Economic Consequences of Income Inequality

Jason Furman
Joseph E. Stiglitz

Traditionally, economics has treated efficiency and equity as separable. The theoretical basis for their separation is the Second Fundamental Theorem of Welfare Economics, which holds that *any* Pareto efficient outcome can be implemented as a competitive equilibrium given the appropriate lump sum taxes and transfers. Policymakers have, for the most part, been happy to respect this separation.

Nowhere is this separation greater than in representative agent models, which have been one of the dominant paradigms for studying business cycles, economic growth, and other macroeconomic phenomena. Not only do these models preclude us from studying the relationship between distribution and efficiency, but they also can be deeply misleading in the way they lead us to assess efficiency. Consider Robert Lucas' (1987) calculation, based on a plausible calibration of an infinitely lived representative agent model, that the utility from eliminating fluctuations in consumption is the equivalent to the utility from a permanent increase in the level of consumption by 0.1 percent—or roughly $20 per person in 1998 and growing with consumption thereafter. Looked at another way, using more calibrations from Lucas (1987), the additional utility from eliminating the business cycle is equivalent to the additional utility from boosting the growth rate by 0.005 percentage points per year. If this is true, then,

even if it were possible, we would choose not to implement a policy that eliminated the business cycle if it, say, entailed lowering the growth rate from 2.50 percent to 2.49 percent. Lucas draws the conclusion that "economic instability of the level we have experienced since the Second World War is a minor problem" and counsels against "devising ever more subtle policies to remove the residual amount of business-cycle risk."

At first this result seems striking. Surely, we would be willing to pay more to avoid times like the Great Depression when the unemployment rate rose above 25 percent or the 1981-83 recession when it rose to 10 percent? Lucas' estimate, however, does not ask this question. There is no unemployment in representative agent models. In these models, the representative worker would have reduced his or her hours in the early 1980s. In reality, several million people lost their jobs. With indivisible employment, risk aversion, and incomplete risk markets,[1] there is a strong presumption that society would be willing, *ex ante*, to give up substantially more average growth to avoid these recessions than is suggested by the Lucas calculation.

The results are also predicated on perfect capital markets. If individuals can borrow and lend at a "market rate of interest," then short spells of unemployment cause only modest changes in lifetime income, and accordingly, only modest changes in consumption. To the extent that these fluctuations are anticipated, there would not even be fluctuations in consumption.

In fact, there is not only a wealth of microeconomic evidence concerning imperfect capital markets, but also ample evidence that these imperfections have macroeconomic effects. Theories of information imperfections developed over the past two decades have provided explanations for these market imperfections and explored their macroeconomic consequences. Capital market imperfections imply that consumption fluctuations induced by business fluctuations are far greater than they would be with perfect capital markets, with correspondingly large negative effects on welfare.[2]

Note that these statements refer only to the welfare costs of economic

fluctuations. Credit and equity rationing may also interact with inequality to affect the volatility of output and also business fluctuations may themselves have consequences for economic growth.

We can go one step further toward reality by introducing heterogeneous agents and an initially unequal distribution of income. In general, we know that with disparate economic circumstances (and preferences), it is impossible to aggregate people's preferences into a single "social welfare function" (Arrow 1963). This means, in particular, that there will be differences in views concerning virtually any aspect of economic policy.

Furthermore, from the perspective of positive analysis, the policies a country adopts, and thus, potentially the mean or the variability of its growth rate, will depend both on the initial distribution of income and on its political institutions. When income is distributed unequally, politics matters.

Research over the last three decades has opened up a number of avenues for thinking about the consequences of inequality. One strand stems from the observation that in practice the government cannot use lump sum taxes and transfers to serve distributional goals, and basing taxes and transfers on income and other variables that are affected by individual action is necessarily distortionary. More profoundly, developments in the economics of information have established that in the presence of imperfect information, which is always, markets are not even constrained Pareto efficient. There is some government action that, while, itself, bound by the same informational limitations, can make at least some people better off while not making anyone worse off (Greenwald and Stiglitz 1986). A direct implication of this result is that distribution does matter; changing the distribution of income or wealth can affect the efficiency of the economy.

In the last decade, macroeconomics has picked up a lot of these points. Non-representative agent models take seriously issues of aggregation. Political economy models consider the links between economics and politics, in which the income distribution is an important variable.

One consequence of this research is that today we have a much better understanding of the effect of inequality on growth than we had 20 years ago. We know that overall inequality is, if anything, harmful to growth, and also, that countries can increase their incomes with relatively little change in inequality. We also understand several channels that could account for this relationship. These channels have important and well-developed implications for many areas of economic policy including education, taxation, financial regulation, and, in the context of developing countries, land reform.

The importance of distributional considerations for the business cycle, and the conduct of monetary policy in particular, has received less attention. Too often we slip back into representative agent models when thinking about monetary policy. Monetary policy decisions are often perceived to be, especially by those making them, highly technical choices about how to balance the twin objectives of price stability and full employment. But monetary policy choices have a huge distributional impact. A movement in interest rates may help bondholders while hurting workers. Furthermore, the choice of monetary policy rules (for example, inflation targeting, the "Taylor Rule," nominal GDP targeting, and so forth) has an impact on the variances of real GDP and inflation, and thus has differential consequences for different groups. Similarly, the institutional design of the central bank (including its degree of independence, the method by which its officials are selected, and the way voting is organized) affects these tradeoffs, and thus can have substantial distributional consequences.[3]

There is a clear link from central bank policy and design to distributional outcomes. But distributional outcomes also in turn affect the business cycle and influence even the aggregate tradeoffs faced by monetary policy. Prolonged unemployment can have long-lasting effects on the natural rate of unemployment through "hysterisis" effects, leading insiders to entrench themselves or workers to lose skills (see Blanchard and Summers 1986 and Layard and others 1991). By the same token, it is plausible that the high employment policies pursued by the United States have had the reverse effect in lowering the natural rate of unemployment (Stiglitz 1997b). In this

paper we explore some of the effects that inequality has on both the natural rate of unemployment and the dynamics of unemployment. These effects arise because inequality affects the incentives, information, and skills of employers, workers, and job seekers.

The plan of this paper is as follows. In the first section, we review some of the theories and evidence on the consequences of inequality for growth. In the second section of the paper, we sketch several models, which imply that greater inequality will lead to greater equilibrium unemployment. We show some very tentative evidence that this is indeed the case. In the third section of the paper we discuss some of the relative distributional consequences of inflation and unemployment. The key insight here is that higher unemployment results in greater inequality; and the marginal contribution to inequality from a given increase in unemployment is rising in the degree of initial inequality. The implication of these relationships is that if there are adverse economic consequences to inequality, then a higher initial degree of inequality will magnify the adverse social and economic consequences of business cycle fluctuations. Thus, the representative agent models, by assuming *no* inequality, may have vastly understated not just the welfare cost of fluctuations, but their consequences for economic policy. In the fourth section, we build on this insight to link the discussion of trends with the discussion of cycles, by discussing the relationship between inequality and the persistence of unemployment and the relationship between inequality and consumption volatility. Finally, in the fifth section we offer some concluding comments about the *intrinsic* importance of inclusion and egalitarian policies.

Income distribution and growth

Traditionally, studies of the relationship between aggregate economic activity and income distribution emphasized the link from output to the distribution of income. The pioneering work of Simon Kuznets (1955) hypothesized that as countries developed, the degree of inequality initially increased and then fell—the inverted-U that has come to be known as the "Kuznets curve." Kuznets' results, however, were based on a sample of only a few data points from three

countries (the United States, the United Kingdom, and Germany). More comprehensive studies, however, find little evidence of a Kuznets curve in the data. In analyzing their new data set on inequality, which is the most comprehensive and carefully constructed, Klaus Deininger and Lyn Squire (1996) find that in 88 spells of decade-long growth, inequality improved in 45 of them and worsened in 43 of them. In most of the cases, the changes in inequality were relatively small and they were uncorrelated with initial income. As a result, according to their findings, the poorest fifth of the population saw their incomes improve in 77 of these 88 spells.

There is also another literature going back to W. Arthur Lewis (1954) and Nicholas Kaldor (1957), which studies the consequences of inequality for income and growth. Research has identified several channels through which inequality might affect growth, the potentially most important of which are briefly surveyed here.

Inequality, saving, and consumption

The traditional theory about the link between income distribution and growth is that greater inequality leads to greater savings, and thus greater capital accumulation and growth. Kaldor (1957) assumes that all wages are consumed and only profits provide resources for additional capital accumulation. In Luigi Pasinetti (1962), this is interpreted as a consumption function in which the rich have a lower marginal propensity to consume than the poor, so that increasing inequality will result in higher savings and thus more rapid capital accumulation and growth. Time has not served this hypothesis well. Since the pioneering contributions of Solow (1957), we no longer think of capital accumulation as being the most important determinant of growth. Still, in the Solow model, capital accumulation is the single most important factor that countries have well-established means of influencing—although these means do not include distributional policy.

The Kaldor-Pasinetti channel is grounded in applying a Keynesian-style aggregate consumption function across households. But, at virtually the same time that these ideas were being developed, Milton

Friedman (1957) showed that this consumption function was the result of confusing total income with permanent income. If the transitory component of income were removed from consumption, Friedman found that savings was proportional to (permanent) income; with a constant marginal propensity to consume, the distribution of income has no effect on savings. These results have, for the most part, held up over the succeeding decades, with one important caveat. Lawrence Kotlikoff and Laurence Summers (1981) have argued that a substantial fraction of wealth is attributable not to life-cycle savings, but to bequests, and there is considerable theory and evidence to suggest that bequests are highly non-linear in lifetime income. (This is consistent with findings of high degrees of inequality, greater than could be accounted for by the life-cycle savings itself.[4]) With convex saving functions, an increase in inequality will increase savings.

In spite of this possibility, the evidence is not only that inequality is not necessary for high savings, but also that it is, in fact, not even associated with it. The experience of East Asia has shown convincingly that even poor peasants can save at high rates. More broadly, cross-country evidence shows that there may be little relationship between inequality and aggregate savings, as evidenced by the scatterplot in Chart 1. This result continues to hold in studies that control for other determinants of savings (see Schmidt-Hebbel and Servén 1997).

Imperfect information and agency costs

The second mechanism by which income distribution can affect growth comes from the economics of information. In the presence of incomplete information, the Second Fundamental Theorem of Welfare Economics will not hold, and there will be some government interventions that represent Pareto improvements over the competitive equilibrium (Greenwald and Stiglitz 1986). This general theorem has specific applications to inequality. Inequalities in wealth will generally require the owner of assets, the "principal," to delegate the use of assets to another, the "agent." In the presence of incomplete information, it will be impossible to write down a complete contract

Chart 1
Inequality and Savings Rates

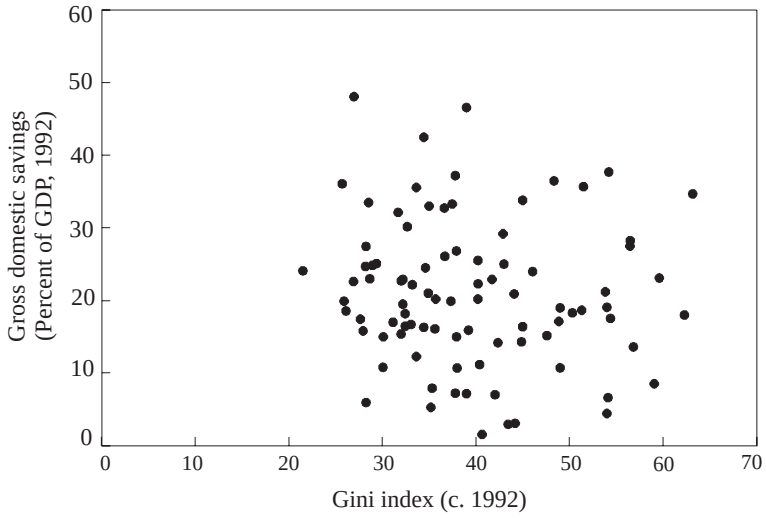

Source: Deininger and Squire (1996) and *World Development Indicators*, World Bank.

to specify what actions the agent should undertake in each contingency or at least impossible to monitor (and thus enforce) that contract. As a result, a principal-agent problem arises, creating what are sometimes called "agency costs." Agency costs and the steps taken to mitigate them—like engaging in costly monitoring—affect efficiency in a number of ways. If the extent of inequality affects the extent of agency costs, then it will also affect the overall efficiency of the economy.[5]

Perhaps the most obvious example of this comes from developing countries, where sharecropping contracts are a widespread way to solve the problem of costly monitoring by ensuring that the tenant has an incentive to work.[6] But by imposing, in effect, a marginal (and average) tax rate of between 33 percent and 50 percent, these contracts lead to undersupply of effort and underinvestment in the land.

If a more equitable distribution of land allowed farmers to buy their land, the result could be an improvement in efficiency.[7]

One of the key elements in this example is credit constraint. Poor farmers do not have collateral and, given imperfect information and incomplete contracts, are unable to borrow money to buy their land. Similar issues arise in more developed countries, and are particularly evident in the context of funding education. Studies of the United States have consistently shown that the real returns to an additional year of schooling are between 5 percent and 15 percent (see Kane and Rouse 1995 and Ashenfelter and Krueger 1994). Some estimates suggest that the returns are even higher in developing countries (Psacharopoulos 1994). Poorer people, however, may be credit constrained from borrowing against expected future income to pay for their education today, especially because human capital cannot act as collateral.

The problems induced by credit constraints are amplified by the absence of insurance and other markets through which individuals can divest themselves of risk. Even where individuals can borrow to pay for their education, the returns on these investments are highly risky, and poor individuals are less willing and able to bear these risks. That provides part of the explanation for why, even where poor individuals do have access to higher education (say, through loans), there is evidence of greater underinvestment on their part and why increases in equality may be efficiency and growth enhancing.

Credit constraints play an important role even in developed countries. Credit constraints help explain fluctuations in agricultural production and in investment in small and medium-sized enterprises; and these fluctuations in turn play an important role in the business cycle.[8] Bruce Greenwald, Michael Salinger, and Joseph Stiglitz (1990) and Stiglitz (1994b) have argued that the business cycle gives rise to large fluctuations in investment in research and development, with adverse effects on productivity growth. They further argue that these fluctuations can best be understood as resulting from a combination of credit constraints (a point which Joseph Schumpter, 1934, had emphasized many years earlier) and imperfections in equity

markets, which limit the ability of firms to spread their risks. More generally, information imperfections and the associated capital market imperfections lead to risk-averse behavior on the part of firms, and this risk-averse behavior in turn helps explain a number of what might otherwise seem anomalous aspects of firm behavior, such as the procyclical nature of inventories (which in traditional models should act as a stabilizer) and the cyclical pattern of hours and employment (Greenwald and Stiglitz 1995).[9]

Many of the elements of these channels from income distribution to growth can be tested directly. Studies have consistently found that ownership and contract arrangements matter for agricultural productivity. One study, for instance, found that, after controlling for other factors, output per acre is 16 percent higher on owner-occupied land than on sharecropped land (cited in World Bank, forthcoming). Studies of education have been less conclusive, because of the potential importance of difficult-to-measure externalities and the pervasive endogeneity. There is, however, at least some evidence that egalitarian policies and growth can form a virtuous circle through their effects on increasing education, which increases equitable growth and in turn leads to further increases in education (Birdsall, Ross, and Sabot 1995).[10]

Fiscal policy

The third channel that income distribution can affect inequality is through fiscal policy. A key assumption of the Second Fundamental Theorem is that the policymaker can use lump sum taxes and transfers to redistribute income. Without perfect information about people's types, however, any redistributive taxes and transfers need to be conditional on variables that are influenced by individual decisions. Thus, they are necessarily distortionary. In standard optimal taxation models, an increase in the degree of inequality will (other things being equal) lead to more progressive—and more distortionary—taxation.[11]

Alberto Alesina and Dani Rodrik (1994) focus on the political economy issues raised by these considerations. They assume the

choice of the tax rate is endogenous to the political process and using the standard majority voting model, the median voter's preference will be decisive. They first show that the worse off the median voter is relative to the mean (their proxy of inequality), the higher the tax rate.[12] They then establish that increasing taxes above the optimal rate (which is determined by balancing the marginal cost of distortionary taxes with the marginal benefit of productive government expenditures) lowers the rate of capital accumulation and thus growth. Thus, their model predicts that greater inequality leads to higher taxes and lower growth, the reduced form of which (inequality is bad for growth) they find in the data.

It is, however, possible that higher inequality leads to lower tax rates, because the rich are able to exert a disproportionately powerful impact on the political system. There is, in the United States, ample evidence of the powerful influence of tiny minorities in debates about, for instance, estate taxes. Their influence may be even greater in non-democracies. Furthermore, many countries may be below the optimal tax threshold, in the region where higher taxes combined with productive spending will lead to higher growth. This may be especially true if the higher spending is used to relax some of the credit constraints described in the previous section, for example, through education spending. As a result, inequality may lead to lower growth for political-fiscal reasons, but because inequality leads to lower taxes and lower taxes lead to slower growth: the opposite of Alesina and Rodrik's explanation. Perotti (1996) finds evidence that this may indeed be the case, with inequality leading to lower personal tax rates (although having an insignificant effect on transfers) and taxes/spending having a positive effect on growth. Still, he finds that the overall political-economy fiscal effect is probably not the reason for the negative relationship between inequality and growth.

Political stability

The fourth potential channel between inequality and growth is through political stability. In contrast to the fiscal policy channel, which involves combining a political mechanism with an economic mechanism, this channel is almost purely political. The intuition is

that greater inequality leads to more political instability, and greater political stability leads to lower growth. Although this mechanism is most plausible in some highly unequal and unstable developing countries, Michael Bruno (1993) has argued persuasively that consensus is essential to undertaking successful macroeconomic stabilization, a consensus that would presumably be facilitated by a greater degree of equality. In this spirit, Rodrik (1998) finds that even in OECD countries the level of inequality (and also the existence of risk-sharing institutions) affects the government's ability to develop sound policies to cope with major shocks, like the 1973 oil shock, and thus their subsequent growth rates.

Slower growth can lead to greater political instability and greater political instability can lead to slower growth. Thus, to test whether inequality affects instability and thus growth it is necessary to jointly estimate the two halves of the relationship. Studies have consistently found that low income, a large share of the population living in poverty, and income inequality all lead to greater political instability (Londregan and Poole 1990, Alesina and Perotti 1996, and Alesina and others 1996). These studies differ in their findings on whether or not political instability leads to lower growth, with Londregan and Poole finding no evidence that it does, while the two Alesina studies and Barro (1996) finding that a higher propensity of government collapse results in lower growth. This is true both for coups as well as constitutional changes that alter the ideological composition of the executive branch.

The overall evidence

As we have noted, there is relatively little movement in inequality over time, and as a result, there is very little evidence about the relationship between inequality and growth from time series data for individual countries.[13] As a result, most of the evidence about the relationship between inequality and growth comes from cross-country regressions. This evidence provides a presumption that, if anything, inequality is bad for growth.

Roland Benabou (1996) surveys 13 papers that contain regressions of growth or investment on inequality. According to his tabulations,

Chart 2
Inequality and Growth

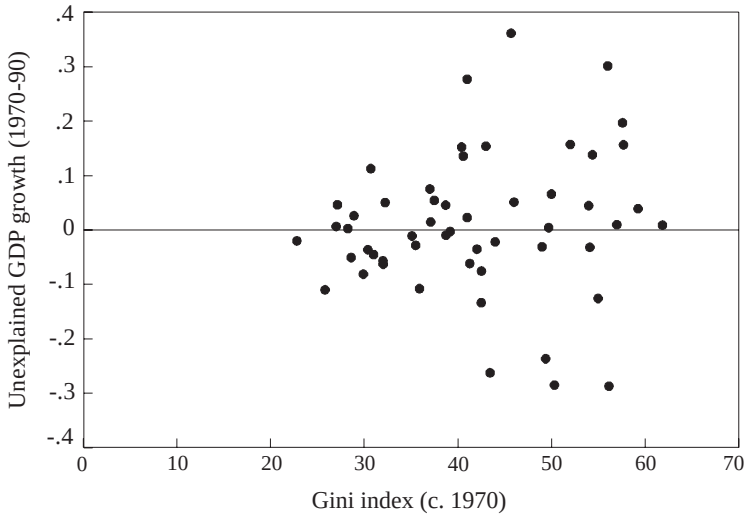

Source: Calculations based on data from Penn World Tables Mark 5.6, Deininger and Squire (1996), and *World Development Indicators*, World Bank.
Note: "Unexplained growth" is the portion of growth from 1970-90 unexplained by a regression on log initial GDP per capita, initial school enrollment, initial fertility rate, and dummy variables for East Asia, Latin America, and Sub-Saharan Africa.

in 10 of these papers, there is a consistent, significant, and negative relationship between inequality and either growth or investment. The other three papers find either negative relationships or essentially no relationship, but do not find results that are consistently significant. Econometrics, however, is only as good as the data and the inequality data employed by these regressions suffer from many inconsistencies of definition, variations in coverage, and lack of comparability. When Deininger and Squire (1996) re-ran many of these regressions with a better data set, they found no relationship between inequality and growth. Chart 2 illustrates this result by showing the relationship (or lack thereof) between the portion of

growth between 1970 and 1990 that is not explained by a regression of growth on standard controls and the Gini index.

At the very least, this suggests that there is little evidence either for Kuznets' or Lewis' hypothesis: inequality is neither necessary for growth nor is it an inevitable consequence of growth. There are obviously policy implications from the observation that those countries that have pursued more egalitarian policies have succeeded in reducing inequality, but have not had to sacrifice much, if anything, in terms of growth. At the very least, the data do not support the view that there is a strong tradeoff between growth and equality.

Of course, if the data had shown that there were a relationship, most of the cross-country regressions themselves do not enable us to distinguish among the possible explanations.[14] But the alternative theories do have different implications for the behavior of "intermediate variables," which allow us to form judgments about the plausibility of alternative mechanisms. There is reasonable evidence for some (increased equality enhances agricultural productivity in developing countries and results in more growth-harming political instability) and reasonable evidence against others (savings does not seem to be systematically related to inequality).

Cross-country evidence, however, may be of limited usefulness in addressing the question of what the adoption of more egalitarian policies will do to the rate of economic growth. The observation that there is no relationship between economic growth and inequality is consistent with the hypothesis that some countries have adopted efficient/egalitarian policies, others efficient/inegalitarian, others and so forth. And, in other words, there exist policies that result in more equitable outcomes with rapid growth. Chart 3 shows the relationship between income inequality and growth, this time without controlling for factors that affect growth. Note that many East Asian countries are clustered in the upper left (greater equality and high growth) and many Latin American countries are in the lower right (greater inequality and low growth). If these differences are the result of the different development strategies they pursued, then we could find policies that are consistent with equality and growth.

Chart 3
Income Inequality and Growth of GDP, 1965-89

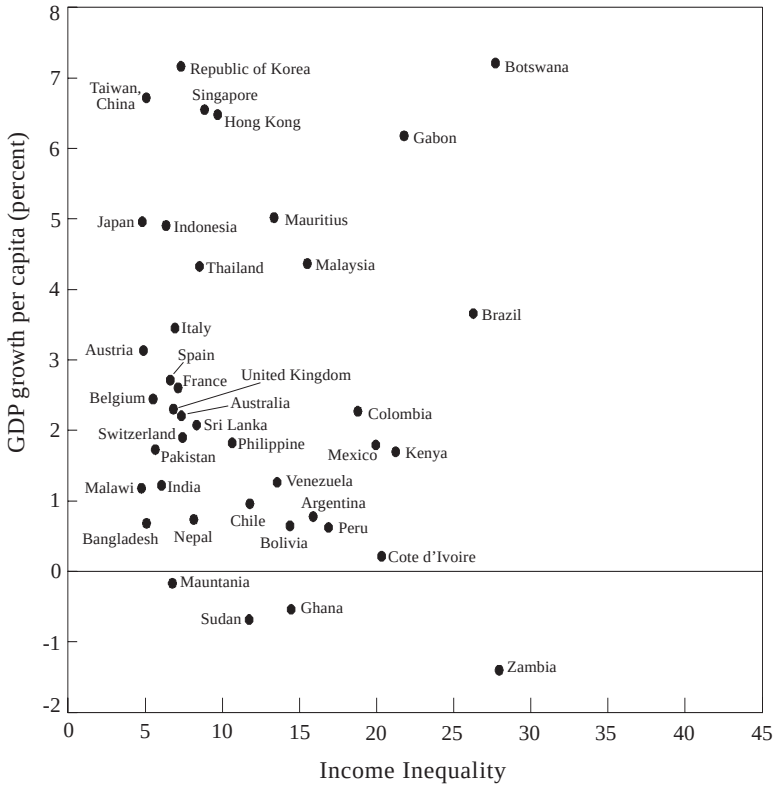

Note: Income inequality is measured by the ratio of the income shares of the richest 20 percent and the poorest 20 percent of the population.
Source: World Bank data.

More important, however, we are interested in what changes in policy do to changes in inequality and changes in growth. Even panel growth regressions do not really address this issue because almost all of the identification of the effect of inequality on growth comes from cross-sectional variation, not time series variation. But again, anecdotal evidence and economic analysis suggest that there is no general

answer to this question. If redistribution is accomplished by, for instance, setting the top marginal tax rate at 90 percent, then it is likely that the income distribution will improve while growth falls. But if, in a developing country, redistribution is accomplished by establishing government credit that allows poor sharecroppers to buy their land, then equality and efficiency could both increase. Thus the most important question is not whether equality increases growth, but how we can make changes that increase both equality and growth.

Inequality and equilibrium unemployment

In this section of the paper, we discuss a much less explored topic, the relationship between income distribution and the equilibrium unemployment rate. Again, there is the possibility of a two-way relationship: unemployment may affect inequality, and inequality may affect unemployment. In this section, we focus on *equilibrium* unemployment and its relationship to unemployment.

In modern macroeconomic analysis, the equilibrium level of unemployment is the non-accelerating inflation rate of unemployment (the NAIRU).[15] This is one of the key variables in delimiting the central tradeoff faced by monetary policymakers: the tradeoff between the level of economic activity and changes in the rate of inflation. Also, if there were a relationship between income distribution and the equilibrium rate of unemployment, it would have a substantial effect on the level of output but would not show up in the cross-country growth regressions discussed above.[16]

The conventional wisdom about the relationship between inequality and equilibrium unemployment was discussed by Paul Krugman (1984) at this conference four years ago. He observed that since the 1970s, the United States had witnessed an increase in inequality with little increase in unemployment, while Europe had seen its unemployment rates rise to double digits with little increase in income inequality. He argued that these were two sides of the same coin: both Europe and the United States were subject to a similar shock. The United States' more flexible labor markets allowed it to respond

Table 1
Unemployment Rates for Different Demographic Groups

Group	1970	1996
Total Civilian	5.0	5.0
Age 16-19	15.3	16.7
Age 20-24	8.2	9.2
Age 25-54	3.4	4.3
Age 55+	2.9	3.4

Source: Bureau of Labor Statistics.

to the shock through wage adjustment, while high minimum wages and other rigidities in European labor markets led to higher unemployment with little change in inequality.

There is some question about whether the stylized facts are as extreme as stated by Krugman, at least for the United States. Table 1 shows that although the overall unemployment rate was the same in 1970 and 1997, the unemployment rate went up for every age group.[17] The success of the United States in keeping its unemployment rate from rising very much was not so much that it had flexible labor markets, but that it was *lucky* with its demographics.

The reason is that the share of high-unemployment groups in the labor force has gone down, most notably that of teenagers. Chart 4 graphs the trend in the unemployment rate and inequality, as measured by the Gini index. The unemployment rate is adjusted in two ways. First, it is a five-year trailing moving average to eliminate business cycle fluctuations. Second, it is demographically adjusted by measuring a weighted average of unemployment rates for different groups by their shares in the labor force in 1970. The graph shows that the two series rise over time. This is not decisive evidence, and there are undoubtedly several other factors at work, but the graph is consistent with the hypothesis that inequality and unemployment move together.

Chart 4
Inequality and Unemployment in the United States

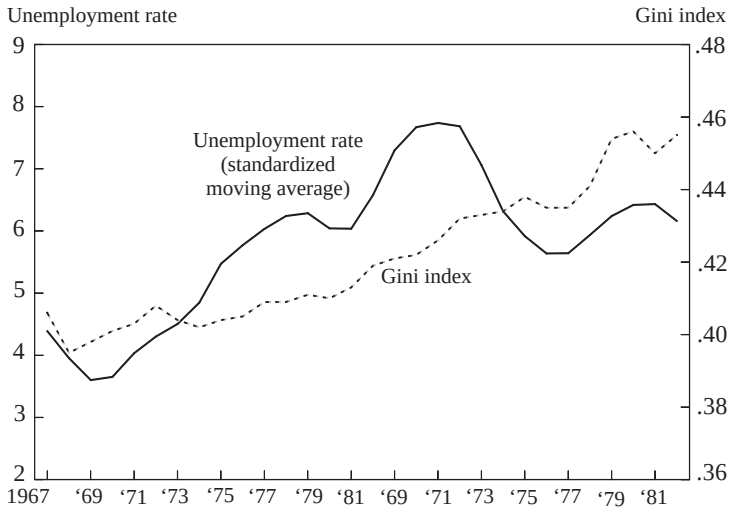

Note: The unemployment rate is a 5-year trailing moving average based on the weighted average of the unemployment rates for different age/sex groups using 1997 labor force shares as weights. The Gini index is adjusted for the redesign of the CPS in 1994.
Source: Calculations based on data from the Bureau of Labor Statistics and the Bureau of the Census.

In what follows, we describe two theories about why this might be the case. The Krugman conjecture is grounded in the theory of standard perfectly competitive labor markets with government distortions of the minimum wage. The competitive labor market model, however, is an inadequate explanation of a variety of phenomena including: the existence of equilibrium unemployment for groups that earn above the minimum wage, the persistence of unemployment, inter-industry wage differentials, the mild procyclicality of real wages, and several other features of labor markets. Other models of the labor market, including efficiency wage models and search models, have done a much better job in explaining these empirical regularities. These models also imply a very different relationship

between inequality and unemployment. In this section of the paper, we sketch two different models, which both imply that greater inequality will increase the equilibrium level of unemployment.

Efficiency wage models of unemployment

Consider first an extension of Carl Shapiro and Stiglitz's (1986) efficiency wage model to incorporate inequality. In the standard model, workers have an incentive to shirk because effort is costly and monitoring imperfect. Each firm will try to encourage effort for their workers by paying above market wages. The unintended consequence is to create unemployment, which, in equilibrium where everyone pays the same wages, creates an incentive for people to exert effort. It is impossible for unemployment to fall to zero because no one would exert effort—any worker caught shirking could find another job instantaneously.

Consider now an extension of this model to include a sector with a competitive labor market. (An algebraic treatment of this extension appears in the Appendix). This market, which is less productive than the efficiency-wage sector, has perfect and costless monitoring. Assume also that on-the-job search is impossible, so that workers wanting jobs in the efficiency-wage sector need to queue up in unemployment and wait to be randomly matched with a job opening. There is free entry between the competitive labor market and unemployment—that is, people are free to quit their jobs and also free to take jobs at the competitive wage. Unlike Shapiro-Stiglitz, unemployment in this model is voluntary search unemployment. But both employment in the competitive sector and unemployment are, in a sense, underemployment because identical workers are being used less productively than they would have been in the higher productivity efficiency-wage sector. Finally, for simplicity of exposition assume infinitely elastic demand for labor in the competitive sector at the competitive wage. Earnings inequality in this model is captured by the difference between the efficiency wage and the competitive wage.

The free entry condition between unemployment and the competitive sector implies that they must be equally desirable, that is, have

the same "value." The value of a job in the competitive sector is fixed at the present discounted value of the competitive wage. Importantly, it does not depend on wages in the efficiency wage sector or the unemployment rate.

The value of being unemployed comes from the possibility of finding a high-paying job in the efficiency-wage sector. This depends on, among other factors, the wage in the employed sector and the probability of finding employment there. This probability, in turn, is falling in the number of unemployed people that are all queuing up for the same jobs.

If the value of a competitive job is greater than the value of unemployment, people will leave unemployment to take competitive jobs. In the process, they will lower the unemployment rate, making it more likely that additional people will find jobs in the efficiency-wage sector, and thus raise the value of unemployment until it is equal to the value of a competitive job. In this extension, as in Shapiro-Stiglitz, there will always be equilibrium unemployment because as the unemployment rate goes to zero, the probability of finding an efficiency-wage job (per instant) goes to infinity, and thus there would be no incentive to work in the competitive sector and no incentive to exert effort in the efficiency wage sector.

With this intuition, we can now examine the effects of a shock that increases wage inequality (that is, raises the efficiency wage relative to the competitive wage) on the equilibrium unemployment rate. Consider a shock (like a decrease in the probability of catching a worker shirking) that raises the wage in the efficiency wage sector. As a result, some people will quit their jobs in the competitive sector and queue up for the now even higher-paying efficiency wage jobs. As a result, the unemployment rate rises.[18]

Another way of explaining the same intuition is to note that the value of a competitive job does not depend on the efficiency wage or the unemployment rate. As a result, any shock (except those that affect the marginal product of labor in the competitive sector or the discount rate) must leave the value of unemployment unchanged.

Therefore, anything that increases the wage in the efficiency-wage sector must simultaneously decrease the chances of finding a job there—that is, raise the rate of unemployment.

This model—in the limiting case with constant productivity in the competitive sector—has a striking implication concerning the impact of progressive taxation. After-tax wages remain completely unchanged. There will be fewer unemployed workers (the unemployment rate actually goes down); labor shifts to the low-wage competitive sector. Thus, in this variant of the model, before-tax inequality increases as unemployment decreases, but after-tax inequality is unrelated to unemployment.

More generally, with a downward sloping demand curve for labor in the competitive sector, the impact on before- and after-tax inequality depends on the relative elasticities of labor demand. With a highly inelastic demand curve for labor in the competitive sector, the new equilibrium may entail a much lower wage there, an increase in before- and after-tax wage inequality, and an increase in the unemployment rate.

Search models with job differentiation and selective hiring

The model just presented is the simplest version of one that combines search and efficiency wage theory. More generally, search theories predict that the greater the *after*-tax inequality, the greater search and thus the greater the equilibrium level of unemployment.[19] The intuition is that unemployment is caused by people prolonging their search to find better jobs. If the differences between jobs were smaller, they would spend less time searching, and thus the unemployment rate would fall. In this case, the inequality that matters most would be "within-group inequality" rather than "between-group inequality." No one with only a high school degree is going to prolong his job search just because the skill premium has risen. But, about two-thirds of the total increase in earnings inequality since 1979 is due to the increasing disparities within groups that have the same sex, race, and educational background (Council of Economic Advisers 1997). The important and unanswered question would be

the magnitude of the importance of this effect for the unemployment rate. The implication of this analysis is that the downward shift in the NAIRU since the early 1980s would have been even greater were it not for the increase in inequality that has occurred over the past two decades.[20]

The model we just sketched has the property that workers are *ex ante* homogeneous, but some are lucky enough to find jobs in the efficiency wage sector. Anything that changes the rewards to working in this sector (like skill-biased technological change) will increase the wages in this sector,[21] and thus earnings inequality and the unemployment rate. The model, however, assumes that workers are homogeneous—the only source of inequality is that some people are lucky enough to find jobs in the efficiency-wage sector. As such, it cannot be used to study the consequences of the other major trend we have witnessed in the last decades: the large increase in the supply of skilled workers.

According to the competitive model, this increase in the supply of skilled workers should drive down the wage premium. This model cannot be used to study equilibrium unemployment, but if this conclusion were put into the efficiency wage model, it is likely that we would also see the unemployment rate fall. Darron Acemoglu (1998), however, has developed a clever theory about why *both* skill-biased technological change and increased supply of skilled workers would lead to greater inequality and greater unemployment. In his model, employers can design jobs for low-skill workers, high-skill workers, or both. Search is costly. When there are relatively few skilled workers or the productivity difference between skilled and unskilled is relatively small, firms will choose a pooling equilibrium in which they post one type of job and accept the first applicant regardless of that applicant's skills. Since the early 1970s, however, there has been a large increase in the supply of skilled workers and a simultaneous increase in the premium paid for educated workers over non-educated workers (see Council of Economic Advisers 1997). As a result of either of these trends, firms would have a greater incentive to post separate jobs and try to match them with the appropriately skilled worker. The result is greater search unemployment

for both groups, lower wages for unskilled workers (who previously filled jobs designed either for skilled or unskilled workers), and greater wage inequality—trends we have indeed witnessed over the last two or three decades.

In addition to being consistent with these overall trends, Acemoglu cites several pieces of evidence that this is a plausible theory to explain developments in the United States over the last decades. This includes evidence that:

—firms are becoming more selective in their recruitment;

— fewer workers report being over- or undereducated for their jobs;

— "middling" occupations are being replaced with better/worse jobs;

— on-the-job training has become even more unequally distributed;

— higher-paying industries in the 1970s have increased their capital-labor ratios, and thus the returns to their skilled employees.

Distributional consequences of inflation and unemployment

We now move from the discussion of equilibrium unemployment and long-term changes in inequality, to the relationship between inequality and unemployment over the business cycle. We are interested in the distributional consequences of changes in demand that affect unemployment or inflation. We argue that the distributional consequences of a change in unemployment depend on the *initial* level of inequality: the more unequal a country, the more inequitable the effect of a shock to aggregate demand. This is only half of the story in thinking about monetary policy; the other half is the distributional consequences of inflation, which we discuss briefly.

Table 2
Unemployment and Earnings in the United States
(1992 – 1997 average)

	Unemployment (percent)	Usual weekly earnings
Civilian (16+)	6.1	
Black	11.6	
Hispanic	9.7	
White	5.3	
<HS Diploma (25+ years)	9.6	
HS Diploma, no college (25+ years)	5.4	
<Bachelors degree (25+ years)	4.4	
College Graduates	2.6	
Farming, Forestry and Fishing	7.9	$282
Operators, Fabricators, and Laborers	9.0	$378
Technical, Sales & Administrative Support	4.9	$428
Precision, Production, Craft & Repair	6.6	$517
Managerial & Professional Specialty	2.6	$695

Source: Bureau of Labor Statistics and calculations by the authors.

Higher unemployment is higher inequality

The business cycle and inequality are linked in a fundamental way through the indivisibility of employment and heterogeneity of economic agents. If a person's employment were continuous, then the economy could adjust to a negative shock by, say, reducing the number of hours that everyone worked by 5 percent. Given, however, that most jobs do not have continuous hours, in practice most of this adjustment is accomplished through, in the extreme case, 5 percent of the employed people losing their jobs while the other 95 percent continue to work normal hours. This is clear in the United States where three-quarters of the variance in variation in total hours worked is explained by changes in the number of people employed, with the remainder accounted for by changes in hours. By imposing

large adjustments on some people while others are relatively unaffected, this indivisibility means that *temporarily increased unemployment increases earnings disparities and thus inequality.*

This effect becomes even stronger when we take into account the initially unequal distribution of earnings. If, as is the case, the job losers are disproportionately drawn from the ranks of lower earners, then inequality will increase by even more. This gives rise to the second result: *the greater the initial degree of inequality, the more a given increase to the unemployment rate increases inequality.*

Table 2 shows evidence on the strength of this effect for the United States. It shows data on unemployment rates and earnings for various race, occupational, and educational groups. There is a clear disparity between the unemployment rates of different groups, disparities that are almost perfectly correlated with the measure of usual weekly earnings.

We can look at the same effect in terms of changes in the unemployment rate. Specifically, we ask what happens to the unemployment rate of a particular group when the overall unemployment rate rises by 1 percentage point. This is measured by running the following regression for each group indexed by *i*:

$$\Delta URATE_{t,i} = B_0 + B_1 \Delta URATE_{t,total}$$

The results are shown in Table 3.[22]

Standard economic theory has little to say on these matters—simply because conventional competitive theory says that there should be no unemployment at all. Real business cycle economists might explain these patterns in terms of different utility functions (presumably those with lower income decide to "enjoy" more leisure in recessions, because of a greater complementarity between interest rates and labor supply!)

Efficiency wage theories and search theories both provide explanations for this phenomenon. In efficiency wage theory, different

Table 3
Percentage Point Change in the Unemployment Rate for a 1 Percentage Point Change in the Overall Unemployment Rate

Civilian (16+)	1.0
Black	1.5
Hispanic	1.4
White	0.9
<HS Diploma (25+ years)	1.3
HS Diploma, no college (25+ years)	1.1
<Bachelors degree (25+ years)	0.9
College Graduates	0.4

Source: Calculations by the authors based on Bureau of Labor Statistics data. Note the regressions use annual data from the first available period to 1997. All of the coefficients are significant at the 1 percent level.

groups have different efficiency wage curves, as illustrated in Chart 5. In this simplified model, wages adjust so that the efficiency per wage unit is the same for all workers; those on a low efficiency wage curve are rationed out of the market. As the demand for labor falls, successive groups get rationed out of the market. More broadly, in an economic downturn, firms become more risk averse (see Greenwald and Stiglitz 1993); firms may feel that they have less information with which to judge the qualities of marginalized groups, and thus investments in hiring these workers are riskier. Moreover, the investment in higher productivity workers tends to be greater, so that laying off these workers involves greater risk. (There is less option value in retaining a low-skilled worker.) Thus, while hiring rates for marginalized workers are likely to be lower, the layoff rates are likely to be higher.

Similarly, search theory suggests that those with weaker search networks may be more adversely affected by a decrease in the density of jobs. This explanation is elaborated in the next section.

Chart 5
Efficiency Wages and Differential Unemployment

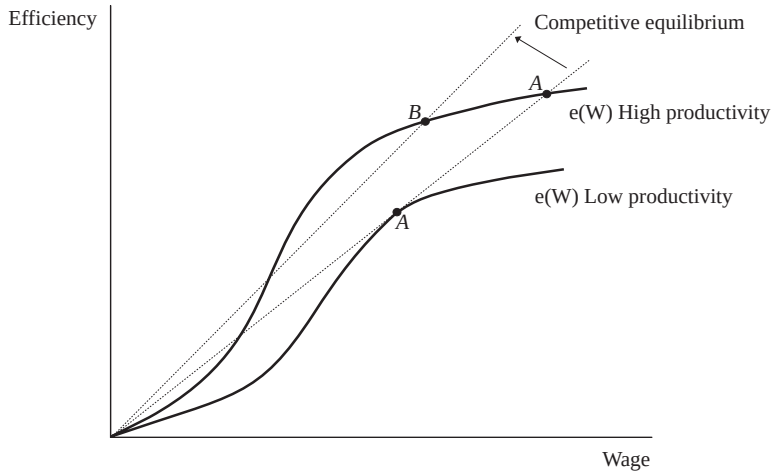

We have now traced out the elements of a vicious cycle between inequality and unemployment: high levels of unemployment lead to high levels of inequality. High levels of inequality lead to high levels of unemployment. This relationship between inequality and labor markets has implications for the persistence of unemployment, and thus for the business cycle.

The distributional consequences of inflation

We have looked at some of the consequences of inequality, unemployment, and their relationship. For the sake of completeness, we would also want to look at the distributional consequences of inflation. We can separate this into several questions. First, what are the costs of inflation? Second, do these costs vary across different income groups? And finally, does the initial degree of inequality affect the magnitude or distribution of these costs?

First, the costs of inflation generally. Although no one doubts that high inflation is costly, studies have consistently shown that there are no significant adverse effects associated with low rates of inflation (see Fischer 1993, Barro 1997, and Bruno and Easterly 1998), and some have even argued that problems may arise as the economy approaches a zero rate of inflation (Akerlof, Dickens, and Perry 1997).

There is very little research on the distribution of the costs of inflation. The conventional wisdom, however, is that the costs of inflation are disproportionately borne by the better off. The lowest quintile of the income distribution gets more of their income from Federal benefits, like Social Security, that are, for the most part, indexed against changes in inflation. In contrast, the top quintile of the income distribution gets a larger fraction of its income from financial assets, which are an imperfect hedge against unexpected inflation. The intuition about the incidence of the costs of anticipated inflation is less clear. This is an important area for more research. It is clear, however, that the distributional costs of inflation are unlikely to be skewed as much toward the poor as the distributional costs of unemployment.

Inequality and the business cycle

Finally, we discuss the relationship between inequality and the business cycle. The relationship between inequality and labor markets also has implications for the persistence of unemployment, and thus for the business cycle. Furthermore, if inequality affects the extent of credit rationing, then it could also affect the degree of aggregate consumption smoothing and thus the volatility of output.

The persistence of unemployment

Search effectiveness and the stigma effects of long-term unemployment. One important factor in determining the hiring rate is how effective workers are in searching for jobs. The ability to learn about jobs, apply for them, and present oneself as a productive worker suited for a particular job can vary greatly across people. The large

literature on hysteresis in European unemployment has documented evidence that people's search effectiveness declines over long spells of unemployment (see Layard and others 1991 for an overview). It also seems that this process is nonlinear: a short spell does little to impair a worker's effectiveness in searching for jobs, but a long spell can have seriously adverse effects. In part this may be the result of the demoralization induced by long spells of unemployment.

Another complementary mechanism that explains the fact that the hiring rate declines over the length of unemployment is adverse selection. Employers do not fully know the productivity of a worker. If the worker has been unemployed for a long time, then this signals that they are not very productive and might make the employer averse to hiring them.[23]

Inequality could exacerbate the consequences of both of these mechanisms, sometimes collectively called the "outsider mechanism," for the search effectiveness of the average person. We have seen that both the level of unemployment and changes in unemployment are distributed very unevenly across races, educational levels, and occupations and that these rates are inversely correlated with earnings—the higher the wages in your group, the less likely you are to be unemployed.

Although these data do not show the gross flows of individuals into and out of employment, it is safe to infer that the disproportionate percentage of the people who lose their jobs in an adverse shock are from disadvantaged groups or have less education. If this observation is combined with the nonlinearities in the "outsider" mechanism, the result could be that the greater the degree of inequality, the more a given shock leads people to become less effective job searchers, demoralized, and stigmatized. This could lead to more persistent changes in unemployment and possibly even to higher equilibrium unemployment. Although we have not written down a formal model, we suspect that the nonlinearities in duration dependence would mean that adverse consequences for disadvantaged groups would more than offset the opposite effect on better-off groups, leading to a higher unemployment rate overall.

Hiring networks. A similar effect could take place through hiring networks. One of the most important ways that people are matched with jobs is through their networks of friends, relatives, and associates. Informal communications can help people learn about job openings, figure out how best to apply for them, and even influence the chance that they are hired. To the extent that these networks are more extensive, the hiring rate will increase and the equilibrium unemployment rate will decrease. If inequality results in greater segmentation of the labor force, it may have an adverse effect on these networks and thus on the ability to effectively get matched with jobs. The result could be a higher equilibrium unemployment rate.

The interaction between hiring networks and inequality could also lead to greater persistence of unemployment. To describe a heuristic model, take the earlier point that when worker-hours go down by, say, 5 percent, it is mostly due to 5 percent of people losing their jobs, with some adjustment in hours. Now imagine that rather than being spread evenly from across the population, these 5 percent are all drawn from a minority group. The result is that these people's hiring networks are seriously affected. Not only are many members of this minority group out of jobs, but they now know even fewer employed people who can help match them with jobs. The result is that their search effectiveness will be diminished, and their unemployment rate will only return to its equilibrium very slowly, if ever. This effect will be partially offset by the fact that the "advantaged" group was spared its share of job reductions. At least in a relatively network hiring model, it can be shown that the former effect outweighs the latter effect and the persistence of unemployment is rising in the degree of inequality.[24]

Liquidity constraints, consumption, and labor participation

In the introduction, we noted that inequality gives rise to agency costs and is related to credit rationing. Direct survey-based measures of aggregate wealth for the United States place it between $10 and $20 trillion in 1993.[25] These numbers represent a plausible range, given that they imply capital-output ratios of 1.5 and 3 respectively. If this wealth had been divided up evenly, the average household

would have had somewhere between $100,000 and $200,000. This is a substantial multiple of the average household's income and would allow them to smooth shocks relatively easily.

The wealth, however, was distributed very unevenly. Only 42 percent of the lowest quintile of households (in terms of income) even had bank accounts; fewer than 5 percent had other interest-earning assets or stocks or mutual funds. Their median net worth in the lowest quintile, excluding home equity, was $949, less than two months worth of income. The median wealth (excluding home equity) for the entire population was $8,283. Even within this group there is substantial heterogeneity of asset holdings, with 26 percent of households in the middle quintile of the income distribution reporting no assets at financial institutions. The median household in the top quintile had a net worth (excluding home equity) of $45,392—representing more than six months of income.

There are several explanations of the stylized facts that most households hold little wealth, and that the wealth-income ratio is rising with income. In addition, access to credit is also rising in wealth and income—a standard implication of equity and credit rationing.

As a result, the huge inequality of wealth in the United States and most other capitalist economies is associated with large fractions of the population having little or no savings to insulate themselves against negative income shocks, particularly against unemployment. For such families, the marginal propensity to consume remains close to 1 for negative shocks and for increases in income that restore them to their previous levels (to their view of their permanent income). The welfare costs of given fluctuations in income is correspondingly larger. This is exacerbated by the fact that much of the incidence of an economic downturn is borne by precisely those who are worst off.

The more people that are credit rationed, and the more job loss is concentrated among credit-rationed people, the closer the economy's aggregate marginal propensity to consume is to 1, and thus demand shocks will result in greater fluctuations in output.[26] The effect of inequality on credit rationing is ambiguous. A mean-

preserving spread of income moves some people up into less credit rationing and other people down into more credit rationing. It is plausible, but far from certain, that the increase in inequality in the United States over the past 20 years has resulted from income increases for the top of the distribution (those who were not credit constrained to start with) and decreases in income for the bottom half of the distribution, pushing some additional people into credit constraints. To the degree that inequality increases credit rationing, the economy will exhibit greater volatility than if there were less inequality and more individuals had wealth they could use as collateral, and therefore, there was less credit rationing.[27]

Credit constraints, and thus potentially inequality, may have a further effect on the cyclical behavior of the economy. As the economy goes into a recession, and the unemployment rate increases, families who do not have wealth to draw upon send their second member into the labor force to search (they "invest" more in search activities), increasing the observed level of unemployment.[28] This "added worker" effect can lead to greater fluctuations of observed unemployment for given fluctuations in output.

Concluding remarks

In this paper, we have given several reasons why even if we did not care directly about inequality, but only about aggregate output, its rate of growth, and variability, we would still be concerned about inequality. There are both strong methodological and policy implications of our analysis.

For macroeconomics, representative agent models have been shown to be misleading, both as to the nature of fluctuations and their welfare costs. At the microeconomic level, the neoclassical models, which assume that issues of efficiency and distribution can be separated, have been shown to be wrong, once imperfections of information and limitations on markets are taken into account.

An understanding of the nature and consequences of fluctuations must take into account the fact that reductions in hours worked are

disproportionately borne by a few individuals (the unemployed), and that there are imperfections in capital markets, which mean that they cannot smooth the reduction in their lifetime income over their whole life. This leads to greater volatility in consumption than would be the case in a neoclassical model. (Similar issues arise at the level of the firm).

We have seen in particular that business fluctuations do have disproportionately adverse effects on those with lower incomes. This observation itself has important implications. Sometimes macroeconomic policy is posed as if there are no tradeoffs. Indeed, a major thrust of the new classical school has been to try to persuade us that there is a vertical Phillips curve and that there is nothing that can be done about the NAIRU. Posed that way, it would seem that macroeconomic management is indeed a technocratic subject, in which technocrats are simply hired to make sure that the economy stays at the NAIRU.

Even if the Phillips curve were vertical, however, there are issues of risks and tradeoffs: there is uncertainty about the level of the NAIRU, there is uncertainty about where the economy is going, and there is uncertainty about the consequences of different policies and their lags. Different groups have different attitudes toward these risks, different perceptions about the nature of these risks, and face different consequences from different policies. (Indeed, as Bayesians with different loss functions, it is not surprising that they have different perceptions about the nature of the tradeoffs.) Thus, macroeconomic policy should not be delegated to technocrats, but involves social choices: macroeconomic policy is an area of collective action, which in democratic societies requires democratic processes. Although there are legitimate issues concerning how best to combine expertise and democratic participation and how, whether, or to what extent to insulate decisionmaking from the day-to-day political processes, it is important that it should be recognized that these are not just matters of technical expertise. Different policies impact different people differently. There is no single "Pareto" optimal policy.[29]

Moreover, government policies can affect the NAIRU. We have shown that there are plausible reasons to believe that an increase in inequality may increase the equilibrium level of unemployment, and may lead to higher levels of transitional (cyclical) unemployment.

Because business fluctuations may increase inequality, and inequality may increase the magnitude of business fluctuations, there is the possibility of multiple equilibria. Governments that simultaneously pursue macroeconomic policies trying to maintain unemployment at low levels (a policy that probably has a more substantial effect on unemployment than microeconomic policies) and active redistributive policies (including through the support of education for disadvantaged groups) may push themselves into a better equilibrium, with sustained lower levels of unemployment, inequality, and instability.

These lower levels of inequality may, in turn, help support faster economic growth. While the econometric evidence on this proposition is not decisive, it is clear that the older view, that greater inequality is associated with faster growth, is not supported by the data.

While we have taken a narrow economic perspective on the consequences of inequality, it should be clear that there are broader concerns. We care about inequality because of its consequences for human well-being, and its broader consequences for the nature of society. Although it is probably true that inequality gives rise to social and political strife and such strife is not conducive to economic performance, we care about that strife not just because of these economic consequences, but for its own sake. Inequality in goods is, moreover, often translated into inequality in political power, especially in those societies that have political processes, which allow the rich and narrow but influential groups, like the media, to have undue influence in determining political outcomes. The sense that the political system is one in which, while each person may have one vote, some votes seem to count more, or are at least more effective in determining outcomes than others, leads to disillusionment with the democratic process and a decrease in participation by those who are effectively partially disenfranchised.

One of the dramatic changes in the U.S. economy over the past quarter century is the marked increase in inequality. While we may not fully know the sources of that increase, and there is a strong indication that the *trend* of increasing inequality has been arrested, it is unlikely that this massive increase will be reversed in the immediate future.

We have argued that there are good reasons to believe that there are adverse economic effects of this inequality, some of which may be masked by other trends in the economy, and that there would be economic gains from active policies, both microeconomic and macroeconomic, that explicitly take into account this increased inequality and try to reverse it. At the very least, the results presented in the first section of the paper present a persuasive argument that the reduction in inequality would not have an adverse economic effect. The positive broader benefits make a commitment to such policies all the more desirable.

Authors' Note: We wish to thank Jessica Seddon and Dennis Tao for excellent research assistance. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent.

Editor's Note: Jason Furman is associated with Harvard University and the World Bank.

Appendix

An Algebraic Treatment of the Efficiency Wage Model with Wage Inequality

This model follows the notation and assumptions of Shapiro-Stiglitz (1986). In particular, the model is set in continuous time. Workers are risk neutral and derive instantaneous utility from wages (w) and disutility from effort (e): $U(w,e) = w - e$. They have a discount rate of r.

There are many perfectly competitive firms, each with the choice of producing with two different technologies. The first is a sophisticated technology, which has a marginal product, assuming the worker exerts effort e, defined by the downward-sloping labor demand function for workers in the efficiency wage sector, $L_E = L_E(w)$ where $L'_{E(.)} < 0$. Workers that shirk produce nothing and get caught with probability q per unit of time. Assume also an exogenous job separation rate of b, due either to the destruction of jobs or workers quitting.

If the worker is employed in the efficiency wage sector, he has a choice of shirking or not shirking, each of which has a different value:

$$\begin{aligned} &\text{shirking:} && rV_E^S = w + (b+q)(V_U - V_E^S) \\ &\text{non shirking:} && rV_E^N = w + b(V_U - V_E^N) \end{aligned} \tag{1}$$

The no-shirking condition requires that $V_E^N \geq V_E^S$ in which case the value of employment V_E will be equal to the value of not shirking, V_E^N.

If the worker is not employed in the efficiency wage sector, then he can choose either a competitive job or unemployment while waiting for an efficiency-wage job:

$$\begin{aligned} &\text{competitive job:} && rV_C = w_C \\ &\text{unemployed:} && rV_U = a(V_E - V_U). \end{aligned} \tag{2}$$

where a is the (endogenous) probability per unit of time of finding a job in the efficiency wage sector which is defined, in steady state, by $aU = bL_E$.

Solving the no-shirking condition (and assuming profit maximization so that it holds with equality) and plugging in the value of unemployment (which equals the value of a competitive job) gives the following solution for the wage:

$$w = w_C + e(r+b+q)/q. \tag{3}$$

Note that, unlike Shapiro-Stiglitz, the wage is independent of the unemployment rate because of the fact that a competitive job has a constant value.

The second key equation is:

$$L_E(w) = (w_C q/be)U. \tag{4}$$

These two equations define the two key endogenous variables, w and U, as functions of the exogenous parameters. It is easy to see that unemployment and wage inequality either move together or else move orthogonally to each other. In particular, raising b (the quit rate), raising e (the cost of effort), or lowering q (the cost of being caught shirking), all lead to higher efficiency wages and higher unemployment. An increase in productivity in the efficiency wage sector shifts out $L_E(.)$ and leads to unchanged wages and higher unemployment. (If there were decreasing returns to scale in the competitive sector, then wages would rise in both sectors and unemployment would rise). Finally, improvements to productivity in the competitive sector, and increase in w_C will drive up the wage in the efficiency wage sector by an equivalent amount, leaving wage inequality unchanged and unemployment lower.

Endnotes

[1]In the presence of incomplete information, it will not be possible to have complete unemployment insurance due to adverse selection (for private arrangements) and moral hazard (for both private and public arrangements).

[2]For some of the relevant literature in these areas see: Asquith and Mullins (1986) on imperfect equity markets; Jaffee and Stiglitz (1990) on credit rationing; Greenwald and Stiglitz (1988, 1993) and the papers in Hubbard (1990) for further evidence and macroeconomic consequences; and Stiglitz and Weiss (1981), Stiglitz (1987a), Myers and Majluf (1984), and Greenwald, Stiglitz, and Weiss (1984) on imperfection information explanations of these phenomena.

[3]See Stiglitz (1997a) for an extended discussion of this issue.

[4]See Stiglitz (1978), Bevan and Stiglitz (1979), Flemming (1979), and Atkinson (1983).

[5]The relationship between inequality and agency costs (and therefore efficiency) is, however, more complex than the above analysis suggests, since there may be increasing returns to scale, which entail agglomerations of capital. When capital is more equally distributed, there may be significant agency costs associated with the creation and managing of these large enterprises.

[6]The share of tenanted land under sharecropping arrangements is 30 percent in Thailand, 50 percent in India, and 90 percent in Bangladesh (Otsuka, Chuma, and Hayami (1992) to name just three countries.

[7]More accurately, (i) a redistribution of land will result in an increase in aggregate output—equivalent to that which would come from reducing taxes from 50 percent to 0; and (ii) after the land redistribution, and the elimination of the agency problem, the economy is (constrained) Pareto efficient, while before, there was a presumption (except in the limiting case of an economy producing a single good) that it is not.

[8]See Calomiris, Hubbard, and Stock (1986) on agricultural production.

[9]While an increase in inequality is likely to increase the inefficiencies resulting from the credit constraints associated with financing education, it may decrease the inefficiencies resulting from these agency costs. For other discussions of the impact of equity constraints on growth, see Greenwald, Kohn, and Stiglitz (1990) and Stiglitz (1994a). Later, we shall argue that increases in inequality may result in greater fluctuations in the economy, and as we have noted, such fluctuations have significant adverse effects on growth.

[10]Indeed, it has been shown that there may be multiple equilibria: in one, high degrees of inequality sustain a poverty trap, while in the other, a more egalitarian distribution of income is also self-sustaining. See, for instance, Stiglitz (forthcoming).

[11]By changing incentives, taxes and transfers do not just alter the post-tax distribution of income, but alter the pre-tax distribution of earnings (see Stiglitz 1987b). Also, the instruments selected for taxation may, in turn, depend on the degree of inequality. As a result, there may be very complex links between inequality, taxation, and the levels of distortion and growth.

[12]The issues are more complex, because the marginal benefits of different kinds of public expenditures may also differ among individuals with different incomes. For a fuller treatment of these issues, see Stiglitz (1986), especially pp. 154-6. For an early discussion of the use of the median voter model to explain equilibrium tax rates, see Romer (1975).

[13]These data are, of course, sufficient to give a strong presumption against any relationship between inequality and the *level* of income.

[14]Worse still, most of the cross-country regressions do not allow us to distinguish causality: there are significant simultaneity problems, as we have noted. Even in the Kuznuts-Lewis models, inequality is both a cause and a consequence of growth. In addition to inequality, many of the other variables on the right-hand side are, in fact, endogenous and the instruments for them are often far from persuasive.

[15]It has been pointed out before that it should actually be called the NIIRU, the non-increasing inflation rate of unemployment.

[16]At the same time, it should be observed that if that were the case, reducing the degree of inequality would allow an increase in the level of employment, which would show up, in the transition to the new equilibrium, as an increase in growth.

[17]There are also important changes across groups when you cut the demographics in other ways. The unemployment rate for men went up while the unemployment rate for women went down; the unemployment rate for blacks went up while that for whites went down. The largest percentage change for any standard demographic group is the unemployment rate for black men age 20+, which nearly tripled between the early 1970s and early 1980s, and remained 40 percent above its 1973 value in 1997 (the series only started to be collected in 1972; the overall unemployment rate was 4.9 percent in 1973, making it similar to 1970).

[18]More accurately, the ratio of job seekers in the unemployment pool to jobs in the high-wage sector increases. The higher wage results in fewer jobs in the high-wage sector. If the demand curve for labor in the efficiency wage sector is relatively inelastic, or the demand curve for labor in the competitive sector is sufficiently elastic, then the number of unemployed—and thus the aggregate unemployment rate—will actually increase.

[19]For some of the literature on search models and their application to unemployment, see Diamond (1982), Mortenson (1986), and Pissarides (1990).

[20]There is an argument that goes the other way, that the increase in inequality has actually been a major factor in shifting the NAIRU down, because it has eroded the bargaining power of workers. It is ironic that many of those who are most complacent about the increase in inequality because they see the positive benefits that result from the decreased bargaining power of workers, something that should not affect a competitive labor market, are most prone to invoke the competitive model in arguing, for instance, against minimum wages.

[21]This is true in an extension of the model with decreasing returns to scale in labor in the competitive sector.

[22]The unemployment rates of different groups move very closely together, so that the

fit in these equations is extremely good. The typical *t*-statistics are 20 or higher and the R^2*s* are almost all above 80 percent.

[23]Offsetting this effect is the fact that as the unemployment spell increases, the individual's reservation wage typically falls. There is evidence that beyond a certain point, this effect is dominated by the effects going in the opposite direction.

[24]In this case it is "unemployment inequality"—either in terms of the means, the variances, or both, of the different groups. We have shown that this is strongly correlated with wage inequality and income inequality. The basic model assumes that unemployed people are matched randomly with one member of their network and get a job if the person they meet has a job. There is also an exogenous and possibly stochastic job separation rate that varies between groups. When there are two groups, it can be shown that as the difference between the job separation rates (or their variances) grows, the persistence of the unemployment rate rises.

[25]Bureau of the Census: Eller and Fraser (1995). All other figures on the distribution of wealth are from this source and, unless stated otherwise, refer to 1993.

[26]This is true, at least, in the simple Keynesian model where the demand multiplier is the reciprocal of the marginal propensity to consume.

[27]In the last decades the United States has also seen a huge increase in the access to and stock of consumer credit, in part as a result of improvements in risk scoring and securitization. The effects of this on the degree of liquidity constraints are, however, ambiguous. As consumers get larger credit lines, they also borrow more, and may not have any more of a buffer of assets plus potential credit to use in order to smooth fluctuations in their income. The rise in the personal bankruptcy rate is consistent with this hypothesis.

[28]For earlier discussions of this effect, see Woytinsky (1940). For a more recent theoretical analysis, see Basu, Genicot, and Stiglitz (1998).

[29]And even if compensation were paid, there is probably not a single "most efficient" policy.

References

Akerlof, George, William Dickens, and George Perry. "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity* 1, 1996, pp. 1-76.

Alesina, Alberto, and Dani Rodrik. "Distributive Politics and Economic Growth," *Quarterly Journal of Economics* 109, 1994, pp. 465-90.

_____, Sule Ozler, Nouriel Roubini, and Phillip Swagel. "Political Instability and Economic Growth," *Journal of Economic Growth* 1(2) (June 1996), pp. 149-87.

Arrow, Kenneth. *Social Choice and Individual Values*. New York: Wiley, 1963.

Ashenfelter, Orley, and Alan Krueger. "Estimates of Economic Returns to Schooling from a New Sample of Twins," *American Economic Review* (December 1994).

Asquith, P., and David Mullins. "Equity Issues and Stock Price Dilution," *Journal of Financial Economics*, 1986.

Atkinson, Anthony. *The Economics of Inequality*, Second Edition. Oxford: Clarendon Press, 1983.

Barro, Robert. "Democracy and Growth," *Journal of Economic Growth* 1(1), 1996, pp. 1-28.

_____. *The Determinants of Economic Growth*. Cambridge: MIT Press, 1997.

Basu, Kaushik, Garance Genicot, and Joseph E. Stiglitz. "Household Labor Supply, Unemployment, and Minimum Wage Legislation." Unpublished work in progress, July 27, 1998.

Benabou, Roland. "Inequality and Growth," in Ben Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual 1996*. Cambridge: MIT Press, 1996.

Bevan, D., and Joseph E. Stiglitz. "Intergenerational Transfers and Inequality," *Greek Economic Review* 1(1), 1979, pp. 8-26.

Birdsall, Nancy, David Ross, and Richard Sabot. "Inequality and Growth Reconsidered: Lessons from East Asia," *World Bank Economic Review* 9(3), 1995, pp. 477-508.

Blanchard, Olivier, and Lawrence Summers. "Hysteresis and the European Unemployment Problem," in Stanley Fischer, ed., *NBER Macroeconomics Annual* 1, (September 1986). Cambridge: MIT Press.

Bruno, Michael. *Crisis, Stabilization, and Economic Reform: Therapy by Consensus*. Oxford: Clarendon Press, 1993.

_____, and William Easterly. "Inflation Crises and Long-Run Growth," *Journal of Monetary Economics* 41(1), 1998, pp. 3-26.

Calomiris, Charles, R. Glenn Hubbard, and James Stock. "The Farm Debt Crisis and Public Policy," *Brookings Papers on Economic Activity* 2, 1986, pp. 441-79.

Council of Economic Advisers. *Economic Report of the President*. Washington, D.C.: Government Printing Office, 1997.

Deininger, Klaus, and Lyn Squire. "A New Data Set Measuring Income Inequality," *World Bank Economic Review* 10(3), (September 1996), pp. 565-91.

Diamond, Peter. "Aggregate Demand Management in Search Equilibrium," *American Economic Review* 90 (October), pp. 881-94.

Eller, T.J., and Wallace Fraser. *Asset Ownership of Households: 1993*. U.S. Bureau of the Census, *Current Population Reports, P70-47*. Washington, D.C.: U.S. Government Printing Office, 1995.

Fischer, Stanley. "The Role of Macroeconomic Factors in Growth," *Journal of Monetary Economics* 32, 1993, pp. 485–512.

Flemming, John. "The Effects of Earnings Inequality, Imperfect Capital Markets, and Dynamic Altruism on the Distribution of Wealth in Lifecycle Models," *Economica* 46, 1979, pp. 363-80.

Friedman, Milton. *A Theory of the Consumption Functions*. Princeton: Princeton University Press, 1957.

Greenwald, Bruce, Meir Kohn, and Joseph E. Stiglitz. "Financial Market Imperfections and

Productivity Growth," *Journal of Economic Behavior and Organization* 13(3), (June 1990), pp. 321–45.

Greenwald, Bruce, Michael Salinger, and Joseph Stiglitz. "Imperfect Capital Markets and Productivity Growth." Paper presented at NBER Conference in Vail, Colorado, April 1990.

_____, and Joseph E. Stiglitz. "Externalities in Markets with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics* 101 (May 1986), pp. 229–64.

_____. "Examining Alternative Macroeconomic Theories." *Brookings Papers on Economic Activity* 1, 1988, pp. 207–70.

_____. "Financial Market Imperfections and Business Cycles," *Quarterly Journal of Economics* 108(1), 1993, pp. 77–114.

_____. "Labor Market Adjustments and the Persistence of Unemployment," *American Economic Review* 85(2), (May 1995), pp. 219–25.

_____, Joseph E. Stiglitz, and Andrew Weiss. "Informational Imperfections in Capital Markets and Macroeconomic Fluctuations," *American Economic Review* 74(2), 1984, pp. 194–99.

Hubbard, R. Glenn. *Asymmetric Information, Corporate Finance, and Investment*. Chicago: University of Chicago Press, 1990.

Jaffee, D., and Joseph E. Stiglitz. "Credit Rationing," in B. Friedman and F. Hahn, eds., *Handbook of Monetary Economics.* Amsterdam: Elsevier Science Publishers, 1990, pp. 837-88.

Kaldor, Nicholas. "Alternative Theories of Distribution," *Review of Economic Studies* 23, 1957, pp. 83-100.

Kane, Thomas, and Cecilia Rouse. "Labor Market Returns to Two- and Four-Year College: Is a Credit a Credit and Do Degrees Matter?" *American Economic Review* 85(3), 1995, pp. 600–14.

Kotlikoff, Laurence, and Lawrence Summers. "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy* 89(4), 1981, pp. 706-32.

Krugman, Paul. "Past and Prospective Causes of High Unemployment," in *Reducing Unemployment: Current Issues and Policy Options.* Papers and proceedings from a Symposium sponsored by the Federal Reserve Bank of Kansas City in Jackson Hole, Wyoming, August 25-27, 1994.

Kuznets, Simon. "Economic Growth and Income Inequality," *American Economic Review* 45, 1955, pp. 1-28.

Layard, Richard, Stephen Nickell, and Richard Jackman. *Unemployment*. Oxford: Oxford University Press, 1991.

Londregan, John, and Keith Poole. "Poverty, the Coup Trap, and the Unconstitutional Seizure of Power," *World Politics* 92, 1991, pp. 1-24.

Lucas, Robert. *Models of Business Cycles*. Cambridge: MIT Press, 1987.

Mortenson, Dale. "Job Search and Labor Market Analysis," in Orley Ashenfelter and Richard Layard, eds., *Handbook of Labor Economics* Volume II. Amsterdam: North Holland, 1986, pp. 849-919.

Murphy, Kevin, Andrei Shleifer, and Robert Vishny. "Industrialization and the Big Push," *Journal of Political Economy* 87, 1989, pp. 1003-25.

Myers, Stewart, and Nicholas Majluf. "Corporate Financing and Investment Decisions When Firms Have Information That Investors Do Not Have," *Journal of Financial Economics* 13(2), 1984, pp. 187-221.

Otsuka, Keijiro, Hiroyuki Chuma, and Yujiro Hayami. "Land and Labor Contracts in Agrarian Economies: Theories and Facts," *Journal of Economic Literature* 30(4), (December 1992), pp. 1965-2018.

Pasinetti, Luigi. "Rate of Profit and Income Distribution in Relation to the Rate of Economic Growth," *Review of Economic Studies* 29, 1962, pp. 267-79.

Perotti, Roberto. "Growth, Income Distribution, and Democracy: What the Data Say," *Journal of Economic Growth* 1(2), (June 1996), pp. 149-188.

Pissarides, Christopher. *Equilibrium Unemployment Theory.* Oxford: Basil Blackwell, 1990.
Psacharopoulos, George. "Returns to Investment in Education: A Global Update," *World Development* 22(9), 1994, pp. 1325-43.
Rodrik, Dani. "Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses," *National Bureau of Economic Research Working Paper No. 6350*, 1998.
Romer, Thomas. "Individual Welfare, Majority Voting, and the Properties of a Linear Income Tax," *Journal of Public Economics* 14, 1975, pp. 163-85.
Schmidt-Hebbel, Klaus, and Luis Servén. "Income Inequality and Aggregate Saving: The Cross-Country Evidence." Forthcoming in *Social Inequality, Values, Growth and the State.* Ann Arbor: Michigan University Press, 1997.
Schumpeter, Joseph. *The Theory of Economic Development*. Cambridge: Harvard University Press, 1934.
Shapiro, Carl, and Joseph E. Stiglitz. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74(June 1986), pp. 433-44.
Solow, Robert. "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics*, (August 1957).
Stiglitz, Joseph E. "Equity, Taxation and Inheritance," in W. Krelle and A.F. Shorrocks, eds., *Personal Income Distribution.* Amsterdam: North-Holland Publishing Company, 1978, pp. 271-303.
_____. *Economics of the Public Sector* Second Edition. New York: Norton, 1986.
_____. "The Causes and Consequences of the Dependence of Quality on Prices," *Journal of Economic Literature* 25 (March 1987), pp. 1-48, (1987a).
_____. "Pareto Efficient and Optimal Taxation and the New Welfare Economics," in Alan Auerbach and Martin Feldstein, eds., *Handbook of Public Economics Volume 2*. Amsterdam: Elsevier Science, (1987b).
_____. "Economic Growth Revisited," *Industrial and Corporate Change* 3(1), (1994a).
_____. "Endogenous Growth and Cycles," in Y. Shionoya and M. Perlman (eds.), *Innovation in Technology, Industries, and Institutions.* Ann Arbor: The University of Michigan Press, pp.121–56, (1994b).
_____. "Central Banking in a Democratic Society," The Tinbergen Lecture, (1997a).
_____. "Reflections on the Natural Rate of Unemployment," *Journal of Economic Perspectives* 11(1), pp. 3-10, (1997b).
_____. "Duality and Development: Some Reflections on Economic Policy," in Gary Saxonhouse and T.N. Srinivasan, eds., *Development, Duality, and the International Economic Regime*. Ann Arbor: University of Michigan Press, forthcoming.
_____, and Andrew Weiss. "Credit Rationing in Markets with Imperfect Information," *American Economic Review* 71(3), 1981.
World Bank. *World Development Report 1998/99: Knowledge for Development*. Oxford: Oxford University Press, forthcoming.
Woytinsky, W. "Additional Workers and the Volume of Unemployment," *Journal of Political Economy* 48, 1940.