NBER WORKING PAPER SERIES

ON THE DURATION OF TRADE

Tibor Besedeš
Thomas J. Prusa

Working Paper 9936
http://www.nber.org/papers/w9936

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2003

We would like to thank the seminar participants at Rutgers University, Louisiana State University, the Midwest International Trade Meetings, the NBER ITI group and also Simon Evenett and Roger Klein for helpful suggestions and comments. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

On the Duration of Trade
Tibor Besedeš and Thomas J. Prusa
NBER Working Paper No. 9936
August 2003
JEL No. F14, F19, C14, C41

ABSTRACT

This paper employs survival analysis to study the duration of US imports. We find that the median duration of exporting a product to the US is very short, on the order to two to four years. Our results also indicate that there is negative duration dependence meaning that if a country is able to survive in the exporting market for the first few years it will face a very small probability of failure and will export the product for a long period of time. This result holds across countries and industries. We find that our results are not only robust to aggregation but are strengthened by aggregation. That is, as we aggregate from product level trade data to SITC industry level trade data the estimated survival increases. We rank countries by their survival experience and show that our rankings are strongly correlated with the rankings in Feenstra and Rose (2002), implying that product cycle followers also experience particularly short duration.

Tibor Besedeš
Louisiana State University
Department of Economics
E.J. Ourso College of Business
Baton Rouge, LA 70803
besedes@lsu.edu

Thomas J. Prusa
Rutgers University
Department of Economics
New Jersey Hall
New Brunswick, NJ 08901-1248
and NBER
prusa@econ.rutgers.edu

1. Introduction

Positive trade theory usually asks questions which address the "who, what, when, and why" of international trade. One question that is not often addressed is "how long?" When countries trade, how long do their trade relationships last? Are they exchanging products over long or short periods of time?

In order to answer these questions we study the duration of trade relationships. To our knowledge we are the first to do so and several important findings emerge from our analysis. First, we demonstrate not only that there is a remarkable amount of entry and exit in the US import market but also that the period of time a country is "in" the market is often fleeting. We find that some trade relationships are long-lived, by which we mean a country exports a product for many years. We find, however, that it is far more common to observe short-lived trade relationships wherein a country trades a product for a few years and then stops. In particular, more than half of all trade relationships are observed for a *single* year and 85 percent are observed for fewer than five years.

Second, because our analysis is performed at a highly disaggregated (7-digit) level we show one must carefully account for censored observations. As we explain censoring issues are more severe when working at the product level because of reclassifications. Once we account for censoring we can estimate survival and hazard rate functions of exporting to the United States. We find the median duration that a country exports a product to the US is very short, somewhere between two and four years depending on the censoring approach employed. Moreover, although duration does vary by source country and region, we show that short-lived relationships characterize trade by most countries. Said differently, we find short relationships are prevalent for both OECD and non-OECD countries although we find systematically higher survival experience for OECD trade relationships.

Third, we find that our results are not only robust to aggregation but in fact are strengthened by aggregation. That is, as we aggregate from product level trade data to SITC industry level trade data the estimated survival decreases. We show that this paradoxical result is related to the unique censoring problems that are only present when the analysis is done at the product level. As a result, it is more straightforward to conduct industry level analysis although doing so makes it somewhat more difficult to interpret what the relationship represents. It should be noted, however, that the results across SITC aggregations are intuitive with higher aggregations having longer durations. The aggregation exercise confirms that our main finding — that trade relationships tend to be quite short — is not an anomaly.

Fourth, our results also indicate that there is negative duration dependence, meaning that the conditional probability of failure decreases as duration increases. Conditional on observing a trade relationship surviving the first five years, the hazard rate for the rest of the sample is just 12 percent. The results suggest the presence of a type of a threshold effect where once a relationship is established and has survived the first few years it is likely the relationship will survive a long time.

All things considered, the results in this paper suggest that more is happening at the micro-level than suggested by either existing theory or empirical studies. For instance, according to the factor proportions theory trade is based on factor endowment differences, and since such differences change gradually trade patterns are likewise expected to evolve slowly. Unhappiness with this implication spurred the development of trade models with richer trade dynamics (Krugman, 1979; Dollar, 1990; Grossman and Helpman, 1991). In this literature Vernon's (1966) seminal product cycle theory is probably the best known. Vernon's model generates a particular pattern for trade relationships. The technological leaders develop and export a product until others learn how to manufacture it and enter the market. As the technology becomes more and more standardized, other countries will begin to produce and export the product. If these follower countries have relatively low costs, they will eventually take over the market and push out the leaders. All of these models, however, imply a fairly predictable pattern of trade dynamics — trade dynamics either evolve slowly or in logical progression from developed to developing countries; none of the models suggest that short-lived trade relationships would be the norm.[1] None of the models suggest that trade relationships would be as fleeting as implied by our results.

Until relatively recently there has been little direct evidence of the product cycle.[2] Gagnon and Rose (1995) and Feenstra and Rose (2000) stand out for their use of disaggregated industry data to look for systematic evidence of trade dynamics. The results from these papers have different implications for trade dynamics. Gagnon and Rose study the persistence of trade balances and obtain results that support the view that trade dynamics are gradual. They find that industries characterized by trade surpluses (deficits) in the early-1960s were extremely likely to also be characterized by trade surpluses (deficits) 25 years later. Gagnon and Rose conclude that "there is little evidence of dynamic patterns" in trade (p. 230).

Feenstra and Rose take a different approach toward identifying the product cycle and provide compelling support for Vernon's hypothesis. Rather the modelling the entire product cycle dynamics, Feenstra and Rose identify the phenomenon by focusing on a single datum: the date that a country first exports a product to the United States. Their work is an excellent example of how one can gain great insight by simplifying the task at hand. Using panel data for US imports, they find certain countries tend to be first to export a product while others tend to begin to export at a later date. They also find that the order in which countries first export a good is consistent with macroeconomic rankings typically associated with technological progress which further

[1] In Grossman and Helpman's (1991) "quality ladder" variant of the product cycle model the leader and follower exchange dominance of exports of a particular good over time, as leaders re-enter the market for a given good by innovating and offering a more advanced version.

[2] The earlier empirical evidence in support of the product cycle was based on indirect measures (Deardorff, 1984; Audretsch, 1987; Cotsomitis et al., 1991).

supports Vernon's hypothesis. Feenstra and Rose conclude that "there is strong evidence of intuitive ordering of countries and commodities, consistent with the product cycle (p. 378)."

Our findings both challenge and complement these earlier results.[3] On the one hand, we find that even though a trade balance perspective may indicate little or no dynamics, each competitor (supplier) in the marketplace may perceive the market to be very fluid and highly dynamic. In addition, our approach shows that entry, whether early or late, can be associated with a highly ephemeral trade relationship. On the other hand, our estimated survival functions do imply that some countries systematically perform better than others. We then take the next logical step and use our estimated survival functions to rank each country's survival experience. We show that our rankings are consistent with Feenstra and Rose's rankings. In other words, our results confirm the notion that there is a product cycle and that leaders not only tend to be the first to the market (Feenstra and Rose) but also tend to have longer spells of service. The results suggest that later in the product cycle products become standardized and competitive market pressures create short-term relationships.

This paper presents, to the best of our knowledge, the first attempt to examine the issue of the duration of trade relationships. The issue is stimulated in part by the findings of Haveman and Hummels (2001), Feenstra and Rose, and Schott (2001) who document that in any given year and for any given product, many countries do not trade, which means there are a lot of zero trade observations.[4] While this is an important realization, none of these previous studies examined whether a country's current state (i.e., being "in" or "out" of the market) was a permanent feature. We show that being in the market is often a transitory phenomenon.[5]

[3] Moenius and Riker (1999)'s findings are also related to this paper. They examine the empirical relevance of the quality ladders model of Grossman and Helpman and confirm most of the volatility implications of the model.

[4] Feenstra and Rose's approach depends on the fact that many of the zero value observations become positive in later years.

[5] Other research on dynamics of trade has focused on examining the evolution of patterns of industrial specialization over time rather than the dynamics of trade per se (Proudman and Redding, 1998; Redding, 2002).

The paper proceeds as follows. The next section presents the data, and offers an example. We also examine the distribution of the length of trade relationships and discuss some complications such as multiples spells and censoring. Section 3 discusses modeling of duration and presents estimation techniques. Section 4 presents estimates of survival and hazard rate functions and discusses differences across countries and products as well as the effects of aggregation. In section 5 we examine the robustness of our results. We show that our findings are not particularly sensitive to changes in how we measure trade relationships, define failure, or level of aggregation. In fact, we find that the median duration is often shorter in these alternative formulations than for our benchmark data. The most significant changes occur when we give greater weight to relationships with large trade values. While a weighted analysis does increase duration, we still find that 50 percent of trade weighted relationships are observed for fewer than five years. In section 6 we relate our findings to Feenstra and Rose's country rankings.

2. Data

Our analysis is based on US import statistics as compiled by Feenstra (1996).[6] From 1972 through 1988 import products were classified according to the seven digit Tariff Schedule of the United States (TS).[7] Since 1989 imports have been classified according to the ten digit Harmonized System (HS). Given that all products were recoded in 1989, we limit our analysis to the period 1972-1988 in order to avoid potential concordance issues.[8] In simplest terms, for each commodity and year we can identify all of the countries from whom the United States purchased. Over the entire period there is a total of about 22,000 different products in the TS

[6] The data is available from http://data.econ.ucdavis.edu/international/.

[7] We only study imports because there is no concordance between the disaggregated import and export codes. In addition, the export data is self-reported making it more likely that exports are misreported.

[8] Many TS products are mapped into multiple HS codes and vice-versa making impossible to discern actual exits. See Schott (2001) for more discussion of the TS-HS concordance problem.

data. On average, in each year we observe import trade for about 10,000 products sourced from about 160 countries.

Trade relationships and spells of service

We are interested in studying the length of time until a country ceases to export a product to the US, an event that we will refer to as a "failure." As a result calendar time is not as important as analysis time, which measures the length of time that a country exports a product to the US. Hence, for each product and country we use the annual data to create *spell* data. By this we mean that if we observe trade in product i from country c from 1976–1980, we would say that the ci^{th} trade relationship has a spell length of five.

Most of our analysis is based on the most disaggregated data available, 7-digit TS data, which means our analysis is at the product level, not the industry level. Inferences are based on trade in tangible products rather than aggregate summaries. Until relatively recently such disaggregated data was not widely available and empirical studies were based on industry classifications, such as the Standard International Trade Classification (SITC). We chose to base our benchmark results on the 7-digit TS data primarily because it does the best job of measuring trade in a specified product. For instance, the TS data reports information on more than 30 different types of ball bearings, differentiating by specialized application (e.g., automobiles), size, and chemistry. By contrast, if we were to perform our analysis using SITC classifications (as we do later in our robustness checks) we would aggregate all of these ball bearing codes along with other similar, or perhaps not so similar, products to create spells of service. For instance, the dozens of ball bearings codes map into a single SITC category "Ball, roller or needle roller bearings" (4 digit code=7491). Further aggregation is possible, say to "Non-electric machinery parts" (3 digit code=749) or "Industrial machinery" (2 digit code=74), but doing so makes it increasingly difficult to interpret the results since each classification includes highly disparate products.

For duration analysis we believe that using highly disaggregated data is imperative. First, the more aggregated is the data, the more the analysis identifies industry trends rather than competitive dynamics at the product level. The fact that country c in industry j has a long duration tells us little about the duration of commodity trade or underlying trade dynamics. For instance, suppose we observe three countries exporting “Industrial machinery” and that each has a long trade duration. This might lead us to surmise that the dynamics among these countries are very similar. This need not be the case. For instance, country c_1 might be uniformly superior: every product it sells could have a long duration. On the other hand, country c_2 might be a dismal failure: the success (i.e., long duration) of one product could hide a multitude of failures (short-lived spells). Finally, country c_3 could be a classic example of Vernon’s product cycle: the observed long duration at the industry level could reflect a progression from simpler products (say, ball bearings) to very complicated products (say, engines). If nothing else, the more aggregated is the data the more cautious we should be interpreting the results; a less munificent view might be that aggregated data makes it impossible to measure dynamics.

Second, if the products are too broadly defined, we cannot expect to see any source countries exit the data. For instance, at the extreme if we aggregate all imports from each country, we will never observe any exit since the United States purchases some product from each source country every year. Simply put, higher aggregations can hide a great deal of dynamics and competition in the import market.

For each product we create a panel of countries which export the product to the US. Table 1 provides an example of what the dataset looks like for a representative product, “Milled Corn (TS=1312000).” The “X”s in the table indicate years in which each country exports the product to the US. As shown, there are countries that export corn every year, such as Canada and Portugal. This pattern illustrates a spell of length 17 years. Another dozen or so countries have one spell. Mexico begins exporting corn in 1974, South Korea in 1976, and Peru in 1984, and all three service the US market in every year after entry. The remaining countries with one spell all export corn for just one year.

There are several countries with two spells. The United States imports corn from Venezuela in every year but 1981. Ecuador exports corn in 1973 and 1974 (the first spell, length 2) and then services the US market for the second time starting in 1979 and continuing through 1988 (the second spell, length 10). The Dominican Republic also has two spells, but both of length 1. Colombia and Italy each have three spells, both starting with two short ones (of length 3 and 1, and 1 and 2, respectively), before entering the US market for the third time (1981 and 1980, respectively) and staying for the remainder of time. Of the remaining countries most have very short spells, with the exception of the Netherlands whose first spell is five years long. However, its remaining four spells are all just one year long.

Once we think of the data in terms of spells our data set has 693,963 observations (see Table 2). The observations have a median (mean) spell length of 1 (2.7) years. We will discuss how aggregation affects our results in more detail in section 5, but we note here that Table 2 also includes statistics when we use the more aggregated SITC industry classification system to construct trade relationships.[9] We have 157,850 observations when we use the 5-digit SITC classification, 90,035 observations when we use the 4-digit level and just 2,484 observations at the 1-digit level. As expected, when we aggregate the data our ability to observe entry and exit diminishes and the mean spell length increases, from 2.7 years in the TS data to 3.9 years in the 5-digit SITC data to 8.4 years in the 1-digit SITC data.

In Figure 1 we plot the distribution for each observed spell length. On the x-axis we plot the observed spell length. On the y-axis we plot the percentage of observations whose observed spell of service is greater than a given length. In the benchmark 7-digit TS data, more than half

[9] SITC imports are also reported in Feenstra (1996).

of all spells are just one year long. About 70 percent of the spells are observed for two or fewer years. About 80 percent of the spells are observed for three or fewer years. On the opposite side of the spectrum less than two percent of all spells are observed for all 17 years. Simply put, there is a plethora of short spells. In Figure 1 we also plot the distribution for alternative measures of trade relationships; these other distributions will be discussed later in the paper.

Multiple spells

Looking at Table 1, one sees that some trade relationships re-occur or have multiple spells of service. That is, a country will service the market, exit, then re-enter the market, and then almost always exit again. As it turns out, about a quarter of trade relationships experience multiple spells of service. About two-thirds of those trade relationships with multiple spells experience just two spells; less than one percent of trade relationships with multiple spells have more than three spells. All told we have 495,763 trade relationships as compared with 693,963 spells of service.

We will begin by treating multiple spells as independent. While this assumption is made primarily in the interest of simplicity in section 5 we explore alternatives and find our results are likely not overly sensitive to this assumption. To get a sense of why this is the case, it is instructive to look at the distribution of spells when we exclude multiple spells. First let's restrict the analysis to those trade relationships with just a single spell. As reported in Table 2, these observations have a median (mean) spell length of 1 (3.2) years; thus, the median is exactly the same as the benchmark data (all 7-digit TS observations) and the mean is just a half a year longer. In Figure 1 we show the entire distribution. In both cases we see that more than half of all spells are just one year long and about 75 percent of the spells are observed for fewer than four years; further, in both distributions only about 5 percent of the spells are observed for 10 or more years.

Alternatively, we can restrict the analysis to just the first spell of a trade relationships. This alternative includes those relationships with just a single spell and also the first spell of those relationships with multiple service spells. As was the case in the first comparison, the distribution is very similar to the benchmark. This similarity among these distributions suggests that the independence assumption is a reasonable starting place.

Censoring

Once we begin to think of the data in terms of spells it becomes apparent that we need to account for censoring in our analysis. By this we mean that we often do not know whether a trade relationship ends because of failure or for some other reason. In practice this means that we are uncertain about the beginning and/or ending date for some trade relationships.

As it turns out, censoring is pretty common in US import data — about half of the trade observations are censored and about 20 percent of spells are censored at one year. The censoring problem comes in two flavors. First, we have no information on trade relationships for the years before the beginning (pre-1972) and after the end of our sample (post-1988). As seen in Table 1 we observe that the US imported corn from the Philippines in 1972 and that this relationship was observed for exactly one year. This trade relationship may have begun in 1972 or it may have begun in some prior year. The most appropriate interpretation is that this relationship had a duration of at least one year. Similarly, we also observe that the US imported corn from Peru from 1984 to 1988. Unfortunately, our data does not continue beyond 1988 and hence we cannot be sure how long that spell ultimately lasted. Once again, the most appropriate interpretation is that this relationship had a duration of at least five years. About 10 percent of spells are observed in 1972, while about 22 percent are observed in 1988.[10] This first type of censoring is typical in survival studies and is incorporated in all of our subsequent analysis.

[10] Remember, less than two percent of all spells are observed both in 1972 and 1988.

The second type of censoring is somewhat unique and stems from the fact that our analysis is done at the product level. US Customs revises product definitions for the tariff codes on an ongoing basis, sometimes splitting a single code into multiple codes and other times combining multiple codes into fewer codes. Unfortunately, information does not exist to allow us to map the old TS codes into the new TS codes. When a code is changed we no longer observe trade in the old product code. But, is this due to the end of a relationship or does it simply mean that trade stopped due to the redefinition? Throughout much of our analysis we choose to be cautious and classify all such "exits" as censored, which means that we interpret reclassified relationships as having duration of at least x years (where x is the number of years where trade in the original code was observed).

An analogous problem exists for the new codes. When a code is changed we begin to observe trade in the new product code. But, is this really the beginning of a new relationship or does it simply reflect the redefinition? Once again for much of our analysis we choose to be cautious and assume the relationship had a duration of at least y years (where y is the number of years where trade in the new code was observed). About 20 percent of the spells are censored due to reclassification. This second type of censoring is a peculiar characteristic of our particular application; it is incorporated into what we refer to as our benchmark estimates.

The issue is that our decision to interpret all product code changes as being censored is very conservative — certainly some of the new codes were truly measuring new products and some of the obsolete codes were measuring products that were no longer traded. In other words, in our benchmark data we have probably classified some spells as being censored when they were truly associated with either an entry or an exit. If more information were available we would be able to identify such births and deaths.[11] Unfortunately, this additional information does

[11] Xiang (2002) uses the changes in the verbal description of 4-digit SIC industry codes to identify new products. Even using relatively aggregated industry data (about 450 industries) Xiang describes a time-consuming and painstaking process.

not exist. This means that our benchmark results will overstate the true duration of the typical trading relationship.

An alternative that we explore is to classify these product code changes as indicating entry and exit. It should be noted, however, that this alternative classification scheme will undoubtedly understate the true duration. As a result, the "true" duration will lie somewhere between these two approaches toward censoring.

3. Modeling Duration

Duration models

Let T denote the time to a failure event. Since time in our analysis is discrete we assume T is a discrete random variable taking on values t_i, $i = 1, 2, \ldots, n$ with probability density function $p(t_i) = \Pr(T = t_i)$, $i = 1, 2, \ldots, n$ where $t_1 < t_2 < \cdots < t_n$. The survival function for a random variable T is given by

$$S(t) = \Pr(T > t) = \sum_{t_i > t} p(t_i).$$

The hazard function is

$$h(t_i) = \Pr(T = t_i \mid T \geq t_i) = \frac{p(t_i)}{S(t_{i-1})}, \; i = 1, 2, \ldots, n$$

where $S(t_0) = 1$. The survival and hazard functions are related through the following expression[12]

$$S(t) = \prod_{t_i < t} [1 - h(t_i)].$$

[12] Keifer (1988) provides a good survey of duration models, while Hosmer and Lemeshow (1998), Lancaster (1990), and Lee (1992) provide good introductions to survival analysis.

Nonparametric estimation

The two functions of interest in this paper are the survivor and hazard functions. To estimate them we will assume we have n independent observations denoted (t_i, c_i), $i = 1, 2, ..., n$, where t_i is the survival time and c_i is the censoring indicator variable C (taking on a value of 1 if failure occurred, 0 otherwise) of observation i. Assume there are $m \leq n$ recorded times of failure. Denote the rank-ordered survival times as $t_{(1)} < t_{(2)} < ... < t_{(m)}$. Let n_i denote the number of subjects at risk of failing at $t_{(i)}$ and let d_i denote the number of observed failures. The Kaplan-Meier product limit estimator of the survivor function is then

$$\widehat{S}(t) = \prod_{t_{(i)} \leq t} \frac{n_i - d_i}{n_i},$$

with the convention that $\widehat{S}(t) = 1$ if $t < t_{(1)}$. The Kaplan-Meier estimator is robust to censoring and uses information from both censored and non-censored observations. The standard error of the Kaplan-Meier estimator is given by Greenwood's formula

$$\widehat{Var}\{\widehat{S}(t)\} = \widehat{S}^2(t) \sum_{i | t_i \leq t} \frac{d_i}{n_i(n_i - d_i)}.$$

However, when calculating confidence intervals the asymptotic variance of $\ln[\ln -\widehat{S}(t)]$ is used instead:

$$\widehat{\sigma}^2_S(t) = \frac{\sum \frac{d_i}{n_i(n_i - d_i)}}{\{\sum \ln(\frac{d_i}{n_i(n_i - d_i)})\}^2}.$$

The hazard function is estimated by taking the ratio of subjects who die in a given year to the number of subjects alive in that year,

$$\widehat{h}(t) = \frac{d_i}{n_i}.$$

4. Empirical Results

7-digit TS data

We begin by examining the benchmark 7-digit TS data dataset. In Table 3 we report the distribution of the Kaplan-Meier estimated survival time. Begin by looking under the columns labelled "TS7-benchmark." The row "Total" shows the overall distribution of survival time. The other rows report the Kaplan-Meier estimates for particular regions and industries.

First note the difference between Kaplan-Meier estimated survival time and the spell lengths reported in Table 2: the Kaplan-Meier estimates are significantly longer. The difference, of course, is that the Kaplan-Meier estimator accounts for censoring while the spell lengths reported in Table 2 do not; this comparison highlights the importance of accounting for censoring when working with product level data.

The second important finding involves how sharply the risk of failure falls. The analysis indicates that the risk of failure is quite high in the early years, but then rapidly falls once a trade relationship survives a threshold duration. The 25^{th} percentile of the distribution is one year, which means that after the first year, more than 25 percent of the spells fail. The median duration is just four years, indicating that a considerable number of observations fail over the following three years. After about four or five years, however, failure becomes a lot less common: the 75^{th} percentile spans the whole data set (which we denoted by a period). In other words, more than 50 percent of the observations fail within the first four years, but over the next 13 years less than one quarter of the observations fail. This suggests that countries exporting to the US face a large conditional probability of failure in the early stages of their trade relationship and a much smaller conditional probability once they have survived a few years.

The estimated overall survival function, $\widehat{S}(t)$ is graphed in Figure 2 together with the 95% confidence interval. The confidence interval is very tight (imperceptible), which is not surprising given the size of the data set used. The survival function is downward sloping with a decreasing slope. The estimated survival function suggests that the hazard rate should be a decreasing function and this is confirmed in Figure 3 where we present the estimated hazard function. As shown, the probability of failing in the first year is 33 percent. The hazard rate between year one and five is an additional 30 percent. The hazard rate for the remaining twelve years is just 12 percent.

Figures 2 and 3 indicate that after a large decrease in the probability of survival over the first few years, the survival function levels off and decreases very little. Both the survival and hazard functions indicate that there is negative duration dependence, meaning that the conditional probability of failure decreases as duration increases. This result indicates the presence of a type of a threshold effect. In other words, it appears that once a relationship is established and has survived the first few years it is likely the relationship will survive a long time. This finding is similar to that documented by Pakes and Ericson (1998) for retail trade establishments.[13]

It is somewhat surprising that the Kaplan-Meier estimate of probability of exporting a product for more than 17 years is so high (0.41) given that we previously showed that less than two percent of all trade relationships span the entire sample (Figure 1). These seemingly inconsistent findings are explained by the prevalence of censoring when using the product level data. Many of the relationships that are observed to end are censored and thus in our benchmark results we do not classify them as failures.

The impact of censoring due to product codes changes can be identified if we estimate the Kaplan-Meier survival function using the modified censoring approach discussed above. Under this alternative scheme we consider all changes in TS codes as indicating births and deaths which means we will observe much more entry and exit in the data. We report the results

[13] Pakes and Ericson interpret their results as supporting Jovanovic's (1982) theory of learning.

in Table 2 and Figure 4.[14] In Table 3 we see that under this modified censoring scheme the median duration falls to just two years as compared with four years in the benchmark data; the 75th percentile is just six years under modified censoring. In other words, at least 75 percent of the spells fail within six years. In Figure 4 we depict the Kaplan-Meier survival function for both the benchmark 7-digit TS data and for the modified censoring data. As expected, the hazard rate is substantially higher during the first years under modified censoring; by year four the survival rate is about 20 percentage points smaller under the modified censoring scheme.[15] Nevertheless, as was the case with the benchmark data, the hazard rate falls sharply after about four or five years and there is little additional failure. Overall, we find that the probability of exporting a product for more than 17 years under modified censoring is only 0.18 — less than half of that of the benchmark data — but still considerably higher than the observed two percent of trade relationships that span the entire sample.

Region and industry survival functions

In Table 3 we also report statistics broken out by individual regions and industries for both the benchmark and modified censoring data. For all regions we observe that a substantial fraction of trade relationships quickly end in failure: the 25th percentile of the distribution is one to two years for all regions. There are differences across regions in the severity of the initial wave of failures with Northern regions tending to outperform Southern regions. When using the benchmark data we find that the Northern regions (as represented by Canada, Asia, and Western Europe) have a median duration of at least six years; of this set, Canada has by far the longest median survival time, which spans the whole data set. In other words, less than half of

[14] Recall that in this modified scheme products traded in 1972 or 1988 are still classified as censored.

[15] As we showed previously, there is a connection between the hazard function and the survival function: when the probability of survival decreases significantly the hazard rate is large; when the probability of survival decreases slightly the hazard rate is small.

the observations on Canadian trade end in failure within the our 17 year sample.

Quite a different story emerges for the Southern regions. Using the benchmark data, the Southern regions have a median duration of no more than two years — half of all trade relationships from these regions will fail within two years.[16] We also point out that the difference between Northern and Southern regions is also observed under the modified censoring scheme, where once again the Northern regions have uniformly higher median survival times.

In Figure 5 we use the benchmark data and estimate the survival functions for individual regions. As one can see the shapes of the survival functions are similar and all reflect the basic dynamics described above. For each region there is a significant probability of failure in the first four or five years of the relationship, but thereafter the hazard falls substantially. Nevertheless, survival varies substantially across source region. In the case of Canada about 60 percent of trade relationships survive 10 or more years. By contrast, in the case of Africa only 20 percent of trade relationships survive 10 or more years.

The results clearly suggest that duration can broadly be thought of as a North versus South story, with Northern sources having longer median survival times. This insight can be drawn more sharply by estimating the Kaplan-Meier survival functions for OECD and non-OECD countries separately. We present the results in Figure 6. On average, the estimated survival for OECD countries is about 10 percentage points higher than for non-OECD countries. Thus while the survival pattern is qualitatively similar across regions, Northern countries clearly fare better.

When we look at survival functions by industry we again see broadly similar trends — a wave of initial failures which are followed by a significant decrease in hazard. As seen in Table 3 and depicted in Figure 7 the 25$^{\text{th}}$ percentile of the distribution is one year for every industry.

[16] The category "Other" contains the following countries that we did not assign to any region: Bermuda, Greenland, St. Pierre Miquelon, and Unknown Partner.

Looking at the benchmark data, with the exception of two industries the median survival time is three years or less. In other words, in most industries more than half of all observations fail within two years.

In two key industries, Machinery (SITC 7) and Miscellaneous Manufactures (SITC 8), the estimated duration is far longer, more than 11 years. An explanation for the distinct difference between these two industries and the others lies at least in part to the apparently high prevalence of product reclassification censoring that occurs in these two industries. Evidence of this can be found when we look at the distribution under our modified censoring scheme where we observe that these two industries have a median survival time of two years, the same as all but one industry (whose median survival time is one year). The higher than average incidence of product codes changes might reflect that these two industries are characterized by the most innovation and shortest product cycles; a dynamism which results in numerous TS codes changes. If this is the case, then the modified censoring scheme probably does a better job capturing the actual dynamics as the code changes probably reflect true entry and exit (i.e., births and deaths). On the other hand, the high incidence of code changes might reflect other factors that are not associated with economically meaningful entry and exit. For instance, the code changes might simply be a way to mitigate the impact of GATT negotiated tariff reductions or expand the scope of the Multifibre Agreement. In this case, the benchmark censoring scheme is probably more appropriate, implying that these two industries really are outliers. A resolution of this issue is beyond the scope of this paper but certainly merits further analysis.

5. Robustness of results

The results indicate that trade relationships are quite short: estimated median survival is between two and four years. We now examine whether our findings are robust to how we measure trade relationships and spells.

Aggregation concerns — SITC level analysis

We begin by exploring whether our finding is possibly due to the highly disaggregated nature of our data. Perhaps 30 different ball-bearing codes represents an overly fine parsing of the data that leads us to observe excessive entry and exit.[17] If so, a trade relationship might be better measured using SITC industry classifications.

To address this concern we also calculated spells of service using SITC industry (revision 2) definitions ranging from the 1-digit to the 5-digit level. As mentioned above, summary statistics for the SITC classifications are given in Table 2 where we show that the median observed spell length is two or fewer years in the 1-digit SITC through the 4-digit SITC data. The distribution of duration time based on the Kaplan-Meier estimates are presented in Table 3; the estimates themselves are depicted in Figure 4.

A number of interesting insights are gained by examining trade relationship using industry-level data. First, note that within the SITC classifications aggregation works exactly as expected. Higher levels of aggregation are associated with longer survival times. As depicted in Figure 4 the survival function for the 4-digit SITC lies above the 5-digit SITC, the 3-digit SITC lies above the 4-digit SITC, and so on. Nevertheless, although aggregation creates longer survival times the estimated survival times remain remarkably brief. There remains a large number of short spells of service until the data is fairly highly aggregated. In particular, as shown in Table 3, the median survival time is two years for the 3-digit, 4-digit, and 5-digit SITC data.

Second, in comparison with the median survival time for the 7-digit TS benchmark data the brevity of duration times for the SITC data is somewhat surprising. Specifically, despite the

[17] It is not obvious that the TS classification system is too fine. To the contrary, Schott (2000) presents evidence that the TS categories may be too broad.

higher degree of aggregation the 2-digit through 5-digit SITC data all have a shorter median survival time than the benchmark 7-digit TS data. As depicted in Figure 4 only the 1-digit SITC data has a *higher* survival function than the 7-digit TS benchmark.

This result seemingly implies that aggregation from products to industries lowers survival. Taken at face value, this is an odd result. This unusual finding, however, has a completely logical explanation. The SITC classification system (revision 2) is unchanged throughout our sample. As a result only the first type of censoring (which is driven by the beginning and end of our sample) is present in the SITC analysis. In other words, it is more appropriate to compare results based on the SITC classifications with the 7-digit TS modified censoring data. As shown in Figure 4 when we do that comparison we get the expected result: survival experience for product level data is lower than for the 5-digit SITC industry level data and aggregation increases estimated survival time.

In effect, the SITC data shows that both approaches toward handling censoring in the 7-digit TS data are imperfect. In an ideal world we would be able to truly identify all births, deaths, and censored product changes. But, since we cannot do so we are left with two imperfect measures: (i) the 7-digit TS benchmark data which overstate the "true" survival experience (since all product code changes are censored) and (ii) the modified censoring data which understate the "true" survival experience (since all product code changes are interpreted as births and deaths).

Moreover, given that survival experience at the product-based level must be shorter than at the 5-digit SITC industry level, the results shown in Figure 4 indicate that the true 7-digit TS survival experience indeed lies between our two measures. In effect, the SITC results imply that the "true" median survival experience at the 7-digit TS level is two to four years and probably is closer to two years. In either case, the evidence is clear: duration is very brief.

Multiple Spells

As mentioned, about a quarter of trade relationships experience multiple spells of service. In the above analysis we have assumed that duration is independent of spell number. We now investigate whether this assumption biases our findings.

To check the sensitivity of our results, we considered several alternative approaches toward the issue of multiple spells. First, we simply limited our analysis to only those relationships with a single spell. In Figure 1 we plotted the distribution for trade relationships with a single spell. And as discussed above there is very little difference between the distributions for single spell and all 7-digit TS data — under this alternative formulation about half the observations have a spell length of just one year and more than 75 percent of the spells last three or fewer years. The Kaplan-Meier estimated survival function for the single spell data is depicted in Figure 8. The survival function for the single spell data has a similar pattern as the benchmark data, namely high hazard in the first few years followed by a levelling off of the survival function. Overall, however, the single spell data exhibits significantly higher survival: about 21 percentage points higher than the benchmark data in year 17.

The similarity of spell length distributions (Figure 1) sharply contrasts with the differences in the Kaplan-Meier estimates (Figure 8). This suggests that the censoring approach taken in the benchmark data plays a role. Specifically, censoring is more common in the single spell data which shifts up the Kaplan-Meier estimates. Moreover, as we discussed in the preceding section, analysis based on the SITC data indicates that the benchmark data overstates the amount of censoring in the data and that the modified censoring approach may yield a more appropriate measure of true survival experience. To get a sense of the magnitude we re-estimate the single spell data using the modified censoring approach and present the Kaplan-Meier estimates in Figure 8. As shown, the median survival time is now 3 years as compared with 2 years when all observations are included. While the single spell data has a higher survival experience, the effect is not as great as in the benchmark approach: about 7 percentage points higher than the benchmark data in year 17.

A second approach we explored was to limit the analysis to only first spells. This alternative includes those relationships with just a single spell and also the first spell of those relationships with multiple service spells. Since the results are generally similar to the single spell results they will not be explicitly discussed.

Another way we addressed the issue of multiple spells was to consider the possibility that some of the reported multiple spells are due to measurement error. Specifically, if the time between spells is short, it may be that the gap is mis-measured and that interpreting the initial spell as "failing" is inappropriate. In this case, it may be more appropriate to interpret the two spells as one longer spell. To allow for such misreporting, we assume that a one-year gap between spells is an error, and hence merge the individual spells and adjust the duration length accordingly. Gaps of two or more years are assumed to be accurate and no merging is done.

The spell length distribution for this gap-adjusted data is depicted in Figure 1. As seen this adjusting for this potential type of measurement error shifts the distribution but short spells remain the norm. In Table 2 we report the summary statistics for the gap-adjusted data. In comparison with the benchmark 7-digit TS data, the average spell length is longer by about a half year.

The Kaplan-Meier estimated survival function for the gap-adjusted data is depicted in Figure 8. As expected, the survival function exhibits less early failure and more failure in later times. For example, after two years the survival function for the benchmark data is 0.57 as compared with 0.66 in the gap-adjusted data. By contrast, the hazard rate for the last twelve years is just 12 percent in the benchmark data but 14 percent in the gap-adjusted data. As was the case with the single spell and first spell alternatives we find that using the modified censoring approach reduces the difference between the gap-adjusted and 7-digit TS data.

Taken together, these results suggest that the independence assumption likely leads to an under-estimate of duration although our analysis indicates that the magnitude of the bias is fairly small for most scenarios. In particular, if one believes the modified censoring approach best captures the true survival experience, the independence assumption does not appear to have a significant impact under any of these scenarios.

Trade Weighted Analysis

The final issue we explore involves putting more weight on the higher valued trade relationships. Our benchmark dataset — and in fact all of the preceding analysis — is unweighted. Small and large trade value relationships receive the same weight. If short spells involve small values and long spells involve large values, an unweighted distribution might overstate the brevity of spells. To examine this issue, we also compute a weighted distribution where each observation is weighed by the value of trade in the first year of the relationship. This alternative formulation de-emphasizes the low-value spells and gives more importance to the high-value spells.

In Figure 1 we plot the spell length distribution under this weighted alternative. Comparing it with the benchmark data we observe a significant shift in the distribution, far greater than any of the other alternatives. In the benchmark data, more than half the observations are observed for just a single year. In the weighted analysis, only about 15 percent of the data is observed for a single year This clearly suggests that lower value trade relationships tend to be shorted lived. Nevertheless, even under the weighted scheme there remain a surfeit of short spells — more than 50 percent of the spells are observed for four or fewer years.

In Figure 8 we present the Kaplan-Meier estimated survival function for the weighted 7-digit TS data. The impact on survival is far greater than under any of the alternatives previously considered. In particular, the impact on duration is more significant than aggregating to the

1-digit industry level or restricting ourselves to single spell relationships.

6. Duration and the product cycle

Feenstra and Rose (2000) rank countries according to the first year they export a product to the US. They find that the order in which countries first export a product to the US is consistent with macroeconomic rankings typically associated with technological progress. In other words, they find that the rankings are consistent with the product cycle.

We use duration methodology to compute country rankings. In other words, rather than base the rankings on the first year of export, we rank countries on their survival experience. In what follows, we report countries ranked according to the probability of exporting a product for more than four years ($\widehat{S}(4)$). We note, however, that given that the estimated survival functions rarely cross, the country rankings are not sensitive to what spell length we use.[18] In Table 4 we report countries rank using the 7-digit TS benchmark data, the 7-digit TS modified censoring data, and also the 5-digit SITC industry classification. We also reproduce the rankings reported by Feenstra and Rose.

The rankings based on estimated survival are similar for all definitions of trade relationships. As reported in Table 5 the pairwise correlations are about 0.90. Nevertheless, there are a couple of anomalies in the rankings based on the 7-digit TS benchmark data. Most notably, Macao has the 10th best survival experience, far higher than for our other two survival rankings or Feenstra and Rose. Sri Lanka, Indonesia, and Haiti are also ranked implausibly high. Trade from each of these countries appears to experience a high incidence of product code reclassifications as their rankings drop sharply under the other survival calculations. These results provide further evidence that censoring is a significant hurdle for analysis at the product level.

[18] The rankings would be virtually the same if we ranked the estimated survival functions in any year.

Overall, the results based on survival experience are quite supportive of those found by Feenstra and Rose. As seen in Table 5 the pairwise rank correlations between the two rankings from Feenstra and Rose (2000), goods and country based, and the rankings based on survival are high. The highest correlation is between Feenstra and Rose's ranking our ranking based on the 5-digit SITC data. This is likely due to the fact that their analysis is also performed at the 5-digit SITC level. Both rankings based on the 7-digit TS data are slightly less correlated with Feenstra and Rose and of the two, the benchmark data is somewhat more highly correlated.

An interesting lesson can be learned by having both Feenstra and Rose and our rankings. Feenstra and Rose show that developed countries tend to be the first in the market. Our results show that these same countries tend to successfully sell their products for a much longer duration than do the followers. Taken together, these results highlight an often overlooked value to being a leader. Namely, as a unique product becomes standardized and production moves to the South, competition among the followers is often fierce. As a result, exports in one year are no guarantee of further sales.

7. Conclusion

The type of highly dynamic trade patterns evidenced in this paper are unanticipated given implications from standard theoretical models and existing empirical work. Our results indicate that typical trade relationships are more fragile than previously thought — or at least, than has been previously discussed in the literature. In particular, we have never seen a discussion that trade relationships may systematically last only a year or two.

This paper is the first one to ask the question of how long trade relationships last. The surprising findings shed light on entry and exit in foreign markets. The results of this paper are not important just for the understanding of dynamics of international trade, but are also important for policy makers. Different policies can have important impacts on trade relationships, from creating them to extending them to ending them. The length of trade relationships adds another dimension to what policy makers must pay attention to.

The results indicate that most trade relationships are very short. We also show that the phenomena are not driven by the highly disaggregated nature of the data set. Aggregating the data results in significantly longer durations only at very high aggregation levels. While the actual duration increases with aggregation, the relative survival experience of countries and industries is not affected by aggregation. There is evidence that the product cycle theory can explain at least some features of the duration of trade as the countries with longest survival times are those who are conventionally defined as technological leaders.

Future work should focus on identifying and discriminating among potential explanations. The product cycle is a reasonable starting point but there are others. Insights and lessons from industrial organization literature are likely to be helpful. For instance, Ericson and Pakes (1995, 1998) develop a dynamic model of market entry and exit with some implications which are observed in this paper. They show that firms are more likely to die in the early years of their lives than later on, a finding we show is true for import trade relationships.[19] Alternatively, our findings could be consistent with a search model where the domestic buyers are searching for the best partners in other countries.

[19] Another potentially interesting starting point is the model developed by Asplund and Nocke (2003).

References

[1] Audretsch, D.B. (1987), "An empirical test of the industry life-cycle," *Welwirtschaftliches Archiv* 123:297–307.

[2] Asplund, M. and V. Nocke (2003), "Firm Turnover in Imperfectly Competitive Markets," mimeo, Department of Economics, University of Pennsylvania.

[3] Cotsomitis, J., C. DeBresson, and A. Kwan (1991) "A reexamination of the technology gap theory of trade," *Welwirtschaftliches Archiv* 127:792–99.

[4] Deardorff, Alan V. (1984), "Testing Trade Theories and Predicting Trade Flows," in R. Jones and P. Kenen (eds.) *The Handbook of International Economics, Vol. I* (Amsterdam: North-Holland).

[5] Gagnon, Joseph and Andrew Rose, (1995) "Dynamic Persistence of Industry Trade Balances: How Pervasive is the Product Cycle?", *Oxford Economic Papers*, 47:229-48.

[6] Grossman, G. and E. Helpman (1991), *Innovation and Growth in the Global Economy*, Cambridge: MIT Press.

[7] Feenstra, R.C. (1996), "US Imports, 1972-1994: Data and Concordances," NBER Working Paper 5515.

[8] Feenstra, R.C. and A.K. Rose (2000), "Putting Things in Order: Trade Dynamics and Product Cycles," *Review of Economics and Statistics*, 82:369-82.

[9] Haveman, J. and D. Hummels (2001), "Alternative Hypotheses and the Volume of Trade: Evidence on the Extent of Specialization," mimeo, Department of Economics, Purdue University.

[10] Hosmer, D.W. and S. Lemeshow (1998), *Applied Survival Analysis: Regression Modeling of Time to Event Data*, New York: John Wiley & Sons, Inc.

[11] Jovanovic, B. (1982) "Selection and the evolution of industry," *Econometrica*, 50:649-70.

[12] Keifer, N.M. (1988), "Economic Duration Data and Hazard Functions," *Journal of Economic Literature* 26:646-79.

[13] Kendall, M. and J. Dickinson Gibbons (1990), *Rank Correlation Methods*, fifth edition, New York: Oxford University Press.

[14] Lancaster, T. (1990), *The Econometric Analysis of Transition Data*, Cambridge: Cambridge University Press.

[15] Lee, E.T. (1992), *Statistical Methods for Survival Data Analysis*, New York: John Wiley & Sons, Inc.

[16] Moenius, J. and D. Riker (1999), “Trade Barriers and the Volatility of Comparative Advantage,” mimeo, Department of Economics, UCSD.

[17] A. Pakes and R. Ericson (1998), “Empirical Implications of Alternative Models of Firms Dynamics,” *Journal of Economic Theory*, 79:1-45.

[18] Proudman, J. and S. Redding (1998), “Persistence and Mobility of International Trade,” Chapter 2 in *Openness and Growth*, Bank of England, 1998.

[19] Rauch, J.E. (1999), “Networks versus Markets in International Trade,” *Journal of International Economics*, 48:7-35.

[20] Redding, Stephen (2002), ”Specialization Dynamics,” *Journal of International Economics*, 58:299-334.

[21] Schott, P. (2000), “Do Countries Specialize?”, mimeo, Yale University, mimeo.

[22] Vernon, Raymond (1966) “International Trade and International Investment in the Product Cycle,” *Quarterly Journal of Economics*, vol. 80, no. 2, May.

Figure 1 - Distribution of Spell Lengths

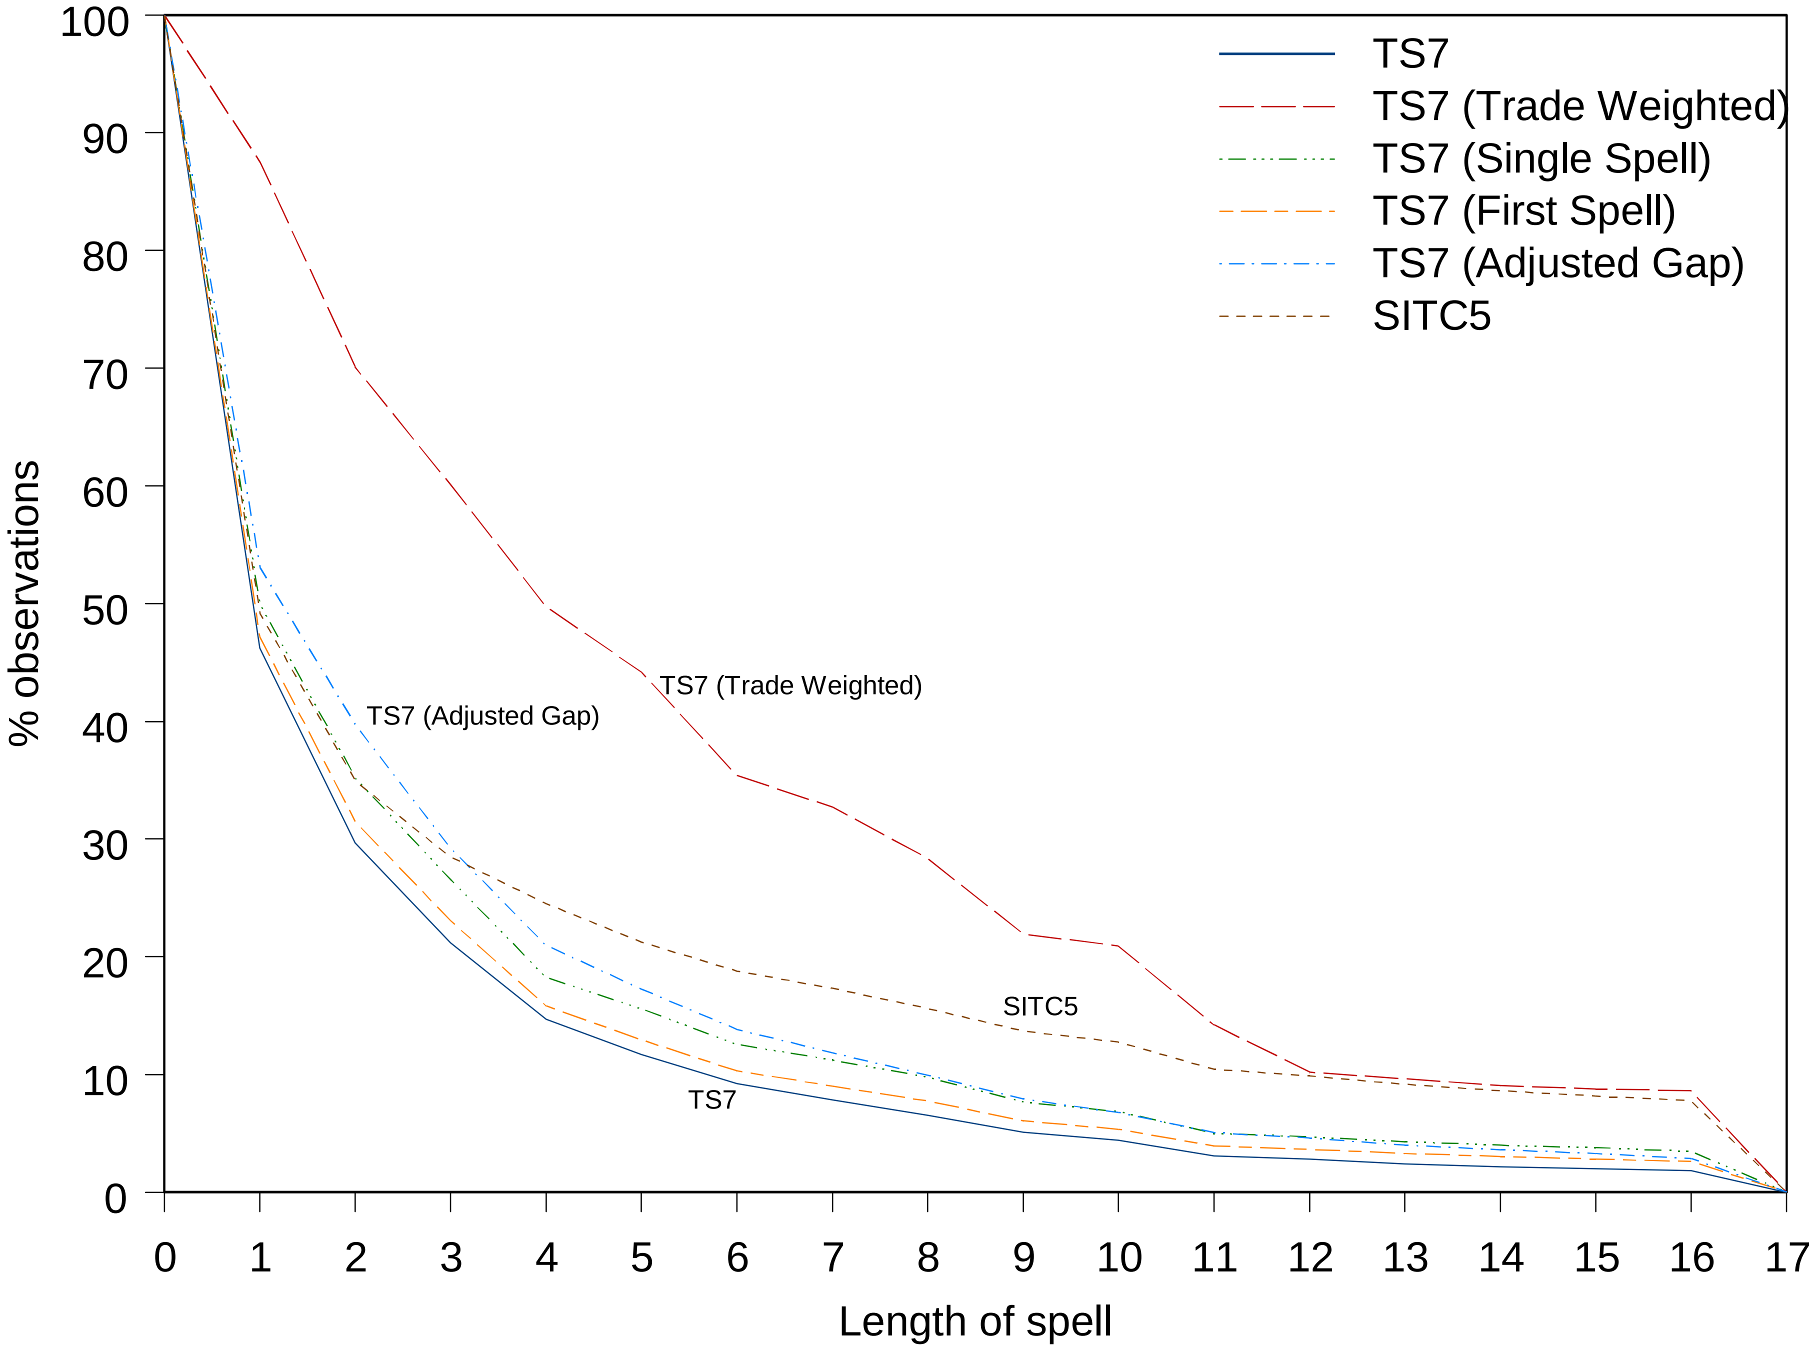

Figure 2 - Survival Function with 95% Confidence Interval

7-digit TSUSA 1972-1989

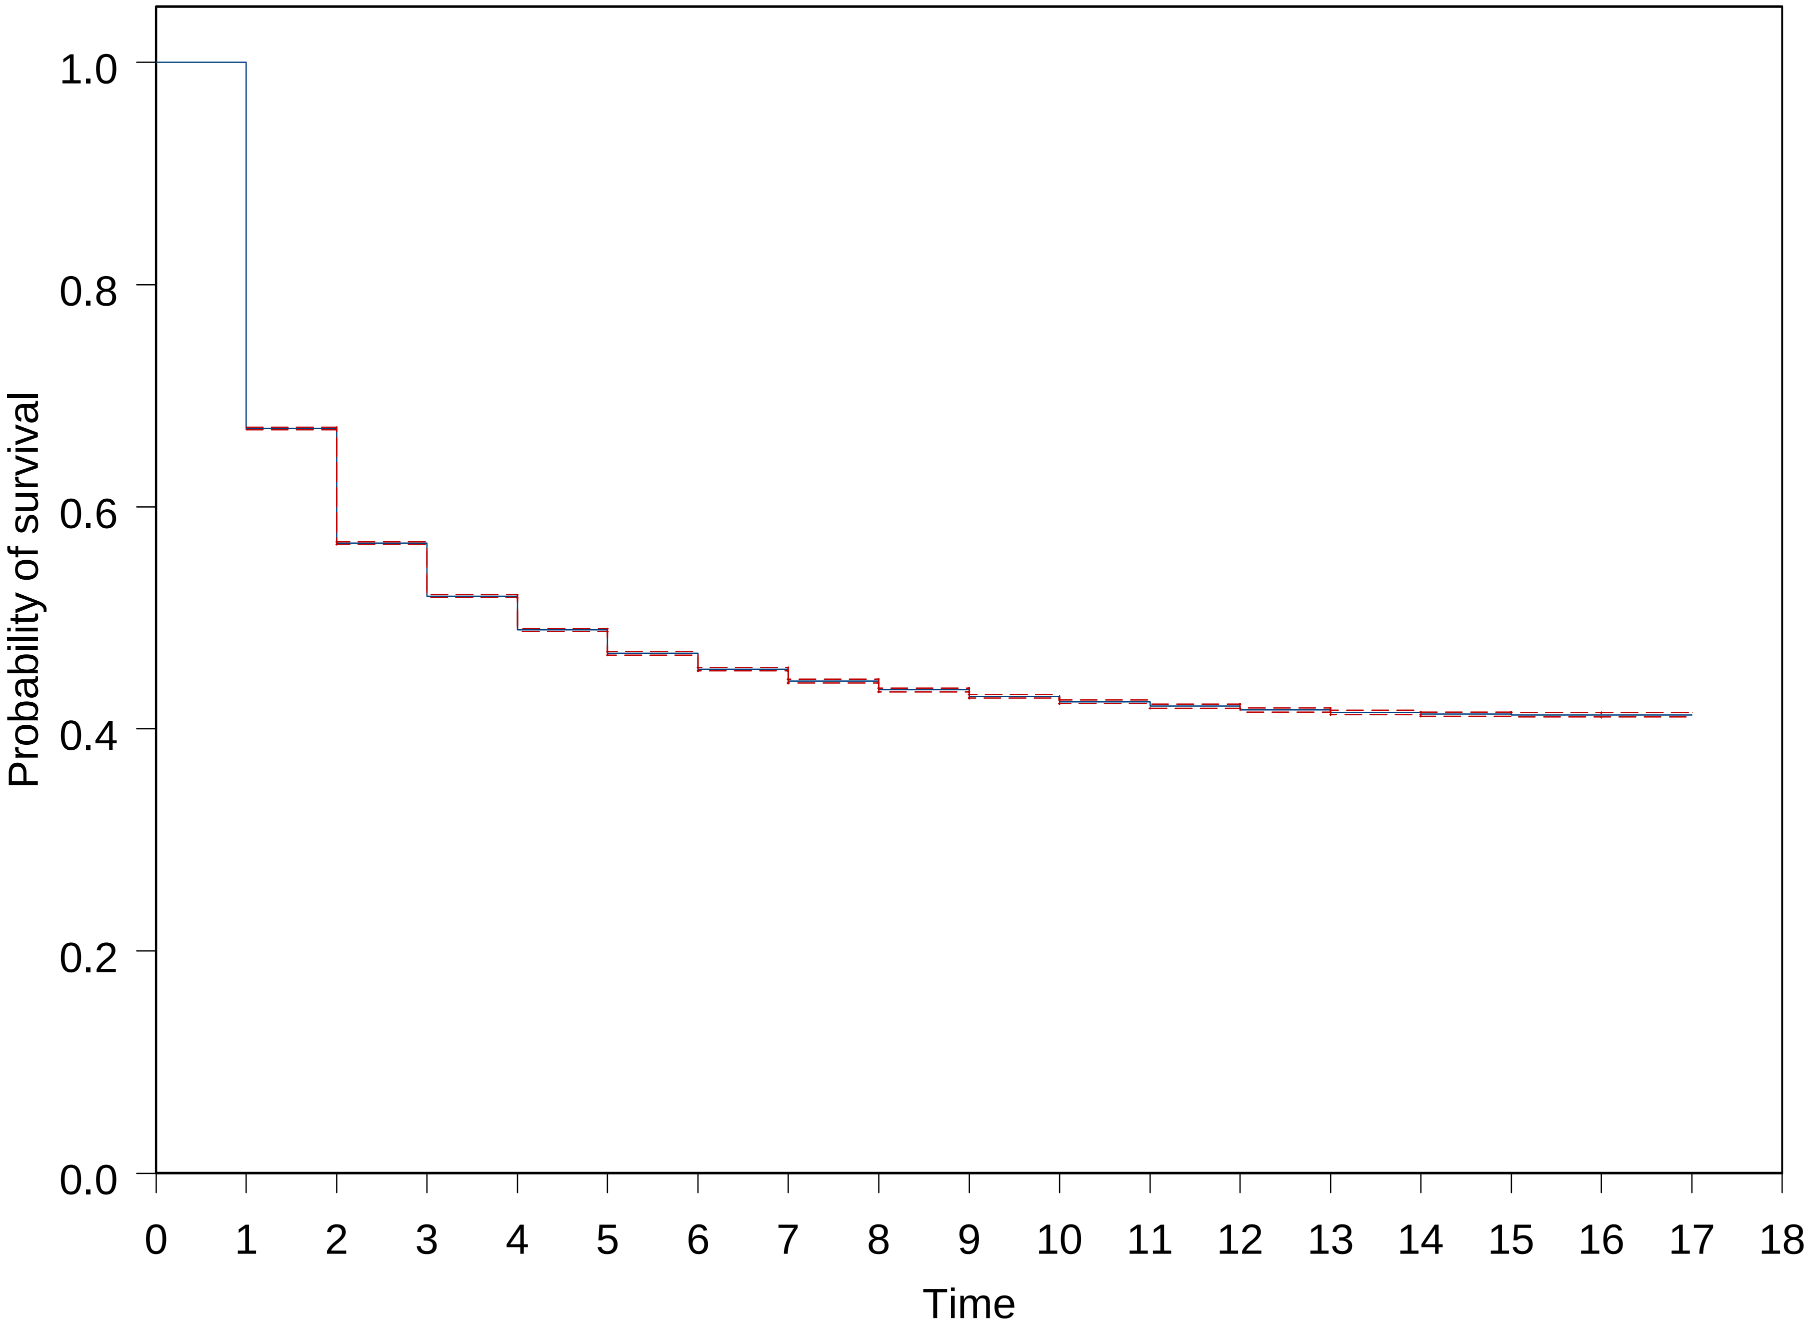

Figure 3 - Overall Hazard Function

7-digit TSUSA 1972-1989

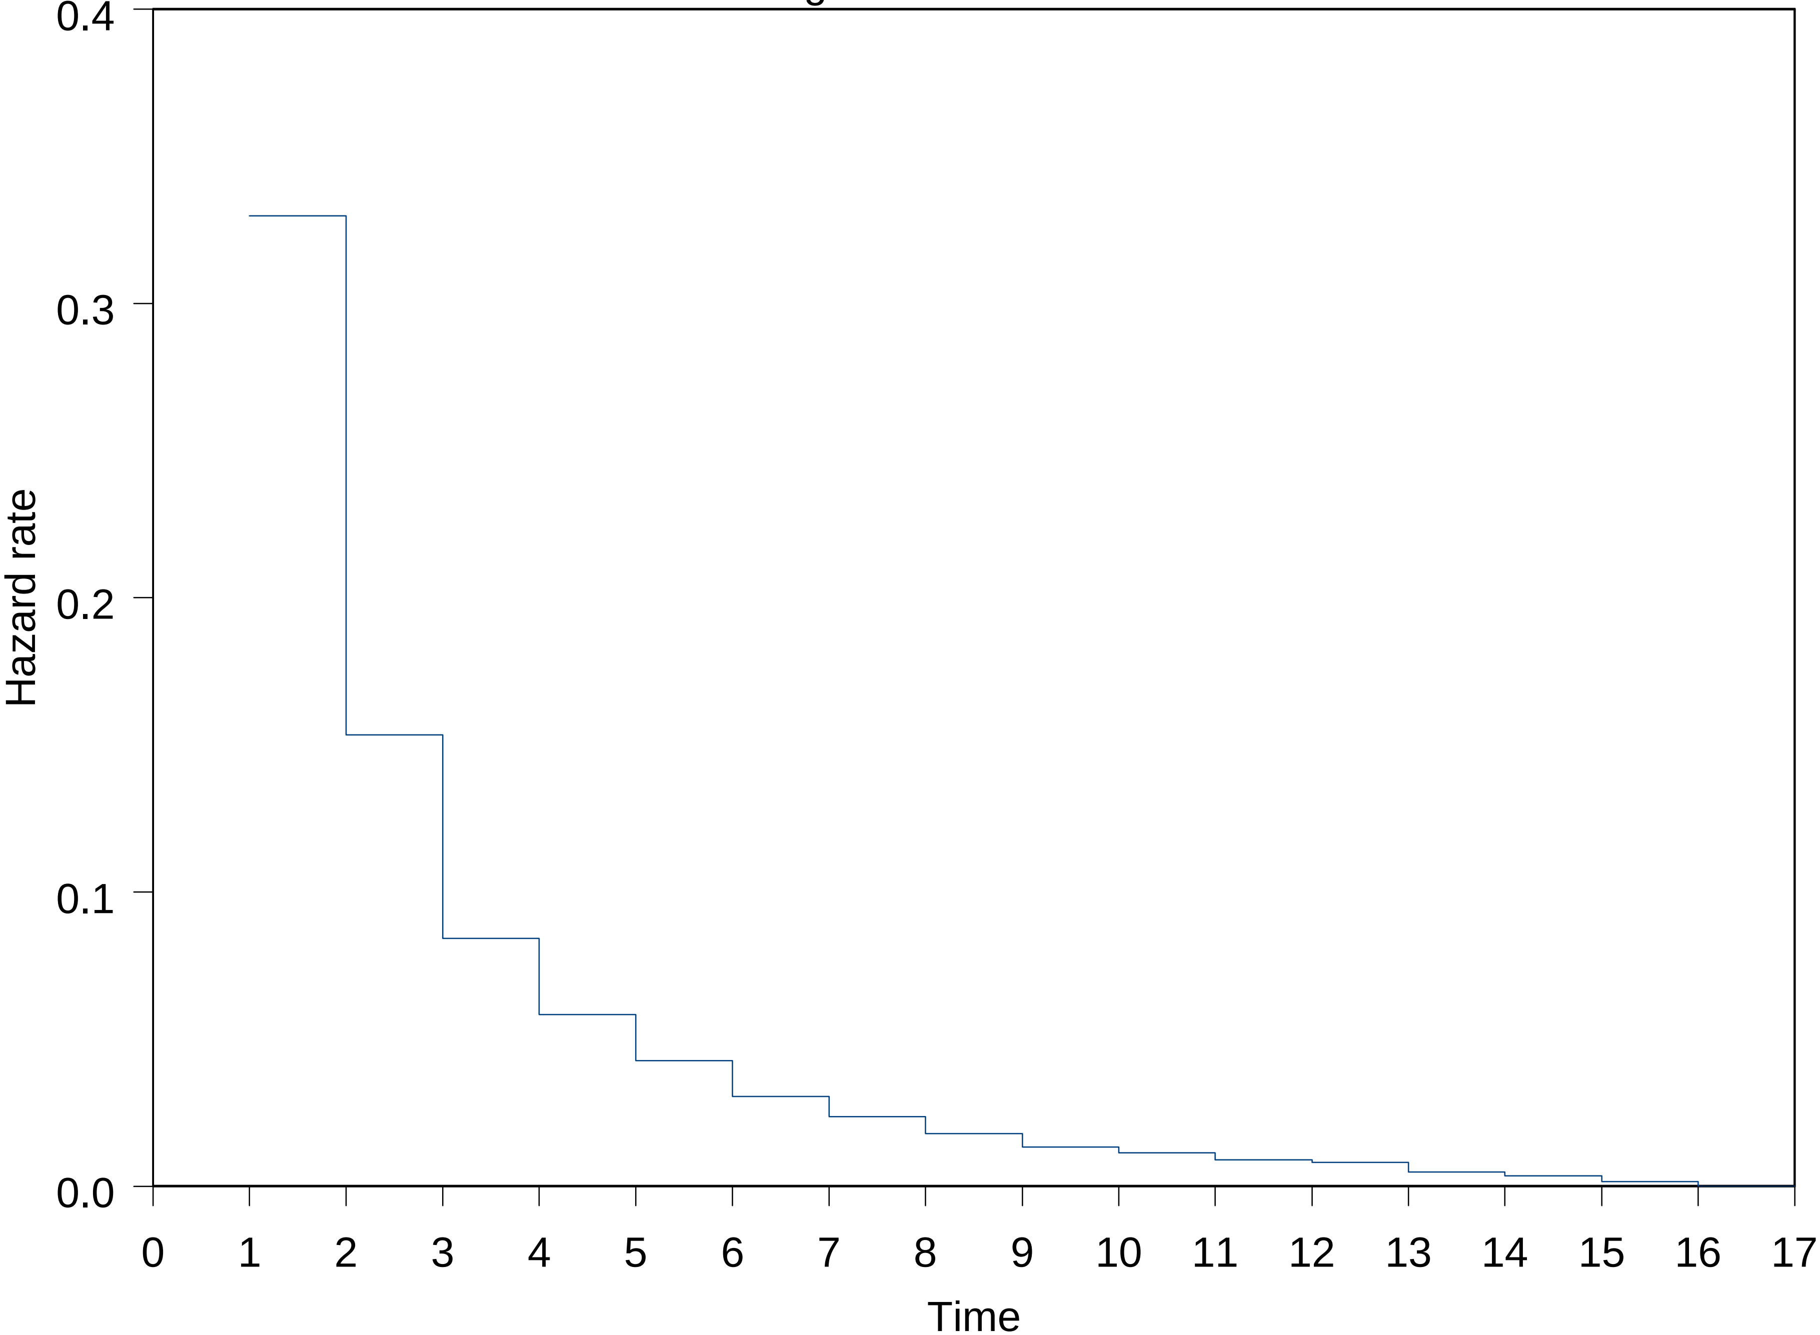

Figure 4 - Survival Functions for TS and SITC Based Data, with Modified Censoring

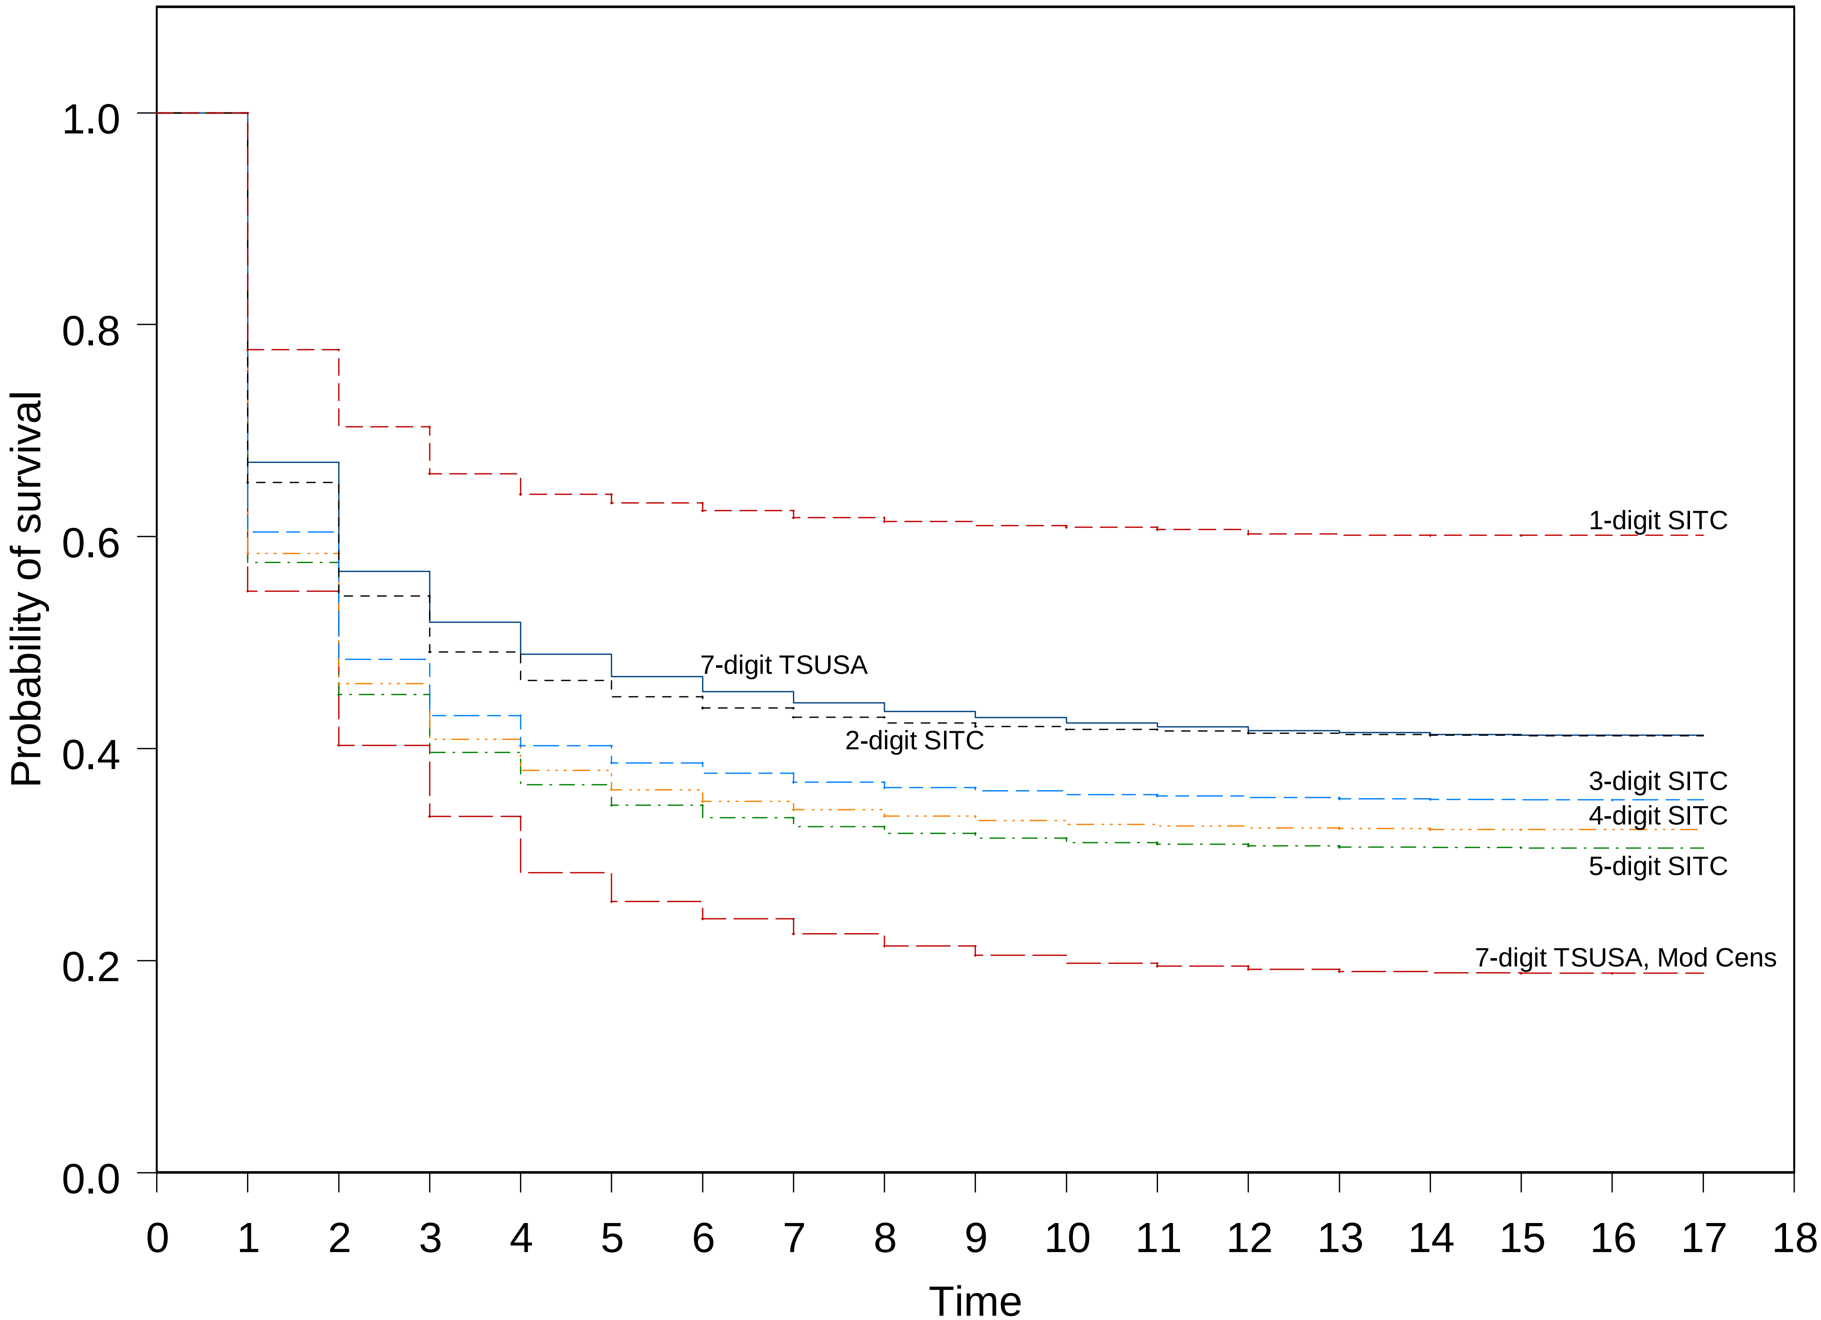

Figure 5 - Survival Functions for Regions

7-digit TSUSA 1972-1989

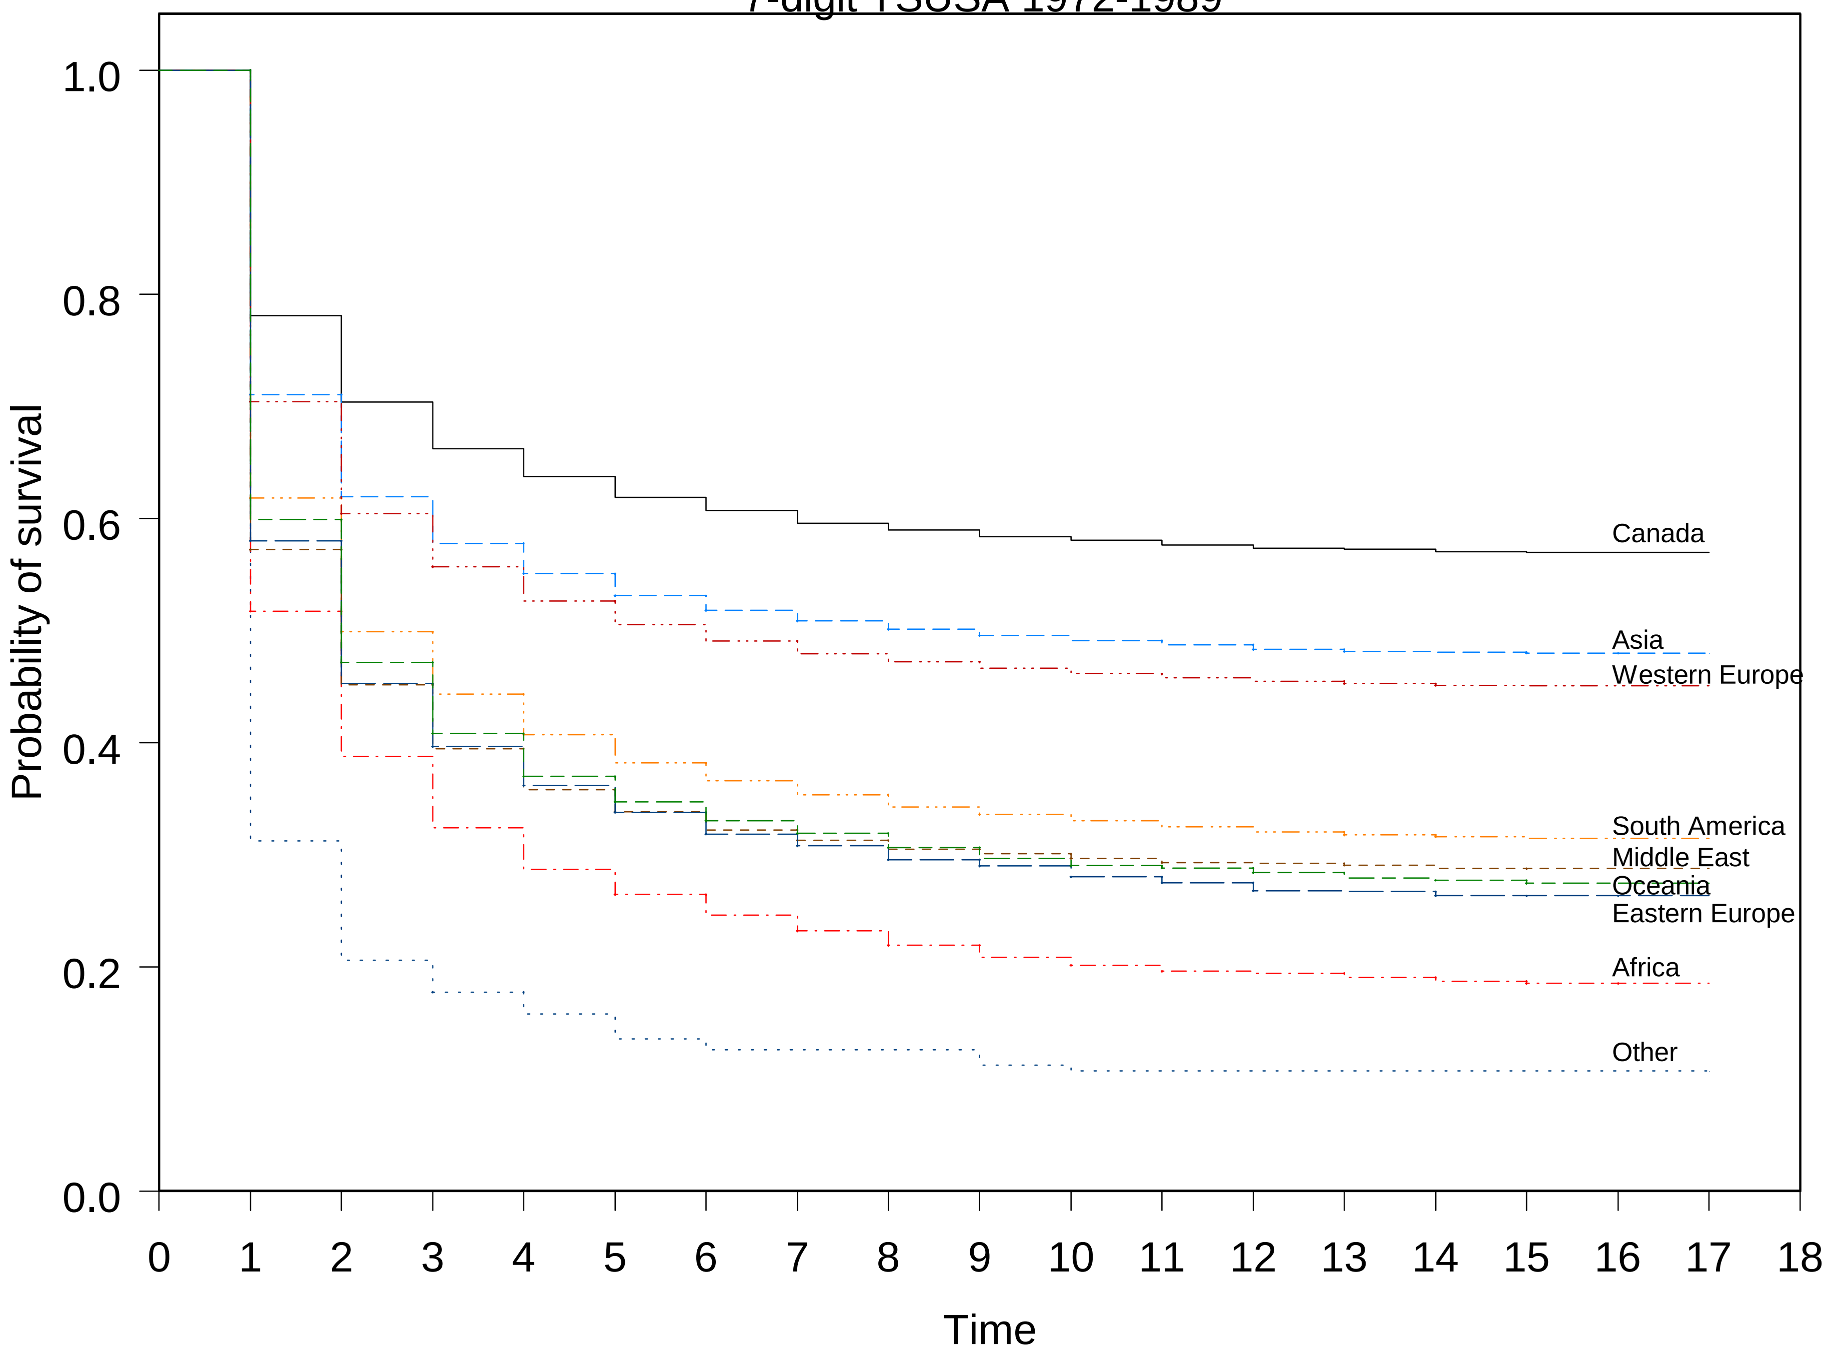

Figure 6 - Survival Functions for OECD and Non-OECD Countries

7-digit TSUSA 1972-1989

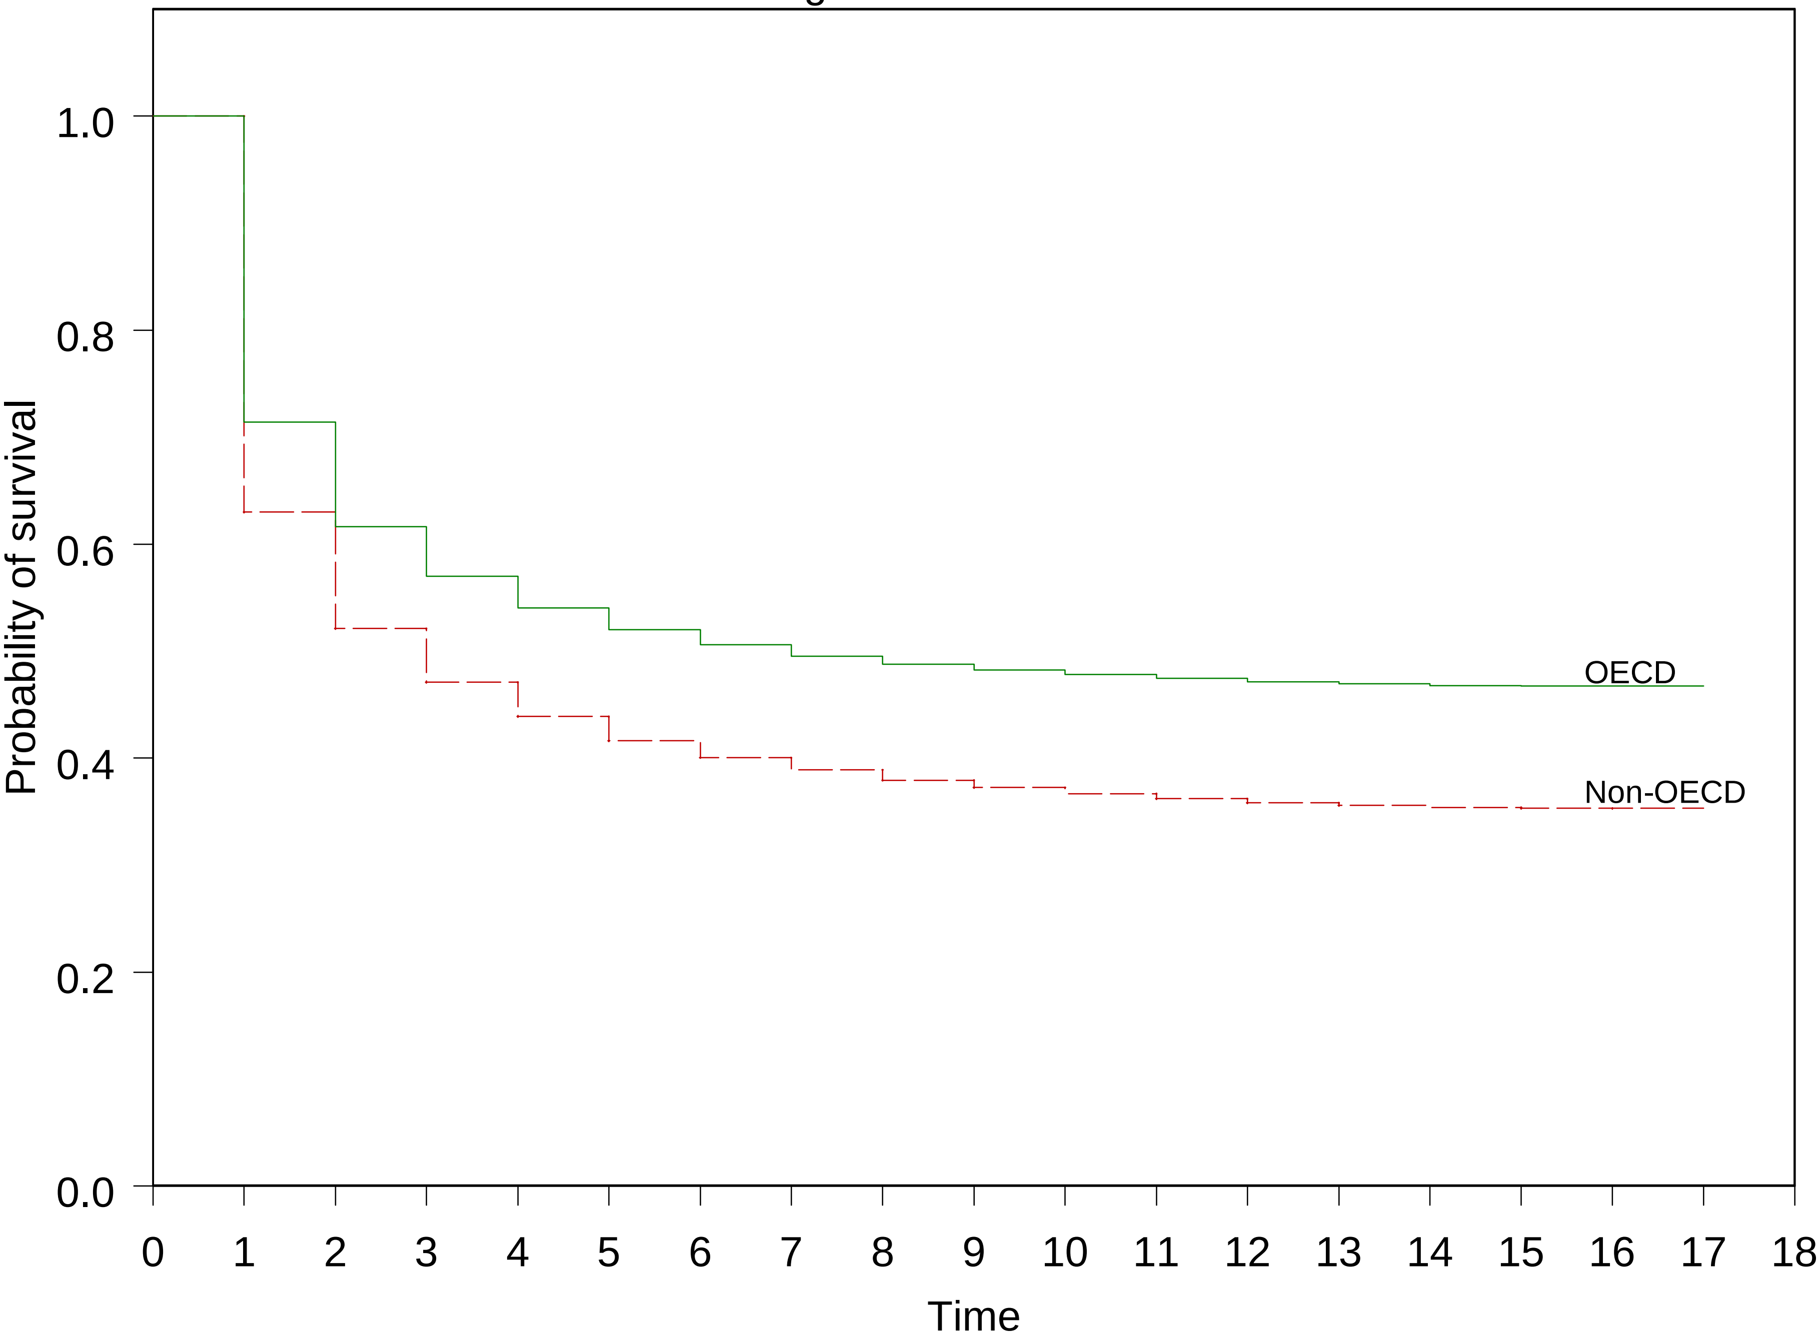

Figure 7 - Survival Functions for 1-digit SITC Industries

7-digit TSUSA 1972-1989

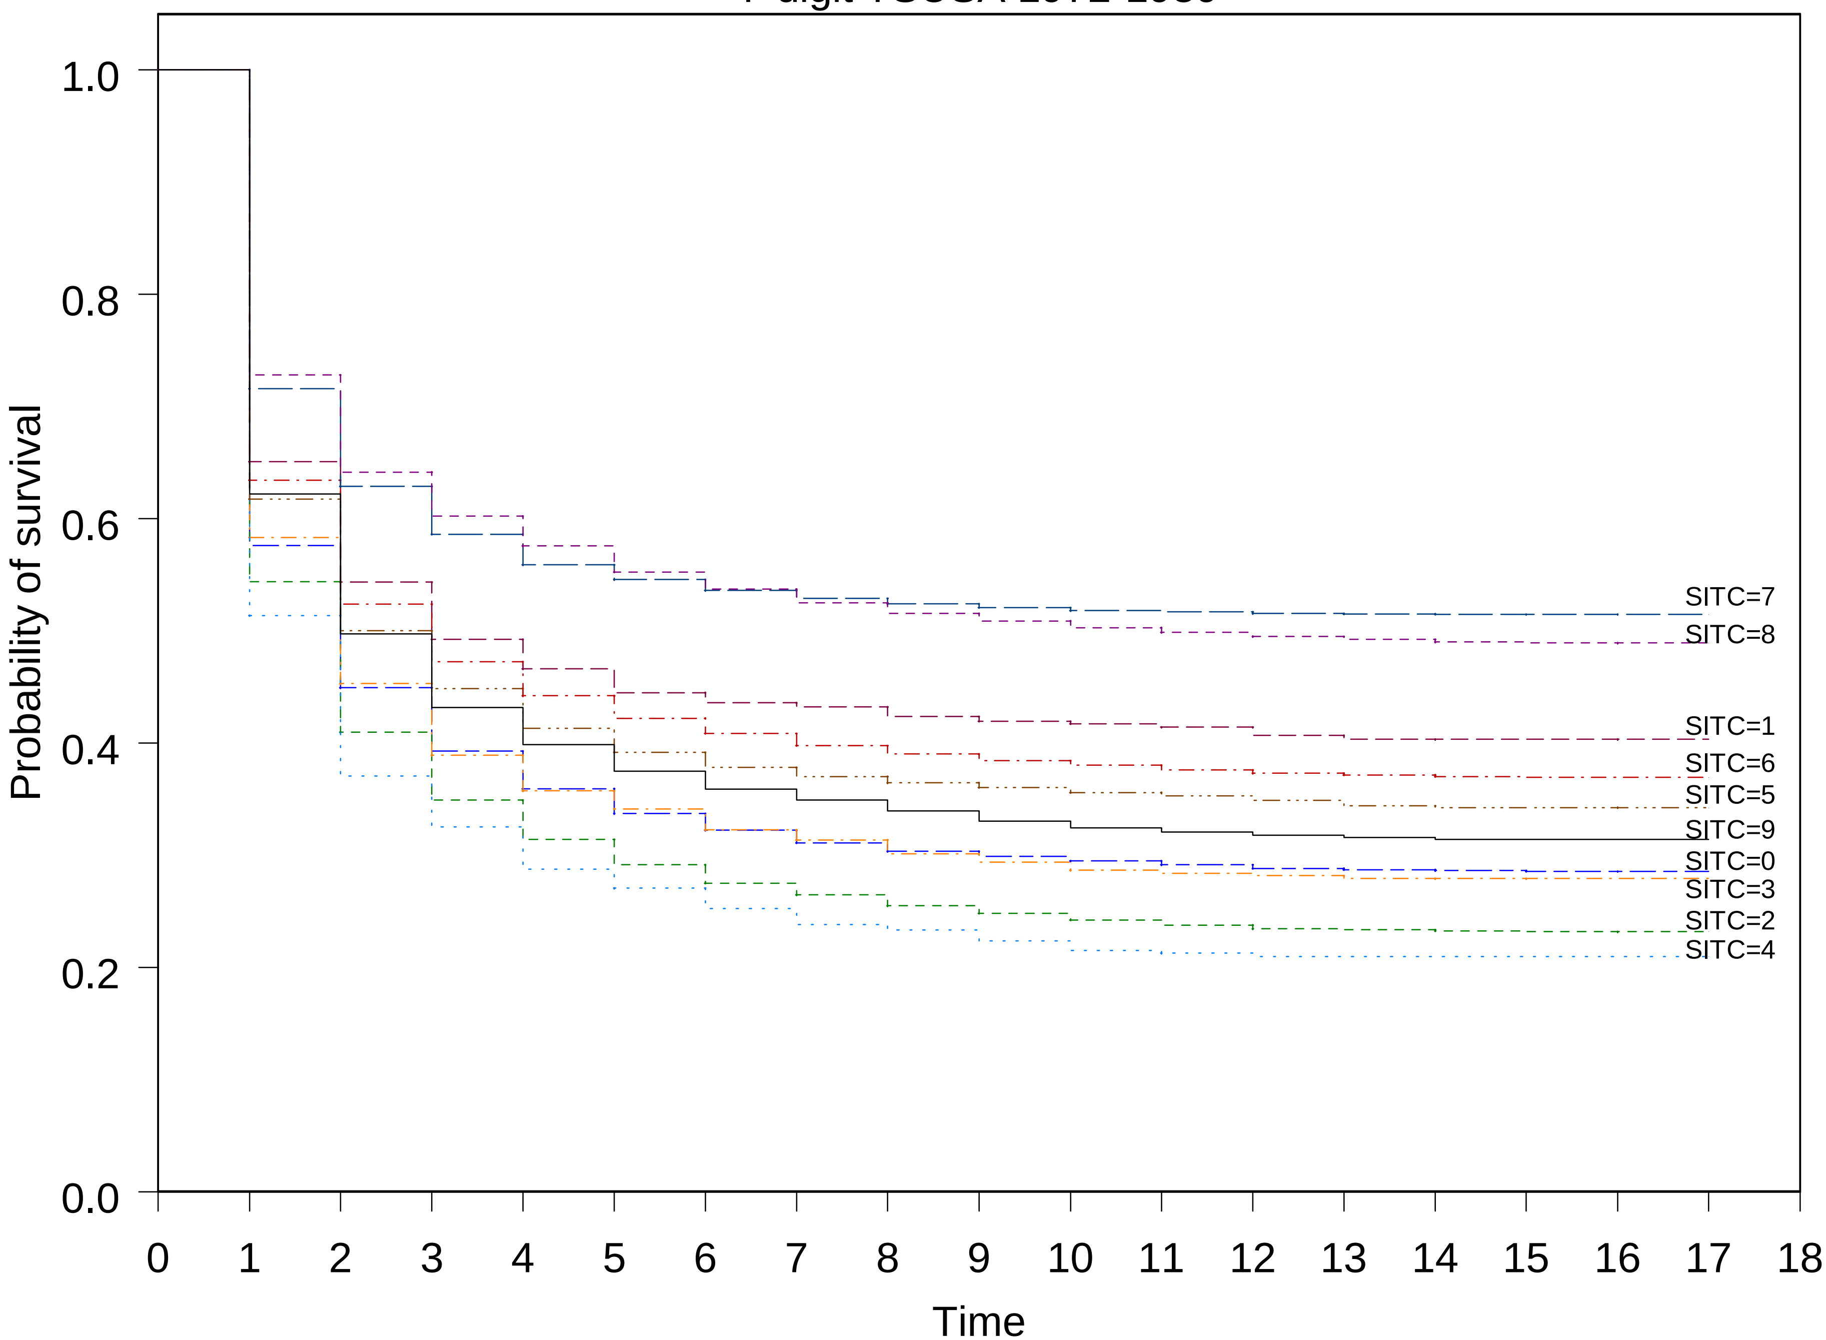

Figure 8 - Survival Functions for Alternative Treatments of 7-digit TSUSA Data

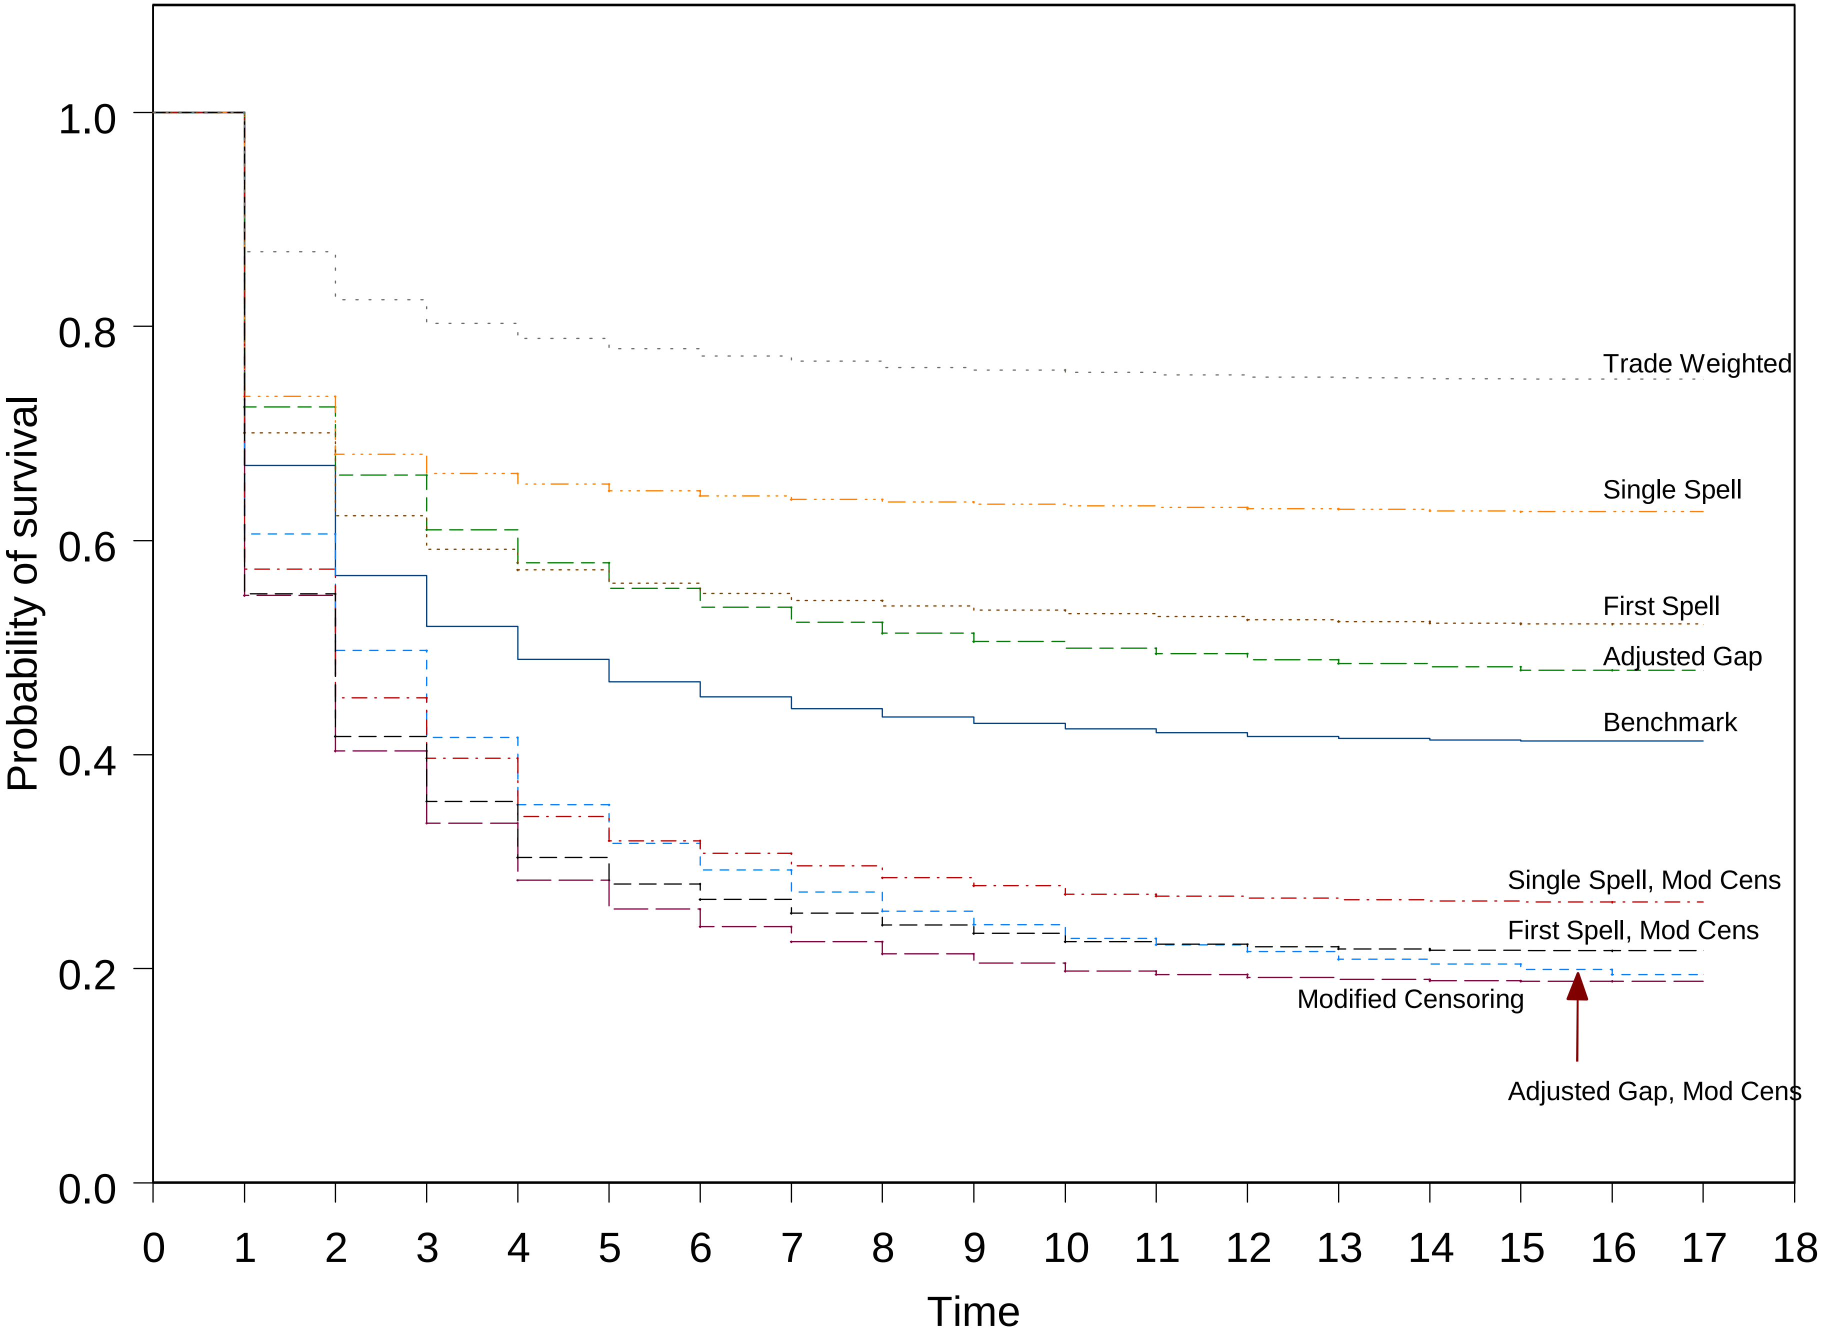

Table 1 - Examples of data

CORN (1312000)

	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	No. Years In	# Spells
Canada	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	**17**	**1**
Portugal	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	**17**	**1**
Mexico			X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	**15**	**1**
South Korea					X	X	X	X	X	X	X	X	X	X	X	X	X	**13**	**1**
Peru													X	X	X	X	X	**5**	**1**
Kenya																X		**1**	**1**
Honduras											X							**1**	**1**
St Kitts Nev							X											**1**	**1**
Hong Kong													X					**1**	**1**
Philippines	X																	**1**	**1**
Belgium-Lux.	X																	**1**	**1**
Greece												X						**1**	**1**
Spain	X																	**1**	**1**
East Germany								X										**1**	**1**
Venezuela	X	X	X	X	X	X	X	X	X		X	X	X	X	X	X	X	**16**	**2**
Ecuador		X	X					X	X	X	X	X	X	X	X	X	X	**12**	**2**
Japan											X	X	X		X			**4**	**2**
Argentina			X											X	X			**3**	**2**
Guatemala									X				X					**2**	**2**
Dominican Rp			X									X						**2**	**2**
India							X										X	**2**	**2**
China								X				X						**2**	**2**
Colombia		X	X	X		X				X	X	X	X	X	X	X	X	**12**	**3**
Italy	X				X	X			X	X	X	X	X	X	X	X	X	**12**	**3**
United Kingdom							X	X			X	X					X	**5**	**3**
Thailand								X		X						X	X	**4**	**3**
Taiwan			X								X						X	**3**	**3**
Germany				X											X		X	**3**	**3**
El Salvador													X		X		X	**3**	**3**
France			X								X	X		X			X	**5**	**4**
Netherlands	X	X	X	X	X			X				X		X		X		**9**	**5**
Brazil	X			X						X	X		X	X			X	**7**	**5**

Table 2 - Summary Statistics

	Observed Spell Length (years)		Total Number
	Mean	Median	of Observations
Benchmark data			
TS7	2.7	1	693,963
Industry-level aggregation			
SITC5	3.9	1	157,441
SITC4	4.2	2	98,035
SITC3	4.7	2	43,480
SITC2	5.5	2	15,257
SITC1	8.4	5	2,445
TS7 Alternatives			
TS7 - Trade Weighted	6.1	4	693,963
TS7 - First Spell	2.9	1	495,763
TS7 - One Spell Only	3.2	1	365,491
TS7 - Gap Adjusted	3.3	2	593,450

Table 3 - Distribution of Kaplan-Meier estimated survival time

	Product Level Data					
	TS7-benchmark			TS7-mod. censoring		
Percentiles of distribution	25%	50%	75%	25%	50%	75%
Total	1	4	.	1	2	6
Regions						
Canada	2	.	.	1	3	.
Asia	1	9	.	1	2	7
Western Europe	1	6	.	1	2	9
South America	1	2	.	1	1	4
Eastern Europe	1	2	.	1	1	3
Oceania	1	2	.	1	1	4
Africa	1	2	6	1	1	2
Middle East	1	2	.	1	1	3
Other	1	1	2	1	1	1
SITC (rev 2) 1-digit industries						
0-Food	1	2	.	1	2	6
1-Beverages/Tobacco	1	3	.	1	2	6
2-Crude Materials	1	2	9	1	2	5
3-Mineral Fuels	1	2	.	1	2	6
4-Animal and Vegetable Oils	1	2	7	1	1	4
5-Chemicals	1	3	.	1	2	7
6-Manufactured Materials	1	3	.	1	2	6
7-Machinery	1	.	.	1	2	8
8-Miscellaneous Manufactures	1	11	.	1	2	5
9-Other	1	2	.	1	2	.

Table 3 - Distribution of Kaplan-Meier estimated survival time

	Industry-level Aggregation														
	SITC5			SITC4			SITC3			SITC2			SITC1		
Percentiles of distribution	25%	50%	75%	25%	50%	75%	25%	50%	75%	25%	50%	75%	25%	50%	75%
Total	1	2	.	1	2	.	1	2	.	1	3	.	2	.	.
Regions															
Canada	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.
Asia	1	2	.	1	3	.	1	3	.	1	7	.	2	.	.
Western Europe	1	4	.	1	8	.	1	.	.	2	.	.	.	.	.
South America	1	2	8	1	2	13	1	2	.	1	4	.	2	.	.
Eastern Europe	1	2	5	1	2	8	1	3	.	1	.	.	13	.	.
Oceania	1	2	.	1	2	.	1	2	.	1	3	.	2	.	.
Africa	1	1	3	1	1	4	1	2	6	1	2	.	1	7	.
Middle East	1	1	4	1	1	4	1	1	4	1	2	.	1	.	.
Other	1	1	1	1	1	2	1	1	2	1	1	3	1	8	.
SITC (rev 2) 1-digit industries															
0-Food	1	2	.	1	2	.	1	2	.	1	4	.	2	.	.
1-Beverages/Tobacco	1	2	.	1	2	.	1	2	.	1	4	.	1	8	.
2-Crude Materials	1	2	10	1	2	.	1	2	.	1	3	.	3	.	.
3-Mineral Fuels	1	2	8	1	2	10	1	2	.	1	2	.	1	3	.
4-Animal and Vegetable Oils	1	2	6	1	2	7	1	2	.	1	2	.	1	3	.
5-Chemicals	1	2	.	1	2	.	1	2	.	1	3	.	1	.	.
6-Manufactured Materials	1	2	.	1	2	.	1	2	.	1	4	.	5	.	.
7-Machinery	1	2	.	1	2	.	1	2	.	1	3	.	2	.	.
8-Miscellaneous Manufactures	1	2	.	1	3	.	1	3	.	1	5	.	.	.	.
9-Other	1	3	.	1	6	.	1	6	.	1	6	.	.	.	.

Table 4 - Country rankings

Country	Kaplan-Meier Survival			Feenstra-Rose	
	TS7	TS7 - Mod. Censoring	SITC 5 digit	Goods Based	Country Based
Japan	1	3	4	4	4
Germany	2	1	2	3	3
Canada	3	2	1	1	1
United Kingdom	4	4	3	2	2
Italy	5	6	6	8	6
Taiwan	6	7	9	13	13
France	7	5	5	5	5
Hong Kong	8	9	16	17	14
South Korea	9	8	13	18	20
Macao	10	62	43	71	70
China	11	13	15	11	25
Switzerland	12	10	10	10	10
Sri Lanka	13	63	52	73	73
Mexico	14	11	8	6	7
Indonesia	15	43	34	46	52
Spain	16	15	14	14	12
Philippines	17	27	29	33	31
Thailand	18	23	25	32	34
Sweden	19	12	12	12	11
Singapore	20	29	28	28	30
Austria	21	19	17	20	17
Haiti	22	55	42	61	39
Portugal	23	25	26	34	24
Netherlands	24	14	7	7	8
Malaysia	25	39	36	39	41
India	26	22	23	24	19
Dominican Rp	27	46	40	38	37
Brazil	28	18	19	15	16
Israel	29	20	20	22	21
Denmark	30	16	18	19	15
Belgium-Lux.	31	17	11	9	9
Mauritius	32	56	47	70	110
Costa Rica	33	68	50	49	54
Turkey	34	32	31	47	50
Falkland Isl	35	69	117	149	150
Poland	36	38	33	37	32
Bangladesh	37	49	91	107	90
Fm Yugoslavia	38	34	30	40	33
Ireland	39	30	22	25	23
Finland	40	31	24	26	26
Romania	41	79	58	50	45
Australia	42	21	21	16	18
Pakistan	43	41	44	57	48
Norway	44	24	27	23	22
Morocco	45	66	45	62	62
Greece	46	50	38	43	35
Nepal	47	105	82	110	89
Guatemala	48	54	51	52	47
Colombia	49	37	35	36	29
Jamaica	50	47	56	51	44
Uruguay	51	73	64	69	65
Hungary	52	52	41	45	49

Table 4 - Country rankings

Country	Kaplan-Meier Survival TS7	Kaplan-Meier Survival TS7 - Mod. Censoring	Kaplan-Meier Survival SITC 5 digit	Feenstra-Rose Goods Based	Feenstra-Rose Country Based
Ecuador	53	57	57	54	56
Cyprus	54	114	65	87	83
Lebanon	55	28	32	78	59
Barbados	56	109	86	84	72
Honduras	57	53	62	60	64
Chile	58	26	48	35	57
Argentina	59	33	39	27	28
Peru	60	67	55	44	46
El Salvador	61	98	63	67	55
New Zealand	62	44	37	31	36
Egypt	63	36	59	56	75
Bolivia	64	83	66	88	77
Czechoslovakia	65	45	54	41	43
Venezuela	66	35	49	30	38
Guyana	67	93	88	82	78
South Africa	68	87	46	21	27
Algeria	69	59	105	101	104
Nicaragua	70	94	71	86	58
Trinidad-Tobago	71	64	72	59	63
Belize	72	81	93	95	92
Tunisia	73	122	84	102	99
Panama	74	90	73	53	53
Reunion	75	75	126	123	148
Zimbabwe	76	51	53	91	123
St Kitts Nev	77	92	70	48	69
Malta	78	121	78	112	71
Bahrain	79	88	121	94	131
Paraguay	80	76	79	114	82
Iceland	81	89	95	68	74
Untd Arab Em	82	70	122	74	98
Afghanistan	83	141	83	111	76
Cote D'Ivoire	84	71	61	72	102
Oman	85	142	135	118	128
Fiji	86	58	74	105	111
Cambodia	87	125	129	148	144
Kenya	88	86	60	64	67
Malawi	89	60	80	136	135
Ghana	90	61	81	97	80
Congo	91	48	76	120	120
Bahamas	92	95	98	65	61
Untd Rp Tanzania	93	40	69	124	88
Angola	94	78	111	133	109
Albania	95	77	109	108	124
St Pierre Miqu	96	74	92	135	142
Zaire	97	42	75	106	103
Fm USSR	98	84	68	29	42
Madagascar	99	65	67	117	106
Seychelles	100	85	77	99	133
Ethiopia	101	72	102	146	97
Bulgaria	102	116	103	66	86
Neth Antilles	103	107	116	63	66
East Germany	104	108	89	42	40

Table 4 - Country rankings

Country	Kaplan-Meier Survival			Feenstra-Rose	
	TS7	TS7 - Mod. Censoring	SITC 5 digit	Goods Based	Country Based
Guinea	105	100	104	92	140
Iraq	106	82	113	139	96
Sierra Leone	107	91	108	116	117
Djibouti	108	112	133	150	149
Cameroon	109	103	97	100	114
Liberia	110	106	119	109	94
Mongolia	111	99	138	89	145
Nigeria	112	80	85	83	81
Rest America Nes	113	137	132	128	141
Togo	114	111	90	143	143
Kuwait	115	126	136	119	115
Myanmar (Burma)	116	102	114	115	119
Sudan	117	120	115	126	138
Qatar	118	115	137	121	136
Zambia	119	101	123	131	121
Senegal	120	117	99	113	129
Mozambique	121	97	94	85	95
Syrn Arab Rp	122	124	110	103	68
Burkina Faso	123	134	101	144	134
Benin	124	96	87	147	122
Saudi Arabia	125	118	106	55	84
Jordan	126	140	134	75	113
Iran	127	133	118	76	51
Bermuda	128	127	130	96	93
Kiribati	129	128	120	80	127
Surinam	130	119	112	90	108
Papua N.Guinea	131	110	107	129	105
Uganda	132	104	96	138	116
New Caledonia	133	135	128	93	101
Gabon	134	123	125	77	125
Central Afr. Rep.	135	113	100	142	147
Greenland	136	139	144	127	91
Mali	137	130	124	98	118
French Guiana	138	138	141	122	132
Rwanda	139	132	127	145	139
Liby Arab Jm	140	129	145	125	100
Laos P.Dem.R	141	149	149	130	107
Mauritania	142	146	150	141	130
Guadeloupe	143	152	140	104	87
Vietnam	144	143	139	137	85
Gibraltar	145	155	147	79	60
St.Helena	146	150	148	134	126
Fm Dem Yemen	147	148	142	81	79
Fm Yemen	148	147	143	132	137
Gambia	149	154	146	140	146
Niger	150	151	131	58	112

Table 5 - Pairwise rank correlations of country rankings

	TS7 Benchmark	TS7 Modified Censoring	SITC 5 digit	Feenstra-Rose Goods Based	Feenstra-Rose Country Based
TS7 Benchmark	1				
TS7 Modified Censoring	0.8792	1			
	(0.0000)				
SITC 5 digit	0.9032	0.913	1		
	(0.0000)	(0.0000)			
Goods Based	0.7568	0.7391	0.8272	1	
	(0.0000)	(0.0000)	(0.0000)		
Country Based	0.766	0.7341	0.8404	0.9089	1
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	

*p-values are reported in parentheses below the correlations