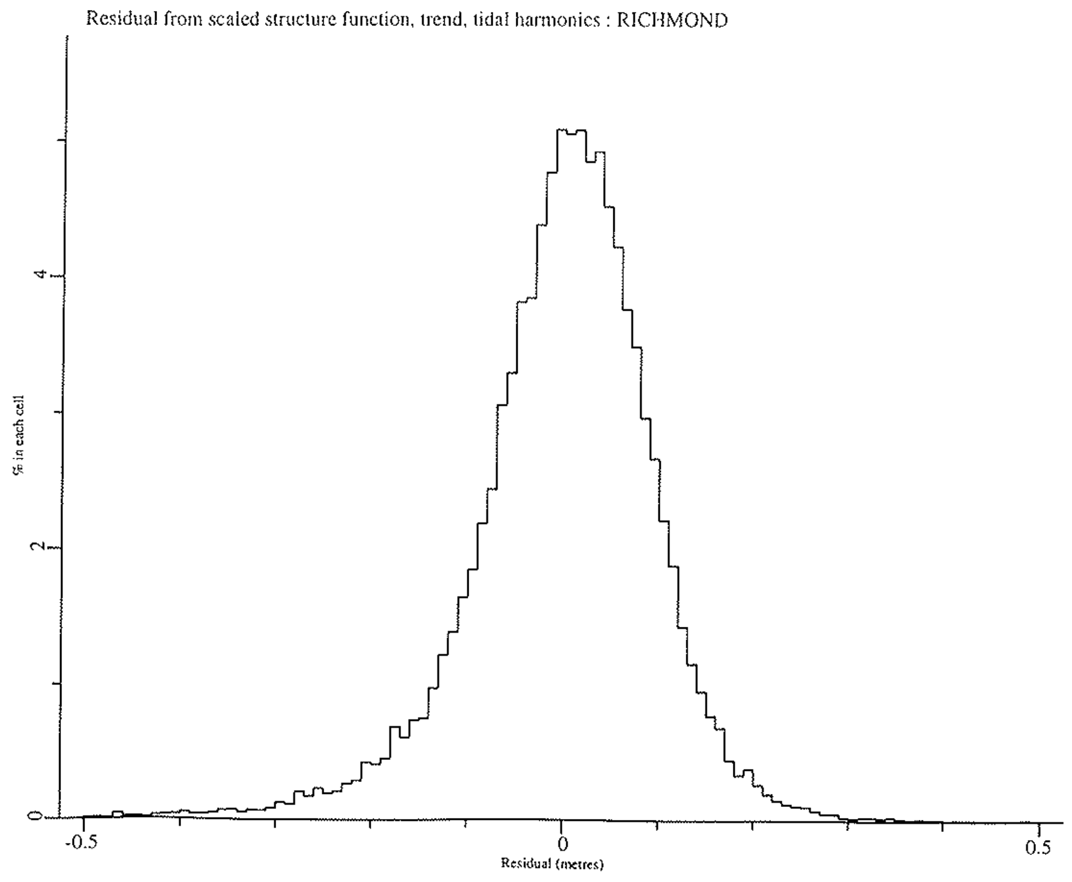

Figure 3.12 Histograms of the residuals from Model 4a.

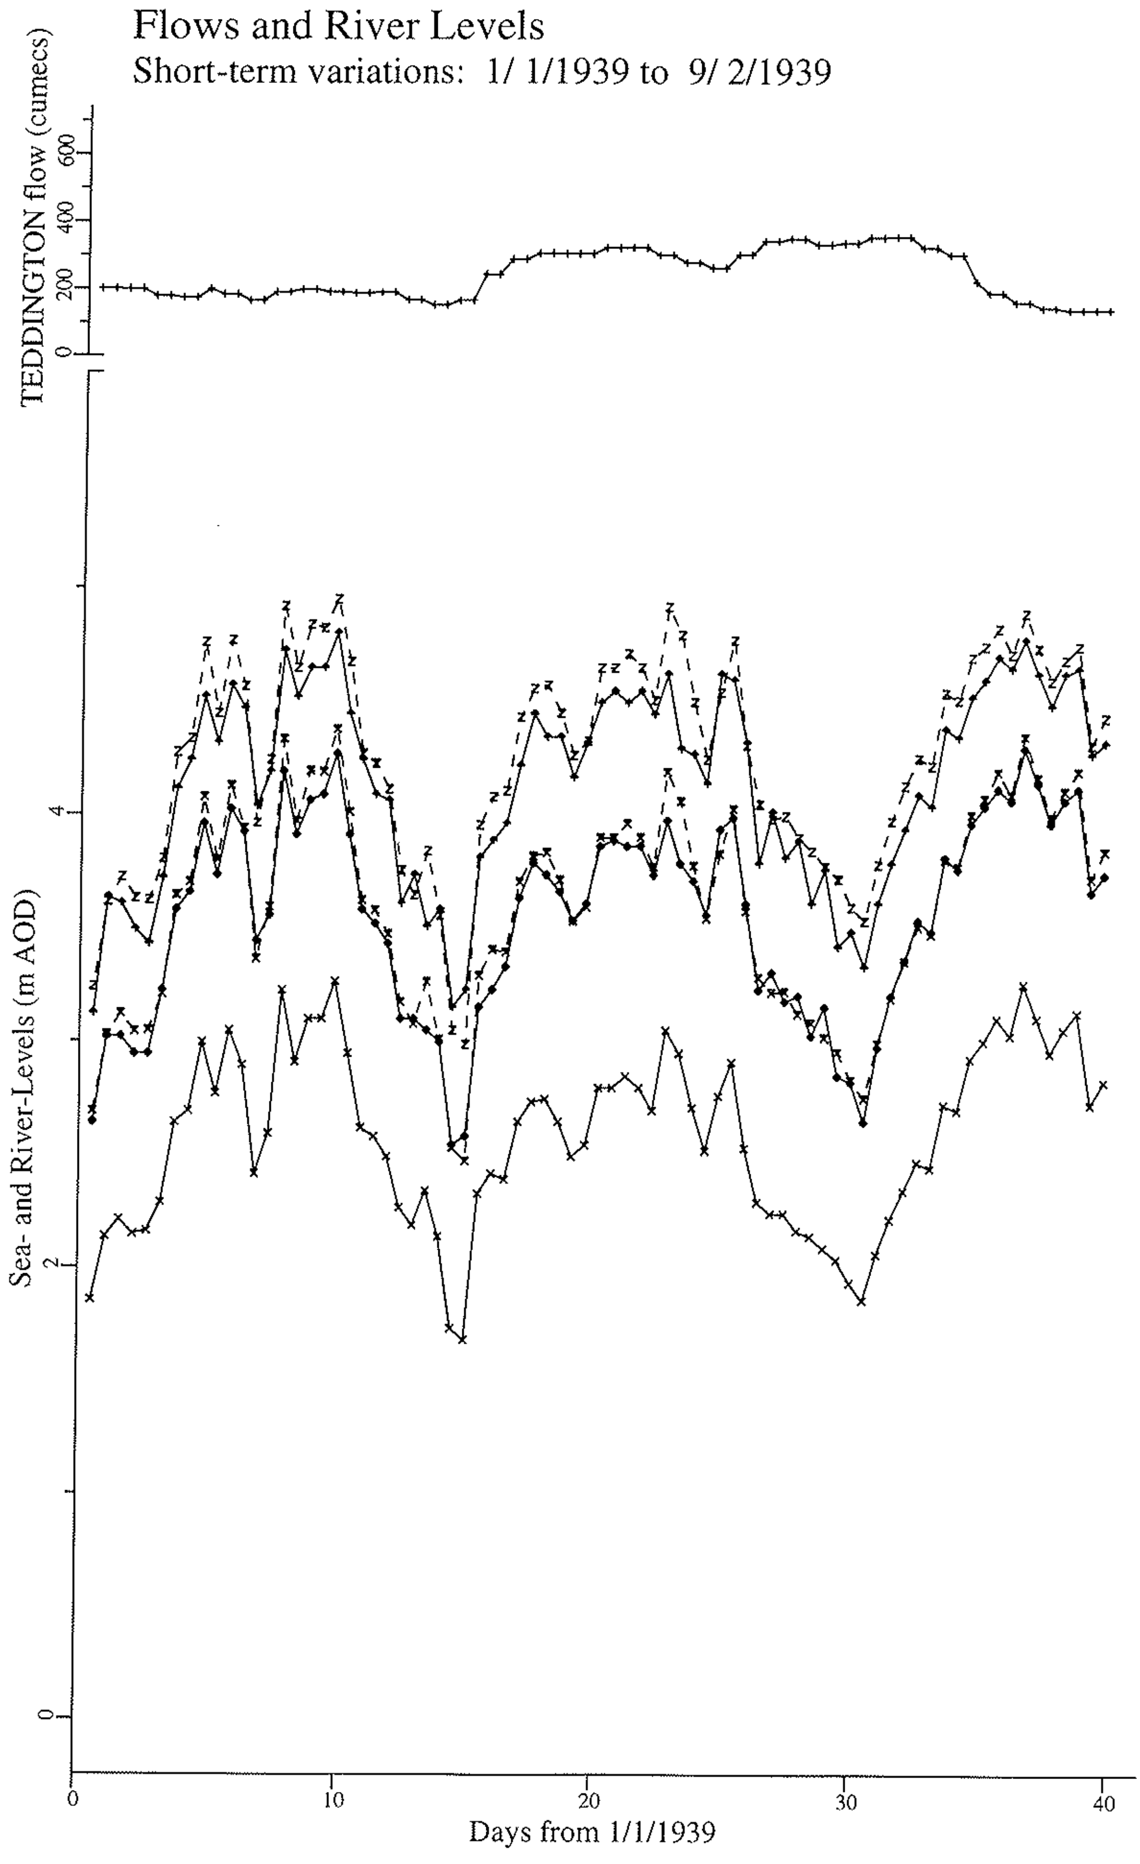

Figure 3.13 Example of the estimated river-levels from the trend-adjusted original structure functions (Model 3).

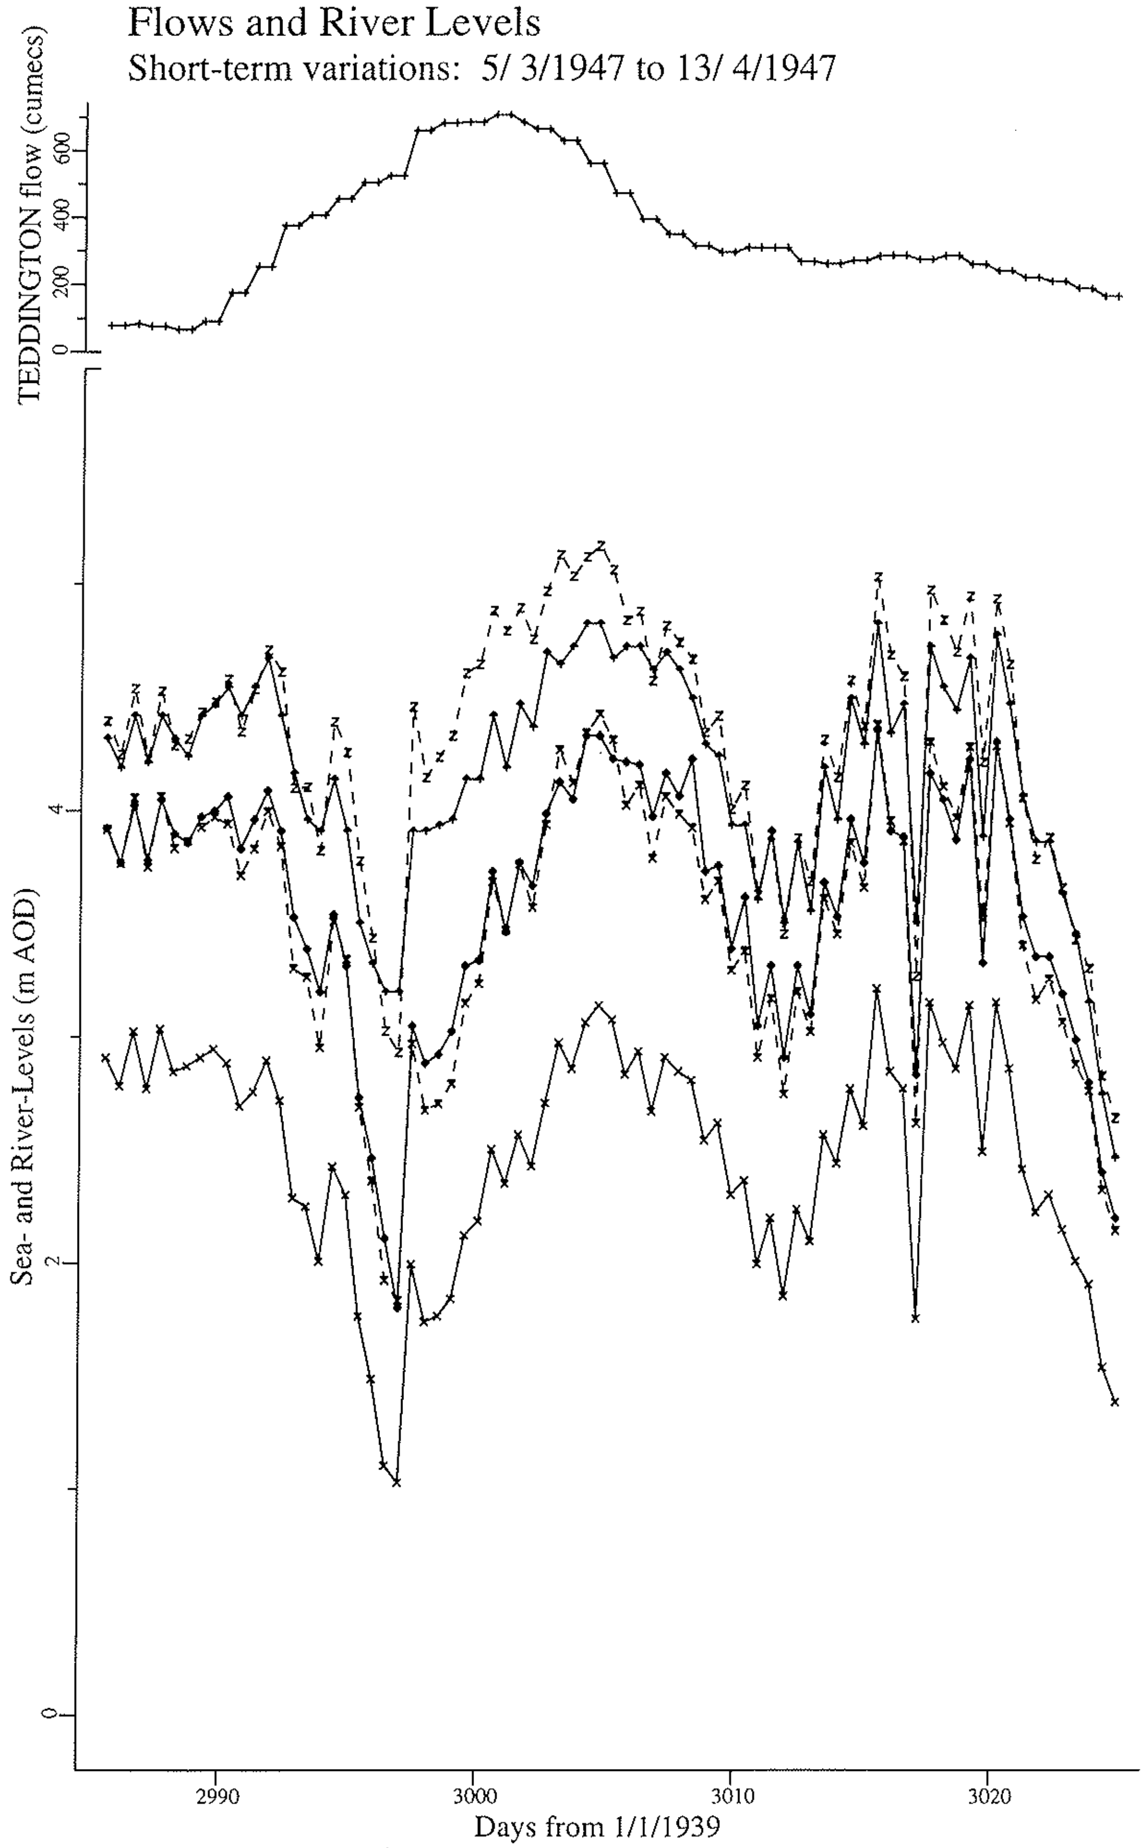

Figure 3.14 Example of the estimated river-levels from the trend-adjusted original structure functions (Model 3).

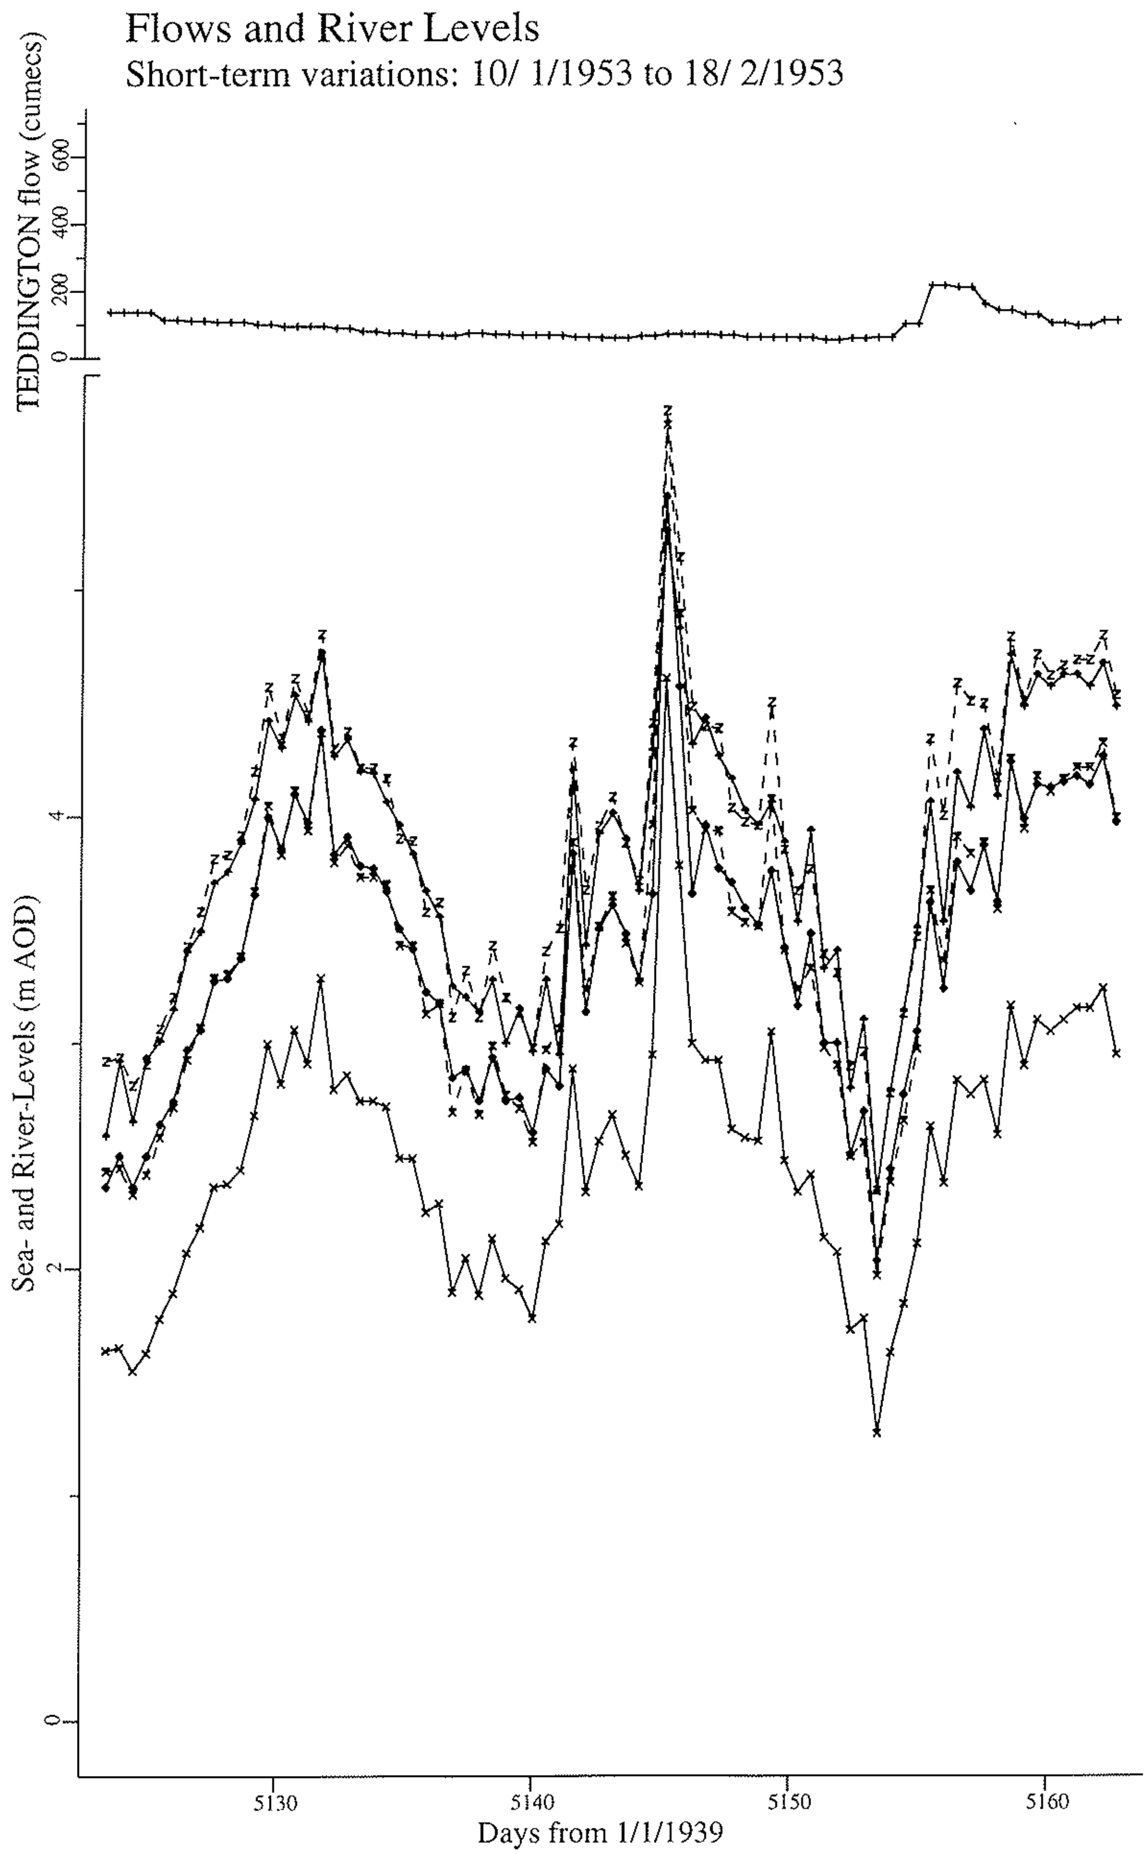

Figure 3.15 Example of the estimated river-levels from the trend-adjusted original structure functions (Model 3).

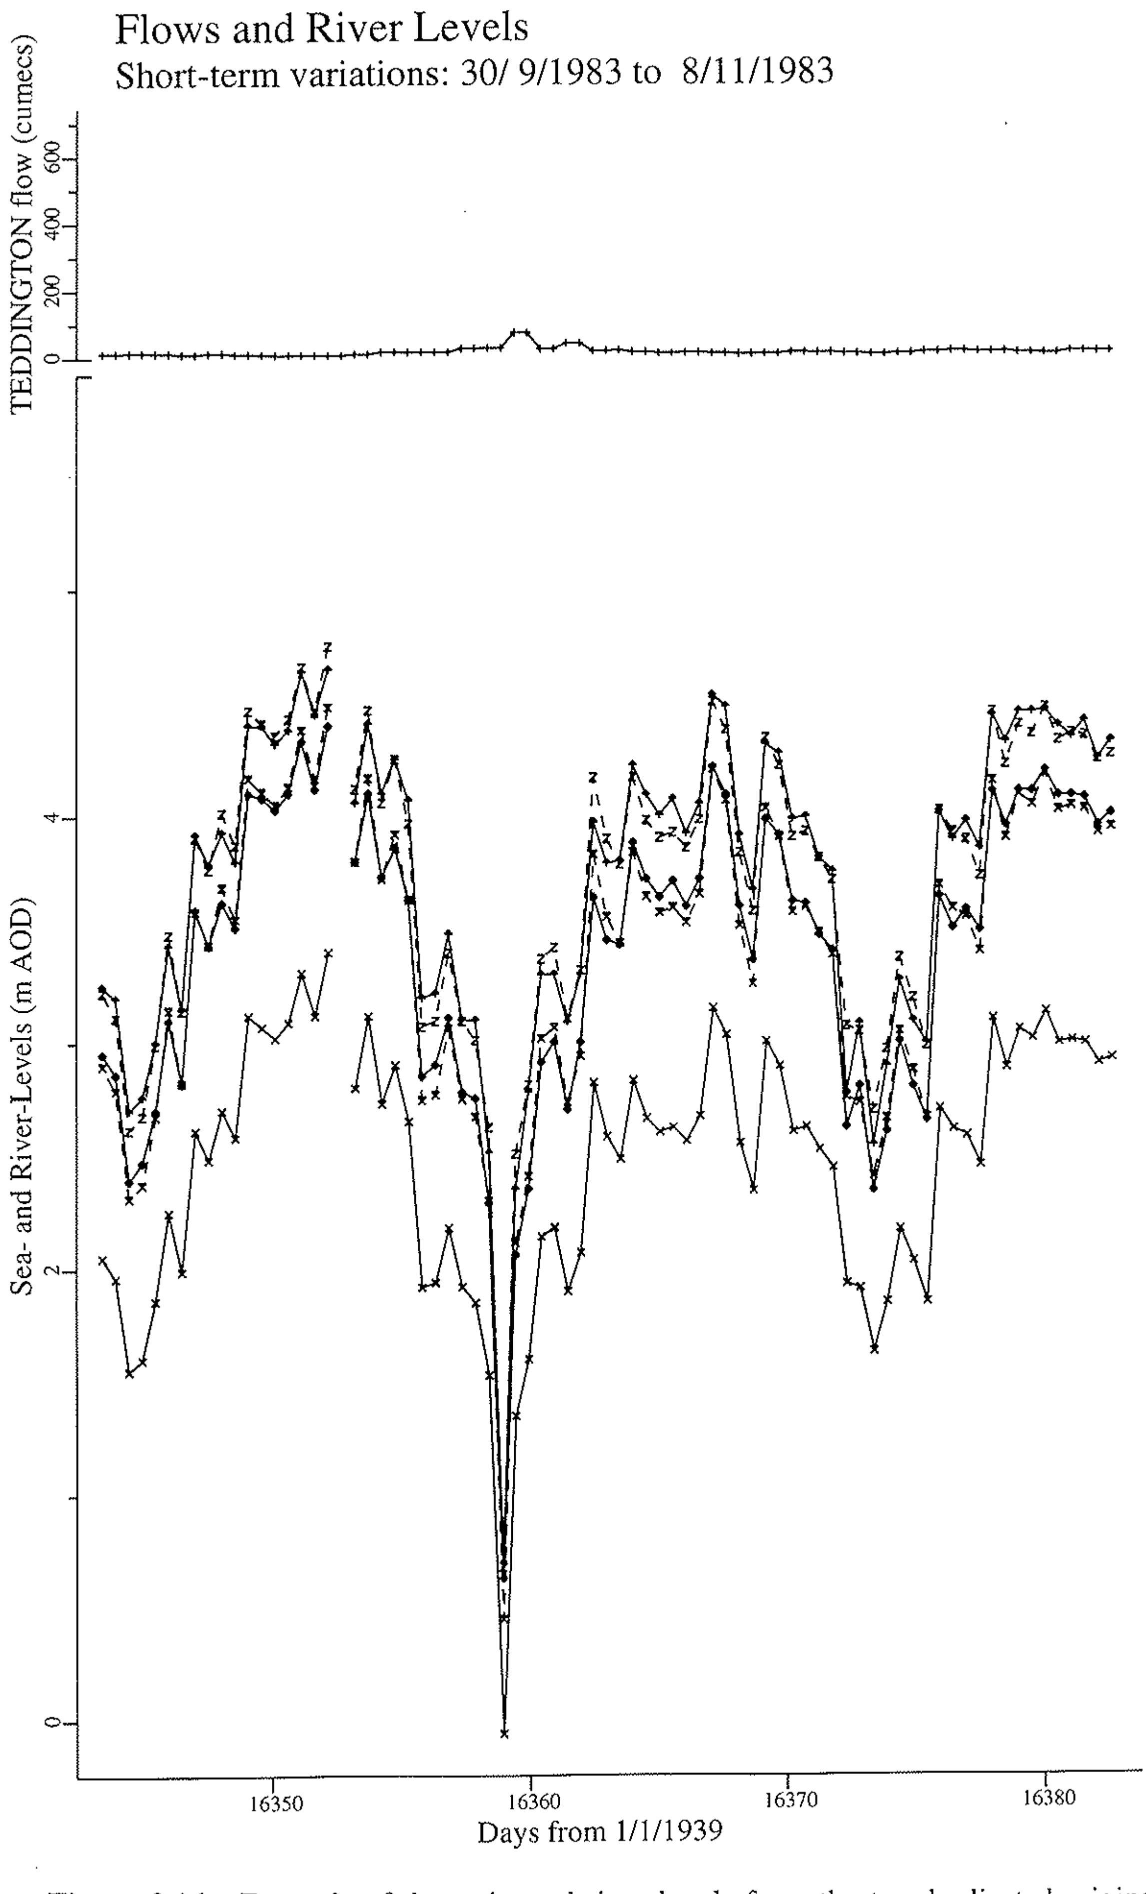

Figure 3.16 Example of the estimated river-levels from the trend-adjusted original structure functions (Model 3).

4. Estimators based on Tidal Residuals

4.1 Introduction

The analysis reported in Section 3 suggests the strong need, in any procedure for estimating tidal-peak river-levels, to take into account in some direct fashion the underlying astronomical tides. While the models employed in Chapter 3 do enable this to be done, the structures of the models could be considered to be unappealing. Thus the collection of seasonal and astronomical terms which appear in the model relate to a correction to a scaled version of Southend levels, rather than to the "astronomical" predictions of sea- or river-level at the individual sites, which are rather better understood.

For the present project, a pre-existing set of astronomical predictions of sea-level and river-level was not available and thus it has been necessary to develop models for these as part of the overall set of models being fitted. Note that this has the advantage that there is then no need to make a choice between using the maximum astronomical prediction for a given tide, or the astronomical prediction at the time of the observed sea-level maximum (if the time had been available). In this particular context one would naturally be thinking of an astronomical prediction covering the whole of the tidal cycle, not just the peaks. Instead, the "astronomical prediction" created here refers directly to a prediction for the maximum observed sea-level, when this is ascribed an equal time-step reference time as described in Section 2.2.

The following notation is introduced in order to describe the estimators of river-level used in this section. On the basis of the models described in Section 3, it seems reasonable to structure models for the sea- and river-level series so that they contain trend, harmonic and flow-related components which contribute in an additive way as follows.

$$\begin{aligned} S_{obs} &= S_{trend} + S_{harm} + S_{flow} + S_{surge}, \\ T_{obs} &= T_{trend} + T_{harm} + T_{flow} + T_{surge}, \\ R_{obs} &= R_{trend} + R_{harm} + R_{flow} + R_{surge}. \end{aligned} \qquad (4.1)$$

Here, as before, S, T and R refer to Southend, Tower Pier and Richmond. The trend, harmonic and flow-related components of these models each entail a number of parameters which occur linearly and which are estimated by least squares. The "surge" components correspond directly to the residuals of the models. However, it seems useful to make an identification here with the concept of surge, since this leads on to considering new estimators of river-level at Tower Pier and Richmond of the following form:

$$\begin{aligned} \text{est (T)} &= T_{trend} + T_{harm} + T_{flow} + m\, S_{surge}, \\ \text{est (R)} &= R_{trend} + R_{harm} + R_{flow} + m\, S_{surge}. \end{aligned} \qquad (4.2)$$

Here m denotes an amplification or reduction factor applied to the observed surge component at Southend in forming the estimated levels for Tower Pier and or for Richmond, with different factors being applied in the two cases. Further, S_{surge}

represents the residual from the model for Southend in the form of Equation (4.1): the other estimated surge components are not used further.

The following subsections describe the fitting of models of the form in Equation (4.1) in Sections 4.2 and 4.3, with models of the form in Equation (4.2) being used in Section 4.4.

4.2 Direct models for sea- and river-levels

The components of the models of the form in Equation (4.1) have been assumed to be of types similar to those already used in Section 3. Thus, taking Richmond as an example, the j'th values in the series representing trend and the contribution from flow have the forms:

$$R_{trend}(j) = a + b\,\{t_j/100\} \qquad (4.3)$$

$$R_{flow}(j) = f_1\,\{Q_j/500\} + f_2\,\{Q_j/500\}^{1/2}, \qquad (4.4)$$

where t_j and Q_j are as described in Section 3.1.

As in Section 3.2, the harmonic components of the model consist of a number of period lengths, although rather more components have been included here. The period lengths included in the models here are listed in Table 4.1. These particular periods have been chosen on the basis of a spectral analysis of the residuals of models containing fewer harmonic components: components were included if they seemed warranted from the analysis for any individual site, although usually the same period lengths were indicated for all three sites. The actual period lengths used to correspond to particular spectral peaks were chosen to be the closest simple combination of the five basic astronomical frequencies, taking into account the aliasing of frequencies associated with the ½ lunar day time-step. There was one major spectral peak, shared by all three sites, for which such a frequency could not be found: this corresponded to a period length of approximately 209 days. Nonetheless, this period has been included in the model since, as discussed in Section 2.2, there is the possibility that non-astronomical cycles might arise because the tidal-peaks, which actually do not occur at an equally spaced time-step, are being treated as if such were the case. It will be seen from Table 4.1 that a range of long-period cyclic components are being fitted: the record length of the observed data is only just sufficient to start to distinguish some of these from each other but, considering the limited objectives of the present study, it was felt worth including them as separate components. Further, the record length is such that long-period cycles of length 40 or more years cannot be distinguished from slowly-changing non-cyclic trends. Again, for present purposes, the assumption of a linear trend with time seemed reasonable with other variations being included in the long-period cycles. However, it would equally well have been possible to adopt a quadratic trend.

As noted in Section 3.2, the Nyquist frequency for the ½ lunar day step-length corresponds to a period of 1.03505 days. Again, the harmonic at this period length is included as the cosine term only, and the harmonic at 1 solar day is aliased with that at a period length of 1.072646 days.

Table 4.1 *Seasonal and tidal harmonics in models for individual sites*

1	45.1204680 years	31	182.6211000 days	61	9.1206980 days
2	18.6125770 years	32	177.8441770 days	62	7.5349717 days
3	9.3062885 years	33	157.2950287 days	63	7.3826450 days
4	6.2041923 years	34	138.1200260 days	64	7.0957910 days
5	4.6531442 years	35	129.6727600 days	65	7.0876508 days
6	3.7225154 years	36	121.7478104 days	66	5.9920797 days
7	3.1020962 years	37	100.2206879 days	67	5.8226180 days
8	2.6589396 years	38	97.3503200 days	68	5.8017168 days
9	2.3265721 years	39	76.9205320 days	69	5.6422262 days
10	1.5510481 years	40	71.2028580 days	70	4.9217640 days
11	1.1632861 years	41	63.1352997 days	71	4.7926000 days
12	16.8725590 years	42	60.2784729 days	72	1.7504646 days
13	8.8499540 years	43	31.8119450 days	73	1.1677692 days
14	6.9175830 years	44	27.5543500 days	74	1.1603530 days
15	5.9980060 years	45	27.0926570 days	75	1.1566753 days
16	4.4249770 years	46	24.5558586 days	76	1.1530229 days
17	2.9499840 years	47	24.3617115 days	77	1.1195149 days
18	2.2124880 years	48	19.6879539 days	78	1.1134601 days
19	1.8054020 years	49	15.3873000 days	79	1.1130767 days
20	1.1273940 years	50	14.7974739 days	80	1.0821808 days
21	1.0169070 years	51	14.7653341 days	81	1.0789838 days
22	365.2422000 days	52	14.7333334 days	82	1.0759759 days
23	328.1625977 days	53	13.7772630 days	83	1.0758057 days
24	255.5074615 days	54	13.6607920 days	84	1.0726460 days
25	222.1146240 days	55	13.6334200 days	85	1.0695052 days
26	217.6259613 days	56	10.2957869 days	86	1.0353652 days
27	209.2442000 days	57	10.0845970 days	87	1.0350500 days
28	205.8858640 days	58	9.6137170 days		
29	201.0527191 days	59	9.5568530 days		
30	192.9905833 days	60	9.1329316 days		

Table 4.2 *Summary of tidal and flow-related models for the individual sites.*

	Site		
	Southend	Tower Pier	Richmond
error criteria (m)			
rmse	0.1759	0.1821	0.1900
min residual	-2.1966	-2.2386	-2.2126
max residual	1.9684	1.8199	1.3501
trend and flow coefficients			
a	2.3129	3.1172	3.4025
b	0.1517	0.3873	0.4197
f_1	0.0134	0.1649	0.1932
f_2	0.2342	0.3703	0.7552
seasonal and tidal amplitudes (m)			
18.61 years	0.0815	0.0663	0.0798
1 year	0.0582	0.0686	0.0784
½ year	0.0153	0.0221	0.0222
27.554 days	0.2015	0.2369	0.2428
14.765 days	0.4750	0.5716	0.5997
13.661 days	0.1127	0.1365	0.1415
7.383 days	0.0521	0.0669	0.0804
1.0758 days	0.0580	0.0761	0.0864
1.0695 days	0.0343	0.0465	0.0554
total amplitude of harmonics			
all harmonics (m)	2.0528	2.2574	2.4748

The results of the models of the form given by Equation (4.1) are summarised in Table 4.2. As in Section 3.2, it is convenient to summarise the harmonic components of the models by quoting only the amplitudes represented by a selection of the terms, together with the total of the amplitudes of all the harmonic components. It should be noted that the models for the individual sites have been fitted entirely separately, and thus there has been no attempt to ensure that amplitudes and phases of the various harmonic components in the individual models follow any preconceived notion of how these should be related.

The results in Table 4.2 may be described as follows. First, the root mean square errors of the models for Tower Pier and Richmond are considerably larger than those reported in Table 3.3: this is as expected, since here the concurrent sea-levels at Southend are not being used as part of the estimators. The largest residuals of the present models are of the order ±2 m. The trends for the individual sites suggested by these models range from 15 cm per century at Southend to 39 and 42 cm per century at Tower Pier and Richmond, respectively, although it should be recalled that "trend" and long-period harmonics are not readily distinguishable.

Second, as in Section 3.1, the difference in estimated sea-levels between occasions when the flow at Teddington is 500 or 0 cumecs is given by $f_1 + f_2$. Because yearly and half-yearly cycles are also being fitted, it is reasonable to interpret $f_1 + f_2$ as being the difference attributable to the different flows if they occurred at the same time of year. In Table 4.2, this flow-effect ranges from 25 cm at Southend, to 53 cm at Tower Pier and 95 cm at Richmond. The finding of a large flow-effect at Southend is perhaps surprising, given the width of the Thames estuary at this point, but the following considerations are relevant.

(a) The eventual peak level at Southend is the outcome of both flow-quantity and momentum effects.

(b) The models estimate a consistent progressive increase to the size of this apparent effect on moving up-river: this includes the site at Gallions, which has only a short record.

(c) The flow-effect coefficients need not be interpreted as representing the existence of an underlying causative mechanism: the coefficients may instead be representing a joint statistical dependence between sea-levels and river flows, in that the same weather mechanisms contribute to both. However, given the different time-scales for variations in river-flow and sea-levels, there is some doubt over whether the joint-dependence explanation could be valid.

One should, of course, consider the question of the statistical significance of these estimated coefficients. As was the case with the models used in Section 3, the residuals of the present models exhibit a moderate degree of serial correlation and there are indications of other patterns in the residuals, as discussed in Section 4.3. Thus assessing the statistical significance of the coefficients is not straightforward. One way of dealing with the problem of serial correlation is to consider a revised form of the same problem in which "pre-whitened" versions of the observed and explanatory variables are used in place of the originals and where the pre-whitening model is constructed so as to remove the serial correlation in the residuals. Note that the results reported in Section 4.3 indicate tests based on this type of approach are also invalid. Nonetheless, it is interesting to note that when this approach was implemented, it was found that the coefficient of the square-root of flow for the Southend model did appear to be significantly different from zero, at about 5½ times its estimated standard deviation (compared with 8 times according to the incorrect standard approach). The coefficient for the second flow term is small for Southend, but it has been retained in the model so as to use the same model-structure for all sites.

For the period-lengths for which results are quoted, which are those with largest amplitude, there is a steady progression up-river of the amplitudes of these components: there is an exception for the 18.61 year cycle but, in the present model, this is not clearly distinguishable from the 16.87 year cycle which is also included. Of course, these amplitudes refer only to cycles in the peak-values, not across the full range of tidal-river behaviour. Note that the total amplitude of the tidal harmonics is calculated as the sum of the amplitudes of the individual components: it reflects a combination of the individual components which does not necessarily occur within the period of record.

Figures 4.1 to 4.4 are revised versions of Figures 2.1 to 2.4, showing the predictions of sea- and river-level obtained from the models here. Figure 4.5 shows the autocorrelation functions of the residuals from the model: specifically, these are the terms subscripted "surge" in the expressions in Equation (4.1). The autocorrelation functions show that the any short-period cyclical behaviour has been successfully modelled by the combination of period-lengths chosen here, but that there is still a moderate amount of non-cyclical serial correlation in the residuals.

4.3 Models for the size of surges

When models of the form discussed in Section 4.2 are fitted by least squares, it follows that the residuals, in this case identified with the surge, will be uncorrelated with the terms for which a linearly occurring multiplication factor has been fitted. This means that the surge components will be uncorrelated with time, with flow and with the identified seasonal and tidal harmonics. However, this does not mean that there is statistical independence between the surge and these quantities. In order to investigate this, it is possible to examine whether the size of the surge component is related to these other variables. One way of doing this is to undertake a least squares analysis of the absolute values of the surge-residuals, and the results from this are reported next. Because the distribution of the absolute values of the surge-residuals is rather skew, one could consider doing an analysis of an alternative transformation of the residuals, but one which still reflects the size of the residuals. While such an analysis will not be reported in detail here, it seems worth recording that essentially similar results to those for the absolute values were obtained when the variable analyzed was the logarithm of a shifted version of the absolute values.

The model-fitting procedures described in Section 4.2 have several outcomes besides the values of the fitted parameters. For example, for the fitted parameters it is possible to compute the series of values S^*_{harm} and S^*_{surge}, being the sum of the harmonic components and the model residuals, S_{harm} and S_{surge}, for the best-fitting model. Similar quantities can be evaluated for the models for Tower Pier and Richmond. It is therefore possible to consider a least squares analysis in which one sees how well the absolute value of the residuals, $|S^*_{surge}|$, can be predicted by a linear model involving, as before, a possible trend with time, a possible relationship to the flow at Teddington and possible seasonal and tidal harmonics. It is also possible to consider using the identified harmonic component, S^*_{harm}, as an additional candidate explanatory variable on the grounds that, if there is any relationship of size of surge to the tidal pattern, it might be expected that it would be most strongly related to the sum of the individual components to the same extent that they influence the overall sea or river-level. When such an analysis was done, the following conclusions were drawn.

(a) The inclusion of a linear-in-time trend term did not result in an improved predictor. This conclusion is possibly surprising in view of the changing precision of recording described in Section 2.2, but it perhaps reflects the fact that this aspect of the data is relatively unimportant.

(b) There did seem to be a relationship between the absolute size of the residual and flow, but the two flow-related terms used in earlier models are not required. The strongest relationship appears to be with the square-root-of-flow term.

(c) While the absolute residual is strongly related to many of the individual seasonal and tidal harmonic components, much the same predictive performance can be achieved by using only two elements in the model to reflect these effects. In particular, the total harmonic component from the model for sea-level, S^*_{harm}, with the addition of a seasonal effect of period equal to a year.

Thus the final form of the model relating the absolute value of the surge-residual to the other available quantities is:

$$\begin{aligned} \text{est}\,\{\,|S^*_{surge}(j)|\,\} = a \quad & + f_2\,\{Q_j\,/500\}^{1/2} \\ & + s_y \sin\{2\pi\, t_j/T\} + c_y \cos\{2\pi\, t_j/T\} \\ & + h\, S^*_{harm}(j)\,, \end{aligned}$$

where j is the tidal-cycle number and other terms are as before. Here T is the period length, in this case equal to one year. The quantities a, f_2, s_y, c_y and h are parameters to be estimated by least-squares. Values for the estimated parameters for the three sites are given in Table 4.3. Note that, for Tower Pier and Richmond, the term S^*_{harm} in the model for Southend is replaced by the total harmonic components T^*_{harm} and R^*_{harm}, respectively. Table 4.3 includes the results for versions of the above model in which the coefficient h is fixed at zero, thus excluding the predicted tide-level from the predictor of the size of the surge component.

The interpretation of Table 4.3 is that there is a strong seasonal effect in the sizes of residuals over the year. If the same moderate size of river flow is assumed, say 100 cumecs, then the mean absolute value of the surge component at Southend varies between 17½ cm in winter and 8 cm in summer. The relative effects of flow and tidal range can be judged by noting that a change from 0 to 500 cumecs in the flow would lead to the estimate of the mean absolute value of the residual being 4.7 cm larger than a typical zero-flow value of 10.7 cm, while the tidal range component might add or subtract 1.3 cm if the expected high-tide level changes by ∓2 m from its central value. It is interesting to note that the analysis using Model 1 indicates that high-tide levels are more predictable when the expected value is high than when it is low, and that this effect is progressively stronger on moving upstream. The results for Model 2, in which the tidal-range effect is removed indicates that the inclusion of this effect is of little immediate benefit in predicting the typical size of the surge effect, at least when judged in terms of the root mean square error with which the absolute value can be predicted. However, if a pre-whitening approach to allowing for the serial correlation in the residuals from this model is adopted, then the coefficient h in Model 1 does appear to be significantly different from zero.

Table 4.3 *Summary of models relating the absolute size of the surge-residuals to tidal, seasonal and flow-related effects.*

Model for size of residuals	Site		
	Southend	Tower Pier	Richmond
Model 1, including tidal effect			
rmse	0.1212m	0.1247m	0.1291m
a	0.1066	0.1082	0.1269
f_2	0.0472	0.0601	0.0228
s_y	-0.0053	-0.0059	-0.0074
c_y	0.0471	0.0421	0.0439
h	-0.0065	-0.0118	-0.0232
Model 2, excluding tidal effect			
rmse	0.1212m	0.1248m	0.1296m
a	0.1066	0.1081	0.1264
f_2	0.0474	0.0603	0.0242
s_y	-0.0050	-0.0055	-0.0065
c_y	0.0473	0.0427	0.0453
h	0.0	0.0	0.0

The models of this section need to be put into context, since they are not themselves directly useful. One possible use for the models here would be to provide a way of supplying weights for a weighted least squares analysis of models of the type used in Section 4.2. It would then be possible to repeat the analysis here using the surge-residuals from the weighted analysis, and one could iterate this process until stability has been reached. A use of somewhat more importance for the present analysis is that it provides a way of assessing the statistical dependence of the components of models of the form given by Equation (4.1). As noted at the beginning of this section, the surge-residual identified in such models by the method of least squares will be uncorrelated with the flow and tidal effects, but the analysis here indicates that there is statistical dependence. In particular, it reveals that there is a dependence of the variability of the residual related strongly to a yearly seasonal cycle, moderately strongly to the river-flows and weakly to tidal component of the initial model. While it would be possible to investigate other aspects of dependence, for example the skewness of the residuals, by a similar procedure to that used here, for present purposes the finding of statistical dependence is sufficient.

4.4 Models for river-levels using the surge-residual at Southend

This section considers models of the type given in Equation (4.2). Here the surge-residuals at Southend, identified by the models discussed in Section 4.2, are used as part of the estimator for river-levels further upstream. The resulting estimator is therefore a simple version of an alternative structure-function approach, as discussed in Section 3.3, in which the estimated river-level would be constructed as a function of three variables: the river-flow at Teddington, the seasonal and astronomical cycle-based predictor of peak river-levels and the "surge" experienced at Southend. To be slightly more specific than earlier, the models being fitted are of the form:

$$\begin{aligned} \text{est (T)} &= T_{trend} + T^{*}_{harm} + T_{flow} + m\, S^{*}_{surge}, \\ \text{est (R)} &= R_{trend} + R^{*}_{harm} + R_{flow} + m\, S^{*}_{surge}, \end{aligned} \qquad (4.5)$$

where, as in Section 4.3, the terms with an asterisk superscript denote values from data-series fitted as part of the estimation procedures described in Section 4.2. For the present least squares fitting scheme, the problem is regarded as having 5 free parameters: the coefficients a and b of the trend term, the coefficients f_1 and f_2 of the flow-effect term and the multiplying factor m for the surge-effect. The trend and flow-effect terms are given by Equations (4.3) and (4.4). While it would have been possible to re-fit the coefficients of the seasonal and tidal harmonics within the present estimator, this was not undertaken because of the computation-time this would have taken.

One advantage of the approach adopted here is that it enables the tidal and seasonal harmonics to be fitted to a larger number of data points, since otherwise these are restricted to those tidal peaks for the target site and Southend are both recording. It turns out that the coefficients for the trend and flow-effect terms fitted for the present model are extremely similar to the values found earlier. In comparison with the models for Tower Pier and Richmond which were fitted in Section 4.2, the estimator here makes use of a single extra term involving the estimated surge at Southend, which one might expect to be closely related to the surge-residuals, T_{surge} and R_{surge}, of the original models.

The results for fitting the present model are shown as those for Model 1 in Table 4.4. The second model in Table 4.4 will be discussed a little later: it is an extended form of the first model in which some interaction terms are included.

Table 4.4 *Summary of the estimators making use of the surge-residual at Southend.*

Model coefficients		Tower Pier	Richmond	Error Criterion	Tower Pier (m)	Richmond (m)
Model 1						
trend and flow coefficients						
	a	3.1169	3.4018	rmse	0.0649	0.0885
	b	0.3838	0.4163	min	-0.5172	-0.6639
	f_1	0.1619	0.1883	max	0.3689	0.6301
	f_2	0.3759	0.7630			
factor for surge-residual at Southend						
	m	0.9658	0.9509			
Model 2						
trend and flow coefficients						
	a	3.1170	3.4012	rmse	0.0648	0.0875
	b	0.3836	0.4172	min	-0.5121	-0.6234
	f_1	0.1605	0.1818	max	0.3678	0.5376
	f_2	0.3765	0.7665			
factor for harmonic						
	h_f	0.0	-0.0142			
factor for surge-residual at Southend						
	m	0.9201	1.0628			
	m_t	0.1349	0.1420			
	m_h	0.0	-0.0458			
	m_f	0.0339	-0.3884			

From the results for Model 1 in Table 4.4, the following two points are notable.

(a) The multiplication factors for the Southend surge-residual are slightly less than one. Thus there is certainly no strong evidence for an amplification of the surge as it moves upstream, even though there was evidence in earlier models (for example in Table 4.2) for an amplification of the tidal and seasonal harmonics. The effect of any observation errors in the identification of the surge residual would be such as to reduce the regression coefficient by a small factor. However, in the present circumstances this effect would not be large enough to lead to the values of m here if the factor applicable (if the true surge-effect was known) is really unity. Thus the values for m found here suggest a definite, if only small, reduction in the size of the contribution made by the surge at the upstream river sites. When an attempt was made to confirm the statistical significance of the difference from unity of this coefficient, somewhat contradictory results were found. On using the pre-whitening approach, the estimated standard error for the coefficient was found to be small enough for the difference here to be significant, but at the same time the estimated values of the coefficients moved rather closer to unity (values of 0.9939 and 0.9813 for Tower Pier and Richmond, respectively).

(b) In comparison with the seasonally and tidally modified versions of the structure-function approach, reported in Table 3.3, the sizes of the estimation errors found here are somewhat smaller. However, this may well be explainable by the fact that the present models make use of a much fuller decomposition of the tidal effect into its astronomically based components. This has not been pursued since there seems to be a strong appeal in the structure of the model used here, compared to that based directly on observed levels at Southend, in which the harmonic components do not have the direct interpretation of being attributable to a given site. Nonetheless, the results here do confirm that there is substantial benefit to be gained over the original structure-function approach, the results for which are given in Table 3.2, by taking into account the underlying tidal cycles.

In order to explore whether the historical record can provide any evidence on the question of whether the separately identified terms in model (4.5) interact in any way, a brief exploration was undertaken of the possible benefits of including interaction terms in the estimator. The model for Richmond which includes interaction effects defines the estimated value to be of the following form:

$$\begin{aligned} \text{est}\,\{R(j)\} = a \; & + b\,\{t_j/100\} \\ & + [1 + h_f\,\{Q_j/500\}^{1/2}]\,R^*_{harm}(j) \\ & + f_1\,\{Q_j/500\} + f_2\,\{Q_j/500\}^{1/2} \\ & + [m + m_t\,\{t_j/100\} + m_h\,R^*_{harm}(j) + m_f\,\{Q_j/500\}^{1/2}\,]\,S^*_{surge}(j). \end{aligned} \tag{4.6}$$

This is Model 2 of Table 4.4. The model in Equation (4.5), called Model 1 in Table 4.4, is a special case of this for which the additional parameters h_f, m_t, m_h and m_f are all zero. The particular interaction terms in this model represent the selection of those tried which led to apparently significant coefficients when judged by the pre-whitening approach. Note that only the square-root-of-flow is included in the interaction effects rather than using the ordinary flow as an alternative or additional term: this seemed to have the larger effect and it did not seem necessary to include both. It would, of course be possible to rearrange the terms in the above expression so that, for example, the term involving the product of the harmonic and surge components (with parameter m_h) could be grouped in the combination of terms which multiply $R^*_{harm}(j)$. However, the present grouping seems a reasonable one. Once again, a similar model is used for Tower Pier, this time involving the fitted tidal and seasonal harmonics for that site.

In the results for Model 2 in Table 4.4, there is some suggestion of a small interaction between the harmonic and flow-effects at Richmond, suggesting a reduction or increase, for relatively high or relatively low tidal peaks respectively, of approximately 2.8 cm in river-levels at Richmond, assuming a tidal range of ±2 m, compared with an average flow-effect of 95 cm for a flow of 500 cumecs. Proportionately smaller effects would be predicted for smaller flows. No effect of a similar kind is apparent in the data for Tower Pier. The interaction effects related to the surge residual seems to be of a larger order of magnitude: these can be described, in terms of the amplification factor for surges, as follows.

(i) The amplification factor shows an increase over the period of record. For Tower Pier, taking a typical flow value of 70 cumecs, and counting the record as 50 years long, from 0.9328 to 1.0002: for an extreme surge of 2 m, this corresponds to an extra 13½ cm between the beginning and end of the record. Similar values for Richmond are amplification factors of 0.9175 and 0.9885 for the surge-residual at the beginning and end of the record, corresponding to an extra 14 cm for a 2 m surge.

(ii) There seems to be no interaction between the underlying tidal conditions, as expressed by the harmonic component, and the surge residual for Tower Pier, but a modest one seems to arise at Richmond. A high underlying tidal level seems to reduce the factor for the surge at Richmond by 0.09 if the level is 2 m above an average tidal-peak level, corresponding to a reduction of 18 cm in the effect of a 2 m surge. Similarly, an underlying tidal peak 2 m below average would increase the factor for the surge by 0.09, corresponding to an increase of 18 cm in the effect of a 2 m surge.

(iii) The effect of interaction between river-flow and surge appears to be different for the two sites. Thus, compared with a zero flow, a river-flow of 500 cumecs leads to a small increase in the factor on the surge at Tower Pier, but to a large decrease at Richmond. Thus, for the beginning of the record and for a medium tide, the factors would be 0.9201 and 0.9550 at Tower Pier for low and high flows respectively, suggesting an increase of 6.6 cm for a 2 m surge, while for Richmond they would be 1.0628 and 0.6744, suggesting a decrease of 77 cm for a 2 m surge. The opposite directions for the interaction at the two sites could well be explainable in terms of the hydrodynamics of the river, so that this is not of itself worrying.

One possible explanation for the apparent increase in the effect of surges, discussed in (i) above, is that this results from the increase in the levels and standards of flood defence over the period of the historical record. Other changes to the estuary and river-channels, such as dredging, would also have had an effect. For all of the analyses in this report, any tidal peak for which the Thames Barrier did in fact operate is eliminated entirely. While this may have had some effect on the present analysis, its effect seems to be neutral.

A comparison of the root mean square errors of the estimators of Models 1 and 2 in Table 4.4 shows that the model incorporating the interactions does not lead to an improvement in the average performance of the estimates of river-level large enough to justify the extra complexity of the procedure. However, as already mentioned, the improvement is significant when judged by the pre-whitening approach: this reflects the fact that small effects can be detected using a sample size of the order of 30,000 observations. The autocorrelation functions of the residuals from the models based on the surge-residual at Southend and on both the surge and interactions (ie. Equations (4.5) and (4.6)) are shown in Figures 4.6 and 4.7. Considering the small reduction in the root mean square errors, it would be expected that these autocorrelation functions would be very similar, as is indeed the case. Both are included to provide some reassurance as to the possible effect of the interaction terms on the autocorrelation functions. Figures 4.6 and 4.7 are comparable with the corresponding results for the direct model of the type in Equation (4.1) shown in Figure 4.5, and with those for original structure functions and modified versions of these shown in Figures 3.1 to 3.8. The higher values in the autocorrelation functions of the residuals for those models which make use of the observed levels at Southend (either directly or indirectly), compared with those that do not, seems to relate to the fact that in these models a lot of the uncorrelated noise in the residuals has been removed by making use of this explanatory variable, meaning that other serially-correlated components of the noise are more apparent.

Figures 4.8 to 4.11 show examples of behaviour of the estimated river-levels produced by the models making using of the observed tidal or surge-residual at Southend. In particular, these plots show the results of the estimator corresponding to Model 2 of Table 4.4. The estimators make use of estimated seasonal and tidal harmonics, flow-effects, estimated values of the surge-residual at Southend and interactions between harmonic and flow-effects and between the surge-residual and both harmonic and flow-effects.

Although, the effects detected here can be important, it is not envisaged that the types of empirical model constructed here on the basis of historical data should be used in practice as structure functions within the types of analysis outlined in Chapter 1. This analysis has given an indication of the types and sizes of the effects to be expected, but it would be preferable in every sense to base any structure functions on the outcome of hydrodynamical modelling of the river, in particular because of the need to define these structure functions for conditions only rarely (or never) experienced during the historical period.

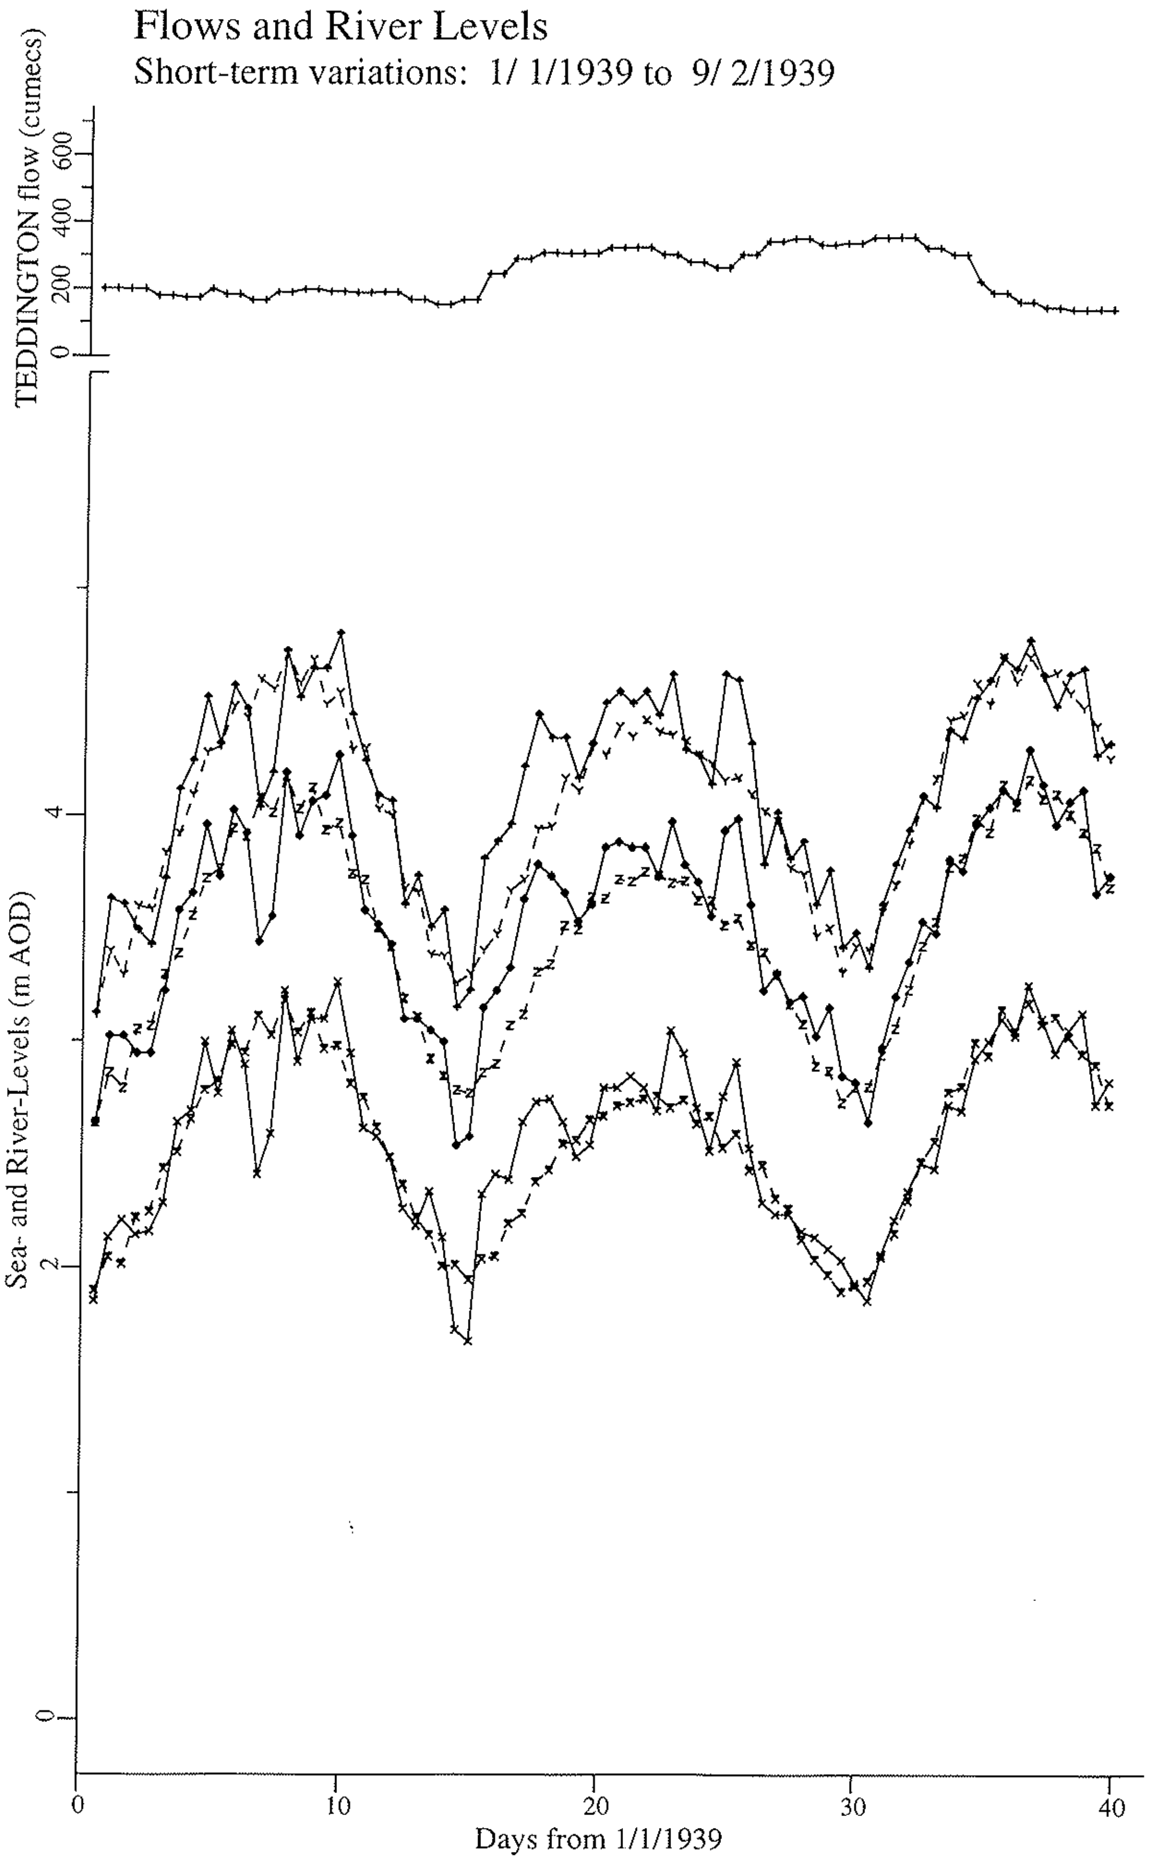

Figure 4.1 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics and on flow-effects (Table 4.1).

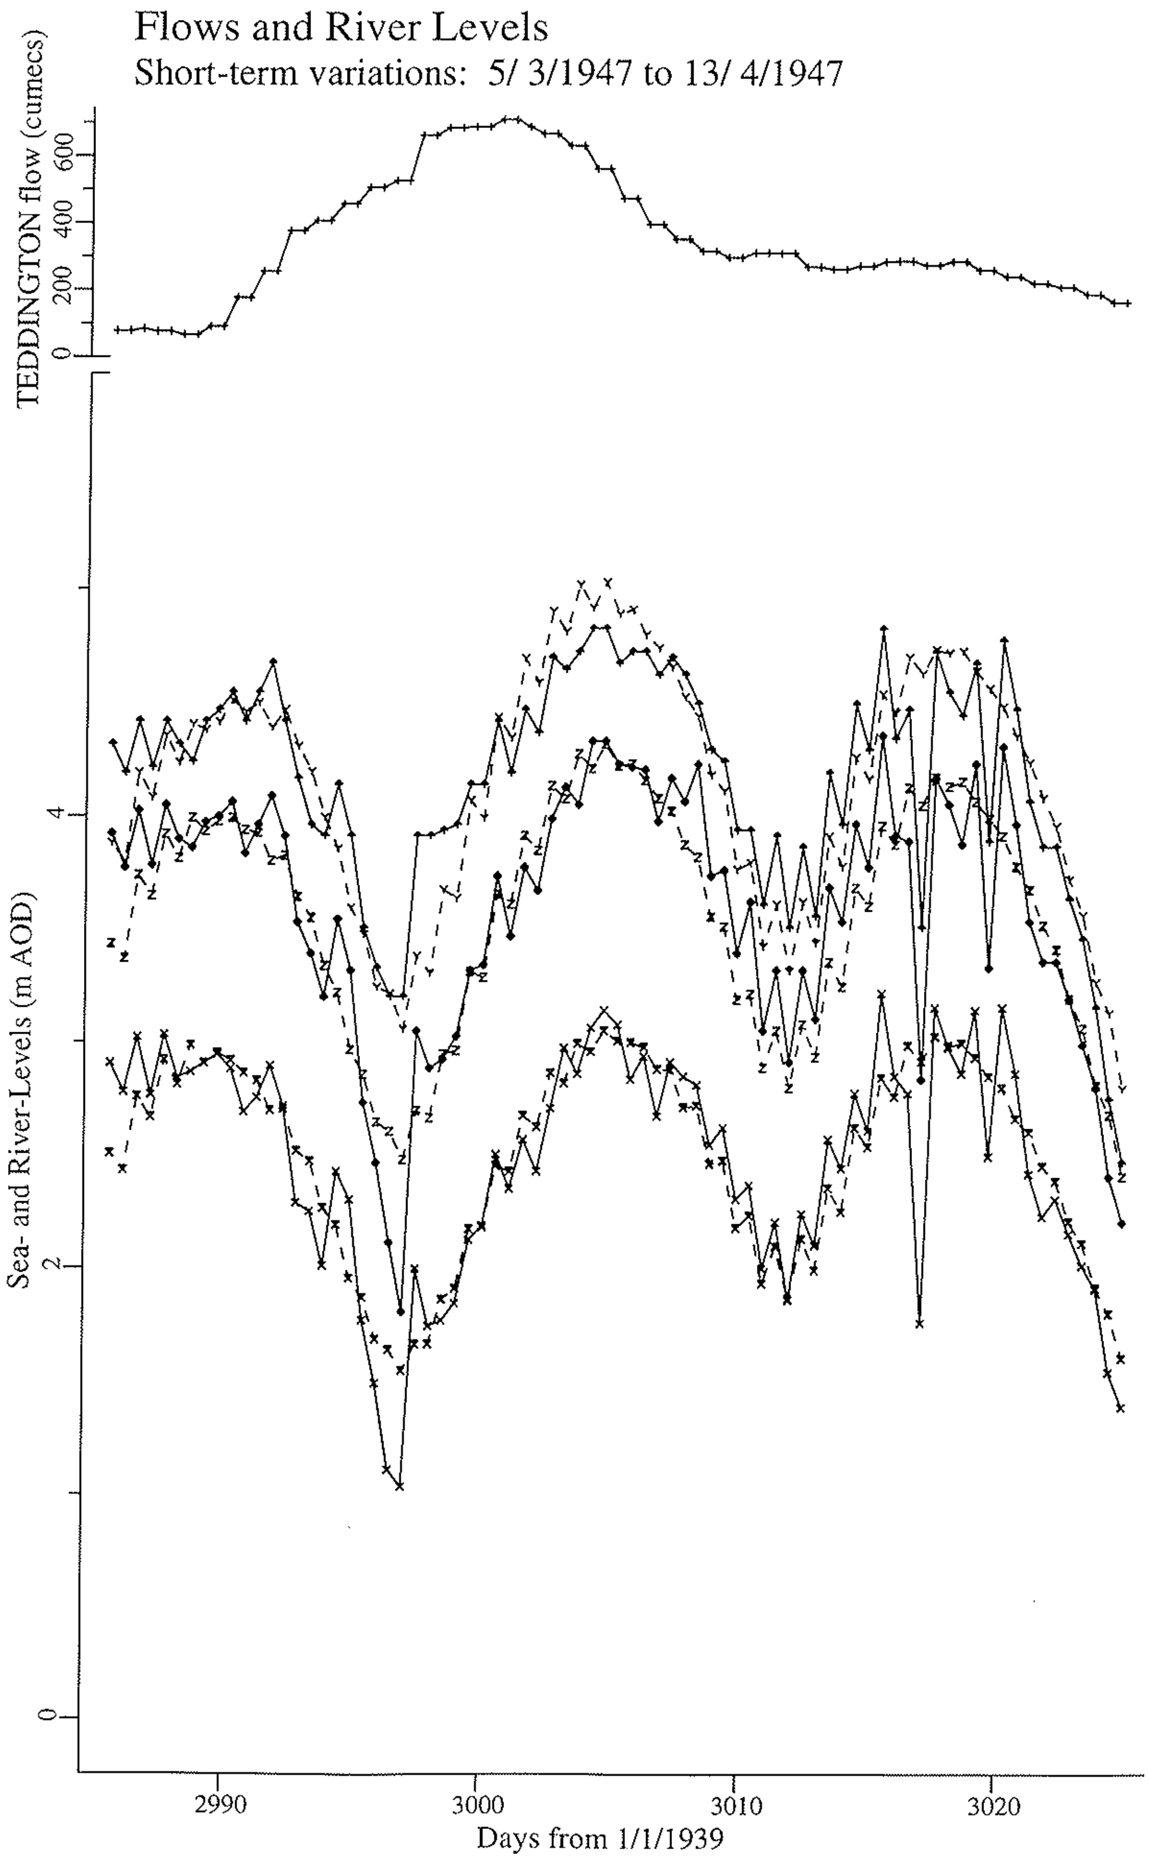

Figure 4.2 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics and on flow-effects (Table 4.1).

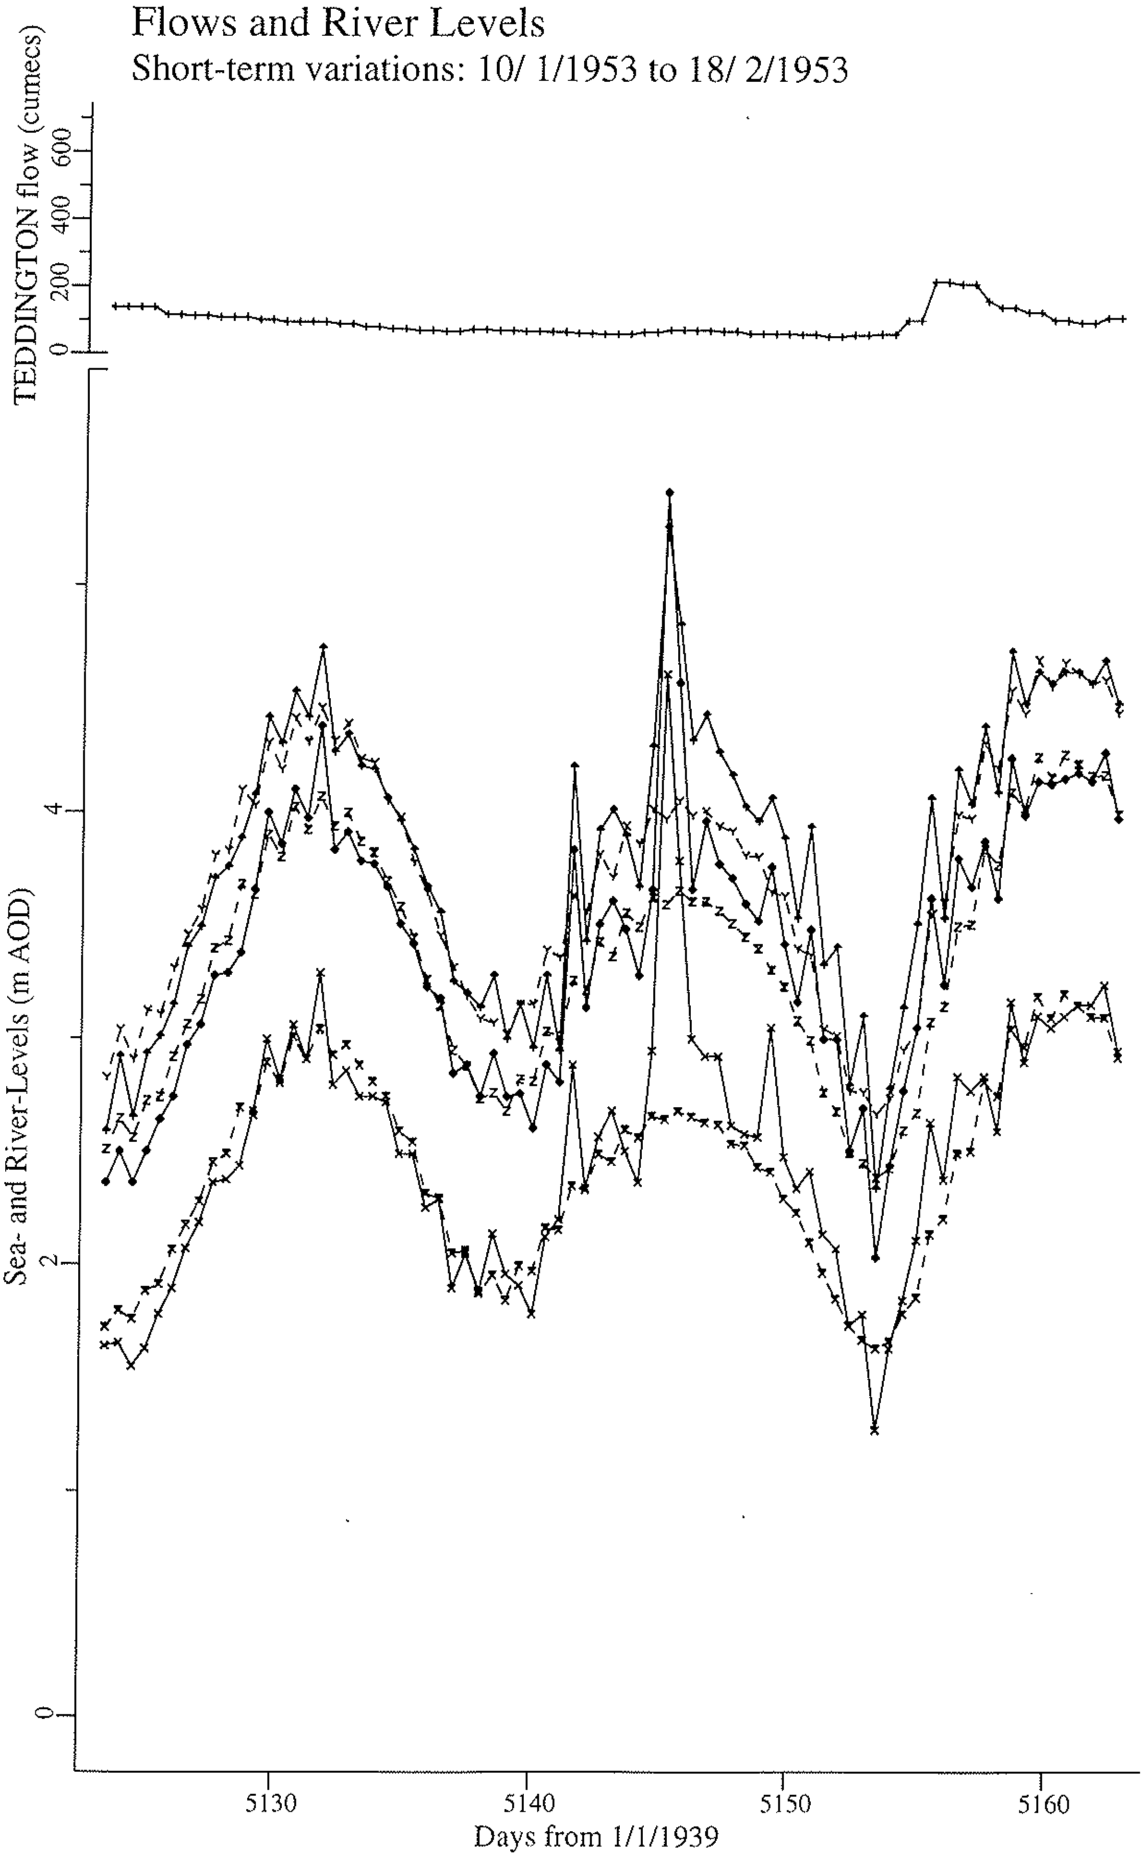

Figure 4.3 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics and on flow-effects (Table 4.1).

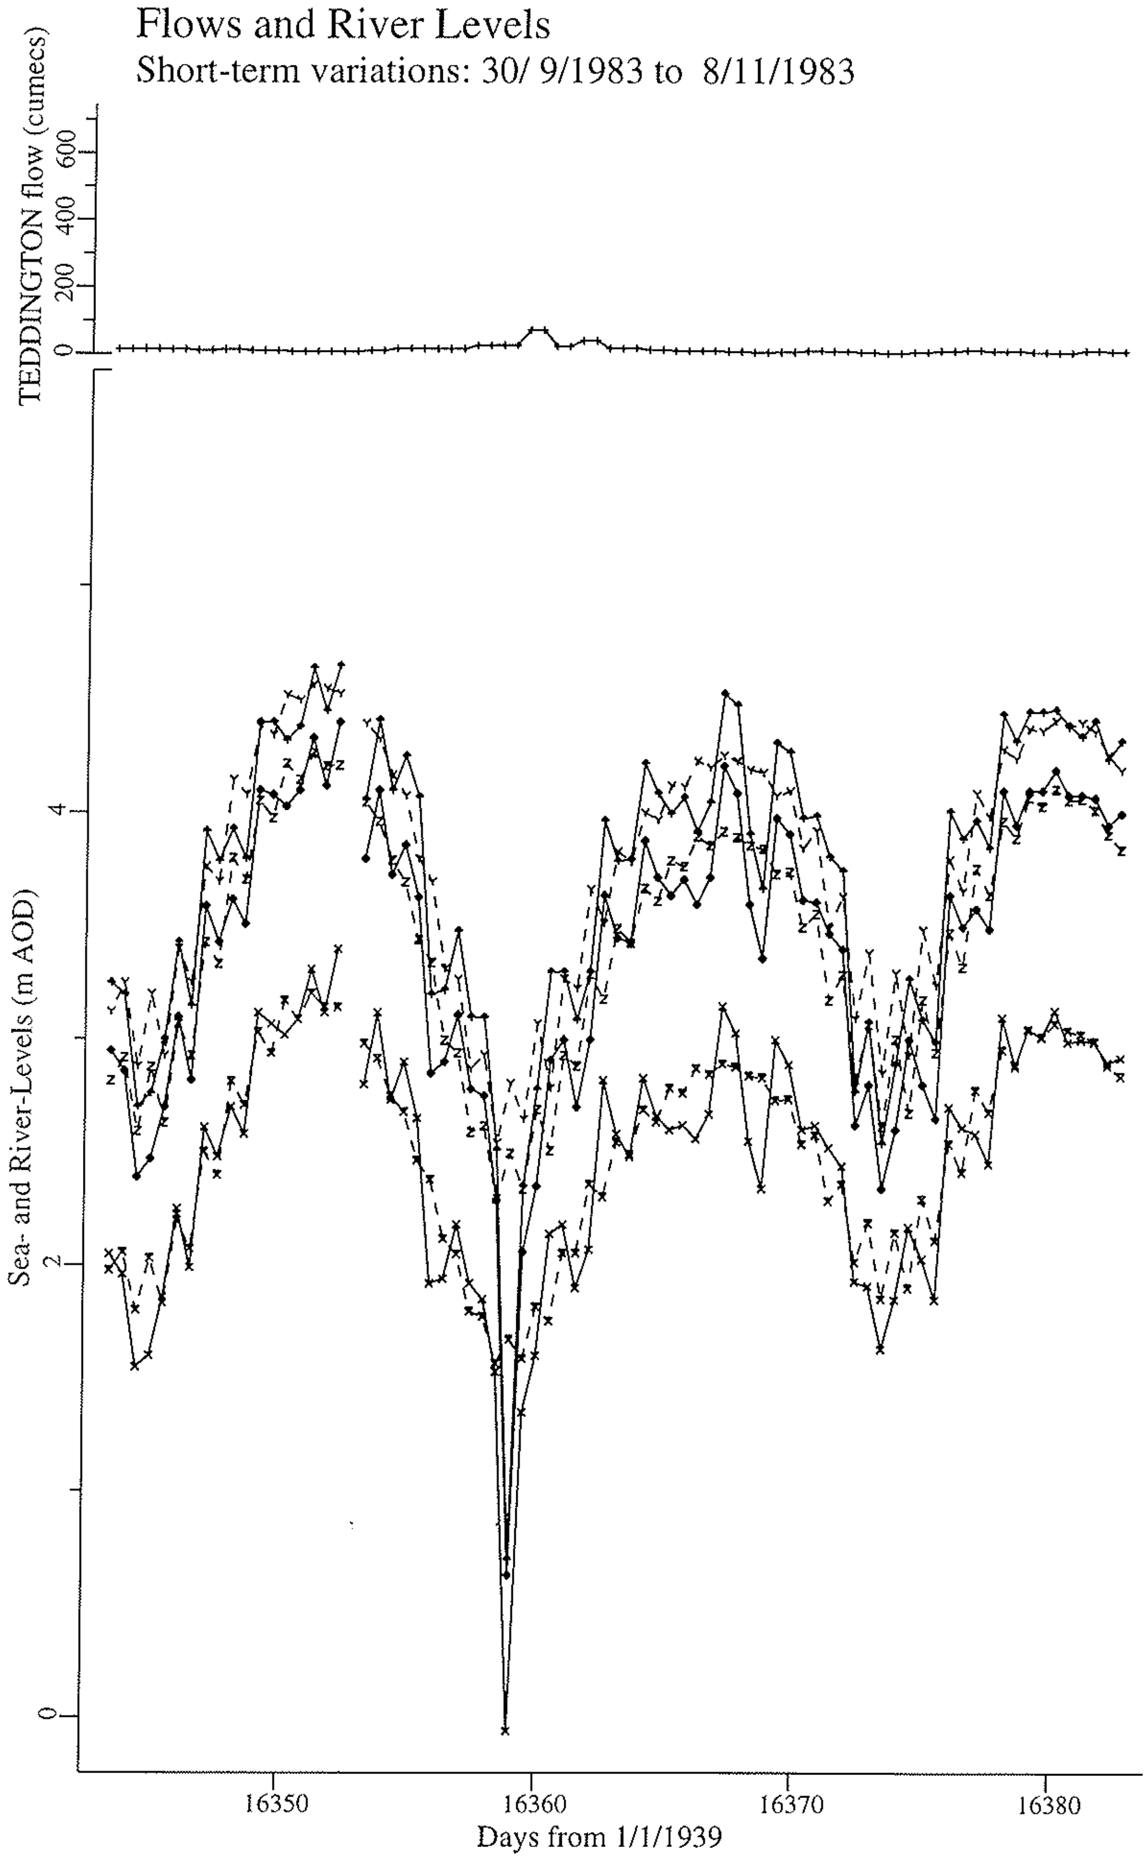

Figure 4.4 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics and on flow-effects (Table 4.1)..

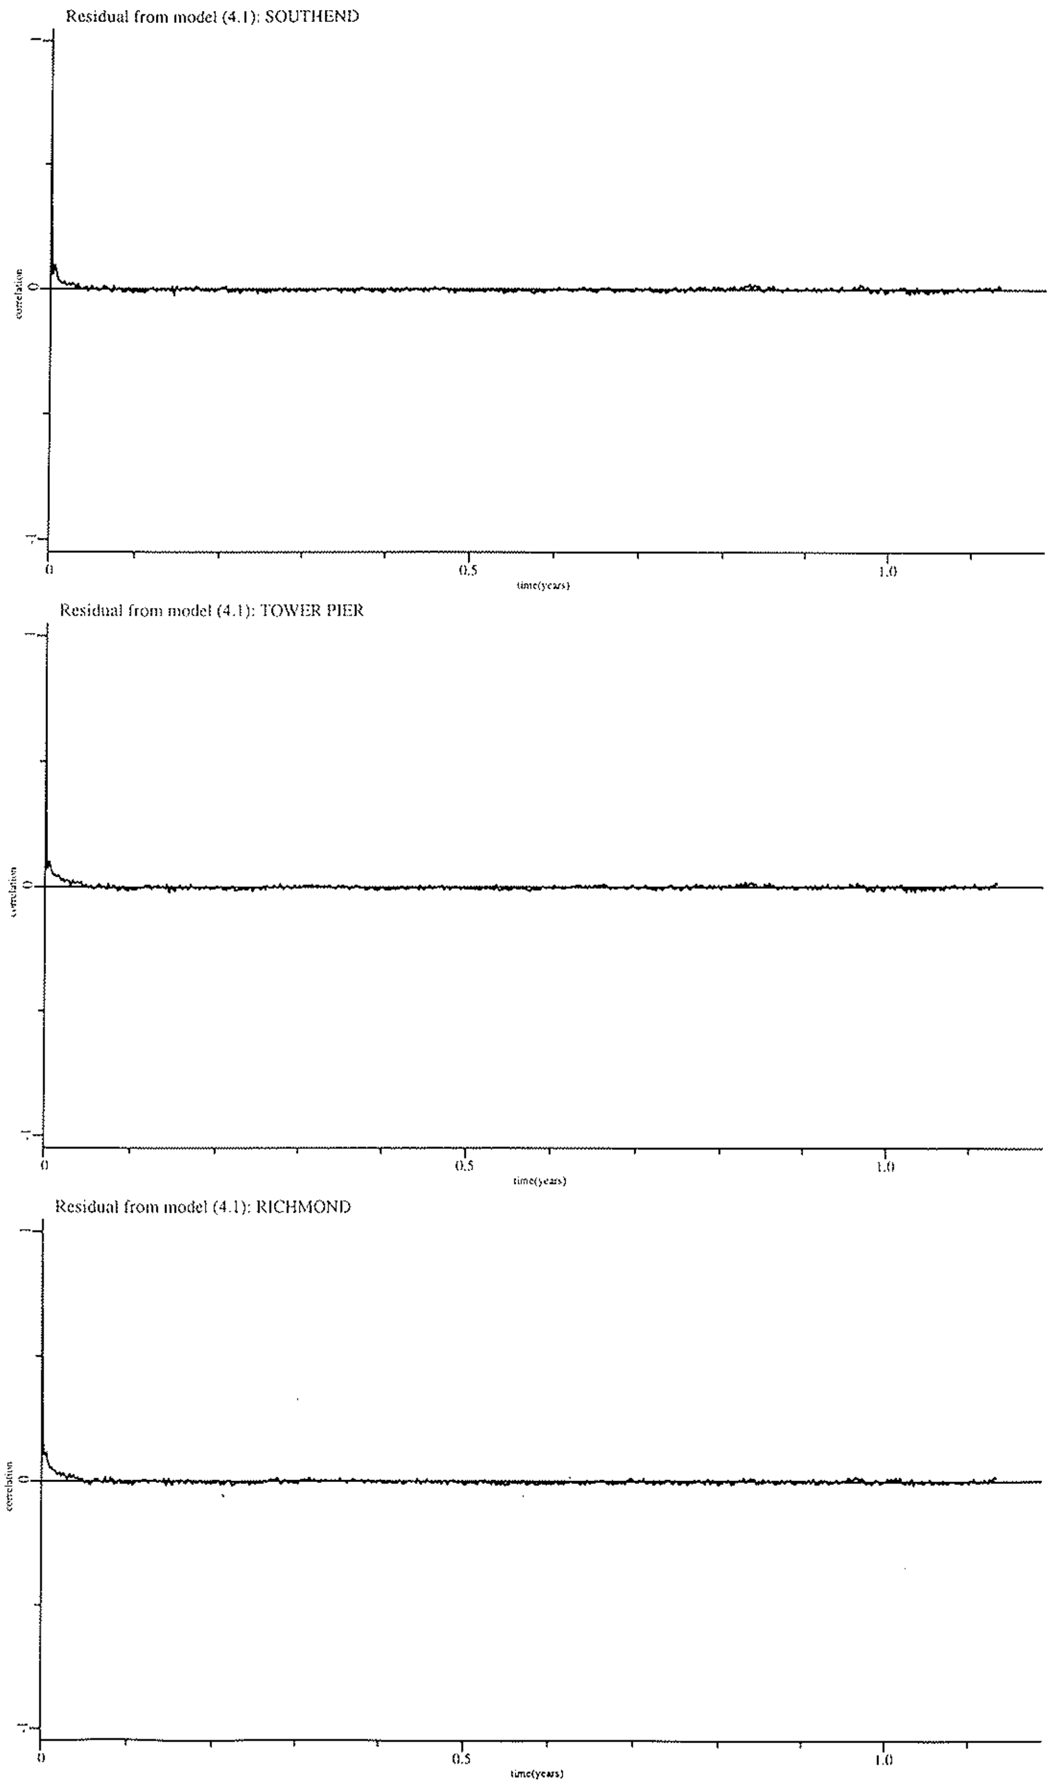

Figure 4.5 Autocorrelation functions of the residuals from the individual site models summarised in Table 4.1.

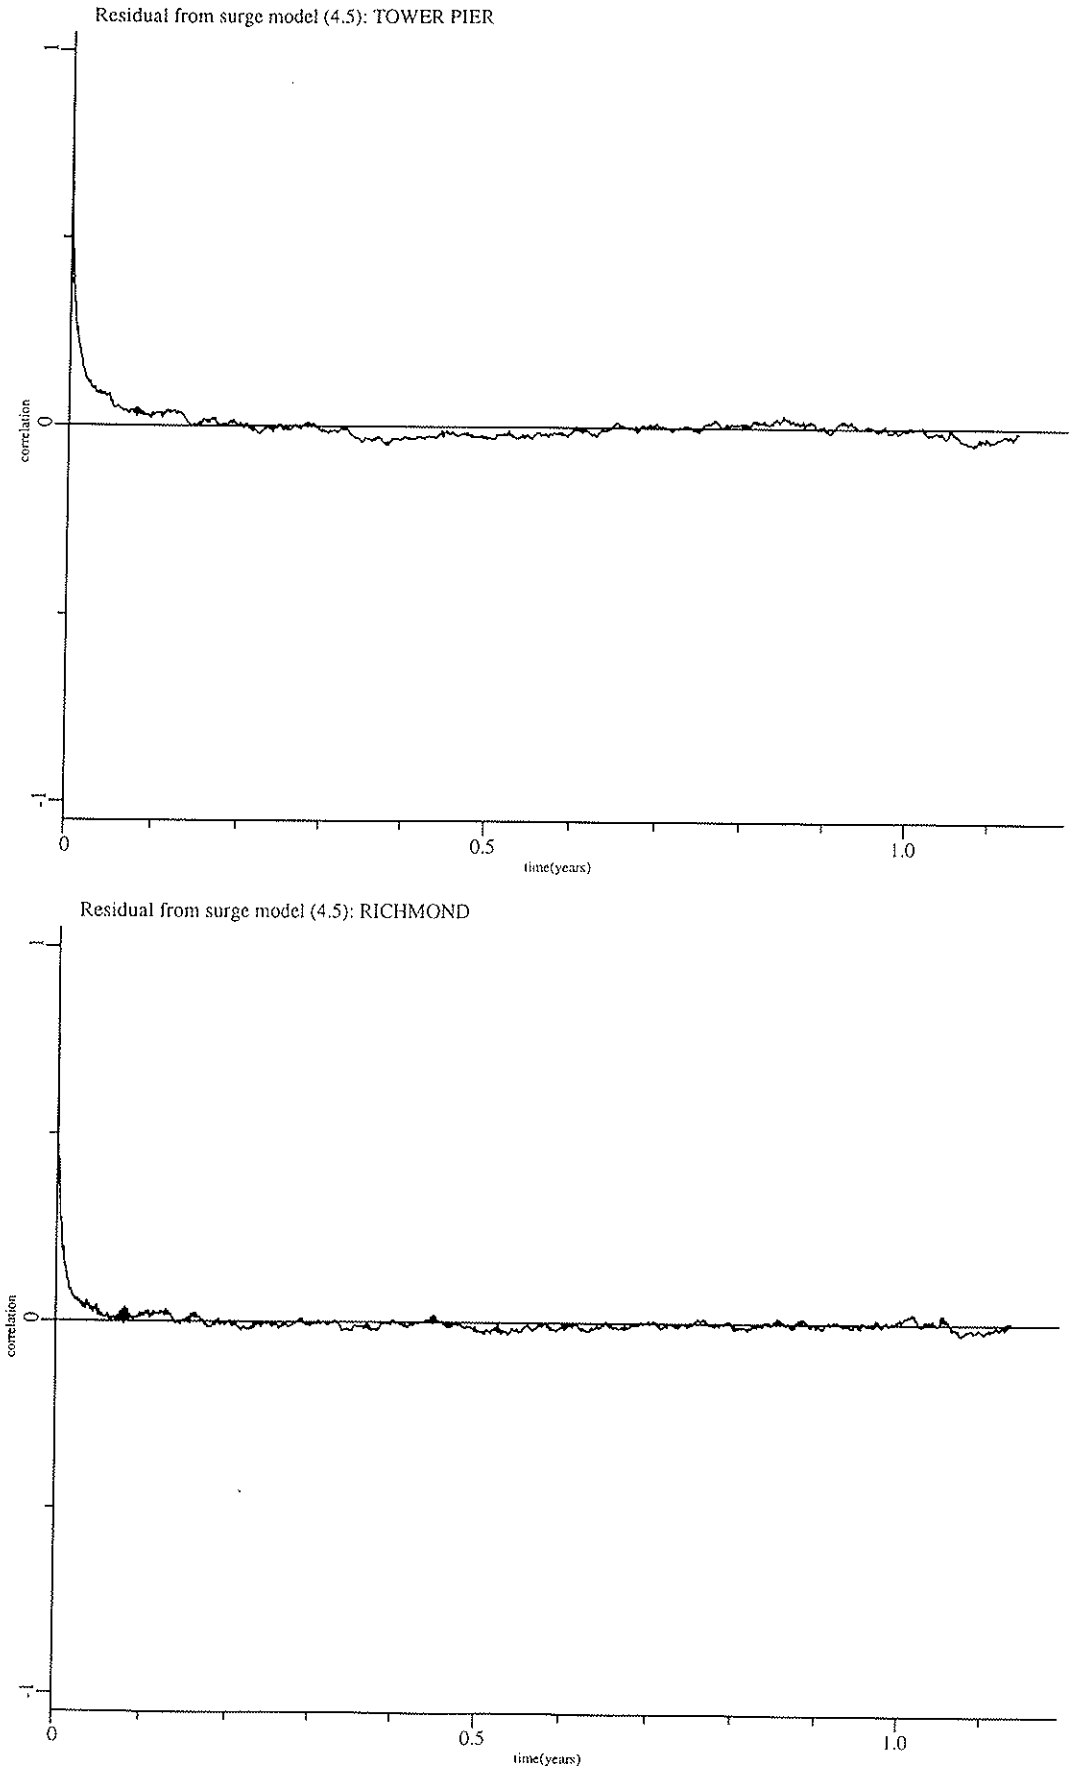

Figure 4.6 Autocorrelation functions of the residuals from the model including the surge (Model 1 of Table 4.4).

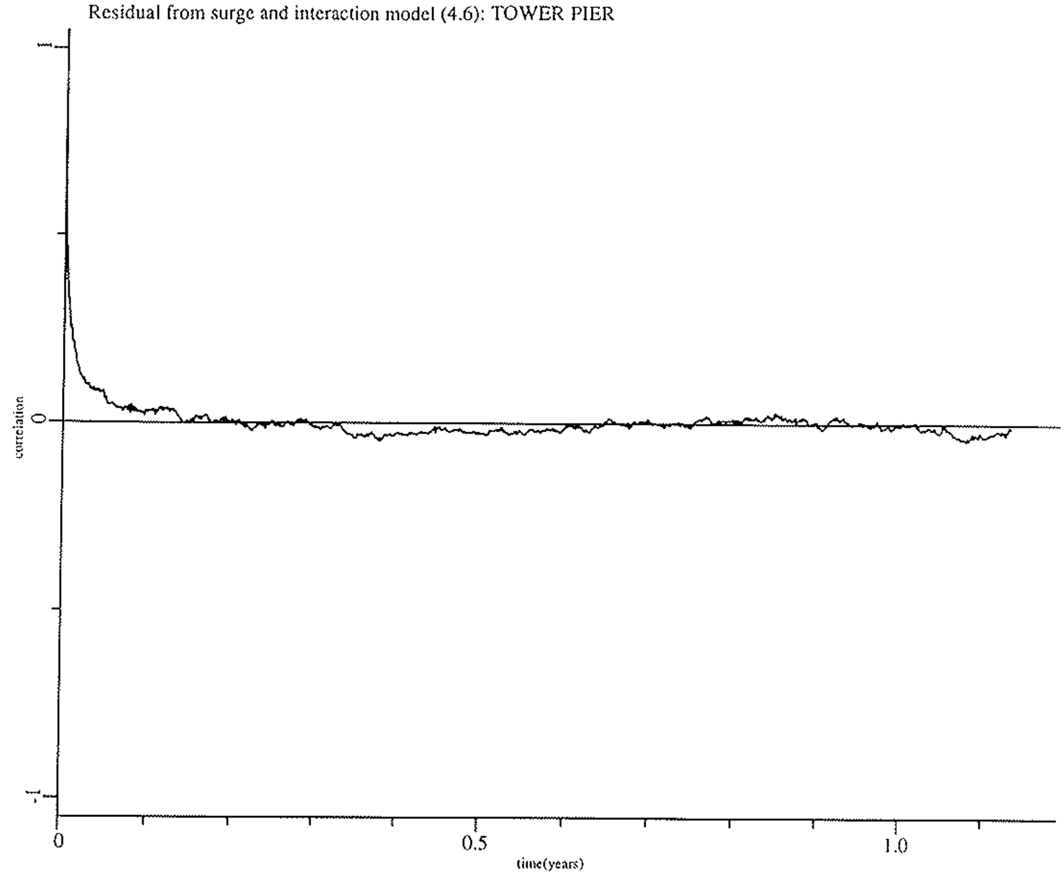

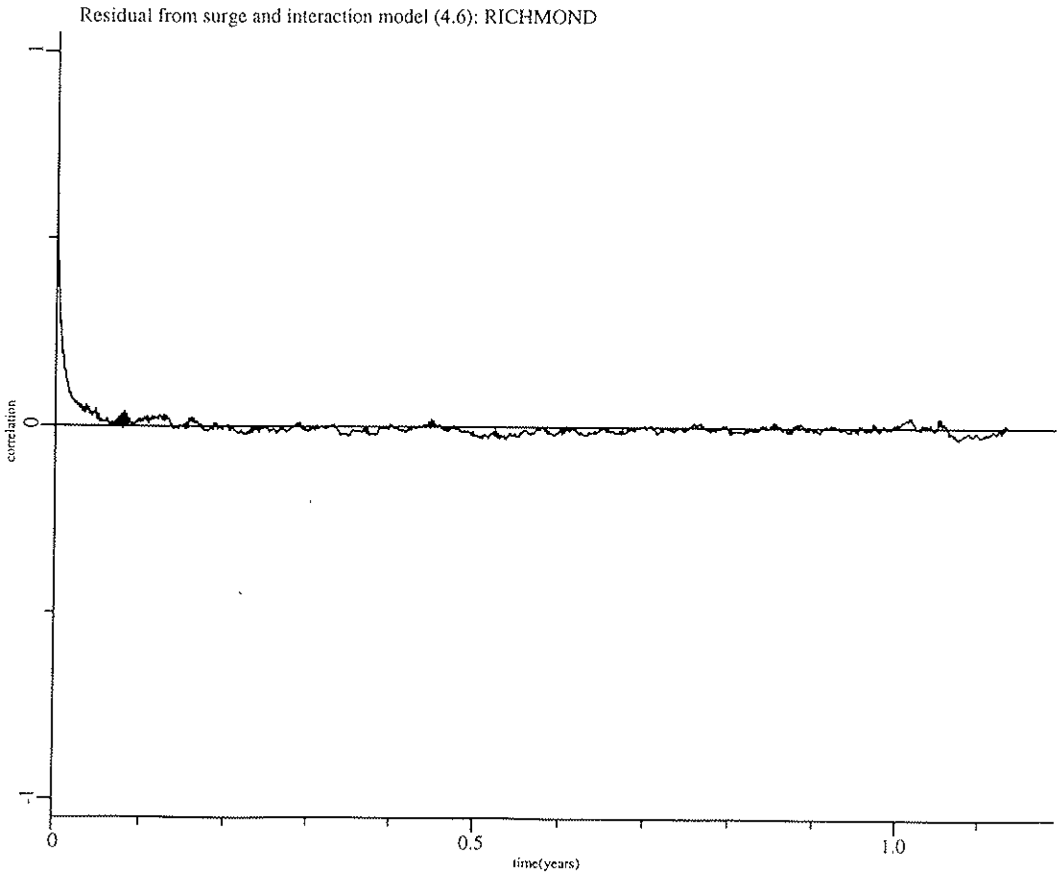

Figure 4.7 Autocorrelation functions of the residuals from the model including the surge and interactions (Model 2 of Table 4.4).

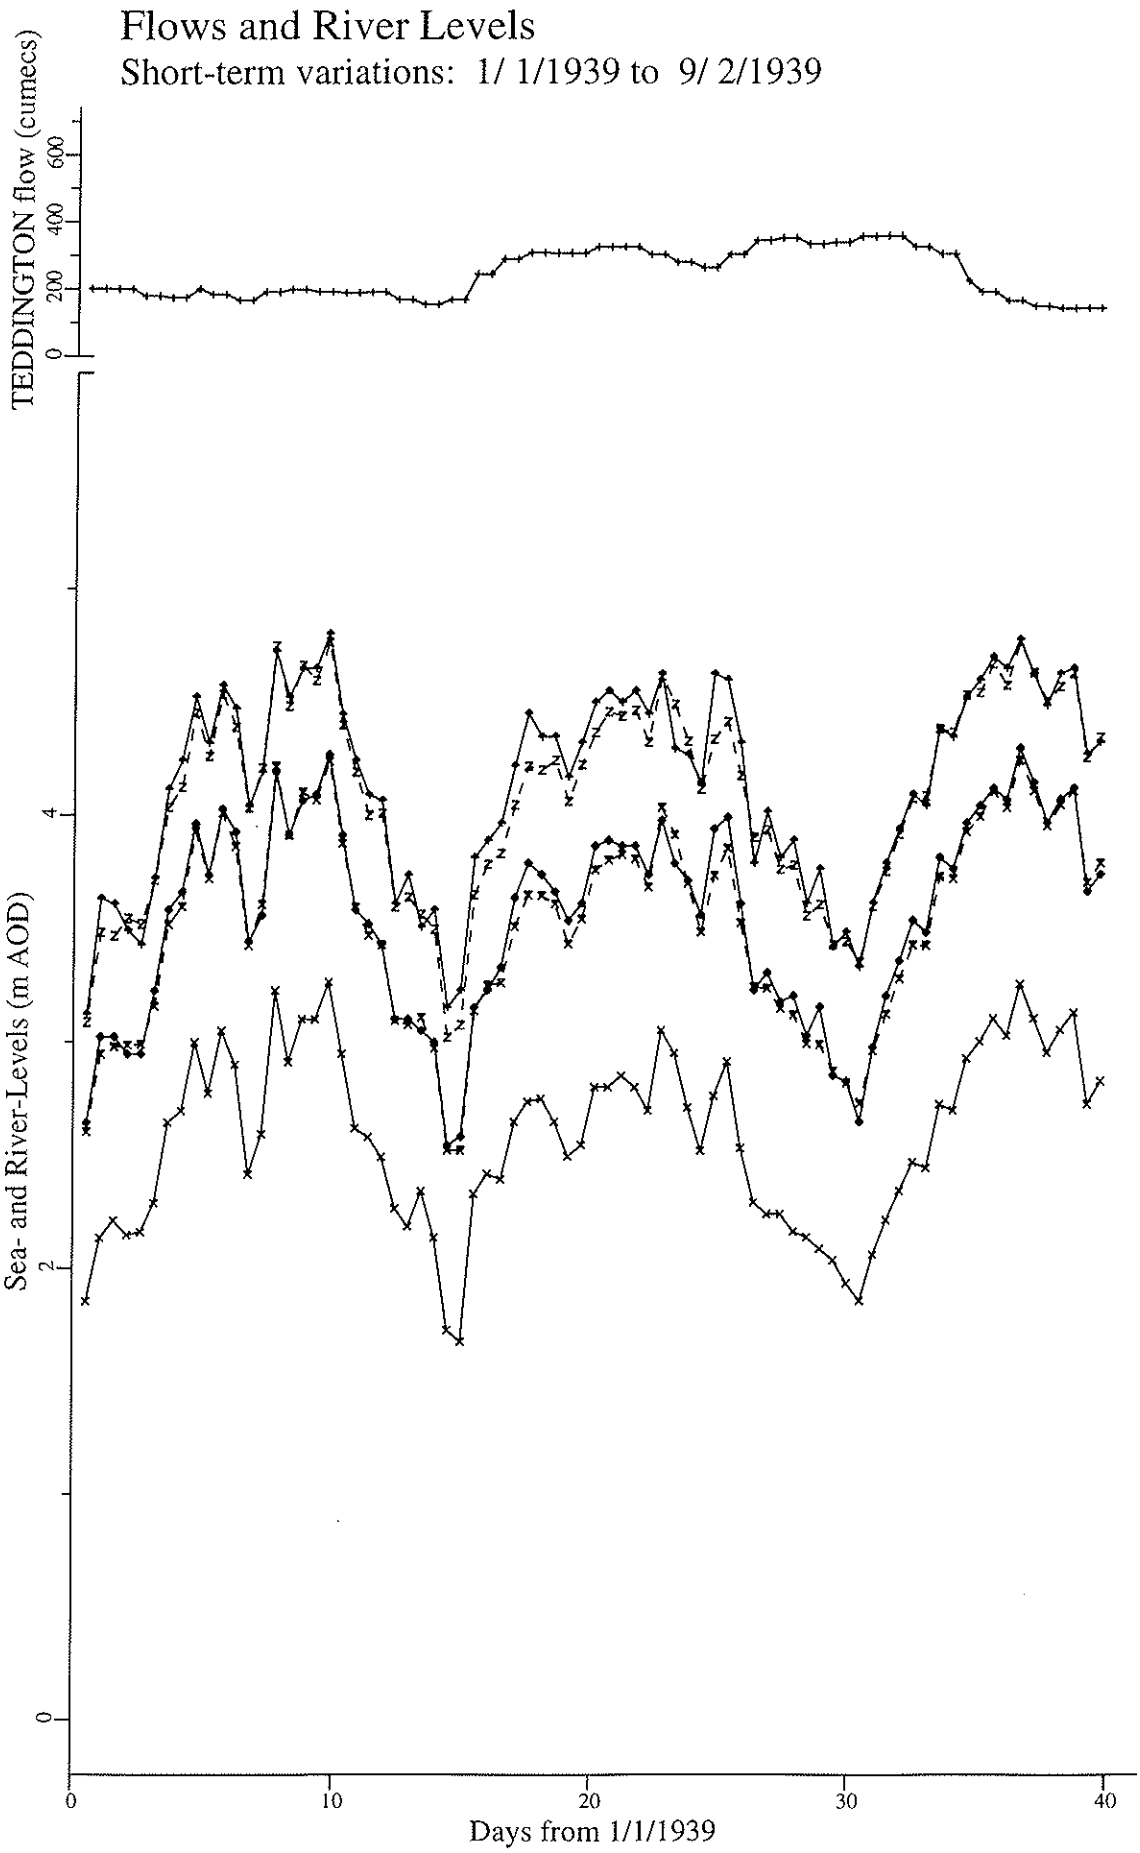

Figure 4.8 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics, flow-effects, surge-residuals and interactions (Model 2 of Table 4.1).

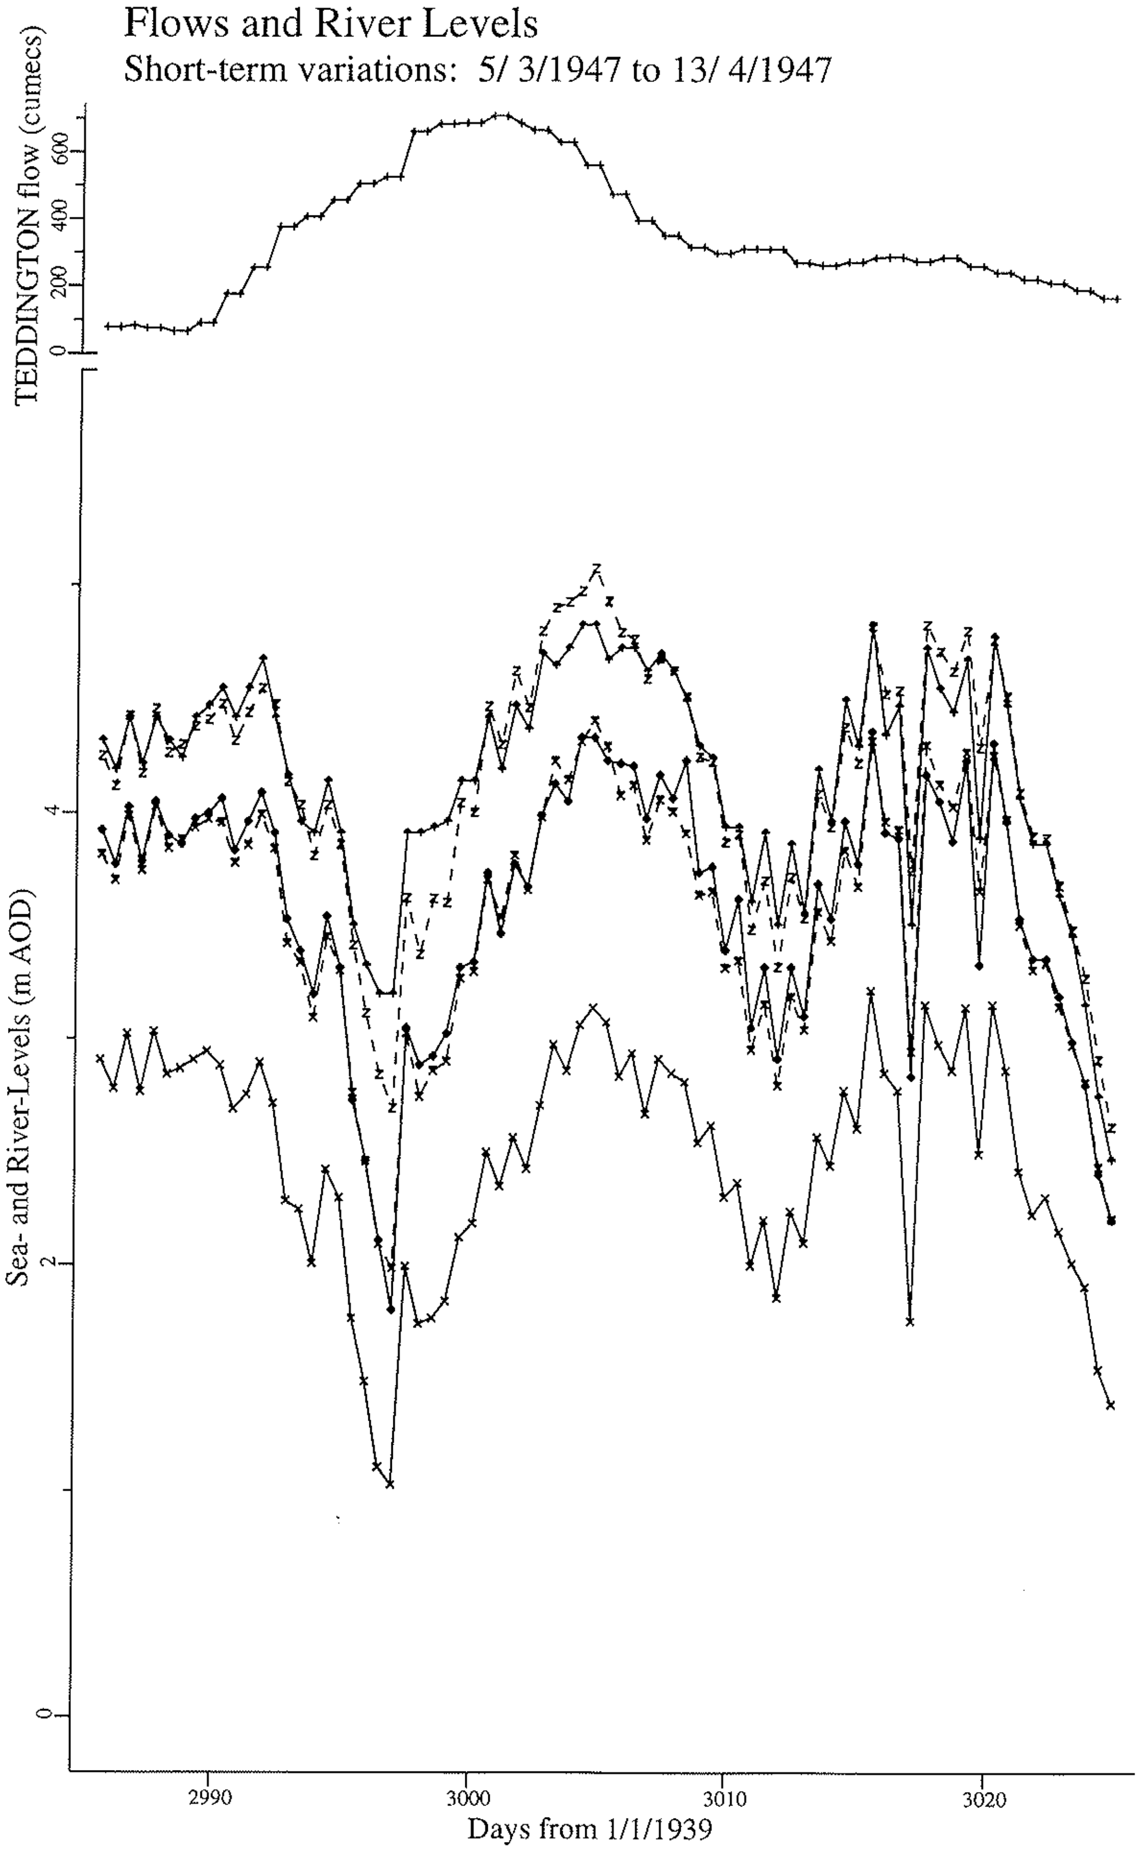

Figure 4.9 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics, flow-effects, surge-residuals and interactions (Model 2 of Table 4.1).

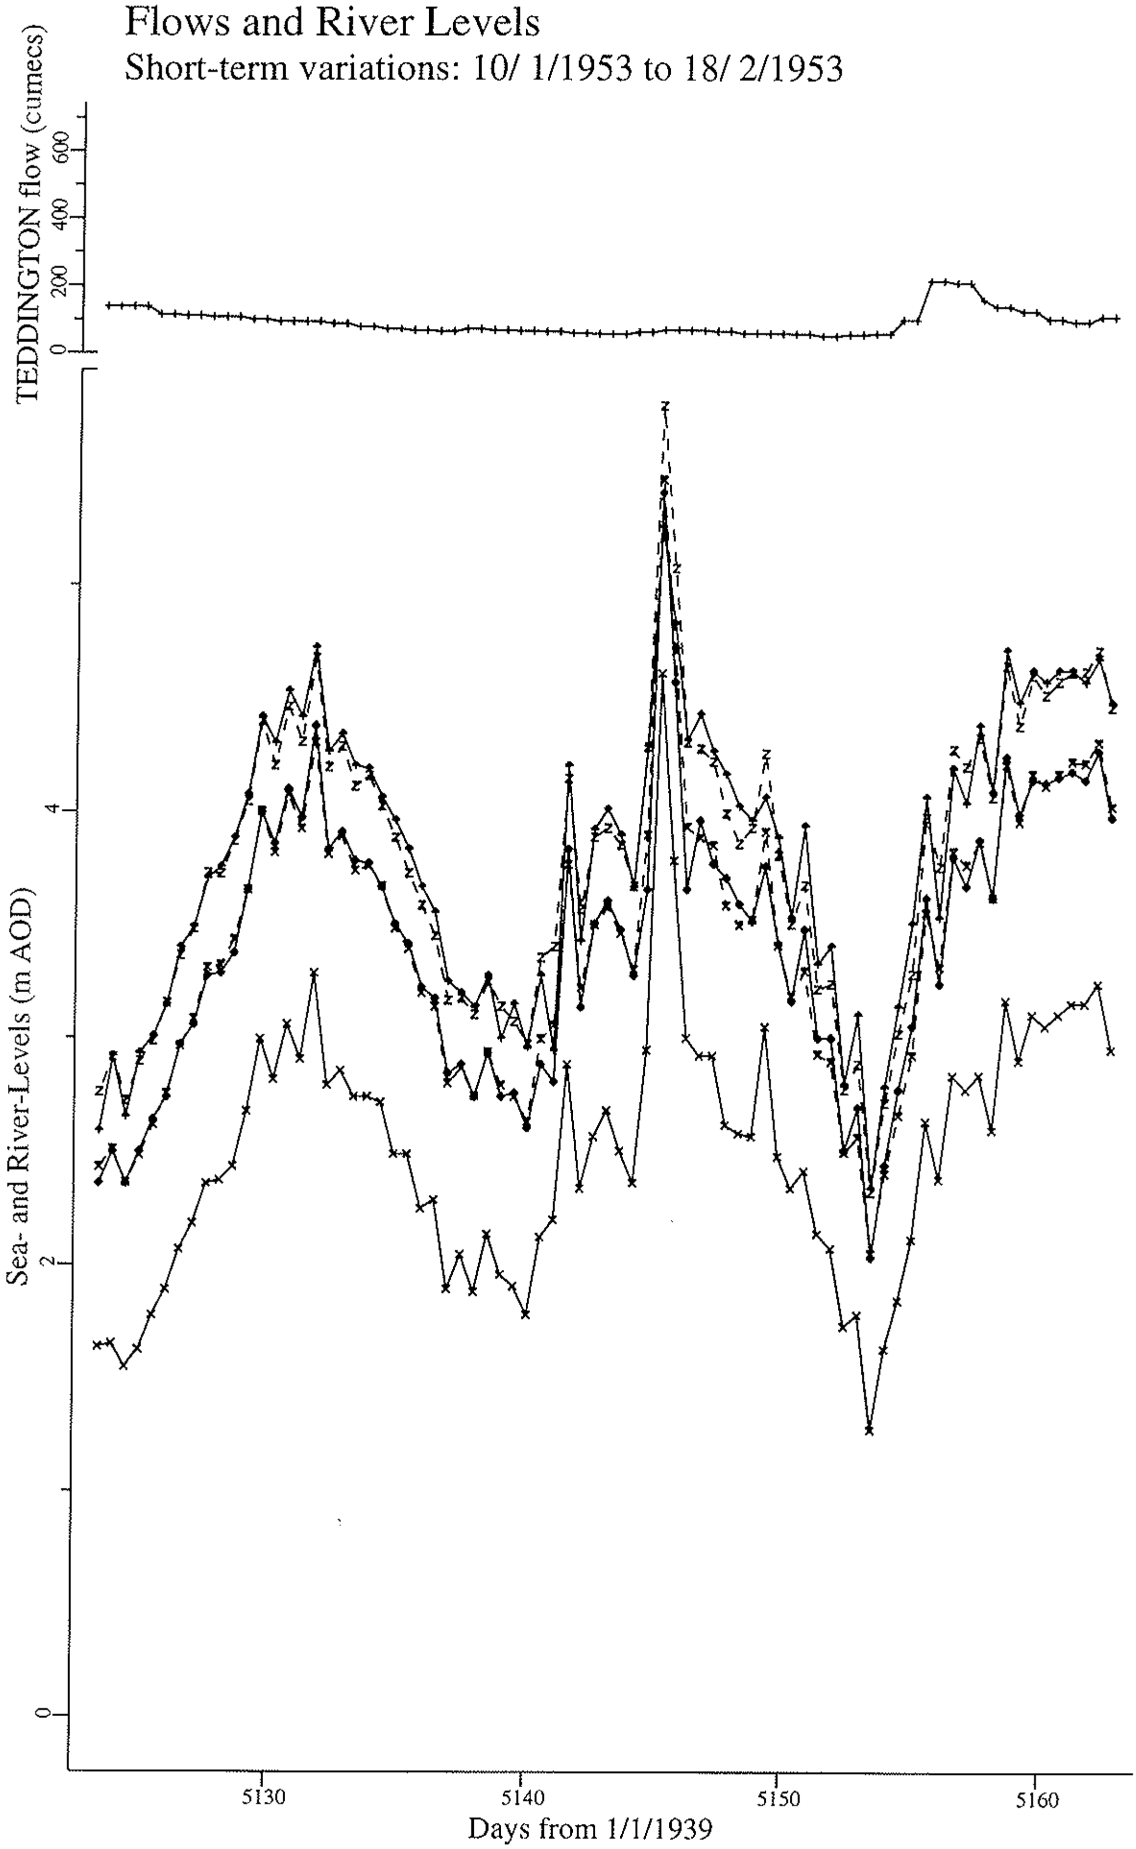

Figure 4.10 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics, flow-effects, surge-residuals and interactions (Model 2 of Table 4.1).

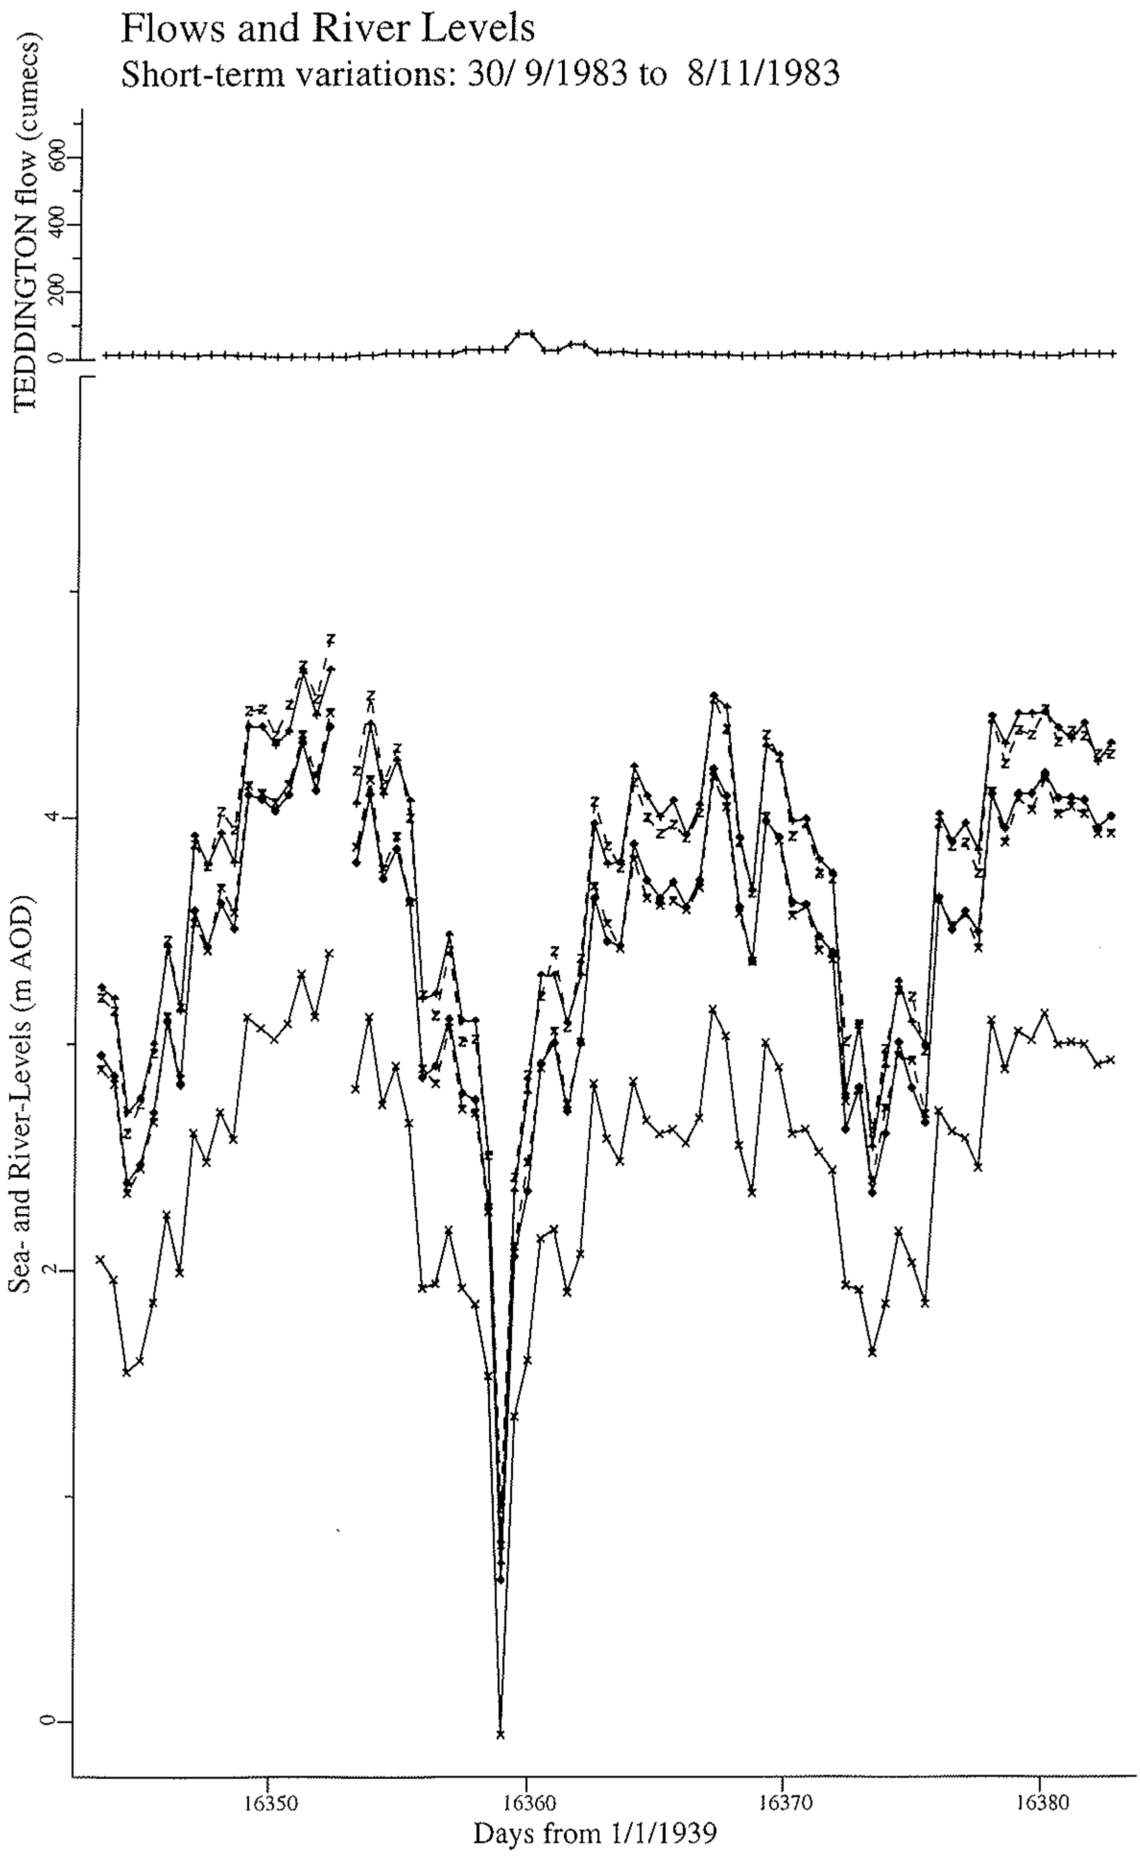

Figure 4.11 Examples of the estimated river-levels from the models based on seasonal and tidal-harmonics, flow-effects, surge-residuals and interactions (Model 2 of Table 4.1).

5. Conclusions

The models in Chapters 3 and 4 have examined the empirical relationships between tidal-peak water-levels at two relatively upstream sites on the Thames Estuary and corresponding tidal-peak values for the sea-level at Southend and river-flows at Teddington. Discussions of the sizes of the effects found have been given earlier. However, given that the principle topic of this report is joint probability analysis, it seems necessary to discuss the conclusions in this context.

It is clear from the outset that the types of empirical analyses of historical data undertaken here cannot give detailed information about the response of the estuary as a whole to the types of conditions of most interest since, by definition these are rarely occurring events and thus there will be little, if any, information in the historical data-set about these conditions. At best, the information obtainable is in the form of over-simplified estimators which one may hope will show the correct tendencies over the set of conditions experienced during the historical period and which, in the absence of anything better, might be extrapolated to a limited extent. In the context here, there is always "something better" which is at least potentially available. In particular, a well-configured and validated numerical hydrodynamic model could be made the basis of explorations of the response of an estuary to a range of conditions. One role of empirical data analysis is then to provide both a check on the performance of a structure-function derived from a complicated model and a check on whether potentially important effects have been missed in defining the variables for the structure function. For example, the analyses of Chapter 3 suggested that, starting from structure functions based on sea-levels at Southend and Teddington flows, improved structure functions should be obtainable by making use of both the astronomical predictions and the actual sea-levels at Southend.

An analysis of the errors experienced in the estimates found by the use of a structure function over a reasonably long historical period could be made the basis of an assessment of the likely size of future errors. In the case of a data-set such as that for the Thames Estuary used in this report, in which there has been significant changes in flood defences over time, there would be a number of difficulties in justifying this in a strict sense. Clearly the observed data would not strictly relate to the conditions assumed within any one structure function and it would not be feasible to define historical changes to the estuary, river channel and flood defences. The inclusion of trend terms in the model, as in Equation (3.1), can be expected to compensate for these changes to only a limited extent. Nonetheless, the empirical analysis of error will probably remain the only source of a nearly realistic assessment of how well a structure function works. The obvious alternative, that of driving a numerical hydrodynamic model with a variety of different inputs corresponding to the same nominal structure-function variables, has the difficulty that a realistic amount of variation might not be incorporated into such experiments. In particular, there may be additional sources of variation, such as other sources of flow into the river channel and the effects of wind speed and direction.

However, the empirical analyses do provide an important source of information about other aspects of the joint probability problem. In particular, the analysis described in Section 4.3 relates to the question of the statistical dependence of the terms in the basic decomposition of the series into trend, seasonal and tidal harmonics, flow-effects and residuals or "surges". If a joint probability model were to be constructed based on such a decomposition it would be essential that this model should reflect the statistical dependence actually occurring, and the historical data-set is the prime source of such information. The results of the analyses in Sections 4.2 and 4.3 are to some extent surprising in that they reveal a strong dependence of both the sea-level at Southend and the identified surge component at Southend on the flows at Teddington, which would ordinarily be thought to have no influence so far downstream in the estuary. Of course there is some doubt about this conclusion, firstly in terms of the validity of the simple analyses in the face of heterogeneous residuals and secondly because the evidence is for a joint dependence between the quantities rather than for a causal relationship. However, besides the possibility of modelling the dependence as found in these analyses, a number of ways of proceeding further remain. Firstly, one could seek some other nearby site to provide sea-levels which would be unaffected by river-flows and hope to build a model based on identifying a "surge" component at this site. Secondly one could seek other sets of data which might explain some of the residual variation: such data-series might be river-flows for other fluvial contributions to the estuary, or wind or pressure data. Thus one might hope to include, explicitly as extensions of the existing models, certain elements presently incorporated into the "surge" component. It might well be that such extra information would both "improve" the model for Southend in terms of having a statistical dependence structure more in accord with intuition, and enable better estimates of river-levels to be constructed for the upstream sites.

As well as the dependence between the surge-residuals for Southend and the other explanatory variables, revealed by the analysis of Section 4.3, another important aspect of statistical dependence is the serial dependence revealed by the autocorrelation functions of the residuals from the various models. For example, the topmost plot in Figure 4.5 reveals a moderately high correlation in the surge-residuals for Southend that extends over 2 or 3 days, while Figures 4.6 and 4.7 indicates that the serial correlations of the residuals for estimates taking account of the surge at Southend are high for 10 to 20 days. It may be that some of this serial correlation can be removed by taking account of additional explanatory variables as discussed above. Unless an explanation via additional variates is found, if a joint probability modelling approach were to be adopted to analysing extreme river-levels these types of serial dependence would need to be incorporated into both the models used and the analytical techniques.

The analyses related to dependence in this report indicate that any joint probability modelling procedures would need to take into account dependencies between all variates, including those involved in any decomposition into tidal harmonics and residuals. Perhaps the one note of hope is that only a weak dependence was found in Section 4.3 between the tidal harmonics and the size of the surge-residual for Southend. It might well be worth exploring the statistical significance of this effect with more rigour than adopted here, since if such dependence could be discounted, or if one were prepared to ignore an effect of such a small size, various simplifications of the modelling procedure would emerge. Within the present project, there has not been time to undertake a similar examination of the dependence of the errors associated with the estimates for upstream sites which make use of the surge at Southend, but there is some hope for a similar weak dependence. If dependence on the harmonic components could be ignored, there would still be major problems in joint probability modelling, but an historical emulation approach could be adopted which would effectively allow analysis of the results of new combinations of conditions, without requiring detailed dependence modelling. While it would be possible to implement this using the types of empirical models derived here, it would be better to use a separately-derived structure function and, since the notation for this is simpler, this case will be assumed in the following brief description.

The steps in a simple analysis based on historical emulation would be as follows, taking Richmond as an example. Analyses would be made of the historical data to achieve decompositions of the following kind, where this means that specific numerical values for the elements of the various series would be found.

$$S_{obs}(j) = S_{trend}(j) + S_{harm}(j) + \boldsymbol{S_{flow}(j)} + \boldsymbol{S_{surge}(j)},$$
$$R_{obs}(j) = R_{trend(2)}(j) + g_R\{S_{harm}(j),\ Q_j,\ S_{surge}(j)\} + \boldsymbol{R_{error}(j)}.$$

Here, terms written in bold italics denote data-series to be derived from a given equation. The term $g_R(.,.,.)$ denotes the structure function based on the tidal components, river-flow and surge. At this stage, the data-series $\{S_{flow}(j)\}$, $\{S_{surge}(j)\}$, $\{Q_j\}$ and $\{R_{error}(j)\}$ are available to be passed on to the subsequent stage of the analysis. In order to take account of possible statistical dependence between this set of quantities, these series are kept in the same order in time internally and new data-series of values of levels at Southend and Richmond are constructed according to these equations:

$$S_{new}(j) = S_{trend(3)}(j) + S_{harm}(j) + \boldsymbol{S_{flow}(j)} + \boldsymbol{S_{surge}(j)},$$
$$R_{new}(j) = R_{trend(3)}(j) + g_R\{S_{harm}(j),\ \boldsymbol{Q_j},\ \boldsymbol{S_{surge}(j)}\} + \boldsymbol{R_{error}(j)},$$

where now the terms in bold italics denote values retained from the first stage. Clearly values for the other series need to be defined. In the case of the trend terms, $S_{trend(3)}(j)$ and $R_{trend(3)}(j)$, these would be set to constant values such as to reflect the conditions to be simulated. Two approaches to the treatment of the harmonic term, $S_{harm}(j)$, could be considered. Firstly, the functional form for this component, in terms of sines and cosines, could be used to generate values for some future time-period. One disadvantage of this is that certain individual harmonic terms, or combinations of these terms, are not well estimated from historical data and thus unrealistic values might be generated by what is essentially an extrapolation procedure. In particular, there is no clear distinction between trend terms and long-period harmonics. Hence it might be advisable to adopt a second procedure in which new harmonic sequences are created by randomly selecting yearly blocks of time from within the historical data-period: then the harmonic terms can either be recreated from the functional form or the values could be copied from the first stage of the analysis. The randomisation might be implemented as a random permutation of the years within the historical period.

The above procedure provides a way in which new series of levels can be created by combining the observed surges with different tidal components than in the original series. Note that the outline given here would need to developed further to cope with missing periods of data in the historical data-set. The reason for creating the new series for levels at Southend, $S_{new}(j)$, is that these would be used to provide the basis, together with Q_j, for deciding for each high-tide whether or not the Thames Barrier would have closed according to whatever closure-rule is being assumed. In the case of barrier-closure, the corresponding level for Richmond would be replaced by some default value. One suggested procedure would be to extract annual maxima of river-levels from the newly created data-set, or from several such data-sets derived from alternative randomisations of years and to combine these with a corresponding set for the historical period: the latter would be essentially just the historical record of levels, but adjusted to take account of barrier closures. An analysis of the combined data set could begin with simple graphical (level against reduced variate plots) treating the overall collection as if they represented independent data-points, although of course they are not statistically independent.

It should be noted that, in the above procedure, the historical values for the series $\{S_{flow}(j)\}$, $\{S_{surge}(j)\}$, $\{Q_j\}$ and $\{R_{error}(j)\}$ are retained in contemporaneous order and thus any between- or within-series dependencies are implicitly retained. Furthermore, since the time-within-a-year for any tide is unchanged any seasonal dependence, such as that found in the size of the surge-residuals, would also be preserved. The procedure makes use of the assumed independence of all the other series from the tidal harmonics by implementing randomisation with respect to this component.

Appendix A Pre-whitening for Regressions

Many of the standard results in statistical theory for regression analysis depend upon the validity of certain basic assumptions. In particular, results about the variability with which parameters are estimated depend upon the assumptions that the residuals of the model are not correlated between observations, and that the residuals arise from random variables which have equal variances.

When formal tests of statistical hypotheses about parameters are made, the further assumption about the distribution of the residuals is usually made, namely that they are Normally distributed: however, the tests made for this report were of an informal type (comparing the parameter-estimates with their estimated variances), and hence this assumption is not considered further. When the assumptions of uncorrelated and equal-variance residuals fail, there are two consequences. The first consequence is that better estimates of the parameters could have been obtained rather than using the ordinary least-squares estimators: however the ordinary estimates still "work", in the sense of being consistent, unbiased estimates of the model parameters. The second consequence is that the estimates provided for the variability of the parameter estimates by the standard approach will be incorrect: the variances will be underestimated if, in some overall sense, there is positive dependence among the residuals.

Within the present report, the analyses reported in detail are all those provided by ordinary least squares, even though strong evidence has been found that the residuals do not have the uncorrelated, equal-variance properties required for ordinary least-squares to provide "optimal" parameter estimates. To a certain extent this choice was made on the basis that the main aim is to provide a preliminary assessment of various aspects of the relationships between tidal-peak levels and flows. Using ordinary least squares has the advantage that it is not dependent on having a "correct" model for the correlation and variance properties of the residuals, while still providing valid estimates of model-parameters. To a certain extent, the tasks of making formal and informal statistical tests for whether an improved model is created by adding extra parameters can be replaced by the more pragmatic approach of asking whether the apparent improvement in model performance is large enough on an ordinary meaningful scale. While somewhat subjective, this pragmatic approach means that internally-generated estimates of parameter-variability are not required from the least-squares procedure.

However, in practice, the above pragmatic approach to deciding model-structure needs to be backed-up by some at least semi-formal assessments of parameter variability. In this context the estimates of parameter variability provided by ordinary least squares can be used, with care, as a way of indicating the relative accuracies with which various parameters are estimated. When a somewhat more formal assessment of uncertainty is required, a simple technique called "pre-whitening" has been used: the effect of this technique is to take account of some aspects of the correlation of the residuals in proving the uncertainty assessment. Suppose that the model for predicting the target variable y_t is to use a linear combination of the vector of explanatory variables x_t: then the model can be written as

$$y_t = \beta^T x_t + \varepsilon_t. \quad (A.1)$$

Suppose that the residuals $\{\varepsilon_t\}$ are serially correlated, but that a simple model for the residual series of the following type can be found. That is, suppose that the new quantities $\{\eta_t\}$ are uncorrelated and have equal variance, where the new quantities are defined by

$$\eta_t = (\varepsilon_t - a_1 \varepsilon_{t-1} - a_2 \varepsilon_{t-2} - ......)/s_t , \quad (A.2)$$

and where $\{a_j\}$ and $\{s_t\}$ have known values. Then, it is possible to define a new set of target variables, $\{y_t^*\}$, and a new set of explanatory variables, $\{x_t^*\}$, by

$$y_t^* = (y_t - a_1 y_{t-1} - a_2 y_{t-2} - ......)/s_t , \quad (A.3)$$

$$x_t^* = (x_t - a_1 x_{t-1} - a_2 x_{t-2} - ......)/s_t . \quad (A.4)$$

It is than clear that, by construction,

$$y_t^* = \beta^T x_t^* + \eta_t, \quad (A.5)$$

which has the form of a least-squares regression in which the residuals $\{\eta_t\}$ are uncorrelated: moreover the parameter-vector β is identical in the two models in Equations (A.1) and (A.5). Depending on the number of coefficients $\{a_j\}$ included in the model of Equation (A.2), a few terms from the beginning of the series will need to be dropped, together with further terms affected by missing data. Overall, this approach is a nearly fully efficient way of estimating parameters for regressions with correlated residuals. Since the new model in Equation (A.5) has uncorrelated equal-variance residuals, applying ordinary least-squares procedures to the new problem will yield valid results. The operations entailed by Equations (A.3) and (A.4) can be described as "pre-whitening".

When applied in this report, the pre-whitening approach has only been used to remove the correlation in the residuals, not to remove the effects of changing variances in the residuals: rather more extensive modelling effort would be required to implement this. The effect of ignoring changes of variance in the residuals should be broadly neutral since overall the variances are averaged-out, unless there is strong dependence of the residual-variances on the explanatory variables.

While superficially attractive and seeming to provide an easy-to-apply method for overcoming the problems with non-standard least squares analyses, the pre-whitening approach is in fact not as immediately applicable as might at first appear. The essence of the approach lies in Equation (A.2) in which it is assumed that all the relevant structure in the residuals can be removed by such a model, or by some type of extension of this model. If such an approach is taken, the results for the regression-analysis will depend on this (partly hidden) model, and thus on its correctness and, strictly speaking, the uncertainty in the estimates of the parameters of the underlying model should be taken into account in the overall model. Thus the approach is probably only suitable as part of a procedure for exploring possible model structures, not for formal analyses. Furthermore, extensive investigations are likely to be required in order to arrive at an appropriate "pre-whitening" step. Thus, besides the question of formulating a model for the residual variances, $\{s_t^2\}$ in the models above, there would be the need to examine the possibility of non-constant serial correlations. For example, in the present study, the serial correlations and the variances of the residuals might both vary seasonally, with the amount of the river-flow entering the estuary, or with the underlying tidal conditions.