

Testing the Heckscher-Ohlin-Vanek Theory with a Natural Experiment

Assaf Zimring*

Abstract

This paper uses the historical episode of the near-elimination of commuting from the West Bank into Israel, which caused a large and rapid expansion of the local labor force in the West Bank, to test the predictions of the Heckscher-Ohlin-Vanek (HOV) model of trade. I use variation between districts in the West Bank to test these predictions, and find strong support for them: Wage changes were not correlated with the size of the shock to the district labor force (Factor Price Insensitivity); Districts that received larger influx of returning commuters shifted production more towards labor intensive industries (Rybczynski effect); And on the consumption side, the data are consistent with the assumption of identical homothetic preferences, which, combined with the production results, supports the Heckscher-Ohlin-Vanek theorem on the factor content of trade.

1 Introduction

According to the Heckscher-Ohlin theory, all else equal, countries will tend to export those goods whose production is intensive in the factors they have in

*assafzim@gmail.com. I thank Doireann Fitzgerald, Kalina Manova, Alan Deardorff, Andrei Levchenko, Kyle Bagwell, Kyle Handley, Sebastian Sotello, Javier Cravino, Kei-Mu Yi, Haggay Etkes, and Ran Abramitzky, for many useful comments. I am grateful to the Upjohn Institute for Employment Research for their generous financial support. Vamika Bajaj and Jialin Liu provided excellent research assistance.

relative abundance. One of the most important formulations of this insight, the Heckscher-Ohlin-Vanek theorem, states this result in terms of the so-called “factors content of trade”: Countries will be net exporters of the services of factors they have in relative abundance, embodied in the goods they trade. The theoretical appeal of the HOV framework has made it one of the pillars of neo-classical international trade theory. However, the long history of its empirical tests gives the theory little, if any, support. As Davis and Weinstein (2001) put it in an important paper: “The prediction [of the HOV theorem] is elegant, intuitive, and spectacularly at odds with the data.” An important reason for these empirical limitations of the theory is that the ‘all else equal’ assumption in the HOV theory is a very strong one indeed. It assumes, essentially, that factor proportions are the only difference between countries. Accordingly, most of the modifications that were introduced into the theory in order to reconcile it with observed trade flows revolved around replacing these assumptions with more general ones, allowing for various kinds of technological differences, differences in preferences, and so on.

In this paper, I take a different approach. Instead of using the theory to explain trade data in a cross-section of countries, where the ‘all else equal’ assumption seems problematic, I use it to explain changes in production and trade patterns in a number of small open economies, following a large, exogenous, and persistent shock to their factor abundance. The use of this “natural experiment” allows me to test the core of the HOV theory without having to take a strong stand on the nature of the differences between the economies under study.¹ In essence, instead of modeling the differences between the economies I study, this empirical setting allows me to “difference them out”. It is, to my knowledge, the first paper to test the HOV formulation of the Heckscher-Ohlin insights using a natural experiment. The results are encouraging for the theory: All three of the relevant predictions of the HOV framework, i.e. Factor Price Indifference, the Rybczynski effect, and the HOV

¹I use the term “natural experiment” only to mean that this historical episode created an exogenous variation in the increase in the size of the labor force. This is not a regression paper, where the natural experiment is used to identify a parameter of the model: the HOV model has no parameters.

theorem, are supported by the data.

The historical episode I use in this study is the near elimination of commuting from the West Bank into Israel in the year 2000. Until that point, around 20% of the labor force in the West Bank commuted to work in Israel on a daily basis. In October 2000, following the outbreak of the Second Intifada, the number of commuters to Israel was severely restricted by the Israeli government, and remained low for many years after. The immediate result of this policy change was a large increase in the effective size of the domestic labor force available for work in the West Bank, making the West Bank substantially more labor-abundant.

Two important features of this historical episode make it especially suitable for testing the predictions of the HOV model. First, throughout the whole period in question, while the movement of people across the border was severely restricted, the movement of goods was not, and the West Bank traded extensively with the world.² Second, large variation in commuting patterns between the different districts of the West Bank before 2000 led to large variation in the size of the shock to their labor force when commuting declined. Moreover, the data show that due to extremely limited mobility between districts, the variation in the size of the initial shock had a persistent effect on the size of the labor force in each district, and it is clearly visible in the data 8 years after commuting declined. It is this variation that I use to test the predictions of the HOV theory.

In the West Bank context, comparing a district which received a large influx of returning commuters to one which received a smaller influx, the HOV theory makes three major predictions. First, the Rybczynski effect predicts that a district which received a larger influx of returning commuters will experience a larger shift in the composition of production towards labor-intensive industries. Second, the Factor Price Insensitivity (FPI) result predicts that such a district will *not* experience slower wage growth, in spite of the larger

²That is not to say that trade was completely free. What is important for our purpose here is that whatever restrictions existed, they didn't change much following the outbreak of the Second Intifada.

increase in labor supply. And third, the HOV theorem predicts that a district that received a larger influx of returning workers will increase its net exports of labor services, embodied in labor-intensive goods, more than a district which received a smaller influx. The first two predictions (Rybczynski effect and FPI) are directly confirmed by the data. Unfortunately, since no data on “imports” and “exports” of districts in the West Bank exist, the third prediction cannot be tested directly. However, it is supported indirectly by data on consumption patterns. Put together, these results provide substantial empirical support for the HOV theory.

What use, however, is a test of the HOV theory, a theory of trade between real countries, that instead of explicitly modeling the real differences between them, finds a way around these differences? To answer this question, one first needs to explain what use is the HOV theory to begin with. The HOV model, if true, can be used for two different purposes. It can explain observed global trade patterns based on global factor endowments, and it can predict the effects that various shocks in an open economy — to terms of trade, to factor endowments, etc. — will have on domestic factor markets and on trade patterns. Most of the literature showed that the basic model performs the former task poorly, and tried to elaborate on the theory to make it more compatible with observed trade flows. This paper shows that even in its most basic version, the HOV model performs the latter task well.

2 Relation to the Literature

This paper contributes to the large literature on empirical tests of the HOV model. Essentially all of this literature, dating back to the famous “paradox” discovered by Leontief (1953), and including the seminal work of Stern and Maskus (1981), Maskus (1985), Bowen, Leamer, and Sveikauskas (1987), and Harrigan (1995), find that the theory, at least in its most basic form, does very poorly in predicting trade patterns. Later work, such as Treffer (1993), Treffer (1995), and Davis and Weinstein (2001) therefore focused on documenting the ways in which observed trade patterns deviate from those predicted by the

HOV theory, and suggested modifications to the theory, such as productivity differences between countries, home bias in consumption, trade costs, and the existence of non-tradeable goods, which greatly improved the predictive power of the model. The main contribution of this paper is that instead of expanding the theory by adding and relaxing assumptions to make it more compatible with observed international trade data, I use an historical episode where the ‘all else equal’ assumption of the HOV model is plausible, but nonetheless the model makes non-trivial predictions.

Bernhofen and Brown (2011) use the natural experiment of Japan’s move from autarky to free trade in the mid-nineteenth century to test what they refer to, following Deardorff (1982), as the general validity of the Heckscher-Ohlin model, or HOD (Heckscher-Ohlin-Deardorff). Relying on numerous sources for factor prices and production techniques, they find empirical support for the main testable prediction of the HOD, which states that evaluated at factor’s autarky prices, the value of the factor content of trade is (weakly) positive. In addition to using more standardized data, from a more recent historical episode, the contribution of this paper relative to Bernhofen and Brown (2011) is that it tests one of the most special versions of the Heckscher-Ohlin model, maintaining all the strong assumptions of the HOV formulation, and not the most general formulation, which Bernhofen and Brown test. This allows for more intuitive results about the trade prediction, and it allows me to test not just the trade prediction, but also the Rybczynski effect, and the Factor Price Insensitivity prediction.

A number of studies tested Heckscher-Ohlin type predictions for different regions within the same country, as I do in this paper. Horiba and Kirkpatrick (1981) perform a cross-sectional test, using Leontief (1953) methodology, for US regions for 1963. Davis, Weinstein, Bradford, and Shimpo (1997) use data from prefectures in Japan to test the predictions of the HOV model for production and for consumption, and find results that are almost perfectly consistent with theory. However, using data on different regions in the same territory raises the issue of the mobility of factors as an alternative explanation for the findings, and thus of interpreting correlations as causality: Did labor flow into

districts with industries that are labor-intensive, or did districts with large labor endowment specialize in labor-intensive sectors? Were wages equalized by the migration of labor to areas with higher labor demand, or by the migration of labor-intensive industries to areas with high labor supply? The contribution of this paper is that the natural experiment I consider does not restrict itself to ex-post statements, but directly demonstrates causality: Labor did not flow into regions in the West Bank with labor-intensive industries, but rather regions that received, for exogenous reasons, a larger influx of labor, shifted their production more towards labor-intensive industries, and exported the increased production of these goods. Michaels (2008) uses the differential effect of the creation of the US Interstate Highway System on different US rural counties. He finds that factor prices changed in a way that is consistent with the prediction of many-goods, two factors, two countries version of Heckscher-Ohlin. Relative to his work, not only am I able to test the predictions of the canonical HOV model, but the shock I study in this paper is much larger and more concentrated in time, thus allowing for better identification.

Another contribution of this paper has to do with the interpretation of the results. A well known issue with some of the empirical literature on the HOV model is that in the absence of a clear alternative theory, it is not obvious how to interpret the results of some of the tests. In particular, a positive correlation between the values of variables as predicted by the theory and the observed values of these variables may not be enough to lend support for the theory. As Davis, Weinstein, Bradford, and Shimpo (1997) explain: “Setting a null that there should be no correlation... could be rejected in most cases, but little comfort can be obtained by rejecting such an absurd proposition.” In this paper, there is an obvious alternative in the form of differential capital flows. In this alternative scenario, larger influxes of returning workers are matched by proportionally larger inflows of capital, thus restoring the original capital to labor ratio. Importantly, under this scenario, the sectoral composition of production, which is the key ingredient in the HOV model, plays no role at all in the absorption of the increase in labor, as each district will simply become a larger replica of its old self.

Since I do not have good data on capital flows, I test the success of the HOV predictions relative to this alternative scenario by creating a counterfactual West Bank, that shares many important features with the real West Bank, but in which sectoral changes in the composition of production do not occur, and therefore the HOV theory, by construction, has no explanatory power. I then use this counterfactual economy as a benchmark against which to evaluate the results of the tests of the HOV model in the real data, and find that systematic sectoral changes toward labor-intensive industries are necessary to account for the absorption of returning workers. Sectoral shifts matter.

This paper is also related to the work of Gandal, Hanson, and Slaughter (2004), who performed “absorption accounting” for the way the immigration wave from the USSR into Israel in the 90’s was absorbed into the labor market. They find that changes in Israel’s output mix did not play a role in absorbing the changes in the size and skill composition of the labor force in Israel. However, lacking a valid counterfactual, they caution against interpreting their findings as causal, and argue that they should only be understood as an accounting of the relative contribution of different elements (local technological change, output mix change, etc.) to the absorption of the new immigrants into the labor market. Hanson and Slaughter (1999) use a similar methodology to analyze output mix changes in the US. in response to immigration waves, and find the output mix changes “broadly match state endowment change”, and that relative FPE holds, which provides indirect support for the output mix change hypothesis. Relative to these studies, the West Bank experience after 2000 provides a much larger change in the size of labor markets, with a more plausibly exogenous source of variation. This allows for sharper tests of the theory, and accordingly, I find that output mix change not only broadly matches changes in labor supply, but it can explain, quantitatively, the absorption of the returning commuters.

This paper also contributes to the very large literature on the effects of immigration on labor markets (see Kerr and Kerr (2011) for a recent survey of the literature). While the historical context of this paper is different from that of most immigration waves, since most of the workers who commuted

into Israel before 2000 have been and remained residents of the West Bank, by highlighting changes to the output mix as a way to absorb new workers into the labor market, this paper can help explain why some studies, such as Card (1990), found immigrants to have little effect on the wages of natives.

Last, but not least, this paper can help in understanding labor markets and their relationship with trade in the West Bank, and the interaction of both with the political conflict—a topic of importance for policy. Several studies, such as Angrist (1995), Angrist (1996), Etkes (2011), Flaig, Siddig, Grethe, Luckmann, and McDonald (2013), and Mansour (2010) have looked at the relatively short-term effects of Israeli policies on Palestinian labor markets. This is the first paper to study the long-term effects of the decline in commuting into Israel, and the first to directly relate labor market conditions in the West Bank to trade.

The rest of the paper is organized as follows: Section 3 describes the historical episode that is used in the paper, and explains what makes it suitable for testing the HOV theory. Section 4 derives the testable predictions of the HOV model that will be tested using the data from the West Bank, and reports the results of these tests, and section 5 concludes. The data I used to compile the variables for this study, and the way they were compiled, are described in the data appendix.

3 Historical Background

3.1 The West Bank Labor Force

Palestinians began commuting from the West Bank into Israel for work almost immediately after the Israeli capture of the West Bank from Jordan in 1967 (See Angrist (1995)), and during the late 1990's about 20 percent of the labor force in the West Bank commuted to work in Israel on a daily basis, most of them using travel permits issued by Israel. (See Etkes (2011) for institutional details.) In September 2000, a wave of violent clashes between Palestinian

rioters and Israeli forces, later known as the Second Intifada³, erupted, and quickly escalated from mass demonstrations to a wave of suicide terrorism by Palestinians and counter-attacks by Israel. As a result, Israel's permits policy changed markedly, and security measures at the border increased, resulting in a very large and, at least so far, permanent drop in the number of commuters into Israel. Figure 1 shows the effects of this change on the Palestinian labor market in the West Bank. From a peak of over 22% in the first three quarters of 2000, the share of commuters out of the total employed persons in the West Bank dropped to less than 6% until 2007, and only inched back to around 8% since. The immediate effect of the substantial decline in the number of commuters was a drop in employment rates from around 75% in 1999 to below 60%. However, by 2007 (marked by a dashed line in the figure) employment rates mostly recovered, though to a somewhat lower level, and it seems reasonable to treat 2007 as the end of the recession in the West Bank. This study therefore focuses on the period between 1999 and 2007.⁴


A key fact about the commuting patterns in the West Bank was the variation between the 10 districts that comprise the West Bank. Table I reports the share of commuters to Israel out of total employed persons by district in 1999, which ranges from 12.8% for Nablus to 41.5% in Salfit. The reason for this variation is mostly geographic. Miaari, Zussman, and Zussman (2012) use Israeli administrative data on work permits issued to Palestinians to analyze commuting patterns at the locality level, and find that distance from Israel has a large and highly significant negative effect on the prevalence of commuting into Israel. The only other variable they report to have any effect is the type of locality, with villagers commuting slightly more than city dwellers.

This large variation in the commuting patterns in the years before the Second Intifada means that the immediate shock to the local labor force differed substantially between districts. However, for the purposes of this study it is not enough that the initial shock differed: It is necessary that the initial shock

³The first Intifada being the one which began in 1987.

⁴Another reason for using two odd years, 1999 and 2007, is to avoid problems related to the bi-annual seasonality of the olives industry, which is a non-negligible part of the West Bank economy.

Figure 1: Commuters to Israel and Employment Rate in the West Bank



Notes: Data are at a quarterly frequency and are taken from the Palestinian Central Bureau of Statistics Labor Force Survey. Commute to Israel is as reported by workers. Share of commuters (left axis) is the number of persons who live in the West Bank and report commuting to work in Israel out of total employed persons who live in the West Bank. Employment rate (right axis) is total employed persons divided by all persons of prime working age (25-55).

had a persistent effect on the size of the labor force in each district. Data from the Palestinian Migration Survey, conducted by the Palestinian Central Bureau of Statistics, suggest that this is indeed the case. Migration between districts in the West Bank was extremely limited: In 2010, 95.3% of the population in the West Bank resided in the same district in which they were born. The breakdown of these numbers by district is reported in Table II. In all but one district, Jericho⁵, the share of residents who were born in the district is between 92.9% and 97.0%. Moreover, only 7.6% of movers reported to have moved for work related reasons⁶, making the size of internal migration motivated by labor market considerations extremely small, and it is therefore likely that the initial shock caused by the Israeli policy change in 2000 had a

⁵The share of non-natives in Jericho, at 34%, is hard to explain, and might be a measurement error. One possible explanation is that Jericho was the first city in the West Bank to be handed over to the Palestinian Authority in 1994, and may have attracted some internal migration as a result. At any rate, given the very low share of movers that move for work reasons, even for Jericho the number of residents who moved there for work purposes is fairly small.

⁶The leading reason for migration is marriage to a person from another district.

Table I: Commuters to Israel in 1999, by District of Residence

	Share of Commuters out of Total Employment
Nablus	12.8%
Ramallah	13.7%
Bethlehem	17.5%
Jericho	21.5%
Tulkarm	23.5%
Tubas	25.1%
Hebron	26.1%
Qalqilya	28.4%
Jenin	31.2%
Salfit	41.5%

Notes: Data are taken from the Palestinian Central Bureau of Statistics Labor Force Survey. Commute to Israel is as reported by workers. Share of commuters is the number of persons who live in a district and report commuting to work in Israel out of total employed persons who live in the district.

persistent effect.

This is confirmed in Figure 2, which plots the share of the workers in each district that used to commute to Israel against the growth in the number of workers in each district following the decline in commuting. As can be seen, the growth in the number of employed persons in each district in the years following the return of the commuters (1999-2007) is very strongly correlated with the share of commuters in that district in 1999.

To test this correlation more formally, I regress the actual growth in employment in each district on the predicted growth rate of employment if the decline in commuting was the only source of variation. Technically, let a_i be the share of commuters out of total employment in 1999 in district i , and a'_i be this share for 2007. Also let g be the population growth rate in the West

Table II: Internal Migration in the West Bank

	% of District Residents, Born in the Same District
Jenin	97.0
Tulkarm	95.0
Nablus	94.9
Qalqilya	96.7
Salfit	95.6
Tubas	92.9
Ramallah	95.8
Jericho	66.1
Bethlehem	97.0
Hebron	96.1


Notes: Data are from the Palestinian Central Bureau of Statistics Migration Survey for the year 2010.

Bank as a whole⁷, L_i be the number of employed persons in district i in 1999, and L'_i their number in 2007. If the decline in commute was the only source of variation in the growth of employment, the growth rate of employment in district i would be $\left(\frac{1-a'_i}{1-a_i}\right)g$. I therefore run a no-constant regression of L'_i on $\left(\frac{1-a'_i}{1-a_i}\right)gL_i$. This regression yields a coefficient of 0.97 (not statistically different from 1 by any reasonable standard), and even more importantly, an $R^2 = 0.93$. Thus, the variation in commuting patterns predicts the variation in the growth rate of total employment very well, and indeed accounts for a very large part of this variation.

Nor did workers in the West Bank increase their commuting to other districts within the West Bank. In 1999, 6.1% of employed persons in the West Bank were commuting to work in a district where they did not live, and in

⁷Taken from the Population Statistics published by the PCBS.

Figure 2: Employment Growth 1999-2007, and Share of Commuters in 1999



Notes: Data are from the Palestinian Central Bureau of Statistics Labor Force Survey. Commute to Israel is as reported by workers.

2007 the number actually dropped somewhat, to 5.6%. These findings are consistent with previous studies which focused on the short-run effects of Israeli policy, such as Mansour (2010), who argues that the labor supply shocks following the decline in commuting to Israel were absorbed locally, because of Israeli restrictions on movements of people within the West Bank. While this low mobility may be surprising, considering the very high mobility that is demonstrated in the large share of commuters into Israel before 2000, it is important to remember that the wage premium for working in Israel was much higher than for commuting within the West Bank. (See for example Angrist (1995).).

At any rate, be the reason for low mobility within the West Bank what it may, all the evidence suggests that mobility between districts was very low, that the return of the Palestinian commuters to the West Bank had a persistent effect on the size of employment in the different districts, and that this effect explains much of the variation in the growth of employment across districts. It is the variation between districts in the size of the shock to their local labor force that is the key to my empirical analysis.

Were commuters similar to workers who did not commute? The West-Bank, by and large, is a low-skill economy with an average of 9.7 years of

education for workers, and less than 12% of workers have any post-secondary education. Thus, it seems reasonable to treat labor as a homogeneous factor. However, if there were important differences between commuters and non-commuters, it would suggest that the shock was not a single shock to aggregate labor supply, but a number of potentially different shocks to different kinds of labor. Indeed, in 1999, commuters were different from non-commuting workers along a few dimensions: They were a bit younger (average age of 31.6 of commuters, and 34.8 for non-commuters), more likely to be from a rural locality (59.1% of commuters were from rural localities, 42.3% of non-commuters), and essentially all men (98.1% of commuters, compared to 76% of non-commuters). Importantly, they also had, on average, one year less of formal education (10 years for non-commuters, 9.1 years for commuters), and a much lower share of commuters had post-secondary education (14.2% of non-commuters, 1.5% for commuters). These numbers suggest that commuters may have been less skilled than the average worker in the West Bank who did not commute. To gauge how important these differences were, I compare the actual mean wage of non-commuters in each district in the West Bank with the mean *predicted* wage of commuters from that district had they not commuted, based on their observable characteristics. Technically, I do so by regressing wages of non-commuting workers on age, age squared, sex, years of formal schooling, type of locality, and district dummy, and based on the coefficients obtained from this regression, I predict the wages of commuters in each district had they not commuted. Table III reports the results of this procedure. When using the full sample of workers (columns 1-3), the mean wage predicted for commuters, based on their observables, is 6.6% lower on average than that of non-commuters. This difference is also seen at the district level, and the difference is statistically significant at the 5% level in six of the nine districts. While statistically significant, the difference is not at all economically significant. This difference narrows further when workers with post-secondary education are removed from the sample (columns 5-7): Excluding workers with post secondary education, commuters' mean predicted wage throughout the West Bank is essentially identical to that of non-commuters,

with a difference of less than 1%. When the data are broken down by district, there is a small difference in favor of commuters, though it is only statistically significant in three districts, and economically meaningless. These considerations suggest that commuters were more similar to the less-educated workers in the West Bank than to the labor force as a whole. I therefore perform all the tests in this paper both for aggregate labor, and, separately, for low-skill workers. The results are equally supportive of the HOV predictions in both specifications.

Table III: Wages of Non-Commuters and Predicted Wages of Commuters in 1999

	Full Sample				Excluding Workers with College Degree			
	Actual Wage - Non-Commuters	Predicted Wage - Commuters	P-value for Difference>0	N	Actual Wage - Non-Commuters	Predicted Wage - Commuters	P-value for Difference≠0	N
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Whole West Bank	63.3	59.6	<0.01	53,489	58.7	58.2	0.33	48,019
Jenin	59.2	56.4	0.02	6,031	53.7	54.8	0.36	5,407
Tulkarm	56.4	54.2	0.12	3,684	49.8	51.8	0.20	3,221
Nablus	58.7	55.6	0.04	6,872	51.5	55.0	0.04	6,038
Qalqilya	55.6	52.8	0.05	3,187	47.9	52.3	0.01	2,830
Salfit/Tubas	57.5	54.1	<0.01	5,927	53.0	53.4	0.73	5,249
Ramallah	73.7	71.3	0.29	7,408	66.3	70.3	0.01	6,690
Jericho	58.2	52.7	<0.01	2,604	55.0	51.8	0.06	2,439
Bethlehem	74.1	71.4	0.18	6,297	68.2	70.5	0.15	5,659
Hebron	62.4	60.5	<0.01	11,479	59.6	59.5	0.98	10,441


Notes: Data are from the Palestinian Central Bureau of Statistics Labor Force Survey. Commute to Israel is as reported by workers. Predicted wage for commuters is calculated based on a wage regression using the sample of workers who live and work in the West Bank, and applying the coefficients from this regression to the commuters in each district.

3.2 Trade in Goods and Services in the West bank

While the movement of workers into Israel was severely restricted, the movement of goods was not, and throughout essentially all of the period in question the West Bank traded extensively. Figure 3 shows imports and exports for the West Bank for the period 1999-2009. The West Bank, like some other developing economies, was and is running a very large trade deficit, equal to about

60% of GDP, funded mostly with large transfers from donating foreign governments and NGO's, and remittances of Palestinians working abroad. The deep recession in the years immediately after the outbreak of the Second Intifada is evident in the trade data. In later years, while exports recovered, imports did not, declining from 82% of GDP in 1999 to around 72% of GDP towards the end of the decade. The larger decline in imports relative to exports is likely at least in part due to the decline in commuting, which funded some of the trade deficit. The importance, or lack thereof, of these changes in trade shares is discussed in detail in Section 4, but what is clear, and is key to the analysis, is that the West Bank was open to trade, and indeed traded extensively throughout the entire period in question.

Figure 3: Imports and Exports as Share of GDP in the West Bank



Notes: Data are from the Palestinian National Accounts. Imports on the left axis, and exports on the right axis.

4 The Predictions of the HOV Model for the West Bank

In its simplest form, the HOV model states that if all countries have access to the same constant returns to scale technology⁸, and have identical homothetic preferences, but have endowments with different ratios of factors of production,

⁸Technology here means a set of possible techniques. A technique is a specific mix of inputs per unit of output.

then they will tend to export those goods that use their abundant factors intensively. This is often stated in terms of the factor content of consumption, production, and trade. First, under free and costless trade, and assuming all countries produce all goods, and have access to the same technology, factor prices will be equalized. Then, if all countries consume all factors, embodied in goods, in the same proportions (due to identical homothetic preferences), but different countries are endowed with different factor proportions, the difference between the uniformity of consumption ratios and heterogeneity in endowment ratios is the factor content of trade.

In the context of the West Bank, the HOV model predicts that, comparing two districts, a larger increase in the size of the labor force will not be absorbed through slower wage growth and the associated implementation of more labor-intensive techniques, but by a larger shift in the composition of production toward more labor-intensive industries while using the same techniques. Moreover, due to identical homothetic preferences, the change in the composition of production will not be absorbed through increased consumption of labor-intensive goods in the district, but through changes in the factor content of trade between districts. Lastly, since the increased production of labor-intensive goods is “exported”, the prices of these goods do not fall (relative to their prices in other districts), and so the Stolper-Samuelson effect is avoided and wages do not decline (relative to wages in other districts). In this section, I show evidence that supports all of these propositions.⁹

4.1 The Composition of Production

I turn now to derive the precise predictions of the HOV model that can be tested using data from the West Bank before and after the decline in commuting. Throughout, I use district level data on gross output and employment for

⁹All of these predictions assume that districts are in the same “cone of diversification”. As an empirical matter, at the level of aggregation I am using, this is not a problem, since all districts had positive production in each of the 15 sectors I am considering in both 1999 and 2007.

15 sectors, which essentially comprise all of the West Bank’s output.¹⁰ Following standard notation, let X_i^t be the 15×1 vector of gross output in each sector in district i in year t , let B_i^t be the $M \times 15$ matrix of cost minimizing direct inputs needed to produce one unit of each good in district i in year t (M being the number of factors of production), and let V_i^t be the $M \times 1$ vector of factor endowments in district i in year t . The HOV theory of production predicts that all districts will use the same technique, that is, that $B_i^t = B^t \forall i, t$. Thus, the HOV model predicts the following relationship between technology, output, and the factor endowment in each district:


$$B^t X_i^t = V_i^t \tag{1}$$

This is not a trivial full employment condition. It states that full employment will be achieved in all districts by using the same production technique, ruling out the possibility that labor-abundant districts will employ more labor-intensive techniques to achieve full employment. As is the case with much of the empirical work on the HOV theory, testing this prediction with West Bank data, either from 1999 or from 2007, would lead us to reject the theory. Figure 4 presents the results of this simple cross-section test of the production side of the HOV theory using 1999 data and 2007 data. Each point in each graph is a district, with the value of the LHS of Equation 1 as its value on the vertical axis, and the value of the RHS as its value on the horizontal axis. If the theory described the data perfectly, all points would have been along the 45-degree line, which is clearly not the case. The cross-section data from the West Bank districts are not compatible with the predictions of the production side of the HOV theory.

Suppose now that districts in the West Bank deviate from this relationship, in a district specific way and in a time specific way. In this case, the HOV

¹⁰The 15 sectors correspond to the upper level structure of the ISIC Rev.4 classification system, minus two sectors: the financial sector, for which no data is available, and Mining and Quarrying, which is very small, and is present in only 2 districts.

Figure 4: A Cross-Section Test of Production in the HOV Theory



Notes: Each observation is a district. Units are 100 workers. Data on the supply difference are from the Labor Force Survey. Data on imputed demand difference are from establishment surveys and the agriculture census.

model of production can be written as:

$$B^t X_i^t = V_i^t + \epsilon_i + \mu_t \quad (2)$$

where ϵ_i and μ_t represent any consideration that may cause a district, or a time period, to systematically deviate from the the HOV model's prediction.¹¹ Presenting the theory this way also makes clear the advantages of using a natural experiment. Much of the work on expanding the HOV model to make it more compatible with the cross-section data focused on making explicit assumptions about the nature of the deviations. I, in contrast, address these potential deviations from the basic model by first taking the difference of Equation 2 between two districts, which yields:

$$B^t(X_i^t - X_j^t) = (V_i^t + \epsilon_i) - (V_j^t + \epsilon_j)$$

and then taking the difference between $t = 0$ and $t = 1$ which, after some

¹¹Time specific deviations can result, among other reasons, from the fact that the data used for the variables on the LHS and the RHS of Equation 2 come from different sources, or because of different levels of labor utilization in different years. It is true that given a time specific deviation it is possible to talk of a true matrix of factor requirements, \hat{B}_t , which can be derived based of the observed B_t matrix, and μ_t . But, since μ_t is not observed, it is not possible to use this \hat{B}_t matrix.

rearranging, yields:

$$B^1(X_i^1 - X_j^1) - B^0(X_i^0 - X_j^0) = (V_i^1 - V_i^0) - (V_j^1 - V_j^0) \quad (3)$$

which is the main prediction to be tested using the data on production from the West Bank. I call the RHS the *relative supply shock*, and the LHS the *imputed relative demand change*, where both demand and supply here refer to the factors market.¹² The RHS is simply the difference in the size of the shock to the factor endowment between district i and district j . The LHS is the difference between the two districts in the change to demand for factors that can be inferred from the observed change in output. Thus, Equation 3 will only hold in the data if districts absorbed larger influxes of labor by changing their output mix, and not if increases in labor endowment were absorbed by shifting to more labor-intensive techniques.¹³

This empirical strategy has a clear intuitive sense. Instead of assuming that districts share the same technologies, or explicitly modeling the differences between them, the natural experiment allows me to “difference out” these differences. I compare a district from before the shock to its factor endowment and after it, in each case using another district to control for time variations that are common to all districts in the West Bank. Since the test is defined over pairs of districts, using 9 districts¹⁴ yields 36 predictions for every factor of production (one for each unordered pair of districts).¹⁵ The results of these tests are encouraging for the theory.

Figure 5 plots the results of the tests for aggregate labor when using all

¹²Davis, Weinstein, Bradford, and Shimpo (1997) call the equivalent of the RHS in their cross-sectional test the *observed* endowment, and the LHS the *predicted* endowment, i.e. the endowment that is predicted by the HOV model, given the observed patterns of production.

¹³It is also possible to use a multiplicative error term in Equation 2, so that the full employment condition is: $B^t X_i^t = V_i^t \times \epsilon_i \times \mu_t$, and re-writing Equation 3 as “ratio-of-ratios” rather than a “difference-in-difference”: $\frac{B^1 X_i^1 / B^0 X_i^0}{B^1 X_j^1 / B^0 X_j^0} = \frac{V_i^1 / V_i^0}{V_j^1 / V_j^0}$. The results are very similar to the results reported here.

¹⁴I combine Tubas and Salfit, the two smallest districts, into one artificial district. Not doing so can raise sample size problems.

¹⁵Though only 8 of them are completely independent of each other, it is important to report all results, to make sure errors tend to cancel out, and not compound.

unordered pairs, and Figure 6 plots the results using Nablus as the base against which all districts are compared.¹⁶ Similar to Figure 4, each point in these graphs is a pair of districts, with the relative supply shocks, $(V_i^1 - V_i^0) - (V_j^1 - V_j^0)$, on the horizontal axis, and the imputed relative demand change, $B^1(X_i^1 - X_j^1) - B^0(X_i^0 - X_j^0)$, on the vertical axis. If the predictions of the HOV model held perfectly in the data, all points would have been on the 45-degree line. While not quite equal in all cases, the correlation between the LHS and RHS of Equation 3 is striking ($\rho = 0.87$ for all unordered pairs, and $\rho = 0.95$ using only Nablus comparisons). However, a high correlation between the LHS and the RHS of Equation 3 is a necessary, but not a sufficient, condition for the results to support the HOV prediction. For example, if all points in Figures 5 and 6 were located exactly on a straight line from the origin, but one with a slope very different from unity, the correlation between LHS and the RHS of Equation 3 would be $\rho = 1$, but in such a case the prediction of equality between the two sides of Equation 3 nonetheless fails. In fact, this would be similar to the results in the simple cross-section test above.


A more suitable way of evaluating how close the observations are to their theoretical prediction (i.e. to the 45-degree line) is a regression analysis. If the theory described the data perfectly, a regression line of the LHS on the RHS of Equation 3 using West Bank data would result in a coefficient of 1, and an R^2 of 100%. In practice, a no-constant regression of the LHS on the RHS of Equation 3 yields a coefficient of 0.96 and an R^2 of 91% for the sample of all unordered pairs, and a coefficient of 0.97 and an R^2 of 95% when Nablus is used as the base for comparisons.

There is clearly more noise at the bottom left corner of Figure 5, with a few points far off the diagonal. This could be as a result of measurement error, which, if it has a fixed component, will have a relatively larger effect on points with lower values.¹⁷ At any rate, even when the sample is restricted to include

¹⁶Results are all in absolute values. Not using the absolute values of the RHS and LHS makes the test look better than when using absolute values. However, it is clear that the only difference that has economic meaning is the absolute value of the difference between districts, not which district is being subtracted from which.

¹⁷It should be kept in mind that the West Bank is a developing economy, with a GDPPC

Figure 5: Actual and Imputed Labor Differences - All Unordered Pairs




Notes: Each observation is a pair of districts, and the results are in absolute values. Units are 100 workers. Data on the supply difference are from the Labor Force Survey. Data on imputed demand difference are from establishment surveys and the agriculture census.

only the bottom left corner, those observations with the value of either the horizontal or the vertical axis being below 100, the results, while not as good, are still fairly consistent with the theory: the coefficient of the regression is 0.89, not statistically different from 1 at 5% confidence level, and the R^2 is 79%. For the rest of the sample (i.e. for all points that have both vertical and horizontal values above 100), the coefficient of the regression is 0.99 and the R^2 is 98%.

Similar, though noisier, results are obtained for workers with high school education or less. Figure 7 reports these results. The correlation between the LHS and RHS of Equation 3 when applied to non-college educated workers is $\rho = 0.86$, and a no-constant regression line yields a coefficient of 0.96 with an R^2 of 88% for all unordered pairs (left panel). When using Nablus as a base for comparison the correlation is $\rho = 0.90$ and the regression line yields a coefficient of 0.95 and an R^2 of 92%.

of 1,731 USD in 1999 and 1,901 USD in 2007, so some measurement error is to be expected.

Figure 6: Actual and Imputed Labor Differences (All Districts v. Nablus)




Notes: Each observation is a pair of districts with Nablus being one of the two. All results are in absolute values. Units are 100 workers. Data on the supply difference are from the Labor Force Survey. Data on imputed demand difference are from establishment surveys and the agriculture census.

4.2 Comparing HOV Predictions to a Simple Neo-classical Growth Model Predictions

While encouraging, the question remains just how substantially these results support the theory. As long as marginal labor productivity is positive, we should expect to find a positive correlation between the supply shock and imputed demand shock. For example, there is no meaningful null hypothesis of zero correlation to reject.¹⁸ I address this concern in two ways. First, it is instructive to compare these correlations to the ones found in studies that use data on a cross-section of countries. Excluding the US, a very large outlier, from the sample, Davis, Weinstein, Bradford, and Shimpo (1997) find a correlation of $\rho = 0.29$ for non-college graduates and a correlation of $\rho = 0.271$ for college graduates. For capital, they consider $\rho = 0.628$ to be a relative success of the model. Moreover, the relative success of the theory’s prediction for capital relative to labor is the case in general, as Trefler (2002) explains: “It is by now well known that the Heckscher-Ohlin-Vanek model performs

¹⁸And so there is no point in reporting the statistical significance of the correlation, which is a hypothesis test against this null.

Figure 7: Actual and Imputed Labor Differences - Low Skill Workers



Notes: Each observation is a pair of districts, and the results are in absolute values. Units are 100 workers. Data on the supply difference are from the Labor Force Survey. Data on imputed demand difference are from establishment surveys and the agriculture census.

reasonably well for natural resources, less well for capital, and poorly for labor.” Thus, the high correlation found here for labor is notable.

Another approach to evaluating the results is to ask how the results of these tests would look if sectoral composition did not play *any* role in the absorption of labor in the West Bank. This is not just a statistical exercise, but also a test of a competing theory. In a simple neo-classical growth model, an increase in the size of the labor force increases the return to capital, and will lead to differential capital flows into districts, which will exactly match the differential inflows of labor. In this case, each district just becomes a larger replica of its old self, and will absorb all the returning commuters without a need to change sectoral composition. In other words, it is at least possible that the absorption of labor was made possible simply through capital accumulation and growth, without any change to sectoral composition or to production technique. Importantly, if that is what happened, the tests that I described above, as well as the wage tests I describe below, would yield perfect results.¹⁹ Since I do not have good data on capital flows, I cannot directly rule out this competing theory.

To address this concern, I proceed in two steps: First, I create a counterfactual West Bank, in which the relative size of the sectors, as measured by

¹⁹Sectoral composition in the different districts did change between 1999 and 2007. However, it is important to determine if these changes were random or did they systematically contribute to the absorption of labor.


output²⁰, remains constant in each district after the shock, but overall output growth is the same as it was in the real data. Thus, in the counterfactual economy, all sectors grow at the same rate, which is the growth rate of total output in the district in the real data. Having created this counterfactual economy, I then perform the same tests on it as I did with the actual data. The difference between the results of the tests when using the actual data and the results when using the counterfactual data is a measure of the importance of sectoral shifts, which are the key difference between the HOV predictions and the neo-classical growth model predictions.

Figures 8 and 9 show these results side by side with the results of the tests on the real data, both for using Nablus as the base of comparison, and for all unordered pairs. The results suggest that sectoral composition did play a role: In the counterfactual West Bank, most of the points lie well below the 45-degree line. The way to interpret this is as follows: If the various districts in the West Bank did not systematically change their sectoral composition after the shock to their labor force (i.e. if the HOV model is wrong), then a test that assumes that they also kept using the same techniques as before (i.e. a test that assumes the HOV model is right) will fit the data poorly. Moreover, the result is what seems like a substantial under-demand for labor in the districts that received large supply shocks: A regression of the counterfactual imputed demand change on the observed supply shock yields a slope of only 0.60 when all differences are taken against Nablus, and of 0.66 for all unordered pairs. The under-demand for labor implies that greater overall growth in total production in districts that received larger labor inflows, by itself, cannot account for the absorption of the labor supply shock, and is not, by itself, the reason for the confirmation of the HOV predictions. It is only when the sectoral composition of production is also taken into account that the tests of the HOV predictions are successful.²¹

²⁰Using labor instead of output as a measure for the size of sectors yield very similar results


²¹One intuitive way to understand this result is through the concept of *comvariance*, developed in Deardorff (1982) in the context of the HO model. A comvariance is a generalization of covariance to the case of three variables. Technically, it is defined as $com(xyz) = \sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})(z_i - \bar{z})$. In the West Bank case the relevant comvariance

Figure 8: Tests of Equation 3 with counterfactual and Real data, All Un-ordered Pairs



Notes: The left panel shows the results of testing the prediction of the HOV production model in the real data, and the right panel shows the results of testing the model with counterfactual data, where the sectoral composition of production is held constant. All results are in absolute values. Units are 100 workers.

Figure 9: Tests of Equation 3 with counterfactual and Real data, Nablus as Base



Notes: The left panel shows the results of testing the prediction of the HOV production model in the real data, and the right panel shows the results of testing the model with counterfactual data, where the sectoral composition of production is held constant. In both Nablus is used as the base against which differences are taken. All results are in absolute values. Units are 100 workers.

is between the labor intensity of an industry, the size of the increase in a district's labor endowment, and the percentage increase in the production in an industry. Since there is a positive comvariance between these three variables in the data, we can make the following comvariance statement: As predicted by the HOV theory, districts that received larger influx of returning workers increased production relatively more, in industries that are relatively more labor-intensive.

Thus, the HOV model fits the data better than a simple neo-classical growth model: Differential capital flows by themselves cannot explain the absorption of labor without the change in sectoral composition of production.

The next section deals directly with the possibility that larger labor supply shocks were absorbed, at least partially, through technique changes, and finds evidence against it.


4.3 Wages

A well-known prediction of the HOV model is the so called “Factor Price Equalization” (FPE) theorem. In the HOV model, output prices fully determine the price of factors. Under free and costless trade, all countries face the same output prices, and so, if they share identical technologies, they will have the same factor prices, regardless of relative abundance of factors in each country. This prediction is clearly rejected by international data, as factor prices obviously differ greatly between countries. It is even rejected in the cross section of the West Bank data, with wages differing across districts. However, the natural experiment in the West Bank allows me to test a weaker version of this result, known as Factor Price Insensitivity (FPI), which is difficult to test using cross-section data. According to FPI, factor prices need not be identical everywhere, but they are nonetheless independent of relative factor endowment. In the context of the natural experiment I study, the prediction of the FPI theorem is that wage changes in a district will not be correlated with the size of the shock to the labor force in that district. An important feature of that prediction is that it has an obvious alternative against which it is tested: That districts that received larger labor supply shock will absorb this excess labor through slower wage growth.

Since lower wages are associated with a shift towards more labor-intensive techniques, this test is complementary to the tests described in the previous section. There, it was established that sectoral shifts had to have played a role in labor absorption. Here it is established that technique changes did not.

Figure 10 shows the percentage growth of average nominal wages in each district between 1999 and 2007,²² and the percentage change in the size of the labor force in the district during the same years. The left panel is for aggregate labor, and the right panel is for non-college educated workers only. The correlation between wage growth and the shock to labor supply is a well-measured zero, and wage increases are not different from each other in a statistically significant way: A district that received a larger influx of labor did not experience slower wage growth. This is consistent with the results in the previous section. Since, as was shown there, the relative increase in labor demand that resulted from changes in the sectoral composition of production matched the increase in labor supply, there was no need for relative wages to fall. It is therefore not surprising that wage changes are not correlated with the size of the shock to the labor force in a district, and indeed are generally similar between districts.

Figure 10: Wage Increases and Growth of Labor Force by District



Notes: Data are from the Labor Force Surveys for the years 1999 and 2007. Black lines are the average of the districts' average wage increases, and dashed lines are the 1.96 standard deviation band.

One concern about this result is positive selection into commuting. If commuters into Israel were, on average, more skilled, then it is possible that after their return, wages went down within every skill level, but average wage in the district was unaffected due to the upgrading of its skill composition. Since the relevant wage for the FPI result is within skill level, such a scenario, if true, would contradict the prediction of FPI. However, there is no evidence that skill

²²The numbers are the total growth, not annual rate of growth.

composition upgrading had an important effect on wages. First, as can be seen in the right panel of Figure 10, wage changes were not correlated with the size of the shock to the labor force even when the sample is restricted to low-skill workers. Second, as was shown in Table III, commuters had lower levels of education, were younger, and their predicted wage, had they not commuted to Israel, was generally similar, or if anything—lower than the wages of stayers. Thus, it is unlikely that positive selection into commuting is driving the results.

4.4 Consumption

So far, I have examined the production side of the HOV model of trade. In this section, I supply indirect evidence for the validity of the HOV theorem on the factor content of trade, which states that countries will be net exporters of the services of factors they have in relative abundance. The direct way to test this prediction in the West Bank would have been to use data on trade at the district level. Unfortunately, such data do not exist. However, since the factor content prediction of the theory is derived by simply taking an arithmetic difference between the results of the HOV theory of production and its results about consumption, I can supply indirect evidence for its validity by testing the validity of the assumptions about consumption in the model. To be sure, having data on trade at the district level would have contributed to this study a great deal, but while the data I do have are limited, they nonetheless give important support to the HOV theorem about the factor content of trade.

Before turning to the formal tests, it is important to note that while the tests of the production and consumption sides of the model are technically separate, they are conceptually related through general equilibrium effects. Even if the main margin of adjustment of production was sectoral changes, unless the increased production of labor-intensive goods was traded away, the price of these goods would have fallen. This, in turn, would have triggered a Stolper-Samuelson effect, which would have led to a decline in wages. Thus, the fact that wages did not fall (or grow more slowly) in districts that received a larger influx of workers, in spite of the fact that they shifted production

more towards labor-intensive industries, already suggests that this increase in production of labor-intensive goods was traded away, and not consumed locally.

Formally, the HOV theorem about the factor content of trade is stated as follows. Denoting country i 's share of consumption of each good out of world consumption of that good by the 15×1 vector s_i , identical homothetic preferences (IHP) imply that all entries in that vector are equal to the share of economy i 's GDP out of world GDP, and it is therefore possible to treat s_i as a scalar.²³ Denoting the factor content of trade for country i by the $1 \times M$ vector F_i (with M being the number of factors), world GDP by Y^w , a district GDP by Y_i , the direct plus indirect factor content of goods with the $15 \times M$ matrix A ²⁴, and keeping all the notation for production from Equation 2, we can state the HOV theorem:

$$AY_i - s_i AY^w = V_i - s_i V^w = F_i \quad (4)$$

In words: The factor content of production, minus the factor content of consumption, equals the factor content of trade. Defining an economy as abundant in factor m if its share of world endowment of it is larger than its share of world GDP, s_i , Equation 4 states that an economy will be a net exporter of the services of the factor in which it is abundant. Since trade data are unavailable, the only part of Equation 4 that can be tested against data, other than the production side, is the assumption of identical homothetic preferences, which gives rise to the existence of the scalar s_i , representing both country i 's share of world GDP and its share of world consumption of each good.

To test whether this assumption is borne out by the data, I use data from

²³This is for the case of balanced economies. If economies are not balanced, their share of world GDP differs from their share of world consumption, which requires adjusting Equation 4 to reflect that imbalance. It makes no difference for the analysis here. Nor does the fact that the West Bank, or any of its districts, was not balanced before or after the shock.

²⁴The matrix A can be stated using the matrix B and the input output matrix of the economy. However, it makes no difference for the analysis here, and unnecessarily complicates the notation. The key thing to note is that $AY_i = BX_i = V_i$.

the Palestinian Household Expenditure Survey, which details the value of expenditure on a few hundred goods and services, grouped into 26 categories, for a representative sample of households in the West Bank. The assumption of IHP translates into the testable prediction of identical expenditure shares for all categories in all districts, independent of income. Unfortunately, the data do not report the district of residence of each household, but only divides households to “Northern West Bank”, “Middle West Bank” and “South West Bank”. Moreover, these data too are only available starting at 2009, so only for the time after the decline in commuting to Israel. Nonetheless, they are still informative. There are only two ways in which the data may yield misleading results for our purposes. The first is if before 2000, expenditure shares were not consistent with IHP, but after the decline in commuting, changes to production patterns led to changes in consumption patterns that are, coincidentally, consistent with IHP. This seems extremely unlikely. Second, it is possible that differences in expenditure shares between districts average out at this higher level of aggregation. While increasing the level of aggregation may mechanically lead to some averaging out of differences, since there are only 3 or 4 districts in each of North, Middle, and South West Bank, it seems unlikely that such differences in expenditure shares balanced each other very effectively. It is much more plausible to interpret the data as supporting the assumptions that households in different districts in the West Bank share similar, homothetic, preferences.

Table IV reports average consumption shares on each of 24 groups of goods, that together cover essentially all of household consumption, in each of the three regions of the West Bank.²⁵

²⁵I exclude the value of home-produced goods, which are likely to be poorly measured and impossible to price, and the value of consumption of alcoholic beverages, since in a predominantly Muslim society there may be severe under-reporting on actual consumption of alcohol. Indeed, the overwhelming majority of households report zero alcohol consumption.

Table IV: Consumption Shares in the West Bank

	Northern West Bank	Middle West Bank	Southern West Bank
Average Total Expenditure (NIS)	4,647	5,032	4,201
Bread and Cereals	5.6%	4.4%	5.8%
Meat and Poultry	7.2%	8.1%	11.4%
Fish and Sea Products	0.5%	0.6%	0.7%
Dairy Products and Eggs	2.7%	2.9%	3.1%
Oils and Fats	1.4%	1.1%	1.2%
Fruits and Nuts	2.7%	2.2%	2.7%
Vegetables, legumes and tubers	3.7%	3.3%	4.8%
Sugar and confectionery	1.7%	1.7%	1.5%
Non-alcoholic beverages	1.6%	1.7%	1.6%
Salt, spices and other food	2.0%	1.7%	1.3%
Take-away food and meals in restaurants	3.6%	2.2%	1.8%
TOTAL FOOD	32.9%	29.7%	36.0%
Clothing and footwear	6.6%	5.6%	6.5%
Housing	6.9%	9.4%	8.4%
Furniture and utensils	4.7%	3.7%	2.9%
Household operations	1.3%	1.2%	1.1%
Medical care	4.5%	2.9%	3.7%
Transport	11.6%	13.0%	12.1%
Education	2.8%	2.9%	2.9%
Recreation	1.3%	1.8%	1.0%
Personal care	2.5%	2.7%	2.2%
Tobacco	4.5%	5.2%	5.7%
Other non-food consumption expenditure	4.8%	3.0%	3.1%
Other than food	0.4%	0.0%	0.0%
Imputed rent	8.6%	12.9%	9.8%
Communication	3.4%	4.2%	2.6%
N	1,189	988	725

Notes: Data are from The Household Expenditure Survey for 2009, conducted by the PCBS.

The resemblance in expenditure shares is remarkable. The correlation in the size of expenditure shares between the areas of the West Bank is given in Table V and is between 0.941 and 0.944. These high correlations are not the result of some categories being inherently much larger than others: The Spearman rank correlations, also reported in Table V, are also very high, ranging from 0.949 to 0.957. This is especially notable, since the value of total expenditure (and so, presumably, income level) differs by 10%-20% between the different parts of the West Bank. Considering the inherent noise in consumption reporting, these results are consistent with the IHP assumption, at least at this level of aggregation.

It is important to understand these results in the context of the results from the previous section. It is because the different parts of the West Bank

Table V: Correlations in Expenditure Shares

Pearson Correlation (ρ_p) and Spearman Rank Correlation (ρ_s)			
	Northern West Bank	Middle West Bank	Southern West Bank
Northern West Bank	1		
Middle West Bank	$\rho_p=0.941$ $\rho_s=0.957$	1	
Southern West Bank	$\rho_p=0.945$ $\rho_s=0.953$	$\rho_p=0.944$ $\rho_s=0.949$	1

shared identical homothetic preferences that increases in the production of labor-intensive goods were traded away and did not result in a relative fall in the prices of labor-intensive goods in districts that received larger labor influxes, which in turn would have led to a relative decline in wages there.

5 Conclusions

The exogenous fall in commuting options from the West Bank to Israel led to a large shock to the size of the labor force in the West Bank, an effect that varied between different districts there. Allowing for district specific deviations from the assumptions of the HOV model about production, I find that the changes in production patterns are consistent with the predictions of the Heckscher-Ohlin-Vanek model of trade. Moreover, these changes allowed for an absorption of the returning commuters, without relative wage changes, as predicted by the FPI theorem. While no direct data on trade exist at the district level, data on consumption is consistent with the assumption of identical homothetic preferences in the West Bank, thus giving indirect support for the HOV theorem on the factor content of trade. Yet, all three of these predictions failed in cross sectional tests conducted in previous studies. What are we then to make of these new findings?

It should come as no surprise that the HOV theory is not a sufficient theory to explain the actual trade patterns we observe in the real world. Germany and China are different in many ways other than the relative supply of different

kinds of labor and capital, and some of these differences affect trade patterns in ways that the HOV theory fails to capture. However, that does not render the theory useless. While the original insights of Eli Heckscher and Bertil Ohlin are not sufficient to explain world trade patterns, the experience of the West Bank suggests that they nonetheless describe a real and important aspect of international trade, and do very well in explaining how an open economy adjusts to shocks to its factor markets by modifying the composition of production, and trading away the excess production. Whatever other forces shape trade patterns, relative factor abundance does play a role.

References

- Angrist, J. D. (1995): “The Economic Returns to Schooling in the West Bank and Gaza Strip,” *The American Economic Review*, 85, 1065–1087.
- (1996): “Short-Run Demand for Palestinian Labor,” *Journal of Labor Economics*, 14, 425–453.
- Bernhofen, D. M., and J. C. Brown (2011): “Testing the General Validity of the Heckscher-Ohlin Theorem: The Natural Experiment of Japan,” *CESifo Working Paper Series*, No. 3586.
- Bowen, H. P., E. E. Leamer, and L. Sveikauskas (1987): “Multicountry, Multifactor Tests of the Factor Abundance Theory,” *American Economic Review*, 77, 791–809.
- Card, D. (1990): “The Impact of the Mariel Boatlift on the Miami Labor Market,” *Industrial and Labor Relations Review*, 43, 245–257.
- Davis, D. R., and D. E. Weinstein (2001): “An Account of Global Factor Trade,” *The American Economic Review*, 91, 1423–1453.
- Davis, D. R., D. E. Weinstein, S. C. Bradford, and K. Shimpo (1997): “Using International and Japanese Regional Data to Determine When the Factor

- Abundance Theory of Trade Works,” *American Economic Review*, 87, 421–446.
- Deardorff, A. V. (1982): “the General Validity of the Heckscher-Ohlin Theorem,” *The American Economic Review*, 72, 683–694.
- Etkes, H. (2011): “The Impact of Employment in Israel on the Palestinian Labor Force,” Bank of Israel Discussion Paper No. 2011.11.
- Flaig, D., K. Siddig, H. Grethe, J. Luckmann, and S. McDonald (2013): “Relaxing Israeli restrictions on Palestinian labour: Who benefits?,” *Economic Modelling*, 31, 143–150.
- Gandal, N., G. H. Hanson, and M. J. Slaughter (2004): “Technology, Trade, and Adjustment to Immigration in Israel,” *European Economic Review*, 48, 403–428.
- Hanson, G. H., and M. J. Slaughter (1999): “The Rybczynski Theorem, Factor-Price Equalization, and Immigration: Evidence from U.S. States,” *NBER Working Paper No. 7074*.
- Harrigan, J. (1995): “Factor Endowments and the International Location of Production: Econometric Evidence for the OECD, 1970–1985,” *Journal of International Economics*, 39, 123–141.
- Horiba, Y., and R. C. Kirkpatrick (1981): “Factor Endowments, Factor Proportions, and the Allocative Efficiency of U.S. Interregional Trade,” *Review of Economics and Statistics*, 63(2), 178–187.
- Kerr, S. P., and W. R. Kerr (2011): “Economic Impact of Immigration: A Survey,” NBER Working Paper 16763.
- Leontief, W. (1953): “Domestic Production and Foreign Trade: The American Capital Position Re-examined,” *Proceedings of the American Philosophical Society*, 97, 331–349.

- Mansour, H. (2010): “The Effects of Labor Supply shocks on Labor Market Outcomes: Evidence from the Israel-Palestinian Conflict,” *Labour Economics*, 17, 930–939.
- Maskus, K. E. (1985): “A Test of The Heckscher-Ohlin-Vanek Theorem: The Leontief Commonplace,” *Journal of International Economics*, 19, 201–212.
- Miaari, S., A. Zussman, and N. Zussman (2012): “Employment Restrictions and Political Violence in the Israeli-Palestinian Conflict,” *Discussion Paper Series, Bank of Israel*.
- Michaels, G. (2008): “The Effect of Trade on the Demand for Skill—Evidence from the Interstate Highway System,” *Review of Economics and Statistics*, 90(4), 683–701.
- Stern, R. M., and K. E. Maskus (1981): “Determinants of the Structure of US Foreign Trade, 1958–76,” *Journal of International Economics*, 11, 207–224.
- Trefler, D. (1993): “International Factor Price Differences: Leontief was Right!,” *Journal of Political Economy*, 101, 961–987.
- (1995): “The Case of the Missing Trade and Other Mysteries,” *The American Economic Review*, 85, 1029–1046.
- (2002): “The Case of the Missing Trade and Other Mysteries: Reply,” *The American Economic Review*, 92, 405–410.

A Data Appendix

Data on wages, employment and commuting are from the Palestinian Labor Force Survey, conducted by the Palestinian Central Bureau of Statistics. The survey is conducted at a quarterly frequency, and is explicitly designed to be geographically representative, based on the population census. In the year 1999, a total of 53,489 persons over the age of 15 were sampled in the West Bank,

and in the year 2007 the number was 52,568. Data on agricultural output are taken from the agricultural censuses of the Palestinian Authority for 1999 and 2007. Data on output of the rest of the sectors come from the firm censuses of 1997 and 2007, and the surveys of manufacturing, services, construction, transportation, and retail trade. All surveys are explicitly designed to be geographically representative. To calculate the unit labor requirements in each industry I divide total gross output in the industry in a year by the number of workers in the industry in that year. Data on consumption are from the Household Expenditure Survey of 2009, conducted by the Palestinian Central Bureau of Statistics, which is explicitly designed to be geographically representative. In 2009, it sampled 2,909 households in the West Bank, detailing expenditure level on 659 items, classified into 26 main groups.