Uncertainty in Forecasts of Long-Run Productivity Growth

Peter Christensen, Kenneth Gillingham, and William Nordhaus*

University of Illinois, Urbana-Champaign

Yale University

Yale University

April 15, 2016

Abstract

Long-run forecasts of productivity growth and related uncertainty have become important in policy analysis. However, a thin long-run forecast literature has led to a practice of using ranges of point estimates to partially bound this uncertainty rather than fully parameterizing uncertainty in long-run growth rates. Building on methods established to use short-run forecasts in the growth literature, we use a sample of expert forecasts to study global and regional productivity growth rates (output per capita) for the periods 2010-2050 and 2010-2100. In our sample, the aggregate expectation for average growth rates between 2010-2100 is 2.06% per year, with a standard deviation of 1.12 percentage points. Results from this method indicate substantially higher uncertainty in global growth rates than what is implied by estimates used in existing policy research that are used to project the damages from global climate change. We discuss the strengths and limitations available estimates and implications for policy decisions.

Keywords: economic growth, forecasting, climate change policy

JEL Classification Numbers:

*e-mails: pchrist@illinois.edu, kenneth.gillingham@yale.edu, william.nordhaus@yale.edu, Address: 431 Mumford Hall, Urbana, IL 61801, telephone: 203-506-6724.

The authors are grateful to The National Science Foundation and Department of Energy for generous support as well as to workshop participants at the Stanford Energy Modeling Forum for helpful comments.

1 Introduction

Public policy analysis increasingly relies upon estimates of long-run productivity growth to estimate the net present value of future benefits from decisions being considered today. Applications range from social security (Lee and Tuljapurkar, 1998) to greenhouse gas emissions charges (Nordhaus, 2014; Arrow et al., 2013). Given the uncertainty in the growth process over long horizons, it is typically not sufficient to rely solely on a point estimate of projected growth rates but rather on stochastic trends. Forecasts of long-run productivity growth involve assumptions about unresolved issues such as the production of technologies and ideas (Brynjolfsson and McAfee, 2012; Gordon, 2012), potential reforms to institutions (Acemoglu et al., 2005), and future returns to education (Freeman, 2010). In the absence of a literature that quantifies uncertainty in productivity growth rates over long horizons, researchers have increasingly attempted to bound the uncertainty using ranges of scenario-based estimates.

In this paper, we develop a method for quantifying the uncertainty in long-run growth rates drawn from a set of long-run forecasts. Our method builds on the use of short-run economic forecasts to estimate parameters in growth research (Thomas, 1999). We use data from the Yale Long-Run Growth Survey, which was designed to measure the central tendency and uncertainty in forecasts of productivity growth for major regions during the 21st Century. Respondents provide probabilistic estimates of growth rates for a defined set of quantiles, allowing us to evaluate functional form assumptions and construct probability density functions for each individual forecast as well as their aggregate forecast. We recover estimates for the mean and standard deviation of average annual GDP growth rates for the world and for major regions for the period 2010-2050 and the period 2010-2100. To our knowledge, this study is the first systematic analysis of long-run productivity forecasts and the first to measure uncertainty.

The forecasting literature draws a distinction between consensus and uncertainty in economic prediction (Zarnowitz, 1992). Consensus is defined in terms of the agreement between individual forecasts, whereas statistical forecast uncertainty is defined by within-forecast probabilities. The current practice of using the range of scenario-based estimates as a proxy for uncertainty is similar to the use of measures of consensus in empirical work on uncertainty. Both differ theoretically and empirically from statistical forecast uncertainty and these differences have a bearing on the validity of these measures as proxies for uncertainty in empirical work (Rich and Tracy, 2010; Giordani and Söderlind, 2003; Zarnowitz and Lambros, 1987). We use the estimates in our survey to compare forecast uncertainty and consensus about long-run growth rates, finding that estimates of within-expert uncertainty is systematically greater than the dispersion in point estimates across experts. This finding is consistent with other studies of short-run economic forecasts and has direct implications for future efforts to quantify and model uncertainty in long-run productivity growth (Zarnowitz and Lambros, 1987; Engelberg et al., 2011).

We next highlight the analysis of long-run growth in climate policy research to examine whether the productivity growth scenarios that are used in policy projections accurately represent uncertainty in long-run productivity growth? We compare the statistical forecast uncertainty from individual and combined forecasts with the range of scenario-based estimates that are widely used in the climate change literature. While they are intended to capture the full range of growth paths that could emerge over the course of the 21st Century, our results indicate that scenario-based estimates dramatically understate the uncertainty implicit in the forecasts developed by a sample of the worlds leading experts on economic growth. The existing scenarios do not bound the high or the low end uncertainty in global growth rates. This finding has critical implications, as we have found that assumptions about productivity growth rates play a dominant role in determining projections of GhG emissions growth, long-run climate damages, and the social cost of carbon (Gillingham et al., 2015).

We discuss the benefits and limitations of relying upon using expert forecasts in long-run growth analysis and compare these parameters to prediction sets obtained using available scenario-based estimates and more recent time series methods. Ours is the first study to comprehensively evaluate the different methods that have been used in modeling the long-run growth process, providing a comparison of the different methodologies and parameters obtained from them. Ultimately, we seek to bring together the best available evidence on long-run productivity growth for use in policy research.

The remaining sections of the paper are organized as follows. The following section presents a basic model of expert forecasts. Section 3 discusses the design of the Yale Long Run Growth Survey and presents summary statistics. Section 4 presents the combined forecast distributions of productivity growth rates for 2010-2050 and 2010-2100, as well as the results of our analysis of forecast uncertainty. Section 5 presents an analysis of differences between forecasts uncertainty and the range of scenario-based estimates that are currently utilized in the analysis of climate change policy. Section 6 concludes.

2 Expert Forecasts

Economic forecasts are an important complement to projections developed using available time series methods. To the extent that expert forecasts are informative beyond extrapolation from past realizations of output growth, they can improve forecasts and our estimates of uncertainty. Expert forecasts may also more flexibly account for the non-stationarity of productivity growth over long horizons. Expert expectations have been widely used to develop forecasts of economic growth and short-run growth forecasts have become key inputs in financial research and monetary policy (Thomas, 1999). Prominent surveys include the Livingston Survey, the Bank of Englands Inflation Report, the Blue Chip Survey of Professional Forecasters, annual forecasts conducted by the National Association of Business Economists, the Goldsmith-Nagan Bond Money Market Letter. While each of these forecasts differs in method and objectives, they are all based on subjective accounts of economic growth over specific forecast horizons and are combined to generate measures of aggregate expectations. The rationale motivating the aggregation of expert forecasts is the observation that, under certain conditions, combinations of expert forecasts generally have smaller forecast MSE than individual forecasts (Wallis, 2011; Bates and Granger, 1969).

2.1 Setup

Consider a set of experts making predictions about the growth of an economic variable at time (t) for a particular region (r) and horizon (h). An expert (i) possesses a continuous forecast distribution ϕ_{itrh} that expresses the probability of occurrence over the support of possible long-run rates for region (r) over horizon (h), from which she can report estimates of growth rates at any given quantile θ_p. An experts quantile estimate (θ_{itrhp}) at probability (p) defines the growth rate that an expert believes will be at least equal to the future growth rate with that probability:

$$\theta_p = X_{(Pr(X \leq x) = p)} \tag{1}$$

A set of two or more quantile estimates can be used to measure the uncertainty in the forecast distribution given by expert (i). An experts underlying forecast distribution (ϕ_{itrh}) can be approximated by fitting a continuous distribution $F_x(x)$ to an empirical quantile distribution $F_x(x; \theta_{p_1} ... \theta_{p_k})_{itrh}$ having k quantile estimates.

An estimate of the combined forecast at any given quantile (θ_p) at time (t) for region (r) and horizon (h) for a sample of n experts is given by:

$$\hat{\theta}_{trhp} = \frac{1}{n} \sum_{i=1}^{n} \theta_{i,t,r,h,p} \tag{2}$$

where $\hat{\theta}_{p_1} ... \hat{\theta}_{p_k}$ are k expected quantile estimates in a sample of n experts. An estimate of the combined forecast distribution for a set of experts can be recovered by fitting a continuous distribution $F_x(x)$ to the empirical quantile distribution $F_x(x; \hat{\theta}_{p_1} ... \hat{\theta}_{p_k})_{trh}$.

$$\min \left(F_x(x) - F_x(x; \hat{\theta}_{p_1} ... \hat{\theta}_{p_k})_{trh} \right) \tag{3}$$

If a nonparametric distribution is used to fit the combined quantile function, then the resulting combined forecast distribution $\hat{\phi}_{itrh}$ is a finite mixture distribution (with $w_i = 0$). Measures of the central tendency and uncertainty in the combined forecast distribution $\hat{\phi}_{itrh}$ of long-run productivity growth are highly relevant for policy analysis.

3 Data and Method

3.1 Survey Instrument and Selection of Participants

Our data come from the Yale Long Run Growth Survey, which was administered in 2014-2015. The survey instrument that was designed to measure expectations and uncertainty about average annual per capita GDP. The instrument consists of 4 primary sections on growth rates (details specified in Appendix A). It also contains a warm-up section requiring probabilistic responses to hard-to-recall facts[1] and provides a historical dataset containing information on growth rates for the period (1900-2010) using data from Barro and Ursúa (2008) and Maddison (2003).

The focal section of the survey asks participants to provide 5 quantile estimates (10th, 25th, 50th, 75th, 90th percentile) for average annual rates of growth of per capita real GDP for 6 world regions and 2 time periods (See Appendix, Figure 7)[2]. We utilize results from the focal section of the survey to study uncertainty in expert distributions. The remaining sections ask respondents about the sources and magnitudes of major shocks to the global economy over the period 2050-2100[3] and near-term predictions for per capita growth during

[1]This section was designed to get respondents in the mode of generating probabilistic estimates, using the same percentiles and structure as the survey questions (Ex. "What was the growth rate of per capita real GDP in the US between 1929-2012?")

[2]Surveys of short-run expectations of economic growth rates have sometimes elicited the probabilities associated with a fixed set of rates. However, anchoring respondents on the selected ranges could lead to significant bias in estimates over long horizons. We address this issue by fixing the quantiles of the distribution and asking experts to forecast rates.

[3]Participants are asked to identify a set of uncertain events that have at least a 10% likelihood of occurring between 2010-2100, for the largest positive and negative shock, to estimate the magnitude of its effect relative to their central estimate of annual growth during the period.

the calendar year following the survey. Participants were asked to provide a self-rank of expertise for each question and across the different regions specified in the survey. They were also asked to comment on the clarity of each survey question and to identify any reference material that was used in the development of their calculations.

Participants were selected for participation through a process of nomination by a panel of peers. The criteria for nomination included contributions to the economic growth literature, familiarity with empirical research on medium and long-run growth, and regional expertise (ie low income countries versus OECD). Participants were selected on the basis of the frequency of nomination. Upon selection, they were contacted by email and provided with a link to the digital Qualtrics survey. 13 participants provided complete survey responses and 2 respondents provided a partial sets. For each completed survey, the resulting dataset records estimates at 5 quantiles for each of 6 regions across 2 temporal horizons. This results in a total sample of 780 estimates.

3.2 Survey Data

Table 1 reports summary statistics for the sample of estimates. Descriptive statistics reveal interesting patterns in expert forecasts across horizons and regions. Central estimates indicate that experts expect higher growth rates for the period 2010-2050 than for the period 2010-2100. The average 50th percentile estimates for world growth rates are 2.66% (to 2050) and 2.10% (to 2100). This difference is consistent across the different quantiles. In fact, in 97% of the responses in the sample, an experts quantile estimate to mid-century is higher than the full-century estimate (mid-century estimates are .56 percentage points higher on average). However, the interquartile range of estimates, on average, is higher for full-century estimates than for mid-century estimates.

The survey separately elicits estimates of average annual growth rates for the World, by income group, and for the United States/China. We find that quantile estimates for China and the low income countries are systematically higher than World rates, suggesting that experts expect that output in developing countries will continue to exceed that of developing countries in the 21st Century. Conditional on quantile and horizon, predicted rates for China are 1.13 percentage points higher than for the World. Predictions for the low income countries (including China) are .63 percentage points higher than for World rates. Predictions for the high income countries and the United States, on the other hand, are systematically lower than global predictions, by .73 and .76 percentage points, respectively. Predictions for the middle income group, which we define as having a 30% share of global income, are not statistically distinguishable from predictions of the global rate.

The consensus in expert quantile estimates, as measured by the raw standard deviation of estimates for a particular quantile-region-horizon group, varies considerably. We find the highest level of agreement in central estimates and the greatest dispersion in the tails. This is true in every region and horizon. Interestingly, the consensus is systematically higher at the upper end of the distribution (90th percentile) than at the lower end of the distribution (10th percentile). On average, the dispersion of 10th percentile estimates is twice as high as that of central estimates, whereas the dispersion of 90th percentile estimates are only 20% higher. Overall, we find the greatest consensus in the full century central estimates for high income countries and the least consensus in estimates of Chinas 10th percentile for the period 2010-2100.

4 Results

4.1 The Combined Forecast Distribution

We utilize the full set of estimates to recover measures of the central tendency and uncertainty for the aggregate forecast distribution $\hat{\phi}_{trh}$ of productivity growth rates across regions (r) and horizons (h). In our sample, F(x) is characterized by the quantile distribution: F(x; $\theta_{p=10}$, $\theta_{p=25}$, $\theta_{p=50}$, $\theta_{p=75}$, $\theta_{p=90}$). We begin by estimating $\hat{\theta}_p$ (Equation 2) for each of the (k=5) quantiles.

Table 2 provides the full set of combined quantile estimates $(\hat{\theta}_{p_1}...\hat{\theta}_{p_k})$, by region and horizon. We report estimates of $(\hat{\theta}_p)$ using the mean, trimmed mean, and median of quantile estimates. Our preferred estimates are the trimmed mean due to robustness to outlier forecasts. These robust estimates are, on average, .01 percentage points lower than mean estimates.

We use the set of estimates to construct an empirical CDF for each combined forecast using the method defined in Equation 3. We test the fit of a number of parametric distributions to the combined quantile distribution $F_x(x; \hat{\theta}_{p_1}...\hat{\theta}_{p_k})_{trh}$ and find that the normal distribution provides the best fit to our estimates[4]. Figure 1 plots the best-fitting normal cdf to the set of combined quantile estimates. We also evaluate the fit of each distribution using a Kolmogorov-Smirnov test and find that, in every case, the fitted normal CDF is not statistically different from the combined quantile distribution (results reported in Table 3).

We define the combined forecast distribution as the cumulative distribution function that minimizes the sum of squared differences between any normal cdf $F_x(x; \mu, \sigma)$ and the combined quantile distribution $F_x(x; \hat{\theta}_{p_1}...\hat{\theta}_{p_k})_{trh}$:

$$\min \left(F_x(x; \mu, \sigma) - F_x(x; \hat{\theta}_{p_1}...\hat{\theta}_{p_k})_{trh} \right) \quad (4)$$

Figure 2 plots the respective probability density functions for each experts forecast of global annual productivity growth for the periods 2010-2050 and 2010-2100. Estimates of the parameters $(\hat{\mu},\ \hat{\sigma})$ from the combined forecasts of growth rates are reported in Table 3. The mean forecast estimate implies an average annual growth rate of 2.06% with a standard deviation of 1.23% during the period 2010-2100. Experts expect a somewhat higher rate of 2.59% during the period 2010-2050, with lower uncertainty (1.13%). Our preferred estimates for $(\hat{\mu})$ and $(\hat{\sigma})$ are the trimmed means, which omit the highest and lowest forecast distributions from the sample[5]. This forecast estimate implies an average annual growth

[4] Results available upon request

[5] Note: The min and max distributions correspond to the min and max quantile estimates for all p

rate of 2.06% with a standard deviation of 1.12 percentage points for the period 2010-2100. For the period 2010-2050, the trimmed mean is 2.54% per year, with a standard deviation of 1.07 percentage points. Figure 3 illustrates the projections for world per capita output (2010-2100) that are implied by the aggregate forecast distribution and compares them with data for the 20th century from Barro and Ursúa (2008).

4.2 Uncertainty Versus Consensus in Long-Run Forecasts

Uncertainty has often been parameterized using the (cross-sectional) variance in point estimates within a sample of forecasters. There is evidence, however, that variance in point estimates across individuals does not capture statistical forecast uncertainty (Rich and Tracy, 2010; Giordani and Söderlind, 2003; Zarnowitz and Lambros, 1987). Wallis (2011) decomposes the variance in the combined forecast distribution $\hat{\phi}_{trh}$ into two constituent terms: (1) the sample mean of within-forecast variance (μ_{σ^2}) and (2) the sample variance of individual point forecasts (s^2_μ). The first of these terms measures aggregate forecast uncertainty while the second measures the level of consensus/dispersion in forecasts. The relationships and co-movement between these measures provide important information about the validity of different measures of uncertainty in empirical work[6].

Table 4 reports the results of this comparison for long-run growth estimates, comparing the level of forecast uncertainty and forecast consensus by region and horizon. We find that the level of within-forecast uncertainty exceeds the level of between-forecast dispersion for every region and horizon in the sample. Therefore, according to this sample of experts, an interpretation of the dispersion of point estimates as a measure of uncertainty in policy applications would dramatically understate the statistical forecast uncertainty (by 25.5% on average). This difference is particularly stark for the middle income countries (where the level of consensus is high) and for low income countries/China (where uncertainty is high).

[6]Prior work has examined consensus and uncertainty by comparing the standard deviation of sample of point estimates (consensus) to the sample average of within-forecast standard deviations (uncertainty), which are the measures that we use here.

While the dispersion of forecasts clearly understates the level of uncertainty in individual forecasts in long-run growth estimates, we find that levels of consensus and uncertainty are highly correlated across region and horizon.

We also evaluate consensus by measuring the frequency with which one experts forecast falls outside the 90% confidence interval of another forecast. In the case of a 50th percentile estimate, this measure can be interpreted as the rate at which one experts central forecast is rejected by another expert. We report these results in Table 5. We find that central predictions made by one forecaster are rejected 10% of the time in the case of mid-century forecasts and 13% of the time in the case of full-century forecasts. Rejection rates are sensitive to the presence of a single extreme (or a highly confident) forecast distribution that rejects everyone. When we trim the minimum and maximum forecast distributions from the sample, the rejection rates drop to 4% for mid-century and 8% for full-century estimates. Predictions for the United States and the high income countries have higher rejection rates, while the middle income and low income countries have lower rates.

5 Productivity Growth Forecasts for the 21st Century: Implications for Climate Change Research

The analysis of climate change policy is particularly dependent upon assumptions about growth rates across the 21st Century due to long time scales across greenhouse gas emissions are linked to economic damages. Modeling efforts in that literature have relied heavily upon estimates of productivity growth rates from 5 subjectively-determined scenarios referred to as the Shared Socioeconomic Pathways (SSPs), which describe demographic and economic trends and provide 100-year forecasts for those parameters (ONeill et al., 2014, 2015). Each of the 5 scenarios includes an estimate of national per capita GDP at 5-year intervals between 2010-2100.

We use the panel of SSP estimates to compute implied annual per capita growth rates that match the definitions described in the Long Run Growth Survey[7] and then compare them to survey forecast distributions. Figure 3 compares the projected growth of global per capita output from the entire range of SSP scenarios to the statistical forecast uncertainty expressed in the combined distribution from expert forecasts. The figure reveals two key issues. First, the differences between SSP scenarios and expert forecasts have a tremendous impact on the projected level of per capita output at the end of the century. Second, despite the fact that the set of SSP projections is intended to capture the full range of possible trajectories of output over the course of the century, the range of SSPs dramatically understates the uncertainty expressed within an average expert's forecast.

These differences are further demonstrated by rejection rates for SSPs, which are presented in Table 6. Mid-century SSP growth rates are rejected by 31% of expert forecast distributions (2 scenarios with the highest growth rates) and by 38% of experts (3 remaining scenarios). Full-century scenarios fall outside the 90th percentile of expert forecast distributions 15% of the time for the 2 most optimistic scenarios and 38% of the time for the other 3. In every case, rejected scenarios fall below the 10th percentile of the associated expert forecast. These results indicate important discrepancies between SSP scenarios and the panel of economic experts. Economic experts are systematically more optimistic than scenario developers in their productivity growth estimates.

6 Conclusion

Public policy research in multiple domains increasingly relies upon forecasts of productivity growth and often explicitly models uncertainty in long-run growth rates. However, reliable estimates of uncertainty in growth rates over the course of the 21st century have not been available in the past. In the absence of this information, researchers have relied upon ranges of forecasts and scenario-based estimates to bound forecast uncertainty. We build upon the short-run forecasting literature to develop and test a set of forecast

[7] See Appendix A for definitions

distributions for long-run growth.

We study long-run productivity growth forecasts using data from experts in the Yale Long Run Growth Survey. To our knowledge, this study provides the most comprehensive analysis of long-run productivity growth forecasts available. Our results indicate that experts generally believe that productivity growth will be higher during the first half of this century, while expressing greater uncertainty in their full-century estimates. Experts tend to expect higher long-run growth rates in China and in low income countries and lower growth rates in higher income countries. We find greater agreement in central estimates than in the extremes of the density distribution – the lower end of the distribution is the area of greatest disagreement in forecasts.

We generate an aggregate forecast distribution for use in policy and uncertainty analysis. For world productivity growth, the aggregate expectation for average annual rates is 2.06% per year between 2010-2100, with a standard deviation of 1.12 percentage points. For the period 2010-2050, the expectation is 2.54% per year, with a standard deviation of 1.07 percentage points. These estimates reveal far greater uncertainty in expert forecast distributions than is commonly implied in policy research, which must be taken very seriously in the climate change literature as uncertainty in productivity growth has a dominant effect on projections of global temperature change, economic damages, and the social cost of carbon (Gillingham et al., 2015).

We demonstrate the importance of these differences using comparison with scenario-based estimates from climate change policy, which significantly understate the uncertainty in individual and aggregate forecasts produced by this sample of economic experts. Furthermore, the SSP scenarios fall outside the 90% confidence intervals given by roughly 1/3 of economic experts in our sample, revealing inconsistencies in assumptions about central tendency as well as uncertainty in long-run productivity growth. We view these results as a step toward resolving critical issue in policy research and advocate the development of a more robust set of estimates that can further guide public decisions.

Figures and Tables

Figure 1: Fit of Combined CDF ($\hat{\phi}_{rh}$) to Combined Quantile Distribution ($(F_x(x; \hat{\theta}_{p_1} ... \hat{\theta}_{p_k})_{rh})$)

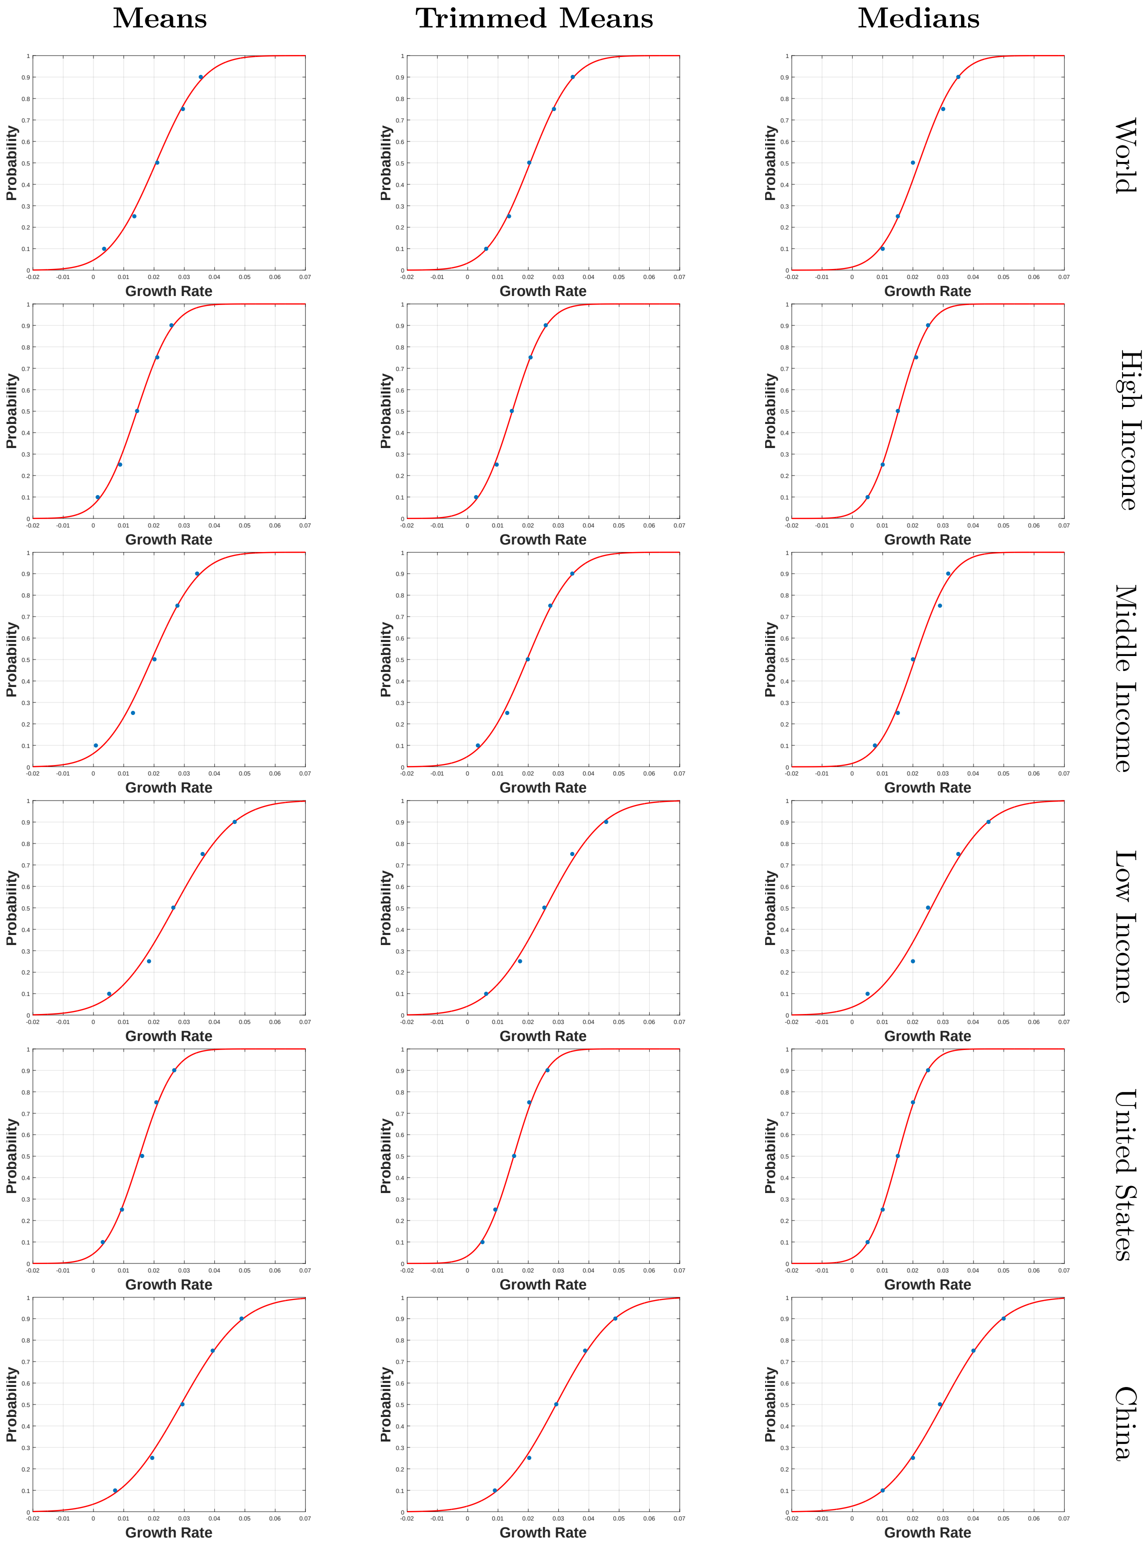

Note: Blue points are combined quantile distribution F(x; $\hat{\theta}_{p=10}$, $\hat{\theta}_{p=25}$, $\hat{\theta}_{p=50}$, $\hat{\theta}_{p=75}$, $\hat{\theta}_{p=90}$) for each quantile and region using mean, trimmed mean, and median of expert estimates. Red lines are the best-fitting normal CDF.

Figure 2: Combined Forecast Distributions ($\hat{\phi}_h$) for Global Growth Rates

Period: 2010-2050

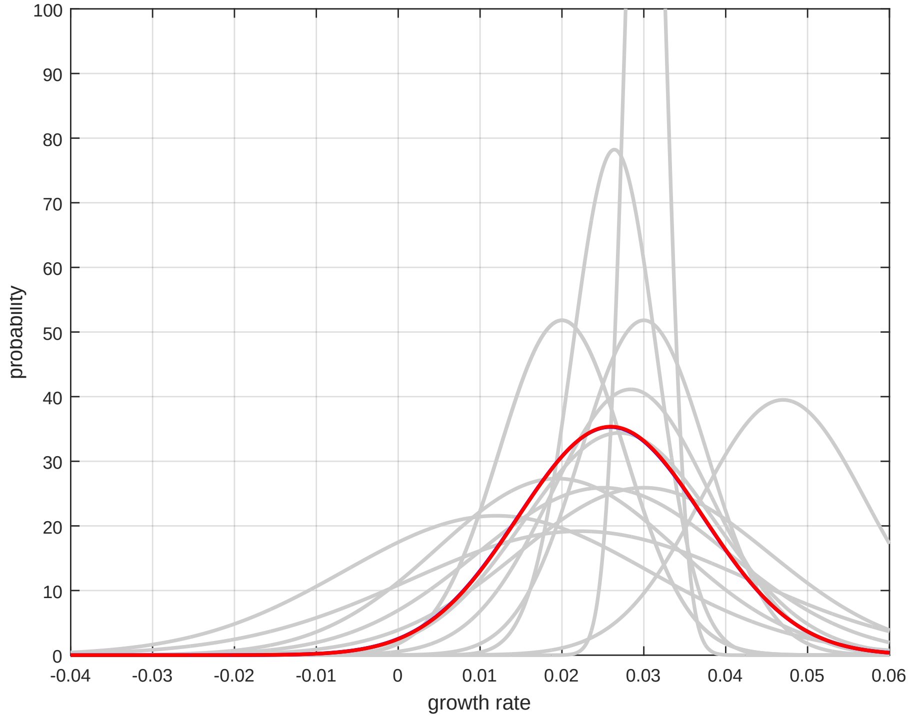

Period: 2010-2100

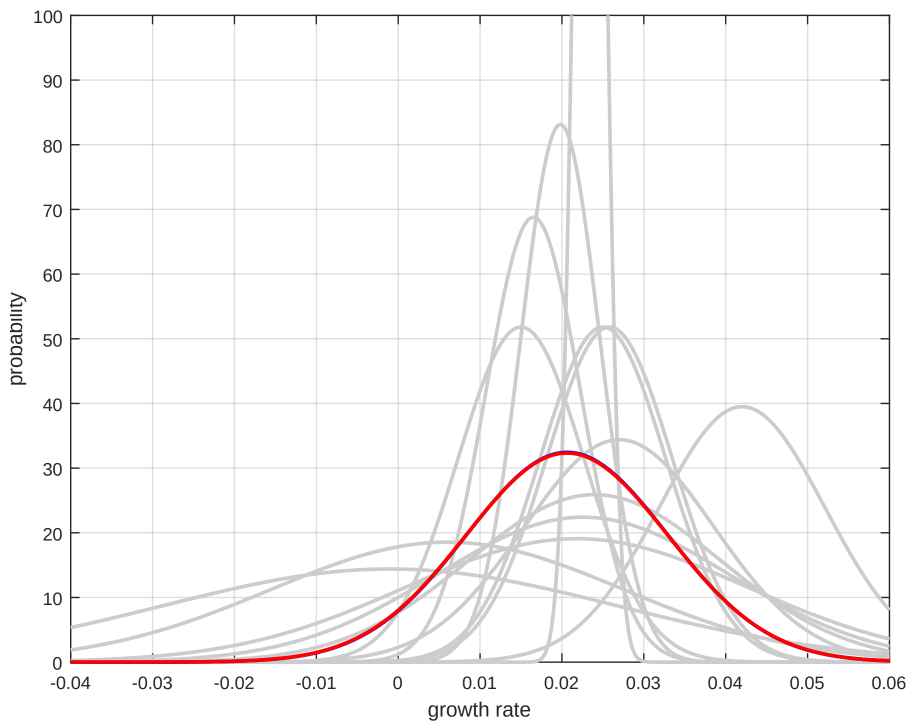

Note: Combined forecast distributions $\hat{\phi}_h$ (red) for global rates based on best-fitting normal pdf using trimmed mean of quantile estimates of expert forecasts. Best-fitting normal pdf of individual forecast distributions shown in grey.

Figure 3: Projections of Annual Global Output to 2100 with Uncertainty

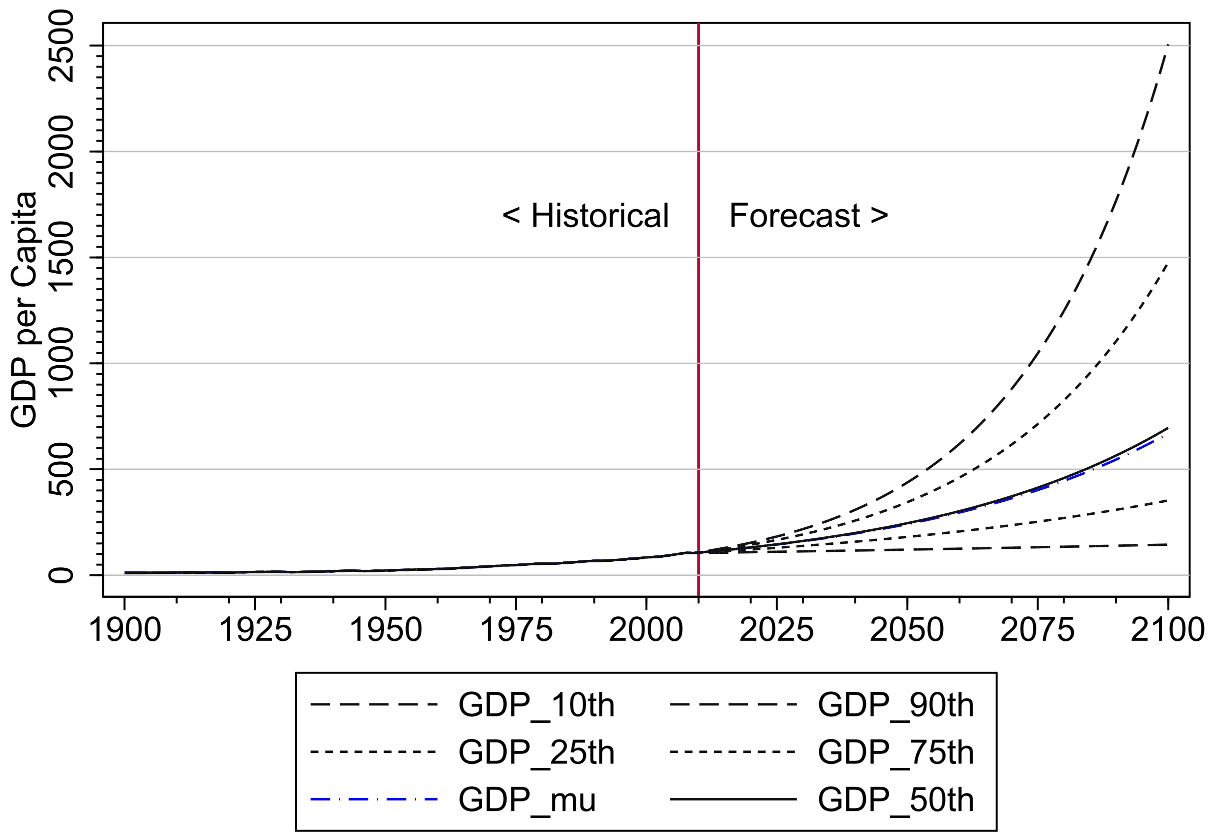

Note: Historical per capita GDP (1900-2010) use data from Barro and Ursua (2008) and indexed by nominal values in year 2006 (year 2006=100). Black dashed lines illustrate per capita growth projection using combined quantile distribution F(x; $\hat{\theta}_{p=10}$, $\hat{\theta}_{p=25}$, $\hat{\theta}_{p=50}$, $\hat{\theta}_{p=75}$, $\hat{\theta}_{p=90}$) for global rates (2010-2100). Blue dashed line illustrates per capita growth projection using $\hat{\mu}$ from best-fitting normal distribution.

Figure 4: Comparison of Projections of Annual Global Output to 2100 with SSP Scenarios

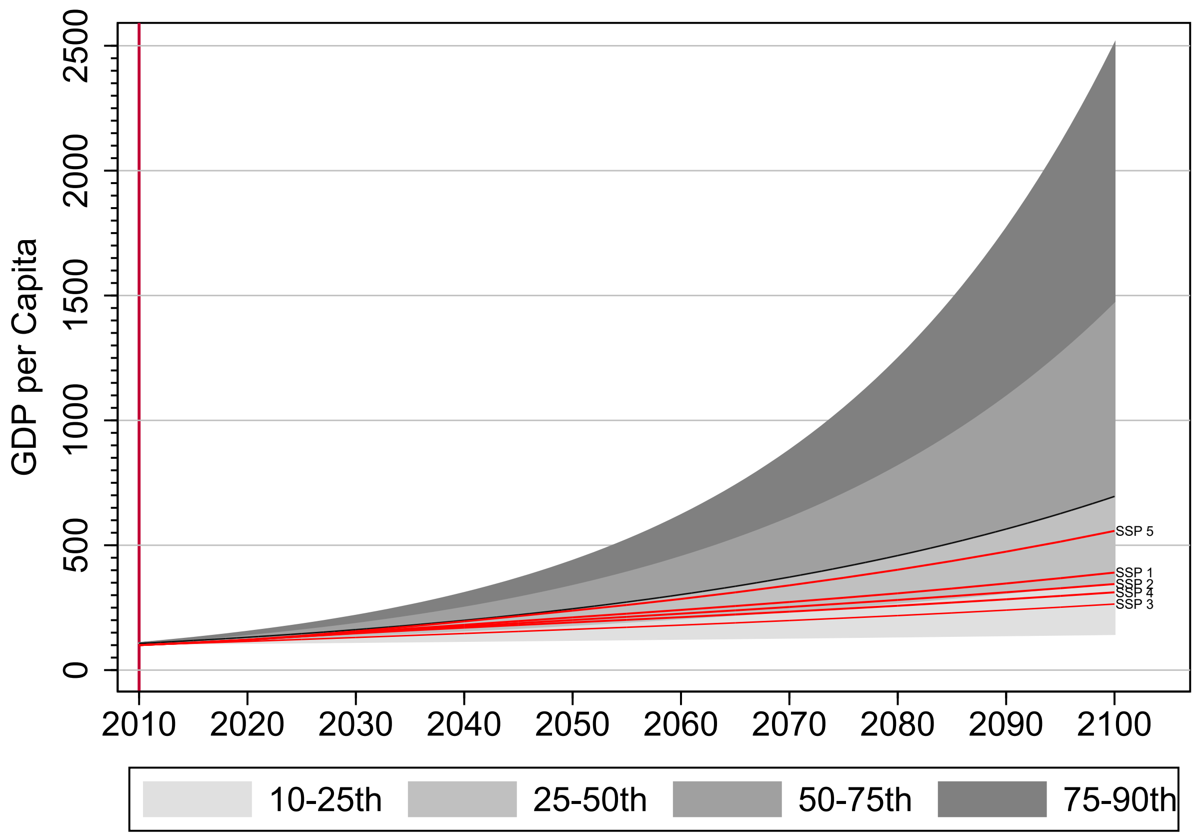

Note: Grey shaded areas illustrate per capita growth projection using combined quantile distribution F(x; $\hat{\theta}_{p=10}$, $\hat{\theta}_{p=25}$, $\hat{\theta}_{p=50}$, $\hat{\theta}_{p=75}$, $\hat{\theta}_{p=90}$) for global rates (2010-2100). Red lines illustrate per capita growth projections using SSP projections of global rates. Global rates are geometric means of national growth rates, weighted by share of global income in 2006.

Table 1: Summary Statistics: Quantile Estimates by Region and Horizon

Region	Statistic	2010-2050					2010-2100				
		10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
World	Mean	1.10	1.86	2.66	3.35	4.01	0.36	1.34	2.10	2.95	3.55
	StDev	1.37	0.97	0.75	0.85	0.92	2.14	1.14	0.84	0.94	1.06
	N	13	13	13	13	13	13	13	13	13	13
High Income	Mean	0.43	1.24	1.76	2.30	2.74	0.14	0.88	1.44	2.11	2.58
	StDev	1.38	0.82	0.68	0.69	0.77	1.55	0.92	0.62	0.73	0.84
	N	13	13	13	13	13	13	13	13	13	13
Middle Income	Mean	0.78	1.73	2.66	3.38	4.13	0.08	1.30	2.02	2.77	3.42
	StDev	1.47	0.91	0.77	0.68	0.77	2.15	0.83	0.81	0.64	0.97
	N	13	13	13	13	13	13	13	13	13	13
Low Income	Mean	1.04	2.27	3.50	4.33	5.18	0.52	1.84	2.64	3.61	4.66
	StDev	1.70	1.25	0.78	0.95	1.26	2.10	1.23	1.10	1.05	1.55
	N	13	13	13	13	13	13	13	13	13	13
United States	Mean	0.53	1.19	1.82	2.23	2.65	0.31	0.94	1.60	2.08	2.66
	StDev	1.18	0.76	0.75	0.69	0.68	1.28	0.76	0.76	0.65	0.78
	N	13	13	13	13	13	13	13	13	13	13
China	Mean	1.51	2.84	4.27	5.24	6.38	0.72	1.94	2.94	3.93	4.89
	StDev	1.83	1.57	1.11	1.18	1.35	2.29	1.59	1.25	1.15	1.38
	N	13	13	13	13	13	13	13	13	13	13

Note: Summary statistics of estimates reported on the Yale Long-Run Growth Survey. See Appendix A for survey response format.

Table 2: Estimates of $\hat{\theta}_p$ by Percentile, Region, and Horizon

Region	Statistic	2010-2050					2010-2100				
		10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Mean		1.10	1.86	2.66	3.35	4.01	0.36	1.34	2.10	2.95	3.55
Trimmed Mean	World	1.17	1.80	2.59	3.23	3.92	0.60	1.36	2.03	2.85	3.47
Median		1.00	2.00	2.80	3.50	4.00	1.00	1.50	2.00	3.00	3.50
Mean		0.43	1.24	1.76	2.30	2.74	0.14	0.88	1.44	2.11	2.58
Trimmed Mean	High Income	0.56	1.23	1.76	2.30	2.75	0.27	0.95	1.46	2.08	2.57
Median		0.50	1.50	1.90	2.50	2.90	0.50	1.00	1.50	2.10	2.50
Mean		0.78	1.73	2.66	3.38	4.13	0.08	1.30	2.02	2.77	3.42
Trimmed Mean	Middle Income	0.93	1.76	2.67	3.36	4.11	0.34	1.30	1.98	2.72	3.45
Median		1.00	1.50	2.70	3.50	4.00	0.75	1.50	2.00	2.90	3.16
Mean		1.04	2.27	3.50	4.33	5.18	0.52	1.84	2.64	3.61	4.66
Trimmed Mean	Low Income	1.05	2.23	3.41	4.25	5.12	0.62	1.72	2.53	3.45	4.57
Median		1.00	2.00	3.27	4.00	5.00	0.50	2.00	2.50	3.50	4.50
Mean		0.53	1.19	1.82	2.23	2.65	0.31	0.94	1.60	2.08	2.66
Trimmed Mean	United States	0.60	1.14	1.75	2.18	2.63	0.49	0.91	1.53	2.04	2.64
Median		0.50	1.00	1.50	2.10	2.50	0.50	1.00	1.50	2.00	2.50
Mean		1.51	2.84	4.27	5.24	6.38	0.72	1.94	2.94	3.93	4.89
Trimmed Mean	China	1.51	2.81	4.23	5.19	6.31	0.89	2.02	2.93	3.87	4.87
Median		1.00	3.00	4.20	5.00	6.00	1.00	2.00	2.90	4.00	5.00

Note: Estimates based on responses reported in the Yale Long-Run Growth Survey. See Appendix A for survey response format.

Table 3: Estimates of $F(x;\hat{\mu},\hat{\sigma})$ and Kolmogorov-Smirnov Test

Region	Statistic	2010-2050				2010-2100			
		$\hat{\mu}$	$\hat{\sigma}$	K-S	K-S(p)	$\hat{\mu}$	$\hat{\sigma}$	K-S	K-S(p)
World	Mean	2.59	1.13	.149	.981	2.06	1.23	.165	.975
	Trimmed Mean	2.54	1.07	.158	.978	2.06	1.12	.160	.977
	Median	2.66	1.16	.166	.975	2.20	1.01	.186	.967
High Income	Mean	1.69	0.88	.156	.979	1.43	0.94	.163	.976
	Trimmed Mean	1.72	0.84	.155	.979	1.47	0.88	.155	.979
	Median	1.86	0.89	.164	.975	1.52	0.79	.169	.974
Middle Income	Mean	2.54	1.29	.143	.983	1.92	1.25	.150	.981
	Trimmed Mean	2.57	1.23	.144	.983	1.96	1.18	.141	.984
	Median	2.54	1.24	.199	.962	2.06	0.96	.209	.946
Low Income	Mean	3.26	1.60	.160	.977	2.66	1.55	.130	.988
	Trimmed Mean	3.21	1.57	.151	.981	2.58	1.49	.120	.992
	Median	3.05	1.54	.157	.978	2.60	1.46	.141	.984
United States	Mean	1.68	0.82	.168	.974	1.52	0.90	.141	.984
	Trimmed Mean	1.66	0.79	.145	.983	1.52	0.84	.167	.974
	Median	1.52	0.79	.169	.974	1.50	0.77	.142	.984
China	Mean	4.05	1.87	.147	.982	2.89	1.60	.143	.984
	Trimmed Mean	4.01	1.85	.147	.982	2.92	1.51	.136	.986
	Median	3.84	1.85	.177	.970	2.98	1.54	.146	.982

Note: Estimates of μ and σ reflect the best-fitting normal distribution using mean, trimmed mean, and median of expert estimates. Results of Kolmogorov-Smirnov test given by K-S (which is KS statistic) and p-value of K-S test. All K-S tests fail to reject (with 95% confidence) the null hypothesis that combined quantile distributions and best-fitting normal distributions are the same.

Table 4: Forecast Uncertainty and Dispersion of Forecasts

Region	2010-2050			2010-2100		
	Uncertainty	Dispersion		Uncertainty	Dispersion	
	mean(σ)	stdev(μ)	stdev(θ_{50})	mean(σ)	stdev(μ)	stdev(θ_{50})
World	1.13	0.83	0.75	1.23	1.04	0.84
High Income	0.88	0.76	0.68	0.94	0.79	0.62
Middle Income	1.29	0.76	0.77	1.25	0.76	0.81
Low Income	1.59	0.95	0.78	1.55	1.10	1.10
United States	0.82	0.75	0.75	0.90	0.70	0.76
China	1.87	1.22	1.11	1.59	1.29	1.25

Note: Forecast uncertainty is measured as mean of expert standard deviations (σ). Dispersion measures (lack of) consensus using the standard deviation of expert means (μ) and the standard deviation of 50th percentile estimates (θ_{50}).

Table 5: Rejection Rates for Expert Forecast Distributions by Region and Horizon

Region	2010-2050			2010-2100		
	50^{th}	75^{th}	90^{th}	50^{th}	75^{th}	90^{th}
World	10%	5%	3%	13%	8%	3%
High Income	18%	12%	4%	16%	9%	2%
Middle Income	6%	4%	0%	12%	8%	1%
Low Income	10%	8%	2%	12%	7%	3%
United States	19%	12%	4%	15%	8%	1%
China	10%	3%	1%	14%	4%	3%

Note: Rejection rates measure the frequency with which one expert's quantile estimate falls outside a different expert's 90% confidence interval.

Table 6: Rejection Rates for Global SSP Scenarios by Horizon

Region		2010-2050					2010-2100				
		SSP1	SSP2	SSP3	SSP4	SSP5	SSP1	SSP2	SSP3	SSP4	SSP5
World	<10 CI	31%	38%	38%	38%	31%	15%	38%	38%	38%	15%
World	>90 CI	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%

Note: Rejection rates measure the frequency with which an SSP scenario falls outside an expert's 90% CI.

Appendix

Appendix A Survey Design

Survey Front Matter:

The font matter in the Yale Long Run Growth survey describes the purpose, protocol, definitions used in the survey. It also provides a link to historical data. We include some of the text from the survey below for clarity.

Purpose The purpose of this survey is to improve our understanding about future economic growth – both the central trends and the uncertainties about those trends. The survey will cover both global growth rates, as well as regional growth rates, over the period 2010 to 2100. The variables chosen correspond to assumptions in major economic integrated assessment models (IAMs) used for climate change policy analysis. There have been no formal surveys of experts on this subject, and we are grateful for your time and effort.

Survey Outline

I. Common problems in expert elicitations
II. Definitions and historical data
III. Survey protocol
IV. Three brief warm-up questions
V. Five substantive questions
VI. Short term predictions
VII. Brief discussion of your reasoning

Definitions
In this survey, output for different regions will be measured using purchasing power parity. That is, output levels are converted using exchange rates that reflect estimates of the relative costs of living in different countries. We will ask about the average annual rate of growth of per capita real GDP in each region and time period. Countries are grouped into high, middle, and low income categories according to statistics for the year 2007 (preceding the recession). Country groupings are assumed constant across time: i.e., the low-income group in 2100 contains the exact same countries as the low-income group in 2007. The historical data and concepts used here are per capita GDP estimates (PPP) from the International Monetary Fund. Note that the income categories defined here (particularly low income) differ from those used elsewhere to increase the balance of world GDP shares across groups.

High Income countries are defined as having a per capita GDP of greater than \$34,000 based on purchasing power parity in 2007 the three countries at the low end of this range are Japan, France, and Germany.

Middle Income countries are defined as having a GDP of \$7,000 - \$34,000 per capita (PPP) in 2007 the three countries at the high end of this range are Taiwan, Italy, and Spain; and the three countries at the low end of this range are Dominican Republic, Bosnia and Herzegovina, and Peru.

Low Income countries are defined as having a GDP of less than $7,000 per capita (PPP) in 2007. The three countries at the high end of this range are Ukraine, El Salvador, and Algeria. For the purposes of this survey, China and India are both considered low income countries.

Historical Data: Using the country-level series constructed by Barro and Ursúa (2008) and Maddison (2003) for the period 1900-2000, we provide historical growth rates and their standard deviations for the World, United States, China as well as the Low, Middle and High Income groups defined below.

Figure 5: Historical Data: Annual Growth Rates By Region

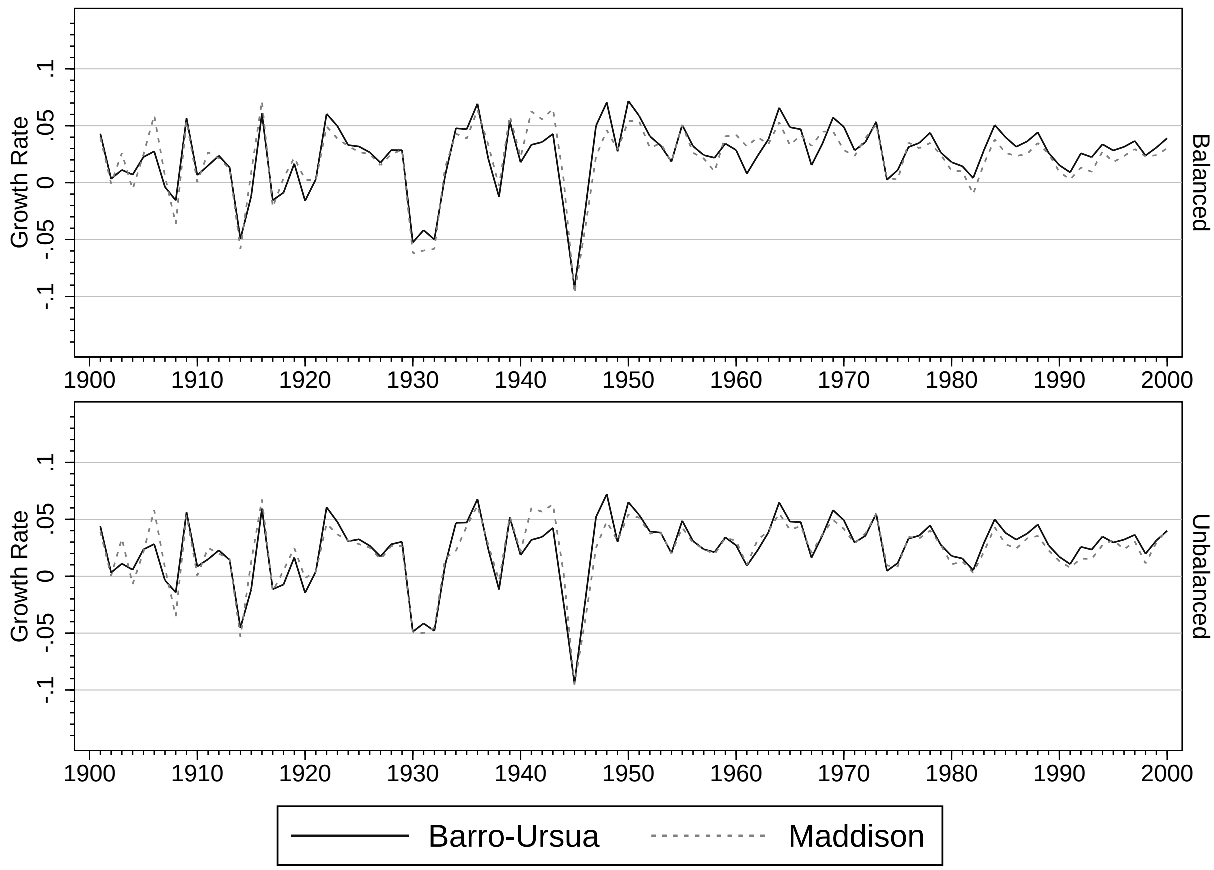

Note: Time series graph is representative of sources provided to experts. Sources: Barro and Ursúa (2008) and Maddison (2003).

Figure 6: Historical Data: Variability in Annual Growth Rates

Table 4: Barro-Ursua Balanced Panel

	World					
	Growth Rate	SD (100,1)	SD (25,1)	SD (25,4)	N	Income Share
1900-1925	0.014	0.0282	0.0275	0.0120	35	0.81
1925-1950	0.015	0.0282	0.0424	0.0120	35	0.81
1950-1975	0.034	0.0282	0.0165	0.0120	35	0.81
1975-2000	0.029	0.0282	0.0110	0.0120	35	0.81

	High Income Group					
	Growth Rate	SD (100,1)	SD (25,1)	SD (25,4)	N	Income Share
1900-1925	0.015	0.0364	0.0369	0.00679	16	0.97
1925-1950	0.020	0.0364	0.0579	0.00679	16	0.97
1950-1975	0.034	0.0364	0.0195	0.00679	16	0.97
1975-2000	0.022	0.0364	0.0135	0.00679	16	0.97

	Middle Income Group					
	Growth Rate	SD (100,1)	SD (25,1)	SD (25,4)	N	Income Share
1900-1925	0.017	0.0298	0.0339	0.00943	14	0.60
1925-1950	0.020	0.0298	0.0414	0.00943	14	0.60
1950-1975	0.039	0.0298	0.0127	0.00943	14	0.60
1975-2000	0.016	0.0298	0.0168	0.00943	14	0.60

	Low Income Group					
	Growth Rate	SD (100,1)	SD (25,1)	SD (25,4)	N	Income Share
1900-1925	0.0071	0.0492	0.0379	0.0253	5	0.77
1925-1950	-0.0057	0.0492	0.0539	0.0253	5	0.77
1950-1975	0.029	0.0492	0.0493	0.0253	5	0.77
1975-2000	0.061	0.0492	0.0257	0.0253	5	0.77

Notes:

(i) (B-U) Barro-Ursua Macroeconomic data available at: rbarro.com/data-sets/

(ii) SD (100,1): standard deviation of 100, 1-year growth rates by income group

(iii) SD (25,1): standard deviation of 25, 1-year growth rates by period and income group

(iv) SD (4,25): standard deviation of 4, 25-year average growth rates by income group

Note: Table of summary statistics is representative of sources provided to experts. Sources: Barro and Ursúa (2008) and Maddison (2003).

Response Format:

Survey participants respond to a digital survey that takes the following form.

Q1. Please provide your central (50th percentile) estimate of the average annual rates of growth of per capita real GDP (percent per year) for the following regions and time periods.

Figure 7: RESPONSE FORMAT: UNCERTAINTY IN GROWTH RATES

	2010-2050					2010-2100				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
World										
United States										
China										
High Income (Includes US)										
Middle Income										
Low Income (includes China and India)										

Appendix B Robustness of Aggregation Method

We test the robustness of our estimates of $\hat{\mu}$ and $\hat{\sigma}$ to an alternate specification.

Alternate Method: Consider a set of experts making predictions about the growth of an economic variable at time (t) for a particular region (r) and horizon (h). An expert (i) possesses a continuous forecast distribution ϕ_{itrh} for long-run rates at time (t) for region (r) and horizon (h), from which she can generate estimates of growth rates at any given quantile θ_p. An experts quantile estimate (θ_{ithp}) at probability (p) defines the growth rate that an expert believes will be at least equal to the future growth rate with that probability:

$$\theta_p = X_{(Pr(X \leq x)=p)} \tag{5}$$

An experts underlying forecast distribution (ϕ_{itrh}) can be approximated by fitting a continuous distribution $F_x(x)$ to the empirical quantile distribution $F_x(x;\theta_p)_{itrh}$.

$$\min\left(F_x(x) - F_x(x;\theta_{p-k})_{trh}\right) \tag{6}$$

We define an expert's forecast distribution as the cumulative distribution function that minimizes the sum of squared differences between any normal cdf $F_x(x;\hat{\mu},\hat{\sigma})$ and the individual quantile distribution $F_x(x;\theta_{p-k})_{itrh}$. For a set of experts, an estimate of the combined $\hat{\mu}$ and $\hat{\sigma}$ at time (t) for region (r) and horizon (h) is given by:

$$F_x(x|\hat{\mu},\hat{\sigma})_{trh} = \frac{1}{n}\sum_{i=1}^{n} F_x(x|\mu,\sigma)_{trh} \tag{7}$$

We obtain estimates of $\hat{\mu}$ and $\hat{\sigma}$ and find very little difference between the combined pdfs produced using our preferred method (Equations 2,3,4) and the alternate method (Equations

5,6,7). Figure 8 compares the results of the two methods by plotting the combined pdfs using mean, trimmed mean, and median estimates for world growth rates for the period 2010-2100. The forecast estimate obtained using our preferred method implies an average annual growth rate of 2.06% with a standard deviation of 1.12 percentage points. The forecast estimate obtained using our preferred method implies an average annual growth rate of 2.06% with a standard deviation of 1.19 percentage points.

Figure 8: Comparison of Combined Forecast Distributions ($\hat{\phi}_h$) for Global Growth Rates

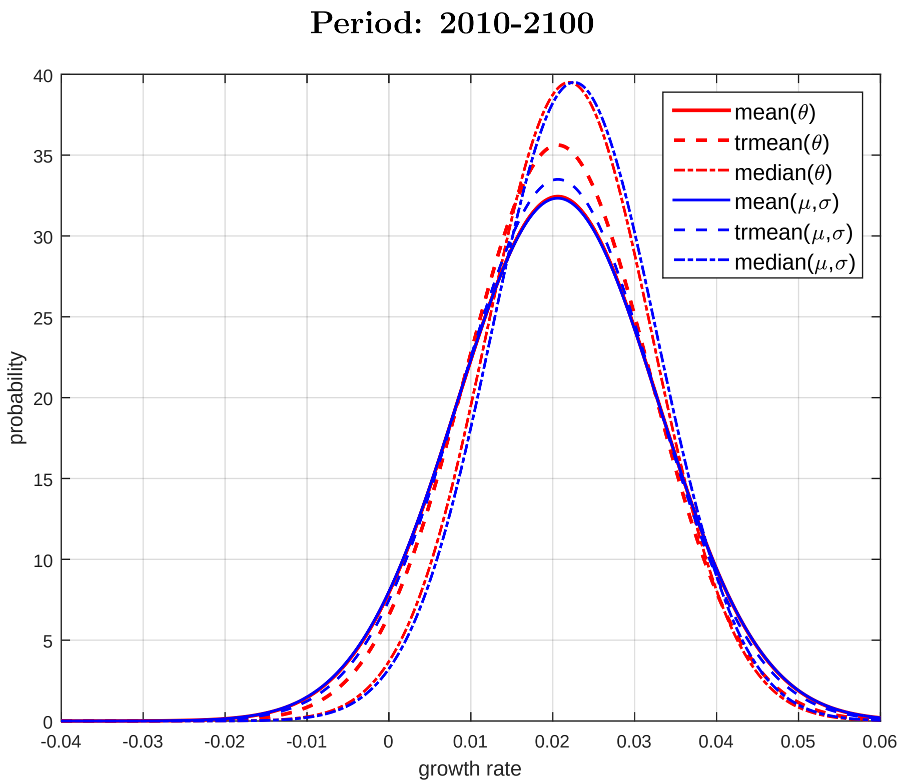

Note: Estimates using mean, trimmed mean, and median from Method 1 (in red) and mean, trimmed mean, and median from Method 2 (in blue).

Appendix C Functional Form Assumptions for Individual Experts

We test the fit of several known parametric functional forms to the probability distributions implied by the 13 full quantile distributions. Figure 9 plots the expert estimates at each of the 5 quantiles against the best-fitting normal distributions for each forecast distribution for global rates 2010-2100. The normal distribution provides a close fit to all but 3 of the individual distributions. Note: this functional form assumption (for individual experts) does not affect the estimates of $\hat{\mu}$ and $\hat{\sigma}$ presented in the primary results section.

Figure 9: Fit of Expert CDF ($\hat{\phi}_i$) to Expert Quantile Distribution ($(F_x(x;\hat{\theta}_{p_1}...\hat{\theta}_{p_k})_i)$) for Global Growth Rates

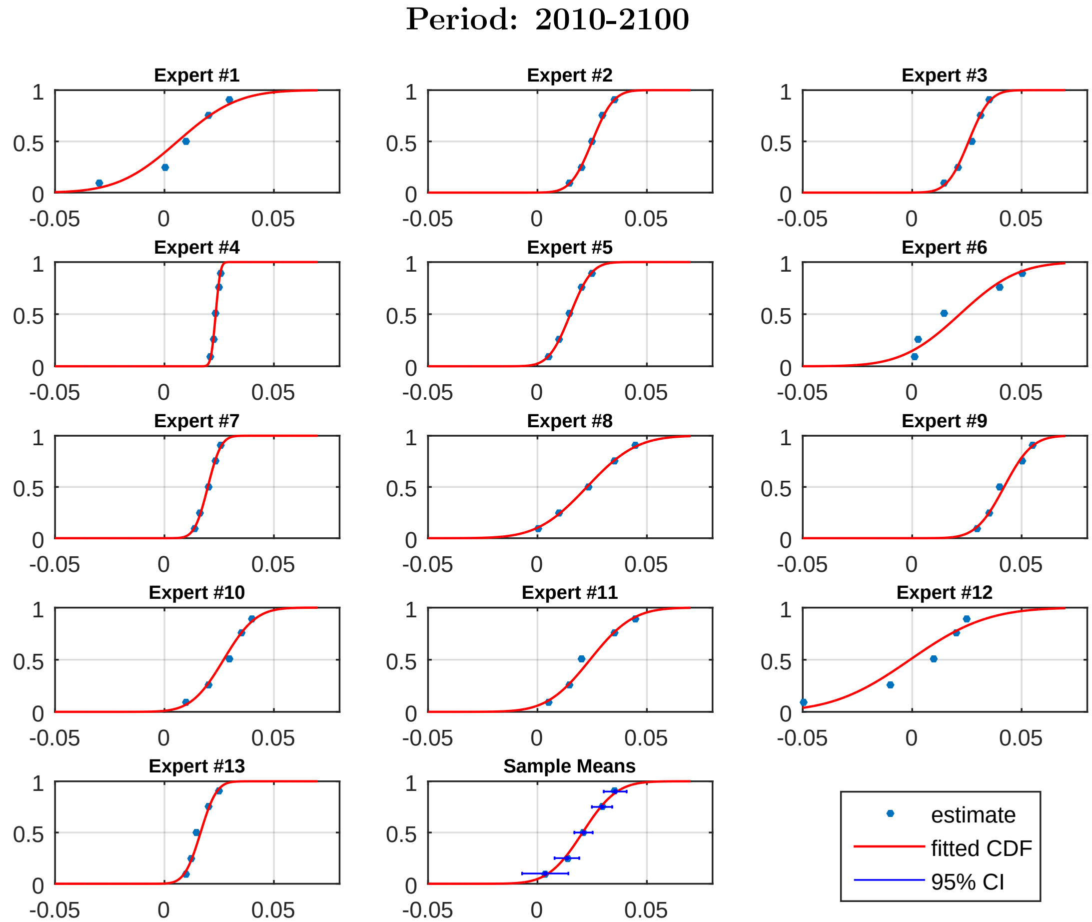

Note: Blue points are the individual quantile distribution F(x; $\hat{\theta}_{p=10}$, $\hat{\theta}_{p=25}$, $\hat{\theta}_{p=50}$, $\hat{\theta}_{p=75}$, $\hat{\theta}_{p=90}$) for expert (i). Red lines are the best-fitting normal CDF.

References

Acemoglu, Daron, Simon Johnson, and James A Robinson, "Institutions as a fundamental cause of long-run growth," *Handbook of economic growth*, 2005, *1*, 385–472.

Arrow, Kenneth, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard Newell, William Nordhaus, Robert Pindyck, William Pizer, Paul Portney et al., "Determining benefits and costs for future generations," *Science*, 2013, *341* (6144), 349–350.

Barro, Robert J and José F Ursúa, "Macroeconomic crises since 1870," Technical Report, National Bureau of Economic Research 2008.

Bates, John M and Clive WJ Granger, "The combination of forecasts," *Or*, 1969, pp. 451–468.

Brynjolfsson, Erik and Andrew McAfee, *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*, Brynjolfsson and McAfee, 2012.

Engelberg, Joseph, Charles F Manski, and Jared Williams, "Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition," *Journal of Applied Econometrics*, 2011, *26* (7), 1059–1078.

Freeman, Richard B, "What does global expansion of higher education mean for the United States?," in "American Universities in a global market," University of Chicago Press, 2010, pp. 373–404.

Gillingham, Kenneth, William D Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, John Reilly, and Paul Sztorc, "Modeling Uncertainty in Climate Change: A Multi-Model Comparison," Technical Report, National Bureau of Economic Research 2015.

Giordani, Paolo and Paul Söderlind, "Inflation forecast uncertainty," *European Economic Review*, 2003, *47* (6), 1037–1059.

Gordon, Robert J, "Is US economic growth over? Faltering innovation confronts the six headwinds," Technical Report, National Bureau of Economic Research 2012.

Lee, Ronald and Shripad Tuljapurkar, "Uncertain demographic futures and social security finances," *The American Economic Review*, 1998, *88* (2), 237–241.

Maddison, Angus, *Historical statistics*, OECD, 2003.

Nordhaus, William D, *A question of balance: Weighing the options on global warming policies*, Yale University Press, 2014.

ONeill, Brian C, Elmar Kriegler, Keywan Riahi, Kristie L Ebi, Stephane Hallegatte, Timothy R Carter, Ritu Mathur, and Detlef P van Vuuren, "A new scenario framework for climate change research: the concept of shared socioeconomic pathways," *Climatic Change*, 2014, *122* (3), 387–400.

_ , _ , Kristie L Ebi, Eric Kemp-Benedict, Keywan Riahi, Dale S Rothman, Bas J van Ruijven, Detlef P van Vuuren, Joern Birkmann, Kasper Kok et al., "The roads ahead: narratives for shared socioeconomic pathways describing world futures in the 21st century," *Global Environmental Change*, 2015.

Rich, Robert and Joseph Tracy, "The relationships among expected inflation, disagreement, and uncertainty: evidence from matched point and density forecasts," *The Review of Economics and Statistics*, 2010, *92* (1), 200–207.

Thomas, Lloyd B, "Survey measures of expected US inflation," *The Journal of Economic Perspectives*, 1999, *13* (4), 125–144.

Wallis, Kenneth F, "Combining forecasts–forty years later," *Applied Financial Economics*, 2011, *21* (1-2), 33–41.

Zarnowitz, Victor, "Consensus and uncertainty in economic prediction," in "Business Cycles: Theory, History, Indicators, and Forecasting," University of Chicago Press, 1992, pp. 492–518.

_ and Louis A Lambros, "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy*, 1987, *95* (3), 591–621.