TRADE POLICY, INCOME RISK, AND WELFARE

Tom Krebs, Pravin Krishna, and William Maloney*

Abstract—This paper develops a framework to study empirically the relationship between trade policy and individual income risk and to evaluate the associated welfare consequences. The analysis proceeds in three steps. First, longitudinal data on workers are used to estimate time-varying individual income risk parameters in various manufacturing sectors. Second, the estimated income risk parameters and data on trade barriers are used to analyze the relationship between trade policy and income risk. Finally, a simple dynamic incomplete-market model is used to assess the corresponding welfare costs. In the implementation of this methodology using Mexican data, we find that trade policy changes have a significant short-run effect on income risk. Further, while the tariff level has an insignificant mean effect, it nevertheless changes the degree to which macroeconomic shocks affect income risk.

I. Introduction

RECENT years have seen an increased integration of countries into the world economy through trade and capital market liberalization. This has led to a parallel surge of interest in the academic and policy literature on the implications of increased openness of countries to cross-border trade in goods and factors. The economic benefits and costs of openness are now being actively debated. While many economists have pointed to the gain in allocational efficiency that results from free international exchange, others have pointed out potential downsides, arguing that openness may lead to an increase in income inequality and, separately, income risk (income volatility). Although there is by now a large empirical literature analyzing the impact of trade openness on wage levels and the distribution of income,[1] an empirical analysis of the effect of trade openness on individual income volatility has so far been lacking.[2] This paper conducts such an empirical investigation and uses the empirical results in conjunction with a simple dynamic general equilibrium model to assess the corresponding welfare effects.

The theoretical literature has suggested various channels through which trade reform might affect individual income risk. For example, lowering trade barriers leads to an increase in foreign competition in the import-competing sectors and is likely to induce a reallocation of capital and labor across firms and sectors. In the short run, the resulting turbulence may raise individual labor income risk.[3] Rodrik (1997), going beyond the short-term reallocational effects of trade reform on income risk, has additionally argued that increased foreign competition following trade reform will increase the elasticity of the goods and the derived labor demand functions. If a higher demand elasticity translates any given shock into larger variations in wages and employment, lower trade barriers may lead to increased individual income risk.[4] It has also been suggested that the world economy is likely to be less volatile than the economy of any single country, which leads to goods prices that are more stable worldwide than in any single autarkic economy. This opens up the possibility that greater openness may reduce the variance in individual incomes. Thus, theoretically, the openness-volatility relationship is ambiguous. The theoretical literature does not offer a strong prior on the sign or magnitude of this relationship.[5]

In this paper, we study empirically the effects of trade policy on individual income risk using the following approach. For each industry (sector), we use longitudinal data on individual earnings to estimate time-varying parameters of individual income risk using a methodology that follows the approach taken by the extensive empirical literature on labor market risk. More specifically, we focus on the variance of (unpredictable) changes of individual income as a measure of income risk and carefully distinguish between transitory and persistent income shocks. The distinction between transitory and persistent income shock is important since workers can effectively "self-insure" against transitory shocks through borrowing or their own savings, which implies that the effects of these types of shocks on workers' consumption and welfare are quite small (Levine & Zame, 2002). In contrast, highly persistent or permanent income shocks have a substantial effect on the present value of future earnings and therefore lead to significant changes in consumption even if workers can borrow or have their own savings (Constantinides & Duffie, 1996; Krebs, 2003,

Received for publication February 14, 2007. Revision accepted for publication October 8, 2008.

* Krebs: Mannheim University; Krishna: Johns Hopkins University and NBER; Maloney: World Bank.

We thank seminar participants at Brown, Columbia, Georgetown, Harvard, International Monetary Fund, Johns Hopkins University, New York University, Princeton University of Pennsylvania, Rutgers, Southern Methodist University, Syracuse University, Texas A&M, World Bank Research Department, the 2004 Summer Meetings of the Econometric Society, and the 2005 meetings of the Society for Economic Dynamics for many helpful comments and suggestions and Jungjin Lee for outstanding research assistance. We also gratefully acknowledge support from the Research Committee and the Office of the Chief Economist for Latin America and the Caribbean at the World Bank.

[1] See, for instance, Lawrence and Slaughter (1993) and Feenstra and Hanson (1996). Feenstra and Hanson (2002) and Goldberg and Pavcnik (2007) provide comprehensive surveys.

[2] On the issue of openness and output volatility, however, see Di Giovanni and Levchenko (2007).

[3] See, for instance, the analysis of policy change by Fernandez and Rodrik (1991), in which ex ante identical workers experience ex post different outcomes since some workers retain their jobs while others are forced to move to other firms. More recently, Melitz (2003) has developed a formal framework in which trade policy changes affecting an entire sector lead to heterogeneous outcomes at the firm level with related empirical evidence provided by Pavcnik (2002).

[4] While Rodrik (1997, 1998) appears to have in mind mostly aggregate volatility, it is easy to see that his arguments equally apply to individual income volatility if there are idiosyncratic shocks to firm-level productivity. See also Traca (2005), who analyzes theoretically the links among globalization, wage dispersion, and volatility.

[5] Clearly, this sign ambiguity does not extend to the short-term reallocational effect of trade policy reforms, which are generally expected to raise income risk. However, we do not have strong priors on the magnitude of this relationship either.

2004). Thus, from a welfare point of view, persistent income shocks matter the most, and we therefore focus on the relationship between trade policy and the persistent component of income risk.[6] More specifically, after obtaining the estimates of the persistent component of income risk for each year and industry, we use these estimates in conjunction with tariff data (as a proxy for trade policy) to study empirically the effect of trade policy on income risk.

In addition to the empirical analysis of the relationship between trade policy and income risk, this paper also provides a quantitative evaluation of the welfare consequences of any changes in income risk brought about by changes in trade policy. If insurance markets and other institutional arrangements for sharing individual income risk are missing (incomplete markets), then changes in income risk will alter consumption volatility and therefore workers' welfare. To find out how income risk is linked to consumption volatility and welfare, we use a dynamic general equilibrium model with incomplete markets in which the consumption-versus-saving choice of workers in the presence of idiosyncratic income risk is explicitly modeled. As is well known, general versions of such models are difficult to solve, and most work in the literature has therefore been computationally intensive. In contrast, in this paper, we rely on an extended version of the incomplete-markets model recently developed and analyzed by Constantinides and Duffie (1996) and Krebs (2004) that is highly tractable but still rich enough to allow a tight link between the econometric framework and the theoretical model. The welfare expressions that we derive theoretically can then be used to translate changes in individual income risk into welfare changes.

Our previous discussion highlights the need for longitudinal information on incomes at a disaggregated level (individual or household)[7] in countries that have undergone discernible (and ideally substantial) changes in their external regime. Unfortunately, countries that maintain detailed longitudinal records on individual incomes have rarely undertaken major trade reforms, and countries that have undertaken extensive trade policy reforms have rarely collected data on individuals of requisite scope and quality. In this paper, however, we focus on one country that satisfies both criteria: Mexico. More specifically, the Mexican economy experienced substantial changes in trade policy in the late 1980s and in the latter half of the 1990s. Moreover, the Mexican government has been collecting individual income data with a panel dimension since the mid-1980s, thus making Mexico an excellent candidate for an application of our general approach.

Our empirical results for the Mexican case can be summarized as follows. First, we find that trade policy changes have a significant short-run effect on income risk for industries with high levels of import penetration. Here a tariff reduction of five percentage points raises the standard deviation of the persistent shocks to income by about 25%. In terms of welfare, we find that this increase in income risk is equivalent to a decrease in lifetime consumption by almost 1% (using a discount factor and degree of risk aversion that are standard in the macroeconomic literature; Cooley, 1995) for workers in the high-import-penetration industries.[8] Second, the effect of the tariff level on income risk is insignificant. Third, while the tariff level has an insignificant mean effect, it nevertheless changes the degree to which macroeconomic shocks affect income risk. For instance, we find that tariff reductions increase the cost of recessions substantially. More specifically, at a tariff level of 10%, a reduction in the growth rate of GDP of 5% is estimated to raise the standard deviation of persistent income shocks by 12%, whereas at a 5% tariff rate, the same reduction in GDP growth increases income risk by 25%. In terms of welfare, this amounts to an increase in the cost of recessions that is equivalent to almost half a percentage point of lifetime consumption. However, our empirical estimates also indicate that tariff reductions decrease individual income risk during economic booms, so that the net welfare cost of tariff reforms is smaller than half a percentage point of lifetime consumption.[9]

At this stage, it is worth pointing out some of the limitations of our analysis. First, we focus exclusively on the link between trade policy and individual income risk, and therefore neglect other channels through which trade policy may affect the economy. More specifically, one would expect trade liberalization to have positive effects on the efficiency of resource allocation and economic growth (the mean of income changes), and these effects are important factors that any comprehensive welfare analysis of trade liberalization ought to take into account. Second, our welfare calculations do not allow the possibility that an increase in income risk might lead to a simultaneous rise in

[6] To see the importance of this distinction more clearly, consider the example of a worker who loses his job due to plant closure or any other "exogenous" event. If the worker quickly finds a new job that pays him as well as the previous job, then his consumption level is not likely to drop by too much either during or after the period of unemployment. If, however, the worker is forced to accept a job that pays him a permanently lower wage (because, for example, job- or occupation-specific human capital has been destroyed), the worker's likely response is to reduce consumption.

[7] It should be clear that our need for longitudinal data follows from our desire to study how trade policy affects the magnitude and frequency of individual income shocks (changes). This is a quite distinct task from that of measuring the impact of trade policy on the distribution of income levels.

[8] Although these are only short-run effects, the fact that we are dealing with permanent income shocks to individual workers means that in this relatively short period, some of the workers get scarred for life. Thus, ex ante, workers are willing to give up a substantial amount of their expected lifetime consumption in return for the elimination of the risk of losing with a trade reform.

[9] Because of space limitations, in this paper we do not attempt to find a precise estimate of this welfare cost, taking into account both the increase in income risk during recessions and the decrease during economic booms. Such an estimate could be found by adopting the methodological approach used in the literature on the welfare cost of business cycles when markets are incomplete. See, for example, Lucas (2003) for more details.

insurance opportunities (endogenous market incompleteness). Finally, the Mexican household survey we use is a rotating panel that follows individual workers for five quarters over time, which means that the panel dimension of our income data is somewhat limited. Thus, our data do not allow us to assess with certainty the persistence of income shocks beyond five quarters. However, a comparison of our estimates of the income risk parameters with existing results that use data sets with a much longer panel dimension suggests that a large fraction of the income shocks we label "persistent" in this paper last for many years (see section IIE for details). In short, the welfare results presented here should not be taken to imply that trade liberalization is welfare reducing; they merely suggest that any comprehensive welfare analysis of trade liberalization ought to take into account the cost of increased labor market risk.

In sum, we develop a general framework that allows us to study empirically the impact of trade reform on individual income risk and evaluate the corresponding welfare effects. We use this framework to study the Mexican economy, which seems well suited for such an analysis. In our empirical implementation of this methodology using longitudinal data on Mexican workers, we find economically significant effects of trade policy on income risk. While the literature has provided us with valuable analyses of the labor market impact of trade policy reforms, it has not thus far examined directly individual income risk or modeled individual income processes in a way that enables welfare analysis of the risk consequences of trade reform. In a similar vein, while several scholars have commented on the potential importance of the link between openness and aggregate volatility in the presence of market incompleteness (Rodrik, 1997),[10] empirical studies have often found the welfare effects of aggregate fluctuations to be quite small (Lucas, 2003).

II. Income Risk

The first stage of our analysis concerns the estimation of individual income risk. Our estimation strategy follows earlier approaches in the literature estimating U.S. labor income risk (Carroll & Samwick, 1997; Gottschalk & Moffitt, 1994; Meghir & Pistaferri, 2004; Storesletten, Telmer, & Yaron, 2004) with some important differences, which we discuss in detail below. As in these papers, we define income risk as the variance of (unpredictable) changes in individual income and distinguish between transitory and persistent income shocks. We eventually focus our attention on the relationship between only persistent shocks (ignoring transitory shocks) and trade policy. This is done for essentially two reasons. First, the transitory term in our econometric specification of the income process will absorb the measurement error in individual income, leading to biased estimates of the magnitude of these shocks. More important, consumption smoothing through borrowing or own saving works well for transitory income shocks (Levine & Zame, 2002), but not when income shocks are highly persistent or permanent (Constantinides & Duffie, 1996; Friedman, 1957; Krebs, 2003, 2004). Thus, highly persistent income shocks have a large effect on consumption volatility and welfare, whereas the effect of transitory shocks is relatively small.

A. *Data*

In Mexico, the National Urban Employment Survey (ENEU) conducts extensive quarterly household interviews in the sixteen major metropolitan areas and is available from 1987 (we use data from 1987 to 1998 in our study). The sample is selected to be geographically and socioeconomically representative. The treatment of sample design, collection, and data cleaning is careful. The survey questionnaire is extensive in scope and covers all standard elements, such as participation in the labor market and earnings. A panel of workers enters the data set each quarter and is tracked for five quarters. We have 44 such panels of 5 quarters each, spanning a total time period of 12 years (48 quarters). We use information on labor market participants between the ages of 16 and 65. Individual panels were constructed by matching workers by their position in an identified household, level of education (years of schooling), age, and sex. Questions referring to labor income refer to income earned in the previous quarter. Workers' earnings include their overall earnings from fixed salary payments, hourly or daily wages, piecemeal work, commissions, tips, and any entrepreneurial earnings earned by the self-employed (in our analysis, observations on wages are winsorized at the 5th and 95th percentiles of the wage distribution in any given quarter).

Table 1 provides the full list of manufacturing industries. Table 2 presents a summary description of the workers surveyed by the ENEU. Other aspects of our ENEU data—the evolution of the mean and variance of earnings and returns to education over time (not presented here but available on request)—match the facts about earnings in the Mexican labor market reported by previous authors.[11]

Data on sectoral trade barriers and other sectoral and macroeconomic variables were obtained from the World Bank. Figure 1 presents tariff data on Mexican industries that are the multilateral tariffs rates applied by Mexico. Subsequent to the signing of the North American Free Trade

[10] Early theoretical analyses of trade patterns and optimal trade policy with aggregate risk and incomplete markets include Eaton and Grossman (1985) and Helpman and Razin (1980), among others. An interesting and somewhat related theoretical literature on international production and trade patterns with incomplete contracting has been developed recently (see Antras, 2003, and Helpman & Grossman, 2002), but it has not (yet) considered explicitly either aggregate or idiosyncratic risk in the economic environment.

[11] See Hanson (2003) for a broad analysis of wage patterns in Mexico in the 1990s based on population census data.

TABLE 1.—INDUSTRY LIST

Industry Code	Description
311	Food products
313	Beverages
321	Textiles
322	Wearing apparel, except footwear
323	Leather products
324	Footwear, except rubber or plastic
331	Wood products, except furniture
332	Furniture, except metal
341	Paper and products
342	Printing and publishing
351	Industrial chemicals
352	Other chemicals
353	Petroleum refineries
356	Plastic products
369	Other nonmetallic mineral products
371	Iron and steel
381	Fabricated metal products
382	Machinery, except electrical
383	Machinery, electric
384	Transport equipment
390	Other manufactured products

FIGURE 1.—INDUSTRY TARIFFS OVER TIME

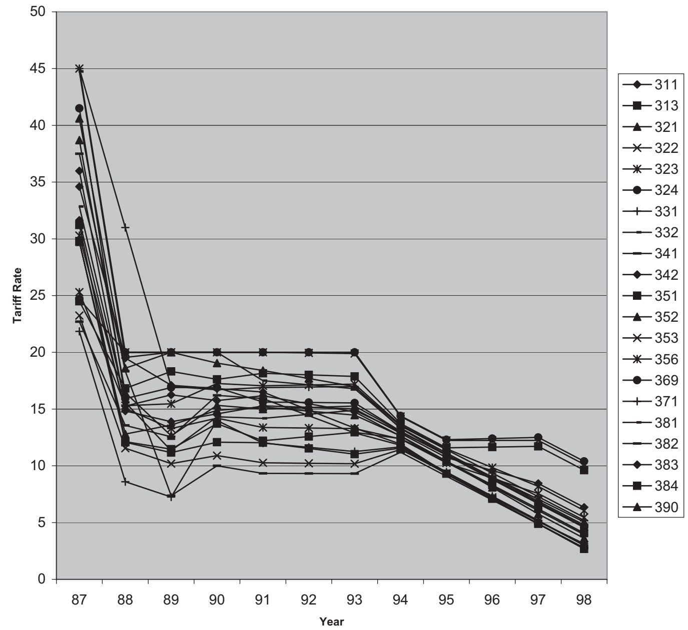

Agreement (NAFTA), the tariff rate are the trade-weighted sum of the multilateral and the NAFTA tariffs.

B. Specification

Our survey data provide us with earnings (wage rate times number of hours worked) of individuals. As in previous empirical work, we assume that the log of this labor income of individual i employed in industry j in period t, $\log y_{ijt}$, is given by

$$\log y_{ijt} = \alpha_{jt} + \beta_t \cdot x_{ijt} + u_{ijt}. \tag{1}$$

TABLE 2.—INDUSTRY AVERAGE CHARACTERISTICS, 1997

Industry	Age	Education	Wage
311	32.11	7.98	14.52
313	31.45	9.76	24.80
321	33.31	8.69	17.09
322	30.02	8.44	13.50
323	29.76	7.82	17.42
324	29.55	7.14	15.66
331	30.83	8.77	14.40
332	30.99	8.31	17.44
341	30.05	8.69	18.31
342	31.68	10.77	23.55
351	34.41	11.93	50.63
352	32.75	11.22	30.06
353	38.54	11.83	41.58
356	30.27	9.16	19.43
369	33.98	7.79	19.27
371	36.31	11.07	47.89
381	32.20	8.85	18.51
382	30.98	10.50	25.91
383	28.81	9.60	23.19
384	29.40	10.12	24.90
390	29.93	9.05	13.92
Mean	31.78	9.40	23.43
s.d.	2.44	1.38	10.79

Note: Age and education are average age and education of the labor force measured in years. Wage denotes the average monthly wage in thousands of pesos.

In equation (1), α_{jt} and β_t denote time-varying coefficients, x_{ijt} is a vector of observable characteristics (such as age and education), and u_{it} is the stochastic component of earnings. The stochastic component u_{ijt} represents individual income changes that are not due to changes in the return to observable worker characteristics. For example, income changes that are caused by an increase in the skill (education) premium are not contained in u_{ijt}. In this sense, u_{ijt} measures the unpredictable part of changes in individual income. Notice that we allow the fixed effects α_{jt} to vary across sectors, but the coefficient β_t is restricted to be equal across sectors. The latter assumption is made in order to ensure that the number of observations is large compared to the number of parameters to be estimated.

We assume that the stochastic term is the sum of two (unobserved) components, a permanent component ω_{ijt} and a transitory component η_{ijt}:

$$u_{ijt} = \omega_{ijt} + \eta_{ijt}. \tag{2}$$

Permanent shocks to income are fully persistent in the sense that the permanent component follows a random walk,

$$\omega_{ij,t+1} = \omega_{ijt} + \varepsilon_{ij,t+1}, \tag{3}$$

where the innovation terms, $\{\varepsilon_{ijt}\}$, are assumed to be independently distributed over time and identically distributed across households. We assume further that $\varepsilon_{ij,t+1} \sim N(0, \sigma^2_{\varepsilon j,t+1})$. The transitory shocks in equation (2) have no persistence, that is, the random variables $\{\eta_{ijt}\}$ are independently distributed over time and identically distributed across households. Clearly, η_{ijt} captures both temporary income shocks and measurement error. We assume that they are normally distributed with zero mean and a variance that is independent of i, but may depend on time or industry: $\eta_{ijt} \sim N(0, \sigma^2_{\eta jt})$.

Our specification for the labor income process is in accordance with the empirical work on U.S. labor income risk. For example, Carroll and Samwick (1997) and Gottschalk and Moffitt (1994) use exactly our specification. Storesletten et al. (2004) assume that the permanent component is an AR(1) process, but estimate an autocorrelation coefficient close to 1 (the random walk case). Finally, some papers have allowed for a third, MA(1), component (see Meghir & Pistaferri, 2004). Notice also that with the exception of Meghir and Pistaferri (2004) and Storesletten et al. (2004), the previous literature has confined attention to the special case of time-independent variances (homoscedastic case). As we discuss in section IIC, the introduction of time variation in the parameters $\sigma^2_{\varepsilon jt}$ and $\sigma^2_{\eta jt}$ makes the estimation of these parameters more challenging.

C. Estimation

Consider the change in the residual of income of individual i between period t and $t + n$:

$$\begin{aligned} \Delta_n u_{ijt} &= u_{ij,t+n} - u_{ijt} \\ &= \varepsilon_{ij,t+1} + \ldots + \varepsilon_{ij,t+n} + \eta_{ij,t+n} - \eta_{ijt}. \end{aligned} \tag{4}$$

We have the following expression for the variance of these income changes:

$$\operatorname{var}[\Delta_n u_{ijt}] = \sigma^2_{\varepsilon j,t+1} + \ldots \sigma^2_{\varepsilon j,t+n} + \sigma^2_{\eta jt} + \sigma^2_{\eta j,t+n}. \tag{5}$$

We use the moment restrictions (5) to estimate the parameters $\sigma^2_{\varepsilon jt}$ and $\sigma^2_{\eta jt}$ using GMM,[12] where the sample analogs to the moment conditions are formed by using the estimates of u_{ijt} obtained as residuals from regressions of labor income on observable characteristics as specified in equation (1)—an approach also used by Meghir and Pistaferri (2004) and Storesletten et al. (2004).[13] Specifically, the estimator is obtained by minimizing

$$\begin{aligned} &\sum_{t,n} (\operatorname{var}[\Delta_n u_{ijt}] \\ &\quad - (\sigma^2_{\varepsilon j,t+1} + \ldots \sigma^2_{\varepsilon j,t+n} + \sigma^2_{\eta jt} + \sigma^2_{\eta j,t+n}))^2. \end{aligned} \tag{6}$$

[12] More specifically, we follow the bulk of the literature and use the equally weighted minimum distance (EWMD) estimator. Altonji and Segal (1996) suggest that the EWMD estimator (identity weighting matrix) is superior to the two-stage GMM estimator (optimal weighting matrix) once small-sample bias is taken into account.

[13] Notice that Meghir and Pistaferri (2004) and Storesletten et al. (2004) exploit additional moment restrictions that follow from the autocovariance function of income changes.

The first-order conditions corresponding to the parameters $\sigma^2_{\varepsilon j,t}$ and $\sigma^2_{\eta j,t}$ are given by

$$\begin{aligned} &\forall\, t : \frac{\partial \Sigma}{\partial \sigma^2_{\varepsilon j,t}} = 0 \\ &\forall\, t : \frac{\partial \Sigma}{\partial \sigma^2_{\eta j,t}} = 0. \end{aligned} \tag{7}$$

Notice that in general there are many more moment conditions (5) than there are parameters to be estimated. More precisely, for each time period t and each industry j, there are two parameters ($\sigma^2_{\varepsilon jt}$ and $\sigma^2_{\eta jt}$), but n moment conditions (5). For example, in our data set on Mexico, for each industry j we have $t = 48$ quarters and $n = 4$ quarters (individuals drop out of the sample after five quarters), and the number of parameters is therefore $2 \times (48)$, whereas the number of moment conditions is approximately $4 \times (48)$.[14] The system is thus overidentified.

Notice also that the objective function (6) is quadratic, which implies that the first-order conditions associated with the corresponding minimum-distance problem are linear in $\sigma^2_{\varepsilon jt}$ and $\sigma^2_{\eta jt}$—a feature that facilitates the estimation substantially. Specifically, the first-order conditions can be organized into a linear equation system,

$$\mathbf{A} \cdot \sigma = \mathbf{b}, \tag{8}$$

where $\sigma = (\sigma^2_{\varepsilon,2} \ldots \sigma^2_{\varepsilon,t} \ldots \sigma^2_{\varepsilon,T}, \sigma^2_{\eta,2} \ldots \sigma^2_{\varepsilon,t} \ldots \sigma^2_{\eta,T})'$ is a $2(T - 1)$-dimensional vector of income parameters (T being the total number of time periods). Estimates of these income parameters can then easily be obtained through matrix inversion: $\sigma = \mathbf{A}^{-1}\mathbf{b}$.

Some intuition for the way in which our approach separates transitory from permanent income shocks can be obtained from the following simple example. Suppose that risk is time invariant, $\sigma^2_{\varepsilon jt} = \sigma^2_{\varepsilon j}$ and $\sigma^2_{\eta jt} = \sigma^2_{\eta j}$, an assumption that has been made by most of the previous empirical literature on income risk. In this case, the moment restrictions (5) become:

$$\operatorname{var}[\Delta_n u_{ijt}] = 2\sigma^2_{\eta j} + n\sigma^2_{\varepsilon j}. \tag{9}$$

Thus, the variance of observed n-period income changes is a linear function of n, where the slope coefficient is equal to $\sigma^2_{\varepsilon j}$. The insight that the random walk component in income implies a linearly increasing income dispersion over time is the basis of the estimation method that several authors have used. For example, Carroll and Samwick (1997) estimate σ^2_{ε} by performing OLS regressions of the left-hand side of equation (9) on n. While the preceding example, with time-invariant parameters, serves to illustrate the intuition

[14] We say "approximate" because toward the very end of the sample period, clearly fewer than $n = 4$ income changes are observed. In the penultimate quarter, for instance, only one income change is observed. However, this does not pose a problem for the estimation of any but the parameters of the very last quarter.

underlying the estimation procedure, it should be clear that our exercise is more general in the sense that it allows arbitrary time variation in the income risk parameters.

D. Estimation Using ENEU Data

The preceding section provided a detailed description of a general econometric methodology that may be used to estimate time-variant income risk parameters given longitudinal data on individual incomes. We note here some additional issues that arise in applying this methodology to our data, with particular emphasis on the type of income risk accounted for by our estimation procedure.

In forming the sample analogs to the moment conditions (5), we use information on all individuals who are present in a given manufacturing industry in both time periods t and $t + n$ (with $n \leq 5$) regardless of their employment status in any intermediate period. In doing so, we pick up shocks to workers who retain their jobs but experience income changes due to changes in their wage rates or the number of hours worked. Moreover, we also account for changes in income experienced by workers who have lost their job in period t but are reemployed in the same industry in some subsequent period $t + n$ (with $n \leq 5$), and this is true even if these workers are unemployed in any intermediate period. In particular, we do account for the long-term earnings losses of a large fraction of displaced workers: all those displaced workers who are reemployed in the same industry but have lost firm- or occupation-specific human capital.[15] In contrast, displaced workers who are reallocated to a different manufacturing industry are not taken into account.[16] However, in our data set, the exclusion of such workers may not cause too much of an underestimation of the income risk parameters since the fraction of displaced manufacturing workers who make the transition from one manufacturing sector to another is very small. Indeed, an examination of reemployment rates for workers who start in manufacturing and go through a period of unemployment suggests that only approximately 10% of these displaced workers undergo a transition from one manufacturing sector to another. Note that this finding is consistent with observations from the United States that most job creation and destruction takes place within industries (see Davis, Haltiwanger, & Schuh, 1996).

Finally, our construction of the sample analogs to the moment conditions (5) could lead to an underestimation of the persistent component of income risk due to the noninclusion of workers undergoing prolonged spells of unemployment (specifically workers who experience unemployment spells exceeding four quarters). However, this is not a severe problem here. As a result of limited unemployment insurance and a very active informal labor market, there are few labor force participants in our survey with extended unemployment durations. Specifically, of those workers looking for work, the proportion who had experienced unemployment durations of four quarters or more was extremely small (less than 0.05% of workers).

Finally, we should mention that the variability in income experienced by workers in our data set derives from both changes in the number of hours worked and changes in the real wage. Real wage changes can be positive or negative. In our data for Mexico, substantial declines in the real wage are quite common. More specifically, Mexico experienced annual declines in the aggregate real wage as high as 25% during our sample period (see Hanson, 2003), implying that the wage rates of some individual workers declined by an even larger amount. Thus, despite the often cited downward rigidity of wages, our sample includes large numbers of workers whose real wages declined substantially.

TABLE 3.—ESTIMATES OF PERSISTENT AND TRANSITORY INCOME SHOCKS, ANNUAL AVERAGES, 1987–1998

Year	σ^2_ε	σ^2_η	Sample Size
87	0.011	0.096	19,136
	(0.003)	(0.002)	
88	0.005	0.101	35,397
	(0.003)	(0.002)	
89	0.004	0.103	28,203
	(0.002)	(0.001)	
90	0.014	0.098	35,167
	(0.002)	(0.001)	
91	0.001	0.103	37,344
	(0.002)	(0.001)	
92	0.006	0.106	54,022
	(0.001)	(0.001)	
93	0.007	0.112	78,741
	(0.001)	(0.001)	
94	0.006	0.110	121,716
	(0.001)	(0.001)	
95	0.014	0.118	164,212
	(0.001)	(0.001)	
96	0.006	0.107	172,766
	(0.001)	(0.001)	
97	0.006	0.104	172,870
	(0.001)	(0.001)	
98	0.008	0.097	158,707
	(0.001)	(0.001)	

Note: Figures shown are annual averages (across industries and quarters) of the point estimates of the persistent shock σ^2_ε and the transitory shock σ^2_η. The figures in parentheses are the averages of the corresponding standard errors. Sample size denotes the numbers of workers surveyed in the respective year.

E. Results

We have individual income data for the time period 1987 to 1998 covering 21 manufacturing sectors in Mexico. Using the methodology outlined above, we estimate the risk parameters σ^2_ε and σ^2_η for each quarter and each manufacturing sector. In tables 3 and 4, we provide the average estimate of σ^2_ε and σ^2_η for each year (averaged across

[15] For the United States, these long-term earnings losses have been estimated to be very large (on average, 25% for high-tenure workers according to Jacobson, LaLonde, and Sullivan, 1993).

[16] This allows us to circumvent the extremely difficult problem of assigning industries (and, importantly, trade policy) to individuals who transit to different industries (specifically, Which tariff rates should we believe the workers to be exposed to: the industry to which they initially belonged or some subsequent industry?). Since we would expect the variance in outcomes to be higher for individuals who switch industries, our estimates of income risk should perhaps be seen as a lower bound.

TABLE 4.—ESTIMATES OF PERSISTENT AND TRANSITORY INCOME SHOCKS, INDUSTRY AVERAGES, 1987–1998

Industry	σ^2_ε	σ^2_η	Industry	σ^2_ε	σ^2_η
311	0.013	0.131	352	0.020	0.111
	(0.0004)	(0.0003)		(0.0025)	(0.0019)
313	0.012	0.088	353	0.002	0.081
	(0.0007)	(0.0005)		(0.0009)	(0.0007)
321	0.005	0.097	356	0.006	0.079
	(0.0006)	(0.0005)		(0.0016)	(0.0011)
322	0.012	0.124	369	0.011	0.113
	(0.0008)	(0.0006)		(0.0014)	(0.0011)
323	0.008	0.107	371	0.003	0.110
	(0.0022)	(0.0015)		(0.0031)	(0.0025)
324	0.004	0.088	381	0.006	0.125
	(0.0002)	(0.0001)		(0.0006)	(0.0004)
331	0.004	0.120	382	−0.002	0.098
	(0.0027)	(0.0020)		(0.0015)	(0.0011)
332	0.019	0.121	383	0.008	0.056
	(0.0017)	(0.0013)		(0.0002)	(0.0002)
341	0.004	0.102	384	0.004	0.073
	(0.0016)	(0.0012)		(0.0002)	(0.0001)
342	0.011	0.134	390	0.005	0.143
	(0.0016)	(0.0012)		(0.0062)	(0.0047)
351	0.012	0.107			
	(0.0029)	(0.0023)			

Note: Figures shown are averages over time of the point estimates of the persistent shock σ^2_ε and the transitory shock σ^2_η for the respective industries. The figures in parentheses are the averages of the corresponding standard errors.

industries) and for each industry (averaged over time), respectively.[17] The mean value (across industries and over time) of the quarterly variance of the persistent shock, σ^2_ε, is estimated to be 0.008, or 0.032 annualized (σ_ε, is estimated to have a mean quarterly value of 0.09 and a mean annualized value of 0.18).[18] As expected, given the extent of measurement error in the income data (see our previous discussion), the estimated variances of transitory shocks are much larger in magnitude. More precisely, the mean value of the annualized variance of transitory shocks is 0.2 (an annual standard deviation of 45%), which is clearly too large to be a true measure of income volatility.

It seems informative to compare our estimates of the permanent component of income risk, σ^2_ε, with the estimates obtained by the extensive empirical literature on U.S. labor market risk using annual income data drawn from the PSID. Most of these studies find an average value of around .0225 for the annual variance σ^2_ε (Carroll & Samwick, 1997; Gottschalk & Moffitt, 1994; Storesletten et al., 2004), with a value of $\sigma^2_\varepsilon = .0324$ being the upper bound (Meghir & Pistaferri, 2004). Assuming that these income shocks are independently and identically distributed (i.i.d.) over time (the maintained random walk assumption), this means that these studies have found a quarterly variance of $\sigma^2_\varepsilon = .0056$, with one study estimating $\sigma^2_\varepsilon = .008$. Thus, the average value of our estimates of permanent income risk is in line with the estimates obtained by the previous literature on U.S. labor market risk, although our estimates lie somewhat on the high end. Notice that our estimates are obtained using a five-quarter rotating panel, whereas Carroll and Samwick (1997), Meghir and Pistaferri (2004), and Storesletten et al. (2004) use the PSID data with a panel dimension of many years. Thus, as long as Mexican workers face similar amounts of permanent labor income risk as U.S. workers (or more), this result suggests that most income shocks we label "permanent" in this paper indeed persist for a very long time.

III. Trade Policy and Income Risk

The procedure outlined in the previous section provides estimates of individual income risk, $\sigma^2_{\varepsilon jt}$, for each industry (manufacturing sector) j and time period (quarter) t. We now use these time-varying, industry-specific estimates in conjunction with observations on trade policy, τ_{jt}, to estimate the relationship of income risk, $\sigma^2_{\varepsilon jt}$, and openness, τ_{jt}, using a linear regression model. In this paper we focus on the permanent component of income risk, σ^2_ε, instead of the transitory component, σ^2_η, for two reasons: transitory income shocks are unlikely to generate substantial consumption volatility, and σ^2_η is likely to contain a large amount of measurement error. Despite these theoretical arguments, it might still be of interest to study the relationship between trade policy and income risk using σ^2_η as a measure of income risk. We therefore also conducted a similar regression analysis (not reported here) for transitory income shock

[17] The averages presented in tables 3 and 4 are merely summary descriptions and do not allow any direct inferences regarding the relationship between trade policy and income risk.

[18] Given that in section III we seek to uncover the relationship between trade policy and income risk using our estimates of the income risk parameters σ_ε, it is also interesting to investigate to what extent these estimates differ across industries and over time after making some adjustment for the fact that there is estimation error. To quantify this variation, we use the methodology of Krueger and Summers (1988). More specifically, we compute a measure of the adjusted standard deviation of the point estimates of the income risk parameters. It turns out that this number (0.018) is over twice the mean value of σ_ε in our sample, indicating that the variation in σ_ε across industries and over time is significant in our exercise.

parameters, σ_{η}^2, but we did not find any statistically significant relationship between transitory shocks to income and trade policy.

A. Specification

We first consider a linear specification that allows industry fixed effects and aggregate time effects:

$$\sigma_{\varepsilon jt}^2 = \alpha_0 + \alpha_j + \alpha_t + \alpha_\tau \tau_{jt} + \alpha_{\delta 1}\Delta\tau_{jt} + \alpha_{\delta 2}\Delta\tau_{jt}D_{jt} + \nu_{jt}. \tag{10}$$

In this equation, we have included on the right-hand side the following variables: τ, the ad valorem sectoral tariff rate; $\Delta\tau$, the change in the tariff over the preceding year; $\Delta\tau D$, the tariff change over the preceding year interacted with an indicator variable that takes the value 1 if the import penetration ratio is greater than its sample median and 0 otherwise[19]; α_j, an industry fixed effect; and α_t, a time dummy that captures general macroeconomic trends in the economy.

The inclusion of industry dummies in equation (10) allows us to control for any fixed industry-specific factors that may affect the level of riskiness of income in that industry. Moreover, the inclusion of time dummies controls for any changes in macroeconomic conditions that affect the level of income risk. While this ensures that our estimation results are not driven by changes in macroeconomic conditions (business cycle effects or long-run structural changes) unrelated to trade policy, it also means that identification of the relationship between $\sigma_{\varepsilon jt}^2$ and τ_{jt} will have to be based on the differential rate of change in trade barriers across sectors over time—or the vector of observations on tariffs in the panel corresponding to equation (10) will be perfectly collinear with the time-dummy vector. This, however, does not pose problems for our estimation since trade barriers in Mexico and their changes over time in fact exhibit substantial cross-sectional variation.[20]

Specification (10) provides the starting point for our econometric analysis. An alternate specification is

$$\sigma_{\varepsilon jt}^2 = \alpha_0 + \alpha_j + \alpha_\tau \tau_{jt} + \alpha_{\delta 1}\Delta\tau_{jt} + \alpha_{\delta 2}\Delta\tau_{jt}D_{jt} + \beta_e \Delta e_t + \beta_g g_t + \phi_e(1+\tau_{jt})\Delta e_t + \phi_g(1+\tau_{jt})g_t + \nu_{jt}. \tag{10$'$}$$

Specification (10′) explores the sources of variation in risk over time by dropping time dummies and including instead the following two macroeconomic variables: Δe, the depreciation of the real exchange rate over the preceding year, and g, the GDP growth rate. Also included are the interaction terms $(1 + \tau)\Delta e$ and $(1 + \tau)g$, which measure the extent to which the relationship between income risk and these macroeconomic factors varies with trade policy.[21]

Several econometric issues arise in the estimation of equations (10) and (10′), most of which we discuss in more detail below (sections IIIC and IIID). At this stage, we note only the following. First, one concern is that the left-hand-side variable, income risk, is estimated and not observed. This is not a substantial problem by itself, as it is well known that while "measurement error" in the dependent variable does reduce precision, it does not bias our estimates. Second, a concern arises from the fact that the estimates of $\sigma_{\varepsilon jt}^2$ have different standard errors across industries, that is, the specification we have described above suffers from a heteroskedasticity problem. Third, since the industries all belong to the same macroeconomic environment, there is a possibility of contemporaneous correlation in their σ's even after controlling for observable macroeconomic factors as in (10′)—Cov $(\nu_{jt}\nu_{j't}) \neq 0$. Finally, serial correlation in income volatility within an industry is a possibility: Cov $(\nu_{jt}\nu_{jt'}) \neq 0$. Given the possible presence of heteroskedasticity, spatial correlation, and serial dependence, consistent estimates of the standard errors associated with the coefficient estimates in equations (10) and (10′) are obtained by using robust estimation techniques.

B. Results

In equation (10), the effect of the tariff level on income risk is given by the coefficient α_τ, and the effect of tariff changes on income risk is given by the coefficient α_δ. Table 5 presents the estimation results (with corresponding estimates obtained using effective rates of protection presented in table 6). We note first that the estimate of α_τ is insignificant, and we are therefore unable to reject that the mean effect of the tariff level on income risk is 0. However, trade policy changes appear to have an economically and statistically significant effect on income risk in industries with high levels of import penetration. In the first column of table 5, we are able to see that this effect derives specifically from sectors with above-median level of import penetration ($D = 1$). Here, ($\hat{\alpha}_{\delta 1} + \hat{\alpha}_{\delta 2} = -0.125$, with an estimated standard error of 0.05). This estimate indicates that lowering the tariff rate by 5% would, for a year, raise σ_ε^2 by .00625 from, for example, .008 (its mean value) to .01425. In terms of the standard deviation σ_ε, this amounts to an increase from .089

[19] Clearly, $\alpha_{\delta 1}$ measures the effect of a trade policy change in sectors that had lower-than-median import penetration both before and after this change, and $\alpha_{\delta 1} + \alpha_{\delta 2}$ correspondingly measures the effect of trade policy changes in sectors that had higher-than-median import penetration over time. This is also true with specification (10′) below.

[20] For instance, in Mexico, tariffs varied between 45% and 20% prior to the trade reforms of 1987 and ranged between 20% and 10% by 1994, implying a variation in tariff changes across sectors that is quite substantial.

[21] Note that the only variable that is interacted with the dummy variable D (representing greater-than-median import penetration) is the change in tariffs, $\Delta\tau_{jt}$. The remaining variables such as exchange rate depreciation, Δe_t, and growth rate of GDP, g_t, are already interacted with the tariff level (which itself has a quite strong within-industry correlation with import penetration). Estimating equation (10′) separately for industries with $D = 0$ and $D = 1$ gave results very similar to those reported here.

TABLE 5.—TRADE POLICY AND INCOME RISK: PANEL ESTIMATES

Variables	σ_ε^2	σ_ε^2
τ	0.043	−0.140
	(0.060)	(0.051)
$\Delta\tau$	−0.035	0.027
	(0.044)	(0.031)
$\Delta\tau \times D_n$	−0.090	−0.109
	(0.047)	(0.047)
Δe		−0.621
		(0.207)
g		−1.208
		(0.414)
$\tau \times \Delta e$		0.539
		(0.184)
$\tau \times g$		1.055
		(0.370)
Time effects	Included	
Industry fixed effects	Included	Included
N	945	945
R^2	0.058	0.044

Note: Figures in parentheses are robust standard error estimates obtained by allowing for heteroskedasticity, contemporaneous correlation of errors across industries, and serial correlation within industries. In the table, τ denotes the tariff rate, g denotes the growth rate of GDP, and Δe denotes the percentage appreciation of the exchange rate.

TABLE 6.—TRADE POLICY AND INCOME RISK: EFFECTIVE RATES OF PROTECTION

Variables	σ_ε^2	σ_ε^2
τ	0.019	−0.108
	(0.043)	(0.046)
$\Delta\tau$	−0.009	0.010
	(0.031)	(0.024)
$\Delta\tau \times D_n$	−0.120	−0.137
	(0.068)	(0.067)
Δe		−0.440
		(0.179)
g		−0.940
		(0.348)
$\tau \times \Delta e$		0.377
		(0.158)
$\tau \times g$		0.814
		(0.310)
Time effects	Included	
Industry fixed effects	Included	Included
N	945	945
R^2	0.058	0.042

Note: Figures in parentheses are robust standard error estimates obtained by allowing for heteroskedasticity, contemporaneous correlation of errors across industries, and serial correlation within industries. In the table, τ denotes the tariff rate, g denotes the growth rate of GDP, and Δe denotes the percentage appreciation of the exchange rate.

to .1193, that is, an increase by more than 30%—a substantial increase in income risk.[22]

Estimates from equation (10′) are presented in the third column of table 5. Note that tariff changes in high-import-penetration sectors continue to have economically and statistically significant effects of magnitude quite similar to those obtained from estimation of equation (10) ($\hat{\alpha}_{\delta 1} + \hat{\alpha}_{\delta 2} = -0.092$, with an estimated standard error of 0.045). More specifically, a 5% reduction in tariffs increases σ_ε^2 from a mean level of .008 to .0126, which in terms of the standard deviation σ_ε amounts to an increase from .089 to .1122 (a 25% increase). Interestingly, the coefficient α_τ is now significant. However, the effect of the tariff level on income risk is now given by ($\alpha_\tau + \phi_e \Delta e + \phi_g g$). After substituting in the mean values of Δe and g from the sample, this estimated sum revealed to be insignificantly different from 0 ($\hat{\alpha}_\tau + \hat{\phi}_e \bar{\Delta} e + \hat{\phi}_g \bar{g} = 0.02$, with an estimated standard error of 0.02). Thus, we are again unable to reject that the mean effect of the tariff level on income risk is zero.[23]

Consider now our estimates of how the tariff level alters the effect of macroeconomic variables on income risk. The coefficient on real exchange rate depreciation, β_e, is estimated to be negative and significant, as is the coefficient on GDP growth, β_g, while the coefficients ϕ_e and ϕ_g relating to the interaction terms, $(1 + \tau)\Delta e$ and $(1 + \tau)g$, are both positive and significant. The extent to which the tariff level alters the effects of exchange rate changes on income risk is given by ϕ_e (as reported in table 4, this parameter is estimated to have a mean value of 0.54 and an estimated standard error of 0.18). Consider a real exchange rate appreciation of 10% under two scenarios—when the tariff rate is 10% and when it is 5%. If the tariff rate is 10%, our estimates indicate that an exchange rate appreciation of 10% (in the preceding year) raises σ_ε^2 from 0.008 to 0.0108 (an increase of just about 35%). In contrast, if the tariff rate is 5% instead, the same appreciation implies an increase in income risk from 0.008 to 0.013 (an increase of over 60%). Similarly, if the growth rate of GDP, g, is lowered by 5%, σ_ε^2 is raised from 0.008 to 0.01 (an increase of over 25%) when the tariff rate is 10% but the same change in g results in a short-run increase in income risk from 0.008 to 0.013 (an increase of over 60%) when the tariff rate is at 5%. Of course, as noted earlier, our empirical estimates also indicate that tariff reductions lead to a corresponding reduction in individual income risk during economic booms. Overall, our estimates suggest that the magnitude of the (short-run) effects of macroeconomic shocks on income risk is significantly altered by the tariff level.

The dependence of the income risk parameter σ_ε^2 on cyclical conditions is observed not only in Mexico; it has also been well documented for the United States (Meghir & Pistaferri, 2004; Storesletten et al., 2004). However, this literature has not studied how trade policy affects this

[22] As we discuss in section IIID, these results are robust to the inclusion of a full set of interactions involving the indicator variable D, that is, including the product of τ and D, as well as the product of the constant term and D (D itself) on the right-hand side. Moreover, the coefficients on $\tau \times D$ and D are themselves insignificantly different from 0 in each of the specifications we examine.

[23] Our estimates of the timing and magnitude of the effect of trade policy changes on measured income shocks (large changes in the year following policy changes and 0 mean effects) also indicate that our results are not being driven by other "unobserved" factors such as skill and sector-biased technical changes that are possibly correlated with trade policy changes. More specifically, we would expect any such changes in technology to have an impact on income in a gradual manner, taking several years for its full impact to be realized. Note also that our own estimates of the returns to education suggest a striking similarity across manufacturing sectors in Mexico, which provides indirect evidence against the view that technological progress in Mexico during the relevant sample period was both skill and sector biased.

dependence of idiosyncratic risk on cyclical conditions. Thus, the estimation results reported in table 5 provide the first empirical evidence that trade liberalization increases the sensitivity of idiosyncratic risk to business cycle conditions. Theoretically, one might speculate that a mechanism similar to the one modeled by Newberry and Stiglitz (1984) is behind our empirical finding. More specifically, Newberry and Stiglitz (1984) argue that a negative productivity shock would have a smaller equilibrium effect on output and employment in a closed economy than an open one, as prices rise with a negative supply shock in the former but are constrained by world prices in the latter. With heterogeneous effects on firms and individuals, the link between macroeconomic downturns and idiosyncratic income risk may therefore also be amplified in more open economies. A more rigorous modeling of this idea within the context of a dynamic general equilibrium model with incomplete markets is an interesting topic for future research.

C. Endogeneity and Selection Bias

One concern that arises in our estimation of equations (10) and (10′) is that tariff rates may not be fully exogenous. Indeed, the theoretical literature on the political economy of trade policy has proposed several hypotheses concerning the endogenous determination of tariffs. Furthermore, a number of empirical studies have explained (partially) the cross-industry variation in tariffs using a number of economic and political variables that vary across industries such as the lobbying strength and employment size of particular sectors.[24] While the literature has not studied (or indeed even suggested) income risk as a determinant of cross-sectional variation in trade policy, the possibility that it might be a relevant determinant of policy makes it potentially problematic. Consider, for instance, an economy in which raising the tariff rate in a sector would in fact lower income risk in that sector. Consider further that the government there is "equity" minded and chooses higher protection levels for industries with intrinsically high levels of income risk, thereby eliminating cross-sectional variation in income risk. If such an economy were studied purely in the cross-section, it may appear that there is no relation between trade policy and income risk even though such a relationship does exist. This type of purely cross-sectional endogeneity, however, is not a problem for our empirical analysis since we follow industries over time. More precisely, the *within* estimator we use is formed by considering changes *within* industries in income risk and tariffs over time, and any endogeneity bias deriving from purely cross-sectionally varying political-economy determinants of trade policy is therefore eliminated.

Along the time dimension, estimation bias could arise if the government attempts to protect vulnerable industries by raising tariff rates for industries that have experienced an increase in income risk. While such endogeneity bias is in principle a matter of concern, at least two facts speak against this view. First, the trade policy changes that we study here were undertaken during major policy reform episodes (both in the late 1980s and under NAFTA), and some observers have argued that the lowering of trade barriers was mainly used by the Mexican government to signal its commitment to overall policy reform (Esquivel & Tornell, 1995). Second, and somewhat related to the first point, in our data virtually no industry experienced a rollback of the liberalization effort once tariff rates had been reduced. Finally, we note that such a pattern of endogeneity would cause only a bias against our reported findings. That is, if such bias exists, the true short-run effect of trade policy changes on income risk is even larger than what we report in this paper. However, it also means that our finding that trade liberalization has no long-run "level effect" could be the result of two opposing effects canceling each other out.[25]

Estimation bias could also arise if systematic changes in nontariff barriers (NTB) reversed the effects of tariff reductions, but we did not take these changes into account. To ensure that this is not the case, we studied the patterns in the use of NTBs in Mexico in the years included in our sample. NTB use in Mexico primarily took the form of antidumping duties in these years, and the antidumping duties were most heavily used in the Basic Metal Products industry.[26] Studying the link between trade policy and income risk using data from the remaining industries did not alter qualitatively or quantitatively any of the reported estimates (see table 7).

Our estimation results could also be biased if there is unobserved heterogeneity among workers and industries, and heterogeneous workers select into different industries. Suppose, for example, that industries with high levels of protection (high tariff levels) are also ones with low job destruction rates. Suppose further that there are two types of workers, good and bad, and that good workers quickly find a new job in the event of job displacement, but bad workers do not. Other things being equal, we would expect bad workers to move to high-protection industries. In this world, high tariff rates lower income risk because they reduce job destruction rates, but they also attract high-risk (bad) workers, leading to a downward bias of our empirical estimates of the relationship between income risk and tariff levels—the coefficient α_τ in equation (10). Thus, it is possible that

[24] See Gawande and Krishna (2003) for a survey discussion.

[25] Notice also that despite the work by Alesina and Drazen (1994) and others, major trade policy reforms are in general rather difficult to understand theoretically once policy is treated as being endogenous. The dominant theory of endogenous trade policy determination—the interest group theory—simply does not predict such dramatic changes in policy. Since the competing strengths of various interest groups are not expected to (and do not) change dramatically over the medium term, the theory predicts stickiness in trade policy over these horizons (consistent with what is observed most of the time). Without theoretical guidance, the choice of suitable "exogenous variables" to help with identification is even more difficult than usual.

[26] See the recent study, "Mexico's Experience with the Use of Anti-Dumping Measures," 2002, http://ctrc.sice.oas.org/geograph/antidumping/mexico.pdf.

TABLE 7.—TRADE POLICY AND INCOME RISK: ROBUSTNESS

Variables	σ_ε^2 (Antidumping Excluded)	σ_ε^2 (1995 Excluded)
τ	−0.139	−0.150
	(0.053)	(0.055)
$\Delta\tau$	0.029	0.028
	(0.032)	(0.032)
$\Delta\tau \times D_n$	−0.112	−0.116
	(0.048)	(0.046)
Δe	−0.608	−0.540
	(0.213)	(0.226)
g	−1.204	−1.303
	(0.423)	(0.466)
$\tau \times \Delta e$	0.527	0.472
	(0.189)	(0.199)
$\tau \times g$	1.050	1.123
	(0.378)	(0.414)
Industry fixed effects	Included	Included
N	809	861
R^2	0.04	0.045

Note: Figures in parentheses are robust standard error estimates obtained by allowing for heteroskedasticity, contemporaneous correlation of errors across industries, and serial correlation within industries. In the table, τ denotes the tariff rate, g denotes the growth rate of GDP, and Δe denotes the percentage appreciation of the exchange rate. In the "AD Excluded" column, observations from industries with high levels of antidumping protection were excluded. In the "1995 Excluded" column, observations from the year 1995 have been excluded. See section IV for a detailed discussion.

our empirical finding that tariff levels have no effect on income risk is simply due to this type of selection bias.[27]

In general, it is difficult to deal with the type of selection bias we have just described. However, further analysis of the data mitigates our concern of the bias due to selection effects. Specifically, we would expect workers with low job-finding rates to be mainly low-ability workers. Using years of schooling as an observable proxy for (unobserved) ability, selection would bias our fixed-effects estimates if tariff changes and changes in mean years of schooling were negatively correlated across industries (that is, industries that experience a lower reduction in protection employ a greater proportion of low-skilled workers). However, in our data set, the correlation between tariff changes and changes in average years of schooling is very small (−0.06) and highly insignificant.

Clearly there could be unobserved ability differences among workers that are uncorrelated with years of schooling, in which case selection bias might still be problematic even if the cross-industry correlation between years of schooling and tariff rates is nil. However, even in this case, we would expect any selection bias to manifest itself in unexplained wage differentials across sectors, at least as long as high-ability workers are paid higher wages. A casual examination of the data, however, suggests that such cross-industry wage differentials are small (at least compared to the differences in the level of income risk across industries and our estimates of changes in these magnitudes following trade policy changes). More precisely, across the manufacturing sectors we study, the mean industry wages are highly correlated with mean educational attainment. That is, the R^2 of a simple cross-sectional regression of average earnings on average worker characteristics is about 0.8 (see the data presented in table 2). Thus unobserved worker characteristics have very little influence on average earnings in an industry, suggesting little selectivity of workers of differing (unobserved) abilities into different manufacturing sectors in our data.

Finally, as always, the presence of relevant omitted variables may bias our results. Specifically, any sector-specific time-varying variables that are also correlated with trade policy changes may result in a bias in our estimates of the effect of trade policy changes on income risk. We leave an exploration of any such linkages for future research.

D. Robustness

We conducted a series of additional estimation exercises to study the robustness of the findings reported here. The effective rate of protection was computed (using the tariff series and input-output matrices for Mexico) and used in place of the raw tariff series in estimating equation (10′). To ensure that the dramatic nominal exchange rate devaluation undertaken by the Mexican authorities at the end of 1994 did not drive our results, equation (10′) was estimated by dropping observations from the year 1995. These results are reported in table 7. As is evident, dropping observations from the year immediately following the exchange rate crisis in Mexico does not alter our results (this should not be too surprising, as we have already taken aggregate output and exchange rate changes into account in the econometric specification).

An additional point concerns the lagged effects of policy changes. Note that we measure tariff changes as the change between the beginning-of-year tariffs of two subsequent years. The corresponding change in income risk measures the average effect over a total of a two-year period; thus, a tariff change implemented at the beginning of 1988 could affect income risk in the last quarter of 1989, and this change in income risk would still be taken into account in our specification, equation (10′). Estimation results with specifications in which we included lagged tariff changes (and other lagged independent variables) on the right-hand side of equation (10′) did not support the inclusion of such lags.

Experimenting with other specifications did not suggest any significant qualitative or quantitative changes in our results concerning the effect of trade policy changes on income risk. Thus, in alternate specifications, we included the indicator variable D (indicating greater than median import penetration) both uninteracted and, separately, by interacting it with both tariff levels and tariff changes. These variables were themselves insignificant, and their inclusion did not result in any change in the previously estimated

[27] If trade liberalization mainly targets high-protection industries, and high-risk workers leave industries that experience large tariff cuts, then this self-selection effect also causes a downward bias of our estimates of the relationship between tariff changes and income risk—the coefficients $\alpha_{\delta 1}$ and $\alpha_{\delta 2}$ in equation (10). Thus, the true short-run effect of trade liberalization might be even larger than the (already substantial) effect reported in table 4.

coefficients. We explored logarithmic specifications trying separately both log (σ_ε^2) and log (σ_ε) on the left-hand side. In general, these specifications did not alter the relationship between trade policy and income risk even if they yielded coefficients on the exchange rate variables that were less significant than in previous runs. Finally, we run our extended specification (10′) using time dummies instead of the macroeconomic variables—while retaining the interactions of GDP growth and exchange rate appreciation with tariffs on the right-hand side. Again, these did not result in any appreciable quantitative or qualitative change in our results concerning the effect of trade policy changes on income risk.

IV. Income Risk and Welfare

The preceding discussion has outlined our approach to estimating the relationship between trade policy and income risk. We now turn to the analysis of the link between income risk and welfare, which is provided by a simple dynamic model with incomplete markets along the lines of Constantinides and Duffie (1996) and Krebs (2004).[28] The model extends the basic insights of the large literature on the permanent income hypothesis to a general equilibrium setting with isoelastic preferences and incomplete markets. It remains tractable enough to permit closed-form solutions for equilibrium consumption and welfare yet is rich enough to provide a tight link to the empirical analysis. Clearly our goal here is not to provide a complete assessment of the effects of income risk on welfare, taking into account all possible channels, but rather to articulate a simple framework that allows us to obtain indicative estimates of welfare change through the income risk channel.

The model features long-lived workers that make consumption-versus-saving choices in the face of uninsurable income shocks. These income shocks are permanent, which implies that "self-insurance" through borrowing or own saving is an ineffective means to smooth out income fluctuations. Indeed the model is set up such that in equilibrium, workers will not self-insure at all. That is, income shocks translate one-to-one into consumption changes. In accordance with Constantinides and Duffie (1996) and Krebs (2004), we consider an exchange economy. Thus, we rule out by assumption any effect of changes in income risk on aggregate output.[29] In this section, we briefly discuss the basic assumptions of the model and state the main welfare results. All derivations are relegated to the supplemental appendix, available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00002.

A. Model

Time is discrete and open-ended. The income of worker i employed in industry j in period t is denoted by $\tilde{y}_{ijt}$. Income is random and defined by an initial level $\tilde{y}_{ij0}$ and the law of motion,

$$\tilde{y}_{ij,t+1} = (1 + \mu_{j,t+1})(1 + \theta_{ij,t+1})\tilde{y}_{it}, \tag{11}$$

where $\mu_{j,t+1}$ is a mean growth-rate effect common across workers in the sector and $\theta_{ij,t+1}$ is an individual-specific shock to the growth rate of income. We assume that log $(1 + \theta_{ij,t+1})$ is normally distributed with time- and industry-dependent variance $\sigma_{j,t+1}^2$. Although the distribution of individual-specific shocks may change over time, the shocks are unpredictable in the sense that current and future shocks are uncorrelated. To ensure that workers are ex ante identical, we also assume that the distribution of shocks is identical across workers.

Each worker begins life with no initial financial wealth. Workers have the opportunity to save at the common interest rate r_t, but they cannot borrow. Hence, the sequential budget constraint of worker i reads

$$\begin{aligned} a_{ij,t+1} &= (1 + r_t)a_{ijt} + \tilde{y}_{ijt} - c_{ijt} \\ a_{ijt} &\geq 0, \quad a_{ij0} = 0. \end{aligned} \tag{12}$$

Here c_{ijt} denotes the consumption of worker i employed in industry j in period t and a_{ijt} his asset holdings at the beginning of period t (excluding interest payments in this period).

Workers have identical preferences that allow a time-additive expected utility representation:

$$U(\{c_{ijt}\}) = E\left[\sum_{t=0}^{\infty} \beta^t u(c_{ijt})\right]. \tag{13}$$

Moreover, we assume that the one-period utility function, u, is given by $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\gamma \neq 1$, or $u(c) = \log c$, that is, preferences exhibit constant degree of relative risk aversion γ.

B. Welfare

In the supplemental appendix, we derive an explicit formula for equilibrium welfare that depends on the preference parameters β and γ and the income parameters μ_{jt} and σ_{jt}^2, where σ_{jt}^2 is the variance of the log-normally distributed income shocks θ. We also show that the variance σ_{jt}^2 of income process (11) can be identified with the variance $\sigma_{\varepsilon jt}^2$ of the permanent component of our empirical specification, equation (1). This provides a tight link between the empir-

[28] Notice that in contrast to Constantinides and Duffie (1996) and Krebs (2004), this paper considers a model with multiple sectors (industries) that differ with respect to the amount of income risk households have to bear. Further, in the current model, workers can save but not borrow.

[29] Krebs (2003) considers a production economy with only permanent income shocks and shows again that self-insurance is highly ineffective. Thus, the result that self-insurance is not very effective does not depend on the zero aggregate saving feature of the endowment economy assumed here.

ical results obtained in section II and the welfare analysis conducted in this section. We now briefly outline and discuss the main welfare results.

For simplicity, assume that the income parameters are time and industry independent: $\mu_{jt} = \mu$ and $\sigma^2_{\varepsilon jt} = \sigma^2_\varepsilon$. Suppose further that trade reform changes the tariff rate from τ to $(1 + \Delta_\tau)\tau$ permanently, and that this change in the rate leads to a corresponding permanent change in income risk from σ^2_ε to $(1 + \Delta_\sigma)\sigma^2_\varepsilon$. Clearly, the change in income risk $\Delta_\sigma\sigma^2_\varepsilon$ corresponds to the long-run effect associated with the level term, $\alpha_\tau\tau$, on the right-hand side of our regression, equation (10). We can find the welfare effect of the change in risk, Δ_σ, by calculating the compensating variation in lifetime consumption, Δ_c. That is, we can ask by how much we have to change consumption in each period and state of the world to compensate the household for the change in income risk. In the supplemental appendix, we show that this compensating differential, expressed as a percentage of lifetime consumption, is given by

$$\Delta_c = \left(\frac{1 - \beta(1+\mu)^{1-\gamma}\exp(.5\gamma(\gamma-1)(1+\Delta_\sigma)\sigma^2_\varepsilon)}{1 - \beta(1+\mu)^{1-\gamma}\exp(.5\gamma(\gamma-1)\sigma^2_\varepsilon)}\right)^{1/1-\gamma} - 1 \ \textit{if}\ \gamma \neq 1 \tag{14}$$

$$\Delta_c = \exp\left(\frac{\beta}{(1-\beta)^2}\frac{\sigma^2_\varepsilon\Delta_\sigma}{2}\right) - 1 \ \textit{if}\ \gamma = 1.$$

Equation (14) shows how to translate long-run changes in labor income risk, Δ_σ, into equivalent changes in lifetime consumption, Δ_c. It provides the answer to the following question: How much lifetime consumption are risk-averse workers willing to give up in return for not having to experience the increase in income risk caused by a change in trade policy? Notice that equation (14) is the result of an ex ante welfare calculation under rational expectations. More specifically, it assumes that workers do not know who will lose and who will gain from trade reform, but they know to what extent trade reform creates winners and losers (the effect of trade reform on the income risk parameters is known ex ante).

The welfare expression (14) assumes that the change in σ^2_ε is permanent. However, we are also interested in the welfare effect of an increase in income risk from σ^2_ε to $(1 + \Delta_\sigma)\sigma^2_\varepsilon$ for n periods. In this case, the welfare effect is given by

$$\Delta_c = \left[\left(\frac{1-x}{1-x'}\right)(1 - x'^{n+1}) + xx'^{n}\right]^{1/\gamma-1} - 1 \ \textit{if}\ \gamma \neq 1$$

$$\Delta_c = \exp\left(\frac{\beta(1-\beta^n)}{2(1-\beta)^2}\sigma^2_\varepsilon\Delta_\sigma\right) - 1 \ \textit{otherwise}, \tag{15}$$

where we introduced the following notation:

$$x = \beta(1+\mu)^{1-\gamma}\exp(.5\gamma(\gamma-1)\sigma^2_\varepsilon)$$

$$x' = \beta(1+\mu)^{1-\gamma}\exp(.5\gamma(\gamma-1)(1+\Delta_\sigma)\sigma^2_\varepsilon).$$

The welfare expressions (14) and (15) have some intuitive properties. First, the welfare effect of a change in income risk is a nonlinear and increasing function of the initial level of income risk. Put differently, if workers are already exposed to a large amount of income risk, then increasing income risk is increasingly costly. This property explains why the welfare effects we find in this paper are so much larger than the welfare cost of business cycles found in the macroeconomic literature (Lucas, 2003). Second, the welfare effects are increasing in the risk-aversion parameter γ: the more risk averse the workers are, the stronger is the welfare effect of a change in income risk. Finally, the welfare effects are the same for all workers regardless of their wealth. This property is the result of the joint assumption of homothetic preferences and an income process defined as in equation (11).

C. *Results*

Welfare expressions (14) and (15) form the basis for our quantitative welfare analysis of trade reform. In order to conduct such an analysis, we need information about the income parameters μ, σ^2_ε, and Δ_σ and the preferences parameters β and γ. Our empirical analysis provides estimates of the income parameters. We estimate an average value of the quarterly variance, σ^2_ε, of .008 (averaged across industries and over time), and this is also the value we use in all welfare calculations reported below. Similarly, we choose a quarterly growth rate $\mu = .005$ to match the average growth rate in aggregate real income in Mexico over the relevant sample period. For the preference parameters, we choose a quarterly discount factor of $\beta = .99$ and a degree of risk aversion of $\gamma = 1$ (log utility) for the baseline economy. These values for the preference parameters are in line with the values used in the macroeconomic literature (Cooley, 1995). However, we also report the welfare results for a higher degree of risk aversion ($\gamma = 2$).

We conduct the following exercises. Starting from a tariff level of $\tau = .10$, roughly the average tariff level in our data set, we consider the welfare consequences of reducing the tariff level to $\tau' = .05$. Our empirical analysis in section III suggests that this tariff reduction has two effects for industries with high import penetration. First, there is a short-run effect that leads to an increase in income risk for one year (four quarters), and in this section we evaluate the welfare cost of this short-run effect. Second, there is an interaction effect, and we report the welfare cost corresponding to this effect as follows. We compute the welfare cost of a short-run increase in income risk following a real exchange rate appreciation of 10%, with the tariff level also at 10%, and then compare this welfare cost of the same exchange rate

Table 8.—Welfare Effects

	Change in σ_ε^2 ($\bar{\sigma}_\varepsilon^2 = 0.008$)	Welfare Change ($\gamma = 1$)	Welfare Change ($\gamma = 2$)
Trade reform			
τ reduced by 5%	0.005	0.98	1.96
	(0.002)	(0.39)	(0.79)
Macroeconomic factors (τ level = 10%)			
g lower by 5%	0.002	0.39	0.78
	(0.001)	(0.20)	(0.40)
e appreciation by 10%	0.003	0.59	1.18
	(0.001)	(0.20)	(0.39)
Macroeconomic factors (τ level = 5%)			
g lower by 5%	0.005	0.98	1.95
	(0.001)	(0.29)	(0.59)
e appreciation by 10%	0.006	1.18	2.36
	(0.002)	(0.40)	(0.80)

Note: Welfare changes are measured in compensating variation terms and denote the change in lifetime consumption necessary to compensate agents for the short-term (one-year) increases in σ_ε^2 (relative to its sample mean of 0.008) that result under the exercises being considered. γ denotes the coefficient of relative risk aversion. Standard errors for the estimated welfare effects were obtained by simulation.

appreciation with the cost that obtains when the prevalent tariff level were 5% instead. Finally, we consider the welfare effect of a change in income risk due to a downturn in the economy, with the growth rate of GDP lowered by 5%, and again see how this is altered if the tariff level were lowered by 5%.

Table 8 reports the effects of a trade reform that lowers tariff rates from 10% to 5% for industries with high levels of import penetration. As indicated in table 8, this would raise σ_ε^2 for one year following the reform from a mean level of 0.008 to 0.013 (here we use our regression results from equation (10′) reported in table 5). The corresponding welfare cost of this change is calculated to be 0.98% of permanent consumption if the coefficient of risk aversion is $\gamma = 1$, and this cost increases to 1.96% of lifetime consumption if we choose $\gamma = 2$ instead.[30] Now consider the indirect effects of trade policy as measured by the interaction terms in equation (10′). As noted above, an exchange rate appreciation of 10% raises σ_ε^2 for a year from 0.008 to 0.011 if the tariff level is 10%. This translates into a welfare cost of 0.59% of lifetime consumption if $\gamma = 1$ and 1.18% if $\gamma = 2$. If the tariff rate were lowered to 5%, however, σ_ε^2 rises to 0.014, and the corresponding welfare costs are 1.18% and 2.36% of lifetime consumption, respectively. Finally, if the tariff rate is 10%, a cyclical downturn in the economy (a drop in g by 5%) raises σ_ε^2 for a year from 0.008 to 0.010, and the corresponding welfare cost is calculated to be 0.39% of lifetime consumption if $\gamma = 1$ and 0.78% with $\gamma = 2$. In contrast, if the tariff rate were lowered to 5%, σ_ε^2 rises to 0.013 instead, and the corresponding welfare costs are 0.98% and 1.96% of lifetime consumption, respectively.

[30] Although the welfare formula (15) is nonlinear in γ, this nonlinearity is not very pronounced for moderate degrees of risk aversion. For example, if $\gamma = 4$, then the welfare cost of this short-run change in σ_ε^2 is 4.16% of lifetime consumption. Notice also that the results reported in table 8 assume $n = 4$, since we use quarterly risk and preference parameters and the increase in income risk lasts four quarters (one year).

Thus, our calculations suggest that both the short-run direct effects of tariff reforms and the indirect effects of the level of the tariff in amplifying the effects of macroeconomic shocks are economically significant.

The limited time-series dimension of our income data might lead us to overestimate the amount of permanent labor income risk Mexican workers face. Consequently, the welfare results reported in table 8 might overstate the true cost of trade liberalization. We therefore also calculate the welfare effects for an economy in which the average income risk, σ_ε^2, and all changes in income risk, $\Delta\sigma_\varepsilon^2$, are scaled down by a factor of 0.7. The factor 0.7 is derived from the fact that the estimate of income risk, σ_ε^2, obtained by previous authors (Carroll & Samwick, 1997; Gottschalk & Moffitt, 1994; Meghir & Pistaferri, 2004; Storesletten et al., 2004) using income data with a very long panel dimension is roughly 70% of our estimate of income risk using a much shorter panel dimension (see our discussion in section IID for details). Using the scaled-down values for income risk and income risk changes, we find the following welfare cost of a 5% tariff reduction for a degree of risk aversion of $\gamma = 1$. First, the one-year increase in income risk immediately following the tariff reduction is equivalent to a decrease in lifetime consumption by .68%. Second, a 5% decline in GDP growth leads to an increase in income risk that is equivalent to a loss of lifetime consumption of .27% before the tariff reduction, and this loss increases to .68% after the tariff reduction (that is, the difference is .41%). Thus, although the welfare costs of trade liberalization are somewhat smaller than for the baseline case, they are still quite substantial.

VI. Conclusions

This paper studies empirically the relationship between trade policy and individual income risk. The analysis proceeds in three steps. First, longitudinal data on Mexican workers are used to estimate individual income risk in various manufacturing sectors. Second, the variation in income risk and trade barriers—both over time and across sectors—is used to arrive at estimates of the relationship between trade policy and individual income risk. Finally, with these estimates of the relationship between trade policy and income risk, we use a simple dynamic general equilibrium model with incomplete markets to obtain estimates of the corresponding welfare effects.

Our findings can be summarized as follows. First, for industries with high levels of import penetration, trade policy changes have a significant short-run effect on income risk. Second, the effect of the tariff level on income risk is insignificant. Third, while the tariff level has an insignificant mean effect, it nevertheless changes the degree to which macroeconomic shocks affect income risk. Finally, the welfare costs associated with the estimated increases in income risk are substantial.

The welfare results reported in this paper have to be interpreted with caution, keeping in mind several limitations of our analysis. More specifically, we focus exclusively on the link between trade policy and individual income risk and do not study how trade reform affects the mean of income growth. Second, our welfare calculations do not allow the possibility that an increase in income risk might lead to a simultaneous rise in insurance opportunities (endogenous market incompleteness). Third, we follow a long-standing tradition in economics and measure risk by the variance (second moment) of the relevant distribution, which is justified if (as assumed in this paper) the economic variables of interest are (log)-normally distributed. Finally, the Mexican household survey we use to implement our general approach is a rotating panel that follows individual workers for five quarters over time, which means that the panel dimension of our income data is somewhat limited. In short, the welfare results presented here do not show that trade liberalization is necessarily costly overall, but they do provide strong evidence that any comprehensive welfare analysis of trade liberalization ought to take into account the cost of increased labor market risk.

REFERENCES

Alesina, A., and A. Drazen, "Why Are Stabilizations Delayed?" *American Economic Review* 81:5 (1991), 1170–1188.

Altonji, J., and L. Segal, "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business and Economic Statistics* 14 (1996), 353–366.

Antras, P., "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics* 118 (2003), 1375–1418.

Carroll, C., and A. Samwick, "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40 (1997), 41–72.

Constantinides, G., and D. Duffie, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104 (1996), 219–240.

Cooley, T., *Frontiers of Business Cycle Research* (Princeton, NJ: Princeton University Press, 1995).

Davis, S., J. Haltiwanger, and S. Schuh, *Job Creation and Destruction* (Cambridge, MA: MIT Press, 1996).

Di Giovanni, J., and A. Levchenko, "Trade Openness and Volatility," this REVIEW 91 (2007), 558–585.

Eaton, J., and G. Grossman, "Tariffs as Insurance: Optimal Commercial Policy When Domestic Markets Are Incomplete," *Canadian Journal of Economics* 18 (1985), 258–272.

Esquivel, G., and A. Tornell, "The Political Economy of Mexico's Entry into NAFTA," NBER working paper no. 5322 (1995).

Feenstra, R., and G. Hanson, "Foreign Investment, Outsourcing and Relative Wages," in R. Feenstra, G. Grossman, and D. Irwin (Eds.), *The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati* (Cambridge, MA: MIT Press, 1996).

——— "Global Production and Inequality: A Survey of Trade and Wages," in E. K. Choi and J. Harrigan (Eds.), *Handbook of International Trade* (London: Basil Blackwell, 2002).

Fernandez, R., and D. Rodrik, "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty," *American Economic Review* 81 (1991), 1146–1155.

Friedman, M., *A Theory of the Consumption Function* (Princeton, NJ: Princeton University Press, 1957).

Gawande, K., and P. Krishna, "The Political Economy of Trade Policy: Empirical Approaches," in E. Choi and J. Harrigan (Eds.), *Handbook of International Trade* (London: Basil Blackwell, 2003).

Goldberg, P., and N. Pavcnik, "Distributional Effects of Trade Liberalization in Developing Countries," *Journal of Economic Literature* 45 (2007), 39–82.

Gottschalk, P., and R. Moffitt, "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Paper on Economic Activity* 1 (1994), 221–273.

Hanson, G., "What Has Happened to Wages in Mexico since NAFTA? Implications for Hemispheric Free Trade," National Bureau of Economic Research working paper, no. 9563 (2003).

Helpman, E., and A. Razin, "Efficient Production under Uncertainty," *American Economic Review* 70 (1980), 716–731.

Helpman, E., and G. Grossman, "Integration versus Outsourcing in Industry Equilibrium," *Quarterly Journal of Economics* 117 (2002), 85–120.

Jacobson, L., R. LaLonde, and D. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review* 83 (1993), 685–709.

Krebs, T., "Human Capital Risk and Economic Growth," *Quarterly Journal of Economics* 11 (2003), 709–744.

——— "Testable Implications of Consumption-Based Asset Pricing Models with Incomplete Markets," *Journal of Mathematical Economics* 40 (2004), 191–206.

Krueger, Alan, and Lawrence Summers, "Efficiency Wages and the Inter-Industry Wage Structure," *Econometrica* 65 (1988), 259–293.

Lawrence, R., and M. Slaughter, "International Trade and American Wages in the 1980s: Giant Sucking Sound or Small Hiccup?" (pp. 161–211), in Martin Neil Baily and Clifford Winston (Eds.), *Brookings Papers on Economic Activity: Microeconomics 2* (Washington, DC: Brookings, 1993).

Levine, D., and W. Zame, "Does Market Incompleteness Matter?" *Econometrica* 71 (2002), 1695–1725.

Lucas, R., "Macroeconomic Priorities," *American Economic Review* 93 (2003), 1–14.

Meghir, C., and L. Pistaferri, "Income Variance Dynamics and Heterogeneity," *Econometrica* 72 (2004), 1–32.

Melitz, M., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71 (2003), 1695–1726.

Newberry, D., and J. Stiglitz, "Pareto-Inferior Trade," *Review of Economic Studies* 51 (1984), 1–12.

Pavcnik, N., "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," *Review of Economic Studies* 69 (2002), 245–276.

Rodrik, D., *Has Globalization Gone Too Far?* (Washington, DC: Institute for International Economics, 1997).

——— "Why Do More Open Economies Have Bigger Governments?" *Journal of Political Economy* 106 (1998), 997–1032.

Storesletten, K., C. Telmer, and A. Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy* 112 (2004), 695–717.

Traca, D., "Globalization, Wage Volatility, and the Welfare of Workers," *Review of International Economics* 13 (2005), 237–249.