Distance, Trade, and Income – The 1967 to 1975 Closing of the Suez Canal as a Natural Experiment*

James Feyrer†
Dartmouth College and NBER

June 24, 2009

Abstract

The closing of the Suez canal in 1967 and its reopening in 1975 had a significant effect on trade routes by sea. For most pairs of countries in the world, the closing and reopening of the canal act as exogenous shocks to sea distance. These exogenous shocks can be used to identify the effect of distance by sea on trade volumes. The time series variation in distance allows for the inclusion of pair effects. The identification of distance effects is therefore more clearly about the impact of transportation costs than typical gravity model estimates. Distance is found to have a significant impact on trade. These trade volume movements can be further exploited to identify the effect of trade on income. Trade is found to have a significant impact on income. Because identification is through changes in sea distance, the effect is coming entirely through trade in goods and not through alternative channels such as technology transfer, tourism, etc. The results should therefore be useful in thinking about the effect of changing trade costs on income, including the reduction of trade barriers.

*Many thanks to Alan Taylor and Reuven Glick for sharing their bilateral trade data. Thanks to Jay Shambaugh, Doug Staiger, Liz Cascio, Doug Irwin and Nina Pavcnik for helpful comments. All errors are my own.

†james.feyrer@dartmouth.edu, Dartmouth College, Department of Economics, 6106 Rockefeller, Hanover, NH 03755-3514. fax:(603)646-2122.

Introduction

The distance between countries has a substantial impact on the volume of trade between them.[1] Why should distance matter? The most obvious answer is that trade is a function of transportation costs (which rise with distance). However, typical gravity model estimations are likely to capture more than just transport costs. Any country pair characteristics that are correlated with distance will bias the coefficient. While some aspects of bilateral relationships like common language, colonial status, etc. can be controlled for, one can never completely eliminate missing variable bias in a cross section. For this reason, the distance coefficient in typical gravity regressions reflects many other aspects of distance beyond pure differences in transportation costs.

This paper will estimate a gravity model of trade using novel variation that more directly targets transportation costs – an exogenous time series shock to distance. On June 5, 1967, at the beginning of the Six Day War, Egypt closed the Suez canal. The canal remained closed for exactly eight years, reopening on June 5, 1975. The Suez Canal divides Africa from Asia and connects the Arabian Sea to the Mediterranean (through the Red Sea). The Suez Canal provides the shortest sea route between Asia and Europe. About 7.5 percent of world trade currently flows through the Canal. The closure of the canal was a substantial shock to world trade. While the impact for countries on the Arabian sea was largest, it had an effect on a substantial number of trade pairs. For most countries in the world, the closure of the Canal can be seen as an exogenous event. The reopening of the canal provides a similar shock in the opposite direction.

This paper will exploit these shocks to identify the effect of distance on trade and further to examine the effect of trade on output. Because there is time series variation, time and bilateral pair controls can be used to insure that all identification comes from the change in distance due to the closure of the Suez Canal. By using variation caused by changes in sea distance, the estimates in this paper are much more closely focused on the pure impact of transportation costs compared to standard gravity estimates. The results suggest that the conventional estimates may overestimate the effect of distance on trade. The elasticities found in this paper are about half those typically found in the literature.

[1] A large literature has been produced testing gravity models of trade. Disdier and Head (2008) collect estimates of the impact of distance on trade from 108 papers.

The second part of the paper will use the gravity model results to generate predictions for the change in aggregate trade at the country level caused by the closing (and reopening) of Suez. The time series variation in these predictions is driven entirely by geography and these predictions make a useful instrument for trade in a regression of income on trade.

The effect of trade on income is of obvious interest and has been explored in numerous papers, but identification has been difficult due to reverse causality.[2] Rodriguez and Rodrik (2000) conclude that none of these papers establish a robust and well identified relationship between trade restrictions and growth. The key difficulty faced in this literature is the lack of exogenous variation in trade or trade policies. Though some papers attempt to use instrumental variables, the instruments tend to violate exclusion restrictions.

One of the more plausible instruments for trade is from Frankel and Romer (1999), who use the distance between countries to predict bilateral trade volumes using the gravity model of trade as a framework. These predicted trade volumes can be aggregated to generate an instrument for aggregate trade for each country that is based on proximity to other countries in the world. The concern with this approach is that proximity may be acting through channels other than trade. Rodriguez and Rodrik (2000) and others show that Frankel and Romer (1999)'s results are not robust to the inclusion of geographic controls in the second stage.[3] For example, countries that are closer to the equator tend to be more remote from other countries. Since proximity to the equator is associated with low incomes it may be that the Frankel and Romer (1999) instruments are picking up this effect rather than trade.

The approach of this paper is similar in the use of geography as an instrument for trade, but with the addition of time series variation provided by the Suez Canal shocks. The time series variation in the trade predictions allows for the inclusion of country dummies in the second stage, controlling for all time invariant income differences. This addresses the concerns of Rodriguez and Rodrik (2000) and others by controlling for static geographic differences and slow moving institutional variables.

This is similar to Feyrer (2009), where the identifying variation comes from the technological improvement in air transport. The income results in this paper differ

[2] Sachs and Warner (1995), Frankel and Romer (1999), Dollar (1992), and Edwards (1998) are some of the more prominent papers finding a positive relationship between trade (or being open to trade) and income.

[3] See also Rodrik, Subramanian and Trebbi (2004) and Irwin and Terviö (2002).

in two important ways. First, Feyrer (2009) examines changes in trade that are slower moving and occur over decades. This paper exploits a short run shock to trade and is therefore more suited to thinking about events and policies that impact trade over the course of years, not decades. The short run nature of the shocks also allows for examining the time path of adjustment to the shocks to distance.

The second important difference is that the variation in distance by sea generated by the closing of Suez is almost certainly identifying the effect of trade in goods. The approach of Feyrer (2009) gets identification from comparing air and sea distances and therefore may be picking up a number of bilateral relationships fostered by easy air travel such as foreign direct investment, trade in services and any other benefits that come from easier movement of people around the globe. Because the variation in this paper relies on sea distance, the effects must be coming through bilateral relationships that change when the distance by sea changes. Trade in goods is the main relationship that fits this description. This paper therefore can more clearly identify the relationship between trade in goods and output.

Changes in sea distance are found to have a significant impact on trade with an elasticity of about 0.4. The adjustment to the distance shock is relatively rapid, with the majority of adjustment occurring over 2 years. The trade movements generated through these distance shocks significantly change income with an elasticity of roughly one half. Because of the unique identification, these results are more directly related to trade in goods than other gravity estimates. This makes the results more applicable to other settings such as estimating the effect of trade policies designed to decrease trade costs.

1 The Six Day War and the Closure of the Suez Canal

The Six Day War was fought between Israel and its neighbors, Egypt, Syria, and Jordan between June 5 and June 10 in 1967. In March of 1967 Egypt expelled the United Nations Emergency Force (UNEF), which had been stationed on the Egypt-Israel border since 1956 helping to enforce the armistice agreement between Israel and Egypt following the Suez Crisis of 1956. The war began on June 5, as Israel launched surprise air strikes which destroyed the majority of the Egyptian Air Forces on the ground.

Though tensions had been high in the region since the Suez Crisis of 1956, the actual outbreak of war was a surprise and the closing of the canal was not anticipated in advance. When the canal closed, fifteen cargo ships known as “The Yellow Fleet” were trapped. They remained in the canal during the entire 8 years of the closure. Since it takes less than a day to transit the canal this suggest there was very little anticipation of the closing beforehand.

At the end of the war, Israel had greatly enlarged the territories under its control. The additions included the Sinai Peninsula and the Gaza strip from Egypt, the West Bank and East Jerusalem from Jordan, and the Golan Heights from Syria. On the Egyptian border, the Suez canal was the cease fire line at the end of the war. The canal had been closed by Egypt at the outbreak of hostilities and remained closed for the next eight years.

In October of 1973, the Yom Kippur War was fought between Israel, Syria, and Egypt. Importantly, Jordan did not take part. Egyptian forces crossed the Suez and attacked Israeli positions in the Sinai Peninsula. Syria staged a simultaneous offensive in the Golan Heights. After taking losses during the first few days, the Israelis counter attacked, retaking the Golan Heights on the northern front and splitting the Egyptian forces in the Sinai, pushing across the Canal. At the time of the UN brokered cease fire Israeli forces were on the west side of the canal and Egyptian forces were on the east side of the Canal.

The ongoing peace negotiations that followed involved reopening the canal. Agreement to reopen the canal was tentatively reached in early 1974. By March 5, 1974, the last of the Israeli troops had withdrawn from the west side of the canal. After fixing war damage and removing mines and munitions the canal reopened on June 5, 1975 eight years to the day of the closure. Unlike the closing, there was roughly a year of advance notice that the canal was to reopen.

2 The Gravity Model

The gravity model has been widely used for almost half a century. The basic idea that trade decreases with the distance between two countries is intuitive and holds up well empirically. This application of the gravity model is particularly straightforward since the nature of the shock is directly to distance. This allows for identifying the effect of distance in a panel of bilateral trade. The inclusion of bilateral pair dummies means that all identification comes from the change in distance caused by the closing of the Suez canal.[4]

Anderson and van Wincoop (2003) develop a theoretical model to derive the gravity model. The basic gravity relationship derived by them is

$$trade_{ijt} = \frac{y_{it}y_{jt}}{y_{wt}} \left(\frac{\tau_{ijt}}{P_i P_j} \right)^{1-\sigma} \tag{1}$$

where $trade_{ijt}$ is bilateral trade between country i and country j, y_i y_j and y_w are the incomes of country i, country j and the world, τ_{ijt} is a bilateral resistance term, and P_i and P_j are country specific multilateral resistance terms. Taking logs,

$$ln(x_{ij}) = ln(y_i) + ln(y_j) - ln(y_w) + (1 - \sigma)(ln(\tau_{ijt}) + ln(P_i) + ln(P_j)). \tag{2}$$

The bilateral resistance term, τ_{ijt}, in Equation (2) encompasses all pair specific barriers to trade such as distance, common language, a shared border, colonial ties, etc. The effect of distance is assumed to be log-linear. The majority of these determinants of bilateral resistance are time invariant and will be controlled for using bilateral pair dummies. The exception is, of course, the change in distance by sea caused by the closing and opening of the Suez Canal. The P and y terms will also be controlled for using country pair dummies. The estimation equation is therefore

$$ln(trade_{ijt}) = \alpha + \gamma_{ij} + \gamma_t + \beta ln(seadist_{ij}) + \epsilon \tag{3}$$

2.1 Data

Trade data was provided by Glick and Taylor (2008) who in turn are using the IMF Direction of Trade (DOT) data. In the DOT data for each bilateral pair in each year there are potentially of four observations – imports and exports are reported from both sides of the pair. An average of these four values is used, except in the case where none of the four is reported. These values are taken as missing.

[4]The distance measures that are commonly used in estimating gravity models are point to point great circle distances. While sea distance occasionally appears in gravity models, it has tended to be in the context of single country or regional studies. Disdier and Head (2008) conduct a meta study of gravity model results and cite the use of sea distance as one differentiator between papers. However the use of sea distance is rare and seems to be limited to regional work. Coulibalya and Fontagne (2005) consider sea distance in an examination of African trade.

Bilateral sea distances were created by the author using raw geographic data. The globe was first split into a matrix of $1x1$ degree squares. The points representing points on land were identified using gridded geographic data from CIESIN.[5] The time needed to travel from any oceanic point on the grid to each of its neighbors was calculated assuming a speed of 20 knots and adding (or subtracting) the speed of the average ocean current along the path. Average ocean current data is from the National Center for Atmospheric Research.[6] The result of these calculations is a complete grid of the water of the globe with information on travel time between any two adjacent points. The grid can be constructed both including and excluding the Suez canal as a valid path. Given any two points in a network of points, the shortest travel time can be found using standard graph theory algorithms.[7] After identifying a primary port for each country all pairwise minimum travel times were calculated from networks with and without the Suez canal as a valid path. For country pairs where the Suez canal is not the shortest path, these two travel times are identical. For country pairs including the Suez canal in the shortest path, the shortest alternative path is calculated. The distance between countries used in the regression is the number of days to make a round trip.

Identifying the location for the primary port for the vast majority of countries was straightforward and for most countries choosing any point along the coast would not change the results. The major potential exceptions to this are the US and Canada, with significant populations on both coasts and massive differences in distance depending on which coast is chosen. For simplicity (and because the east-west distribution of economic activity in the US and Canada can be seen as an outcome) the trade of the US and Canada with all partners was split with 80 percent attributed to the east coast and 20 to the west coast for all years. This is roughly based on the US east-west population distribution for 1970, the middle of the sample. In effect, the US and Canada are each split in two with regards to the trade regressions, with each country in the world trading with each coast independently based on appropriate sea distances. When generating predicted trade shares for the US and Canada, the trade with both halves are summed. Choosing just the east coast sea distances, changing the relative east-west weights, or even

[5] http://sedac.ciesin.columbia.edu/povmap/ds_global.jsp

[6] Meehl (1980), http://dss.ucar.edu/datasets/ds280.0/

[7] Specifically, Djikstra's algorithm as implemented in the Perl module Boost-Graph-1.4 http://search.cpan.org/ dburdick/Boost-Graph-1.2/Graph.pm.

removing all observations including the US and Canada has no significant effect on the results.

Because countries need to abut the sea in order to be located on the oceanic grid, the sample excludes landlocked countries. Oil exporters were also left out of the sample because they have atypical trade patterns and have an almost mechanical relationship between the value of trade and income.

The handling of Israel and Egypt is potentially tricky for two reasons. First, they were the main participants in the war and they each have ports on both sides of the canal. This can be handled in a few different ways. First, they could be treated the same as other countries. Second, they could be removed from the sample entirely. Third, their distances can be coded as if the canal never closed since they have ports on both sides of the canal. Options one and three are ultimately quite similar since the primary ports for both Egypt and Israel are on the Mediterranean Sea. The distance shock to their trade is therefore relatively small. In any case, none of these options has any impact on the results.

Jordan is much more problematic. Jordan participated in the Six Day War (but not the Yom Kippur war) and lost a substantial amount of territory. Unlike Egypt and Israel, Jordan only has ports on the Arabian Sea side of the Suez canal and trades heavily with Europe. Bilateral pairs including Jordan therefore have the largest shocks to distance in the sample. All analyses will be presented excluding Jordan and Jordan will be discussed in detail in a later section.

The panel is unbalanced and only pairs that have data point in the periods before, during, and after the closing of the canal are included in the analysis. Using a balanced panel of country pairs reduces the sample size by nearly one half. Using just the balanced panel does not change the results significantly though it does tend to increase standard errors.

For all the results that follow the sample will be comprised of trade and income for the year 1959 through 1984. This provides for 8 full years before the closing and 8 full years after the reopening, matching the 8 years of the closure.

To present the results graphically, I will collapse the data into three periods, 1) 1966 and earlier, all full years before the closure, 2) 1968-1974, the six complete years with the canal closed, and 3) 1976 and later, all full years after the canal reopened. The partial years (1968 and 1975) are dropped in the collapsed data.

Figure 1: Average bilateral trade residuals grouped by Suez Distance Increase

Source: IMF direction of trade database, author's calculations.
Residuals from a regression with country pair and year dummies.

3 Did the Closure of the Suez Canal Reduce Trade?

Figure 1 shows the average of residuals of the natural log of bilateral trade grouped by the size of the distance shock caused by the closure of Suez. The residuals are from a regression of the natural log of bilateral trade against a full set of time and bilateral pair dummies. For these graphs the sample is limited to country pairs with continuous data from 1959 to 1982. The vertical lines represent the closing and opening of the Suez Canal. There is a clear drop in trade during the closure and the fall is larger for the groups with more extreme shocks. These graphs also suggest that the impact on trade takes several years to reach its peak. In later sections, this time dynamic will be explored more formally.

Figure 2 is a scatter plot analogous to the gravity model estimation described in the previous section. On the x-axis is the log difference between the distance by sea when Suez is closed versus when it is open. Country pairs whose shortest sea routes do not use the Suez Canal (and therefore experience no shock) are omitted from this graph for clarity. About 23 percent of bilateral pairs representing 10 percent of the trade in the sample have the Suez canal as the shortest sea route. The y-axis is the change in log trade. The change in log trade is the difference between average log trade over the periods before, during, and after the closure of the canal. The

Figure 2: Log change in bilateral trade versus Suez Distance Change

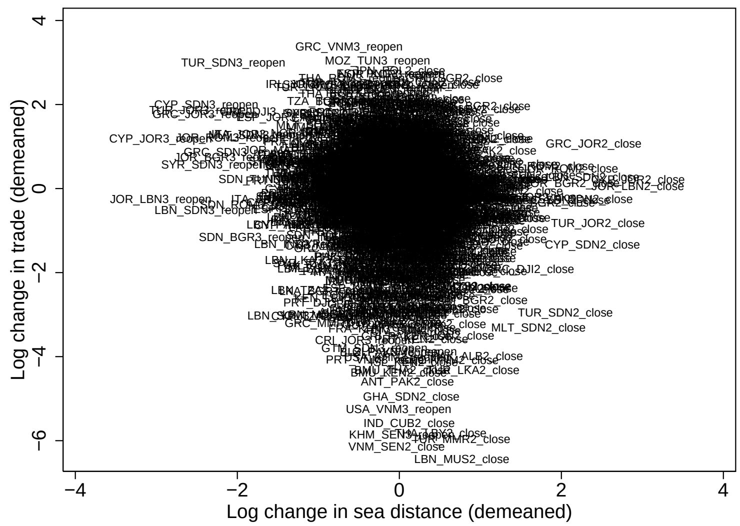

Trade change based on average trade for three periods, 1959-1966, 1968-1974, 1976-1983

partial years (1968 and 1975) are not included in these averages. Larger shocks to distance are associated with slower trade growth after the closure and more rapid trade growth after the reopening. The distribution of shocks is clearly quite skewed, with a small set of countries in the Indian Ocean and the Arabian Sea having the largest shocks. Jordan, in particular, has large distance shocks. This is potentially problematic as Jordan was a participant in the Six Day War and lost a substantial amount of territory (the West Bank and East Jerusalem). Regressions will be presented excluding Jordan from the analysis.

Table 1 shows the results of running panel regression of log trade against sea distance (essentially estimating equation 3). All the regressions include a full set of bilateral pair and year dummies. The bilateral dummies control for the time invariant factors that are typically included in gravity regressions such as common borders, colonial relationships, etc. All identification of the effect of distance on trade is coming from the change in sea distance caused by the closing (and reopening) of the Suez canal. Table 1 also includes regressions where the opening of Suez and the closing of Suez are treated as different shocks. It may be possible that the two shocks had different effects.

Table 1 shows that the elasticity of trade with respect to sea distance is approximately 0.2. A ten percent increase in sea distance reduces trade by about two percent. Though pairs including Jordan have relatively large shocks to sea distance

Table 1: Trade Versus Sea Distance with the Closure of Suez 67-75

VARIABLES	(1) ln(trade)	(2) ln(trade)	(3) ln(trade) xJOR	(4) ln(trade) xJOR
ln(sea distance)	-0.217***		-0.170**	
	(0.054)		(0.070)	
ln(sea dist) (1967)		-0.251***		-0.342***
		(0.066)		(0.081)
ln(sea dist) (1974)		-0.194***		-0.055
		(0.072)		(0.087)
Constant	15.640***	16.298***	15.519***	16.162***
	(0.155)	(0.305)	(0.198)	(0.392)
Observations	68804	68804	67807	67807
R-squared	0.866	0.866	0.867	0.867

*** p<0.01, ** p<0.05, * p<0.1
All regressions include a set of country pair and year dummies.
Standard errors clustered by country pair

(see Figure 2) the elasticity excluding Jordan is not significantly different. The comparison of the up and down shocks does change with the exclusion of Jordan. For the full sample, the positive and negative shocks are not significantly different from each other and both match the elasticity of the combined estimate. For the sample without Jordan, the closing of Suez is unchanged, but the reopening has a smaller point estimate and the difference between the positive and negative shocks is significant.

The estimated elasticity of trade with respect to distance of about 0.2 is relatively small compared to standard gravity model estimates on distance. In an extensive meta study of 103 gravity model studies Disdier and Head (2008) find an average elasticity of about 0.9.

Table 2 shows the results of more conventional gravity model estimation on the same data set used for Table 1. The regressions in this table include individual country dummies, not country pair dummies so the identification is from the cross section as well as the time series. The results are near the center of the results collected in Disdier and Head (2008). The lower coefficients found in Table 1 are therefore being driven by the use of time series variation and not anything inherent in the data set.

There are reasons to think that the traditional estimates are overstated. Typical

Table 2: Trade Versus Sea Distance with the Closure of Suez 67-75

	(1)	(2)	(3)
VARIABLES	ln(trade)	ln(trade)	ln(trade)
ln(air distance)	-1.107***		-0.775***
	(0.029)		(0.057)
ln(sea distance)		-0.999***	-0.344***
		(0.029)	(0.053)
Constant	17.961***	10.881***	15.935***
	(1.183)	(1.193)	(1.230)
Observations	68804	68804	68804
R-squared	0.714	0.708	0.716

*** p<0.01, ** p<0.05, * p<0.1
All regressions include country and year dummies.
Standard errors clustered by country pair

gravity model regressions are run in a cross section with controls for characteristics of the pair such as a shared border, a shared language, or a colonial relationship. Obviously no set of controls can account for all the potential causes of bilateral resistance to trade and the coefficient on distance in such a regression may suffer from missing variable bias if distance is correlated with the missing variables.

These results are also estimating something different from typical gravity model papers. This paper is looking at the elasticity of trade with respect to changes in sea distance, not the point to point distances typically included in gravity regressions. The shock to distance will not affect all trade, just trade carried by sea, while the bilateral trade measures being used accounts for all trade. Column (3) of Table 2 includes both air and sea distance and finds an elasticity of -0.344 for the sea distance. This is comparable to the estimate of -0.217 from column (1) of Table 1 where the pair fixed effects account for the air distances between the pairs.

3.1 Impulse Response Functions

There are also good reasons to think that the estimates of Table 1 are understated. It seems likely that trade did not adjust immediately to the closing of Suez. Figure 1 suggests that trade took about 3 years to reach its low point after the closing and a similar amount of time to reach a new high after the reopening. Since the regressions from Table 1 are essentially comparing means of log trade from the three different periods, the full effect will only be reflected in the coefficient estimates if there is no adjustment path.

The time series nature of the data allows for looking at the time series of the path of trade after the shock. Because the shock is exogenous, the estimation of the time path can be accomplished by including a series of lags of the shock in the regression. The basic specification is:

$$\Delta y_{it} = \alpha + \sum_{k=0}^{M} \beta_k \Delta sea\ distance_{i,t-k} + \gamma_t + \epsilon_{it} \tag{4}$$

where Δy_{it} is the change in log income per capita (or the change in log trade), M is the number of lags, $\Delta sea\ distance_{i,t-k}$ is the trade weighted average change in sea distance for country i in year $t-k$, γ_t is a set of year dummies, and ϵ_{it} is an error term. A full set of year dummies can also be included, giving each country an individual trend. Doing so does not change the results in any significant way. Standard errors are clustered at the country level in all regressions.

The impulse response functions shown in the rest of the paper are constructed by summing the β coefficients from estimating equation (4). The response in the contemporaneous period is β_0, for the second period $\beta_0 + \beta_1$, and so on up to the total number of estimated lags.

$$response_t = \sum_{k=0}^{t} \beta_k \tag{5}$$

In each case the standard error of the sum is calculated. All impulse response function graphs include error bands of two standard errors.

Figure 3 plots the time path of trade after a permanent shock to sea distance. The magnitude is analogous to the elasticity estimates from Table 1. The response function suggests that it takes roughly two years for the shock to have its full impact with a long run elasticity of about 0.4, which is very similar to the sea distance results of column (3) in Table 2 and larger than the Table 1 regressions where the initial low elasticities are essentially averaged with the long run elasticities.

There may also be reasons to think that shock of opening and closing are not identical. The closing of the canal was a surprise (fifteen ships were caught inside the canal) while the time frame for reopening of the canal was known roughly a year beforehand. Figure 4 shows separate impulse response functions for the opening

Figure 3: The Response of bilateral Trade to Suez Distance Shocks

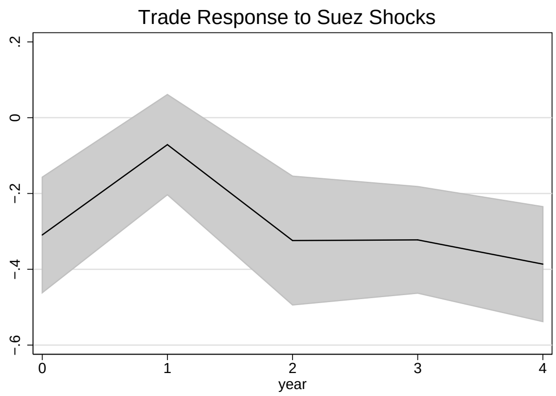

Shaded bands represent plus or minus two standard errors

Figure 4: The Response of Trade to Suez Distance Shocks

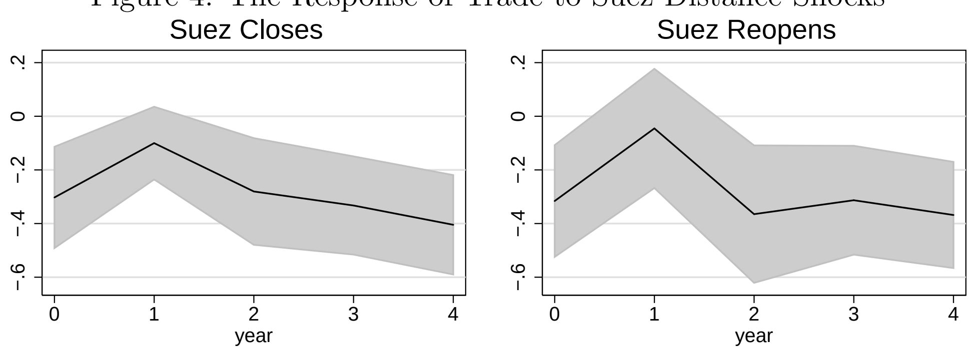

Shaded bands represent plus or minus two standard errors

and closing of Suez. They are both drawn representing a positive shock to distance for comparative purposes. The opening and closing of the canal do not appear to generate substantially different time paths for trade. Both the up and down shocks generate an elasticity of roughly 0.4 when the full effect is in place.

3.2 What about Jordan?

As discussed earlier, trade involving Jordan experienced a particularly large shock to distance. At the same time, Jordan had direct consequences from the Six Day War which caused the closure in the first place. In order to see if Jordan is driving the time path of trade in response to the shock, the impulse response functions were also drawn using a sample excluding pairs that include Jordan. Figure 5 shows the overall impulse response function for the non-Jordan sample. The shape and magnitude are very similar to Figure 3 which includes Jordanian trade.

Figure 6 separates the effects of closing and opening the Suez canal for the sample without Jordan. Again, the exclusion of Jordan does not seem to significantly change the shape and magnitude of the impulse response. This suggests that Jordan's responses to the Suez shocks, though large, are similar in timing to rest of the sample. If we thought that there was a disparate effect on Jordan related to the war itself we would not expect to see the effect of including Jordan to be symmetric. In fact, the separate up and down results of Table 1 suggest that the inclusion of Jordan is more important to the results of opening the canal, not the closing.

3.3 Conclusions for Trade and Suez

The closure and reopening of the Suez Canal appear to be useful shocks for thinking about changes in the costs of trade between nations. The long run elasticity of trade with respect to sea distance is roughly 0.4, with the adjustment process taking two to three years. The response to the reopening of the Suez Canal appears to be the same in magnitude and time path as the closing of the canal. This suggests that the true effect of transport costs on trade is being identified and not some indirect effect of the Six Day War. These results are robust to the exclusion of Jordan and the symmetry of responses with and without Jordan suggest that the primary cause of trade movements in Jordan after the Six Day war was the closing of the canal, not the direct effect of involvement in the war.

Figure 5: The Response of Trade to Suez Distance Shocks - No Jordan

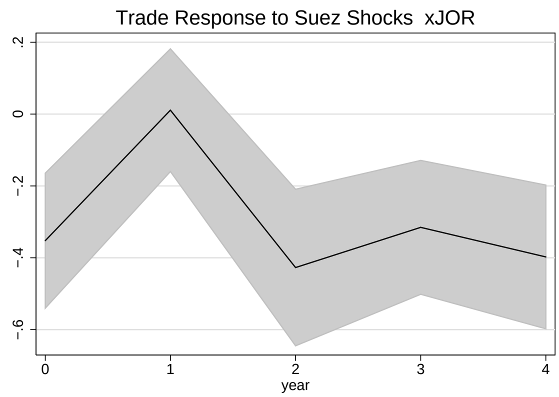

Shaded bands represent plus or minus two standard errors

Figure 6: The Response of Trade to Suez Distance Shocks

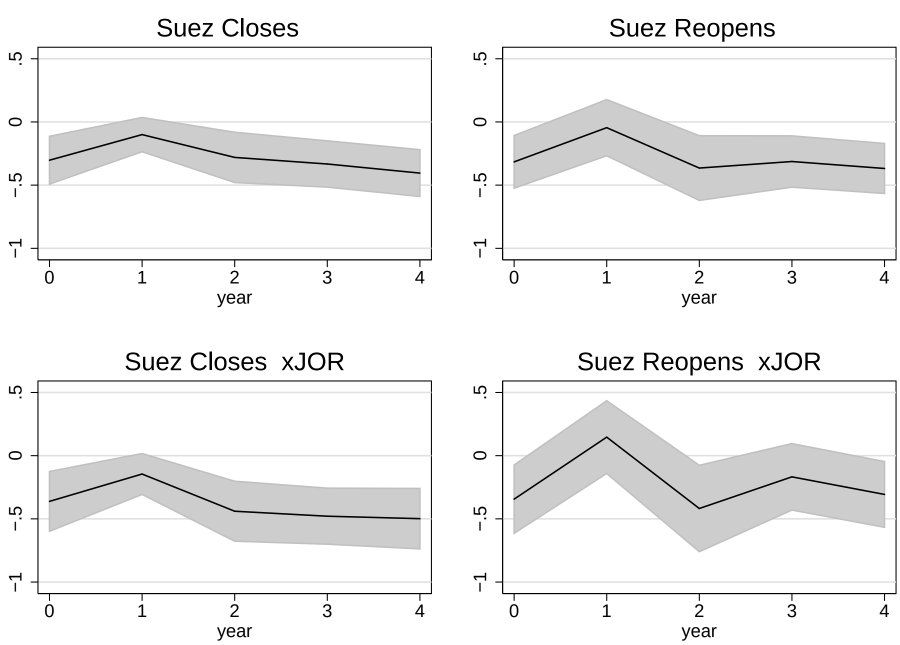

Shaded bands represent plus or minus two standard errors

4 Trade and Income

The previous section establishes that the closing and reopening of the Suez Canal effected bilateral trade between partners whose shortest sea route is through the Suez Canal. For any individual country these changes in distance were exogenous and generated entirely through differences in geography. Different countries were differentially effected depending on their geography and pre-existing trade patterns. These shocks to trade can therefore be used to identify the impact of changes in trade on income at the aggregate country level.

4.1 Predicting Aggregate Trade

The coefficients reported in Table 1 can be used to construct predicted values for bilateral trade for each pair of countries for each year. The predicted values are derived from equation (3) and are therefore comprised of a time effect, a bilateral pair effect and the distance effect. These predicted trade volumes can be summed in order to arrive at a prediction for aggregate trade in each country for each year.

These predictions can be made out of sample. As long as there is a single observation of bilateral trade between two countries, an estimate for the bilateral pair can be generated in every year since distance is always available. This has the advantage of keeping the set of bilateral pairs constant over time for the predicted trade, avoiding the problem of changes to aggregate trade driven by the appearance and disappearance of trade data for a particular pair.

Because the goal is to instrument the actual trade share with the predicted trade share in a regression of trade on per capita GDP, these out of sample predictions create some difficulties because there are observations where there is a predicted trade value, but not an actual trade value. This matters because the instruments and observations of trade volumes need to be matched for the IV regressions.

Two different methods are used to deal with these holes. First, the missing values of real trade are imputed using a full set of country pair and time dummies. These imputations are based entirely on information that is controlled for in the second stage and should not affect the results. They are only necessary to keep the scaling of the actual changes in trade consistent. In order to confirm that these imputations are not driving the results I can also generate results where the sample is restricted to country pairs with a full panel of observations. This eliminates out of sample predictions and imputations at the cost of losing almost half of the trade observations.

Following Frankel and Romer (1999), unlogged versions of these bilateral relationships are summed to obtain a prediction for total trade for each country. The actual trade figures are similarly summed to arrive at a value for total trade.

$$\begin{aligned} predicted\ trade_{it} &= \sum_{i \neq j} e^{\gamma_t + \gamma_{i,j} + ln(sea\ distance_{ijt}) * \beta} \\ &= e^{\gamma_t} \sum_{i \neq j} e^{\gamma_{ij}} e^{ln(sea\ distance_{ijt}) * \beta} \end{aligned} \tag{6}$$

For the predictions using country-pair dummies, the country pair effects act as weights in an average of distances. Because the country level regressions will include country and time fixed effects, all the identification will be from the within country variation over time. None of the identifying time variation is generated from the bilateral or time effects.

Additionally, I can construct a simpler and somewhat more transparent set of instruments. A weighted average of the distance change across all trading partners (using the average trade over the whole sample as the weight) will give me the average log distance change per unit of trade for each country. If I run a regression using this average log distance change over the change in log aggregate trade I should get a coefficient that is approximately equal to the β from the bilateral level regression.

$$Suez\_Shock_i = (\overline{trade_i})^{-1} \sum_{i \neq j} (ln(seadist_{noSuez}) - ln(seadist_{Suez})) * \overline{trade_{ij}} \tag{7}$$

Table 3 lists the countries with the twenty largest shocks as calculated by equation 7. The list is obviously very regional, with the countries closest to the canal having the largest shocks of over 50 percent followed by Pakistan and India with about a 30 percent increase in average distance. Many east Asian countries experience a shock in the 10 percent range.

Both the predicted trade from the trade regression and the weighted average of the changes in distance derive all of their idiosyncratic variation from the opening and the closing of the Suez Canal. Since the Suez canal shocks are exogenous with respect to the majority of individual countries, this variation should provide a useful instrument for investigating the impact of trade on GDP.

Table 3: Trade weighted Increase in Sea Distance from Suez Closure

Country	Trade Weighted Distance Increase
Jordan	95.9
Sudan	70.2
Djibouti	58.6
Somalia	43.9
Pakistan	31.4
India	29.9
Kenya	23.1
Sri Lanka	22.2
Tanzania	19.9
Lebanon	13.9
Malaysia	13.3
Madagascar	13.0
Cambodia	12.2
Myanmar	11.2
Mauritius	11.2
Romania	11.1
South Korea	11.0
Singapore	10.7
Vietnam	10.7
Thailand	9.8

Figure 7: The relationship between output and trade

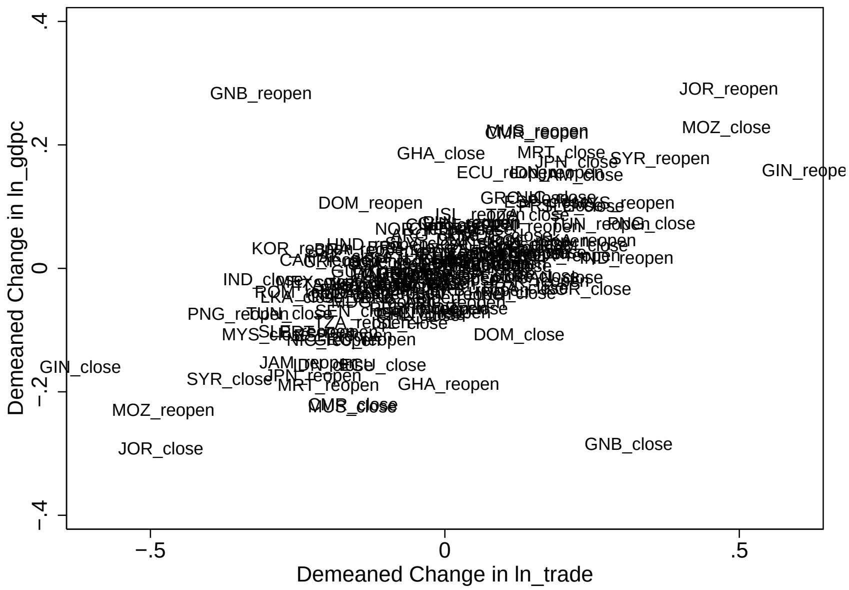

4.2 OLS Regression of Income on Trade

It is well known that trade and GDP are very highly correlated in the time series. Figure 7 shows a scatter plot of changes in trade versus changes in GDP per capita over the three major periods of this investigation. Both variables have been demeaned by country and time so this is a visual representation of a regression in differences with the inclusion of both time dummies and individual country time trends.

Table 4 shows the results of regressing trade on GDP per capita in a regression with a set of country and time dummies. The dependent variable is real per capita income from the Penn World Tables and the key independent variable is the volume of trade from the DOT database described earlier summed at the individual country level. There is obviously a strong and significant relationship between trade and income with an elasticity of about one third. The exclusion of Jordan makes no difference in the outcome, though Jordan's trade change is one of the extreme values over this period.

The OLS regressions are, of course, unidentified since we do not know the direction of causality. In the next sections, instruments based on the shock to trade from the closure of the Suez canal will be used to establish a causal link between

Table 4: Trade versus GDP per capita

VARIABLES	(1) ln(gdpc)	(2) ln(gdpc) xJOR
ln(trade)	0.334***	0.334***
	(0.037)	(0.038)
Constant	1.028	0.943
	(0.823)	(0.781)
Observations	1957	1931
R-squared	0.983	0.983

*** p<0.01, ** p<0.05, * p<0.1
Regressions include country pair and year dummies.
Standard errors clustered by country

trade and output.

4.3 First Stage

Do the instruments provide a good prediction for aggregate trade? Figure 8 shows a scatter plot of actual trade changes versus the average distance change caused by the closure and reopening of Suez. The top panel includes Jordan and the bottom panel is the same scatter excluding Jordan. Table 5 shows the results of regressions of this relationship in the full panel of data.

In all cases there appears to be a strong first stage whether or not Jordan is included or not. The t-statistics imply an F-stat of over 5.5 of the instruments on the instrumented variables in all four columns with much higher values when Jordan is included. This is obviously consistent with the earlier part of the paper which established that the closure of the Suez Canal affected trade. It is not surprising that the relationship between the closing of Suez and trade holds once the data is aggregated.

Figure 8 does illustrate that Jordan is an extreme value in this analysis, at least in the sense of having the largest treatment. The regressions do, however, suggest that Jordan is not an outlier in the sense of moving point estimates. The point estimates are not significantly different when Jordan is excluded and the results are still significant (but the standard errors are substantially higher). Jordan is clearly on the regression line that is found when it is excluded.

Figure 8: Log change in trade versus Suez Distance Shock

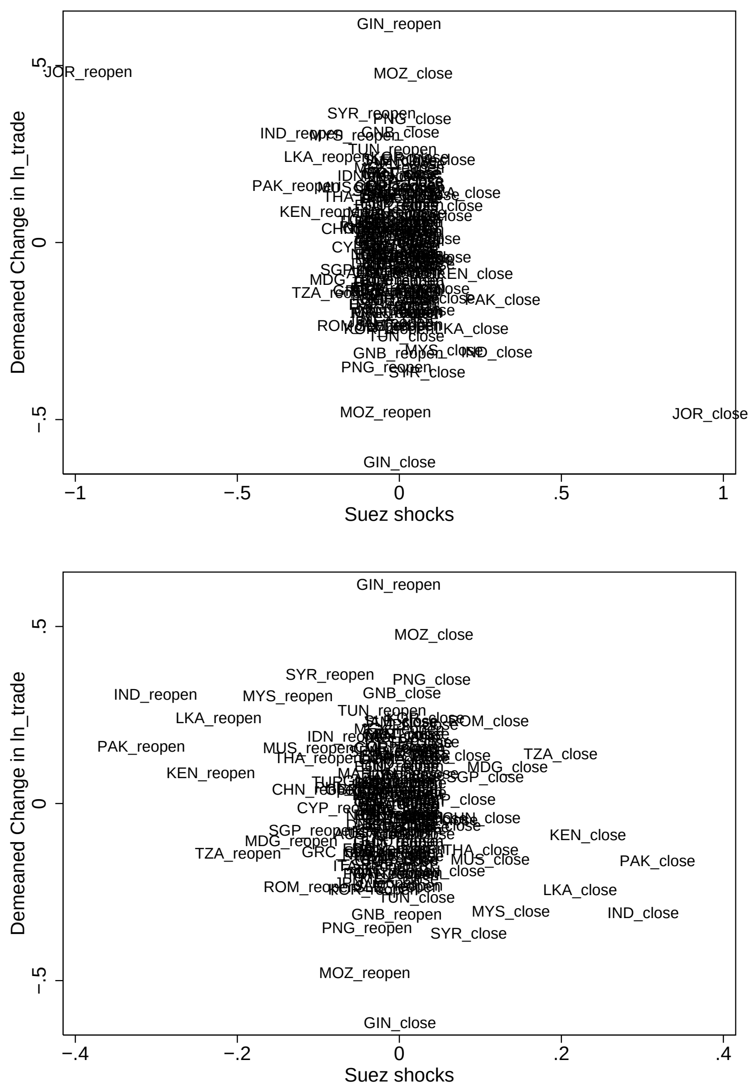

Table 5: First Stage Regressions – Predicted Trade versus Actual Trade

	(1)	(2)	(3)	(4)
VARIABLES	ln(trade)	ln(trade)	ln(trade) xJOR	ln(trade) xJOR
ln(Predicted Trade)	2.614***		2.963**	
	(0.399)		(1.245)	
Suez Shock		-0.584***		-0.669**
		(0.088)		(0.275)
Constant	-33.265***	20.610***	-42.775	20.541***
	(8.214)	(0.040)	(27.209)	(0.050)
Observations	1957	1957	1931	1931
R-squared	0.978	0.978	0.979	0.979

*** p<0.01, ** p<0.05, * p<0.1
All regressions include a set of country pair and year dummies.
Standard errors clustered by country

Table 6: Reduced Form – The Effect of Predicted Trade on GDP per capita

	(1)	(2)	(3)	(4)
VARIABLES	ln(gdpc)	ln(gdpc)	ln(gdpc) xJOR	ln(gdpc) xJOR
ln(Predicted Trade)	1.341***		1.040*	
	(0.228)		(0.621)	
Suez Shock		-0.304***		-0.212
		(0.052)		(0.133)
Constant	-19.802***	7.831***	-14.355	7.794***
	(4.697)	(0.025)	(13.568)	(0.026)
Observations	1957	1957	1931	1931
R-squared	0.974	0.974	0.974	0.974

*** p<0.01, ** p<0.05, * p<0.1
All regressions include a set of country pair and year dummies.
Standard errors clustered by country

Table 7: IV Regressions – The Effect of Trade on GDP per capita

	(1)	(2)	(3)	(4)
VARIABLES	ln(gdpc)	ln(gdpc)	ln(gdpc) xJOR	ln(gdpc) xJOR
ln(trade)	0.513***	0.520***	0.351**	0.318*
	(0.100)	(0.110)	(0.170)	(0.168)
Constant	-3.364	-3.542	2.211	2.715
	(2.485)	(2.729)	(2.666)	(2.617)
Observations	1957	1957	1931	1931
R-squared	0.980	0.980	0.983	0.983

*** p<0.01, ** p<0.05, * p<0.1
All regressions include a set of country pair and year dummies.
Standard errors clustered by country

4.4 The Reduced form

Figure 9 shows a scatter plot of the log change in GDP per capita versus the average distance change caused by the closure and reopening of Suez. The top panel includes Jordan and the bottom panel is the same scatter excluding Jordan. Table 6 shows the same relationship in regressions on the full panel of data.

Once again, the exclusion of Jordan raises the standard errors without a substantial change in the point estimates. Unlike the first stage regressions the exclusion of Jordan increases the standard errors enough that the results are only marginally significant without the Jordanian data points.

4.5 IV results

Table 7 shows the results of IV regression where actual trade is instrumented with the predicted values of trade derived earlier and the average distance change caused by the closing and opening of Suez. Columns (1) and (3) are instrumented with predicted trade and columns (2) and (4) are instrumented with the trade weighted average sea distance. These correspond to the columns in Tables 5 and 6.

The results for the sample including Jordan are highly significant and suggest an elasticity of about one half between trade and income. Excluding Jordan the elasticity is reduced to 0.35 and the standard errors rise. The regressions remain significant, though only at the 5 and 10 percent level. The results from the sample without Jordan, unlike the full sample are also fragile. Changes to the specification that would tend to reduce power generally cause increases in the standard errors and render the estimates insignificant. Examples include restricting the sample to years 1960 through 1982 (instead of 1959 and 1983), or estimating on the balanced panel. The results with Jordan are robust to these different samples.

An elasticity of 0.5 is smaller than the elasticities found in Frankel and Romer (1999), but within their error bands and is very similar to the elasticities found in Feyrer (2009) which relies on the rise in the relative importance of air travel for identification. This second comparison is interesting because the use of air travel allows for things other than trade such as movements of people to play a role. Because the identification is coming from the change in sea distance, these estimates are much more clearly identifying the effect of trade in goods and not integration in general.

Figure 9: Log change in GDP per capita versus Suez Distance Shock

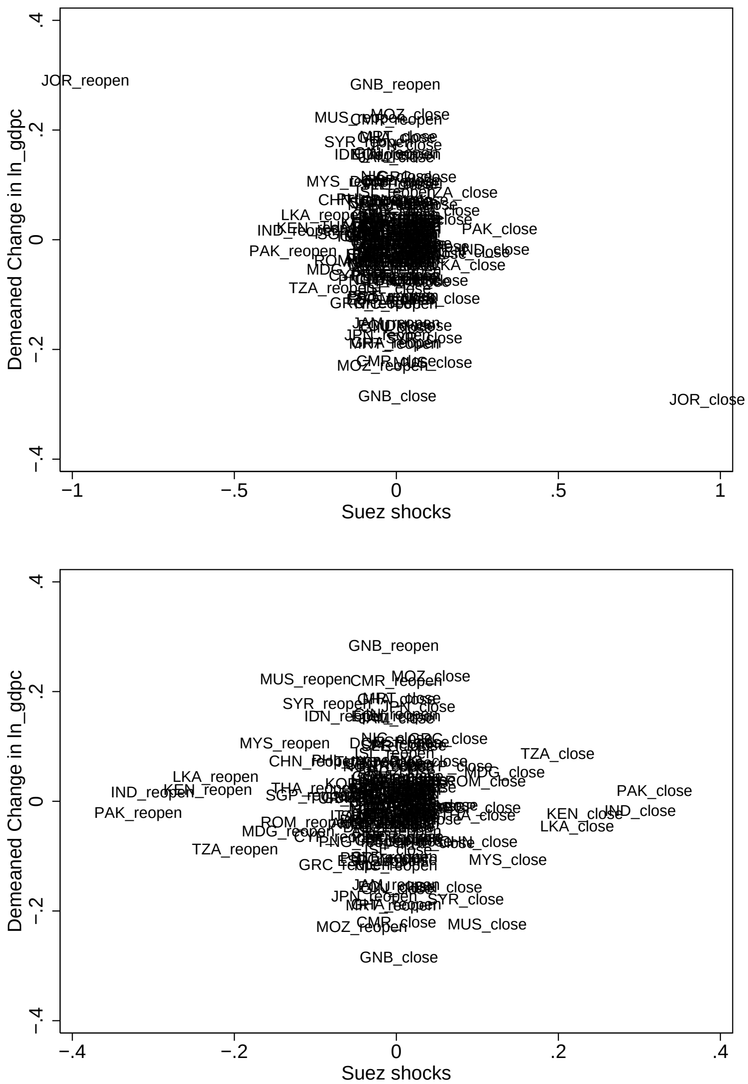

4.6 Impulse Response functions at the country level

The earlier impulse response functions were drawn for data at the country pair level. The same exercise is possible for the aggregated country level data. The advantage of this approach is that we can draw the time path of the shocks on output as well as trade. Figure 10 shows the impulse response function of country level trade to the shock of closing Suez, where the shock is measured as the trade weighted average of the increase in distance caused by the closing of Suez described in equation (7). Figure 11 shows the impact of the same shock on GDP. Figure 12 shows the impact of a shock to trade where trade is instrumented by the shocks to distance caused by Suez used in the first two graphs. One can think of these response functions as dynamic versions of the first stage, reduced form, and IV regressions presented in Tables 5, 6, and 7. The magnitudes can be though of as elasticities.

These graphs allow us to think about the eventual magnitude of the impact of closing Suez. The trade response graph provides the elasticity of trade with regards to an increase in sea distance. This is an aggregated version of Table 1 and Figure 3. The 3 year elasticity of 0.5 is not significantly different from the result in Figure 3 of about 0.4. The response of GDP to the Suez shock has a similar shape.

The impulse response function estimated by instrumenting actual trade with the shock to Suez is useful because it provides the correct scaling of the effect of trade increases on GDP per capita. This graph is essentially the reduced form response function (Figure 11) scaled by the first stage response function (Figure 10). The first stage response function shows that the trade response has significant time lags. The reduced form mirrors this. The IV results can distinguish between lags in the impact of the shock to GDP that come from the slowness of the trade response and true lags in the response of GDP to changes in trade. Figure 12 suggests that the effect on GDP of a shock to trade is rapid, with 2/3 of the ultimate effect happening within the contemporaneous period. The long run elasticity of trade to GDP is roughly 0.8, at the high end of estimates by Feyrer (2009), but with overlapping confidence intervals.

Similar graphs can be drawn excluding Jordan from the sample. Figure 13 shows the response of trade and GDP to the shocks to Suez with and without including Jordan in the sample. As in the formal econometric results, the trade response functions are very similar whether or not Jordan is included and the error bands clearly overlap. For the sample without Jordan, the effect on trade is significantly

Figure 10: The Response of Trade to Suez Distance Shocks

Shaded bands represent plus or minus two standard errors

Figure 11: The Response of GDP to Suez Distance Shocks

Shaded bands represent plus or minus two standard errors

Figure 12: The Response of GDP to Trade Shocks (IV)

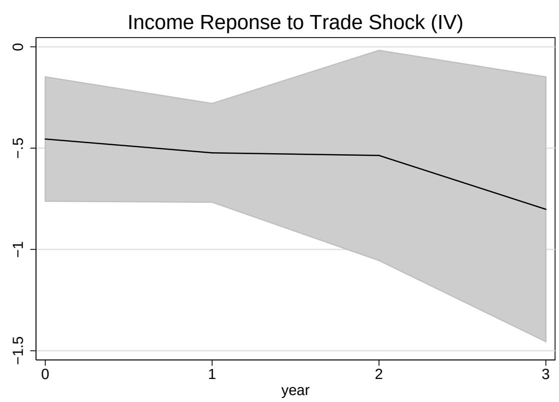

Shaded bands represent plus or minus two standard errors

Figure 13: The Response of Trade and GDP to Suez Distance Shocks

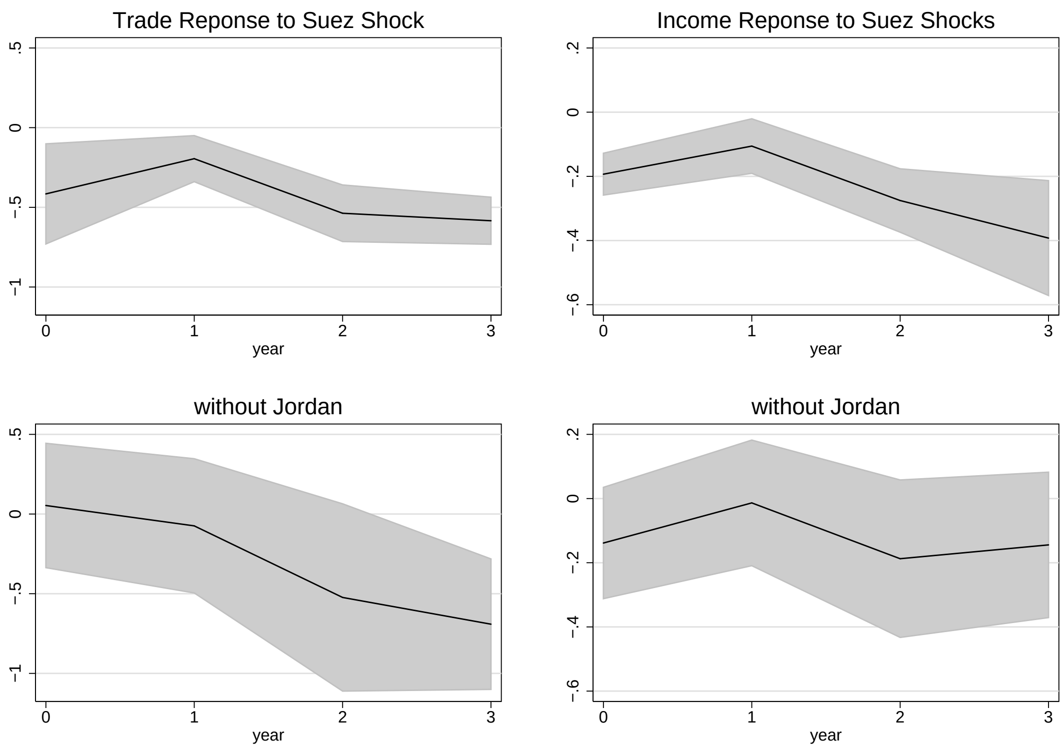

Shaded bands represent plus or minus two standard errors

different from zero three years after the shock. For income the error bands also overlap, but the decrease in precision makes the effect of the shock insignificantly different from zero. The point estimate for the effect of trade on GDP is also somewhat smaller for the non Jordan sample, at something less than 0.2 rather than 0.4. Scaling this to the first stage results this suggests an elasticity of roughly 0.4 for trade and income for the non-Suez sample, about half of the full sample value.

5 Is Jordan an Informative Data Point?

Jordan is clearly an important data point and the strength of the results varies based on whether Jordan is included. It is therefore important to address whether Jordan is a useful and usable data point. The problem with including Jordan is twofold. First, Jordan has an extreme treatment value that is over twice as large as any other country. Second and more importantly, Jordan was involved in the hostilities that precipitated the closing in the first place. The fact that Jordan was on the losing side of the Six Day War may depress Jordan's economic activity at the same instant that the Suez Canal was closed. The key question is whether Jordan saw a reduction in trade and output from 1967 to 1975 because of the shock to trade or because of the aftermath of war.

One way to examine this question is to use the fact that identification in this paper is actually coming from two shocks that are of equal and opposite sign, the closing of the Suez canal in 1967, and the reopening of the canal in 1975. If the events of the Six Day War make Jordan a suspect data point for the closing of Suez, it is not clear that this is true for the reopening. Figure 14 shows the time path of trade and GDP per capita for Jordan before, during and after the closure of Suez. The data presented is from residuals from aggregate regressions of the natural log of trade and GDP per capita against a full set of time and country dummies.

The time paths do not suggest that Jordan faced a single shock in 1967 due to the war. To the contrary, the negative impact on trade and income seems perfectly timed to the closing and reopening of the canal. Even if we think that the closing of the canal corresponded to large negative war shocks to Jordan, the positive shock of the reopening looks remarkably similar and occurred without other coincident events.

Figure 14: Trade and GDP in Jordan and the closing of Suez

Residuals from a regression with country and year dummies

Figures 4 and 6, presented earlier, tell a similar story. The shape of the trade responses to the closing and reopening of Suez are roughly symmetrical whether Jordan is included in the analysis or not. If there were dramatic differences between the two when Jordan was included, this would suggest that Jordan's response was differentially affected by one shock versus the other. This would be exactly what we would expect to see if Jordan's response is being driven directly by the war and not by the trade effects of the closing of Suez. The fact that we see no signs of asymmetry when Suez is included suggests that the impact of closing the canal on Jordan's trade is being driven by the same trade cost changes seen by everyone else in the sample.

5.1 Conclusions for Trade and Income

The use of shocks to trade generated by the closing and reopening of the Suez Canal provide clean identification on the impact of trade on income. The results suggest that increases in trade volumes generated by decreases in trade costs generate higher income per capita. The elasticity of income with respect to trade appears to be about one half, with the effects of increasing trade occurring rapidly. Lags in the effect of decreasing trade costs are on the order of two to three years with most of the delay coming in response of trade volumes to trade costs.

The results are much stronger when the country most effected by the closing of the Suez Canal, Jordan, is included in the sample. Excluding Jordan leads to much higher standard errors and in some cases, the results become insignificant. The point estimates, however, are largely insensitive to the inclusion or exclusion of Jordan. Jordan appears to be on the regression line in most specifications. Jordan also appears to have symmetrical responses to the opening and closing of Suez, suggesting that the shocks to trade may be treated as exogenous for Jordan despite Jordan having been a participant in the Six Day War.

6 Conclusions

This paper attempts to use the shock provided by the temporary closure of the Suez Canal as a natural experiment. The movements in trade costs generated by closing Suez can be usefully thought of as an exogenous shock effecting most countries in the world. This shock is useful for identifying the impact of trade costs on trade and furthermore the effect of trade on income. To summarize, the Suez Canal had a significant and robust affect on bilateral trade patterns. Aggregating these changes to trade suggests that trade has a significant affect on output.

The nature of the canal shock makes it unique. First, the shock was sudden and short term. We have precise dates when the shocks took place. Second, the shocks are very precisely targeted at trade by sea. Generally when we consider instruments for trade, they can potentially act through channels that go beyond trade. Since the variation in this paper is being provided by the Suez closure, any channels other than trade need to involve bilateral relationships between countries that involve travel by sea. It is hard to imagine anything other than trade in goods than fits this description.

The ability to get clean identification on the effect of trade on goods on output is potentially useful when considering the effect of policies designed to reduce trade costs between nations. This paper suggests that while activities that are related to trade such as foreign direct investment and multinational participation may be important, simple increases in the raw volume of trade may increase income. This may be useful in evaluating policies intended to increase trade such as tariff reductions.

References

Anderson, James E. and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, March 2003, *93* (1), 170–192.

Baldwin, Richard and Daria Taglioni, "Gravity for Dummies and Dummirs for Gravity Equations," Working Paper 12516, National Bureau of Economic Research September 2006.

Coulibalya, Souleymane and Lionel Fontagne, "South – South Trade: Geography Matters," *Journal of African Economies*, 2005, *15* (2), 313–341.

Disdier, A.C. and K. Head, "The Puzzling Persistence of the Distance Effect on Bilateral Trade," *The Review of Economics and Statistics*, 2008, *90* (1), 37–48.

Dollar, David, "Outward-Oriented Developing Economies Really Do Grow More Rapidly: Evidence from 95 LDCs, 1976-1985," *Economic Development and Cultural Change*, 1992, *40* (3), 523.

Edwards, Sebastian, "Openness, Productivity and Growth: What Do We Really Know?," *The Economic Journal*, 1998, *108* (447), 383–398.

Feyrer, James, "Trade and Income – Exploiting Time Series in Geography," Working Paper 14910, National Bureau of Economic Research April 2009.

Frankel, Jeffrey and David Romer, "Does Trade Cause Growth?," *American Economic Review*, June 1999, *89* (3), 379–99.

Glick, R. and A.M. Taylor, "Collateral Damage: Trade Disruption and the Economic Impact of War," *Review of Economics and Statistics*, 2008. forthcoming.

Heston, Alan, Robert Summers, and Bettina Aten, "Penn World Table Version 6.2," Technical Report, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania September 2006.

Irwin, D.A. and M. Terviö, "Does trade raise income? Evidence from the twentieth century," *Journal of International Economics*, 2002, *58* (1), 1–18.

Meehl, G.A., "Observed World Ocean Seasonal Surface Currents on a 5-degree Grid," *NCAR Technical Note*, 1980.

Rodriguez, Francisco and Dani Rodrik, "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence," *NBER Macroeconomics Annual*, 2000, *15*, 261–325.

Rodrik, D., A. Subramanian, and F. Trebbi, "Institutions Rule: The Primacy of Institutions Over Geography and Integration in Economic Development," *Journal of Economic Growth*, 2004, *9* (2), 131–165.

Sachs, Jeffrey D. and Andrew Warner, "Economic Reform and the Process of Global Integration," *Brookings Papers on Economic Activity*, 1995, *1*, 1–118.

Summers, R. and A. Heston, "The Penn World Tables Mark 5.5 (updated supplement to The Penn World Tables (Mark 5): An Expanded Set of International Comparisons 1950-1988)," *Quarterly Journal Of Economics*, May 1991, *106*, 327–368.