THE GEOGRAPHICAL FOUNDATIONS

OF

DRY BULK SHIPPING

RISTO LAULAJAINEN

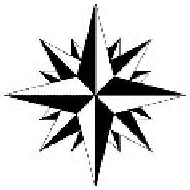

GOTHENBURG SCHOOL OF BUSINESS, ECONOMICS AND LAW
PO BOX 630, SE-405 30 GOTHENBURG, SWEDEN

June 2006

The data has been donated by
(in alphabetical order)

Baltic Exchange, London

Clarkson Research Studies, London

Drewry Shipping Consultants, London

Lloyd's Marine Intelligence Unit, London

The publication has been financed by

Lighthouse, Gothenburg

Front cover courtesy of Rundqvists Boktryckeri, Gothenburg

Distribution by Risto Laulajainen
Department of Human and Economic Geography
Gothenburg School of Business, Economics and Law
Box 630, SE-405 30 Gothenburg, Sweden
E-mail risto.laulajainen@geography.gu.se

Distributes also
Laulajainen, Risto – Holgersson, Thomas – Strömberg, Ann-Brith
Operating Panamax Dry Bulk Carriers on the Seven Seas
Gothenburg 2001

EXECUTIVE SUMMARY

Laulajainen, Risto, **The Geographical Foundations of Dry Bulk Shipping**. XIV + 262 pages. Available free of charge atwww.ristolaulajainen.com

A comprehensive presentation of the dry bulk industry with emphasis on geographical interaction. Describes 117,000 ship movements in cargo and ballast by size segment (large and small capesize, panamax, handysize) in the mid-1990s. Merges 40,000 voyage and trip fixtures and develops simple, i.e. user-friendly, freight rate functions for practically all possible routes 1993–2002. Compares by size segment, with the help of a Simulator, the daily revenue which ships collect during a calendar year once they have taken the first cargo in a particular region, ten in all. This revenue is a linear function of the ratio of demanded and available tonnage (Tonnage Balance), weighted by sailing distances from/to a discharging/loading region (Revenue Potential). The function, called Revenue Gradient, may be *the first published effort ever* to explain local, as contrasted to global, freight rates in general terms. There is also a level difference between Atlantic and Pacific rates, usually in the Atlantic's favor, believed to reflect oligopolistic tendencies in the Pacific. When the function is calibrated by ship segment and year, and the share of these two variables out of the revenue range is estimated, roughly 10 pct of the range in daily revenue can be explained by the initial positioning of ships, i.e. simple geography. The simulator consequently has potential for strategic planning and can be developed into a powerful operational tool.

When some dramatic event upsets the current order and a reorganization of trade flows takes place, freight rates follow. A related paper, based on the Tonnage Balance and Revenue Gradient concepts, shows how the relative rate levels can be estimated in advance. Probably *another, first ever, achievement*.

The central theme is complemented by excursions to important but so far insufficiently illuminated problems. Among them are estimates about the share of retail (voyage and trip) and wholesale (COA and time) markets, estimates about the share of cargo legs covered by published fixtures, estimates about the actual and "optimal" workloads, whether in cargo or ballast, an attempt to find a general explanation to the ratio of fronthaul and backhaul rates (led to the related paper above), and attempts to cluster routes by their rate levels, or alternatively their cyclic patterns, possibly related to the export or import character of the loading region.

Approximately 120 pages of appendices provide a wealth of tabular and oral material about the dry bulk industry. They will open many a "black box", so typical in research reports. Some information is interesting, perhaps unique, for its own sake – for example, cargo, ballast and fixture matrices by ship segment, and all the freight rate functions used, often with scatterplots. For typographical reasons, the regional mesh is usually ten macroregions which, however, rests on a finer mesh of 37 regions.

Keywords: Atlantic/Pacific, freight rate, loading region, revenue gradient, strategic planning, tonnage balance, trade matrix

Some terminology

Leg = Travel made by a vessel between two ports or parts of a port.
Cargo resp. ballast leg = Leg made in cargo resp. in ballast (to cargo, to ballast).
Trip = Chronological total of all cargo and ballast legs from the previous discharging port to the last discharging port, possibly complemented with waiting in any of them. In a circular liner-type trip any port will do as the simultaneous first/last port.
Itinerary = Chronological total of all trips by a vessel during a time period, usually year.
Route = Geographical space between two ports or regions, used for traffic between them. Can also refer to the traffic itself.
Trade = Route for a specific commodity.
Multiporting = Cargo collected or distributed at more than one port.
Cargo sequence = Loading – cargoing – discharging sequence of activities.
Cargo cycle, also full sequence = Waiting - ballasting – loading – cargoing – discharging sequence of activities. Waiting can alternatively follow discharging, or be excluded, because it is not indispensable for chartering (speculative ballasting) and because it can be lengthy.

Published at www.ristolaulajainen.com

CONTENTS

TABLES

FIGURES

PREFACE

"Retirement is looming and this report may well be my last book-size study." That is what I wrote four years ago when putting the finishing touches to the predecessor of this volume. How wrong a man can be. The intervening time could well have been only two years had I not been tempted by various tasks in the field of Financial Geography and had I disposed of proper financing. It took some time to convince the Donators of the databases that this was a serious academic study and that their ordinary business might benefit when somebody ploughed the field at depth and looked at what there was to be found.

The bait of the sales pitch was a functional relationship between the freight revenue a ship earns during a year and the character of the region where it takes its first cargo. That sounds absurd because, having discharged the first cargo, the ship is free to sail almost anywhere to pick up a new one. Why should there be a functional relationship? Because, having discharged, the ship may find itself in a position where few lucrative cargoes are readily available. It has to ballast and will spend some of the revenue just earned for that purpose. There are many ways a ship can make the most of the situation, only it cannot choose among cargoes at will in a competitive world but is constrained by the activity of other ships. This being the case it may be equally fruitful to replace explicit decision-making with a Monte Carlo simulator. The simulator cannot specify the credentials of the optimal ship but it is perfectly capable of finding out an average one. This is sufficient as long as the relationship holds in changing circumstances.

The answer cannot be found by deductive reasoning. It is necessary to specify the circumstances and conduct experiments. The original study was about the panamax segment in 1997. It was conceivable that circumstances would be so different in some other year that the 1997 relationship would not hold, or that another segment might react differently whatever the year. It was hypothesized in particular that the slope of the relationship would be related to the level of freight rates. When rates are high there is more scope for differentiation, and vice versa. These deliberations framed the empirical implementation. The year 1997 was average, 1998 rather disastrous, and one had to go back to 1995 to see real profits. That made three years in all, about as much as could be reasonably handled. Actually three years exceeded my handling capacity and the handysize segment was only analyzed properly in 1997. With hindsight, it might have been better to select more recent years. The possibility was never considered seriously. I was still an unwritten leaf in the Donators' books and asking for the most recent data might have aroused suspicions. Again with hindsight, no harm appears to have been done by working with data a few years older.

Once the data had arrived its gravitational pull made itself rapidly felt. It was difficult to focus exclusively on the actual task, to find if a general relationship exists. There were so many related questions that either gave background or might offer a foundation for further analysis, and which had never received proper attention because the empirical data was not easily available. The share of ballasting, the regional balance between loading and discharging ships, the minimum workload of dry bulk shipments, the identification of fronthaul–backhaul routes with relative rate levels, and the clustering of route-specific freight rates into geographically meaningful groups, may be the most promising ones.

Raising these questions in an empirical and not only conceptual context implies a fresh focus in macroeconomic shipping studies. These have been dominated by economists, if not outright econometricians, and their interest lies in time series, preferably global ones. Such series are artifacts. They do not exist in the real life which is made up of charters of individual ships. If it is a spot charter (voyage or trip), the route, the ship's current position, and possibly the intended commodity play central roles. If it is a time charter, the vessel characteristics become paramount. Only in this latter case, after the characteristics have been abstracted to a common reference level, there is substance for the term "global". We stay away from global. Our business is regional. We see regional differences and try to make sense out of them. It is something geographers have always tried to do.

A large project such as this one is always obliged to numerous individuals who take a personal interest in its success. Some might be called freelancers because they put their expertise at disposal out of intellectual curiosity. The donators also had commercial motives. They are organizations and the identity of the person who made the basic decision, "Yes, we will participate", is unknown to me. I know only the contact persons at the executive and managerial level, and their assistants and secretaries, who were prepared to supply data, more data, and still some data – and to explain its idiosyncrasies, again and again. And I wish, cordially, to thank them all although I identify only a few: *Susan Oatway* at Drewry's, *Wally Mandryk* at LMIU, *Cliff Taylor* at Clarkson's, *Anthony Humberston* at Baltic Exchange and *Mark Glasgow* at Island Navigation.

At the technical level, I have leaned on the experience from the original study. As previously, *Olle Westman* and *Timo Laulajainen* have been most helpful in technical details. I have simplified the analytical frame and hopefully made the report more accessible to the general reader. It has not been subjected to the standard linguistic check but should be legible nevertheless.

The product was posted on the Internet in June 2005 to make it accessible to the shipping community as quickly as possible. The busy executive will find the essential information in the two articles published almost simultaneously with this book (see References). *The Swedish Shipowners' Association* have made a publication grant available within the framework of *Lighthouse*, a co-project between the Gothenburg School of Business, Economics and Law and Chalmers University of Technology.

"It all has been great fun. Whether I would do it once again – I do not know. Once bitten, twice shy. No promises this time." so did I write exactly one year ago. Today, however, *The Geographical Foundations of Tanker (Dirty) Shipping* is well underway.

Gothenburg, 9 June 2006
Risto Laulajainen
Professor Emeritus, Economic Geography

”Unfortunately, even Royal Commissions are unable to alter the essential facts of Geography ...”
Sargent (1930), 19.

1 INTRODUCTION

1.1 Theory

Pragmatic, empirical research has forecasting as it ultimate goal. It need not be forecasting in the time perspective, about economic development, future exchange rates, the outcome of parlamentary elections, or climatic change. It may simply be the desire to understand, to sense structures and to create a framework. It is the times-old goal of all science. Representatives of abstract sciences such as philosophy and mathematics may be indifferent to the everyday use of their results. The ability to structure and to understand is the real thing, application is secondary. Applied scientists see matters differently. Structuring and understanding is important because it will upgrade performance, enhance the ablity to claim that given such-and-such conditions, resources and course of action it is reasonable to expect a certain outcome, a typical forecasting setup. It follows that interest is turned to structures and processes which are interesting from the practical angle. We belong to this cohort of scholars.

Most forcasting sources from past experience (Granger 1989; Schwartz 1996; Stopford 1997, 500-512; Ford 1975). But it is not always possible. Biology and physics, for example, may be so innovative that historical precedents hardly exist. Societal development has so many unknowns and so many courses of action that the historical past is of doubtful futuristic value. Analysts try, nevertheless. They simplify the problem and come forth with alternative scenarios. Many aspects of societal development are conditioned by technology. Where technological development is evolutionary rather than revolutionary the prospects to arrive at useful forecasts are radically enhanced. Shipping is such an industry.

Shipping has a long past. Its physical environment has remained practically unchanged. The continents and waterbodies are where they have always been. Excluding few sea canals, most notably the Panama and Suez Canals, their mutual connections have not changed much. What change there has been, has been the substitution of approximate great circle tracks for prevailing winds with the coming of steam power. The exploitation of natural resources has changed more. Some have got depleted or partially destroyed by careless use. Others have been discovered or simply put to use, following technological advance and increased demand. Such measures take time, however, before they have a global impact (Stopford 1997, 253-289). And the underlying truth has remained unchanged: natural resources at large are distributed in proportion to land masses, the continental shelf included (Charvet 1985, 64-86; Giraud 1983, 40-50). What varies markedly is their mix and accessibility. Human resources are different. There concentration is the rule. But once an agglomeration has got established the location is very stable and only time is needed for it to develop the full potential. The end result is a locational imbalance between resource use and their availability, an imbalance likely to persist. The imbalance is corrected by transportation flows which, logically, must also be relatively stable. Therefore it is possible to speak of the geographical foundations of shipping, and dry bulk shipping in particular, because the wet bulk trades are more concentrated in space and historically more volatile.

Dry bulk shipping is widely considered to be close to atomistic competition for many reasons. The market is believed to consist of a multitude of small actors, whether shipowners or charterers, too numerous and small to influence the price, i.e. freight rates, substantially. Rather, the rate level is set by the aggregate action of all market participants. Their action is greatly facilitated by the ready availability of market information and the relative ease of overall entry and exit, manifested by the vibrant market for second-hand vessels. It is believed further that entry to geographical submarkets, i.e. trades, is without friction. But there are also caveats. Some chartering is not for single but several cargos, or lengthy periods. There are anecdotes in the trade press about the concentration of market power. There is the analytical inconvenience that instead of a single market clearing price routes have their specific freight rates, i.e., local prices. These have a ready explanation in the tonnage shortage of loading areas and oversupply of discharging areas, a direct consequence of imbalances between resource availability and use. Does this mean that there is no single market but several local ones? Does it follow that the premises of atomistic competition apply in each of them with different force? At least it is to be expected that differences in rates between routes can be tied to the relative size of tonnage imbalances and their mutual locations, i.e., distances. If this idea is corroborated and the rate structure has stability over time, a kind of theory with some forecasting ability has been found. The theory can be used for evaluating the efficiency of the market, a traditional topic in economic research, only this time from a geographical perspective. And if the efficiency of the market cannot be established attempts to benefit from its inefficiency will result.

That, in all simplity, is our theory. The shape in which it is outlined above is an idealized reconstruction. The true origins are quite different.

1.2 Origins

The origin of this report is a three-year project to create an operational tool which ship managers could use when selecting between chartering alternatives (Laulajainen, Holgersson and Strömberg, or LHS 2001). It soon turned out that the information available in the public domain was too sparse and partially unreliable to allow the completion of the idea. And information in the private sphere was too dispersed to be collected and far too valuable to be released.

> As a flamboyant broker characterized the business: ”The essence is to lie, milk as much information as possible and release as little of it as possible”.

Attention was consequently turned to the maximization of ship revenue through route optimization, an inherently geographical problem. Bazaraa’s algorithm (Bazaraa et al. 1993) was used on two external conditions: that each route (loading-discharging pair of ports) had at least 1.0 cargo each week, alternatively that the minimum traffic density was only 0.1 cargo per week. Freight rates were annual averages (Table 1.1). Handling charges at port were overlooked. Distances were fixed. In short, it was a very simple deterministic optimization. Variability to the results was introduced by repeating the optimization 21 times, each time adding a new constraint which prohibited the previous solution. The overall result was a stepwise worsening object function (LHS 2001, Table 5.4 and Fig. 5.4).

This left the more fundamental question about differences in freight rate levels and, indirectly, market inefficiency open. The itineraries included loading ports of widely differing freight levels. Routes beginning from large export regions with a tonnage shortage ended, almost by necessity, in large import regions with a tonnage surplus and depressed outward rates, or ballasting. It was possible to redo everything with a different port as a starting point. But that was not a real solution because any itinerary was basically a loop with unchanged revenue irrespective of the identitiy of the first port. Counting the port frequencies of the best itineraries did not help either because the frequencies accumulated heavily in few ports irrespective of their contribution to the revenue. The deterministic approach necessarily supressed the competitive elements of the marketplace. It was thought advisable to take the effort closer to the way shipping markets were believed to function. Newbuilding and second-hand markets were then left outside the analysis. The tool was a stochastic market Simulator, based on identified panamax cargo flows in 1997. The simulation was by vessel and the results were recorded, after a warm-up period, from the day when loading started in the first port and continued to the end of the year. With several years hindsight it is not at all certain that the decision was a good one because the long and uncontrolled simulation period introduced a lot of noise, possibly white but anyway. It should also be observed that we try to measure market inefficiency, a macroeconomic problem, by microeconomic means. All the same, the Simulator, with some modifications, will be used now for the second time although in a more extensive setting.

Table 1.1 Freight rates used in the optimization

	ECNA	ECSA	CONT	SAF	IND	FE	AUST	WCNA	WCSA	CHL
ECNA	7600	10000	13381		12814	17260			7000	
ECSA	10150	7000	14369		13024	16904			13430	
CONT	6400	8000	6000		12141	12632			12500	
KAMS	8290		11407	12000						
SAF		9000	10576			14507				
IND			9600		6000	11069				
FE	5000	7000	6305		11600	6314	6000	8200		
AUST	8500	5000	8880	6000	11687	13766	7000	7000		7000
WCNA		8300	9406		12403	13625				
CHL	7000									

Source: LHS (2001) Table 5.3.

The Simulator is very simple, conceptually and structurally. Ports are consolidated into regions (Fig. 1.1). Ships within a segment and cargos are of equal size. New freight rates are calculated for each route each Monday. Cargos are generated for each route every weekday. Each uncommitted (free) ship scans within a radius of two legs of any type, calculates the profitability of each identified cargo by the approximate formula (for exact one, see LHS 2001, 213, 230), and ranks them (bids for them):

Daily Rate = (DR-i + DR-j) / (TM-ij + TM-jk)
in which
DR-i = Rate-ij * Cargos-j * TB-j * TM-ij
DR-j = Rate-jk * Cargos-k * TB-k * TM-jk
TM-ij = Sailtime-ij + Porttime-j
TM-jk = Sailtime-jk + Porttime-k
TB-j = $(\text{Loads-j} / \text{Arriv-j})^{0.5}$
TB-k = $(\text{Loads-k} / \text{Arriv-k})^{0.5}$
i, j, k = region identifiers

One cargo from the most ”profitable” port (actually region) is allocated to the ship which placed the bid. With several equal bids the allocation is made randomly from the uniform distribution. If a ship fails in all its bids it remains in the port where it happens to be, the last discharging port. The rationality of waiting in a port as compared to ballasting to a more promising location is evaluated daily. Having been allocated a cargo, a ship becomes committed. It ballasts to the loading port, if not aready there, loads, sails to the discharging port and unloads there. Loading and discharging are possible only on weekdays but a ship need never wait for a berth to load or discharge. Arrival to a port is possible also on Sunday. Each activity, rating, cargo generating, sailing, loading and discharging contains a random element. For rates, the random element is taken from the normal distribution with the rate estimate as mean and one-half of the SEE as standard deviation. Cargos are generated from the Poisson distribution with the average number of cargos per day as the only parameter. The times used for sailing, loading and discharging are taken from truncated gamma distributions based on average times spent for each activity (norm times). These random times are decided at the time when the bid is accepted. The only elements which lack randomness are freight rates during a warm-up period. Structurally, the simulator is disaggregated into modules: handling time, finding freight rates, generating cargos, comparing routes, allocating cargos, and handling ships. The modular structure helps in understanding how the model works and implementing changes. Further detail is available in LHS (2001), 113-145.

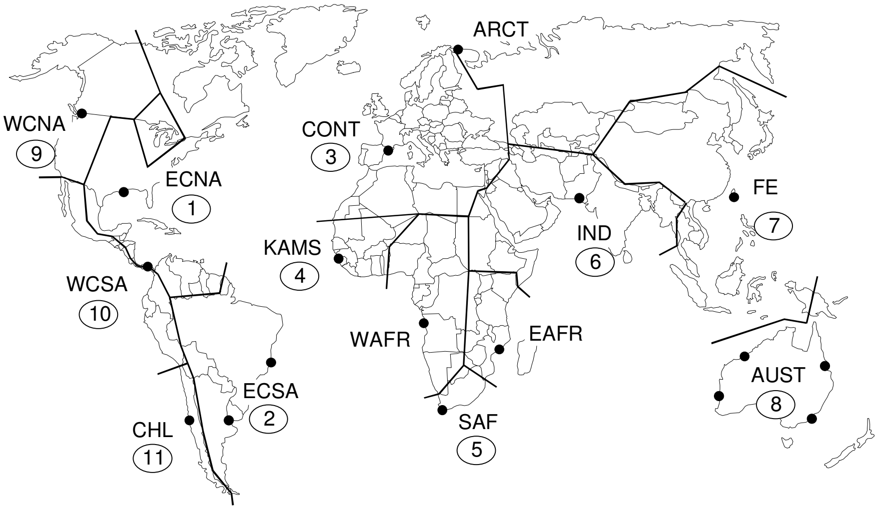

Figure 1.1 Regions and their reference points

Source: LHS 2001, Fig. 3.5.

The analysis of the results rapidly converged to classifying the ships by the loading region of their first cargo. Such a choice may appear strange but is entirely rational. Cargo legs with high rates are often followed by those with low rates or outright ballasting, and the other way round. Therefore it is important that analytical calculations are extended beyond the first cargo leg, the normal time perspective of a ship manager making chartering decisions. He/she probably has the next move in mind but avoids a firm commitment because more cargos will become available with time and deadlines in the distant future are risky. But he/she has a mental picture about the next-move possibilities, a picture which an analyst consolidates into a matrix of transfer probabilities, and it affects his/her decision. These probabilities can be chained to long sequencies, itineraries. When weighted by freight rates they give an estimate of future revenues available for ships taking their first cargo in this particular region. The difficulty, of course, is the great number of alternative itineraries. An average panamax ship makes about 7 cargo legs a year and, with 10 regions to select, the total number of leg combinations, i.e., itineraries is 86,400 (LHS 2001, 51). A short cut through this maze of competing alternatives is to let the Simulator make the choice. The choice can be deterministic but it is probably more fruitful to integrate stochastic and competitive elements because the Real World is full of them.

Among the numerous attempts for explanation one gave promising although unexpected results. The daily freight revenue during a year of ships taking their first cargo in a particular region was a linear function of the the region’s identification number plus a dummy indicating whether the region was a net exporter or importer (LHS 2001, Fig. 6.4). We label this basic function as **Revenue Gradient**. The use of the identification number as an argument was immediately replaced by a ratio called Revenue Potential, which did the same job.

$$RP\text{-}i = \sum [(DC\text{-}i * ST\text{-}ij - DB\text{-}i * S\text{-}ij) / (AC\text{-}i * ST\text{-}hi + AB\text{-}i * ST\text{-}hi)],$$

totalled over all trading connections, in which:
RP-i = Revenue Potential
DC-i = Departures in Cargo
DB-i = Departures in Ballast (technical variable)
AC-i = Arrivals in Cargo
AB-i = Arrivals in Ballast
ST-ij and ST-hi = Sailing Time between regions
h, i, j = region identifiers.

The ratio tells that revenue is generated by departures in cargo, to be related to all arrivals, i.e., tonnage demand and supply balance and through it freight rates. The subtraction of DB from DC extends the range of observations and thereby facilitates explanation but does not change the rank order nor the sign. Weighting by sailing time gives an idea of a region' s relative closeness to its trading partners, a component of its revenue-generating ability. It follows that typical export or import regions differ primarily by their geographical location, a split between the Atlantic and Pacific basins. The split is handled by a dummy which replaces the, by now, defunct export/import dummy (Fig. 1.2). The fit is almost perfect (R-sqr, adj = 0.939).

Against this background the classification by the first loading region becomes understandable. The results also appear reasonable although the dummy differentiates the Atlantic and Pacific Spheres with only $692/day, much less than the average difference between Atlantic and Pacific loadings $2,922/day (calculated from monthly figures). This is so because the former figure is a mixture of both. Otherwise it can be speculated that the relationship reflects first and foremost differences in outgoing rates between regions (Table 1.1). To an extent this is true. Outgoing rates from the east coast of Americas tend to be higher than from the west coast or Australia, to the same destinations. But the Continent and South Africa are more in line with the Australian rates than the east coast rates and still the relationship holds. The possibility to find new cargos after unloading and the length of ballasting legs is the explanation. Exactly these facts have been included in the Revenue Potential and cannot be seen from plain freight rates alone.

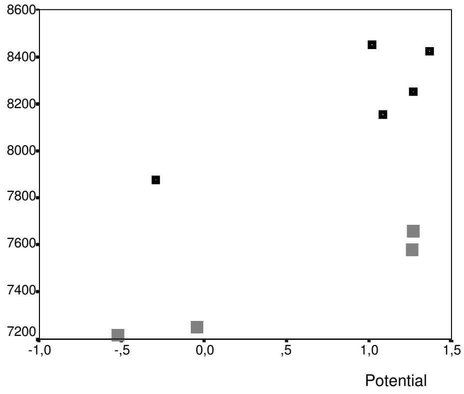

Figure 1.2 Daily revenue ($/day) as function of Revenue Potential

Legend: Small marker = Atlantic; large marker = Pacific. Only regions 1-9.
Source: LHS 2001, Fig. 6.7.

Taken at face value the existence of a gradient is unexpected, to say the least. In a reportedly atomistic market with superficially perfect mobility of production factors it simply should not exist. Still it does. It is coined ”inefficiency”, here and elsewhere (Laulajainen 2003; 2007). Whether this is the correct term can be questioned. In classical theory, efficient market under polypolistic/perfect/atomistic competition means that marginal cost equals marginal revenue which is equal to price (freight rate). The principle is clear but its application in shipping is complicated because marginal cost varies heavily in time, and both it and marginal revenue contain fixed items (Evans 1994). Fuel cost, affecting optimum speed and through it a ship’s revenue earning capacity, is the obvious example of a variable cost item. Port time is a major fixed, or almost, item in operational calculations. On some routes and ship segments ballast bonus is routine, fixed for a particular charter but variable otherwise. The closest thing to an empirical real-world test may originate from *Glen* (1990) who compared four tanker routes by ship size segment as to gross profit and its variation. The thesis was that structural change in the market catalyzes different response by route and ship segment. That, exacly, was the case. Does it mean that the tanker market was not polypolistic as some believe, or was the test inappropriate for such a question. It was not fully appropriate but the variation in profit and the differences between routes and segments were too large for the efficient market hypothesis to fit in comfortably.

To prove our case, the regularity of the Gradient is a crucial feature. As will emerge later on the regularity is not as complete as Figure 1.2 suggests. But it is firm enough to support our case. And the case is that the panamax market 1997 was not efficient. Had it been efficient in the classical sense of the word the marginal costs should show a pattern similar to the Daily Revenues. These Revenues are simply trip charter rates averaged over all types of ship activity. No figures about marginal costs have been shown and their reliable derivation by region may be a difficult task. But it is extremely unlikely that if such a cost pattern will ever be derived that it will show the kind of regional regularity as the Revenue Gradient does. This being the case the thesis that the gradient supports an inefficient market hypothesis is a plausible interpretation. It can be expanded to a case where no regularity, and consequently no gradient, exists. If so, perhaps the organizing principle is something else than the Revenue Potential. But whichever the organizing principle, the outcome should be either a horizontal line of Daily Revenues, or a scatterplot in which the marginal costs approximately reproduce Daily Revenues.

Whatever the gradient's theoretical underpinnings, it exists and is too pronounced to be wiped away by simplistic arguments. Some have been raised in a previous article and found insufficient (Laulajainen 2003). Far more relevant is the question **whether a broadly similar Revenue Gradient will emerge also in other years and dry bulk segments** than 1997 and panamax alone. A rapid check was conducted within a year of the main report (Laulajainen 2002). It applied to 1998 and was based on data from the Drewry Monthly (Appendix 1.1). The selection of 1998 was based on the turnround of the freight rate differential between the Atlantic and Pacific Spheres (Fig. 1.3), prompted by the economic crisis in the Asia-Pacific Rim, and the reduced draught of the Panama Canal, due to a prolonged drought (el Niño effect; Dewitte 2001). It could be hypothesized that such rare and dramatic events would leave a pronounced imprint on the Gradient and possibly unseat it completely. This was not quite what happened. The Gradient was affected but it was not unseated. The weakness of the test was that the cargo flows were not estimated from ship movements but were based on published freight contracts (fixtures). It is well known that fixtures do not reflect accurately the underlying cargo flows (e.g., LHS 2001, Fig. 4.1). To test the significance of this fact some radical assumptions were made about the dislocation of contemporary cargo flows following, for example, an imaginary closure of the Panama Canal. Again, the relationship was affected but not unseated. This still left open the possibility of a radical change in the rate structure, a research avenue too wide to be followed at the moment. The overall gut feeling was that the crucial thing is the rate structure, the relative levels of rates, and that very substantial shifts in cargo flows are necessary before the structure is upset. This is a fortunate observation in so far that the volume of data for reliable conclusions about freight rates is much smaller than about cargo flows.

> The content of the Atlantic/Pacific Ratio of monthly trip rates for loadings in the respective ports is understood to be as follows. The underlying rates are broker and consultant estimates based on actual fixtures. Loadings in Indian Ocean ports and South Africa are not included. Australia is part of the Pacific. One or two heavily trafficed routes in each direction and size class are considered sufficient for calculating the ratio. The exact identification of these routes is avoided although the guess is not too difficult. In any case, the routes are less representative than the fairly comprehensive calculations which this report uses for calculating rate indices. The ratio is currently calculated for new 160,000 dwt capesizes, 70,000 dwt panamaxes and 40,000 dwt handymaxes. It is possible that smaller ship sizes have been used in the past. The capesize ratio is actually calculated from earnings rather than fixtures because a large share of the tonnage is owned or time-chartered. The ratios show a broadly similar pattern but their Pearson correlations are not too impressive (Fig. 1.3). They are only weakly correlated with the Baltic Freight Index which will play a central role in organizing the sprawling rate quotations into useful rate functions.

Handymax 0.725 / 1.861 / 1.336 / 0.269

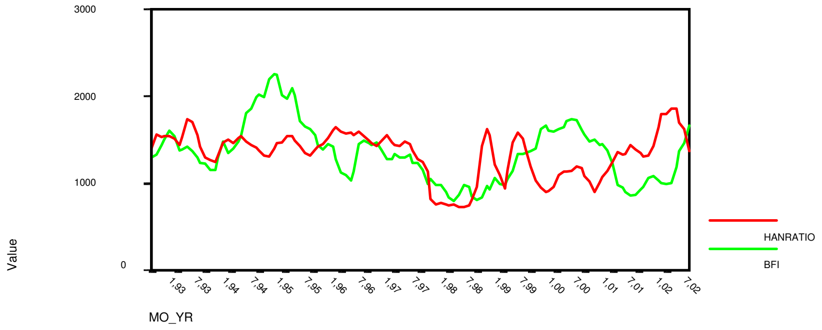

Panamax 0.619 / 2.021 / 1.289 / 0.260

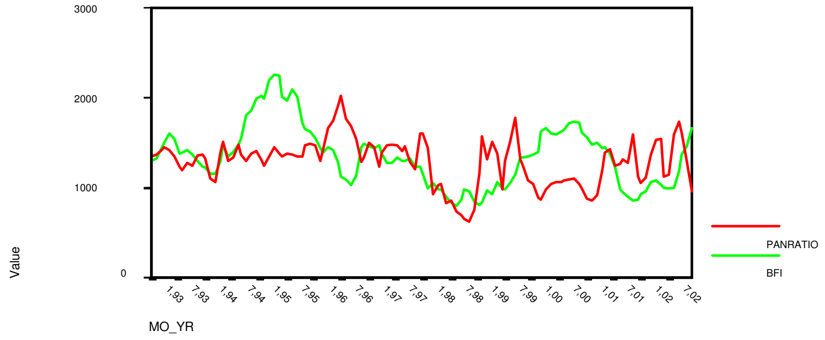

Capesize 0.771 / 1.863 / 1.187 / 0.194

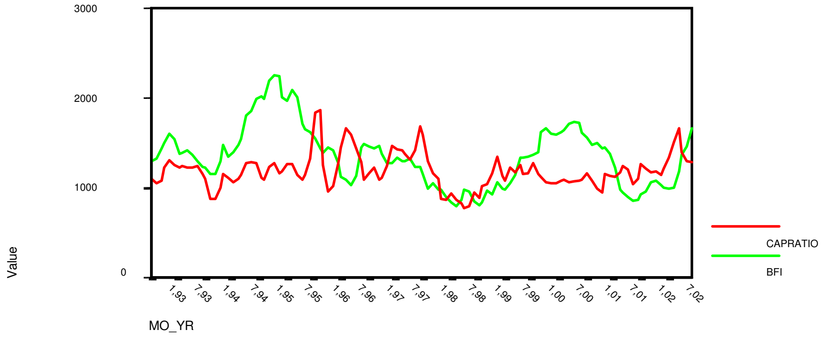

Figure 1.3 Atlantic/Pacific Ratios, 1993-2002

Legend: Minimum / maximum / mean / std deviation. BFI 799 / 2258 / 1349 / 337. Atlantic/Pacific values multiplied by 1,000.

Notes: Pearson correlation coefficients: Han/Pan 0.721, Han/Cap 0.528, Pan/Cap 0.555; Han/BFI 0.183, Pan/BFI 0.092, Cap/BFI 0.109; (BFI = Baltic Freight Index).

Source: Courtesy of Clarkson Research Studies.

The current effort is continuation of this unpublished study. Thanks to the donations by the *Drewry Shipping Consultants* (fixtures) and *Lloyd's Marine Intelligence Unit* (ship movements) it is possible to extend the study of freight rate structure over 1993-2002 and of cargo flows over 1995 (high rates) and 1997-1998 (average and low rates). The rate level may have significance because the scope for a steeper Gradient is larger when rates are higher. The fixture data goes down to the 10,000 dwt ship size while the movement data ends at the 25,000 dwt mark. The latter donator had the definite, and as it turned out very wise, opinion that ship movements below 25,000 dwt would be far too numerous to be handled in a reasonable time. Consequently the total number of movement data lines remained at 335,000. This drew the lower limit also for the fixtures, about 50,000. Both datasets needed much manual handling to be really usable and finally two years of handysize movements were left unfinished. It was not only the exessive time spent but also the worsening quality below the 35,000 dwt mark which led to this decision. When, in addition, incomplete and superfluous lines (incl. time charter) had been excluded some 117,000 and 40,000 lines were left as the core data base.

1.3 Previous research

The purpose of this literature review is to get some background for our effort, which may be rather unique. Seasoned academics and consultants declined to have seen anything like it, not even in their own well-protected vaults. Perhaps so, although commercial secrecy can be very strict. *Ronen* (1981) comments about the Management Information System of an anonymous shipping company, intended to rationalize its spot market chartering and really succeeding in this task:

> ’For quite a few years the mere existence of the MIS system was kept secret due to its uniqueness.”(Ronen 1981, 130)

From that angle, an excursion to previous research appears waste of time. From a more enlightened angle, it goes without saying that many of the problems we are going to face will be identical with the problems faced by previous scholars and it makes good sense to to get acquainted with them in advance. In line with this approach only studies which will be of direct use, either by giving background or offering important operational parameters, will be discussed. In other words, this overview is not an antology.

Four types of studies are of relevance here:

- historical accounts about strategic chartering
- operative decisions in
  - industrial surroundings
  - free spot market
- freight rates
- port operations.

Historical accounts are offered by *Burley* (1960), *Sargent* (1930), *Isserlis* (1938), *Manners* (1971), *Couper* (1972), *Lundgren* (1996) and *Nossum* (1996). The first three refer to the time when the British merchant fleet made up one half and over of the global one. Its mainstay was outward shipments of coal and homeward shipments of raw materials and food, particularly grain. Because each outward trip could not be balanced by a matching homeward trip cross-trade became necessary. *Burley* describes the cross-trade of New South Wales coal in about 1850-1914. Its export rested on the serious imbalance in Australia's foreign trade. The country was building its infrastructure and imported large amounts of capital goods from Britain. Own exports did not fill the cargo space for a direct backhaul and the surplus tonnage had to look for return cargos farther afield. These were available in Asia, on the American west coast and the Pacific islands. Rather than ballasting there a coal cargo was taken from NSW. The coal was of good quality and mined close to export ports but its competitiveness depended on a low freight rate. In Asia, between Calcutta and Shanghai, it was able to push back British coal but faced subsequently keen competition by Japanese and Indian colliers. On the American west coast, the first destination was San Francisco where grain cargos for Europe were available. When grain fields gave way to orchards and competition by Canadian coal mounted, the loss of business could be replaced by Chilean nitrate exports, and when this traffic withered because of the nitrogen synthesis the Pacific islands, from the Philippines to Hawaii and Tahiti, offered a partial relief. The lasting remedy, however, came from Australia's wheat export which finally absorbed the surplus tonnage.

Sargent's analysis about British shipping in 1927 belongs to the evergreens of the geographical shipping litterature. The world is divided into seven areas: Europe, South Africa, Australia, India, Far East, North America and South America. The traffic between them is quantified, split into laden and ballast traffic, and further into British and foreign flags. The reasons for the frequent imbalance in trade and tonnage are scrutinized carefully and the ensuing cross-trades estimated to achieve a better understanding of the global circulation of British vessels. The conclusion is that, sooner or later, they all return home. The effect of the Suez and Panama Canals to trading and shipping patterns is discussed, with a particular view to the competition between European and US east coast manufacturers. Parallel to these general features there is much geographical detail about weather conditions and their influence on ship speed and sailing routes, about bunkering stations, and the need to call on way home at Canary Islands for bunkers and telegraphic advice about the discharging port. The study is exceptional also in making comparisons with 1912, discussed in the first edition (1918). Technical development has made some of the detail obsolete but left the principles intact.

Isserlis is more interested in freight rate indices and the contribution of shipping to British invisible exports than geography as such. The index data is from 1869-1936 on 211 homeward and 112 outward trades (called routes). Calculations of freight income concern the year 1935 when 12,491 applications for freight subsidy offer a reliable and practically complete dataset about cargo trips on British keels. The geographical content is in the trades for which the origin, destination and type of cargo are given, sometimes in amazing detail of port and commodity. The flip side of the detail is that the time series are on several occasions only few years and often fragmentary. The application data from 1935 maintains the detail, and even adds to it. The vessels are split into those of 3,000 gross register tons and above, and those below this mark, the latter being used primarily in European traffic. Both the freight income and the weight of cargo is given by trade. The geographical detail allows, at need, the derivation of distances, and from them, approximate daily freight rates. Although the study applies only to British vessels and the internal dynamics of the year remain concealed it is regrettable that the data may never have been used in a geographical analysis.

Manners' monograph about the global iron ore trade contains all the ingredients of a sophisticated shipping study: freight rates on the major routes, corresponding distances, turnround times in the major loading and discharging ports, and some estimates of port charges. All what is needed is their integration into a numerical system, a tall order for the computer technology of the time. For computing maybe, but not for naval architecture which was developing ever larger, multipurpose vessels. The novel opportunities offered by them are thoroughly discussed and the logic has not lost anything of its persuasiveness during the past decades.

Couper's focus is in the post-WWII bulk carrier trade in fairly regular shuttle service under long-term contracts. Traditional tramping remained the realm of spot-chartered smaller tonnage which benefited from the seasonality of many cargos, iron ore, coal, timber, and particularly grain. Ocean-side storage capacity and powerful ice breakers have levelled off most of the seasonality which survives only in the agricultural exports of the southern hemisphere. But the fundamentals have not changed so much after all, e.g., the temporal profiles and relative locations of freight rates from US Gulf to India and from La Plata to the Continent were practically the same in 1997 (weekly) as during 1963-1967 (monthly, p. 96).

Lundgren tracks trends in freight rates, ship sizes and transport volumes since the early part of the 19th century with emphasis on the post-WWII years. Volumes have grown and rates declined, although not always linearly. Rates, in particular, have declined so low, absolutely and relative to commodity prices, that the end of an era is in sight. The driving force has been technical progress in the shape of larger vessels and deeper ports, and the logistical outcome has been a reorganization of trade flows, short interregional flows having lost absolutely and relative to long intercontinental ones. The gains have been achieved at the expense of a financially weakened shipping industry plagued by excessive swings in profitability and capacity.

Nossum follows similar trails but because his contribution is a book and not an article he can be more specific. There are volumes of descriptive geography investigating, commodity by commodity, the type and size of ships transporting them, identifying the export countries and loading ports with their tonnages, and repeating the same about the import countries and discharging ports. When the locations are few he also gives flow matrices. This is normally done at 5-year intervals. For our study, most of the material is superfluous, except by shedding light on cargo-port combinations. But if we would restructure our effort and put emphasis on the level of freight rates and the emergence of trade flows, in a kind of equilibrium exercise, the situation would be completely different. Then, Nossum's statistics would constitute the foundation of the effort.

Desriptive material naturally abounds in consultant publications. The most well-known for its geographical content may be World Bulk Trades, and its co-publication World Bulk Fleet, published annually by *Fernresearch*. World Bulk Trades displays in tables and figures the flows of major bulk cargos, iron ore, coal, grain, bauxite & alumina and phosrock, between major world regions. Tonnages are complemented by tonne-miles and average distances. This publication is based mainly on LMIU data, supplemented by the Publisher's own intelligence and public trade statistics, goes currently down to the 60,000 dwt level and allows, at least in principle, the following of individual vessels by cargo, port, turnround time and routing (Hammer 2005). Since the conglomerate comprises also a brokerage with a steady flow of fixtures, Fearnresearch has both the data and the tools for sophisticated geographical analysis.

An inherent potential of similar, or possibly larger, magnitude is available also at *Lloyd's Marine Intelligence Unit* (LMIU), our partner. LMIU is an independent company but originally part of the Lloyd's organization with insurance and ship register as its main lines of business. Shipping intelligence, knowledge about vessel condition, whereabouts, activity and working environment, were essential for the successful writing of marine insurance. A network of 2,000 agents and subagents worldwide and a strong brand name make LMIU an unrivalled data source. The lower limit for this study was agreed at 25,000 dwt but the company did discuss, at complete ease, the alternative limit of 10,000 dwt, also. The limiting factor then is the rapid deterioration of the quality of data at smaller sizes. The vessels are less visible, their port calls more frequent, and the ports smaller and more out-of-way, problems faced by all intelligence agencies. The other, academically important, difference between LMIU and Fernresearch is that LMIU does not publish general flow statistics for everybody to consume but only prepares individual reports to professional customers.

Official maritime statistics are less useful. The last worldwide figures covered the years 1972-1975 and came from print with a considerable time lag (Maritime Transport Study 1980).

The other type of data needed in this report is about freight rates, available in fixtures, abstracting verbos charterparties. Fixture data is available at brokers who sell it futher to consultants. All large consultants procure fixtures from numerous brokers and an outsider cannot see any real difference between them in the amount and quality of data available. The difference is preferably how that data is published. *Drewry's* monthly and quarterly publications are ideal for this study. They give the actual fixtures, consolidate each month rates of some twenty dry bulk trades, list important exporters and importers, give figures about fleet size, newbuildings and scrappings, and general background data such as breakeven rates and canal fees. As usual, there is a short section of market commentary. They offer most of what is needed, in a concise and reader-friendly form.

This takes us to the way **how ships and fleets are operated** in practice. A most entertaining treatise orginates from *Utstein Kloster* (1948; 1952), a narration about the travels of a small all-round tramp in the pre-satellite, non-standardized world. Having read it, one begins to understand how many pieces must fall in place before a trip becomes a success. *Packard's* (1996) routing calculations provide it with a welcome envelope. A practical conclusion from these reports is that it is unwise to commit a ship on the spot market very far ahead. The freight rate is less of a problem, the schedule instead is. Arrival too early to a loading port is tantamount to unpaid waiting and arrival too late will lead to damages and possibly cancellation. This ”very far ahead” is operationalized in this study as ”beyond two legs”, whether in cargo or ballast. At a higher level of abstraction, *Fleming* (1978) compares, in a hypothetical situation, the relative efficiencies of shuttle service, combined carrier, and shipowner consortium vis-a-vis simple spot charter which well gives maximum flexibility but is a poor hedge against a market slump. He also remainds that at the end of a long ballast leg there usually are cargos paying high freight rates, the very idea of this study.

At a more generalizing and systematic level, the two review articles by *Ronen* (1983; 1993) are most worthwhile reading. Ronen sources from his experience as a consultant to industrial corporations with tied fleets, but the message has relevance also in the spot market. In ship routing, problems are poorly structured, weather, congestion, strikes in ports, breakdowns of equipment and engine bring havoc to schedules, all cargos are not known in advance, destinations are changed when at sea, ships are individuals and must be handled as such, capital investments matter more than operational decisions, industry is conservative and resistant to outside advice, most analytical work is conducted for owned company fleets and tonnage hired under long-term contract. Standard vehicle routing methods cannot be fully applied, and routing & scheduling problems are intimately connected with those of fleet size and structure, themselves intertwined with strategic chartering decisions, made with certain markets in mind, each consisting of specific routes with specific parameters about route length, port facilities, cargo balance and freight rates. In liner traffic, cargo density (tonnes/shore km) is a key parameter because it affects service frequency (Gilman 1975; Jansson and Shneerson 1985). Itinerary configuration matters. For example, there is a maximum acceptable ratio between the collection or delivery range and the coast-to-coast distance. These aspects have relevance for bulkers, too, when they call at numerous ports as in the grain trade, or run liner-type operations as many handysizes do. Although managerial articles seldom address their examples in decoded geographical terms, no great imagination is needed to sense real-world equivalents (e.g., Berg-Andreassen 1998; Xie et al., 2000).

At the aggregate evel, this leads to fleet time use and productivity. Globally, we can consult the studies by *Kalindaga* (1990) and *Stopford* (1993). Both are about tankers which may have shorter port times than is possible for dry bulk carriers of similar size. Otherwise, the ”only” difference is between the classification setup and shipping cycle (Table 1.2). In general terms, 80-90 pct of the fleet is in operative use, effective or not, and the rest inoperative. Waiting can then be considered operative because vessels are most likely bidding for a new charter. The split is important because the Simulator comprises only operative use (e.g., Table 4.3). It is also interesting to note that inefficient use can be almost twice as important as the conventional indicators of inactivity.

Table 1.2 Tanker time use, early 1991 and 1988

Stopford		days		pct	Kalindaga		mill. dwt	pct
Trading			288.8	79.1	Effective use		163.2	71.1
Port		40.5				n.a.		
Loading	16.0					n.a.		
Discharging	23.7					n.a.		
Bunkering	0.8					n.a.		
Sailing		248.3				n.a.		
Cargo	137.0					n.a.		
Ballast	111.3					n.a.		
					Ineffective use		41.9	18.2
					Slow steaming	25.0		
					Part cargo	16.9		
Non-trading			76.2	20.9	Non-trading		24.7	10.7
Waiting	7.3				Waiting	2.4		
Storage	32.4				Storage	6.9		
Repairs	24.7				n.a.			
Incidents	6.8				n.a.			
Laidup	4.9				Laidup	5.4		
Total			365	100.0	Supply		229.8	100.0

Sources: Stopford (1993), 28; Kalindaga (1990), Table.

Information about speed in different conditions is most useful (Stopford 1993). There is no substantial difference between average speeds in cargo and ballast, 11.4 and 11.7 knots (p. 348), respectively, but the ranges during short and long trips are quite different, those below 1,000 and over 10,000 nm measuring 7.3-10.4 and 12.1-12.9 knots, respectively (p. 75). The latter observation is explained by the time lost on heavily-trafficed shipping lanes and at port approaches. The similarity of speed whether in cargo or ballast leads to the conclusion that a ship is nearly as costly to run whether empty or full. When it is sailing, yes. In port, the day cost is only one-third of a sea day. Some ship itineraries are disclosed port-by-port but not analyzed in detail, save for few examples about multiporting. They may have catalyzed the remark that operational patterns are much more complex than anticipated (p. 342, note 5).

The insight gained from the 1998 year study with its radically changed trade flows was that the flows alone cannot upset the Revenue Gradient (Appendix 1.1). For that is needed a change in the **freight rates**, or rather their relations, too. As will be shown in a separate study, the relative rate of a route depends on the distance-weighted tonnage balances of both the loading and discharging regions (Laulajainen 2006). The absolute revenue is a separate question because it rests ultimately on shipowners' need to earn a decent interest on investment. This study does not go that far. All we look for is how previous studies have solved our current problems, such as the handling of ballast bonus in trip charters, unclosed ballast legs in voyage charters, the conversion of cargo tonnes into dwt, and so on.

Pettersen Strandenes (1984) and *Kavussanos and Alizadeh* (2002) have used voyage charters, convertered to monthly trip rates. *Glen and Rogers* (1997) and *Tvedt* (2003) have sharpened them to weekly figures, based on information from three renowned consultants. Dr. Alizadeh explains how the conversion happens (Alizadeh 2003):

> I have used spot earnings, which are published by Clarkson Research Studies ... these are based on ... voyage charter revenue minus the voyage cost on a round trip basis for ... different routes.
>
> ... essentially the average earnings for a particular size vessel over a month and hence can be compared to time-charter earnings ... one should note that time-charter rates are also average earnings of a particular size vessel over a month.
>
> ... the way Clarkson calculates the earnings for spot market also takes into account any extra Ballast Bonus, which is paid even for voyage charter contract.

The related question is how it is possible to give worldwide rates for particular ship segments, recalling the widely varying rates of individual routes. The answer is, by weighting them with the number of fixtures. How well these fixtures then cover the whole fixture population is seldom discussed in the litterature. *Glen et al.* (1981) point out, nevertheless, that charters tend to occur in clusters and do not always provide continuous time series. Another question which has been raised recently is whether the US dollar should be replaced by the Japanese yen (Tvedt 2003, 228).

Research on freight rates is dominated by economists. They look at the factors influencing rates, the relative forecasting ability of spot and time rates, their stationarity, cointegration, term structure, optimal split for a risk-averse shipowner, and so on. *Veenstra* (1999a) gives an outstanding and up-to-date introduction to the thinking. It is unusual that a shipping economist descends the geographical scale and gets involved with particular routes, except for substantiating some general topic such as grain shipments. *Binkley and Harrer* (1981) and *Park and Koo* (2000) have made this choice, the first with the help of some 9,350 voyage fixtures from major exporting regions to major importing ones during 1972-1976, and the latter by using annual rate series from the US export regions to a great number of importers at four-year intervals beginning in 1987. Binkley and Harrer resembles our study because of the geographical detail, more than 1,000 ports are aggregated into 16 exporting and 34 importing regions. There are dummies for seasonal effects and handling terms. Ship age, not available in published fixtures, is overlooked. The results in a nutshell are that larger cargos and higher traffic densities give lower freight rates, because of better port facilities and backhaul possibilities, whereas an increase in distance will raise them although at a decreasing rate. Seasons play a subdued role. Park and Koo is, from our point, very similar. Shipment size and frequency behave as above. Dummies for the 4 export and 12 import regions, four types of grain, and the seasons are mostly significant with correct signs. Multiporting gets a significant coefficient in the first two years but not thereafter, which happens to be our study period, also. *Kavussanos* (1996b), building on the research tradition of *Tinbergen* (1934) and *Armington* (1969), and using published UN data (Maritime Transport Study 1980), estimates bilateral export price elasticies for dry bulk cargo between 30 trading regions in the mid 1970s. The effort is midway between international trade and shipping studies, by geographical breakdown and cargo selection shipping, but by paradigm choice international trade. It is important to note that when trading regions happen to be parts of large continental countries, changes in internal transports may affect the split of export between ports without a noteworthy change in export price (Hauser 1986; Ungar 1973).

This study adopts the shipowner angle. What happens in **ports** is interesting only when it directly affects ships. A port's technical standard, the time spent there, and the attached costs are the main themes. Most of them will be abstracted away in this study. The depth of the channel and the berths lose their relevance because the area units will be large enough to have most kind of port and the analyses are conducted by ship segment with the dwt as a classification variable.

Loading and discharging rates will be used as proxies for cargo handling time. The remaining time, for red tape, cargo & berth queuing, bunkering, etc. is difficult to generalize and therefore overlooked. Port costs are abstracted away by using trip charter but the time spent at port remains intact. The workings of the real world are too detailed to be caught by our coarse mesh and the detailed study by *Heaver and Studer* (1972) about the turnround time of 1,305 grain ships in Vancoucer during the 1960s should be convincing in that respect. It is fortunate, of course, if the spectrum can be reduced to some simple rules. Laytime, the days reserved for loading and discharging, appear a possibility. *Mokia and Dinwoodie* (2002) have analyzed over 2,000 tanker charterparties for ocean trips during 1997-9. It is possible to choose between laytimes of varying length by selecting a suitable charterparty, a longer laytime being compensated by a higher freight rate. In practice, the time actually used need not tally with the time decreed in the charterparty at all. Overflows are particularly frequent where shallow ports call for lightering, where congestion is endemic, or where environmental legislation necessitates extra vigilance. Generalizations are not easy. There are ports with slow official handling rates but which are very efficient when speeding money is paid to the longshoremen. Indian ports are avoided by the larger, more efficient ships because delayed investment and inflexible labor practices lead to long turnround times (Bose 2000). Sporadic delays appear in most large bulk ports: Vancouver, Prince Rupert, Baltimore, Norfolk, Newcastle, Lagos, Dalrymple Bay. (Daniels 1997; Evans 1994, 323; Goodman and Lenze 1988, 140; Marsh and Morrison 2005). Breakdowns in infrastructure, strikes, weather, problems in competing ports, illustrate the spectrum of reasons virtually impossible to forecast.

How the port standard, efficiency and cost level affect its catchment area, domestically (hinterland) and abroad (foreland), is of little consequence here. The time perspective of this study is too short for identifying change and if the low standard of Indian ports, for example, hampers its export that will be manifest in the cargo flows but without any true point of comparison. *de Lombaerde and Verbeke* (1989) give contrasting evidence but only within NW Europe and based on a technique which is not readily applicable in global scale. Changes in ocean freight rates, port & canal charges, and channel depths must be very substantial to have an impact to global trade flows. Evidence from the US grain exports, at home and abroad, suggests this (Fuller et al. 1984; Koo et al. 1988; Viscencio-Brambilla and Fuller 1986; 1987). So does the freezing of the Baltic to Finnish bulk exports (Säntti 1952). The geographical scales are simply too different. Only a complete closure of the Panama (or Suez) Canal would be likely to have dramatic effects. But such events are outside this study.

1.4 Market structure

The structure of the market is of great significance for the interpretation and relevance of our results. Reference was made above to three characteristics of an atomistic market: No buyer or seller is large enough to influence the price, all are price takers. Entry to and exit from the market is easy. Information flows freely. Although no real market can fully meet such conditions the dry bulk market is considered by many to be sufficiently close to the ideal to make atomistic competition a plausible paradigm (e.g., Metaxas 1972; Chang and Chang 1996; Kavussanos 1996a; Tvedt 2003). On the other hand, there is academic evidence to the contrary also, originating from both the wet and dry bulk trades (e.g, Binkley and Harrer 1981, Jonnala et al. 2002; Rinman and Nilsson 2004; Serghiou and Zannetos 1982). Large shipowners and charterers exert considerable market power. The share of the spot market out of the total is not constant but varies between 2 and 50 pct, depending on the phase of the cycle. Pronounced rate peaks and creeks are explained by concentrated changes in demand and supply upon this uncommitted sector. Examples from the trade press will follow but first will we look at the prime indicator of market concentration, the number and relative size of market participants, i.e., shipowners and charterers.

Shipowners are covered by LMIU Vessel Data and charterers by Drewry Fixture Data, both from 1997-8. LMIU's data lists, among others, ship name, its dwt, and the contact organization (Table 1.3). This organization is assumed to be the owner. And here the problems start. Many shipping companies are multinational and split organizationally into national or functional subsidiaries with suggestive names. Others have originally been family enterprises and been subsequently divided between heirs. They have the same family name but can be differentiated by the top person's first name, for example. It is very difficult for an outsider to decide whether either type of company is an operational entity or not. To indicate the uncertainty the number of owners is given by a range. Even that is rather inaccurate because owners of one or few ships often leave the operational management to management companies which can have 20-30 vessels of varying size on their roster. It is as well then to gauge market concentration by simply listing the largest actor. Their shares are around the 5 pct mark, too low to flatly disqualify the atomistic market concept. But that is only part of the story.

Table 1.3 Dry bulk vessels – shipowner's angle, about 1997-8

Segment	Market			Largest owner			
	Owners	Ships	Dwt (1000)	Name	Ships	Dwt (1000)	Pct
Handy	854-918	2663	94,675	COSCO	65	5,038	5.3
Panamax	456-491	1150	76,763	COSCO	55	3,713	4.8
Cape small	159-165	362	44,375	Good	12	1,602	3.6
Cape large	121-125	330	59,342	Bergesen	12	2,995	5.0

Notes: Owners are either operational companies (upper limit) or consolidated firms (lower limit). Very few ships lack a known owner. Ship segments will be explained in Subchapter 1.5.

Source: LMIU Vessel Data.

Shipowners charter additional vessels to gain economies of scale in operations and to economize with capital. They also organize in pools to bid for large freight contracts (Haralambides 1996). In the liner business, conferences and alliances are the vehicles for increased market control (Batchelor 1997). The following examples are partially from the wet trades but are thought-provoking, nevertheless.

> The top ten dry bulk companies contol 15.6 pct of the world fleet. COSCO is the largest of them with about 13 mill. dwt and 260 vessels. When only owned fleet is included the figures shrink to 4.0 mill. dwt and 80 vessels. Management shares trail closely ownership although the topical companies may be different. (O'Mahony 1999)
>
> Pan Ocean Shipping Co. had in 2002 115 chartered handysize bulkers against 40 owned. The corresponding figures for panamax and capesize segments were 42 – 5 and 2 – 7. (www.panocean.com/fleet_top.htm)
>
> 'World-Wide Shipping ... controls more than 120 vessels [mostly liquid cargo] and a chunk of the global oil and gas transport business." (Gimbel 2004)
>
> Storli has a global market share of 18 pct in chemical carriers and Stolt-Nielsen has 22 pct. (Barnes 1995)

The charterer's angle is more diffuse (Table 1.4). There is the usual jungle of conglomerates and their subsidiaries (Jupe et al. 1996). Charterer's identitity may be disclosed only by nationality, may not be disclosed at all, or may even be unknown at the time when the electronic data file goes to the printing shop. Unknown identities are usual in voyage fixtures. To give an idea about the importance of the totally unidentified charterers, a separate line is added to the table. The entries in the data file may be inconsistent. Perhaps they were not made by the same person or she may have used different shorthand expressions on different occasions. Then there are the usual typographical errors which can be fatal when company names consist of 3-4 capital letters only. Unfortunately, there is no objective way to indicate the degree of uncertainty and therefore no range is given. No attention is paid to the different length of charter periods either. As previously, the largest actor is listed and, again, their shares cluster around the 5 pct mark. Particularly in the capesize segments there may be more actors of about the same size than elsewhere. If the size segment is focused on one major commodity (iron ore, coal, grain) the shares will easily double if not triple (Jupe et al. 1996, 7-8). That matters more in the particular commodity market than shipping in general because dry bulk shipowners specialize seldom in a single commodity. As above, anecdotal evidence is more dramatic:

> In the VLCC market there are 8 oil companies against 60 shipowners. The companies know where the own ships are and where the third party's ships are, and where the cargos are. They have the information advantage. When oil company vessels happen to cluster unintentionally peaks and throughs of demand will develop. In the aframax market, fewer and stronger owners exert greater influence, which results in aframax rates being often about the same as VLCC rates, in spite of the smaller size and investment. (Gray and Owen 1996)
>
> Three largest iron ore suppliers, CVRD, Rio Tinto and BHP Billiton controlled in 2004 over 70 pct of global trade. (Marsh 2004)
>
> 'The company [Cargill] ... now controls 25 percent of all US grain exports." (Daniel 2004)

Table 1.4 Dry bulk fixtures – charterer's angle, 1997-8

	All charterers			Largest charterer			
	Chart.	Fixtures	Dwt (1000)	Name	Fixtures	Dwt (1000)	Pct
Handy	330	2194	79,573	Dreyfuss	132	4,629	5.8
unnamed		202	7,134				
Panamax	370	5232	354,297	Pan Ocean	274	18,330	5.2
unnamed		264	17,574				
Cape small	224	1719	218,117	Coeclerici	75	8,815	4.0
unnamed		35	4,402				
Cape large	122	790	130,115	BHP	39	6,420	4.9
unnamed		14	2,298				

Notes: Dreyfuss is an agricultural wholesaler, Pan Ocean and Coeclerici shipping & chartering groups, and BHP a mining company. Ship segments will be explained in Subchapter 1.5.
Source: Drewry Monthly Statistics, 1997-8.

Entry and exit barriers become tangible in the purchase and sales prices of ships, the costs of creating an organization, and the proceedings available from the selling of a going company. There are liquid markets for most size and kind of ship, new, secondhand and scrap. This is particularly true for bulk carriers. Prices of newbuildings are moderate compared with prices of major industrial plants and shopping centers. Secondhand vessels are under normal market conditions available at half price and scrapping yards pay 5-10 pct of the value of a newbuilding (Table 1.5). What makes the market sticky is that the delivery of a standard newbuilding takes two years or so, again under normal market conditions. The costs of setting up business from scratch accrue during a lengthy period and are difficult to quantify. The availability of management companies, however, offers a soft start. It is risky to generalize about the sale of a going company because its health and the business cycle play in, but considering the fragmentation of the industry and the inevitable consolidation there is no shortage of buyers.

Table 1.5 Vessel prices ($mill.), 2Q 1998

Segm., 1000 dwt	Newb.	Secondhand		Scrapped	
		5 yr	10 yr	Notes	Drewry
Handy, 40-45	22.0	17.5	nr	1.7	n.a.
Panamax, 65-75	26.5	19.0	14.3	2.8	2.0
Cape, 130-170	39.5	29.0	17.8	6.2	3.5

Notes: Scrap value from equation: Ldt/GT/Dwt = 1.0/2.0/R in which R = 1.63, 1.80, 1.90 for handysizes, panamaxes and capesizes, respectively. Indian prices. Ship segments will be explained in Subchapter 1.5.
Sources: LMIU Vessel Data; Drewry Monthly Statistics, 1999 (July), Fig. 5; Marriot and Oatway 1998, Tables 4.6, 6.5 and 6.7; Packard 1997, Fig. 7.1.

There are no serious barriers for the flow of general information and technical expertise, whereas specific information, about ship whereabouts and availability for example, may be extremely difficult to come by.

Table 1.6 Approximate share of the wholesale market (COA and Time charter), 1993-2002

	Fixtures, number					**Time,** 1000 days				
	Handy	Panam	CapeS	CapeL	Total	Handy	Panam	CapeS	CapeL	Total
Voyage	5032	7230	5681	2460	20403	169.0	257.3	163.3	84.1	673.7
COA	19	99	144	61	323					
Trip	5675	12118	1310	1110	20213	236.2	496.9	57.0	56.9	847.1
Time	694	2968	470	410	4542	132.7	556.8	123.0	138.5	950.9
Total	11401	22316	7461	3980	45158	537.9	1311.0	343.3	279.5	2471.7
Voyage, raw						169.0	257.3	163.3	84.1	673.7
Trip/Voyage						1.240	1.152	1.515	1.499	
Voyage, treated						209.6	296.4	247.4	126.1	879.8
COA/Voy						0.025	0.057	0.108	0.113	
Voyage, treated						209.6	296.4	247.4	126.1	879.8
Trip						236.2	496.9	57.0	56.9	847.1
Retail						445.8	793.3	304.4	183.0	1726.9
COA						5.2	16.9	26.7	14.2	63.0
Time						132.7	556.8	123.0	138.5	950.9
Wholesale						137.9	573.7	149.7	152.7	1013.9
pct						23.6	42.0	33.0	45.5	37.0
Total						583.7	1367.0	454.1	335.7	2740.8

Notes: Month = 30 days. Time charters and COAs are often given by ranges rather than point estimates for periods and tonnages, respectively. In such case the midpoint is selected. Fixture numbers here and elsewhere may differ slightly from each other, being taken from different work files. The files became rapidly too large for convenient handling and were simplified when the work progressed. At the same time fixtures for vessels below 25,000 dwt and those with vague or incomplete data were (belatedly) discovered and excluded from the analysis. The end product, files usable for the estimation of freight rate functions and comprising only voyage and trip fixtures, have about 40,330 observations or 11 pct less than the original input.

Sources: Drewry Fixture Data.

Doubts about the relevance of an atomistic market as a paradigm thus remain although clear oligopolistic features can be sensed only in certain submarkets. But there is another type of complication. The Simulator is constructed for a single-trip spot market. Cargos are generated, they are allocated competitively, and immediately transported. There is no speculative element and everything is completely transparent. That is not quite as it is in the Real World where time charters, geographically vague contracts for three months to several years, and Contracts of Affreightment (COA), geographically and volumewise specific contracts for large transport needs over lengthy periods, play an important role. They are concluded to enhance logistical security, control transaction costs, or just to speculate in the freight market (Hale and Vanags 1989; Kavussanos and Alizadeh 2002; Pirrong 1993;

Veenstra 1999b). Because the contract time comfortably exceeds the time needed for completing voyage and trip charters it is pertinent to talk about the wholesale freight market. It comprises time charters and COAs, which do not fall well in line with the Simulator.

The main difficulty of producing a reliable estimate of its size is to make charters in cargo tonnes and months commensurate. Distances involved in voyage charter (only in cargo) are first converted to time and thereafter treated by a multiplier which makes them approximately comparable with trip charter time (incl. paid ballasting time). The transportation work of a COA can be estimated by selecting the most likely ship size and calculating the time needed for the task, whereafter all times are aggregated. But almost the same accuracy will be achieved by estimating the share of COAs out of all voyage cargo tonnes and applying the percentage to the total voyage time. There still is the inaccuracy that a time charter collects revenue all the time after delivery whereas voyage and trip charters are for one trip only, to be followed by a period when no revenue is earned. With these reservations, the wholesale market is easily one third of the total market (Table 1.6).

The size of the wholesale market appears to dowgrade seriously the practical value of this study. It is for two reasons. First, hardly nothing is known about the geographical deployment of time-chartered ships. Second, a time charter rate is not tied to any route or trade. The instinctive analytical response to the existence of time charters, insensitive to route-specific freight rates, is to scale down the transportation flows proportionally. The weakness of this solution is that time-chartered ships need not distribute between routes in proportion to the traffic available, particularly if they have been hired for a specific route (steel plant) or sales campaign covering a given geographical area (grain trader). Only if they have been chartered speculatively is a geographical bias unlikely. A flexible charterer with a taste for risk taking will pursue strict revenue maximization by employing the vessel in the retail market to best advantage. The disturbing thing is that we know virtually nothing about the extent of these alternatives.

1.5 Cargos and ships

Ocean traffic originates from the need to transfer cargos from one part of the world to another and all shipping research is connected, directly or indirectly, with the ensuing trade flows. To an extent the flows can be derived from foreign trade statistics. Dry bulk shipping deals primarily with iron ore, coal and grain (inc. soya). Upon that come 10-15 ”small” bulk cargos such as fertilizers, oil seeds & cakes, sugar, bauxite, metal concentrates and steel (World Bulk Trades 1997, 34, 37, 40). It is quite possible to derive their approximate annual trade flows from foreign trade statistics. It may be possible to sharpen the picture to the quarterly and even monthly level, at least for some commodities. It is further possible to make rough estimates about the distribution of monthly, quarterly and annual shipments between weeks. But trade statistics are sometimes badly inconsistent with maritime statistics (if available), e.g., they are inaccurate temporally and do not account for swaps. The split between ports may be unavailable in practice. That matters when countries have ports facing several oceans.

Such information is available in the LMIU Movement Data which gives the ports visited. It also discloses the size of vessel, an important item because shiploads vary greatly depending on the type of cargo (Table 1.7). It is therefore advisable to access data about ship movements directly, when this possibility exists. Recommendable as this solution is, it also has disadvantages. The foremost is the volume of data, 335,000 lines of movement during three years for the entire bulk sector above

25,000 dwt, to be pared down to good 50,000 cargo legs and as many ballast legs. There is no direct information about the type of cargo, or even its existence. In a generalizing study as ours the type of cargo may not be so important, the ship is assumed to be indifferent to that, but the split between cargo and ballast legs is obligatory. That split must be made by the analyst and it is based entirely on the general cargo structure of the ports. The task is fairly easy for large ships which concentrate on the three major cargos but difficult for smaller vessels which carry the full range.

Charterer preferences are for certain sizes of cargo and shipowners respond by offering corresponding sizes of vessel. These are habitually classified as handysize, panamax and capesize, with formal limits between the segments at 50,000 and 80,000 dwt (Fig. 1.4). It is also customary to refer to the subsegment above 40,000 dwt as handymax while the 50-55,000 dwt range of vessels built mostly in the 1970s and earlier lacks a name. Overlooking these subsegments, a typical handysize might be about 35,000 dwt, panamax 70,000 dwt and capesize 150,000 dwt (Table 1.8). The Simulator is constructed for one size of ship with standard cargo. It follows that each segment will be simulated separately.

Table 1.7 Average ship size for some commodities, 1993-2002

Commodity	Fixtures	Dwt
Sugar	1,129	21,554
Fertilizer & raw mat.	1,066	34,392
Minerals	109	34,638
Steel	503	36,815
Agriproducts	692	41,677
Grain	8,115	52,547
Bauxite & aluminium	346	54,514
Petrocoke	120	55,876
Coal	5,579	104,383
Iron ore	4,992	133,609

Notes: Fixtures below 25,000 dwt also included. Cargo tonnes converted to dwt as explained in Appendix 3.2.
Source: Drewry Fixture Data.

Table 1.8 Average ship size (dwt), about 1997-8

	Table 1.3	Table 1.4	Simulator
Handy	35,600	36,300	36,000
Panamax	66,800	67,700	67,000
Cape small	122,600	126,900	125,000
Cape large	179,800	164,700	170,000

Note: Sizes for the Simulator are purely notional.
Sources: Tables 1.3 and 1.4.

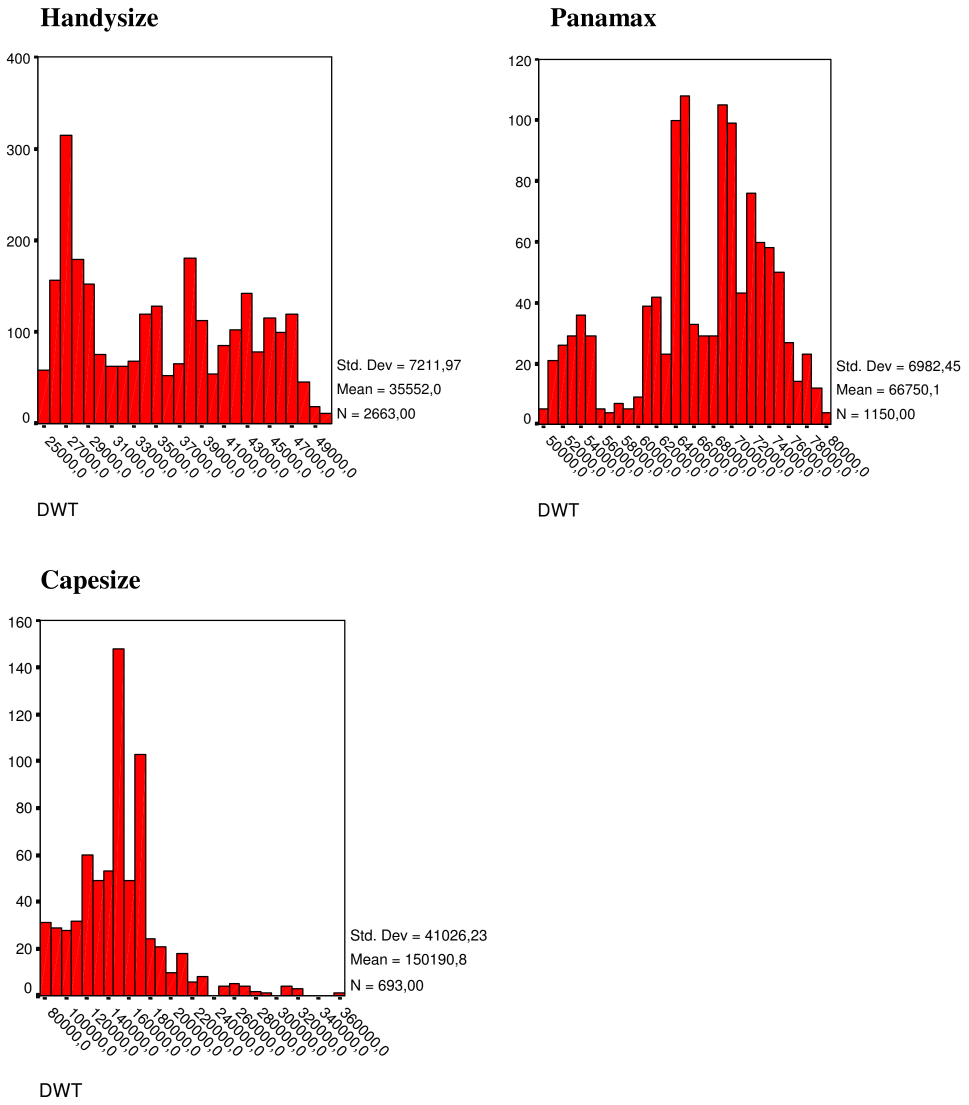

Figure 1.4 Ship size by segment, about 1997-8

Source: LMIU Vessel Data.

1.6 Geographical space

There is also another reason for the existence of ship size segments than charterer preferences, the ability of the ships to pass the Panama and Suez Canals fully loaded, or at all. The Panama Canal differentiates effectively the panamax and capesize segments. The locks allow a maximum breadth (beam) of 32 feet and a maximum length of 295 feet, both at waterline. The channel is dredged to 12.3 m at normal water level which, however, cannot be maintained during prolonged droughts. These measures translate to about 80,000 dwt. The Suez Canal constrains only by the depth of its channel, 21 m, sufficient for sizes up to 150,000 dwt with full cargo and for all sizes thereafter when in ballast (Fairplay Ports Guide, 1314; Lloyd's Ports of the World, 25). That splits the capesize segment into two, small and large, for operational purposes. The limit is less sharp than the one drawn by the Panama Canal because there are no locks, the bottom is sand and the shape of the ship's hull affects its exact draught. The limit affects particularly the largest ore carriers which float deep for maximum stability at open sea. Sea canals hardly matter for the split between handysize and panamax. The Kiel Canal with its 11 m depth is a theoretical possibility but the distance saved is only some 200 nm. The St Lawrence Seaway to/from the Great Lakes is more relevant allowing only the handysize and even it only with a cargo of 18,000 tonnes, to be topped up at a St. Lawrence port. In practice, ships entering the water system seldom exceed 30,000 dwt. Their movements are classified in this study as calls at a St. Lawrence port and their fixtures are entirely overlooked.

The different ability of ship segments to use the main sea canals means that their distance matrices differ. Handysizes and panamaxes always use the Kiel Canal. Both use the Panama Canal when it offers the shortest distance. Small capesizes always use the Suez Canal when it offers the shortest distance whereas large capesizes use it only in ballast. The shortest distance principle is a simplification, of course, because canal fees make non-use economic when the extra distance is not excessive or freight rates are low (Table 1.9). How excessive and how long, cannot be decided unambiguously because discounts are commonplace for frequent users. A rate difference of $0.50-1.00 per cargo tonne (out of a total of $12.80) from Hampton Roads & Richards Bay to Far East for a 130,000 dwt capesize may give some guidance.

Table 1.9 Canal transit costs (1000 $), 1997

Dwt	Panama	Suez
40,000	67	154
65,000	102	178
140,000	n.a.	229

Source: Marriot and Oatway 1998, Table 6.10.

The distance matrices are compiled from *Caney and Reynolds* (1995) and complemented by *Tanimoto* (1992) and *Aldworth* (1997). The entries so obtained are only approximate. The main reason is that there are no highways on the open sea, only broad tracks moving with seasons and anticipated weather (Cockett 1997, 274-281). Tracks in opposite directions between a given port pair need not be the same because of sea currents and prevailing winds. This may be the reason why some port pairs have different distances on different pages. Large capesizes may also be constrained by the shallowness and difficult navigation of some natural waterways. For example, they cannot enter the Baltic with full cargo and possibly prefer the Lombok Strait for the Malacca Strait, irrespective of the additional 1,200 nm (Couper 1983, 155). It follows that different sources may report distances which differ by up to 500 nm within a total of 7,000 nm. Differences of 100 nm are commonplace. Such ambiquity is one reason why no time correction will be made for the passage of canals. Another is the substantial error element imbedded in port times (Subchapter 3.1.3). When making mental comparisons of distances it is advisable to have a desk globe available. Two-dimensional maps, particularly in the Mercator projection, give easily a distorted idea of true distances worldwide. Full distance tables are given in Appendix 1.2 and reduced ones in Appendix 1.3.

The distances are measured between reference ports for wider areas, regions, 37 in all (Fig. 1.5). Their boundaries normally follow political ones except where a country is very large or spans several distinct waterbodies (Canada, USA, Russia, Australia, China, Brazil, Mexico, Colombia, France, Spain, Italy, Germany, Egypt, Saudi Arabia, India). The reference port is at an approximate point of gravity for ocean traffic and has an entry in published distance tables. As an overall principle this works adequately although distortions occasionally occur. For example, Guayaquil is the reference port of Region 424 when most of sugar exports and grain imports go via Buenaventura in handysizes and iron ore is shipped in capesizes through ports in southern Peru. Region 127 is a similar case. Kamsar is the reference port, in reality used by panamaxes for bauxite exports whereas iron ore is shipped in capesizes from Nouadhibou in the northermost corner, and handysizes visit the southern ports.

> Short distances, often within regions, are problematic because the measurement error becomes large percentagewise. The possible inaccuracy is perhaps one day when the mesh is 37 regions and 1-3 days when it is 10 macroregions. Short distances emerge more often in movements than fixtures because within-region fixtures are comparatively rare. Smaller ship sizes are more likely to have short distances than larger ones. The problem is most acute in the context of multiporting because of the inevitable deviations and delays at entry and departure. Fortunately, that complication will not be explicit in this study because multiporting will be abstracted away. Distances differentiated by the major traffic is a plausible solution in serious cases. For example, the routes Narvik – Hamburg, Hamburg – ARA and Narvik – Dunkirk benefit from 1100 nm instead of the standard 500 nm. Some other exceptions are given in the Distance Tables. The effect depends also on the region's relative position in the overall port sequence. At the end of a multiporting section it matters whether the next, long leg is in the overall travel direction or not. Bahia Blanca is a typical case and related to the destination in the Atlantic respective Pacific Sphere.

Thirty-seven regions means that a complete interaction matrix has 1,369 cells. The number is rather large and recommends downsizing. The downsizing can even be enforced by the underlying data. That is the case with trip charters, an analytical cornerstone, which often use extensive geographical ranges such as Skaw – Cape Passero and southern Japan – Singapore. In practice, the analysis is conducted at two rosters: 37 regions for a basic coding and orientative experiments, and 9 or 10 macroregions for the actual analysis. For freight rates there are 9 macroregions (regions for short), and for ship movements and the Simulator there are 10. The small number of fixtures from/to Region 10 made its consolidation with Region 9 necessary. The downsizing means that interaction matrices have only 81 and 100 cells, respectively. These cells represent routes between origin and destination (O-D) regions. Macroregions measure normally about 2,500 nm or less, from regional reference point to another, with the Continent (Region 3) as an exception. Great care is exercised to find realistic internal distances for the macroregions and these occasionally vary by ship segment.

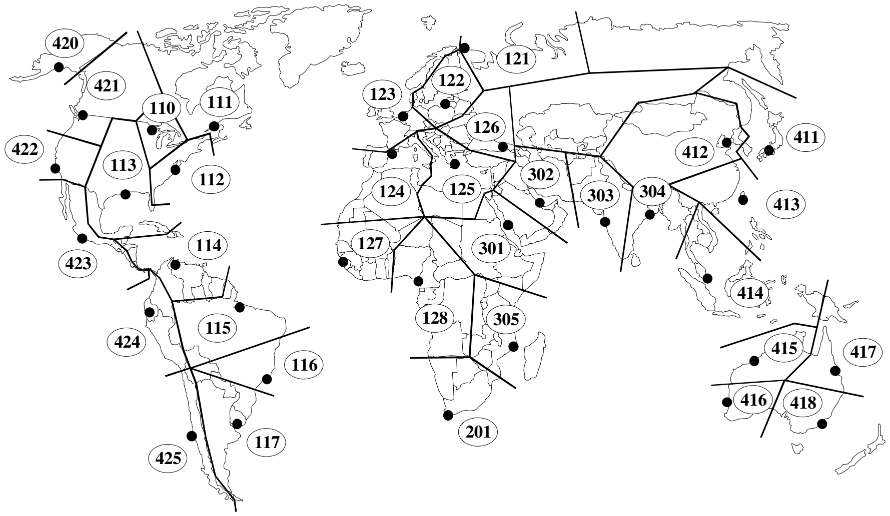

Figure 1.5 Regions and their reference points

Reference points

420 Anchorage
421 Vancouver
422 Los Angeles
423 Manzanillo
424 Guayaquil
425 Valparaiso

110 Duluth
111 Baie Comeau
112 Norfolk
113 New Orleans
114 Maracaibo
115 Belem
116 Rio de Janeiro
117 Buenos Aires

121 Murmansk
122 Klaipeda
123 Rotterdam
124 Barcelona
127 Conakry
128 Douala
201 Capetown

125 Piraeus
126 Novorossisk
301 Jeddah
302 Bahrain
303 Mumbai
305 Mocambique

304 Calcutta
414 Singapore
415 Port Hedland
416 Fremantle

411 Osaka
412 Dalian
413 Kaohsiung
417 Mackay
418 Newcastle
419 Honolulu

Source: Author.

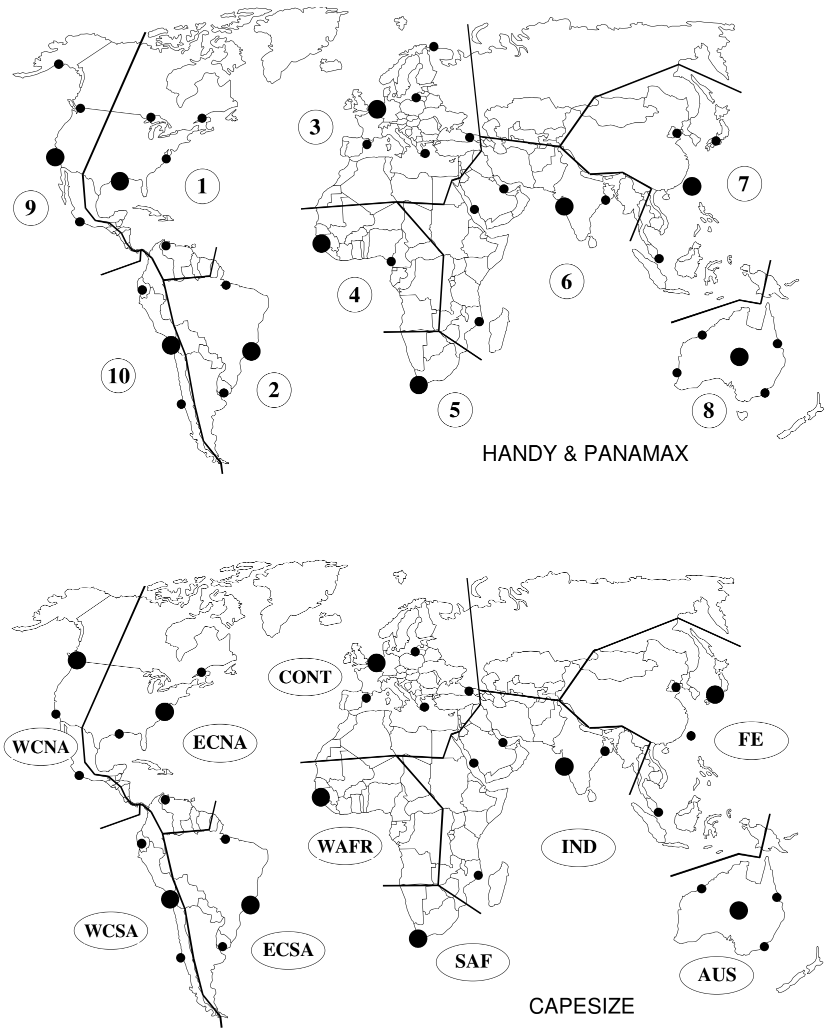

Figure 1.6 Macroregions and their reference points

Notes: Great Lakes only in handy segment. Red Sea not in large cape segment. Hawaii belongs to Region 9.
Source: Author.

1.7 Implementation

The next two chapters develop the input data for the Simulator. The bulk consists of ship movements and freight rates. Upon them come a number of operational parameters, a byproduct of fixture data. Ship movements are derived from LMIU Movement Data 1995 & 1997-8, and freight rates from Drewry Fixture Data 1993-2002.

Data about ship movements was originally collected to support the writing of marine insurance at Lloyd's of London. It covers in principle each movement of all the vessels of commercial importance. The result, from our angle, is a mountain of unassorted raw data, badly in need of editing. The first task consequently is to convert the data into a shape from which O-D matrices can be produced. That means a file in which the itinerary of each ship consists of a sequence of alternating loading and discharging ports. No other port visits are allowed and multiporting is reduced to the first loading and last discharging port. Each ship segment and each calendar year has its own file. These files constitute the raw data for half of the Simulator input.

In addition, they offer a unique base for plain description and simple, orientative analyses. Figures will be prepared about the number of arrivals and departures in cargo by ship segment and region, the number of cargo and ballast legs between and within macroregions by ship segment, and the workload between and within the major oceans, Atlantic, Pacific and Indian, by ship segment. The workload is split between cargo and ballast and the attached average distances are recorded. The observed workload is compared with an approximate "optimum" derived by linear programming and the difference is split into time and cargo & location components.

Freight rates are available in fixtures, abridged charterparties released to the public domain. Except for the rate itself there is information about the loading & unloading alternatively delivery & redelivery ports/areas, ship age, ship size alternatively cargo size, the time range for loading (laytime) or ship delivery, loading & unloading rates/times, speed and fuel consumption in fair weather conditions. This supplementary information is an invaluable source of simulation parameters. Relating fixtures to cargo legs gives a numerical idea how well they represent the underlying market.

The main line of rate analysis, however, is the conversion of published fixtures into usable freight rates. The basic split is between voyage and period charter and the later divides further into trip and time charter. Voyage and trip charter are almost exclusively for single transport tasks, "trips" in our terminology (see page VI). They are loosely called the spot market because most charters are for immediate execution and constitute our data base. Its use is complicated by the fact that voyage and trip charters differ in units of measurement and the way how costs are split between the charterer and the ship. It is a major task to consolidate them into a homogeneous data set. This set is used to derive freight rate functions for each route, about 150 in all. The first set of functions is for the whole period 1993-2002, but being too inaccurate for use in the Simulator, is developed further into biannual sets.

These sets are all what is needed in the Simulator. But, from a wider angle, it is important to have an idea about the rate structure at large. Do certain rate series behave in the same way? Is the covariance so pronounced that classification becomes meaningful? Do functions influence each other, i.e., is there a lag effet? Can the familiar concepts of fronthaul and backhaul be operationalized in a way which permits generalizations about their relative rate levels? These are extensive questions and all what we can do here is to make short excursions to get a taste of the soup.

The threads are pulled together in the Simulator. The original version has been modified to allow the use of more accurate functions for estimating rates and generating cargos. The details are made available in two appendices and the motivations given in three subchapters, about rate functions, cargo generation and the ships doing the job. A simulator is an inherently dangerous analytical tool because there are no automatic checks and balances. It is up to the analyst to create them. The most potential traps are discussed shortly and found impotent. The first results are less unambiguous than desired. The simple Revenue Gradient functions well in certain segments and years but not in others. The output is then pooled alternatively by segment or year, the desired generalization is achieved and the impact of geography measured.

2 SHIP MOVEMENTS

2.1 Preparing data

The main purpose of handling ship movements is to get an idea of the number of cargo legs between all region pairs during a year, to be used as input in the Simulator. In this purpose a chronological file of all-inclusive ship movements is first converted to one comprising only port visits for loading and discharging, and thereafter converted into an Origin-Destination (O-D) matrix. The latter step requires that loading and discharging ports alternate. To achieve this it is not possible to have a unbroken string of either loading or discharging ports. In such a situation, other ports than the first loading and last discharging port of a trip will be deleted, implying a sad loss of information. In addition to the main purpose, much descriptive material will be given because this size of data handling is comparatively rare and still more rare is its availability in the public domain. Three sets of matrices will be discussed: **Original** which gives the number of visits per port; **Stripped** which gives the number of ocean bulk (as opposed to liner) legs between regions; **Final** which gives the number of ocean bulk legs of the polypolistic market. The raw data underlying these matrices discloses the extent and intensity of trade areas of individual ports and port groups. They can also be used for estimating the average lengths of cargo and ballast legs and the actual and minimum transportation work, assuming homogeneous commodity in a cross-sectional situation (the classical transportation problem). The latter will show the relativity of efficiency studies based on layups and average speeds. These are, of course, important but the order of magnitude will be different.

The discussion is based on the data collected and marketed by *Lloyd's Marine Intelligence Unit* (LMIU) and called here LMIU Movement Data. It is collected by about 2,000 Lloyd's agents and subagents all over the world and is the most comprehensive database of commercial ship movements available, covering in principle all international traffic plus national traffic between different water bodies. In practice there are restrictions. Historically, collection of data on individual ship movements at industrial ports was problematic in Taiwan and South Korea, and incomplete in Indonesia and China. By contrast, coastal traffic within a country is occasionally included. In the USA, traffic between the West Coast and the Gulf is recorded, at least when carried by ships in ocean traffic. So is the traffic between Glensanda and British & Continental ports, Tarragona and Palma de Mallorca, Tubarao and Santos, Weipa and Gladstone, Trombetas and Vila do Conde, Sorel and Havre St. Pierre, Port Hedland and Port Kembla. This does not yet mean that all international movements are long-distance. Just look at shipments from Cozumel to USG, Narvik to Bremen, Seven Islands to Baltimore, Mina Saqr to Shuaiba and Vostochnyy to Japan.

The database actually consists of three types of file, one for vessels, one for ports and regions, and one for actual movements, linked with each other through identification codes. The set used in this study is a combination of the three. A simplified example is displayed in Appendix 2.1A. The set includes all movements of dry bulk carriers at and above 25,000 dwt in the years 1995, 1997 and

1998, almost 335,000 movement lines (Table 2.1, Total A). The selection of years has been explained in Chapter 1.

The volume of the data called for special attention and there were some doubts in the beginning whether it exceeded the analyst's possibilities in the first place. The total workload ultimately amounted to about 1,000 hours and it was important to scale the original data down to manageable proportions at an early stage. The technical device was the Excel compilator. An Excel file has at most about 65,000 lines. That necessitated the division of the original Handy segment into subsegments HandyS, HandyM and HandyL, a division that was kept until the Standby file (below). The Practical Steps are called Downsizing, Coding and Weeding:

Downsizing	
LMIU to Move	Consolidate parts of LMIU files about movements, ports and vessels to a new file with port and vessel names in addition to their identification numbers.
Move to Original (Total A)	Delete all lines indicating plain passage (Panama Canal, Suez Canal, Cape Finisterre, Gibraltar, Skaw, etc).
Original to Standby (Total B)	Delete all lines indicating newbuildings, scrappings, bunkering, waiting and pilot stations. Delete columns Vessel name, Move ref-no. (running), Place id-no., Move qualification, Move type 2, Arrival estimated, Departure estimated, Edition date and Line no. Replace YoB by Ship age (YoB – Year = Age).
Coding	
Standby to Readied (Total C)	Add regional codes, port code and week id-no. Delete all wet (oil) cargos and all lines for 2nd, 3rd, etc. loading ports and for n^{th-1}, n^{th-2}, etc. discharging ports in a trip. Check logical sequence 1-0-1-0 by vessel. Add codes for traffic to be subsequently deleted.
Weeding	
Readied to Stripped	Delete columns Move type 1, Arrival qualification, Departure date, Departure qualification. Convert the one-column port sequence into a two-column variant, the one column for loading ports and the other for discharging ones (Appendix 2.1C). Delete all lines indicating within-port moves, unessential within-region moves, liner-type, oil and specialized traffic.
Stripped to Final	Delete all shuttle and tied traffic.

The crucial and most labor intensive phase was the conversion of Standby files into Readied files. That was due to the need to decide whether a port was visited for loading, discharging, or both, a piece of information which is not collected by LMIU, plus the unavoidable gaps and errors of the data. To survive, an analyst must have real stamina and a solid knowledge of old-time production geography, the kind of knowledge which has very little demand in modern academic research. The drudgery notwithstanding, it is not wasted time to handle about 250,000 lines of port visits. One can see patterns to develop and gets a better understanding of the shipping world within few months than many a professional during a lifetime. It was possible to handle about 2,000 lines/day in the Cape, Panamax and HandyL segments but only 1,000 lines/day in the HandyM and HandyS segments because of the growing complexity of itineraries. To keep the job manageable the Handy segment was handled only in 1997. It was about the same size as the Panamax segment in 1995 and 1997-8.

Table 2.1 Dry bulk movement data, 1995, 1997-8

Segment	dwt	1995	1997-8	Total A	Total B	Total C
HandyS	25,000 to 29,999	18,087	43,334	61,421	58,043	14,514
HandyM	30,000 to 39,999	20,228	45,046	65,274	61,595	18,036
HandyL	40,000 to 49,999	13,202	39,609	52,811	48,875	15,155
Panamax	50,000 to 79,999	17,623	41,894	59,517	54,818	44,133
CapeS	80,000 to 149,999	6,985	13,574	20,559	15,308	13,665
CapeL	150,000 to 350,000	4,445	12,250	16,695	14,508	11,646
Total		80,570	195,707	276,277	253,147	117,149
Originally		95,926	239,015	334,941		

Note: Total C for handysizes in 1997 only.

There were 24.6 pct more movement lines in 1997-8 on average than in 1995, about 10 pct-units more than there was ship capacity. Both exceeded the growth of voyage and trip chartering by a wide margin (Table 2.2). Exports and imports (tonnage, not tonnage-miles) of iron ore, coal and grain by major trading partners remained stagnant. It all suggests that movements were recorded less fully in 1995 than 1997-8. Particularly the absence of older ships in 1995 is striking. Freights were high and every available vessel should have been employed. The only rational explanation is that LMIU succeeded in expanding and intensifying its intelligence network from 1995 to 1998. Other evidence suggests that the upgrading was particularly marked in the Far East.

Table 2.2 Vessels and chartering, 1995, 1997-8

		Vessels at mid-year mill. dwt		
Segment	dwt	1995	1997-8	Chg
Handy	30,000 to 49,999	62.2	69.4	11.6
Panamax	50,000 to 79,999	57.4	64.5	12.3
CapeS	80,000 to 149,999	34.1	32.5	- 4.7
CapeL	150,000 to 350,000	31.9	45.7	43.3
All segm		185.6	212.1	14.3

	Chartering					
	Voyage (mill. dwct)			Trip (mill. dwt)		
	1995	1997-8	Chg	1995	1997-8	Chg
All segments	176.4	189.4	7.4	149.0	151.3	1.5

Notes: Handy segment truncated. Figures can vary by issue reflecting corrections to preliminary estimates. Change in percent from 1995 to 1997-8 average.

Source: Drewry Monthly, various issues (May-July).

The discussion follows broadly the Practical Steps outlined above: Deletion of plain passages and non-cargo visits, Vessel type, Port coding, Deletion of wet cargos, Fill-up of gaps and adding of within-port ballast legs, Deletion of 2nd, 3rd, etc. loading ports and n^{th-1}, n^{th-2}, etc. discharging ports of a trip, Itinerary principles, Deletion of liner-type itineraries, within-port movements, shuttle-traffic and recognizable tied traffic. The order of deletion may be controversial. In particular, lines should not be deleted before port coding because passings and ostensible bunkering visits give clues about loading and discharging ports. The benefit of downsizing the files at an early date outweighs this aspect, however.

Plain passage, including direction, is recorded at the Panama and Suez Canals plus some focal locations with lively ship traffic. Skaw, Dover Strait, Cape Finisterre, Gibraltar, Algericas, Ceuta, Canakkale, Istanbul and Thursday Island are among them. The heavy concentration in Europe draws attention. One would have expected also Malacca Strait, Sunda Strait, Cape of Good Hope and possibly Cape Horn. The both Straits flounder on the reluctance of the neighboring countries to allow observation. Cape of Good Hope is covered adequately by bunkering visits in the area's ports. Cape Horn has only little traffic.

Non-cargo lines include inactivity, newbuilding, scrapping, plain bunkering and taking/leaving a pilot. For inactivity there is a special identifier. Locations for temporary waiting are comparatively few and well known: Annapolis Anchorage, Chennai (Madras) Roads, Chittagong Roads, Colombo Roads, Everingen, Falmouth Bay, Fujairah Anchorage, Mutsure and Tees Bay, although roadsteads can also be used for lightering. Waiting is more common by capesizes and panamaxes than by handysizes (cf., Table 4.10). Newbuildings leaving shipyards are those vessels which make their first move or have an age of 0 years. It happens occasionally that vessels have negative age, possibly due to erroneous launching year or delayed completion. Such lines may not indicate first moves but even fifth or sixth. They are, nevertheless, considered newbuildings. Two locations dominate the scrapping industry, Adang and Gadani Beach. Yard visits for repairs are difficult to identify and little effort was spent on them. Theoretically, they have a special identifier but very little was seen about it.

Bunkering can be a headache, partly because the need is frequent and partly because the purpose is to combine it with port visits in general. There are several ports which are visited mostly for bunkers: Algeciras, Ceuta, Falmouth, Gibraltar, Honolulu, Las Palmas and Tenerife. To them can be added some anchorages and roadsteads: Colombo Roads, Everingen, Falmouth Bay and Fujairah Anchorage. Then there are ports where bunkering is routine but which also are important cargo ports, at least for some ship sizes: Busan, Capetown, Los Angeles, Rotterdam and Singapore (handysizes). The only way to screen plain bunkering visits from other activity is port time. Fortunately, port times in these major ports are quite reliable, which means that the appropriate columns cannot be deleted before port coding. The main rule is that bunkering in handy and panamax segments can be completed in 0-1 days, weekends excluded. This rule is only approximate, easily seen on the Great Lakes where typical "terminal ports" like Thunder Bay have 0-day visits. Ostensibly loading is so efficient that a 25-30,000 dwt vessel can be handled within hours. It must be admitted, however, that "thruway ports" like Detroit have proportionally much more 0-day visits. Two days belongs to the grey zone and any time beyond that involves probably loading, discharging and/or waiting. Ships in liner-type traffic challenge these rules and have very short port visits, in Singapore, Los Angeles, Durban and Busan for example. Then it becomes impossible to keep bunkering visits and genuine cargo visits apart from each other. Very short visits are occasionally made for taking or leaving a pilot, by panamaxes in Port Angeles and Victoria (BC), for example.

Vessel type LMIU's data for dry bulk vessels differentiates between seven types at the 3-digit level, one overall and six specialized to varying degrees (Table 2.3). The most specialized type is BWC with a voluminous hull and automated cargo handling equipment. Specialization is visible in repetitive itineraries. The type is used almost exclusively by Japanese. Wood chip is also transported by BBU and BCB vessels but the economies are then much worse. BCB and BCE vessels have pneumatic, air-tight cargo handling equipment to minimize pollution. That does not hamper their use also for other purposes. The same degree of pollution prevention can be achieved quite easily on BBU vessels, and BCB and BCE types have fallen out of favor. They are lumped together with the run-of-the-mill BBU vessel. BOR type has a small hull adapted to the small volume of ore cargos, strong enough to withstand rough seas and forced loading. The type has poor economies when used for voluminous cargos

Table 2.3 Total B (from Table 2.1) disaggeregated

	1995	1997	1998	Total	1995	1997	1998	Total
	HandyS				**Panamax**			
BBU	15,370	17,871	18,643	51,884	14,360	16,377	17,742	48,479
BCE	1,755	1,895	1,816	5,466	222	180	146	548
BCE	0	95	80	175	0	0	0	0
BOR	119	129	97	345	267	197	142	606
BWC	28	0	10	38	149	122	143	414
CBO	12	18	14	44	1,480	1,595	1,595	4,670
COO	0	52	39	91	41	29	31	101
Total	17,284	20,060	20,699	58,043	16,519	18,500	19,799	54,818
	HandyM				**CapeS**			
BBU	16,724	18,779	18,537	54,040	4,605	4,227	4,047	12,879
BCB	2,224	2,295	2,156	6,675	2	0	4	6
BCE	0	9	20	29	0	0	0	0
BOR	21	8	0	29	321	299	293	913
BWC	208	264	211	683	0	0	0	0
CBO	0	34	23	57	939	248	162	1,349
COO	0	41	41	82	111	33	17	161
Total	19,177	21,430	20,988	61,595	5,978	4,807	4,523	15,308
	HandyL				**CapeL**			
BBU	9,130	13,923	15,626	38,679	2,986	4,377	4,836	12,199
BCB	1,631	1,709	2,123	5,463	0	0	0	0
BCE	0	0	0	0	0	0	0	0
BOR	38	54	54	146	550	527	521	1,598
BWC	1,051	1,331	1,345	3,727	0	0	0	0
CBO	263	301	296	860	101	14	1	116
COO	0	0	0	0	275	183	137	595
Total	12,113	17,318	19,444	48,875	3,912	5,101	5,495	14,508

Legend: BBU standard bulker; BCB clinker; BCE cement; BOR ore; BWC chips, CBO oil and dry bulk; COO oil and ore.

Source: LMIU Vessel Data.

such as grain. CBO and COO were built to minimize the length of ballast legs by alternating between dry and wet cargos. The time needed to clean the holds after dirty oil cargos was in practice longer than anticipated, which constrained use and spoiled economies. Also these types are seldom built nowadays. Because this study is about dry bulk shipping wet-cargo legs will be excluded. What remains are the standard BBU vessels. The segment looks homogeneous from a distance but at closer look the picture changes. There are specialized vessels which, however, are not indicated as such. Paper carriers are an example. A subtle feature is a vessels's own cranes, important when lightering at out-of-way roadsteads and ports, in the NW corder of India for example. There is nothing we can do to this manyfold in practice.

Vessel type can occasionally help in defining the cargo status of a leg. For example, the panamax CBO vessel LLP 150150 plying between Tarragona and Palma de Mallorca probably carried water from the mainland to Mallorca's tourist industry during the season, when the island's resident population is quadrupled. BBU capesizes LLP 3365 and 3367 plied between Valdez and Cherry Point, two oil ports, ostensibly carrying supplies from the Pugent Sound area to Alaskan oil fields. In reality, they had been converted to tankers and carried crude to an oil refinery. The conversion only had not been registered by LMIU, a rare slip. The same seems to have happened to the panamax LLP 177645 which was used as a car carrier although registered as standard BBU. Liner-type visits to Toyohashi and other car export ports raised doubts. There were some other suspicious cases but their investigation was seen as waste of resources and consequently dropped.

Port coding Lloyd's agents do not inform about the cargo status of an arriving or departing dry bulk vessel. The market demand for such information does not warrant it. There would also be the complication of handling part cargos. In oil shipments information about the cargo status is available. The practical conclusion is that the cargo status of a leg must be determined by the analyst on the basis of indirect evidence. It was thought initially that files where the cargo status (loaded/ballast) of a leg is indicated would practically solve the construction of O-D matrices. With hindsight the idea was absurd as will become apparent below.

Consistency in coding was enhanced by starting from the cape (largest and most reliable) segments and implementing the same technique in the panamax segment. The handy segment, by contrast, was handled in reverse order, from small to large vessels, the small ones being considered the most difficult ones with many minor ports. The total number ultimately rose to about 1,600. Each segment was divided into separate files by year. Within a file, the lines were sorted by port, each port was given a regional code and a tag which suggested its role as exporter or importer in that size segment (1 = export; 0 = import; 2 = both). Empty cells during the first input indicated ignorance.

When the lines were rearranged by ship and move-id into itineraries it was comparatively easy to make adjustments in cargo status, which often depended on the context, i.e., on the previous/next port and the assumed cargo, the latter being influenced by the ship size. For example, Xingang and Shimazu/Suzuki are known for their coke shipments. It was logical to assume that movements between them and the USA were in cargo. Similarily, movements from Slite or Barcelona to Florida or US Gulf implied clinker cargos. Familiarity with general trading patterns helped. Handbooks gave no clue whether there were cargos from the Continent (Antwerp, Terneuzen, Yjmuiden, Balticum) to Montreal and the Great Lakes. The Continental port sequence suggested it because these topical ports were often the last ones in a general west-east sequence. Midwest consumption patterns supported the idea. There is a large population descending from Germans and Poles who are avid consumers of rye bread and beer, for which barley is needed. It was further assumed in all handy segments that there were scrap and fertilizer cargos from the Continent to Montreal and the Great Lakes. Brokers found the idea acceptable. But such reasoning was not always possible. Short single-port stops in South Africa were very popular among vessels sailing from Far East to South America. Unfortunately, it was impossible to know whether they were genuine port visits or just bunkering, particularly if they had previously visited Singapore.

The likely cargos to be handled at a port were derived from port guides and various handbooks, Lloyd's Ports of the World (Emery) and Fairplay Ports Guide (Watson) with data from about 1995 being the most useful sources (Appendix 2.2). But even these are only as good as the information delivered by the port authorities. There is no cargo-specific information about tonnages, sometimes no clue which of the listed cargos are exports and which are imports, nor specification of any cargo in general. The last alternative is typical for small ports. Private ports, power stations, steel mills and export terminals may not be listed at all. The port guides were complemented with more specialized publications such as handbooks for iron ore and coal ports, steel mills, fertilizer and alumina plants. The use of handbooks to identify plant locations helps only to a degree, however. The locations are too many and some inland plants like Siilinjärvi (fertilizer) can have extensive overseas trade (through Kotka).

The technique worked satisfactorily at single-commodity ports, often midsize, handling the major bulk cargos, iron ore, coal and grain. Amazing accuracy was sometimes possible. We know for example that Sepetiba is for coal and bauxite imports whereas Sepetiba Terminal (Guaiba Island) is for iron ore exports. Praia Mole imports coal but the adjacent Tubarao exports iron ore. Problems arose at small and very large ports. Port directors readily list even very small ports, down to piers and roadsteads, but seldom inform about their cargos. The port code must then be based on the area's general economic characteristics. Sumatera and Kalimantan, for example, are sources of logs, sawn goods, plywood and minerals. Large ports, such as Rotterdam, Houston, Hampton Roads and Baltimore handle too many cargos to allow the simplistic split export or import. Quality differences contribute. Coking coal and thermal coal are two very different products, as to quality and use. So are bread grain and feed grain, a difference reflected in the 50 pct higher price of the former. A port can be an importer of the one and exporter of the other. Richards Bay highlights the possibilities:

Segment	Exports	Imports
Cape	coal	coking coal, petcoke, potash
Panamax	coal, fertilizer	alumina, phosrock, potash, sulphur
Handy	cement, sands, ores	alumina, phosrock, potash, sulphur

When a large port is part of a large city it is helpful to know that these routinely import thermal coal, cement, salt, gravel and sand, and export scrap.

These principles are valid in ocean traffic. In coastal traffic, only single-commodity ports offer a realistic chance. It is almost certain that visits to Cedros Island (Morro Redondo) involve export of salt, those to Cozumel Island export of building materials, those to Glensanda export of aggregate and those to Port Hawkesbury and Port Tupper export of gypsum or import of coal. Otherwise, the possibilities are often too numerous for indirect reasoning. Either one has to ask shippers directly, which is impossible in practice, or one must make rigid assumptions such as, every other leg in cargo and every other in ballast. If the first guess is incorrect the continuation has a good chance of being incorrect, too. The very realistic possibility of multiporting adds to the confusion. Therefore, when a ship moves between closeby ports which are not of the single-commodity type it is as well to lump them together. That has been a usual solution in the area stretching from the Vancouver Island to the Columbia River, for example.

Wet legs Vessels constructed for both dry and wet bulk cargos (combis) were mentioned above (Table 2.3). Their holds are suited for both types of cargo and because the geographies of dry and wet trades are different the flexibility helps to minimise ballasting. The difficulty is cleaning the holds after a “dirty” (crude or heavy fuel oil) cargo. It takes time, becomes expensive and discourages frequent shifts, most likely triggered off by a change in relative freight rates. This suggests infrequent shifts, an opinion which did not tally with empirical observations in the studied years. Such vessels did exist but there were also numerous others alternating continuously between dry and wet cargos. Infrequent or frequent, the cleanup has made combis unpopular, new ones are seldom built and the existing ones are rather old. They operate primarily in the Atlantic, often in the Caribbean and Mediterranean, seldom in the Indian Ocean and never in the Pacific. They are less likely to come to the open charter market.

The task, in all simplicity, is to identify wet cargo legs and their attached ballast legs, and delete them from the file. Cargo legs are conceptually easier to handle than ballast legs and we begin with them. Their identification is tantamount to identifying oil ports. This principle is not watertight, however, because of the attached dry bulk goods. A refinery at an importing crude oil port needs chemicals for its processes and produces sulphur prills and petroleum coke as residuals. An exporting crude oil port, in turn, needs drilling equipment such as drill mud, and possibly ships salt in addition to oil, salt layers being a condition for the existence of crude oil reserves in the first place. The identification of a leg as dry-cargo or wet-cargo is consequently subject to error, actually grave error because the two types of flow probably move in opposite directions. Fortunately enough, LMIU had prepared a data file about combi movements with wet cargos (WetFile). The information about panamax combis was available also for this author.

This WetFile does not fully match the general movement file but discrepancies, usually omissions, are not serious. The intervening legs in the movement file can be partially identified as standard dry-cargo legs, for example Tubarao-Ghent. Partially they resemble liner trade or, still more accurately, company trade between operational centers (drilling equipment, refinery chemicals, inter-depot transfers). Panamaxes often ply such routes, paying regular visits to great many ports before coming back to a recognized crude oil export port. It is impossible to define any cargo status and the legs are considered wet-cargo. Irrespective of the wet or dry character they are company business and accessible to shipping companies only on a time charter basis if at all.

The itinerary revolving around Arzew, an Algerian port, helps in making matters tangible. Arzew exports crude oil and salt, both of which have a market in the US Gulf and Netherlands. It probably imports supplies from the Gulf as well. The WetFile suggests that cargos from Arzew to USG and Netherlands (Houston, Freeport, Corpus Christi, Terneuzen) are oil. Return cargos to the Continent can be oil products (selected alternative) as well as grain. Ventspils would be a similar case exporting both crude and fertilizers.

Arzew
Houston
New Orleans
Lisbon
Antwerp
Coryton
New York
Arzew
continues

continued
Terneuzen
Antwerp
Portugal
New York
Hampton Roads
Terneuzen
Arzew
Houston
continues

continued
Lake Charles
La Pallice
Algeciras
Port de Bouc
Tarragona
Arzew
Antwerp
Arzew
Immingham

Each time a port is identified with a wet cargo, it is marked with X in the movement file. Thereafter the lines, except for the first and last ones, are erased from the file. But also the first and last lines can be erased, and each of the variants will lead to a different set of ballast legs. Examples shed light on the details, although Exampel 4 is actually irrelevant because all indication about wet legs has been erased.

	Start		**Exp 1**		**Exp 2**		**Exp 3**		**Exp 4**	
Port	Cargo	Reg	Cargo	Reg	Cargo	Reg	Cargo	Reg	Cargo	Reg
Tubarao	1	116	1	116	1	116	1	116	1	124
Piombino	0	124	0	124	0	124	0	124	0	124
Arzew	X	124	X	124			X	124		
Terneuzen	X	123								
Immingham	X	123								
New York	X	112								
Houston	X	113								
Aruba	X	114	X	114	X	114				
Tubarao	1	116	1	116	1	116	1	116	1	116
Ghent	0	123	0	123	0	123	0	123	0	123
Ballast legs	124	124	124	124	124	114	124	124	124	116
	114	116	114	116	114	116	124	116		

Example 1 is more truthful than the others by giving the ballast legs to the beginning and from the end of the wet-leg string. But are these ballast legs relevant in a dry bulk study because they originate from an alien although integrated trade? An impossible question and in Example 4 they have been omitted. The choice is relevant for the number of legs and average distances. No answer is fully unambiguous because the dry-bulk itinerary might have developed differently without the wet legs. Example 2 tries to rationalize Example 1 by erasing the line which leads to a ballast leg of zero length on condition that intraregional distances are assumed to be zeros. Example 3 is inferior to Example 1 because it gives partially erroneous information about the dry-bulk ends of the ballast legs. About the distances it gives, by chance, an identical answer as Example 4.

It is obvious that there is no best answer to the “correct” ballasting distance nor the number of ballasting legs. This report follows Example 2, which gives the “correct” frequencies and a fair estimate of the distance. It also gives the correct dry-bulk regions of the ballasting legs. That information will not be used, however. When a Standby file is converted to a Readied file from which O-D matrices are calculated, each line with a X cargo identification (Example 2) is deleted. It means that only “pure” dry bulk trade is left. The solution is not identical with Example 4 where the X cargo identifications have well been deleted but the “faulty” ballast leg 124-116 has been left intact. Also this will disappear when the O-D file is prepared. Whether the solution is preferable to Example 4 is, again, a matter of judgement. Whichever the case, the practical significance is not too large because the X lines number only about 500 or 1.1 pct in the panamax segment where they are comparatively numerous.

These principles are applied only to wet trades. Gaps due to non-reporting (below) have been handled according to Example 4 rather than Example 2. The prime reason is economizing of effort, and the implications were realized only when the Cape and Panamax files had been Readied. A secondary reason is the difficulty to decide whether there is a real gap or just vessel inactivity.

Gaps Deletions make the movement files smaller but that is partially compensated by lines added. They originate from obvious gaps in the data. LMIU arranges ship movements into a chronological order by arrival dates, supported by geographical logic when dates are inaccurate (below), and gives each port visit a running identification number (move-id). The data file comprises only movements known to LMIU. Their consecutive numbering does not mean that each movement has been recorded.

> There are also plain errors. For example, Stockholm, Oslo, London, Nice and Casablanca are inserted into shuttle traffic between Poland and Finland (LLP 164201, 1997). There are no time gaps for visiting such outlying ports. The clerk obviously had no idea about geography.

LMIU does not add speculative lines. The result is incomplete itineraries, even about capesize vessels. There may be four consecutive lines about an iron ore export port at one-month intervals but not a hint about discharging ports. These may be private industrial ports where LMIU does not have an agent. Such gaps appear preferably in Indonesia, Korea, Taiwan and China. The Atlantic Sphere is covered much better.

Dates have corresponding shortcomings. The estimated arrival time is put equal to the estimated departure time from the previous port and the departure time equals the arrival time to the next port when no better information is available. The user is alerted to the approximate character with notes. This is the main reason why dates have been replaced by weeks in the Readied file.

In this study, apparently missing lines are added when reasonable assumptions can be made about the loading/discharging port/region and the arrival time. The purpose is clear: The inaccuracy thereby introduced is smaller than if the problem is ignored. Discharging ports/regions can usually be deduced from other, disclosed ones because many vessels trade for lengthy periods for the same country and even on the same routes. Passings of Gibraltar, Singapore, etc., available in the file Move, give welcome clues about general patterns. The arrival time is put approximately half way the neighboring times, the week first and thereafter the date, which in ocean traffic is the last day of the appropriate week. The time gap must be sufficiently wide for an additional port visit and the width influences the choice of ports/regions when several alternatives are available. On the other hand, a very wide gap, ostensibly comprising several port visits, is left intact. The same applies to gaps in the beginning and end of a year. These gaps observe the conclusions made about Example 4 (above).

Then there is a purely technical need to add lines. The need appears at ports where a visit involves both discharging and loading. First one cargo is discharged and thereafter another one loaded, but there is only one line in the data. One line makes sense when the interest is in ship movements and no attention is paid to cargo status. But an O-D matrix can be produced only when there are separate lines for discharging and loading. It means that one line with appropriate movement identification number must be added.

Everything included, the added lines amount to 2-5 pct of the total, depending on the segment and year (Table 2.4). Large capesizes have the best record which deteriorates towards the panamax segment, and rises again in the handy segments. There are several plausible explanations. Panamaxes are more prone to make short trips than capesizes which means that they are more likely to stay in the secretive Far East. They are also more likely to ply in coastal traffic and thereby escape the LMIU's network by definition. Recording in general is more likely to be incomplete in shuttle traffic, be it only for the frequency of port calls. But there are also opposite forces and they gain importance in the handy segments. Gaps are much easier to discover in a shuttle itinerary than a liner-type itinerary which is common in the handy segment. Handysizes, when not in liner-type traffic, practice less multiporting than panamaxes. Whatever the truth, this is still a partial picture because some movements must have escaped LMIU's net completely.

The lines deleted from CBO and COO itineraries were counted only in the panamax segment. Based on them, the overall upper limit of deleted lines can be estimated at 0.25 times the number of added lines. In relative terms and recalling the data's global coverage, an overall error percentage below 5 pct is very modest.

Table 2.4 Number of added and deleted lines by segment and year

Segment	Year	Total	Added	Pct	Wet ballast
Cape large	1995	3,201	26	0.8	
	1997	4,117	52	1.3	
	1998	4,328	38	0.9	
Cape small	1995	4,969	68	1.4	
	1997	4,537	125	2.8	
	1998	4,159	147	3.5	
Cape total		25,311	456	1.8	
Panamax	1995	13,738	587	4.3	136
	1997	14,794	670	4.5	177
	1998	15,601	767	4.9	186
Panamax total		44,133	2,024	4.6	499
HandyL	1997	15,155	333	2.2	
HandyM	1997	18,036	273	1.5	
HandyS	1997	14,514	375	2.6	
Handy total		47,705	981	2.1	
Overall total		117,149	3,461	3.0	

Source: Author.

Multiporting Coding completed, vessels were arranged by ship-id and move-id. The ideal pattern was that export and import ports alternated in a 1-0-1-0 sequence. This was more an exception than a rule, however. One reason was multiporting. It arises out of the collection and distribution of partial cargos at the both ends of a trip, due partially to shallow ports. Large ship size obviously promotes multiporting, other things being equal, and old ships are more prone to practise it because of their low capital cost. When the ocean leg in cargo is long multiporting sections are quite insignificant in relation to the total trip length and can be safely ignored.

Japan with its long coastline has a port for each major city, and many minor ones as well. It follows that ships discharging there often do it in 2-3 ports. One of the ports can actually be in Korea. The sequence may follow the vessel's overall route, from south to north when approaching from Australia, etc. But there are too many exceptions to encourage generalization. When the size of consignment is not a constraint port draft may limit free movement. Grain ships taking main cargo on the

Mississippi River between Baton Rouge and New Orleans get it topped at the Southwest Passage. It is standard practice in Argentina to start loading grain upriver La Plata and take the last consignments on the Patagonian coast, in Bahia Blanca or Necochea. Iron ore cargos for Swedish Steel are routinely discharged first in Oxelösund because the Luleå port is too shallow even for panamaxes when fully loaded. Capesizes bound for China with iron ore often discharge part cargo in Beilun (Ningpo) and finish the trip in Baoshan (Shanghai).

When multiporting occurs the first loading port and the last discharging port are kept and the intervening ports are deleted. First and last port can be defined either timewise or geographically (most distant from each other). Normally, although not always, the principles coincide. When they do not, the geographical alternative is adopted unless it leads to grossly unreasonable results timewise. That looks simple but there are analytical complications. First, the regional coding of the recorded loading and discharging ports becomes somewhat hapharzad, particularly in the Continent. Second, it is difficult to separate multiporting from short cargo (and attached ballast) legs. When a long ocean trip has been completed it may be necessary to “park” the vessel in short-term tasks until attractive ocean trips become available. Since there is no direct information about the cargo status of a leg it is many times impossible to decide whether a short leg is part of an ocean trip or a separate undertaking.

The third complication is that multiporting also occurs in intercontinental scale, and particularly with large capesizes. Having loaded in Hampton Roads, Brazil, or Australia, whether eastbound or westbound, they complement cargo in South Africa, at Richards Bay or Saldanha Bay. When eastbound this is very usual. Alternatively, Santa Marta replaces Richards Bay and Ponta da Madeira or Mormugao does the same for Saldanha Bay. Villanueva is often the first discharging port for iron ore cargos from Australia to Japan. Yanbu is a frequent port of call by handysizes between Australia and Continent. Less usual is the topping of iron ore cargos from St. Lawrence or Chile to the Far East in Brazil respective Australia. Coal from Puerto Bolivar to the Continent, topped with iron ore in Nouadhibou, is another rare possibility. Intercontinental multiporting is a possibility also for panamaxes. Hay Point via Richards Bay to Oxelösund and Luleå is a comparatively frequent trip. When the multiporting section makes up a large part of the trip the technique of deleting all but the first loading and last discharging port becomes less attractive. Its use depends entirely on the lack of a better alternative.

The deletion of evidence about multiporting has the negative consequence that some organizatorial connections will disappear simultaneously. These would be helpful when separating tied traffic from the rest (below) and analyzing simulation results (Subchapter 4.6).

Itineraries can be classified by the sequence of loading and discharging ports, their repetition in space and time, the length of leg or frequency of service, and possibly the number of ports if not of liner type. As in port coding, the segments were handled in the following order: CapeL, CapeS, Panamax, HandyS, HandyM, HandyL. That may have affected the quality of the result because of the learning effect.

Three or four basic itinerary patterns have been discovered:

StandardS	1---------------------------0	non-repetitive, long or variable legs
StandardM	1-1-1------------------0-0-0	non-repetitive, port clusters (multiporting)
Shuttle	1-0	repetitive, normally short legs
Liner-type	1-0-1-0-1-0-1--------0-1-0-1-0-1	repetitive, many short legs, loading and discharging alternate

Their importance varies by segment:

Pattern	**Handy**	**Panamax**	**Cape**
Standard simple	frequent	frequent	frequent
Standard multiport	milk rounds	frequent (grain)	often 2-port
Shuttle	exists	exists	
Liner-type	frequent	exists	

Multiporting and liner-type itineraries are close cousins. Operations easily assume liner-type features along a long coastline (Japan, West Coast South America, Mediterranean, Middle East Gulf, US Atlantic Coast) and when consignments are small. Small-volume commodities such as non-ferrous ores support the case. But the same itinerary can be considered normal multiporting when long ocean legs intervene. Knowledge of the transported commodity is helpful in making the choice. So are visits in typical container ports like Felixtowe and Hook of Holland. Yet, the choice is difficult enough as this author realised when working with Scandinavia, an area he thought to be fully familiar with. A special kind of liner-type itinerary is "milk round", circular trip with lengthy multiporting sections for which it may be impossible to decide the starting/end point. The general direction is rather obvious but exceptions occur when cargos are not ready or port conditions so dictate. An itinerary in the Indonesian archipelago has a good chance of being a milk round. Liner-type becomes conventional liner when a time schedule is announced, the port sequence is strictly adhered to, and loading and discharging mix freely. This cannot be seen from the LMIU Movement Data.

Shuttle traffic necessitates that a given pair of ports (or consistent port groups, usual at aluminium companies) appears sequentially (no intervening ports) at least 4-5 times during a year. These vessels normally have local (within-region) movement patterns. Local patterns are also much easier to identify than global ones which need time to develop. If only the loading or discharging port remains constant the traffic is not genuine shuttle, although it can very well be traffic tied to a manufacturer, mining or trading company. Tied traffic is common on lively trafficced routes, e.g., from Australia and New Orleans (to Mexico). Its exact identification is difficult in practice, except when the same narrowly specialized port (Cozumel Island, Glensanda, Morro Redondo) is the origin of all the cargo legs.

When figuring out itineraries one meets three problems: wet (oil) cargos, gaps and multiporting. All these have been addressed above and we can continue by finding the 1-0-1-0 sequencies. Ports are crucial for success: no gaps, no superfluous ones, correct sequence. When these conditions are met it is always possible to identify probable cargo and ballast legs. The principles are identical in both cases. The intial port coding (above) will be checked and, if necessary, corrected. Normally, ports are about correct. The whereabouts of a vessel on the globe is seldom in doubt, although old ships are used for odd jobs in odd countries. The cargo legs and the direction of traffic can be problematic. The technique is to identify a certain port pair, e.g. Necochea to Rotterdam, and chain from it. If another certain port pair synchronizing with the first one, e.g. New Orleans to Damietta, can subsequently be found it partially validates the decision. And so on, although without some guesswork it would be absolutely impossible to get any result at all. It is also helpful to realise that ship movements often have directional logic. When a ship returns it is more likely to pick up a new cargo rather than deliver one. In conclusion, it is possible to create a firm idea about the operational area of each vessel but not necessarily about its detailed itinerary. But because the main purpose here is to create O-D matrices rather than itineraries the accuracy is acceptable.

The basic technique was applied almost exclusively in the cape, panamax and large handy segments. In the medium and small handy segments it was only partially useful. The reason was liner-type itineraries. These appeared occasionally among panamaxes and large handysizes but became a problem in the medium and small handy segments. The explanation is the small size of ships and the numerous commodities which they carry. The size allows visits to small ports about which little information is available. The commodities are often so-called minor bulk commodities such as cement, fertilizer, nonferrous ores and concentrates, sugar, china clay, fertilizer raw materials, chips, logs and paper & pulp, with plants and shipping ports everywhere. The frequent co-location of a cement and fertilizer plant is particularly troublesome. Cement plants have their own distribution systems and fairly regular deliveries. It is very difficult to keep these shipments apart from non-programmed, genuine bulk shipments. There are simply too many cargo-leg alternatives. Short cargo legs make the split between cargo and ballast still more difficult. In short, liner-type handy itineraries become often pure guesses. The difficulty of the task gained crisis proportions and led to uncoordinated experiments. Uncoordination is horrendously unscientific but the effort to repeat work already done was also excessive. Remedy is sought in a forthcoming paragraph where liner-type operations are simply excluded from the analysis. The motivation is that they are incompatible with the traditional concept of "tramp" which in modern times has been renamed euphemistically as "dry bulk carrier". The following discussion is consequently superfluous here but may be useful in some future analysis where the focus is on liner-type and pure liner shipping.

Four techniques were experimented:

1) Indicate every other leg as cargo and every other as ballast. The solution will lead to some very erroneous decisions although, on average, it should work adequately (Appendix 2.3).
2) Find the turning point of the trip (e.g., Hakata in Japan), indicate it as the loading port and the first port of the next region (e.g., Los Angeles) as the discharging port. Delete the rest.
3) Indicate recognized loading and discharging ports as such and delete the rest. The philosophy is that there are typical loading and discharging regions and their numerous ports are only a variant of multiporting by typical tramps.
4) Indicate a region's first port as the discharging port and last port as the loading port. Delete the rest. The idea is that there is an even flow of traffic all along the chain and that each region has something to import from the preceding region and something to export to the following one.

All over the board, some attempt was made to identify long legs as cargo legs. This is reasonable because short cargo legs and long ballast legs do not make sense. But it was a guess, anyway.

There was a gut feeling that each technique had merit depending on the specific situation, although it was very difficult to be specific about the "situation" when it turned up. It was impossible to combine the techniques in the same itinerary and their choice affected the main diagonal of the O-D matrix. The impact on off-diagonal macroregions, by contrast, was modest, whichever the technique. It follows that segment comparability is best when attention is paid only to off-diagonal macro flows. That may be unacceptable, however, because within-macroregion and even within-region legs offer welcome parking lots, particularly in the Continent and Far East.

The purpose of deleting legs from the **Readied files** is to come closer to the polypolistic spot market which is the real topic of this study. The deletions are made best when the one-column port sequence has been converted into a two-column variant and the port-to-port legs can be read horizontally (Appendix 2.1C). The conversion involves a technical detail which affects the comparability of the years, not decisively but anyway. Ports preceding Antwerpen and following Taranto are not known because they are outside the calender year. If also the next year would be analyzed Taranto would have another port pair there. Legs spanning two calendar years were not deleted at the turn of 1997/8, to save work. At the both ends of 1995 they are unknown and this impairs comparability. The following deletions are made for technical or conceptual reasons.

Technical deletions include all wet legs, all within-port legs, except the one for ballasting, and other short, within-region legs which are unessential for the study. In some size segments the deletions were implemented already during the coding or itinerary building stage. The rest is done now, which goes a long way towards making all segments comparable. Examples of legs which are deleted are Puget Sound area, Vancouver-Columbia River, Durban-Richards Bay, Houston-Tampa, Nikolayevsk-Odessa, Yokohama-Nagoya and Kaohsiung-Taichung. By comparison, Narvik-Bremen, Gdynia-Mäntyluoto, Houston-Tampico, Samarinda-Jakarta, Vostochnyy-Fukuyama, and attached ballast legs, are retained. The deletions are made to arrive to the **Stripped files**.

Conceptual deletions include liner-type and specialized traffic. Liner-type traffic does not fit the dry bulk concept at all. The BWC vessels do but their advanced specialization makes them uneconomic for dry bulk cargos at large. It follows that liner-type and BWC lines will be deleted. Deletions are done by vessel and apply to the whole itinerary or its part. For example, the deletion of Esquivel-Mosjoen (alumina) legs at one vessel does not mean that similar legs are excluded at other vessels. A separate column with appropriate itinerary codes is inserted for identification. When the O-D matrix is prepared all lines with the appropriate itinerary code and their attached ballast legs will be deleted. The deletions are made to arrive to the **Stripped files**. The topical O-D matrices are in Appendix 2.4.

Tied traffic is conducted by owned or long-term chartered vessels. It can well be dry bulk but falls outside the polypolistic spot market concept. That comes manifest in the stability of either loading or discharging ports, or both, under long periods. Frequent service between the same closeby ports is called shuttle traffic. Shuttle is easy to identify but the general case can be problematic because a calender year may well truncate long sequences. Uncertain cases are given the benefit of doubt and 4-5 port pairs are considered sufficient for inclusion. No difference is made between within-region and between-region traffic. Traffic coordinated by charterers, perhaps of different nationality, and agreements to fill up vessel capacity are also considered tied. Large steel companies are typical cases:

Port	**Company**
Port Kembla & Whyalla & Newcastle	BHP Steel
Immingham & Port Talbot & Tees	British Steel
Kwangyang & Pohang	Posco
Mizushima & Chiba	Kawasaki
Wakayama & Kokura & Kashima	Sumitomo
Tobata & Nagoya & Kizarasu & Oita	Nippon Steel
Kawasaki & Fukuyama	NKK
Kakogawa & Kobe	Kobe Steel
Kure	Nisshin Steel

To them is added other heavy industry belonging to the same conglomerates. The Muroran cement plant (Nippon Steel) is a typical case. In combination, these conglomerates easily originate 30-45 pct of the capesize traffic worldwide, and much more arriving to their respective countries. Such agglomerations are less usual in Europe, or their shipping patterns are less concentrated (cf., Timmermann and McConville 1996, 46). These additional deletions are made to arrive to the **Final files**. The O-D matrices are in Appendix 2.5.

To recapitulate, the deletions work as follows:

	Cape	**Pana**	**Handy**
Stripped			
Wet cargos	yes	yes	yes
Legs within port/closeby ports	yes	yes	yes
Liner-type		yes	yes
BWC		yes	yes
Final			
Shuttle	yes	yes	yes
Other tied	yes	yes	yes

Evaluation When codings and deletions have been completed and the result converted to a matrix it easily becomes an undisputed truth. We know better. Such truth is very relative. The result is not intended for analyzing individual itineraries but to gain a holistic idea about interaction patterns. It does not matter, then, that many within-region legs are identified incorrectly as long as the flows between macroregions are approximately correct, what they probably are. Recalling the way how multiporting has been handled the matrices are best kept at the 10 x 10 macroregional level.

It is always a matter of judgement how much detail should be included. Some itineraries are complicated and interesting in their own right, revealing the complexities of practical shipping. Here, the focus is on the overall view. This is the classical difference between micro and macro views, business economics and economics. As the economist assumes the rational man, with some modifications, we assume the typical shipowner/manager and project his/her activities as the 'typical vessel".

It has by now become clear that eventual cargo/ballast code (1/0) available in the original data would not be sufficient for fully computerized handling. Complementary manual handling would be necessary. It requires certain qualifications, a mental picture of distances, the relative location of ports, and the likely cargos involved. Handbooks tell much but the best source is probably a senior broker. That is a good reason to integrate consultancy and brokerage.

2.2 Origin-destination analysis

The analysis will make use of three sets of matrices: Original, Stripped and Final, produced from the appropriate files. Original matrices include all ports visited for whatever purpose by all types of vessel possibly carrying dry bulk goods. They are used for plain description. Stripped matrices include only ports visited for dry bulk handling, exclusive specialized wood chips carriers and liner-type operations. These matrices are also used for description, but mostly for independent analyses, and as input for the cape segments in the simulation. Final matrices include only bulk legs of the polypolistic market. These legs are used as input for the panamax and handy segments in the simulation.

All the matrices will be annual although occasionally weeks 104-155 make up year 1995, weeks 209-260 year 1997, and weeks 261-312 year 1998 (the first full week in 1993 is week 1). The vessel segments will be CapeL, CapeS, Panamax and Handy, the two first often combined into Cape. Cargo and ballast legs will get their own matrices. The matrices will display both regions and macroregions, depending on the case. There are too many matrices to be published *in toto*. Therefore, they are used selectively, to highlight special features and to create an overall mental image. The year 1997 is used as the base for comparisons and tests of similarity.

2.2.1 Largest ports

The **Original files** are used for getting acquantained with the largest individual ports. This is the only occasion when a systematic discussion is taken to that level. It does not serve any deeper purpose but is thought to make the ensuing abstractions more tangible. The aggregate statistics first (Table 2.5). The number of ports and port visits roughly doubles by moving from the cape, to the panamax, and finally to the handy segment; whereas the average vessel size halves. Multiplication of port visits and average vessel size suggests that the transportation work done by capesizes and panamaxes is about the same, whereas that by handysizes is 50 pct larger. Interesting as the products are it should be recalled that handysizes have comparatively much liner-type traffic with numerous port visits. Also the average leg distances are shorter at handysizes than the other segments. Ultimately it turns out, however, that the estimate is reasonably close to the mark (Table 2.10).

Table 2.5 Overview of Original files, 1997

	Cape	Panamax	Handy
Ports, number	353	881	1,478
Port visits, number	9,654	20,294	59,128
Top 30, pct of visits	59	35	23
Vessel dwt, average	152,106	65,862	37,098
Visits * avg dwt, mill.	1,468	1,337	2,194

Note: The corrections mentioned in the text below are not yet implemented.
Source: LMIU Movement Data.

The Top 30 ports by number of visits in each segment are taken for closer scrutiny (Table 2.6). Neither the concept of port, nor of visit, is very exact. Important subports and roadsteads are sometimes recorded separately. Familiar cases are Europoort & Rotterdam, Bremen & Brake, Chennai & Chennai Roads. We do not interfere with subports but delete roadsteads. Once in the port, movements between its parts and the roadstead may be recorded as separate visits but under the same label. This is the case with Vancouver and the same used to apply to Richards Bay until questioning by this author prompted a change (LHS 2001, 40). In Vancouver, there appears to be roughly 1.5 phantom ‘visits” per one real visit in the cape segment and 2.0 in the panamax and handy segments. The counts in Table 2.6 are scaled down accordingly. Plotted against the rank, the count of port visits forms a concave curve, steeper at first and then gently sloping. The initial slope is quite steep in the handy and panamax segments. Thereafter the scatterplot of the panamax and cape segments are close to each other, suggesting that the segments are close replicas and quite distinct from the handy segment. Such a conclusion is premature as the comparison of the individual ports shows. In other words, there are very few really important ports.

The Top 30 comprise in total 62 ports. Their joint belonging to the three segments gives a rapid idea about geographical specialization (Table 2.7). Almost two-thirds of the ports are loners. They are

Table 2.6 The 30 largest ports (visits) by ship segment, 1997

Rank	Cape Port	Visit	Panamax Port	Visit	Handy Port	Visit
1	Singapore	438	New Orleans	894	New Orleans	1358
2	Europoort	399	Singapore	763	Vancouver	708
3	Port Hedland	370	Newcastle (AUS)	411	Columbia River	656
4	Dampier	315	Taichung	360	Antwerp	614
5	Tubarao	272	Gibraltar	356	Kaohsiung	589
6	Richards Bay	263	Richards Bay	317	Hongkong	580
7	Newcastle (AUS)	231	Vancouver	270	Singapore	571
8	Hampton Roads	205	Kaohsiung	224	Richards Bay	554
9	Kisarazu	205	Hay Point	223	Incheon	531
10	Ponta da Madeira	193	Los Angeles	221	Visakhapatnam	461
11	Gibraltar	189	Gladstone	203	Shanghai	452
12	Hay Point	189	Hampton Roads	176	Durban	427
13	Oita	174	Columbia River	165	Santos	420
14	Fukuyama	166	Kamsar	165	Nagoya	411
15	Kashima	165	Ghent	161	Houston	410
16	Gwangyang	161	Santos	160	Port Klang	410
17	Port Walcott	161	Dampier	156	Chennai	396
18	Nagoya	159	Xingang	153	Chiba	361
19	Saldanha Bay	153	Mobile	152	Taichung	357
20	Pohang	152	San Lorenzo	152	Osaka	347
21	Mizushima	147	Tubarao	151	Ko Sichang	343
22	Kaohsiung	146	Incheon	150	Haldia	324
23	Vancouver	116	Houston	145	Newcastle (AUS)	317
24	Gladstone	112	Weipa	140	Fremantle	312
25	Kawasaki	109	Paranagua	136	Capetown	301
26	Dunkirk	105	Europoort	131	Philadelphia	300
27	Port Kembla	101	Port M. Bin Qasim	128	San Lorenzo	296
28	Ymuiden	101	Hongkong	126	Paranagua	292
29	Port Cartier	100	Kashima	125	Jakarta	291
30	Kakogawa	98	Baltimore	123	Rotterdam	287

Source: LMIU Movement Data.

Table 2.7 Cohabitation of the 30 largest ports by segment, 1997

Segments	Cases	Segments	Cases
C	16	C-P	8
P	8	P-H	9
H	15	C-H	1
		C-P-H	5

Source: LMIU Movement Data.

important only in one segment. That speaks strongly for the specialization of dry bulk shipping, particularly in the cape segment in which the Top 30 account for 59 pct of visits (Table 2.5). The individuality of the handy segment originates from minor bulk commodities and comparative focusing on less industrialized countries, both leading to smaller consignments and shallower ports. Cohabitation succeeds best with neighboring segments. The only port with a cape-handy combination is Nagoya which has a large steel mill and is, in addition, a populous city with a manysided manufacturing industry. Among the five ports (Kaohsiung, Newcastle, Richards Bay, Singapore and Vancouver) which appear in each segment, Singapore benefits from its bunkering role. Kaohsiung is the main bulk port of Taiwan. The three others are manysided bulk ports specializing particularly in coal and grain. In each segment bunkering visits figure prominently. Gibraltar & Algeciras gain their rank almost exclusively from bunkering which is very important also in Singapore, Los Angeles and Capetown.

A short characterization of the individual ports follows. The **cape** export ports ship either iron ore (Port Hedland, Dampier, Tubarao, Ponta da Madeira, Saldanha Bay and Port Cartier) or coal (Richards Bay, Newcastle, Hampton Roads, Hay Point, Vancouver, Gladstone and Port Kembla) from Australia, Brazil, South Africa and Canada. The import ports serve steel plants in Japan, Korea, Taiwan, Continent and Australia. In 2004 (time of writing), China would have joined the pack. The **panamax** export ports are partially the same as in the cape segment, with the coke port Xingang as a new member. To them come grain & soya ports New Orleans, Vancouver, Columbia River, Santos, San Lorenzo, Houston and Paranagua, and two bauxite ports Kamsar and Weipa. It is noteworthy that all the large grain ports are in America, either North or South. The import ports locate in Korea, Continent, Pakistan and Japan. Two US ports, Baltimore and Mobile are both exporters and importers. The **handy** segment is the most varied of the three. The large grain & soya ports are complemented by Ko Sichang (tapioca) and Fremantle. Also the new iron ore and coal ports Visakhapatnam, Durban, Chennai and Haldia are dotting the coasts of the Indian Ocean. Most import ports are again found in Continent, China, Taiwan, Japan and Korea. But also Malaysia, US Atlantic Coast and Indonesia are represented by one port each. Antwerp, Port Klang and Jakarta, in particular, derive their rank from liner-type traffic. Antwerp is the closest port of NW Europe to the Far East and the Americas. Port Klang is, together with Singapore, a gateway to Southeast Asia.

2.2.2 Movement patterns

The **Stripped files** are the main data body for studying vessel movements (Appendix 2.4). They have been stripped from oil cargos and liner-type traffic, legs from newbuilding yards and to scrapping sites, and within-port movements, but they include all shuttle and tied traffic. They are what a layman would probably understand by the term "dry bulk traffic". It is the same body of traffic which was the topic of the original study and counted there 7,132 cargo legs (LHS 2001, Table 5.1). The current estimate 7,185 legs is practically the same.

Theoretically, ship movements including only cargo and ballast legs should split 50/50 between the two. Multiporting has the capacity to upset the theoretical balance but it was cleansed away with the specific purpose to arrive to a 1-0-1-0 sequence. Still, the even split has not been achieved but there is a consistent and not always insignificant surplus of cargo legs (Table 2.8). The likely explanation is twofold. When oil cargos were excluded from the Readied files one ballast leg was substituted for two. There may also have been a subconscious tendency to begin a calendar year with a cargo rather than ballast leg. Whatever the explanation, the difference will not upset the forthcoming conclusions and often it will have no identifiable effect whatsoever.

Table 2.8 Stripped O-D leg numbers

Segments	Year	Cargo	Ballast	Crg/Ball
Cape	1995	3,602	3,540	1.018
	1997	4,072	4,017	1.014
	1998	3,793	3,770	1.006
Panamax	1995	6,240	5,906	1.057
	1997	7,185	6,701	1.072
	1998	7,168	6,807	1.053
Handy	1997	19,227	18,476	1.041

Source: Appendix 2.4.

There are altogether 10 Stripped cargo files, three for the large capesizes, small capesizes and panamaxes each (1995, 1997, 1998) and one for the handysizes (1997), plus a similar number of ballast files. They all display different movement patterns but it is helpful if some differences can be considered unessential because that will simplify the discussion. The mesh of 10 macroregions gives a better overview than 37 regions with many almost empty cells (Fig. 2.1). The split between cargo and ballast cannot be overlooked. When only cells with at least one pct of the legs are considered an essentially similar pattern comes out whichever the segment, a strong case for the stabilizing influence of basic geography. An Index comparing the observed cell entries with uniformly distributed (expected) entries is more informative (Fig. 2.1, Notes). The Cargo Index becomes smaller with decreasing ship size, i.e., small cargos are less clustered than large ones. By contrast, the Ballast Index is largest in the handy segment, implying that ballast legs are exceptionally short. Most are but there are also long ballast legs, although normally unable to reach the 1 pct threshold. Whether they are real or the result of inaccurate itinerary construction is open to speculation. Be that as it may, their impact will be felt when calculating the workload in Subchapter 2.2.4.

But are the years within a segment so different that each needs separate commentary? And is it necessary to keep the large and small cape segments separate? Pearson correlation coefficients with cell frequencies as observations between all segment and year combinations shed light on the matter (Table 2.9). As expected, the correlations are highest within a segment and decline when moving to more distant segments. Traffic patterns rest on the firm foundation of regional supply and demand. Trade policies finetune the flows but the underlying capacities of sellers and buyers/consumers are not easily changed. The scope for differences is clearly larger between segments who meet partially different needs and face different movement constraints. The situation can be expressed as a formal function by plotting the average correlations between segments against the differences in average vessel size (Fig. 2.2). Observations which have a handy component locate at a lower level than the rest which underlines the segment’s uniqueness, a familiar fact for shipping analysts. The conclusion is that it is comparatively safe to focus exclusively on the year 1997 but that it is important to keep the segments apart, also small and large capesizes.

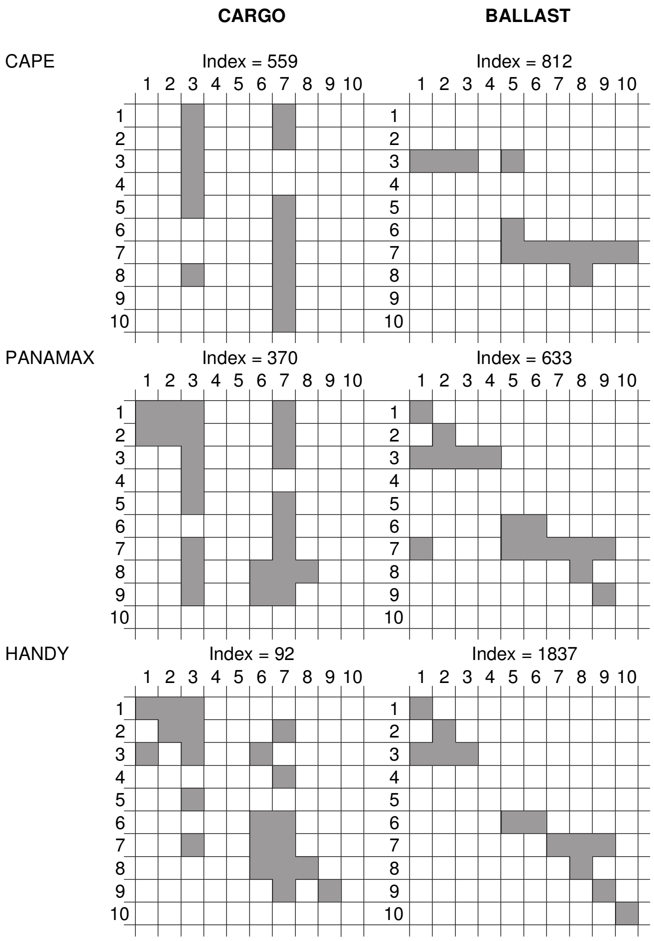

CARGO **BALLAST**

CAPE

Index = 559

	1	2	3	4	5	6	7	8	9	10
1			▨				▨			
2			▨				▨			
3			▨							
4			▨							
5			▨							
6							▨			
7							▨			
8			▨				▨			
9							▨			
10							▨			

Index = 812

	1	2	3	4	5	6	7	8	9	10
1										
2										
3	▨	▨	▨		▨					
4										
5										
6					▨					
7					▨	▨	▨	▨	▨	▨
8								▨		
9										
10										

PANAMAX

Index = 370

	1	2	3	4	5	6	7	8	9	10
1	▨	▨	▨				▨			
2	▨	▨	▨				▨			
3			▨				▨			
4			▨							
5			▨				▨			
6							▨			
7			▨				▨			
8			▨			▨	▨	▨		
9			▨			▨	▨			
10										

Index = 633

	1	2	3	4	5	6	7	8	9	10
1	▨									
2		▨								
3	▨	▨	▨	▨						
4										
5										
6					▨	▨				
7	▨				▨	▨	▨	▨	▨	
8								▨		
9									▨	
10										

HANDY

Index = 92

	1	2	3	4	5	6	7	8	9	10
1	▨	▨	▨							
2		▨	▨				▨			
3	▨		▨			▨				
4							▨			
5			▨							
6						▨	▨			
7			▨			▨	▨			
8						▨	▨	▨		
9							▨		▨	
10										

Index = 1837

	1	2	3	4	5	6	7	8	9	10
1	▨									
2		▨								
3	▨	▨	▨							
4										
5										
6					▨	▨				
7							▨	▨	▨	
8								▨		
9									▨	
10										▨

Figure 2.1 Flow patterns, 1997

Notes: Shaded cells have at least 1 pct of legs. Index = $\Sigma(\text{Obs-Exp})^2/1000*\text{Exp}$.
Source: Laulajainen 2007, Fig. 3. Courtesy of Elsevier Ltd.

Table 2.9 Pearson correlation coefficients for cargo legs between macroregions

	CL-95	CL-97	CL-98	CS-95	CS-97	CS-98	P-95	P-97	P-98	H-97
CL-95	1.000	0.993	0.989	0.762	0.792	0.853	0.622	0.692	0.683	0.262
CL-97	0.993	1.000	0.994	0.754	0.818	0.872	0.663	0.722	0.706	0.277
CL-98	0.989	0.994	1.000	0.790	0.847	0.894	0.664	0.725	0.722	0.293
CS-95	0.726	0.754	0.790	1.000	0.975	0.940	0.805	0.791	0.788	0.431
CS-97	0.792	0.818	0.847	0.975	1.000	0.979	0.806	0.817	0.804	0.434
CS-98	0.853	0.872	0.894	0.940	0.979	1.000	0.742	0.781	0.773	0.420
P-95	0.622	0.663	0.664	0.805	0.806	0.742	1.000	0.962	0.932	0.600
P-97	0.692	0.722	0.725	0.791	0.817	0.781	0.962	1.000	0.972	0.684
P-98	0.683	0.706	0.722	0.788	0.804	0.773	0.932	0.972	1.000	0.681
H-97	0.262	0.277	0.293	0.431	0.434	0.420	0.600	0.684	0.681	1.000

Notes: Stripped files in the cape segments and Final files in the panamax and handy segments.
All correlations significant at the 1 pct level (two-tailed).
Source: LMIU Movement Data.

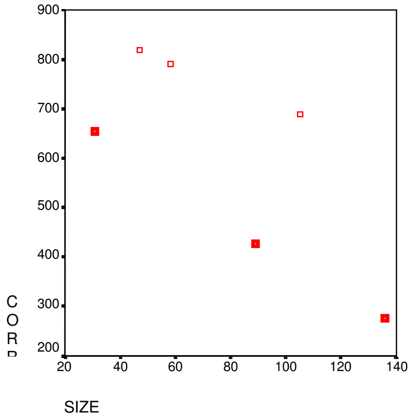

Figure 2.2 Difference in size, difference in movement

Notes: Large marker indicates observation with a handysize component.
Correl = 994 - 3.258 Size - 263 Handy; R-sqr(adj) = 0.979; SEE = 31.4
Sources: Tables 1.8 and 2.10.

2.2.3 Tonnage balances

Cell frequencies bring us to the regional tonnage balances. Arrivals and departures, when aggregated over longer time periods or over all regions, always match by mathematical necessity. There is a complete balance. That is not an issue. The issue is whether they do it also within the time horizon of charterers and shipowners. We will not descend to this detailed level but content ourselves with annual figures by region (37-region mesh), which can be quite informative in their own right (Appendix 2.6). The following accounting balance applies:

arrivals in cargo departures in cargo
arrivals in ballast departures in ballast
arrivals total = departures total.

So much is clear. But unexpectedly, two other equations also hold quite well:

arrivals in cargo = departures in ballast and
arrivals in ballast = departures in cargo.

These equations tell that it is difficult to find a new cargo or an empty vessel within the same region at short notice. Otherwise discharged vessels would not leave in ballast and empty ships would not be chartered outside the region. This conclusion applies to all segments. Fortunately enough, many vessels need not ballast very far because the regions are of moderate size. But purely formally there is a lack of temporal coordination and it leads, again formally, to regional polarization, regions with tonnage surplus and regions with tonnage shortage.

That polarization can be made tangible in many ways. We begin with histograms where all arrivals, alternatively departures in cargo, by region are on the horizontal axis and cell numbers on the vertical one (Fig. 2.3). The polarization grows with ship size as can be expected. Large vessels ply comparatively few routes whereas handysizes, as the name reveals, are versatile. When the departures in cargo are given as a percentage of arrivals the content changes. It is less about concentration as such but whether regions are clearcut exporters and importers, or not. The answer is obvious: In the cape segment they are, in the handy segment they are not, and the panamax segment is in-between. Capesizes cannot find so many cargos of suitable size and ports of suitable depth, the other segments can. For capesizes, the following opinion holds to an extent (Wijnolst and Wergeland 1996, 316):

> "... most bulk trades are shuttle business, where the carrier goes laden one way and
> ballast on the same return leg."

That is not the whole truth, however, when individual vessels are the units of observation. There is much shuttle traffic but it does not dominate. From a general angle, such statements are unfortunate because they intimidate further research. It was very close that the original research (LHS 2001) was never started because of this particular formulation (‘most bulk trades’).

Arrivals

Handy 0.90 **Panamax 1.08** **Cape 1.63**

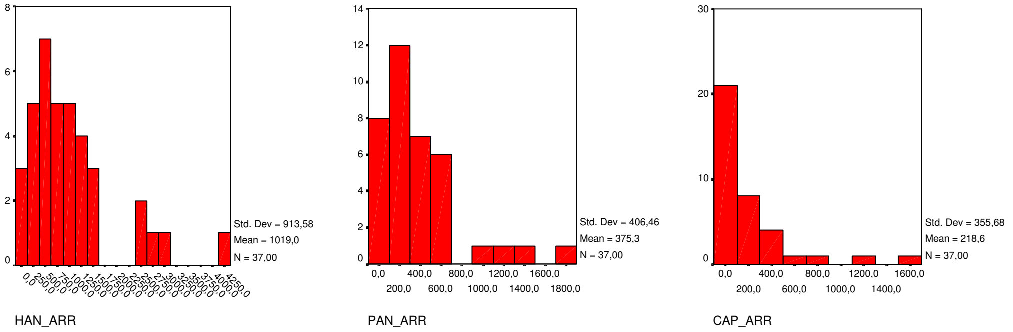

Departures in cargo

Handy 0.74 **Panamax 1.05** **Cape 1.67**

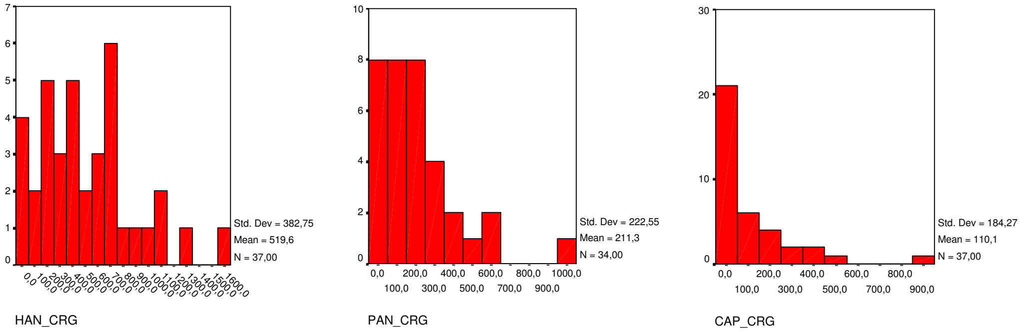

Cargo out of arrivals, pct

Handy 0.36 ; (57) **Panamax 0.57 ; (59)** **Cape 0.73 ; (62)**

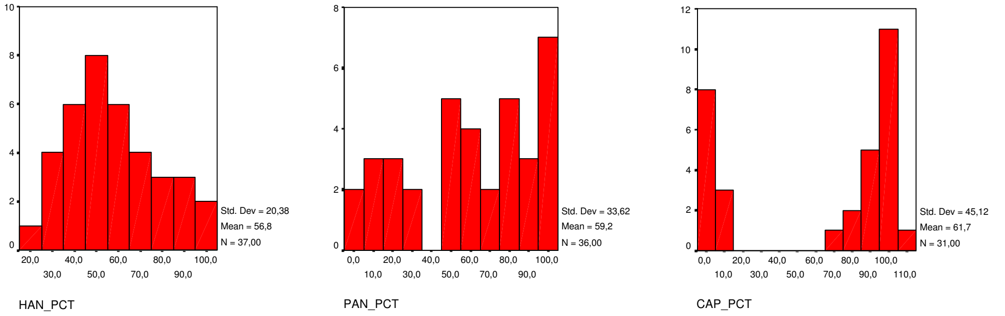

Figure 2.3 Arrivals and departures by region (37), 1997

Legend: Standard deviation / mean; (pct).
Source: LMIU Movement Data.

The observation about the immediate non-availability of cargos and empty vessels is important but not disastrous logistically because many regions are quite small. If a cargo or an empty vessel can be found in a neighboring region it will mitigate the polarisation. Converting the histograms to thematic maps will shed light on that possibility (Fig. 2.4). Again, the answer varies by segment. In the cape segment, the possibility is slight, everywhere. The split between exporting and importing regions is very sharp and these group into vast continental areas. That happens in spite of iron ore shipments from Narvik and Nouadhibou, too small to be visible here. In the Far East, such within-region traffic simply does not exist. Another noteworthy feature is that the Middle East Gulf (MEG) comes out as an importing region, albeit small one in this segment. The front of the traditional industrialized world has been broken there. The panamax segment is less extreme. The Baltic and North American Atlantic coast have a balanced mixture of exports and imports. So have India, Southeast Asia, northern China and the long coastline from Mexico to Chile. Fertilizers and coal from the Baltic, alumina to eastern Canada, construction materials to eastern USA, grain and coal to India and coke from northern China are among the commodities which stand for the better balance. The trend continues in the handy segment. Most of the Americas, in particular, is quite balanced. Pronounced exporters remain in the marginal regions, Australia, South Africa, Argentina, British Columbia & Columbia River and Great Lakes. The industrial core areas of Continent & Far East, and the Middle East, are importers as previously. Also the commodity mixes are similar: grain, iron ore, coal (not to Middle East), other raw materials. From the shipowner angle, the chances of finding closeby cargos and empty vessels are clearly best in the handy segment.

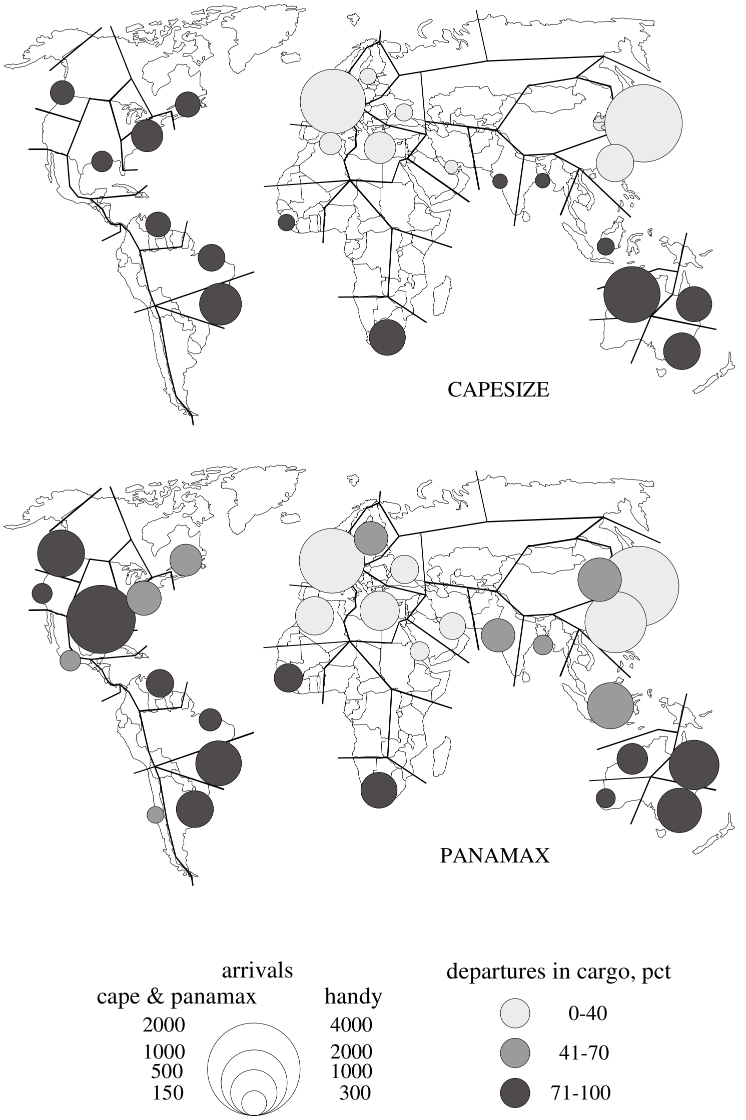

Figure 2.4 Arrivals and departures by segment and region, 1997

Source: LMIU Movement Data.

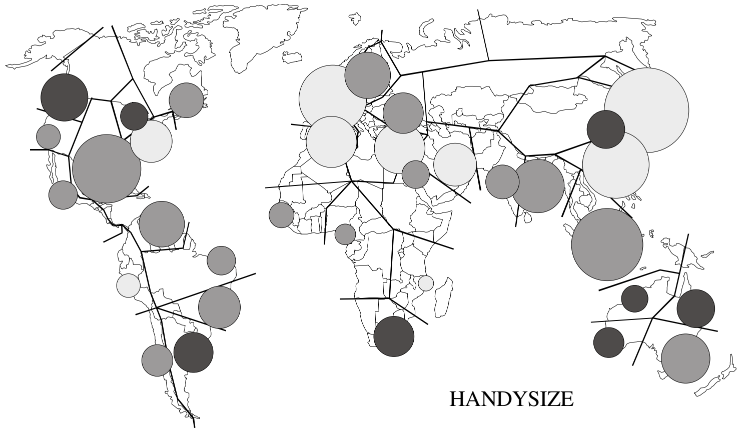

A sound idea of proportions has been gained but they are for a calendar year which is a very long period operationally. The Simulator uses week as the basic time unit and the ensuing smaller frequencies are easier to memorize and put into a context. The conversion is made by dividing the cell frequencies of the O-D macroregion matrices by 50. In other words, is assumed that a year consists of 50 weeks, a reasonable assumption recalling that our years have all the time been approximations of calendar years and that a year contains several holiday periods, not necessarily simultaneous in every corner of world. To facilitate the overview, results are given as stylized "maps" (graphs) rather than matrices (Figs. 2.5 and 2.6). Macroregions are placed in approximate locations in relation to each other, titled by abbreviated names (Fig. 1.6), and connected by arrows proportional to the volume of the flow. Each ship segment has a minimum displayed flow, selected so that there is approximately the same number of flows per segment, which facilitates visual comparison. The graphs qualify also as rough checks about the tonnage balances by macroregion. It is possible to estimate, for example, that the panamax balance in SAF is approximately -0.4 which is the difference of displayed incoming ballast (+) and outgoing cargo (-) flows. We look at the graphs by segment, cargo first and ballast thereafter.

Figures 2.5 and 2.6 Cargo and ballast legs/week on major routes, 1997 (pp. 56-57)

Notes: Minimum frequencies displayed: cape = 1.0, panamax = 2.0, handy = 4.5. Flow widths in the handy segment one half of those in the cape and panamax segments.

Source: LMIU Movement Data.

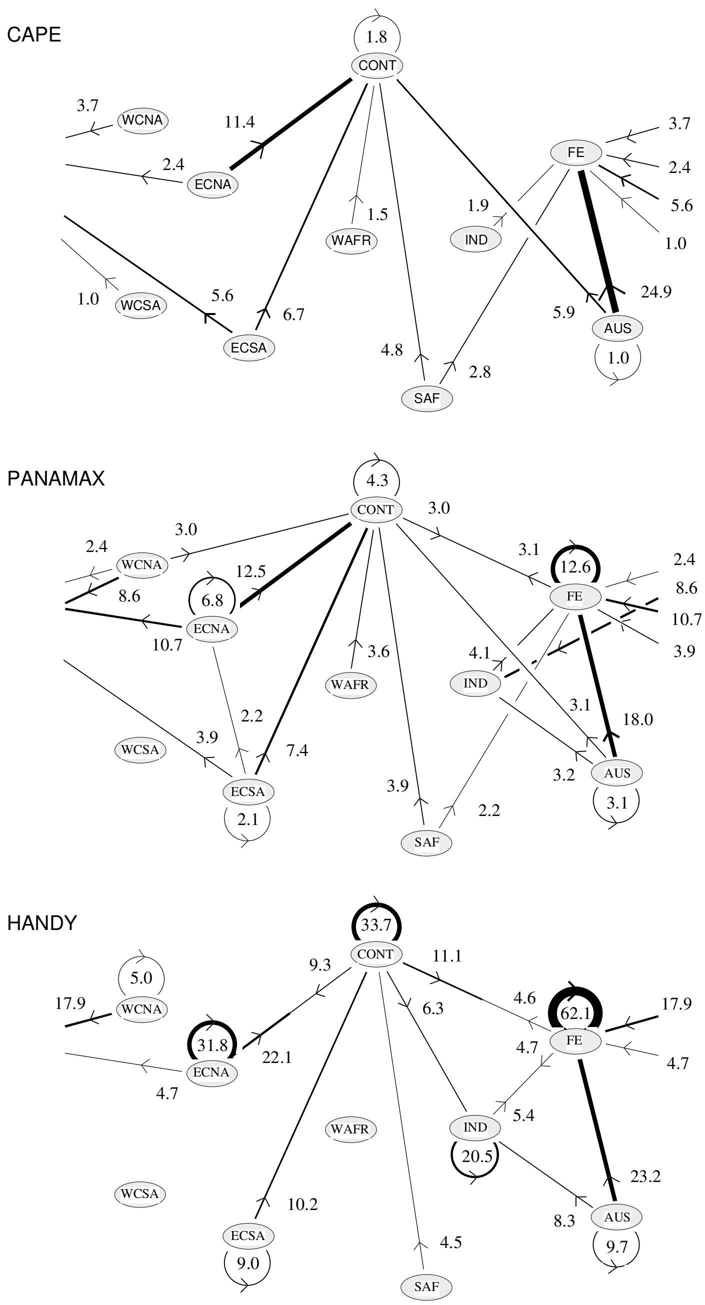

CARGO

CAPE

PANAMAX

HANDY

BALLAST

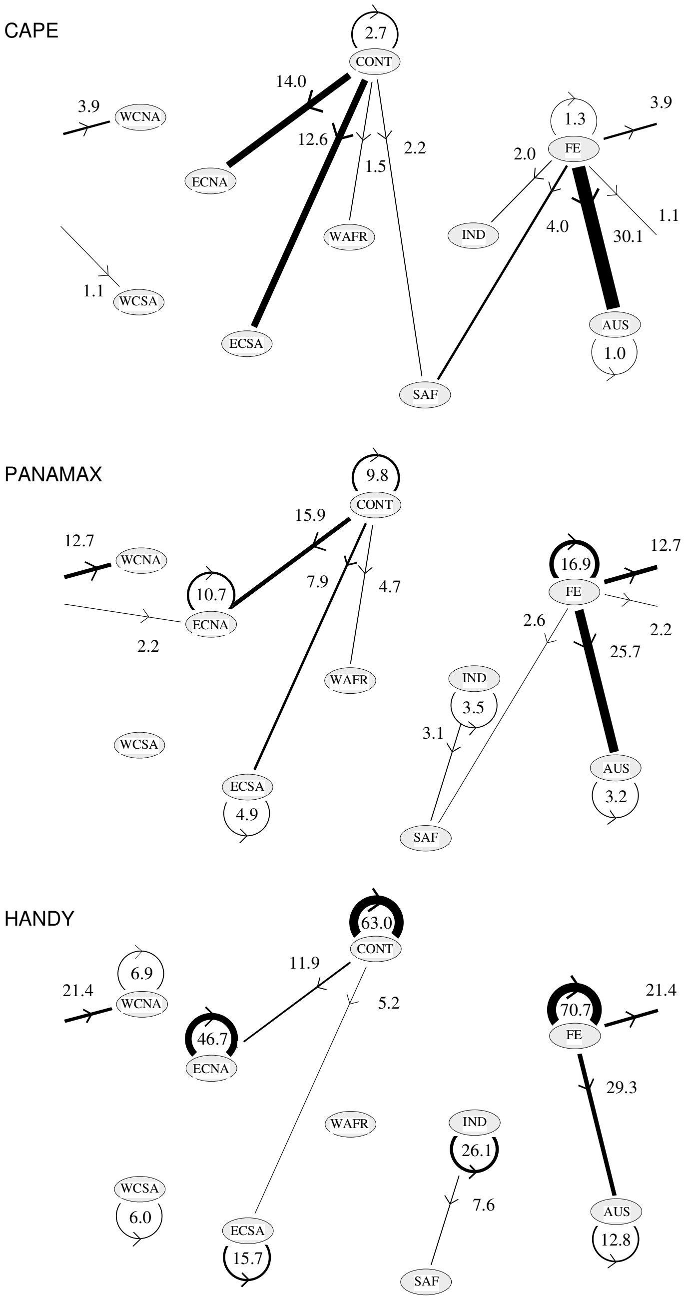
CAPE
2.7
CONT
14.0
3.9
WCNA
12.6
1.5
2.2
1.3
3.9
FE
2.0
ECNA
WAFR
IND
4.0
30.1
1.1
1.1
WCSA
ECSA
AUS
1.0
SAF
PANAMAX
9.8
CONT
15.9
12.7
WCNA
10.7
ECNA
7.9
4.7
16.9
12.7
FE
2.2
2.6
2.2
25.7
2.2
WAFR
IND
3.5
3.1
WCSA
ECSA
4.9
AUS
3.2
SAF
HANDY
63.0
CONT
11.9
6.9
21.4
WCNA
5.2
46.7
ECNA
70.7
21.4
FE
29.3
WAFR
IND
26.1
7.6
WCSA
6.0
ECSA
15.7
AUS
12.8
SAF

Capesize cargo flows are practically all ocean legs and end to the Far East or Continent. The Far East receives cargos from all macroregions except the Continent and West Africa, and the Continent from all macroregions except the Far East and the American West Coast. Australia dominates the Far East whereas the Continent divides its sourcing more evenly, particularly among the southern macroregions. Cargo flows originating from the Southern Hemisphere, US Gulf and Hampton Roads offer ample possibilities for transferring tonnage between the large operative theaters, Atlantic and Pacific. It is possible, for example, to take an occasional iron ore cargo from Narvik to MEG or the Far East; coal cargos from South Africa to Brazil or from Australia to Chile. The capesize ballast flows are simple return trips of the largest cargo routes although different in volume. They easily create an impression that most dry bulk traffic is simple shuttle.

That impression, as pointed out above, is incorrect and the actual complexity begins to unravel in the panamax segment. There cargos continue to converge upon the Continent and Far East, but they are directed toward India and ECNA, also. Australia and North America maintain their relative importance for the Far East and Continent, respectively. The Continent and Far East have opposite and balanced cargo flows. This is the key route for comparing the relative freight levels in the Atlantic and Pacific Spheres (Appendix 3.14). Cargo trips within macroregions (internal trips) become frequent. Only WCSA stays seemingly isolated. The panamax ballast flows are as in the cape sector, simple return legs of the largest cargo routes, although sometimes of widely different size.

The handy segment is now, as on many other occasions, different from the others. The higher threshold for displaying the flows contributes to that impression but only marginally. The Continent and Far East cease to be the dominant destinations. The number of opposing cargo flows increases. Internal cargo flows increase also. India, in particular, emerges out of thin air. It is partially coastal traffic and partially traffic between India proper and MEG. Construction materials are the mainstays of internal traffic in ECNA and Continent. To that come the feeder traffic from the Great Lakes grain ports to the St. Lawrence River, and grain, bauxite & alumina shipments in the Caribbean. Coal, iron ore and fertilizers are moved along the Indian coast and between India proper and MEG. In ECSA, shuttle traffic with iron ore and bauxite creates many cargo legs, completed by swaps of wheat against soya & oilcakes between Argentina and Brazil. In WCNA, transports by the forest and construction industries dominate. Australia includes New Zealand and the internal traffic comprises a wide range of products. Also Far East has a wide product range and the insular location of most national economies combined with short distances contributes. In transoceanic traffic, short routes are relatively important, for example ECNA-Continent in both directions, Vancouver to Far East, and Continent to India, particularly fertilizer shipments from the Black Sea. Although much of the handysize cargo traffic is short-distance there is no problem to take a 25,000 dwt vessel round the globe in as few legs as a 65,000 dwt panamax carrier. The cargos may be partially different but the itinerary is about the same. The handy ballast flows are as in the other sectors, tied to the largest cargo routes.

2.2.4 Workload

As a final step we take a deviation from our main theme, the interaction of cargo and ballast legs and look at the workload of the dry bulk industry. It is not a central topic in this report but because the scarcity of reliable figures is sometimes regretted in the literature and the figures are available here at a minimal effort, we use the opportunity. At the same time, the workload will give a tangible idea about the size of the market and the industry's efficiency, to be estimated by linear programming. We also have the opportunity to compare the effect of different regional meshes. The basic calculation is to multiply the leg frequencies in an O-D matrix with corresponding distances (Table 2.10). That is made by segment and separately for cargo and ballast legs. Comparability between segments is achieved by weighting with average dwt.

Table 2.10 Aggregate workload based on regions, 1997

Segment	Workload, 1000 nm			Dwt * Cargoload		Average distance, nm		
	Cargo	Ballast	Crg/Ball	avg dwt	bn nm	Cargo	Ballast	Crg/Ball
Handy	69,282	38,274	1.810	35,412	2,453	3,603	2,072	1.739
Panamax	37,400	21,870	1.710	66,837	2,500	5,205	3,264	1.595
CapeS	10,555	8,600	1.227	125,293	1,322	5,060	4,212	1.201
CapeL	13,301	8,562	1.553	177,592	2,362	6,697	4,335	1.545
CapeAll	23,856	17,162	1.390	150,875	3,599	5,859	4,272	1.371

Note: Cape cargoloads calculated from this table. Discrepancy in the cape segment originates from approximate dwts. Average dwts from Stripped files. Note their difference with Table 2.5.
Source: LMIU Movement Data.

The overall weighted cargoload is 8,552 bn nm (bn = 10^9), a very large figure, possibly in the same class as container ships of at least 25,000 dwt although less than corresponding tankers. The segments are not too different in spite that the leg totals relate roughly as 5:2:1 (Table 2.8). That gives perspective to a suggestion to extend this study down to the 10,000 dwt level. When the split of the cape segment into small and large capesizes is overlooked the details are logical: Cargoload relative to ballastload becomes smaller with increasing vessel size as does average cargodistance relative to ballastdistance. That tallies well with Figure 2.3 which suggestes that the probability to find a cargo or an empty vessel nearby is better when the cargo and vessel are small than when they are large, a good reason why ballast bonus is less common in the handy than panamax segment (Table 3.15). The difference in closeness is naturally connected with the well-known fact that large vessels have a competitive advantage over small ones on long routes with few port visits. These statements do not hold very well for the small capesizes, however, which have a remarkably short average cargo leg, shorter than that of panamaxes. The fact simply is that this size class is popular also in comparatively short trades such as Canada/Caribbean-Continent and Australia-Japan/Taiwan. The large capesizes, by contrast, focus on long trades such as Brazil-Far East and Australia-Continent.

These calculations are based on the geographical mesh of 37 regions. The mesh of the 10 macroregions is preferable, however, if results are not seriously affected. The optimization by linear programming will be simpler. An idea about the loss of accuracy when using the 10-region mesh, particularly in the Simulation, is also valuable. The workloads by segment, aggregated over regions, are not too different whether calculated from the finer or coarser mesh (Table 2.11). The largest difference 6.6 pct is in the ballastload of handysizes and fully acceptable. This encouraging result gets bruised when the segment totals are disaggregated by region (Table 2.12). Many differences exceed 10 pct, and the largest one is 61 pct. At closer look, 13 out of 16 large differences accumulate in ballastloads. That is natural because ballastdistances are shorter than cargodistances and modest absolute inaccuracies easily become large in relative terms. It is also so that 5 out of these 13 originate from regions with modest numbers of ballasting arrivals. Further investigation discloses that the main

Table 2.11 Aggregate workload based on regions and macroregions, 1997

Segment	Cargo, 1000 nm			Ballast, 1000 nm			Avg cargo dist, nm		
	Region	Macro	Reg/Mac	Region	Macro	Reg/Mac	Region	Macro	Reg/Mac
Handy	69,282	66,090	1,048	38,274	35,899	1.066	3,603	3,437	1,048
Panamax	37,400	36,724	1.018	21,870	21,360	1.024	5,205	5,111	1,018
CapeS	10,555	10,306	1.024	8,600	8,318	1.034	5,060	4,941	1,024
CapeL	13,301	13,562	0.981	8,562	8,921	0.960	6,697	6,829	0,981
CapeAll	23,856	23,868	0.999	17,162	17,239	0.996	5,859	5,861	1,000

Source: LMIU Movement Data.

Table 2.12 Workload disaggregated by regions and macroregions, 1997

	Handy			**Panamax**			**Cape**		
Cargo	Reg	M-reg	Pct	Reg	M-reg	Pct	Reg	M-reg	Pct
1	12451	11683	1.066	9695	9789	0.990	4376	3964	1.104
2	6730	6152	1.094	5131	5001	1.026	5285	5327	0.992
3	12547	12473	1.006	2668	2590	1.030	461	480	0.960
4	1461	1382	1.057	852	848	1.005	246	241	1.021
5	3758	3833	0.980	2143	2107	1.017	2600	2621	0.992
6	3612	3318	1.089	1228	1123	1.093	546	594	0.919
7	8332	6660	1.251	3345	2961	1.130	413	418	0.988
8	10156	9951	1.021	7049	6785	1.039	8310	8687	0.957
9	8218	8597	0.956	5106	5350	0.954	1089	1034	1.053
10	2017	2042	0.988	183	170	1.076	531	501	1.060
Total	69282	66091	1.048	37400	36724	1.018	23857	23867	1.000
Ballast	Reg	M-reg	Pct	Reg	M-reg	Pct	Reg	M-reg	Pct
1	3044	1955	1.557	625	503	1.243	37	23	1.609
2	971	688	1.411	298	284	1.049	40	37	1.081
3	7599	6604	1.151	7284	7043	1.034	7249	6917	1.048
4	567	490	1.157	13	10	1.300	0	0	
5	591	633	0.934	66	63	1.048	4	3	
6	5871	5590	1.050	1567	1479	1.059	388	352	1.102
7	16847	16433	1.025	11589	11357	1.020	9310	9759	0.954
8	875	1403	0.624	121	301	0.402	81	82	0.988
9	1101	1080	1.019	240	219	1.096	55	66	0.833
10	808	1023	0.790	67	101	0.663	0	0	
Total	38274	35899	1.066	21870	21360	1.024	17164	17239	0.996

Source: LMIU Movement Data.

source of inaccuracy are legs within macroregions, i.e. the shortest possible legs. This is seen by calculating weighted internal distances (legs as weights) and comparing these with macroregion distance tables. One can generalize fairly confidently that when the workload based on regions exceeds the workload based on macroregions the weighted distance exceeds the entry in the distance table, and the other way round. This is rather obvious. Whether measures should be taken is less obvious.

The absolute differences are usually less than 1,000 nm and often much less. If measures are considered necessary a correction of all internal distances which are the source of most large differences becomes necessary. The descreptive benefits would hardly be in proper relation to the labor input. The sign and approximate size of the bias are known and that should be sufficient. What comes to the Simulation, experience from former runs (LHS 2001) indicates that vessels are unlikely to remain consistently on one particular route. This means that whatever bias there may be it is unlikely to accumulate in space.

The geographical perspective is given by cargo- and ballastloads (legs * distance) within and between the three oceans, Atlantic, Pacific and Indian (including Western Australia). Segment comparability can be enhanced by substituting average vessel size for leg. And a monetary market potential can be created by substituting time for distance and daily freight rate for vessel size. That measure would apply only to cargo flows because no freight revenue is allocated to ballast legs in this study.

Atlantic and Pacific are the main operational theaters and Atlantic is the larger of the two (Fig. 2.7). That was back in 1997. Today, China's economic growth may have turned the scales in Pacific's favor. Indian Ocean cannot compare with these two. Its area is much smaller and the countries bordering it are less developed economically. Underdevelopment becomes manifest in very modest internal flows, external flows dominating in each of the six sub"markets", three size segments split into cargo and ballast. The Atlantic and Pacific Spheres have much interaction in cargo, compared to their internal flows. The interaction is not balanced, however, and the flows from the Atlantic are in each size segment roughly twice the opposing flows. There it is ample reason to speak about front- and backhauls. This characterization has only one flaw, it does not differentiate between alternative routes. There are hardly any backhauls through the Panama Canal but it does not necessarily mean that the scales are balanced at the Suez Canal. The explanation is routing via the Cape of Good Hope. Ballasting between the Atlantic and Pacific is minimal which suggests that the allocation of empty vessels is quite efficient. This conclusion does not imply that the use of loaded vessels is inefficient. They only face more binding constraints in relation to cargo origin and destination than empty vessels. For a chartered vessel the philosophy that any cargo or any loading port is equally suitable does not hold. And for a loaded vessel the same negation applies to any discharging port. The cargo constraint also includes the time window in which the delivery must take place. It is a very severe constraint.

The joint effect of time window, cargo quality, and location can be estimated by comparing the actual workload to the workload unconstrained by these factors. This minimum is the solution of the classical transportation problem. Here, the distance variables are physical distances but freight revenues (rate * days) would be an alternative. But only a technical alternative, because a pronounced shift of transport flows would most likely lead to changes in freight rates. We work with physical distances and optimize by segment, separately for cargo and ballast legs (Table 2.13). The both optima are practically identical, suggesting that plain shuttle service would give a very efficient result. In technical jargon, the solution surface around the global optimum tapers off slowly and the existence of multiple optima is not far fetched. The reason, of course, is the pattern of between-region distances. The actual flows, by contrast, need not be too similar (Appendix 2.7). In the handy ballast solution, for example, about 15 pct of ballast legs are not mirror images of the cargo solution.

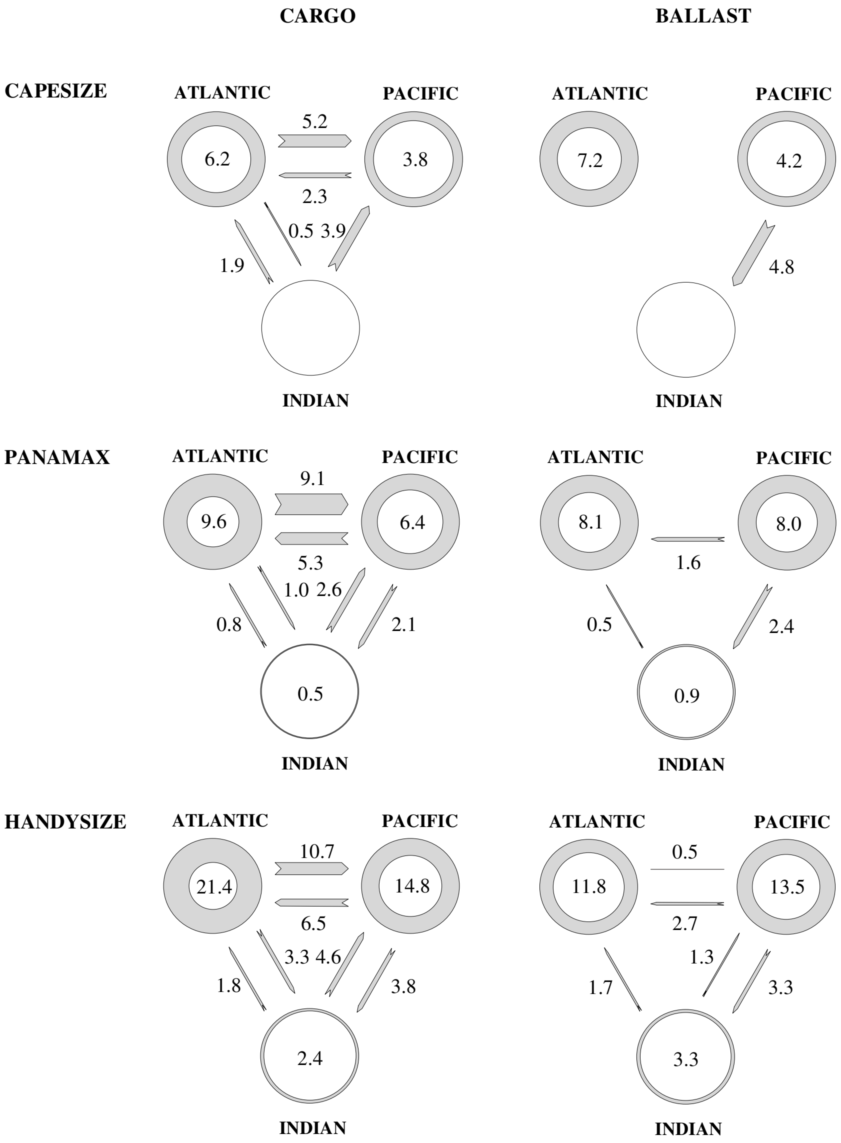

Figure 2.7 Workload within and between the Atlantic, Pacific and Indian Oceans, 1997

Notes: Atlantic comprises regions 111-128, Indian Ocean regions 301-305 and 415-416, and Pacific regions 411-414 and 417-425. Region 201 is allocated to Atlantic or Indian Ocean depending whether the flow is west- or eastbound. Its internal flow is allocated to the Indian Ocean. Minimum flow displayed is 0.5 mill. legmiles (nautical). Handysize has different scale than capesize and panamax.

Source: LMIU Movement Data.

Table 2.13 Observed and optimized workload based on macroregions, 1997

Segment	Cargo, 1000 nm			Ballast, 1000 nm			Avg cargo dist, nm	
	Observ	Optim	Obs/Opt	Observ	Optim	Obs/Opt	Observ	Optim
Handy	66,090	29,984	2.204	35,900	29,777	1.206	3,437	1,559
Panamax	36,724	20,676	1.776	21,360	20,248	1.055	5,111	2,878
CapeS	10,305	8,087	1.274	8,320	8,075	1.030	4,940	3,877
CapeL	13,562	8,762	1.548	8,920	8,788	1.015	6,829	4,412
CapeAll	23,867	16,849	1.417	17,240	16,863	1.022	5,861	4,138

Notes: Optimization manually (Dorfman et al. 1958, 106-117). Iteration discontinued when change in the object function minimal. Obs/Opt for average cargo and ballast distances identical with workloads.
Source: LMIU Movement Data.

The observed ballastload is quite close to the optimum in the panamax and cape segments but not in the handy segment. The closeness means that ballast movements are mostly made within the same or between neighboring macroregions (Fig. 2.1). When this is not case there is reason to suspect factors such as bias originating from a large mesh, heterogeneous fleet, tied traffic, and continental imbalances in the supply and demand of empty tonnage. None of these reasons helps in explaining the large ballastload surplus in the handy segment. Without bias the surplus would be still larger (Table 2.12). Fleet heterogeneity has been handled. Continental imbalances in ballasting are about the same as in the other segments (Fig. 2.7). Whether they are real or the result of inaccurate itinerary construction is formally irrelevant (Appendix 2.4). They are all what is available.

Table 2.14 Apparent effect of time, cargo and location constraints

Segment	Cargo	Ballast		Differ	Diff/Surp
	Obs/Opt	Obs/Opt	Surplus	Crg-Ball	
Handy	2.204	1.206	0.206	0.998	5
Panamax	1.776	1.055	0.055	0.721	13
CapeS	1.274	1.030	0.030	0.244	8
CapeL	1.548	1.015	0.015	0.533	36
CapeAll	1.417	1.022	0.022	0.395	18

Legend: Cargo = time & cargo & location; Ballast = time; Difference = cargo & location.
Source: Table 2.13.

The observed cargoloads are, by contrast, roughly 40-120 pct larger than the corresponding optima. They are the numerical expression of a joint time & cargo & location constraint. If the surpluses observed in the ballastloads are assumed to reflect only the time constraint, then the difference between the cargoload and ballastload surpluses equals the joint cargo & location constraint (Table 2.14). Not entirely however, because there are good commercial reasons, such as risk diversification and price differentiation, to source the same commodity from various quarters. Therefore, the effect of the said constraints must be labeled as "apparent". With this reservation, the joint cargo & location constraint dominates in all segments and the relative dominance grows with increasing vessel size.

If desired, the analysis can be sharpened by disaggregating the year into 26 biweekly periods and conducting corresponding optimizations for cargoloads and ballastloads, to be subsequently aggerated back to annual figures. The shorter time period would largely eliminate the time window effect and leave only cargo and location constraints intact.

3 FREIGHT RATES

3.1 Fixtures

3.1.1 What they are

Fixtures are compact copies of charterparties, released by ship brokers with the consent of the contracting parties to the public domain. Their share out of the original total has been guesstimated as 30 pct, a reasonably accurate figure (Table 3.5). The figure is higher in large than small size segments and at the 10,000 dwt mark there are not many fixtures left. The undisclosed share can be used internally, e.g., when a brokerage and a consultancy are parts of the same company. No broker covers the total freight market and it is obvious that there is some geographical bias. Brokers in London, Athens and Oslo are better informed about Atlantic matters while their peers in Tokyo, Hongkong and Singapore cover better the Pacific and Indian Ocean. Major brokers have a considerable overlap, however, and there is no reason to assume that any of them would not cover the globe adequately. They release the public share either directly or through data vendors whose main business is normally consulting. An academic has access to that part of the data only. Our data originates from *Drewry Shipping Consultants*.

The fixture data is not immediately usable. Items of different information appear in the same field (column), labeled as ”comments” or ”narrative”. For efficient use they must be separated into fields of their own and this takes time. The same applies to ranges, minimum and maximum values need to be placed into separate fields. A particular sign can have different interpretation in different files and reports. There are typos, more frequent in printed than electronic files. Most are obvious and need only correction, tenfold loading rate, unchanged year from December to January and suchlike. All this information is usable. But there are also numerous fixtures which must be rejected wholesale. The process was long and rather disorderly because all cases did not surface during a formal search but emerged gradually although at a diminishing rate (some detail in Appendix 3.1). In the meanwhile, the seemingly ready files had been used for orientative tabulations. Because files became rapidly oversized for available computing facilities they were pruned of seemingly superfluous information. Some of these cleanups were premature and return had occasionally to be made to older files. The result is that frequencies can vary slightly between tables, the final authoritative count being found in the context of freight rates (Subchapter 3.2). Care has been taken that the discrepancies do not affect conclusions drawn. Overall, 5,498 rejections were made, 11 pct of the grand total available in the data source. A shrinkage of this magnitude must be observed when planning future studies.

Four types of fixture are relevant for this study: voyage, Contract of Affreightment (COA), trip and time. Voyage and trip charters are mostly for single, COAs and time charters for multiple transportation tasks. COA is simply a repeated voyage charter during a time period. Time charter is a lengthy and geographically vague trip charter. Period charter is the joint term for trip and time. There is some overlap timewise. The longest trip charter took 96 days to complete and the shortest time charters in the data were for one month. The longest time charter was for 6 years, and the longest COA covered a period of 4 years. All types of charter offered options for extending the charter period. A geographically distinct option was considered a separate fixture.

Time charters are made for a number of months/years, usually with explicit delivery areas. They do not specify ship movements and are consequently excluded here except in some overview tables (e.g., Tables 1.6 and 3.1). The specific needs of the charterer may, but need not, differ from the market at large. Certain is that time-chartered vessels compete for the same cargos as vessels sailing under other charters. A similar competitive situation exists between voyage charters and COAs. Because their geographical detail is the same they are consolidated under the voyage charter logo.

Table 3.1 Usable fixtures, 1993-2002

Charter	Handy	Panamax	CapeS	CapeL	First count	Final count	Original
Voyage	5,032	7,230	5,681	2,460	20,403	20,368	23,082
Trip	5,568	12,018	1,297	1,099	19,982	19,963	22,531
Time	694	2,968	470	410	4,542	4,542	4,758
Total	11,294	22,216	7,448	3,969	44,927	44,873	50,371

Note: Original as a memory item includes fixtures excluded from the analysis.
Source: Drewry Monthly 1993-2002, yellow pages.

The essentials of voyage and trip charters are best given by two uncomplicated examples (Drewry Monthly, July 2002, 32, 34):

Voyage charter
Richards Bay to Rotterdam (coal)

Reported	Vessel	Cargo	Charterer	Terms	Laydays	Rate
12/06/02	Bulk Asia	160,000	Swiss Marine	Scale/25000 Shinc FIO	15-20.07.02	$5.55

Trip charter
Cape Passero/ECSA/Skaw-Passero range

Reported	Vessel	Built	Dwt	Speed/Cons.	Charterer	Delivery period	Rate
21/06/02	Fu Hua	1997	72,437	14.5/-	Cargill	25-30.06.02	$6250

For our purposes, the main differences are two. In voyage charter the freight rate is per cargo tonne, in trip charter it is per day. To the trip rate is in many trades with an endemic tonnage shortage added a fixed amount (ballast bonus), up to one-half of the total freight revenue. In voyage charter the ship pays all the costs connected with the transport, in trip charter it escapes the fuel (bunkers), port charges, possibly loading and discharging costs (depends on charter) and canal fees which are paid by the charterer. It follows that the total freight revenue of a given transport task is lower when trip charter is used. But when a voyage charter is converted to a trip charter equivalent the freight revenue is the same, or almost. The choice between voyage and trip, therefore, is not primarily monetary. Uncertainty about the exact loading and discharging ports at the time of concluding the deal, the risk of unexpected bottlenecks in ports and narrows, and the administrative capacity to handle the paperwork called for by the voyage charter also play a role.

In the examples above, the voyage charter specifies the ports. In some others it might give options such as Antwerp or Rotterdam, or several ports. But it would still be quite specific. That specificity is only apparent, however, because only the cargo leg is disclosed. The distance which the vessel must cover to reach the loading port influences the rate but remains concealed. The trip charter, by contrast, is openly vague. It tells that delivery takes place when the ship passes Cape Passero and that redelivery will be somewhere between Cape Passero and Skaw. The preceding discharging port must have been in the eastern Mediterranean or Black Sea. But the current discharging port can be anywhere from the Baltic to the Black Sea. The loading port/s appear to be on the eastern coast of South America. It is a long coastline and a clue about the cargo might help focusing on a part of it. That clue comes from the charterer, a major trader of grain and agricultural products. The size of the vessel almost certainly excludes sugar cargos, i.e. northeastern Brazil. Remain wheat, corn (maize), soya, meals and oil cakes. Their shipping ports locate mostly between Rio de Janeiro and Bahia Blanca. It is still a long stretch of coast although not longer than the Skaw-Passero range. An analyst with a background in broking enjoys here a competitive advantage. Geographical vagueness is one reason why regions rather than individual ports are the geographical roster of this study. Still, the definition above is very research friendly. Redelivery could also be in the Atlantic which would not exclude North America, and the operational area might be worldwide. Such fixtures are not usable in this study.

Laydays and delivery period define the time period during which the vessel must be at the charterer’s disposal. The charter must be concluded before loading can begin or delivery take place. If reported to the market the delay is normally a week or so and laydays have probably not yet started to accumulate. But some 5 pct of all fixtures are reported later than that, even much later (retros). It is then up to the analyst to decide which date to use. The reporting date appears to be the standard choice. Here it is modified when appropriate. The workload of this revamp was considerable and it is doubtful whether the outcome justifies it.

The allocation of trip fixtures to the appropriate ship segment is trivial because the vessels’s dwt is disclosed. In voyage fixtures dwt is replaced by the weight of the cargo, abbreviated dwct. The use of a multiplier, related to the cargo’s volume weight, is then necessary to arrive at an approximate dwt (Appendix 3.2). The inaccuracy can become serious when a vessel takes a part cargo. This becomes dramatically to day when the one and same vessel is discovered in different size segments. There is also the minor although systematic inaccuracy that cargo does not include the weight of the crew, supplies and bunkers which together make 5-10 pct of a full cargo.

Voyage and trip fixtures disclose several important operational parameters. Their discussion will be postponed for a while.

Trip fixtures offer a better analytical foundation than voyage fixtures because they abstract from most of the geographically varying cost items, port & cargo handling charges and canal fees in particular. Against this perspective, the geographical vagueness of trip fixtures is less important. Voyage and trip fixtures split approximately 50/50 overall although certain cargos, ship sizes and routes tilt for the one or the other (Appendix 3.3). For example, grain shipments from the US Gulf are usually made under the voyage charter and it is the preferred alternative in published iron ore fixtures of large capesizes, too. By contrast, panamax shipments from Australia and WCNA to Far East are mostly under trip charter. Most major routes are well supplied by both, however, whereas minor ones are not. If, in addition, the period 1993-2002 is split into subperiods the shortage of fixtures on some routes becomes acute. Their relative shortage from/to Region 10 (WCSA) catalyzed its consolidation with Region 9 (WCNA) in all freight rate calculations (Appendix 3.4). With hindsight, the decision was too heavy-handed but, once implemented, very difficult to reverse.

3.1.2 Coverage

Shipping research rests on the often unspoken assumption that published fixtures are a fair representation of the freight rates paid for actual cargo movements. Such an assumption contains some important unknowns.

- The size of owned fleets of non-shipping (industrial) companies is seldom mentioned and because ownership is not always overt the question may never get a reliable answer. Logically, when fixtures are converted to freight rates and rates are combined with movements into a system, part of the system will rest on a shaky ground. Does it matter? Nobody knows for certain. It can be claimed that the going rate level reflects the opportunity cost of industrial companies. But it cannot be proved that these companies actually observe the opportunity cost in their everyday decision making.
- Published fixtures are not a random sample of the total charter population. They are a judgemental sample and that can occasionally mean that the information is seriously biased. For example, in the Drewry data there are 99 scrap fixtures in 1996 but only 5 in 1999 and 2000; or 269 sugar fixtures in 2002 but only 25 in 1997.
- Estimates about the published fraction apply to the total market and are seldom detailed by route. One also gets an impression that they are more hunches than results of careful surveys.
- Many charters are purely speculative, concluded only to be sold at a higher price later on (relets), and never intended to lead to the delivery of the vessel. This author has seen fixtures of the same vessel on the same route at few days intervals where the new rate is $500 or 5 pct higher than the old one. One seasoned broker puts the share of relets at 80 pct of the total, roughly in line with the practice at commodity exchanges.

Irrespective of these reservations it is generally thought that published fixtures give a fair idea of the market at large and because this opinion originates from people who know also the unpublished market, their word carries weight. In any case, the published fixtures are the only available data source for this study about freight rates and some operational parameters. They will be related to the number of cargo legs below.

The next query is whether the fixtures from Drewry are an unbiased sample of fixtures generally available for the international broker community. Drewry is an independent consultancy without an attached brokerage. The formulation emphasizes ocean shipping. Coastal (short-sea) shipping is mostly outside our interest. Three simple tests are given.

First test The Baltic Exchange (Baltic) published on 18 August 2000 18 fixtures, 8 for voyage and 10 for trip charters (Collins 2000, Fig. 2.1). Drewry's figures were 11 voyages and 12 trips. All voyages were the same as at Baltic whereas only 7 trips matched. The trip fixtures originating from Baltic, but not Drewry, quote vessels which appear regularly among published fixtures. Two nameless ('’vessel'’) Drewry voyage fixtures can be identified with Baltic with the help of other characteristics. Even in the geographically constrained London market a fair percentage of fixtures appear to escape the net of a major player.

Second test The original study was based on data from three vendors: Drewry, Lloyd's and Maritime Research (New York). The databases were consolidated although their compilation principles differed somewhat (Tables 3.2 and 3.3). Drewry appeared in 79 pct of voyage fixtures and in 88 pct of trip fixtures. There is no reason to think (as some interviewed persons speculated) that the London market and the New York market would be markedly independent.

Third test Continent and Far East are the largest destination areas for dry bulk shipping. If Drewry's data is badly biased towards Continent that should be visible in the number of fixtures (Table 3.4). The possibility cannot be outruled. The superiority of Far East should be more pronounced, particularly in the cape segment. But Far East is also the macroregion where tied traffic is believed to be more common than elsewhere. We investigate the controversy by comparing the fixture data with Final cargo legs in 1997 (Appendix 3.4). These legs approximate the polypolistic (spot ?) market and correspond best with voyage and trip fixtures. The comparison (pct, column totals) shows that there is no systematic difference between Continent and Far East.

Table 3.2 Compatibility of voyage fixtures, 1997

	Vendor combination							
	D	L	M	DL	DM	LM	DLM	Total
Lloyd's register	13	53	9	11	112	44	335	577
No register	75	26	115	1	362	4	37	620
Total	*88*	*79*	*124*	*12*	*474*	*48*	*372*	*1197*

Source: LHS 2001, Table 4.1.

Table 3.3 Compatibility of trip fixtures, 1997

	Vendor combination							
	D	L	M	DL	DM	LM	DLM	Total
Direct	12	35	17	7	136	24	392	623
Indirect	14	39	20	2	207	33	489	804
Total	*26*	*74*	*37*	*9*	*343*	*57*	*881*	*1427*

Source: LHS 2001, Table 4.2.

Table 3.4 Fixtures discharging in Continent and Far East, 1993-2002

	Handy	Panamax	CapeS	CapeL	Total
Continent	3,859	6,299	4,219	1,887	16,264
Far East	3,347	10,983	2,396	1,622	18,348
Total	7,206	17,282	6,615	3,509	34,612

Source: Appendix 3.3.

The comparison is extended to all routes with decent frequencies, and it comprises both Final and Stripped legs (Appendix 3.4). Stripped legs are given as a memory item because they are the data which are more likely to be quoted in unsophisticated research. The geographical mesh is macroregions. The comparison has two rationales. First, to get an idea about the share of cargo legs for which the freight rate is known. Is the share large enough for the rates to be considered reliable? Second, to investigate the usability of fixtures as substitutes of cargo legs. If they are found to be fair substitutes much effort can be saved. Fixtures are an integral part of any comprehensive shipping study, they are more easily derived than cargo legs, and every fixture represents basically a physical cargo leg. The matrices can be studied from two angles: row, column and grand totals, and cell patterns, i.e., the mutual locations and frequencies of non-zero cells. The focus here is on the grand totals from which the most reliable shares can be derived. They are given for all years and size segments (Table 3.5), full matrices only for the year 1997.

Table 3.5 Fixtures and cargo legs compared overall

Segment	Year	Fixtures	Final files		Stripped files		Avg percent		
			Legs	Fixt/Legs	Legs	Fixt/Legs	Fin	Str	Diff
CapeL	1995	191	773	24.7	1,393	13.7			
	1997	326	1,102	29.6	1,986	16.4			
	1998	339	1,063	31.9	2,187	15.5	29	15	14
CapeS	1995	966	1,583	61.0	2,209	43.7			
	1997	840	1,457	57.7	2,086	40.3			
	1998	751	1,217	61.7	1,811	41.5	60	42	18
CapeAll	1995	1,157	2,356	49.1	3,602	32.1			
	1997	1,166	2,559	45.6	4,072	28.6			
	1998	1,090	2,280	47.8	3,798	28.7	48	30	18
Panamax	1995	1,988	5,539	35.9	6,240	31.9			
	1997	2,221	6,289	35.3	7,185	30.9			
	1998	2,387	6,149	38.9	7,168	33.3	37	32	5
HandyL	1997	219	4,793	4.6	5,522	4.0			
HandyM	1997	536	6,517	8.2	7,516	7.1			
HandyS	1997	311	5,585	5.6	6,189	5.0			
HandyAll	1997	1,066	16,895	6.3	19,227	5.5	6	5	1
Total	1997	4,453	25,743	17.3	30,457	14.6	17	15	2

Sources: Drewry Fixture Data; LMIU Movement Data.

A surprisingly large share of the cargo legs are covered by fixtures. Using the Final files as the reduction base the rounded segment averages are 30, 60, 50, 35 and 5 pct. And using the Stripped files, they are 15, 40, 30, 30 and 5 pct. Apparently a large percentage of vessels release at least some deals to the public domain. An estimate can be based on the voyage and trip fixtures in 1997 (Table 3.6). From the about 4,500 fixtures some 3,500 can be identified by vessel name, the rest are mostly voyage fixtures denoted only as “vessel/s” or “tonnage”. Ortographical typos contribute to non-identification. About one half of vessels with recorded movements released at least one fixture. That does not mean that the vessels have been compared one by one (cf., Table 3.14), only that the totals have been compared. The identified vessels released on average 2.0 fixtures. That is not bad when an average vessel made 7-9 trips (Stripped). The both tables support each other and the estimates are apparently in correct size class. However, that is only a formal result and the reason is the speculative element of the market.

Table 3.6 Fixtures compared with vessels, 1997

Segment	Vessels Movem	Released	R/M	Fixtures All	Identif	I/A	Fixt/Vess
CapeL	294	138	0.47	326	294	0.90	2.13
CapeS	291	289	0.99	840	699	0.83	2.42
Panamax	983	761	0.77	2221	1837	0.83	2.41
Handy	2108	589	0.28	1066	746	0.70	1.27
Total	3676	1777	0.48	4453	3576	0.80	2.01

Sources: LMIU Movement Data; Drewry Fixture Data; Tables 3.5 and 4.4.

If the estimate of 80 pct for fixtures which do not lead to a physical transport is correct the percentages of Table 3.5 must be reduced to a fifth. The margin of error is thus very wide and the only way to reduce it is to solicite the broker community's help. Such an effort is out of question in this study and its practical usefulness can also be questioned. As pointed out above, consequtive fixtures for the same vessel on the same route used to differ by about 5 pct ($500). Such a difference is within the measurement error of this study. It is immaterial then whether the released fixture leads to a physical transport or not. It is a fair estimate of the market at that particular time and usable as such.

The evaluation of individual macroregions and routes is less straightforward because the frequencies and percentages vary within wide limits. An attempt will be made, nevertheless. Two groups of cells are identified, those where the coverage (percent) is based on a sound number of fixtures and cargo legs, and those where a sound number of cargo legs generates too few fixtures. What is "sound" depends on the context, i.e., the grand total of cargo legs and their overall coverage in the particular segment. The topical cells in the "Pct" and "Final" matrices in Appendix 3.4 are given in fat style. The identification is apparently impressionistic but sufficient for our purpose. The first group gives an idea of the "acceptable" foundation of the study, the share of cargo legs with a rate idea better than a guess. The second group are the "doubtful" routes with an acute shortage of fixtures. Why this is so, is open for speculation. The traffic may be broked locally and be uninteresting for the global broker community. The traffic may be mostly tied, although there is no way how that share can be identified reliably in this study. Worst of all, the codification of the legs may be incorrect. This possibility cannot be discounted because the share of the doubtful routes grows with declining vessel size (Table 3.7). Finally, the calculations apply to a single year only and ignore traffic under time charter, unless tied, almost 40 pct of the market (Table 1.6).

Table 3.7 Soundness of fixture coverage (Final files), 1997

	Percent of legs Acceptable	In-between	Doubtful	Total legs
CapeL	87.2	12.8	0.0	1,102
CapeS	89.2	10.4	0.4	1,457
Panamax	85.7	5.2	9.1	6,289
Handy	76.1	2.2	21.7	16,895

Source: Appendix 3.4.

The idea that fixture matrices can be a substitute for cargo leg matrices has some merit. But considering the wide variation of coverage, even in those cells which belong to the acceptable group, this is an emergency solution, nothing to be recommended (Appendix 1.1). If it is resorted to, each cell must be treated by an appropriate multiplier, derived from some reliable and up-to-date source.

3.1.3 Source of parameters

Fixtures are a ready source of several, operationally and analytically important parameters: ship size, loading and discharging rates, speed, age, fuel consumption, and number of ports (Table 3.8).

Table 3.8 Occurrence of important parameters in fixtures, percent

	Handy	Panamax	Cape	First count
Voyage fixture				
Load/Disch separated	83.4	53.4	38.3	54.8
Load/Disch lumpsum	11.5	36.8	27.6	26.9
Trip fixture				
Speed	71.9	78.6	60.0	75.3
Age	97.2	99.5	99.7	98.9
Fuel	60.1	67.9	44.1	63.5

Source: Table 3.1.

Size is measured in deadweight tons (dwt) or deadweight cargotonnes (dwct). The former gives ship's total carrying capacity and is used in trip charters, the latter gives its cargo carrying capacity and is used in voyage charters. They are practically always available, even when the ship is not identified ("vessel") because the desired size is a key parameter of the deal. The size measure primarily used in this study is dwt because trip charters use it. Dwct can be converted to dwt with the help of stowage factors (Appendix 3.2). Differences in average dwt (converted or not) by segment and loading region between trip and voyage charter are small (Table 3.9). That can be used as indirect evidence about the frictionless mobility between the two. Ship characteristics derived from the one type of fixture will consequently apply to the whole fixture population. The fixtures are sufficiently numerous in all loading regions, Region 4 possibly excepted, that the dwt-figures can be considered reliable. The Atlantic loading regions tend to have smaller ships than Pacific regions in the handy and small cape segments, but the difference is not large.

Port time is given either as a lumpsum or separately for **loading and discharging** (jointly: handling). In the latter case it may be necessary to derive it from loading and unloading rates. This inconvenience notwithstanding separate rates are preferable because they are needed in the Simulator. About one half of the fixtures offer this possibility (Table 3.8). Lumpsum is available in an additional quarter of the fixture population. It can be used only for estimating freight rates. Both estimates are likely to be too low because there must be a berth available, the ship must be fastened to it, and paperwork be cleared before cargo handling can begin. The handling time is a function of port standard, ship size, cargo, and whether the activity is loading or discharging. Technical efficiency, good labor relations and large ships obviously speed operations. Given identical port standard loading is effected up to twice as fast as unloading, but because the standard varies this rule-of-thumb is not very helpful here (Table 3.10). The number of working hours influences total time and is tied to port practice and the particular charterparty used. It is one aspect of laytime, the time during which the vessel should be handled (Collins 2000, 203-208, 213-216; Coulson 1991, 65-66). Grain is often shipped Shex (Sundays and holidays excluded) whereas coal and iron ore go Shinc, which means about 15 pct more working hours. When loading and discharging are given separately the unit may not be days but a rate, tonnes handled per day. In the fixture example above, 25,000 tonnes per day suggests that *Bulk Asia* will be unloaded in 6.4 days. The size of cargo must then be divided by this rate to arrive at the time estimate. Hopefully the rate is in tonnes because terms such as Berth Terms, Custom of Port (COP), Customary Quick Despatch (CQD), or Scale make the information unusable here. Handbooks are available for consultation but their figures may well be too optimistic.

There are also unused sources: laytime for the handling time, given in voyage fixtures, and total port time, available as the difference of ship arrival and departure dates in LMIU Movement Data. The former went unnoticed by the author and is quite complicated anyway; the reliability of the latter can often be questioned (Subchapter 2.1).

Table 3.9 Average dwt (1,000) by loading region, 1993-2002

Loading reg	Handy			Panamax			CapeS			CapeL		
	Voy	Trip	All	Voy	Trip	All	Voy	Trip	All	Voy	Trip	All
1	34	38	35	69	68	69	119	124	120	161	163	162
2	34	37	35	67	67	67	125	132	127	167	168	167
3	34	37	36	67	67	67	118	123	119	161	163	162
4	46	36	42	60	67	63	116	129	116	155	165	158
5	34	36	36	65	68	67	130	133	131	166	166	166
6	31	38	34	61	68	67	120	129	123	162	160	161
7	38	39	39	64	69	68	135	131	134	165	166	166
8	40	38	39	65	69	68	132	137	133	164	166	165
9	35	37	36	67	67	67	128	134	130	164	167	166
All	35	37	36	67	68	68	125	131	126	165	166	166

Note: Region 10 consolidated with Region 9 due to small frequencies.
Source: Author.

Speed (knots or nm/h) is indispensable for converting distances into sea days. The figure given for *Fu Hua* (above) refers to sailing in fair weather in open sea, tantamount to almost 350 nm in 24 hours. But a deduction, perhaps 1 knot, must be made for foul weather, narrows and port approaches.

Ship age and fuel consumption do not have the role of operational variables in this study but are potentially useful in analyzing the results. Old ships have lower capital costs but higher operation costs. They face less demand and lower freight rates. At a given rate, new vessels are chartered first. When the market is weak old vessels may be unemployable (Tamvakis and Thanopoulou 2000). Age, then, may function as a proxy for general neglect better captured by frequent changes of owner, name, flag and classification society, all beyond our possibilities of measurement (Timmerman and McConville 1996). Logically, route-specific rate ideas by consultants normally specify ship size and age. Our data originates from trip fixtures (Table 3.11). Small capesizes and handysizes are the old ships, large capesizes the new ones. There has apparently been a shift from handysizes to panamaxes and from small capesizes to large capesizes. The observation has significance for the functional relationship we are interested in. By contrast, the difference between the Atlantic and Pacific loading regions remains subdued.

Table 3.10 Loading and discharging times by region, 1993-2002

Region	Loading			Discharging		
	Handy	Pana	Cape	Handy	Pana	Cape
111	4.8	4.4	3.7	7.5	3.7	6.1
112	5.0	3.2	3.4	10.0	5.6	4.8
113	5.0	5.0	5.0	8.6	4.6	3.2
114	2.8	2.9	3.9	11.4	10.0	10.0
115	10.9	2.3	3.1	13.5	5.0	5.0
116	7.6	6.9	3.4	12.6	5.9	5.3
117	5.9	8.4	8.9	9.8	10.0	10.0
121	5.0	5.1	5.0	10.0	5.0	5.0
122	6.2	7.1	7.8	13.5	8.6	4.2
123	4.7	7.6	4.3	8.1	3.8	4.7
124	4.4	5.6	10.7	10.5	5.5	5.0
125	7.4	6.8	7.0	8.6	7.6	5.0
126	6.4	9.7	10.0	15.9	5.8	6.5
127	1.9	2.3	3.4	13.8	20.0	20.0
128	6.7	10.0	10.0	13.3	20.0	20.0
201	5.2	1.8	2.9	8.4	9.2	7.8
301	3.7	5.0	5.0	13.6	16.6	25.1
302	5.9	3.8	5.0	13.2	14.6	12.0
303	8.7	2.4	5.9	17.7	15.7	10.4
304	5.7	15.0	2.0	12.6	12.1	19.4
305	18.8	15.0	10.0	16.4	25.0	25.0
411	5.2	4.4	5.0	11.9	9.1	4.6
412	8.7	8.8	5.6	11.9	13.4	7.3
413	18.0	4.7	5.0	14.7	9.6	4.7
414	7.0	3.8	5.2	12.0	11.0	8.4
415	4.0	2.4	2.4	8.0	5.0	8.0
416	4.9	4.7	5.0	8.0	8.0	8.0
417	2.4	2.3	3.3	8.0	2.0	8.0
418	3.2	2.7	3.5	4.3	6.0	5.4
419	5.4	5.0	5.0	8.0	8.0	8.0
420	9.4	5.0	5.0	8.0	8.0	8.0
421	5.0	5.0	4.3	8.9	8.0	8.0
422	5.7	4.2	6.4	10.0	9.0	9.0
423	11.0	5.0	5.0	10.0	7.5	5.2
424	17.5	3.0	4.4	9.1	6.2	43.3
425	2.9	3.7	3.4	5.7	5.0	9.0
Total	5.6	4.7	3.7	11.2	8.4	5.6

Source: Author.

Fuel consumption obviously influences the freight rate but because lower consumption is achieved by more advanced technology the higher capital cost will partially offset this advantage. Therefore, fuel consumption plays a role primarily in the comparison of voyage and trip rates because its price has varied during the study period within a range of $85-165/tonne IFO in Rotterdam. For these reasons the display of fuel consumption is meaningless. It also has the tendency to destroy the significance of both dwt and age variables in a regression analysis explaining daily freight rate or similar. When the estimation, nevertheless, succeeds the coefficient is 0.40-0.45.

Table 3.11 Average age by loading region, 1993-2002

Loading reg	Handy	Pana	CapeS	CapeL
1	14	10	15	5
2	16	12	14	6
3	13	11	16	5
4	15	12	15	3
5	15	9	15	6
6	15	11	17	11
7	12	9	14	7
8	11	8	11	6
9	12	9	12	6
All	13	10	14	6

Source: Author.

By contrast, an analysis with dwt and age as independent variables is meaningful because they will be used in explaining the Daily Revenue derived from the Simulator (Subchapter 4.6). The explanation is not too successful there, because of multicollinearity, but succeeds here. The dependent variable is C$/Day (below) and the hypothesis is that an increase of ship size will raise the rate whereas an increase of age will lower it. There is also a third independent variable, weekly BFI, to account for movements in the underlying market. The analysis is in two steps. The estimation is first conducted by the 47 routes which have a reasonable number of observations. The results are blurred (Appendix 3.5). Dwt coefficients have often wrong sign, are statistically insignificant and small in absolute terms. Age behaves differently. Its coefficients can be small but have always a correct sign and lack seldom statistical significance. Hopeless estimates are ignored and weighted (by fixtures) averages calculated for the rest (Table 3.12). When dwt and age pull in the same direction the effect is not trivial. Consider the most elementary situation, Atlantic loading regions with smaller and older ships confronting Pacific loading regions. The constellation works against the paradigm of this study which is exactly as it should be.

These observations will be modified when periods of high and low rates are juxtaposed. When rates are high (1995) there is more scope for variation than when they are low (1998). Only few routes have a sufficient number of fixtures and statistically significant coefficients for a test. The coefficients are related to the 1993-2002 period (= 100) and given as index points. Both variables support the hypothesis with approximate overall index ratios (95:98) 2:1 for dwt and 3:1 for age (Table 3.13).

Table 3.12 Dwt and age coefficients by segment, 1993-2002

Segment	Fixtures	Pct	Coefficient, weighted Dwt	Age
Handy	4,310	78	0.124	-94
Panamax	11,462	95	0.041	-110
CapeS	907	70	0.094	-326
CapeL	834	76	0.089	-250

Note: Coefficients weighted by number of fixtures per route.
Source: Appendix 3.5.

Table 3.13 Dwt and age coefficients on some routes, 1995 and 1998

Route	Coefficient Dwt 1993-2002	Coefficient Age 1993-2002	Coefficient Dwt 1995	Coefficient Dwt 1998	Coefficient Age 1995	Coefficient Age 1998	Index Dwt 1995	Index Dwt 1998	Index Age 1995	Index Age 1998
TH-37	0.112	-79	0.215	0.079	-167	-46	1.92	0.71	2.11	0.58
TP-13	0.087	-115	0.035	0.100	-137	-42	0.40	1.15	1.19	0.37
TP-17	-0.040	-105	0.212	0.107	-181	-28	-5.30	-2.68	1.72	0.27
TP-23	0.142	-124	0.213	0.139	-109	-74	1.50	0.98	0.88	0.60
TP-67	0.074	-127	0.170	0.051	-142	-124	2.30	0.69	1.12	0.98
TP-77	0.046	-147	0.110	0.080	-215	-39	2.39	1.74	1.46	0.27
TP-87	0.043	-141	0.243	0.080	-103	-53	5.65	1.86	0.73	0.38
TP-97	0.062	-63	-0.001	0.060	-153	-16	-0.02	0.97	2.43	0.25
TCL-27	0.111	-324	0.051	0.138	-615	-208	0.46	1.24	1.90	0.64
Avg. of positive							2.09	1.17	1.50	0.48

Source: Author.

The fixtures still contain an aspect which goes easily unnoticed but is extremely important for deriving internally consistent freight rate estimates in the next subchapter, the **number of ports**. More ports means higher port charges and a longer total time because of deviations and setup times at ports. When freight rates are unrelated to time or contain fixed elements their allocation between more days gives a lower day rate. Voyage rate, lumpsum and ballast bonus in trip rates are unrelated to time and there the possibility exists. It is important to obtain an idea of the possible distortionary effect which the varying number of ports in fixtures may cause. The best situation, of course, is if there is no variation.

The test is quite laborious and therefore only the panamax fixtures in 1997 are used for a check. Panamaxes is the segment where the ballast bonus plays the largest role (Table 3.15). Grain shipments by panamaxes are usually voyage chartered. The number of loading and discharging ports can be large. It is not unusual that a vessel loading grain in Argentina visits 4-5 ports, first on the La Plata River and thereafter on the Atlantic coast. Multiple loading ports are usual also in the US Gulf. When discharging in the Mediterranean or Japan the ship easily visits 3-4 ports. The first question is whether fixtures in the public domain are biased as to the number of ports. They probably are because they apply more often than not to comparatively new ships. This can be easily verified by looking at the monthly rate quotations for major trades published by consultants. The age bias matters because older ships with lower capital costs visit more ports than younger ones. Ships of 5 years and less can very well visit only one port at the both ends of a trip. That, however, is not the whole truth. Multi-porting also leads to higher rates.

Voyage fixtures can be very specific. Two ports instead of one in the panamax grain trade from USG to Japan adds $1.50 at the loading and $0.50-0.75 at the discharging end to a rate of $20, i.e., 10 pct. Similar examples are legio and the surcharge can well be 15-20 pct. The encouraging feature is that the options given in fixtures are so specific. It is possible to record the base case, one loading and discharging port, and overlook the rest. Trip fixtures give such information less often because the number of port visits is not the shipowner's concern. The relevant information can be found only by comparing specific fixtures with the ensuing physical movements. The comparison gives unexpected results. Out of 1,241 trip fixtures, 470 cannot be tied to a particular ship at all. These fixtures are known by ship name and dwt whereas movements are identified best by the IMO (hull) number (Lloyd's Register of Ships 1997). Ships change name too often for reliable identification. Out of the 771 remaining vessels only 643 can be tied to a particular movement through their name, dwt, reporting and arrival dates. The first layday had, unfortunately, been deleted from the file. For 98 fixtures the tieup is impossible. Most such fixtures must have been replaced by relets, although the 15 pct "shrinkage" is much smaller than generally assumed (see above). The remaining 30 problem fixtures identify only the country and if its coasts face different oceans the information cannot be used. The 643 usable fixtures still suffer from the geographical inaccuracy of the trip charters. The approximate character of the result hardly affects conclusions, however (Table 3.14).

Table 3.14 Average number of ports per panamax trip fixture, 1997

Loading region	Trips	Average
1	134	2.6
2	79	3.1
3	30	2.7
4	6	2.0
5	30	2.2
6	36	2.2
7	47	2.3
8	146	2.3
9	135	2.4
Total	643	2.5

Note: Average of 2.0 ports means that there is only one loading and one discharging port.
Sources: Drewry Fixture Data; LMIU Movement Data.

There are on average 2-3 ports per trip. The highest figures are in Regions 1 to 3. Cargo legs departing from them and covered by our data last normally 25-50 days. Each additional port may add 2-3 days and reduce the daily revenue by 1.5 pct (ballast bonus 20 pct of total revenue). The potential to destroy the functional relationship of our interest is there but it is not threatening. It must also be seen in relation to the uncertainty of the port times at large.

3.2 From fixture to freight rate

The fixtures need to be converted to freight rates suited for the Simulator. It means in practice that all elements of the freight bill must be added and allocated to the days when the ship is in cargo, either sailing or at port loading and discharging (cargo time). Allocation is between cargo days because information about them is fairly complete which is not the case with ballasting time. All that sounds simple but can be complicated in practice.

Voyage fixtures appear easier because they give the freight rate per cargo tonne with contracted tonnage (alternatively lumpsum freight), indicate loading and discharging ports and in 55-80 pct of cases the time to be spent in them, either directly or as loading and discharging rates per day. But there is a hook because brokers claim that the rate is not a function of the cargo leg alone but the preceding (and possibly following) ballast leg also. The preceding ballast leg is known at the time of chartering but is not released in the fixtures. No less problematic is the inclusion of bunkers, port charges (and canal fees) because these items vary in time and space.

Trip rates give the freight rate per day, which is ideal, but to that must be added a fixed ballast bonus where it exists. In loading areas with an endemic shortage of tonnage, such as the US Gulf, La Plata, South Africa, Australia and the Vancouver – Columbia River area, ballast bonus is usual in panamax and handy segments (Table 3.15). It follows that the conventional rate given as $/day must be multiplied by the number days from ship delivery to redelivery, to make it additive with the ballast bonus. Since delivery and redelivery locations are often defined as extensive areas or port ranges an extra element of error is introduced (LHS 2001, 94-95). An attempt to minimize the error is made by randomly allocating these fixtures among regions in proportions known from segmentized voyage fixtures (Appendix 3.6). Since the proportions need not be the same for voyage and trip fixtures it is uncertain whether any good is made thereby, however. ECSA is a typical case, understood to consist of regions 116 and 117. The destinations of their export cargos have approximately similar destinations but only approximately.

Table 3.15 Travel/Cargo ratios and ballast bonuses, trip fixture, 1993-2002

Loading	Travel/Cargo ratio				BB/Total freight			
region	Handy	Panam	CapeS	CapeL	Handy	Panam	CapeS	CapeL
1	1.11	1.14	1.41	1.33	0.06	0.21	0.06	0.00
2	1.26	1.37	1.56	1.48	0.10	0.17	0.04	0.02
3	1.05	1.07	1.17	1.08	negl.	negl.	0.00	0.00
4	1.33	1.67	2.00	2.00	0.01	0.02	0.00	0.03
5	1.18	1.49	1.82	1.93	0.04	0.09	0.04	0.02
6	1.08	1.22	1.50	1.36	neg	0.01	0.01	0.00
7	1.08	1.17	1.11	1.10	0.01	0.01	0.00	0.00
8	1.51	1.58	1.62	1.48	0.07	0.06	0.01	0.01
9	1.57	1.56	1.67	1.63	0.13	0.11	0.01	0.01
All	1.18	1.34	1.55	1.48	0.05	0.12	0.03	0.01

Legend: T = travel from delivery to redelivery; C = travel in cargo; negl. = below 0.5 pct.
Source: Author.

The cargo time comprises both sailing and cargo handling. Because trip fixtures do not contain any information about port time the information is taken from voyage fixtures, specifically their regional averages, derived from the said 55-80 pct of voyage fixtures (Table 3.10). The averages differ easily by 2-3 days in neighboring regions, possibly comprising a joint delivery or redelivery area. The differences can originate from port (even pier) individuality but equally well they originate from the type of cargo handled. Grain cargos, for example, take twice as much time per tonne as ore and coal, and trip fixtures do not contain information about the cargo. It is easy to comprehend the degree of inaccuracy thereby introduced. This is another reason, in addition to ship routing and multiporting, why the inclusion of an additional day for canal passage is considered an unnecessary complication (below). The situation is compeletely different when charter alternatives with known ports and cargos are compared.

The total freight revenue needs to be allocated to the cargo days, whether in port or sailing. About 45 pct of all trip fixtures have only delivery and redelivery areas/locations which, logically must also be loading and discharging areas. Or if they are not (Skaw, Cape Passero and similar), they can still be tied to a specific area. These constitute no problem. But the rest have one or more intermediary locations, like *Fu Hua* above. The fixture does not indicate which alternative is the loading respective discharging location. It is up to the analyst to make the choice (Fig. 3.1). Fortunately, in most cases the choice is obvious. But it is not always so. For example, when a handysize is delivered in Continent, will sail via West Africa and is redelivered in Far East, it is impossible to know which leg is cargo and which is ballast. Or, perhaps both legs are cargo. It can be fertilizers from Continent and tropical logs or agriproducts from West Africa, for example. Practically all trip fixtures for West and East Africa in the handy segment are such enigmas. Somehow the cargo days are defined, however, and the appropriate rate is derived. Numerical examples are given in Appendix 3.7.

The end result is two sets of freight rates, the one for voyage and the other for trip charter. They include partially different cost items and are, therefore, not directly comparable. The simplest alternative would be to use only trip fixtures and overlook voyage fixtures. After all, trip fixtures are the reduction base for every single voyage fixture, in one way or another. And when there are only voyage fixtures on a route, refuge is by necessity made to a substitute trip route. For a given transportation task trip and voyage rates are practically identical, claim brokers, when the same cost items are considered. What else do voyage fixtures bring to the study than increased variance? A fair question. The answer is that voyage fixtures cover partially different transportation tasks, be it routes or commodities, than trip fixtures. It is practically impossible, for example, to find a handysize trip fixture from Kamsar to US Gulf (Appendix 3.3). Most iron ore shipments from Narvik and Nouadhibou to the Continent are of the voyage type. And so on. There are routes on which the total number of fixtures is modest, even during a decade. Voyage fixtures make a welcome contribution, fully realized when route populations are disaggregated by year, month and week. Therefore, it is desirable to make use of both sets of fixture and consolidate them, if at all possible. The technique rests on the assumption that the cargo time of a given route is the same for voyage and trip fixture. This is not strictly true because port times vary by cargo. But it makes possible to convert the total voyage freight revenue into $/day. Whereafter the trip rate $/day is divided by it and a conversion ration T/V is derived. This is the principle. The actual calculations are rather different (below).

The ideal situation would be to juxtapose fixtures in a given ship segment and on a given route against each other on the same day. It would give an unpeccable conversion ratio. Unfortunately, that cannot be done. Most basically, there are not enough fixtures for each day. In the original study, the problem was solved by approximating the temporal midpoint of fixtures on a route and using the corresponding point estimates to find the conversion ratio. The practice was possible because the freight rates during 1997 had a steady declining trend. Now there are pronounced rate cycles. The cycles in different size segments and routes coincide broadly but not well enough to recommend wholesale consolidation (Fig. 3.2).

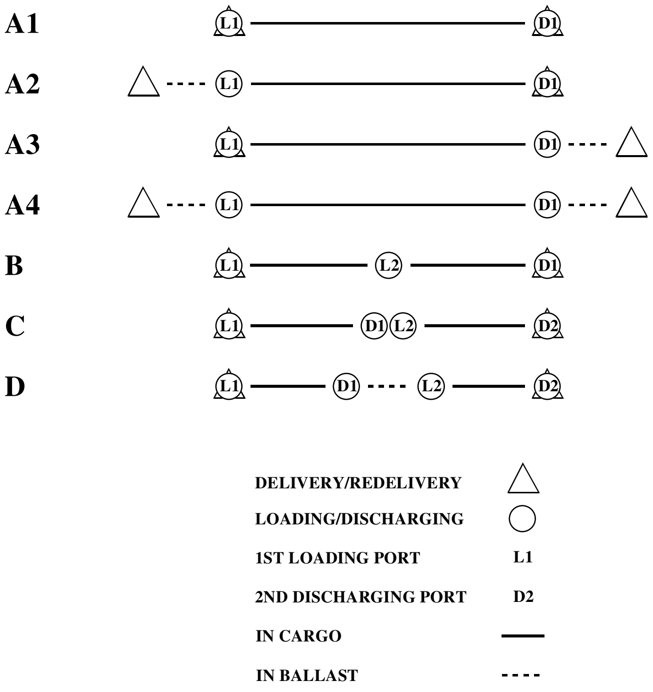

Figure 3.1 Cargo and ballast leg alternatives for trip fixtures

Note: Alternatives A and B display one trip L1-D1 each. Alternatives C and D display two trips L1-D1 and L2-D2 each.
Source: Author.

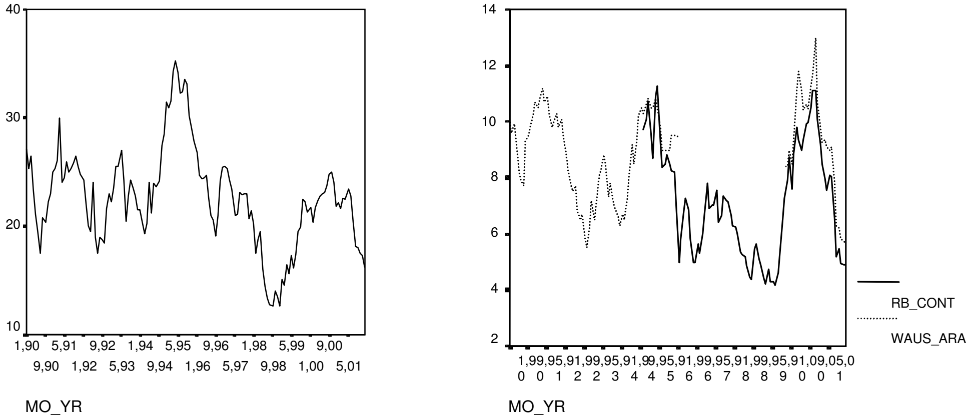

Panamax
Grain US Gulf to Japan

Capesize
Iron ore W Australia to ARA
Coal Richards Bay to Continent

Figure 3.2 Voyage freight rates, 1990-2000

Source: Drewry Shipping Consultants.

Therefore, the solution is to tie the segment- and route-specific rates to the general (weekly) rate level, calculate appropriate functions and juxtapose voyage and trip rate estimates with each other. The general rate level is measured by the Baltic Exchange's overall dry bulk freight index BFI. Since *Baltic Exchange* handles 30-40 pct of global dry bulk chartering the index can claim some representativeness (Shelley 2003). The index has been restructured many times over since its introduction in 1985 and may therefore appear to be a less than ideal market indicator (Kavussanos 2002, 679). Handysizes were removed from it in April 1997 and capesizes in early 1999 leaving it a pure panamax index (Solman 1998). It must be understood, however, that freight rates are reasonably correlated and the selection of weights for an index, particularly within a segment, is therefore not of decisive importance (Glen and Rogers 1997). BFI's undeniable virtue is that it exists, has predictive capability, although not needed here, and fares hardly worse in accuracy than the many other items of our empirical data (Chang and Chang 1996). The function is the simplest possible: \$/day = a + b * BFI. It is calculated separately for trip and voyage fixtures on each route (region pair). At least five fixtures, with few exceptions, are needed for a function. Otherwise they are dumped into a Rest group. It was soon discovered that neigboring regions could be consolidated into larger macroregions without much loss of information. In fact, we have been using these macroregions all the time. The following discussion is consequently about them although the term “region” is used for simplicity.

> To economize presentation, a codification system will be used in the continuation. First digit indicates fixture type, second ship segment, third loading and fourth discharging region. The codes are as follows: T = trip, V = voyage, C = consolidated, H = handysize, P = panamax, CS = capesize small, CL = capesize large. For example; TH-13 means trip rate/function for handysize from East Coast North America to Continent.

The ideal case is that the voyage function lies above the trip function and rises less steeply percentagewise, although not necessarily in absolute terms. The rationale is that tonnage shortage does not affect fuel price, port charges and canal fees. This is logical. The assumption holds in the examples of Figure 3.3. The uncomfortable features are the large variation at a given BFI value and occasional outliers. The variation can be reduced dramatically by measuring distances between ports (where possible) rather than regions and performing a full-scale regression analysis with all the ship and port characteristics known to influence rates. That would be beyond the point, however, because the application will be in the Simulator where all the ships are standardized and distances are between regions. Outliers were studied case by case and excluded from the analysis if found in some way non-representative. The resulting parameters appear promising (Table 3.16).

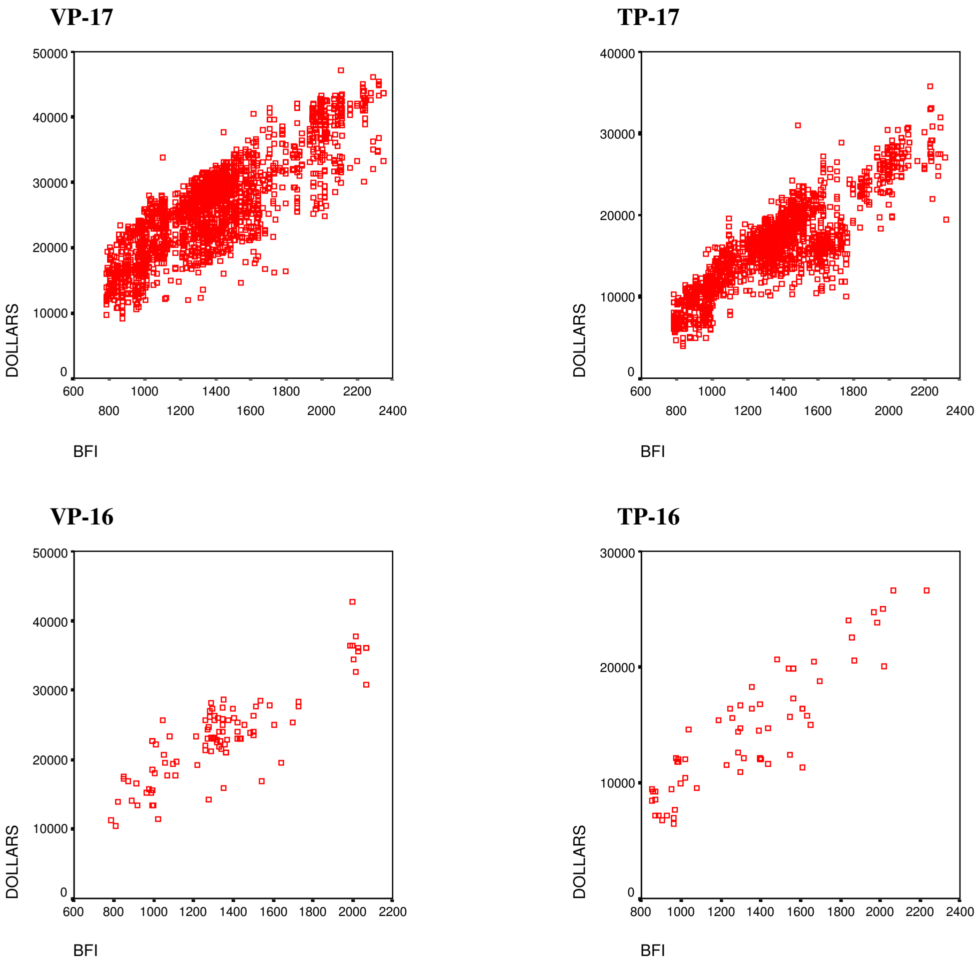

Figure 3.3 Well-behaving rate functions, 1993-2002

Source: Author.

Table 3.16 Statistical indicators for Figure 3.3

	VP-17	TP-17	Trip/Voy	VP-16	TP-16	Trip/Voy
Intercept	3,079	-2,093		1,052	-2,992	
Coefficient	16.751	13.375		16.674	13.016	
SEE	3,980	2,705		3,252	2,422	
R-sqr (adj)	0.682	0.741		0.735	0.793	
Observ.	2,115	1,641		98	64	
Estimate at						
BFI-750	15,642	7,938	0.51	13,558	6,770	0.50
BFI-2400	43,281	30,007	0.69	41,070	28,246	0.69

Note: Outliers excluded.
Source: Author.

So far so good. But these subpopulations are exceptionally well-behaving. Particularly the T/V ratios should be observed, lower at cycle bottom and higher at cycle top, as hypothesized. Many others are far worse. Sometimes it is possible to identify an apparent reason (Fig. 3.4). But several cases resemble a mosquito swarm with low correlations and even negative slope coefficients, ostensible a mixture of inaccurate distances and widely varying local conditions. A small number of fixtures contributes to the confusion although that is not a rule cut in stone as TP-51 and VCS-21 show (Appendix 3.8).

Figure 3.4 Problematic rate functions, 1993-2002 (p. 84)

Comments:

VH-11	Outliers above of diverse origin. No apparent reason.
VH-33	Three outliers above. Inappropriate distances. Excluded.
VH-63	Dense band in the middle. Sugar fixtures.
VCS-33	Mostly iron ore from Narvik. Too rigid distances.
TCS-53	Four outliers below. From 1995 (cycle top).
TCL-83	Four outliers below at BFI-2000/2400 and one above at BFI-800. No apparent reason.

Source: Author.

VH-11

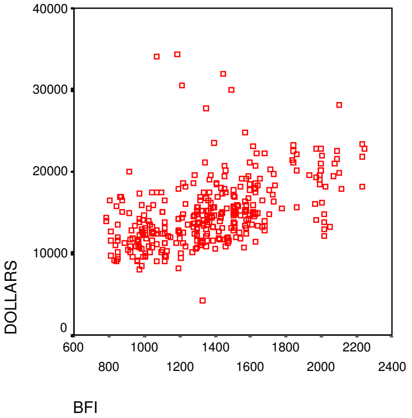

VH-33

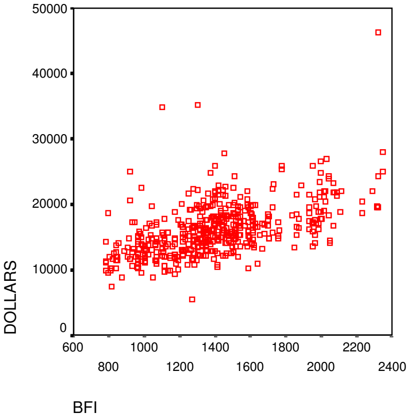

VH-63

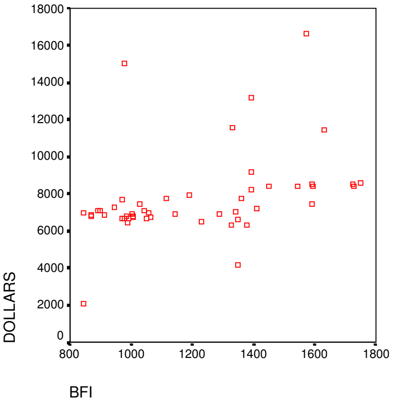

VCS-33

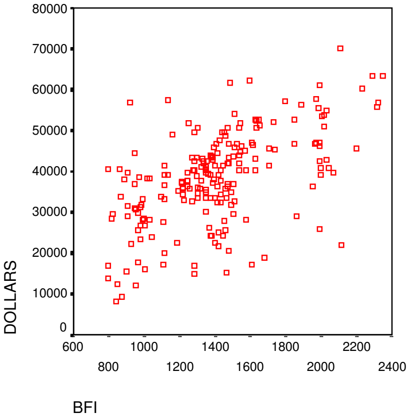

TCS-53

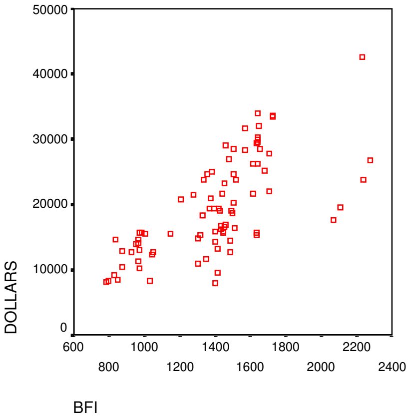

TCL-83

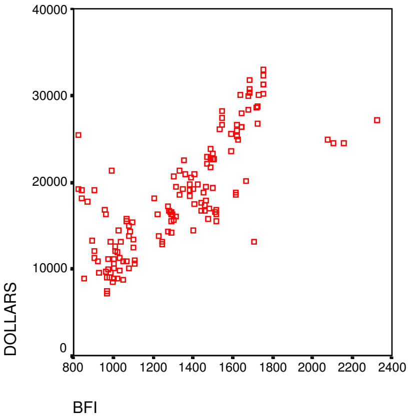

Where T/V ratios are hopeless the calculated functions are replaced by administered alternatives. These alternatives are functions taken from other routes and size classes or they are independent calculations. One should be careful when applying alternatives, however. Port charges vary widely, route lengths do the same, only some routes use canals, the type of cargo influences port time, functions are affected by the clustering of observations along the x-axis, and so on. Here, two sets of calculations are offered, the one by the author and the other by Drewry. The core results are in Table 3.17 and some details in Appendix 3.9.

Table 3.17 Trip/Voyage (T/V) comparisons

Segment	Cargo	Loads	Discharges	Cycle bottom		Cycle top	
				T_E	T_M	T_E	T_M
Table 3.16							
Panamax	all	ECNA	Far East		0.51		0.69
		ECNA	India & MEG		0.50		0.69
Author							
Panamax	grain	USG	Taiwan	0.50	0.52	0.79	0.69
CapeS	iron ore	Dampier	Rotterdam	0.31	0.49	0.68	0.60
Drewry							
Handy	grain	USG	Venezuela			0.77	
Handymax	grain	USG	Algeria			0.78	
	scrap	USWC	Kwangyang			0.79	
Panamax	grain	USG	Kobe			0.80	
			ARA			0.84	
		USNP	Kobe			0.84	
CapeS	coal	USAC	Rotterdam			0.80	
		Richards Bay	Rotterdam			0.81	
		Newcastle	Rotterdam			0.78	
			Kwangyang			0.78	
	iron ore	Narvik	Rotterdam			0.71	
		Tubarao	Rotterdam			0.79	
			Beilun & Baoshan			0.76	
		Saldanha Bay	Beilun & Baoshan			0.77	
		Dampier	Rotterdam			0.74	
			Beilun & Baoshan			0.69	

Legend: T_E = Trip equivalent after deductions; T_M = Trip measured from fixtures.
Note: In Drewry's set Dampier to Rotterdam is via Cape.

The input values in the author's calculations and Drewry's differ, the main difference being that the author's estimates of port charges are 2-3 times higher than Drewry's. For bunkers, the items largely balance each other out. The route Dampier-Rotterdam becomes longer via Cape (Drewry) but this is well outbalanced by the saving of Suez Canal fees. The result is a 0.06 higher T/V ratio. The ratios based on estimated trip rates (voyage rates minus cost items) at cycle top scatter around the 0.75 mark, two grain trades excepted, irrespective of ship size, cargo or route. At cycle bottom, most readings are close to the 0.50 mark and the author’s estimate for Dampier-Rotterdam at T_E also approaches this mark when the trip is routed via the Cape. In short, there is remarkable consistency for conversion factors 0.50 and 0.75 at cycle bottom and top, respectively.

The question is, whether these factors should be applied overall or whether the factors derived from the functions are preferable. The answer must be flexible, pay attention to the functions when they make sense (R-sqr > 0.300) and resort to judgement otherwise. Unfortunately, neither alternative turned out to be practicable. It is useless to dwell on the futile attempts and return will be made to the rationale of the exercise: To find out a way how the swarm of voyage fixtures is transferred upon the swarm of trip fixtures so that a maximum overlap is achieved. In most cases it is not enough that the voyage fixtures will be transferred a fixed distance in vertical direction. Rather the higher fixtures must be shifted a longer distance than the lower ones to achieve an acceptable fit. After some experimentation it was thought that the following formula was the solution:

$$T/V = M/R * E_1 + E_2/R * BFI$$

in which

E_1 and E_2 = Experimental constants for the difference of T/V ratio at cycle top and bottom (BFI-2400 and BFI-750), to be found separately for each route. E_1 is used for transferring the function vertically without changing its slope whereas E_2 tilts the slope.

BFI = BFI quotation for the topical week

M = Overall BFI minimum (750; exact 782)

R = Overall BFI range (1650; exact 2347-782=1565).

For example:	Voyage rate	T/V ratio		Trip equiv.
	10,000	0.4545 * 0.45 + 0.45 / 1650 * 1200	(0.53)	5,318
	20,000	0.4545 * 0.45 + 0.45 / 1650 * 2200	(0.80)	16,091

In other words, each voyage fixture gets an individual trip equivalent. The trip equivalents plus genuine trip rates are consolidated for each route and called in this and the next subchapter Consolidated Rates (C$/day or C_DOLLAR in some figures). Later on, when their use has become routine, the simple variant $/day is adopted. The conversion is called heuristic because of the heavy experimentation.

Several things should be observed. There may be too few trip fixtures for a proper conversion. Three trip fixtures is sufficient for an attempt, the philosophy being that a ”weak” implicit function is better than none. But if there are no fixtures, another similar route is used as a substitute. It is sufficient to know (or guess) the experimental constants (Appendix 3.10). The formula functions best when most of the fixtures are for a given port pair and ship size, more usual in the cape segments than the rest (Fig. 3.5). Experimentation with the constants to find a satisfactory fit can be a lengthy process and the result leave much to be desired. In four cases the formula is reduced into $T/V = E_3 * BFI$. Routes with 1-3 fixtures of the voyage or trip type only are lumped together into a Rest group, almost tantamount to rejection.

The fit is essentially visual although it can be supported by the rule that trip and voyage rate averages should be approximately equal after a successful conversion. The condition is that both fixtures distribute along the BFI-axis roughly in equal proportions, which is not always the case. Visual observation can get seriously impaired when the number of observations is large because hundreds of markers easily disguise each other.

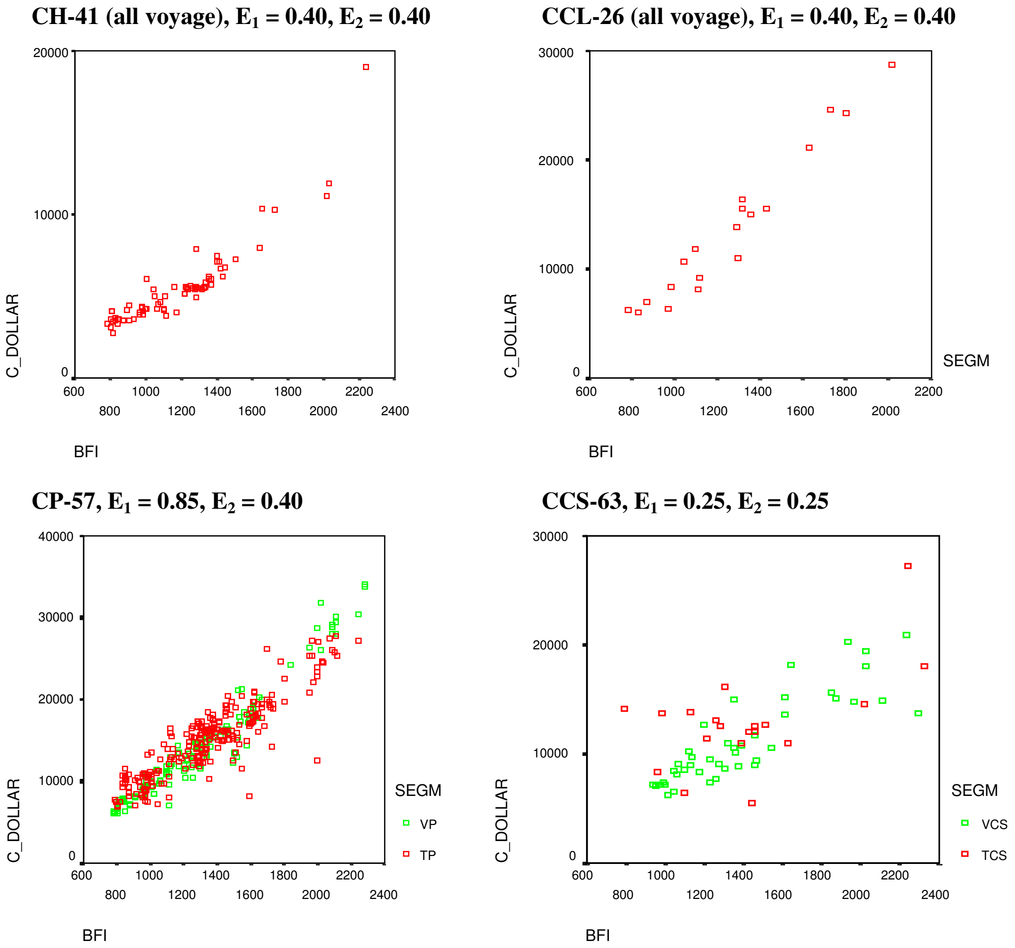

Figure 3.5 Heuristic conversion, examples

Source: Author.

The shortcomings of heuristic conversion suggest a strictly formal approach. The approach selected guarantees that the converted voyage rates have an identical function as the trip rates. It follows that the average of converted voyage fixtures is the same as of trip fixtures, if the both types of fixture distribute evenly along the BFI-range and on both sides of the joint function. This rule applies to the function as a whole but does not exclude local bias (Fig. 3.6).

The following formula performs the transfer:

$$CV = V + TP - VP$$

in which

CV = Converted voyage rate \$/Day

V = Original voyage rate expressed as \$/Day

TP = a + b * BFI (trip rate function for the topical route)

VP = c + d * BFI (voyage rate function for the topical route)

BFI = BFI quotation for the topical week

a, b, c, d = Parameters of the topical route estimated from fixture data

For example, the highest voyage rate on Route CP-17 in Week 114 works out as follows:

Voyage rate		Trip equiv.
47,189	-2,093 + 13.375 * 2,106 - 3,079 - 16.751 * 2,106	34,908

As in the heuristic conversion, each voyage rate gets an individual trip rate equivalent.

This approach has its own shortcomings. Some functions are too ”odd” to be usable at all, have a negative slope coefficient for example. Non-conforming observations must be rejected or the function be borrowed from another route. The conversion does not minimize the sum of squared distances as the OLS does, which usually means a larger variance than in the heuristic approach, faraway outliers of voyage origin, and numerous low, occasionally negative, estimates. When the data set will be disaggregated into shorter time units, weeks for example, such estimates will introduce a great deal of spurious variance. Odd estimates can be rejected but the rejections cannot be too many lest the credibility of the result be jeopardized. Rejections at this stage also add to editorial workload because of changed frequencies, for example. The scatterplots are less enjoyable, sometimes outright horrible (Fig. 3.6). That is not fatal, however, because the formal indicator SEE, to be used in the Simulator for delimiting the range within which freight rates are randomly generated, can be scaled down if necessary.

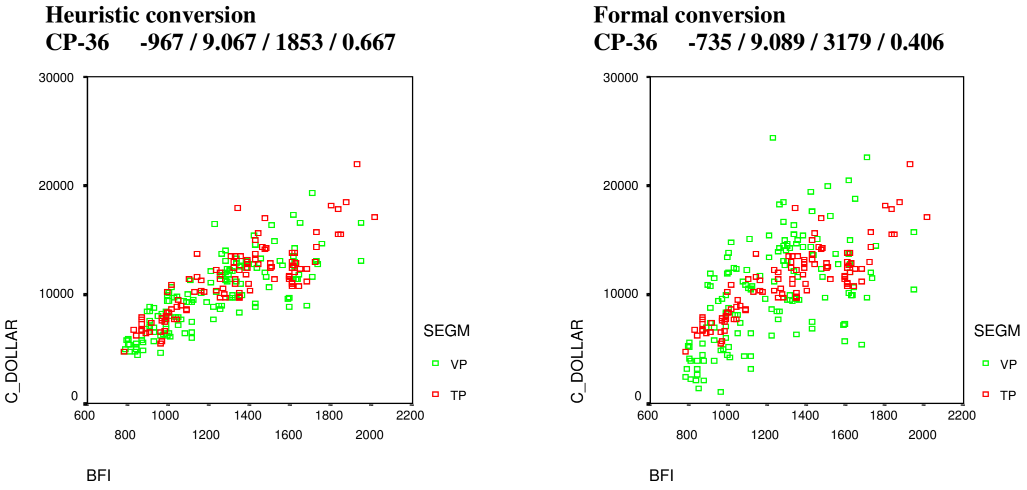

Figure 3.6 Heuristic and formal conversion, problematic example

Legend: Intercept / Coefficient / SEE / R-sqr (adj.).
Source: Author.

Both conversions, when tested with the 51 routes (inc. Rest group) of the panamax segment, are very similar in performance on three scores, averages, estimates and coefficients. The formal conversion is inferior on two scores, SEE and R-sqr (adj.). In more detail:

- Averages of conversions deviate from each other 6-10 pct on five routes and 11-15 pct on two routes, all with a small number of observations.
- Estimates of conversions at BFI-1350 deviate from each other 6-10 pct on six routes and 11-15 pct on two routes, all with a small number of observations.
- All coefficients are positive in both conversions.
- Three coefficients lack statistical significance in both conversions.
- R-sqr (adj.) is higher on 49 routes in the heuristic conversion.
- SEE is higher, up to twice higher, on 49 routes in the formal conversion.

The lower SEE in particular is reflected in more accurate point estimates. That decides the selection in favor of the heuristic conversion. The consolidated rates derived through it will constitute the data base for all subsequent freight rate analyses. The overall impression is favorable (Appendix 3.10) and so are the details:

- A fair share of the variation in rates (R-sqr, adj.) is explained in most cases. Only 12 and 44 routes (inc. Rest group) out of 152 remain below the 0.250 and 0.500 marks, respectively. These shares are for R-sqrs adjusted for the degrees of freedom, radically lower than the non-adjusted figures if the number of observations is small. Logically, most low R-sqrs originate from routes with comparatively few observations. Practically all are for handysizes and panamaxes whereas capesizes are widely overrepresented among the high R-sqrs. The reason is obvious enough. Capesize routes have few loading and discharging ports and the cargos are almost exclusively iron ore, coal and grain. Many factors contributing to the wide dispersion of rates are consequently eliminated. Exactly the opposite is the case within the smaller size segments. The correspondence between the cargos and the setup of the BFI also plays a role. It was expected that within-region routes and those discharging in the Continent would score low R-sqrs because of short distances with large relative errors and the highly fragmented coastline of Europe, respectively. Little evidence is found for either hypothesis.
- The slope coefficient is always positive. It is statistically insignificant at the 5 pct level on only 10 routes out of 152. These cases have either few observations or low R-sqrs or both. It need not mean that the measurement is incorrect. It may as well mean that the topical market is comparatively indifferent to the general rate level (backhauls) or simply lacks transparency. Overall, the rates are clearly related to the BFI which can be safely used as a driving force in the Simulator.
- The standard error of estimate (SEE), related to the approximate average daily BFI quotation BFI-1350 (exact BFI-1348), exceeds 0.30 on only five routes, whereas 51 routes out of 152 fall within the range 0.21-0.29. In other words, the rate estimates to be derived in the Simulator with the help of the route functions will have an acceptable although not small dispersion. This duly reflects the colorfulness of the Real World.

3.3 Anticipating Simulator results

At this stage it is possible to create a tentative idea whether the hypothesis underlying this study will be substantiated or not. Simplified to the utmost, the hypothesis states that loadings within the Atlantic Sphere give a higher freight revenue than those in the Pacific Sphere. The Atlantic/Pacific freight rate ratio is a way to make the idea tangible (Fig. 1.3). But an advance notice is preferably based on the actual data to be used to test the hypothesis. For this purpose, two panamax routes are selected, the one loading on the Atlantic and the other on the Pacific, their consolidated rates will be averaged by year, and compared with each other. The test routes are CP-17 and CP-87. The former loads in a region which scores high on the original 1997 Revenue Gradient and the latter is its opposite. About one half of fixtures on the former are of voyage origin whereas trip fixtures dominate on the latter. Both routes have the largest number of fixtures in their respective Spheres, 3,756 and 2,026, respectively. CP-17 is flanked by TP-17 to see how well the consolidated figures match with the reduction base. The estimates from the original study (LHS 2001) are also given for the sake of comparison. Their derivation was not as sophisticated as it is now. Yet, their comparison with the current estimates gives an indication about the accuracy of our work in general (Table 3.18).

Table 3.18 Average rates (C$/Day) on P-17 and P-87

Year	TP-17	CP-17	Diff TP/CP	CP-87	Diff 17/87	Atl/Pac
1993	18,456	17,194	1,262	15,557	1,637	1,336
1994	19,559	18,163	1,396	15,853	2,310	1,330
1995	25,581	26,826	-1,245	23,539	3,287	1,374
1996	17,818	16,273	1,545	14,616	1,657	1,597
1997	17,527	16,337	1,190	14,631	1,706	1,413
LHS		17,260		13,766	3,494	
1998	8,643	8,585	58	10,116	-1,531	0,931
1999	11,181	10,925	256	10,641	284	1,346
2000	15,559	15,934	-375	16,984	-1,050	1,020
2001	12,546	12,727	-181	13,613	-886	1,205
2002	13,810	13,726	84	12,156	1,570	1,338

Note: TP-17 estimates calculated from the trip function.
Sources: Author; Clarkson Research Studies; LHS 2001, Table 5.3.

Several conclusions can be drawn. The difference between TP-17 and CP-17 can be sufficiently large to affect results. When positive, it works against the hypothesis and the other way round. A positive difference between CP-17 and CP-87, by contrast, works for the hypothesis. This latter, crucial difference was in the original study double what it is here. It follows that the Revenue Gradient will not be as steep as originally. Anyway, it supports our hypothesis in seven years and works against it in three. The support appears to be connected with the Atlantic/Pacific rate ratio. How else, CP-17 and CP-87 are its largest components. The functional relationship is quite firm and it does not matter much whether it is calculated from the difference or the ratio CP-17/CP-87 (Fig. 3.7). The firmness is fortunate because it allows drawing approximate conclusions also about other years than 1995, 1997 and 1998. Equally fortunate, 1997 and 1998 are at the opposite ends of the gradient and may allow some wide-sweeping speculation about the geographical order of dry bulk shipping in general.

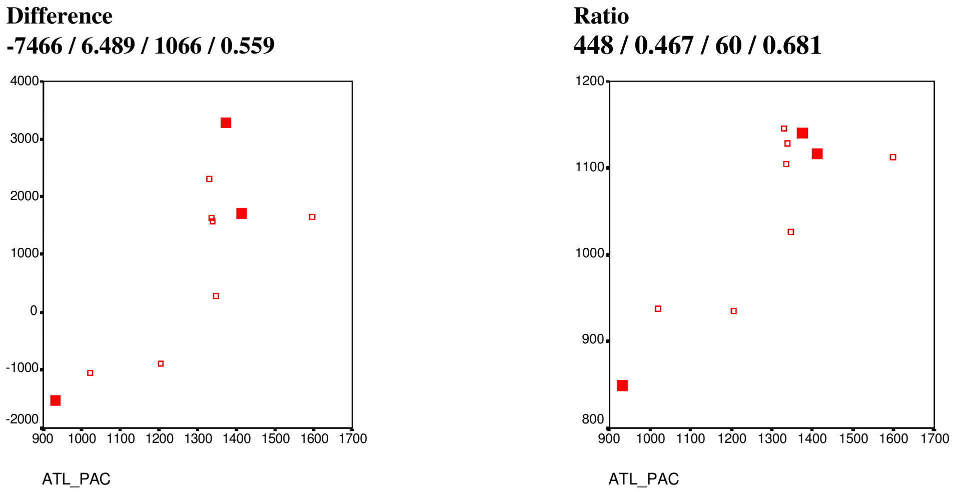
Difference
-7466 / 6.489 / 1066 / 0.559

Ratio
448 / 0.467 / 60 / 0.681

Figure 3.7 Comparison of CP-17 and CP-87 (C$/Day) as function of Atlantic/Pacific Ratio, 1993-2002

Legend: Intercept / Coefficient / SEE / R-sqr (adj.).
Note: Atl/Pac values multiplied by 1,000.
Source: Author.

This looks reasonably good but the estimates to be used in the Simulator will not be calculated from rate averages but the functions of the consolidated rate sets using weekly BFI values as argument. The resulting rates are called in this subchapter ‘estimated” (E$/Day). Then the picture changes radically (Table 3.19).

- Differences between TP-17 and CP-17 become much smaller. This is welcome.
- Differences between CP-17 and CP-87 become smaller during 1993-7. This works against the hypothesis. During 1998-2002 they are erratic.
- Differences between CP-17 and CP-87 become erratic vis-a-vis the Atlantic/Pacific Ratio. The scatterplot is not displayed.

The last observation is fatal because the trustworthiness of the Atlantic/Pacific Ratio is beyond reasonable doubt. Comparison with Table 3.18, where differences are calculated from ‘raw” rates without the intermediation of rate functions, suggests that the rate functions calculated from the whole 1993-2002 period are not accurate enough. The solution is to disaggregate the period and calculate separate functions for each subperiod. Two-year subperiods are selected. Some are dominated by an overall trend, declining during 1995-8, and increasing during 1999-2000. It is reasonable to assume that at least during these periods the forces underlying the rate functions remained more or less stable.

The R-sqrs (adj.) give some support for that (Table 3.20). When the biannual functions are applied to the consolidated rates most of the variance is reduced away (Fig. 3.8). Annual averages calculated from these new estimates change the picture, once again (Table 3.21). The comparison between TP-17 and CP-17 has at this stage become irrelevant and is omitted.

Table 3.19 Average rates (E$/Day) by heuristic function P-17 and P-87

Year	TP-17	CP-17	Diff TP/CP	CP-87	Diff 17/87	Atl/Pac
1993	16,645	16,640	5	15,360	1,280	1.336
1994	18,545	18,150	394	16,252	1,899	1.330
1995	24,309	25,577	-1,268	23,004	2,573	1.374
1996	15,724	15,441	281	14,608	833	1.597
1997	16,017	15,768	249	14,596	1,173	1.413
LHS		17,260		13,766	3,494	
1998	10,640	9,743	897	9,551	192	0.931
1999	11,643	10,895	748	11,309	-414	1.346
2000	19,508	19,816	-309	17,806	2,010	1.020
2001	13,783	13,448	335	12,990	458	1.205
2002	13,181	12,545	636	12,583	-37	1.338

Note: TP-17 estimates calculated from the trip function.
Sources: Author; Clarkson Research Studies; LHS 2001, Table 5.3.

Table 3.20 Biannual rate (C$/Day) functions for CP-17 and CP-87

Period			CP-17					CP-87		
	Obs	Interc	Coeff	SEE	R-sqr	Obs	Interc	Coeff	SEE	R-sqr
1993-4	753	-5,268	15.887	2,103	0.741	467	1,766	9.726	1,632	0.580
1995-6	1,050	-4,562	15.889	2,986	0.803	545	-3,441	13.498	2,209	0.844
1997-8	840	-10,264	19.786	2.037	0.814	451	-1,629	12.309	1,738	0.698
1999-0	626	1,108	9.391	1,872	0.730	300	-1,665	11.630	2,059	0.744
2001-2	487	1,974	9.715	1,779	0.656	263	-4,123	14.215	1,837	0.800
1993-2	3756	-5,080	15.570	2,700	0.799	2,026	-2,425	12.740	1,979	0.832

Source: Author.

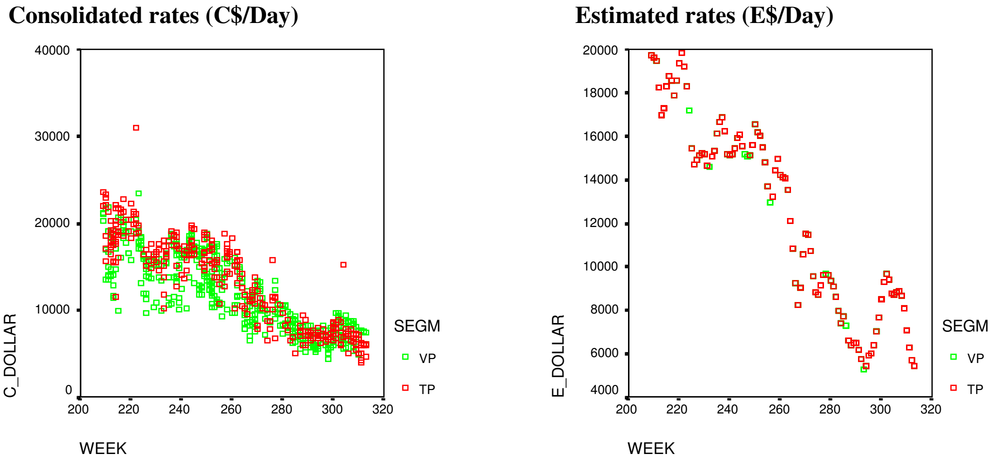

Consolidated rates (C$/Day)

Estimated rates (E$/Day)

Figure 3.8 From consolidated to estimated rates, CP-17, 1997-8

Source: Author.

Table 3.21 Average rates (E$/Day) by biannual function CP-17 and CP-87

Year	CP-17	CP-87	Differ	Ratio	Atl/Pac
1993	16,894	15,347	1,547	1.101	1.336
1994	18,430	16,027	2,403	1.150	1.330
1995	26,717	23,505	3,212	1.137	1.374
1996	16,374	14,628	1,746	1.119	1.597
1997	16,255	14,787	1,468	1.099	1.413
LHS	17,260	13,766	3,494		
1998	8,603	9,931	-1,328	0.866	0.931
1999	10,741	10,892	-151	0.986	1.346
2000	16,129	16,807	-678	0.960	1.020
2001	13,532	13,082	450	1.034	1.205
2002	12,969	12,624	345	1.027	1.338

Sources: Author; Clarkson Research Studies; LHS 2001, Table 5.3.

The measures taken have the desired effect. The function between the Atlantic/Pacific Ratio and the average rates of CP-17 and CP-87 is reestablished (Fig. 3.9). The relationship is not as close as when it was calculated from consolidated rates directly, but this is what can be expected. It is sufficiently firm for careful generalizations and that may be what can reasonably be expected from an aggregate indicator. The exercise also drives home a practical necessity. The rate functions to be used in the Simulator must be calculated from the biannual data rather than the whole period 1993-2002. That, of course, is possible only on routes which have enough fixtures (Appendix 3.11). Elsewhere the functions of Appendix 3.10 must be used.

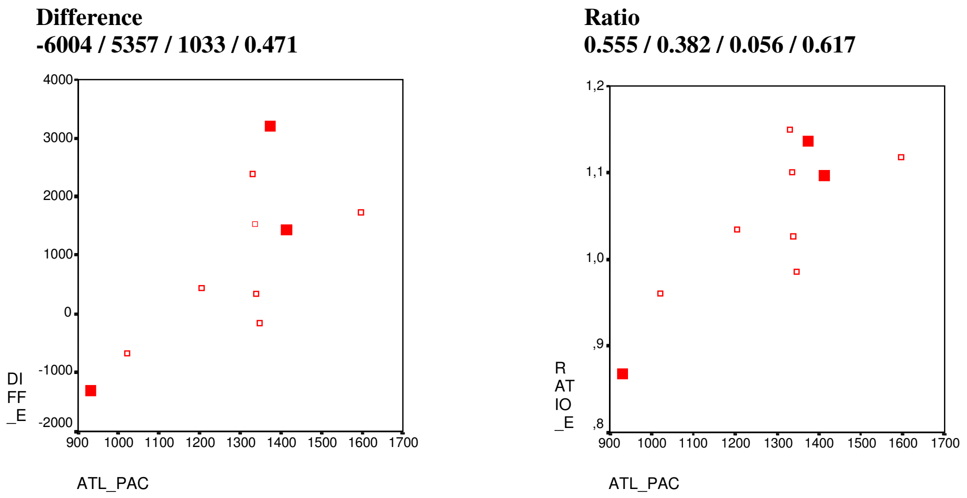

Difference
-6004 / 5357 / 1033 / 0.471

Ratio
0.555 / 0.382 / 0.056 / 0.617

Figure 3.9 Comparison of CP-17 and CP-87 (E$/Day) as function of Atlantic/Pacific Ratio, 1993-2002

Legend: Intercept / Coefficient / SEE / R-sqr (adj.).
Note: Atl/Pac values multiplied by 1,000.
Source: Author.

The comparison of CP-17 and CP-87 has given valuable information about the reliability of the measurement and the way how to conduct the technical calculations. The comparison also supports the hypothesis which this study tries to verify. But the conclusions apply only to two panamax routes among a total of 152, divided between four segments. It is valuable if the tentative idea can be extended to the remaining 150. That cannot be done meaningfully by route because they are so many. Nor is the hypothesis about individual routes but about routes loading in a given region, ten in number. The test must be scaled accordingly.

The test is based on the 1993-2002 functions (E$/Day) at BFI-1350, the approximate mean of weekly BFI quotations. Each route within a segment is given a point estimate and a weighted (by observations) average is calculated for each loading region. The averages for the original study are calculated in the same way, except that the point estimates are taken directly from a table (LHS 2001, Table 5.3). The observations for the original study are cargo legs (LHS 2001, Table 3.4) and for the current data they are fixtures. The comparability, therefore, is not perfect but because the results make sense no more effort is spent for the matter. Within each segment the highest regional average is given an index score 100 and the other averages are related to it. This makes regions and segments commeasurable. There are substantial differences in the number of observations between regions but since these are irrelevant in this context they are not displayed. Three comparisons are made, one for handysize and panamax, one for the two capesizes (small and large), and one for panamax in 1997 and

1993-2002 (Fig. 3.10). It is important to realize that the freight rates are calculated by using cargo days as a reduction base and include also revenue earned during a paid ballast leg.

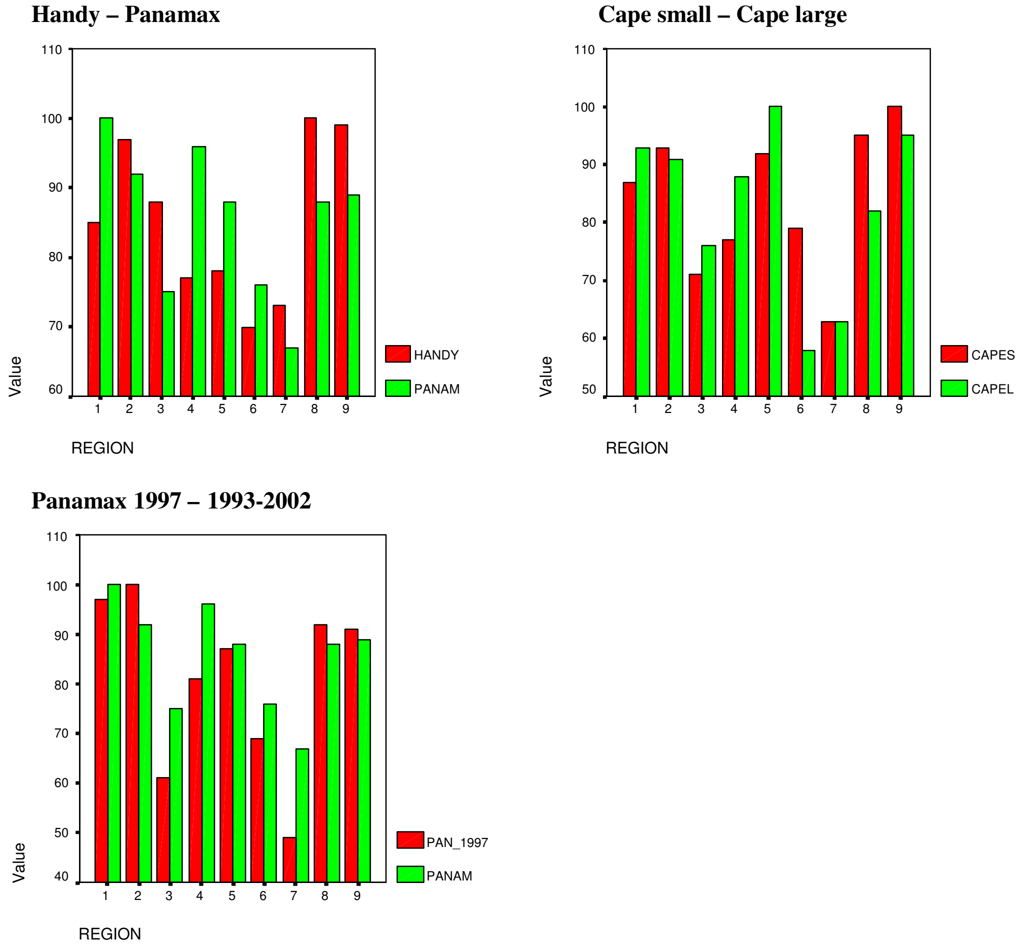

Handy – Panamax

Cape small – Cape large

Panamax 1997 – 1993-2002

Figure 3.10 Freight rates by loading region at BFI-1350

Sources: Author; LHS 2001, Table 5.3.

We begin with the panamax indices. Regions with small and large identification numbers, which also happen to represent the high and low values of the Revenue Gradient, are remarkably similar in the both data sets. So is Region 5, a midpoint in the function as in the Atlantic/Pacific Ratio. The large import Regions 3 and 7 have much higher scores in the current than the 1997 data. The reason is simple. The 1997 data had comparatively few fixtures from these regions and particularly the within-region rates were more guesses (far too low) than anything else. Alas, they represented comparatively large cargo flows. It is difficult to draw a general conclusion from all this. But a speculative scatterplot can be prepared, displaying the rates (rather than index scores) against the Revenue Potential derived in the original study for 1997 (LHS 2001, Appendix 6.4). Both plots are only approximate. But with this reservation, the general picture is encouraging (Fig. 3.11). There is an apparent Revenue Gradient. The import Regions 3, 6 and 7 are in the lower left-hand corner. Among the rest, positions shift with the exception of Region 1 which ranks high in both figures. To the 1993-2002 function can also be added a significant dichotomous variable claiming that rates are $850/day higher in the Atlantic than Pacific Sphere. That insight comes at the expense of a higher SEE and lower R-sqr (adj.), however. But more important than such detail is the overall support given to the hypothesis. From now on, things cannot go entirely wrong.

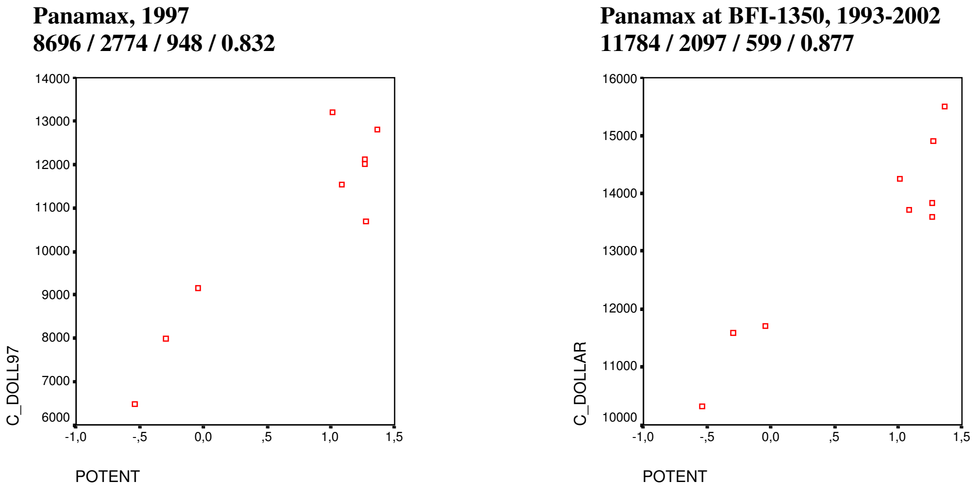
Panamax, 1997
8696 / 2774 / 948 / 0.832

Panamax at BFI-1350, 1993-2002
11784 / 2097 / 599 / 0.877

Figure 3.11 Regional average (C$/Day) as function of Revenue Potential, tentative idea

Legend: Intercept / Coefficient / SEE / R-sqr (adj.).
Sources: Author; LHS 2001, Fig. 5.3 and Appendix 6.4.

A similar test is impossible in the other segments because there is yet no information about the Revenue Potential (Appendix 4.5). Rough estimates based on fixtures rather than cargo legs are possible although superfluous in this study. But given a shortage of resources, time and money, they would be a possibility. Here it is sufficient to compare the index scores. In the handy segment, the relative position of the import Regions 6 and 7 should not change markedly, and Regions 4 and 5 will obviously move down in the vertical direction to about the same level. Regions 2, 8 and 9 will stay in the upper right-hand corner and Regions 1 and 3 constitute an intermediary group. If there are no marked changes in the Revenue Potential (there are) it is quite possible that no clear functional relationship will emerge. In the cape segments, the most likely outcome is that the import Regions 3 and 6 will form one group as previously, whereas Region 7 may lack loadings. The mutual positions within the export group may change, the Pacific loadings overshadowing the Atlantic ones and overturning the state of affairs as we know it from the panamax and handy segments (cf., Tables 4.8 and 4.9).

3.4 Rate structure

3.4.1 Clustering

Rate structure is about the mutual relations of freight rates on different routes. In the original study the outcome was a tree-like graph (LHS 2001, Fig. 4.13) based on eight simultaneous rate equations, each for a group of routes (LHS 2001, Table 4.14). The right-hand variables included the lagged value of the left-hand variable, except for the USAC/USG – Far East route which was given externally (Root; here CP-17). The data were 52 weekly index scores complemented with linear interpolations where genuine observations were lacking (LHS 2001, Fig. 4.11). Grouping was a technique to minimize the number of lacking observations but has proved to be a risky solution because routes have different rate levels and cyclicity, and because the proportions of different routes are likely to vary in time (LHS 2001, Table 4.13). Indices were used to cleanse the data from a strong trend which threatened to overwhelm everything else. They were calculated by dividing the weekly group rates by the corresponding global rate. The drawback was that estimates could not be converted back to freight rates. The R-sqrs seldom exceeded 0.5 but most coefficients had logical signs and were significant at the 5 pct level. The equations were not tested for co-integration, which should have been done (Veenstra 1999a, 146-153). The core message was that there were two rate clusters, the first having its loading ports in the Atlantic Sphere and the other in the Pacific Sphere. The split came to day also in the negative sign of the base route (Root) in Pacific equations.

The first decision now is whether the time dimension should be overlooked alltogether. If the earlier core result was the split into Atlantic and Pacific loading regions it should be possible to derive it without time series and lagged variables. It is possible, for example, to calculate rate estimates for a range of BFI-values by using the functions in Appendix 3.10. The set BFI-1000, BFI-1400 and BFI-2000 covers the empirical range of fixtures quite well and the point estimates can be used for clustering the routes. Not all the 150 routes which have been specified previously, however, but only those with a reasonable number of observations, say, those with at least 50 fixtures. That gives a total of 88, distributed between segments as follows: 37 in handy, 23 in panamax, 18 in small cape and 10 in large cape. The segments are analyzed separately because of the difference in ship size and correspondingly rates. There is also an additional although unused criterion, R-sqr, which gives an idea of the share of variance accounted for by the topical function. If R-sqr is low, say, below 0.25 that casts doubt over the reliability of the point estimates and, by extension, the route's correct classification. Six such routes exist in the handy segment and they can be excluded without too much ado if so desired. It should be observed, however, that a low R-sqr can also mean that rates do not react to change in BFI in general. That can happen on backhaul routes where shipowners are anxious to return to lucrative loading areas with little concern about the way they do it, loaded or not.

The clustering technique used is the Average Linkage Between Groups with squared Euclidean distances as a similarity measure. A certain structure can be identified in every segment but it is nowhere particularly geographical (Appendix 3.12). The conclusion is that the three point estimates based on BFI are not well suited to unravel a geographical rate structure for dry bulk routes. Attention is consequently turned to time series. There the existence of time units without observations becomes important.

The problem crystallizes in the number of fixtures on a route and their distribution in time. The largest flow CP-17 has 3,756 fixtures (Appendix 3.13) and the maximum number of time units is 520 (weeks). Also 2-week and 4-week periods are useful timeframes. For each time unit is calculated a weighted average freight rate and corresponding BFI using weekly fixture numbers are weights. If the fixtures are distributed evenly between time units there will be about 29, 14 and 7 fixtures for each 4-, 2- and 1-week period. But most routes have far less fixtures, these are seldom distributed evenly, and occasionally a pronounced seasonal pattern is discovered, for example on route CP-27. It follows that even with 1,000 fixtures there is a good chance that some time units will have no fixtures at all (zero cells; Fig. 3.12, Note). When this happens a surrogate estimate is calculated from the consolidated rate function (Appendix 3.10). Each surrogate entry downgrades the value of the data and to avoid ending into a spurious world, routes with more than 8 pct surrogates are excluded from the analysis. Converted to frequencies, that means a maximum of 10, 20 and 40 surrogates for 4-, 2- and 1-week time periods. It means in practice that only 4- and 2-week periods, based on actual fixtures, can be used for structural analysis.

Before going further it is useful to compare the routes of 4-week periods with routes for which monthly rates, trip or voyage, are routinely published in the trade press (Table 3.22). The discrepancy is considerable. In the handy and small cape segments trade press covers more routes than this study whereas the opposite is true in the panamax and large cape segments. Where coverage is superficially the same, trade press may have 2-4 rates per route. The trade press simply meets a different need than this study and therefore the difference. Trade press figures are point estimates (without cross-sectional variance) for a narrow size and age range of ships plying between given ports or port ranges, usually more narrow than the regions of this study. Trade press makes no attempt to convert voyage rates to trip rates but sticks to the custom of the industry. The rates are administered estimates based on actual fixtures. In months without fixtures, no rates are published.

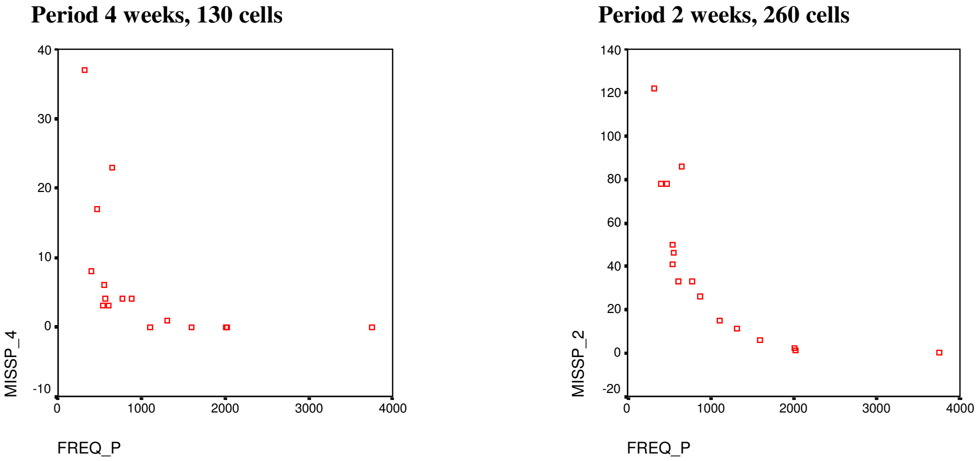
Period 4 weeks, 130 cells

Period 2 weeks, 260 cells

Figure 3.12 Empty cells per panamax route as function of fixture numbers

Note: The cutoff point for 4 weeks is roughly 400 fixtures and for 2 weeks about 1,000 fixtures.
Source: Author.

Table 3.22 Routes qualifying for 4- and 2-week analyses

Handy				Panamax				CapeS				CapeL			
Route	4 wk	2 wk	D	Route	4 wk	2 wk	D	Route	4 wk	2 wk	D	Route	4 wk	2 wk	D
CH-13	Y	Y	2	CP-13	Y	Y	1	CCS-13	Y	Y	4	CCL-23	Y	N	
CH-23	Y	Y		CP-17	Y	Y	3	CCS-23	Y	N	1	CCL-27	Y	N	
CH-33	Y	N	1	CP-23	Y	Y	1	CCS-27	Y	N	2	CCL-83	Y	N	
CH-37	Y	N	2	CP-37	Y	N		CCS-53	Y	Y	1				
				CP-53	Y	N		CCS-57	Y	N	1				
				CP-57	Y	N		CCS-83	Y	N	2				
CH-11	N	N	2	CP-67	Y	N		CCS-87	Y	N	3				
CH-53	N	N	1	CP-73	Y	N									
CH-67	N	N	1	CP-77	Y	Y									
CH-87	N	N	1	CP-83	Y	N		CCS-17	N	N	1				
CH-97	N	N	1	CP-87	Y	Y	1	CCS-33	N	N	1				
				CP-93	Y	N		CCS-37	N	N	1				
				CP-97	Y	Y	2	CCS-43	N	N	1				
								CCS-73	N	N	1				
								CCS-97	N	N	1				

Leged: D = Rates published in Drewry Monthly; N = no; Y = yes.
Source: Appendix 3.13.

When plotted against corresponding BFIs the 4-week rate averages give reasonably compact scatterplots and consistent functional parameters (Fig. 3.13; Table 3.23). The functions are used to calculate rate estimates at BFI-1350. These and the parameters are used to arrange the panamax and small cape routes into heuristic clusters. Handy and large cape routes are too few to make clustering worthwhile. The seven small cape routes split nicely into Atlantic and Pacific destinations. This is against conventional wisdom which differentiates between Atlantic and Pacific loadings. The difference between the groups is that the Pacific functions are more sensitive to changes in the BFI than the Atlantic functions. When BFI-800 is taken as the approximate bottom quotation for the BFI then the both groups of trajectory start from about the same level, but when BFI increases the Pacific functions reach higher. This is a remarkable exception to the general rule that Atlantic rates are most of the time higher than Pacific rates (Fig. 1.3). The panamax segment is more complicated. It is grouped into five clusters which is much for 13 routes. The boundaries are also less sharp than in the small cape segment. The two last clusters can be described as backhauls of a kind, to the Continent, although only CP-73 is a genuine backhaul. The mid cluster has two comparatively short hauls to the Far East. The two first clusters cannot be described unambiguously in geographical terms. They are more like the small cape clusters, differentiated by the sensitiveness to changes in the BFI.

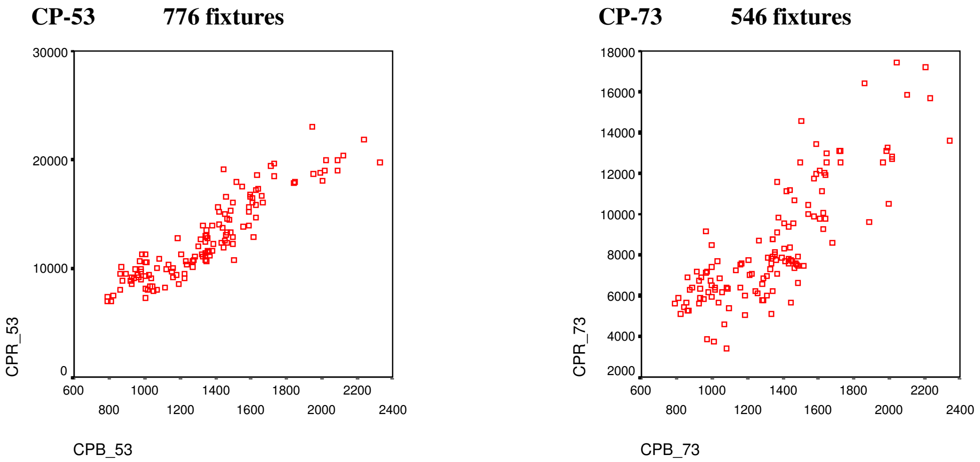
CP-53 **776 fixtures**

CP-73 **546 fixtures**

Figure 3.13 Two scatterplots of 4-week periods

Legend: Vertical axis = C$/period; horizontal axis = BFI/period.
Source: Author.

Table 3.23 Parameters of the 4-week period

Panamax Route	Interc	Coeff	SEE	Rate	CapeS Route	Interc	Coeff	SEE	Rate
CP-13	-4,135	14.210	1,348	15,049	CCS-13	-7,171	17.883	2.206	16,971
CP-17	-4,414	14.997	1,764	15,832	CCS-23	343	11.772	2,414	16,235
CP-57	-4,865	14.852	1,499	15,185	CCS-53	-6,651	17.363	2,566	16,789
					CCS-83	-4,739	15.044	2,722	15,570
CP-23	-1,918	11.844	1,543	14,071	CCS-27	-10,914	22.603	2,339	19,600
CP-87	-2,574	12.773	1,160	14,670	CCS-57	-9,444	21.257	2,622	19,253
CP-97	-1,515	12.142	1,280	14,877	CCS-87	-8,658	21.295	2,525	20,090
CP-37	-2,335	10.848	1,852	12,310					
CP-67	-1,712	10.044	1,391	11,847					
CP-77	-1,385	9.409	1,229	11,317					
CP-53	-1,002	10.198	1,536	12,765					
CP-83	-999	8.696	1,539	10,741					
CP-93	582	8.170	1,610	11,612					
CP-73	-1,156	7.170	1,658	8,524					

Note: Rate at BFI-1350.
Source: Author.

3.4.2 Cyclicity

Clustering routes by rate function parameters or point estimates based on the BFI is not the only possibility. Rate cycles are also available. The route sample is the same as used for the BFI-based point estimates with at least 50 fixtures per route, divided into ship segments. This time point estimates are not derived for a few BFI values but calculated for the 520 weeks during 1993-2002. The apparent weakness of these estimates is that they all are directly related to the same variable, the BFI. To downphase this feature the rates are converted into index scores by dividing them routewise by the weighted weekly averages of the segment, with fixture numbers as weights. This weighted series approximates the submarket at large. The variation of index scores within a segment then reflects the varying rate levels by route and the changes of the BFI but not in a strict one-to-one way. Tendencies for clustering are investigated by calculating Pearson correlation coefficients between all routes and comparing them with each other. The absolute values of the coefficients are very high, normally above 0.900 and often 1.000. The sign, however, can be either positive or negative and this characteristic runs through all the routes within a segment. In other words, if CH-11 is negatively correlated with CH-12, CH-13, etc. these other routes are positively correlated with each other. And if CH-11 is positively correlated with CH-21, CH-23, etc. these other routes are, again, positively correlated with each other. Consequently it is possible to split the routes within a segment into two distinctive groups, a "positive" group with a similar cyclic pattern as the underlying 1-week rate estimates and a "negative" group with an inverse, anticyclic pattern (Fig. 3.15). The groups are so compact and well apart that no formal clustering technique is necessary.

The examples have a simple cyclic/anticyclic logic. The staple cargo on these particular fronthauls is grain. It is logical that grain cargos loaded on the St. Lawrence River and in US Gulf follow similar rate patterns irrespective of whether the destination is Continent or Far East. By contrast, grain cargos loaded in Southern Brazil and the La Plata area have the harvest period half a year earlier which is reflected in the main shipping season and rates. In many other cases it is difficult to find a plausible explanation but that need not mean that there is none. It may only be a longer chain of reasoning, or the answer may be beyond our knowledge. Geographical order may hide in the split into loading and discharging regions, or the conventional fronthaul/backhaul thinking irrespective that the simplistic idea of opposing unbalanced cargo flows does not work very well (below). When evaluating these ideas one should note that the cyclic/anticyclic pattern need not split 50/50 by segment, even approximately. The reason is the varying number of fixtures (weights) by route.

The split into loading and discharging regions has some merit (Table 3.24). The routes themselves do not change pattern, of course. They are only split into two by activity and then arranged by macroregion. Loadings in Region 1 follow the cyclic pattern quite well in all segments whereas Region 7 is its direct opposite. These are the most important net exporter and importer, respectively, which establishes a budding "theory". Region 3 is the next largest net importer and predominantly anticyclic, raising hopes. Region 5 is more cyclic than anticyclic and so is Region 2 although only in the larger segments. Both are net exporters and the inclination makes sense. The anticyclic smaller segments in Region 2 can be explained by the agricultural exports and location in the Southern hemisphere. Region 6 is rather balanced and so are the cyclic and anticyclic elements. But then comes an apparent anticlimax. Regions 8 and 9 are typical net exporters, yet mostly anticyclic. But so is their most important trading partner Region 7. In conclusion, the loading paradigm makes sense. The discharging angle has merit only in few regions. Region 1 has mostly handy routes and there the anticyclic pattern is more common. These routes are typical backhauls with small volumes. Region 3 is also more anticyclic than cyclic and there this observation covers all but the large cape segment. The other regions are comparatively balanced.

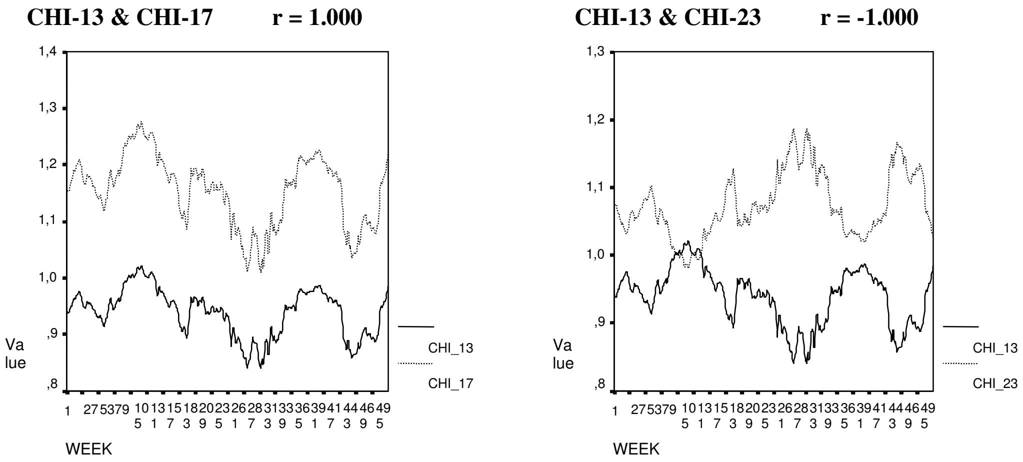

Figure 3.14 Positive and negative correlations of index scores

Notes: Solid = CHI-13, Dotted = CHI-17, CHI-23. Values divided by 10,000 for better display.
Source: Author.

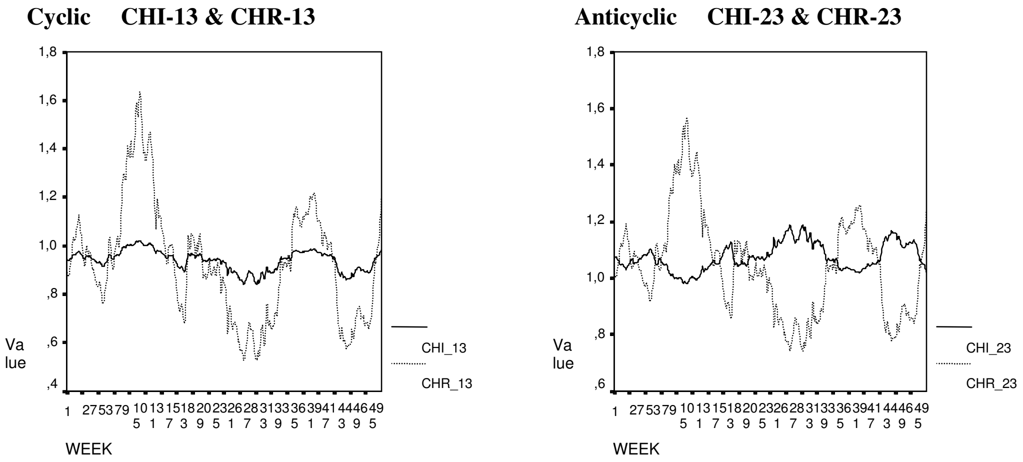

Figure 3.15 Cyclic and anticyclic patterns of index scores and underlying rate estimates

Notes: Solid = Index, Dotted = Rate. Values divided by 10,000 for better display.
Source: Author.

The fronthaul/backhaul paradigm, even when supported by data, can be of partial relevance only. There are only 13 two-way route pairs and 10 of them are in the handy segment. In 7 cases out of 10 the handy pairs show both cyclic and anticyclic patterns while the other segments show it in two cases out of three. The evidence is only tentative, not conclusive.

Table 3.24 Cyclic/anticyclic index patterns

Loading						Discharging					
Reg	Rte	Handy	Pana	CapeS	CapeL	Reg	Rte	Handy	Pana	CapeS	CapeL
1	11	anti				1	11	anti			
	12	cycle					21	anti	cycle		
	13	cycle	cycle	cycle	cycle		31	anti			
	15	cycle					41	cycle			
	16	cycle	cycle	cycle			51	anti			
	17	cycle	cycle	cycle	anti		71	anti	anti		
	19	cycle					91	cycle			
2	21	anti	cycle			2	12	cycle			
	23	anti	anti	anti	cycle		32	anti			
	26	anti	cycle	cycle		3	13	cycle	cycle	cycle	cycle
	27	anti	anti	cycle	cycle		23	anti	anti	anti	cycle
	29	cycle					33	cycle	anti	anti	
3	31	anti					43		anti	anti	
	32	anti					53	anti	anti	cycle	cycle
	33	cycle	anti	anti			63	anti		anti	
	36	cycle	anti	anti			73	anti	anti	anti	cycle
	37	anti	cycle				83	anti	anti	anti	anti
	39	anti					93	cycle	anti	anti	
4	41	cycle				5	15	cycle			
	43		anti	anti		6	16	cycle	cycle	cycle	
5	51	anti					26	anti	cycle	cycle	
	53	anti	anti	cycle	cycle		36	cycle	anti	anti	
	57	cycle	cycle	cycle	cycle		66	cycle			
6	63	anti		anti			76	anti			
	66	cycle					86	cycle	anti		
	67	cycle	anti	cycle			96	anti	anti		
7	71	anti	anti			7	17	cycle	cycle	cycle	anti
	73	anti	anti	anti	cycle		27	anti	anti	cycle	cycle
	76	anti					37	anti	cycle		
	77	anti	anti				57	cycle	cycle	cycle	cycle
8	83	anti	anti	anti	anti		67	cycle	anti	cycle	
	86	cycle	anti				77	anti	anti		
	87	anti	anti	cycle	cycle		87	anti	anti	cycle	cycle
9	91	cycle					97	anti	anti	anti	anti
	93	cycle	anti	anti		9	19	cycle			
	96	anti	anti				29	cycle			
	97	anti	anti	anti	anti		39	anti			
	99	cycle					99	cycle			
Cycle		17	7	9	7						
Anti		20	16	9	3						

Source: Author.

3.4.3 Fronthaul/backhaul

The fronthaul/backhaul dichotomy belongs to the evergreen concepts of shipping. The fronthaul volume is larger and commands full freight rate. Perhaps even more because there is not enough cargo for all the tonnage on the backhaul and that must be compensated for. The backhaul rate is lower because of this imbalance and some vessels prefer to ballast rather than spend time with cargos which may cover only operational costs. So runs the mantra. In practice, many ships prefer to seek return cargos, or any cargos, elsewhere, as noticed already by classical authors. Only those vessels which succeed in finding an acceptable backhaul cargo draw our attention. Could there be a firm relationship between the relative sizes of opposite cargo flows and their relative rate levels? The idea is not entirely blue although whatever relationship there may be it is unlikely to be very firm (Fig. 3.16).

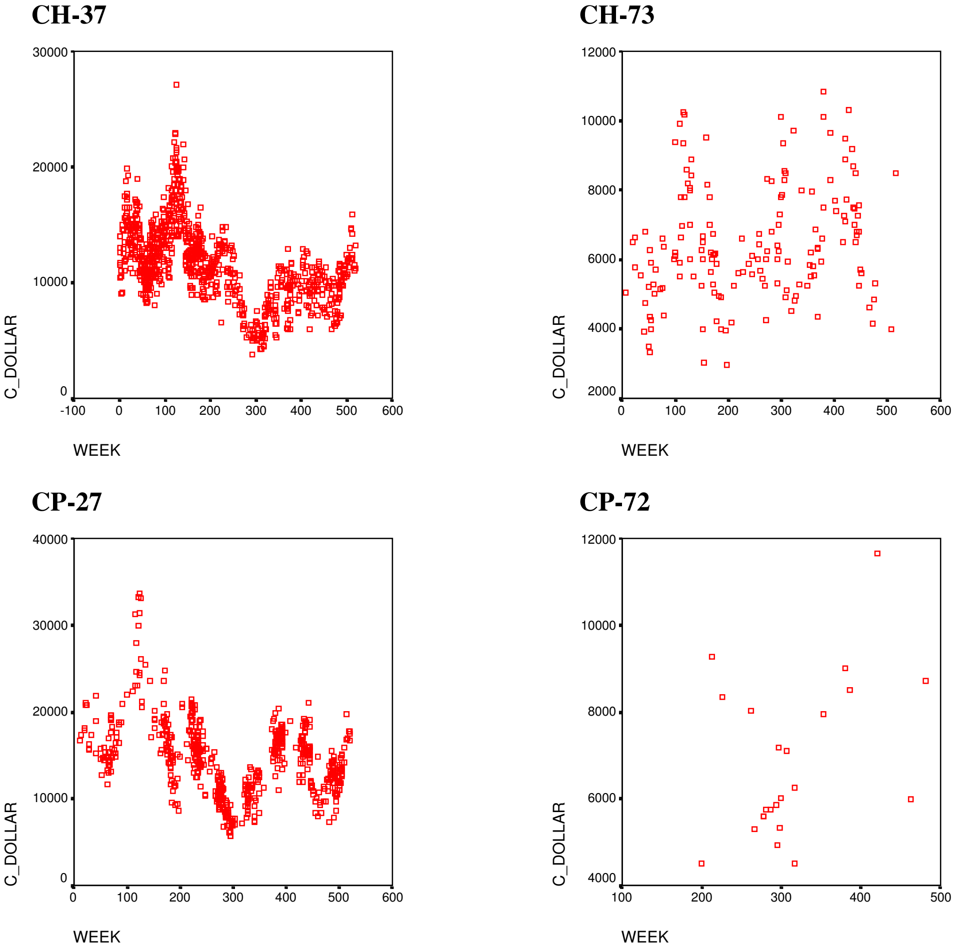

CH-37

CH-73

CP-27

CP-72

Figure 3.16 Fronthaul/backhaul rates, 1993-2002

Source: Author.

The thinking can be rephrased in few statements. A-B and B-A are a fronthaul/backhaul pair. The fronthaul has larger volume and higher rate than the backhaul. When the fronthaul/backhaul ratios of cargo flows and freight rates are compared by route a functional relationship will emerge. Internal flows do not qualify by definition. With a 10 x 10 matrix the theoretical maximum is 45 pairs. The actual figures are much less. Many large fronthauls have no meaningful backhaul and the tonnage must look for employment elsewhere. Our requirement is at least 15 cargo legs a year. It is hardly meaningful to speak of a backhaul if the volume falls below that. Large capesizes do not a have a single pair and small capesizes have only two pairs. Their volumes are modest and the very concept is a non-issue in these size segments. Remain panamax and handy segments. The cargo and rate ratios vary within wide ranges (Appendix 3.14). Particularly the ratio of cargo legs can be uncomfortably large but that cannot be brushed aside. Opposite cargo flows simply are very unbalanced between some regions. More disturbing is that the strict setup A-B and B-A is not always relevant. It happens, for example, that the rates from Regions 1 to 3 are higher than the return rates although the volumes are lower. For the pair CP-37/CP-73 the cargo ratio even turns round without a noteworthy effect on the rate ratio. Such pairs are obviously part of a wider tonnage circulation. The correct reduction base then is Tonnage Balance or their combination. The idea is developed further in a separate paper (Laulajainen 2006). Here, the rigid concept about A-B and B-A is followed to its logical end by running simple regressions of the the fronthaul/backhaul rate ratio with the cargo ratio as an argument. There, indeed, is a meaningful relationship between the two ratios although it is not very intimate (Table 3.25). The small number of pairs must then be noted, of course, and a single maverick pair in a function such as CP-1995 will have a substantial impact.

Table 3.25 Fronthaul/backhaul rate ratios regressed on cargo ratios

Segment	Pairs	R-sqr (adj)	SEE	Coefficient	Interc.
CH-1997	21	0.333	0.218	0.0721 (0.00)	1.205
CP-1995	8	0.051	0.396	0.0160 (0.29)	1.661
	7	0.387	0.332	0.0296 (0.08)	1.615
CP-1997	11	0.371	0.339	0.0435 (0.04)	1.177
CP-1998	10	0.273	0.359	0.0473 (0.07)	1.034

Note: Each backhaul has at least 15 cargo legs. Significance in parentheses.
Source: Appendix 3.14.

4 SIMULATION

4.1 New Simulator

The new Simulator differs from the original version (LHS 2001, 113-145) in two respects: 1) rate estimation and 2) cargo generation when pronounced seasonality or trend exists. In the original study, a new rate was estimated from previous rates by an autoregressive technique, and seasonality in cargo generation was achieved by using a parabola. In the current study, a new freight rate is derived from a linear regression function with the BFI as an argument, and the parabola is replaced by a series of weekly means (below). Both amendments are believed to make the analysis more accurate. A parabola, understandably, is a rather crude estimate of most seasonality. Autoregressive estimation needs an uninrupted time series which in many cases can be achieved only by interpolating missing values and merging routes. The latter is a risky techniqe as will be seen below.

The Simulator needs great amounts of input data (Appendix 4.2):

- Sixteen parameters for counting time, defining the size of the run and the statistical distributions.
- Weekly BFIs to be used as an argument in the rate functions.
- Freight rates currently in use on each route (RTENEW).
- Freight rate functions for updating them (RTEFCN).
- Functions for cargo generation for each route
  - the standard ones for general use (CRGFCN),
  - the special ones for certain routes with a pronounced trend or seasonal pattern (CRGSSN),
- Stock of unallocated cargos by loading region (CRGSTK).
- Record of cargos generated during the past 6 weekdays on each route (CRGLST).
- Record of ships leaving in cargo during the past 6 weekdays on each route (LOADS).
- Record of ships arriving during the past 7 days on each route (ARRIV).
- Norm times for loading a ship by segment and region (LOADTM).
- Norm times for discharging a ship by segment and region (DISCTM).
- Norm times for sailing a cargo leg by segment (SAILCG).
- Norm times for sailing a ballast leg by segment (SAILBL).
- Ship identification, location, current and next activity (SHPLST).

Some of the input are self-evident, others benefit from explanations. Rate functions, cargo generation and the ship list are among the others. All variables describing a state (in contrast to a flow), such as current freight rates, cargo stock, and ship location & activity need intial values. These aspects are taken up in the following Subchapters.

4.2 Freight rate functions

The Simulator needs three items about the freight rate functions, intercept, BFI-coefficient and standard error of estimate (SEE), the last for estimating the random element. It was discovered in Chapter 3 that the overall result is likely to be affected by the way the freight rate functions are estimated. Specifically, an estimate based on the time period 1993-2002 is likely to work against the hypothesis whereas estimates based on two-year periods are likely to support it. Basing the analysis on the two-year estimates consequently appears a kind of qualified cheating? Not really, because these estimates reflect the BFI-index setup, exchange rates and other idiosyncracies at the topical time. Their application brings with it a technical puzzle, however. The particular 2-year periods used here are 1995-6 and 1997-8. About the former, only the 1995 observations are used. During 1995 the weekly BFI-index fluctuated between (about) 1550 and 2350 whereas during 1997-8 the range was between 1500 and 800, meaning that point estimates at a given BFI-value will be higher in 1995 than 1997-8. That need not be incorrect as such because the estimates refer to different times and there is no law or rule which says that actual rates must reflect the BFI index exactly, which is a very general indicator. Fortunately, there is hardly any index overlap between the years and the possibility materializes only during three weeks during 1997-8 and never in 1995.

When there are at least five observations on a route in a year, then the 1995-6 and 1997-8 scatterplots are produced and functions calculated. If the results make sense, i.e., each scatterplot suggests a straight functional line, the coefficient is positive and significant at the 5 pct, occasionally 10 pct, level, the SEE is about the same or lower than in the 1993-2002 function, then these two-year functions will be used instead of the ten-year function. Routes which do not have five observations even during the ten-year period are lumped together into a Rest group and a ten-year function is derived for it. The only requirement for the ten-year functions is that the coefficient is positive. Should that not be the case then the average rate is used. It should be emphasized that although some of this finetuning may have a noticeable impact on the results, most of it will not. In many cases the difference between two-year and ten-year functions is not very substantial and the routes with a small number of rate observations tend to have a small number of cargos, too.

The smallest number of observations for calculating a function, five, is very little for covering a ten-year period, and particularly when the observations happen to accumulate to a part of it or represent different routes. These adverse conditions are ameliorated if the route has been trafficed only during the subperiod from which the observations originate. Such cases do exist but they are not the rule. The release of fixtures to the public domain has preferably been erratic in relation to the traffic. When there is little traffic the ensuing vagueness of the function hardly matters, except that some estimated rates become unrealistically small at low BFI-values (Table 4.1). The technical reason is a large negative intercept term, seemingly outbalanced by a larger positive coefficient. But when the argument becomes smaller the coefficient term declines, while the intercept term remains unchanged. At worst, negative estimates are produced and the simulation is terminated. The problem is solved by eliminating from the underlying data observations (outliers) which lead to the oversized parameters. Specifically:

> Panamax 1997-8 rate function 7 to 8: -24,328 + 25.910 * BFI (SEE = 1,909; six observations) changed to: 1,279 + 6.478 * BFI (SEE = 500; five observtions). CapeL 1995 rate function 3 to 6: -21,773 + 27.799 * BFI (SEE = 4,493; twelve observations) changed to –11,954 + 19.625 * BFI (SEE = 1,715; ten observations). In both cases the old function for 1993-2002 gives very low or negative estimates when BFI declines

below 1,000. These were the only occasions when input data had to be changed to make the simulation technically feasible.

It is difficult to think about a freight rate which a ship operator would consider too high. By contrast, he/she will accept rates which do not cover operating costs only in exceptional circumstances, for example, when an old ship has been fully depreciated and when it is transfered between the Spheres. The Simulator is not constructed to handle such details and it barrs an abnormally low rate simply by introducing an obligatory minimum rate (Table 4.1). This safeguard appears to have come to heavy use in the CapeS segment. The impression is deceptive in so far that the great majority of these routes has very little traffic.

Table 4.1 Bulk carrier operating costs and freight rates

Segment	Dwt	Built	FairRate	OpCost	MinRate	Min 1998
Capesize	160,000	1999	14,500	5,675	4,500	3,412 (1)
Cape small					4,500	2,276 (20)
Panamax	70,000	1992	10,350	4,700	3,500	3,312 (1)
Handymax	40,000	1992	9,250	4,275	3,000	as handysize
Handysize	25,000	1992	8,000	4,000	2,800	3,020 (2)

Notes: FairRate for 12-month time charter. MinRate used in the simulation. Min 1998 refers to the lowest calculated rate (week 293, BFI = 783, without random effect) and on how many routes with some traffic it appeared.

Sources: Drewry Monthly, January 1997, Table 1, p. 5; Appendix 2.4; Author.

4.3 Cargo generation

There are two alternative movement files, Stripped and Final. The Final Files correspond better to the spot market and are therefore preferable. Unfortunately, it proved impossible to estimate the size of tied traffic arriving to Atlantic ports in the capesize segment. It is known that a large share arrived on owned or long-term chartered keels (Jupe et al. 1996, 56-57, 72-75) but because a given vessel seldom frequented the same few ports for longer time periods there was no empirical basis for making the same kind of subtractions as in the Pacific Sphere, Japan in particular. It is better, then, to acknowledge this fact openly and base simulations in the cape segments on the appropriate Stripped Files. Strict comparability with the other segments is thereby lost. The loss is more apparent than real, however, because tied traffic is not as common in the other segments and because short-distance shuttle traffic is usual by small capesizes. It should not be forgotten either that there are some very substantial charterers in the grain trade using panamax and handymax vessels and there is no reliable way to differentiate their long-term chartering from the spot market. In short, it must be acknowledged that the results do not refer to the real spot market but are a kind of vaguely defined “as if” or apparent spot market. Whether this makes any great difference is open to conjecture. Suppose that the tied, by ownership or long-term charter, traffic would disappear, vessels as well as cargos. Would it matter for the rest? Physically, no. Psychologically, probably yes, because the overhang of tied traffic would not be there. But it may be impossible to say in which way the overhang exerts its influence. Its absence would imply a smaller number of observations and greater difficulty of drawing firm conclusions from them. Therefore, apparent or real, whatever emerges cannot be simply wiped away by refering to the partial irrelevance of the spot-market concept.

The cargo generator uses the daily ship arrivals as a proxy for loadings, although for better legibility the following figures are given by week. In many regions and segments the arrivals are abundant and even throughout the year. In some others there may be none for several weeks. When there is no pronounced trend or seasonal pattern it is sufficient to approximate the stream with its daily average (total divided by 312 working days). This is the preferred solution because an average is easy to calculate and give as input to the Simulator. It is the only parameter of the Poisson distribution, widely used for estimating the probability that an event takes place during a time unit, particularly when the probability is small. When the probability grows the Poisson distribution approaches the normal distribution and continues to apply. Whatever its advantages, the average is an approximation only. The fluctuation around it is often considerable in percentage terms. That may be due to the nature of the market and its logistical infrastructure. There are vacation periods and long holidays as seen in the few arrivals at year end. The beginning and end of the time series seem to leave many loadings unrecorded. Incomplete weeks are a partial explanation and the observed loadings may flow normally when a series bridges two consecutive years (Fig. 4.1).

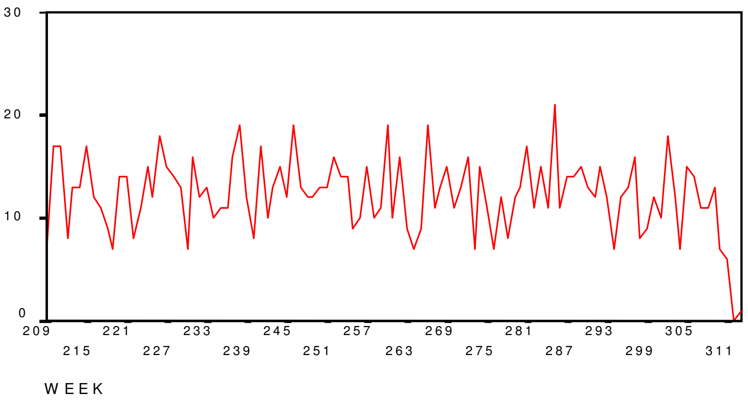

Figure 4.1 Small capesize loadings from Region 8, 1997-8

Source: LMIU Movement Data.

The average is not always applicable, however. There may be a pronounced sloping trend (average is tantamount to a horizontal trend), practically always linear. It may be due to a sustained change in the competitiveness of a ship segment or trading nation, or in the demand for an important commodity. Such changes are seldom rapid enough to be identifiable during a calendar year, our time perspective, but need several years to develop. All the better because then the average may still be an adequate measure of the cargo flow. This is less likely when a severe economic crisis or a healthy boom creates a short-term “trend”, and particularly when it coincides with the calendar year. When a trend is so marked that it cannot be overlooked, it is estimated with the week identification number as an argument, the weekly values are estimated, and subsequently converted to daily figures.

Then there are routes with a heavy seasonal component. The component may have its peak in the northern summer and be convex, or it will occur in the southern summer, be split by the yearend, and appear concave in the frame of a calendar year. Neither the northern peak nor the southern creek need to locate in the middle of the year. A peak may further be split into two because of the vacation period – or some unidentified reason. The seasonal component may lay upon a sloping trend. All that means that a mathematical formula is unlikely to capture the situation very well. The situation is handled by using a 5-week moving average from which the daily cargo figures are derived (Fig. 4.2). The moving average smooths excessive variation and reproduces the major features. The observations lost at the beginning and end of the year (not 1997/8) are replaced by estimates.

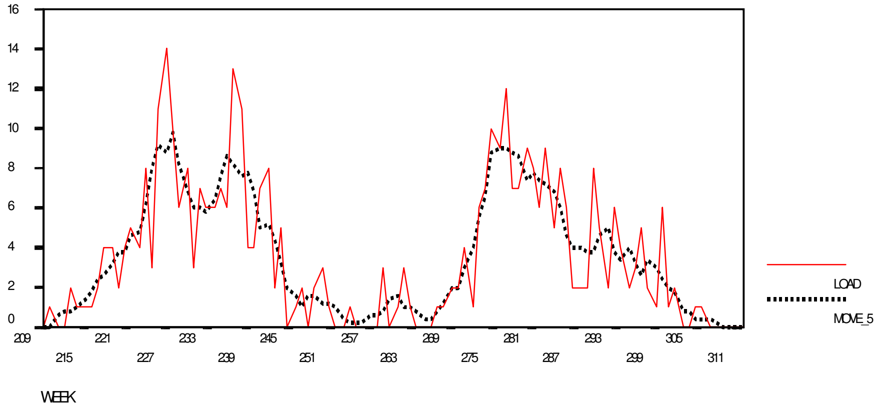

Figure 4.2 Panamax loadings on Route 27, 1997-8

Note: Dotted line 5-week moving average.
Source: Author.

This may sound absurd. The weekly data is smoothed, converted to daily figures, and the Poisson distribution is applied. Would it not be more straightforward to use the raw data as such? In the context of freight rates the variation around a trend is natural because each point estimate is based on a number of observations. Now there is only one observation for each week but it is created by a random process (Fig. 4.1) and represents 6 weekdays which also are subject to a random process. It would be inconsistent to use plain raw data on these routes when it is not used elsewhere. But if consistency is the only problem, why not use raw data overall? Primarily for the sake of saved effort. Each simulation is based on a 10 x 10 trade matrix, 100 cells in all. Even in the large capesize segment, which has the simplest structure, there are 25-30 non-zero cells a year. All the segments and all the years imply roughly 550 non-zero cells with more than 52 weekly observations each, a total of more than 28,000. Most of this workload is avoided by using the average and the linear trend. Only the 5-week moving average requires the complete procedure but such routes are few.

How can these routes be identified? The first step is to look at the aggergate loadings by macroregion. When trends and seasonal components emerge it may be helpful to look at the dischargings, also. Sometimes the topical feature is repeated there, although by no means always. When it is repeated, attention is turned to the corresponding routes. And when it is not repeated, intuition is used. It originates from general familiarity with the topical industries and regions, and has been accumulated during decades of professional experience. There is no shortcut to that kind of knowledge. In practice, only 24 annual routes (out of about 550) were investigated, most of them in the panamax segment. The aggregate loadings simply did not give reason for additional effort. Of these 24 routes, 10 necessitated a linear trend or moving average, 9 of them in the panamax segment. In the following we will take a closer look at them.

A very clear case is the convex trajectory of panamax shipments from Region 2 in 1997-8, with the apex approximately in May-June (Fig. 4.3). A similar but not quite as clear peak appears in the panamax shipments from Region 1 in November. The apexes appear to be connected with the harvest seasons, lagging it by 3-4 months. There is no corresponding peakedness in the main discharging Regions 3 and 7 where the effect of grain shipments is diluted by other cargos (Fig. 4.4). But when the focus is shifted from aggregate loadings and dischargings to individual routes the peaks reappear (Route 17), possibly in an accentuated form (Routes 23 and 27; Fig. 4.5). Smaller peaks can also be surmised on Routes 16 and 26 although there the smaller number of shipments does not allow firm conclusions. But the year 1995 is already different: peaked in Region 2 but not in Region 1. In the handy segment, there is no trace of the harvesting season although its occurrence would appear natural. Where the seasonal pattern in loadings is pronounced it is carried over to arrivals in ballast, in turn moderated by arrivals in cargo (Fig. 4.6). The logic was discussed at some length in Chapter 2 (cf., Appendix 2.6).

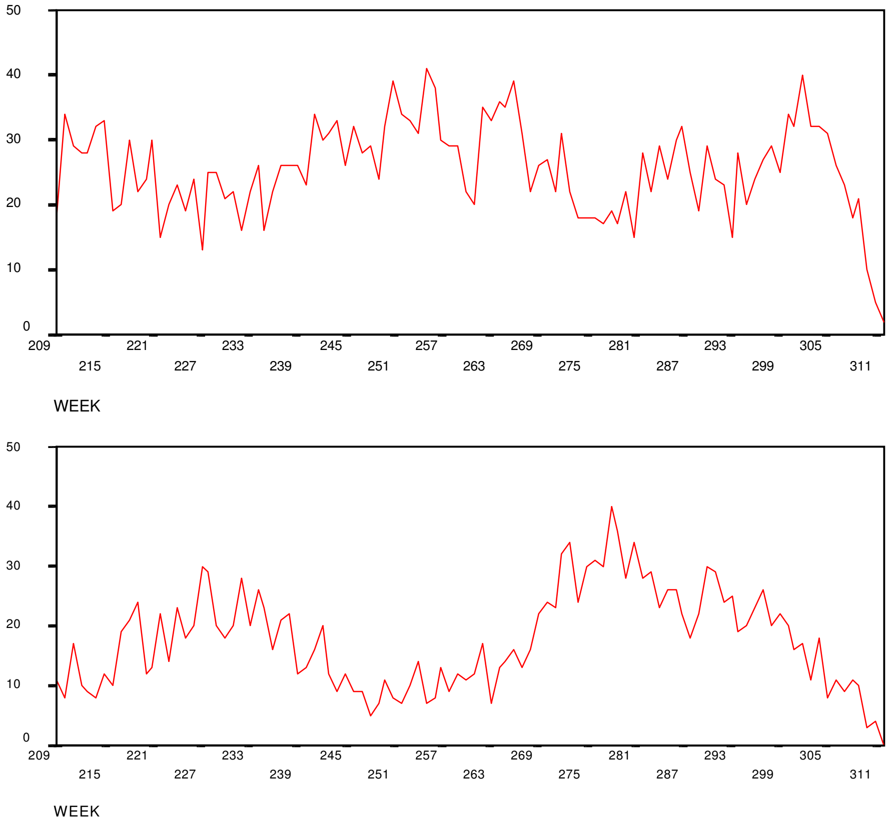

Figure 4.3 Panamax loadings from Regions 1 and 2 (upper and lower), 1997-8

Source: LMIU Movement Data.

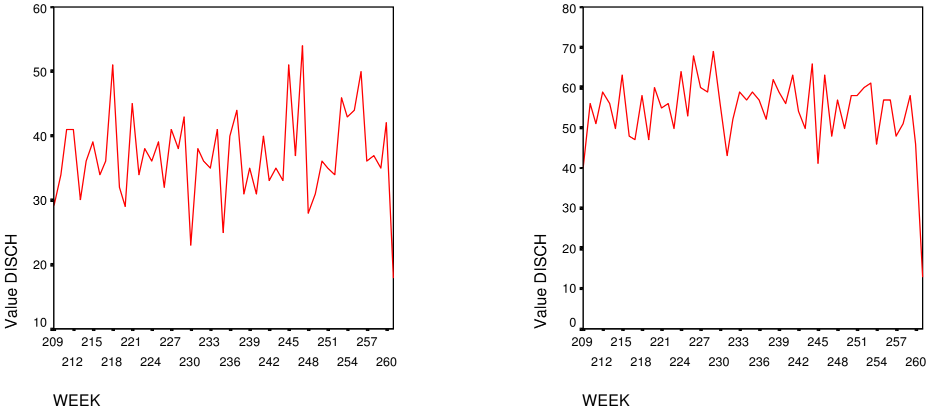

Figure 4.4 Panamax dischargings in Regions 3 and 7 (left and right), 1997

Source: LMIU Movement Data.

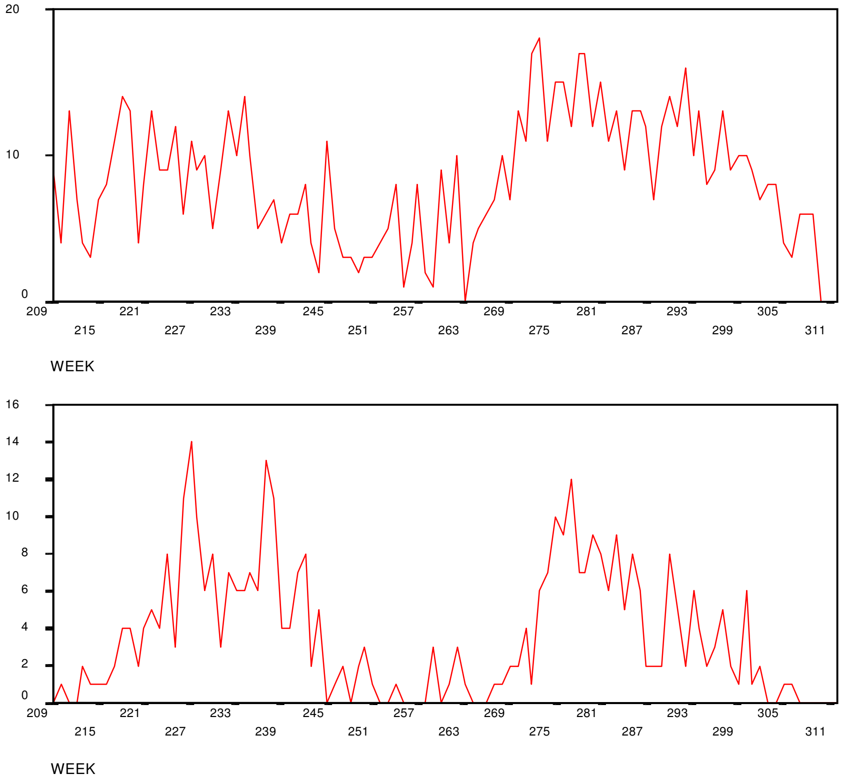

Figure 4.5 Panamax loadings on Routes 23 and 27 (upper and lower), 1997-8

Source: LMIU Movement Data.

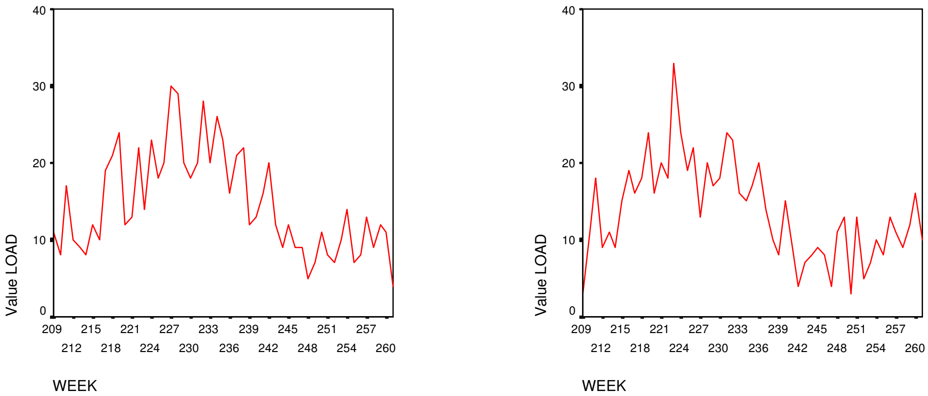

Figure 4.6 Panamax loadings (left) from and ballast arrivals (right) to Region 2, 1997

Note: Arrivals in cargo were on average 4 per week throughout the year.
Source: LMIU Movement Data.

Pronounced trends are rare and what looks like a solid trend may level off in a longer time perspective (Fig. 4.7). And when there appears to be a sustained trend its explanation may prove difficult. Take the small capesize loadings from Region 1, dominated by iron ore and coal and clearly on the decline (Fig. 4.8). Were small capesizes losing ground to more competitive segments, large capesizes in particular, or was consumption declining? Large capesize loadings from the same region did not increase. Perhaps dischargings were declining, most likely in Region 3, the standard small capesize destination. Why there? Either because Region 3 was restructuring purchases or because its consumption was really declining. The facts were that the small capesize flow 13 declined whereas the flows 53 and 83 remained practically unchanged. The large capesize flow 13 increased but that was outbalanced by a similar decrease in the flow 83. The flow 53 increased only slightly. In balance, restructuring took place in the CapeL segment whereas the fall in consumption fell upon the CapeS segment.

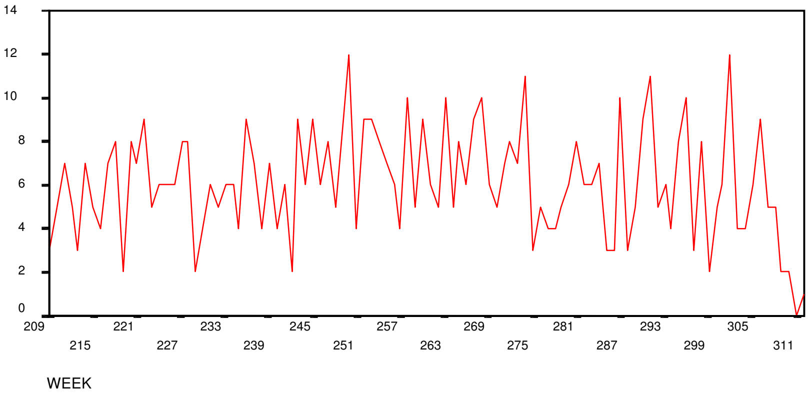

Figure 4.7 Panamax loadings from Region 5, 1997-8

Source: LMIU Movement Data.

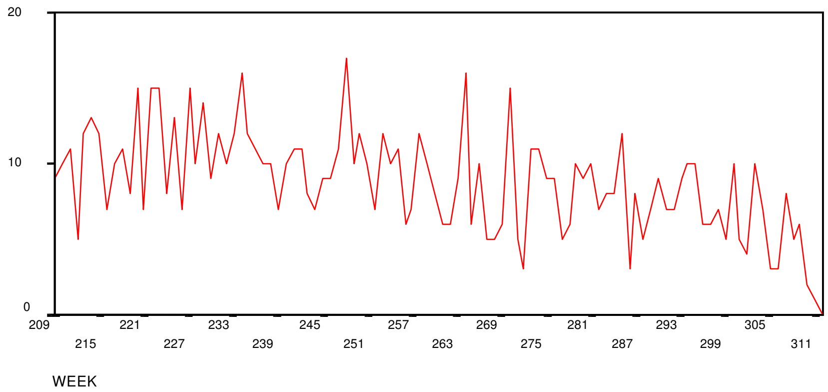

Figure 4.8 Small capesize loadings from Region 1, 1997-8

Source: LMIU Movement Data.

The practical question is when should a trend or seasonal pattern be accounted for. For seasonal patterns the answer is clear. They should be accounted for when they are repetitive, which is not always the case. For trends there is no clearcut answer. There are modest trends which can be visualized in aggregate figures but for which it is difficult to find an explanation. Declining loadings by handysizes from Regions 8 and 9 are examples. Increasing loadings by panamaxes on Routes 21 and 22 during 1997-8 are others. Then there are trends which can be traced back to a particular event, such as the declining handysize and panamax cargos on Route 37 in 1997, due to the economic crisis in Southeast Asia. For practical reasons a trend should be visible within a calendar year which is the time period of simulations, assuming that a calendar year is long enough to allow the observation of a true trend. This invalidates the panamax trends on Routes 21 and 22 and constrains the small capesize trend on Route 13 to year 1998 only.

In general, it makes sense to observe the actual events as closely as practicable as long as reference is made to a particular time period. But when the purpose is to explain principles (“geographical fundamentals”) temporal idiosyncracies create white noise and it may be better to smooth over them. The Revenue Gradient is not very sensitive to the volume of trade flows but reacts to changes in freight rates (Appendix 1.1). Therefore we have paid comparatively much attention to freight rates and feel also now justified to pay less attention to trade flows. Only a few routes are considered worth the extra workload of non-constant flows and the choice falls mostly on a seasonal pattern (Table 4.2). The “common” feature is that there is no genuine common feature. If there is a pronounced temporal pattern in one year one cannot be certain that the same type of pattern will develop in the next year. That, obviously, raises question marks about the validity of any generalization.

To keep the Simulator versions at a minimum the input data contains a mock route of plain zeros in such cases when a seasonal pattern is not warranted.

Table 4.2 Routes with pronounced non-constant cargo flow

Route		1995	1997	1998
Panamax	13			decline
	17		season	season
	23	season	decline	season
	27	season	season	season
CapeS	13			decline

Source: Author.

In the no-man’s territory between cargo generation and ship list are LOADS and ARRIV. Loads is the number of vessels departing in cargo during the past 6 weekdays and Arrivals comprises all arriving vessels whether in cargo or ballast [LHS 2001, 246 spelled “all arrivals” incorrectly as “arrivals in cargo” but the Simulator followed the current practice]. The initial totals equal the annual averages, rounded to the nearest integer, and are subsequently allocated to the contributing days with the help of random numbers from a uniform distribution (www.mrs.umn.edu). Their ratio is Tonnage Balance, an extremely important element for evaluating the attractivity of cargos, because it provides a recent history of demand and supply.

4.4 Ship list

The first thing in constructing the SHPLST is to decide the number of vessels. This is based on the time needed for transporting all cargos from the loading ports to discharging ports, all attached cargo handling, and the necessary ballasting to new loading ports. Some slack, in the form of waiting, is also needed because the cargo generation is subject to variation. The calculations can be refined by the experience gained in the original study (LHS 2001, Table 5.8).

The key parameters of the original study were: speed 11.5 knots (275 nm a day), all loading and discharging 5 days. There were a total of 5,747 cargos to be transported within a year (LHS 2001, Table 5.2). It was estimated that 700 vessels would be enough to do the job and that initially 38 of them (5 pct) would be uncommited, waiting for a cargo. It turned out that the number of vessels was insufficient and at the end of the simulation, including a 42-day warm-up period, 746 cargos or 13 pct of those generated remained unallocated (LHS 2001, 148). The dilemma was solved by reducing the number of cargos by 10 pct to 5,172, an easier solution than creating some 70-100 additional vessels. That gave an average figure of 7.4 cargos per vessel.

The numbers of the current study for the corresponding data set, panamax 1997, are slightly different: speed 12.5 knots (300 nm a day), and loading and discharging differentiated by macroregion with global averages 5 and 8 days, respectively. The greater speed and the slower discharging broadly outbalance each other. The number of cargos is 6,289, a 9 pct increase. The simple arithmetic then is that with the new parameters 850 vessels are needed (6,289 divided by 7.4). It is 16 pct more than if the need were based on the time spent for sailing and cargo handling only, i.e., without waiting. This percentage is important because we have no previous study to be used as a reduction base for the other ship segments and years. The relevant data are consolidated in Table 4.3.

Table 4.3 Calculatory time budgets (days) and the number of vessels in the simulations

		Cargo	Ballast	Load	Disch	Sea	Port	Total	Ships	+16 %	Simul.
CapeL	1995	31,069	21,591	4,559	7,074	52,660	11,633	64,293	177	205	210
	1997	45,204	29,756	6,547	10,028	74,960	16,575	91,535	251	292	300
	1998	42,405	30,372	6,518	10,056	72,777	16,574	89,351	245	285	290
CapeS	1995	36,503	27,665	7,790	11,827	64,168	19,617	83,785	230	267	270
	1997	34,240	27,584	7,277	11,271	61,824	18,548	80,372	221	256	260
	1998	29,665	23,074	6,247	9,825	52,739	16,072	68,811	189	219	220
Panamax	1995	102,365	62,993	27,011	43,946	165,358	70,957	236,315	649	753	750
	1997	117,966	66,024	30,681	51,835	183,990	82,516	266,506	732	849	850
	1998	110,325	58,539	30,798	48,291	168,864	79,089	247,953	681	790	800
Handy	1997	210,783	111,116	104,479	191,459	321,899	295,938	617,837	1,697	1,969	2000

Note: Ships = Total/364.
Source: Author.

The continuation is tailored with a view of minimizing the workload, a relevant aspect when there are almost 6,000 vessels to be handled. The procedure is based on the assumption that the annual averages will give an adequate estimate of the intial state of the system. The alternative would be to take the first move of each vessel in the appropriate movement file until the number needed in the simulation is reached, and work these moves up to a usable ShipList file. An experiment was made with large capesizes in 1995. The ships emerged from the data file at a slower pace than expected and involved much manual work. During weeks 104-111 (until 19 February 1995) the screening of 387 moves identified 202 vessels when 210 were needed in the simulation. It is doubtful whether mid-February still qualifies as the beginning of the year and whether the manual effort is justified.

The choice, therefore, is for the average conditions. The first step is to allocate the vessels between activities. It is assumed that there will be 5 pct of uncommited vessels, due to the uneven generation of cargos and possibly oversupply of vessels. It must be admitted, however, that the number of vessels does not tally with the freight market: tight in 1995 and easing towards 1998 (Table 4.4). The only sensible explanation is extensive nonreporting of vessel movements (but not vessels themselves) in 1995 (cf., Tables 2.1 and 2.2).

Table 4.4 Number of vessels, 1995, 1997, 1998

Segment	File	Observed			Calculated			Obs/Calc		
		1995	1997	1998	1995	1997	1998	1995	1997	1998
CapeL	Strip	236	294	292	205	292	285	1.15	1.01	1.02
CapeS	Strip	337	291	275	267	256	219	1.26	1.14	1.26
Panamax	Final	912	983	994	753	849	790	1.21	1.16	1.26
Handy	Final		2,108			1,969			1.07	

Note: LMIU Vessel Data 1997-8 gives: 330 large capesizes; 362 small capesizes; 1,150 panamaxes; 2,663 handysizes. The discrepancy is due to gaps in reporting, docking and inactivity.
Sources: LMIU Movement Data; Table 4.3.

The remaining 95 pct are divided between cargoing, ballasting, loading and discharging in proportion to the times actually used for the respective activity (Table 4.3). When that has been decided each activity is divided between regions and routes in proportion of the number of legs in the appropriate matrix cells. Loadings and dischargings are taken from the cargo matrix, its row and column totals, respectively. The weakness of this deterministic allocation is that regions and routes with a low level of activity loose the chance to be represented, as they would have if a random element were included.

Ships arriving to a region in ballast to load and ships already loading there will have discharging regions in the average proportion of that region. About discharging ships no such information is needed because they will become available only when the activity is finished.

The distances differ by ship segment because of the size restrictions of the Panama and Suez Canals and the different points of gravity of traffic (Fig. 1.6). In addition, the distances in cargo and ballast are different for large capesizes which cannot pass the Suez Canal fully loaded. In all cases, the average speed of 12.5 knots (300 nm a day) is used to arrive at sailing times.

Loading and discharging times depend on the commodity setup and the technical standard of ports. The third, potentially relevant factor, the number of loading and discharging ports visited during a trip, hardly applies because the average figure does not vary too much between regions (Table 3.14). No effort is made to estimate the attached red tape. The regional (3-digit) figures of Table 3.10 in 1997 are converted to macroregional figures by weighting them by regional (3-digit) loadings and dischargings and then rounding off to the closest full day. Where weights are lacking (capesize dischargings in Regions 4 and 10) the estimates are taken directly from Table 3.10. Capesize loadtime in Region 3 is changed manually for the simulation from 9 to 4 days.

Because all initial values for times represent a snapshot of a time continuum they reflect two things: the total time of the activity and the share still left. It has been established that the total times follow a modified gamma distribution (LHS 2001, 137-140). The time still left is a function of the random sample taken from that distribution and the time so far spent in the activity (sailing, loading, discharging). Both are unknown in the beginning of the simulation. Two things are known, however. First, the Random (estimated) Time (RT) is at most five days longer than the Norm Time (NT) based on 11.5 (here 12.5) knots' travel speed and the empirical loading and discharging times. This rule-of-thumb is not sensitive to the length of NT and it gives an upper limit which no intial value can exceed. Second, the percentage of RT below the NT is smaller than the percentage above it. This matters if all the vessels depart simultaneously, which they do not. Without knowledge of their departure times the rational assumption is that they depart at constant intervals (Sunday is then overlooked). Imposing gamma distributed RTs upon a steady flow of departures appears to lead to another steady flow, i.e. uniform distribution. In short, it is difficult to get satisfactory randomness into the RT, particularly because several activities and routes have the same NT. There is a danger that short, alternatively long, RTs will accumulate in a certain activity/route (periodicity). That can hardly be completely avoided except by resorting to laborious random number generation. The practical solution is that for a constant Norm Time there will be some intial times which exceed it and the rest is distributed evenly between all the appropriate ships. The result can be as in Figure 4.9. Whatever the shortcomings, their effect can be ameliorated by a long warm-up period in the beginning of the simulation. The 42 days of the original study will be extended to 63 days, i.e from six to nine weeks, out of the total of 427 days. This also exceeds comfortably the longest NT 49 days.

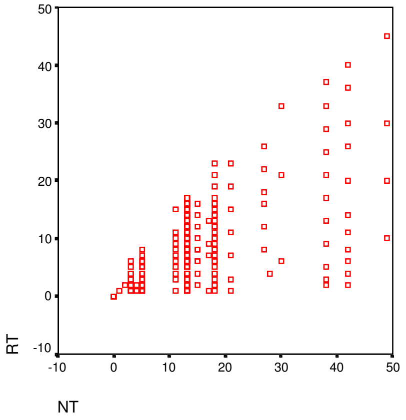

Figure 4.9 Random Time as function of Norm Time, CapeL-95

Source: Author.

4.5 Using the Simulator

The Simulator exists basically in only one version. Most parameters can be handled with the help of input data. Variables in the dimension statements (REAL and INTEGER) are an exception because dimensions can be given to the Fortran compiler only in numeric form. This applies also to temporary (auxiliary) variables and their loops such as PICK, ALLOR, SHPAL and SHPALC, which are not defined in the input statements. Too large a dimension should not do any harm because the program sets a vector or matrix to zeros before anything is stored in it and handles only those parts which contain relevant data. The input and output statements, instead, should always be scaled to the size of the problem at hand.

A capesize run with 260 ships took 1-2 minutes, a panamax run with 800 ships about 5 minutes, and a handysize run with 2,000 ships about 69 minutes, all in the Linux operative environment with a 350 Mhz, 3 GB Dell PC.

The Simulator had been tested carefully in the original project (LHS 2001, 124-142) and the amendments were comparatively minor. First, the autoregressive freight rate generation module was modified to handle a simple linear function with BFI as the argument. Second, the modules specifying parabolic cargo generation functions were replaced by a more general module which accepts any number of routes with weekly cargo generation functions (averages). Third, the number of macroregions was reduced from 11 to 10. Fourth, the sailing time was disaggregated into cargoing and ballasting with appropriate distance matrices.

The simulation was monitored by day. The following variables, vectors and matrices were recorded: day, weekday and week, cargos generated (daily), cargos left unallocated (= overflow; daily), tonnage balance (= loadings/arrivals averaged over 6-days; daily), freight rates (weekly), ship diaries (after warm-up; daily), ship revenue (after warm-up; annual). There were also remarks when the cargo inventory or fleet got exhausted and when a cargo was beyond the economic reach of uncommitted (free) vessels. All this created a large number of output lines, almost 100,000 in the Handy segment.

The general impression was that the simulation reproduced the real world adequately, i.e., in a way which was unlikely to distort the purpose of the study. Remarks about exhausted cargo inventory and fleet appeared occasionally. Cargos beyond the economic reach were more frequent, even weekly, events. There is nothing dramatic in them. A within-region task (short distance, low rate) and distant uncommitted vessels are the ingredients. Cargos might also remain unallocated in larger numbers than appeared realistic. In the Real World that would lead to higher rates, greater sailing speeds, and postponing of yard visits. The Simulator lacks appropriate modules and the practicable way was to reduce the number of cargos generated (see below). Unrealistically low rates were controlled by refined rate functions and minimum rates (see above). Rates without the random effect undercut the minimum rate mostly in the small capesize segment in late 1998 and on routes with little traffic. The effect of the random element was not studied in detail. It was considered more fruitful to calculate the rates without the random element and prepare appropriate line graphs. These behaved well, meaning that the trajectories of the main routes did not cross each other in a major way or at all. In other words, the rate structure did not change which also appears realistic in a year's perspective.

The number of cargos generated had been subjected to a detailed analysis in the original report with satisfactory results (LHS 2001, Appendix 6.3). It was considered prudent, however, to conduct a new check and O-D cargo matrices were produced from the simulated itineraries and compared with corresponding input matrices. The number of cargos generated tended to exceed the input but not seriously (Table 4.5, left field). This was due to the random effect as such but also the fact that the input was given by matrix cell which necessitated some rounding off. For a better overview, the simulated matrices were calibrated by a multiplier so that their grand totals equalled the input matrices (Appendix 4.3). They were also subjected to a χ^2-test which disclosed two conversion errors (from xls-file to txt-file) in the input data. After appropriate corrections and reruns all matrices passed the χ^2-test at the 5 pct level of significance (Table 4.5). There are formally 99 degrees of freedom but because many cells have no or few cargos only those cell-pairs (simulated & input) were considered whose combined frequency is at least 10.

Table 4.5 χ^2-test (0.05) of simulated cargo legs

Segment	Simul/Observ (pct)			χ^2			5-pct level		χ^2
	95	97	98	95	97	98	χ^2	df > 9	df 99
CapeL	105.5	106.7	104.4	16.08	13.52	14.58	28.87	18	123,23
CapeS	93.0	93.9	95.3	28.01	18.17	23.84	37.65	25	123,23
Panamax	104.8	105.3	107.2	41.79	29.39	37.99	75.62	57	123,23
Handy		103.4			45.03		105.27	83	123,23

Notes: Cargo generation in the CapeS segment scaled to 90 pct of the nominal input.
Source: Appendices 2.4, 2.5 and 4.3; Author.

Thereafter only one problem remained, the overflow of cargos in 1998. The size of the overflow is measured by the number of unallocated cargos on the last simulation day because the accumulation is gradual rather than a sudden outburst. The apparent reason is a declining freight rate level which pushes more cargos beyond the economic reach of uncommitted ships. The overflow is compared with the average daily cargo generation. The unallocated figure mostly stays well below the daily generation, except at small capesizes in 1998 where it approaches two daily generations (Table 4.6). That happens after the cargo generation parameters have been reduced by 10 pct (cargo coefficient). Why are the small capesizes different from the rest? The combined effect of rounding error and the distribution of small frequencies over several possible routes is a plausible explanation.

Table 4.6 Overflow in a nutshell

Segment	Cargo coeff.			Unallocated Day 426			Cargos/day			Surpl. ships %		
	95	97	98	95	97	98	95	97	98	95	97	98
CapeL	1.00	1.00	1.00	1	1	5	4.46	6.37	6.35	18.6	19.5	18.4
CapeS	0.90	0.90	0.90	1	2	10	7.08	6.69	5.80	17.4	17.6	16.4
Panamax	1.00	1.00	1.00	1	2	9	17.75	20.16	19.71	15.6	16.1	17.5
Handy		1.00			5			54.15			17.9	

Source: Appendices 2.2 and 2.3; Table 4.3; Author.

It appears that the overflow can be easily corrected by adding ships, say, 5 at a time and observing the effect. It is not quite that simple, as some experiments showed. The CapeS segment was given five additional ships in 1998. Unexpectedly, the overflow increased from 10 to 16. Then the cargo coefficient was raised from 0.90 to 1.00 and the overflow, logically, increased further but only to 17. The total sailing time needed for their delivery followed suit: 138, 208 and 303 days, respectively. In short, the experiment led nowhere. It also raised the question whether still more ships would contribute to the research effort. A slight overflow as now gives a tangible feeling of the workings of the Simulator. Without it, some other indicator would be desirable, the number of uncommitted ships, for example. Such balancing is unessential for the purpose of the study, the comparison of the macroregions as first loading regions. Their mutual position within a ship segment and year is unlikely to get affected by a modest overflow, perhaps connected with a similar reduction of the cargostock. One might preferably go in the opposite direction, to double the size of the cargostock and the fleet, and by so doing increase the number of observations also for lightly trafficed macroregions. The conclusion is that no amendments are necessary and the simulation results are accepted at face value.

4.6 Results

The formal task of the study is to find out whether there is a close functional relationship between the Daily Revenue (DayRev) and Revenue Potential (RevPot or RP), as discovered in the original study (LHS 2001, Fig. 6.7). The Daily Revenue is calculated for each ship separately whereafter it is averaged over all the ships which happen to have their first loading port in a particular macroregion. Since there are 10 macroregions there will, theoretically, be each year 10 observations for each ship segment. The Revenue Potential is calculated for these same regions.

To achieve a decent comparability between ships, the first loading region is defined as the first region in which it is possible to record a loading day as from the 63rd simulation day onwards. It means that for most ships there will be several days in the beginning of the recorded simulation period which are not accounted for the daily revenue. They are not revenue-days. Whether loading is the best activity to trigger off the accounting period can be discussed. The preceding ballast leg would be at least as good because it starts the standard cargo cycle: chartering-ballasting-loading-cargoing-discharging. The accounting period ends with the end of the simulation. It would be logical and, indeed, more accurate to end the accounting period with some completed activity. Programming that refinement would be comparatively complicated and is overlooked. Because cargo cycles differ in length and begin randomly it would be impossible anyway to achieve full comparability. The accounting period thus comprises both types of day: when a vessel collects revenue - during loading, cargoing and discharging, and when it does not – during ballasting and waiting. Ballasting and waiting are not excluded because the former bridges the geographical and the latter the temporal gap between cargos.

Ballasting is central to the whole philosophy. It is well known that the daily freight rate varies widely between routes and that this variation is intimately connected with the demand and supply of tonnage at the loading end. When demand exceeds "natural" supply, i.e. tonnage discharging there, ballasting from surplus regions becomes necessary. This raises the rate, either directly or through a ballast bonus. The opposite applies when demand falls short of supply and discharging tonnage must ballast to deficit regions. Example: the panamax rate on Route 13 is much higher than in the opposite direction. But the logic does not end there and the possibility to get a new cargo at the discharging end of a cargo leg also affects the rate. If possibilities there are slim ballasting becomes necessary which depresses the rate, and the other way round. Example: the panamax rate on Route 73 is lower than on Route 79.

The demand/supply balance (Tonnage Balance) at the loading end is easy to calculate since it comprises only four aggregated flows from 9 other regions: arrivals in cargo, arrivals in ballast, departures in cargo and departures in ballast. At the discharging end, the corresponding four flows apply but they need to be calculated for each of the 9 regions because the classification principle (first loading region) does not specify the discharging region. The calculations are considered excessive here and the discharging end is overlooked. However, if the analysis is sharpened to routes instead of loading regions extended calculations cannot be avoided. We will return to the topic in another context (Laulajainen 2006).

There is also the time dimension, easily overlooked. It has at least two angles. First, the length of a cargo leg affects the revenue collected. If the freight rate is high, a long trip is welcome. But if the rate is low, a long leg may imply a high opportunity cost. Unfortunately, what is high and what is low cannot be answered in a general way. Only one thing is certain, a poor cargo is normally better than no cargo at all. But as the Simulator is run, with a minimum number of vessels as input (Table 4.3), all the ships are fully occupied most of time. It boils down to the conclusion that long legs are desirable when revenue is collected and undesirable when it is not. The second angle is that an averaged Tonnage Balance is a good proxy only if its elements are not too different in size. That is most likely when the number of arrivals and departures is large. When either number is very small the Tonnage Balance is likely to vary wildly (Fig. 4.10). This variation is smoothed out in the annual average.

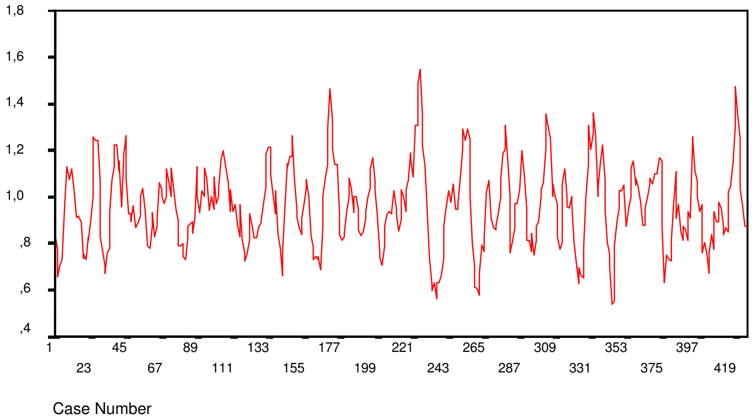

Figure 4.10 Panamax daily Tonnage Balance in Region 9, 1997

Note: Annual tonnage balance = 0.956.
Source: Simulation.

The discussion leads to the following form of the Revenue Potential (RevPot, RP), thus named because it gives the geographical frames to a certain level of revenue:

RP-i = ∑ [(DC-i * ST-ij – DB-i * S-ij) / (AC-i * ST-hi + AB-i * ST-hi)],
totalled over all trading connections, in which:
RP-i = Revenue Potential
DC-i = Departures in Cargo
DB-i = Departures in Ballast (technical variable)
AC-i = Arrivals in Cargo
AB-i = Arrivals in Ballast
ST-ij and ST-hi = Sailing Time between regions
h, i, j = region identifiers.

DC reflects tonnage demand and the total of AC and AB reflects tonnage supply. DB is actually unnecessary and can be omitted. Its only role is to widen the range of RP and facilitate visual understanding. RP with DB (RP-1) and without DB (RP-2) are closely correlated, a scatterplot displaying a slightly concave rising trajectory. Loading and discharging times are omitted for simplicity. Straight distance, without decay, is used because the freight rate remains unchanged throughout the leg.

Since the implication is that RevPot will somehow explain the variation in DayRev it is appropriate to remind that the cargo allocation module of the Simulator uses Tonnage Balance, a close cousing of RevPot, as a weight for available cargos when bidding for them. Sailing Time is also included. The bidding function tries to anticipate what might happen when the immediate task has been completed by considering two conseqtutive legs rather than just one. Does not all this lead to a kind of circular reasoning, Tonnage Balance and Sailing Time explaining themselves? Not directly. Bidding was done daily for simulated, not observed, cargos and a daily tonnage balance can be a far cry from the annual average. Freight rates applied only to cargo legs and the attempt to account for the attached ballast legs was just that, an attempt, because at the bidding stage a ballast leg following rather than preceding a cargo leg is still hypothetical. Ballast legs included in the RP-formula are, by contrast, empirical although at the aggregate level. There is also the practical aspect that annual data can be approximated in advance and estimates prepared from it directly, without cumbersome simulation. If the cargo allocation mechanism of the Simulator really reproduces the actual market behavior then, by using the RP-formula, considerable savings in research effort can be made.

At this stage it is useful to compare the original study with the current one (Table 4.7; Fig. 4.11). The comparison is not too favorable, recalling that the same person is responsible for the both sets of result. The main difference is in the way the freight rates were/are calculated. Previously they were autoregressive time series because the original purpose was to construct a forecasting and decision-making model for ship managers. Now the rates are based on real-time, or almost, exchange quotations which is a far more efficient way of using available data.

Table 4.7 Input for Figure 4.11

Region	Original			Current			Atl
	Freq	DayRev	RP-1	Frq	DayRev	RP-1	
1	182	8,426	1.368	192	9,695	1.266	1
2	94	8,453	1.016	85	9,669	1.028	1
3	59	7,878	-0.295	44	9,318	-0.367	1
4	18	8,254	1.272	20	9,343	1.123	1
5	30	8,156	1.084	35	9,226	1.003	1
6	31	7,251	-0.043	41	8,867	-0.108	0
7	75	7,269	-0.540	120	8,968	-0.460	0
8	117	7,581	1.262	156	8,969	1.444	0
9	92	7,661	1.266	154	9,266	1.100	0
Total	698			847			

Note: Too few observations in Regions 10 and 11.
Sources: LHS 2001, Appendix 6.4; Author.

Specifically, some routes were consolidated in the original study in an inappropriate way (LHS 2001, Tables 4.13 and 5.5). Consolidation could not be avoided because an autoregressive model needs an observation for each time unit (week) and that observation was usually not available on lightly trafficed routes. An interpolated observation was a stopgap measure and could not be used extensively (LHS 2001, Fig. 4.11). Consolidation was the only way. Its riskiness was well understood but the technicalities of autoregression were too strong. The approximate effects (current data) are visualized in Figure 4.12 (p. 126). For example, Kamsar (Region 4) had been consolidated with ECNA and ECSA (Regions 1 and 2) before simulation which gave it $200 lower DayRev when $300-350 is closer to reality. Australia (Region 8) had been consolidated with WCNA (Region 9) which gave it only $80 lower DayRev when $300 is close to reality. SAF (Region 5), a borderline case, locates better halfway between the Atlantic and Pacific. When such inaccuracies are corrected the beautiful split into Atlantic and Pacific loadings gets destroyed and an intermediate zone is created. Figure 4.12 discloses also another shortcoming in measurement and it applies to the current study as much as the original one. Revenue Potential is calculated for the whole year and is in that perspective constant. Freight rates, by contrast, change continuously (weekly) influencing cargo allocation, and it is this continuum which is sampled several times a year to arrive at the Daily Revenue.

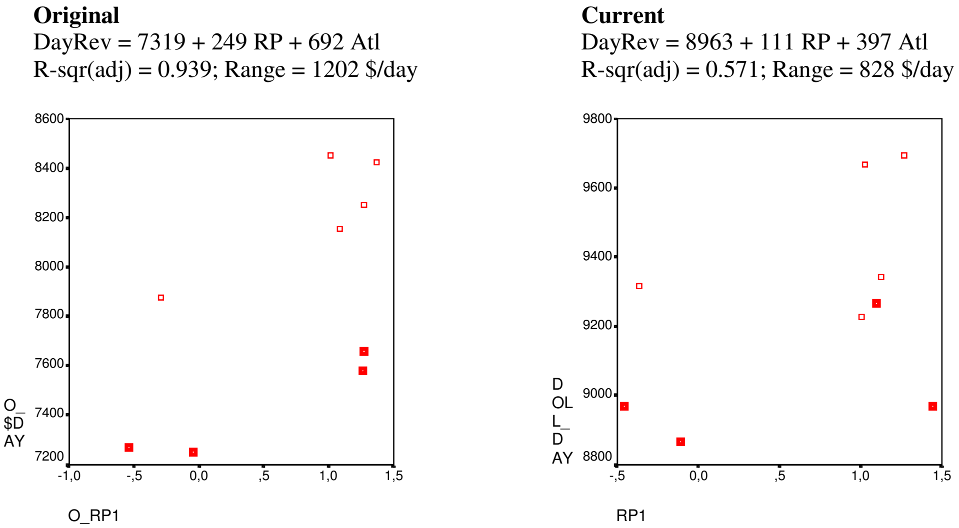

Original
DayRev = 7319 + 249 RP + 692 Atl
R-sqr(adj) = 0.939; Range = 1202 $/day

Current
DayRev = 8963 + 111 RP + 397 Atl
R-sqr(adj) = 0.571; Range = 828 $/day

Figure 4.11 Panamax Daily Revenue as function of Revenue Potential, 1997

Legend: Small marker = Atlantic; large marker = Pacific.
Source: Table 4.7; Laulajainen 2007, Fig. 2. Courtesy of Elsevier Ltd.

Do these observations mean that the original result was seriously wrong? To find out, the analysis is repeated for each segment and year for which data is available. Only the numerical results are given but they are telling enough by their inconsistency (Table 4.8). To iterate the most negative features only, low R-sqrs or F-statistics, insignificant coefficients and unstable signs. These facts can be partially attributed to the small number of macroregions, between 7 and 10. The standard solution then is to pool the data and possibly look for additional explanatory variables.

Table 4.8 Daily Revenue regressed on Revenue Potential and Atlantic/Pacific dummy

Segment	Year	Obs	R-sqr (adj)	F-stat	SEE	Coeff RP-1	Coeff Atl	Signif RP-1	Signif Atl	Interc
CapeL	1995	7	0.977	127	119	304	1325	0.03	0.00	19,527
	1997	8	0.045	1.2	718	386	508	0.28	0.38	12,855
	1998	7	-0.031	0.9	559	-2	576	1.00	0.28	8,532
CapeS	1995	9	0.281	2.6	646	529	359	0.12	0.45	14,407
	1997	8	0.458	4.0	249	312	263	0.07	0.21	9,091
	1998	9	-0.209	0.3	596	248	-28	0.47	0.95	6,190
Panamax	1995	9	0.903	38.3	174	290	845	0.02	0.00	14,401
	1997	9	0.571	6.3	193	111	397	0.28	0.02	8,963
	1998	10	0.513	5.8	95	126	95	0.02	0.17	6,036
Handy	1997	10	0.654	9.5	98	210	151	0.03	0.06	6,777

Note: Only macroregions with at least 4 ships.
Source: Appendix 4.5; Laulajainen 2007, Table 1. Courtesy of Elsevier Ltd.

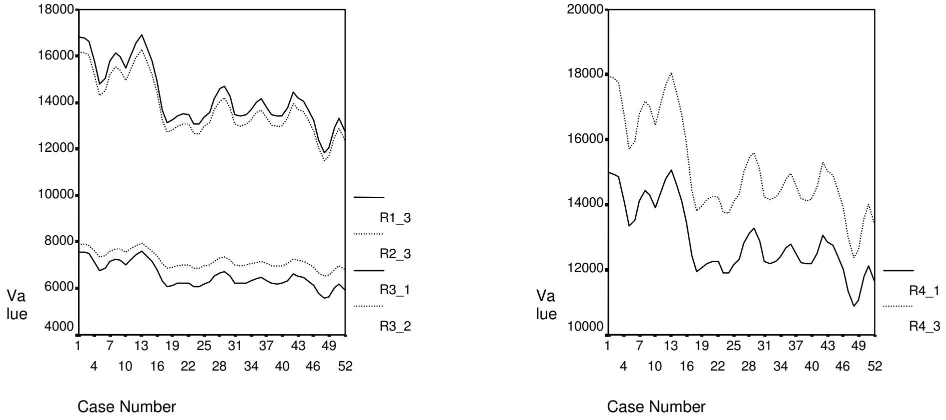

Appropriate
ECNA/ECSA – Continent, round trip

Inappropriate
Kamsar – ECNA/Continent

Australia/WCNA – Far East

Australia/WCNA - Continent

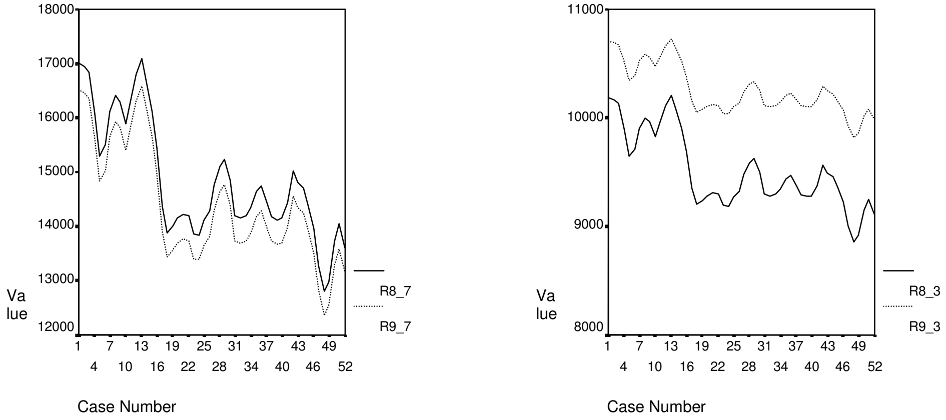

Figure 4.12 Examples of route consolidation in the original study

Legend: Dotted lines: upper figure ECSA (left), CONT (right); lower figure WCNA.
Note: Based on the rate functions (without random effect) used in the simulation.
Sources: LHS 2001, Table 4.13; Author.

The first step is to consolidate all segments and years into one data set. This decision brings with it the introduction of five new dummies (Atlantic exists already), to account for the rate differences between segments and years (Fig. 4.13). The panamax segment and the year 1997 are chosen as bases which leads to the following dummies: D-CL, D-CS, D-H, D-95 and D-98. It is known further that an increasing vessel size leads to a higher albeit diminishing freight rate whereas increasing age leads to a lower rate. This information is based on fixtures and may not hold for aggregate data (Tables 3.12 and 3.13). It is worth a try, however, and the variables Age and Dwt are added. The new angle makes the results quite palatable although not yet fully satisfactory (Table 4.9, Full size).

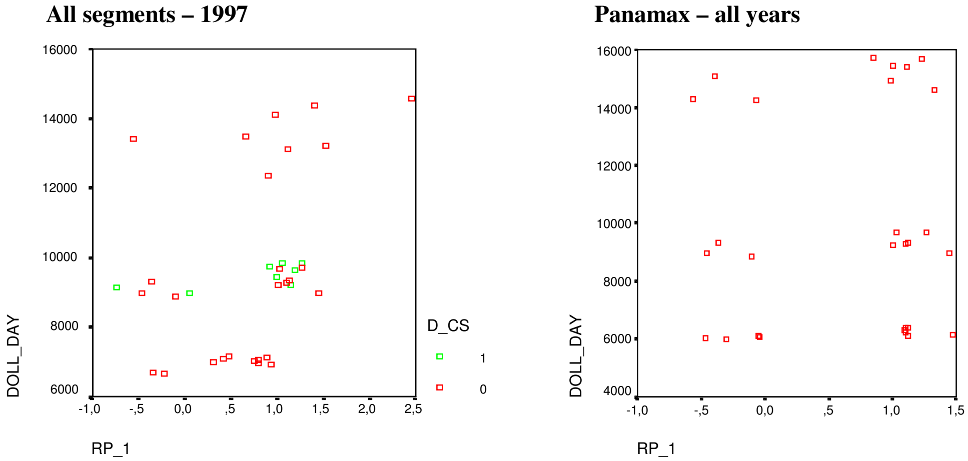

All segments – 1997

Panamax – all years

Figure 4.13 Examples of Daily Revenue by segment and year

Note: Small capesizes indicated by green (light) color.
Source: Appendix 4.5.

Specifically, the capesize segment dummies are either insignificant or have an unconventional (negative) sign. They and the Dwt-variable obviously cannot coexist in the same model. It would make sense to exclude the offending dummies and keep the variable measured on ratio scale but the significance of either Revenue Potential, the geographically most interesting variable, or Dwt itself would be destroyed. The CL-dummy is consequently kept but the CS-dummy is excluded because small capesizes have almost the same Daily Revenue as the panamax segment. The Age-variable has only modest variance across regions, an insignificant coefficient, and gets excluded (Table 3.11). The exclusions lead to the Lean Model which is acceptable. All coefficients are significant, their signs are correct and absolute values make sense.

The formal task has been completed and it is time to proceed to the finishing chapter. An attractive but premature idea. The fact is that at times there has been an uncomfortable feeling that the discussion is about a bogus problem, the annual freight revenue of ships taking their first load in a certain region. Why have we not discussed the freight revenue of the first trip, instead? Because we have been intrigued by the combinatorial character of ship routing and thought that a good start would put the ship on a better itinerary (path dependency) than otherwise and, although the alternatives have been far too numerous to be tracked down in detail, the overall result would give indirect testimony. That was the idea. But if a superior freight revenue, and attached Revenue Gradient, is mostly due to the first cargo sequence, then it should be just as well to focus on that sequence alone. The reasoning carries weight because the median share of the first trip, ballasting and waiting included, is 12-17 pct of the total recorded time, the exact figure depending on the segment and year. And when its effect has been subtracted from the total, the rest stands for the path dependency only.

Table 4.9 Daily Revenues regressed by alternative models and data

Model	Full set DayRev		Lean DayRev		Lean FstRev		Lean RstRev	
Observations	86		86		86		86	
R-sqr (adj)	0.978		0,977		0,866		0.974	
F-statistic	428		601		92		527	
Standar error of estimate	654		675		1435		744	
Regression coefficient								
Revenue Potential	259	(0.03)	335	(0.00)	448	(0.07)	302	(0.02)
Atlantic	493	(0.00)	444	(0.00)	425	(0.19)	353	(0.04)
Age	-24	(0.61)	na	na	na	na	na	na
Dwt	31	(0.01)	na	na	na	na	na	na
D-CL	632	(0.61)	3,992	(0.00)	4,419	(0.00)	3,924	(0.00)
D-CS	-1,650	(0.03)	na	na	na	na	na	na
D-H	-1,332	(0.00)	-2,356	(0.00)	-2,260	(0.00)	-2,363	(0.00)
D-95	6,062	(0.00)	6,050	(0.00)	4,687	(0.00)	6,300	(0.00)
D-98	-3,621	(0.00)	-3,562	(0.00)	-2,147	(0.00)	-3,779	(0.00)
Intercept	7,124		8,897		9,041		8,898	
Range								
DayRev	15,681		15,681		16,903		16,396	
RevPot	3.240		3.240		3.240		3.240	
Mean DayRev	10,851		10,851		11,224		10,761	
GeoImpact, pct	8.5		9.8		11.1		8.1	
(RevPot only, $	839		1085		1452		978)	
(pct	5.4		6.9		8.6		6.0)	

Notes: Only macroregions with at least 4 ships. Significance in parentheses. GeoImpact = 100 * Abs ($\text{Range}_{\text{RevPot}}$ * RevPot + Atlantic) / $\text{Range}_{\text{DayRev}}$.
Source: Author.

To find out, the full data set is split into the First Trip and the Rest, and the Lean runs are repeated with FstRev and RstRev as the dependent variables. The DayRev includes all kinds of activity, ballasting, loading, cargoing, discharging and waiting, and the Rest does the same by inference. It follows that the First Trip must be tailored accordingly. That creates slight inaccuracy because the recording period always begins with loading and only the first cargo sequence (loading-cargoing-discharging) with subsequent waiting can be measured unambiguously. By contrast, the preceding ballasting leg is unknown and the average ballasting time of the segment and year to that specific loading region must be used, instead (Appendices 2.3 and 2.4; Table 4.10).

The hypothesis is that in regions with high rates the FstRev will exceed RstRev whereas in regions with low rates it is the other way round. That is logical when DayRev supports a Revenue Gradient and when most of it originates from the First Trip. The reality is too varied to suit this simplistic idea, however (Appendix 4.5). Generalization succeeds better when regions are aggregated by segment and year. Then FstRev tends to be higher than DayRev, as anticipated, and even a trend can be identified: FstRev upgrades with time relative to DayRev while RstRev downgrades (Table 4.10). No explanation can be offered for the trend. But the next observation is a partial contradiction. RstRev hardly deviates from DayRev in 1995 and 1997 and support comes only in 1998. These observations are only indicative, however. The only reliable way to find out is to repeat the regressions with First Trip and Rest data (FstRev and RstRev; Table 4.9). The fit is worse with FstRev than DayRev or RstRev and this feature will remain all through the forthcoming runs. FstRev simply has less scope for smoothing, visible also in a scatterplot (Fig. 4.14). That, however, is not the point but the fact that the parameters are about the same in each of the three runs which means that the Revenue Gradient applies for each of them. DayRev and RstRev, in particular, will be quite similar to each other in all the relevant alternatives (Table 4.11). Even DayRev and FstRev are similar to each other, which is unexpected considering that once the first cargo leg has been completed there need be no lucrative, or any, cargos readily available. The correlation coefficients suggest something else. On average, the net effect may tilt slightly in FstRev's favor, meaning that the Gradient is on a higher level and steeper. Indeed, mean FstRev is 4 pct higher than DayRev or RstRev and the geographical variables Revenue Potential and Atlantic have a larger total effect (Table 4.9).

Table 4.10 Simulated times and relative Daily Revenues

	Ballasting			Waiting			FstRev			RstRev		
Segm/Year	95	97	98	95	97	98	95	97	98	95	97	98
CapeL	17.61	17.12	16.35	4.02	4.72	6.59	93	104	130	101	99	94
CapeS	14.61	14.34	14.82	8.45	9.17	6.75	106	105	128	99	99	95
Panamax	12.72	11.75	10.82	3.81	3.49	2.40	90	99	126	101	100	96
Handy		6.86			3.97			106			99	

Notes: Time in days; DayRev = 100.
Source: Author.

Table 4.11 Important Pearson correlations between regression variables

	Full set	CapeL	CapeS	Panam	Handy	1995	1997	1998
DayRev, RstRev	0.998	0.998	0.998	0.999	0.939	0.994	0.995	0.990
DayRev, FstRev	0.929	0.915	0.932	0.939	-0.065	0.824	0.929	0.919
FstRev, RstRev	0.905	0.887	0.912	0.923	-0.396	0.764	0.888	0.867
RevPot, Atlantic	0.141	0.118	0.102	0.174	0.351	0.238	0.127	0.053

Source: Author

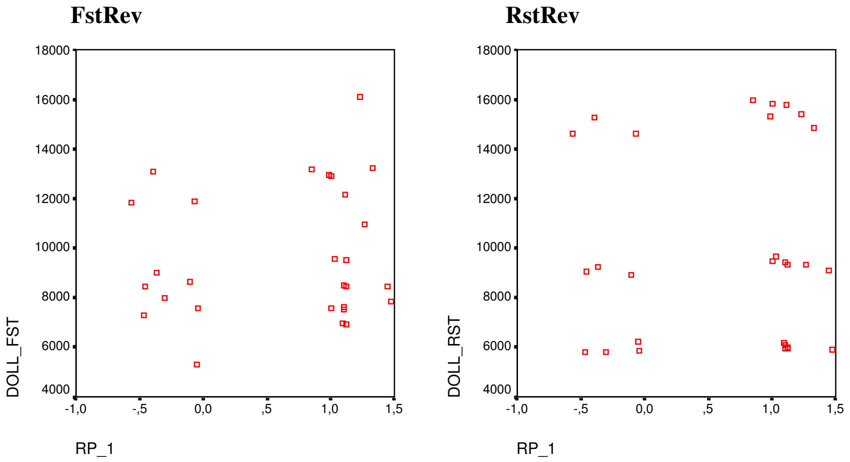

Figure 4.14 Panamax Daily Revenue by year

Source: Author.

The outcome suggests that the Simulator has been an unnecessarily complicated tool to the purpose it has been used for. No greater truth has been discovered with its help. We know it now. But we did not know it when we started. We might have sensed some of it by deduction. A monte carlo simulator which follows known transfer probabilities (movement matrices) cannot possibly give results very different from these probabilities. But we were unable to track down all the myriads of alternatives and weigh them properly. Now we have measured their aggregate effect with reasonable accuracy. Without the Simulator it had not been possible. The ship diaries, a data source not readily available elsewhere, have given us a firm idea about the itineraries, the geographical core of all shipping, and the time spent for different activities in a given demand/supply situation. We have not analyzed this data thoroughly and hope to return to it in some other context.

This project rests on two hypotheses:

- Geography matters for revenue generation.
- It matters more, relatively speaking, when freight rates are high.

That geography matters is confirmed by the acceptable, if not high, significance of the both geography variables, Revenue Potential and Atlantic dummy. The importance of the Atlantic/Pacific split is standard fare among shipping professionals (Fig. 1.3). The importance of tonnage demand and supply is also obvious. Its measurement may have been less common because the data is more difficult to collect. Now we will integrate these two measures and relate them to the rate level.

The integration is easy. The regression model is additive and the two variables hardly correlate with each other (Table 4.11). It means that their coefficients, suitably weighted, are additive and the total gives the aggregate geographical impact (GeoImpact). The Atlantic-variable is a dummy and measures simply the difference between the Atlantic and Pacific Spheres. There is no theoretical reason for the freight rates to be in all circumstances higher in the former (positive sign). It can also be the other way round. Therefore, the absolute value is taken of the Atlantic coefficient. Representing a dummy, it needs no weighting, whereas the coefficient of Revenue Potential is weighted by the variable's range. Their total is divided by the range of the Daily Revenue and expressed as a percentage. Range is considered a proper weight because it defines the playroom available for a ship manager. When absolutely necessary (Handy segment, below) it can be replaced by some other reduction base, such as mean or intercept. This percentage varies now within 8-11 pct between the runs with the FstRev run in the higher end. **Geography matters that much, about 10 pct.** Revenue Potential stands for most of it.

The link between the GeoImpact and different rate levels cannot be established from the aggregate data which has to be disaggregated by year. At the same time it will be disaggregated by segment also, recalling the importance of ship size for practical ship management. The disaggregation means in practice three one-page tables which are best given as an appendix (Appendix 4.6). All the runs have a similar setup, using RevPot, Atlantic, segment and year dummies as independent variables. It would be possible to substitute Dwt and the D-CS dummy for the current set of variables in the year runs without essentially affecting either the R-sqr or coefficient significancies (Fig. 4.15). That possibility is not used because it would not contribute to the goal of this study.

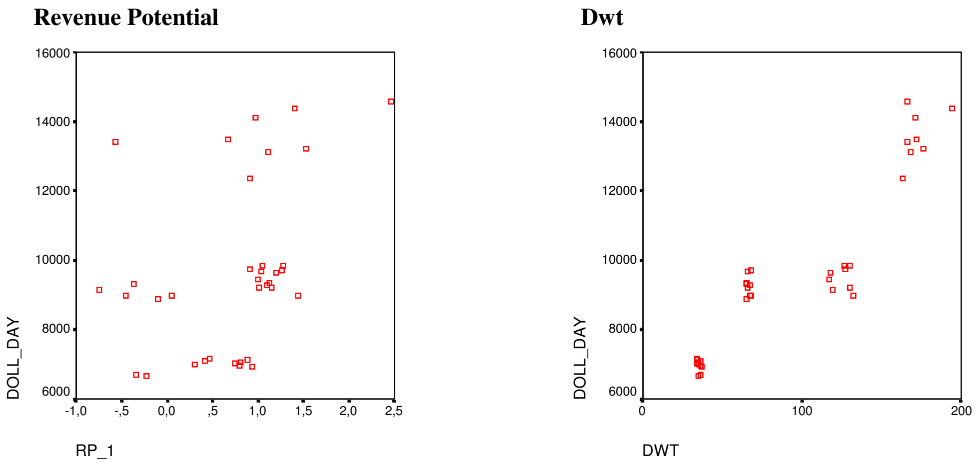

Revenue Potential

Dwt

Figure 4.15 Daily Revenue against Revenue Potential and Dwt, 1997

Source: Author.

The quality of the results deteriorates with disaggregation, meaning a larger share of insignificant coefficients, generally less stable coefficients, and occasionally coefficients with an illogical sign. All that is probably unavoidable because the number of observations per run is at least halved, which their greater homogeneity is unable to compensate. The outcome is worst with FstRev and there is no point to comment on it at all. The results of the DayRev and RstRev runs are similar enough to be discussed simultaneously.

The coefficients, when compared pairwise, are in the same size class and tend to decline with the mean freight rate. Within the segment runs, the only seriously insignificant coefficient is Atlantic in the DayRev CapeS segment. The negative sign of the coefficient, indicating higher rates in the Pacific than Atlantic, may originate from the insignificance although it may also reflect the actual state of affairs. It is impossible to know for certain on the basis of available evidence. Within the year runs, insignificant geography-coefficients are a problem in 1998. Recalling the very low level of freight rates that may reflect the actual state of affairs. There simply was no scope for serious rate differentiation and erratic coefficients are then a logical outcome.

The estimation of GeoImpact contains a formal inconsistency. In the Handy segment, the comparison is to mean and not range as elsewhere. The reason is the narrowness of the range, due to the lack of such high-profile export and import regions which are typical for the other segments. Since the mean and range are not too different in the other segments the decision should be defensible, nevertheless. It leads to estimates which decline with ship size or mean freight rate. The answer to the second hypothesis rests on the interpretation of the insignificant coefficients in 1998. If the insignificance is interpreted in the orthodox way, viz. that the coefficients do not differ from zero, then the GeoImpact is zero, too. And then the sequence from 1995 to 1998 declines with the rate level, or almost because the "RevPot only, pct" line in the RstRev runs has an anomaly in 1997. The overall evidence, however, is for the substatiation of the second hypothesis. **Geography matters more when rates are high.**

As to the "absolute" size of percentages, they fluctuate around the 10 pct mark, plus minus 4 pct, in the segment runs and within a 10-15 pct range in the year runs. The plain RevPot figures can be still lower although it should be observed that the exclusion of Atlantic from a run is likely to raise the RevPot coefficient. From a practical angle, such levels are both too low and appropriate. They may be too low to offer an incentive for serious efforts to upgrade the routing, although the practical advice of how to do it is very tangible – "start from an Atlantic region with a serious tonnage shortage and then apply Bazaraa's algorithm". And they may be appropriate because a higher level would have aroused attention long ago and led to corrective measures. For the sake of comparison, a basic vehicle routing algorithm normally gives a 15 pct shorter driving distance for a fleet of 20 delivery vans compared to intuitive reasoning and experimentation. We are in the correct size class and that gives credence to our results.

”It is this stability over a fairly long period ... and such conditions on the whole change but slowly.”
Sargent (1930), 120.

5 DISCUSSION

The Discussion contains several themes. The ability of Revenue Potential to explain Daily Revenue, made tangible by Revenue Gradient, is the central topic and most of the effort and space is dedicated to it. Success in that effort paves the way for the application of the concept in a related context, the forecasting of freight rates after a fundamental rearrangement of trade flows (Laulajainen 2006). Success also there presses home a fundamental truth: **Distance-weighted tonnage balance is the key for all freight rates and, through them, profitability**. New light is shed upon several, in our context, subsidiary topics: the numerical split into retail and wholesale markets, the share of cargo legs covered by published fixtures, the comparison of actual trade matrices with those displaying minimum transport workloads, and attempts to form route clusters on the basis of freight rates. Practically all data has been “raw”, allowing great flexibility but also forcing ad-hoc solutions in situations where few examples existed. The allocation of cost items in time and geographical space, and the large-scale conversion of voyage rates into trip rate equivalents, have been the most challenging ones. Practical possibilities to handle such questions are influenced by the form in which the “raw” data is available, and that may constrain the size of problems amenable for analysis. Few simple, but potentially expensive, measures would have been extremely helpful.

The overall theme of the study has been the observation that the average daily freight revenue during a year (Daily Revenue or DayRev) by ships taking their first cargo in a particular region is a linear function of the region’s distance-weighted Tonnage Balance (Revenue Potential), plus a dummy indicating whether the region is located in the Atlantic or Pacific Sphere. Weighting measures the region’s closeness to its trading partners. The function has a rising slope which gives it the title Revenue Gradient. The immediate reaction was that the gradient is mostly due to the first cargo sequence. The inference carries weight, when the First Trip, ballasting and waiting included, typically accounts for 15 pct of the total recorded time. The First Trip certainly influences the rest, but less than one might think. Its average daily revenue (FstRev) is only 4 pct higher than during the rest of the recorded year (RstRev), and when the geographical variables Revenue Potential and Atlantic dummy are related to the range of Daily Revenue, the percentage is only a few pct-units above the overall figure of 10 pct.

Conceptually, the existence of a Revenue Gradient is unexpected, at least from the conventional economic angle. The market is reasonably atomistic and the key production factors, ships, appear to be perfectly mobile. Consultants and trade press inundate the market with information. Within this framework, a local shortage of vessels stimulates a higher rate and vice versa, which leads to differentiated local price levels, sufficient to give rise to the Revenue Gradient. In other words, there are numerous local markets and their atomistic features may be in a constant flux. This insight sparked off the original study, which tried to show that the surplus tonnage available in neighboring macroregions influences the local spot rate. The idea did not work in that simplistic shape but, when reworked to the current conceptual framework, it was promising at the aggregate level (LHS 2001, Fig. 6.7). Our current effort is an extension of that early work. There are also other potentially relevant factors such as variation in ship size (within a segment), age, multiporting, cargo handling terms, oligopolistic tendencies and newbuildings, but most seem to be relevant at the operational rather than aggregate (macroregion) level. Ship size and age have been tried in regression equations but with limited success. Multiporting might face a similar fate. The modest survey about panamaxes in 1997 did not reveal pronounced differences between macroregions (Table 3.13), although most of the vessels were comparatively new and not very keen to follow the practice. Handling terms are visible only in voyage fixtures and abstracted away in the conversion to trip equivalent. Oligopolistic tendencies among charterers were in general stronger in the Pacific than Atlantic loading regions, and this manifested itself in lower rates through the Atlantic/Pacific dummy. The concentration of newbuildings in the Far East contributed, although the effect was imbedded in the Tonnage Balance as well.

The conclusion then is that change in the Revenue Potential is a powerful argument for the Revenue Gradient to emerge. The explanation offered previously was that the market is inefficient (Laulajainen 2003). Routing and scheduling are typical combinatorial problems in which the great number of feasible solutions renders the finding of good overall (global) solutions extremely difficult. This leaves room for seemingly illogical pricing, labeled as inefficiency. The label may be too challenging. It can be oligopoly as well. Or everything can be as it should, given local imbalances in tonnage demand and supply (Laulajainen 2006). Call it myopic or opportunistic pricing. When hard arguments are insufficient, one turns to soft ones. The early interpretation of the Cobb–Douglas production function with its two components, labor and capital, is illuminating. What the function could not explain was claimed to originate from technical development. It would have been far better to find a suitable indicator for this development and measure it directly, as also happened later on. We have no need for soft explanations. Revenue Gradient is the expression of Revenue Potential, a very tangible measure consolidating relative tonnage availability and distance.

Against this background it is not so astonishing that empirical calculations do not support the Revenue Gradient hypothesis unconditionally (Table 4.8). External expertise forecasted that the hypothesis would fare well in the capesize and panamax segments but not among the handysizes. This author was rather confident about the handysizes but could see no good reason why the simplistic capesize segments would fall into line. When tested with the DayRev data by segment and year, the R-sqr is, in four cases out of ten, all capesize runs, so low that the outcome can be labeled as a failure. And when 10-pct statistical significance is taken as a criterion, in only five cases the coefficient of Revenue Potential clearly supports the hypothesis. Four of the non-supporting ones are in the capesize segments. The handysize segment is studied in only one year and cannot, therefore, give full-good evidence. The panamax segment stands the test best, but the year 1997 also has an insignificant coefficient. The small number of observations can be used as a partial excuse, however. The remedy against too few observations is to pool the data. It follows that dummies are added to account for level differences between the segments and years, and that most of the variance in the Daily Revenue is accounted for by them. High to very high R-sqrs emerge and cannot be used as a selection criterion. The coefficients remain: their signs, significancies and relative sizes. Overall, they give a logical and reliable picture. The hypothesis passes the test with an acceptable grade.

The test is conducted by using macroregions as units of observation. In the original report, most tests were based on the 700 panamax vessels rather than the nine macroregions. The practical consequence was that the R-sqrs fell from about 0.94 to about 0.15. The regression coefficients, by contrast, were not seriously affected. It is to be expected that the current results would change in a corresponding way if vessels were subjected for macroregions. There is no reason for such a change, however. The emphasis in the original report was on individual vessels, preferably the best ones, whereas here it is on the average vessels, for reasons given when Bazaraa's optimizing algorithm was abandoned and the simulator selected as the analytical tool (Subchapter 1.2).

The formal task set for the study, to substantiate or reject the validity of the Revenue Gradient, has been achieved. The Revenue Gradient is substantiated. It rests on the distance-weighted Tonnage Balance, which reflects the locational imbalance between resource use and availability, an imbalance which is geographical and likely to last. Hence the title, "The Geographical Foundations of Dry Bulk Shipping". The importance of the geographical impact is measured by multiplying the ranges of the Revenue Potential and the Atlantic/Pacific dummy by their (absolute) regression coefficients and comparing the product total with the range of the dependent variable. The GeoImpact indicator falls in the segment runs roughly within a 6–15 pct range, $400–$1,800 in monetary terms (Appendix 4.6). That is the formal end of the study.

But there is always a world beyond the immediate horizon. The result is likely to hold in conditions approximately similar to those in which it has been derived. An academic is not satisfied with that. She/he is on the lookout for a general solution. Suppose that by some streak of fate the current distribution of resource availability and use would be completely upset. Such things have happened in the historical past. The "discovery" of new continents and their integration with the economies of the colonizing countries was a major upheaval, although its realization took time. The substitution of oil for coal and the demise of coal mining in Western Europe have taken place during this author's lifetime. It is utterly unrealistic to think that reorganizations of similar magnitude would not take place in the future also. Would they be large enough to upset the global dry bulk rate structure as we know it today? We do not know but we can make assumptions and test our ideas in this Virtual World.

We make assumptions about resource availability and use. We derive the minimum distance transportation flows by linear programming. We know that they must be augmented by some surcharge because of commodity non-homogeneity and time & location constraints in ship move-ments (Tables 2.15 and 2.16). How the surcharge is distributed between matrix cells is a question to be addressed when the time arrives. Trader and broker opinions might be a good start. We calculate Revenue Potentials, and there the routes diverge. The first one makes use of the functions developed in this study. They give only the total revenue accruing to an average ship which takes its first cargo in a particular macroregion, and leaves the way open for the question of how the revenue accumulates during the year. That disclosure is connected with ships' transfer probabilities between routes (matrix cells), and the relative rate levels are derived by goal programming. The formal task is to distribute the global freight revenue between the routes in such a way that the row constraints from the Revenue Gradient are honored. It may not work out exactly like this but the general direction should be correct. The second route is much simpler, to bring the Revenue Gradient idea down to the route level. The details will be discussed in a separate study due to be published in the near future (Laulajainen 2006).

The original study failed to show that the surplus tonnage available in neighboring macroregions contributes to the explanation of freight rates within a two-week time period. An unpublished lag experiment, connected to this report, came to the same conclusion. Both tests used historical, rather coarse information and are therefore only conditionally valid. An analyst working in real time could have purchased data about ships currently in port, waiting–loading–discharging, and satellite pictures locating vessels en route with an accuracy of 50 nm, at a cost of $6,000 per picture (in 2000) – the daily operating cost of a capesize vessel. That belongs to the past. Today, more accurate information is available at the fraction of the cost. The mandatory Automatic Identification System continuously signals a ship's name, position, destination, course and speed over a VHF radio channel (Anonymous 2002). The purpose is to avoid collisions and identify the origin of oil spills, but the same information can be used for commercial intelligence as well. The short range of VHF technology is a handicap but not an insurmountable one. That still leaves us short of full information because cargoes remain unknown, as are ship availabilities. But long strides have been made in a short time to help those who wish to make use of them.

Then the "minor" topics, which are minor only in this report. In reality, each of them is a field of its own and worthy of specialized study. One of them is the split into retail and wholesale markets, i.e. voyage and trip charters in contrast with COAs and time charters, which helps in putting the study in context. The analysis has been conducted as if the whole dry bulk sector operated in an undifferentiated single-trip spot market. Deductions made for within-port, specialized, liner-type and shuttle traffic modify this blunt statement to a degree, but not fundamentally. The concept of a unified spot market is, of course, untrue. There are owned fleets and there are long-term chartered vessels, the wholesale market. The latter share has been estimated as 37 pct of the sector total (Table 1.6), whereas no attempt has been made to measure the share of owned fleets. If a figure must be put on them, our guess is 10 pct overall, which leaves 53 pct for the spot market. It would have been possible to make a corresponding, undifferentiated downsizing to our calculations. Unfortunately, all route-specific variation which reasonably must exist would have escaped measurement, exactly as it does now. Therefore, what is the use of downsizing? Moreover, there is no reason to think that the results would be much different with the market split into a black-box-like wholesale market and a transparent retail market.

Another taxonomic problem is the share of cargo legs covered by published fixtures. Any answer is impregnated with qualifications. Segmentwise figures based on Stripped files (without within-port, specialized and liner-type traffic) are 15, 40, 30, 5 and 15 pct (large and small capesize, panamax, handysize and all; Table 3.5). Distributed between routes, the variation becomes very large, from 0 pct to over 100 pct. Long, heavily-trafficked routes have high percentages because they are used as benchmarks, whereas short, particularly within-region, routes are their opposite. This information is derived from one-half of the active fleet. More problematic is the fact that there is no direct evidence that a published fixture leads to a physical movement. Quite on the contrary, there are very many speculative fixtures, some claim up to 80 pct, which would tally with derivative exchanges.

Our calculations about the minimum workload have little in common with traditional efficiency estimates. The traditional way is to seek slack in the fleet's time budget, layup, inactivity, excessive port time, waiting, slow steaming, part cargo, etc. Some of it is genuine inefficiency, but much is better termed "surplus capacity". From our geographical angle, inefficiency is connected with uncoordinated physical movement. Assuming that ships within a size segment are equally suited, or almost, for all kind of transport tasks, and that trading partners are broadly indifferent to each other's identity, the inefficiency can be estimated by comparing the actual movement pattern (in cargo and ballast) with the minimum cost pattern. The minimum cost pattern is the solution to the classical transportation problem. We have solved it on an annual basis and got a very coarse estimate because most cargoes actually need to be delivered in a time window. That dilemma can be solved by disaggregating the annual problem into quarterly, monthly, biweekly or weekly subproblems. Our choice in a real-world (non-academic) situation would be a set of biweekly problems. Relevant cost data would obviously be the next problem. Plain distance has been used in this study but a case can be made for using freight rates, instead. Since annual tonnage balances would be replaced by biweekly ones, rates would necessarily follow suit. A step towards dynamics had been taken.

Freight rate clusters are a far cry from normal economic thinking which operates with segment-specific or outright global rates. Clustering was attempted originally with the help of recursive equations (LHS 2001, Table 4.14). The result was a tree-like graph which divided route groups into two, with loading ports either in the Atlantic or Pacific, a familiar split although in an unexpected context (LHS 2001, Fig. 4.13). A new attempt has been waged in this report by correlating rate indices pairwise. Two very distinct groups have emerged with diametrically opposite cyclical patterns. These patterns appear to be connected with the character of loading regions either as a net exporter or importer, modified by the region's location in the northern or southern hemisphere if the exports are dominated by agriproducts. Since clustering based on time series is not an everyday analytical problem and definitely on the margin of this author's competence, both results should be considered as tentative.

The total number of fixtures to be converted into freight rates has been 40,000, of which one-half are voyage fixtures necessitating another conversion into trip equivalents. The conversions call for a certain technical skill and what started as a traditional shipping project rapidly assumed features of an accounting exercise. As corporations have their Generally Accepted Accounting Principles (GAAP), the shipping industry needs its Generally Accepted Rating Principles (GARP). The question is not trivial because, exactly as in conventional accounting, different principles give different results and that matters. And exactly as in conventional accounting, much depends on how cost items are allocated in time. In this study, all cost items have been allocated to the time when the ship is in cargo, whether loading, sailing or discharging. The data distributes among some 150 routes in four size segments, and all kinds of ships coexist in this once-in-lifetime batch. Its size, heterogeneity and abrupt appearance ruled out the piecewise analysis of typical ships as practiced by consultants and a rapid, uncomplicated, mass-production technique, basically heuristics, had to be developed. Judging by the scatterplots the results are acceptable, nevertheless (Appendix 3.8). The wide variation of rates at a given value of argument is primarily an eyesore as long as a sound function can be derived. When the variation originates from a heterogeneous fleet or varying local conditions, it can be calibrated by changing the SEE. And when it originates from a time-related parameter change, the estimation can be conducted from appropriate subperiods. These exact techniques have also been used.

Shortcomings in data have been pointed out when they have emerged. These comments should be seen in perspective. Shipping is an industry where practically all worldwide interaction data is collected by the private sector. The United Nations tried, but abandoned it two decades ago. The private sector cannot operate by fiat, commercial reasoning and diplomacy are its tools. It cannot use tax money either but must be self-financing. The very existence of this report is a potent testimony to its success. But everything human is imperfect and our databases are no exception. The data about ship movements suffer from the lack of a ship's cargo status. An analyst must spend much time in error-prone reasoning. It would be fully possible organizationally, with reservation for part cargoes, to collect this information in the field. The lack is entirely due to commercial reasons; the market is not prepared to pay for the service. In the fixture data, the shortcoming is both industrywide and banal. Various types of information, including alphabetical information, is given in one single data column (Narrative, Comments or similar). They need to be diveded between separate columns and possibly converted into a numerical form. This report may be of some help in the effort.

REFERENCES

Aldworth, Paul (ed) (1997), *Lloyd's Maritime Atlas*, 19th ed. LLP, London.

Alizadeh-Masoodian, Amir H. (2003), letter dated 23 January.

Anonymous (2002), July deadline for AIS. *Lloyd's Ship Manager*, January/February, 50.

Armington, P.S. (1969), A theory of demand for products distinguished by place of production. *IMF Staff Papers 6*, 159-178.

Barnes, Hilary (1995), Storm over tax regime. *Financial Times* - Norway, 20 November, IV.

Batchelor, Charles (1997), Choppy waters ahead. *Financial Times*, 25 April, 15.

Bazaraa, M.S., Sherali, H.D. and Shetty, C.M. (1993), *Nonlinear Programming. Theory and Algorithms.* Wiley, New York.

Berg-Andreassen, Jan A. (1998), A portfolio approach to strategic chartering decisions. *Maritime Policy and Management 25* (4), 375-389.

Binkley, James K. and Harrer, Bruce (1981), Major determinants of ocean freight rates for grains: an econometric analysis. *American Journal of Agricultural Economics 63* (1), 47-57.

Bose, Kunal (2000), Stormy seas on horizon as labour resists privatisation. *Financial Times*, 17 March, 15.

Burki, S. J. (2001) An Asian stampede. *Financial Times*, 11 June, 13.

Burley, K.H. (1960), The overseas trade in New South Wales coal and the British shipping industry, 1860-1914. *The Economic Record* (August), 393-413.

Caney, R.W. and Reynolds, J.E. (1995), *Reed's Marine Distance Tables*, 8th edition. Thomas Reed Publications, East Molesey, Surrey.

Chang, Young-Tae and Chang, Hak-Bong (1996), Predictability of the dry bulk shipping market by BIFFEX. *Maritime Policy and Management 23* (2), 103-114.

Chavet, Jean-Paul (1985), *Les greniers du monde.* Economica, Paris.

Cockett, N. (1997), *Neil Cockett on Bunkers*. Lloyd's of London Press, London.

Collins, Nick (2000), *An Essential Guide to Chartering and the Dry Freight Market*. Clarkson Research Studies, London.

Coulson, E.C. (1991), *A Guide for Tanker Brokers.* Clarkson Research Studies, London.

Couper, A.D. (1972), *The Geography of Sea Transport*, Chapter 6 Conventional dry cargo shipping, 90-129. Hutchinson, London.

Couper, Alastair (ed) (1983), *The Times Atlas of the Oceans*. Times Books, London.

Cunningham, Kevin (ed) (1997a), *International Metals Databook*, 1st edition. CRU Publishing, London.

----- (1997b), *World Fertilizer Plant List & Atlas*, 11th ed. British Sulphur Publishing, London.

Daniel, Caroline (2004), Château Cargill throws open its hallowed halls. *Financial Times*, 25 February, 4.

Daniels, Alan (1997), Thirty grain ships delayed in Vancouver. *Lloyd's List*, 7 February, 1.

Dewitte, Olivier (2001), Le point sur les phénomènes d'*el Niño*, de *la Niña* et l'oscillation australe. *Bulletin de la Société géographique de Liège 40* (1), 15-32.

Dorfman, R., Samuelson, P. and Solow, R. (1958), *Linear Programming & Economic Analysis*, International student edition. McGraw-Hill, New York.

The Drewry Monthly (1993-2002), Drewry Shipping Consultants, London.

Emery, Chris (ed), 1998), *Lloyd's Ports of the World*. LLP, London.

Evans, J.J. (1994) An analysis of efficiency of the bulk shipping markets. *Maritime Policy and Management 21* (4), 311-329.

Fleming, D.K. (1978), The concept of flexibility. *GeoJournal 2* (2), 111-116.

Ford, David A. (1975), Shang inquiry as an alternative to Delphi: some experimental findings. *Technological Forecasting and Social Change 7*, 139-164.

Fuller, Stephen, Makus, Larry and Gallimore, William (1984), Effects of increasing Panama Canal toll rates on US grain exports. *Southern Journal of Agricultural Economics 16* (2), 9-19.

Gilman, S (1975), The choice of ship size on deep-sea general cargo routes. *Maritime Studies and Management 3* (2), 95-102.

Gimbel, Florian (2004), A perfect mix of oil and water. *Financial Times*, 25 October, 7.

Giraud, Pierre-Noël (1983), *Geopolitique des ressources minieres*. Economica, Paris.

Glen, David (1990) The emergence of differentiation in the oil tanker market, 1970-1978. *Maritime Policy and Management* 17 (4), 289-312.

Glen, D., Owen, M. and van der Meer, R. (1981), Spot and time charter rates for tankers, 1970-77. *Journal of Transport Economics and Policy 15* (1), 45-58.

Glen, D.R. and Rogers, P. (1997), Does weight matter? A statistical analysis of the SSY capesize index. *Maritime Policy and Management 24* (4), 351-364.

Goodman, Allen C. and Lenze, David G. (1988), Port facilities and the international coal market. *Transportation Research 22A* (2), 137-144.

Granger, C.W.J. (1989), *Forecasting in Business and Economics*, 2nd ed. Academic Press, San Diego.

Gray, Tony and Owen, Philip (1996), Why VLCC owners must pull together. *Lloyd's List*, 23 October, 3.

Hale, C and Vanags A. (1989), Spot and period rates in the dry dulk market: some test for the period 1980-1986. *Journal of Transport Economics and Policy 23* (3), 281-291.

Hammer, Jarle (2005), Director, phone interview, 20 January. Fearnresearch, Oslo.

Haralambides, Hercules E. (1996), The economics of bulk shipping pools. *Maritime Policy and Management 23* (3), 221-237.

Hauser, Robert J. (1986) Competitive forces in the U.S. inland grain transport industry. *Logistics and Transportation Review 22* (2), 158-183.

Heaver, Trevor D. and Studer, Keith R. (1972), Ship size and turnround time. *Journal of Transport Economics and Policy 6* (1), 32-50.

Isserlis, L. (1938), Tramp shipping cargoes and freights. *Journal of the Royal Statistical Society 101* (1), 53-146.

Jansson, J.O. and Shneerson, Dan (1985), Economies of trade density in liner shipping and optimal pricing. *Journal of Transport Economics and Policy 19* (1), 7-22.

Jones Bob (ed) (1997), *Iron and Manganese Ore Databook*, 2nd ed. Metal Bulletin Books, Worcester Park, Surrey.

Jonnala, Sneha, Fuller, Stephen, and Bessler, David (2002), A GARCH approach to modelling ocean grain freight rates. *International Journal of Maritime Economics 4*, 103-125.

Jupe, Malcom and others (1996), *Dry Bulk Freight Rates and Chartering, Players, Strategy and the Market*, October. Drewry Shipping Consultants, London.

Kalindaga, Y.C. (1990), Estimation of capacity utilization in world shipping. *Maritime Policy and Management 17* (1), 41-67.

Kavussanos, Manolis G. (1996a), Comparisons of [freight rate] volatility in the dry-cargo sector. *Journal of Transport Economics and Policy 30* (1), 67-82.

----- (1996b), Highly disaggregated models of seaborne trade. An empirical model for bilateral dry-cargo trade flows in the World economy. *Maritime Policy and Management 23* (1), 27-43.

----- (2002), Business risk measurement and management in the cargo carrying sector of the shipping industry. Grammenos, Costas Th. (ed), *The Handbook of Maritime Economics and Business*, 661-692. LLP, London.

Kavussanos, Manolis G. and Alizadeh-M, Amir H. (2002), The expectation hypothesis of the term structure and risk premia in dry bulk shipping freight markets. *Journal of Transport Economics and Policy 36* (3), 267-204.

Koo, Won W., Thompson, Stanley R. and Larson, Donald W. (1988), Effects of ocean freight rate changes on the U.S. grain distribution system. *The Logistics and Transportation Review 24* (1), 85-100.

Lacroux, Eliane (ed) (2002), *World Cement Directory*, Volume II, 11th ed, 154-189. Cembureau, Brussels.

Laulajainen, Risto (2002), Geographical inefficiencies in the panamax dry bulk market. Unpublished manuscript, March.

----- (2003), Dry bulk market inefficiency, *Lloyd's Shipping Economist 45* (December), 25-27.

----- (2006), A static theory of dry bulk freight rates by route. *Maritime Economics and Policy 33*.

----- (2007), Dry bulk shipping market inefficiency, the wide perspective. *Journal of Transport Geography 15*.

Laulajainen, Risto, Holgersson, Thomas and Strömberg, Ann-Brith (LHS) (2001), *Operating Panamax Dry Bulk Carriers on the Seven Seas*. Gothenburg School of Economics and Commercial Law.

Lloyd's Register of Ships 1997-98 (1997), Lloyd's Register of Shipping, London.

de Lombaerde P. and Verbeke, A. (1989), Assessing international seaport competition: a tool for strategic decision making. *International Journal of Transport Economics 16* (2), 175-192.

Lundgren, Nils-Gustav (1996), Bulk trade and maritime transport costs. *Resources Policy 22* (1/2), 5-32.

Manners, Gerald (1971), *The Changing World Market for Iron Ore*, Chapter 9 Changes in ocean transport, 173-200. Johns Hopkins Press, Baltimore.

Marriot, Ian and Oatway, Susan (1998), *The Drewry Dry Bulk Quarterly*, July issue. Drewry Shipping Consultants, London.

Marsh, Peter (2004), High prices have helped put the steel industry back on its feet. But a secure future is still under construction. *Financial Times*, 29 September, 11.

Marsh, Virginia and Morrison, Kevin (2005), Feeding the Chinese giant. *Financial Times* – World Ports, 23 May, 2.

Metaxas, B.N. (1972), The future of the tramp shipping industry. *Journal of Transport Economics and Policy 6* (3), 271-280.

Mokia, Z. and Dinwoodie, J. (2002), Spatial aspects of tanker lay-times. *Journal of Transport Geography 10* (1), 39-49.

Nossum, Birger (1996), *The Evolution of Dry Bulk Shipping 1945-1990*. Fearnley and Egers, Oslo.

O'Mahony, Hugh (1999), Carrying the can. *International Bulk Journal*, June, 11-14.

Packard, W. V. (1996), *Voyage Estimating*. Fairplay Publications, Coulsdon.

----- (1997), *Sale and Purchase*, 2nd ed, reprint. Fairplay Publications, Coulsdon.

----- (1998), *Sea-trading, Cargos*, reprint. Fairplay Publications, Coulsdon.

Park, Joon Je and Koo, Won W: (2000), An econometric analysis of ocean freight rates for grain shipments from the United States to major importing Countries. *Journal of the Transport Research Forum 43* (2), 85-100.

Pettersen Strandenes, Siri (1984), Price determination in the time charter and second hand markets. Working Paper MU 06. Norges Handelshøyskole, Bergen.

Pirrong, , S.C. (1993), Contracting practices in bulk shipping markets: a transactions cost explanation. *Journal of Law and Economics 36* (October), 937-976.

Prokopy, Steven (ed) (2003), *The 2004 North American Cement Directory*. Primedia, Chicago.

Results of the Maritime Transport Study for the years 1972-1975 (Maritime Transport Study 1980). Analysis by type of goods moved during 1972-1975 between regions of the world. Commodity Trade (By Sea) Statistics. Statistical Papers, Series D, Vol. XXIII-XXIV, No. 2. Department of International Economic and Social Affairs. United Nations, New York, 1980.

Rinman, Thorsten and Nilsson, Rolf P. (2004), Trampsjöfart. Så fungerar fraktmarknaden. *Sjöfartens bok*, 21-26. Svensk Sjöfarts Tidning, Göteborg.

Ronen, David (1981), The potential of commercial intelligence in the ocean shipping industry. *Information & Management 4*, 127-131.

----- (1983), Cargo ships routing and scheduling: survey of models and problems. *European Journal of Operational Research 12*, 119-126.

----- (1993), Ship scheduling: the last decade. *European Journal of Operational Research 71*, 325-333.

Säntti, Auvo A. (1952), *Railway Traffic in Finland from Centres of Population to Export Ports in 1948*. Turun Yliopiston julkaisuja A XIII.

Sargent, A.J. (1930), *Seaways of the Empire*, 2nd ed. A.C. Black, London.

Schwartz, Peter (1996), *The Art of the Long View*, paperback. Doubleday.

Scurr, Joy (ed) (1999), *Major Coal and Ore Ports*, 5th ed. SSY Consultancy nad Research, London.

Serghiou, Serghios S. and Zannetos, Zenon S. (1982), The level and structure of single voyage freight rates in the short run. *Transportation Science 16* (1), 19-44.

Serjeantson, Richard (ed) (1997), *Iron and Steel Works of the World*, 12th ed. Metal Bulletin Books, Worcester Park, Surrey.

Sewell, Tom (1999), *Grain Carriage by Sea*. LLP, London.

Shelley, Toby (2003), China's rapid expansion boosts world shipping industry. *Financial Times*, 25 November, 3.

Shiprepair Directory (1998), A Motor Ship Supplement, September.

Solman, Paul (1998), Liffe to alter Biffex dry cargo futures. *Financial Times*, 28 October, 28.

Stopford, Martin (1993), *VLCC Tracking Survey*. Clarkson Research Studies, London.

----- (1997), *Maritime Economics*. Routledge, London.

Tamvakis, Michael N. and Thanopoulou, Helen A. (2000), Does quality pay? The case of the dry bulk market. *Transportation Research Part E 36*, 297-307.

Tanimoto, Hironori (ed) (1992), *Distance Tables for World Shipping*, 8th edition. The Japan Shipping Exchange, Tokyo.

Timmerman, K.W. and McConville, J. (1996), An analysis of the quality and redistribution of dry capesize tonnage. *Maritime Policy and Management 23* (1), 45-53.

Tinbergen, Jan and Buys, B.G.F. (1934), Tonnage and freight. *De Nederlandsche Conjunctuur*, (March), 23-35. In: L.H. Klaassen, L.M. Koyck and H.J. Witteveen (eds), Jan Tinbergen Selected Papers, 93-111. North-Holland, Amsterdam.

Tvedt, Jostein (2003), A new perspective on price dynamics of the dry bulk market. *Maritime Policy and Management 30* (3), 221-230.

UNCTAD (1997), *Recent and Planned Changes in Production Capacity for Bauxite, Alumina and Aluminium*. 23 December.

Ungar, A. (1973) Market areas for ocean shipping. *International Journal of Physical Distribution*, 155-160.

Utstein Kloster, K. (1948; 1952), *Oversjøisk trampfart I-II*. Gyldendal, Oslo.

Watson, Charles (ed) (1998), *Fairplay Ports Guide 1999-2000*, Volumes 1-4. Fairplay Publications, Coulsdon.

Wijnolst, Niko and Wergeland, Tor (1996), *Shipping*. Delft University Press.

World Bulk Fleet (various years), January and July. Fernresearch, Olso.

World Bulk Trades (various years), Fernresearch, Oslo.

Veenstra, Albert Willem (1999a), *Quantitative Analysis of Shipping Markets*. Delft University Press.

----- (1999b), The term structure of ocean freight rates. *Maritime Policy and Management 26* (3), 279-294.

Viscencio-Brambilla, Hector and Fuller, Stephen (1986), Effects of port user fees on export-grain flow patterns. *Southern Journal of Agricultural Economics 18* (2), 25-37.

----- (1987), Estimated effect of deepened US Gulf ports on export-grain flow patterns and logistics costs. *Logistics and Transportation Review 23* (2), 139-153.

Xie, Xinlian, Wang, Tengfei and Chen, Daisong (2000), A dynamic model and algorithm for fleet planning. *Maritime Policy and Management 27* (1), 53-63.