ESSAYS ON ECONOMIC POLICY:
INCOME INEQUALITY AND HEALTH INSURANCE

by

ERIC DOVIAK

A dissertation submitted to the Graduate Faculty in Economics in partial fulfilment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2010

This manuscript has been read and accepted for the
Graduate Faculty in Economics in satisfaction of the
dissertation requirement for the degree of Doctor of Philosophy.

Professor John Devereux

Date

Chair of Examining Commitee

Professor Merih Uctum

Date

Executive Officer

Professor John Devereux

Professor Michael Grossman

Professor Merih Uctum

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

ESSAYS ON ECONOMIC POLICY:
INCOME INEQUALITY AND HEALTH INSURANCE

by

Eric Doviak

Advisor: Professor John Devereux

This dissertation contains economic analyses of two critical social issues facing the United States at the dawn of the 21st century: income inequality and the affordability of health insurance.

The chapter on income inequality uses the Solow Model of economic growth to model the evolution of inequality over time. In steady state, differences in household saving rates generate differences in household capital income. Households that save more accumulate more capital and have higher steady-state income. Tax policy affects the distribution of income through its influence on household saving rates. Increasing the tax rate on labor income causes a greater percentage decrease in the steady-state saving rates of relatively low savers, thus increasing pre-tax income inequality. Conversely, increasing the tax rate on capital income reduces pre-tax income inequality because it causes a greater percentage decrease in the steady-state saving rates of relatively high savers. Empirical tests of the model using data from the March Current Population Survey and NBER's TAXSIM model suggest that higher taxes on wage income are associated with higher levels of income inequality. A high degree of correlation among the average marginal tax

rates prevents us from drawing inferences about the effect that taxation of capital income has on inequality.

The chapter on health insurance examines states' efforts to make health insurance more accessible and affordable to small employers by restricting insurers' ability to set premium rates on the basis of health status and other factors which predict a group's future medical needs. The chapter presents evidence that rating restrictions reduce health insurance coverage rates and increase market concentration in the insurance industry. From the perspective of a public policymaker however, such reforms may still be desirable if they increase the ability of less healthy individuals to obtain and afford health insurance coverage.

ACKNOWLEDGEMENTS

This dissertation would not be possible without the incredible support and encouragement of my family. At the times when pursuit of a doctoral degree seemed interminable, they encouraged me to push onward. Most importantly, I would like to thank them for their patience when I missed important events and holidays, so that I could spend time working on this dissertation.

I would also like to thank Profs. Devereux, Grossman and Uctum for reading my work and for providing valuable suggestions, comments and criticisms. Prof. Grossman also deserves special thanks for his patience, support and dedication to students. I would also like to thank Prof. Uctum for providing moral support and for consistently providing me with a good schedule of classes to teach at Brooklyn College.

I also owe a debt of gratitude to Pennsylvania State Senator Rob Wonderling and his staff for sponsoring the chapter on health insurance market regulation. In particular, I would like to thank John Basial for carefully reading and editing an early version of the chapter and for generously providing thoughts and suggestions on the chapter on income inequality.

Finally, I would like to thank the seminar participants at the Graduate Center, St. Lawrence University and the International School of Economics at Tbilisi State University for reading my work and providing many helpful suggestions.

Contents

1 Introduction **1**

2 A Contribution to the Theory of Income Inequality **9**

2.1 Literature Review 16

2.2 Using the Solow Model to Explain Income Inequality 28

2.3 Empirical Tests 41

2.4 Conclusion 63

3 The Economic Effects of Health Insurance Rating Restrictions **67**

3.1 Background 69

3.2 How Restrictions Affect the Demand for Health Insurance 70

3.2.1 Access and Affordability 71

3.2.2 How Community Rating Affects Coverage 73

3.3 How Restrictions Affect the Supply of Health Insurance 78

3.3.1 Increasing Returns to Scale 79

3.3.2 How Community Rating Affects Market Concentration 80

3.4 How Coverage and Concentration Affect Each Other 81

3.5 Empirical Evidence 82

3.5.1 Data 83

3.5.2 Dependent Variables 84

3.5.3 Explanatory Variables 85

3.5.4 Empirical Results 95

3.5.5 Discussion 97

3.6 Conclusion 102

Bibliography **104**

List of Tables

1.1 Correlation Matrix – Fictional Data from Obenchain 3

1.2 OLS Regression Results – Fictional Data from Obenchain 3

2.1 Simulation: Effect of a Tax on Capital Income 38

2.2 Simulation: Effect of a Tax on Labor Income 39

2.3 Correlation Matrix: Data in Levels 46

2.4 Correlation Matrix: Data in Log Differences 47

2.5 OLS Regression on CV of Log Income 49

2.6 Fixed Effects Regression on CV of Log Income 51

2.7 OLS Regression on Log Differenced CV of Log Income 53

2.8 OLS Regression on Average Log Income 57

2.9 Fixed Effects Regression on Average Log Income 59

2.10 OLS Regression on Log Differenced Average Log Income 61

3.1 Correlation Matrix 88

3.2 Logit Regression on Employment-Based Coverage Rate 89

3.3 Logit Regression on Market Share of Five Largest Insurers 92

List of Figures

1.1 Ridge Coefficients and Relative Mean Squared Error 5

2.1 Ridge Trace of OLS Regression on CV of Log Income 50

2.2 Ridge Trace of Fixed Effects Regression on CV of Log Income 52

2.3 Ridge Trace of OLS Regression on Log Differenced CV of Log Income 54

2.4 Ridge Trace of OLS Regression on Average Log Income 58

2.5 Ridge Trace of Fixed Effects Regression on Average Log Income 60

2.6 Ridge Trace of OLS Regression on Log Differenced Average Log Income 62

3.1 Ridge Trace of Logit Regression on Employment-Based Coverage Rate 91

3.2 Ridge Trace of Logit Regression on Market Share of Five Largest Insurers 94

Chapter 1

Introduction

This dissertation presents two essays on current public policy issues in the United States: income inequality and the affordability of health insurance. The essays are mutually independent, but empirical analysis of both issues – and public policy in general – requires a common regression technique: ridge regression.

This chapter provides an overview of ridge regression and argues that researchers should consider using it when working with tax and regulatory variables because those variables tend to be highly correlated with each other. For example, the chapter on income inequality shows that states with relatively high average marginal tax rates on wage income also tend to have relatively high average marginal tax rates on dividends and interest received. Similarly, states that impose relatively tight restriction on an insurers' ability to set premium rates on the basis of health status also tend to impose relatively tight restrictions on age and gender-based rating.

In an extreme case of multicollinearity, the moment matrix (i.e. $X^T X$) of the regression equation (i.e. $\widehat{\beta} = \left(X^T X\right)^{-1} X^T Y$) is exactly singular, so the rank of the matrix is less than the number of regressors. More commonly however, the moment matrix is ill-conditioned, but not of reduced rank.

When working with tax and regulatory variables, it is important to carefully examine the effect of multicollinearity because when two explanatory variables in a regression model are positively correlated, their regression coefficients will be negatively correlated and one of the OLS coefficients might have the "wrong" sign.

A strict interpretation of the OLS coefficients would therefore lead the researcher to conclude that a particular tax or regulatory variable has no effect or the effect opposite of the true effect on the outcome of interest. Ridge regression helps diagnose such problems by tracing the paths of the coefficients as they are shrunk towards zero.

To better understand the effect of multicollinearity, we'll use a fictional dataset that Obenchain provides (2004, chap. 2) and place it in a public policy context. Suppose that the variables in his dataset represent growth rates and cigarette and alcohol tax rates by state. Suppose also that there are strong theoretical reasons to believe that raising tax rates on cigarettes and alcohol increases a state's rate of economic growth.

Given the correlations in Table 1.1 and our theoretical reasons for believing that cigarette and alcohol taxation accelerate economic growth, one would expect that the coefficients on the tax rates would be positive in a regression on the growth rate. This is not the case however. Table 1.2

Table 1.1
Correlation Matrix – Fictional Data from Obenchain

Table 1.1 – Correlation Matrix

	growth rate	alcohol tax rate	cigarette tax rate
growth rate	1.00		
alcohol tax rate	0.94	1.00	
cigarette tax rate	0.79	0.87	1.00

Table 1.2
OLS Regression Results – Fictional Data from Obenchain

Table 1.2 – Ordinary Least Squares
Dependent Variable: Growth Rate

alcohol tax rate	1.03	***
standard error	0.24	
cigarette tax rate	−0.11	
standard error	0.24	
most likely Q-shape	0.0	
***p-value < 0.01		

shows that the coefficient on the alcohol tax rate is positive and statistically significant, while the coefficient on the cigarette tax rate is negative and is not statistically significant.

Because the two tax rates are positively correlated with each other, their regression coefficients are negatively correlated with each other. Moreover, the coefficient on the alcohol tax rate should be smaller and the coefficient on the cigarette tax rate should be positive.

Both principal components regression and ridge regression address multicollinearity by effectively reducing the number of dimensions along which the dependent variable is regressed. This produces biased estimates of the full-set of regression coefficients, but can reduce the mean-squared error associated with the regression coefficients and correct "wrong" signs.

The difference between ridge regression and principal components arises in the way they reduce the dimensionality. Principal components analysis explicitly reduces the number of dimensions in the regression problem by compressing several variables into one or more composite variables, whereas two-parameter ridge regression effectively reduces the number of dimensions by applying a "shrinkage factor" to each of the regressors. The two-parameter ridge shrinkage factor is:

$$\delta_j = \frac{1}{1 + k\lambda_j^{Q-1}} \tag{1.1}$$

where λ_j is the eigenvalue associated with parameter j, Q is a "shape parameter" that determines the path that regression coefficients take through likelihood space as they are shrunk toward zero. By definition:

$$0 \leq \delta_j \leq 1. \tag{1.2}$$

Figure 1.1: Ridge Coefficients and Relative Mean Squared Error

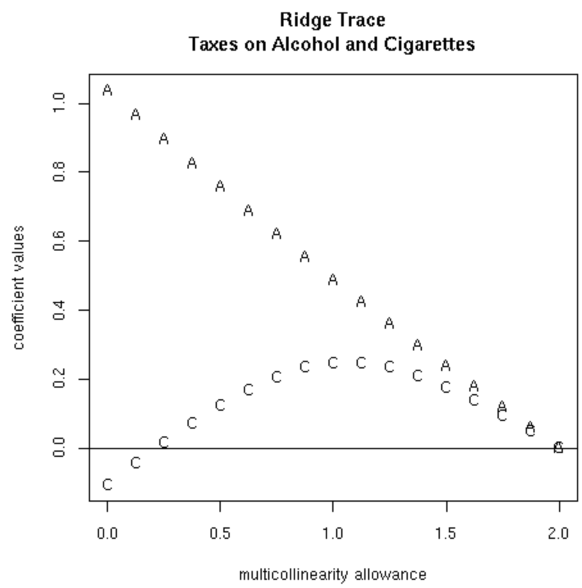

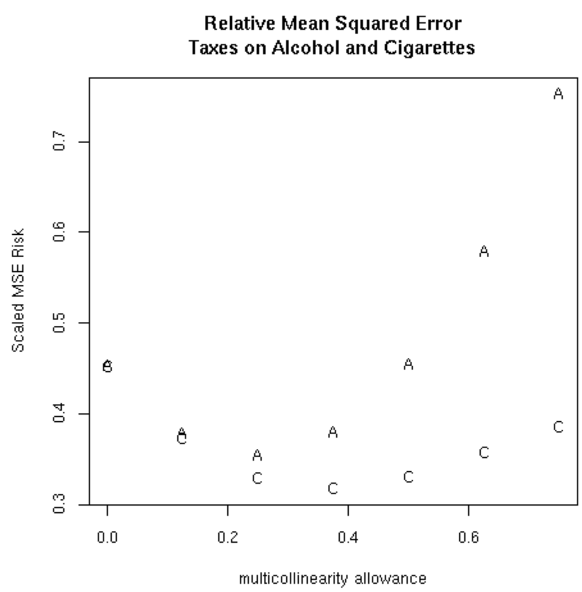

The extent of shrinkage is measured by:

$$M = R - \delta_1 - \delta_2 - \cdots - \delta_R \tag{1.3}$$

where M is the “multicollinearity allowance” and R is the number of regressors.

Principal components regression is a special case of two-parameter ridge regression because each shrinkage factor is equal to either one or zero when $Q = -\infty$ and $k > 0$.[1] In this special case, integer values are subtracted from the number of regressors and the estimated two-parameter ridge regression coefficients are identical to the principal components regression coefficients (Obenchain, 2004, chap. 3).

Obenchain (1975) shows that the Q-shape most likely to achieve overall minimum MSE risk in estimating regression coefficients is the one which maximizes:

$$COS(Q) = \frac{\sum |r_j{}^o| \lambda_j{}^{(1-Q)/2}}{\sqrt{\sum r_j{}^{o2} \sum \lambda_j{}^{(1-Q)}}} \tag{1.4}$$

where $r_j{}^o$ is the correlation between the dependent variable and the j-th column of H, the $N \times R$ matrix of standardized principal coordinates of the centered regressor matrix. (H is obtained from the singular value decomposition of the centered regressor matrix). Given the most-likely value of Q, one can – in theory – minimize MSE risk by controlling the extent of shrinkage with the k parameter. In practice however, one must keep other considerations in mind. For example, one

[1]There is an exception to this rule. In the rare case that one of the eigenvalues is exactly equal to one, one of the shrinkage factors will be equal to 0.5.

would not want to shrink the coefficients so far that they lie outside the 95 percent confidence interval.

In two-parameter ridge regression, the vector of estimated coefficients is given by:

$$\widehat{\beta} = \left(X^T X + k \cdot G\Lambda^Q G^T\right)^{-1} X^T Y \tag{1.5}$$

where X is the centered regressor matrix, Y is the centered vector of response variables, Λ^Q is an $R \times R$ diagonal matrix of the centered regressor matrix's ordered singular values raised to the Q-th power and G is the $R \times R$ matrix of principal axis direction cosines of the centered regressor matrix. (Like the H matrix, G is also obtained from the singular value decomposition. It's important to note that $GG^T = I$, so that if $Q = 0$ the regression takes the form of Hoerl and Kennard's (1970) classic ridge regression).

Note that when $k = 0$, the estimated coefficients are the ordinary least squares (OLS) estimates.

When working with shrunken estimates from ridge regression, one cannot use conventional t-statistics to test the null hypothesis that the true coefficient value is zero because the expected value of the numerator in the t-statistic is not equal to zero (Vinod, 1976). Obenchain (1977) however proves that the OLS t-statistic equals the Ridge t-statistic and argues that practitioners should center their confidence intervals on the OLS estimates.

As an extra diagnostic for our model, it is worthwhile to examine the ridge trace displays to see if the coefficients retain the same sign as they are shrunk towards zero. Some of the coefficients may change sign if the Q-shape parameter is not equal to one. When $Q = 1$ (as it is in some of the figures in this dissertation), the ridge traces form straight lines because the coefficients have stable relative magnitudes (Obenchain, 2004, chap. 3).

In other cases however, the coefficients do change sign, so before concluding that the coefficient on a tax or regulatory variable is negative and not statistically significant (like the coefficient on cigarette tax rates in the example presented here), we must examine the ridge trace to see if it changes sign as it is shrunk towards zero.

Finally, researchers often want an estimate of the effect of one variable on a dependent variable of interest, with other variables held constant. In principal, the researcher could obtain such an estimate by selecting the set of coefficients that either minimize MSE risk or reduce MSE risk without shrinking the coefficients so far that some of them lie outside of their 95 percent confidence intervals. In practice however, the large number of criteria that the researcher must consider make it difficult to select a particular set of coefficients.

Nonetheless, ridge regression provides a useful diagnostic to check for sign stability. As the tax-growth example shows, the researcher should not conclude that the effect of cigarette taxation on growth is negative after controlling for the effect of alcohol taxation.

Chapter 2

A Contribution to the Theory of Income Inequality

Much of the economic literature allocated to the topic of income distribution has asserted that there is a tradeoff between economic growth and income inequality, which implies that society must accept a larger degree of inequality in exchange for a faster rate of economic growth.

However, few of the articles in the literature contain a theoretical model that explains the source of income inequality, the persistence of inequality and the rate of growth. Some previous attempts, assert an initial uneven distribution of income (Stiglitz, 1969; Alesina and Rodrik, 1994; Perotti, 1993; Galor and Zeira, 1993). Such models don't explain income inequality, they take it as given. In a similar vein, Persson and Tabellini (1994) and Meckl and Zink (2004) assume that inequality arises from an uneven distribution of skills, while Loury (1981) models the distribution of income as a function of ability and inheritance. The trouble with such models is that they implicitly assume that income inequality is inevitable.

To explain income inequality, a model must be capable of explaining how inequality can arise from an initially egalitarian distribution of income and how equality can arise from an initially unequal distribution of income.

To the extent that income inequality is an outcome in the models developed by Alesina and Rodrik, by Persson and Tabellini and by Perotti that outcome depends on a political decision (the choice of tax rate), not the sum of the economic decisions of individuals in the economy.

Although the model developed in this chapter also assumes that the choice of tax rate affects the degree of income inequality, there are five important differences.

1. In the model developed in this chapter, pre-tax income inequality and the level of economic development (as measured by average pre-tax income) are jointly determined. Taxes, which are assumed to be exogenous, affect the degree of pre-tax income inequality and the level of average pre-tax income. By contrast, other models assume that pre-tax income inequality determines society's choice of tax rate and then examine how the chosen tax rate affects the rate of economic growth (Alesina and Rodrik, 1994; Persson and Tabellini, 1994; Perotti, 1993).

2. This chapter examines the effect that taxes on capital income, labor income and total income have on income inequality, whereas other models focus exclusively on the taxation of capital income (Alesina and Rodrik, 1994; Persson and Tabellini, 1994) or total income (Perotti, 1993). The theoretical model developed in this chapter shows that different forms of taxa-

tion imply different relationships between the degree of income inequality and the level of development.

3. In the model developed in this chapter, income inequality is not inevitable. Instead it assumes that a household's steady-state income depends on its own saving decisions, not an initial factor endowment, endowment of skills, endowment of ability or inheritance.

4. The theoretical model developed in this chapter examines mean and the coefficient of variation of income, a Lorenz-consistent measure of income inequality (Ray, 1998, chap. 6). Other models do not use a Lorenz-consistent measure (Meckl and Zink, 2004; Alesina and Rodrik, 1994; Persson and Tabellini, 1994; Perotti, 1993; Galor and Zeira, 1993).

5. Like the models of Stiglitz and of Meckl and Zink, the model developed in this chapter is based on the classic model of economic growth developed by Solow (1956). Consequently, the model developed in this chapter examines the relationship between income inequality and an economy's level of economic development, not its rate of growth. By contrast, other authors have used endogenous growth models (Alesina and Rodrik, 1994) and overlapping generations models (Persson and Tabellini, 1994; Galor and Zeira, 1993) to relate income inequality to the rate of economic growth.

To emphasize the novelty of the approach developed here, this chapter first reviews a sample of the vast literature on the relationship between the distribution of income and economic growth. This chapter then introduces income inequality directly into the Solow Model in a very simple fashion and derives intuitive results. Attempts are also made to test the model empirically.

The key to the model's simplicity is the assumption that steady-state income inequality depends entirely on household saving rates (defined as the percentage of gross income that is saved). Relatively high-saving households accumulate more capital than relatively low-saving households and therefore have higher steady state income.

The model also makes the simplifying assumptions that household saving rates do not depend on real interest rates and that household labor supply is perfectly inelastic. Relaxation of these assumptions is left for future research.

Households are assumed to save out of disposable income, so – when the government's budget must be balanced – raising tax rates reduces household saving rates and causes the economy to converge to a lower steady-state level of average income. Different forms of taxation have different effects on inequality however.

Increasing the marginal tax rate on capital income causes the economy to converge to a more egalitarian distribution of income because relatively high-saving households face a larger percentage decrease in disposable income than relatively low-saving households. Consequently, relatively high savers see a greater percentage decrease in their saving rates, steady-state capital stocks and income than relatively low savers. As the economy converges to its new steady state, the degree of income inequality narrows.

Increasing the marginal tax rate on labor income increases the degree of inequality because it causes a greater percentage decrease in the saving rates of relatively low savers. As the economy converges to its new steady state, the degree of income inequality widens because relatively low savers see a greater percentage decrease in their capital stocks and income than relatively high savers.

Given the model, we should expect to find a positive correlation between taxes on labor income and inequality and a negative correlation between taxes on capital income and inequality. We should also expect to find a negative correlation between tax rates and the level of average income.

In practice however, five difficulties arise in testing the model. First, a tax on wage income is both a tax on labor income and a tax on the returns to human capital, so its empirical sign in a regression on income inequality could be either positive or negative. Second, a high degree of correlation among the various tax rates in our dataset hinders our ability use regression analysis to determine the empirical signs of taxes on capital income.

Third, this is a model of economic development and income inequality. It is not a model of household saving behavior. The model makes the simplifying assumption that households save a constant share of their disposable income and it assumes away the possibility that tax policy will affect households' decisons about what percentage of their disposable income they should save. In practice, lower tax rates on capital income may be associated with higher levels of economic development if reducing the tax rate on capital income induces households to save a larger share of their disposable income.

Fourth, in order to obtain a steady-state solution, the model assumes that the government's budget is balanced at all times. This assumption implies that tax increases reduce the economy's saving rate because tax increases reflect higher government consumption. Relaxing this assumption helps explain the positive correlation between tax rates and average income. When the government's budget does not have to be balanced, the economy's saving rate is an increasing function of the tax rate.

Finally, we need a sample of economies for which both income and tax data were collected in a consistent manner. Use of US state data overcomes the last difficulty, but creates an additional problem: labor migration violates the model's assumption of inelastic labor supply. In fact, Feldstein and Vaillant Wrobel found that gross wages are particularly responsive to changes in the tax structure because individuals may migrate from one state to another until the after-tax wages of equivalent individuals are equal.

Nonetheless, empirical tests of provide some support for the theoretical model's prediction of the effects of taxation on income inequality. The tests used the NBER-TAXSIM's average marginal state income tax rates and the coefficients of variation of full-time workers' adjusted gross income, which were computed from the March Current Population Survey (CPS).

The empirical tests suggest that higher wage taxes are associated with a greater degree of inequality, even when the regression model includes educational inequality as a control for the fact that a tax on wages taxes the returns to both labor and human capital. The coefficient is positive across specifications and is statistically significant from zero when the data is examined in levels, even after controlling for state fixed effects. When the data is differenced, the coefficient on wage taxes is positive, but not statistically significant. If taxation of wage income taxes the returns to labor more than the returns to human capital, then these results provide some support for the theoretical model.

The collinearity of the tax rates and the various forms of capital income taxation prevent us from discerning the empirical sign of taxes on capital income. When the data is examined in levels, taxes on long-term capital gains, which exhibit the lowest degree of correlation with the other tax rates, are associated with lower inequality (as predicted by the theoretical model), but the coefficient is only statistically significant from zero when fixed state and year effects are excluded from the regressions. When the data is differenced, the coefficient is positive, but not statistically significant.

By contrast, the theoretical model performs poorly when examining average income. Lower average marginal tax rates on long-term capital gains are associated with higher average income, but most of the other tax rates exhibit a positive or insignificant correlation with average income. The positive correlations may reflect the fact that the economy's saving rate is an increasing function of tax rate when government budgets do not have to be balanced. The negative correlation between tax rates on long-term gains probably reflects a substitution effect whereby households increase the share of disposable income that they save when a lower tax rate raises the relative return to saving.

The rest of this chapter is organized as follows: Section 2.1 reviews some of the economic literature on the topic and argues that any empirical study of the relationship between economic development and income inequality should be grounded in a theoretical model that explains both the level of economic development and income inequality. Section 2.2 develops such a theoretical model from the Solow Model of economic growth. Section 2.3 presents the results of empirical tests of the model. Section 2.4 concludes.

2.1. Literature Review

Kuznets (1955) first suggested that inequality initially increases and then falls as per capita income rises. Reviewing evidence from the United States, England and Germany he hypothesized that the rise and fall in inequality reflects movement of workers from a relatively low paying sector to a high paying sector. Specifically, he hypothesized that it occurred as workers moved from agriculture to industry. Robinson (1976) formalized his hypothesis mathematically and showed that the only assumptions necessary to derive a U-shaped relationship between income inequality and average income are that the economy has two sectors and that there is a monotonic increase in the share of one sector over time.

Chenery et al. (1974) and Ahluwalia (1976) provided critical support for Kuznets's hypothesis when they ran cross-sectional regressions with quadratic per capita income variables. In a more recent study, List and Gallet (1999) found that cubic per capita income variables were positive and statistically significant in a panel dataset of 71 countries over the period 1962 to 1992 and argued that inequality may increase at higher levels of economic development as workers move from industrial jobs to service sector jobs.

It is important to note however that List and Gallet's specification tests reject the null hypothesis that all country fixed effects are equal to zero, which implies that a universal relationship between development and inequality does not exist. Such a finding compels us to derive a theoretical model of the relationship from a standard model of economic growth and development, so that we can trace several possible inequality-development paths that an economy may follow.

To this author's knowledge, the models developed by Stiglitz (1969) and Meckl and Zink (2004) are the only models of the relationship between income inequality and economic development that are based on the Solow Model. We will first review these two models and then compare them with theoretical models that are based on theories of endogenous growth (Alesina and Rodrik, 1994), overlapping-generations (Persson and Tabellini, 1994; Galor and Zeira, 1993; Loury, 1981), "trickle-down growth" (Perotti, 1993) and political tolerance for inequality (Hirschman and Rothschild, 1973). This chapter will also review some of the microeconomic-based analyses of the relationship between inequality and development, such as the analyses of Becker and Chiswick (1966), of Chiswick (1971) and of Newhouse (1971).

In his basic model, Stiglitz (1969) divides society into groups of constant proportions that differ in their initial endowment of capital and assumes that all groups have the same marginal propensity to save[1]. The latter assumption is particularly important because it implies that there will be complete equalization of wealth and income when the economy converges to a stable steady-state.

In his article, Stiglitz also considers other savings functions, but only one generates a two-class economy. That case however depends on an unstable equilibrium for the lower class and if that unstable equilibrium were disturbed by an increase in the capital-labor ratio, the incomes of the lower class would rise and economy would converge to an egalitarian steady-state.

Egalitarian steady-states are not the norm in Solow-based models however. Meckl and Zink (2004) show that differences in inherent ability among workers are sufficient to ensure that income inequality will exist in steady state. Specifically, Meckl and Zink used the Solow Model to examine

[1]There is no human capital in Stiglitz's model.

the effect that capital accumulation has on the wage structure. In their model, individuals invest in acquiring skills (human capital) if their ability level exceeds a critical threshold that depends on the interest rate (the cost of borrowing). Meckl and Zink then observe that, as an economy accumulates physical capital, the rate of return on capital falls, which reduces interest rates and induces a larger share of the population to acquire skills. Consequently, the heterogeneity of skilled individuals rises as the economy accumulates physical capital and the heterogeneity of unskilled individuals falls.

Because workers are paid their marginal value product of labor, wage inequality rises among skilled workers and falls among unskilled workers. Overall inequality, as measured by the average wage of skilled workers relative to the average wage of unskilled workers, first falls and then rises. The casual reader will quickly notice that their model appears to predict the opposite of what Kuznets's hypothesis predicts. The more observant reader will also notice that the average wage of skilled workers relative to the average wage of unskilled workers is not a Lorenz-consistent measure of inequality, which hampers our ability to test their model empirically.

Meckl and Zink are not alone in their omission of a Lorenz-consistent measure. In other models, such as the one developed by Alesina and Rodrik (1994), income inequality is measured by the ratio of median to mean income. Such a measure is not Lorenz-consistent because a direct transfer of income from the richest households to the poorest would leave the ratio of median to mean income unchanged. In defense of Meckl and Zink, one might point out that they used the Gini coefficient to measure intra-group wage inequality, even though they didn't use it to measure overall wage inequality.

In their defense, one might also point out that the theoretical models of Meckl and Zink and of Alesina and Rodrik at least attempt to provide a measure of income inequality. Other authors do not provide an explicit measure. Persson and Tabellini (1994) leave the reader to infer that the degree of inequality depends on the ratio of median to mean skills. One cannot even measure the degree of inequality in the models of Perotti (1993) and of Galor and Zeira (1993).

Another similarity between the Solow-based models of inequality and development (discussed above) and the ones developed by Alesina and Rodrik, by Persson and Tabellini, by Perotti, by Galor and Zeira and by Loury (1981) is that they assume an initial uneven distribution of income or skills. Because such models assume inequality, they are unable to explain the source of income inequality.

One difference between the Solow-based models and other models is that there is no political process in the Solow-based models. In other models, the unequal distribution of income affects the political process through the choice of a tax rate on capital income. For example, in the models developed by Alesina and Rodrik and by Persson and Tabellini, the chosen tax rate is the one preferred by the median voter and the median voter's preferred tax rate is an increasing function of the disparity between the median and average incomes.

Specifically, Alesina and Rodrik incorporate income inequality into an an endogenous growth model by assuming an initial uneven distribution of factor endowments. In their model, government services are financed by taxing the capital stock. Because government services are a productive factor, the growth rate is a function of the tax rate. The tax rate that maximizes the growth rate is the one preferred by individuals whose income comes solely from capital. Relatively more capital-poor individuals prefer a higher tax rate (and a slower growth rate). Using the median voter theorem, they assume that the chosen tax rate is the one preferred by the median voter.

Because they take the difference between the median and average incomes as the measure inequality and assume that the median is less than the average, they conclude that a larger degree of inequality is associated with a slower growth rate.

The crucial assumption of their model is that all individuals have the same rate of time preference. This assumption simplifies their model because it allows the wealth and income distributions to remain constant over time (so that the median voter also remains the same across time). If the wealth and income distributions remain constant over time however, then it is unclear how inequality could have arisen in the first place.

By contrast, Persson and Tabellini (1994) assume that inequality arises from an uneven distribution of skills. In their model, all individuals have the same saving rate, so higher-skill individuals accumulate more capital. Specifically, they use an overlapping generations model with a proportional capital income tax, which they refer to more generally as a “public policy variable” because it indexes the ability of individuals to appropriate the returns on their investments. They also assume that the revenue from the tax is distributed evenly across all individuals.

In their model, the chosen tax rate is the one preferred by the median voter. If the median voter is less skilled than average, then he/she prefers a redistributive policy that reduces the growth rate. Like Alesina and Rodrik’s model, Persson and Tabellini’s model implies that faster growth rates occur when the median skill level is close to the mean, whereas a larger disparity results in more redistribution and less growth.

Although their model explains how income inequality could arise and how a greater degree of inequality could be associated with lower rates of economic growth, its capital accumulation path is poorly defined. In their model, the “old” consume all of the capital that they accumulated while they were “young” (net of redistribution). To generate growth, they assume that the average stock of capital accumulated “by the previous generation has a positive externality on the income of the newborn generation” (Persson and Tabellini, 1994, p. 602).

Other models also lack a well-defined capital accumulation path. For example, Perotti (1993) uses “trickle-down growth” to model the relationship between growth and inequality. In his model, individuals belong to one of three income groups (classes). Individuals cannot borrow, so they can only invest in education if their after-tax income in the initial period exceeds the cost of investment. Individuals who invest receive higher income in the subsequent period and their education has an externality that raises the income of all other individuals, some of whom can then invest in the subsequent period.

In his model, taxes are levied in proportion to pre-tax income and the tax revenue is evenly redistributed. Perotti assumes that the political process selects the median voter’s preferred tax rate. The chosen tax rate determines whether or not the growth process stops before all classes have invested in education. If all classes invest in education, steady state income is maximized.

Because the choice of tax rate depends on the economy’s initial conditions, Perotti’s model is capable of generating a cross-sectional inverted-U shaped relationship between average income and inequality. The growth process may not begin in a poor economy that starts with a very low degree of income inequality. A poor economy with a high degree of income inequality may begin the growth process, but the middle class may stop the growth process after they have invested. In a more egalitarian (but not poor) economy, all classes eventually invest in education and maximize steady state income.

One shortcoming of Perotti's model is that both education and the "externality" associated with education exhibit constant marginal returns. Most economic models assume diminishing marginal returns. Another shortcoming stems from the fact that Perotti designed his model to match the empirical findings of Kuznets (1955); Chenery et al. (1974) and Ahluwalia (1976) who found that income inequality is lower at low and high levels of economic development than it is at intermediate levels. One should test a theory.

One common shortcoming of these three models of an inequality-determined tax rate is that they only consider one form of taxation. Alesina and Rodrik (1994) and Persson and Tabellini (1994) only consider the case of a tax on capital, while Perotti only considers the case of a proportional income tax. The implications of their models may have been different if they considered other forms of taxation.

Even if we accept the notion that the degree of inequality determines the tax rate, it is not necessarily obvious that the tax rate chosen would be the one preferred by the median voter. After all, the "rich" would have more resources to spend on lobbying for a favorable tax rate and parties with more to gain (or more to lose) would spend more resources on lobbying.

Alternatively, the chosen tax rate may reflect society's tolerance for inequality, in which case there's an ambiguous relationship between inequality and development. Hirschman and Rothschild (1973) point out that tolerance for inequality may rise as the economy enters a growth phase if an individual's utility increases when a similarly situated individual's income increases. They refer to this as a "tunnel effect" and argue that it arises when an individual's expectation of future income depends in part on the income of his/her neighbors. If a neighbor's fortunes improve, he/she may also expect to experience an improvement in fortune, which increases his/her tolerance for inequality in the course of economic development.

Hirschman and Rothschild caution however that tolerance for inequality may wane if the benefits of economic development increasingly accrue to a small number of decision-makers, as would occur if the economy became "more oligopolized and bureaucratized" or if the benefits accrue to one ethnic or religious group. They also argue that there may be no tolerance for inequality if members of society believe that the economic system is a zero-sum game in which the gain of a neighbor will inevitably result in a loss for them.

Hirschman and Rothschild conclude that developing countries that possess a strong "tunnel effect" are relatively more politically stable. However developing countries are relatively less politically stable when their societies are intolerant of inequality.

Although there is no mention of redistributive taxes in their paper, one could draw the analogy that developing countries with a strong "tunnel effect" are less likely to adopt redistributive taxes or have lower rates of redistributive taxation. In this sense, Hirschman and Rothschild's model also implies that the distribution of income affects the rate of growth. Other models share this assumption.

In a dynastic model of how income distribution affects economic growth, Galor and Zeira (1993) assume that individuals can borrow or lend and an individual's inherited wealth determines whether or not he or she will invest in education, which exhibits constant marginal returns. An individual's investment decision determines his/her lifetime wealth and how much he/she will bequeath to his/her offspring.

To relate the distribution of wealth to economic growth, they assume that the lenders receive a lower interest rate than the rate at which borrowers pay. Since borrowing is costly, individuals who inherit more wealth need to borrow less in order to finance their investment. According to Galor and Zeira, credit market imperfections are not sufficient to affect growth in the long run however because the distribution of wealth would converge to a unique, ergodic distribution if individuals could invest in education incrementally.

To generate a model in which an initial unequal wealth distribution affects the rate of economic growth in the long run, Galor and Zeira assume that investment in education is indivisible. This assumption implies that the long-run level of average income and wealth depends on the initial number of individuals who inherit enough wealth to both invest in education and bequeath enough wealth to their offspring, so that all future generations of their dynasty will invest in education. Consequently, an economy's long-run level of economic development depends on its initial distribution of wealth.

The growth implications of their model hang delicately on the assumption of indivisible education. In fact, Galor and Zeira based their model on the one that Loury (1981) developed, but in Loury's model, education exhibits diminishing marginal returns and there are no capital markets to finance educational investment. Consequently, income inequality persists in Loury's model.

Importantly however, Loury shows the conditions under which public provision of education increases average income and reduces income variance.

There are no capital markets in Loury's model because he assumes that parents altruistically finance the education of their children and low-income parents cannot legally obligate their children to repay debts incurred for their training. Consequently, a high-income parent invests more in the training of his/her child than a low-income parent, but faces a lower marginal return on the investment. Under such circumstances, the absence of intertemporal borrowing and lending reduces the welfare of both families.

Loury adds risk to his model with the assumption that parents do not know their child's ability level when they decide how much to invest. As a result, the return on a given investment is variable. Permanently taxing and redistributing the income of the next generation would reduce the risk associated with investment however. Loury shows that risk aversion implies that a sufficiently small degree of redistribution would enhance welfare because the "insurance effect" of redistribution is stronger than "excess burden effect" of the tax.

With the exception of Loury's model and the Solow-based models, all of the models discussed above assume that the income distribution affects either the level of economic development or the rate of economic growth. Empirically however, the correlation between inequality and growth is ambiguous.

Alesina and Rodrik (1994) found a negative correlation between income inequality (as measured by the Gini coefficient) and subsequent growth, while Persson and Tabellini (1994) found a positive correlation between the income share of the middle quintile (which is inversely related to the degree of inequality) and subsequent growth. Refuting such findings, Forbes (2000) argues that previous studies of the relationship between inequality and growth found an inverse relationship because they omitted regional dummy variables. She differenced the dataset developed by Deininger and Squire (1996) and found that the short-run effect of an increase in inequality is an increase in growth. Turning to non-parametric methods, Banerjee and Duflo (2003) found an inverted-U shaped relationship between inequality and growth. More precisely, they found that a slower growth rate follows an increase in inequality as well as a decrease in inequality.

Given the ambiguous correlation between inequality and growth and the assumptions authors have made to generate theoretical models of the relationship between them (e.g. income distributions that remain constant over time and poorly defined capital accumulation processes), it seems sensible to focus our attention on the level of development (instead of the rate of growth) and develop a model in which the degree of inequality and the level of economic development are jointly determined.

In fact, microeconomic-based analyses also assert joint determination of inequality and development. For example, Becker and Chiswick (1966) argue that income inequality reflects inequality in education if higher ability people have better access to financing for their education. Conversely, income inequality does not reflect inequality in education if people who have more ability have less access to financing for their education.

Using data from the 1960 U.S. Census, they found positive correlations between inequality in adult male incomes, rates of return to schooling and schooling inequality. They also found that those variables are all negatively correlated with average schooling and average income. On the basis of these correlations, they concluded that the 1960 Census data suggests that income and schooling inequality reduce a state's level of economic development.

Chiswick (1971) however argues on theoretical grounds that – ceteris paribus – earnings inequality should increase with economic development because relative inequality is an increasing function of the average level of investment. He based his argument on a decomposition of the variance of income however. Had he used the coefficient of variation (a Lorenz-consistent measure of inequality), he might have obtained a different result because average income (which of course is the denominator of the coefficient of variation) is also an increasing function of the average level of investment.

Newhouse (1971) decomposed human capital into general training and specific training and applied Becker and Chiswick's theoretical framework to an empirical analysis of income inequality in US cities. Specifically, they assumed that different industries provide different amounts of specific training, but the firms within a given industry provide the same amount of specific training. Under such conditions, the industry mix in a local labor market determines the distribution of specific training within a local labor market. Industry mix can also determine the distribution of general training if more educated workers are more mobile and move to labor markets that offer jobs for highly-skilled workers. Using 1960 income data from the US Treasury Department and 1960 industry data from the US Census Bureau, they found that their simple hypothesis explained 88 percent of the variance in the proportion of incomes in 15 size classes.

While these studies provide valuable insights into the relationships between development and inequality, any assertion that there is a relationship between development and inequality requires a theoretical model of the relationship in which gross income inequality and the level of economic development are jointly determined. This chapter attempts to build the requisite model from Solow's classic growth model.

2.2. Using the Solow Model to Explain Income Inequality

The greatest difficulty in designing a theoretical model that explains both an economy's level of development and its degree of inequality is overcoming the fact that the Solow Model is not designed to handle issues of income distribution. The basic model simply aggregates capital and labor.

A convenient way to introduce income inequality into the Solow Model is to assume that households have different saving rates, so that they accumulate slightly different amounts of capital. When capital is unevenly distributed, household incomes vary.

Exogenous changes in household saving rates (such as autonomous decisions to save more/less or changes induced by tax policy) affect both the economy's level of economic development (as measured by average pre-tax income) and the economy's degree of inequality (as measured by the coefficient of variation of pre-tax income). In contrast to the articles discussed in the literature review, this chapter's measure of income inequality (the coefficient of variation) is Lorenz-consistent (Ray, 1998, chap. 6).

Although the model explains how tax policy affects the degree of pre-tax income inequality and the level of average pre-tax income, it does not include any description of the process by which a particular tax rate is chosen or how the tax revenue is spent. As discussed in the literature review, the median voter theorem is a poor simplification of the political process when some households have more political power than others and/or when the chosen tax rate reflects society's tolerance for inequality.

On the expenditure side, the model assumes that government purchases are equal to tax revenue. The assumption of a balanced budget is necessary to obtain a steady-state solution.

The model also assumes that government purchases are pure consumption and therefore do not enter the production function. One could argue that government spending on education adds to the economy's stock of human capital and therefore contributes to production. In this case, the assumption that government spending is a form of consumption is poor. However one can think of many other cases in which government spending does not contribute to production. For example, government purchases may provide protection from foreign attack, but national security only prevents destruction of the economy. One cannot use national security to produce goods and services. Other forms of government spending provide households with positive utility (such as parks and recreational facilities), but the welfare implications of such spending are outside the scope of a model that seeks to describe the relationship between income distribution and economic development.

The purpose of this model is to provide insight into the relationship between economic development and income inequality. Instead of assuming that inequality affects development, inequality and development are jointly determined within the model. It accomplishes this task by linking differences in household saving rates (which determines the steady state degree of income inequality) to the overall saving rate (which determines the steady state level of average income). The model also provides insight into the way in which tax policy affects the distribution of income. When a tax rate changes, the coefficient of variation will monotonically increase or decrease during the transition to its new steady state level.

The model does not attempt to explain why some households have a higher marginal propensity to save than others, but one can easily imagine several possible explanations, some of which imply that income inequality is inevitable and some of which imply that it is not. If households save for investment in human capital, then relatively higher ability households may save more because they face a higher marginal return from investment in education. In such a case, income inequality is inevitable. On the other hand, if differences in household saving rates reflect differences in rates of time preference, then the steady-state level of income inequality depends on households's own decisions. In such a case, simply convincing households to save and invest at the same rate would completely eliminate income inequality in the long run.

Importantly, this chapter examines the coefficient of variation of gross income as opposed to net income because gross income inequality is the primary source of income inequality and because tax rates could be set to completely equalize net income.

To derive the coefficient of variation of gross income, we'll start by assuming that output, Y, is produced using various forms of capital, K_j, and labor, L, according to the linearly homogenous production function:

$$\begin{aligned} Y &= F(K_1, K_2, \cdots, K_M, L) \\ Y &\equiv L \cdot \bar{y} = L \cdot f(\overline{k_1}, \overline{k_2}, \cdots, \overline{k_M}, 1) \end{aligned} \tag{2.1}$$

where $\bar{y} \equiv Y/L$ is income per worker and $\overline{k_j} \equiv K_j/N$ is type-j capital per worker.

For simplicity, we have omitted technological progress from the production function because the absence of technological progress does not affect the key implication of our model: that differences in household saving rates are the source of income inequality. Regardless of whether technology is labor-augmenting or capital-augmenting (or both), the coefficient of variation derived below would be the same.

Finally, we'll also assume that each household contributes one unit of labor to the production of output (i.e. $L = N$ where N is the number of households) and that there is no growth in the number of households. The assumption of a constant number of households is necessary to obtain a closed form solution for the steady-state degree of income inequality.

Given these assumptions, average household income, $\bar{y} \equiv Y/N$, is given by:

$$\bar{y} = w + r_1\overline{k_1} + r_2\overline{k_2} + \cdots + r_M\overline{k_M} \tag{2.2}$$

where w is the wage rate and r_j is the rental rate on capital of type j. If factors of production are paid their marginal value products, then the rental rate on type-j capital is given by:

$$r_j = \frac{\partial F}{\partial K_j} \equiv f_j \tag{2.3}$$

and the wage rate paid to labor is given by:

$$\begin{aligned} w &= \frac{\partial F}{\partial L} \equiv f - \overline{k_1} f_1 - \overline{k_2} f_2 - \cdots \overline{k_M} f_M \\ &= \bar{y}(1 - \alpha_1 - \alpha_2 - \cdots - \alpha_M) \end{aligned} \tag{2.4}$$

where α_j is defined as type-j capital's share of income (i.e. $\alpha_j \equiv \overline{k_j} f_j / \bar{y}$).

When individual households do not possess equal amounts of capital, households with more capital receive more income and households with less capital receive less income:

$$y_i = w + r_1 k_{i,1} + r_2 k_{i,2} + \cdots + r_M k_{i,M}. \tag{2.5}$$

Making use of the capital shares we can express household i's income as:

$$y_i = \bar{y}\left(1 - \sum_{j=1}^{M} \alpha_j + \sum_{j=1}^{M} \alpha_j \frac{k_{i,j}}{\overline{k_j}}\right). \tag{2.6}$$

which implies that the coefficient of variation of gross income is given by:

$$cv(y) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\sum_{j=1}^{M} \alpha_j \frac{k_{i,j}}{\overline{k_j}}\right)^2 - \left(\sum_{j=1}^{M} \alpha_j\right)^2} \tag{2.7}$$

It is easiest to analyze such an expression when the economy is in steady state. (i.e. $\bar{y} = y_{ss}$ and all $\overline{k_j} = k_{ss,j}$), but to do so we first need to know the *individual household saving rates* that are consistent with an individual household possessing more or less capital than average *in steady state*. After all, if households all had the same saving rate, but started off with different initial endowments of capital, then depreciation of the capital stock would ensure that each household possesses the same amount of capital in steady state (Stiglitz, 1969). Therefore, income inequality can only exist in steady state if individual households have different saving rates.

Households save out of their disposable income, y_i^D, but taxes affect their saving rates because the rate at which a household saves and invests in type-j capital is its ratio of saving to pre-tax income. In this chapter, we'll assume that each household saves constant fractions, $\sigma_{i,j}$, of their disposable income, so that – in the absence of taxes – their marginal saving rates are equal to their average saving rates.

This assumption obviously ignores the fact that taxation of capital income (which is effectively a tax on the returns to saving) may cause households to reduce their marginal saving rates via a substitution effect, whereby households substitute out of future consumption and into present consumption (Musgrave, 1959). Nonetheless, this dissertation leaves such an analysis to future research.

In the presence of taxes, a household's average saving rates, $\sigma'_{i,j}$ and the economy's average saving rates, s'_j are less than the marginal saving rates. Using τ_L to denote the rate at which labor income is taxed and τ_j to denote the rate at which type-j capital income is taxed, we can make use of 2.6 to express the (average) rate at which household i saves and invests in type-h capital as:

$$\sigma'_{i,h} \equiv \frac{\sigma_{i,h} \cdot y_i^D}{y_i} \equiv \sigma_{i,h} \left(\frac{\left(1 - \sum_{j=1}^{M} \alpha_j\right)(1-\tau_L) + \sum_{j=1}^{M} \alpha_j \left(1-\tau_j\right) \frac{k_{i,j}}{\overline{k_j}}}{\left(1 - \sum_{j=1}^{M} \alpha_j\right) + \sum_{j=1}^{M} \alpha_j \frac{k_{i,j}}{\overline{k_j}}} \right) \tag{2.8}$$

Similarly, making use of 2.6 and 2.8 and assuming that the government's budget is balanced, we can express the economy's (average) type-*h* saving rate as:

$$s_h' \equiv \frac{\sum_{i=1}^{N} \sigma_{i,h}' \cdot y_i}{N\bar{y}} \equiv \left(\left(1 - \sum_{j=1}^{M} \alpha_j\right)(1-\tau_L)\frac{1}{N}\sum_{i=1}^{N}\sigma_{i,h}\right) + \left(\sum_{j=1}^{M}\alpha_j\left(1-\tau_j\right)\frac{1}{N}\sum_{i=1}^{N}\sigma_{i,h}\frac{k_{i,j}}{\overline{k_j}}\right) \quad (2.9)$$

According to the Solow Model, all type-*h* saving, $s_h'Y \equiv \sum_{i=1}^{N} \sigma_{i,h}' \cdot y_i$, is invested in new type-*h* capital, $I_h = s_h'Y$, so the net change in the type-*h* capital stock is equal to investment minus depreciation:

$$\dot{K}_h = s_h'Y - \delta_h K_h \quad (2.10)$$

where δ_h represents the rate at which type-*h* capital depreciates. Because we have assumed away technological progress and labor force growth, there is no growth of the capital stock in steady state (i.e. $\dot{K}_{ss,h} = 0$). In steady state, the average type-*h* capital stock is equal to the steady state stock of type-*h* capital per household (i.e. $\overline{k_h} = k_{ss,h}$) and:

$$k_{ss,h} = \frac{s_{ss,h}' y_{ss}}{\delta_h}. \quad (2.11)$$

Similarly, the net change in a household's type-*h* capital stock is also equal to investment minus depreciation:

$$\dot{k}_{i,h} = \sigma_{i,h}' y_i - \delta_h k_{i,h}. \quad (2.12)$$

When the economy is in steady state, all households' must also be in steady state (i.e. $\dot{k}_{i,ss,h} = 0$) and the ratio of household *i*'s steady-state type-*h* capital stock to the economy's steady-state type-*h* capital stock depends on household *i*'s saving rates and the economy's saving rates:

$$\frac{k_{i,ss,h}}{k_{ss,h}} = \frac{\sigma'_{i,ss,h}}{s'_{ss,h}} \left(\frac{1 - \sum_{j=1}^{M} \alpha_j}{1 - \sum_{j=1}^{M} \alpha_j \frac{\sigma'_{i,ss,j}}{s'_{ss,j}}} \right) \tag{2.13}$$

which allows us to write the steady-state coefficient of variation entirely in terms of the saving rates:

$$cv(y_{ss}) = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left(\frac{1 - \sum_{j=1}^{M} \alpha_j}{1 - \sum_{j=1}^{M} \alpha_j \frac{\sigma'_{i,ss,j}}{s'_{ss,j}}} \cdot \sum_{j=1}^{M} \alpha_j \frac{\sigma'_{i,ss,j}}{s'_{ss,j}} \right)^2 - \left(\sum_{j=1}^{M} \alpha_j \right)^2}. \tag{2.14}$$

Without loss of generality, we can examine the impact of a change in tax policy (or any other public policy that affects household saving rates) by restricting our attention to the case where there are only two factors of production: capital, K, and labor, L.

Because households save out of their disposable income, tax policy reduces households' average saving rates and changes in tax rates affect the distribution of income through their effect on household capital accumulation. The effect that changes in tax rates have on steady-state income inequality can be measured with the percentage change in the coefficient of variation, which depends on the percentage changes in household saving rates and the percentage change in the economy's saving rates. When there are only two factors of production, the percentage change in the coefficient of variation is given by:

$$\frac{\partial cv(y_{ss})/\partial\tau}{cv(y_{ss})} = \left(\frac{\alpha\,(1-\alpha)}{cv(y_{ss})}\right)^2 \frac{1}{N}\sum_{i=1}^{N} \frac{\frac{s'_{ss}}{\sigma'_{i,ss}}\left(\frac{\partial\sigma'_{i,ss}/\partial\tau}{\sigma'_{i,ss}} - \frac{\partial s'_{ss}/\partial\tau}{s'_{ss}}\right)}{\left(\frac{s'_{ss}}{\sigma'_{i,ss}}-\alpha\right)^3}. \tag{2.15}$$

Because taxes reduce household saving rates and the economy's saving rate, all forms of taxation reduce steady-state average income. Different forms of taxation have different effects on inequality however.

When a tax is levied on total income (i.e. the case in which: $\tau = \tau_L = \tau_K$), a higher tax rate reduces each household's saving rate and the economy's saving rate by the same proportion.

$$\frac{\partial\sigma'_i/\partial\tau}{\sigma'_i} = \frac{\partial s'/\partial\tau}{s'} = \frac{-1}{1-\tau}$$

Consequently, a tax on total income has no effect on income inequality. One may also notice that Lorenz-consistency requires that a tax on total income will leave the coefficient of variation unchanged.

The impact of a tax on the returns to capital and the impact of a tax on labor income depend on the distribution of household saving rates, we need to use a simulation to obtain the percentage change in the coefficient of variation in response to a change in those tax rates.

To perform the necessary simulation, $10,000$ values of σ_i were randomly drawn from a log-normal distribution with $mean(ln(\sigma)) = ln(0.10)$ and $sd(ln(\sigma)) = 0.10$. In the presence of taxes, the steady-state capital ratio $k_{i,ss}/k_{ss}$ depends on the average saving rates $\sigma'_{i,ss}$ and s'_{ss} (and vice versa), so these values had to be obtained through recursive substitution. Since average income depends on the functional form of the production function, a simulation requires the selection of a production function. A Cobb-Douglas production function was chosen.

Tables 2.1 and 2.2 present the results of the simulation. Table 2.1 assumes that $\tau_L = 0$ and simulates the effect of changes in the tax on capital income. Table 2.2 assumes that $\tau_K = 0$ and simulates the effect of changes in the tax on labor income.

Because the tax on capital reduces every household's saving rate (by reducing disposable income), each household's steady-state capital stock and steady-state income level falls. This effect is reflected in the negative values of the percentage change in average income in Table 2.1a.

The decrease in steady-state income is not proportionate among households however. Relatively low-saving households see a smaller percentage decrease in their saving rates than relatively high-saving households because relatively low savers receive a smaller share of their income from returns on capital than relatively higher savers.

Because the percentage decrease in their saving rates is smaller, relatively low savers experience a smaller percentage decrease in their capital stock than relatively high savers as the economy converges to the new steady state. Therefore, the variance in household capital stocks declines and the economy converges to a more egalitarian distribution of income. This effect is reflected in the negative values of the percentage change in the coefficient of variation in Table 2.1b.

Table 2.1
Simulation: Effect of a Tax on Capital Income

Table 2.1a

Average Income

	$\tau_K = 0.15$	$\tau_K = 0.30$	$\tau_K = 0.45$
$\alpha = 0.25$	1.247	1.230	1.213
$\alpha = 0.35$	1.417	1.374	1.329
$\alpha = 0.45$	1.679	1.578	1.475

Percentage Change in Average Income

	$\tau_K = 0.15$	$\tau_K = 0.30$	$\tau_K = 0.45$
$\alpha = 0.25$	−0.087	−0.091	−0.095
$\alpha = 0.35$	−0.201	−0.212	−0.226
$\alpha = 0.45$	−0.398	−0.429	−0.465

Table 2.1b

Coefficient of Variation

	$\tau_K = 0.15$	$\tau_K = 0.30$	$\tau_K = 0.45$
$\alpha = 0.25$	0.029	0.028	0.026
$\alpha = 0.35$	0.044	0.042	0.040
$\alpha = 0.45$	0.066	0.061	0.056

Percentage Change in the Coefficient of Variation

	$\tau_K = 0.15$	$\tau_K = 0.30$	$\tau_K = 0.45$
$\alpha = 0.25$	−0.435	−0.469	−0.508
$\alpha = 0.35$	−0.733	−0.815	−0.913
$\alpha = 0.45$	−1.123	−1.128	−1.489

Table 2.2
Simulation: Effect of a Tax on Labor Income

Table 2.2a

Average Income

	$\tau_L = 0.15$	$\tau_L = 0.30$	$\tau_L = 0.45$
$\alpha = 0.25$	1.214	1.161	1.102
$\alpha = 0.35$	1.382	1.300	1.213
$\alpha = 0.45$	1.659	1.537	1.414

Percentage Change in Average Income

	$\tau_L = 0.15$	$\tau_L = 0.30$	$\tau_L = 0.45$
$\alpha = 0.25$	−0.281	−0.321	−0.375
$\alpha = 0.35$	−0.386	−0.432	−0.491
$\alpha = 0.45$	−0.487	−0.534	−0.592

Table 2.2b

Coefficient of Variation

	$\tau_L = 0.15$	$\tau_L = 0.30$	$\tau_L = 0.45$
$\alpha = 0.25$	0.031	0.033	0.036
$\alpha = 0.35$	0.050	0.054	0.060
$\alpha = 0.45$	0.076	0.084	0.096

Percentage Change in the Coefficient of Variation

	$\tau_L = 0.15$	$\tau_L = 0.30$	$\tau_L = 0.45$
$\alpha = 0.25$	0.515	0.680	0.940
$\alpha = 0.35$	0.821	1.047	1.389
$\alpha = 0.45$	1.207	1.506	1.964

In the case of a tax on capital income, there is a tradeoff between development and equality. Reducing the tax on capital increases average income, but it also increases inequality. By contrast, when the same simulation was run using a tax on labor income, no tradeoff between development and equality arose. Reducing the tax on labor income increases average income and reduces inequality.

Just as we saw in the case of a tax on capital income, raising a tax on labor income reduces every household's disposable income and therefore reduces every household's saving rate, which reduces every household's steady-state capital stock and the economy's steady-state average income. The decrease in average income is reflected in the negative values of the percentage change in average income in Table 2.2a.

In contrast to the case of a tax on capital income, raising a tax on labor income causes relatively low-saving households to experience a larger percentage decrease in their saving rates than relatively high-saving households, because labor income constitutes a larger share of the relatively low savers' total income than it does of the total income of relatively high savers.

Because they experience greater percentage declines in their saving rates, relatively low-saving households experience a greater percentage decrease in their capital stocks and a greater percentage decrease in income than relatively high-saving households as the economy converges to its new steady state.

Consequently, the economy converges to a more unequal distribution of income when the tax rate on labor income increases. The complementary relationship between development and equality in the case of a tax on labor income is reflected in the positive values of the percentage change in the coefficient of variation in Table 2.2b.

2.3. Empirical Tests

If the model is valid, then we should find a negative correlation between taxation and average income, we should find a positive correlation between taxes on labor income and inequality and we should find a negative correlation between taxes on capital income and inequality.

As mentioned in the introduction to this chapter however, five difficulties inhibit any test of the model. First, there is no pure tax on labor income. Taxes on wage income tax both labor and human capital, so we cannot predict the empirical sign of the coefficient on wage taxes when inequality is the dependent variable. In an attempt to control for the effect of educational inequality on income inequality, the coefficient of variation of years of education has been included in the regression models. Although the inclusion of a measure of educational inequality does not decompose the tax on wage income into its labor and human capital components, it does account for differences in human capital accumulation.

Multicollinearity among the tax rates also inhibits our ability to determine what effect taxes have on inequality and development. One difficulty associated with multicollinearity is that the correlation between two regression coefficients has a sign opposite to the sign of their corresponding explanatory variables, so it is difficult to know what the "true" sign of the regression coefficients is. A second difficulty is that multicollinearity generates large standard errors of the regression coefficients, which encourages to us to conclude that one or several coefficients are not statistically significant from zero even though their “true” effect on the coefficient of variation is strong.

One common method of dealing with multicollinearity is to use principal components regression. As discussed in Chapter 1, Obenchain (2004, chap. 3) shows that principal components regression is a special case of two-parameter ridge regression. Ridge regression shrinks the coefficient estimates towards zero, but reduces the variance around those parameter estimates. In their classic work on ridge regression, Hoerl and Kennard (1970) show that there always exists a biased estimate of the coefficient vector that minimizes the mean squared error (thus optimally balancing variance and bias).

A third problem that arises in examining the empirical relationship between taxation, development and inequality is that we need a data set that contains income and tax data that was collected in a consistent manner for each of the economies in the sample. This difficulty can be overcome by using US state data, but at the cost of computing income inequality statistics from small state samples. This is a price worth paying however if the gains from consistent data collection exceed the losses associated with small sample sizes.

US state data may also inhibit our ability to test the model however because Americans may migrate from one state to another in response to differences in tax structure. This violates the model’s assumption of inelastic labor supply. In fact, Feldstein and Vaillant Wrobel (1998) found that US gross wages adjust quickly to changes in the state tax structure.

In their study, Feldstein and Vaillant Wrobel hypothesized that the ability of Americans to migrate from one state to another should equalize the after-tax incomes of equivalent individuals in the long run. To test their hypothesis, they created a composite "net average tax rate" from federal and state average income tax rates, federal marginal income tax rates, state average sales tax rates and average property tax rates and found that gross wages adjust rapidly to changes in the tax structure.

The tax structure may also affect households' saving decisions. This creates a fourth obstacle to testing the model because the model makes the simplifying assumption that households save a constant share of their disposable income. Because capital must be accumulated over time, present saving postpones consumption into the future. Consequently, taxing the returns to capital is the equivalent of taxing future consumption, which encourages households to substitute out of future consumption and into present consumption (Musgrave, 1959). The functional form of the household utility function then determines whether or not such substitution affects the share of disposable income that households choose to save.

Finally, the assumption of a balanced government budget was necessary to obtain a steady-state solution to theoretical model. This assumption implies that tax increases are reflected in increased government consumption, so the economy's saving rate is a decreasing function of the tax rates in the theoretical model.

This creates a fifth problem when testing the empirical model because the economy's saving rate is an increasing function of tax rates when the government's budget does not have to be balanced. Dropping the balanced-budget assumption implies that higher tax rates should be associated with higher average income because increasing the economy's saving rate should cause the economy to converge to higher steady-state average income.

To obtain the data necessary to empirically test this chapter's model of the relationship between inequality, development and taxation (which reduces saving and the steady state level of economic development), the coefficients of variation of full-time workers' adjusted gross income were computed from the March Current Population Survey (CPS) and combined with average marginal state income tax rates from the NBER's TAXSIM model.

Specifically, the weighted mean and variance of adjusted gross income of people who worked at least 1400 hours in the reference year were computed from the 1992 to 2004 March CPS samples. To preserve the confidentiality of its respondents, the Census Bureau does not report income in excess of $100,000, so the weighted mean and variance had to be recovered by assuming that the log of income is distributed normally and by calculating the moments of the truncated distribution.

The assumption that the log of income is distributed normally prevents us from testing the model directly. Instead of regressing the coefficient of variation of gross income on tax rates, we instead have to regress the coefficient of variation of the log of gross income on tax rates.

The tax rates used in the analysis were the combined federal and state average marginal tax rates computed from the NBER-TAXSIM model's 1995 nationally representative sample. To compute the average marginal tax rates, NBER-TAXSIM calculates each individual's tax liability and his/her tax liability if his/her income was one percent higher. The average marginal tax rate is the average ratio of the additional tax liability to the additional income. According to Feenberg (2005), using the same sample across states and years creates a set of tax rates that does not depend on the economic conditions within a state and can be used as an instrument for the actual tax rate.

To account for the possibility that changes in tax rates affect the degree of inequality with a lag, the tax rates for each year were paired with the coefficient of variation of the log of gross income in the following year.

As shown in the correlation matrices in Tables 2.3b and 2.4b, the tax rates in the NBER-TAXSIM model are highly correlated with one another both in levels and in differences. As discussed in Chapter 1, this multicollinearity inhibits our ability to determine the true signs of the coefficients on the various taxes because when two explanatory variables are positively correlated with each other their coefficients will be negatively correlated with each other.

Multicollinearity may therefore cause us to conclude that raising one of the tax rates may reduce income inequality or average income, when in reality the opposite would occur. In fact, Figures 2.1, 2.4 and 2.6 show that some of the coefficients on forms of capital taxation do change signs as they are shrunk towards zero.

To test the theoretical model, the mean and the coefficient of variation of the log of gross income were regressed on the tax rates in levels with and without state and year dummies. The regressions were also run using log differences of the data. Examining the data in both levels and in log differences provides a check for robustness across specifications.

In an attempt to control for the effect of past inequality and the fact that a tax on wage income implicitly taxes both the returns to labor and the returns to human capital, each state's coefficient of variation of the number of years of education (computed from the same March CPS sample) was included in the inequality regressions. Similarly, states with a more educated population should

Table 2.3: Correlation Matrix: Data in Levels

Table 2.3a

	mean ln(adj. gross income)	CV ln(adj. gross income)	mean years of education	CV years of education
mean ln(adjusted gross income)	1.00			
CV ln(adjusted gross income)	0.50 ***	1.00		
mean years of education	0.62 ***	0.31 ***	1.00	
CV years of education	0.03	0.29 ***	−0.05	1.00
tax on wages	0.27 ***	0.21 ***	0.31 ***	0.06 *
tax on interest received	0.35 ***	0.25 ***	0.30 ***	−0.02
tax on dividends received	0.27 ***	0.22 ***	0.29 ***	0.01
tax on long-term gains	−0.48 ***	−0.29 ***	−0.13 ***	0.14 ***
tax on mortgage interest paid	−0.09 **	−0.14 ***	−0.24 ***	−0.08 **
tax on pensions received	0.18 ***	0.16 ***	0.28 ***	−0.03

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Table 2.3b

	tax on wages	tax on interest received	tax on dividends received	tax on long-term gains	tax on mortgage interest paid	tax on pensions received
tax on wages	1.00					
tax on interest received	0.82 ***	1.00				
tax on dividends	0.80 ***	0.92 ***	1.00			
tax on long-term gains	0.01	−0.20 ***	−0.01	1.00		
tax on mortgage interest paid	−0.87 ***	−0.67 ***	−0.64 ***	−0.03	1.00	
tax on pensions	0.84 ***	0.74 ***	0.74 ***	0.02	−0.75 ***	1.00

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Table 2.4: Correlation Matrix: Data in Log Differences

Table 2.4a

	Δ ln(mean ln(adj. gross income))		Δ ln(CV ln(adj. gross income))		Δ ln(mean years of education)		Δ ln(CV years of education)	
Δ ln(mean ln(adjusted gross income))	1.00							
Δ ln(CV ln(adjusted gross income))	−0.10	**	1.00					
Δ ln(mean years of education)	0.39	***	−0.03		1.00			
Δ ln(CV years of education)	−0.07	*	0.06		−0.28	***	1.00	
Δ ln(tax on wages)	0.03		0.04		−0.06		0.10	**
Δ ln(tax on interest received)	−0.03		0.13	***	−0.10	**	0.15	***
Δ ln(tax on dividends received)	−0.09	**	0.04		−0.05		0.11	***
Δ ln(tax on long-term gains)	−0.31	***	0.05		−0.11	***	0.08	**
Δ ln(tax on mortgage interest paid)	0.03		0.02		−0.10	***	0.12	***
Δ ln(tax on pensions received)	−0.07	*	−0.01		−0.13		0.08	*

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

Table 2.4b

	Δ ln(tax on wages)		Δ ln(tax on interest received)		Δ ln(tax on dividends received)		Δ ln(tax on long-term gains)		Δ ln(tax on mortgage interest paid)		Δ ln(tax on pensions received)
Δ ln(tax on wages)	1.00										
Δ ln(tax on interest received)	0.47	***	1.00								
Δ ln(tax on dividends)	0.52	***	0.53	***	1.00						
Δ ln(tax on long-term gains)	−0.14	***	0.11	***	0.12	***	1.00				
Δ ln(tax on mortgage interest paid)	0.89	***	0.40	***	0.34	***	−0.10	***	1.00		
Δ ln(tax on pensions)	0.59	***	0.31	***	0.34	***	0.24	***	0.57	***	1.00

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

have higher average income, so the average number of years of education was included in the average income regressions.

As shown in Tables 2.5, 2.6 and 2.7, the model's predicted positive correlation between income inequality and taxes on labor income appears in the positive coefficient on the tax on wages. The OLS coefficient is also statistically significant from zero when the data is examined in levels (Tables 2.5 and 2.6), but not when examined in differences (Table 2.7). The inclusion of educational inequality does not upset this pattern. Moreover, the ridge trace displays (Figures 2.1, 2.2 and 2.3) show that the coefficient on wage taxes (denoted by "W") retains its positive sign when the coefficients are shrunk towards zero.

If taxation of wage income taxes the returns to labor more than the returns to human capital, then the regression results offer a fair amount of support for the theoretical model's prediction that taxation of labor income increases income inequality. The evidence on the effect of capital income taxation is less encouraging however.

Taxes on interest received and dividends – the two forms of capital income that correspond most closely to the economic definition of returns to capital – are highly correlated both in levels and in differences with each other and with the tax on wages, so it is difficult to disentangle the effect that they have on our measure of inequality. For this reason, it is better to focus on the tax on long-term gains because it exhibits the lowest degree of correlation with the other tax rates.

When the data is examined in levels, the coefficient on the tax on long-term gains is negative – as the model predicts – and the coefficient (denoted by "L") retains its sign as the coefficients are shrunk towards zero. However, the coefficient is only statistically significant from zero when

Table 2.5
OLS Regression on CV of Log Income

Table 2.5 – OLS Regression Results
dependent variable: CV ln(adjusted gross income)

CV years of education			0.073	***			0.072	***
standard error			0.008				0.008	
scaled mean squared error			6.6				6.4	
tax on wages	0.049	***	0.036	**	0.031	***	0.028	***
standard error	0.016		0.015		0.005		0.005	
scaled mean squared error	24.8		25.0		2.8		2.8	
tax on interest received	−0.019		−0.012					
standard error	0.016		0.015					
scaled mean squared error	24.4		24.5					
tax on dividends received	0.030	**	0.026	*				
standard error	0.014		0.013					
scaled mean squared error	20.3		20.3					
tax on long-term gains	−0.025	***	−0.028	***	−0.023	***	−0.027	***
standard error	0.003		0.003		0.003		0.003	
scaled mean squared error	1.2		1.2		0.8		0.8	
tax on mortgage interest paid	0.016		0.021	**				
standard error	0.010		0.009					
scaled mean squared error	9.9		9.9					
tax on pensions received	−0.011		0.001					
standard error	0.009		0.009					
scaled mean squared error	8.3		8.5					
most likely Q-shape	0.0		0.5		0.5		1.0	
F-statistic	22.10	***	36.33	***	99.44	***	101.20	***
observations	663		663		663		663	

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

Figure 2.1: Ridge Trace of OLS Regression on CV of Log Income

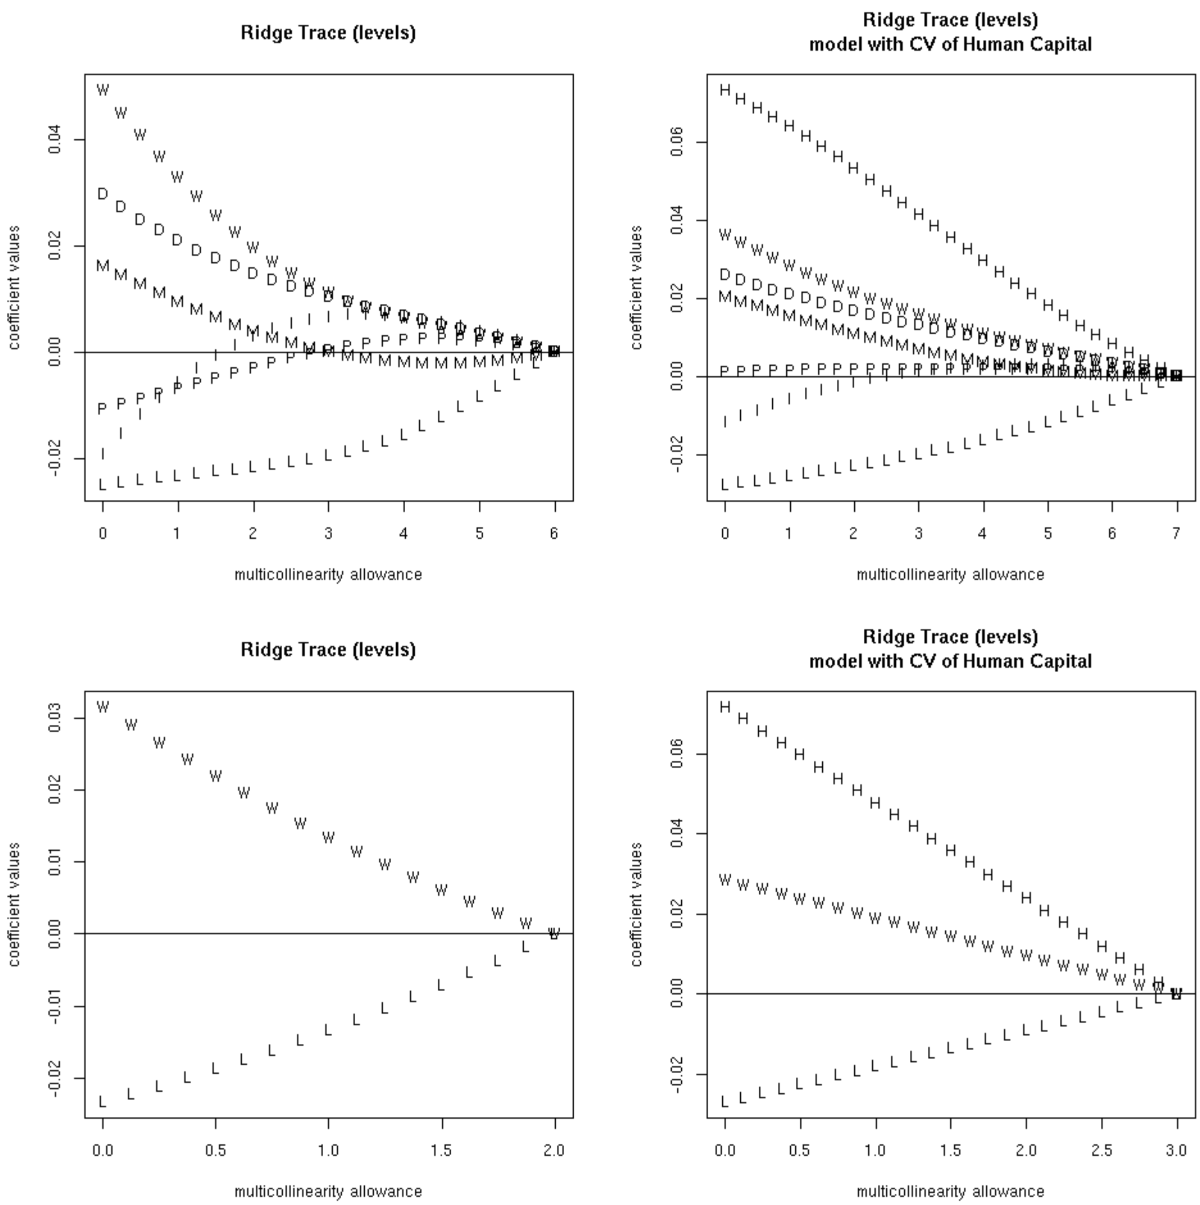

Table 2.6
Fixed Effects Regression on CV of Log Income

Table 2.6 – OLS Regression Results, Fixed State and Year Effects
dependent variable: CV ln(adjusted gross income)

CV years of education			0.025	**			0.025	**
standard error			0.011				0.011	
scaled mean squared error			23.6				23.6	
tax on wages	0.078	*	0.085	**	0.063	*	0.069	**
standard error	0.041		0.041		0.034		0.034	
scaled mean squared error	306.9		308.6		209.9		211.3	
tax on interest received	−0.008		−0.009					
standard error	0.030		0.030					
scaled mean squared error	157.5		157.5					
tax on dividends received	0.002		0.001					
standard error	0.037		0.037					
scaled mean squared error	242.2		242.3					
tax on long-term gains	−0.029		−0.031		−0.031		−0.033	
standard error	0.026		0.026		0.025		0.025	
scaled mean squared error	119.8		119.9		112.3		112.4	
tax on mortgage interest paid	0.033		0.033					
standard error	0.048		0.048					
scaled mean squared error	411.1		411.1					
tax on pensions received	0.001		0.000					
standard error	0.022		0.022					
scaled mean squared error	83.0		83.1					
most likely Q-shape	0.5		1.0		1.0		1.0	
F-statistic	11.53	***	11.50	***	12.32	***	12.28	***
observations	663		663		663		663	

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

The coefficient values and associated statistics for 62 state and year dummy variables are not reported.

Figure 2.2: Ridge Trace of Fixed Effects Regression on CV of Log Income

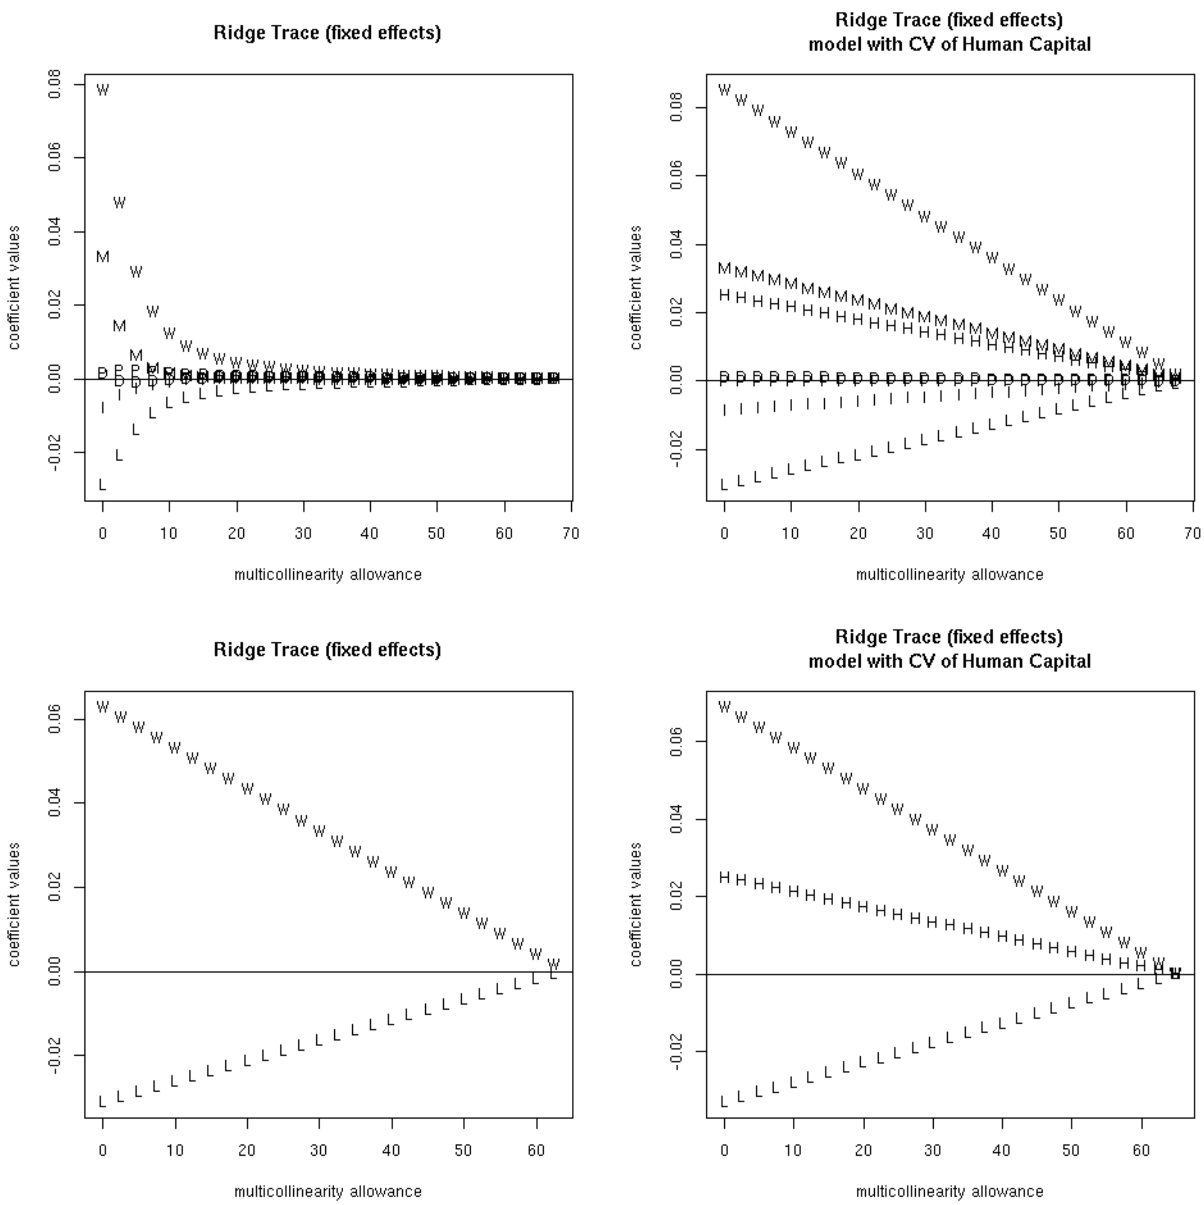

Table 2.7
OLS Regression on Log Differenced CV of Log Income

Table 2.7 – OLS Regression Results
dependent variable: Δ ln(CV ln(adjusted gross income))

Δ ln(CV years of education)			0.038				0.044
standard error			0.034				0.034
scaled mean squared error			0.6				0.5
Δ ln(tax on wages)	0.325		0.340		0.095		0.083
standard error	0.227		0.227		0.084		0.084
scaled mean squared error	24.7		24.8		3.3		3.3
Δ ln(tax on interest received)	0.192	***	0.186	***			
standard error	0.061		0.061				
scaled mean squared error	1.8		1.8				
Δ ln(tax on dividends received)	−0.094		−0.098				
standard error	0.075		0.075				
scaled mean squared error	2.7		2.7				
Δ ln(tax on long-term gains)	0.038		0.036		0.031		0.028
standard error	0.026		0.026		0.023		0.023
scaled mean squared error	0.3		0.3		0.2		0.3
Δ ln(tax on mortgage interest paid)	−0.218		−0.237				
standard error	0.186		0.187				
scaled mean squared error	16.6		16.8				
Δ ln(tax on pensions received)	−0.141		−0.138				
standard error	0.097		0.097				
scaled mean squared error	4.5		4.5				
most likely Q-shape	1.0		1.0		1.0		1.0
F-statistic	3.25	***	2.92	***	2.73	*	2.23
observations	612		612		612		612

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

Figure 2.3: Ridge Trace of OLS Regression on Log Differenced CV of Log Income

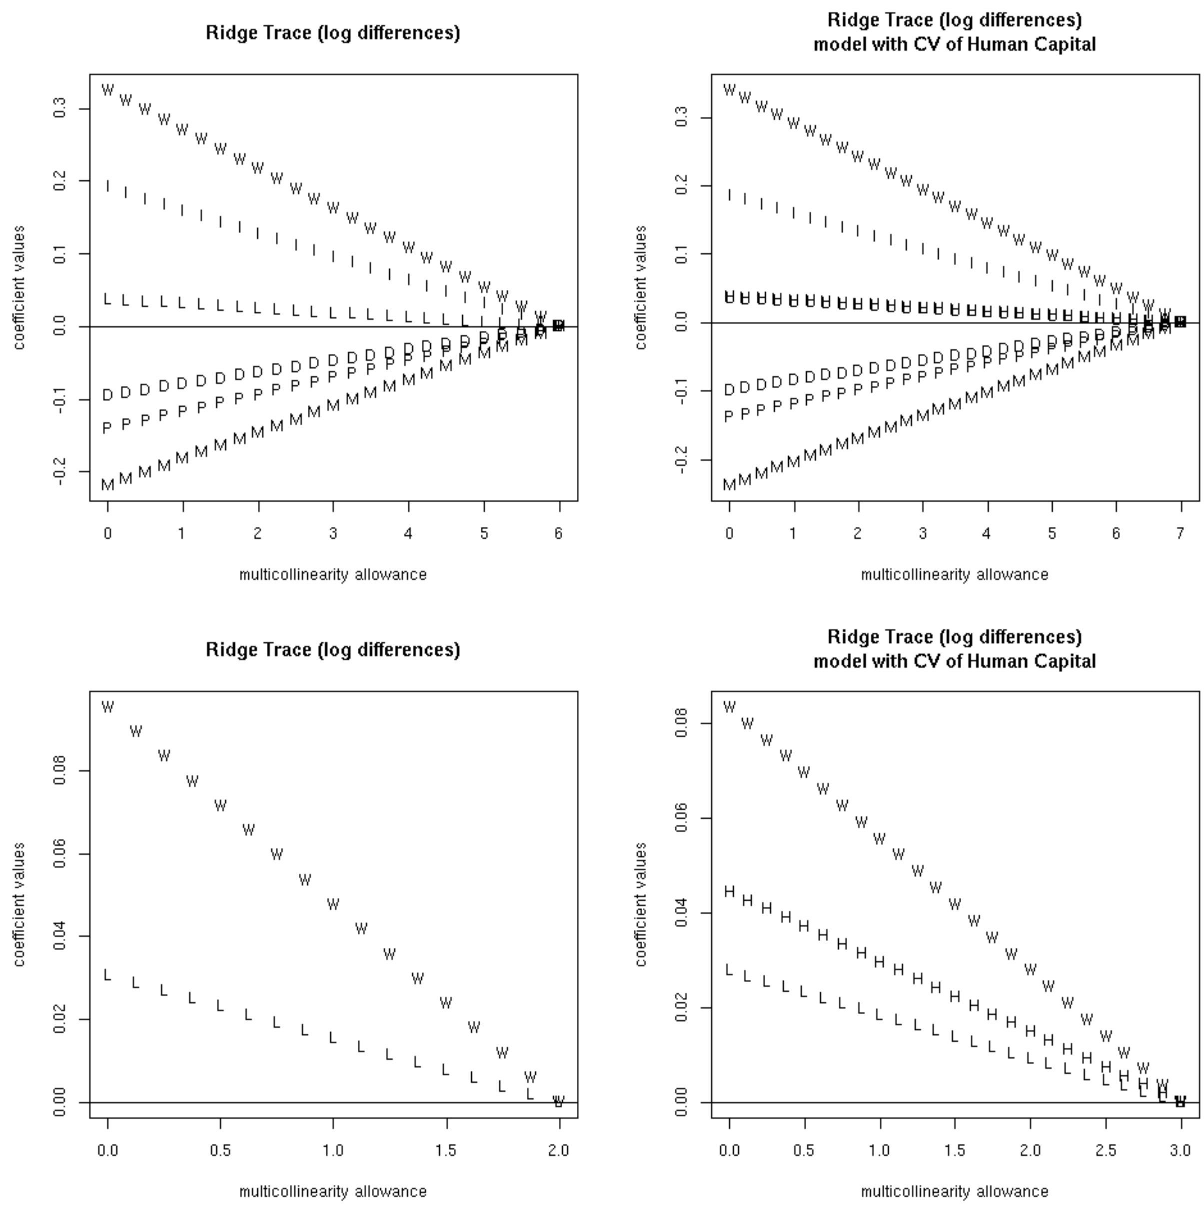

we exclude the fixed state and year effects from the regression model. The coefficient becomes positive (though not statistically significant) when the data is differenced.

The empirical evidence on the relationship between taxation and average income is more difficult to interpret. As mentioned above, the simplifying assumptions of the model present two obstacles. The first is that tax structure may affect households' saving decisions. When saving is taxed (via a tax on the returns to capital), households substitute out of saving (i.e. future consumption) and into present consumption (Musgrave, 1959). Such substitution may reduce the the share of disposable income that households choose to save. If it does, then we may find a negative correlation between average income and average marginal tax rates on capital income.

The assumption that government budgets are always balanced is also problematic. The theoretical model's saving rate is the household saving rate, which is equal to the economy's saving rate when the government budget is balanced. When the government budget is not balanced however, we must add the government saving rate to the household saving rate to obtain the economy's saving rate.

When there are two factors of production, the government's saving rate (expressed as a share of total income) is simply:

$$s'_{gov't} \equiv \frac{T-G}{Y} \equiv (1-\alpha)\,\tau_L + \alpha\tau_K - \frac{G}{Y}$$

where T represents total taxes and G represents government consumption. The economy's saving rate then becomes:

$$s'_{econ} \equiv (1-\alpha)\left(\tau_L + (1-\tau_L)\frac{1}{N}\sum_{i=1}^{N}\sigma_i\right) + \alpha\left(\tau_K + (1-\tau_K)\frac{1}{N}\sum_{i=1}^{N}\sigma_i\frac{k_i}{\bar{k}}\right) - \frac{G}{Y}$$

and is an increasing function of the tax rates, τ_L and τ_K.

Relaxing the assumption of a balanced government budget and the assumption of constant shares of disposable income saved helps explain the empirical relationships between tax rates and average income. Once again, the coefficients on the tax on wages and the tax on long-term gains tend to be statistically significant in most of the regressions and tend to retain their signs across specifications and as the coefficients are shrunk towards zero.

The coefficient on the tax on wages is positive and statistically significant from zero when the data is examined in levels (Tables 2.8 and 2.9), but not when examined in differences (Table 2.10). The positive correlation between the tax on wages and average income contradicts the theoretical model, but relaxing the balanced-budget assumption explains the contradiction.

The coefficient on the tax on long-term gains is negative when the data is examined in levels and when it is examined in differences. It is statistically significant from zero across specifications except when average years of education is excluded from the regression models with fixed state and year effects.

The negative correlation between the tax on long-term gains and average income appears to provide theoretical support for the model, but it is at odds with the correlation that we would

Table 2.8
OLS Regression on Average Log Income

Table 2.8 – OLS Regression Results dependent variable: mean ln(adjusted gross income)								
mean years of education			0.183	***			0.188	***
standard error			0.009				0.010	
scaled mean squared error			0.0				0.0	
tax on wages	3.598	***	2.841	***	1.237	***	0.488	***
standard error	0.416		0.328		0.147		0.123	
scaled mean squared error	24.8		25.1		2.8		3.1	
tax on interest received	−0.076		0.195					
standard error	0.413		0.324					
scaled mean squared error	24.4		24.5					
tax on dividends received	0.309		−0.099					
standard error	0.376		0.296					
scaled mean squared error	20.3		20.4					
tax on long-term gains	−1.193	***	−0.994	***	−1.213	***	−1.034	***
standard error	0.090		0.071		0.080		0.065	
scaled mean squared error	1.2		1.2		0.8		0.8	
tax on mortgage interest paid	2.191	***	1.948	***				
standard error	0.263		0.206					
scaled mean squared error	9.9		9.9					
tax on pensions	−0.450	*	−0.568	***				
standard error	0.241		0.189					
scaled mean squared error	8.3		8.4					
most likely Q-shape	0.5		0.5		0.0		0.5	
F-statistic	83.03	***	181.23	***	295.48	***	421.92	***
observations	663		663		663		663	

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

Figure 2.4: Ridge Trace of OLS Regression on Average Log Income

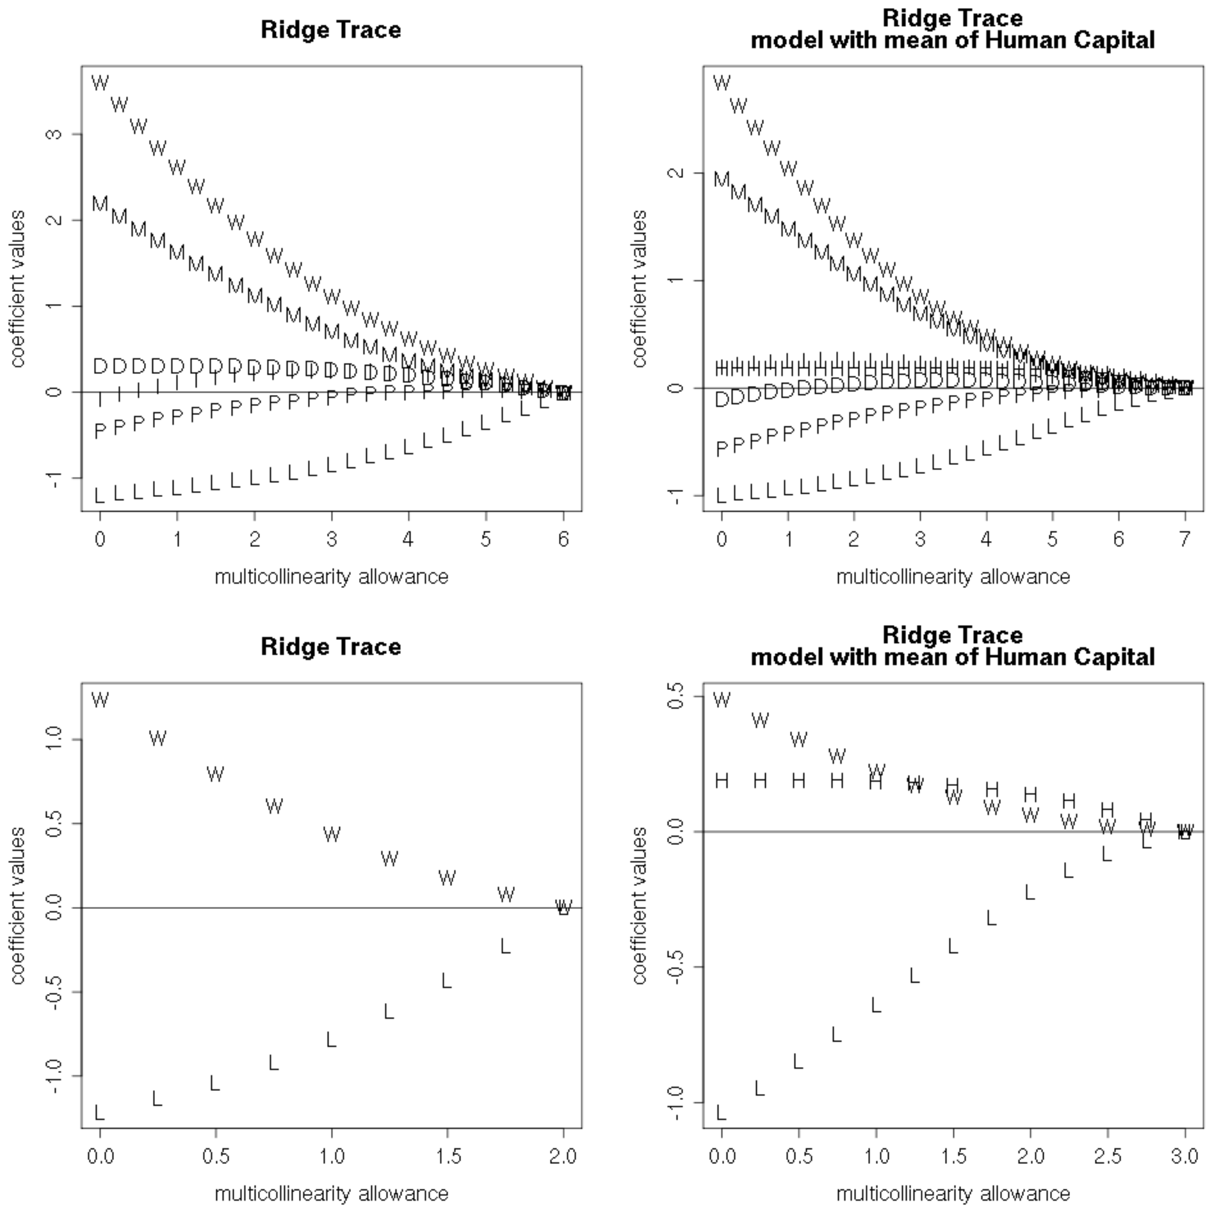

Table 2.9
Fixed Effects Regression on Average Log Income

Table 2.9 – OLS Regression Results, Fixed State and Year Effects
dependent variable: mean ln(adjusted gross income)

mean years of education			0.142	***			0.143	***
standard error			0.012				0.011	
scaled mean squared error			0.1				0.1	
tax on wages	2.170	***	1.188	**	1.487	***	0.750	
standard error	0.669		0.603		0.554		0.497	
scaled mean squared error	306.9		312.4		209.9		212.9	
tax on interest received	−0.565		−0.206					
standard error	0.479		0.429					
scaled mean squared error	157.5		158.3					
tax on dividends received	0.508		0.218					
standard error	0.594		0.531					
scaled mean squared error	242.2		242.7					
tax on long-term gains	−0.592		−0.637	*	−0.572		−0.616	*
standard error	0.418		0.373		0.405		0.361	
scaled mean squared error	119.8		119.8		112.3		112.3	
tax on mortgage interest paid	0.850		0.501					
standard error	0.774		0.692					
scaled mean squared error	411.1		411.8					
tax on pensions received	−0.545		−0.437					
standard error	0.348		0.311					
scaled mean squared error	83.1		83.1					
most likely Q-shape	0.5		−5.0		1.0		−5.0	
F-statistic	67.67	***	85.70	***	71.65	***	91.14	***
observations	663		663		663		663	

*p-value < 0.10, **p-value < 0.05, ***p-value < 0.01

The coefficient values and associated statistics for 62 state and year dummy variables are not reported.

Figure 2.5: Ridge Trace of Fixed Effects Regression on Average Log Income

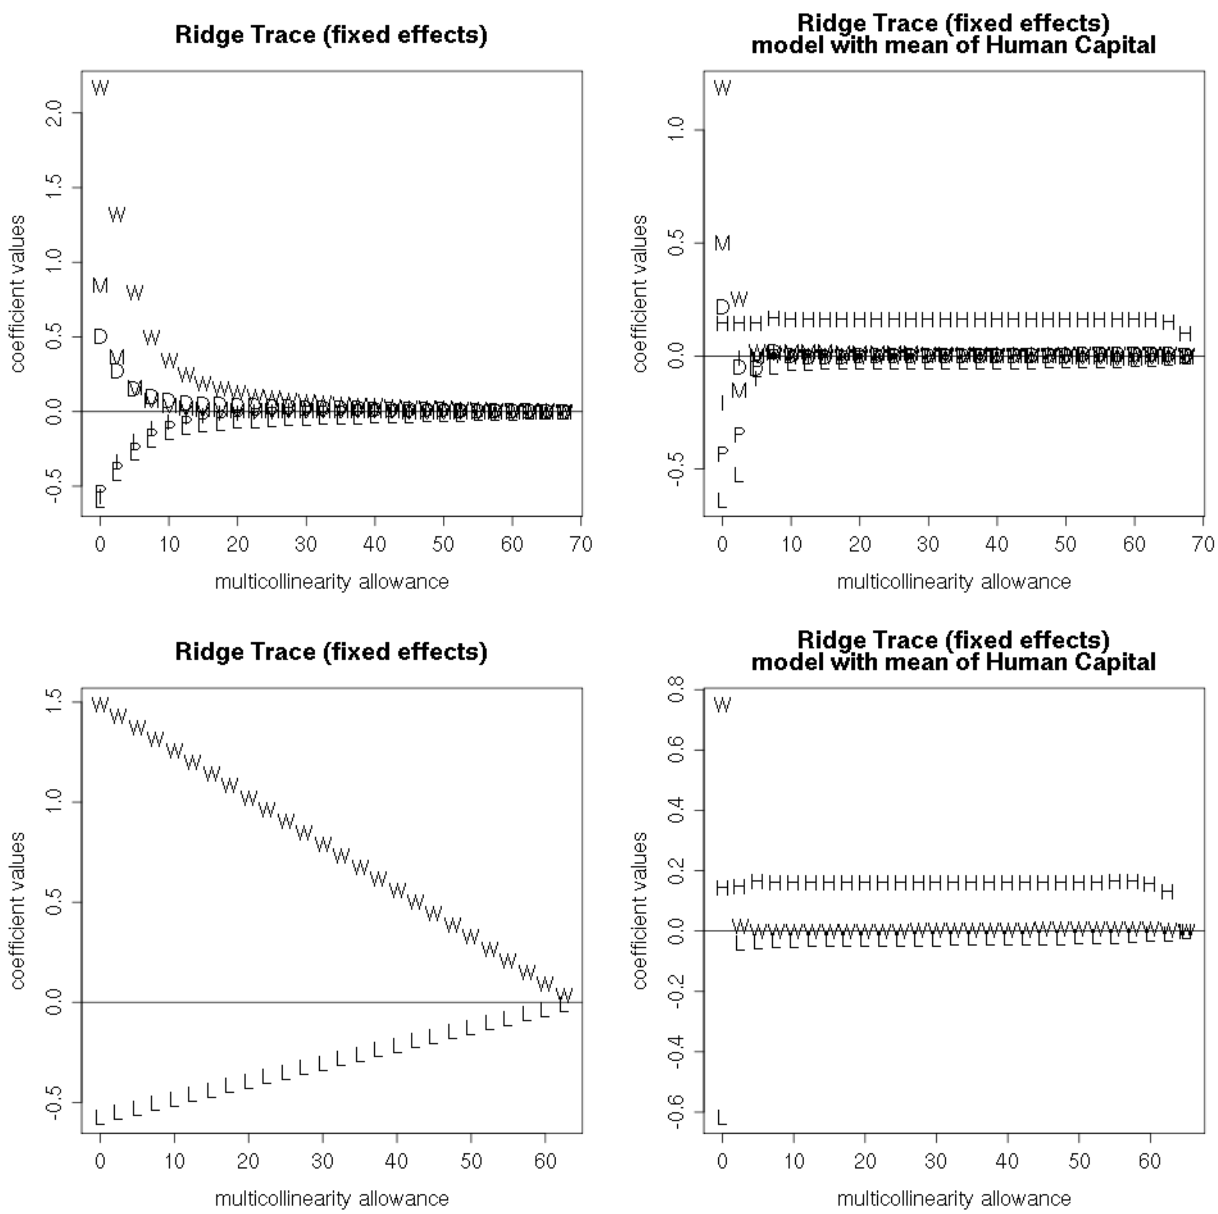

Table 2.10
OLS Regression on Log Differenced Average Log Income

Table 2.10 – OLS Regression Results
dependent variable: Δ ln(mean ln(adjusted gross income))

Δ ln(mean years of educ.)			0.155	***			0.150	***
standard error			0.015				0.015	
scaled mean squared error			16.9				16.6	
Δ ln(tax on wages)	0.001		−0.016		−0.002		0.003	
standard error	0.020		0.018		0.007		0.007	
scaled mean squared error	24.7		24.9		3.3		3.3	
Δ ln(tax on interest rec.)	0.005		0.008					
standard error	0.005		0.005					
scaled mean squared error	1.8		1.8					
Δ ln(tax on dividends rec.)	−0.010		−0.009					
standard error	0.007		0.006					
scaled mean squared error	2.7		2.7					
Δ ln(tax on long-term gains)	−0.016	***	−0.014	***	−0.016	***	−0.014	***
standard error	0.002		0.002		0.002		0.002	
scaled mean squared error	0.3		0.318		0.2		0.3	
Δ ln(tax on mort. int. paid)	0.001		0.017					
standard error	0.016		0.015					
scaled mean squared error	16.6		16.8					
Δ ln(tax on pensions rec.)	0.001		0.006					
standard error	0.008		0.008					
scaled mean squared error	4.5		4.5					
most likely Q-shape	−5.0		−5.0		−1.5		1.0	
F-statistic	13.39	***	30.44	***	64.15	***	87.16	***
observations	612		612		612		612	

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Figure 2.6: Ridge Trace of OLS Regression on Log Differenced Average Log Income

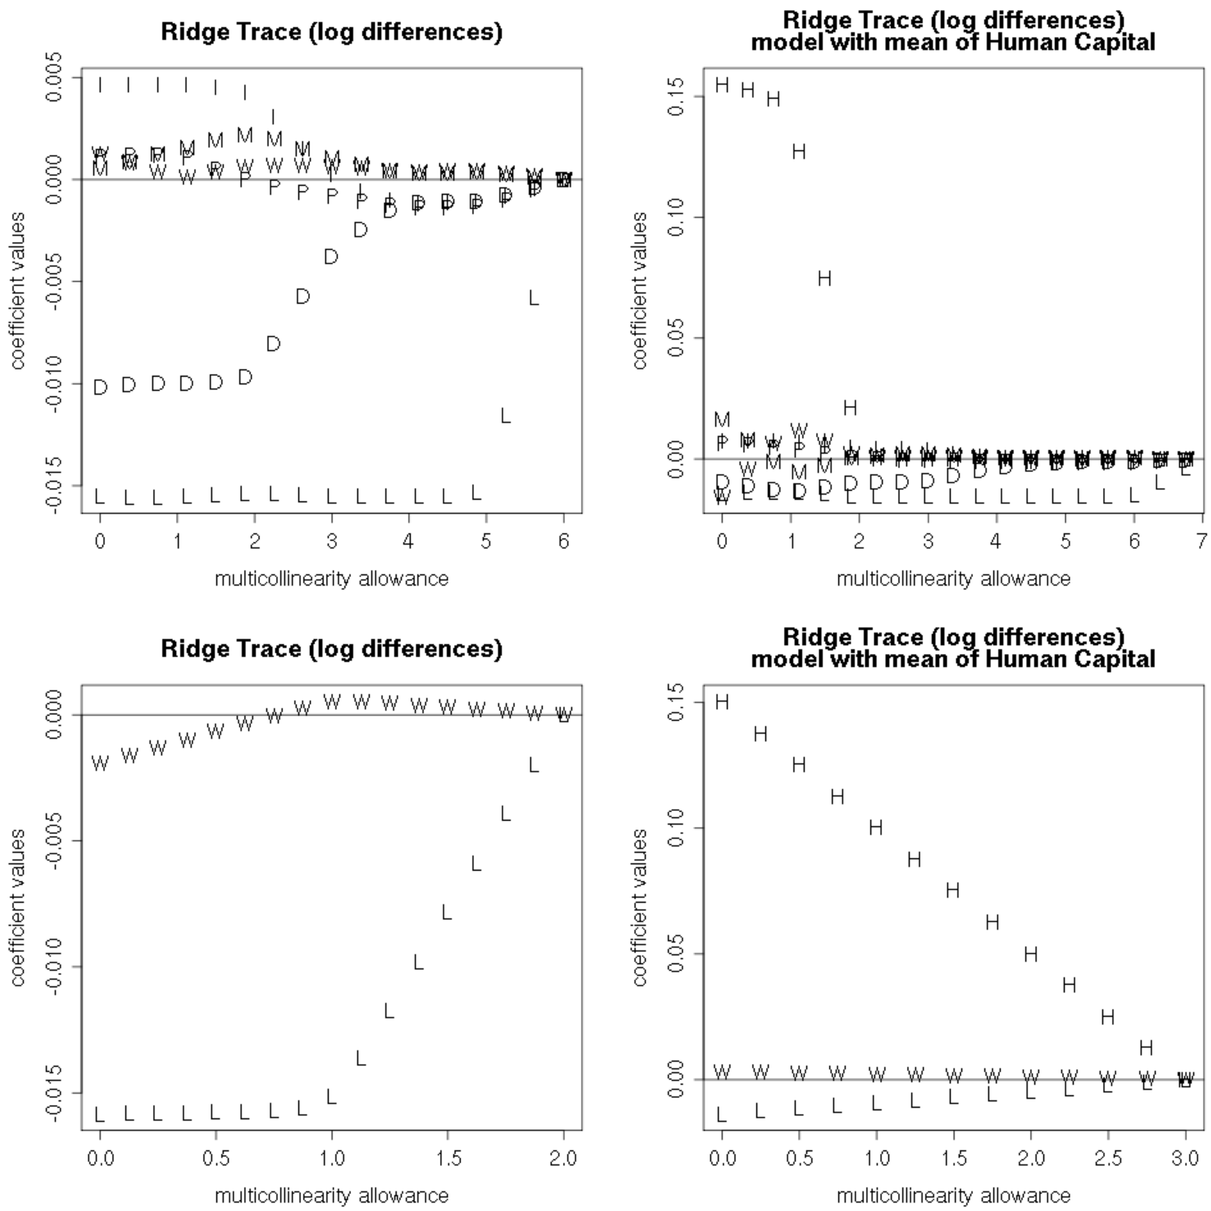

expect to find when we drop the balanced-budget assumption. Therefore we should not seek an explanation from the theoretical model. Instead, the ability of the tax structure to affect households' saving decisions provides the most likely cause of the negative correlation between the tax on long-term gains and average income.

By reducing the relative reward to saving, taxes on long-term gains may induce households to reduce the share of disposable income that they save, thus reducing the economy's saving rate and causing the economy to converge to a lower steady-state level of average income. Some limited support for this explanation can also be found in the negative coefficients on pension income taxation, which are statistically significant when state and year fixed effects are excluded from the regression model.

Finally, it's important to note that the empirical evidence suggests that higher taxes on wage income and lower taxes on long-term capital gains are associated with both higher average income and higher income inequality. Policymakers may therefore want to consider the social implications of higher income inequality before raising taxes on wage income and reducing taxes on long-term capital gains in an attempt to increase the steady-state level of average income.

2.4. Conclusion

Much of the previous literature on the relationship between development and inequality assumes that the distribution of income determines either the rate of economic growth or the level of economic development. Such models focus exclusively on a single tax rate and assume that the adopted tax rate is the one preferred by the median voter (Alesina and Rodrik, 1994; Persson and Tabellini, 1994; Perotti, 1993).

By contrast, this chapter has used Solow's (1956) classic model of economic growth to develop a simple model in which an economy's steady-state level of economic development and its steady-state degree of income inequality are jointly determined. Importantly, the measure of income inequality used in this chapter, the coefficient of variation, is Lorenz-consistent.

This chapter has also explored the effects that taxes on total income, taxes on the returns to capital and taxes on labor income have on development and inequality. Lowering tax rates on the returns to capital increases both an economy's level of economic development and its degree of income inequality. By contrast, lowering tax rates on labor income causes the economy to converge to a higher steady-state level of average income and reduces income inequality.

Several directions for future research are apparent. A more general model would include the income and substitution effects of changes in tax rates on household saving behavior.

Nonetheless, the simple model developed in this chapter finds empirical support from regressions run using data from the March CPS and the NBER's TAXSIM model. The regression results suggest that higher taxes on wage income are associated with greater income inequality. Such a finding is consistent with the theoretical model if we assume that taxation of wage income taxes the returns to labor more than it taxes the returns to human capital.

Discerning the effect that taxation of returns to capital has on inequality is more difficult however because the NBER-TAXSIM's set of average marginal state income tax rates contains a large number of taxes that could be considered taxes on the returns to capital and many of them are highly correlated with each other. However, focusing on the tax rate that is least correlated with other tax variables, the tax on long-term gains, provides some limited support for the theoretical model. Higher tax rates on long-term gains are associated with lower income inequality when fixed state and year effects are excluded from the regression model.

By contrast, the theoretical model's predictions of the effect that taxation has on average income only find support from the empirical evidence when some of the simplifying assumptions are relaxed.

Relaxing the balanced-budget assumption implies that increasing a tax rate will increase the economy's saving rate and cause the economy to converge to a higher steady-state level of average income. This is reflected in the positive correlation between average income and the average marginal tax rate on wages.

Relaxing the assumption that households save a constant share of their disposable income implies that households may increase the share of disposable income that they save in response to a reduction in the tax rate on long-term gains, which causes the economy to converge to a higher steady-state level of average income. This is reflected in the negative correlation between average income and average marginal tax rate on long-term gains.

Worryingly, tax policies that are associated with higher levels of average income are also associated with higher degrees of income inequality.

If we use average health status as a proxy for average standard of living, then one should be concerned about the rising degree of income inequality in America because countries that have greater degrees of income inequality tend to have lower average health status. Babones (2008), for example, finds that higher income inequality is associated with lower life expectancy, higher infant mortality and higher murder rates.

In fact, coupling this chapter's model of inequality and development with the positive correlation between health and schooling (Grossman, 1975) might explain the correlation between inequality and health status. If we interpret capital in the theoretical model as human capital (acquired through education), then the model implies that countries which favor educational investment by low-income people should have lower income inequality and better health status.

It is important therefore to find out how public policy can encourage low-income people to invest in their own human capital. Future research should explore the question of who invests in education, who moves up the income ladder, what characteristics people who move up the income ladder have and if there's a way to "give" those characteristics to people who otherwise would not move up the income ladder.

Chapter 3

The Economic Effects of Health Insurance Rating Restrictions

Since the early 1990s, 48 states have – to varying degrees – reformed their small-group insurance markets with the desire to make health insurance more accessible and affordable to small employers. Pennsylvania and Michigan are the only states that have not. The small group reforms have made health insurance more accessible to older and sicker groups, but the second goal of affordability remains elusive. Instead, health insurance premiums are rising quickly due to rapidly rising spending on health care.

One small-group health insurance market reform, community rating, restricts insurers' ability to set premium rates on the basis of health status, age, gender mix and other factors which predict a group's future medical needs. The only allowable rating factors under a pure community rating regime are: geography, plan design and family composition.

After a brief review of the small-group reforms, this chapter, which was sponsored by the Office of Pennsylvania State Senator Robert C. Wonderling, discusses how the small-group reforms affect the demand for and the supply of health insurance. In particular, this chapter examines the way in which health insurance coverage rates (the percentage of people covered by health insurance) and market concentration (the degree to which one insurer or a few insurers dominate the market) affect each other.

The effect that small group reforms have on consumers is commonly discussed in terms of access and affordability. Unfortunately, access cannot be measured and there are no reliable studies of affordability [1]. The one measure of demand which can be used to study the effect of small group reforms is the coverage rate. Some studies suggest that community rating does not significantly affect the coverage rate. Others, including this one, suggest that community rating reduces the coverage rate.

The problem with using coverage rates to study the demand for insurance is that coverage rates tell us nothing about *who* is covered by insurance. In other words, even though pure community rating reduces coverage rates, a community rating regime may still be desirable if it enables the people who need insurance the most – the old and the sick – to obtain coverage.

On the supply side, pure community rating increases the market share of the largest insurers in a state. Because pure community rating makes coverage more affordable to high-risk individuals, insurers would have to spread risk over a larger number of insured lives if pure community rating

[1] Critics of community rating often point to a GAO (2001b) study which found that health insurance premiums are more expensive in community rated states and therefore less affordable. As discussed in more detail below, the study's measure of premium rates contained considerable bias and, consequently, its conclusion is unreliable.

were enacted. Since community rating will not increase coverage rates, insurers will have to spread risk by capturing a larger market share.

3.1. Background

Before Congress passed the Health Insurance Portability and Accountability Act (HIPAA) in 1996, many states had already enacted a variety of measures to make the insurance market more "consumer-friendly" to small groups. Among the most popular measures were guaranteed issue, guaranteed renewal and limitations on the length of time an insurer could exclude coverage of pre-existing conditions. All of which were incorporated into HIPAA.

As their names suggest, guaranteed issue requires insurers to offer coverage to any group that wants it, while guaranteed renewal prevents an insurer from refusing to renew a policy simply because a group's claims increase. Under HIPAA, insurers can exclude coverage of pre-existing conditions that were diagnosed or treated six months prior to the enrollment date, but only for 12 months after the enrollment date. HIPAA gives states the flexibility to impose more stringent requirements.

As of June 2000, 35 states had restricted insurers' ability to set premium rates on the basis of health status, 12 states had imposed either pure or modified community rating and four states had no restrictions on premium rates[2] (GAO, 2001b, p. 17-18).

[2]The District of Columbia, Hawaii, Michigan and Pennsylvania do not restrict the methodologies used to set premium rates, however the District of Columbia and Hawaii have enacted other small group market reforms.

The states that only restrict variation based on health status generally adopted the National Association of Insurance Commissioners' (NAIC) model law or a variant of it. Under the NAIC model law, an insurer cannot charge one group more than twice the rate it charges another solely on the basis of health status[3].

Under a pure community rating regime, like the one in New York, insurers can only set premium rates on the basis of geography, plan design and family composition. States generally do not mandate pure community rating however. New Jersey, for example, prohibits rating on the basis of health status, but allows insurers to rate on the basis of age, gender and geography, so long as the highest rate charged to a small group does not exceed two times the lowest rate.

3.2. How Restrictions Affect the Demand for Health Insurance

Reducing the number of uninsured workers depends critically on making health insurance more accessible and affordable to small employers because larger employers are far more likely to offer health benefits to their employees than small employers (GAO, 2001a, p. 8) and when small employers do provide benefits they typically pay the same premium that large employers pay, but employees of small businesses receive fewer benefits and face higher out-of-pocket expenses (GAO, 2001b, p. 7).

To evaluate the effect that small group market reforms had on small employers, three criteria are commonly used:

[3] More specifically, the NAIC model requires that premiums charged to similar small employers – called a "class of business" – may not deviate from the average – called the "index rate" – by more than 25 percent and the highest index rate cannot exceed the lowest by more than 20 percent.

- **access –** the range of health insurance options available to older and sicker groups and individuals
- **affordability –** the ability of groups and individuals to pay for health insurance
- **coverage rate –** the percentage of the population covered by health insurance

The small group reforms that states undertook in the last decade have improved access to health insurance, but have not improved affordability. Unfortunately, reliable data and studies on access and affordability are hard to find. Consequently, this chapter must limit its discussion of these two criteria to anecdotal evidence.

The coverage rate, on the other hand, is a criterion on which reliable data exists and it can be used to estimate the effects of small group legislation. This study concludes that pure community rating reduces the coverage rate, while others suggest that community rating will not significantly affect the coverage rate. In either case, coverage rates provide no information on the availability and affordability of health insurance to the highest-risk groups – the groups that need insurance the most.

3.2.1. Access and Affordability

According to survey research conducted by Wonderling (2003b), there is widespread agreement among state insurance officials[4] that guaranteed issue greatly increased the access of older and sicker groups and individuals. A few state officials claimed that rating restrictions also improved

[4] The Office of Pennsylvania State Senator Robert C. Wonderling conducted a telephone survey of all 50 state insurance departments to find out how their small group regulations evolved since HIPAA was passed and to assess their opinions about how the legislation affected their small group markets.

access. In their opinion, guaranteed renewal cannot be enforced without rating restrictions; otherwise insurers can price a group out of the market if just one member of the group gets sick.

While access has improved, there is also widespread agreement among the same officials that the small-group market reforms did nothing to address the issue of affordability.

Some officials claim that rating restrictions have made insurance less affordable. In their opinion, rating restrictions, which force younger and healthier groups to pay higher premiums and subsidize the insurance premiums of older and sicker groups, discourage younger and healthier groups from participating in the insurance pool. As the young and healthy drop out, the insured population becomes older and sicker and, as a consequence, insurance becomes ever less affordable for those who remain.

Other officials claim that increases in medical costs are primarily to blame for the lack of affordability of health insurance premiums, but these officials all concede that it is possible that a decline in coverage among the young and healthy could contribute to rising health care premiums.

This debate is difficult to resolve because affordability is difficult to measure. Measurement of affordability requires knowledge of the price of small-group health premiums, average incomes and average medical price levels. While the Bureau of Economic Analysis publishes quarterly estimates of state personal income and the Centers for Medicare and Medicaid Services has developed a Geographic Practice Cost Index, there are no state-level indices of health insurance premiums.

In a 2001 study, the GAO reported that average single premiums and average family premiums for fully insured small employer plans are, respectively, six and seven percent higher (after adjusting for geographic disparities in provider costs) in states that prohibit insurers from using health status as a rating factor than in states with NAIC rating bands or no restriction at all (GAO, 2001b, p. 19-22).

The GAO's study contained one critical flaw however. In a 1996 study, the GAO concluded that mandated benefits increase the cost of health insurance (1996, p. 8-17). "Cost estimates are higher in states with more mandated benefits and in states that mandate more costly benefits" (1996, p. 4). When, in 2001b, the GAO compared average premiums among states however, it did not consider differences in industry mix, plan type or plan benefit levels (2001b, footnote no. 25, p. 21).

The lack of adjustment for plan benefit levels introduces upward bias into the average premium in community rated states. Five of the 12 states which prohibit rating on health status also mandate over 20 specific benefits. One of those five, Maryland, mandated 39 specific benefits – the highest in the country (GAO, 1996, p. 9).

3.2.2. How Community Rating Affects Coverage

While the effect of small-group reforms on access and affordability has not been reliably measured, coverage rates[5] can be used to develop an understanding of how small group reforms affect the demand for insurance.

Because uninsured individuals often work for small businesses, the empirical research in this study focuses heavily on employment-based coverage rates[6]. The regression results reported be-

[5] The March Current Population Survey (CPS), which is conducted annually by the Bureau of the Census, asks households if they had insurance in the last year and if so, from what source they obtained insurance (i.e. public programs or private insurance; if they had private insurance, they were also asked if they obtained it from an employer).

[6] Admittedly, the employment-based coverage rate includes many individuals who are enrolled through a large group, but if the employment-based coverage rate is highly correlated with the unknown rate of coverage through small businesses, it can be used to estimate the effect of small-group market reforms on health insurance coverage.

low indicate that imposing restrictions on insurers' use of age, health and gender when setting premiums reduces employment-based coverage rates, but restricting insurers' ability to rate on the basis of industry and group size increases coverage rates. (The degree to which the coverage rates are affected depends of course on the tightness of the rating restrictions).

While the empirical evidence discussed in this chapter suggests that restricting rating on the basis of age and health status will adversely affect employment-based coverage rates, the data is unable to tell us anything about how demographic rating, community rating or NAIC rating bands affect the age and health distributions of groups that have coverage.

It is however reasonable to assume that restricting insurers' ability to set premium rates on age and health status makes health insurance more affordable to older and sicker groups (the people who need coverage most). Although community rating may be good social policy, Feldman (1987) points out that the price subsidization implicit in community rating (where younger, healthier workers subsidize the premium rates of older, sicker workers) is inefficient and that simply transferring income from healthy to sick workers would make each group at least as well off.

Moreover, if they are risk adverse, then younger, healthier people would not want to completely forgo health insurance coverage. In fact, some evidence from the "prospect theory" of utility suggests that individuals will not drop coverage at all. According to prospect theory, an individual's utility depends on perceived gains and losses as well as on the risky and riskless components of a decision, both of which the individual evaluates differently.

Using data from the RAND Health Insurance Experiment, Marquis and Holmer (1996) tested prospect theory and found that there is some inertia to individuals' choice of health insurance plans.

The plan that they currently hold is dearer to them than an equivalent plan with which they have no experience. They also found that individuals who are fully insured require a larger reduction in premium to induce them to switch to a plan which provides them with less coverage than the premium increase individuals who are not fully insured would accept to switch to a fully insured plan.

Marquis and Holmer's study suggests that a move to community rating won't cause coverage rates to fall significantly in the near-term even if premiums rise, but their study does not suggest that community rating will have no long run effect on coverage.

One state insurance official told Wonderling (2003b) that the age distribution of the insured population could increase even if no young and healthy workers drop coverage. In his opinion, if the rating restrictions encourage younger workers to wait longer before entering their first health insurance plan, the age distribution could increase through attrition.

That's a particular danger because increases in the age distribution of insured individuals could cause an "adverse selection death spiral."

Adverse selection arises in the health insurance industry when an insurer is unable to distinguish between people who are naturally more susceptible to disease than others. Community rating coupled with guaranteed issue and guaranteed renewal force this situation upon insurers.

A death spiral occurs when the coverage rate falls and the average age of policy-holders rises, pushing up premiums and discouraging healthy people from buying insurance until the only people who will pay for insurance are those who have the highest expected medical expenses.

Buchmueller and Dinardo (1999) tie together the questions of whether or not community rating will cause a death spiral and whether or not a risk-adverse individual would forgo health insurance coverage. They hypothesized that under a community rating regime:

- a risk-adverse individual would seek some coverage, but might not choose to be fully insured since he cannot buy insurance at an actuarially fair price and
- to avoid a death spiral, younger low-risk "individuals [must] not abandon coverage altogether, but purchase (cheaper) less complete coverage" (1999, p. 8-9).

Because New York and Connecticut reformed their small group and individual markets in 1993 while Pennsylvania did not and because Connecticut's reforms were milder than New York's[7], Buchmueller and Dinardo saw in these three states a natural experiment.

Examining evidence from the small group and individual markets in Pennsylvania, New York and Connecticut they found that after 1993 the percentage of covered individuals in small groups did not fall in New York relative to Pennsylvania or Connecticut. They also found that the age distribution of adults with both small group and individual health insurance shifted toward older individuals in New York, but no more than in Pennsylvania (1999, p. 14-19).

To explore the question of whether individuals are buying cheaper, less complete coverage, Buchmueller and Dinardo assumed that HMOs' rationing of care and their limitation of a patient's choice of doctors and treatment options represents the purchase of cheaper, less complete coverage. They found that New York's reforms significantly affected the structure of its insurance market.

[7]New York and Connecticut both enacted community rating, guaranteed issue and guaranteed renewal in the small group and individual markets. The main difference between the two was that New York imposed pure community rating while Connecticut's modified community rating places no restrictions on age and gender rating.

After 1993, HMO penetration of the small group and individual markets increased in New York far more than it did in Pennsylvania and Connecticut (1999, p. 21-22).

Responding to the critics of community rating, they pointed out that if some groups and individuals are switching from traditional coverage to HMO plans, some insurance companies may observe evidence of an adverse selection death spiral, even though no death spiral is actually occurring (1999, p. 9). They concluded that community rating did not reduce the percentage of insured individuals in New York, but the reforms did not succeed in increasing the percentage of insured individuals either (1999, p. 23).

The main advantage of Buchmueller and Dinardo's research was its test of how consumers respond to a change in the price of insurance. Unfortunately, their research only compared three states.

In a broader study, Herring and Pauly (2006) examined the effect that an individual's risk of incurring medical expenditures has on the premium rates that they pay for coverage through the individual health insurance market, the effect of community rating on premium rates and the effect that rating restrictions have on an individual's probability of obtaining coverage through the individual market.

Although they find that high-risk individuals in states without rating restrictions tend to pay higher premiums than comparable high-risk individuals in states with rating restrictions, the difference is not statistically significant. They also found that high-risk individuals in states without rating restrictions tended to have a lower probability of purchasing insurance than comparable high-risk individuals in states with rating restrictions.

Herring and Pauly did not find a strong positive correlation between an individual's risk status and his/her probability of obtaining coverage however and, on the basis of this finding, they also conclude that community rating does not cause an "adverse selection death spiral." They hypothesize that community rating does not substantially reduce the coverage rate because guaranteed renewal enables high-risk individuals who obtained coverage prior to becoming high-risk to retain premium rates that do not reflect their health status. Consequently, a high degree of pooling occurs in the absence of community rating.

The use of microdata to study the effect of community rating provides good insight into the effects that health insurance market regulation has on the premium rates and the probability of obtaining coverage, however Herring and Pauly's study does not provide much insight into the reasons why high-risk individuals face higher premiums and have a lower probability of obtaining coverage.

As discussed in the next section, premiums may rise in community rated states if rating restrictions increase market concentration in the insurance industry. Unfortunately, Herring and Pauly did not study the supply of health insurance.

3.3. How Restrictions Affect the Supply of Health Insurance

Switching from demographic rating to community rating doesn't just affect coverage rates, it also affects the insurance industry, possibly forcing some small insurers out of the market and increasing the market share of the five largest insurers.

Because insurers' average costs decline as they issue more policies, insurers face increasing returns to scale. To explain the effect that restrictions have on the insurance industry, this section will first explain how increasing returns to scale apply to the insurance industry and its relevance to rating restrictions. Then this section will then discuss how rating restrictions affect market concentration.

3.3.1. Increasing Returns to Scale

When an insurer facing increasing returns to scale doubles the value of its outstanding policies, its risk increases less than twofold because the insurer's expected loss ratio (the ratio of claims paid to premiums earned) converges toward the true average, which reduces the risk that an insurer will have to pay an unexpectedly high amount of claims.

Small and medium-size insurers both face increasing returns to scale, while large insurers face constant returns to scale because large insurers' loss ratios have already achieved a high degree of convergence to the true population average.

Community rating forces insurers to reduce the premiums that higher-risk individuals pay for insurance coverage and raise the premiums that lower-risk individuals pay for insurance. Community rating therefore encourages higher-risk individuals to buy coverage and may discourage lower-risk individuals from buying coverage, thus increasing risk to the insurer.

To maintain the same level of profitability after the enactment of community rating, insurers that face increasing returns to scale must reduce their average costs (in terms of risk) by spreading risk over a larger pool of individuals. If the demand for insurance remains constant or falls after the enactment of community rating, insurers must spread risk by capturing a larger market share, thus increasing the degree of market concentration as small insurers either merge with other insurers or drop out of the market (Chollet et al., 2000).

Greater market concentration is not necessarily undesirable, if the market remains competitive enough that insurers' lower average costs (in terms of risk) render lower premiums and greater coverage. Alternatively, the greater degree of market power associated with larger market share may enable insurers to raise their premiums. This "monopoly" effect could outweigh the effect that increased returns to scale has on premium rates and (all else equal) reduce the number of insured individuals.

3.3.2. How Community Rating Affects Market Concentration

Many state insurance officials told Wonderling (2003b) that some insurers dropped out of their markets after small-group reforms passed. This was also true of Pennsylvania however, which lost three plans even though its reforms were no more stringent than those required by HIPAA. Only one official told them that small group reforms leveled the playing field among insurers and created a more competitive market, but several said that their markets remained viable despite the decline in the number of insurers.

In an study of the effect that regulation of pre-existing exclusion periods has on market concentration, Chollet et al. (2000) found that higher market concentration in the small-group market is associated with longer pre-existing condition exclusion periods, but they did not find comparable statistically significant results for the individual market and on that basis they concluded that shortening pre-existing condition exclusion periods does not encourage individuals to drop coverage until their insurance needs change.

3.4. How Coverage and Concentration Affect Each Other

The employment-based coverage rate might also affect market concentration. Higher coverage rates could increase the market share of the largest insurers if employers on the margin of purchasing insurance purchase policies disproportionately from large insurers. Conversely, higher coverage rates could decrease market concentration if small insurers sell disproportionately more marginal policies than large insurers.

The latter is more likely to be the case. Due to increasing returns to scale, each extra policy a small insurer issues reduces the average cost (in terms of risk) of its outstanding policies more than a large insurer reduces its average cost by issuing an additional policy. Therefore small insurers have greater incentive to seek new business than large insurers do.

Market concentration might also affect coverage. In a competitive small group insurance market, employers compare premium rates among insurers and select the best offer. As the largest insurers gain increasingly larger shares of the market however, they gain monopoly power and become increasingly able to charge high premium rates that discourage employers from offering their employees a health plan, thus reducing the employment-based coverage rate.

3.5. Empirical Evidence

This section presents the methodology of the empirical study of small group reforms. The study was conducted with two goals in mind:

- to predict the effect that a move towards or away from community rating will have on a state's employment-based coverage rate and insurance industry and
- to develop an empirical framework that enables policymakers to understand how small group reforms[8] have on employment-based coverage rates, market concentration and the interaction between employment-based coverage rates and market concentration.

Of particular interest, the results indicate that restricting insurers' ability to rate on the basis of age, gender and health status is associated with lower employment-based health insurance coverage rates, while restricting insurers' ability to rate on the basis of group size and industrial classification is associated with higher coverage rates. However, the only statistically significant variables in the regressions on coverage rates are rating restrictions on health status and rating restrictions on group size.

The regression results also indicate that restricting insurers' ability to rate on the basis of age and health status increases the market share of the state's five largest insurers, while rating restrictions on the use of gender, group size and industrial classification reduces market concentration. Once again, not all of the regulatory variables have a statistically significant effect on market

[8]Small group reforms include: community rating, NAIC rating bands and limiting pre-existing condition exclusion periods.

concentration. The ones that do are rating restrictions on health status and rating restrictions on industrial classification.

Longer limits on the length of time during which insurers may deny coverage of pre-existing conditions increases market concentration. Finally, higher coverage rates are associated with lower market concentration.

3.5.1. Data

This study analyzes both the supply side and the demand side of the health insurance market with data from 49 states[9] and the District of Columbia between 1997 and 2002. It estimates the effect that variations in employment, rating restrictions and pre-existing condition exclusions have on the employment-based coverage rate and the market share of the five largest insurers in each state.

To perforn the analysis, this study uses annual data on health insurance coverage rates, the insurance industry, employment, population and small group regulation. The data was obtained from a variety of sources. The data on state population and health insurance coverage rates were obtained from the US Census Bureau's Historical Health Insurance Tables. The number of insurers in each state and the total premiums earned and claims paid by each insurer was obtained from data that the NAIC provided to Wonderling (2003a). Data on employment was obtained from the US Bureau of Labor Statistics.

[9] Wonderling (2003b) was unable to obtain reliable information on the small group market in Illinois.

Finally, data on regulation of the small group market was compiled by Wonderling (2003b) from a telephone survey of all 50 state insurance departments and from information that the NAIC and the Blue Cross and Blue Shield Association provided to him.

3.5.2. Dependent Variables

To measure the effect of small-group market regulation on health insurance coverage and market concentration, the two dependent variables used in the regression analysis were the percentage of people under age 65 covered by employment-based private health insurance and the market share of the five largest insurers. Because these variables are percentages, they were expressed as the log of an odds ratio and analyzed with two-step weighted least squares logistic regression.

Admittedly, the Census Bureau's employment-based coverage rate includes many individuals who are enrolled through a large group. Assuming however that small business constitutes a relatively constant share of private sector employment from state to state, the employment-based coverage rate should be highly correlated with the unknown rate of coverage obtained through small employers because:

- the large-group health insurance market is not regulated by the individual states and therefore the rate of coverage through large employers should be relatively constant from state to state and because
- almost all private businesses that employ more than 50 employees provide health benefits, so variation in employment-based coverage rates among states should be explained by variations in regulation, employment, demography and market concentration among the states.

In section 3.4, it was observed that market concentration and employment-based coverage should affect each other. The degree of market concentration should affect the employment-based coverage rate if a low degree of market concentration lowers premium rates and induces more employers to purchase or retain coverage for their employees. Coverage rates should affect the degree of market concentration if small insurers sell disproportionately more new policies than large insurers and face a disproportionately large share of lost opportunities for new business when the demand for insurance dries up.

To account for these twin possibilities, two-stage least squares was employed in the regression analysis.

3.5.3. Explanatory Variables

To explain inter-state variations in coverage and market concentration, the regression analysis used variations in regulation, employment and demography.

The regulatory variables of particular interest were the restrictions on the use of rating factors when setting premium rates. Specifically, restrictions on the use of health status, age, sex, industrial classification and group size were considered. The measure of restrictiveness was the minimum allowable rate divided by the maximum allowable rate. For example, if a state prohibits variations in premium rates on the basis of age in excess of a four-to one-ratio, one was divided by four to get 0.25. If the state has no rating restriction, the variable was set to zero. Conversely, if the state requires insurers to use a community rate, the variable was set to one.

For NAIC model legislation, the permissible downward variation from the index rate was divided by the permissible upward deviation from the index rate to calculate the rate variation within a class of business, then that number was divided by one over the maximum rate variation between classes of business. For example, if rates must fall within 25 percent of the index rate within a class of business and the highest index rate may not exceed lowest by more than 20 percent, then: $(0.75/1.25) \cdot (1/1.2) = 0.5$.

Some states also impose a constraint on overall rate variation. To account for this feature of state regulation, a composite variable was created that is the product of all rating restrictions. For states that do not impose a restriction on a rating factor, the value of that factor and the composite variable were both set to zero.

For example, a state might prohibit rating on the basis of health, but allow insurers to use any other rating factor they like so long as the overall rate variation does not exceed a two-to-one ratio. In such cases, the composite variable was set equal to one half and the variables describing the restrictiveness on use of age, sex, industrial classification and group-size were each set equal to the fourth root of one half. Although specific bands on individual rating factors do not exist in such cases, the existence of a composite band imposes an effective band on other rating factors.

This approach has the drawback of increasing the correlation among the regulatory variables; however ignoring the effect of a composite rate band on insurers' ability to set premium rates would misrepresent the policy decisions of state legislatures.

Another regulatory variable used in the regression is the total number of months that pre-existing conditions can be excluded from health care coverage. The value of the variable is the number of months before enrollment when such conditions are treated or diagnosed plus the number of months from the date of enrollment that the condition can be excluded from coverage.

Because one would expect to find more insurers (and therefore less market concentration) in more populous states, the log of the population under age 65 was used in the regressions.

Similarly, one would expect to find higher employment-based coverage rates in states where a larger fraction of the population under age 65 is employed. For example, suppose that two states have the same number of people under age 65, but children constitute a relatively larger share of the population in one of the states. The cost of health insurance would impose a relatively larger burden on the working population in the state with more children.

One would therefore expect to find lower rates of employment-based coverage in states that have a lower total employment as a fraction of the total population under age 65. Although this variable is not the labor force participation rate, it is closely related to it and (for better or for worse) is dubbed the “participation rate” in the regression tables.

One might also suspect that consumers have better access to insurers when they are located in more densely populated areas. Market concentration might also be lower in more densely populated states if urban areas can support more insurers. To test the hypotheses that higher population density is associated with higher coverage rates and less market concentration, the log of population density (as measured by the state’s population divided by its total land area in square miles) was included in the regressions.

Table 3.1
Correlation Matrix

Table 3.1a

	employ.-based coverage		"top 5" market share		rating restr. on health		rating restr. on age	
employ.-based cov.	1.00							
"top 5" mkt. share	−0.27	***	1.00					
restr. on health	−0.04		0.14	**	1.00			
restr. on age	−0.07		0.07		0.64	***	1.00	
restr. on sex	0.02		0.10		0.74	***	0.83	***
restr. on industry	0.05		0.07		0.64	***	0.58	***
restr. on group-size	0.04		0.02		0.79	***	0.78	***
months pre-ex. exc.	−0.26	***	−0.04		−0.46	***	−0.33	***

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Table 3.1b

	rating restr. on sex		rating restr. on industry		rating restr. on group-size		months pre-existing excluded
restr. on sex	1.00						
restr. on industry	0.65	***	1.00				
restr. on group-size	0.82	***	0.66	***	1.00		
months pre-ex. exc.	−0.43	***	−0.33	***	−0.42	***	1.00

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Table 3.2
Logit Regression on Employment-Based Coverage Rate

Table 3.2 – Two-Step Weighted Least Squares Logit Model
Dependent Variable: Employment-Based Coverage Rate (Expressed as Log of Odds Ratio)

	"naive" regression		reduced form		Hausman test		2SLS structural	
restr. on health rating	−0.264	***	−0.344	***	−0.228	***	−0.229	***
standard error	0.071		0.069		0.071		0.071	
scaled mean squared error	0.009		0.009		0.010		0.010	
restr. on age rating	−0.081		−0.106		−0.078		−0.079	
standard error	0.067		0.067		0.066		0.067	
scaled mean squared error	0.008		0.008		0.008		0.008	
restr. on sex rating	−0.014		−0.006		−0.023		−0.032	
standard error	0.064		0.063		0.063		0.063	
scaled mean squared error	0.008		0.008		0.008		0.008	
restr. on industry rating	0.059		0.086	**	0.057		0.057	
standard error	0.039		0.039		0.039		0.039	
scaled mean squared error	0.003		0.003		0.003		0.003	
restr. on group-size rating	0.123	**	0.137	**	0.104	*	0.110	*
standard error	0.056		0.055		0.056		0.056	
scaled mean squared error	0.006		0.006		0.006		0.006	
months can exc. pre-exist.	−0.013	***	−0.014	***	−0.012	***	−0.012	***
standard error	0.003		0.003		0.003		0.003	
scaled mean squared error	0.000		0.000		0.000		0.000	

continued on the next page

Table 3.2 continued – Two-Step Weighted Least Squares Logit Model

Dependent Variable: Employment-Based Coverage Rate (Expressed as Log of Odds Ratio)

	"naive" regression		reduced form		Hausman test		2SLS structural	
"top 5" share (log odds)	−0.019	***						
standard error	0.004							
scaled mean squared error	0.000							
predicted "top 5" share					−0.026	***	−0.024	***
standard error					0.005		0.004	
scaled mean squared error					0.000		0.000	
residual of "top 5" share					−0.010	*		
standard error					0.005			
scaled mean squared error					0.000			
log of pop. under age 65			0.070	***				
standard error			0.013					
scaled mean squared error			0.000					
"participation rate"	4.154	***	4.707	***	4.215	***	4.258	***
standard error	0.337		0.347		0.334		0.335	
scaled mean squared error	0.216		0.233		0.217		0.216	
log of pop. density	0.015	*	0.023	***	0.008		0.007	
standard error	0.008		0.008		0.008		0.008	
scaled mean squared error	0.000		0.000		0.000		0.000	
most likely Q-shape	0.5		1.0		0.5		0.5	
F-statistic	32.5	***	33.47	***	30.47	***	33.47	***
observations	282		282		282		282	

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Figure 3.1: Ridge Trace of Logit Regression on Employment-Based Coverage Rate

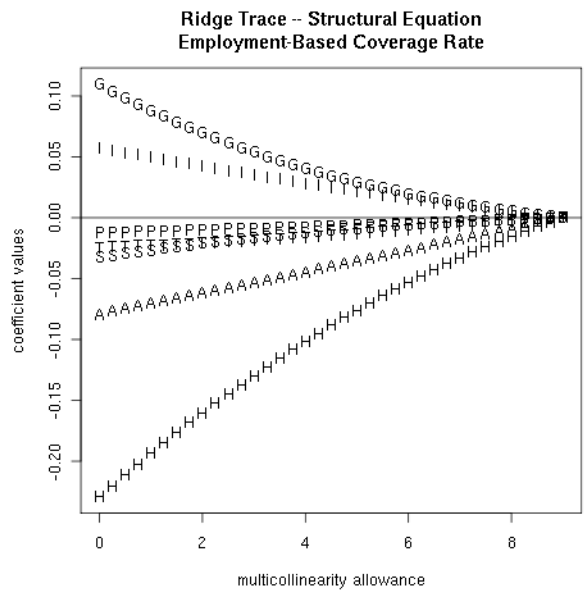

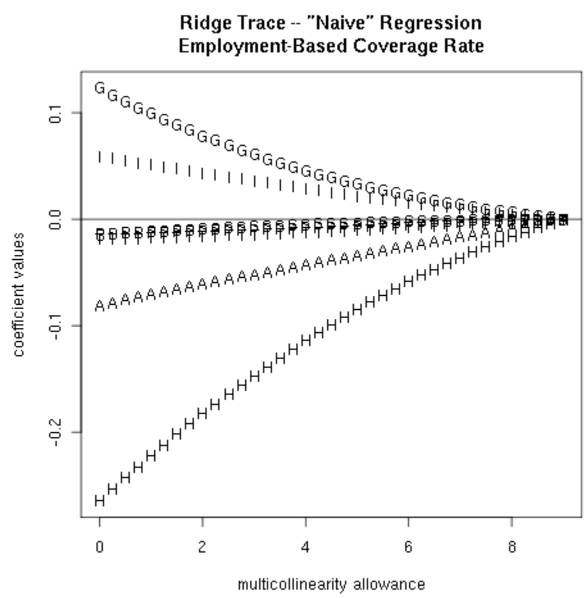

Table 3.3
Logit Regression on Market Share of Five Largest Insurers

Table 3.3 – Two-Step Weighted Least Squares Logit Model

Dependent Variable: Market Share of Five Largest Insurers (Expressed as Log of Odds Ratio)

	"naive" regression		reduced form		Hausman test		2SLS structural	
restr. on health rating	3.428	***	4.788	***	3.061	***	3.423	***
standard error	0.906		0.879		0.928		0.933	
scaled mean squared error	1.612		1.519		1.629		1.606	
restr. on age rating	0.643		1.032		0.422		0.612	
standard error	0.853		0.848		0.865		0.871	
scaled mean squared error	0.947		0.968		0.956		0.998	
restr. on sex rating	−0.733		−1.109		−1.057		−1.131	
standard error	0.803		0.804		0.798		0.803	
scaled mean squared error	0.937		0.876		0.884		0.876	
restr. on industry rating	−1.118	**	−1.179	**	−0.790		−0.837	*
standard error	0.504		0.499		0.502		0.505	
scaled mean squared error	0.348		0.322		0.344		0.337	
restr. on group-size rating	−0.524		−1.093		−0.293		−0.549	
standard error	0.713		0.704		0.713		0.717	
scaled mean squared error	0.710		0.634		0.673		0.660	
months can exc. pre-exist.	0.052		0.075	**	0.009		0.019	
standard error	0.038		0.036		0.040		0.041	
scaled mean squared error	0.002		0.002		0.002		0.002	

continued on the next page

Table 3.3 continued – Two-Step Weighted Least Squares Logit Model

Dependent Variable: Market Share of Five Largest Insurers (Expressed as Log of Odds Ratio)

	"naive" regression		reduced form		Hausman test		2SLS structural	
emp.-based cov. (log odds)	−2.217	***						
standard error	0.597							
scaled mean squared error	0.592							
predicted emp.-based cov.					−4.091	***	−3.974	***
standard error					0.932		0.937	
scaled mean squared error					1.365		1.306	
residual of emp.-based cov.					−0.685			
standard error					0.765			
scaled mean squared error					0.858			
log of pop. under age 65	−2.739	***	−2.875	***	−2.611	***	−2.596	***
standard error	0.159		0.163		0.158		0.159	
scaled mean squared error	0.083		0.084		0.078		0.077	
"participation rate"			−18.705	***				
standard error			4.412					
scaled mean squared error			28.941					
log of pop. density	−0.460	***	−0.659	***	−0.498	***	−0.567	***
standard error	0.105		0.104		0.105		0.106	
scaled mean squared error	0.021		0.022		0.021		0.021	
most likely Q-shape	1.0		1.0		1.0		1.0	
F-statistic	25.68	***	26.23	***	24.1	***	26.23	***
observations	282		282		282		282	

*p-value $<$ 0.10, **p-value $<$ 0.05, ***p-value $<$ 0.01

Figure 3.2: Ridge Trace of Logit Regression on Market Share of Five Largest Insurers

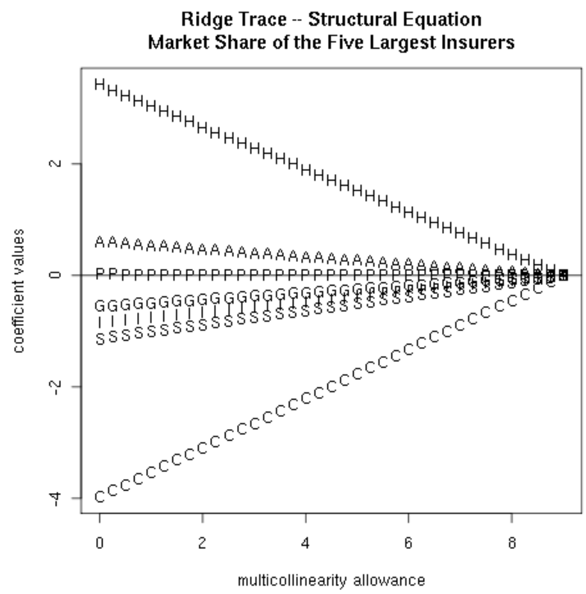

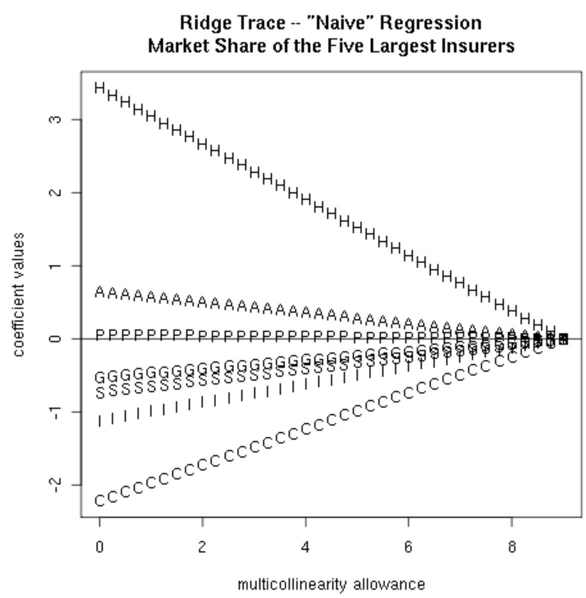

3.5.4. Empirical Results

Tables 3.2 and 3.3 show that tighter restrictions on the use health status when setting premiums are associated with lower employment-based coverage rates and higher market share of the five largest employers and the effects are statistically significant. In fact, the coefficient on health restrictions is much larger (in absolute value) than any of the other regulatory variables, indicating that restricting insurers' ability to rate on the basis of health status has a substantial impact on coverage and market concentration.

The effects of other rating restrictions are more mixed however. Restrictions on the use of age as a rating factor is also associated with lower coverage rates and higher market concentration, but the effects are not statistically significant.

Restrictions on the use of gender seems to reduce both coverage rates and market concentration, but once again the effects are not statistically significant.

Some of the other rating restrictions are statistically significant however. Restricting the use of industrial classification and group size as rating factors appears to reduce the market share of the five largest insurers and raise employment-based coverage rates, but the effect of restrictions on the use of industrial classification is only significant in a few of the regressions on market concentration, while the effect of restrictions on the use of group-size is only significant in the regressions on the employment-based coverage rate.

Longer time periods during which insurers can exclude coverage of pre-existing conditions has a negative and statistically significant effect on employment-based coverage rates, but does not have a significant effect on market concentration. These results suggest that longer exclusion periods are harmful to less healthy individuals, while having little or no effect on the insurance industry.

Since the regulatory variables are highly correlated with each other, it is important to see if the signs of the coefficients change as they are shrunk towards zero. As discussed in Chapter 1, when two explanatory variables are correlated, their regression coefficients will exhibit correlation of the opposite sign. In such a case, multicollinearity may cause us to conclude that the effect of one of the regulatory variables increases the coverage rate (or market concentration) even though its true effect on the dependent variable is negative.

Fortunately, multicollinearity does not appear to affect the empirical signs in this case. As shown in Figures 3.1 and 3.2 all of the coefficients retain their sign as they are shrunk towards zero.

Finally, employment-based coverage rates and market concentration appear to be negatively correlated with each other and the effects are statistically significant. Prior to drawing this conclusion however, we must first ensure that simultaneity bias does not affect the regression results.

As discussed in Section 3.4, if small insurers sell a disproportionately large share of marginal insurance policies than large insurers, then market concentration should fall as employment-based coverage rates rise. Moreover, lower market concentration is an indicator of a more competitive market. When there are more suppliers of health insurance, the equilibrium price of insurance should be lower and the percentage of people covered by employment-based coverage should be higher (than in states with high degrees of market concentration).

In short, theory suggests that coverage and market concentration are endogenous to each other. Therefore we cannot simply regress coverage rates on market concentration (and vice versa) because if they are endogenous to each other, we will obtain inconsistent estimates of their coefficient values.

The third columns of Tables 3.2 and 3.3 show a statistically significant correlation between the residual of the five largest insurers' market share and employment-based coverage rates, but do not show a statistically significant correlation between the residual of the coverage rates and market concentration. In other words, coverage rates appear to be exogenous to market concentration, but the reverse is not true.

Ideally, we would also have run the similar regressions with state-fixed effects to account for the state-specific differences in coverage rates and market concentration. Unfortunately however, the regulatory variables are often perfectly correlated with the state dummy variables because changes in state law take time to enact. Under such circumstances the moment matrix (i.e. $X^T X$) is exactly singular and regression is impossible.

3.5.5. Discussion

The regressions described above effectively tested three hypotheses about employment-based coverage rates:

- the hypothesis that the higher the percentage of the population under age 65 that is employed (hereafter: the "participation rate") the higher the employment-based coverage rates will be,

- the hypothesis that more densely populated states have higher employment-based coverage rates because population clusters provide consumers with better access to insurers and/or because more densely populated states have less agriculture and therefore higher employment-based coverage rates
- the hypothesis that tighter regulation of insurers' rating practices and limitation of pre-existing condition exclusion periods reduces employment-based coverage rates.

The hypotheses of how employment variables, population variables and regulatory variables affect employment-based coverage rates were broadly in line with the regression results. Higher participation rates and population density enable a higher percentage of the population to obtain health insurance coverage through employment (although the effect of population density is not always statistically significant).

The regression results also indicate that tighter restrictions on the use health status when setting premiums are associated with lower employment-based coverage rates. The rating restrictions on age and sex are also associated with lower coverage rates, although their effects are not statistically significant.

It is important to point out that if community rating does reduce coverage rates (as the regression results suggest they do), the lower coverage rates are not necessarily undesirable. If coverage rates fall because some younger, healthier workers elect to forgo insurance coverage rather than paying high premiums to subsidize the price of insurance for older, sicker workers, then health insurance may become more affordable to those who need it the most – the old and the sick. The young and the healthy may not have coverage, but that group is also less likely to need health insurance coverage.

By contrast, rating restrictions on the use of group size and industrial classification are associated with higher employment-based coverage rates (although the effect of restrictions on industry rating is not always statistically significant).

Small groups' insurance policies generally carry a higher loading charge[10] than the policies of large groups. Small groups therefore pay more than large groups for plans with equivalent benefits, which may be one of the reasons why the employees of small firms are more likely to lack insurance coverage (GAO, 2001a, p. 8-9; GAO, 2001b, p. 7-16). The GAO also found that the likelihood of being uninsured varies by industry. Construction workers and agricultural workers are the two groups most likely to be uninsured (GAO, 2001a, p. 9-10).

Prohibiting insurers from setting premium rates on the basis of group size and industrial classification would therefore make insurance more affordable to the groups mostly likely to lack insurance.

Limiting the length of time that insurers may exclude pre-existing conditions may also increase the percentage of people covered by employment-based health insurance as the regression results indicate that longer exclusion periods are associated with lower employment-based coverage rates and the effect is statistically significant.

[10] A loading charge reflects the cost of administering a health plan. It's equal to one minus the ratio of benefits to premiums.

Turning to the supply of health insurance, the regressions effectively tested three hypotheses about market concentration:

- the hypothesis that larger populations support a higher number of insurers and thereby limits the domination of the health insurance market by the five largest insurers,
- the hypothesis that more densely populated states support a higher number of insurers because urban areas can support more insurance agencies, which dampens market concentration and finally
- the hypothesis that stricter state regulation of insurers discourages insurance companies from operating in a state, thus increasing market concentration.

The regression results suggest that larger state populations and higher population density are associated with lower market share of the five largest insurers, as predicted.

As expected, tighter restrictions on the use of age and health status as rating factors are associated with higher market concentration, while constraining the use of industrial class, group-size and gender mix is associated with lower market concentration. (It should be noted however that the effect of only two regulatory variables is statistically significant from zero: restrictions on the use of health status and industrial class).

Longer pre-existing condition exclusion periods also increase market concentration, although the effect is not statistically significant. While unexpected, this effect is consistent with the results that Chollet et al. (2000) obtained when they examined the small-group market.

Finally, the regression results suggest that the coverage rate has a negative and statistically significant effect on the market share of the five largest insurers, but the coverage rate appears to be exogenous to market concentration, indicating that the unexplained component of coverage rates does not have a statistically significant effect on market concentration. Consequently, we cannot conclude that employers on the margin of purchasing insurance, purchase policies disproportionately from small insurers.

The unexplained component of market concentration does appear to have a negative and statistically significant effect on coverage rates however, indicating that marginal increases in market concentration reduce coverage rates.

A low degree of market concentration is an indicator of a more competitive market. It is often asserted that a more competitive market forces insurers to offer employers lower insurance premiums, thereby inducing employers to purchase (or retain) insurance coverage for their employees.

It is however theoretically possible that a less competitive market increases insurance coverage rates. Because insurers face increasing returns to scale, each extra policy an insurance company issues costs the company less (in terms of incremental risk) than the previous one they issued. Large insurers can therefore offer lower premiums than small insurers. Provided that no insurer has an absolute monopoly, high market share of the five largest insurers might increase employment-based coverage rates.

This is not the case however. The regression results presented in Section 3.5 indicate that higher market concentration is associated with lower employment-based coverage rates.

3.6. Conclusion

This chapter has shown that restricting insurers' ability to set premium rates on the basis of health status reduces the percentage of people who receive health insurance coverage from an employer and increases the market share of the five largest insurers within a state.

As a normative matter, one could argue that such findings make a strong case against community rating. Such an argument implicitly assumes however that the goal of public policy should be to maximize the employment-based coverage rate and ignores the important question of *who* is covered by insurance.

In fact, Herring and Pauly (2006) find that rating restrictions do enable some "high-risk uninsureds" to obtain individual coverage, but they also find that there is a slight decrease in coverage rates as more "low-risk" individuals elect to forego coverage.

Proponents of community rating could therefore argue that community rating is desirable – despite the fact that it reduces coverage rates – because it helps older and sicker workers afford coverage. As mentioned in Section 3.4 however, requiring younger and healthier workers to subsidize the premium rates of older and sicker workers is less efficient than a transfer of income Feldman (1987).

Nonetheless, there is considerable pooling of premium rates in states that do not regulate the individual insurance market (Herring and Pauly, 2006). One can therefore infer that there is even greater pooling in the small group market, where price subsidization occurs because firms employ both younger and healthier workers and older and sicker workers.

The pooling implicit in any employer-sponsored insurance plan may also explain why the magnitude of the rating restriction coefficients are so much smaller in the regressions on coverage rates than they are in the regressions on market concentration.

Consequently, mandating community rating may reduce the competitiveness of the small-group health insurance market, but it is unlikely to cause a substantial reduction in employment-based coverage rates. Most importantly, restricting insurers' ability to rate on the basis of health status is likely to meet the social policy objective of making insurance more affordable to older and sicker workers.

Bibliography

M. S. Ahluwalia. “Income Distribution and Development: Some Stylized Facts”. *American Economic Review*, 66(2):128–135, May 1976.

A. Alesina and D. Rodrik. “Distributive Politics and Economic Growth”. *Quarterly Journal of Economics*, 109(2):465–490, May 1994.

S. J. Babones. “Income Inequality and Population Health: Correlation and Causality”. *Social Science and Medicine*, 66(7):1614–1626, Apr. 2008.

A. V. Banerjee and E. Duflo. “Inequality and Growth: What Can the Data Say?”. *Journal of Economic Growth*, 8(3):267–299, 2003.

G. S. Becker and B. R. Chiswick. “The Economics of Education: Education and the Distribution of Earnings”. *American Economic Review*, 56(1/2):358–369, Mar. 1966.

T. Buchmueller and J. Dinardo. “Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania and Connecticut”. *NBER Working Paper*, (6872), Jan. 1999.

H. Chenery, M. S. Ahluwalia, C. L. G. Bell, J. H. Duloy, and R. Jolly. *Redistribution with Growth*. Oxford University Press, Aug. 1974.

B. R. Chiswick. "Earnings Inequality and Economic Development". *Quarterly Journal of Economics*, 85(1):21–39, Feb. 1971.

D. J. Chollet, A. M. Kirk, and K. I. Simon. "The Impact of Access Regulation on Health Insurance Market Structure". *US Dept. of Health and Human Services*, 20 Oct. 2000.

K. Deininger and L. Squire. "A New Data Set Measuring Income Inequality". *World Bank Economic Review*, 10(3):565–591, Sept. 1996.

D. Feenberg. "Average Marginal State Income Tax Rates 1977+", 1 Jan. 2005. URL `http://www.nber.org/~taxsim/state-marginal/`.

R. Feldman. "Health Insurance in the United States: Is Market Failure Unavoidable?". *Journal of Risk and Insurance*, 54(2):298–313, June 1987.

M. Feldstein and M. Vaillant Wrobel. "Can State Taxes Redistribute Income?". *Journal of Public Economics*, 68(3):369–396, 1 June 1998.

K. J. Forbes. "A Reassessment of the Relationship between Inequality and Growth". *American Economic Review*, 90(4):869–887, Sept. 2000.

O. Galor and J. Zeira. "Income Distribution and Macroeconomics". *Review of Economic Studies*, 60(1):35–52, Jan. 1993.

GAO. "Health Insurance Regulation: Varying State Requirements Affect Cost of Insurance". *General Accounting Office*, (GAO/HEHS-96-161), Aug. 1996.

GAO. “Health Insurance: Characteristics and Trends in the Uninsured Population”. *General Accounting Office*, (GAO-01-507T), 13 Mar. 2001a.

GAO. “Private Health Insurance: Small Employers Continue to Face Challenges in Providing Coverage”. *General Accounting Office*, (GAO-02-8), Oct. 2001b.

M. Grossman. *“Household Production and Consumption”*, chapter The Correlation between Health and Schooling, pages 147–211. National Bureau of Economic Research, 1975.

B. Herring and M. V. Pauly. “The Effect of State Community Rating Regulations on Premiums and Coverage in the Individual Health Insurance Market”. *National Bureau of Economic Research Working Paper Series*, (12504), Aug. 2006. URL `http://aspe.hhs.gov/daltcp/reports/2006/HIMregs.htm`.

A. O. Hirschman and M. Rothschild. “The Changing Tolerance for Income Inequality in the Course of Economic Development”. *Quarterly Journal of Economics*, 87(4):544–566, Nov. 1973.

A. E. Hoerl and R. W. Kennard. “Ridge Regression: Biased Estimation for Nonorthogonal Problems”. *Technometrics*, 12(1):55–67, Feb. 1970.

S. Kuznets. “Economic Growth and Income Inequality”. *American Economic Review*, 45(1):1–28, Mar. 1955.

J. A. List and C. A. Gallet. “The Kuznets Curve: What Happens After the Inverted-U?”. *Review of Development Economics*, 3(2):200–206, June 1999.

G. C. Loury. “Intergenerational Transfers and the Distribution of Earnings”. *Econometrica*, 49(4): 843–867, July 1981.

M. S. Marquis and M. R. Holmer. “Alternative Models of Choice Under Uncertainty and Demand for Health Insurance Coverage”. *Review of Economics and Statistics*, 78(3):421–427, Aug. 1996.

J. Meckl and S. Zink. “Solow and Heterogenous Labor: A Neoclassical Explanation of Wage Inequality”. *Economic Journal*, 114:825–843, Oct. 2004.

R. A. Musgrave. *The Theory of Public Finance*. McGraw-Hill, 1959.

J. P. Newhouse. “A Simple Hypothesis of Income Distribution”. *Journal of Human Resources*, 6 (1):51–74, Winter 1971.

R. L. Obenchain. “Ridge Analysis Following a Preliminary Test of the Shrunken Hypothesis”. *Technometrics*, 17(4):431–441, Nov. 1975.

R. L. Obenchain. “Classical F-Tests and Confidence Regions for Ridge Regression”. *Technometrics*, 19(4):429–439, Nov. 1977.

R. L. Obenchain. Shrinkage Regression: ridge, BLUP, Bayes, spline & Stein. Mar. 2004. URL `http://members.iquest.net/~softrx/rxadobe.htm`.

R. Perotti. “Political Equilibrium, Income Distribution, and Growth”. *Review of Economic Studies*, 60(4):755–776, Oct. 1993.

T. Persson and G. Tabellini. “Is Inequality Harmful for Growth?”. *American Economic Review*, 84(3):600–621, June 1994.

D. Ray. *Development Economics*. Princeton University Press, 1998.

S. Robinson. "A Note on the U Hypothesis Relating Income Inequality and Economic Development". *American Economic Review*, 66(3):437–440, June 1976.

R. M. Solow. "A Contribution to the Theory of Economic Growth". *Quarterly Journal of Economics*, 70(1):65–94, Feb. 1956.

J. E. Stiglitz. "Distribution of Income and Wealth Among Individuals". *Econometrica*, 37(3): 382–397, Aug. 1969.

US Bureau of Labor Statistics. "State and Area Employment, Hours, and Earnings: Statewide: Total Nonfarm: All Employees". URL `http://www.bls.gov/sae/home.htm`.

US Census Bureau. "Historical Health Insurance Tables: Table HI-6: Health Insurance Coverage Status and Type of Coverage by State – People Under 65: 1987 to 2005". URL `http://www.census.gov/hhes/www/hlthins/historic/hlthin05/hihistt6.html`.

H. D. Vinod. "Application of New Ridge Regression Methods to a Study of Bell System Scale Economies". *Journal of the American Statistical Association*, 71(356):835–841, Dec. 1976.

R. C. Wonderling. "Number of Insurers in Each State, Total Premiums Earned and Claims Paid by Each Insurer". data provided to the Office of PA State Senator Robert C. Wonderling by the National Association of Insurance Commissioners, July 2003a.

R. C. Wonderling. "Telephone Survey Research on Small Group Health Insurance Markets". unpublished manuscript from the Office of PA State Senator Robert C. Wonderling, July 2003b.