Appendix A
The Theory of Policy Cleavages

This appendix expands on the section "Theory of Policy Preferences" in chapter 3, where we summarized some standard economic models of policy cleavages as a prelude to our empirical analysis.

Trade Policy and Income

The key empirical implication of Stolper-Samuelson intuition is that the wage effects of trade-induced changes in product prices tend to depend on their sector bias. Any change that initially increases profits in a particular sector tends to raise the economywide wage for factor(s) employed relatively intensively in that sector. See Deardorff (1994) for a comprehensive survey of the many different theoretical statements of this theorem.

The Stolper-Samuelson discussion in chapter 3 presumed that all industries are tradable. However, in the United States today, more people work either in retail trade or in government at all levels than in all of manufacturing. Surely the logic of the Stolper-Samuelson model does not hold for countries like the United States, in which nearly 80 percent of the labor force is employed in largely nontradable sectors? Actually, it can very easily. With nontraded sectors, so long as the number of tradable goods is at least equal to the number of primary factors, national wages are still determined by the zero-profit conditions of the tradable sectors only. Nontraded product prices are endogenously determined by non-

Figure A.1 World wage pool: Manufacturing wages and population, 1994

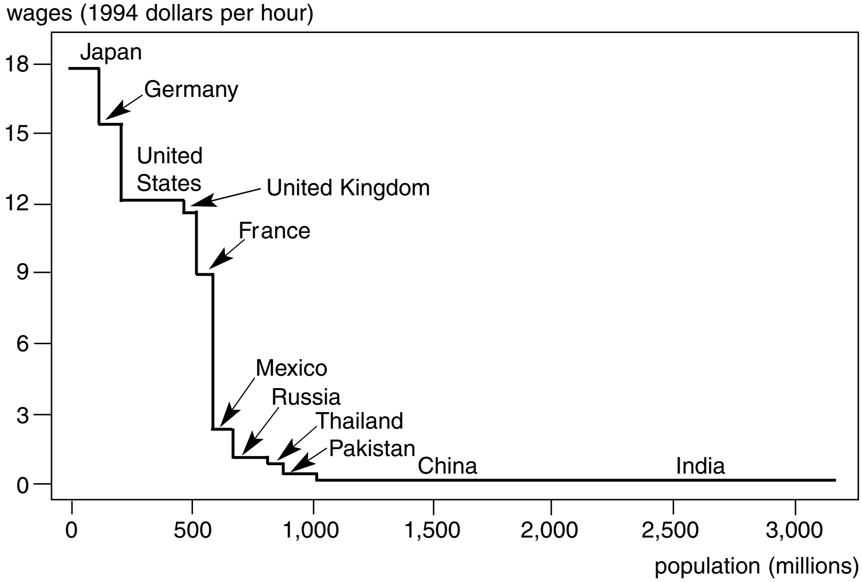

Source: Leamer (1998, updated).

traded production technology and by national wages, where these national wages are set by the product prices and technology levels in the tradable industries only.

What empirical evidence is there that the United States is well endowed with skilled labor relative to the rest of the world? Figure A.1 provides an illustration. The diagram shows the world wage pool in 1994 by plotting average manufacturing wages against population for many of the world's nations. The vertical axis shows each country's wage level in real 1994 US dollars; the horizontal axis shows each country's total population. Figure A.1 clearly shows dramatic differences: the world has some relatively small, high-wage countries like the United States along with many relatively large, low-wage countries (such as China). These cross-country differences in wages can result from many forces, but Leamer (1984) and others have shown that one important force is cross-country differences in relative labor supplies.

As for the implication that the United States protects its less-skill-intensive industries, figure A.2 shows evidence of this. For 1988, figure A.2 plots the level of industry tariff rates (measured as customs duties collected as a share of imports f.o.b.) on industry skill intensity measured as the ratio of nonproduction to production employment. The clear message of figure A.2 is that US tariffs and transportation costs were highest

Figure A.2 US tariffs are higher in less-skill-intensive sectors

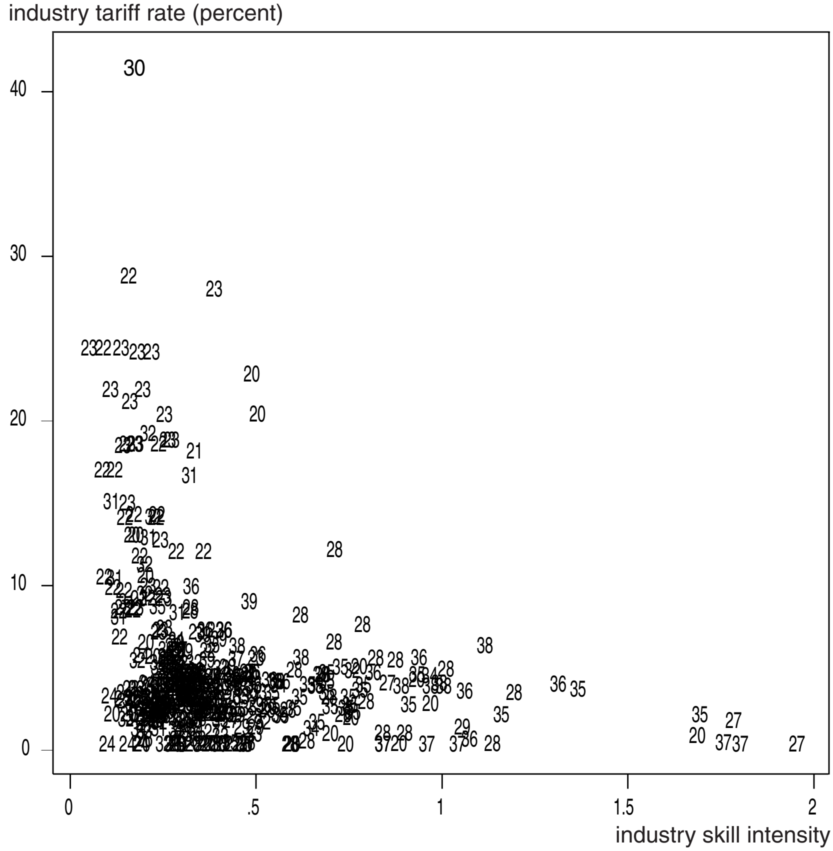

Note: Each observation is a four-digit SIC industry from the year 1988; for brevity, only the two-digit SIC code for each observation is reported. Industry skill intensity is the relative employment of nonproduction to production workers. Tariff rates are duties collected as a share of the customs value of imports. For readability, two zero-tariff skill-intensive industries are omitted: SIC 2721 (skill intensity of 4.83) and SIC 2731 (skill intensity of 3.25).

Source: Haskel and Slaughter (2000, figure 1a).

in the less-skill-intensive sectors, particularly in textiles (SIC 22), apparel (23), and footwear (31).[1]

Our analysis of trade policy preferences should not be interpreted as a direct test between the HO and RV models. A direct test would require data such as intersectoral factor movements and factor prices. There is

1. To examine the sector bias of barrier levels more formally, Haskel and Slaughter (2000) also regressed the levels of tariffs on the shares of more-skilled and less-skilled labor in total industry costs for 1974, 1979, and 1988. In every case tariffs were significantly concentrated in less-skill-intensive sectors.

nothing preventing preferences from being consistent with both models, not just one. The RV model can be characterized as a short-run version of the more long-run HO model. If there are meaningful barriers preventing worker mobility across industries, these barriers are likely to matter less over time. For example, it may be costly to lose industry-specific human capital today, but over many years this loss matters much less. In terms of the theory, Mayer (1974), Mussa (1974), and Neary (1978) compare wage changes in the two models, and Mussa (1978) formalizes how with intersectoral mobility costs an RV short run gradually becomes an HO long run. In reality, each model might be relevant over different time horizons. If individuals evaluate both short-run and long-run effects of trade liberalization, then trade policy preferences might depend on both factor type and industry of employment.[2]

Immigration Policy and Income

There is a good deal of empirical evidence on the skill mix of US immigrants in recent decades. For example, Borjas et al. (1997, 6) report that "on average, immigrants have fewer years of schooling than natives—a difference that has grown over the past two decades, as the mean years of schooling of the immigration population increased less rapidly than the mean years of schooling of natives. As a result, the immigrant contribution to the supply of skills has become increasingly concentrated in the lower educational categories." This skills gap between immigrants and natives does not address other interesting facts about the distribution of skills among immigrants. For example, Borjas et al. (1997, 7) show that the skill distribution of US immigration has been somewhat bimodal, with concentrations at both the high-skill and low-skill ends of the distribution.

To better understand immigration in the HO framework, start with two key assumptions. First, above and beyond the intersectoral factor mobility discussed earlier, there is interregional labor mobility as well: thanks to sufficient mobility of natives (and immigrants on arrival), there are no geographically segmented "local" labor markets. The second key assumption is that there are more tradable products (i.e., sectors) than primary factors of production, with products differentiated by their factor intensities. Multiple products are essential for establishing many fundamental trade theory results, such as comparative advantage.

With these assumptions, in equilibrium a country chooses (via the decentralized optimization of firms) the output mix that maximizes

2. Another way both models might accurately describe the economy is that within some time frame specificity might vary across units in the economy (such as industries or factor types). Thus, within the same time frame both the HO and RV models might apply, each to different parts of the economy. Alt et al. (1999) find support for this perspective in their study of firm lobbying behavior in Norway.

national income subject to the constraints of world product prices, national factor supplies, and national technology. This output mix consists of both the products that actually get produced and the quantities of production. In turn, this output mix helps determine the country's national factor prices. The general intuition is that the technology parameters and world price for each produced sector help determine national wages. The standard case was discussed in the previous section: when the country makes at least as many products as the number of primary factors, equilibrium wages are a function of just the world prices and technology parameters of the produced sectors. These wages do not depend on the prices and technology of the nonproduced sectors. They also do not depend directly on the level of endowments (only indirectly through the endowments' role in selecting the product mix).

Immigration's wage effects depend on the initial product mix, on the size of the immigration shock, and on whether the country's economy is large enough to have any influence on world product prices. Consider the standard case, where the initial output mix is sufficiently diversified that wages depend on just world prices and technology.

In this case, "sufficiently small" immigration is absorbed by the country changing its output mix as predicted by the Rybczynski (1955) theorem: the same products are produced, but output tends to increase in the less-skill-intensive sectors and to decrease in the more-skill-intensive sectors. How exactly does this happen? With the change in factor supplies available to hire, firms will have an incentive to produce more output of those products that employ relatively intensively the now more-abundant factors. Whether wages change depends on whether the country is big or small. If the country is small, world prices do not change and thus there are no wage effects. The cross-industry shifts in output, thanks to differences in factor intensity across industries, generate economywide shifts in factor demand that just match the economywide change in factor supplies. Leamer and Levinsohn (1995) call this insensitivity of national wages to changes in national factor supplies the factor-price-insensitivity (FPI) theorem. On the other hand, if the country is large, then wages do change via the Stolper-Samuelson process: the relative price of less-skill-intensive products declines, which lowers wages for less-skilled workers and raises wages for more-skilled workers.

With "sufficiently large" immigration shocks, national wages do change. Large enough shocks induce the country to make a different set of products, which entails a different set of world prices and technology parameters and thus different wages. This absorption of large shocks via changes in both output mix and wages holds whether the country is large or small: in either case wage inequality rises.

It is worth emphasizing the somewhat counterintuitive idea of FPI, with which immigrants might generate zero wage pressures thanks to changes in output mix. Do these changes actually happen in the real world? Little empirical research has examined this (see Hanson and Slaughter 2000). Much of the empirical work on the labor-market impacts of immigration have, implicitly or explicitly, used the factor-proportions or area-studies frameworks. Studies using the former include those of Borjas et al. (1996; 1997),which calculate immigration-induced shifts in national factor proportions and then infer the resulting national wage changes. A large number of studies have used the area-studies framework, including Card (1990), Altonji and Card (1991), and LaLonde and Topel (1991). These studies generally test for correlations between immigrant flows into local labor markets and local native wages. Again, the key difference between these two frameworks is the geographic extent of immigrants' wage effects. In a national labor market, immigrants' wage pressures spread beyond gateway communities. Natives can leave gateway communities when immigrants arrive; immigrants can move on to other communities; or natives can choose not to enter gateway communities as planned preimmigration. In cases between these two extremes, immigrants affect wages everywhere, but to a greater extent in gateway labor markets.

Other Considerations: Trade Policy and Asset Values

Many assets do fit easily into the HO framework. Some assets are currently employed by firms as factors of production—for example, machine tools and office buildings. These assets earn rates of return that are determined just as wages are, as outlined above. In turn, these rates of return are an important determinant of asset prices.[3] Another kind of asset that fits into standard trade models is currently produced goods such as automobiles. Their domestic prices are set just like those of other nonasset products, as some combination of world prices and domestic trade barriers.

The assumption that construction adds very little to housing stocks is borne out in the data. The US Census Bureau estimates that on 1 July 1996, the total US housing stock was 110 million housing units (US Census Bureau, Estimates of Housing Units, Households, Households, by Age of Householder, and Persons per Household [August 1997], http://www.census.gov/population/www/estimates/housing.html [October 1997]). In 1996, approximately 1.3 million new homes were constructed nationwide. According to averages from the 1980s, approximately 0.3 million existing homes would have become uninhabitable that year due to demolition, disasters, and other causes. Thus, the net construction rate

---

3. There is a well-developed literature analyzing how these productive assets accumulate over time in open trading economies. See, for example, the surveys of Findlay (1984) and Smith (1984).

in 1996 was about 1 million new homes—0.9 percent of the existing stock. Also, the Census Bureau estimates that nationwide in 1996, an average of 8.33 months passed from the time a residential construction permit was issued to the time construction was completed.

For the average US household in 1990, the gross value of the primary residence accounted for nearly 90 percent of total household assets (Caplin et al. 1997). For the median US homeowner in all age groups in 1986, housing equity accounted for more than half of his or her total wealth (Skinner 1994). So even though the NES survey data cover only one asset, it is the single largest asset for a significant share of the population.

Research in the regional-economics literature has documented an empirical link between local industry mix and local housing prices. Case and Mayer (1996) found that in the Boston area during the 1980s, average house prices rose less in housing jurisdictions with a larger share of residents employed in manufacturing in 1980. They hypothesize that this empirical link reflects "displaced manufacturing workers. . .reducing their demand for housing" (391).

Note that the link between trade and asset values operates independently of trade's effect on labor incomes. People's economic welfare depends on both income and asset holdings, and freer trade might affect these two channels differently. Consider a more-skilled homeowner in Gary, Indiana, a city with production very concentrated in a sector with comparative disadvantage, steel. This person might support freer trade through the income channel but oppose it through the asset channel. We distinguish between these two links in the data analysis.

Appendix B
Data Description

The survey data presented in chapter 2 are from the Public Opinion Databank at the Roper Center for Public Opinion Research. The original sources for each survey are listed with each question.

The data used in the analysis in chapter 3 were obtained from the following sources.

- National Bureau of Economic Research, Manufacturing Industry Database, http://www.nber.org/nberces.
- Sapiro, Virginia, Steven J. Rosenstone, Warren E. Miller, and the National Election Studies. 1998. American National Election Studies, 1948-97 [CD-ROM], ed. ICPSR. Ann Arbor, MI: Inter-University Consortium for Political and Social Research [producer and distributor].
- United States Bureau of the Census. 1992. Census of Manufactures, 1992. Washington: US Government Printing Office.
- United States Bureau of the Census. 1994. State and Metropolitan Area Data Book, 1994. Washington: US Government Printing Office.
- United States Department of Commerce, Bureau of Economic Analysis, Survey of Current Business and Trade and Employment, various years. Washington: US Government Printing Office.
- United States Department of Labor, Bureau of Labor Statistics. 1992, 1994, and 1996. Unpublished tabulations from the current population survey, table A-26, 1992, 1994, and 1996 annual averages.

- United States International Trade Commission. 1997. Unpublished data file on US tariffs collected in 1992.

Sector Net Export Share

To construct this variable, we obtained Robert C. Feenstra's data from the NBER on 1992 manufacturing exports, imports, and value of shipments at the four-digit Standard Industrial Classification (SIC) (revision 2) level. To cover all truly tradable sectors, we obtained similar data for agriculture and tradable services from various BEA sources. All these data were concorded to the 1980 Census Industrial Classification (CIC) industries, and then for each industry we calculated *Sector Net Export Share* as exports minus imports divided by value of shipments. For all nontradable sectors we set this variable equal to zero.

Sector Tariff

To construct this variable, we obtained data from the ITC on 1992 tariff duties collected and customs-value imports at the four-digit SIC (revision 3) level.[1] We concorded these tariff and import values to the 1980 CIC industries, and then for each industry we calculated an effective tariff rate by dividing tariffs value by imports value.

High-Immigration MSA

First, we defined local labor markets two ways: by a combination of metropolitan statistical areas (MSAs) and counties, and by states. In our MSA/county definition, each MSA (with all its constituent cities and counties) is a separate labor market; for individuals living outside an MSA, the labor market is the county of residence. Following the extensive use of MSAs in area-analysis studies and Bartel's (1989) finding that immigrants arrive mostly into cities, we prefer the MSA/county definition but tried states for robustness. Second, for each definition of local labor markets we tried three different definitions of a high-immigration labor market: 5 percent, 10 percent, and 20 percent shares of immigrants in the local population. These immigration and labor force data are from the 1990 census as reported by the US Census Bureau (*State and Metropolitan Area Data Book,* 1994). Altogether, for each of our six primary measures we construct a dichotomous variable, *High-Immigration MSA,* equal to 1 for residents in high-immigration labor markets. In the tables we report

1. We thank Michael Ferrantino at the US International Trade Commission for helping us obtain these data.

results for our preferred measure, the MSA/county 10 percent definition. Alternative measures are discussed in the robustness checks reviewed in chapter 3.[2]

Noneconomic Variables

Gender is a dichotomous variable equal to 1 for females. *Age* is a continuous variable (for the trade analysis it is divided into three dichotomous measures—Age 18-29, Age 30-44, and Age 45-59—and a residual category). *Race* is a dichotomous variable equal to 1 if the respondent is African-American. For ethnicity we constructed the dichotomous variable *Hispanic*, equal to 1 if the individual self-identifies with a Hispanic ethnic group. *Immigrant* is a dichotomous variable equal to 1 if the respondent or his or her parents were immigrants into the United States. *Party Identification* is a categorical variable ranging from 1 for "strong Democrat" to 7 for "strong Republican." Finally, *Ideology* is a categorical variable ranging from 1 for "extremely liberal" to 7 for "extremely conservative." In addition to these variables, for certain specifications in the robustness checks we included additional regressors, which are discussed in the text.

2. In 1990, immigrants accounted for 7.9 percent of the US population. Thus, our 5 percent cutoff might seem too low, but for the sake of completeness we tried it anyway. Also, the 1990 census MSA data are organized by 1990 MSA definitions, but the 1992 NES survey locates individuals by 1980 MSA definitions. Using unpublished information on 1980-90 MSA changes obtained from Census Bureau officials, we corrected discrepancies as best we could.

Appendix C
Multiple Imputation Methodology

The data constructed for analysis in this study are not fully observed. The sources of "missingness" range from survey respondents' refusal to answer particular questions in some surveys to the suppression of data by national agencies to avoid identifying particular commercial establishments. Incomplete data, whatever the source, can create a number of serious problems for making valid statistical inferences.

For example, the most common approach in the social sciences to multivariate analyses of incomplete data is to drop cases that have any missing data and analyze only the complete cases. This standard approach, known as listwise deletion, can create two major problems. One is the inefficiency of throwing away information relevant to the statistical inferences being made. The other is that inferences from listwise deletion estimation can be biased if the observed data differ systematically from the unobserved data. In this study, since "missingness" in particular variables ranged from none to nearly 75 percent, inefficiency was clearly a concern. Moreover, *ex ante* there was little reason to believe that data were missing completely at random, so employing a listwise deletion approach risked bias as well.

Alternatives to listwise deletion for dealing with missing data have been developed in recent years. The most general and extensively researched approach is multiple imputation (King et al. 2001; Schafer 1997; Little and Rubin 1987; Rubin 1987). Multiple imputation makes a much weaker assumption than listwise deletion about the process generating the missing data. Rather than assuming that the unobserved data are missing completely at random, multiple imputation is consistent and gives correct uncertainty estimates if the data are missing randomly conditional on the data included in the imputation procedures.[1] Moreover, multiple imputation offers important advantages over ad hoc procedures for dealing with missing data. Imputing sample averages on a variable-by-variable basis biases estimates and standard errors toward zero. Imputing predicted values from regression models tends to inflate sample correlations and thus bias estimates away from zero. Given all these advantages of multiple imputation, we use this estimation methodology.

The remainder of this appendix provides the specific details of the imputation procedures used in the trade and immigration analyses.

Trade

After the variables described in chapter 3 were constructed and combined into individual-level data sets for each cross-sectional survey, there was a significant amount of missing data. In the NES survey some individuals did not report either occupation, education, or industry of employment, which prevented the construction of some of the factor-income trade-exposure variables for these people. The most serious missing-data problem arose from the homeowners' exposure variables that were constructed for the asset analysis we presented in the section Asset Ownership and Trade Policy Preferences. In fact, the structure of the missing-data problems differed between the 1992 data, for which the asset analysis was conducted, and the 1996 data, for which it was not. Consequently, we used slightly different procedures for the imputations in each year and describe them separately.

For 1992, the first step in our multiple imputation procedures was to impute missing observations for *County Exposure 1* and *County Exposure*

1. The multiple imputation procedures used in this study actually require that two conditions be met. First, as discussed in the text, the probability that a data cell is missing may depend on observed data included in the imputation model but must be independent of unobserved data. In the imputation literature, this assumption is known as Missing at Random (MAR). Note that this assumption is weaker than assuming that the data are Missing Completely at Random (MCAR), which means that the probability that a data cell is missing does not depend on any data, whether observed or not. Further, the analyst can make the MAR assumption more reasonable by including a large number of variables in the imputation model. For example, even if less-skilled respondents in the trade opinion analysis are less likely to report their occupations, the MAR assumption is not violated if the other observed data correlated with skill account for this aspect of the missingness mechanism. The second condition is that the parameters describing the data are distinct from those describing the missingness mechanism in the data. Schafer contends that in many situations similar to the analyses in this study distinctness is a reasonable assumption, as knowing the data parameters provides little information about the parameters describing the patterns of missingness in the data set (1997, 11). If the missingness problem meets these two conditions, it is called ignorable and the imputation methods used in this study are appropriate. These assumptions have been shown to be reasonable in many studies similar to the one here that include rich data sets with many covariates to include in the imputation model.

2. Imputations were based on dozens of variables selected for their sample correlation with the missing variables. For example, for imputation of *County Exposure 1* and *County Exposure 2*, county employment in textiles was used because it had one of the highest sample correlations with the county-exposure variables. Altogether, 10 complete county data sets were imputed.

The exact algorithm used for these imputations is a data augmentation method known by the acronym IP because it involves two key steps: the imputation step and the posterior step. The goal of the imputation procedure is to estimate a set of parameters (means and variance/covariances of all the variables) that can be used to create the 10 imputed data sets. In this application, it is assumed that the data have a joint multivariate normal distribution. Consequently, IP employs an iterative sampling scheme in which in the first step imputations are drawn from the multivariate normal conditional predictive distribution of the missing data. This distribution depends on the observed data and the assumed or current value of the complete data parameters. In the second step, a new value of the complete data parameters is drawn from its posterior distribution, which is conditioned on the observed data and the current values of the imputations for the missing data. This posterior step is a simulation from the normal inverted-Wishart distribution. Repeating this iterative sampling scheme produces stochastic subsequences that converge on the stationary predictive distribution for the missing values and the stationary posterior distribution of the complete data parameters.[2]

For the county data set we ran 5,000 iterations of IP and then ran 1,000 more creating an imputed data set every 100 iterations of these last 1,000. The preliminary iterations ensure that sequences have converged to their stationary distributions. After creating the 10 complete county data sets, we merged the NES survey data (including the constructed skill measures and industry measures) with *County Exposure 1* and *County Exposure 2*. The resulting 10 data sets still had substantial amounts of missing individual-level data, however. Consequently, for each of these 10 data sets we ran separate iterations of IP in order to impute values for the missing survey data. We found that 2,100 preliminary iterations were more than sufficient for these data sets. An imputation was saved on the last iteration of each of the 10 cases to create the 10 final data sets with no missing data at all. Each of these final data sets for 1992 contains 1,736 observations, equal to the actual number of individuals in the NES survey either supporting or opposing more trade restrictions.[3] Also, each data set contains

2. See section 3.4 in Schafer (1997) for a complete description of IP. Also, our methodology assumes that all variables are normally distributed. To make the data fit this assumption more closely, we redefined each county-level variable to equal the natural log of the variable plus one.

3. All the main results reported are qualitatively the same for the case where imputations are also made by treating as missing data the fact that some respondents did not express

the exact same nonimputed information (i.e., all observations for the variable *Trade Opinion* plus the nonimputed observations for all the trade-exposure variables). They differ only in their imputed values for missing data.

For 1996, the first step in the multiple imputation procedure was to create imputations in the missing-data cells for all the individual-level variables (since no county-level analysis was conducted for these data, it was not necessary to break down the imputation procedures to two levels of aggregation). We based the imputations for the 1996 data on 26 variables. These variables included all those used in the analysis as well as additional information that would be helpful in predicting the missing data. Altogether we imputed 10 complete individual-level data sets. The final data sets contain 846 completed observations, equal to the actual number of individuals in the NES survey either supporting or opposing more trade restrictions. The imputation model was multivariate normal with a slight ridge prior. The algorithm used to implement this model is known by the acronym EMis, because to generate imputations it combines a well-known expectation maximization missing-data algorithm with a round of importance sampling. King et al. (2001) provide a complete explanation of the use of this algorithm for missing-data problems.

The second step in the multiple imputation analysis was to run various logit models separately on each of the 10 final data sets for each year, and the final step was to combine the 10 sets of estimation results for each specification to obtain a single set of estimated parameter means and variances. The single set of estimated means is simply the arithmetic average of the 10 different estimation results. The single set of estimated variances consists of two parts. The "within" component is simply the arithmetic average of the 10 estimated variances. This accounts for the ordinary within-sample variation. The "between" component is the variance of the estimated parameter means among the imputed data sets. See King et al. (2001) and Schafer (1997) for a complete description of the final multiple-imputation step.

Immigration

After the variables described in chapter 3 were constructed and combined into individual-level data sets for each cross-sectional survey, there was a

a trade policy opinion. For this analysis the multiple imputation procedures created 10 data sets of 2,485 observations, equal to the total number of respondents in the NES survey. In addition, all the main results are qualitatively the same using the listwise deletion method for missing data. Finally, we explicitly modeled whether or not the respondent offered an opinion simultaneously with the determinants of the opinions given using a Heckman probit selection model. Our skill and industry findings were robust to this specification, and the analysis revealed significantly lower levels of political awareness among those who did not answer the question.

significant amount of missing data. Across the range of models estimated, when we simply dropped observations with any missing data we generally lost 25 to 45 percent of the total observations.

The first step in the multiple imputation procedure was to create imputations in the missing-data cells for all the variables. We based the imputations for the 1992, 1994, and 1996 data on 36, 28, and 26 variables, respectively, selected from the NES surveys. These variables included all those used in the analysis as well as additional information from each survey that would be helpful in predicting the missing data. Altogether we imputed 10 complete individual-level data sets for each year. The final data sets contain completed observations equal to the actual number of individuals in each NES survey. The imputation model was multivariate normal with a slight ridge prior. The algorithm used to implement this model is known by the acronym EMis, because to generate imputations it combines a well-known expectation maximization missing data algorithm with a round of importance sampling. King et al. (2001) provide a complete explanation of the use of this algorithm for missing data problems.

The second step in the multiple imputation analysis was to run various ordered probit models separately on each of the 10 final data sets for each survey year. The last multiple imputation step was to combine the 10 sets of estimation results for each specification.

Appendix D
Further Evidence of the Skills-Preferences Cleavage

The skills-preferences cleavage described in chapter 3 is corroborated by surveys other than the NES surveys. Some surveys in the data we present in chapter 2 report responses broken down by skill measures such as educational attainment or household income. Survey responses grouped solely by skills must be interpreted cautiously: these groups alone do not control for other possible sources of preferences cleavages that might be correlated with skills, and sample differences across skill groupings need not be statistically significant. That said, these surveys support our findings on the skills-preferences cleavage originally presented in the multivariate analysis of Scheve and Slaughter (2001a, b).

Here are three examples. All repeat questions of the "is trade good or bad" variety, like those presented in chapter 2. The pattern is clear: among respondents with less education, opposition to trade seems much stronger.

Question: Do you think it should be the policy of the country to restrict foreign imports in order to protect jobs and domestic industries, or do you think there should be no restrictions on the sale of foreign products in order to permit the widest choice and the lowest prices for the consumer?

"Restrict" responses:

Those with advanced degree: 53%
Those with college degree: 61%
Those with some college: 68%
Those with high school degree: 77%
Those with less than high school: 73%

Source: Los Angeles Times, February 1992

Question:	What do you think foreign trade means for America? Do you see foreign trade more as an opportunity for economic growth through increased US exports, or a threat to the economy from foreign imports?
"Opportunity" responses:	
	Those with advanced degree: 73% Those with college degree: 75% Those with some college: 57% Those with high school degree: 42% Those with less than high school: 37%

Source: CNN/*USA Today*, November 1994

Question:	As you may know, with freer trade, jobs are often lost due to imports from other countries, while new jobs are created when the US exports more products to other countries. I'd like you to imagine in one industry some jobs are lost because of foreign competition, while in a different industry an equal number are created, but these new jobs pay higher wages. Which of the following statements do you agree with the most? A: Even if the new jobs that come from freer trade pay higher wages, overall it is not worth all the disruption of people losing their jobs. B: It is better to have the higher-paying jobs, and the people who lost their jobs can eventually find new ones.
"A" responses:	Those with college degree or higher: 33% Those with high school degree: 65% Those with less than high school: 66%

Source: Program on International Policy Attitudes, October 1999