On the Welfare Gains of Growth and Welfare Costs of Inequality[*]

Juan-Carlos Córdoba[†]

Rice University and Carnegie Mellon University

Geneviève Verdier[‡]

Texas A&M University

First Version: June 2005.

This version: July 2005.

This note extends Lucas' (1987) analysis to assess welfare gains of economic growth and welfare costs of consumption inequality, both within and across countries. We find that the welfare costs of inequality are significantly larger than the gains of economic growth. While the gains of economic growth are equivalent to a permanent increase of 26% in per-capita consumption, the costs of within-country and cross-country inequality are equivalent to a permanent reduction in per-capita consumption of 45% and 90% respectively. A benevolent planner would accept a negative growth rate of 1% (instead of the baseline positive rate of 2.1%) in exchange for the elimination of all within-country inequality. The gains of economic growth are equivalent to those of reducing within-country inequality by approximately 1/3.

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

JEL Classification: E1, E2, D3

[*]Acknowledgements.

[†]Email address: jcordoba@rice.edu.

[‡]Email address: gverdier@econmail.tamu.edu.

1. Introduction

Consider the following thought experiment: a 1900 social planner is given the choice between an ever-increasing consumption path, similar to the one observed between 1900 and 2000, or constant consumption but the elimination of all consumption inequality. Which option would have the greater impact on social welfare? Economists often argue that inequality is a necessary evil since it provides incentives conducive to economic growth. In addition, growth, even if unequal, does increase the size of the pie for everyone. If so, one should expect that the welfare gains from growth should dwarf any welfare costs of inequality, and that if any trade-off exists, societies should be willing to forgo a lot of equality for additional growth. One should then conclude that even if inequality has risen within and across countries, we must be better off on average than 100 years ago. In this note, we argue that in fact it appears that the trade off is much less favorable to growth, and present evidence that the welfare costs of inequality are much larger than previously thought. Our answer to the thought experiment is that eliminating inequality would be the better option.

This note assesses the quantitative welfare gains of economic growth and welfare costs of inequality. For this purpose, we extend the framework proposed by Lucas (1987) in two dimensions. First, we consider a general isoelastic utility function instead of the log utility used by Lucas to assess the welfare gains from economic growth. This simple extension is important because the welfare gains from economic growth significantly depend on the intertemporal elasticity of substitution (IES), and the existing empirical literature typically regards this elasticity to be lower than one[1]. In fact, the welfare gains from economic growth are lower than the costs of business cycles if the IES is sufficiently low. For commonly used values of the IES, however,

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

the gains from economic growth are not insignificant but perhaps smaller than one may have thought. As in Lucas (1987), welfare measures are defined as the percentage change in consumption, uniform across all dates, individuals and values of the shocks, required to leave a benevolent social planner indifferent between a baseline and an alternative consumption path[2]. We find that the welfare gains from one additional point of economic growth per year range from 1.3% to 17%, which are lower than the 20% reported by Lucas (1987). For a per-capita consumption growth rate of 2.1% — the 1960-2000 average for the 108 Penn World Tables economies analyzed here — we find that the welfare gain from total economic growth ranges from 10% to 48%.

The second extension is to introduce consumption heterogeneity across individuals. This extension requires the choice of a social welfare function — weights attached to the utility of each member of society. We consider that the natural benchmark is to weigh everyone equally. Due to the concavity of the utility function, any inequality is thus socially costly. To determine how costly it might be, we use information about the distribution of consumption across individuals in the US from Krueger and Perri (2002), and about distribution of consumption per capita across countries from the Penn World Tables.

We find that the welfare costs of inequality are significantly — and perhaps surprisingly — larger than the gains from economic growth. They range from 12% to 90% for within-country inequality, and from 40% to almost 100% for cross-country inequality, for various values of the intertemporal elasticity of substitution. For an intermediate value of the IES of 1/5, these figures imply that a social planner would give up all economic growth and even accept a negative consumption growth rate of -1% per year in exchange for eliminating all within-country inequality. Alternatively, the welfare losses from eliminating economic growth can be compensated by reducing within-country inequality by approximately 1/3.

[2]Details and some caveats are given below.

These results are explained by the large dispersion of consumption across individuals both within countries and across countries. While the standard deviation for the log of aggregate consumption is around 1 to 2%, the standard deviation of the cross sectional distribution of log consumptions is around 50% within countries and around 100% across countries. Furthermore, social costs and gains turn out to be exponential functions of the variance of log consumption and the IES. For the planner, consumption varies due to aggregate risk but also due to consumption inequality across individuals. This second source of variation completely overshadows the first.

This note is organized as follows: Section 2 presents the theoretical framework and derives welfare measures, Section 3 describes the data, Section 4 presents the results, and Section 5 concludes.

2. The Framework

2.1. Lucas' Framework

Lucas (1987) studied the welfare of a representative individual in a country facing a random consumption sequence. The individual welfare and consumption path are described by

$$U = E\left\{\sum_{t=0}^{\infty}\beta^t \frac{1}{1-\gamma}\left(c_t^{1-\gamma}-1\right)\right\},\ 0<\beta<1,\ \gamma>0, \tag{1}$$

and

$$\ln c_t \sim N(\widetilde{c}_t, \sigma_z^2), \tag{2}$$

where $\widetilde{c}_t = \ln\left[(1+\lambda)(1+\mu)^t e^{-\frac{1}{2}\sigma_z^2}\right]$, and $1/\gamma$ is the intertemporal elasticity of substitution. Equation (2) implies that mean consumption is $(1+\lambda)(1+\mu)^t$, μ is the growth rate of consumption, and λ is a proportional shift on mean consumption affecting all periods. Moreover, the coefficient of variation of consumption is $\sqrt{e^{\sigma_z^2}-1}$.

Substituting (2) into (1) and simplifying, one obtains

$$U = U(\lambda, \mu, \sigma_z^2) \equiv \frac{(1+\lambda)^{1-\gamma}\, e^{-\gamma(1-\gamma)\sigma_z^2/2}}{1-\gamma} \frac{1}{1-\beta\,(1+\mu)^{(1-\gamma)}} - \frac{1}{1-\gamma}\frac{1}{1-\beta}. \tag{3}$$

provided that $\beta\,(1+\mu)^{(1-\gamma)} < 1$.

Given baseline levels of growth, μ_0, and consumption variance, $\sigma_{z_0}^2$, Lucas defines the welfare costs of business cycles and of an alternative growth rate μ as the values of λ_z and λ_μ which satisfy

$$U(\lambda_z, \mu_0, 0) = U(0, \mu_0, \sigma_{z_0}^2) \tag{4}$$

and

$$U(\lambda_\mu, \mu, \sigma_{z_0}^2) = U(0, \mu_0, \sigma_{z_0}^2) \tag{5}$$

respectively.

The system of equations (1) to (5) can be interpreted in two alternative ways. The first interpretation regards (1) as the expected lifetime utility of a representative individual whose consumption path follows a log normal distribution with trend. In this interpretation, σ_z^2 determines the riskiness of consumption. A second interpretation regards (1) as the average welfare of economies in a world where the cross-sectional distribution of consumptions is described by (2). Under the second interpretation, individual countries may face risk but there is no aggregate uncertainty at the world level. Under this interpretation σ_z^2 measures consumption dispersion, dispersion that may be random, but it could also be purely deterministic. Under the first interpretation of the welfare function λ_z measures the cost of uncertainty for the representative consumer, or the cost of business cycles. Under the second interpretation of the problem, λ_z measures the welfare cost of autarky, or the cost of missing markets to share country risk. More generally, λ_z measures the cost of consumption dispersion.

2.2. An Extended Version

One can extend this basic framework to incorporate different sources of cross-sectional dispersion of individual consumption. Specifically, one can regard the world-wide dispersion of individual consumptions as being partly due to cross-country differences, and partly due to within-country differences. For this purpose, assume that the cross-sectional distribution of individual consumptions in the world is described by the log-normal distribution

$$\ln c_t \sim N(\widehat{c}_t, \sigma_x^2 + \sigma_y^2),$$

where σ_x^2 is the variance associated with the cross-country dispersion of per-capita consumption, and σ_y^2 is the variance associated with within-country dispersion. One can also incorporate business cycles into this baseline deterministic description by allowing *i.i.d* shocks, as in Lucas. In that case, the distribution of consumptions could be described by

$$\ln c_t \sim N(\overline{c_t}, \sigma_w^2 + \sigma_x^2 + \sigma_y^2) \tag{6}$$

where σ_w^2 is the variance associated with business cycles, and $\overline{c_t} = \ln[(1+\lambda)\,(1+\mu)^t\ e^{-\frac{1}{2}\left(\sigma_w^2+\sigma_x^2+\sigma_y^2\right)}]$. Substituting (6) into (1), and simplifying produces

$$U(\lambda, \mu, \sigma_w^2 + \sigma_x^2 + \sigma_y^2) = \frac{(1+\lambda)^{1-\gamma}\ e^{-\gamma(1-\gamma)\left(\sigma_w^2+\sigma_x^2+\sigma_y^2\right)/2}}{1-\gamma} \times \frac{1}{1-\beta\,(1+\mu)^{(1-\gamma)}} - \frac{1}{1-\gamma}\frac{1}{1-\beta} \tag{7}$$

With the changes introduced, $U(0, \mu, \sigma_w^2+\sigma_x^2+\sigma_y^2)$ can be interpreted as the world-wide welfare function, a function that weighs everyone equally. Alternatively, it could be interpreted as the expected lifetime utility of a new-born child whose country of birth is randomly selected. Similarly, $U(0, \mu, \sigma_w^2 + \sigma_y^2)$ could be interpreted as the welfare of a particular country as long as σ_w^2 and σ_y^2 are properly chosen to represent that country.

In order to perform welfare comparisons, we need to compute a baseline welfare level, U_0, determined by a baseline set of parameters $[\mu_0, \sigma^2_{w0}, \sigma^2_{x0}, \sigma^2_{y0}]$. Let $\sigma^2_0 \equiv \sigma^2_{w0} + \sigma^2_{x0} + \sigma^2_{y0}$. The baseline welfare level satisfies

$$U_0 = U(0, \mu_0, \sigma^2_0) \tag{8}$$

2.2.1. Standard Welfare Measures

Following Lucas (1987), we can define four relevant welfare cost measures, λ_μ, λ_w, λ_x and λ_y, as the solutions to the following equations:

$$\begin{aligned} U(\lambda_\mu, 0, \sigma^2_0) &= U_0 \\ U(\lambda_w, \mu_0, \sigma^2_{x0} + \sigma^2_{y0}) &= U_0 \\ U(\lambda_x, \mu_0, \sigma^2_{w0} + \sigma^2_{y0}) &= U_0 \\ U(\lambda_y, \mu_0, \sigma^2_{w0} + \sigma^2_{x0}) &= U_0 \end{aligned}$$

The $\lambda's$ are thus defined as the percentage change of consumption in the alternative consumption path, uniform across all periods, individuals, and shocks, required to leave the world planner indifferent between the baseline consumption path and the alternative consumption path. Notice that one expects $\lambda_\mu > 0$ and $(\lambda_w, \lambda_x, \lambda_y) < 0$. These measures can be interpreted in the following way: λ_μ is the welfare *gain* of economic growth, λ_w is the welfare *cost* of business cycles, λ_x is the welfare *cost* of consumption dispersion across countries, and λ_y is the welfare *cost* of consumption dispersion within a country.

The equations above have the following simple solutions, using (7):

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

and

$$\lambda_i = e^{-\gamma\sigma^2_i/2} - 1 \text{ for } i = \{w, x, y\}\,. \tag{10}$$

The assumptions about isoelastic utility and log normal distribution provide welfare measures that depend only on a single relevant parameter. For example, λ_w does not depend on μ, σ_x^2, or σ_y^2. This implies that welfare gains or loses are as identical for the world social planner as for the local or country-specific social planner. Notice also that λ_μ is strictly decreasing in γ while λ_i increases exponentially with γ. This implies that the welfare gains of economic growth can be made arbitrarily small and the welfare costs of business cycles arbitrarily large by increasing γ. The reason for this is that a larger γ increases the concavity of the utility function which makes growth less attractive, as in produces a steep rather than a smooth consumption profile, and makes consumption dispersion more costly.

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

$$U(\lambda_{1\%}, \mu_0 - 0.01, \sigma_0^2) = U_0.$$

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

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

2.2.2. Alternative Welfare Measures

The welfare measures defined above may be difficult to compare to each other. For example, λ_μ is a compensation rate based on a consumption distribution path with no trend while λ_y is a compensation rate on a consumption distribution path with positive trend. Thus, if λ_μ were equal to $-\lambda_y$, for example, it would actually imply that the welfare gains from economic growth were lower than the welfare costs of within-country inequality. The following welfare measures have the advantage of being defined as compensations rates based on the same consumption path. Define

$\widehat{\lambda}_\mu$, $\widehat{\lambda}_w$, $\widehat{\lambda}_x$ and $\widehat{\lambda}_y$, as the solution to the following equations:

$$\begin{aligned}
U(0,0,\sigma_0^2) &= U(\widehat{\lambda}_\mu, \mu_0, \sigma_0^2) \\
U(0,\mu_0,\sigma_{x0}^2+\sigma_{y0}^2) &= U(\widehat{\lambda}_w, \mu_0, \sigma_0^2) \\
U(0,\mu_0,\sigma_{w0}^2+\sigma_{y0}^2) &= U(\widehat{\lambda}_x, \mu_0, \sigma_0^2) \\
U(0,\mu_0,\sigma_{w0}^2+\sigma_{x0}^2) &= U(\widehat{\lambda}_y, \mu_0, \sigma_0^2)
\end{aligned}$$

The $\widehat{\lambda}'s$ are thus defined as the percentage change of the baseline consumption path, uniform across all periods, individuals, and shocks, required to leave the planner indifferent between the modified baseline consumption path and the alternative consumption path. Notice that one expects now $\widehat{\lambda}_\mu < 0$ and $(\widehat{\lambda}_w, \widehat{\lambda}_x, \widehat{\lambda}_y) > 0$. These measures can be interpreted in the following way: $\widehat{\lambda}_\mu$ is the welfare *cost* of zero growth, λ_w is the welfare *gain* of no business cycles, λ_x is the welfare *gain* of no consumption dispersion across countries, and λ_y is the welfare *gain* of no consumption dispersion within a country.

The equations above have the following simple solutions, using (7):

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

and

(13)
$$\widehat{\lambda}_i = \frac{1}{1+\widehat{\lambda}_i} - 1 \text{ for } i = \{w, x, y\}\,.$$

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

2.2.3. Welfare Measures in Terms of Growth Rates

An alternative way to compensate the planner for alternative consumption paths is through changes in the growth rate of consumption, μ, rather than proportional changes in consumption, λ. One can define welfare costs measures in terms of growth rates, μ_w, μ_x and μ_y, as follows:

$$\begin{aligned} U(0, \mu_0 + \mu_w, \sigma_{x0}^2 + \sigma_{y0}^2) &= U_0 \\ U(0, \mu_0 + \mu_x, \sigma_{w0}^2 + \sigma_{y0}^2) &= U_0 \\ U(0, \mu_0 + \mu_y, \sigma_{w0}^2 + \sigma_{x0}^2) &= U_0 \end{aligned}$$

According to these definitions, μ_w are the additional points of economic growth that the planner would be willing to accept in exchange for eliminating business cycles. Naturally, $\mu_w < 0$. Similarly, μ_y and μ_x are the additional points of economic growth that the planner would be willing to accept to in exchange for eliminating within-country and cross-country consumption inequality respectively. Using (7) and the definitions above, these values of μ satisfy:

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

2.2.4. Welfare Gains of Growth in Terms of Inequality

It is also interesting to measure the gains of economic growth in terms of an inequality-equivalent rate. It can be defined as:

$$U(0, 0, \sigma_0^2(1 - \theta)) = U_0.$$

where $\theta \times 100$ is the percentage reduction in total inequality, as defined by the total variance of consumption, that a planner would be willing to accept in exchange for zero economic growth. An alternative is to measure the gain only as percentage of within-country inequality, a more relevant measure for a local social planner, as follows:

$$U(0, 0, \sigma_{w0}^2 + \sigma_{x0}^2 + \sigma_{y0}^2(1 - \theta_y)) = U_0.$$

Using these definitions and (7), θ and θ_y satisfy:

$$\theta = \frac{2\left[\ln(1-\beta) - \ln\left(1-\beta(1+\mu_0)^{(1-\gamma)}\right)\right]}{\gamma(1-\gamma)\sigma_0^2}, \tag{15}$$

$$\theta_y = \frac{2\left[\ln(1-\beta) - \ln\left(1-\beta(1+\mu_0)^{(1-\gamma)}\right)\right]}{\gamma(1-\gamma)\sigma_{y0}^2}. \tag{16}$$

2.2.5. Net Gains from Economic Growth?

It is often argued in the literature that inequality is partly the result of providing individuals with the proper incentives conducting to economic growth. It is thus natural to ask what would be the welfare consequences of eliminating all inequality and at the same time eliminating all economic growth. The welfare consequences of this experiment are given by λ^*, defined as:

$$U(\lambda^*, 0, 0) = U_0.$$

Thus, λ^* is the welfare cost (gain if negative) of eliminating all growth and all inequality. A more relevant measure for a local social planner is λ_y^* as defined by:

$$U(\lambda_y^*, 0, \sigma_w^2 + \sigma_x^2) = U_0$$

λ_y^* is the welfare cost (gain if negative) of eliminating all growth and all within-country inequality. Using the definitions above and (7), one finds that

$$\lambda^* = e^{-\gamma\sigma_0^2/2}\left[\frac{1-\beta}{1-\beta(1+\mu_0)^{(1-\gamma)}}\right]^{\frac{1}{1-\gamma}} - 1 \tag{17}$$

$$\lambda_y^* = e^{-\gamma\sigma_y^2/2}\left[\frac{1-\beta}{1-\beta(1+\mu_0)^{(1-\gamma)}}\right]^{\frac{1}{1-\gamma}} - 1 \tag{18}$$

2.2.6. The Growth-Inequality Trade-off

Finally, we can also compute a growth-inequality trade off as follows:

$$U(0, \mu, \sigma^2) = U_0$$

This indifference curve traces all combination of growth rates and consumption variances of alternative consumption paths that produce the same welfare as the baseline situation. The slope of this map determines the willingness of the planner to trade inequality for growth. Solving for this indifference map produces

$$\sigma^2 = \frac{2}{\gamma(1-\gamma)} \ln\left[\frac{1-\beta(1+\mu_0)^{(1-\gamma)}}{1-\beta(1+\mu)^{(1-\gamma)}}\right] + \sigma_0^2. \tag{19}$$

3. Data

Figure 1 shows the unweighed and weighed averages of the log of consumption per capita in 108 economies between 1960 and 2000. The data is from the Penn World Tables 6.1. Population size from each country is used to compute weighed averages. The average growth rate of yearly per-capita consumption is 2.3% and its standard deviation 1.17% for the unweighed average, and 2.1% and 0.99% respectively for the weighed average.

Figure 2 shows the unweighed and weighed standard deviations of the log of per-capita consumption across countries from 1960 to 2000. The unweighed and weighed averages provide different characterization about the evolution of this cross-country dispersion. The unweighed average suggests a significant increase of consumption dispersion during this 40-year period. However, the weighed average suggests that the dispersion has remained roughly constant over the long term, and in fact has actually decreased during the past 25 years, after increasing significantly during the 1960-1975 period.

Evidence about the dispersion of per-capita consumption within particular countries is scarce. There are is some evidence about the dispersion of income per-capita within countries (e.g. Bourguignon and Morrison (2002), Sala-i-Martin (2002)), but they do not provide evidence about the dispersion of per-capita consumption. The differences between consumption dispersion and income dispersion may be significant. Krueger and Perri (2002, Figure 1) provide some estimates about the dispersion of per-capita consumption and income for the United States. They find that the standard deviation of log per-capita consumption has been roughly constant in a 25 year period at around 0.48. The dispersion of per-capita income is at around twice as large.

The previous evidence suggests the following choice of parameters for our exercise: the baseline average per-capita consumption growth rate is $\mu_0 = 2.1$, the standard deviation of per-capita consumption over time is $\sigma_w = 0.01$, the standard deviation of cross-country per-capita consumption is $\sigma_x = 1$ and the standard deviation of within-country per-capita consumption is $\sigma_y = 0.5$. As for the parameter γ, we consider different possible values, in particular the same ones considered by Lucas (1987.) Finally, $\beta = 0.95$ is assumed.

4. Results

Table 1 reports the welfare measures described in Section 2.2.1 for the parameters described above and for different values of γ. Recall that λ_μ is the welfare gain of economic growth, λ_w is the welfare cost of business cycles, λ_x is the welfare cost of consumption dispersion across countries, and λ_y is the welfare cost of consumption dispersion within a country. The table reproduces the now well-known result that the welfare costs of business cycles, λ_w, are very small. Similarly, as reported by Lucas (1987), the welfare gains from economic growth, as described by λ_μ and $\lambda_{1\%}$, are substantial for γ close to 1. The gains, however, significantly decrease as γ increases. For example, Lucas reports a 20% gain from 1 additional point of economic growth, but the Table shows that if γ takes an intermediate value of 5, the gains are 8.5% instead.

The novel and perhaps surprising result reported in Table 1 is the large welfare costs associated with consumption inequality, both within and across countries. These costs range from around 12% to 92% for within-country inequality, and from around 40% to almost 100% for cross-country inequality. These figures are just staggering and reveal how important inequality is for aggregate welfare. 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, homogeneity of consumers, or his choice of matching the risk implied in aggregate rather than individual data (see Lucas (2003) or Barlevy (2005) for a survey of this literature). When some of these assumptions are relaxed, the cost of business cycle can reach 3-4% for households with no wealth (e.g. Beaudry and Pages (2001) or Krusell and Smith (1999)) or be as large as 12% with less restrictive preferences (e.g. Tallarini (2000)). But even these estimates of the costs of business cycles seem small relative to the costs of inequality. It is also remarkable that the welfare costs of inequality are systematically larger than the benefits of economic growth, which themselves are substantial.

Table 2 reports the alternative welfare measures described in Section 2.2.2. Recall that $\widehat{\lambda}_{\mu}$ is the welfare cost of zero growth, λ_w is the welfare gain of no business cycles, λ_x is the welfare gain of no consumption dispersion across countries, and λ_y is the welfare gain of no consumption dispersion within a country. The message is similar to the one in Table 1, but the magnitudes, particularly regarding the welfare consequences of inequality, are even more staggering. These welfare measures can be interpreted as proportional taxes (if negative) or proportional subsidies (if positive) on the baseline consumption path that leaves the planner indifferent between the after-tax (subsidy) baseline consumption path and the alternative path. Consider $\gamma = 5$ for example: eliminating economic growth is equivalent to introducing a permanent 20% tax on consumption, eliminating within-country is equivalent to introducing a permanent 86% subsidy on consumption, and eliminating cross-country inequality is equivalent to introducing a permanent 1,118% subsidy on consumption (!)

Table 3 reports welfare measures in terms of growth rates as described in Section 2.2.3. Again for $\gamma = 5$, $\mu_y = -3.1\%$, which means that the planner would be willing to accept a reduction of 3.1 points in the growth rate of consumption in exchange for eliminating all within-country inequality. In that case, the growth rate would be -1% instead of the 2.1% in the benchmark. Alternatively, $\theta_y = 0.36$ means that the planner would be willing to trade a 36% reduction in the within-country variance of consumption in exchange for zero economic growth.

Finally, Figure 3 shows the various combinations of growth and inequality that keep social welfare constant as computed in (19) for various values of γ. This picture is consistent with the previous discussion. The slope of these indifference curves measures the planner's willingness to trade inequality for growth. This willingness to trade inequality for growth falls with the intertemporal elasticity of substitution $1/\gamma$. As the coefficient of risk aversion γ rises, the planner requires more compensation for higher growth because the representative consumer is weary of more variable consumption paths. Evaluated at $\mu = 0.021$, $\beta = 0.95$ and $\gamma = 1.001$, this slope is -0.002 but falls to -27.23 for $\gamma = 5$.

Why is inequality so costly? One way to explain the magnitude of the cost is to interpret the welfare function $U(\lambda, \mu, \sigma^2)$ as the lifetime utility of a newborn child. If given a choice before birth, what level of growth and cross-country inequality would such a child choose behind this Rawlsian 'veil of ignorance' knowing that his country of birth is random? Although growth is clearly welfare-enhancing, this hypothetical child faces a non-trivial probability of poverty, which reduces expected utility. As risk aversion increases, increases in inequality must be compensated with much higher growth.

5. Final Comments

This note emphasizes the importance of better understanding a macroeconomic problem that has not received enough attention in the literature. Lucas' (1987) results first suggested that the value added of better understanding economic growth far exceeded the value added of research on business cycles. Fortunately, research in economic growth has became increasingly more important in recent years, and this shift in focus is probably due in part to Lucas' calculations.

This note highlights the importance of another macroeconomic issue: consumption inequality. A surprising result is that the potential gains associated with a better understanding of inequality surpass those from understanding economic growth, which are themselves large. Naturally, these potential gains are just upper bounds to the actual gains. However, even if the actual gains are much smaller, these large upper bounds suggests that these issues deserve much more attention than they currently receive.

Figure 1: Log of Average World Per Capita Consumption, 1960-2000

Source: Penn World Tables 6.1 and authors calculations

Figure 2: Standard Deviation of Log of Average World Per Capita Consumption, 1960-2000

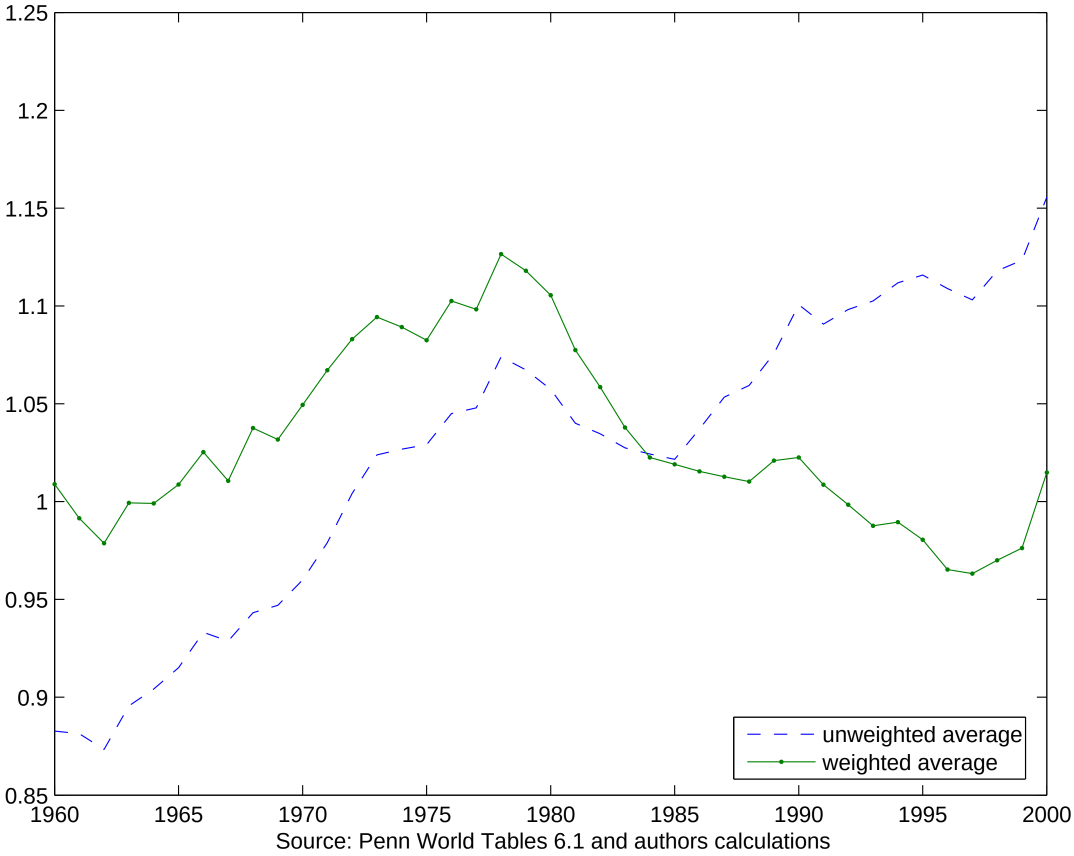

Table 1: Welfare Measures - Standard Formulation

	γ				
	1.0001	2	5	10	20
λ_μ	0.4842	0.3908	0.2594	0.1741	0.1095
$\lambda_{1\%}$	0.2056	0.1525	0.0853	0.0480	0.0242
λ_w	-0.0001	-0.0001	-0.0002	-0.0005	-0.0010
λ_x	-0.3935	-0.6321	-0.9179	-0.9933	-1.0000
λ_y	-0.1175	-0.2212	-0.4647	-0.7135	-0.9179

Table 2: Welfare Measures - Alternative Formulation

	γ				
	1.0001	2	5	10	20
$\widehat{\lambda}_\mu$	-0.3262	-0.2810	-0.2060	-0.1483	-0.0987
$\widehat{\lambda}_{1\%}$	-0.1705	-0.1323	-0.0786	-0.0458	-0.0236
$\widehat{\lambda}_w$	0.0001	0.0001	0.0002	0.0005	0.0010
$\widehat{\lambda}_x$	0.6488	1.7181	11.180	148.25	22025
$\widehat{\lambda}_y$	0.1331	0.2840	0.8681	2.4904	11.180

Figure 3: Growth-Inequality Trade-Off

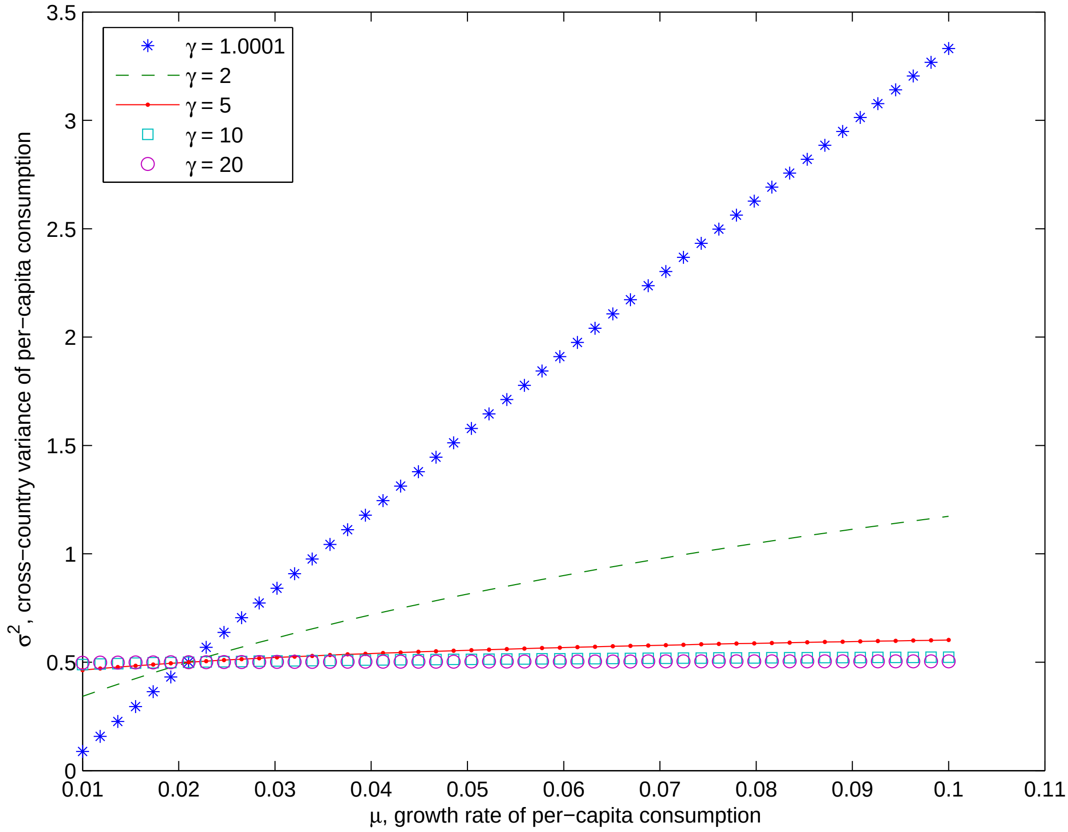

Table 3: Welfare Measures - Growth and Inequality Equivalents

	γ				
	1.0001	2	5	10	20
μ_w	-0.0000	-0.0000	-0.0000	-0.0001	-0.0006
μ_x	-0.0265	-0.0461	-0.0337	-0.0267	-0.0237
μ_y	-0.0067	-0.0166	-0.0312	-0.0267	-0.237
θ	0.3463	0.1447	0.0405	0.0141	0.0046
θ_y	3.1586	1.3195	0.3690	0.1284	0.0416
λ^*	-0.5254	-0.8578	-0.9958	-1.0000	-1.0000
λ_y^*	0.3098	0.0832	-0.3259	-0.6636	-0.9089

References

Beaudry, P. and Pages, C. (2001), "The Costs of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review*, 45: 1545-1572.

Bourguignon, Francois and Christian Morrison (2002), "Inequality among world citizens: 1820-1992," *American Economic Review*, 92(4), September

Campbell, J. and Mankiw, G. (1989), "Consumption, income, and interest rates: reinterpreting the time series evidence." In O. Blanchard and S. Fisher (eds), *NBER Macroeconomic Annual 1989:* 185-216. Cambridge, Mass.: NBER.

Hall, R.E. (1988), "Intertemporal substitution in consumption," *Journal of Political Economy*, 96: 339-57.

Krueger, Dirk; and Perri, Fabrizio (2002), "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," NBER Working Paper 9202.

Krusell, P., and Smith, A. (2002), "Inequality among world citizens: 1820-1992," *American Economic Review*, 92(4), September.

Lucas, Robert E, *Model of Business Cycles*, Basil Blackwell, New York, 1987.

Sala-I-Martin (2002), "The Disturbing "Rise" of World Income Inequality," NBER Working paper 8904, April 2002.

Tallarini, T. (2000), "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics,* 45: 507-532.