

Article

Estimating the probability of a lost decade for U.S. and global equity

The Journal of **Financial Perspectives**

EY Global Financial Services Institute

July 2013 | Volume 1 – Issue 2



EY

Building a better
working world

The EY Global Financial Services Institute brings together world-renowned thought leaders and practitioners from top-tier academic institutions, global financial services firms, public policy organizations and regulators to develop solutions to the most pertinent issues facing the financial services industry.

The Journal of Financial Perspectives aims to become the medium of choice for senior financial services executives from banking and capital markets, asset management and insurance, as well as academics and policy-makers who wish to keep abreast of the latest ideas from some of the world's foremost thought leaders in financial services. To achieve this objective, a board comprising of leading academic scholars and respected financial executives has been established to solicit articles that not only make genuine contributions to the most important topics, but are also practical in their focus. *The Journal* will be published three times a year.

gfsi.ey.com

The articles, information and reports (the articles) contained within *The Journal* are generic and represent the views and opinions of their authors. The articles produced by authors external to EY do not necessarily represent the views or opinions of EYGM Limited nor any other member of the global EY organization. The articles produced by EY contain general commentary and do not contain tailored specific advice and should not be regarded as comprehensive or sufficient for making decisions, nor should be used in place of professional advice. Accordingly, neither EYGM Limited nor any other member of the global EY organization accepts responsibility for loss arising from any action taken or not taken by those receiving *The Journal*.

Editorial

Editor

Shahin Shojai
EY LLP

Advisory Editors

Dai Bedford
EY LLP
Shaun Crawford
EY LLP
Carmine DiSibio
EY LLP

Ratan Engineer
EY LLP
David Gittleson
EY LLP
Bill Schlich
EY LLP

Special Advisory Editors

Ben Golub
Blackrock
Anthony Neoh
Bank of China
Kishore Ponnaveolu
Metlife

Anthony M. Santomero
Citigroup
Nick Silitch
Prudential Financial

Editorial Board

Viral V. Acharya
New York University
John Armour
University of Oxford
Tom Baker
University of Pennsylvania
Law School
Philip Booth
Cass Business School and IEA
José Manuel Campa
IESE Business School
Kalok Chan
Hong Kong University of Science
and Technology
J. David Cummins
Temple University
Allen Ferrell
Harvard Law School
Thierry Foucault
HEC Paris
Roland Füss
University of St. Gallen
Giampaolo Gabbi
SDA Bocconi
Boris Groysberg
Harvard Business School
Scott E. Harrington
The Wharton School
Paul M. Healy
Harvard Business School
Jun-Koo Kang
Nanyang Business School
Takao Kobayashi
Aoyama Gakuin University
Howard Kunreuther
The Wharton School
Deborah J. Lucas
Massachusetts Institute of Technology

Massimo Massa
INSEAD
Patricia A. McCoy
University of Connecticut
School of Law
Tim Morris
University of Oxford
John M. Mulvey
Princeton University
Richard D. Phillips
Georgia State University
Patrice Poncet
ESSEC Business School
Michael R. Powers
Tsinghua University
Andreas Richter
Ludwig-Maximilians-Universitaet
Philip Rawlings
Queen Mary, University of London
Roberta Romano
Yale Law School
Hato Schmeiser
University of St. Gallen
Peter Swan
University of New South Wales
Paola Musile Tanzi
SDA Bocconi
Rolf Tilmes
EBS University
Marno Verbeek
Erasmus University
Ingo Walter
New York University
Bernard Yeung
National University of Singapore

Executive summary

Estimating the probability of a lost decade for U.S. and global equity

by **Blake LeBaron**, Abram L. and Thelma Sachar Chair of International Economics at the International Business School, Brandeis University

This paper estimates the probability of equity investments losing value over a ten year period. It suggests that while lost decades are often treated as events that are extremely rare, they are not be as unlikely as many believe. Using data from U.S. and international markets over a very long period it finds that the likelihood of a lost decade is around 7%. This paper also highlights the importance of volatility in stock markets, even over the long term.

Estimating the probability of a lost decade for U.S. and global equity

Blake LeBaron

Abram L. and Thelma Sachar Chair of International Economics, International Business School, Brandeis University

Abstract

This paper estimates the probability of a “lost decade,” where equity investments lose value over a 10-year period. The findings are a reminder that equity investments are risky even over longer time periods, and investors should take this into consideration when making portfolio choices. It also introduces a simple method to allow the reader to combine beliefs about long-run stock returns along with computer simulated return distributions. Finally, the results for the U.S. are augmented with international data which strengthen the case for large long horizon risk.

Introduction

It is often assumed that in the long-run equity markets always generate positive returns for investors. This paper looks at this question using data from U.S. and international markets, concentrating on real and nominal losses over a decade. While the probability of generating a loss over a decade might not easily fit into many standard asset pricing models, it does give investors an intuitive measure of downside risk. This measure is similar in spirit to value-at-risk calculations in that it will be concentrated on a specific aspect of the return distribution.

There is a vast literature exploring the long-range properties of financial market data. The exploration of the equity premium is the most extensive.¹ This paper considers equity returns alone, and only a single feature of their distribution. There are several reasons for doing this. As already discussed, the first reason is that decade losses can be an interesting measure of long-run risk. Second, it avoids having to estimate or proxy for risk-free rates of return in early periods when this might be difficult. Finally, it keeps the bootstrap methodology relatively simple and assumption free.

A long series of U.S. returns is constructed by merging two annual data sets. This set of returns is used with a bootstrapping technique to estimate decade length return distributions, and from these decade loss probabilities. Though small, these loss probabilities are probably higher than what most investors expect. The methodology also allows for readers to estimate loss probabilities using their own beliefs about future expected returns. In the final section, the U.S.-centered data perspective will be augmented using the cross-sectional country return series from Dimson et al. (2002) and Dimson et al. (2008). This dataset contains a comprehensive asset return cross-section from 18 industrialized countries starting at the beginning of the 20th century. These series are used as outside information that should influence investors' beliefs about the U.S. empirical experience.

Return summary

The long return history is built by merging two stock return datasets. The first is the monthly return series described in Schwert (1990), which extends back to 1802. The second is the annual series from 1871-2010, constructed by Shiller, and used in Shiller (2000).² Shiller's dataset also includes inflation series from 1872 onward. This is augmented with inflation series obtained from "Measuring Worth."³ From 1918 onward, the stock returns are from the S&P composite portfolio. From 1872-1917 the stock market information is from the indices created in Cowles and Associates (1939) to track aggregate stock market movements. Earlier returns were assembled by Schwert to best track aggregate market movements. They involve several different sources, and become relatively narrow indices as one moves back in time.⁴ In the earliest period they are mostly bank stocks, and in later periods they include railroad stocks. There are obvious survivorship biases in these indices. Also, many are constructed as monthly averages from bid and ask prices, making precise time series analysis at higher frequencies difficult.⁵

Figure 1 gives an overall picture of lost decades. The y-axis plots the total decade return including dividends for the 10 years ending on the year given by the x-axis. Both nominal, and inflation adjusted real returns are plotted. For the decade ending in 2009 both real and nominal returns are negative, indicating that investors would have lost value on their equity investments. In this figure, lost decades appear to be relatively rare. The decades ending in 1858, 1939, and 1940 are the only other decades with negative nominal equity returns. If one considers real returns, then several others appear. The most recent of these would be the decades ending in the late 1970s and early 1980s, where large U.S. inflation adds a heavy cost to a relatively flat market. This picture is interesting, but not informative as to how likely these events are. The analysis below addresses this question.

The first two rows of Table 1 provide summary statistics for the annual holding period returns used to construct the previous

1 There are many collections of classic papers on this topic. These include Goetzmann and Ibbotson (2006), Mehra (2008), and most recently Hammond et al. (2011). Many books have covered the topic using extensive long-range data sets both U.S. and international. These include Cornell (1999), Dimson et al. (2002) and Siegel (2002). Also, looking at the most recent decade alone, Arnott and West (2010) show that investment losses are far from uniform across various asset classes.

2 Both of these datasets are available at the authors' websites.

3 See <http://www.measuringworth.com/> for full information on the methodology behind the early inflation estimates.

4 Also, dividend yields can only be approximated for the earliest samples from the first part of the 19th century.

5 See Schwert (1990) for detailed discussions.

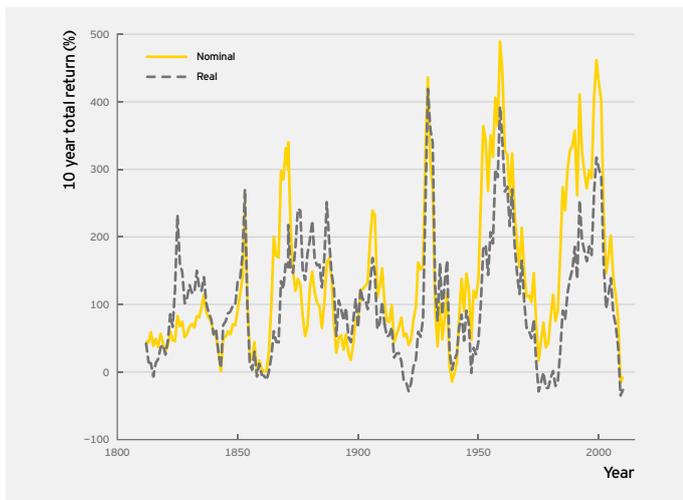


Figure 1: Preceding decade returns

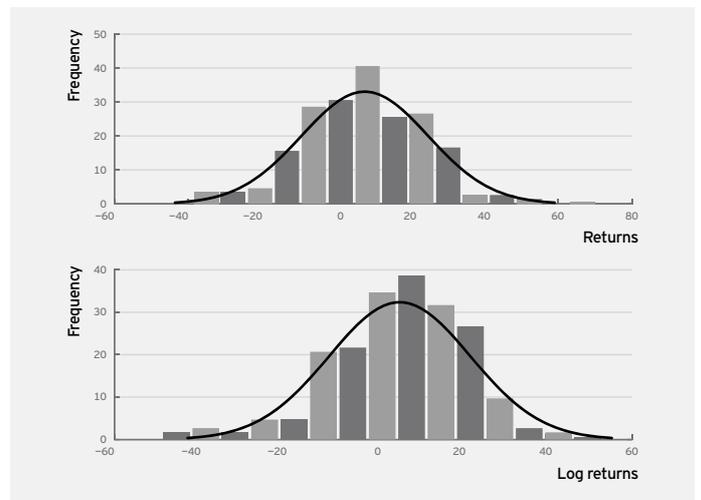


Figure 2: Annual return densities

Series	Mean	Standard deviation	Skewness	Kurtosis	Normal (p-value)
Nominal returns	9.1	16.8	0.18	3.79	0.04
Real returns	7.8	17.1	0.10	3.68	0.09
Nominal log	7.5	15.9	-0.53	4.55	0.00
Real log	6.2	16.4	-0.53	3.91	0.00
Nominal (1802-1871)	6.9	14.0	1.18	7.34	0.00
Nominal (1872-2010)	10.2	18.1	-0.13	3.13	0.78
Real (1802-1871)	7.1	15.6	0.64	5.90	0.00
Real (1872-2010)	8.1	17.8	-0.10	2.99	0.89

Table 1: Summary statistics

Annual arithmetic and geometric returns. Sample length is 209 years covering returns from years ending 1802-2010. All returns include dividends. Mean and Std. estimates are in annual percentages. Kurtosis estimates the kurtosis in the return distribution. This value would be 3 for a normal distribution. Normal refers to a Jarque-Bera test for normality. The value reported is the p-value for the normal null hypothesis. The period 1802-1871 uses the data from the Schwert (1990), series. The period 1872-2010 uses the dataset contained in the Shiller (2000), annual series.

figure. The mean nominal and real annual returns of 9.1 and 7.8% respectively, represent this 209-year U.S. history of equity returns. The table also reports the standard deviation of about 17% per year for both series. This estimate will not be surprising to most investors familiar with the properties of long-range return series. The table also presents skewness and kurtosis levels for these series. These are quick tests of whether a normal distribution would be a reasonable approximation. Skewness is

near zero, and kurtosis near 4, which is larger than its value of 3 for a normal distribution. The last column in the table presents a simple test for normality, the Jarque-Bera test. The values given are the p-values corresponding to the normal null hypothesis with unknown mean and variance.⁶ The first two rows indicate a weak rejection of normality for the nominal returns, and a borderline p-value of 0.09 for the real returns.

Long horizon returns should not be expected to follow a normal distribution since they are compounded short horizon returns,

$$r_{t,A} = \prod_{j=0}^{t-1} (1 + r_{t-j}) - 1 \quad (1)$$

Where r_t are monthly returns. The logged annual returns would then be

$$\log(1 + r_{t,A}) = \sum_{j=0}^{t-1} \log(1 + r_{t-j}) \quad (2)$$

If these annual logged returns are independent with finite second moments, then geometric, or log returns at longer horizons must approach normality. Lines 3 and 4 of table 1 report the same set of summary statistics for the logged annual nominal and

⁶ Since small sample properties of the Jarque-Bera test are questionable, the p-values report results from a 100,000 length monte-carlo.

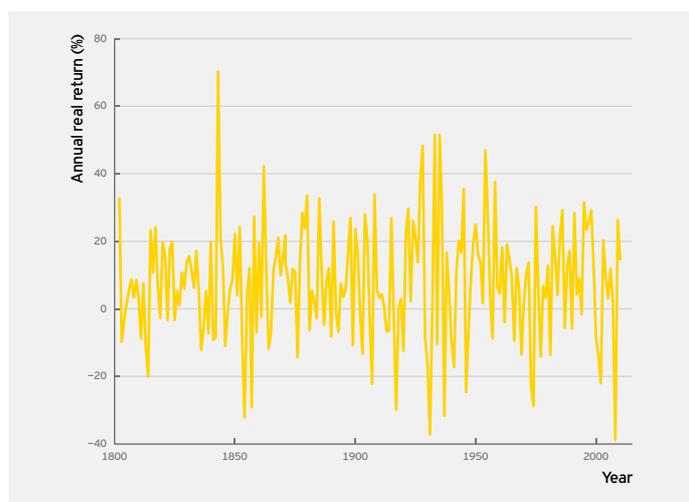


Figure 3: Annual returns

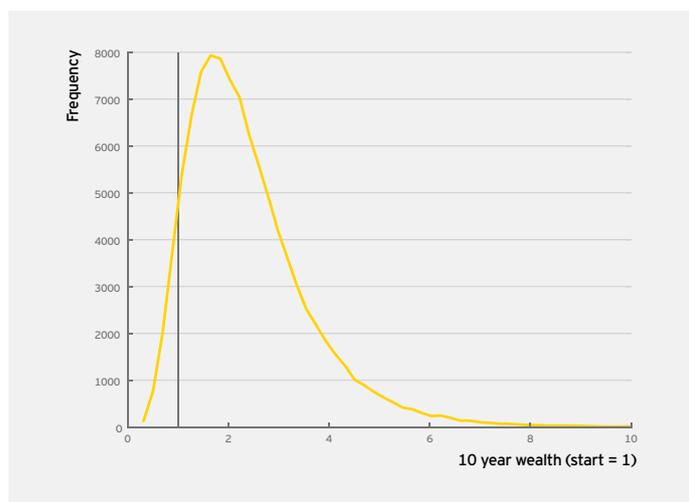


Figure 4: 10-year portfolio values

real returns. Results are similar to those for the holding period returns, except for a moderate increase in kurtosis. This shows up in the Jarque-Bera statistics yielding a p-value of zero in both cases. A visual test for normality is presented in figure 2, which shows the density for the nominal holding period and log returns superimposed with a normal density. The figures suggest that the

practical deviations from normality in these series may be quite small. This will be checked in the later simulations.

The last four rows in table 1 present subsample estimates using 1871 as a break date. This is the date when the Schwert series ends, and the Shiller series begins. Also, a plot of the full return time series is presented in figure 3. The figure shows no obvious visual changes in the series. The table shows that the earlier subsample contains a reduced mean and standard deviation as compared to the full sample. It also displays much larger kurtosis. In the latter periods kurtosis falls enough so that the Jarque-Bera test is unable to reject normality for both real and nominal returns.⁷

Independent returns

Many of the results presented here will be based on a form of bootstrapping where the 209 years of return data are expanded out to 250,000 years by repeatedly drawing years from the original 209 with replacement until a new very long series is built. This is unusual for the bootstrap which usually redraws samples with length less than or equal to that of the original sample. The difference here is that by scrambling the time order we are creating novel decade periods which did not occur in the original series, so the bootstrap is adding to our information beyond the original 209 data points. This does depend critically on assumptions about return independence. Assumptions about independence will be weakened in the next section.

Figure 4 uses the simulated decade returns to generate a histogram of ending wealth levels for an investor starting with a one dollar investment at the beginning of a decade. Two important features should be noted. First, the area to the left of the black line indicates the lost decades, or periods when the investment lost money. Although small, the area is not insignificant. The well-known near log normality of this distribution drives the strong right skew which is evident in the figure. At decade lengths it is interesting to note that probabilities are not insignificant for wealth doubling, or even increasing six-fold.

⁷ One might initially suspect that there are more tail events driving the larger kurtosis in the first subsample, but this does not seem to be the case. In looking at figure 3 one sees similar tail behavior, but a higher concentration of smaller return years (in absolute value) in the earlier part of the dataset. This might be driven by the fact that this is a relatively narrow index, and some years may not have many large information events in the industries covered.

Table 2 presents estimates of decade loss probabilities from the long time series sampling exercise with T=250,000 years, using overlapping decades. The first column and row correspond to the simulation presented in the previous figure. This is a bootstrap using the actual returns in the data, and yields a probability of 0.072.⁸ It is generated from the nominal returns with a mean of 9.1% per year. The second column repeats this estimation for real returns where the mean return falls to 7.7% per year. The probability of a lost decade in this case rises to 0.12. The second line in the table reports bootstrap standard errors for each of the probability estimates. These values are calculated by drawing a new set of 209 annual returns with replacement from the original data, and then using these as a sample for the expansion to the long T sampling methodology for decade return estimation as is done for the actual returns. The bootstrap procedure is repeated 1,000 times, and the standard deviation across this simulations is shown in parenthesis. For the first two return assumptions standard errors of 0.034 and 0.047 show that the estimated probabilities still contain a large amount of uncertainty. This is a reminder that with 200 years of data we will not be able to make strong statements about the tails of decade return distributions.

The sample mean returns from the data may not be the best estimate of expected returns for future decades. This paper does not take a strong stand on what future returns should be, but other return assumptions can easily be applied to the simulations.⁹ The last three columns in table 2 estimate loss probabilities assuming annual mean returns of 10, 6, and 4% respectively. The sample mean is subtracted from the historical series, and one of the three expected return assumptions is added to the returns. It is clear that the estimated loss probabilities change dramatically as assumptions about returns change. For example, if investors only expect a 6% return, then they should prepare for a loss over a decade occurring with probability 0.186. For optimists, expecting a ten% annual return, the decade loss probability falls to 5%.

8 Since this estimator is built from a bootstrapping experiment there will be some uncertainty from the finite length of the bootstrap itself. This can be kept small by keeping the simulation sizes large. The standard error for this estimate will be approximately $\sqrt{p(1-p)/n}$ using the asymptotic approximation to the binomial trial. For $p = 0.07$ this yields an estimate of 0.0016 which was confirmed with a slightly lower value of 0.0015 from a monte-carlo experiment which used the longer sample, and overlapping decades. Therefore, T is sufficiently large to ignore the error coming from the bootstrap.

9 Investors may not agree with either of these long-run expected return levels. See Hammond et al. (2011) for many different perspectives on this. Also, Welch has conducted survey results that are reported on his website <http://research.ivo-welch.info/equpdate-results2009.html>.

Simulation	Annual return				
	9.1% (Nominal)	7.7% (Real)	10%	6%	4%
Bootstrap	0.072 (0.034)	0.120 (0.047)	0.052 (0.028)	0.186 (0.061)	0.300 (0.075)
Bootstrap: 1872-2010	0.067 (0.042)	0.122 (0.060)	0.073 (0.045)	0.218 (0.082)	0.332 (0.098)
Log-normal	0.069 (0.036)	0.118 (0.048)	0.047 (0.028)	0.191 (0.062)	0.313 (0.078)

Table 2: Probability of decade loss – independent returns

Probabilities of a loss in equity portfolios compounded with dividends over decades. All mean returns are adjusted to the given annual mean return shown in the column heading. Bootstrap uses actual returns with readjusted means. Returns are drawn independently from the set of all annual returns. Log Normal uses estimated means and standard deviations along with a standard normal distribution to estimate the decade losses.

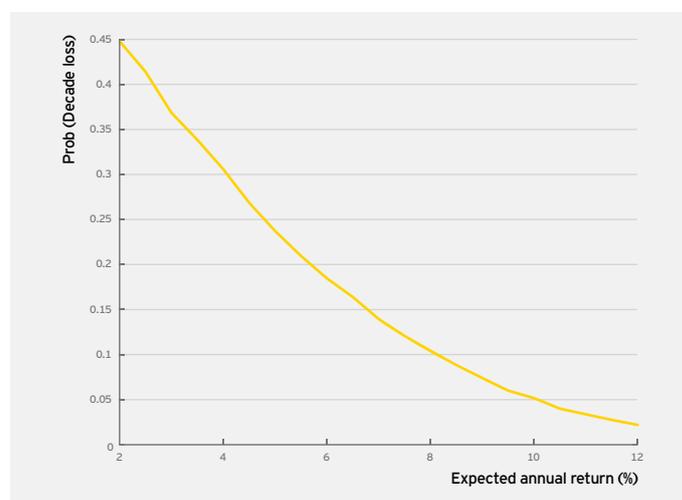


Figure 5: Decade loss probabilities

Figure 5 shows the impact of various return assumptions on the expected decade loss probabilities. Readers can quickly put their own return assumption on the x-axis to assess their long-term chances of a loss. The probabilities are calculated for a discrete set of expected returns using the same bootstrap long-range sampling methods used in table 2. For example, at an expected annual return of 8%, the decade loss probability is approximately 10%. This falls to 2.5% for assumed long-range returns of 12%.

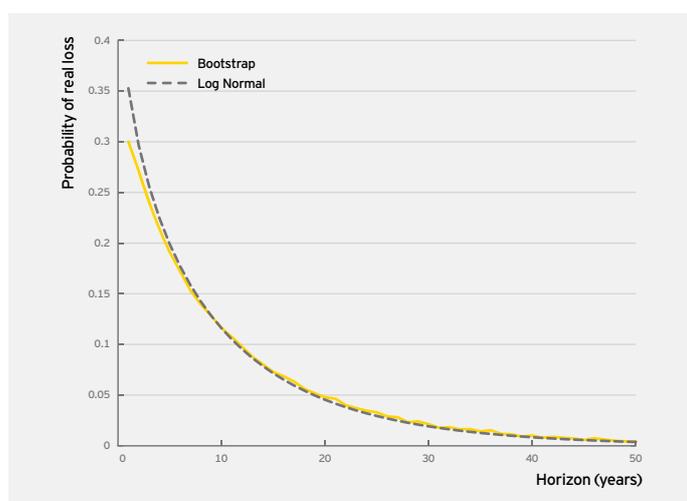


Figure 6: Loss probabilities for increasing horizons

The next two rows in table 2 check whether the results are driven by returns coming from the earliest part of the sample. The sample is restricted to the period from 1872-2010, and the same loss probabilities are estimated. The decade loss probabilities of 0.067 and 0.122 for the real and nominal returns show little change from the full sample. The mean adjusted returns are again used in the last three columns. For these simulations the means will again be adjusted to the given values, but other aspects of the distributions will be different, reflecting the later sample period.

The results presented so far have made minimal assumptions for the return distributions. The final results in table 2 assume that long horizon returns are log normal. In this case, estimating the loss probability involves only estimating the mean and standard deviation for log returns, expanding these to decade length, and then using a normal cumulative distribution function (CDF). Results of this estimation are given in the row labeled “log normal.” The numbers do not change substantially from the corresponding bootstrap values. For example, for real returns, the probability goes from 0.120 for the data to 0.118 under the log normal assumption. The numbers in parenthesis correspond to bootstrap standard errors estimated by redrawing the annual returns, and using the new series for mean and standard deviation estimation. In other words, it recreates the estimation error on the two moments which go into the log normal probability estimates.

The 10-year horizon explored in table 2 may be viewed as a short horizon for some long-term investors. Figure 6 displays loss probability estimates over a range of time horizons using real returns. These values are estimated using the methods from table 2 with both the bootstrap and log normal assumptions. As the horizon increases, the probabilities fall as expected. However, even for a 20-year horizon, the point estimate of the probability of a real loss is still near 5%, which is not zero. One has to move to the 40- and 50-year horizons to consider the loss probabilities as negligible. Also, beyond the very short horizons, normality becomes a good approximation for the return series.

Dependent returns

Up to this point returns have been assumed to be independent over time. If long-range returns were dependent this could change the estimated risk of decade returns. Figure 7 presents the estimated autocorrelations for the real and nominal returns with asymptotic 95% confidence bands around the uncorrelated null hypothesis. Evidence for any correlation in these series is very weak. Combining the first five autocorrelations into a Ljung/Box test yields a p-value of 0.07 for the nominal returns, and 0.16 for the real returns. Figure 8 looks at another measure of long-range dependence, the variance ratio. Assuming σ^2 is the variance of annual log returns, independence gives m period return variances of $m\sigma^2$. Deviations from this are a direct measure of how risk is increasing with horizon lengths.¹⁰ The figure shows the variance ratio declining, indicating some long-range mean reversion, but the asymptotic confidence bands remind us that the sample length is still too short to say anything significant about these values.¹¹

The results on dependence show some weak indications of mean reversion. A cautious tester should consider some form of dependent bootstrapping to explore the possible impact of this. Two different methods for generating dependence in the annual returns series will be used here. First, a parametric bootstrap is performed, using an estimated autoregressive model with five lags, an AR(5), for the returns process. The model and residuals are estimated, and the residuals are then redrawn with replacement and used along with the parameter estimates to

10 See Lo and MacKinlay (1988) and Poterba and Summers (1988) for early examples of this.
11 The 95% confidence bands are generated as in Lo and MacKinlay (1988).

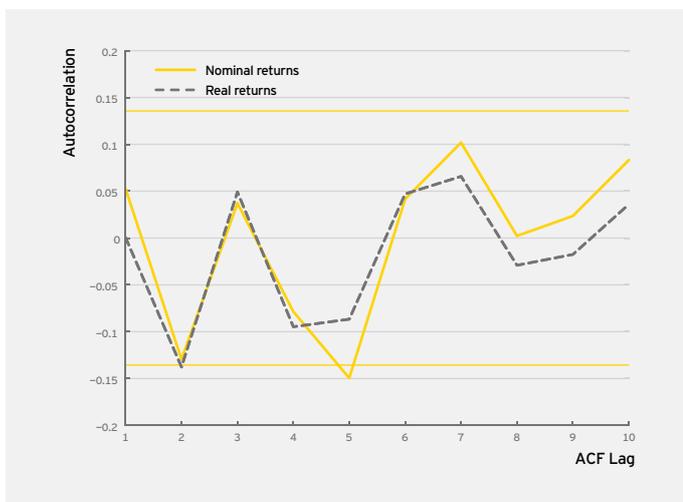


Figure 7: Annual return autocorrelations
Autocorrelations and 95% confidence bounds around uncorrelated null (zero).

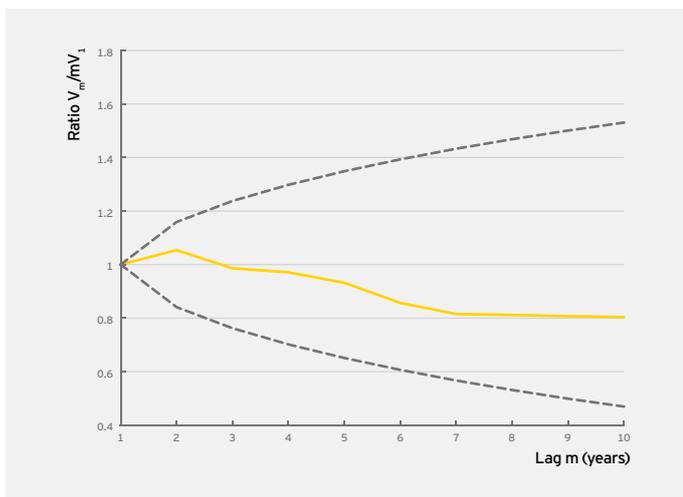


Figure 8: Variance ratios
Estimated variance ratios V_m/mV_1 . Estimate and asymptotic 95% confidence band under IID null (ratio = 1).

Simulation	Real (7.7%)	Nominal (9.1%)
Independent	0.120	0.072
Autoregressive: AR(5)	0.081	0.042
Stationary	0.106	0.051

Table 3: Probability of decade loss – dependent returns

Autoregressive estimates an autoregressive model with 5 lags on the return series, and uses the estimated parameters and resampled residuals to generate dependent data. The stationary bootstrap draws random blocks from the original time series.

generate a new simulated time series of length $T=250,000$.¹² The independent loss probabilities are repeated in the first row of table 3 for comparison. The second row presents the estimated loss probabilities using the AR(5) null model, using overlapping decades on the long sample as in table 2. There is some reduction for both the real returns, and the nominal returns series, indicating the data has some weak information about long-run mean reversion. However, the probabilities are still not trivial.

A second simulation experiment uses the stationary bootstrap to replicate the dependence in a nonparametric fashion.¹³ This form of bootstrap draws returns in contiguous blocks where block length is controlled by a random variable, X_t which is 1 with probability λ , and 0 with probability $1-\lambda$. Assume a new time series is being constructed at the current point t , and is drawn from point τ in the original series. The next point, $t+1$, will come from $\tau + 1$, if $X_t = 0$, and will come from a new point $\hat{\tau}$, if $X_t = 1$. This generates a series containing blocks of varying lengths from the old series, where the lengths are controlled by the behavior of X_t .¹⁴ Results for this simulation are given in the second line of the table. For these runs, $\lambda = 0.2$, which gives an average block length of five years.¹⁵ The values are close to those from the AR(5) simulation, and indicate that these two methods may be replicating a similar amount of dependence for the long-range returns.

12 See Efron and Tibshirani (1993) for a basic description of the parametric bootstrap, and Maddala and Li (1996) for financial applications. Kilian and Berkowitz (2000) is a useful survey which covers many issues on modeling dependence in time series. Given the reported autocorrelations, the AR(5) would appear to be an overfit model. Kilian (2001) shows that when bootstrapping standard errors, the dangers of under parameterization outweigh those of over parameterization.

13 See Politis and Romano (1994) for the original derivations. Also, see Sullivan et al. (1999) for a financial application.

14 The varying block lengths follow a geometric distribution with mean $1/\lambda$.

15 Values for 10 and 20 years have also been tried, generating similar results.

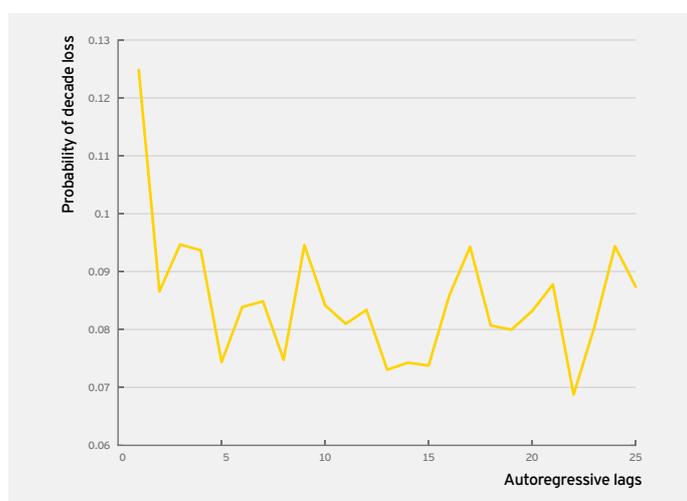


Figure 9: Decade losses over varying autoregressive models

Figure 9 explores the possibility that weaker long-range dependence might yield evidence for lower decade loss probabilities. The parametric AR simulation is run for lags of 1 through 25, and the loss probabilities, estimated using the previous methodology are plotted. This figure uses the merged real returns as the starting data series. There is an early sharp drop off in probabilities as the lags increase from 1 to 5. However, at this point the probabilities stabilize near 0.07–0.08, which is consistent with table 3. Adding further lags appears unlikely to dramatically reduce estimated long-range losses. This is consistent with the evidence that long-range mean reversion is weak.

International cross-section

Up to this point, the analysis has concentrated only on long-range series built from U.S. returns and inflation series. Series used in Dimson et al. (2002) provide a useful long-range cross-section for comparison.¹⁶ The series are annual, and extend from 1900 though 2010 for 111 years of data. Only real equity returns will be used here.

¹⁶ These series are available from Morningstar. Another important study on the international cross-section is Jorion and Goetzmann (1999).

Table 4 presents the results for the international returns. It includes all the countries in the dataset along with value and equal weighted portfolios. The first two columns record the annual mean and standard deviation for the logged returns. The column labeled “Bootstrap” repeats the long resampling procedure, taking each series out to 250,000 observations to estimate the decade loss probability. The column labeled “Normal” uses the independent log normal return assumption, with the sample means and standard deviations.

The results from table 2 reported a decade loss probability for real returns in the U.S. of 0.12. Comparison of this number with values in table 4 shows that relative to the rest of the world, the U.S. is a safer country than most when it comes to long-run tail risk. Using the bootstrap estimator, countries vary from a high of 0.396 for Italy to a low of 0.108 for Australia. Similar to the previous results, the normal approximations do not have a large impact on the results, continuing to support the idea that normality is a good approximation at long horizons.

A graphical summary of this table is given in figure 10, which displays a histogram of the bootstrapped loss probabilities from table 4. This gives a clear picture of where the U.S. falls in terms of long-run risk. Also, it shows that if investors are going to use this data to adjust their beliefs about risk in the U.S., they should increase their risk assessment. Finally, the last two lines in table 4 present results for both value and equal weighted global portfolios. These results show surprisingly small reductions in risk from either form of diversification. The equal weighted portfolio reduces the decade loss probability to 0.10, which is lower than the individual countries, as it has to be, but the gain is small. Furthermore, given various impediments to international investing over the early parts of this sample, the feasibility of achieving these returns should be viewed with some skepticism.

Table 5 performs some additional experiments, exploring the cross-sectional dimensions of the real return data, and how it impacts investors. All the experiments are bootstrap simulations done with 10,000 replications. For calculating the loss probabilities, each simulation uses the normal approximation using estimated means and standard deviations.

The first experiment, labeled “Country draw,” assumes a random draw of new returns data which is structured by country. First,

Country	Mean	Standard deviation	Bootstrap	Normal
Australia	0.072	0.179	0.108	0.103
Belgium	0.025	0.228	0.360	0.364
Canada	0.057	0.166	0.136	0.139
Denmark	0.050	0.187	0.194	0.200
Finland	0.052	0.276	0.268	0.274
France	0.030	0.226	0.329	0.337
Germany	0.030	0.345	0.326	0.391
Ireland	0.037	0.233	0.287	0.309
Italy	0.020	0.291	0.396	0.415
Japan	0.037	0.326	0.306	0.358
Netherlands	0.049	0.203	0.220	0.225
Norway	0.041	0.238	0.285	0.292
New Zealand	0.057	0.184	0.157	0.166
South Africa	0.071	0.204	0.132	0.137
Spain	0.035	0.209	0.291	0.299
Sweden	0.061	0.215	0.179	0.184
Switzerland	0.041	0.189	0.238	0.244
United Kingdom	0.052	0.193	0.186	0.197
U.S.	0.061	0.199	0.166	0.167
World	0.053	0.173	0.165	0.165
Equal weighted	0.061	0.148	0.100	0.097

Table 4: Probability of decade loss – country real equity returns

Summary statistics and decade loss probabilities for developed country returns. Bootstrap repeats methods from table 2, and Normal assumes a log normal return distribution.

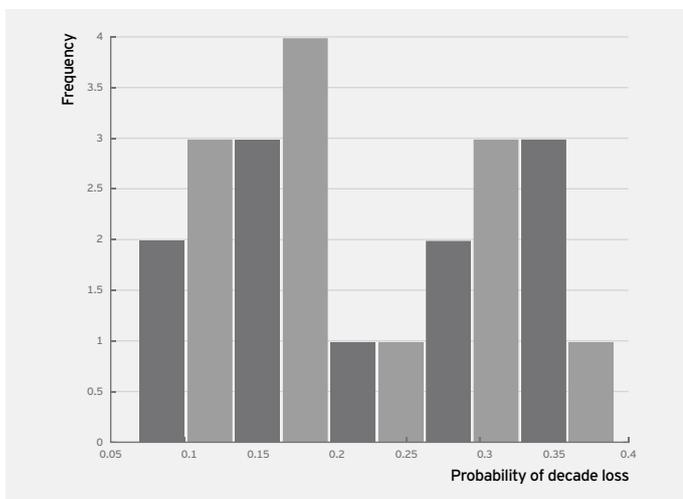


Figure 10: Country loss histogram

Experiment	q0.10	q0.50	q0.90
Country draw	0.099	0.245	0.445
Full sample	0.136	0.260	0.415
AR(2)	0.127	0.252	0.347
AR(5)	0.102	0.233	0.350

Table 5: Probability of decade loss – cross-section experiments

Columns represent quantiles from 10,000 bootstrapped distributions for estimated decade loss probabilities. Country draw assumes investors face a random draw of a country. Full sample pools all the return data into one large sample. AR(2) and AR(5) assume autoregressive models with two and five lags respectively.

one of the countries is drawn at random from the pool. Then its returns are redrawn with replacement giving a new sample which is used to estimate the mean and standard deviation. The investor here is viewing the countries as different, but could potentially face data that looks like any one of them going into the future. The columns report the 0.1, 0.5, 0.9 quantiles for the decade loss probability distribution. The median value of 0.245 is consistent with the cross-sectional results, and the graphical information in figure 10, all of which show that the probability of a lost decade is large in the international returns series. The bootstrap runs generate a large amount of dispersion as shown by the 0.1 and 0.9 quantiles, which are estimated at 0.099 and 0.445.

The second row of the table, labeled “Full sample,” pools the entire dataset into one set of returns, and then draws country samples from this pooled population. This implements a null hypothesis that all country returns come from the same population. Pooling all the data reduces the dispersion from the separate country sampling method as is seen in the narrowing of the extreme quantiles to 0.136 and 0.415. However, the median value of 0.260 does not change much by moving to the pooled sample.

The last two rows test the impact of dependence on the results. None of the international returns series show much evidence for return autocorrelation, but the importance of this is tested by repeating the parametric bootstrap that was used before with the U.S. returns. Two models, an AR(2), and an AR(5) are used as a simple dependent null hypothesis. These are estimated once on each country, and then simulated for T=250,000 years to estimate the decade loss probability. The quantiles represent the cross-section across the countries. The results again show little change from the earlier results that assumed independence in

the returns. The median loss probabilities for the two dependent cases are 0.252 and 0.233 for the AR(2) and AR(5) respectively. Evidence from the international returns again suggest that long-range return dependence should not reduce investors' beliefs about the riskiness of decade returns.

Summary

Lost decades are often treated as a kind of "black swan" event that are almost impossible. Results in this paper show that while they are a tail event, they may not be as far out in the tail as the popular press would have us think. Allowing the data to speak directly in an independent bootstrap, with two centuries of return time series, the estimate of a portfolio loss over a decade is about 7%. For the investor concerned with real returns the results are more depressing, with decade loss probabilities of 12%. The bootstrap methodology is not dependent on distributional assumptions about annual returns. However, in the reported estimates, normality assumptions for annual returns do not have a major impact on decade length results.

The estimated loss probabilities are checked for robustness in two ways. First, the independent null hypothesis is weakened by using several methods for simulated dependent long-range returns. In all the tests there is only a small reduction in the decade loss probability, which is consistent with the very weak mean reversion present in aggregate long-range returns. Second, the U.S. experience is compared with international equity data using several different tests. Consistent with other research, global data do not give U.S. investors any increased confidence in terms of risk. On the contrary, long-run results across the globe consistently appear riskier in terms of decade losses in real equity returns.

The simple message here is that stock markets are volatile. Even in the long-run volatility is still important. These results emphasize that 10-year periods where an equity portfolio loses value in either real or nominal terms should be an event on which investors put some weight when making their investment decisions.

References

- Arnott, R. and West, J. (2010)**, "Was it really a lost decade?", Index Universe .
- Cornell, B. (1999)**, *The Equity Risk Premium: The Long-Run Future of the Stock Market*, Wiley, New York, NY.
- Cowles, III, A. and Associates (1939)**, *Common Stock Indexes*, number 3 in "Cowles Commission Monograph", 2nd edn, Principia Press, Bloomington, Indiana.
- Dimson, E., Marsh, P. & Staunton, M. (2002)**, *Triumph of the Optimists: 101 Years of Global Investment Returns*, Princeton University Press, Princeton, New Jersey.
- Dimson, E., Marsh, P. & Staunton, M. (2008)**, *The worldwide equity premium: A smaller puzzle*, in R. Mehra, ed., "The Handbook of the Equity Risk Premium", Elsevier, Amsterdam, Netherlands.
- Efron, B. & Tibshirani, R. (1993)**, *An Introduction to the Bootstrap*, Chapman and Hall, New York.
- Goetzmann, W. N. & Ibbotson, R. G. (2006)**, *The Equity Risk Premium*, Oxford University Press, Oxford, UK.
- Jorion, P. & Goetzmann, W. N. (1999)**, "Global stock markets in the twentieth century", *Journal of Finance* 54, 953–980.
- Hammond, P. B., Leibowitz, M. L. & Siegel, L. B. (2011)**, *Rethinking the Equity Risk Premium*, CFA Institute.
- Kilian, L. (2001)**, "Impulse response analysis in vector autoregressions with unknown lag order", *Journal of Forecasting* 20, 161–179.
- Kilian, L. & Berkowitz, J. (2000)**, "Recent developments in bootstrapping time series", *Econometric Reviews* 19(1), 1–48.
- Lo, A. W. & MacKinlay, A. C. (1988)**, "Stock prices do not follow random walks: Evidence from a simple specification test", *Review of Financial Studies* 1, 41–66.
- Maddala, G. S. & Li, H. (1996)**, *Bootstrap based tests in financial models*, in G. S. Maddala & C. R. Rao, eds, "Handbook of Statistics", Vol. 14, North-Holland, Amsterdam, pp. 463–488.
- Mehra, R., ed. (2008)**, *Handbook of the Equity Risk Premium*, Elsevier, Amsterdam, the Netherlands.
- Politis, D. & Romano, J. (1994)**, "The stationary bootstrap", *Journal of the American Statistical Association* 89, 1303–1313.
- Poterba, J. M. & Summers, L. H. (1988)**, "Mean reversion in stock prices: Evidence and implications", *Journal of Financial Economics* 22, 27–59.
- Schwert, G. W. (1990)**, "Indexes of United States stock prices from 1802-1987", *Journal of Business* 63, 399–426.
- Shiller, R. J. (2000)**, *Irrational Exuberance*, Princeton University Press, Princeton, NJ.
- Siegel, J. J. (2002)**, *Stocks for the Long-run: The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies*, 3rd edn, McGraw-Hill, New York, NY.
- Sullivan, R., Timmerman, A. & White, H. (1999)**, "Data-snooping, technical trading rule performance and the bootstrap", *Journal of Finance* 54, 1647–1691.

About EY

EY is a global leader in assurance, tax, transaction and advisory services. The insights and quality services we deliver help build trust and confidence in the capital markets and in economies the world over. We develop outstanding leaders who team to deliver on our promises to all of our stakeholders. In so doing, we play a critical role in building a better working world for our people, for our clients and for our communities.

EY refers to the global organization and may refer to one or more of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. For more information about our organization, please visit ey.com.

© 2013 EYGM Limited.
All Rights Reserved.

EYG No. FP0006



In line with EY's commitment to minimize its impact on the environment, this document has been printed on paper with a high recycled content.

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, or other professional advice. Please refer to your advisors for specific advice.

ey.com

The articles, information and reports (the articles) contained within *The Journal* are generic and represent the views and opinions of their authors. The articles produced by authors external to EY do not necessarily represent the views or opinions of EYGM Limited nor any other member of the global EY organization. The articles produced by EY contain general commentary and do not contain tailored specific advice and should not be regarded as comprehensive or sufficient for making decisions, nor should be used in place of professional advice. Accordingly, neither EYGM Limited nor any other member of the global EY organization accepts responsibility for loss arising from any action taken or not taken by those receiving *The Journal*.

Accredited by the American Economic Association

ISSN 2049-8640
