Inequality and Growth: Some Welfare Calculations*

Juan Carlos Córdoba[†] and Geneviève Verdier[‡]

June 18, 2007

Abstract

The main lotteries individuals face during their lifetime is first, their place of birth, and the second, their parents. Children that are ex-ante identical in all respects start their life with vast differences in resources and opportunities, depending on which country and what family they are born in. How much consumption level and growth would a new born child be willing to give up in order to avoid these birth lotteries? This paper shows that a child may well be willing to give up all growth to avoid birthplace risk, and a large fraction of growth, if not all, to avoid family risk. The critical elements for the 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 risk involved at birth, which is staggering. Our calculations suggest a research agenda that treats not only growth but also inequality as a priority.

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

JEL Classification: E1, E2, D3

*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, IMF, the 2006 NBER Summer Institute, and the 2006 Midwest Macro Meetings for helpful comments and suggestions. A longer version of this paper previously circulated under the title 'Lucas vs Lucas: On Inequality and Growth'. This research was conducted while the first author visited Carnegie Mellon University. Any errors or omissions are our own.

[†]Department of Economics, Rice University, 6100 Main Street, Houston TX, 77005, e-mail: `jcordoba@rice.edu`.

[‡]International Monetary Fund, Washington DC, 20431, e-mail: `gverdier@imf.org`.

1 Introduction

The two main lotteries individuals face during their lifetime are their place of birth and their parents. Children that are ex-ante identical in all respects start their lives with vast differences in resources and opportunities, depending on which country and family they are born into. How much consumption, in level and growth, would a newborn child be willing to give up in order to avoid these birth lotteries? In this paper we show that a newborn may well be willing to give up all growth to avoid birthplace risk, and a large fraction of growth, if not all, to avoid family risk.

We compute potential welfare gains and costs of economic growth and inequality using a tractable heterogenous-agent model. In the spirit of Lucas (1987), we postulate a process for individual consumption that allows simple aggregation and closed-form solutions, and define social welfare as the average welfare of individuals in the society or, alternatively, as the ex-ante welfare of a newborn. In our model, inequality and risk at birth have symmetric welfare effects, and as a result welfare measures of inequality can be interpreted as welfare measures of risk, and vice versa. In the model, inequality and risk are completely characterized by the standard deviation of the cross-sectional log of consumptions. Welfare costs and gains are defined as permanent consumption compensation rates required to leave the consumer indifferent between an initial situation and an alternative situation with more or less risk and/or growth. We use the model to study social welfare at the country and worldwide levels.

Our calculations confirm that changes in the growth rate of the economy have important welfare consequences, but this depends on the intertemporal elasticity of substitution (IES). For plausible values of this elasticity, we estimate that the welfare gains from growth are significantly lower than previous estimates. To assess the welfare burden of inequality, we consider different welfare measures. An upper bound for the potential welfare gains of redistribution is obtained by eliminating inequality completely. We compute this and other welfare measures, such as the net consequences of simultaneously eliminating growth and inequality, or the marginal rate of substitution between inequality and growth[1]. All calculations lead to a similar conclusion. Welfare measures associated with inequality, or alternatively, with risk at birth, are of the same order of magnitude as those

[1]Further exercises are provided in the working paper version of this paper.

associated with growth. Measures of cross-country inequality, or birthplace risk, systematically dominate measures of growth, and for some plausible parametrizations, within-country inequality, or family risk, also dominates growth. Overall, our calculations reveal that the social burden of inequality is significant. Two key elements for these results are time discounting and risk aversion. Both factors attenuate the role of growth for welfare while risk aversion amplifies the benefits of more equal outcomes. A third key factor is the amount of risk involved at birth, which is staggering.

To illustrate our results, we present two telling examples of different historical social choices between inequality and growth. In a first exercise we use our utilitarian framework to compare the experience of the United States with that of Scandinavian countries. Although per capita income is 25 per cent higher in the U.S. than in Scandinavian countries, inequality is much lower in that region. We find that welfare differences between Americans and Scandinavians are much lower than those implied by per-capita income differences, and that it is plausible than Scandinavians enjoy higher social welfare than Americans. We also compare welfare levels in East and West Germany at the end of the World War II.

This paper only shows that the *maximum potential* benefit from redistributive policies is large, but is completely silent on the actual *realizable* potential gains[2]. These results, however, suggest that inequality matters greatly for social welfare, likely as much as growth, which challenges a conventional wisdom in the profession. For example, Lucas (2004) asserts that "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." This paper makes clear that growth does not necessarily dwarf inequality and that further research is needed to correctly quantify the potential welfare gains, if any, of altering growth and/or inequality. In fact, our calculations suggest a research agenda that treats not only growth, but also inequality as a priority.

This paper is organized as follows: Section 2 describes the theoretical framework and defines welfare measures, Section 3 calibrates the model, Section 4 reports the welfare results, Section 5 provides additional illustrations, and Section 6 concludes.

[2]See the working paper version of this paper for an attempt at assessing the actual realizable gains.

2 The Model

The model has two components: a path for the distribution of consumption and a social welfare function.

2.1 The 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:

$$\ln c_t = \rho \ln c_{t-1} + (1-\rho)(a+bt) + \sigma_\eta \eta_t + \sigma_\epsilon \epsilon_t, \tag{1}$$

where $\rho \in [0,1)$ determines 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:

$$\ln c_t \sim N\left(a+bt, \sigma_x^2 + \sigma_y^2\right), \tag{2}$$

where

$$\sigma_x^2 \equiv \frac{\sigma_\eta^2}{1-\rho^2}; \ \sigma_y^2 \equiv \frac{\sigma_\epsilon^2}{1-\rho^2}. \tag{3}$$

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

$$\ln c_0 \sim N\left(a, \sigma_x^2 + \sigma_y^2\right). \tag{4}$$

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 — inequality associated with individual factors — and σ_x^2 measures cross-country inequality — inequality associated with country-specific factors. This framework can be used to analyze a single country if $\sigma_x^2 = 0$, or a group of countries populated by identical individuals if $\sigma_y^2 = 0$[3].

Parameters a and b are chosen so that aggregate worldwide consumption at time t, Ec_t, equals $(1+\lambda)(1+\mu)^t$. This requires defining $a \equiv \ln(1+\lambda) - \frac{1}{2}\left(\sigma_x^2 + \sigma_y^2\right)$ and $b \equiv \ln(1+\mu)$. In this specification, μ is the growth rate of worldwide consumption, and $1+\lambda$ is the "intersect". In other words, $1+\lambda$ 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 of alternative consumption paths, and is set to 0 in the baseline case.

2.2 Individual and Social Welfare

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

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

where E_0 is the mathematical expectation conditional on time-0 information, and u is a momentary utility function, identical for all individuals. Our extension of Lucas' framework does require the choice of a social welfare function and we consider that the natural benchmark is to weigh everyone equally.[4] We define social welfare as the average welfare of the society:

$$W \equiv \int U(c) dF(c), \tag{6}$$

[3] Our assumption that σ_x is constant abstracts from issues of convergence or catch-up. As shown below, the relevant concept for world welfare is that of σ-convergence (which characterizes the variance of the cross-sectional distribution of log incomes) rather than β-convergence. The evidence on σ-convergence is mixed. Barro and Sala-i-Martin (2004) have documented σ-convergence across U.S. states, Japanese prefectures and European regions, but Durlauf and Quah (1999) find divergence for a larger set of countries. For our set of 108 countries and for the period 1960-2000 we do not observe σ-convergence. See Section 3 below.

[4] Lucas (2004, last paragraph) suggests that everyone should be 'equally valued' in the social welfare function.

where $F(c)$ is the fraction of the population with consumption below or equal to c at time 0. F is defined by (4). There are least two alternative interpretations of W. First, W is a standard utilitarian social welfare function. The second interpretation of W is as 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 a particular country — if $\sigma_x = 0$.

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

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

Substituting for the definitions of a and b produces[5]:

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

provided that $1 > \beta\,(1+\mu)^{1-\gamma}$ which is assumed to be the case. 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 easy to check that inequality reduces social welfare ($\frac{\partial W}{\partial \sigma^2} < 0$) and growth increases social welfare ($\frac{\partial W}{\partial \mu} > 0$). 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 inequality aversion equals the coefficient of risk aversion, and the inverse of the intertemporal elasticity of substitution (Atkinson (1970)). In other words, the planner penalizes consumption dispersion across individuals in the same way that individuals penalize consumption dispersion across states and time.

In order to perform welfare comparisons, one needs a baseline welfare level, W_0, determined by

[5] Appendix 1 provides all the results for the log-utility case $\gamma = 1$.

a baseline set of parameters $[\mu_0, \sigma^2_{x0}, \sigma^2_{y0}]$. Let $\sigma^2_0 \equiv \sigma^2_{x0} + \sigma^2_{y0}$. The baseline welfare level is given by:

$$W_0 = W(0, \mu_0, \sigma^2_0) \tag{8}$$

The following subsections use the above framework to define a number of welfare measures. A first set of measures defines welfare costs and gains in terms of λ, i.e., in terms of proportional increases in consumption required to leave the social planner indifferent between a baseline and an alternative situation; a second welfare measure is the marginal rate of substitution between inequality and growth; a third measure is obtained by quantifying the welfare costs of eliminating all growth and all inequality simultaneously.

2.3 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:

$$\begin{aligned}
W_0 &= W(\lambda_\mu, 0, \sigma^2_0), \\
W_0 &= W(\lambda_x, \mu_0, \sigma^2_0 - \sigma^2_{x0}), \\
W_0 &= W(\lambda_y, \mu_0, \sigma^2_0 - \sigma^2_{y0}). \\
W_0 &= W(\lambda_0, \mu_0, 0).
\end{aligned}$$

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. Thus, for example, λ_μ is the proportional change in consumption, uniform across all periods, countries, and individuals, required to leave the world planner indifferent between the baseline consumption path and a path with no growth but higher initial consumption level 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* of 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 welfare cost of total inequality.[6] Using (7) and (8), λ_i for $i \in \{\mu, x, y, 0\}$ can be solved as:

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

and

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

These formulas reveal three key 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 but not on μ or σ_x^2. This is a consequence of assuming an isoelastic utility function and a log-normal distribution of consumption. An 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 $\sigma_x > 0$) or a country (so that $\sigma_x = 0$).

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 (e.g. Kocherlakota 1996). Third, the absolute value of 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. The last two results imply that the welfare gains of 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, $\lambda_{1\%}$, defined as:

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

[6] These welfare measures 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. Using measures defined as compensation rates on the same consumption baseline path only makes our conclusion stronger. See the working paper version of this paper for more details.

Using (7) and (8), $\lambda_{1\%}$ can be solved as:

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

2.4 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 $\overline{\lambda}_x$, $\overline{\lambda}_y$, and $\overline{\lambda}_0$ defined as:

$$\begin{aligned} W_0 &= W(\overline{\lambda}_x, 0, \sigma_0^2 - \sigma_{x0}^2), \\ W_0 &= W(\overline{\lambda}_y, 0, \sigma_0^2 - \sigma_{y0}^2), \\ W_0 &= W(\overline{\lambda}_0, 0, 0). \end{aligned}$$

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

$$\overline{\lambda}_i = e^{-\gamma\sigma_i^2/2} \left[\frac{1-\beta}{1-\beta\left(1+\mu_0\right)^{(1-\gamma)}} \right]^{\frac{1}{1-\gamma}} - 1 \text{ for } i = \{x, y, 0\} \tag{12}$$

2.5 The Shadow Price of Inequality

A measure of the welfare costs of inequality is given by the marginal willingness of the planner to substitute growth for inequality or consumption level for inequality. The social marginal rate of substitution between inequality and growth — the shadow price of inequality — can be defined as the partial derivative $\frac{\partial \mu}{\partial \sigma_i}$ obtained from (7) evaluated at the baseline parameters $(\sigma_{i0}, \mu_0, \lambda_0)$:

$$MRS^i_{growth} = \frac{\partial \mu}{\partial \sigma_i} = \sigma_{i0}\gamma \frac{\left(1+\mu_0\right)^{\gamma} - \beta\left(1+\mu_0\right)}{\beta} \text{ for } i = \{x, y, 0\}.$$

Note that MRS^i_{growth} increases linearly with the degree of inequality and exponentially with γ.

The social marginal rate of substitution between inequality and consumption level can be defined by $\frac{\partial\lambda}{\partial\sigma_i}$ and it is given by:

$$MRS^i_{level} = \frac{\partial\lambda}{\partial\sigma_i} = \sigma_{i0}\gamma \text{ for } i = \{x, y, 0\}.$$

Alternatively, these marginal rates of substitutions can be expressed in terms of Gini coefficients (G) instead of standard deviations. For log-normal distributions:

$$G = 2\Phi\left(\frac{\sigma}{\sqrt{2}}\right) - 1 \tag{13}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution. Therefore, $\frac{\partial G}{\partial\sigma} = \frac{1}{\sqrt{\pi}}e^{-\left(\frac{\sigma}{2}\right)^2}$ and alternative marginal rates of substitution can be defined as:

$$\overline{MRS}^i_{growth} = \frac{\partial\mu}{\partial G_i} = \frac{\partial\mu}{\partial\sigma_i}\frac{\partial\sigma_i}{\partial G_i} = MRS^i_{growth}\sqrt{\pi}e^{\left(\frac{\sigma}{2}\right)^2},$$

$$\overline{MRS}^i_{level} = \frac{\partial\lambda}{\partial G_i} = \frac{\partial\lambda}{\partial\sigma_i}\frac{\partial\sigma_i}{\partial G_i} = MRS^i_2\sqrt{\pi}e^{\left(\frac{\sigma}{2}\right)^2} \text{ for } i = \{x, y, 0\}.$$

Note that if inequality reduction has little social value then the social indifference curve in the (σ, μ) space should be close to flat and these marginal rates of substitutions should be close to zero.[7]

3 Calibration

In order to compute welfare measures we need to calibrate the parameters μ_0, σ^2_{x0}, σ^2_{y0}, γ and β. We set μ_0 equal to 2.1% which is the annual growth rate of the weighted average of cross-country consumption between 1960 and 2000, according to the Penn World Tables.[8] The data includes 108 countries, and weights are given by their relative population size. To calibrate cross-country inequality σ_{x0}, we use the weighted average of the standard deviations of the log of per-capita consumption across countries from 1960 to 2000. This dispersion has remained roughly constant

[7] Lucas' (2004) implicitly postulates that $MRS^i_1 = 0$.

[8] Lucas (2004) suggests that the relevant growth rate is the one observed since the industrial revolution. Using Clark's (2005) series of real wages produces an annual growth rate of 1.4%. For the U.S the corresponding growth rate is 1.9% during the last 130 years. These alternative calibrations would strengthen our results.

around 1 during that period and therefore we choose $\sigma_{x0} = 1$.

International evidence about the dispersion of consumption within countries is scarce. Some authors report measures of income dispersion within countries (e.g. 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 about 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 since the early 1970s at around 0.48. The dispersion of per-capita income is approximately twice as large. Based on this evidence, we choose $\sigma_{y0} = 0.5$.

For the parameter γ, we consider a range of values between 0.5 and 20. As reported below, the exact value of γ is critical for the welfare results. Unfortunately, there is little consensus in the literature regarding its value. Authors such as Hall (1988) and Cambell and Mankiw (1989) find the IES to be close to zero, or γ large. This is also the view of a number of authors who regard the equity premium puzzle as indicative of a large γ (see Kocherlakota 1996, page 52). In the other hand, the real business cycles literature typically consider values of γ closer to 1 or 2. For example, Kydland and Prescott (1982) use $\gamma = 1.5$. Beaudry and van Wincoop (1996) suggest that the IES is close to 1. Given the lack of consensus on γ, we study the same range of values considered by Lucas (1987), between 1 and 20, plus an additional value, $\gamma = 0.5$, for illustration purposes.

For the discount factor, Lucas (1987) uses $\beta = 0.95$. This is consistent, for example, with the value employed by King, Plosser and Rebelo (1988). Kydland and Prescott (1982) use instead a value of β closer to 0.96. In order to be conservative on the welfare calculations below, β is set to 0.96. The main findings of the paper are strengthen if $\beta = 0.95$ is used instead. All parameter values are summarized in Table 1.

4 Results

Table 2 reports the welfare measures defined in Section 2.3 for different values of γ (risk aversion). Recall that λ_μ is the welfare gain from economic growth, $\lambda_{1\%}$ the welfare gain from one percentage point of 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, however, significantly decrease as γ increases. For example, the welfare gain of an additional point of economic growth is 26% for $\gamma = 1$ (Lucas reports 20% for $\gamma = 1$ and $\beta = 0.95$), but the table shows that if γ takes an intermediate value of 5, the gains are 8.67% instead. The total gains from economic growth range from 10% for $\gamma = 20$ to 39% for $\gamma = 1$.

The novel and surprising result reported in Table 2 is the large welfare cost associated with consumption inequality, both within and across countries. The costs of within-country inequality range from around 12% for $\gamma = 1$ to around 92% for $\gamma = 20$, and the costs of cross-country inequality are between 40% and almost 100%. We find that the gains of economic growth are smaller than the cost of total inequality if $\gamma \geq 1.38$, smaller than the cost of cross-country inequality if $\gamma \geq 1.58$, and smaller than the cost of within-country inequality if $\gamma \geq 3.65$.

The figures reported in Table 2 reveal the importance of inequality for aggregate welfare. In recent years, many authors have argued that Lucas' original calculation understates the true cost of business cycles because of his assumptions on preferences, homogeneity of consumers, or his choice of matching the risk implied in aggregate rather than individual data (see Barlevy 2005 for a survey of this literature). When some of these assumptions are changed, 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 12% for alternative preferences (e.g. Tallarini 2000), or be larger for poor countries (Pallage and Robe 2003). But even these estimates of the costs of business cycles seem small relative to the costs of inequality.

Consider next the welfare measures defined in Section 2.4 and reported in Table 3 under the labels $\overline{\lambda}_i$, $i = \{x, y, 0\}$. These measures assess the welfare consequences of eliminating all growth and inequality simultaneously. If inequality only has second order effects on welfare, as suggested by Lucas (2004), then this experiment should result in a significant welfare loss since growth has first order effects. However, the results in Table 3 show instead that large welfare gains are possible. For example, $\overline{\lambda}_y = -0.30$ for $\gamma = 5$ means a 30% welfare gain of eliminating inequality and growth. More generally, eliminating within-country inequality and growth is welfare enhancing if $\gamma \geq 2.8$.

But in addition, inequality has a large first-order effect on welfare for any plausible value γ. For example, for $\gamma = 1$, eliminating all growth would produce a welfare loss of 64% (λ_μ in Table 2), but simultaneously eliminating within-country inequality reduces this loss to 45% (λ_y in Table 3).

The social marginal rate of substitution between inequality and growth is possibly a more relevant welfare measure. It was defined in Section 2.5 and is reported in Table 3 under the labels MRS^i_{growth} and $\overline{MRS}^i_{growth}$ for $i = \{x, y, 0\}$. These slopes measure the willingness of the planner to trade inequality for growth. For a better understanding of these figures, consider a planner trying to reduce the Gini coefficient of within-country inequality by 10 percentage points ($\triangle G_y = -0.10$). How many points of economic growth would the planner be willing to trade? This number is labelled $\overline{MRS}^y_{growth}$ in Table 3. Our calculations imply that the planner is willing to trade 1.2 percentage points of economic growth ($\triangle\mu = 0.012$) if $\gamma = 2$ or 6.3 points if $\gamma = 5$. That is, the planner is willing to make a considerable sacrifice in growth to reduce the Gini coefficient. Contrary to Lucas' (2004) suggestion, these figures clearly show that growth does not dwarf inequality.

Finally, Table 3 also reports the social marginal rate of substitution between inequality and consumption levels (MRS^i_{level} and $\overline{MRS}^i_{level}$) defined in Section 2.5. This is likely a more relevant tradeoff because growth rates tend to be similar across countries despite different degrees of inequality while consumption levels vary widely. According to this measure, the planner is also willing to pay a large price to reduce inequality. For example, for a permanent reduction of 10 percentage points in the Gini coefficient of within-country inequality, the planner would give up almost 19% ($= \overline{MRS}^y_{level} \times 0.10$) of consumption in all periods if $\gamma = 2$ or 47% ($= \overline{MRS}^y_{level} \times 0.10$) if $\gamma = 5$.[9]

Why is inequality so costly? Three factors underlie the quantitative results. First, most of the gains of economic growth occur in the future while the costs of inequality are borne every period. Second, inequality is large both within and across countries. Third, commonly used values of the IES imply substantial social aversion to any source of consumption dispersion, in particular social aversion to inequality. One way to explain the magnitude of the cost is to interpret the welfare function $W(\lambda, \mu, \sigma^2)$ as the lifetime utility of a newborn child. If given a choice before birth, what

[9]It is worth noting that our measures of the marginal rates of substitution would reflect not only *willingness* to trade inequality for consumption level of growth but also the technological tradeoff *if* the observed social choices are optimal. In that case, the choices lie both on the social indifference curve and on the production possibility frontier.

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 substantially reduces expected utility.

How can we translate these calculations into real-world implications for social choices between inequality and growth? The next section provides two illustrations of social welfare across different regimes.

5 Two Illustrative Examples

To illustrate some implications of the utilitarian framework described in Section 2, this section reports the results from two exercises. The first exercise compares social welfare in the US and Scandinavian countries. The second exercise compares social welfare in East Germany and West Germany by the end of World War II.

5.1 U.S. versus Scandinavia

The most popular measure of social welfare is income per-capita. According to this measure, social welfare is around 25% larger in the U.S. than in Scandinavian countries (Denmark, Norway, and Sweden). However, inequality is also significantly larger in the U.S. According to our social welfare function (7), this differential in the levels of inequality should increase social welfare in Scandinavian countries relative to the U.S. By how much? Are the inequality differences sufficiently large as to reverse the ranking of social welfare implied by the income per-capita measure?

Let W_{US} and W_{SC} be the level of social welfare in the U.S. (US) and in Scandinavian countries (SC) respectively. According to (7), $W_i = W(\lambda_i, \mu_i, \sigma_i^2)$ for $i = \{US, SC\}$. Normalizing $\lambda_{US} = 0$ then $1 + \lambda_{SC}$ is the average consumption level of Scandinavia relative to the U.S. Define λ_{US}^* by the equation:

$$W(\lambda_{US}^*, \mu_{US}, \sigma_{US}^2) = W_{SC}.$$

According to this definition, $\lambda_{US}^* \times 100$ is the percentage increase (decrease if $\lambda_{US}^* < 0$) in U.S. average consumption that would render social welfare equal to that of Scandinavia. Using (7), λ_{US}^*

can be solved as:

$$1+\lambda_{US}^{*}=\left(1+\lambda_{S}\right)e^{\gamma\left(\sigma_{US}^{2}-\sigma_{SC}^{2}\right)/2}\left[\frac{\left(1-\beta\left(1+\mu_{US}\right)^{1-\gamma}\right)}{\left(1-\beta\left(1+\mu_{SC}\right)^{1-\gamma}\right)}\right]^{\frac{1}{1-\gamma}}.$$

To compute λ_{US}^{*} one needs the respective growth rates, μ_i, the variances of log-consumption, σ_i^2, and the relative initial mean consumption levels $1+\lambda_{SC}$. According to the Penn World Tables 6.1, average per-capita income in the U.S. was around 25% higher than in Scandinavian countries both in 1990 and 2000. The relative stability of this difference means that the U.S. and Scandinavia have very similar growth rates. This simplifies the formula above to $1+\lambda_{US}^{*}\simeq\left(1+\lambda_{S}\right)e^{\gamma\left(\sigma_{US}^{2}-\sigma_{SC}^{2}\right)/2}$. These figures also suggest that $1+\lambda_{SC}=\frac{1}{1.25}=0.8$.

Regarding differences in inequality, we assume $\sigma_{US}=0.48$ as documented in the previous section. As for σ_{SC}, Aaberge *et al.* (2000, Table 2) report Gini coefficients of disposable incomes for the U.S. and Scandinavian countries for different years up to 1990. They find a Gini coefficient of 0.346 for the U.S. and an average Gini of 0.2173 for Scandinavian countries in 1990. Assuming that disposable income is log-normal distributed, Gini coefficients can be transformed into standard deviations using equation (13), which can be rewritten as:

$$\sigma=\sqrt{2}\Phi^{-1}\left(\frac{1+Gini}{2}\right). \tag{14}$$

Solving this equation gives a standard deviation of 0.634 for the U.S. and of 0.39 for Scandinavian countries, or a ratio of $\frac{\sigma_{US}}{\sigma_{SC}}=1.63$. Assuming that this ratio also applies to the standard deviation of consumptions, it follows that $\sigma_{SC}=\sigma_{US}/1.63\simeq0.30$.[10] Under these assumptions, $1+\lambda_{US}^{*}\simeq 0.8\cdot e^{\gamma\left(0.48^{2}-0.3^{2}\right)/2}$.

Table 4 reports λ_{US}^{*} for the different values of γ considered in the previous section. The most surprising result in the table is that for large but plausible values of γ social welfare in Scandinavian countries is larger than in the U.S. For example, if $\gamma=5$ then Scandinavian welfare is 13.64% ($\lambda_{US}^{*}=0.1364$) higher than it is in the U.S. More precisely, W_{SC} is larger than W_{US} if γ is larger than 3.18.

[10] Note that we do not assume that the standard deviation of consumption and disposable income are the same, but instead that their ratio is the same in the U.S. and in Scandinavia.

A second observation is that for any plausible value of γ welfare differences are significantly smaller than the ones implied by per-capita incomes. For example, while Scandinavian income is 80% of US income, adjusting for inequality increases this fraction to 85.88% ($= 1 - \lambda^*_{US}$) if $\gamma = 1$, or to 92.06% if $\gamma = 2$.

5.2 Social Welfare in East and West Germany at the end of the war

The meltdown of most communist regimes around the globe was in part due to their inability to produce sufficient economic growth. For example, after almost 45 years of communist ruling per-capita income in East Germany (EG) was less than half the income in West Germany (WG) by 1990. By this measure, a communist regime for East Germany was clearly a mistake. But it is also the case that the premise of communist regimes is to reduce inequality. In fact, by 1990 inequality in East Germany, as measured by the log variance of income, was around half of the one in West Germany. Given that equality has important social value in our utilitarian framework, these figures raise the intriguing (and certainly controversial) question of whether a communist regime for East Germany could have been actually a sensible social choice by the end of World War II.

To investigate this question, denote $W_i = W(\lambda_i, \mu_i, \sigma_i^2)$ for $i = \{EG, WG\}$. Define λ^*_W by the equation:

$$W(\lambda^*_{WG}, \mu_{WG}, \sigma^2_{WG}) = W_{EG}.$$

According to this definition, $\lambda^*_{WG} \times 100$ is the percentage increase (decrease if $\lambda^*_{WG} < 0$) of WG consumption that would render social welfare equal to that of EG. Using (7), λ^*_{WG} can be solved as:

$$1 + \lambda^*_{WG} = \left[\frac{1 - \beta \left(1 + \mu_{WG}\right)^{1-\gamma}}{1 - \beta \left(1 + \mu_{EG}\right)^{1-\gamma}} \right]^{\frac{1}{1-\gamma}} e^{\gamma\left(\sigma^2_{WG} - \sigma^2_{EG}\right)/2}.$$

For the quantitative exercise, we assume $\mu_{WG} = 0.021$ and $\sigma_{WG} = 0.48$, the same values used for the U.S. in the previous illustration. Moreover, 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. Assuming that the ratio of standard deviation of consumptions is the same as that of incomes then $\sigma_{EG} = \sigma_{WG} \cdot \sqrt{\frac{0.115}{0.23}} \simeq 0.34$.

The only remaining parameter is μ_{EG}. According to Biewen (2000, Tables 1 and 2) mean income in West Germany was 1686.6 and 782 in west Germany in 1990, a ratio of 2.15. This ratio is partly confirmed by Burda and Hunt (2001, Table 3) who report a ratio of 2.32. This ratio exaggerates the existing gap right before the fall of the Berlin Wall in 1989 because GDP in East Germany dropped dramatically after reunification. According to Burda and Hunt (Table 3), GDP in East Germany fell by 15.6% and 22.7% in 1990 and 1991 respectively while GDP in West Germany grew by 5.7% and 4.6% during the same period. A better estimate would be $\frac{y_{west1989}}{y_{east1989}} = 2.32\frac{(1-0.156)*(1-0.227)}{1.057*1.046} = 1.37$. Using data from the The Penn World Table Mark 5.6, this ratio was 1.95 in 1970 but only 1.28 in 1988. Overall, these numbers suggest a range for the ratio of per-capita consumptions in 1989 of 1.4 to 1.9. Assuming a ratio of per-capita consumptions of 1 in 1945 and 1.65 in 1989, the ratio of annual growth rates is given by $\frac{1+\mu_{WG}}{1+\mu_{EG}} = 1.65^{1/44}$ or $\mu_{EG} = 0.00945$.

Table 4 reports λ^*_{WG} for the different values of γ considered in the previous Section. Surprisingly, social welfare in East Germany could have been larger than in West Germany if γ is large. For example, if $\gamma = 5$ then EG welfare was 19.51% ($\lambda^*_{WG} = 0.1951$) higher than in WG. [11] More precisely, W_{EG} is larger than W_{WG} if γ is larger than 2.82.

6 Final Comments

Children that are identical in all respects will start their life with vast differences in resources and opportunities, depending on which country and what family they are born in. How much consumption level and growth would a new born child be willing to give up in order to avoid these birth lotteries? Lucas' (2004) suggests this child would give up very little. Even if the child is born poor, economic growth would help him overcome poverty. In this paper we show that, on the contrary, this child may well be willing to give up all growth to avoid birthplace risk, and a large fraction of growth, if not all, 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 staggering.

[11] An implicit assumption in the exercise is that inequality can be reduced quickly. This is probably a good description of communist regimes at the time that quickly implemented populist policies such as land reform, and abolished private property.

The contribution of this article 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 help rationalize why societies may not always find it optimal to adopt growth-enhancing institutions, particularly when inequality is large and those institutions may foster inequality.

Societies commonly face major choices between equality on one hand and efficiency and growth on the other hand. The degree of progressivity of the tax system, for example, reveals the 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 in terms of equality. Similarly, the extent of law enforcement, illustrated for example in the 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 institutional and policy choices made by different societies requires the proper assessment of the welfare implications of these choices, and in particular, a careful consideration of the welfare gains of 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 policy evaluations.

The results we present may seem more relevant for within-country discussions regarding inequality and growth. This is mostly because cross-country redistribution policies seem difficult to achieve. In addition, the many growth-enhancing policies put in place by development agencies seem undistinguishable from policies that would seek to reduce inequality across countries. However, one could in fact imagine policies that could reduce consumption dispersion across countries without directly affecting long-run growth rates. For example, some of the United Nations Millennium Development Goals do not seek to directly affect growth but to reduce extreme poverty and hunger. These policies are of particularly high social value according to our calculations.

We believe that our work provides an important step in objectively evaluating the costs of inequality in a well-understood welfare framework, and we think our calculations support the case for a research agenda that treats not only growth, but also inequality as a priority.

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),
$\sigma_y = 0.5$ (standard deviation of individual consumption within countries),
$\gamma = \{0.5, 1, 2, 5, 10, 20\}$ (coefficient of risk aversion),
$\beta = 0.96$ (discount factor).

Table 2: Standard Welfare Measures

γ	0.5	1	2	5	10	20
λ_μ	0.7811	0.6467	0.4936	0.3066	0.1983	0.1215
$\lambda_{1\%}$	0.3430	0.2665	0.1844	0.0949	0.0512	0.0252
λ_x	-0.2212	0.3965	-0.6321	-0.9179	-0.9933	-1.000
λ_y	-0.0606	-0.1186	-0.2212	-0.4647	-0.7135	-0.9179
λ_0	-0.2684	-0.4681	-0.7135	-0.9561	-0.9981	-1.000

Note: λ_μ is the welfare gain of 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 consumption inequality; λ_0 is the welfare cost of total inequality.

Table 3: Welfare Measures - Growth and Inequality Equivalents

γ	0.5	1	2	5	10	20
$\overline{\lambda}_x$	0.3871	-0.0012	-0.4505	-0.8927	-0.9919	-0.9999
$\overline{\lambda}_y$	0.6731	0.4532	0.1632	-0.3006	-0.6567	-0.9079
$\overline{\lambda}_0$	0.3030	-0.1186	-0.5721	-0.9426	-0.9977	-0.9999
MRS^x_{growth}	0.0158	0.0425	0.1298	0.6737	2.6129	11.1499
MRS^y_{growth}	0.0079	0.0213	0.0649	0.3368	1.3064	5.5750
MRS^0_{growth}	0.0176	0.0476	0.1451	0.7532	2.9213	12.466
$\overline{MRS}^x_{growth}$	0.0359	0.0968	0.2953	1.5332	5.9466	25.376
$\overline{MRS}^y_{growth}$	0.0149	0.0401	0.1224	0.6355	2.4650	10.519
$\overline{MRS}^0_{growth}$	0.0333	0.0897	0.2737	1.4211	5.5118	23.520
MRS^x_{level}	0.5000	1.0000	2.0000	5.0000	10.000	20.000
MRS^y_{level}	0.2500	0.5000	1.0000	2.5000	5.0000	10.000
MRS^0_{level}	0.5590	1.1180	2.2361	5.5902	11.180	22.360
$\overline{MRS}^x_{level}$	1.1379	2.2759	4.5518	11.379	22.759	45.518
$\overline{MRS}^y_{level}$	0.4717	0.9434	1.8868	4.7169	9.4338	18.868
$\overline{MRS}^0_{level}$	1.0547	2.1095	4.2189	10.547	21.095	42.189

Note: $\overline{\lambda}_x$, $\overline{\lambda}_y$ and $\overline{\lambda}_0$ is the welfare cost (gain if negative) of eliminating growth and cross-country inequality, growth and within-country inequality, and growth and total inequality respectively; $\text{MRS}^i_{growth} = \frac{\partial \sigma}{\partial \mu}$ and $\text{MRS}^i_{level} = \frac{\partial \sigma}{\partial \lambda}$ for i={x,y,0}; $\overline{MRS}^i_{growth} = \frac{\partial \mu}{\partial G_i}$ and $MRS^i_{level} = \frac{\partial \lambda}{\partial G_i}$ for i={x,y,0}.

Table 4: Social Welfare Across Countries and Regimes

γ	0.5	1	2	5	10	20
λ^*_{US}	-0.1714	-0.1418	-0.0794	0.1364	0.6142	2.2572
λ^*_{WG}	-0.2654	-0.1939	-0.0803	0.1951	0.6707	2.0566

Note: λ^*_{US} is the welfare lag (lead if negative) of the US relative to Scandinavian Countries; λ^*_{WG} is the welfare lag (lead if negative) of West Germany relative to East Germany in 1945.

References

[1] Aaberge et al. (2002), "Income Inequality and Income Mobility in the Scandinavian Countries Compared to the United States." Review of Income and Wealth, 48(4).

[2] Atkinson, Anthony B. (1970) "On the Measurement of Inequality." *Journal of Economic Theory* 2: 244-263.

[3] Barro, R. J. and Sala-i-Martin, X. (2004) *Economic Growth*, 2nd Ed., MIT Press, Cambridge, MA, U.S.A.

[4] Barlevy, G. (2005) "The Costs of Business Cycles and the Benefits of Stabilization." *Economic Perspectives*, 1Q: 1-18.

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

[6] Beaudry, P. and van Wincoop, Eric. (1996) "The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data," *Economica*, 63: 495-512.

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

[8] 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.

[9] Córdoba, J.C. and Verdier, G. (2007) "Lucas vs. Lucas: On Inequality and Growth." IMF Working Paper 07/17, 39 pages.

[10] Durlauf, S. N. and Quah, D. T. (1999) "The New Empirics of Economic Growth" in J. B. Taylor & M. Woodford (ed.) Handbook of Macroeconomics, chapter 4, pages 235-308.

[11] HALL, R.E. *(1988)* “Intertemporal substitution in consumption.” *Journal of Political Economy*, 96: 339-57.

[12] HESTON, A. SUMMERS, R. AND ATEN, B. (2002). Penn World Table Version 6.1, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

[13] KING, R., PLOSSER, C. AND REBELO, S. (1988) ”Production, Growth and Business Cycles: I. The Basic Neoclassical Model.” *Journal of Monetary Economics* 21: 195-232.

[14] KOCHERLAKOTA, NARAYANA R. (1996) “The Equity Premium; It’s Still a Puzzle.” *Journal of Economic Literature*, 34(1): 42-71.

[15] KRUEGER, DIRK; AND PERRI, FABRIZIO (2002) “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory.” NBER Working Paper 9202.

[16] KRUSELL, P., AND SMITH, A. (1999) “On the Welfare Effects of Eliminating Business Cycles”, *Review of Economic Dynamics*, 2(1): 245-272.

[17] KYDLAND, F. AND PRESCOTT, E. (1982) ”Time to build and aggregate fluctuations.” *Econometrica* 50: 1345-1371.

[18] LUCAS, ROBERT E. (2004) ”The Industrial Revolution: Past and Future.” *The Region* (May), Federal Reserve Bank of Minneapolist, pages 5-20.

[19] LUCAS, ROBERT E. (1987) *Model of Business Cycles*, Basil Blackwell, New York.

[20] PALLAGE, S. AND ROBE, MICHEL A. (2003) ”On the Welfare Cost of Economic Fluctuations in Developing Countries.” *International Economic Review*, 44(2): 677-698.

[21] SALA-I-MARTIN (2002) “The Disturbing “Rise” of World Income Inequality,” NBER Working paper 8904, April 2002.

[22] TALLARINI, T. (2000) “Risk-Sensitive Real Business Cycles,” *Journal of Monetary Economics,* 45: 507-532.

Appendix 1: The Log-utility case

For the log-utility case, social welfare is given by:

$$\begin{aligned} W &= EE_0\left[\sum_{t=0}^{\infty}\beta^t \ln c_t\right] = \sum_{t=0}^{\infty}\beta^t E\ln c_t \\ &= \sum_{t=0}^{\infty}\beta^t (a+bt) = \frac{a}{1-\beta} + b\sum_{t=0}^{\infty}\beta^t t \\ &= \frac{a}{1-\beta} + b\frac{\beta}{(1-\beta)^2} \end{aligned}$$

or, using the definitions of a and b:

$$W = \frac{\ln(1+\lambda) - \frac{1}{2}\left(\sigma_x^2+\sigma_y^2\right)}{1-\beta} + \frac{\beta\ln(1+\mu)}{(1-\beta)^2}.$$

Moreover, welfare measures are given by:

$$\begin{aligned} \lambda_\mu &= (1+\mu_0)^{\frac{\beta}{1-\beta}} - 1, \\ \lambda_{1\%} &= \left(\frac{1+\mu_0}{1+\mu_0-0.01}\right)^{\frac{\beta}{1-\beta}} - 1, \\ \mu_i &= (1+\mu_0)\left(e^{-\frac{1}{2}\frac{1-\beta}{\beta}\sigma_i^2} - 1\right), \\ \theta_i &= \frac{2\beta}{(1-\beta)\sigma_i^2}\ln(1+\mu_0), \\ \overline{\lambda}_i &= (1+\mu_0)^{\frac{\beta}{1-\beta}} e^{-\frac{1}{2}\sigma_i^2} - 1, \\ MRS_1^i &= \frac{\partial\mu}{\partial\sigma_i} = \sigma_{i0}(1+\mu_0)\frac{1-\beta}{\beta} \text{ for } i = \{x,y,0\}, \\ MRS_2^i &= \frac{\partial\lambda}{\partial\sigma_i} = \sigma_{i0} \text{ for } i = \{x,y,0\}. \end{aligned}$$