



WP/07/17

IMF Working Paper

Lucas vs. Lucas: On Inequality and Growth

*Juan Carlos Córdoba and
Geneviève Verdier*

IMF Working Paper

IMF Institute

Lucas vs. Lucas: On Inequality and Growth

Prepared by Juan-Carlos Córdoba and Geneviève Verdier¹

Authorized for distribution by Marc Quintyn

January 2007

Abstract

This Working Paper should not be reported as representing the views of the IMF

The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the authors and are published to elicit comments and to further debate.

Lucas (2004) asserts that “Of the tendencies that are harmful to sound economics, the most seductive, and in my opinion the most poisonous, is to focus on questions of distribution... The potential for improving the lives of poor people by finding different ways of distributing current production is *nothing* compared to the apparently limitless potential of increasing production.” In this paper we evaluate this claim using an extended version of Lucas’ (1987) welfare-evaluation framework. Surprisingly, we find that the welfare costs of inequality outweigh the benefits of growth in most cases. These calculations support the case for a research agenda that treats not only growth but also inequality as a priority.

JEL Classification Numbers: E1, E2, D3

Keywords: Welfare costs, business cycles, economic growth, inequality.

Authors’ email addresses: jcordoba@rice.edu and gverdier@imf.org

¹ We would like to thank Per Krusell, Rui Castro, Dror Goldberg, Marvin Goodfriend, David Levine, Adi Mayer, Camille Nelson, and seminar participants at the Universities of Pittsburgh, Texas A&M, Rochester, the International Monetary Fund, the 2006 National Bureau of Economic Research Summer Institute, and the 2006 Midwest Macro Meetings for helpful comments and suggestions. Juan Carlos Córdoba is an Assistant Professor at Rice University. Geneviève Verdier is an Economist in the IMF Institute. This research was conducted, in part, while the first author visited Carnegie Mellon University.

Contents	Page
I. Introduction	3
II. Welfare Analysis	8
A. The Model	8
Distribution of consumption	8
Individual and social welfare	9
Standard welfare measures	11
Alternative welfare measures.....	13
Welfare measures in terms of growth rates.....	14
Welfare gains from growth in terms of inequality.....	14
Net welfare consequences of growth and inequality	15
The shadow price of inequality.....	15
B. Calibration.....	16
C. Results	17
III. Optimal Inequality	20
A. The Planner's Problem.....	20
B. Optimal Choices.....	21
C. Calibration.....	23
D. Results.....	25
IV. Final Comments.....	26
References.....	36
Tables	
1. Parameter Values	33
2. Welfare Measures—Standard Formulation	33
3. Welfare Measures—Alternative Formulation.....	33
4. Welfare Measures—Growth and Inequality Equivalents	34
5. Optimal Inequality, Growth, and Consumption Levels.....	35
Figures	
1. Log of Average World Per Capita Consumption, 1960–2000.....	29
2. Standard Deviation of Log of Average World Per Capita Consumption, 1960–2000	30
3. Social Indifference Curves.....	31
4. Optimal Inequality and Growth	32

I. INTRODUCTION

In the 2003 Annual Report of the Federal Reserve Bank of Minneapolis, Robert Lucas asserts:

Of the tendencies that are harmful to sound economics, the most seductive, and in my opinion the most poisonous, is to focus on questions of distribution. In this very minute, a child is being born to an American family and another child, *equally valued by God* [italics ours], is being born to a family in India. The resources of all kinds that will be at the disposal of this new American will be on the order of 15 times the resources available to his Indian brother. This seems to us a terrible wrong, justifying direct corrective action, and perhaps some actions of this kind can and should be taken. But of the vast increase in the well-being of hundreds of millions of people that has occurred in the 200-year course of the industrial revolution to date, virtually none of it can be attributed to the direct redistribution of resources from rich to poor. The potential for improving the lives of poor people by finding different ways of distributing current production is *nothing* [italics in the original] compared to the apparently limitless potential of increasing production.

Lucas (2004)

The positive statements in this assertion are likely uncontroversial, but the normative conclusion is not. Lucas stresses the overwhelming importance of economic growth and dismisses redistributive issues as “poisonous” and quantitatively unimportant—basing his conclusion on the observation that economic growth helps to eliminate poverty. This view of the costs of inequality and the benefits of growth has profound implications. If correct, it provides quantitative support for eliminating or limiting costly redistributive programs and institutions. Moreover, if economic growth is all that matters for aggregate welfare, then societies should embrace growth-enhancing institutions and policies, ignoring the potentially adverse but likely minor distributive consequences.

The view that economic growth is the major determinant of social welfare is supported by Lucas’ own calculations; Lucas (1987) shocked the economics profession by showing how insignificant are the potential welfare benefits from eliminating business cycles, particularly when compared to the potential gains from increasing economic growth. Lucas’ view is not universally shared, however. Economists have considered distributional issues to be of crucial importance at least since Ricardo wrote, in a famous letter to Malthus:

Political Economy you think is an enquiry into the nature and causes of wealth—I think it should rather be called an enquiry into the laws which determine the division of the produce of industry amongst the classes who concur in its formation. No law can be laid down respecting quantity, but a tolerably correct one can be laid down respecting proportions. Every day I am

more satisfied that the former enquiry is vain and delusive, and the latter only the true objects of the science.

Ricardo, ed. Sraffa (1951, pp. 278–279)²

A correct assessment of the potential benefits of redistributive policies must be based on proper welfare analysis, such as the one that Lucas himself used (1987) to evaluate the costs of business cycles and the benefits of economic growth. Welfare analysis evaluates costs and gains in present-value terms and takes into account features of individual preferences, particularly time discounting, aversion to consumption risk and dispersion, and inequality levels—factors that may all significantly enhance the role of redistribution in promoting social welfare and reduce the role of economic growth.

The present paper evaluates Lucas' (2004) assertion using a version of Lucas' own (1987) welfare evaluation framework. We find that the potential welfare gains from reducing inequality are significant, and may actually exceed the potential gains from economic growth. This implies that maximizing social welfare is not equivalent to maximizing economic growth. Our findings are consistent with the idea that actual institutional choices reflect, to an important extent, a compromise between inequality and efficiency (Okun 1975), based on the premise that inequality has a major effect on aggregate welfare.

Lucas (1987) studies the welfare of a representative consumer who has an isoelastic utility function and a log-normal distribution of consumption. Lucas defines welfare costs in terms of the permanent compensation rate on consumption that would be required to leave the consumer indifferent between the current situation and an alternative ideal situation.³ He reports measures of the cost of business cycles for different degrees of intertemporal substitutability, but unfortunately he only reports welfare measures of economic growth for the log-utility case.

We extend his framework in three ways: we evaluate welfare gains from growth not only for the log case but also for when the intertemporal elasticity of substitution (IES) differs from one; we introduce consumption inequality across individuals; and we assess optimal inequality and optimal growth by introducing technological restrictions—a production possibility frontier for inequality, consumption levels, and growth—into the basic welfare framework.

First, we consider different degrees of intertemporal substitutability to assess the social benefits and costs of inequality and growth. This simple extension is important because welfare measures depend significantly on the intertemporal elasticity of substitution, but the

² Ricardo's view is partly explained, as Lucas (2002) points out, by the fact that classical economists did not suspect that economic growth would last.

³ Details and some caveats are given below.

empirical literature has typically regarded the IES as lower than one.⁴ The welfare gains of economic growth are smaller than the costs of business cycles if the IES is sufficiently low. For commonly used values of the IES, the gains of economic growth are significant but perhaps smaller than previously thought.

From this part of the analysis we find that the gains in welfare resulting from one additional percentage point of annual economic growth range from 2 percent to 21 percent, for the various values of the IES considered by Lucas (1987). These gains are typically lower than the 20 percent reported by Lucas. At a per-capita consumption growth rate of 2.1 percent—which was the 1960–2000 average for the 108 economies analyzed here—the welfare gains from total economic growth range from 7.6 percent to 51 percent.

Our second extension of the Lucas framework is to introduce consumption dispersion across individuals. This extension requires the potentially controversial choice of a social welfare function. Sources of controversy include the very existence of a social welfare function, ordinal versus cardinal ordering, and aggregation (Sen 1970). Fortunately, Lucas’ assertion suggests a natural choice for the welfare function. First, his quantitative claim presumes the existence of a cardinal ordering of preferences that makes interpersonal comparison of utilities possible. Second, he proposes that everyone should be “equally valued” in the social welfare function. These two observations suggest the use of a classical utilitarian or Benthamite social welfare function, which weighs everyone’s welfare equally.

We further assume, following standard practice in macroeconomics, that all individuals have identical preferences. This assumption is suggested by Lucas’ idea of treating the American and Indian children as “brothers” (and gives further support to the use of a utilitarian welfare function).

Our model thus assumes that inequality arises from unequal opportunities rather than differences in tastes. Because the individual’s utility function is concave, our social welfare function implies that any inequality is socially costly. By construction, the social welfare function penalizes consumption differences across individuals in the same way that individuals penalize consumption differences across states and time.

We use the extended model to compute the welfare costs of within-country and cross-country inequality. Measures of within-country inequality in consumption are based on US estimates from Krueger and Perri (2002). Measures of consumption inequality across countries are based on the Penn World Tables. Inequality in our model is fully characterized by the standard deviation of cross-sectional consumption. The cost of inequality is defined as the permanent compensation rate on consumption that is required to leave the social planner

⁴ Two classic examples are Hall (1988) and Campbell and Mankiw (1989), who find the IES to be close to zero.

indifferent between the observed situation and an ideal situation with no cross-sectional dispersion.

We find that the welfare costs of inequality are large and, perhaps surprisingly, probably larger than the gains from economic growth. For various values of the intertemporal elasticity of substitution, they range from 12 percent to 91 percent of current welfare for within-country inequality, and from 40 percent to almost 100 percent of current welfare for cross-country inequality.

Why do the potential welfare costs of inequality seem larger than the potential gains from growth? There are three reasons. First, most of the gains of economic growth occur in the future while the costs of inequality are borne during every period. Second, inequality is large. While the standard deviation of the log of aggregate consumption is around 1 to 2 percent (which explains the small welfare costs of business cycles uncovered by Lucas), the standard deviation of the cross-sectional distribution of the log of consumption is around 50 percent within countries and around 100 percent across countries. These estimates of inequality imply that consumption dispersion is larger than mean consumption both within and across countries. Third, commonly used values of the IES imply substantial social aversion to any source of consumption dispersion, and particularly inequality.

An alternative way to understand why inequality is so costly is to interpret the welfare function as the lifetime utility of a newborn child.⁵ Under this interpretation, the social planner cares only about newborns and values all children equally, including the “American child and his Indian brother.” Given a choice before birth, what level of growth and inequality would such a child choose, knowing that his place of birth is random? Although growth is clearly welfare-enhancing, this hypothetical child faces huge consumption risk, and a nontrivial probability of poverty, which substantially reduces his expected utility.

To gain further insight into the cost of inequality we perform counterfactual experiments. We ask: (i) how much growth would society be willing to give up for perfect equality? (ii) what level of inequality would compensate for the lack of economic growth? and (iii) what would be the welfare consequences of eliminating all growth and inequality simultaneously?

Consider these questions for an intermediate value of the IES of 1/2. In response to the first question, we find that a planner would give up 1.62 points of economic growth (out of 2.1 points) to eliminate all within-country inequality, and 4.49 points of economic growth to eliminate all cross-country inequality. In answer to the second question, a reduction of 34 percent in cross-country inequality, or of 136 percent in within-country inequality, would

⁵ Not only Lucas (2004) but Harsanyi (1953, 1955) and Vickrey (1960) provide similar interpretations of the social welfare function.

compensate for the lack of growth. On the third question, we find that eliminating all growth and cross-country inequality improves welfare—by 48 percent—but that eliminating all growth and within-country inequality reduces welfare—by 9 percent.

Overall, these results suggest that cross-country inequality is the major determinant of worldwide welfare, and that for an individual country's welfare, within-country inequality is as important a determinant as growth.

An alternative way to evaluate Lucas' assertion that the benefits from growth outweigh the benefits from reducing inequality is to compute the shadow price of inequality in terms of growth. This shadow price is the marginal willingness to substitute inequality for growth $\frac{\partial \mu}{\partial \sigma} = \frac{\partial \text{growth}}{\partial \text{inequality}}$. If Lucas is right, this price must be close to zero. Moreover, if social

decisions are nearly optimal then this shadow price provides not only an assessment of the social willingness to trade inequality for growth, but also an assessment of what is technologically feasible. We find, however, that the price of inequality is typically far from zero. For various values of the IES, it ranges from 0.026 to 6.55 for within-country inequality, and from 0.0521 to 13.11 for cross-country inequality.

These basic welfare measures provide only upper bounds to the actual gains and costs because they ignore the costs of redistributive policies and, in particular, they overlook potential tradeoffs between efficiency and inequality. Specifically, the estimated welfare gains from economic growth ignore the possibility that additional growth may foster inequality. Similarly, the estimated welfare costs of inequality ignore the real possibility that eliminating inequality may adversely affect incentives, and reduce economic growth and the level of consumption.

Our third extension to Lucas (1987) is to introduce technological constraints. For this purpose, we restrict our analysis to within-country inequality. We postulate a simple reduced-form technology that seeks to capture the major tradeoffs that a social planner faces at the country level. The technological frontier is defined in the space of inequality, growth, and consumption levels. We provide two different calibrations of this frontier using US and Scandinavian data from Aaberge and others (2002), and data from West and East Germany from various sources. We use the model to compute optimal levels of inequality and growth, along with the initial consumption level.

We find that the current US values of inequality, growth, and consumption are close to their optimal values if the IES is around 1/2. However, if this elasticity is smaller, as various studies suggest, then there is excessive inequality in the US.⁶ For example, if the IES is 1/5,

⁶ A major reason to believe that this elasticity is small is the equity-premium puzzle. See Kocherlakota (1996, p. 52).

then the optimal level of inequality in the US is close to that observed in Scandinavian countries, and the overall welfare cost of maintaining current suboptimal choices is around 15 percent.

In conclusion, we show that the gross and net costs of inequality are likely to be large. Instead of Lucas or Ricardo's, our findings support Okun's view:

I am wandering away from my usual concerns briefly to discuss an even more nagging and pervasive tradeoff, that between inequality and efficiency. It is, in my view, our biggest socioeconomic tradeoff, and it plagues us in dozens of dimensions of social policy.

Okun (1975, p. 2)

The paper is organized as follows. Chapter II develops a model for welfare analysis based on Lucas (1987), derives welfare measures, calibrates the parameters, and reports results; Chapter III extends the model of Chapter II by introducing technological restrictions, derives optimal choices, calibrates the parameters, and reports results; and Chapter IV concludes.

II. WELFARE ANALYSIS

A. The Model

Distribution of consumption

Consider a world composed of a large number of countries (a continuum), with each country equally populated by a large number of individuals (also a continuum). The size of the world population is normalized to 1. The time t consumption of a particular individual in the world is described by the autoregressive process:

$$\ln c_t = \rho \ln c_{t-1} + (1 - \rho)(a + bt) + \sigma_\eta \eta_t + \sigma_\varepsilon \varepsilon_t \quad (1)$$

where $\rho \in [0,1)$ is the persistence of consumption, a and b are constants, η_t is a country-specific shock, and ε_t is an individual-specific shock, both assumed to be independently drawn from a standard normal distribution. Individuals in the same country share the same draw of η_t . Under these assumptions, the unconditional distribution of consumption satisfies:

$$\ln c_t \sim N(a + bt, \sigma_x^2 + \sigma_y^2) \quad (2)$$

where:

$$\sigma_x^2 \equiv \frac{\sigma_\eta^2}{1 - \rho^2}; \sigma_y^2 \equiv \frac{\sigma_\varepsilon^2}{1 - \rho^2} \quad (3)$$

We assume that the law of large numbers holds, so that there is no aggregate uncertainty at the worldwide level. Furthermore, we assume that the initial distribution of consumption is given by the unconditional distribution evaluated at time $t=0$:

$$\ln c_0 \sim N(a, \sigma_x^2 + \sigma_y^2) \quad (4)$$

These assumptions imply that the worldwide distribution of consumption at any point in time is described by (2). According to this cross-sectional interpretation of (2), σ_y^2 measures the degree of within-country inequality—that is, inequality associated with individual factors—and σ_x^2 measures cross-country inequality—that is, inequality associated with country-specific factors. This framework can be used to analyze developments in a single country if $\sigma_x^2 = 0$, or in a group of countries populated by identical individuals if $\sigma_y^2 = 0$.

Parameters a and b are chosen so that aggregate worldwide consumption at time t , $E(c_t)$, equals $(1 + \lambda)(1 + \mu)^t$. This requires $a = \ln(1 + \lambda) - \frac{1}{2}(\sigma_x^2 + \sigma_y^2)$ and $b = \ln(1 + \lambda)$. In this specification, μ is the growth rate of worldwide consumption, and λ is the worldwide level of consumption at time 0, which determines the level of consumption in any subsequent period. λ is used below to measure the welfare gains and costs in different experiments, and it is set to 0 in the baseline case.

Individual and social welfare

The welfare of an individual with initial consumption c_0 is described by the expected utility function

$$U(c_0) = E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right] \quad (5)$$

where E_0 is the mathematical expectation conditional on time $t=0$ information, and u is a momentary utility function, identical for all individuals. We define social welfare as the average welfare of the society

$$W = \int U(c) dF(c) \quad (6)$$

where $F(c)$ is the fraction of the population with consumption below or equal to c at time 0. F is implicitly defined by (4).

There are at least two alternative interpretations of W . First, W is a standard utilitarian social welfare function that weighs everyone's welfare equally. Second, W is the expected welfare of a newborn child, EU . Furthermore, depending on the values of σ_x and σ_y , W refers to the world—if $\sigma_x > 0$ and $\sigma_y > 0$ —or to a particular country—if $\sigma_x = 0$. As noted in Chapter 1, our choice of the social welfare function is inspired by Lucas' (2004) suggestion that a social planner (“God”) places equal weight (“equally valued”) on individuals who are similar at birth (“American and Indian brothers”). The difference between individuals is to be found not in different tastes but in different endowments of resources (“The resources of all kinds that will be at the disposal of this new American will be on the order of 15 times the resources available to his Indian brother”).

We further simplify the problem by assuming that $u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}$, where $1/\gamma > 0$ is the intertemporal elasticity of substitution and γ is the coefficient of relative risk aversion. Given that F is given by (4), W satisfies:

$$\begin{aligned} W = EU &= EE_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right] = \sum_{t=0}^{\infty} \beta^t \frac{Ec_t^{1-\gamma} - 1}{1-\gamma} \\ &= \frac{1}{1-\gamma} \sum_{t=0}^{\infty} \beta^t \left(e^{(1-\gamma)(a+bt) + \frac{1}{2}(1-\gamma)^2(\sigma_x^2 + \sigma_y^2)t} - 1 \right) \end{aligned}$$

Substituting for the definitions of a and b , and simplifying, produces

$$W = W(\lambda, \mu, \sigma_x^2 + \sigma_y^2; \gamma, \beta) \equiv \frac{(1+\lambda)^{1-\gamma} e^{-\gamma(1-\gamma)(\sigma_x^2 + \sigma_y^2)/2}}{(1-\gamma)(1-\beta(1+\mu)^{1-\gamma})} - \frac{1}{1-\gamma} \frac{1}{1-\beta} \quad (7)$$

Accordingly, social welfare depends on the initial level of consumption, λ , the growth of consumption, μ , the total dispersion of consumption, measured by $\sigma_x^2 + \sigma_y^2$, and preference parameters γ and β . It is easy to check that inequality reduces social welfare $\left(\frac{\partial W}{\partial \sigma^2} < 0 \right)$ and growth increases social welfare $\left(\frac{\partial W}{\partial \mu} > 0 \right)$.

Notice that the degree of persistence, ρ , only affects social welfare through its effect on consumption dispersion, $\sigma_x^2 + \sigma_y^2$. Thus, social mobility only matters for social welfare to the extent that it affects the level of inequality. Furthermore, the social welfare function implicitly assumes that the coefficient of aversion to inequality equals the coefficient of

aversion to risk, and is the inverse of the intertemporal elasticity of substitution (Atkinson 1970).⁷ In other words, the planner penalizes consumption dispersion across individuals in the same way that individuals penalize consumption dispersion across states and time.

To perform welfare comparisons, we need to compute a baseline welfare level, W_0 , determined by a baseline set of parameters $[\mu_0, \sigma_{x0}^2, \sigma_{y0}^2]$. Let $\sigma_0^2 \equiv \sigma_{x0}^2 + \sigma_{y0}^2$. The baseline welfare level satisfies:

$$W_0 = W(0, \mu_0, \sigma_0^2) \tag{8}$$

The following subsections use the above framework to define a number of welfare measures. We first define welfare costs and gains as the proportional increase in consumption that is needed to leave the social planner indifferent between a baseline and an alternative situation. We then present the welfare costs of inequality as the amount of growth a planner would give up to eliminate inequality. Finally, we present some preliminary measures of the tradeoff between inequality and growth—first, by quantifying the welfare costs of eliminating all growth and all inequality simultaneously; second, in the form of a marginal rate of substitution between inequality and growth.

Standard welfare measures

Following Lucas (1987), we define four measures of social gains (or costs, if negative)— λ_μ , λ_x , λ_y , and λ_0 —as solutions to the following equations:

$$\begin{aligned} W_0 &= W(\lambda_\mu, 0, \sigma_0^2) \\ W_0 &= W(\lambda_x, \mu_0, \sigma_0^2 - \sigma_{x0}^2) \\ W_0 &= W(\lambda_y, \mu_0, \sigma_0^2 - \sigma_{y0}^2) \\ W_0 &= W(\lambda_0, \mu_0, 0) \end{aligned}$$

These values of λ are the proportional changes in consumption that would be required to leave the world planner indifferent between the baseline consumption path with welfare W_0 , and an alternative consumption path with welfare W . Notice that the λ 's affect the consumption path of the alternative situation, not the baseline situation. Thus, for example,

⁷ See also Auerbach and Hassett (1999) and Kaplow (2003) for a discussion and generalizations, and Harsanyi (1975) for a critique of the generalizations.

λ_μ is the proportional change in consumption, uniform across all periods, countries, and individuals, that is required to leave the world planner indifferent between the baseline consumption path and a path with no growth (but a higher initial consumption level, as measured by λ_μ). Notice also that $\lambda_\mu > 0$ and $\lambda_x, \lambda_y, \lambda_0 < 0$.

These measures can be interpreted in the following way: λ_μ is the welfare *gain* from economic growth, λ_x is the welfare *cost* of cross-country inequality, λ_y is the welfare *cost* of within-country inequality, and λ_0 is the welfare *cost* of total inequality. The equations above have the following simple solutions:

$$\lambda_\mu = \left[\frac{1-\beta}{1-\beta(1+\mu_0)^{1-\gamma}} \right] - 1 \quad (9)$$

and

$$\lambda_i = e^{-\gamma\sigma_i^2/2} - 1 \text{ for } i = \{x, y, 0\} \quad (10)$$

These formulas reveal three important properties of the welfare measures. First, the welfare measures depend only on the single relevant parameter. For example, λ_y depends only on σ_y^2 and not on μ or σ_x^2 . This is mainly a consequence of assuming an isoelastic utility function and a log-normal distribution of consumption. An important implication of this feature is that the welfare costs of within-country inequality and gains of economic growth are the same regardless of whether the society is the world (so that $\sigma_x^2 > 0$) or a country (so that $\sigma_x^2 = 0$).

Second, λ_μ is strictly decreasing in γ : growth is less attractive if consumers are less willing to substitute consumption intertemporally, a result that is well known in the risk-free rate puzzle literature (Kocherlakota 1996). Third, λ_i increases exponentially with γ (more risk-averse planners would gain more from inequality reduction). This result follows naturally from the fact that a more concave momentary utility function makes any dispersion of consumption more costly. These last two results imply that the welfare gains from economic growth can be made arbitrarily small and the welfare costs of inequality arbitrarily large by increasing the coefficient of risk aversion γ .

Another interesting welfare measure is the gain associated with one additional percentage point of economic growth, $\lambda_{1\%}$. This is defined as:

$$W_0 = W(\lambda_{1\%}, \mu_0 - 0.01, \sigma_0^2)$$

Using (7), $\lambda_{1\%}$ is given by

$$\lambda_{1\%} = \left[\frac{1 - \beta(1 + \mu_0 - 0.01)^{1-\gamma}}{1 - \beta(1 + \mu_0)^{1-\gamma}} \right]^{\frac{1}{1-\gamma}} - 1 \quad (11)$$

Alternative welfare measures

The welfare measures just defined cannot be easily compared to each other. For example, λ_μ is a compensation rate on a flat consumption path while λ_y is a compensation rate on a growing consumption path.⁸

By contrast, the following welfare measures have the advantage of being defined as compensation rates on the same, baseline, consumption path. Define $\hat{\lambda}_\mu, \hat{\lambda}_x, \hat{\lambda}_y, \hat{\lambda}_0$ as the solutions to the following equations:

$$\begin{aligned} W(\hat{\lambda}_\mu, \mu_0, \sigma_0^2) &= W(0, 0, \sigma_0^2) \\ W(\hat{\lambda}_x, \mu_0, \sigma_0^2) &= W(0, \mu_0, \sigma_0^2 - \sigma_{x0}^2) \\ W(\hat{\lambda}_y, \mu_0, \sigma_0^2) &= W(0, \mu_0, \sigma_0^2 - \sigma_{y0}^2) \\ W(\hat{\lambda}_0, \mu_0, \sigma_0^2) &= W(0, \mu_0, 0) \end{aligned}$$

The $\hat{\lambda}$'s can thus be expressed as percentage changes relative to the baseline consumption path rather than the alternative consumption path. These definitions imply that $\hat{\lambda}_\mu < 0$ and $\hat{\lambda}_x, \hat{\lambda}_y, \hat{\lambda}_0 > 0$. The measures can be interpreted in the following way: $\hat{\lambda}_\mu$ is the welfare *cost* of no growth, $\hat{\lambda}_y$ is the welfare *gain* of eliminating within-country inequality, $\hat{\lambda}_x$ is the welfare *gain* of eliminating cross-country inequality, and $\hat{\lambda}_0$ is the welfare *gain* of eliminating all inequality. The equations above have the following simple solutions:

$$\hat{\lambda}_\mu = \frac{1}{1 + \lambda_\mu} - 1 \quad (12)$$

⁸ Thus, for example, a case in which $\lambda_\mu = -\lambda_y$ would imply that the welfare gains of economic growth are smaller than the welfare costs of within-country inequality.

and

$$\hat{\lambda}_i = \frac{1}{1 + \lambda_i} - 1 \text{ for } i = \{x, y, 0\} \quad (13)$$

and similarly for $\hat{\lambda}_{1\%}$.

Welfare measures in terms of growth rates

Another way to compensate the planner for alternative consumption paths is through changes in the growth rate of consumption, μ , rather than through changes in the level of consumption, λ . One can define welfare-cost measures in terms of growth rates, μ_x , μ_y , and μ_0 , as follows:

$$W_0 = W(0, \mu_0 + \mu_x, \sigma_0^2 - \sigma_{x0}^2)$$

$$W_0 = W(0, \mu_0 + \mu_y, \sigma_0^2 - \sigma_{y0}^2)$$

$$W_0 = W(0, \mu_0 + \mu_0, 0)$$

According to these definitions, μ_y represents the additional percentage points of economic growth that the planner would be willing to give up in exchange for eliminating within-country inequality. Naturally, $\mu_y < 0$. Similarly, μ_x and μ_0 are the percentage points of economic growth that the planner would be willing to give up in exchange for eliminating cross-country and total consumption inequality respectively. Using (7) and the definitions above, these values of μ satisfy:

$$\mu_i = \left[\frac{1}{\beta} \left(1 - e^{\gamma(1-\gamma)\sigma_i^2/2} \left(1 - \beta(1 + \mu_0)^{1-\gamma} \right) \right)^{\frac{1}{1-\gamma}} - (1 + \mu_0) \right] \text{ for } i = \{x, y, 0\} \quad (14)$$

Welfare gains from growth in terms of inequality

An alternative way to compensate the planner for the absence of economic growth is through changes in inequality rather than consumption. For this purpose, we define the rates θ_x , θ_y , and θ_0 as follows:

$$W_0 = W(0, 0, \sigma_0^2 - \theta_x \sigma_{x0}^2)$$

$$W_0 = W(0, 0, \sigma_0^2 - \theta_y \sigma_{y0}^2)$$

$$W_0 = W(0, 0, \sigma_0^2 - \theta_0 \sigma_0^2)$$

According to these definitions, $\theta_x \times 100$ is the percentage reduction in inequality (cross-country, within-country, or total) that a planner would be willing to give up in exchange for zero economic growth. Using (7), these measures can be computed as:

$$\theta_i = \frac{2 \left[\ln(1 - \beta) - \ln(1 - \beta(1 + \mu_0)^{1-\gamma}) \right]}{\gamma(1 - \gamma) \sigma_i^2} \text{ for } i = \{x, y, 0\} \quad (15)$$

Net welfare consequences of growth and inequality

It is often argued in the literature that inequality is partly the result of providing individuals with incentives conducive to economic growth. It is therefore natural to wonder what the consequences would be of eliminating inequality and economic growth at the same time. The welfare consequences of these experiments are given by $\bar{\lambda}_x$, $\bar{\lambda}_y$, and $\bar{\lambda}_0$, defined as:

$$W_0 = W(\bar{\lambda}_x, 0, \sigma_0^2 - \sigma_{x0}^2)$$

$$W_0 = W(\bar{\lambda}_y, 0, \sigma_0^2 - \sigma_{y0}^2)$$

$$W_0 = W(\bar{\lambda}_0, 0, 0)$$

Thus, $\bar{\lambda}$ is a *net* welfare gain (cost if negative) of the current inequality-growth combination. Using the definitions above and (7), it follows that:

$$\bar{\lambda}_i = e^{-\gamma \sigma_i^2 / 2} \left[\frac{1 - \beta}{1 - \beta(1 + \mu_0)^{1-\gamma}} \right]^{\frac{1}{1-\gamma}} - 1 \text{ for } i = \{x, y, 0\}$$

The shadow price of inequality

One way to interpret Lucas' assertion is that a society should be willing to trade only a small reduction in growth for a one-point reduction in inequality. If so, the social marginal rate of substitution between inequality and growth must be close to zero. We can compute this marginal rate around (σ_{i0}, μ_0) using (7), holding $W = W_0$ and $\lambda = 0$. We find that

$$MRS_1^i = \frac{\partial \mu}{\partial \sigma_i} = \sigma_{i0} \gamma \frac{(1 + \mu_0)^\gamma - \beta(1 + \mu_0)}{\beta} \text{ for } i = \{x, y, 0\}$$

Lucas' assertion suggests that $MRS_1^i = 0$. Notice that MRS_1^i increases linearly with the degree of inequality, and exponentially with γ . One can alternatively compute the marginal rate of substitution between inequality and the consumption level, rather than growth. This shadow price is given by:

$$MRS_2^i = \frac{\partial \lambda}{\partial \sigma_i} = \sigma_{i0} \text{ for } i = \{x, y, 0\}$$

B. Calibration

In order to compute welfare measures we need to calibrate the growth rate of per-capita consumption μ_0 , the measures of inequality σ_{x0}^2 and σ_{y0}^2 , the risk-aversion coefficient γ , and the discount factor β . Figure 1 shows the unweighted and weighted averages of the log of per-capita consumption in 108 economies between 1960 and 2000. The data are from the Penn World Tables 6.1. The population size of each country is used to compute weighted averages. The unweighted average growth rate of yearly per-capita consumption is 2.3 percent and its standard deviation is 1.17 percent; the weighted average growth rate is 2.1 percent, and its standard deviation 0.99 percent. This evidence suggests that $\mu_0 = 2.1\%$.

Figure 2 shows the unweighted and weighted standard deviations of the log of per-capita consumption across countries from 1960 to 2000. The unweighted average suggests that consumption became significantly more dispersed during this 40-year period. However, the weighted average suggests that the dispersion remained roughly constant over the long term, and actually decreased during the past 25 years, after rising significantly during 1960-75. Figure 2 thus suggests that $\sigma_{x0} = 1$.

International evidence about the dispersion of consumption within countries is scarce. Some authors report measures of income dispersion within countries (for example, Bourguignon and Morrison 2002; Sala-i-Martin 2002), but not of consumption dispersion. Differences in dispersion between consumption and income may be significant. Krueger and Perri (2002, Figure 1), provide some estimates of the dispersion of individual consumption and income for the United States. They find that the standard deviation of the log of individual consumption, controlling for age and race, has been roughly constant over a 25-year period at around 0.48. The dispersion of per-capita income is about twice as large. Based on this evidence, we choose $\sigma_{y0} = 0.5$.

Numerous studies present estimates of γ but there is little agreement about its proper value. At one extreme, Hall (1988) and Campbell and Mankiw (1989) find the IES to be close to zero (or $\gamma \cong \infty$).⁹ At the other extreme, Beaudry and van Wincoop (1996) suggest that the IES is close to 1 ($\gamma \cong 1$). We follow Lucas (1987) and compute welfare measures for different values of the IES $\frac{1}{\gamma}$ or for $\gamma \in [1, 2, 5, 10, 20]$.

Finally, for the discount factor, we adopt the same value used by Lucas (1987): $\beta = 0.95$.

All parameter values are summarized in Table 1.

C. Results

Table 2 reports the welfare measures defined in Section A above for the parameters described there and for different values of γ (risk aversion) and β (the discount factor). Recall that λ_μ is the welfare gain from economic growth, λ_x is the welfare cost of cross-country inequality, λ_y is the welfare cost of within-country inequality, and λ_0 is the welfare cost of total inequality. The table reproduces Lucas' (1987) finding that the welfare gains from economic growth, as described by λ_μ and $\lambda_{1\%}$, are substantial if γ is close to 1. These gains significantly decrease as γ increases, however. For example, Lucas reports a 20 percent gain from one additional point of economic growth, but Table 2 shows that if γ takes an intermediate value of 5, the gains are 8.5 percent instead. Total gains from economic growth as measured by μ are between 11 percent and 48 percent.

The novel and surprising result reported in Table 2 is the large welfare costs that are associated with consumption inequality, both within and across countries. The costs of inequality range from around 12 percent to 92 percent for within-country inequality, and from around 40 percent to almost 100 percent for cross-country inequality. In fact, the gains from economic growth are smaller than the cost of total inequality if $\gamma \geq 1.11$, smaller than the cost of cross-country inequality if $\gamma \geq 1.28$, and smaller than the cost of within-country inequality if $\gamma \geq 3.2$.

Overall, the results reported in Table 2 reveal how important inequality is for aggregate welfare, and contradict Lucas' (2004) assertion that the benefits of growth dwarf the costs of inequality. In recent years, many authors have argued that Lucas' original calculation understates the true cost of business cycles because of his restrictive assumptions on preferences, the homogeneity of consumers, or his decision to match risk implied in

⁹ Many researchers regard the equity premium puzzle as evidence that γ is large. See Kocherlakota (1996, p. 52), for references.

aggregate rather than individual data (see Lucas 2002 or Barlevy 2005 for surveys of this literature). When some of these assumptions are relaxed, the cost of the business cycle can reach 3–4 percent of baseline consumption for households with no wealth (for example Beaudry and Pages 2001; Krusell and Smith 1999) or be as large as 12 percent with less restrictive preferences (for example Tallarini 2000). Even these estimates of the costs of business cycles seem small, however, relative to the costs of inequality.

Table 3 reports the alternative welfare measures that were introduced in Section A above. Recall that $\hat{\lambda}_\mu$ is the welfare cost of no growth, and $\hat{\lambda}_x$ is the welfare gain of eliminating cross-country inequality, and $\hat{\lambda}_y$ is the welfare gain of eliminating within-country inequality.

The message from these results is similar to that in Table 2, but the magnitudes, particularly of the welfare consequences of inequality, are even greater. These welfare measures can be interpreted as the proportional taxes (if negative) or proportional subsidies (if positive) on the baseline consumption path that would leave the planner indifferent between the after-tax (subsidy) baseline consumption path and the alternative path. Consider $\gamma = 2$, for example. Eliminating economic growth is equivalent to introducing a permanent 28 percent tax on consumption; eliminating within-country inequality is equivalent to introducing a permanent 28 percent subsidy on consumption; and eliminating cross-country inequality is equivalent to introducing a permanent 249 percent subsidy on consumption!

The major disadvantage of these welfare measures is that they ignore potential trade-offs between efficiency and inequality. Higher economic growth may cause more inequality. On the other hand, inequality may foster growth. For example, the prospect of appropriating higher returns without the fear of redistributive taxation may spur investment, effort, and innovation. The next chapter addresses this concern by introducing a technological tradeoff between inequality and consumption growth and levels into the welfare framework. In the rest of the present section, we provide an indirect assessment by studying the planner's willingness to substitute equality for growth.

A first measure of the equality-growth tradeoff is given by the reduction in growth that would perfectly offset the gains of eliminating inequality. This welfare measure was described in Section A and its values are reported in Table 4 under the labels $\mu_i, i = \{x, y, 0\}$ for different levels of risk aversion. For example, $\mu_y = -3.1$ for $\gamma = 5$ means that the planner would accept a reduction of 3.1 points in the growth rate of consumption in exchange for eliminating all within-country inequality. Since the baseline growth rate is 2.1 percent, the new growth rate would be -1 percent. All the μ -welfare measures are negative, in most cases implying negative net growth rates, which again reveals the enormous social costs of inequality.

A second measure of this tradeoff is given by the reduction in inequality that would compensate the planner for the elimination of growth. This welfare measure was introduced in Section A and its values are reported in Table 4 under the labels θ_i , $i = \{x, y, 0\}$, again for different levels of risk aversion. For example, the finding that $\theta_y = 36.9\%$ for $\gamma = 5$ means that a 36.9 percent reduction in within-country inequality would be sufficient to compensate the planner for a total lack of growth. The fact that most θ 's are below 37 percent suggests that relatively small reductions in inequality would compensate for the lack of growth.

A third measure of the equality-growth tradeoff is the welfare consequences of eliminating growth and inequality simultaneously. This welfare measure was described in Section A and its values are reported in Table 4 under the labels $\bar{\lambda}_i$, $i = \{x, y, 0\}$. For example, $\bar{\lambda}_y = -33\%$ for $\gamma = 5$ means that a large welfare gain ensues from eliminating inequality while simultaneously stopping growth. The fact that most $\bar{\lambda}_i$'s are negative suggests that there is too much inequality relative to growth.¹⁰

A final measure of the tradeoff between equality and growth is given by the social marginal rate of substitution between inequality and growth, as defined in Section A and reported in Table 4 under the label MRS_1^i for $i = \{x, y, 0\}$. Given our framework, we can construct a social indifference curve between inequality and growth, defined by $W(0, \mu, \sigma^2) = W_0$.¹¹ The slope of this indifference curve reflects the willingness of the planner to trade inequality for growth. The curve traces all combinations of growth rates and consumption dispersion for the alternative consumption paths that produce the same welfare as the baseline situation.

The top panel of Figure 3 illustrates these indifference curves for various values of γ . According to Lucas (2004), the shadow price of inequality must be close to zero, or alternatively the social indifference curve must be flat. Figure 3 shows, however, that this is

not generally the case. For example, for $\gamma = 5$, the result $MRS_1^y = \frac{\partial \mu}{\partial \sigma_y} = 0.377$ means that

the shadow price of one point of inequality, measured by σ , is around $\frac{1}{3}$ of a point of

¹⁰ For $\gamma \leq 2.5$, eliminating only within-country inequality and growth is detrimental to welfare. However, even in that case a reduction in inequality has a significant first-order effect on welfare. For example, for $\gamma = 1$, eliminating growth would imply a loss of 51 percent, but simultaneously eliminating within-country inequality reduces this loss to 33 percent.

¹¹ Solving for this indifference map produces $\sigma^2 = \frac{2}{\gamma(1-\gamma)} \ln \left[\frac{1-\beta(1+\mu_0)^{1-\gamma}}{1-\beta(1+\mu)^{1-\gamma}} \right] + \sigma_0^2$.

economic growth. Even for $\gamma = 1$ the shadow price of inequality differs significantly from zero.

Finally, Table 4 reports the marginal rate of substitution between inequality and the consumption level (MRS_2^i), as defined in Section A. This is a more relevant tradeoff for policymakers, as we discuss in the next section, because although countries with different degrees of inequality tend to have similar growth rates, their consumption levels vary widely. We find that this shadow price of inequality is also different from zero, and in most cases significantly larger than one. For example, the result that $MRS_1^y = \frac{\partial \lambda}{\partial \sigma_y} = 2.5$ for $\gamma = 5$ means that the planner would be willing to permanently give up 2.5 percent of consumption for a permanent reduction of one point of inequality.

Why is inequality so costly? One way to explain this is to interpret the welfare function $W(\lambda, \mu\sigma^2)$ as the lifetime utility of a newborn child. As described in Chapter 1, as risk aversion rises, increases in inequality must be compensated with much higher growth.

The last two welfare measures also reflect the technological tradeoff that is implied if the observed social choices are optimal. In this case, the observed choices lie on both the social indifference curve and the production possibility frontier.¹² This is illustrated in the bottom panel of Figure 3.

If social choices are suboptimal, however, we need to specify technological constraints—that is, we need to specify what levels of growth, inequality, and consumption are actually feasible—in order to determine *actual* rather than just *potential* welfare costs and gains. Chapter III addresses this issue.

III. OPTIMAL INEQUALITY

A. The Planner's Problem

In this chapter we seek to assess the optimal level of inequality by introducing technological restrictions into the welfare framework of the previous chapter. We assume that social preferences over the consumption level $1 + \lambda$, the consumption growth rate μ , and the inequality level σ^2 are described by the social welfare function (7). For convenience we rewrite this function, dropping the constant, as

¹² The optimal choice of growth and inequality is the one that both maximizes welfare—and therefore is on the social indifference curve—and is feasible—i.e., is on the production possibility frontier.

$$W = W(\lambda, \mu, \sigma_x^2 + \sigma_y^2; \gamma, \beta) \equiv \frac{(1 + \lambda)^{1-\gamma} e^{-\gamma(1-\gamma)(\sigma_x^2 + \sigma_y^2)/2}}{(1-\gamma)(1-\beta(1+\mu)^{1-\gamma})} \quad (16)$$

Furthermore, we postulate that the technological frontier is described by the function

$$\sigma = A(1 + \mu)^{\varepsilon_1} (1 + \lambda)^{\varepsilon_2}, \quad A > 0, \varepsilon_1 \geq 0, \varepsilon_2 \geq 0 \quad (17)$$

This function describes the technological tradeoffs between inequality, growth, and consumption levels. It is a reduced-form approximation of what is likely to be a very complex determination of consumption levels, growth, and inequality. For example, within-country inequality may be the result of providing proper incentives to individuals who possess private information about their own abilities (Atkinson and Lucas 1992). The socially optimal policy may induce significant inequality in order to elicit private information. A full solution of such a model is beyond the scope of this paper, and we only provide some suggestive results.¹³

At the level of cross-country inequality, the reason for a technological trade-off among inequality, growth, and consumption levels is less clear. Private information issues that could explain the need for within-country inequality do not apply at the international level. When information is public, a worldwide planner could in principle design a nondistortionary tax scheme. The major technological restriction at the worldwide level seems to be an enforcement problem: a worldwide social planner would be unable to enforce any progressive cross-country tax schemes.¹⁴ If no worldwide redistribution is feasible, then current cross-country inequality is the constrained efficient solution and not much more can be said. None of the large potential welfare gains documented in the previous section could be realized.

Since we lack any knowledge about the shape of (17) for cross-country inequality, we focus the remaining analysis on within-country inequality.

B. Optimal Choices

The planner's problem is to choose values of σ , μ , and λ in order to maximize W subject to (17). Substituting (17) into (16) produces

¹³ Solving a version of such a model would require us to address major theoretical and computational issues. One major problem is getting around the immization result to obtain a stationary distribution (see Sleet and Yeltekin 2005; Phelan 2003; and Farhi and Werning 2005 for recent contributions).

¹⁴ A Pareto improvement could be obtained by allowing countries to share their aggregate risk so that consumption at the country level would grow deterministically. Since consumption is highly persistent, the resulting cross-country dispersion of consumption with perfect risk-sharing would be similar to the current distribution. Thus, the major gains from a worldwide planner's perspective would not come from sharing country risk but from redistributing initial resources.

$$V(\lambda, \mu) = \frac{(1+\lambda)^{1-\gamma} e^{-\gamma(1-\gamma)A^2(1+\mu)^{2\varepsilon_1}(1+\lambda)^{2\varepsilon_2/2}}}{(1-\gamma)(1-\beta(1+\mu)^{1-\gamma})}$$

Maximizing this function over λ and μ respectively provides the following optimality conditions:

$$1 = \varepsilon\gamma A^2 (1+\mu)^{2\varepsilon_1} (1+\lambda^*)^{2\varepsilon_2} = \varepsilon_2\gamma\sigma^{*2} \quad (18)$$

The first condition determines the optimal degree of inequality, given by

$$\sigma^{*2} = \frac{1}{\varepsilon_2\gamma} \quad (19)$$

This expression states that optimal inequality depends positively on the intertemporal elasticity of substitution, $\frac{1}{\gamma}$, and negatively on the elasticity of inequality with respect to the level of consumption. The discount factor and, in particular, ε_1 , play no role in determining optimal inequality. Substituting (20) into (19) and solving for $1+\mu^*$ produces:

$$1+\mu^* = \left[\beta \left(1 + \frac{\varepsilon^1}{\varepsilon_2} \right) \right]^{\frac{1}{\gamma-1}} \quad (20)$$

This expression states that if $\gamma > 1$ —the empirically relevant case—then the optimal growth rate depends positively on β and ε_2 and negatively on ε_1 . Moreover, $\beta \left(1 + \frac{\varepsilon^1}{\varepsilon_2} \right) > 1$ is required for positive growth. If this condition is satisfied, then the optimal growth rate depends negatively on γ . Thus, the growth rate is higher if a society is more patient and less risk averse.

The optimal consumption level λ^* is found by substituting (19) and (20) into (17) and solving for λ^* . The solution is:

$$1+\lambda^* = \left[\frac{\sigma^*}{(1+\mu^*)^{\varepsilon_1} A} \right]^{\frac{1}{\varepsilon_2}} \quad (21)$$

Finally, optimal social welfare is given by $W^* = W(\lambda^*, \mu^*, \sigma^*)$. We can compute the welfare cost associated with the baseline choices, λ^c , as the solution $W(0, \mu_0, \sigma^2) = (1 + \lambda^c)^{1-\gamma} W^*$. Thus,

$$1 + \lambda^c = \left[\frac{W(0, \mu_0, \sigma^2)}{W^*} \right]^{\frac{1}{1-\gamma}} \quad (22)$$

C. Calibration

To compute optimal inequality we only need an estimate of ε_2 , the elasticity of inequality with respect to consumption. It follows from (17) that $\varepsilon_2 = \frac{\partial \ln \sigma}{\partial \ln(1 + \lambda)}$. This elasticity could be computed using the percentage difference in inequality relative to the percentage difference in consumption per capita for two countries or country groups with similar growth rates. A case that has been well documented is that of the US versus Scandinavian countries (Denmark, Norway, and Sweden). According to Penn World Table 6.1, per capita income in the US was around 24 percent higher than that in the Scandinavian countries, on average, in 1990, and around 25 percent higher in 2000, suggesting that the US and these Scandinavian countries have shared similar growth rates. These figures suggest that an estimate of $\lambda = 25\%$ or $\Delta(1 + \lambda) = \ln 1.25$ would be a reasonable estimate of the percentage difference in consumption.

Regarding differences in inequality, Aaberge and others (2002, Table 2), report Gini coefficients of disposable income for the US and Scandinavian countries for different years up to 1990. They find Gini coefficients of 0.346 for the US and 0.2173 for the Scandinavian countries on average in 1990. Assuming that disposable income is log-normally distributed, Gini coefficients can be transformed into standard deviations using the formula

$Gini = 2\Phi\left(\frac{\sigma}{\sqrt{2}}\right) - 1$, where Φ is the standard normal distribution. Thus

$$\sigma = \sqrt{2}\Phi^{-1}\left(\frac{1 + Gini}{2}\right)$$

Solving this equation gives the standard deviation of disposable income as 0.634 for the US and 0.39 for the Scandinavian countries, or a ratio of 1.63 between the two countries. Assuming that this ratio also applies to the distribution of consumption, one obtains $\Delta \ln \sigma = \ln 1.63$.

Given the estimated values for $\Delta \ln \sigma$ and $\Delta(1 + \lambda)$, we can approximate

$\varepsilon_2 = \frac{\partial \ln \sigma}{\partial \ln(1 + \lambda)} \cong \frac{\ln 1.63}{\ln 1.25} = 2.19$. Again, this elasticity and γ are all that is needed to compute optimal inequality.

Estimates of A , the constant in equation (17), and ε_1 , the elasticity of inequality with respect to consumption level, are needed to compute the optimal choices λ^* and μ^* and the welfare costs of current choices, λ^c . To compute A , we assume that the observed levels of σ , μ , and λ , $(\sigma_{y_0}, \mu_0, 0)$, lie on the technological frontier. In that case, A can be solved from (17), given ε_1 , as $A = \sigma_{y_0} (1 + \mu_0)^{-\varepsilon_1}$.

We provide two different estimates of ε_1 . The first of these assumes that the baseline growth rate is optimal. This assumption is likely to provide a robust estimate of ε_1 , since long-term growth does not seem to differ much across countries (Klenow and Rodriguez-Clare 2004), suggesting that current growth rates are approximately optimal. In this case, equation (20) can be used to solve for ε_1 as:

$$\varepsilon_1^a = \frac{\varepsilon_2}{\beta^{-1} (1 + \mu_0)^{\gamma-1} - 1}$$

This way of estimating ε_1 allows us to compute λ^* and λ^c , which are now denoted λ_a^* and λ_a^c to stress their dependence on ε_1^a , using formulas (21) and (22). By construction, the optimal growth rate is μ_0 .

An alternative way to estimate ε_1 is to note that it corresponds to an elasticity, that is,

$\varepsilon_1 = \frac{\partial \ln \sigma}{\partial \ln(1 + \mu)}$ as derived from (17). One way to calibrate this parameter would be to turn to

the empirical research on inequality and growth. The estimates in this literature often suggest that the relationship between inequality and growth is negative, but these results have been challenged.¹⁵ The elasticity could also be computed as the percentage difference in inequality

¹⁵ See Benabou (1996) for a survey. For example, Alesina and Rodrik (1994), as well as Persson and Tabellini (1994), present evidence that the relationship between inequality and growth is negative but Partridge (1997) disputes the robustness of their findings. Deninger and Squire (1998) argue that data in previous studies of this relationship were poor. Using improved measures, they find that asset inequality reduces long-term growth but that it is detrimental to the income growth of the poor rather than the rich. Forbes (2000) argues that these studies suffer from omitted-variable bias and that the relationship is positive.

relative to the percentage difference in gross growth rates for two similar countries with initially similar levels of consumption per capita (so that $\Delta \ln(1 + \lambda) = 0$).

Since there still does not seem to be a consensus on the sign or magnitude of ε_1 in the empirical literature, we prefer to use the natural experiment offered by the 1945 separation of East and West Germany. These former countries probably had very similar levels of per capita consumption and inequality at the end of the Second World War. The radically different social regimes that were imposed after the war caused a major divergence of per capita consumption and inequality over a period of 44 years. According to Biewen (2000, Table 2), by 1990 the log-variance of income per capita (σ_y^2) was 0.23 in West Germany and

0.1150 in East Germany. This implies a ratio of $\frac{\sigma_{ywest}}{\sigma_{yeast}} = 1.41$. Assuming the same ratio for

the standard deviation of consumption, then $\Delta \ln \sigma = \ln 1.41$.

Moreover, according to Biewen (2000, Tables 1 and 2), in 1990 mean income was 2.15 times higher in West Germany than in East Germany. This ratio is not unlike that provided by Burda and Hunt (2001, Table 3), at 2.32. It exaggerates the income gap just before the fall of the Berlin wall in 1989, because GDP in East Germany dropped dramatically after the two countries were reunified. According to Burda and Hunt (2001, Table 3), GDP in East Germany fell by 15.6 percent and 22.7 percent in 1990 and 1991 respectively, while GDP in West Germany grew by 5.7 percent and 4.6 percent during the same years. A better estimate of the ratio of mean income in West Germany to that in East Germany would be

$\frac{y_{west1989}}{y_{east1989}} = 2.32 \frac{(1 - 0.156) \times (1.0.227)}{1.057 \times 1.046} = 1.37$. Using data from Penn World Table Mark 5.6,

the ratio was 1.95 in 1970 but only 1.28 in 1988.

Overall, these numbers suggest a range for the ratio of the two countries' per-capita consumption in 1989 of between 1.4 and 1.9. Assuming a ratio of per-capita consumption of 1 in 1945 and 1.65 in 1989, the ratio of their annual growth rates is given by

$\frac{1 + \mu_{west}}{1 + \mu_{east}} = 1.65^{1/44}$. This implies that $\Delta \ln(1 + \mu) = \ln \frac{1 + \mu_{west}}{1 + \mu_{east}}$.

Thus, our second estimate of ε_1 is $\varepsilon_1^b = \frac{\Delta \ln \sigma}{\Delta \ln(1 + \mu)} = 30$.

D. Results

The first row of Table 5 reports the optimal levels of inequality for different degrees of risk aversion. The values range from 0.67 for $\gamma = 1$ to 0.15 for $\gamma = 20$. They suggest that the current degree of inequality in the US is optimal if the coefficient of risk aversion is around 2. However, if the coefficient is larger than 2, as the equity-premium puzzle suggests, then

the current level of inequality is excessive. If $\gamma = 5$, for example, then the optimal inequality level is more like that of Scandinavian countries. For the US, Figure 4 illustrates both the optimal and the actual choices for inequality and growth. At the actual choice, the production possibility frontier and the social indifference curve for inequality and growth do not touch, indicating that US consumers on average would benefit from lower inequality.

Table 5 also reports the calibrated values of ε_1^a (see row 2). These range from 3.89 for $\gamma = 20$ to 41.4 for $\gamma = 1$. The table also reports the values of λ^* and λ^c , associated with ε_1^a , labelled λ_a^* and λ_a^c respectively. By construction, the optimal growth rate for this case is $\mu_a^* = \mu_0 = 2.1\%$, regardless of γ . Consistently with the previous finding, if γ is around 2 then current US choices are approximately optimal and the welfare costs are close to zero. But at other values of γ , the welfare costs of current choices are significant. If $\gamma = 5$, then reducing inequality to an optimal level of $\sigma^* = 0.3$ would cost a reduction in consumption (λ_a^*) of 20 percent but would yield a welfare gain (λ_a^c) of 15 percent.

Finally, Table 5 reports the values of μ^* , λ^* , and λ^c associated with $\varepsilon_1^b = 0.30$, labelled μ_b^* , λ_b^* and λ_b^c respectively (see last three rows). As γ goes to 1, it becomes optimal to sacrifice all consumption ($\lambda_b^* \rightarrow -1$) in order to obtain an infinite growth rate of output (μ_b^*). Notice that $\varepsilon_1^a \approx \varepsilon_1^b$ for $\gamma = 2$. This reinforces our previous finding that current US choices for inequality, growth, and consumption are close to optimal if $\gamma = 2$. But if $\gamma \neq 2$, the current US choices have significant welfare costs. If $\gamma = 5$, the optimal inequality level is still similar to that of Scandinavian countries, at $\sigma^* = 0.3$, but now the tradeoff is between inequality and growth rather than between inequality and the consumption level. In fact, at $\gamma = 5$, the optimal policy would require consumption levels to stay around their baseline ($\lambda_b^* = -0.0108$ is close to 0), but to reduce inequality would imply a fall in the growth rate from 2.1 percent (μ_a^*) to 0.48 percent (μ_b^*).

These findings contradict Lucas's suggestion that the gains from economic growth always dwarf the costs of inequality. It is very plausible that the burden of inequality is large enough to merit some sacrifice in economic growth.

IV. FINAL COMMENTS

Children who are otherwise identical will start their lives with vast differences in resources and opportunities, depending on which country and which family they are born in. How much consumption and economic growth would a newborn child be willing to give up in order to avoid these birth lotteries? Lucas (2004) suggests that this child would give up very

little, because even if the child is born poor, economic growth would help him or her overcome poverty.

Our findings in this paper show that, on the contrary, the child may be willing to give up all growth, in order to avoid birthplace risk, and a large fraction of growth, if not all, in order to avoid family risk. The critical elements for our results are time discounting and risk aversion. Both factors downplay the role of growth for welfare, while risk aversion enhances the benefits of more equal outcomes. A third key factor is the size of the risk involved at birth, which is enormous.

The contribution of this paper is to quantify the social cost of inequality under the most standard assumptions made in macroeconomic theory. Our results suggest that societies could greatly benefit from reducing inequality. They also suggest the existence of a “big tradeoff” between inequality and efficiency, as described by Okun (1975), and help to explain why societies may not always find it best to adopt growth-enhancing institutions, particularly when inequality is large and those institutions may foster further inequality.

Societies commonly face major choices between equality on one hand and efficiency and growth on the other. The degree of progressivity of the tax system, for example, reveals a society’s willingness to trade equality for efficiency. Other examples are trade liberalization and labor market reforms, which are often regarded as beneficial for economic efficiency but detrimental to equality. Similarly, the extent of law enforcement, illustrated for example in efforts to crack down on tax evasion or informal markets, is influenced by distributional concerns at the expense of efficiency and growth. Migration policies are also strongly influenced by this tradeoff. Any correct evaluation of the institutional and policy choices made by different societies requires a proper assessment of the welfare implications of these choices, and in particular, a careful weighing of the welfare gains from efficiency and growth against the welfare costs of more inequality. Our results suggest that inequality concerns are of the utmost importance and should be explicitly considered in any aggregate evaluation of institutions.

Public discussions of the costs of inequality and the value of redistributive policies are often framed in political rather than academic terms and can lead to mistaken impressions of their effects on social welfare.¹⁶ We believe that our work provides an important first step in objectively evaluating the costs of inequality in a well-understood welfare framework.

A necessary caveat in interpreting our results is that we have not fully specified the micro-foundations of the technological restrictions for inequality, growth, and consumption levels.

¹⁶ Glaeser (2005) argues that social attitudes toward inequality in the US and Europe are shaped much less by reality than by the use of ideological language on all sides of the political spectrum.

We postulate a reduced-form technology, and calibrate it using natural experiments rather than cross-country regressions. Some estimates of the inequality-growth tradeoff appear in the empirical literature on inequality and growth, but there is still no consensus on this relationship. Since most countries seem to share the same long-run growth rate, we suspect that the main tradeoff is not between inequality and growth, but between inequality and consumption *levels*. In future research, we hope to improve our measurement of these technological constraints using panel data for inequality and consumption levels.

Finally, our exercise suggests that the following is the proper ranking of issues in macroeconomics from the point of view of their potential social welfare impact: cross-country inequality, within-country inequality, economic growth, and business cycles.

Figure 1. Log of Average World Per Capita Consumption, 1960–2000




Figure 2. Standard Deviation of Log of Average World Per Capita Consumption, 1960–2000




Figure 3. Social Indifference Curves




Figure 4. Optimal Inequality and Growth




Table 1. Parameter Values

$\mu_0 = 2.1\%$ Average per capita consumption growth rate
$\sigma_x = 1$ Standard deviation of per capita consumption across countries
$\sigma_x = 1$ Standard deviation of individual consumption within countries
$\gamma = [1, 2, 5, 10, 20]$ Coefficient of risk aversion
$\beta = 0.95$ Discount factor

Table 2. Welfare Measures—Standard Formulation

γ	1	2	5	10	20
λ_μ	0.4830	0.3908	0.2594	0.1741	0.1095
$\lambda_{1\%}$	0.2049	0.1525	0.0853	0.0480	0.0242
λ_x	-0.3965	-0.6321	-0.9179	-0.9933	-1.000
λ_y	-0.1186	-0.2212	-0.4647	-0.7135	-0.9179
λ_0	-0.4681	-0.7135	-0.9561	-0.9981	-1.000

Notes: λ_μ is the welfare gain from economic growth; $\lambda_{1\%}$ is the welfare gain associated with one additional percentage point of growth; λ_x is the welfare cost of cross-country consumption inequality; λ_y is the welfare cost of within-country inequality; λ_0 is the welfare cost of total inequality.

Table 3. Welfare Measures—Alternative Formulation

γ	1	2	5	10	20
$\hat{\lambda}_\mu$	-0.3257	-0.2810	-0.2060	-0.1483	-0.0987
$\hat{\lambda}_{1\%}$	-0.1701	-0.1323	-0.0786	-0.0458	-0.0237
$\hat{\lambda}_x$	0.6570	1.718	11.18	147.4	22052
$\hat{\lambda}_y$	0.1346	0.2840	0.8682	2.4903	11.182
$\hat{\lambda}_0$	0.8800	2.4903	21.759	517.01	268336

Notes: $\hat{\lambda}_\mu$ is the welfare cost of zero growth; $\hat{\lambda}_x$ is the welfare gain of no cross-country consumption inequality; $\hat{\lambda}_y$ is the welfare gain of no within-country consumption inequality; $\hat{\lambda}_0$ is the welfare gain of no inequality.

Table 4. Welfare Measures—Growth and Inequality Equivalents

γ	1	2	5	10	20
μ_x	-0.0268	-0.0461	-0.0337	-0.0267	-0.0237
μ_y	-0.0068	-0.0166	-0.0312	-0.0267	-0.0237
μ_0	-0.0334	-0.0517	-0.0337	-0.0267	-0.0237
θ_x	0.7803	0.3299	0.0923	0.0321	0.0104
θ_y	3.1212	1.3195	0.3690	0.1284	0.0416
θ_0	0.6242	0.2639	0.0738	0.0257	0.0083
$\bar{\lambda}_x$	-0.1050	-0.4884	-0.8966	-0.9921	-0.999
$\bar{\lambda}_y$	0.3071	0.0832	-0.3259	-0.6636	-0.9086
$\bar{\lambda}_0$	-0.2112	-0.6015	-0.9447	-0.9977	-0.9999
MRS_1^x	0.0545	0.1526	0.7345	2.7479	11.4822
MRS_1^y	0.0272	0.0763	0.3772	1.3739	5.7411
MRS_1^0	0.0609	0.1706	0.8212	3.0722	12.8375
MRS_2^x	1.0000	2.0000	5.0000	10.0000	20.0000
MRS_2^y	0.5000	1.0000	2.5000	5.0000	10.0000
MRS_2^0	1.1292	2.2361	5.5902	11.180	22.360

Notes: μ_x, μ_y, μ_0 are the points of economic growth that the planner would be willing to give up in exchange for eliminating cross-country inequality, within-country inequality, and total inequality; $\theta_x, \theta_y, \theta_0$ are the percentage reductions in cross-country inequality, within-country inequality, and total inequality that a planner would be willing to accept in exchange for zero economic growth; $\bar{\lambda}_x, \bar{\lambda}_y, \bar{\lambda}_0$ are the welfare costs (gains if negative) of simultaneously eliminating growth and cross-country inequality, or growth and within-country inequality, or growth and total inequality; $MRS_1^i = \frac{\partial \sigma}{\partial \mu}$ and $MRS_2^i = \frac{\partial \sigma}{\partial \lambda}$ for $\{x, y, 0\}$.

Table 5. Optimal Inequality, Growth, and Consumption Levels

γ	1	2	5	10	20
σ^*	0.6724	0.4778	0.3022	0.2137	0.1511
ε_1^a	41.44	29.30	15.22	8.14	3.89
μ_a^*	0.0210	0.0210	0.0210	0.0210	0.0210
λ_a^*	0.1448	-0.0205	-0.2054	-0.3217	-0.4210
λ_a^c	-0.0327	-0.0010	-0.1536	-0.4693	-0.8919
ε_1^b	30.00	30.00	30.00	30.00	30.00
μ_b^*	∞	0.0194	0.0048	0.0021	0.0010
λ_b^*	-1.0000	0.0014	-0.0108	-0.1242	-0.2408
λ_b^c	-0.9885	-0.0012	-0.2072	-0.5336	-0.8517

Notes: σ^* is optimal inequality; $\varepsilon_1^a = \frac{\partial \ln \sigma}{\partial \ln(1+\lambda)}$; μ_a^* = optimal growth; λ_a^* = optimal consumption level;

and λ_a^c = welfare cost of baseline, with values computed assuming $\sigma^* = \sigma_{y0}$; $\varepsilon_1^b = \frac{\partial \ln \sigma}{\partial \ln(1+\lambda)}$.

μ_b^* = optimal growth; λ_b^* = optimal consumption level; and λ_b^c = welfare cost of baseline, with values computed using German data.

References

- Aaberge, Rolf, Anders Björklund, Markus Jäntti, Mårten Palme, Peder J. Pedersen, Nina Smith, and Tom Wennemo, 2002, "Income Inequality and Income Mobility in the Scandinavian Countries Compared to the United States," *Review of Income and Wealth*, Vol. 48, pp. 443–69.
- Alesina, Alberto and Dani Rodrik, 1994, "Distributive Politics and Economic Growth," *Quarterly Journal of Economics*, Vol. 109, pp. 465–90.
- Atkinson, Andrew, and Lucas, Robert E., 1992, "On Efficient Distribution with Private Information," *Review of Economic Studies*, Vol. 59, pp.427–53.
- Auerbach, Alan J., and Kevin A. Hassett, 1999, "A New Measure of Horizontal Equity," NBER Working Paper No. 7035 (Cambridge, MA: National Bureau of Economic Research).
- Barlevy, Gady, 2005, "The Costs of Business Cycles and the Benefits of Stabilization," *Economic Perspectives*, Vol. 1Q, pp. 1–18.
- Beaudry, Paul, and Eric van Wincoop, 1996, "The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data," *Economica*, Vol. 63, pp. 495–512.
- Beaudry, Paul, and Carmen Pages, 2001, "The Costs of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review*, Vol. 45, pp. 1545–72.
- Bénabou, Roland, 1996, "Inequality and Growth," in *NBER Macroeconomics Annual*, O.Blanchard and S. Fisher, eds., (Cambridge, MA: Massachusetts Institute of Technology Press), pp. 11–74.
- Biewen, Martin, 2000, "Income Inequality in Germany during the 1980s and 1990s," *Review of Income and Wealth*, Vol. 46, pp. 1–19.
- Bourguignon, Francois, and Christian Morrison, 2002, "Inequality among World Citizens: 1820–1992," *American Economic Review*, Vol. 92, pp. 727–44.
- Burda, Michael, and Jennifer Hunt, 2001, "From Reunification to Economic Integration: Productivity and the Labor Market in Eastern Germany," *Brookings Papers on Economic Activity*, Vol. 2, pp. 1–92.

- Campbell, John, and N. Gregory Mankiw, 1989, "Consumption, Income, and Interest Rates: Reinterpreting the Time-series Evidence" in *NBER Macroeconomics Annual*, O. Blanchard and S. Fisher, eds., (Cambridge, MA: Massachusetts Institute of Technology Press), pp. 185–216.
- Deninger, Klaus, and Lyn Squire, 1998, "New Ways of Looking at Old Issues: Inequality and Growth," *Journal of Development Economics*, Vol. 57, pp. 259–87.
- Farhi, Emmanuel, and Ivan Werning, 2005, "Inequality, Social Discounting, and State Taxation," NBER Working Paper No. 11408 (Cambridge, MA: National Bureau of Economic Research).
- Forbes, Kristin J., 2000, "A Reassessment of the Relationship between Inequality and Growth," *American Economic Review*, Vol. 90, pp. 869–87.
- Glaeser, Edward L., 2005, "Inequality," NBER Working Paper No. 11511 (Cambridge, MA: National Bureau of Economic Research).
- Hall, Robert E., 1988, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, Vol. 96, pp. 339–57.
- Harsanyi, John C., 1953, "Cardinal Utility in Welfare Economics and in the Theory of Risk Taking," *Journal of Political Economy*, Vol. 61, pp. 434–35.
- _____, 1955, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy*, Vol. 63, pp. 309–21.
- _____, 1975, "Nonlinear Social Welfare Functions: Do Welfare Economists Have a Special Exemption From Bayesian Rationality?" *Theory and Decision*, Vol. 6, pp. 311–32.
- Kaplow, Louis, 2003, "Concavity of Utility, Concavity of Welfare, and Redistribution of Income," NBER Working Paper No. 10005 (Cambridge, MA: National Bureau of Economic Research).
- Klenow, Peter J., and Andrés Rodríguez-Clare, 2004, "Externalities and Growth," NBER Working Paper No. 11009 (Cambridge, MA: National Bureau of Economic Research).
- Kocherlakota, Narayana R., 1996, "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, Vol. 34, pp. 42–71.

- Krueger, Dirk, and Fabrizio Perri, 2002, "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," NBER Working Paper No. 9202 (Cambridge, MA: National Bureau of Economic Research).
- Krusell, Paul, and Anthony A. Smith, 1999, "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, Vol. 2, pp. 245–72.
- Lucas, Robert E., 1987, *Models of Business Cycles* (New York: Basil Blackwell).
- _____, 1988, "On the Mechanics of Economic Development," *Journal of Monetary Economics*, Vol. 22, pp. 3–42.
- _____, 2002, *Lectures on Economic Growth* (Cambridge, MA: Harvard University Press).
- _____, 2004, "The Industrial Revolution: Past and Future," *The Region* (2003 Annual Report of the Federal Reserve Bank of Minneapolis), pp. 5–20.
- Okun, Arthur M., 1975, *Equality and Efficiency: The Big Tradeoff* (Washington, DC: The Brookings Institution).
- Partridge, Mark D., 1997, "Is Inequality Harmful for Growth? Comment," *American Economic Review*, Vol. 87, pp. 1019–1032.
- Persson, Torsten, and Guido Tabellini, 1994, "Is Inequality Harmful for Growth?" *American Economic Review*, Vol. 84, pp. 600–21.
- Phelan, Christopher, 2003, "Opportunity and Social Mobility," Staff Report, Federal Reserve Bank of Minneapolis.
- Ricardo, David, 1951, *Works and Correspondence*, ed. by Piero Sraffa, 10 vols. (Cambridge: Cambridge University Press).
- Sala-i-Martin, Xavier, 2002, "The Disturbing 'Rise' of World Income Inequality," NBER Working Paper No. 8904 (April) (Cambridge, MA: National Bureau of Economic Research).
- Sen, Amartya K., 1970, *Collective Choice and Social Welfare* (San Francisco: Hoden Day).
- Sleet, Christopher, and Sevin Yeltekin, 2005, "Social Credibility, Social Patience, and Long Run Inequality," Mimeo (July).

Tallarini, Thomas D., 2000, "Risk-sensitive Real Business Cycles," *Journal of Monetary Economics*, Vol. 45, pp. 507–532.

Vickrey, William, 1960, "Utility, Strategy, and Social Decision Rules," *Quarterly Journal of Economics*, Vol. 74, pp. 507–35.