rijkswaterstaat
dienst verkeerskunde

cambridge systematics europe b.v.

REPORT: 7

PART: 1

TRAVEL DEMAND MODELS

CAMBRIDGE SYSTEMATICS EUROPE b.v.

ZUIDVLEUGEL STUDY

REPORT 7

TRAVEL DEMAND MODELS

CONTENTS

PAGE

1. INTRODUCTION AND OVERVIEW

This is the seventh report of the Zuidvleugel study. This study was commissioned by the Dienst Verkeerskunde (Traffic Studies Office) of the Rijkswaterstaat division of the Ministry of Transport and Public Works. It was carried out by Cambridge Systematics Europe b.v. in conjunction with staff of the Dienst Verkeerskunde. The aim of the study was to provide a travel demand forecasting capability for the Zuidvleugel area in the West of the Netherlands.

This report describes the methods that were used to estimate the demand models and presents the models themselves. Previous reports in this series have covered:

1- the design proposed for the demand models;
2- the computer software used for forecasting;
3- the computer software used for estimating the models;
4- the sample design for the home interview survey;
5- preliminary analysis of the home interview survey;
6- expansion of the home interview survey.

Report 4 is written in Dutch; the other reports are in English. Report 8 and 9 produced simultaneously with this report cover the implementation of the models for forecasting and the addition of car ownership and driving licence models to the system. Reports 10 and 11 (also in Dutch) describe applications of the system.

This report is organised in four main Chapters with an initial Summary. Chapter 2 describes some relevant aspects of the planning context. Chapter 3 gives some information about the data on which the estimation was based. The fourth chapter discusses a number of issues in model estimation that had to be resolved before the actual estimations could be made. Finally, Chapter 5 presents the models themselves: this is the main part of the report, to which the other chapters are preparatory.

The model system described in Chapter 5 is large and complicated. For this reason, the following section gives an overview of the system and the main ideas that were important in its development.

1.1 INTRODUCTION.

This chapter gives a broad overview of the development of the travel demand models for a study in the Zuidvleugel ("South Wing") of the Randstad conurbation in the west of the Netherlands. The study was carried out by the Dienst Verkeerskunde of the Rijkswaterstaat, the Traffic and Transport Engineering Division of the Ministry of Transport and Public Works. Several other agencies also participated in the study. In particular D.H.V. Raadgevend Ingenieursbureau and the Nederlandse Stichting voor Statistiek together were responsible for the home interview survey, for which advice was also received from the Centraal Bureau voor de Statistiek. The model development on which this report concentrates was the responsibility of Cambridge Systematics.

The work of the study was done in the years 1977-1980. In 1977 the concepts and planning of the study were worked out and towards the end of that year the home interview survey was conducted. Delays were caused by substantial coding errors found in the survey and by difficulties in computer processing of the large networks of the study area. Therefore, it was not until September 1979 that model development could begin, this stage lasting until June 1980. Software was then prepared to apply and validate the models at both aggregate and disaggregate levels. Policy analysis using the models began in 1981.

The objectives of the study are described in the following section. The third section briefly outlines the reasons for the choice of the disaggregate modelling methodology employed. Section four lists the data sources used to support this methodology and section five the innovative representation of travel by "tours" instead of the more common "trip" representation. Sections six and seven respectively describe the model structures developed and the methods adopted for estimating the unknown parameters. Section eight gives an outline of the aggregation process. The final sections describe the implementation of the models for forecasting and the validation that was possible.

1.2 CONTEXT AND AIMS OF THE STUDY

Current model development for the Dienst Verkeerskunde is specifically designed to develop policy analysis methods on the one hand for forecasting traffic flows for strategic studies and on the other hand for predicting the transport effects of general policy measures (i.e. policies less specific to location). In the first case, where flows must be forecast depending on alternative transport or land-use policies (e.g. construction of new roads, improvement of existing roads and land-use development plans) or on differing secular develpments (e.g. spatial and/or socio-economic "scenarios" for the future), forecasts (in terms of traffic assignments) discriminating clearly between such changes must be available. Methodology for such discriminating forecasts must inevitably distinguish between different travel purposes and travel modes. Note that the personal and political valuations of different purposes and modes may well be different. In the second case, the general policy measures under consideration must be able to be expressed in variables incorporated in the predictive models. Examples of such variables are the changes in the distribution of population or employment or variables directly related to transport such as travel times, overall parking restraints, public transport pricing measures or petrol tax changes. In these cases an assignment is not normally required and more important is the speed with which forecasts of the effects of policy with respect to mode choice, journey length, frequency, etc. can be evaluated. Both types of model are required for application in the Structuurschema Verkeer en Vervoer (National Long-Term Transporatation Plan) and in the Meerjarenplan Personenvervoer (Five Year Plan for Personal Travel). For the long term Plan forecasts have to be made for the whole country, and the overall model system for this plan can be expected to be rather coarse in its zoning and networks. This national model can then be completed by the addition of regional models (such as the Zuidvleugel model).

Because of this complex context, the Zuidvleugel Study contains a wide range of methodological and practical developments that are to some extent diverse but nevertheless inter-related as they serve diverse but inter-related objectives. These objectives can be summarised under two main groups. First, it was necessary to develop models forecasting travel demand specifically for the Zuidvleugel area. This need arose because of specific issues within the area, and in particular issues caused by migration from the Rotterdam areas to Noord Brabant. The limited capacity of the bridges and tunnels crossing the rivers separating those areas played an important rôle here. Second, the Dienst Verkeerskunde aimed in this study to extend its modelling capabilities and to enlarge the range of methodologies at its disposal. The Zuidvleugel models were required to be able to forecast both at a level suitable for evaluating infrastructure investment or management policies in a region and at a level suitable for global policy analysis. These two rather different requirements for policy analysis required that the methodology adopted for the study be particularly flexible.

1.3 CHOICE OF METHODOLOGY

A range of alternatives were considered before it was decided to base the demand models for this study on disaggregate analysis. The context and aims of the study described in the previous section give the most important considerations behind this decision.

There were already in existence regional models based on data collected many years earlier (and in some cases in other countries). The development of a new model for a new area suggested an opportunity for estimation using up-to-date local data and using the best available methodology. The extent to which "state-of-the-art" models could be applied was of course limited by the need to produce a reliable working model without undertaking extensive fundamental research. Essentially this restricted consideration to models of the logit form. The properties of this model are well-known (and are discussed again in Section 1.6).

The application of disaggregate methodology is particularly suitable for the tasks required in this study. Disaggregate models can be used, as in a number of previous studies, as the basis for zonal forecasting models of the classical type. Such models are suitable for preparing vehicle trip tables for assignment. Disaggregate models can also be used, through the technique of sample enumeration, to make more general forecasts for the whole area. Thus disaggregate modelling can be used to achieve both of the major modelling tasks of the study.

Given the diverse objectives of the study and the need to develop models responsive to many different policies over a wide area, it is inevitable that a large-scale model would be required. This large scale of development is unusual in disaggregate modelling, which is normally characterised by cheapness and simplicity since a model can be developed specifically for a single policy issue. In this case, however, so many issues were to be addressed that comprehensive model system was essential.

1.4 DATA

The data from which the models were estimated came from three main sources. The home interview survey has already been mentioned; this survey was supported by transport networks and a detailed set of data describing the characteristics of the zones into which the study area was divided.

The study area was defined contain nearly all of the province of Zuid Holland and the western part of Noord Brabant. This area contains the cities of The Hague and Rotterdam and then to the south as far as the Belgian frontier. About three million people, nearly a quarter of the total population of the country, live in the area, which was defined to minimise the number of boundary crossings - that is, to make the area as much as possible self-contained.

The home interview survey was administered to a very carefully selected sample of approximately 3000 households living in this study area. The sample was selected to over-represent households of a particular interest to the study: long-distance commuters, car owners and public transport users. The methods used to select this sample are described in Report 4 of the study. Special methods were devised for the expansion of this data to allow not only the estimation of study area total values for the various statistics required, but also confidence limits for those estimates. The survey itself was primarily a one-day travel diary, for all household members over four years old. In addition, a wealth of socio-economic data was collected describing each member of the household and the vehicles owned. Data was also collected of a five-year history of the location of the household itself and of the workplaces of its members: this was used in a separate study.

Separate transport networks were prepared for the highway network, for public transport and for "slow" modes (walk, cycle and moped). For the highway networks, paths through the network were built minimising a conventional generalised cost measure incorporating time and cost. Distance, time and tolls were measured along these paths for input to the model estimation. For the public transport network, a generalised cost measure was defined incorporating the various travel impedances (walking time, waiting time, etc.) and paths were built minimising this measure: the various components were measured along these paths. A special program was written to calculate public transport fares, in an attempt (which proved successful) to obtain good estimates of elasticities with respect to fares by measuring the fares accurately. For both highway and public transport modes, paths were built separately for morning peak, evening peak and off-peak conditions. For slow modes, however, only one set of paths was built, minimising distance, and distance was the sole measure for this mode input into model estimations.

The networks were constructed at two levels of detail. The locations of the households surveyed were coded to an accuracy of 500 metres, and for connections to the households this level of detail was employed. In other cases, however, locations were used only at a zonal level of accuracy, so that the networks also had to be properly defined for this level of detail. In the event, the "home" end of trips was defined at 500 metres accuracy; the "other" end was defined only at zonal accuracy.

Severe problems arose in network processing, which caused a delay in the schedule of the study. These problems arose partly from the large size of the study area and the complexity of the networks inside it and partly from the need to transfer software to the new computer being used for the study. In the end, however, these problems were overcome.

The final important data set was collection of characteristics for each of the 319 internal and 77 external zones defined for the study. Such characteristics as employment and population were broken down into classifications of employment types and age groups in the population. Land use data was included, as was vehicle ownership, but data on driving licences was not available at this level.

1.5. THE UNIT OF TRAVEL: TOURS.

The purpose of the models being developed was to forecast travel for policy analyses which usually require the forecasts in terms of person-trips. For example, perhaps the most important policy analyses for this study are the traffic assignments which inevitably require as input vehicle-trips, or equivalently, driver-trips. Similarly, nearly all other analyses can exploit forecasts in the form of trips. The familiar four-stage Urban Transportation Planning models therefore use trips as their fundamental unit for measuring travel.

However, if consideration is given to the reason leading people to travel, the adoption of trips as a basic unit appears less satisfactory. The key point recognised by so many writers on the behavioural basis of travel is that, in the overwhelming majority of cases, travel is a derived demand. People do not usually travel for the fun of it, but because their daily schedule requires them to perform some activity (work, shopping, etc.) which can only or can best be done away from home. It is eminently reasonable to suppose that the ultimate approach to predicting travel is therefore first to make some prediction of the acivity patterns of each person within a household, contingent on all the constraints that each person faces, and then to predict what travel is necessary to permit activities outside the home. This approach is however not yet feasible, since activity models are not yet sufficiently developed to form a basis for reliable travel forecasting.

It is nevertheless possible to take some note of fundamental research on activity patterns in the design of practical models, and this was attempted in the Zuidvleugel study. The idea is to bring the models one step closer to our best understanding of the behavioural basis of travel. This idea had two practical consequences for the models themselves: first, that a larger number of purposes than usual were separately identified and independently modelled; and second, that the units of travel were changed from trips to tours. The first change is simple enough in principle, and clearly allows more insightful modelling of, for example, social and recreational travel by treating them separately. It also fits with the separation of purposes for policy reasons mentioned earlier. The second change, however, requires more justification and explanation.

A tour is defined to be all the travelling between a traveller's leaving home and his or her next return. When an activity is performed away from home, in the majority of tours (83%) the traveller goes directly to the required destination, performs the activity, and returns directly home, thus making two "trips", in the familiar definition. Ignoring for a moment the 17% of more complicated tours in which more than one destination is visited, a tour is therefore exactly the amount of travel necessary for a single non-home activity. It is reasonable to suppose that the traveller considers the total inconvenience of the tour in deciding whether, where or when to perform the activity and how the necessary travel (in both directions) should be performed. This direct connection between the characteristics of the tour and the traveller's decision concerning the activity means that we can expect more rational modelling at many points in the forecasting system.

An important example is the treatment of non-home-based travel (15% of the trips in this area). A non-home-based trip arises when more than one destination is visited on a tour, when the connections between these destinations form non-home-based trips. It is difficult to determine exactly why each trip has been necessary (i.e. what the activity was that caused an additional trip to be made) and to make reasonable links for mode choice and even destination choice between these trips and with the home-based part of the tour.

In a tour model, non-home-based travel arises in the form of "detours" (an apt word) from the simple there-and-back norm. Given a series of destinations on a tour, it is often reasonable to label one of them as the "primary" destination and to treat travel to other destinations as being conditional on travel to the primary destination. This detour approach to non-home-based travel is clearly a simplification, but it seems a constructive first step to a comprehensive model. In this study, the primary destination was taken to be the workplace, if one of the destinations on the tour was a workplace (the usual workplace taking priority); otherwise the destination at which most time was spent was defined as the primary destination.

In addition to these improvements gained by including the whole tour in modelling the choice between alternative modes, destinations, etc., and the improved treatment of non-home-based travel, the tour basis gave the Zuidvleugel models a further important improvement over trip models in preserving the linkages affecting the whole tour. For example, mode choice was found to be constrained to be the same for most tours (97%) since a traveller leaving home by car or bicycle normally takes the vehicle with him until he returns home, although switches between walking, public transport and car passenger are of course possible (1½%). Such constraints are very difficult to incorporate in trip models. Similarly time-of-day constraints (that, for example, a return trip must start a reasonable time after the outbound trip) are difficult to incorporate.

The discussion above summaries the rationale for the choice of tours as the basis for the Zuidvleugel models. More detailed information is given in Report 5 of the study and by Weisbrod & Daly*.

1.6 MODEL STRUCTURE

The scope of the modelling defined for this study was, as indicated, to predict frequency, destination and mode for personal travel for all purposes by the residents of the study area. The problem remained of determining the form and structure of model appropriate for these three decisions.

It was decided to use models of the "logit" form (with extensions) for all of the models within the system. The typical form of this model is well-known:

$$\log P_c = V_c - V_K = V_c - \log \sum_{k \in K} \exp V_k \qquad (1)$$

giving the log probability of choosing a single clement c from a set of K in terms of a measure of attractiveness V_k of each member of the set. The crucial assumptions on which this model is based are:

- the symmetry of the alternatives k within the set K;**

---

* G. Weisbrod and A.J. Daly, "The Primary Destination Tour Approach to Travel Demand Modelling", Cambridge Systematics Inc. (1979). (unpublished).

** This assumption of symmetry is often described as the "Independence from Irrelevant Alternatives" (IIA) property. In a utility-maximising approach the assumption is more accurately characterised as that of "Independent and Identically Distributed" (IID) unmeasured components.

- that the set K itself can be precisely defined;
- that the attractiveness functions V_k can be properly specified.

The first assumption has been the subject of much successful research in recent years (see, for example, Manski and Westin*). The simplest method developed in that research allows limited relaxation of the symmetry assumption by defining logit models of the form (1) over subsets and groupings of the fundamental choice set. This nested logit model was applied extensively in the Zuidvleugel Study, as will be described.

Other model forms (e.g. the "probit" or normal model) have been applied in a limited way in recent years, their main advantage is in dealing with failures of the first assumption on which the logit model is based. Since progress in removing the most restrictive aspects of this assumption is possible within the logit framework (as described) and since further progress with the solution of structural problems is perhaps not the most urgent improvement needed in the modelling, better treatment of structural problems was not considered to be worth the computational difficulties caused by such models.

The second assumption is finally receiving more attention from researchers. The question of what alternatives are actually open to the travellers and the parallel one of how we can predict these have led to the definition of models in which the set K is randomly variable between different travellers. These models, together with some commonsense determination of circumstances in which some alternatives are certainly not available to some travellers (e.g. car driving for children) allows a small amount of progress to be made, but there is much scope for further work on this point. For the Zuidvleugel Study, the sophisticated structures with random choice sets were not applied, but attempts were made to apply absolute restrictions wherever possible.

The third assumption is now also receiving more research attention. The variables included and the form in which they appear have been questioned and constructive experiments have been made. The state of the art is not however sufficient well advanced to allow reliable measurement of forecasting of such variables as habit. For the Zuidvleugel Study an important extension was made in the form of the attractiveness function, allowing a much better estimation of the parameters of destination choice models.

---

* C.F. Manski and D. Westin, Theoretical and Conceptual Developments in Demand Modelling, in "Behavioural Travel Modelling", Croom Helm, London, 1979.

The various decisions to be modelled, for each purpose,were:

- frequency, i.e. making zero, 1,2,3, etc. tours in the survey day;
- destination, which was represented at a zonal level as a choice among the 319 internal and 83 external zones;
- mode, in which eight modes were recognised (as shown in Table 1).

The reasons leading to the choice of the model structure actually implemented are described in Report 1 of the Study. In general these decisions were taken on the basis of a priori expectations about the proper structure that had to be verified by subsequent empirical tests. Some minor modifications were made to the structures on the basis of the empirical findings.

The general model structure is shown in Figure 1. In the figure, the linking of models by a solid line indicates that the lower model is conditional on the higher model (for example, that main mode choice is represented as conditional on destination choice). The dashed lines indicate the incorporation of a "logsum" variable from the lower model in the higher. This variable (defined by V^K in equation (1) is shown in the utility theory of the logit model to be the total expected utility from making a choice in a given choice set. It gives therefore a kind of summary of the attractiveness of the choices at a lower level that is approriate for inclusion in a model at a higher level. The utility theory further shows that the logsum inclusion makes the whole model consistent with a theory of utility maximisation. It must be noted, however, that features common to the lower level choice are not incorporated in a logsum variable*.

The structure involving frequency models is somewhat unusual. It was found, however, that the choice whether or not to make the first tour in a day for a given purpose was significantly different from the choice whether or not to make repeated tours for the same purpose. However, it did prove possible to model the choices of whether or not to make a further tour, when a given number (1,2,3...) had already been made, by the same model whatever the number of tours already made. This structure is a simple extension of an idea applied by Sheffi**. The stop/repeat model is conditional on the 0/1+ model in the sense that if 0 is chosen in the 0/1+ model, the stop/repeat models is irrelevant. In principle, a logsum connection should also be present in the opposite direction, but this was not found to be worth the additional complexity involved.

---

* K. Sobel, "Travel Demand Forecasting with the Nested Multionamial Logit Model". TRB 1980

** Y.Sheffi,"Estimating Choice Probabilities among Nested Alternatives" , Transportation Research, 1979.

Table 1: Modes Observed

Mode Group	Mode	Mean Tour Length (km)	Share%	Group Share%
Car	Driver	21.4	21.4	29.5
Pass.	household Car	22.1	6.3	
Pass.	other Car		1.8	
Public Transport	Bus/Tram	25.6	3.7	5.6
	Train/Metro		1.9	
"Slow"	Walk	3.7	32.3	63.1
	Bicycle		27.8	
	Moped		3.0	
Other	Various	--	1.8	1.8
Total	All	10.4	100.0	100.0

Figure 1: General Structure of Zuidvleugel Demand Models

(for tours of each purpose)

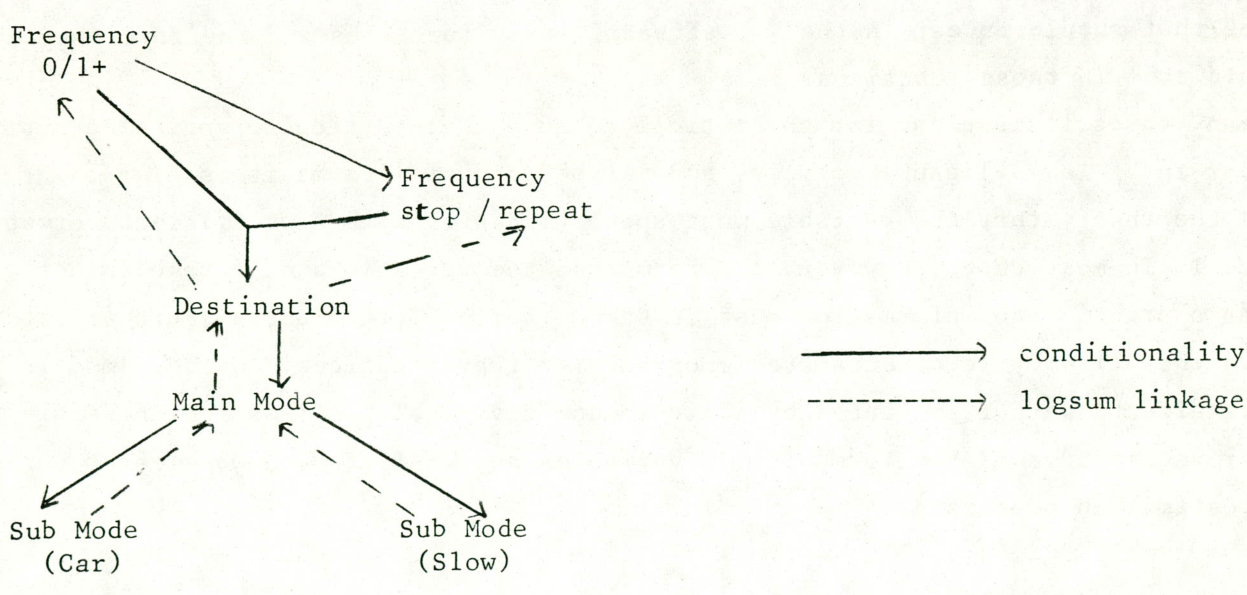

The destination model is conditional on both frequency models in the sense that as many destination choices are to be made as there are tours. For all purposes other than work and education, destination and mode were combined as a joint model. In principle, this need not be seen as a structural change, since the joint model is equivalent to a nested structure with a logsum connection if the logsum coefficient is constrained to 1.0. In practice, however, the change is really structural, since the joint model allows simultaneous estimation of mode and destination parameters. This simultaneous estimation has important advantages of statistical efficiency that are not otherwise available with the present computer software, and (for these tour purposes) it was adopted for that reason in this study.

The remainder of the structure is reasonabl straightforward, the unusually complicated mode structure reflecting the importance of a wider range of modes than is usual in industrialised countries.

1.7 ESTIMATION METHOD.

The structural consideration outlined in the previous section allowed the specification for each tour purpose of a series of logit models, interacting with each other through the use of logsum variables as described. The remaining task in developing fully specified models was to determine for each model the variables that should appear in the V_k attractiveness functions and the form they should take in those functions.
In many cases it is clear for theoretical reasons that particular variables must appear in V_k (e.g. logsums, prices) and in the case of logsums it is also clear from the theory that the variable must appear linearly with a coefficient between 0 and 1. In most cases, however, the form that the variable should take is not clear a priori and information must be taken from empirical tests. Further, it is known that life-cycle effects are important for travel choices: in these models these effects were represented by socio-economic variables. It was not possible to determine in advance the form of these variables, and in fact much was learnt during the estimation process.

It was therefore necessary to make a large number of trials of models with alternative forms of V_k functions. For this project software was developed to facilitate the estimation of large numbers of logit models and their detailed analysis and comparison with each other. The opportunity was taken to complete a standard suite of software for logit modelling, which is described in Report 2 of the Study.

The usual form of the V_k function in logit models is linear

$$V_k = x_{k1}\theta_1 + x_{k2}\theta_2 + \ldots + x_{kr}\theta_r \tag{2}$$

in which x are observed data and θ are unknown parameters to be estimated. This assumption of linearity in the parameters has been questioned a number of times. Note, however, that the x variables can be any possible transformation of observed data. To generalise the functional form beyond this point is normally not important, which is fortunate since the process of estimation becomes difficult for technical reasons if V_k is non-linear in θ. For this study, it was considered important to allow one particular form of non-linearity, and the programs have been amended to allow this extension, which can be of importance to other studies. The extension is to allow the incorporation within the logit model of multiple attraction variables. In previous studies it was necessary to estimate the parameters for these variables separately from those of separation variables, and it can be expected that biases can arise from such a procedure. The methods used for this extension of the methodology are described in a separate paper*.

In a study with so many zones (402), the estimation of a good destination choice model with all the zones present would be computationally infeasible. Fortunately it has been found that unbiased estimates can be made with a random sample of alternatives**. In this case, for efficiency, it was necessary to sample the "important" (i.e. nearby) alternatives more heavily than the more distant ones; this requires adjustments in the estimation process. Further sampling was done to reduce the number of slow mode tours appearing in the joint mode-destination choice estimations; again corrections must be implemented. These sampling procedures reduced the computational burden significantly.

1.8 AGGREGATION.

In the models presented in the previous section all the available information was used at the highest possible level of detail in an attempt to obtain the best possible model. The applicability of these models for forecasting is clearly limited by the availability of data at the appropriate level of detail. One

---

* A.J. Daly, "Estimating Choice Models Containing Attraction Variables", Transportation Research, 1982.

** D.McFadden, "Modelling the Choice of Residential Location", conference Spatial Interaction Theory and Planning Models Bastad, Sweden, 1977.

of the main applications of the model is, as mentioned, at an overall level in a sample enumeration procedure and there exists a suitable sample, with all the required data, in the home interview survey itself. For the other main application, to a classical zonal forecasting system, data restrictions play an important rôle.

The zonal data set that was available was unusually detailed but it was inevitable that some data items in the home interview were imperfectly represented, apart from the simple loss of accuracy due to averaging:

- some data (e.g. driving licences) was simply not available at zonal level;
- some data (e.g. car ownership) was available only in different form(total number of cars in zone, not information about the number of car-owning or multiple car-owning households).

For these reasons the form of the models had to be revised.

In moving from disaggregate to zonal level, there is of course a loss of information. This is particularly marked with variables such as age or sex which separate the population into groups with quite different behaviour but which have very little variation at a zonal level. This problem was reduced by introducing market segments into the zonal models. For computational reasons it is undesirable to introduce too many segments, and a limit of 20 was fixed.Within the 20 segments it was possible to distinguish age, sex, licence holding and car ownership, thus greatly improving the accuracy of the zonal models.

Two alternative approaches are available for obtaining zonal models. The approach more simple in concept is to replace the disaggregate variables with their nearest zonal equivalents and re-estimate each model. The alternative approach is to take the view that the coefficients obtained in the most disaggregate version of the model are the best that can be obtained: modification to the model in order to use zonal level variables should therefore preserve these coefficients as much as possible. The disaggregate coefficients can be preserved by calculating a composite variable for each alternative which approximates as closely as possible the disaggregate V_k function, but using the zonal variables. The model should then be re-estimated using the composite variables, when the parameters for the composite variable should be between 0 and 1, because a model with aggregate data has a smaller scale (or reduced sensitivity to explanatory) variables relative to the same model with disaggregate variables.

For this study a combination of these approaches was used as appeared appropriate in the circumstances of a particular model. In many cases the replacement of disaggregate variables by their zonal equivalents was not straightforward, and in such cases the use of a composite function was unsatisfactory. It could, however, be used for a part of the utility function. In other cases, however, the theoretically preferable approach of a composite function turned out to be practicable, and it was then adopted.

Thus an entire system of zonal models was developed in parallel to the fully disaggregate model system. It was decided not to develop zonal models for slow sub-mode choice, because these models mainly function with detailed socio-economic variables, and because nearly all slow mode travel is in any case intra-zonal. A total of 32 models was developed (c.f. 37 fully disaggregate models); these are reported in full in Chapter 5.

1.9 IMPLEMENTATION.

The sample enumeration and zonal model systems were implemented separately.

The sample enumeration system was implemented in a specially written program using the most disaggregate forms of the models. The objective of this program is to permit rapid evaluation of a range of policy options, and it has been designed accordingly to run quickly and easily on the computer system available. The sample enumeration program is linked with other programs allowing the user to specify tables for detailed analysis of the policy. The tables may be expanded to study-area totals, and the sampling error is given along with the value in each cell of the tables.

The zonal system requires for its implementation matrix manipulation on a large scale. A matrix management program "JMODEL" was written, giving in standard FORTRAN many of the capabilities of the well known matrix management program UMODEL, which unfortunately was not available on the Univac machine. The JMODEL program is fully documented in Report 3 of the Study.

The zonal model system uses JMODEL as its basis. It is designed to be "friendly" to the user, but it nevertheless requires the provision and manipulation of substantial data sets. JMODEL matrix management is instrumental in keeping control of all this data. Both implementation systems are fully documented in Report 8 of the Study.

1.10 VALIDATION.

The models developed in this study are intended for forecasting the effects of policy or secular developments at a future time. It is therefore not possible in principle to validate the essence of the model (responsiveness to changes) at the present time. Nevertheless it is possible to make tests of the reasonableness of the model's performance and a number of such tests were carried out.

The first series of tests concerned comparisons of model "predictions" for the base year (end 1977) with figures derived from the home interview survey. In principle such tests are simply checks that the models have been properly coded into the application programs and this is in fact an important check when dealing with complicated models. In practice there are several reasons why the estimated models could differ from home interview findings that could only be found by a subsequent check. For example, there could be:

- geographical or other variations in travel intensity not incorporated in the models;
- error filters (necessary for modelling) introducing biases;
- "other" modes (e.g. taxis and motorcycles) were excluded from the modelling;
- biases deliberately introduced into the home interview (which therefore appear in the models) and which are removed by expanding the home interview.

For these reasons it is always desirable to make a series of checks between model and home interview. In the case of a tour model such as developed in this study, comparisons are further necessary to deal with modal and geographical variation in non-home-based travel, which is not fully represented in the models.

The results of these checks were encouraging. For overall parameters, such as modal shares, tour lengths and frequencies, agreement was almost exact. For more local effects, there was more variation, but the model figures generally remained well within the confidence interval for the home interview survey. The conclusion was clear that the model reproduced rather closely the behaviour observed in the home interview.

Comparisons of model predictions with "independent" data proved more difficult. The aim was to make the main validation of the model at a screenline, the Hollands Diep, a major waterway to the south of Rotterdam, and roadside interviews and a train survey were held there. On analysis of this survey it became apparent, however, that the travel observed at the Hollands Diep was extremely atypical of the study area as a whole. The major difference was in the tour lengths: 95% of the Hollands Diep traffic was travelling more than 50 km, compared with 4% in the home interview. This naturally meant that much of the Hollands Diep traffic did not remain in the study area. Further problems arose with the definition of purposes and modes: nearly 10% of the Hollands Diep travellers were in a "car" holding 5 or more people. When these problems had been resolved, however, the difference between the model and the screenline survey was in the range of 15-20%. This difference is just what we should expect from a model based on a home interview survey, in which some travel is inevitably lost.

Because of these difficulties with the Hollands Diep survey, further comparisions were made on other screenlines around Rotterdam and later around Den Haag. Two principal difficulties arose here out of the fact that the screenline surveys had not been intended for this purpose but were part of routine surveys or of other studies. These problems were that it was difficult to separate out freight traffic, or, more importantly, external personal traffic from this data which was largely collected with automatic counters. No complete cordon survey was made for the Zuidvleugel Study, and although of course freight and external traffic must be forecast in addition to the travel covered by the models described here, these forecasts were not available at the time the validation was done. Thus the comparisions were somewhat approximate. The general indication, however, was that a larger fraction of traffic (20-30%) appeared to be missed in the home interview than was suggested by the Hollands Diep survey.

Because of these results, intensive investigations were made of the correctness of both the home interview and screenline surveys. A number of errors were corrected in the screenlines, but although local weaknesses were found in the home interview survey, comparisons with other home interviews showed it to be on the whole sound. It was concluded that home interviews of this conventional type simply fail to record a fraction of travel and it is necessary to include screenline factors in order to match traffic counts.

In summary, however, although the validation process was less severe than could have been wished there was no reason to suspect that the models estimated contain significant unreliabilities.

After the estimation of the models described in this Report, applications were undertaken. The computer systems implementing these applications are described in Report 8 and some of the results are described in Reports 10 and 11 of this series.

CHAPTER 2. BACKGROUND

The function of this Chapter is to give background information about the Zuidvleugel Study and about the methods applied in the study to predict passenger travel.

2.1. THE ZUIDVLEUGEL STUDY

The Zuidvleugel Study, like most major regional analyses, was prompted by a variety of problems. The study concentrates on the Zuidvleugel (Southern part) of the Randstad conurbation in the western Netherlands. In the Zuidvleugel area, shortages of housing, recreational areas, and industrial land, have increased migration to the nearby rural areas, particularly to the areas in the province of Noord Brabant just south of the major rivers. Because residential relocation has occurred at a higher rate than shift in employment, passenger trip lengths have increased. Accompanying this trend has been an important rise in the level of car ownership, partly in response to the new travel needs of the population and partly due to rising incomes. The increased peak-period congestion that has resulted is aggravated by the limited capacity of the river crossings connecting the major employment concentration with the growing residential areas.

Because of these developments it became essential to develop methods that could indicate the size and character of the traffic flows in the Zuidvleugel area. Of particular interest was the relationship between the provinces of Zuid Holland and Noord Brabant. In addition it was desirable to gain insight into the short-term effects of policy changes that could be implemented quickly. Specifically, Cambridge Systematics was commissioned to:

- develop a prediction methodology to allow long-term forecasts to be made of the travel demand impacts of altering the transportation infrastructure, operating policies (e.g., pricing), or land-use changes; and
- develop a prediction methodology for conveniently carrying out short-term policy tests designed to explore the consequences of immediately applicable changes to the system.

In conjunction with other studies by DVK, these methodologies were restricted to predictions of personal travel by residents of the area. The modelling analysis on which the prediction methods developed by CS are based is the chief subject of this report.

The Zuidvleugel area covers about 2000 square kilometers in parts of the Netherlands provinces of Zuid Holland and Noord Brabant. This area contains the second and third largest cities in the nation: Rotterdam (by far the world's largest and busiest seaport) and The Hague (the seat of government). Figure 2 shows the location of the study area with in the Netherlands and some of the important features within the Zuidvleugel area.

The boundaries of the study area were selected to minimize the flow of commuter travel across them and to account for existing natural and administrative borders wherever possible. In 1977, the population of the Zuidvleugel area was about 3.178.000, including a work force of approximately 1.180.000. Employment in the area totalled about 1.115.000 jobs. Over four million (round trip) tours were made on a typical weekday during the survey period, and were distributed across the modes of travel as shown in Table 2 (taken from Report 6). The mean (straight-line) length of the tours was 7.2 km; tours to work were on the whole much longer, averaging 12.9 km. Detailed analyses of the travel patterns observed in the home interview survey are given in Reports 5 and 6.

2.2 ANALYSIS APPROACH

The principal objective of CS work for the Zuidvleugel Study was the forecasting of travel demand. For this study, it was not planned to undertake fundamental research, and the methods adopted are therefore essentially the most advanced that have been proven sound in practice. Nevertheless it was possible to make several important steps forward with the methodology in the course of the study, always

Table 2: Study Area Mode Split (expanded)

Train of Metro	1.9%	
Bus or Tram	4.5%	
Public Transport		6.4%
Car Driver	21.4%	
Car Passenger	7.7%	
Car		29.1%
Bicycle	26.8%	
Moped	3.7%	
Walk	32.1%	
"Slow"		62.6%
Other	1.9%	
Total	100.0%	

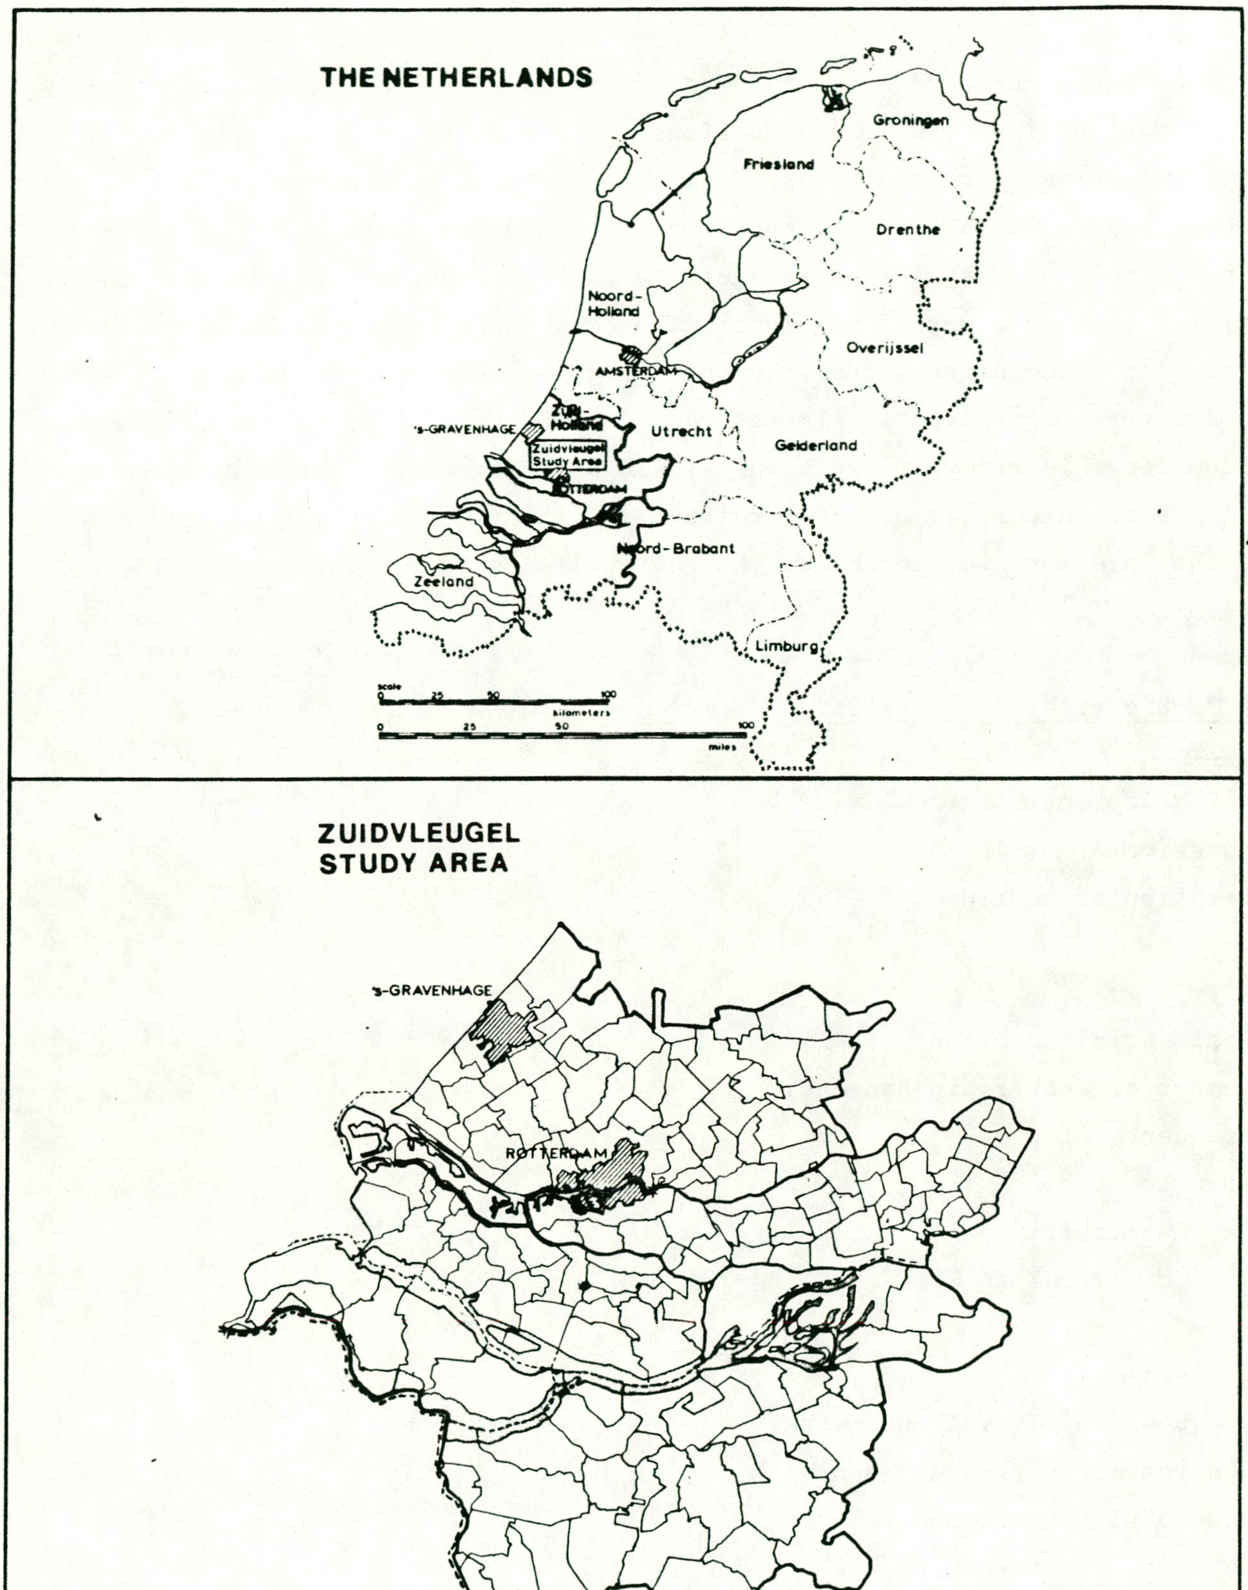

Figure 2: The Zuidvleugel Study Area

remaining within the framework developed in numerous previous studies.

The travel behaviour studied in this report forms a part of the huge comples of family behaviour involving many decisions that have little to do with travel. A complete representation of these decisions is beyond the scope of current knowledge and would in any case be too complicated for a study aiming to produce practical forecasts. What is clear is that the complex of decisions forms a hierarchy, in that more fundamental, longer-term decisions can be considered more or less independently from the shorter-term decisions which, however, they influence very significantly. To set the scope of this study in context, a more complete hierarchy of decisions is shown in Table 3. As explained in Report 1, the sub-hierarchy of inter-related decisions covered by the study are those that are most influenced by the policy issues to be addressed by the study.

Table 3: Hierarchy Travel-Related Decisions

Urban Development
- Location of Employment
- Location of Housing

Mobility
- Residential Location
- Number of Workers in Household
- Frequency of Work*
- Work Location*
- Car Ownership**
- Mode of Travel to Work*

Travel
- Frequency of Travel, non-work*
- Destination of Travel, non-work*
- Mode of Travel, non-work*
- Sub-mode (or Access/Egress** Mode of Travel)
- Route of Travel**

---

* Included in Zuidvleugel models reported here.

** Subjects of continuing or planned future study.

For reasons described in Report 1 of this study, the basic form used for discrete choice modelling was the "nested logit model". Some properties of this model are described in Report 1. For some sub-models (specifically those describing travel frequency), a slightly different form, the "ordered logistic model", which consists of a sequence of binary logit models was applied.

As described in Report 5, the travel observed in the home interview survey was interpreted in terms of "tours", continuous journeys beginning and ending at home. Following the arguments of that report, the models described here approximate the tours by neglecting all destinations other than the "primary" destination. Parallel studies, not reported here, are extending the analysis to "secondary" destinations for the 17 percent of tours that visit more than one non-home destination.

The models described in this report are therefore models of tour frequency, tour (primary) destination choice and tour mode choice.

2.3 FORECASTING METHODS

As noted above, it was required that this study should derive two separate forecasting procedures, one for application to the intermediate effects of short-term policies, the other for application of new infrastructure. This requirement had an important consequences for the estimation of the models used for forecasting: in fact, two separate sets of models were produced.

For many short-term policies, we apply the technique of sample enumeration to the sample of households selected by the home interview. This technique can be used to study short term policies. The technique is very simple to apply and produces unbiased aggregate forecasts. It requires a sample that is adequately representative for the policy under consideration. For general policy issues (such as pricing, for example), the household interview sample can be considered adequate in the short-term. For such analysis, sample enumeration can be applied using models containing all the data from the home interview survey, thus exploiting to the full information that was collected.

For other policy issues, the home interview sample is not sufficiently representative. Such issues are particularly those concerned with local policies or those with longer-term effects: infrastructure construction naturally has both these properties. In particular, the need to produce traffic assignment is a crucial issue. It is in theory possible to apply the sample enumeration procedure in such cases, but a more appropriate sample must be obtained, probably through an artificial sample generation technique. For this study, however, it was decided to develop an aggregate forecasting model system using restricted data - specifically, the data that was available for the whole of the study area and that could be forecast for future years. Such data is typically available on a zonal basis, and the models were therefore applied for aggregate forecasting using socio-economic and geographic segmentation on the basis of the 319 internal and 77 external zones defined for the study.

Accordingly, each sub-model was developed in two versions: one using the full data available from the home interview survey, the other using only data that can be forecast at a zonal level. More details of the data differences between the two sets of models are given in Chapter 3. The two sets of models are reported in parallel in Chapter 5.

2.4 SOFTWARE FOR ANALYSIS AND FORECASTING.

For a major data processing project such as the Zuidvleugel Study, a large number of computer programs are naturally required. Specifically, software is needed

- to process and check the survey data;
- to expand the home interview data;
- to make tabulations from the survey data;
- to process the transportation networks;
- to collate the network data with the survey data;
- to estimate and check the logit models;
- to implement the sample enumeration procedure;
- to implement the zonal forecast procedure.

In addition, many other minor programs were written to perform ancillary processing.

Survey data processing implies checking, assignments of zonal locations and produciton of reformatted files, in particular the derivation of tours. For this study, numerous programs were written on an ad hoc basis. In a parallel study, however, software is being prepared that will allow these tasks to be performed on a more systematic and consistent basis.

Software was written to expand the home interview survey. These programs are described in Report 6. Tabulations were consistently made using the SPSS package, which became available on the Univac machine during the study.

Processing the transportation networks for model estimation required the derivation of "level-of-service" (accessibility) measures for complete tours from networks defined in terms of links and modes. This process is essentially familiar and straightforward, but for this study it became and enormous problems because of the huge networks under study and very poor software available on the Univac. Three completely distinct processes were used for the three distinct main modes.

- The car networks processed primarily using the VERMOD programs. This system of programs gives a comprehensive processing capability for highway networks. Because of the size of the networks, it was necessary to simplify large numbers of less important roads into "access links" connecting ultimate origins and destinations to the more important modes. The program (called ROUTPV) used to form these links was not entirely satisfactory and is being replaced.
- The "slow mode" networks were also too large to process in their entirety with the VERMOD programs, not is it reasonable in this case to approximate the minor road and paths, since these are exactly the most important roads for slow mode use. The networks were therefore processed in sections using VERMOD and the resulting partial level-of-service matrices merged using a specifically written program.
- The public transport networks were the most awkward. Again the software available was unsatisfactory, and new software is being prepared. A particular problem was the calculation of fares for public transport tours. In the study area, fourteen different systems were in use at the time of the survey and much care was taken to reproduce these different fares properly.

It is no exaggeration to say that preparing the software and processing the networks was by far the most time-consuming part of the whole study. Many of the problems encountered can be alleviated by improving the programs used, and, as noted, new programs are in several cases being prepared. The size of the networks was also a problem, and although the size of the networks can be reduced (using programs which are also being prepared), the network to be considered when analysis is conducted in detail over a large region will always be substantial.

New programs were also required to merge the information from the home interview survey with level-of-service data from the networks. The problems here are caused exclusively by the huge network data files (the public transport data contains rather more than 10 million items), and to deal with these files long processing sequences had to be devised.

Programs for estimating and checking the logit models were prepared specifically for this project by adapting Cambridge Systematics software used in previous studies and introducing new features into those programs. The logit programs are fully documented in Report 3.

An entirely new program was written to implement the sample enumeration procedure. This program is fully documented in Report 8.

The zonal forecasting programs are based on a general matrix processing program called JMODEL, which was written specifically for this study, but which has wide applications to other projects. This program is documented in Report 2 of the Study.

CHAPTER 3. DATA FOR ESTIMATION

Data describing travel and urban activities in the Zuidvleugel area were assembled to fulfil two basic purposes in the modelling* : to enable the parameters of the demand models to be estimated and to validate the demand models that result from the estimation process. For estimating model parameters, two types of data are necessary: data describing the characteristics of particular travel decisions made by individuals (mode, destination, origin, time-of-day, purpose, etc.), and data describing factors which may contribute to the determination of the individual's travel decisions (level-of-service, socio-economic status, travel and activity opportunities, etc.). These two types of data were derived from three basic sources. Information about individual travel decisions and the socio-economic status of travellers was collected in the course of the home interview survey, which is described in Section 3.1. Travel level-of-service data were derived from network representations of the transportation systems in the Zuidvleugel study area and are descibed in Section 3.2. Opportunities for travel and descriptions of the Zuidvleugel area activity system are contained in a set of zonal-based data, assembled from existing data sources, and discussed in Section 3.3.

3.1. TOUR, PERSON, AND HOUSEHOLD-BASED DATA.

The home interview survey was conducted between September 1977 and January 1978. The sample design, described in Report 4, was stratified to increase the variance of a number of key independent variables, particularly the degree of urbanization, the quality of public transport level-of-service, and the number of commuters in the sample. A total of 2949 households were surveyed.

The survey contained three types of forms on which information was collected: the household (or main) form, the person form, and the trip form. Children between the ages of four and fifteen were given a shortened version of the person form and those below the age of four were not required to complete either person or trip forms.

The completed survey forms were then coded, keypunched, and processed to "link" the trips (eliminating as separate stops those with purposes "change mode" and

---

* Also, as presented in Report 6, the data allowed a general picture of travel to to be given.

"park car"). Subsequent to this trip linking, a long series of data checking and editing steps discovered and corrected a surprisingly large number of errors in the data*.

Following the very lengthy correction process, it was possible to combine trips to form the tours that were the basis for the modelling. This process and the analysis on which it was based in reported in the first two chapters of Report 5. The subsequent chapters of Report 5 give analyses and tabulation of the home interview data (in the form of tours) from the point of view of the proposed modelling. A number of conclusions were drawn that are used in the models.

Other data were also collected in the Zuidvleugel home interview survey for use in other studies. Some of the items describe residential location changes over time, workplace changes over time, and the detailed nature of automobile holdings (make, model, model year, and whether the vehicle was purchased new or used).

3.2 NETWORK DATA.

The required network data are level-of-service measures for tours made in or from the Zuidvleugel study area. This information is, of course, critical to the development of all the passenger travel demand models developed since the level-of-service variables are exactly those influenced by transport policy decisions. In order to develop the required information, the networks are processed with a variety of programs as outlined in Section 2.4. In this section the information derived by these processes is discussed.

The networks were defined at two levels of detail. All surveyed households in each 500 metre square were assigned to a (single) household centroid. Further, every internal and external zone was assigned to a centroid node. The household and zone centroids were then linked in several directions to the closest nodes in the network. The origin of a tour (a household) could thus be connected to the network with two levels of accuracy: to the nearest 500 metres or (much less accurately) to zonal accuracy**. Both these levels were used in model development. For both zonal and fully disaggregate models, however, the destinations of tours were only represented to zonal accuracy***, because of the effort that would have been re-

---

* Investigations of the causes of these errors have led to recommendations for improvements in data collection, coding and processing which are being implemented in new studies.

** The zones, although less accurate than 500 m, are nevertheless rather smaller than is common in regional studies.

*** With the exception of a few preliminary model versions that did not use the networks.

quired to code the eleven thousand actual destinations (and virtually unlimited potential destinations) to greater accuracy.

For the car network, three level-of-service parameters were derived: time, distance and toll costs (which were rare). Networks processed separately for three time periods (morning peak, evening peak, and off-peak) to represent the variation of congestion during the day. In each case a full matrix was calculated giving level-of-service for every possible tour in both directions at both levels of accuracy.

These level-of-service measures were taken from paths calculated by the programs to minimize the "generalized cost" of travelling from the origin to the destination. That is, a route was found through the network that minimized a specific combination of time, distance and cost and the level-of-service measures used in the modelling were those relating to this individual path. Considerable attention was given and new methods were developed to find the approriate combination of time and cost for calculating the path, and also to ensure that the paths themselves were reasonable.

For the slow network, a single parameter was derived: the distance between origin and destination. Because congestion was not felt to be an important influence, a single path was built for each origin-destination pair, and the distance along that path applied to all potential tours irrespective of the time of day.

The most difficult problems arose with the public transport networks, mainly because of the inadequate software available, Because of the computer and analyst time required to run this software, a full matrix could not be developed. Instead, paths were built from all zones to all zones and 500 metre household nodes. For trips from household nodes to zones, it was assumed that the level of service parameters were the same* as for the trip in the opposite direction. Because services vary during the day, different parameters were developed for the same three time periods as for the car network.

The appropriate parameters for combining walking, waiting and in-vehicle times together with interchange penalties in a linear function were more difficult to determine for the public transport networks than for the car networks. This was because public transport networks naturally generate more level-of-service parame-

---

* With appropriate permutations for access links and waiting times - See Table 4.

ters than car networks and also because of the difficulty of using the software. Further, it gradually became apparent that whatever weights were used, the network suggested more interchanges than were in practice made. This problem could only have been rectified by a fundamental recoding of the networks, and this was not possible. Results regarding the importance of interchanges derived from the models are therefore suspect.

The complete list of level-of-service parameters derived from the networks is given in Table 4. Several million data items were generated, and several computer tapes are required to store it.

TABLE 4. Level-of-service measures

Car (varies by time of day)

Travel time, origin to destination
Travel time, destination to origin
Distance, origin to destination
Distance, destination to origin
Tolls, origin to destination
Tolls, destination to origin

Slow (constant)

Distance, round trip

Public Transport (varies by time of day)

Walk time, origin to first stop
Walk time, last stop to origin (return)
Walk time, interchanges outbound
Walk time, interchanges return
Walk time, last stop to destination
Walk time, destination to first stop (return)
First Headway, outbound
First Headway, return
Headway at interchanges, outbound
Headway at interchanges, return
Last Headway, outbound
Last Headway, return
In-vehicle travel time, outbound
In-vehicle travel time, return
Number of interchanges, outbound
Number of interchanges, return
Full fare
Pass fare
Concession fare
Concession Pass fare

Units on the final files, times are measured in minutes, distances in kilometres, and costs in guilders.

3.3 ZONAL DATA

Data describing each zone in the Zuidvleugel study area was assembled from available sources. The data items included are listed in Table 5.

TABLE 5. Zonal data

Households
Cars
Population
Population, 0-11 years
Population, 12-18 years
Population, 18+ years
Workforce
Total employment
Agricultural employment
Industrial employment
Construction employment
Service employment
Retail employment
Retail establishments (number)
Students, 0-18 years
Students, 18+ years
Intrazonal travel distance (km)
Parking cost, before 9 AM arrival (guilders)
Parking cost, after 9 AM arrival (guilders)
Average income class (low, medium, high)
Total area (sq. km.)
Sport area (sq. km.)
Water area (sq. km.)
Wooded area (sq. km.)
Natural area (sq. km.)
Developed area (sq. km.)
Other area (sq. km.)
Total recreational area (sq. km.)

Number of recreation establishments receiving 10000-250000 visitors annual
Number of recreation establishments receiving 250000-500000 visitors annual
Number of recreation establishments receiving 500000-1000000 visitors annual
Number of recreation establishments receiving more than one million visitors annually

4. ANALYSIS ISSUES

It was decided from the outset that the demand models for the Zuidvleugel study would be based primarily on the logit model. This choice was made because it was known to have given reasonable results in previous studies and to be reliable in practice. That is, since the study was intended to develop practical forecasting methods, it was not appropriate to attempt basic research on model forms beyond the current state of the art.

Nevertheless, in a number of cases it was found necessary to implement marginal improvements to the methodology available at the outset. These marginal improvements, though not radical, represent together a substantial advance in methodology, particularly in the practical aspects of applying recent theoretical findings. These improvements are the subject of this chapter.

Perhaps the most significant single development implemented in the study was the use of "tours" rather than "trips" as the basic unit of travel. This departure is discussed in detail in Report 5 of the study. All the models described in the current report are thus models of travel represented as tours.

In this chapter, the following analysis issues are discussed in successive sections:

- model structure and linking of sub-models;
- availability of alternatives;
- use of attraction variables;
- sampling of observations and alternatives;
- aggregation.

A final section describes the goodness-of-fit measures used to assess the quality of the models developed.

4.1. MODEL STRUCTURE

Report 1 of the Zuidvleugel Study foresaw a model structure primarily comprising a linked sequence of "nested" * logit models. These models will be described briefly, noting some of the problems that arise in practical im-

---

* the literature also uses "structured", "sequential", "hierarchical" or "tree logit".

plementations, and the solutions adopted in this study. Report 1 also foresaw that detailed amendments would be necessary to the general structure, and these are also described.

4.1.1 THE NESTED LOGIT MODEL

A nested logit model structure is built up to a series of logit models of the usual form:

$$\log P_{c/C} = V_c - V_C \qquad (3)$$

$$\text{where } V_C = \log \sum_{j \in C} \exp V_j$$

giving the probability $P_{c/C}$ of making a specific choice c from the set of available alternatives C as a function of "attractiveness" functions V_j defined for each alternative j:

$$V_j = \theta_1 x_{1j} + \ldots \theta_r x_{rj} \qquad (4)$$

where $x_{1j},\ldots,x_{rj}$ are observed characteristics of the alternatives of the traveller, or are functions of observed characteristics; and $\theta_1,\ldots,\theta_r$ are parameters to be estimated.

A logit model may be derived from several different theoretical bases. The most useful basis for work up to the present has been that the model (3) is the result of utility maximisation by the travellers. That is, as postulated by economic theory, each individual selects (subject to the constraints of his time and money budgets*), the alternative that gives the greatest satisfaction. In the theory V_j (4) represents the analyst's best approximation to the true attractiveness U_j:

$$U_j = V_j + \varepsilon_j \qquad (5)$$

where ε_j represents the error in the approximation process. This error arises because of the omission or approximation of relevant factors that are impossible (or too expensive) to measure or to measure accurately, and because of in-

---

* The effects of constraints on choice are discussed in Section 4.2.

ter-personal variations in taste (i,e, in the θ parameters).

The true attractiveness U_j experienced by the traveller may be interpreted as the realised utility resulting from the choice of this alternative. Since travel choices in general involve the consumption of the scarce resources of time and money, it is to be expected that times and money prices as well as income may appear as variables in these functions. There is therefore little restriction imposed by utility theory on the form of the attractiveness functions. Characteristics of both the traveller and the alternatives may well play a proper rôle in these functions.

It is neither necessary not desirable to model all aspects of the traveller's choices. Some choices are taken as given by the models described in this report; other choices are too detailed to model in this study*. All the choices described here are conditional on "higher" level decisions and incorporate "lower" level choices. Moreover, the choices described here have their own hierarchy, as described in Report 1.

Thus in the approximation (5) of utility, not only does V_j incorporate the approximations mentioned previously, but also its distribution is influenced by the travellers choice from a number of sub-alternatives. This fact suggests that the appropriate distributional form for the error ε_j (in 4.3) should be the "extreme value" distribution:

$$\Pr\{\varepsilon_j \leq t\} = \exp\left(-e^{-t}\right). \qquad (6)$$

It is not intended to advance the argument that this form (6). is unequivocally the appropriate distribution to use for this purpose, but to suggest that this distribution has some theoretical merit and can be considered as at least as attractive as other forms (e.g. normal distributions) that might be employed.

Rather than the specific distributional forms for these distributions of ε_j which in any case have little effect on forecasting, the important assumptions concern the interdependence or otherwise of the distributions of $\varepsilon_1 \ldots \varepsilon_r$. It can be proved that if each has a distribution of this form (6) and if these distributions are independent and have equal variance then the

---

* See Section 2.2.

model (3) is obtained*. It is this "independent, identically distributed" (I.I.D.) assumption that has the important consequences of uniform cross-elasticity that is a restriction of the simple logit model**.

In recent years considerable effort has been devoted to finding models that do not suffer from the cross-elasticity restriction of the logit model. Considerable success has been obtained in developing theoretical models free from this restiction, rather less in establishing the practicality of these models; two problems may be cited. The first is that these more general models are difficult: it is a challenging problem in numerical analysis to find workable means of estimating the unknown parameters, both because of the complexity of the formulae and because of the awkward behaviour of the functions. Of course, this problem can be expected to be reduced in importance gradually over time as methods improve. The second problem, however, is more basic: a more general model inevitably has more unknown structural parameters which require estimating. Apart from the requirement for increasing the data needed to estimate the model, the question arises of whether the resources devoted to estimating these structural parameters could be better devoted to measuring more characteristics of the traveller's choice situation.

For the Zuidvleugel Study, it was essential to choose a model form that was practicable without basic research. The logit form (3) was used as the basis, with attention given to remedying a number of its limitations:

- cross-elasticity restrictions were reduced (see below);
- consideration was given to the possibly limited availability of alternatives in the set C (see section 4.2);
- limited non-linearity in the attractiveness function V (4) was introduced (see section 4.3).

This approach switched the emphasis of the work from structural considerations to more practical assessments of the importance of measured characteristics of traveller and alternatives. The resulting models show the richness of detail that was achieved by this approach.

Two main steps were taken to improve the cross-elasticity problem. One, applied

---

* See, for example, Domencich and McFadden, Urban Travel Demand: A Behavioral Analysis, North Holland, Amsterdam, 1975.

** This restriction is, in summary, that the (point) cross-elasticities of all other alternatives with respect to a characteristic (e.g. price) of one alternative are equal.

only to the frequency models, was the use of ordered logistic models: these models are described in section 4.1.3. The more important was the use of nested logit models throughout the system. The principle of nested logit models is that although the assumption of equal cross-elasticity is unreasonable over all the dimensions of choice, there are sub-sets of the choices for which this assumption is tenable. The approach is then to use the simple logit model within these sub-sets, and to link the resulting series of models with the characteristic "logsum" variable.

The logsum is the variable V_C defined in the basic logit model equation (3), where the reason for the choice of name is obvious. Apart from its key rôle in that equation, the essential property of the logsum is that if the actual attractiveness of the alternatives is distributed as indicated by (5) and (6), then the attractiveness of the alternative selected (i.e. the maximum of the set) is given by:

$$U_C = V_C + \varepsilon_C \qquad (7)$$

and ε_C has also the distribution (6). Thus the logsum gives a summary of the attractiveness of the whole choice set C.

The nested logit model exploits this summary property of the logsum variable by predicting choice within a subset of alternatives (for which equal cross-elasticity may reasonably be assumed) and then using the logsum as a component of the attractiveness of the subset in a "higher level" model which includes the subset as a single choice. The theory indicates a restriction on the value of the coefficient θ which the logsum variable may have to be in the range $0 < \theta \leq 1$. If the value of θ is 1 then the representation is two levels is equivalent to a single level model. The hierarchical structure proposed for the Zuidvleugel models (Report 1) clearly leads itself to representation in this way.

The nested logit model thus predicts a complicated choice (say of mode and destination) as a process of two stages (or more). A model is developed of choice within a limited subset of the choices (say of mode) of the form (3). The logsum V_C from that model is then used in a further model (of destination) which has attractiveness functions V_j^* containing these lower level logsums:

$$V_j^* = \theta_1 V_C + \theta_2 x_{2j} + \dots + \theta_r x_{rj} \qquad (8)$$

and Θ_1 has the restriction $0 < \Theta_1 \leq 1$. The process may be continued by calculating further logsums V_C^* and entering these in a higher level model.

In practice, a number of considerations other that these theoretical points are also important. These considerations are discussed in the following section.

4.1.2 LOGSUM VARIABLES AND INCLUSIVE UTILITY

For a logit model predicting choice between a number of alternatives it is possible to use the logsum variable, as described in the previous section, as a summary of the utility of all the alternatives. This summary utility can then be used as a variable in a model of choice over a wider set of alternatives, thus relaxing the cross - elasticity restrictions of the logit models. Three practical difficulties severely limit the applicability of this procedure.

The first is that the calculation and use of logsum variables is a complicated and time-consuming process. A forecasting system is always a balance between the conflicting needs of realism and complication, and the additional burden imposed by the calculation of logsums is entirely unwelcome. It is necessary to be certain that this complexity is balanced by substantial improvement in the accuracy of the model before it is possible to justify the investment of study resources in logsum calculations rather than other improvements to the model.

The second problem is that a sequence of linked logit models for two or more levels is actually a single nested logit model for the choice over all the alternatives at all of the levels. Estimation of this model ought therefore to be simultaneous over all the levels in order to exploit all the information. At present, however, no computer software exists that can perform this simultaneous estimation. Instead,a consistent estimation procedure is used in which estimation is performed at the "lowest" level first, logsums calculated from that estimation are used as explanatory variables for estimation at the second level, and so on. Although it can be shown that this procedure is asymptotically unbiased, it is clear that it is not "efficient" in the sense that the lowest level estimations are made in ignorance of the second level alternatives. Until new software is developed, this inefficient procedure is the only practical method available for estimating hierarchical models.

The third problem is that variables describing lower level choices that have equal values across those choices cannot appear in the lower level model but must be introduced (along with the logsum) as characteristics of the lower level choice subset in the model at a higher level. A clear example of such a case be a model of mode choice in which public transport alternatives were taken as a subset. If the fares for all public transport modes were the same (as happens in some areas) the the coefficient of fare could not be identified in the lower choice model. The fare could, however, be introduced as a variable applying to all public transport modes at the second level. Apart from this type of difficulty in application, the issue in estimating the model appears to be identify the variables that are common between lower level alternatives. Most are not as clear as the example given.

In particular, the key difficulty arises from variables that are unequal but highly correlated across the lower level alternatives. Including such variables in the lower level model will result in large standard errors of their estimated coefficients. Thus when variables are formally distinct but highly correlated the inefficiency of the sequential procedure at the lower level may well give poor estimates of their parameters which are propagated to higher level models in the form of inappropriate values for crucial parameters in the logsum variables*. The key example of this problem concerns the variables describing journey length (e.g. times, distance, costs), which are obviously highly correlated across different modes and may therefore acquire inappropriate parameters if estimation is only over mode choice.

Since the only practical technique currently available for estimation hierarchical models is the "bottom-up" process described, and since this process inevitably leaves the possibility of the problems described there is good reason to avoid hierarchical modelling if at all possible. Section 4.1.4 describes the approach taken for this study, in which hierarchies were used in some cases (believed to be the most important) and joint models in other cases.

4.1.3 ORDERED LOGISTIC MODELS

This study took a rather different approach to modelling travel frequency than that taken in most previous studies, whether aggregate on disaggregate. Instead of representing mean trip or tour rates, a model was developed of the choice

---

* In such circumstances the parameter estimated for the logsum variable will obviously be suspect.

of tour frequency.

Travellers decide each day on the number of tours they will make (for each purpose), depending on their social and economic needs and desires, and the opp rtunities available to them. This is exactly the same type of choice represented by the other models (e.g. of mode choice), and there is no reason whatever not to model tour frequency by the same type of model used in the other choices.

A great advantage of using the same type of model in all choice situations is that the models can then be made consistent with each other within the theoretical framework adopted. In the case of logit models within the utility maximising framework, consistency requires the introduction of "logsum" variables in frequency models, when these are at the "top" of a hierarchy of logit models. Such variables are in principle relevant to tour frequency choice because the logsums derived from destination choice models can be interpreted as measures of accessibility.

A further advantage of the formulation of frequency models in the choice framework is that a proper distinction can be made between the choices of whether or not to make the first tour and whether or not to make a subsequent tour. In terms of the behaviour of the traveller, the first tour (requiring a decision to take part in an activity) is quite different from subsequent tours (a matter of the organisation of the time and place of taking part in the activity).

The appropriate logit implementation of this type of thinking is the ordered logistic model. These ordered models represent first the choice between no tours and making one or more. Conditional on making one or more tours, the choice between making exactly one and making two or more is represented by the second model, and so on. In principle, a logsum from the subsequent choice models should be included in each model, but in practice this was too complicated to implement.

A further simplification was to use the same model to represent choice between one tour and two-or-more, two tours and three-or-more, and subsequent choices. Thus for each purpose, two models were estimated: choice between no tour and one-or-more, and choice given a positive number of tours, making no more tours and making at least one more. This simplification was made primarily because of the small number of people observed making more than one tour for a given purpose, and seems appropriate to the data available.

4.1.4 EMPIRICAL AMENDENTS TO STRUCTURE

The model structure proposed in Report 1 was based on the choice hierarchy: frequency - destination - mode chòice. It was not envisaged that the simplicity of the structure would necessarily be preserved through the process of estimation with real data. In the event a number of minor structural amendments were necessary, and these are the subject of this section.

Slow Mode Choice

In mode choice modelling a more detailed structure was introduced. The mode choice observed in the area (described in detail in Report 5) is characteristic of the Netherlands in the predominance of "slow" modes (walk, cycle, moped), the consequency of topography, urban structure and tradition. The modelling of this study was intended to produce forecasts of traffic on strategic highways,so that the policy focus was primarily on car traffic and secondarily on public transport. Studies of this type in other countries have typically neglected "slow" modes, but in the Zuidvleugel area the importance in sheer numbers of these modes made their inclusion essential, at least as alternatives to the other modes.

Consideration of the structure of a choice model including slow modes reveals a number of unmeasurable common characteristics of these modes, particularly exposure to weather and the use of one's own physical effort, that are not shared with other imporant modes. Many measurable characteristics such as speed and price are also very similar among "slow" modes relative to other modes. These considerations suggest a structure in which choice among slow modes is considered at a low level in the hierarchy, and choice between them and other modes at a higher level.

These suggestions from analysis of the home interview data were reinforced by modelling resuls, which showed significant improvements in modelling accuracy to result from the structuring. These results are reported in Chapter 5.

Thus a series of models of choice among slow modes was estimated. Although distance naturally plays an important rôle in the choice among these modes, the other variables found were primarily socio-economic, with age particularly prominent. Although these variables are very interesting in giving insight into the users of the different slow modes, they are not the variables affected by the policies for which the model system was designed. The primary function of the slow mode models therefore to give (through the logsum) a non-linear representation of increasing disutility of slow mode travel with increasing distance, conditioned by the socio-economic variables.

The objectives of the zonal models are still more directly focussed on forecasting car traffic. Further, the socio-economic variables with which the slow mode models operate are not in general available at zonal level. There is little point in a series of models giving only a non-linear distance function, which can be estimated directly in any case. Slow mode models were therefore not aggregated to the zonal level, and in those models the choice "slow" appears without further subdivision.

Car Mode Choice

The household interview survey recognised four different car modes:

- drive, with passenger(s);
- drive, alone;
- passenger, in car owned by the household;
- passenger, in other car.

There was, however, strong evidence of a problem of questionnaire design causing confusion between the two driver modes. No distinctions were made between these two modes in the modelling.

Although it appeared that the questionnaire design problem might extend to the distinction between passenger modes, this was not clear and no further evidence could be otained. In the models presented in the report no distinction was made between the two passenger modes, but subsequent models estimated on the Zuidvleugel data have made this distinction.

The structure of driver-passenger choice with the remainder of mode choice is not so clear an issue as with slow modes. Nevertheless there are several common characteristics and empirical tests were therefore made. These showed that the structuring with driver-passenger choice at a lower level gave a significantly better explanation of the observations than the more simple model in which driver and passenger appeared as choices along with slow modes and public transport in the main mode choice model.

An exception to the above finding was in the education tours model. Here it was found that very few of the travellers had licences (not surprisingly) and that of those that had licences almost none travelled as car passengers. It was in these circumstances impossible to develop a driver-passenger model for education tours. The main mode choice for education tours thus has four alternatives: drive, passenger, public transport, slow.

Note that there was no question of eliminating the distinction between car modes in the zonal models, as was done for slow modes. Particularly in the zonal models, where assignment is an important output, the prediction of car occupancy given by the car mode models is a necessary part of the forecast system.

Mode and Destination-Choice

Existing travel demand models for large metropolitan areas usually places destination choice at a "higher" (more fundamental) level than mode choice, although there has been until recently little chance of obtaining evidence to support the tradition. For some purpose (e.g. work, education, social) it is difficult to believe that this hierarchy is wrong, but for other purposes (recreation, shopping) the alternative structure with destination conditional on mode is more readily acceptable.

Study resources did not permit an extensive investigation of the structural issue. Preliminary tests showed that a structure with mode choice conditional on destination choice gave difficulties in estimating the time and cost parameters, because of the problem described in Section 4.2.1. The decision was therefore taken to estimate "joint" models in which mode and destination choice appeared at the same level in the hierarchy and their parameters could be estimated simultaneously by the software available.

An exception to the use of joint models was made for work and education tours. In these cases it was felt a priori that mode choice must be conditional on destination choice, and hierarchical structures were therefore implemented.

Frequency Choice

The basic structure of ordered logistic models used for frequency choice was described in Section 4.1.3. An exception to the structure had to be made on empirical grounds, however, in the case of education tours.

First, it was found that a distinction needed to be made between full-time students (including children) and other adults. The latter group obviously make many fewer education tours, and almost none in that group were observed to make two education tours in the day. A simple model for that group predicting only the choice between no tour and one tour was therefore estimated.

For full-time students , the model predicting choice between no tour and one or more presented no particular problem. Almost no students, however, were observed making three tours or more. The secondary frequency model for this group was therefore simply of choice between one tour and two tours, rather than the (theoretically) infinite sequence of choice assumed for other travel purposes.

For the zonal models, numbers of resident students were not available per zone. The distinction therefore had to be between children (under 18) and adults. This was slightly inferior to the distinction made in the fully disaggregate case.

4.2 AVAILABILITY OF ALTERNATIVES

A logit model is usually presented as the formalisation of a choice among a number of "available" alternatives. That is, a preliminary process is envisaged during which a choice set is generated, from which choice set the traveller makes his final selection, e.g. by choosing the alternative with the highest utility. Much modern work has stressed the importance of constraints in determining behaviour, and has pointed out that many travellers do not perceive a "choice" at all - they have no alternative to the tour they made. This work, together with the experience in previous studies of mode and other choices that variables describing availability are particularly important, led to a careful study being made of the issues of availability arising in the Zuidvleugel system.

An important initial point is that the concept of availability is not always entirely clear. The man who cycled to work, when asked "was a car available?", may respond either:

yes, but my wife wanted to use it;

or

no, because my wife wanted to use it.

The distinction between these two cases is clearly very fine, and the question of what might have happened if it had been raining heavily is not advanced very much by the answers to the questions. Second, a detailed study of the actual availability of travel alternatives as perceived by travellers would have involved extensive research beyond the scope of the Zuidvleugel work. For these reasons few questions on availability were asked in the home interview survey.

The uncertainly of availability goes further than the questionnaire dificulty of obtaining the traveller's true situation. There is in reality no clear dividing line between a choice being unavailable and a choice that is very unattractive. The uncertainly is increased when the data for analysis is restricted to variables which can be observed and forecast reliably.

An explicit model of availability, from which the output would be the probabilistic set, appears to be beyond present knowledge. It is essential, however, to take account of uncertain availability issues, at least implicitly, in developing the models. This implicit treatment of availability means that the models are not "choice" models in the usual sense of the word, but become models predicting availability and choice together.

Two separate means may be used for introducing availability into lógit models. The first occurs when a choice may be eliminated with a high degree of certainty, the second when the best that can be done is to condition the probability that a choice will be available.

4.2.1 CERTAIN UNAVAILABILITY

Despite the uncertain nature of availability, discussed above, there are a number of circumstances in which it is reasonable certain that an alternative is not available. The most obvious case is when the absence of a car or a driving licence makes driving unavailable. The introduction of such variables into the model gives substantial improvements and clearly as many as possible should be found and used.

A necessary condition for the use of deterministic constraints on availability is that it should be clear that the choice is conditional on the variables, and not vice versa. For example, it would not be reasonable to present the choice of mode for journey to work as conditional on bicycle ownership: for such an important and frequent journey it is more reasonable to suppose that a traveller would be willing to buy a bicycle if he felt that this was the best mode for his journey.

A second condition for the use of deterministic constraints is that they really should imply the unavailability of the alternative. For example, the hypothesis that a necessity to take packages made public transport unavailable would have to be abandoned if a significant number (5 to 10 percent) of travellers taking packages were observed to use public transport. The variable would be more appropriately handled as an indication of conditional availability. In principle, a single counter-example should cause the abandonment of deterministic constraints. In practice, a certain degree of realism is necessary, and to obtain the simplicity given by deterministic constraints it is worth accepting that a small number of observations are not consistent with the model.

In practice, the deterministic constraints used were car ownership and driving licence possession, both of which were required before the car driver mode was considered available; and the existence of a path in the public transport network,to make that mode available. Paths were not built if either the origin or destination was very distant (more than 40 minutes) from the public transport network, or if the journey was too short to build a path because both origin and destination were connected to the same node. These variables caused few problems in the sense that very few travellers had selected alternatives that were given as unavailable by these constraints.

In previous studies (e.g. SIGMO), cars had been considered unavailable for day-time non-work travel if they were used for work travel. In this study, this approach was not taken for a number of reasons. First, it is not always the case that work mode choice desisions are made without regard to non-work travel needs. This interconnection is therefore better expressed as a conditional availability. Second, the complexity of scheduling within each household is often difficult to follow with real data, making it often unclear exactly when cars were or were not in use. Third, in forecasting, the need to maintain links between work and non-work models would add significantly to the complexity of the model system.

In a logit model, the existence of a deterministic constraint making an alternative unavailable is equivalent to assigning a utility of minus infinity to that alternative. This in turn is equivalent to the omission of the alternative from the choice set C used in calculating the logsum in equation (3). The same reduced choice set logsum calculation is appropriate both for calculating the probability as in equation (3) and in using the logsum V_C as a summary measure of utility in higher level choices.

Note that when there is a variable that implies certain unavailability, it is preferable to recognise this explicitly in the specification of the model. Simply allowing the logit model to assign a large negative utility runs the risk of computer difficulties.

4.2.2 CONDITIONAL AVAILABILITY

Most variables influencing the availability of alternatives do so only in a manner that cannot be made certain by the analyst. For example, a traveller may be unaware of the existence of a shopping centre, thus making that centre unavailable for shopping tours, but this cannot be determined by the analyst with only limited insight into the traveller's mind. All that can be done is to recognise that variables may be found that influence availability (e.g. distance from the home), and to include these in the proper form in the models.

Conceptually, the modelling may be considered in two stages: first a model of availability and then, conditional on availability, a model of choice among the available alternatives. This structure with two-stage conditional modelling is very similar to the conditional modelling used, for example, for destination and mode choice. A reasonable approach for the problem of availability as for the other conditional choices is to use hierarchical logit models.

Two difficulties arise in representing availability and choice in this structure. The first is that, as noted is Section 4.1, software for estimating a hierarchical model is not currently available. The second in that, for reasons noted above, information about actual availabilities is uncertain and was collected only to a limited extent. This second problem means that even the sequential estimation used as an approximation for the hierarchical destination and mode models is not available. The approaches considered for this study were therefore to write new software allowing proper estimation of an availability-choice structure, or to find a means of incorporating the most important effects within simple structures, The latter approach was selected to avoid the necessity of undertaking basic research.

In this approach variables are introduced into the utility function of the logit models which do not represent utility in the usual sense but are connected with availability. It can be shown* that this approach is an approximation to the precise form of model that would be appropriate for the hypothesized structure. In most practical cases the approximation is good.

A wide range of variables connected with availability have therefore been introduced into the models. In particular, the models for mode choice contain where possible variables indicating the competition for cars within the household. In addition, the models of destination choice contain distance variables to allow representation of decreasing availability (because of decreasing information) with increasing distance. These distance variables are additional to variables representing the increasing disutility of travel with increasing distance (such as travel times or "logsums" from lower level models).

The approach of modelling availability implicity through the utility functions has been used in previous studies, although it was not always realised that the variables in question were in fact representations of conditional availability rather than utility.

4.3 ATTRACTION VARIABLES

A further amendment to the standard form of logit function that is desirable in the interest of accurate modelling is the introduction of multiple attraction variables. These variables differ from the usual variables in a utility function in that they represent the quantity or size of an alternative, rather than its quality or utility.

In previous studies (e.g. SIGMO) it had been necessary to estimate the relative values of such attraction variables in a preliminary model, usually a linear regression. A failure to estimate attraction variables simultaneously with the accessibility measures can easily give rise to biases when attraction variables are correlated with accessibility (positively or negatively). For this study an improved procedure allowing simultaneous estimation was adopted.

The approach adopted is described by Daly**. In that paper it is explained why

---

* Notes ZPS-79-12 and ZPS-79-15.

** A.J. Daly "Estimating Choice Models Containing Attraction Variables" Transportation Research, 1982.