Risksharing Within the United States

WHAT HAVE FINANCIAL MARKETS AND FISCAL FEDERALISM ACCOMPLISHED? [1]

Stefano Athanasoulis
Yale University

Eric van Wincoop
Federal Reserve Bank of New York

April 1998

[1]We would like to thank Joshua Coval, Jordi Gali, Marco del Negro, Greg Hess, Holger Sieg, Bent Sorensen, Oved Yosha, and seminar participants at the "Regional Macroeconomics and Finance" conference at Tel-Aviv University and at the 1998 AEA meetings in Chicago for useful comments and suggestions. We particularly thank Bent Sorensen and Oved Yosha for many useful discussions and for making available to us data from Asdrubali, Sorensen and Yosha (1996). We would like to thank Laura Brookins and Sairah Burki for excellent research assistance. The views expressed in the paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Abstract

We document aggregate income growth uncertainty at the state level, and the extent to which this uncertainty is reduced by risksharing through financial markets and federal fiscal policy. A methodology is adopted that is closely connected to the empirical growth literature. It does not rely on assumptions about a model or stochastic process of income. This is important because estimated gains from international risksharing have been found to be very sensitive to the assumed model or income process. We only make assumptions about the information set used to predict growth, which is sufficient (and necessary) to compute income growth uncertainty. Estimates of growth uncertainty for a representative state are obtained at horizons from one to twenty-six years from panel regressions of state income growth, in deviation from national growth, on variables in the information set. This "residual risk" can be fully diversified at a zero risk premium. We show that potential welfare gains from risksharing depend on a weighted average of the variance of residual risk at different horizons. Three measures of income are considered: gross state product, pre-tax state income, and disposable state income. The difference between the first two is associated with financial markets (net asset income), while the difference between the last two is associated with federal fiscal policy (federal taxes and transfers). At a 26-year horizon we find that 71% of potential welfare gains from risksharing have already been achieved; 60% through financial markets and 11% through federal fiscal policy.

1 Introduction

Households all over the world are faced with substantial uncertainty about the long term growth of the country in which they reside. Athanasoulis, Shiller and van Wincoop (1998) find that over a period of three decades it is quite likely that per capita GDP of the best performing OECD country doubles relative to that of the worst performing country, purely by surprise. Among a larger set of forty nine countries this is true almost with probability one. Despite the large uncertainty about growth, the degree of risksharing among countries is very limited: consumption growth uncertainty is about as large as GDP growth uncertainty.[1] Although international financial markets are becoming more integrated[2], there still is a strong "home bias" in equity portfolios. The most recent numbers on international portfolio diversification, reported by Tesar and Werner (1997), show that in 1996 US and Japanese residents invested respectively 90% and 95% of their stock portfolio at home.

In this paper we will evaluate the extent of risksharing within the United States. This is of interest by itself as it sheds light on what financial markets and fiscal federalism have accomplished within the US. It also provides a good benchmark against which to place the international evidence in perspective. We would like to know whether the low degree of risksharing among countries is due to barriers that are specific to international borders. To this end it is useful to evaluate the extent of risksharing within a country. There are no capital controls among the states, there are no language barriers, there is a common regulatory framework, there are common accounting standards, and there is a single currency. If the extent of risksharing is limited even within a country, we may have to look for other explanations, unrelated to geographic borders.

Our interest is not in risksharing at the micro level of individuals or households, which has been the focus of studies such as Altug and Miller (1990), Cochrane (1991), Hayashi, Altonji and Kotlikoff (1996), Mace(1991), Miller and Sieg (1997), and Townsend (1994). Full risksharing at that level is generally rejected. For moral hazard reasons one might not expect financial markets to provide insurance against individual-specific risks. Financial markets should be able to provide insurance against aggregate risk factors, which are not under the control of households. Examples are uncertainty about aggregate income of a region (geographic risk), aggregate income of an age cohort (demographic

[1]See Athanasoulis and van Wincoop (1997).

[2]Since the early 1980s gross capital flows among industrialized countries have quadrupled as a fraction of GDP.

risk), or aggregate income of an industry (occupational risk). In this paper we focus on aggregate income growth uncertainty at the level of US states. Our goal is to measure this geographic risk and the extent to which it is reduced by financial markets and federal fiscal policy.

Risk is measured as the standard deviation (or variance) of income growth uncertainty at horizons from one to twenty-six years. We use measures of aggregate income in a state before and after risksharing from a nice data set constructed by Asdrubali, Sorensen and Yosha (1996).[3] Income before risksharing is measured as per capita gross state product (GSP). This is the output of the state and does not include any transfers among states that can be associated with risksharing. The first measure of income after risksharing is per capita pre-tax state income, which is close to the state equivalent of GNP. It includes net asset income. A second measure of income after risksharing is disposable state income, which subtracts from pre-tax state income net payments (taxes less transfers) from the state to the federal government. We use these income measures to determine the reduction in the standard deviation (or variance) of income growth uncertainty accomplished by financial markets and federal fiscal policy. Using standard CRRA preferences, we also compute the potential welfare gain from risksharing based on GSP and the fraction of that gain that has already been achieved through financial markets and federal fiscal policy. Although it is ultimately welfare that matters, we believe simple standard deviations of income growth uncertainty before and after risksharing are at least as informative because they do not rely on particular assumptions about preferences.[4]

We focus on income growth uncertainty that in theory can be diversified away through risksharing. This is uncertainty about income growth of a state in deviation from national growth. National income growth uncertainty cannot be diversified through risksharing within the United States. One approach to measure income growth uncertainty is to specify a particular model or statistical process for income. If one knew the correct model or income process, this would be the preferred approach. However, no one knows this. Assuming a particular process can lead to misleading results. van Wincoop (1998) shows that an important reason estimates of welfare gains from international risksharing vary so much in the literature is the sensitivity to the assumed model or

[3]We would like to thank Bent Sorensen and Oved Yosha for making available to us data from Asdrubali, Sorensen and Yosha (1996).

[4]One problem with utility functions that are commonly adopted in the welfare gains literature is that they are subject to the equity premium puzzle.

income process.[5]

We therefore adopt a more model-independent approach to compute income growth uncertainty, whereby we only make assumptions about the information set used to predict growth. The component of growth orthogonal to this information set is the unpredictable part. The information set does not tie down the process that income growth follows[6], but knowing the information set is sufficient (and necessary) to compute income growth uncertainty. We construct a large state-level information set, focusing on variables that have been found to be important at the national level in the empirical growth literature. Although one can of course argue about what variables should be in such an information set, we believe that our findings are quite robust. After including the three variables with most predictive power, adding an additional six variables does not affect the estimated risk very much. Moreover, separate from the estimated level of risk, our estimates of the percentage risk-reduction achieved by financial markets and federal fiscal policy are almost the same with an "empty" information set as with a full information set of nine variables.

The remainder of the paper is organized as follows. The next section discusses previous work on intranational risksharing and relates our approach to that literature. Section 3 discusses the methodology used to measure risk and the extent of risksharing. Section 4 applies the results to the United States. After describing the income measures and the information set, it reports output growth uncertainty at horizons from one to twenty six years, and the potential welfare gain from sharing this risk. The extent to which this risk has been diversified through financial markets and federal fiscal policy is discussed in section 5. The final section concludes.

[5]Estimated gains from international risksharing, measured as the welfare equivalent permanent percentage increase in expected income, range from 0.1% to 100%. van Wincoop (1998) shows that the estimated potential welfare gain among OECD countries at a 50-year horizon is 3% when assuming an AR process in growth rates, and drops to a very small 0.2% when only adding a global cointegration term. For other papers computing welfare gains from international risksharing see Athanasoulis and van Wincoop (1997), Backus, Kehoe and Kydland (1992), Cole and Obstfeld (1991), Kose (1995), Lewis (1996), Mendoza (1995), Obstfeld (1994a,b), Shiller and Athanasoulis (1995), Tesar (1995), and van Wincoop (1994,1996).

[6]To see this, consider an information set consisting of current per capita income and the most recent growth rate, both in deviation from the national average. To predict deviations from national growth, this information set is consistent with an AR(1) process in levels, an AR(2) process in levels, an AR(1) process in growth rates, and an AR(1) in growth rates with a global cointegration term added.

2 Previous Literature on Intranational Risksharing

Since the paper is focused on risksharing among regions within a country, we will limit ourselves to a discussion of the literature in that area, leaving aside findings on the extent of risksharing at the micro level among individuals or households. One problem that immediately arises is that the term "risksharing" itself has been used in different ways. The definition we will adopt in the main body of this paper is the sharing of uncertainty about future income, which we refer to as R. It has been used this way in most of the literature on welfare gains from international risksharing as well as in the political economy literature.[7] On the other hand, Cochrane (1991), Mace (1991) and Townsend (1994), from here on CMT, use the term risksharing in a broader sense, referring to complete risksharing as the Arrow Debreu allocation. This amounts to both R and intertemporal substitution of consumption through borrowing and lending. We refer to this as AD.[8] Asdrubali, Sorensen and Yosha (1996), Bayoumi and Klein (1997), Hess and Shin (1997,1998), and van Wincoop (1995) can be seen as papers on the extent of intranational risksharing in the sense of AD. Coval and Moskovitz (1997), Crucini (1998), Del Negro (1997), and Huberman (1997) are papers on intranational risksharing in the sense of R used in this paper.

These papers use a wide variety of risksharing metrics and do not always lead to the same conclusion. Asdrubali, Sorensen and Yosha (1996) estimate panel regressions of consumption growth of individual states on GSP growth and time dummies. According to the CMT test the coefficient on income growth should be zero under perfect risksharing. The authors use one minus the coefficient on income growth as a measure of risksharing in the sense of AD. According to this definition they find that 75% of income growth is "smoothed" at a one year horizon and 47% at a 10 year horizon. Bayoumi and Klein (1997) estimate the same equation for Canadian provinces, but allow the coefficient on income growth to differ across provinces. They also find small coefficients on the income growth variable.[9] An important contribution of Asdrubali, Sorensen and Yosha (1996) is the quantification of the channels that contribute to the reduction in

[7] See for example Persson and Tabellini (1996).

[8] It is a bit unfortunate that the word risksharing has also been used in the broader sense because intertemporal substitution of consumption can take place in an environment without risk. It amounts to the smoothing of consumption relative to the *expected* path of income.

[9] Although they generally assume this coefficient to vary across provinces, they estimate it to be 0.15 when it is assumed equal across provinces. This means that 85% of income growth is "smoothed" at the one-year horizon, adopting the terminology in Asdrubali, Sorensen and Yosha (1996).

the cross sectional covariance between consumption and income growth. They consider capital market smoothing, federal smoothing and credit market smoothing.

Hess and Shin (1997) also address the risksharing question in the context of the CMT test. They use micro data, regressing household food consumption growth on national consumption growth, the deviation between state and national consumption growth, and the deviation between consumption growth of members of the same industry and national consumption growth. The last two coefficients, which reflect the extent of risksharing across states and across industries, are both positive and significant. Perfect risksharing is rejected, but the coefficients are not very big, ranging from 0.2 to 0.4.

Hess and Shin (1998) and van Wincoop (1995) report another risksharing measure, the correlation of consumption growth across regions in comparison to the cross region output correlation. Although this measure is not based on the CMT test it can nonetheless still be interpreted as an AD risksharing measure. Backus, Kehoe and Kydland (1992) find that a model with complete assets markets, in the sense of AD, implies cross country consumption correlations that are significantly above income correlations. The opposite is the case in international data, a phenomenon they have called the "quantity anomaly". Hess and Shin (1998) find that the quantity anomaly also holds among US states, while van Wincoop (1995) finds that the correlation of consumption across Japanese prefectures is about the same as the correlation of output across the prefectures. From these results one would conclude that the extent of intranational risksharing is still quite limited.

The remaining four papers consider measures of risksharing in the R-sense. Crucini (1998) and Del Negro (1997) regress state consumption growth on national and state-specific innovations in the present discounted value of income. Crucini (1998) assumes that a fraction of agents consume state permanent income and the remainder consume national permanent income. The coefficient on national permanent income growth can then be interpreted as a measure of risksharing. In order to compute innovations in state and national permanent income the authors assume either i.i.d. state income growth or a VAR including state and national income growth.[10] Both authors conclude that the degree of risksharing within the US is substantial. Crucini (1998) finds that about 90% of consumption growth is associated with innovations in national rather than state-specific permanent income.[11]

[10]Crucini (1998) also considers a univariate autoregressive process.

[11]He finds the same for Canadian provinces.

Coval and Moscovitz (1997) and Huberman (1997) consider a more direct measure of risksharing among states by looking at portfolio diversification. They report data on respectively portfolios of US investment managers and households' holdings of regional Bell companies. Both papers report that even within the US there is a geographical home bias.

All of these risksharing measures are interesting on their own and provide different perspectives on the risksharing question. But there are four limitations of these measures that we would like to discuss. First, they do not provide estimates of risk agents are exposed to, and the extent to which this risk is reduced through financial markets and other channels. We believe that measuring risk and risk reduction is important in order to evaluate what financial markets and federal fiscal policy have actually accomplished. It may be that portfolios are not completely diversified even within the United States because a limited extent of diversification achieves most of the attainable welfare gains.

Second, and this is related to the first point, without computing the reduction in income uncertainty achieved through risksharing conclusions about the extent of risksharing can only be drawn in the context of specific models. Interpretations will generally be sensitive to the particular model adopted. In Crucini (1998) the interpretation that the coefficient on permanent national income growth is a measure of risksharing relies on a model where some agents engage in perfect risksharing while others consume their permanent income. Asdrubaly, Sorensen and Yosha (1996) and Bayoumi and Klein (1997) regress consumption growth on current income growth rather than permanent income growth. In order to interpret the results, Bayoumi and Klein (1997) adopt a model whereby a fraction of agents simply consume their current income, while others engage in perfect risksharing (in the sense of AD). Then one minus the estimated coefficient on current income growth is equal to the fraction of individuals that engages in risksharing. These interpretations are only valid in the context of the particular model adopted.[12] The models above are very stylized and conclusions may change in the context of richer models.[13] Similar problems arise when comparing cross region consumption and income correlations. van Wincoop (1995) interprets the low cross region consumption correla-

[12]If the same model as in Crucini (1998) is adopted, the interpretation of the regressions in Asdrubaly, Sorensen and Yosha (1996) and Bayoumi and Klein (1997) would, among other things, depend on the persistence of shocks, which affect the relationship between current and permanent income growth.

[13]One possibility is that there are three types of individuals, consuming respectively current income, permanent state income and permanent national income. In that case one may want to regress consumption growth on all three income measures. But even that is a rather rudimentary model.

tion in Japan in the context of a model in which some agents are "capital owners" and are perfectly diversified, while others only earn wages and are not diversified. While such a model does reasonably well in accounting for the stylized facts, there may be many other models that do so as well.

Third, the risksharing measures in Crucini (1998) and Del Negro (1997) rely on estimated innovations in permanent income. It is natural to focus on permanent income because individuals are concerned with risk associated with income at all future dates. But to obtain innovations in permanent income, the authors estimate a specific stochastic process of income. As mentioned in the introduction, and by Crucini (1998), results are likely to be sensitive to assumptions made about the income process.

Finally, most papers in this literature rely on consumption data. For the United States consumption data are not available at both the micro and state levels.[14] Therefore proxies have been used involving food consumption at the micro level and retail sales at the state level. Another problem, particularly in the context of the methodology adopted in this paper, is that consumption may also be affected by "taste" shocks.[15] In that case state consumption growth uncertainty relative to GSP growth uncertainty is not a good measure of risksharing even if consumption is accurately measured. For these reasons we avoid using consumption data in this paper to measure the extent of risksharing.

3 Measures of Risksharing

We consider a country with I regions. Let y_{it} be per capita income before risksharing of region i during year t. We define the growth rate of region i's income from t to $t+s$ as $g_{i,t,t+s} = ln\ y_{i,t+s} - ln\ y_{it}$, and write

$$g_{i,t,t+s} = \alpha + \lambda_s z_{it} + \epsilon_{i,t,t+s} \tag{1}$$

Here $\alpha + \lambda_s z_{it}$ and $\epsilon_{i,t,t+s}$ are the predictable and unpredictable components of growth, respectively. α is a constant. The vector z_{it} contains the information set available at time t to predict future growth.

[14]Better data are available for Canada and Japan, used by Crucini (1997) and van Wincoop (1995).

[15]Stockman and Tesar (1995) attribute low cross country consumption correlations to such taste shocks. In this context, Sorensen and Yosha (1998) criticize using cross country or cross state consumption correlations as a measure of risksharing.

It is important to stress that (1) is not a statistical process for income growth. Saying that we can write income growth as the sum of a predictable and an unpredictable component is not imposing any structure. The only structure that is imposed is that the predictable component is written as a linear function of variables in the information set. While it is important to agree on what variables are in the information set, the linearity is less restrictive than it seems. Let us say we choose the investment rate and current per capita income as variables that are in the information set. Then the vector z_{it} can also include non-linear terms, such as the investment rate squared, and non-linear interaction terms between the investment rate and per capita income. We will consider such non-linear terms in section 5.[16]

Regions remain exposed to national income risk even after complete risksharing among each other. We therefore focus on the uncertainty about a region's growth in deviation from national growth. Defining national variables with a superscript N, from (1) we can write the deviation from national growth as

$$g_{i,t,t+s} - g^N_{t,t+s} = \lambda_s(z_{it} - z^N_t) + u_{i,t,t+s} \tag{2}$$

where $u_{i,t,t+s} = \epsilon_{i,t,t+s} - \epsilon^N_{t,t+s}$ is referred to as "residual risk".[17]

The total region i growth innovation from t to $t+s$ is $\epsilon^N_{t,t+s} + u_{i,t,t+s}$. When regions engage in complete risksharing, exposure to the residual risk $u_{i,t,t+s}$ is eliminated. When all regions are equally risk-averse they will be identically exposed to national risk $\epsilon^N_{t,t+s}$ after risksharing. We therefore focus on the extent to which risksharing reduces uncertainty about the residual $u_{i,t,t+s}$.

While regions are not exposed to residual risk under complete risksharing, regions whose output is relatively uncertain, or highly correlated with national output, have to pay a premium to achieve this risksharing, while others receive a premium. In finance terms, one can think of $cov(\epsilon_{i,t,t+s}, \epsilon^N_{t,t+s})/var(\epsilon^N_{t,t+s})$ as the beta associated with a region's growth innovation.[18] Regions whose beta is larger than one have to pay a premium to

[16] The assumption that α is independent of the region is not restrictive. If α were to differ across regions, there must be a reason why we expect different growth rates across regions. This can only occur because of some prior information that we have, which should be captured by the information set z_{it}.

[17] For a region specific variable x_i, we define the national counterpart as $x^N = \sum_{j=1}^I \theta_j x_j$. The weights are given by $\theta_i = n_{it+s}E_t\ y_{it+s}/(\sum_{j=1}^I n_{jt+s}E_t\ y_{jt+s})$, where n_j is the state j population. These weights are based on the expected relative size of aggregate state income at $t+s$.

[18] For the US, Del Negro (1997) computes the beta for individual states.

diversify away their risk, while regions whose beta is less than one receive a premium. The sum of these premia across regions is zero.

For sufficiently long growth intervals there is only one observation per region, so that we are not able to compute the variance of residual risk for individual regions.[19] This problem arises because we have not assumed a specific process for income growth. If we had, this process could be estimated and the variance of residual risk at all horizons could be computed for each region separately. As discussed in the introduction, we would like to avoid this because results are known to be very sensitive to the assumed process. We also want to avoid the strong and unrealistic assumption that all regions have the same standard deviation of residual risk and the same "beta" (equal to one).

The compromise is to focus on an artificial "representative" region with "average" residual risk. We assume that the innovations $u_{i,t,t+s}$, with non-overlapping intervals $[t, t+s]$, are all independent draws from the distribution of residual risk of the representative region. This is attractive because it gives us many observations for the representative region, even if there is only one observation per region. Notice that this is an unconditional distribution. Conditional on the region, distributions of the innovations $u_{i,t,t+s}$ are generally different. These conditional distributions cannot be estimated if we have only one observation per region. We assume that residual risk of the representative region has a $N(0, \sigma_s^2)$ distribution and is independent of the national innovation $\epsilon_{t,t+s}^N$. The "beta" of the representative region is therefore one.

We use the representative region's standard deviation σ_s of residual risk as a measure of diversifiable risk. As we know from the finance literature, uncertainty about the deviation from the market return can be fully diversified at a zero risk premium (without affecting the expected return) when an asset has a "beta" of one. In our context, the residual risk of the representative region can be diversified at zero cost, without changing the expected growth rate. The extent of risksharing accomplished by financial markets and federal fiscal policy is measured by the percentage drop in the variance or standard deviation of residual risk based on income before and after risksharing.

Estimates of $u_{i,t,t+s}$ are obtained as follows. After choosing an information set we estimate (2) with a panel regression, using all non-overlapping intervals of horizon s in the sample (starting with the most recent). We estimate this with OLS, separately for each horizon s. For a given region these innovations should be uncorrelated across non-overlapping intervals as new innovations are orthogonal to past information. For a given

[19] The same problem arises with regards to the average variance of residual risk across regions. For long horizons we have only one observation on that moment as well.

interval, the innovations are not necessarily uncorrelated across all regions. For example, the residuals are positively correlated across two regions with relatively high exposure to national income shocks. But we generally cannot estimate these correlations since for sufficiently long horizons we have only one growth observation per region. Since we do not have any priors, and on average the residuals are uncorrelated across regions[20], it is natural to estimate (2) with OLS. No matter what the actual correlation structure, the estimates are always unbiased.

At this point it is useful to point out some differences with regressions in the empirical growth literature. First, the growth literature often includes contemporaneous right hand side variables. This leads to well known causality problems. We avoid that because our interest in growth uncertainty forces us to use variables that are in the information set at the time growth predictions are made. Second, the growth literature estimates (1), before subtracting aggregate (national or global) counterparts. The difference stems from our focus on residual risk. Estimating (2) has the advantage that the residuals are on average uncorrelated. A final difference is that the growth literature generally estimates cross section regressions. For long horizons we do the same, but for shorter horizons we estimate panel regressions, using all non-overlapping intervals of a given length.

The estimate for σ_s^2 is $\hat{u}'\hat{u}/(IH - K)$, where $\hat{u}$ is the stacked vector of estimated residuals, H is the number of observations per region and K is the number of variables in the information set. The variance associated with this point estimate is $\frac{2\sigma_s^4}{(IH-K)}$, where σ_s is replaced by its estimate.[21] We compute the variance reduction as the percentage drop in the estimated variance when comparing income before and after risksharing. Appendix A derives the precision (confidence interval) associated with the variance reduction estimate.

We now briefly discuss the welfare metric of risksharing. Details can be found in Appendices B and C. We assume that the expected growth rate of the representative region is the same as the expected national growth rate. This has two advantages. First, under this assumption the price of a claim on the representative region's per capita income is equal to the price of a claim on the per capita national income, so that under perfect risksharing the representative region consumes national per capita income. This

[20]This is true in the limit for a large number of regions. For example, consider I regions of equal size and equal variance σ^2 of the residuals. By construction $lim_{I\to\infty} \sum_{i\neq j} \frac{1}{I-1} corr(u_{i,t,t+s}, u_{j,t,t+s}) = lim_{I\to\infty} \frac{1}{I-1} cov(\sum_{i\neq j} u_{i,t,t+s}, u_{j,t,t+s})/\sigma^2$. Using $\sum_{i=1}^{I} u_{i,t,t+s} = 0$ this is equal to $-lim_{I\to\infty} \frac{1}{I-1} = 0$.

[21]This uses the fact that $\hat{u}'\hat{u}/(\sigma_s^2)$ has a chi-square distribution with $IH - K$ degrees of freedom.

is shown in Appendix B. Second, the assumption has the additional advantage that there are no gains from intertemporal consumption smoothing, so that we can focus exclusively on risksharing in the R sense.

We adopt standard CRRA preferences

$$V = E_0 \sum_{t=1}^{T} e^{-\beta t} \frac{(c_{it})^{1-\gamma}}{1-\gamma} \tag{3}$$

where E_0 denotes the expectation at time zero and γ is the constant rate of relative risk aversion.

Let y^B be the per capita income stream of the representative region before risksharing, and y^N the national per capita income stream. To compute potential gains from complete risksharing, we need to evaluate expected utility before risksharing, when residents of the representative region consume y^B, to that after risksharing, when they consume y^N.[22] The potential gain from risksharing can then be computed as the permanent percentage increase in expected income that yields a welfare improvement identical to what can be achieved through risksharing. In Appendix D it is shown that this welfare gain can be written as

$$0.5\gamma \sum_{t=1}^{T} \frac{e^{-(r_t-\bar{\mu}_t)}}{\sum_{s=1}^{T} e^{-(r_s-\bar{\mu}_s)}} \sigma_t^2 \tag{4}$$

where r_t is the t-period real interest rate on a risk-free bond, and $\bar{\mu}_t$ is the risk-adjusted growth rate of the representative region ($\bar{\mu}_t = \mu_t - 0.5\gamma\sigma_t^2$). The gain is equal to the rate of relative risk-aversion, divided by two, times a weighted average of the variance of residual risk, at all horizons. This reflects the fact that residual risk can be completely diversified. The weights depend on the difference between the risk-free interest rate and risk-adjusted growth rate at various horizons.

Using our estimates for σ_t^2 at all horizons, we can compute the gain from risksharing for the representative region. In order to determine the actual extent of risksharing by financial markets and federal fiscal policy, we compute the welfare gain (4) based on income before and after risksharing. We can then compute the percentage of potential welfare gains that have already been achieved by financial markets and federal fiscal policy.

[22] We implicitly assume an endowment economy, an assumption that is commonly made in the international risksharing literature. This ignores additional gains that are specific to production economies. For example, Obstfeld (1994a) has pointed out that there may be a switch to high risk, high expected return technologies when the risk is shared among countries. He shows that this raises the growth rate of countries.

4 Application to US Data

We now apply results from the previous section to US data from 1963 to 1990. We first discuss measures of income before and after risksharing and the information set used to determine the predictable component of growth. After that we quantify the diversifiable risk. The next section discusses the extent of risksharing through the two channels. Details on the data can be found in Appendix D.

4.1 Income Measures

The measure of income before risksharing is per capita gross state product (GSP), the state equivalent of GDP. We call this income before risksharing because it doesn't include net asset income. However, even without asset trade there is risksharing through migration across states. If a state suffers a negative shock and real wages decline, individuals can move to states that pay higher real wages.[23] This also reduces the risk for those who do not move because labor migration reduces real wage differentials across states. Here we do not attempt to quantify risksharing through migration. We focus on the extent of risksharing through financial markets and federal fiscal policy.[24]

The measure of income after risksharing is per capita pre-tax state income, obtained from Asdrubali, Sorensen and Yosha (1996). Since this income measure is described in detail in their paper, we will be relatively brief here. It is equal to total income by state residents and state and local governments, before federal taxes and transfers. It is interpreted as a measure of state income in the absence of intervention by the federal government. Since spending by state and local governments benefits residents from the state, their income is included in pre-tax state income. Because pre-tax state income includes net asset income, we consider it a measure of income after risksharing through financial markets.

[23] Barro and Sala-i-Martin (1991,1992b) provide evidence of a large and significant relationship between migration to US states and per capita personal income. Eichengreen (1993) regresses immigration to states on the lagged growth rate of wages (minus the national growth rate) and also finds a large and significant coefficient. Blanchard and Katz (1992) find that migration plays a very important role in the US as a buffer against business cycle shocks.

[24] The results are particularly relevant for individuals for which possible income gains from moving do not outweigh the cost of migration. This includes those on the margin, who are indifferent. The cost of moving also includes non-monetary costs, such as moving away from family, friends, a familiar environment, and desirable weather.

There is one difference between a state equivalent of GNP and pre-tax state income. GNP includes gross retained earnings, while pre-tax state income does not. About 80% of gross retained earnings is used to replenish the depreciated capital stock (consumption of fixed capital), while the rest is undistributed profits. Not including retained earnings is not necessarily a problem. Retained earnings is an investment that contributes to dividends in the future, which are captured in pre-tax state income. For consistency it would be best to also subtract retained earnings from GSP, so that we only measure the component of profits that is paid out to investors. The problem is that retained earnings data are not available at the state level. In most of the paper we will therefore use GSP as our measure of income before risksharing. In sensitivity analysis in section 5.4 we will construct an estimate of retained earnings at the state level and subtract it from GSP.

The final income measure is per capita disposable state income, also from Asdrubali, Sorensen and Yosha (1996). The difference between disposable state income and pre-tax state income captures the role of the federal government. It adds federal direct transfers to individuals in the state, plus federal grants to the government of the state, and subtracts total federal taxes raised in the state. For all income measures we assume a common price deflator across states. We do not need to compute this deflator since it drops out in eqn. (2). We will use estimates of state price deflators in the sensitivity analysis.

4.2 The Information Set

In order to select variables to be included in the information set, the empirical growth literature provides a useful guideline. The problem however, as mentioned in Sala-i-Martin (1994), is that in this enormous literature over 50 variables have been found to be correlated with growth in at least one regression. This is not surprising since there is significant multi-collinearity among the explanatory variables. Nonetheless this multi-collinearity problem is less of a concern to us than it is to those that attempt to identify which variables affect growth. We are after all only interested in the residual variance, the part of growth that cannot be predicted.

In contrast to the long list of variables usually included in country growth regressions, state level growth regressions, such as Barro and Sala-i-Martin (1991,1992a,b), generally only include per capita income (and sometimes migration) on the right hand side. Therefore the country growth literature is a more useful guide for the construction of our information set. Based on findings reported by Barro and Sala-i-Martin (1995), we construct the following "base information set":[25] (1) the log of initial per capita GSP, (2) the five year lagged population growth rate, (3) the ratio of government expenditure to GSP, (4) the ratio of investment to output in the manufacturing sector[26] , (5) the fertility rate (births per thousand women), (6) average migration as a fraction of population during the most recent 10-year period, (7) the one year lagged growth rate of per capita GSP (change in logs), (8) total expenditure on education per student in average daily attendance, and (9) total expenditure on education divided by the total school age population.[27] We omit some of the variables that Barro and Sala-i-Martin (1995) use in their country growth regressions that are clearly not relevant for our sample: political instability, the black market premium, the tariff rate, and terms of trade growth. We also omit life expectancy, which is an important variable in country growth regressions, but not available at the state level. We include the migration rate in the information set, which is generally not used in country regressions, but is likely to be more important within countries. Barro and Sala-i-Martin (1991) find the migration rate to be positive and significant in a cross section regression of US state growth rates on per capita income and the migration rate.

In order to compare the risk associated with different income measures, for each income measure we use the same information set to compute the variance of residual risk. We will also consider subsets of the "base information set" and adding non-linear transformations of the variables.

4.3 Measuring Risk

For the pre-risksharing income measure, GSP, the first row of Table 1 reports the standard deviation of residual risk for the representative state. Over a 26-year period, with the base information set, the standard deviation is 12.8%. With an expected annual growth rate of 1.8% (the average over our sample), and controlling for national growth uncertainty, this implies a 95% confidence interval of 26-year growth of 33% to 85%. This rather substantial uncertainty would be completely eliminated under complete riskshar-

[25] For income growth over the interval $[t, t+s]$, the variables in the information set are measured at time t.

[26] Total investment is not available at the state level.

[27] Enrollment rates and education attainment are often used as education variables in country growth regressions. Because of incomplete data, we use the alternative education variables in (8) and (9) at the state level.

ing. Even at a five year interval residual growth uncertainty is substantial. The standard deviation is 8.1%. With an expected 1.8% annual growth rate, and controlling for national growth uncertainty, the 95% confidence interval for 5-year growth is -7% to +25%. The standard errors associated with this risk measure are all quite small.

Another way of quantifying risk is to compute the welfare gain from diversifying residual risk. In order to implement the formula (4) for the welfare gain from complete risksharing, we need to make assumptions about the risk-free interest rate, the risk adjusted growth rate and the coefficient of relative risk aversion. The rate of relative risk-aversion is assumed to be three. This is the average of estimates in Friend and Blume (1975). We consider it to be a good consensus estimate overall. Since from (4) the welfare gain is proportional to γ, the gain would be one third of what we report if one believes in log-utility, and three times of what we report if one believes the coefficient of relative risk-aversion is nine. The risk-aversion parameter does not affect our estimate of the percentage of the total potential gains that has already been achieved by financial markets and federal fiscal policy. It only affects the level of the potential gain itself.

The risk-adjusted growth rate is $\bar{\mu}_t = \mu_t - 0.5\gamma\sigma_t^2$. We find that the annualized risk premium, $0.5\gamma\sigma_t^2/t$, is somewhere between 0.001 (0.1%) and 0.002 (0.2%). The average annual growth rate from 1950 until 1990 in the US is 1.8%. We set the risk-adjusted growth rate at a constant 1.6% annually: $\bar{\mu}_t = 0.016t$.

A measure of the risk-free rate r_t is constructed as follows. Over the period 1889 to 1978 Mehra and Prescott (1985) find an average annual real interest rate of 0.8% on relatively riskless three month US Treasury securities. In order to obtain an average term structure, we use data from 1947 to 1985 from McCulloch (1990) on the zero-coupon yield curve for Treasury bonds. r_t/t is set equal to 0.008, plus the average difference between the t year yield and the 3-month yield.

Based on this parameterization, the last three columns of Table 1 report the welfare gain for the representative state from complete risksharing. With the base information set the welfare gain is 1.6% at a 26-year horizon. The potential benefits from risksharing are clearly large. They amount to $1200 per year per household in 1995 dollars.[28] If $\gamma = 35$, which Mankiw and Zeldes (1991) find is necessary to account for the equity premium, the welfare gain rises to a stunning 18.7%, or $14,000 per year per household. Since it is hard to know what the correct rate of relative risk-aversion is, we believe the confidence intervals for growth rates reported above are at least as informative.

[28]The $1,200 is 1.6% of GDP in 1995, divided by the number of households in 1995 (from the Statistical Abstract of the United States).

The next row of Table 1 reports the results for an empty information set, in which case all deviations from national growth are considered to be unpredictable. This tells us something about the predictive power of the information set. Although we have found that there is subtantial uncertainty left, the base information set nonetheless has significant predictive power, particularly at long horizons. At a 15-year horizon the welfare gain is 49% higher under the empty information set when compared to the base information set. For the 26-year horizon it is 42% higher.[29] There is less predictive power at short horizons. At a 5-year horizon the welfare gain is only 10% higher under an empty information set.

The remaining rows of the table report the risk-measures for the "best" 1, 2, and 3 variables in the information set, defined as the variables with the most predictive power as measured by the welfare gain at the 26-year horizon. The first best variable is government spending as a fraction of GSP. The second best is the lagged one year growth rate. The third best is migration as a fraction of the population. If we consider the difference between the welfare gain based on the empty information set and the base information set as a measure of predictive power, so that the base information set accounts for 100% of the predictive power, 80% of that comes from the three "best" variables. Adding more variables to the information set does not significantly improve predictive power.

It may be surprising that initial per capita GSP does not appear as one of the most important variables in the information set since initial income is almost always the most important variable in country growth regressions. Barro and Sala-i-Martin (1992) provide evidence of convergence among the states. While they find the convergence parameter to be significant for almost all decades up to 1970 (their data starts in 1880), it is insignificant for the most recent intervals 1970-80 and 1980-88. One possible explanation is that most convergence has already taken place by now.[30]

[29]These numbers are not sensitive to the assumed rate of relative risk-aversion.

[30]In the same spirit, Easterly et. al. (1993) find that growth rate differences across US states and across European provinces have lower persistence than country growth differentials. This is attributed to the higher degree of technological convergence at the intranational level. They find that persistence of growth rate differences among US states has dropped significantly over time from the 1950s to the 1980s, going from large positive persistence to large negative persistence.

5 What Have Financial Markets and Fiscal Federalism Accomplished?

5.1 The Role of Financial Markets

Figures 1 to 3 provide information on the extent of risksharing through financial markets. Figure 1 shows the percent variance reduction as a function of the length of the horizon, together with a 95% confidence interval. At a 26-year horizon there is a 66% variance reduction, with a 95% confidence interval of 50% to 82%. We can clearly reject either no risksharing or complete risksharing. The point estimate indicates that about two thirds of the risk is shared through financial markets. The variance reduction is not very sensitive to the horizon. The average variance reduction over all horizons is 57%. This stands in stark contrast to the limited extent of risksharing among countries. Athanasoulis and van Wincoop (1997) find that at the national level, the standard deviation of residual risk based on GDP is about the same as that based on consumption.[31]

Figure 2 shows the percent standard deviation reduction, which is on average 35%. One would expect this to be smaller than the variance reduction because $\frac{d\sigma 2}{\sigma^2} \approx 2\frac{d\sigma}{\sigma}$ for marginal changes in risk. Which of the two is most relevant, the standard deviation reduction or variance reduction, depends on the utility function. Under standard assumptions about preferences, such as (3), it is the variance that matters, as is apparent from (4). Epstein and Zin (1990) construct preferences with "first-order risk aversion", in which case the risk premium on a small gamble is proportional to the standard deviation rather than the variance of the payoff. While we are not going to resolve what type of preferences are more relevant, the perceived extent of risksharing depends on this in an important way. Based on the standard deviation, about two thirds of risk is not shared. Based on the variance, about two thirds of risk is shared.

For standard CRRA preferences (3) the thick lines in Figure 3 show the welfare gain from risksharing based on GSP and pre-tax state income, as a function of the horizon T. The shaded area between the two lines shows how much of the potential gain from risksharing has already been achieved by financial markets. At the 26-year horizon 60% of the potential welfare gain has been achieved through financial markets. Not surprisingly, this number is similar to what we found for the variance reduction because the welfare gain is proportional to a weighted average of the variance of residual risk

[31] They do not compute results for GNP, which would make it more comparable to this paper.

at all horizons. Just as with the variance reduction, the percentage of potential welfare gains achieved is not very sensitive to the horizon.

These numbers indicate that financial markets have already achieved a substantial portion of potential gains from risksharing. This may seem surprising because these markets trade claims on only a small component of wealth. There is no trade in claims on human capital and non-corporate profits. On top of this, Coval and Moscovitz (1997) and Huberman (1997) report evidence of an intranational home bias with regards to tradable securities.[32] But, as was mentioned earlier, it is possible that a limited extent of diversification achieves most of the attainable welfare gains. Moreover, trade in claims on human capital may not be important when most of the residual risk in GSP shows up in profits rather than wages.

In order to shed some more light on this we compute total asset income of state residents in the absence of risksharing as $\Pi = GSP - W$, which we simply refer to as profits. Here W is total wages, broadly defined as the sum of employee compensation, other labor income (mostly benefits) and proprietor's income.[33] W therefore also includes income from non-corporate business. Π is the sum of rental income, corporate profits and interest paid by businesses. We compute actual before tax asset income (after risksharing) as the difference between pre-tax state income and W, and refer to this as dividends.

We then compute the unpredictable components of per capita growth rates of wages, profits and dividends, all in deviation from national growth rates. The same base information set is used. At a 26-year horizon the standard deviation of the per capita wage growth rate is 8.3%, while the standard deviation of per capita profits growth is 22.2%. Most of the risk shows up in profits. Since there are traded claims on profits, it may not be surprising to see so much risksharing. However, the standard deviation of per capita dividends growth is 17.6%. If everyone held the same portfolio their dividend income would be perfectly correlated and we would find a standard deviation of zero for residual dividend growth. Although dividends are less volatile than profits, a standard deviation of 17.6% is still large. We find that to a large extent the achieved gain from risksharing

[32] Coval and Moskowitz (1997) find that the average distance between the location of US investment managers and the US firms in which they invest is 1654 miles, while the average distance under full diversification would be 1814 miles. The intranational portfolio bias is almost certainly larger for individual investors. Huberman (1997) reports a home bias in households' holdings of regional Bell companies.

[33] The data are from the personal income CD-ROM from the Bureau of Economic Analysis. W is a gross income measure, before personal contributions to social security or any other taxes are subtracted.

is associated with a lower correlation after risksharing between asset income growth and wage growth. At the 26-year horizon, the correlation between wage growth and profits growth is 0.69, while the correlation between wage growth and dividend growth is only 0.01.

Using the profits data we can obtain one other interesting insight. We have reported that 60% of potential welfare gains from risksharing have already been achieved through financial markets. It is of interest to know how large that number would be if agents held identical portfolios. We assume that agents receive wages (same as before) plus per capita national profits. In that case the standard deviation of residual income growth uncertainty drops to 4.8% at the 26-year horizon, down from 7.5% for pre-tax state income. With this income measure we find that 81% of potential welfare gains at the 26-year horizon would be achieved, versus the 60% that is already accomplished.[34] Although we find that financial markets have already achieved a lot, the fact that we could do even a little better is consistent with findings by Huberman (1997) and Coval and Moscovitz (1997) of a portfolio "home bias" within the United States.

5.2 The Role of Federal Fiscal Policy

Figures 4 and 5 describe the extent of risksharing though federal fiscal policy. Figure 4 shows the percent variance reduction from pre-tax state income to disposable state income. It is on average about 27% and significant for almost all horizons. For horizons from 3 to 13 years the extent of risksharing through federal fiscal policy is most extensive, with a variance reduction of about 36%. Figure 5 shows that the standard deviation reduction is again smaller, an average of 15% over all horizons.

There is a substantial literature that has described the role of federal fiscal policy through regressions of fiscal variables (or disposable state income) on per capita GSP or per capita state personal income.[35] These regressions have been done both in levels, reflecting the long term redistributive role of federal fiscal policy, and in changes,

[34] A possible drawback of this experiment is that it ignores that some states are richer than others. We have also considered an alternative, whereby per capita dividends received are still perfectly correlated across states but the level of per capita dividends varies according to how rich a state is. The ratio of one state's per capita dividends to that of another is assumed to be constant over time and equal to the average over the sample. This alternative experiment has little effect on the results though: 80% of potential welfare gains at the 26-year horizon would be achieved through financial markets.

[35] See Asdrubali, Sorensen and Yosha (1996), Atkeson and Bayoumi (1993), Bayoumi and Masson (1995), Melitz and Zumer (1997), Sala-i-Martin and Sachs (1992), and von Hagen (1992).

capturing fiscal stabilization. From the point of view of risksharing, the stabilization role of fiscal policy is of most interest. Estimates of the net federal transfer to a state in response to a dollar income loss (relative to other states) range from 10 cents (von Hagen (1992)) to 31 cents (Bayoumi and Masson (1995)). If we take the average of this, 20 cents[36], we would conclude that for every dollar drop in pre-tax state income, relative to its national counterpart, disposable state income drops only 80 cents. Assuming that this applies to both expected and unexpected changes in income, the standard deviation of income growth is reduced by 20%. This is consistent with our results in Figure 5, where the reported standard deviation reduction from pre-tax state income growth to disposable state income growth is about 20% for horizons up to 13 years. For longer horizons the standard deviation reduction is a bit smaller, about 10% on average.

5.3 Total Risksharing

Figure 6 shows the welfare gain from risksharing for all three income measures. The dotted area and the area with the slanted lines represent the component of the potential gain from risksharing achieved by financial markets and federal fiscal policy, respectively. At the 26-year horizon 60% of potential welfare gains from risksharing have been achieved through financial markets, and an additional 11% through federal fiscal policy, bringing the total to a very large 71%. Figures 7 and 8 show the total variance and standard deviation reduction when comparing GSP to disposable state income. The average variance reduction is 75%, while the average standard deviation reduction is 45%.

5.4 Sensitivity Analysis

Table 2 reports results from sensitivity analysis, focusing on the conclusion that 71% of potential welfare gains have already been achieved. The first four rows consider results based on smaller information sets: the "empty" information set and information sets of the one, two and three variables with most predictive power (discussed in section 4.3). The results are not very sensitive to the the size of the information set. The estimate of the percentage of welfare gains already achieved is only slightly lower for the empty information set as for the entire information set of nine variables.[37]

[36] This is what Obstfeld and Perri (1998) refer to as a rough consensus from the literature.

[37] As reported in Table 1, the level of the potential gains is obviously higher for an empty information set.

We next consider adding to the base information set variables that are non-linear functions of the existing variables. Although we have tried many such non-linear relationships, we report only eight examples in Table 2. The message is clear. Adding different functional forms has practically no effect on the results.

The next two rows respectively raise and lower the risk-free interest rate by one percent annually. As can be seen from (4) this only affects the weights attached to income growth uncertainty at different horizons. But since the estimated variance reduction from risksharing does not depend a lot on the horizon, the percentage of welfare gains achieved depends very little on the assumed interest rate.

The next row shows the results when we deflate the state income measures by state price deflators. We use state CPI deflators constructed by Del Negro (1997) from cost of living indices for metropolitan areas, other urban areas, and rural areas.[38] Using these state price deflators has little affect on the results. The percentage of welfare gains achieved remains at 71% at the 26-year horizon.

Finally, we consider a different measure of income before risksharing, GSP minus our estimate of retained earnings. As discussed above, it is appropriate to subtract retained earnings from GSP, but data are not available at the state level. Gross retained earnings is equal to the sum of depreciation (consumption of fixed capital) and undistributed profits. We construct separate estimates of both. National data for depreciation are distributed to the states based on their share in the national capital stock.[39] This implicitly assumes that all states have the same rate of depreciation. Undistributed profits is a smaller part of gross retained earnings (about 20%), but more difficult to estimate at the state level. We use state profits data to distribute national undistributed profits to the states.[40] The results from this experiment are shown in the last row of Table 2. The estimated percentage of welfare gains achieved rises a bit at the 26-year horizon, to 76%.

[38]We would like to thank Marco Del Negro for making these data available to us. Because this data set starts in 1970, we use national CPI inflation rates for the period 1963-1969 of our sample.

[39]State capital stock data are from the data set constructed by Munnell (1990) and updated by Jeff Fuhrer at the Federal Reserve Bank of Boston. The data set starts in 1970. For earlier years in the sample we assume that a state's share in the national capital stock is the same as that in 1970.

[40]To be precise, for the first year of the sample, 1963, national undistributed profits are allocated to the states based on the relative size of their profits. We then estimate changes over time in undistributed state profits. We regress the change in undistributed national profits, divided by last year's national profits, on the percentage change in national profits. The result from this regression is used at the state level to compute changes in undistributed state profits from changes in total state profits. The change in national undistributed profits that is not explained by the regression is distributed to the states based on the level of their profits (as opposed to changes in profits).

Because depreciation is very smooth, the net state product (subtracting depreciation) is more volatile than GSP, so that the potential gain from risksharing is a bit larger and therefore also the estimated percentage of gains achieved. The exact opposite is the case when subtracting undistributed profits. Although we do not observe state corporate profits and dividends, the latter will be less volatile. Substracting undistributed profits from GSP should therefore reduce its volatility. If we had a better estimate of undistributed state profits our estimate of the percentage of welfare gains received would therefore probably be closer to the original 71%.

6 Conclusion

Using an approach that does not rely on a particular model or income process, we have documented aggregate income growth uncertainty at the state level, and the extent to which this uncertainty is reduced by risksharing through financial markets and federal fiscal policy. We estimate that 71% of potential welfare gains from risksharing have already been achieved; 60% through financial markets and 11% through federal fiscal policy. However, we would like to caution against drawing too strong conclusions about the extent of risksharing within the United States. The welfare metric we used is based on standard CRRA expected utility preferences, whereby gains from risksharing are proportional to the variance rather than the standard deviation of diversifiable income growth uncertainty. Conclusions change substantially if we express risk in terms of the standard deviation of growth uncertainty. In that case we found that on average 55% of residual growth uncertainty has *not* been diversified away. Financial markets reduce the standard deviation of residual income growth uncertainty by only 35%.

Appendix A: Precision of Risksharing Measure

Let $\hat{\sigma}_B^2$ and $\hat{\sigma}_A^2$ be the estimated variance of residual risk before and after risksharing. In this Appendix we discuss how we compute the precision associated with the risk reduction measure $1-(\hat{\sigma}_A^2/\hat{\sigma}_B^2)$. We first linearize the expression as $\alpha_1\hat{\sigma}_A^2+\alpha_2\hat{\sigma}_B^2$, where we treat $\alpha_1=-1/\hat{\sigma}_B^2$ and $\alpha_2=\hat{\sigma}_A^2/\hat{\sigma}_B^4$ as constants, substituting the point estimates. We then compute $var(\alpha_1\hat{\sigma}_A^2+\alpha_2\hat{\sigma}_B^2)=\alpha_1^2 var(\hat{\sigma}_A^2)+\alpha_2^2 var(\hat{\sigma}_A^2)+2\alpha_1\alpha_2 cov(\hat{\sigma}_A^2,\hat{\sigma}_B^2)$. As discussed in the text, $var(\hat{\sigma}_B^2)=\frac{2\sigma_B^4}{(IH-K)}$, and similarly for $var(\hat{\sigma}_A^2)$. To obtain an estimate of the covariance between $\hat{\sigma}_A^2$ and $\hat{\sigma}_B^2$, we multiply the product of the estimated standard deviations of $\hat{\sigma}_A^2$ and $\hat{\sigma}_B^2$ by the correlation between v_B and v_A, the vectors containing the squared estimated residuals. Once we have computed the variance associated with the risk reduction estimate, we approximate the 95% confidence interval as two standard errors below and above the point estimate.

Appendix B

Let *rep* denote the representative region. We will prove that the price of a claim on *rep*'s per capita income stream relative to the price of a claim on the national per capita income stream is 1 when (i) $E_0\ y_{rep,t}=E_0\ y_t^W$, (ii) $cov(\epsilon_{rep,0,t},\epsilon_{0,t}^N)=var(\epsilon_{0,t}^N)$ for $t=1,..,T$ (*rep*'s β of growth rates is one). To do this, all we need to show is that for any horizon t the relative price of a claim on *rep*'s per capita income stream is one. There are n_{it} (population) claims on region i's per capita income in period t, each with a payoff of y_{it} and a price of p_i. Consider an investor in any region, who at time 0 invests a total of Y_t in period t claims ($\sum_{t=1}^T Y_t$ is the investor's period 0 revenue from selling all claims on her own income). The investor maximizes $E_0\left(\sum_{i=1}^I q_i y_{it}\right)^{1-\gamma}/(1-\gamma)$, s.t. $\sum_{i=1}^I q_i p_i=Y_t$. Here q_i is the quantity of claims on state i income purchased. The first order conditions with respect to the q_i's are:

$$E_0\left(\sum_{i=1}^I q_i y_{it}\right)^{-\gamma} y_{jt}=\nu p_j \quad j=1,..,I$$

The price of a claim on the per capita national income is $\sum_{i=1}^I \frac{p_i n_{it}}{n_t}$. Using the first order conditions, and the fact that in equilibrium $\frac{q_i}{\sum_{j=1}^I q_j}=\frac{n_{it}}{n_t}$, the relative price of a claim on *rep*'s per capita income is:

$$\frac{p_{rep}}{\sum_{i=1}^I p_i n_{it}/n_t}=\frac{E_0\ (y_t^N)^{-\gamma} y_{rep,t}}{E_0\ (y_t^N)^{1-\gamma}}$$

$y_{rep,t}=(E_0\ y_{rep,t})e^{-0.5var(\epsilon_{rep,0,t})+\epsilon_{rep,0,t}}$, where $\epsilon_{rep,0,t}$ has by assumption a normal distribution. We make the commonly made approximation in the risksharing literature

(Lewis (1996), van Wincoop (1994)) that the distribution of $ln\ y_t^N$ is the same as that of $\sum_{i=1}^N \theta_{i,0,t} ln\ y_{it}$, which is normal with standard deviation σ_{Nt}^2. With that approximation $y_t^N = (E_0\ y_t^N) e^{-0.5\sigma_N^2 + \epsilon_{0,t}^N}$. Substituting these expressions for y_t^N and $y_{rep,t}$ into the formula for the relative price, it follows immediately from assumptions (i) and (ii) that the relative price is one.

Appendix C: Welfare Gains

In this appendix we compute the gains from risksharing for the representative region. Since $\epsilon_{rep,t,t+s} = u_{rep,t,t+s} + \epsilon_{t,t+s}^N$, with $u_{rep,t,t+s}$ and $\epsilon_{t,t+s}^N$ normally distributed and independent, we can write $y_{rep,t} = (E_0 y_{rep,t}) e^{u_{rep,t,t+s} + \epsilon_{t,t+s}^N - 0.5\sigma_t^2 - 0.5\sigma_{Nt}^2}$, where $\sigma_{Nt}^2 = var(\epsilon_{t,t+s}^N)$. Since we assume that $E_0 y_{rep,t} = E_0 y_t^N$, expected utility before risksharing is

$$\sum_{t=1}^{T} e^{-\beta t} \frac{(E_0\ y_t^N)^{1-\gamma}}{1-\gamma} e^{-0.5\gamma(1-\gamma)(\sigma_t^2 + \sigma_{Nt}^2)} \tag{5}$$

After risksharing consumption is equal to y_t^N. We make the same commonly made approximation as in Appendix B that the distribution of $ln\ y_t^N$ is the same as that of $\sum_{i=1}^N \theta_{i,0,t} ln\ y_{it}$, which is normal with standard deviation σ_{Nt}^2. Expected utility after risksharing then becomes

$$\sum_{t=1}^{T} e^{-\beta t} \frac{(E_0\ y_t^N)^{1-\gamma}}{1-\gamma} e^{-0.5\gamma(1-\gamma)\sigma_{Nt}^2} \tag{6}$$

From (5) and (6) the permanent percentage increase in consumption that leads to a welfare gain equivalent to what can be achieved through risksharing is equal to

$$gain = \left[\sum_{t=1}^{T} \Omega_t\ e^{-0.5\gamma(1-\gamma)\sigma_t^2} \right]^{1/(1-\gamma)} - 1 \tag{7}$$

The weights Ω_t are defined as

$$\Omega_t = \frac{e^{-\beta t}\ (\bar{y}_{rep,t})^{1-\gamma}}{\sum_{s=1}^T e^{-\beta s}\ (\bar{y}_{rep,s})^{1-\gamma}} \tag{8}$$

where $\bar{y}_{rep,t} = (E_0\ y_t^N) e^{-0.5\gamma\sigma_{rep,t}^2}$ is the certainty equivalent of rep's income at time t. Here $\sigma_{rep,t}^2 = var(\epsilon_{rep,t}) = \sigma_t^2 + \sigma_{Nt}^2$.

The weights Ω_t can be simplified by introducing the following notation. Let μ_t be the expected growth rate of national per capita income over [0,t]: $e^{\mu_t} = \frac{E_0\ y_t^N}{y_0^N}$. The certainty equivalent of rep's growth rate is $\bar{\mu}_t = \mu_t - 0.5\gamma\sigma^2_{rep,t}$. From the Euler equation, the t period real interest rate on a risk-free bond is $r_t = \beta t + \gamma\bar{\mu}_t$. With this notation,

$$\Omega_t = \frac{e^{-(r_t - \bar{\mu}_t)}}{\sum_{s=1}^{T} e^{-(r_s - \bar{\mu}_s)}} \tag{9}$$

So the appropriate discount factor is the difference between the risk-free real interest rate and the risk adjusted growth rate.

While we use (7) to compute the welfare gain for the representative region, the somewhat more intuitive expression in the text, which numerically is very close to (7), can be obtained using the approximations $e^x \approx 1 + x$ and $(1+x)^a \approx 1 + ax$ for x close to zero. This results in (4).

Appendix D: The Data

We use annual data from 1963 to 1990. Below we list the data, the source for the data and any observations which are missing.

Measures of Income

These are all obtained from *Asdrubali, Sorensen and Yosha (1996).*

- Gross State Product: value added of the industries of a state.
- Pre-tax State Income: total income by state residents and state and local governments, before federal taxes and transfers. It is also equal to gross state product, minus gross retained earnings (including capital consumption), plus net factor income.
- Disposable State Income: pre-tax state income, plus federal direct transfers to individuals in the state, plus federal grants to the government of the state, minus total federal taxes raised in the state.

Data to Construct the Information Set

For some of the variables in the information set there are missing observations. Since these series are quite smooth, we interpolate the two closest data points. We indicate these missing observations below.

- Capital Expenditures: from the *Annual Survey of Manufactures*. Missing observations for all states: 1979-81. Observations missing for particular states: 1970 Minnesota, 1971 Minnesota.

- Education Expenditure per student in Average Daily Attendance: from *Statistical Abstract of the United States*, US Department of Commerce, Bureau of the Census. Missing observations for all states: 65, 67, 69, 71, 72, 82, 83, 94. Observations missing for particular states: 1975 Georgia, 1975 Maine, 1975 Massachussettes, 1975 New Jersey, 1976 Massachussettes, 1978 Wisconsin, 1979 Alaska, 1979 California, 1979 Florida, 1979 Georgia, 1979 Illinois, 1979 Wisconsin, 1980 Arizona, 1980 Florida.

- State Government Expenditures: from *State Government Finances* 1947-1990.

- Fertility (births per 1,000 women): from *Statistical Abstract of the United States*.

- Net migration: from *Current Population Reports*. We have observations for net migration over the periods 1960-1970, 1970-1979, 1980-1988. We divide the net migration by total population over the period to obtain the percentage net migration. Missing observation: 1970-1979 Nebraska.

- School Age Population-children attending school between the ages of 5 and 17: from *Statistical Abstract of the United States*. Missing observations for all states: 1963, 64, 69. To get 63,64 interpolated 60 to 65.

- Total Expenditure on Education: from *Statistical Abstract of the United States*. Missing observations for all states: 1965, 67, 69, 71-72, 76, 82-83.

- Population: from *Regional Data, Haver Analytics*. We use observations for 1957-1990.

- Value Added by Manufactures: from the *Annual Survey of Manufactures*. Missing observations for all states: 1979-81. Observations missing for particular states: 1970 Minnesota, 1971 Minnesota.

References

Altug, S. and R.A. Miller, 1990, Household Choices in Equilibrium, *Econometrica* 58, no. 3, 543-70.

Asdrubali, P., B.E. Sorensen and O. Yosha, 1996, Channels of Interstate Risksharing: United States 1963-1990, *Quarterly Journal of Economics* 111, 1081-1110.

Athanasoulis, S., and E. van Wincoop, 1997, Growth Uncertainty and Risksharing, Staff Reports 30, Federal Reserve Bank of New York.

Athanasoulis, S., R. Shiller and E. van Wincoop, 1998, Macro Markets and Financial Security, working paper, Federal Reserve Bank of New York.

Atkeson, A. and T. Bayoumi, 1993, Do Private Capital Markets Insure Regional Risk? Evidence from the United States and Europe, *Open Economies Review* 4, 303-24.

Backus, D.K., Kehoe, P.J., and F.E. Kydland, 1992, International Real Business Cycles, *Journal of Political Economy* 100, 745-775.

Barro, R.J. and Sala-i-Martin, X., 1991, Convergence Across States and Regions, *Brookings Papers on Economic Activity* 1, 107-82.

Barro, R.J. and Sala-i-Martin, X., 1992a, Convergence, *Journal of Political Economy* 100, 223-251.

Barro, R.J. and Sala-i-Martin, X., 1992b, Regional Growth and Migration: A Japan-United States Comparison, *Journal of the Japanese and International Economies* 6, 312-346.

Barro, R.J. and Sala-i-Martin, X., 1995, *Economic Growth*, McGraw-Hill, New York.

Bayoumi, T. and M.W. Klein, 1997, A Provincial View of Economic Integration, *IMF Staff Papers* 44, no. 4, 534-556.

Bayoumi, T. and P.R. Masson, 1995, Fiscal Flows in the United States and Canada: Lessons for Monetary Union in Europe, *European Economic Review* 39, 253-274.

Blanchard, O.J. and L.F. Katz, 1992, Regional Evolutions, *Brookings Papers on Economic Activity* 1:1992.

Cochrane, J., 1991, A Simple Test of Full Consumption Insurance, *Journal of Political Economy* 99, 957-976.

Cole, H.L. and M. Obstfeld, 1991, Commodity Trade and International Risksharing: How Much do Financial Markets Matter?, *Journal of Monetary Economics* 28, 3-24.

Coval, J.D. and T. Moskowitz, 1997, Home Bias at Home: Local Equity Preference in Domestic Portfolios, working paper, University of Michigan Business School.

Crucini, M., 1998, On International and National Dimensions of Risksharing, *Review of Economics and Statistics*, forthcoming.

Del Negro, M., 1997, Aggregate Risksharing Across US States, working paper, Yale University.

Easterly, W., M. Kremer, L. Pritchett and L.H. Summers, 1993, Good Policy or Good Luck? Country Growth performance and Temporary Shocks, *Journal of Monetary Economics* 32, no. 3, 459-484.

Eichengreen, B., 1993, Labor Markets and European Monetary Unification, in: P.R. Masson and M.P. Taylor, eds., *Policy Issues in the Operation of Currency Unions*, Cambridge University Press.

Epstein, L.G. and S.E. Zin, 1990, 'First-Order' Risk Aversion and the Equity Premium Puzzle, *Journal of Monetary Economics* 26, 387-407.

Friend, I. and M.E. Blume, 1975, The Demand for Risky Assets. *American Economic Review* 65, 900-922.

Hayashi, F., J. Altonji, and L. Kotlikoff, 1996, Risk-Sharing Between and Within Families, *Econometrica* 64, no. 2, 261-94.

Hess, G.D. and K. Shin, 1997, Risk Sharing by Households Within and Across Regions and Industries, working paper, University of Cambridge.

Hess, G.D. and K. Shin, 1998, Intranational Business Cycles in the United States, *Journal of International Economics*, forthcoming.

Huberman, G., 1997, Familiarity Breeds Investment, working paper, Columbia Business School.

Kose, M.A., 1995, Welfare Implications of International Risk Sharing, working paper, Brandeis University.

Lewis, K.K., 1996, Consumption, Stock Returns, and the Gains from International

Risksharing. NBER Working Paper no. 5410.

Mace, B.J., 1991, Full Insurance in the Presence of Aggregate Uncertainty, *Journal of Political Economy* 99, 928-56.

Mankiw, N.G. and S.P. Zeldes, 1991, The Consumption of Stockholders and Non-stockholders, *Journal of Financial Economics* 29, 97-112.

McCulloch, J.H., 1990, Appendix B to The Term Structure of Interest Rates, by R. Shiller, in B.Friedman and F. Hahn, eds., *The Handbook of Monetary Economics* V. I, New York, North-Holland.

Mehra, R. and E.C. Prescott, 1985, The Equity Premium: A Puzzle, *Journal of Monetary Economics* 15, 145-161.

Melitz, J.. and F. Zumer, 1997, Regional Redistribution and Stabilization by the Center in Canada, France, the UK and the US, working paper, CREST-INSEE.

Mendoza, E., 1995, The Terms of Trade, the Real Exchange Rate and Economic Fluctuations, *International Economic Review* 36, 101-1337.

Miller, B. and H. Sieg, 1997, A Microeconomic Comparison of Household Behavior Between Countries, *Journal of Business and Economic Statistics* 15 (2), 237-253.

Munnell, A., 1990, How Does Public Infrastructure Affect Regional Economic Performance?, in Munnell, A., ed., *Is There a Shortfall in Public Capital Investment?*, Federal Reserve Bank of Boston.

Obstfeld, M., 1994a, Risk-Taking, Global Diversification and Growth, *American Economic Review* 84, 1310-1329.

Obstfeld, M., 1994b, Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability, *European Economic Review* 38, 1471-86.

Obstfeld, M. and G. Perri, 1998, Regional Nonadjustment and Fiscal Policy, *Economic Policy*, forthcoming.

Persson, T. and G. Tabellini, 1996, Federal Fiscal Institutions: Risk Sharing and Moral Hazard, *Econometrica* 64(3), 623-46.

Sala-I-Martin, X., 1994, Economic Growth: Cross-Sectional Regressions and the Empirics of Economic Growth, *European Economic Review* 38(3-4), 739-747.

Sala-i-Martin, X. and J. Sachs, 1992, Fiscal Federalism and Optimum Currency Areas: Evidence for Europe from the United States, in Canzoneri, M., Masson, P. and V. Grilli (eds.), *Establishing a Central Bank: Issues in Europe and Lessons from the US*, Cambridge University Press, London.

Shiller, R.J. and S. Athanasoulis, 1995, World Income Components: Measuring and Exploiting International Risk Sharing Opportunities., NBER working paper 5095.

Sorensen, B.E. and O. Yosha, 1997, Income and Consumption Smoothing among US States: Regions or Clubs?, working paper, Brown University.

Sorensen, B.E. and O. Yosha, 1998, International Risksharing and European Monetary Unification, *Journal of International Economics*, forthcoming.

Stockman, A.C. and L. Tesar, 1995, Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Co-Movements, *American Economic Review* 85(1), 168-185.

Tesar, L.L., and I.M. Werner, 1997, The Internationalization of Securities Markets Since the 1987 Crash, working paper, University of Michigan.

Tesar, L.L., 1995, Evaluating Gains from International Risksharing, *Carnegie-Rochester Conference Series on Public Policy* 42, 95-143.

Townsend, 1994, Risk and Insurance in Village India, *Econometria* 62, 539-91.

van Wincoop, E., 1994, Welfare Gains from International Risksharing, *Journal of Monetary Economics* 34, 175-200.

van Wincoop, E., 1995, Regional Risksharing, *European Economic Review* 37, 1545-1567.

van Wincoop, E., 1996, A Multi-Country Real Business Cycle Model with Heterogeneous Agents, *Scandinavian Journal of Economics* 98, 233-251.

van Wincoop, E., 1998, How Big are the Potential Welfare Gains from International Risksharing?, *Journal of International Economics*, forthcoming.

von Hagen, J., 1992, Fiscal Arrangements in a Monetary Union: Evidence from the US, in Fair, D. and C. de Bossieu (eds.), *Fiscal Policy, taxation, and the Financial System in an Increasingly Integrated Europe*, Kluwer, Boston.

# variables info. set	Standard Deviation Residual Risk (σ_s)			Welfare Gain Horizon T		
	$s = 5$	$s = 15$	$s = 26$	$T = 5$	$T = 15$	$T = 26$
9 ("base info. set")	8.1 (0.4)	9.8 (1.1)	12.8 (1.4)	0.68	1.36	1.60
0	8.4 (0.4)	14.0 (1.4)	15.2 (1.5)	0.75	2.03	2.28
1	8.3 (0.4)	13.3 (1.3)	14.1 (1.4)	0.74	1.82	2.04
2	8.3 (0.4)	12.7 (1.3)	13.6 (1.4)	0.71	1.64	1.85
3	8.2 (0.4)	11.0 (1.1)	13.7 (1.4)	0.70	1.54	1.74

Table 1: Measures of Diversifiable Risk

Notes : The table reports the standard deviation of residual risk for the representative state over a period $[t, t+s]$, for $s = 5, 15, 26$. Standard errors are in brackets. Another reported measure of risk for the representative state is the welfare gain from risksharing, which is equal to the gain from completely eliminating residual growth uncertainty for all years up to an horizon T, with results reported for $T = 5, 15, 26$. The welfare gain represents the potential gain from risksharing for the representative state, measured as the permanent percentage increase in income that yields an equivalent welfare improvement. The table reports results for the "base information set" of 9 variables, for an empty information set (unpredictable deviations from national growth), and for information sets consisting of the 1, 2, and 3 variables leading to the lowest welfare gain for $T = 26$. These variables are, in order of importance, the ratio of government spending to GSP, the lagged one year growth rate, and the immigration rate.

	Percent Welfare Gain Achieved		
	$T = 5$	$T = 15$	$T = 26$
number of variables in info. set:			
0	68.0	70.7	68.3
1	67.5	68.2	65.0
2	68.4	68.9	65.3
3	68.5	67.4	63.8
9 ("base info. set")	69.2	71.8	71.2
"base info. set" plus			
$(log\ GSP)^2$	68.7	72.0	72.0
$log(G/Y)$	70.0	72.1	71.5
$(G/Y)^2$	68.7	71.5	71.1
$(migration)^2$	69.1	71.4	71.2
$(lag\ growth)^2$	68.2	71.3	70.8
$(lag\ growth) * (G/Y)$	68.7	71.9	71.1
$(laggrowth) * (migration)$	69.1	71.4	71.2
$(migration) * (G/Y)$	69.1	71.2	70.3
Δr_t=.01t	69.1	71.6	71.3
Δr_t=-.01t	69.6	71.7	71.3
Deflate by state CPI's	65.1	72.0	70.8
adjustment for retained earnings	76.4	75.5	75.6

Table 2: Sensitivity Analysis

Notes : The table reports the percentage of the total potential welfare gain from risksharing that is achieved through financial markets and federal fiscal policy for horizons of five, fifteen and twenty-six years. The first four rows consider smaller information sets: the "empty" information set and information sets of the one, two and three variables with most predictive power (discussed in section 4.3). The next set of rows add non-linear functions of the variables in the base information set. The following two rows show results from raising and lowering the risk-free interest rate by one percent annually. The next row reports results when the income measures are deflated by state CPI deflators. The final row is based on a different measure of income before risksharing: GSP minus our estimate of retained earnings (see text).

Figure 1 Percent Variance Reduction

Gross State Product to Pre-tax State Income *

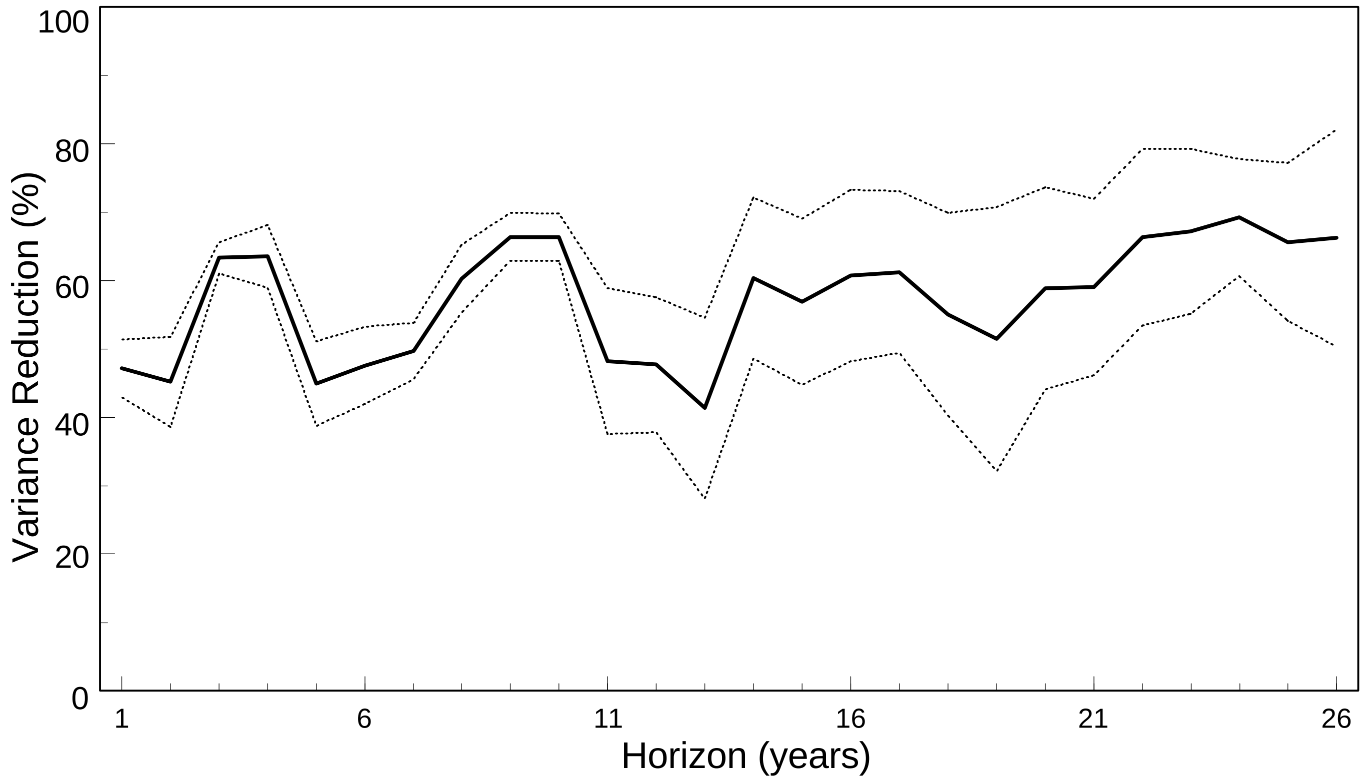

Figure 2 Percent Standard Deviation Reduction

Gross State Product to Pre-tax State Income *

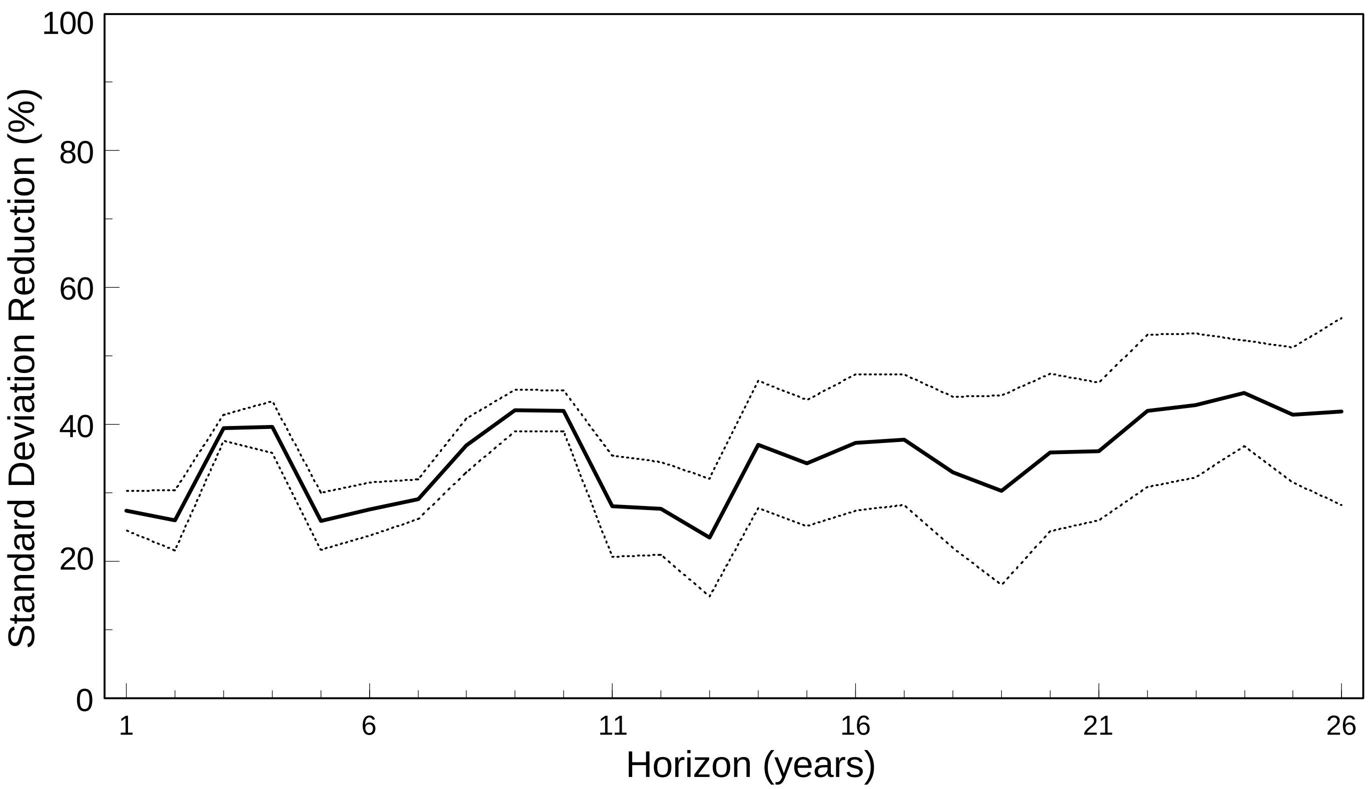

* The outer lines represent a 95% confidence interval. Pre-tax state income is total income by state residents and state and local governments, before federal taxes and transfers.

Figure 3 Welfare Gain from Risksharing *

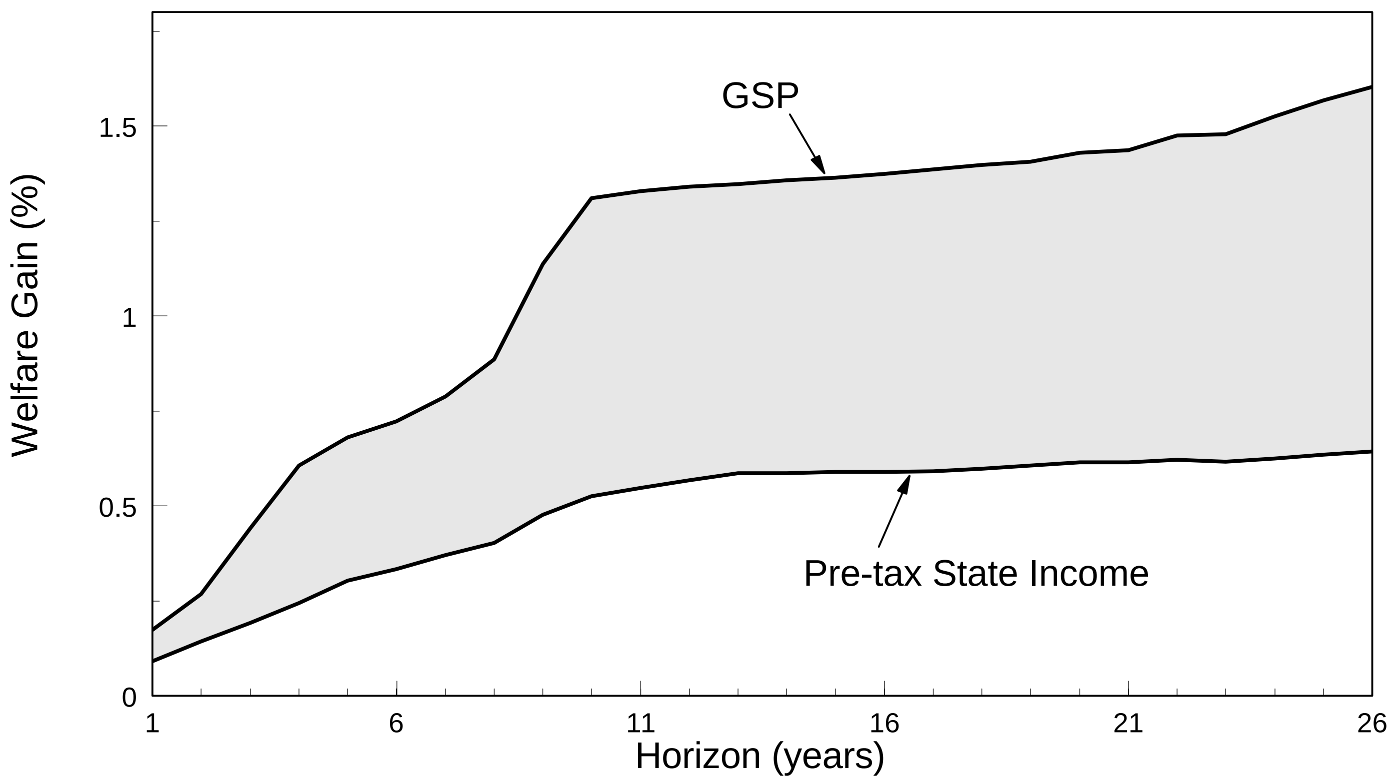

* The welfare gain represents the potential gain from risksharing for a representative state, measured as the permanent percentage increase in income that yields an equivalent welfare improvement. The shaded area represents the component of the potential gain from risksharing with respect to GSP that is achieved through financial markets.

Figure 4 Percent Variance Reduction

Pre-tax State Income to Disposable State Income *

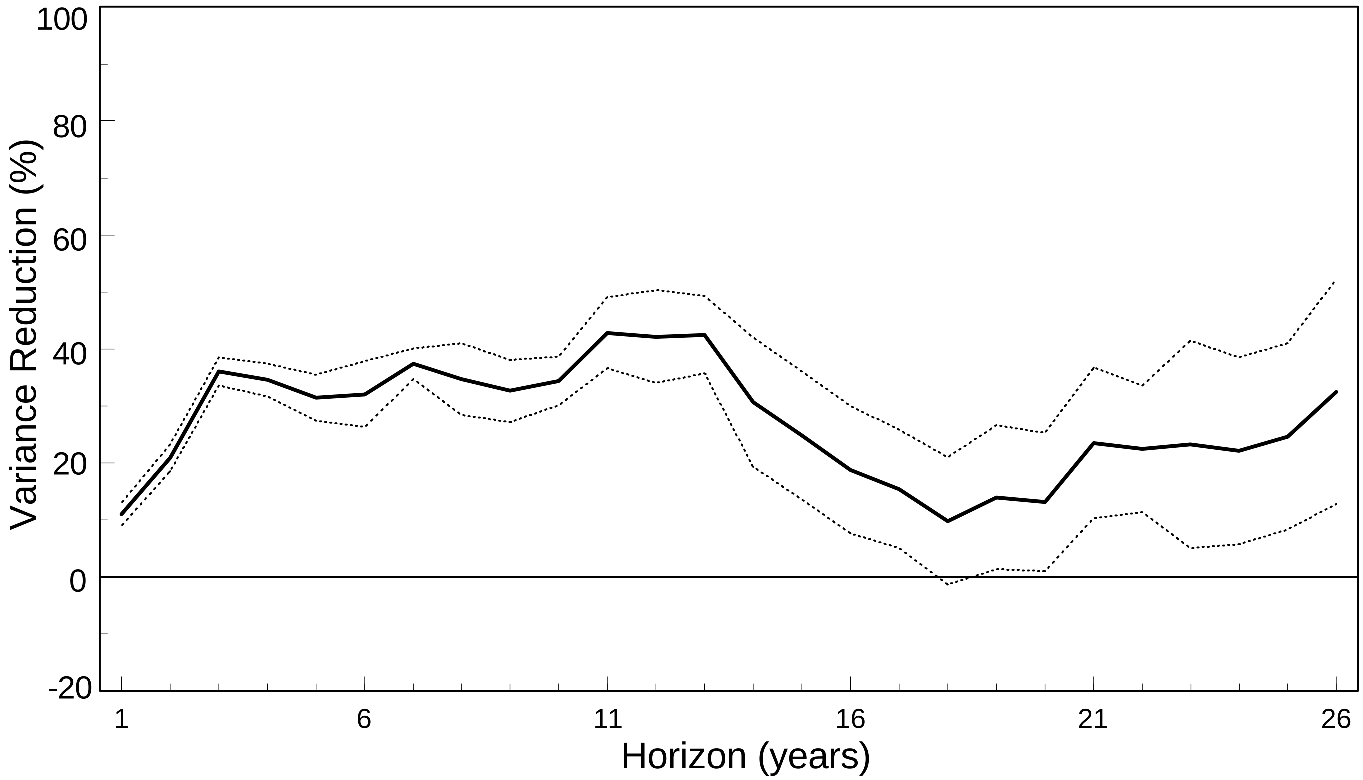

Figure 5 Percent Standard Deviation Reduction

Pre-tax State Income to Disposable State Income *

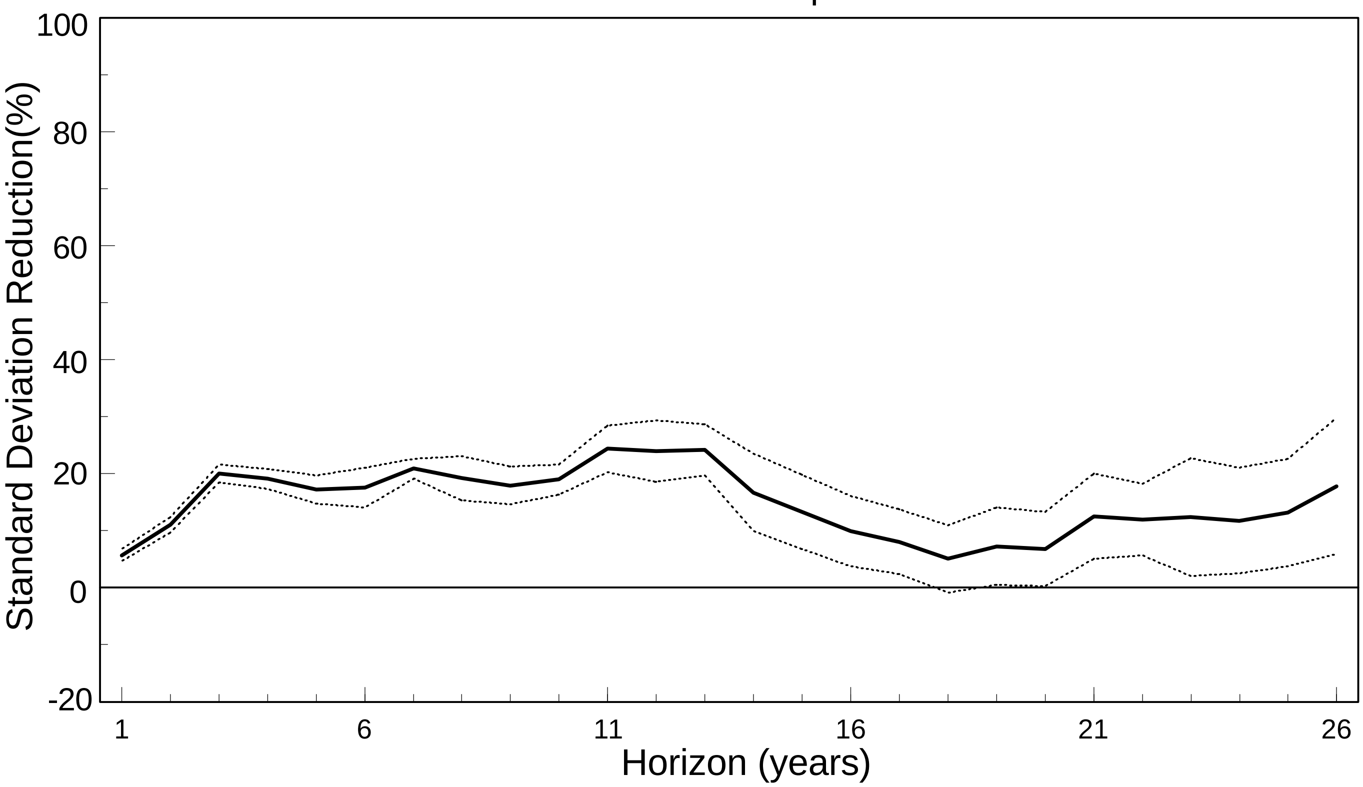

* The outer lines represent a 95% confidence interval. Pre-tax state income is total income by state residents and state and local governments, before federal taxes and transfers. Disposable state income is total income by state residents and state and local governments, after federal taxes and transfers.

Figure 6 Welfare Gain from Risksharing *

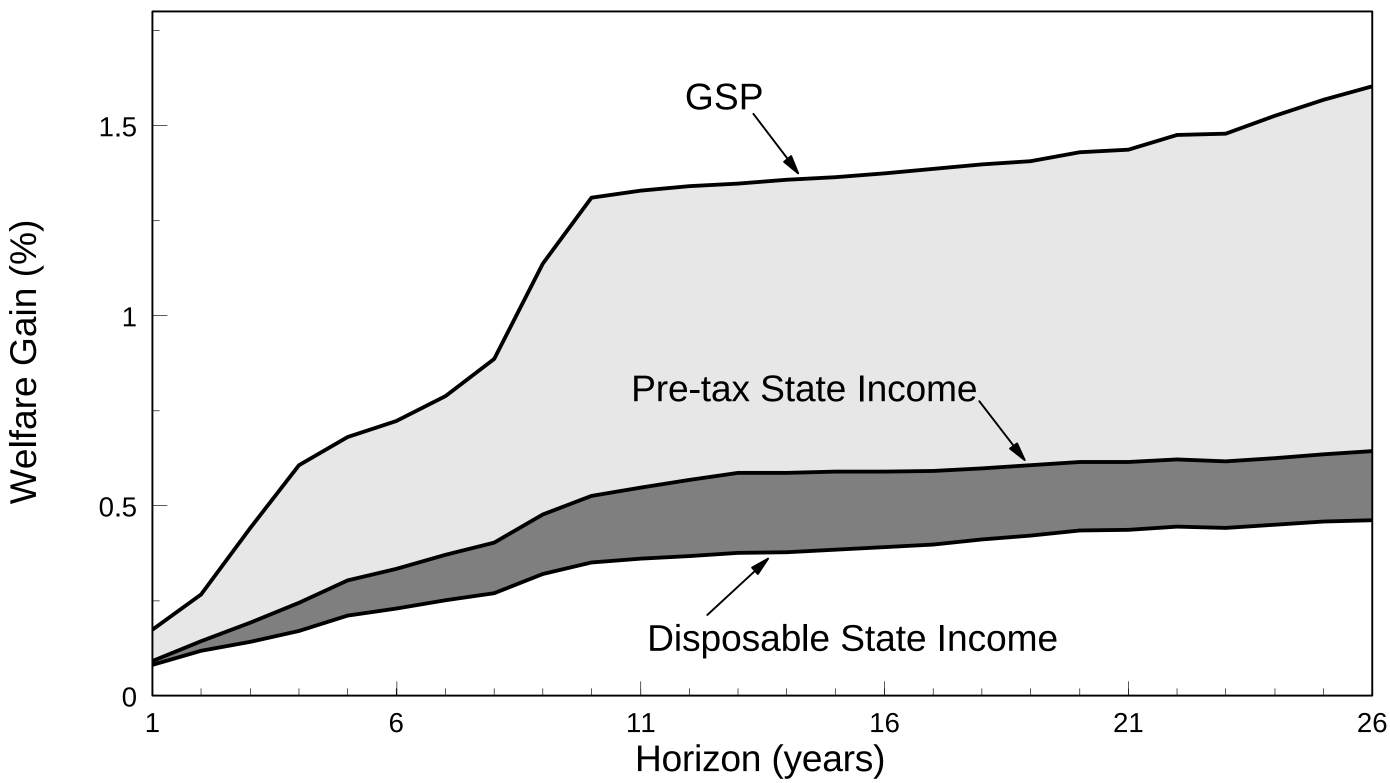

* The welfare gain represents the potential gain from risk-sharing for a representative state, measured as the permanent percentage increase in income that yields an equivalent welfare improvement. The dotted area and the area with slanted lines represent the component of the potential gain from risksharing achieved by financial markets and federal fiscal policy, respectively.

Figure 7 Percent Variance Reduction

Gross State Product to Disposable State Income *

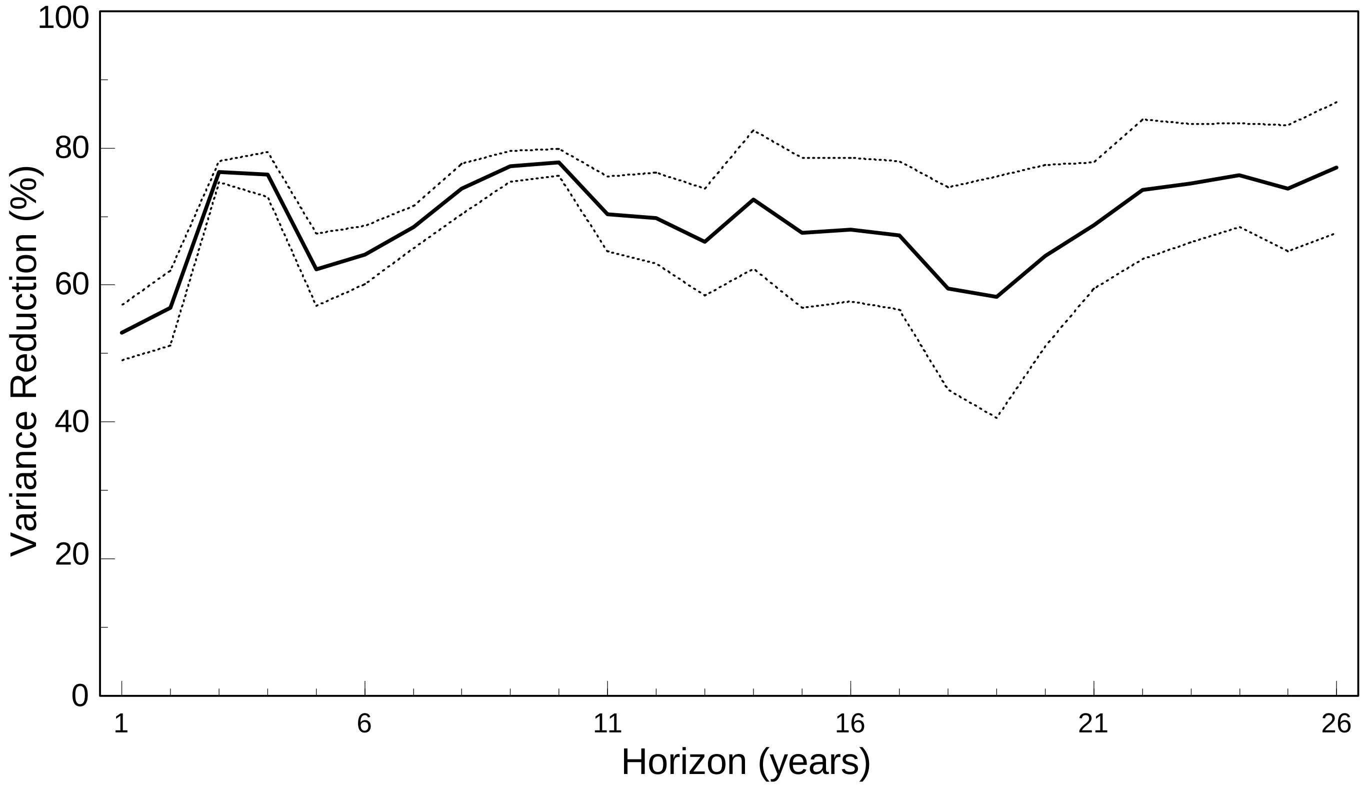

Figure 8 Percent Standard Deviation Reduction

Gross State Product to Disposable State Income *

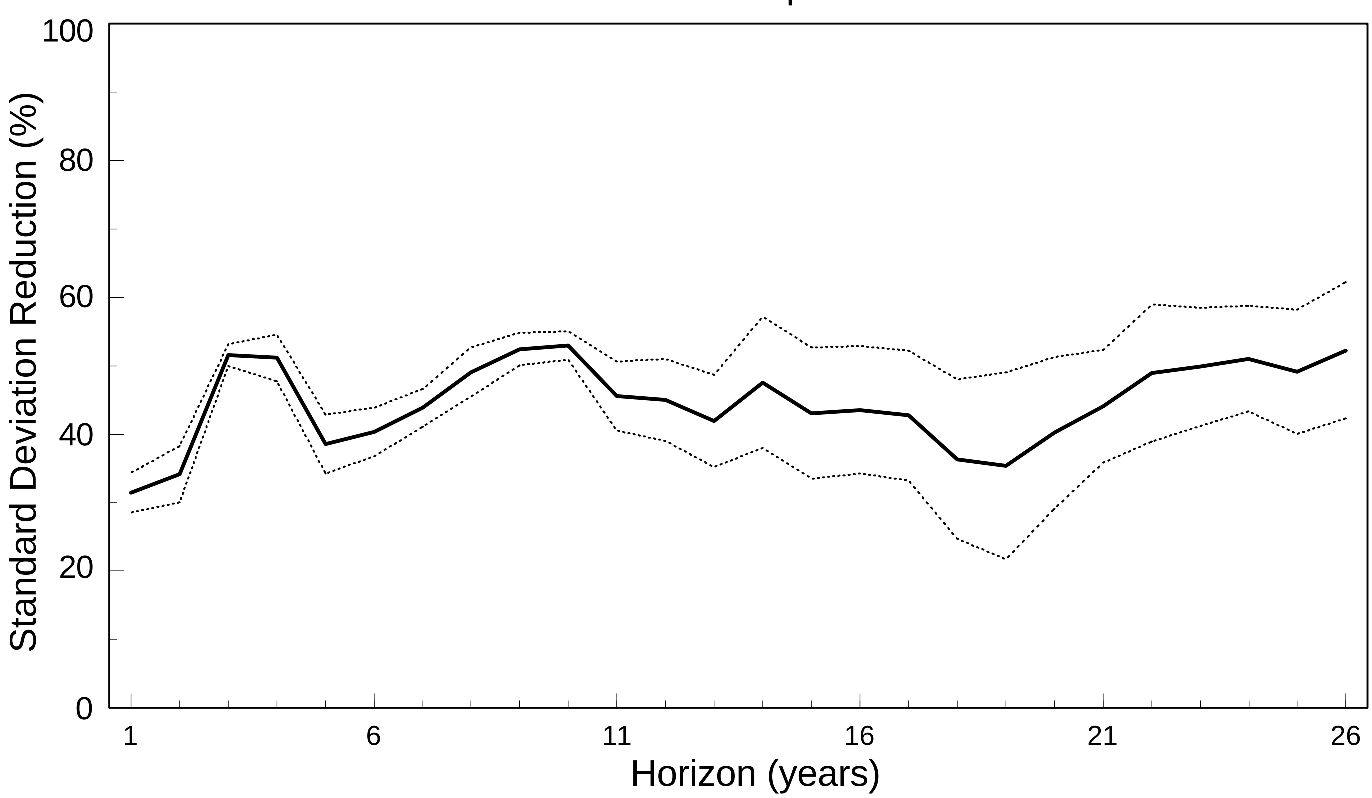

* The outer lines represent a 95% confidence interval. Disposable state income is total income by state residents and state and local governments, after federal taxes and transfers.