Journal of Financial Economics 127 (2018) 248–263

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Four centuries of return predictability☆

Benjamin Golez[a], Peter Koudijs[b,c,*]

[a] University of Notre Dame, 256 Mendoza College of Business, Notre Dame, IN 46556, USA
[b] Stanford University, 655 Knight Way, Graduate School of Business, Stanford, CA 94305, USA
[c] National Bureau of Economic Research, 1050 Massachusetts Ave., Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 16 March 2016
Revised 16 January 2017
Accepted 17 January 2017
Available online 19 December 2017

JEL classification:
G12
G17
N2

Keywords:
Dividend-to-price ratio
Return predictability
Dividend growth predictability

ABSTRACT

We combine annual stock market data for the most important equity markets of the last four centuries: the Netherlands and UK (1629–1812), UK (1813–1870), and US (1871–2015). We show that dividend yields are stationary and consistently forecast returns. The documented predictability holds for annual and multi-annual horizons and works both in- and out-of-sample, providing strong evidence that expected returns in stock markets are time-varying. In part, this variation is related to the business cycle, with expected returns increasing in recessions. We also find that, except for the period after 1945, dividend yields predict dividend growth rates.

1. Introduction

A large body of work suggests that price multiples, such as the dividend-to-price ratio, predict stock returns.[1] As a result, modern asset pricing theory increasingly incorporates time-varying expected returns (Campbell and Cochrane, 1999; Bansal and Yaron, 2004; Albuquerque et al., 2016; among others). The majority of the empirical work underpinning these findings uses US stock market data going back to 1926 or 1945.

In this paper, we examine whether dividend yields predict returns in a long sample that covers four centuries of data, going back to the stock market's earliest years in the 17th century. In doing so, we provide an out-of-sample test, asking whether results that hold in the recent US period are generalizable to other times and places (Schwert, 1990). Through the extended time series, we also obtain

☆ We thank Bill Schwert (the editor), an anonymous referee, Jules van Binsbergen (discussant), Jacob Boudoukh, John Campbell (discussant), Anna Cieslak, Zhi Da, Nicholas Hirschey (discussant), Arvind Krishnamurthy, Hanno Lustig, Stijn van Nieuwerburgh (discussant), Monika Piazzesi, Martin Schneider, Robert Stambaugh (discussant), Jessica Wachter, Michael Weber (discussant), Ivo Welch, and seminar participants at University of Notre Dame, University of California, San Diego, Stanford University, European Central Bank, University of Konstanz, Università della Svizzera Italiana in Lugano, Switzerland, Nova School of Business and Economics, Bocconi University, University of Zurich, 2014 Stanford Institute in Theoretical Economics (SITE) Conference, 2015 Rodney L. White Center for Financial Research Conference, 2015 Western Finance Association Conference, 2015 Red Rock Conference, 2016 American Economic Association Conference, and 2016 European Winter Finance Conference for comments and suggestions. We thank Gary Shea for generously sharing his data about 18th century British companies and John Turner for his advice on how to use British stock price information from the 19th century. George Jiang, Manisha Goswami, and Cathy Quiambao provided valuable research assistance.

* Corresponding author at: Stanford University, 655 Knight Way, Graduate School of Business, Stanford, CA 94305, USA.
E-mail addresses: bgolez@nd.edu (B. Golez), koudijs@stanford.edu (P. Koudijs).

[1] Amongst other work, this includes Shiller (1981), LeRoy and Porter (1981), Campbell and Shiller (1988b), Fama and French (1988), Cochrane (2008), and Binsbergen and Koijen (2010).

https://doi.org/10.1016/j.jfineco.2017.12.007
0304-405X/

more statistical power to reject the null of no return predictability (Stambaugh, 1999; Boudoukh et al., 2008).

We assemble annual stock market data for the most important equity markets of the last four centuries: the Netherlands and UK (1629–1812), UK (1813–1870), and US (1871–2015). We analyze the data for each subperiod individually and for the sample as a whole. We consider the most basic form of predictability analyzed by Campbell and Shiller (1988b); Fama and French (1988), and Cochrane (1992, 2008) and ask whether the dividend-to-price ratio forecasts returns. Many other variables could predict returns as well. The purpose of this paper is not to come up with the best possible forecasting model. Instead, we take the most common predictive variable in the literature and evaluate its forecasting power in an extended sample. Additional variables would only strengthen the evidence for return predictability.

Our results confirm that dividend yields predict returns. With 384 annual observations, we have sufficient statistical power to reject the null hypothesis of no return predictability, both over short and long horizons. Our estimates are stable across the different markets and periods we consider. Moreover, returns are forecastable in real time, and in- and out-of-sample estimates have a similar fit (Goyal and Welch, 2008). This confirms that expected returns are time-varying. In line with earlier literature (Fama and French, 1989), we find that the documented predictability is related to the business cycle with expected returns increasing in recessions. Our results also indicate that the dividend-to-price ratio predicts dividend growth rates, but only for the period before 1945. In the final section of the paper, we conjecture what could explain the reduced ability of the dividend yield to forecast dividend growth rates in the recent period.

An alternative to using a longer sample is to look at a cross section of countries.[2] Data for most countries typically span relatively short time periods, and markets exhibit a high degree of co-movement, especially in recent decades. This reduces the statistical power to reject the null of no return predictability. In contrast, extending the data backward adds independent variation to the data. At the same time, many key characteristics of modern financial markets, such as separation of ownership and control and the ability to freely trade shares, were already present in the 17th century.

Apart from a large body of literature on return predictability, our paper is related to other studies analyzing stock markets in earlier periods. Schwert (1990) combines secondary sources to reconstruct a US stock index for the 19th century and finds that both the volatility and seasonal patterns of stock returns are similar in the 19th and 20th centuries. Relying on primary sources, Goetzmann et al., (2001) estimate a new index for the New York stock market between 1815 and 1925. They find little evidence for return predictability, but data limitations force them to approximate dividends for the period before 1870. Chen (2009) analyzes the predictability of returns and dividend growth for the US between 1872 and 2005 and documents that returns are largely unpredictable before 1926 or 1945. At the same time, dividend growth rates are strongly forecastable.[3] In comparison with these studies, we use a much longer sample period with complete data for dividend payments. Bris et al., (forthcoming) analyze six hundred years of dividend and price data for the Bazacle Company in France. They find evidence that both time-varying risk premia and changes in expected dividends affect share prices. In comparison, we analyze the aggregate market. Both studies find that the dividend-to-price ratio fluctuates around a long-run average of 5% in the early data.

The paper proceeds as follows. Section 2 provides a brief description of the data and discusses a number of summary statistics. Section 3 provides the main results. Sections 3.1 and 3.2 document the predictability of returns and dividends growth rates from dividend yields. Sections 3.3 to 3.5 provide additional statistical tests. Section 4 examines the link between recessions and predictability. Section 5 discusses the results and concludes.

2. Data

We extend the annual time series of stock prices and dividends back to 1629 using the most important financial market of a specific period. For 1629 through 1812, we look at the equity markets in the Netherlands and the UK. Between 1813 and 1870, we focus on Great Britain exclusively. For the period after 1870, we rely on US market data. The indices we use weigh individual stocks by their market capitalization. End-of-year dividends are obtained by summing dividends within the year.

2.1. Data description

We provide a brief description of the markets we analyze and the main sources of the data. Details are in the Online Appendix.

2.1.1. Netherlands and UK: 1629–1813

Amsterdam was arguably the most important financial center of the 17th and 18th centuries. It was closely integrated with the market in London and featured trade in the largest English securities (Neal, 1990). Although technologically less advanced, the market functioned similarly to today. Harrison (1998) provides evidence that the time series properties of returns in these markets were similar to more recent periods. Koudijs (2016) shows that stock prices responded to the arrival of news in an efficient way. The paper's calculations suggest that trading costs in the 1770s and 1780s were comparable to those on the NYSE between 1993 and 2005. We take the perspective of an Amsterdam investor who held a value-weighted portfolio of Dutch and English securities. We convert prices and dividends of English securities from pounds sterling to Dutch guilders. Because both countries were on metallic standards, exchange rate fluctuations were relatively small with an annual standard deviation of 3%.

[2] See Campbell and Shiller (2005) and Rangvid, Schmeling, and Schrimpf (2014). Santa-Clara and Maio (2015) perform a similar analysis on a cross section of portfolios.

[3] Schwert (2003) and Goyal and Welch (2003) provide similar evidence.

Information is available for up to nine securities, representing the near universe of traded equities in Amsterdam and London. All companies had limited liability for their shareholders and were widely traded. We collect all information available in the secondary literature, and supplement it with original archival material whenever possible. Years exist in which the information for a particular stock is missing. Data coverage for the five largest securities is nearly complete (Online Appendix, Section A.1 has details). These include the Dutch East India Company (VOC, 1629–1794), the Dutch West Indies Company (WIC, 1719–1791), the (United) British East India Company (EIC, 1692–1813), the Bank of England (BoE, 1696–1813, then a private bank with strong ties to the government), and the South Sea Company (SSC, 1711-1813). Before 1692, our data set consists of the VOC alone, which was by far the largest company in the Netherlands and the UK at that time. For example, in 1696, after both the EIC and BoE are added to the index, the VOC still accounts for 63% of total market capitalization.

Initially, dividend payments were highly irregular. Companies started making annual dividend payments only at the end of the 17th century. For the VOC, this seems to have been related to shareholders having different preferences for dividends. Depending on which shareholder group dominated in a particular year, dividends were either very high or foregone altogether (Lent and Sgourev, 2013). Following the literature on cyclically adjusted price-earnings ratios (Campbell and Shiller, 1988a; 2005), we use a ten-year trailing average to smooth out dividend payments before 1700. Returns are calculated using actual payout. To ensure that the adjustment of dividends does not drive our findings, we also report results using post-1700 data only.

To gauge the size of the market, we compare the aggregate value of the securities in our data with the Gross Domestic Product (GDP) of the province of Holland, the richest and most populous region of the Netherlands (including cities such as Amsterdam and Rotterdam), for which comprehensive GDP estimates are available (Zanden and Leeuwen, 2012). Using information from Bowen (1989) and Wright (1997), we estimate the fraction of English securities held by Dutch investors. Market capitalization to GDP ranged from 15% (during the 1630s and again in the early 1800s) to 64% (during the 1720s). In comparison, US stock market capitalization amounted to 39% of GDP in 1913 and 152% in 1999 (Rajan and Zingales, 2003).

2.1.2. UK: 1813–1870

For the period between 1813 and 1870, we focus on the UK market. After the Napoleonic Wars, London became the financial center of the world, with the United Kingdom its largest economy. Starting in the 1810s, many new equities were issued. Initially, these were mainly canals and insurances companies. Later on, banks and railroad companies became the most important issuers of new stocks. The period includes the so-called Railroad Manias of the 1830s and 1840s. In contrast to the earlier period, securities issued between 1810 and 1855 generally had full shareholder liability (Hickson et al., 2011). After 1855, it became possible to issue shares with limited liability, but many banks and insurance companies continued to maintain full liability until the end of the sample period.

We use, starting in 1825, the value-weighted stock market index constructed by Acheson et al., (2009) that includes all frequently traded domestic equities in London. We extend their series backward to 1813 using the same source material and methodology (details are in the Online Appendix, Section A.2). The final index covers between 50 (1813) and 250 (1870) different securities. Total market capitalization accounted for between 10% and 30% of English GDP.

2.1.3. US: 1871–2015

To facilitate comparison with the existing literature, we rely on the US stock market data starting in 1871 using data from Amit Goyal's website.[4] For the period between 1871 and 1925, these data come from Cowles (1939), covering between 50 (1871) and 258 (1925) securities. From 1926, the data are based on the Standard and Poosr's (S&P) 500 index provided by the Center for Research in Security Prices (CRSP). Before 1957, this was actually the S&P 90.

In alternative estimates, we focus on the London market for the entire 19th century [relying on data from Grossman (2002) for the period after 1870], switching to the US only after 1900 when it became the world's largest economy with New York as an international financial center. This has the additional advantage that London featured many more securities (~750) than New York in this period. In the Online Appendix (Table OA.2), we show that our results are robust to this alternative approach. In the same appendix, we also show that results are similar when using the broad based CRSP index after 1925 that, again, is composed of a substantially larger number of securities (522 in 1926).

2.1.4. Inflation, business cycles, and the risk-free rate

To account for changes in purchasing power, we obtain data on inflation. For the Netherlands and UK period (1629–1809), we rely on information from the International Institute of Social History.[5] Clark (2015) provides price data for the UK period (1813–1870 or 1900). For the US period, we use the Consumer Price Index (CPI) from Robert Shiller's webpage.[6]

We also use data on recessions. For the US period, we rely on the standard National Bureau of Economic Research (NBER) chronology of expansions and contractions. The data are monthly. We classify a year as a recession if at least six months in that year are characterized as a downturn. For the period 1700–1870, we focus on recessions in the UK. We rely on peak and through dates from Ashton (1959), Gayer et al., (1953), and Rostow (1972) (details are in the Online Appendix, Section A.4). UK recession data are annual. We let recessions start in the year following a peak and end in the year of a trough. No business cycle information is available for the 17th century.

Finally, we collect information on the risk-free rate. For most of the early periods, there were no liquid short-term government securities. Before 1871, we use returns

[4] See http://www.hec.unil.ch/agoyal/.

[5] See http://www.iisg.nl/hpw/data.php#netherlands.

[6] See http://www.econ.yale.edu/~shiller/.

on Dutch and English annuities or consols (details are in the Online Appendix, Sections A.1–A.3). Information for the years before 1650 and between 1720 and 1727 is missing. Between 1871 and 1919, we use US interest rates on short-term commercial paper and long-term government bonds to estimate the risk-free rate (details on the estimation procedure are in the Online Appendix, Section A.3). Starting in 1920, we use the rate on US T-bills. In sum, before 1920 we have only an estimate of the risk-free rate; the data are incomplete, and the underlying securities were exposed to interest rate and sometimes even default risk. We therefore run all our regressions on raw instead of excess returns.

2.2. Summary statistics

Fig. 1 plots annual real returns, dividend growth rates and dividend yields. Table 1 reports summary statistics, in both nominal (Panel A) and real (Panel B) terms. To facilitate comparison with the existing literature, we divide the data into four periods: 1629–1812 (Netherlands and UK),

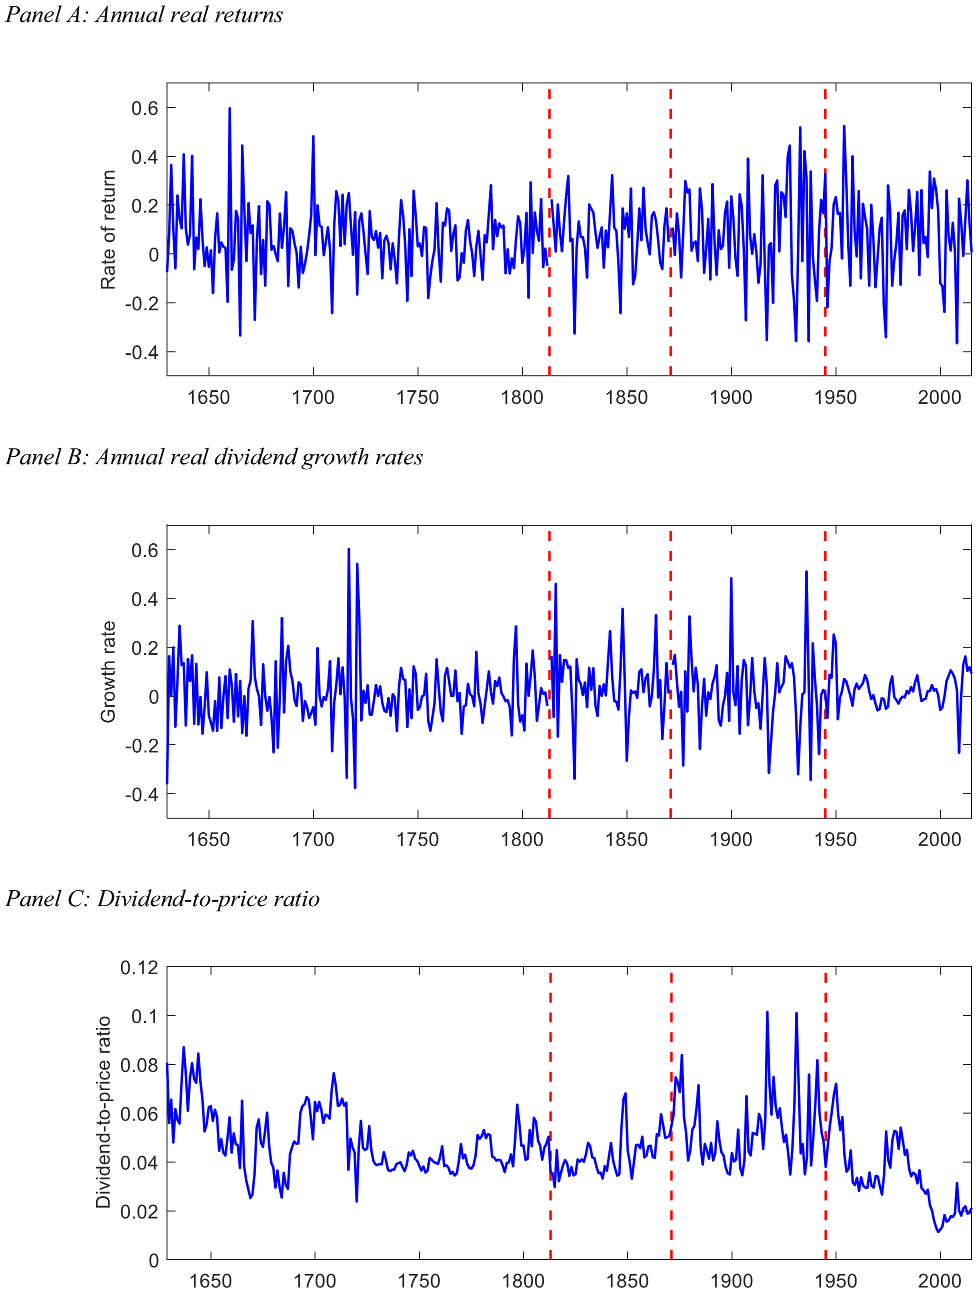

Fig. 1. Returns, dividend growth rates, and the dividend-to-price ratio**.** This figure plots real returns, real dividend growth rates, and the dividend-to-price ratio over time. The data are annual. Dashed vertical lines denote the time periods: Netherlands and UK (1629–1812), UK (1813–1870), US early (1871–1945), and US recent (1945–2015).

Table 1
Summary statistics.
This table reports summary statistics for the annual variables in nominal (Panel A) and real terms (Panel B). Column 1 reports the statistics for the 1629–1812 period based on the Netherlands and UK data. Annual dividend growth rates (DG) and the dividend-to-price ratio (DP) before 1700 are based on ten-year trailing averages of nominal or real dividends. Column 2 reports statistics for the 1700–1812 period separately. Column 3 reports the same statistics for the UK period, 1813–1870. Columns 4 and 5 present the statistics for the US before and after 1945. Column 6 reports the statistics based on the full sample. RET is annual returns. DY/RET is the ratio of the dividend yield (D_t/P_{t-1}) to returns. The risk-free rate (RF) is the return on government securities before 1870, the average return on commercial paper and government securities between 1871 and 1920, and the return on T-bills thereafter. The inflation rate comes historical consumer price indices. Details are in the Online Appendix. The Sharpe ratio is calculated assuming zero variation in the risk-free rate. We calculate standard errors for the AR (1) coefficients as 1/sqrt (T). The corresponding *t*-statistics are in parentheses below the AR (1) coefficients. We present augmented Dickey-Fuller (ADF) tests for the presence of a unit root. *** and * denote statistical significance of the ADF at the 1% and 10% levels, respectively.

	Netherlands and UK 1629–1812	Netherlands and UK 1700–1812	UK 1813–1870	US 1871–1945	US 1945–2015	Full period 1629–2015
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Nominal data						
RET (percent)	6.11	5.64	6.91	9.05	12.00	7.87
Standard deviation (percent)	10.94	9.31	8.90	19.76	16.87	14.11
AR(1)	−0.11	−0.05	−0.03	0.04	−0.04	−0.01
t-statistic	(−1.51)	(−0.57)	(−0.24)	(0.36)	(−0.30)	(−0.14)
DY/RET	0.79	0.86	0.63	0.59	0.30	0.59
RF (percent)	3.35	3.16	4.21	2.57	4.06	3.47
Risk Premium (percent)	2.76	2.48	2.69	6.48	7.94	4.40
Sharpe Ratio	0.25	0.27	0.30	0.33	0.47	0.31
Inflation (percent)	0.49	0.57	−0.67	0.73	3.79	0.96
DG (percent)	0.90	0.96	3.40	2.46	6.38	2.57
Standard deviation (percent)	11.32	10.74	11.24	15.44	6.81	11.72
AR(1)	−0.21	−0.24	−0.05	0.22	0.44	0.03
t-statistic	(−2.79)	(−2.57)	(−0.37)	(1.88)	(3.67)	(0.56)
Panel B: Real data						
RET (percent)	6.15	5.51	8.29	8.42	8.13	7.27
Standard deviation (percent)	13.36	11.06	12.20	19.11	17.28	15.20
AR(1)	−0.10	0.03	0.13	0.01	0.01	−0.02
t-statistic	(−1.37)	(0.27)	(0.99)	(0.10)	(0.12)	(−0.47)
DY/RET	0.78	0.87	0.53	0.63	0.43	0.63
DG (percent)	0.71	0.88	4.61	1.89	2.58	1.86
Standard deviation (percent)	12.75	12.74	13.28	15.09	7.01	12.53
AR(1)	−0.10	−0.10	−0.07	0.07	0.42	−0.01
t-statistic	(−1.33)	(−1.11)	(−0.52)	(0.59)	(3.50)	(−0.16)
DP (percent)	4.84	4.57	4.27	5.34	3.39	4.59
Standard deviation (percent)	1.24	0.96	0.79	1.45	1.41	1.42
AR(1)	0.78	0.77	0.65	0.47	0.90	0.78
ADF	−4.08***	−3.41*	−4.92***	−4.90***	−2.38	−6.28***

1813–1870 (UK), 1871–1945 (US early), and 1945–2015 (US recent).

Annual nominal returns are 8% on average and vary between 6% in the early data and 12% in the recent US period. Inflation, however, is much higher after 1945, and in real terms average returns across periods are more similar, varying between 6% and 8%. In all our main tests, we use real data. Returns are more volatile after 1870–the standard deviation of real returns is roughly 50% higher compared with the earlier periods. Throughout all centuries, returns exhibit low autocorrelation. Weak evidence exists for negative autocorrelation in the earliest period, possibly indicative of measurement error. However, the AR(1) coefficient is not statistically significant. Finally, the data indicate that, for the period as a whole, most of the real return to investors has come from dividend yields, with 37% coming from price appreciation. This has changed in the most recent period, when price appreciation accounts for 57% of real returns.

Our estimates indicate that the annual risk premium was relatively low in the early part of our data, 2–3%, in comparison to 6–8% in the US after 1870. These estimates need to be interpreted with caution, though, as we use a crude proxy for the risk-free rate before 1920. Also, because returns are more volatile in the later part of the data, Sharpe ratios differ much less across periods, although the recent US period (1945–2015) still has a Sharpe ratio that is relatively high at 0.47, compared with 0.31 for the sample as a whole.

Annual real dividend growth rates are around 2% on average. In the 17th and 18th centuries they are below 1%. They start to pick up in the 19th century, with an average growth rate of 5%, falling to about 2% in the later part of the data. The volatility of dividend growth rates changes significantly after 1945. Before, its standard deviation lies between 13% and 15%. After 1945, this drops by approximately half. At the same time, dividend growth becomes much more persistent. Initially, the AR (1) coefficient for real dividend growth is negative and statistically insignificant. For the period after 1945, it becomes positive at 0.42 and significant with a *t*-statistic of 3.50.

The dividend-to-price ratio is stationary over the sample as a whole, fluctuating around a long-term average of close to 5% – an augmented Dickey-Fuller test rejects a unit root at the 1% confidence level. Nevertheless, Fig. 1 reveals that dividend yields have been falling since 1950 and became more persistent, with an AR (1) coefficient of 0.90, compared with 0.78 for the sample as a whole. The recent period is the only subsample in which we cannot reject the presence of a unit root.

3. Results

In this section, we present the key results of the paper.

3.1. Preliminary evidence

We start by previewing the key results of the paper in Fig. 2. Panel A plots the dividend-to-price ratio and the subsequent five-year real return. If dividend yields predict returns, we would expect that the two lines tend to rise and fall together. Panel B indicates whether specific data points help in predicting returns. The figure is based on regressions predicting one-, three-, and five-year real log returns using the log dividend-to-price ratio. The final points on the graph match the corresponding full-sample R-squareds. The lines indicate where the fit improves or declines (details are in the caption of Fig. 2). Shaded areas indicate recessions (information on recessions is available only after 1700).

Panel A confirms that the dividend-to-price ratio tends to increase before a period of high returns (and tends to decrease before a period of low returns), suggesting that returns are predictable by the dividend-to-price ratio. This is also apparent from Panel B, which shows that the fit of the predictive model improves in all four subperiods. With the exception of the UK period (1813–1870), the fit of the model improves at longer horizons.

Panel A also suggests that return predictability from dividend yields is related to the business cycle. Both the dividend-to-price ratio and subsequent returns tend to increase in recessions. This pattern is strongest for the early US period. Panel B confirms that recession periods contribute significantly to the fit of the predictive model. For example, the Depressions of 1709–1712 ["among the worst of the century" according to Ashton (1959, p. 141)], 1873–1878 and 1929–1933 all lead to a higher R-squared. In comparison, World War I and World War II do not add to the fit of the model. The 1990s, when exceptionally high returns coincide with a decrease in the dividend-to-price ratio, do not either.

Panels C and D present the same information for dividend growth rates. Evidence exists for dividend yields predicting dividend growth rates in all but the final period, with the caveat that predictability in the Netherlands and UK period (1629–1812) is concentrated in the 17th century, for which the data are largely based on a single company (the VOC). Panel D indicates that the fit of the model improves during the two world wars, indicating that the higher dividend-to-price ratio in these periods reflected a change in expected growth rates. Panel D also suggests that the increase in dividend yields during the Great Depression of the 1930s partly reflected a downward revision of expected dividend growth rates. In contrast to the return predictability results, the full-sample fit of the model is similar across different horizons.

In the next subsections, we analyze these results more formally.

3.2. Main regressions

In Table 2, we report return and dividend growth predictability results over a one-year horizon and evaluate statistical significance using Newey and West (1987) *t*-statistics with one lag. The results show that in the sample as a whole, and in the four subperiods individually, dividend yields predict returns with a positive coefficient. This holds true when we take logs. This relation is highly statistically significant for the sample as a whole (*t*-statistic of 3.17, 2.81 when we take logs). For certain subperiods, *t*-statistics fall below two, in particular in the early US period (1870–1945), but the difference with the full sample coefficient estimates is economically small (especially for the log results) and not statistically different from zero.

A different picture emerges from the regressions of dividend growth rates on the dividend-to-price ratio. For the sample as a whole, dividend yields negatively predict dividend growth rates (*t*-statistic of −4.82, −4.09 when we take logs), but this is largely driven by the period before 1945. In the recent US period (1945–2015), little evidence exists that dividend growth rates can be predicted using the dividend-to-price ratio. The predictive coefficient is close to zero, and the difference with the rest of the sample is statistically significant (*t*-statistic of 3.10, 4.12 when we take logs).

In the Online Appendix we replicate Table 2 using different variable definitions and specifications. Table OA.1 uses nominal data, Table OA.2 uses the UK stock market until 1900 and broad-based CRSP index (rather than the S&P 500 index) from 1925 onward, and Table OA.3 adds lagged returns and dividend growth rates as additional predictors. Results are qualitatively similar to our main findings in Table 2. In Table OA.4, to remedy the effects of possible measurement error, we use a trailing two-year rolling average of the dividend-to-price ratio as a predictor. This leads to more return predictability in the early US period (1870–1945), while the evidence for dividend growth predictability in the Netherlands and UK period becomes weaker. The full sample estimates are virtually unchanged though.

3.3. Refined statistical tests

Nelson and Kim (1993), Stambaugh (1999) and others argue that return predictive coefficients reported in the previous section could overstate return predictability. In particular, a negative correlation exists between innovations in dividend yields and the errors in the predictive regression for returns, leading to an upward bias in the estimated coefficient. This bias is more pronounced when the negative correlation is stronger, the sample period is shorter, and the dividend-to-price ratio is highly persistent.

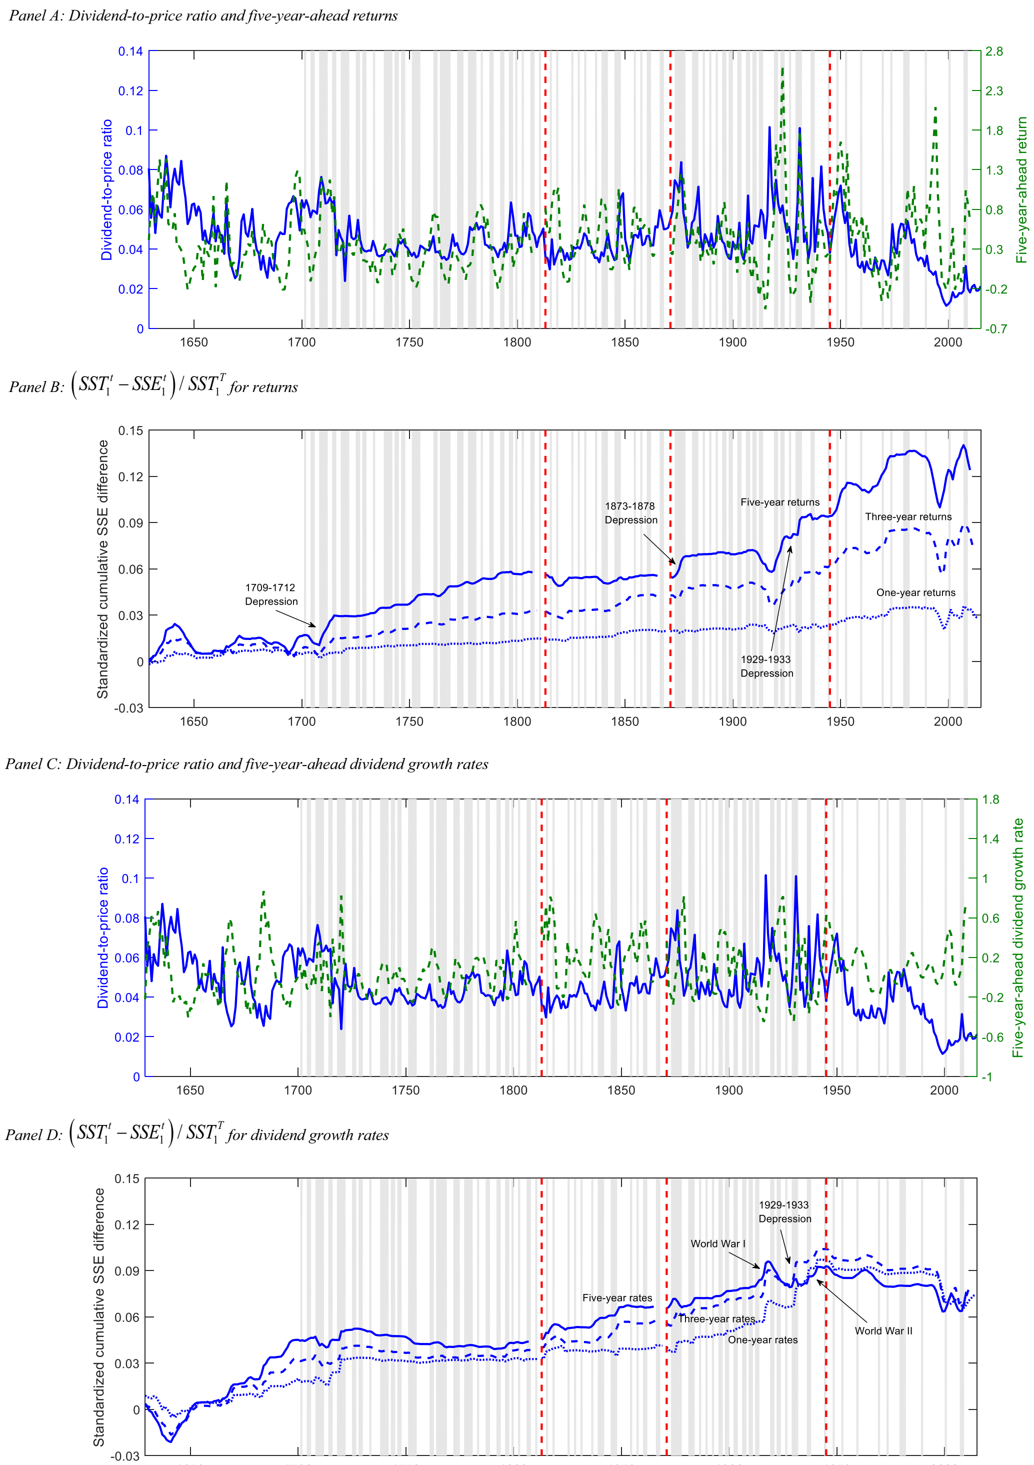

Fig. 2. Dividend-to-price ratio, recessions and in-sample predictive regressions. In Panels A and C, we plot the dividend-to-price ratio along with the five-year-ahead returns and dividend growth rates. Panels B and D track when the fit of the (full sample) predictive regression improves or declines. We plot the difference in the cumulative sum of squared errors between a model with a constant only and a model with a constant and the dividend-to-price ratio. The difference is normalized by the total sum of squared errors from the constant-only model, so that the last observation corresponds to the in-sample R-squared. Apart from five-year returns, we present this for one- and three-year returns. Each line in Panels B and D plots $(SST_1^t - SSE_1^t)/SST_1^T$, where $SST_1^t = \sum_{\tau=1}^{t}(y_\tau - \bar{y})^2$, $SSE_1^t = \sum_{\tau=1}^{t}(y_\tau - \hat{y}_\tau)^2$, and $SST_1^T = \sum_{\tau=1}^{T}(y_\tau - \bar{y})^2$, with y_τ the one-, three-, or five-year-ahead return (or dividend growth rate), $\hat{y}_\tau$ the return (dividend growth rate) predicted by a constant and dp_τ using coefficients estimated on the full sample, and $\bar{y}$ the full sample mean return (dividend growth rate). Dashed vertical lines denote the time periods: Netherlands and UK (1629–1812), UK (1813–1870), US early (1871–1945), and US recent (1945–2015). Shaded areas indicate recessions (data for recessions are available only after 1700).

Table 2
Return and dividend growth predictability.
This table reports ordinary least squares estimates of regressing annual real returns and dividend growth rates on the lagged dividend-to-price ratio. Lowercase letters are logs of corresponding capital letters. All regressions include a constant (not reported). Below the estimated coefficients (in parentheses) are Newey and West (1987) t-statistics with one lag. In brackets are the t-statistics for the difference of the estimated coefficient from the rest of the sample (based on a full-period regression with an interaction term).

	Netherlands and UK 1629–1812	UK 1813–1870	US 1871–1945	US 1945–2015	Full period 1700–2015	Full period 1629–2015
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: RET_{t+1}						
DP_t	2.72	4.32	1.50	3.14	2.00	1.96
t-statistic	(3.13)	(2.01)	(1.11)	(2.30)	(2.71)	(3.17)
Difference (t-statistic)	[0.76]	[1.08]	[−0.41]	[0.56]	[−0.23]	
R^2	0.06	0.08	0.01	0.07	0.03	0.03
Dependent variable: DG_{t+1}						
DP_t	−2.25	−3.86	−7.03	−0.11	−2.79	−2.64
t-statistic	(−2.36)	(−1.36)	(−7.80)	(−0.12)	(−4.45)	(−4.82)
Difference (t-statistic)	[0.44]	[−0.46]	[−4.99]	[3.10]	[−0.41]	
R^2	0.05	0.05	0.45	0.00	0.09	0.09
Dependent variable: ret_{t+1}						
dp_t	0.12	0.22	0.09	0.09	0.08	0.07
t-statistic	(2.97)	(2.21)	(1.14)	(2.03)	(2.47)	(2.81)
Difference (t-statistic)	[0.96]	[1.44]	[0.10]	[−0.29]	[−0.27]	
R^2	0.06	0.10	0.01	0.06	0.03	0.03
Dependent variable: dg_{t+1}						
dp_t	−0.12	−0.16	−0.42	−0.01	−0.10	−0.10
t-statistic	(−2.73)	(−1.26)	(−6.83)	(−0.40)	(−3.45)	(−4.09)
Difference (t-statistic)	[−0.64]	[−0.50]	[−5.39]	[4.12]	[0.37]	
R^2	0.06	0.05	0.48	0.00	0.07	0.07

In this subsection, we follow Cochrane (2008) to adjust test statistics for this issue and to develop additional tests.

Starting point for the analysis is a simple present value relation linking prices, returns and dividend growth rates. Define log return $ret_{t+1} = \log[(P_{t+1} + D_{t+1})/P_t]$, log dividend growth rate $dg_{t+1} = \log(D_{t+1}/D_t)$, and log dividend-to-price ratio $dp_t = \log(D_t/P_t)$. Using a first-order Taylor expansion around the long-run mean of the dividend-to-price ratio $\overline{dp}$, Campbell and Shiller (1988b) show that the dividend-to-price ratio can be approximated by

$$dp_t \simeq r_{t+1} - dg_{t+1} + \rho dp_{t+1}, \tag{1}$$

where all variables are demeaned and $\rho = \exp(-\overline{dp})/(1 + \exp(-\overline{dp}))$ is the linearization constant. Eq. (1) shows that a high dividend-to-price ratio is related to high future returns or low future dividend growth rates, or a high future dividend-to-price ratio, or some combination. This can be formulated in terms of the following three predictive regressions:

$$r_{t+1} \simeq \alpha_r + \beta_r dp_t + \varepsilon^r_{t+1}, \tag{2}$$

$$dg_{t+1} \simeq \alpha_{dg} + \beta_{dg} dp_t + \varepsilon^{dg}_{t+1}, \tag{3}$$

and

$$dp_{t+1} \simeq \alpha_{dp} + \beta_{dp} dp_t + \varepsilon^{dp}_{t+1}, \tag{4}$$

where the coefficients are linked by the approximate identity

$$\beta_r - \beta_{dg} + \rho\beta_{dp} \simeq 1 \tag{5}$$

and errors from Eqs. (2)–(4) are linked according to

$$\varepsilon^r_{t+1} - \varepsilon^{dg}_{t+1} + \rho\varepsilon^{dp}_{t+1} \simeq 0. \tag{6}$$

Eq. (5) implies that we can always infer the third predictive coefficient from the other two. In Table 3, we run all three predictive regressions and report both the direct and implied coefficients. The identity holds well empirically. In most periods, the direct and indirect estimates are almost identical. The exception is the earliest period, when the relative difference between estimates is approximately 15%. This is likely a result of the smoothing procedure we use for dividends before 1700. For the period 1700–2015 direct and implied coefficients are indistinguishable.

The return predictive coefficient is upward biased if the errors in Eqs. (2) and (4) are negatively correlated: $corr(\varepsilon^r_{t+1}, \varepsilon^{dp}_{t+1}) < 0$. Table 3 indicates that, in the recent period, this correlation is negative and high at −0.91. In the full sample, the correlation is less negative, but still substantial at −0.67. Following Cochrane (2008), we use Monte Carlo simulations to construct test-statistics that take the resulting bias into account. We simulate the data using the empirical estimate for β_{dp} and impose $\beta_{ret} = 0$. β_{dg} follows from the restriction in Eq. (5). We rely on the sample covariance matrix of ε^{dp} and ε^{dg}, and we let ε^{ret} follow from Eq. (6). The null hypothesis of no return predictability corresponds to the following system of equations:

$$\begin{pmatrix} ret_{t+1} \\ dg_{t+1} \\ dp_{t+1} \end{pmatrix} = \begin{pmatrix} 0 \\ \rho\beta_{dp} - 1 \\ \beta_{dp} \end{pmatrix} dp_t + \begin{pmatrix} \varepsilon^{dg}_{t+1} - \rho\varepsilon^{dp}_{t+1} \\ \varepsilon^{dg}_{t+1} \\ \varepsilon^{dp}_{t+1} \end{pmatrix}. \tag{7}$$

We simulate 50 thousand data sets matching the length of the sample period. For each data set, we then re-estimate Eqs. (2)–(4). This yields a distribution of β_{ret} coefficients

Table 3
Return and dividend growth predictability: additional tests.

This table reports ordinary least squares estimates of regressing annual real returns, dividend growth rates, and the dividend-to-price ratio on the lagged dividend-to-price ratio. All variables are in logs. We calculate implied coefficients using the identity $\beta_{ret} - \beta_{dg} + \rho\beta_{dp} = 1$ and $\rho = \exp(-\overline{dp})/(1+\exp(-\overline{dp}))$. The correlation of errors is the correlation between innovations in the dividend-to-price ratio and errors in the predictive regression for returns or dividend growth rates. The p-values are based on Newey and West (1987) t-statistics with one lag, p-value (N-W), or Monte Carlo simulations, p-value (Sim.), where Direct tests for the presence of return or dividend growth predictability, and Implied infers return (dividend growth rate) predictability from the lack of dividend growth (return) predictability. The long-run coefficient is implied from the short-run coefficient using $\beta_x^{lr} = \beta_x/(1-\rho\beta_{dp})$. The p-values for the long-run coefficient are based on the Delta method, p-value (Delta m.), and Monte Carlo simulations, p-value (Sim., Direct). All regressions include a constant (not reported).

	Netherlands and UK 1629–1812	UK 1813–1870	US 1871–1945	US 1945–2015	Full period 1700–2015	Full period 1629–2015
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: ret_{t+1}						
dp_t	0.12	0.22	0.09	0.09	0.08	0.07
Implied coefficient	0.14	0.22	0.08	0.09	0.08	0.08
Correlation of errors	−0.48	−0.46	−0.85	−0.91	−0.72	−0.67
p-value (N-W)	.00	.03	.26	.05	.01	.01
p-value (Sim., Direct)	.00	.03	.24	.24	.01	.00
p-value (Sim., Implied)	.00	.00	.04	.00	.00	.00
Long-run coefficient	0.46	0.57	0.17	0.90	0.43	0.41
p-value (Delta m.)	.00	.01	.24	.00	.00	.00
p-value (Sim., Direct)	.00	.00	.19	.00	.00	.00
Dependent variable: dg_{t+1}						
dp_t	−0.12	−0.16	−0.42	−0.01	−0.10	−0.10
Implied coefficient	−0.14	−0.16	−0.42	−0.01	−0.10	−0.11
Correlation of errors	0.62	0.59	0.44	0.23	0.53	0.54
p-value (N-W)	.01	.21	.00	.69	.00	.00
p-value (Sim., Direct)	.01	.11	.00	.40	.00	.00
p-value (Sim., Implied)	.00	.02	.00	.14	.00	.00
Long-run coefficient	−0.47	−0.43	−0.83	−0.10	−0.55	−0.55
p-value (Delta m.)	.00	.05	.00	.58	.00	.00
p-value (Sim., Direct)	.00	.04	.00	.32	.00	.00
Dependent variable: dp_{t+1}						
dp_t	0.77	0.65	0.52	0.93	0.86	0.85
Implied coefficient	0.79	0.64	0.52	0.93	0.86	0.86
p-value (N-W)	.00	.00	.00	.00	.00	.00

that we use to evaluate the statistical significance of our empirical estimates.

The p-value (Sim., Direct) in Table 3 is the fraction of return coefficients from Monte Carlo simulations that are larger than the empirical estimate. Consistent with Stambaugh (1999) and Cochrane (2008), once the bias is accounted for, only weak evidence exists for return predictability in the recent U.S period. The simulated p-value is .24 in comparison with .05 based on the Newey-West standard error. In the early US period, both p-values are insignificant. For the earlier samples, however, the simulated and Newey-West p-values are virtually the same and close to zero, consistent with a lower correlation of errors. Most important, for the sample as a whole, the simulated p-value is again close to zero. Thus, the long time period, in combination with a lower correlation of errors and less persistence in the dividend-to-price ratio, yields sufficient statistical power to strongly reject the null of no return predictability.

By the same logic, the dividend growth coefficient is downward-biased if the errors in Eqs. (3) and (4) are positively correlated: $corr(\varepsilon_{t+1}^{dg}, \varepsilon_{t+1}^{dp}) > 0$. Table 3 indicates that the correlation is relatively low in the recent period at 0.23, but it gradually increases to 0.62 going back in time. We thus run similar simulations for dividend growth rates, except that we now impose $\beta_{dg} = 0$ and let β_{ret} follow from Eq. (5). Also, differently from before, we now rely on the sample covariance matrix of ε^{dp} and ε^{ret}, and we infer ε^{dg} from Eq. (6). The p-values (Sim., Direct) in Table 3 are the fraction of dividend growth coefficients from Monte Carlo simulations that are smaller than the empirical estimate. These p-values are qualitatively similar to the ones calculated from the Newey-West standard errors. The dividend-to-price ratio generally predicts dividend growth rates before 1945, but not afterward.

Cochrane (2008) shows that one can also test for the predictability of returns by looking at (the absence of) dividend growth predictability. The intuition is that if the dividend-to-price ratio is stationary, it should predict either returns or dividend growth rates, or both. If the forecastability of dividend growth rates is weak, dividend yields have to predict returns. Using the distribution of coefficients simulated under the null of no return predictability (Eq. (7)), we report the fraction of dividend growth coefficients from Monte Carlo simulations that are larger than the estimate in Table 3 (p-value Sim., Implied). The equivalent p-value for dividend growth rates is the fraction of return coefficients simulated under the null of no dividend growth predictability that are smaller than their empirical estimate. Consistent with Cochrane (2008), we strongly reject the null that returns are not predictable for each individual period, including the recent US period. Accordingly, the p-value associated with the full sample estimate of β_{ret} remains close to zero. This approach also strengthens the evidence for dividend growth predictability, with p-values close to zero for all but the recent period in which we reject it with a p-value of .14.

Finally, we calculate the long-run estimates implied by the coefficients from Table 3. Iterating Eq. (1) forward and excluding rational bubbles, the dividend-to-price ratio can be expressed as an infinite sum of discounted returns and dividend growth rates (because the relation holds ex-ante and ex-post, an expectations operator can be added):

$$dp_t \simeq E_t \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} - E_t \sum_{j=1}^{\infty} \rho^{j-1} dg_{t+j}. \tag{8}$$

Thus, ultimately, any variation in the dividend-to-price ratio is related to future changes in expected returns or expected dividend growth rates, or a combination of both. We can express Eq. (8) in terms of regression coefficients [details are in Cochrane (2008)]:

$$\beta_r^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} \beta_{dp}^{j-1} \beta_r = \frac{\beta_r}{1-\rho\beta_{dp}}, \tag{9}$$

and

$$\beta_{dg}^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} \beta_{dp}^{j-1} \beta_{dg} = \frac{\beta_{dg}}{1-\rho\beta_{dp}}, \tag{10}$$

where $\beta_r^{lr} - \beta_{dg}^{lr} \simeq 1$. The two long-run coefficients capture the fraction of the variance of the dividend-to-price ratio that can be attributed to time-varying expected returns or dividend growth rates.

Table 3 reports the two long-run coefficients and corresponding p-values, derived either from Newey-West standard errors (using the Delta method) or simulations. Because we use direct estimates and do no impose the restriction from Eq. (5), the two coefficients do not necessarily add up to one. With the exception of the early US period, dividend yields significantly predict returns. Dividend growth rates are predictable until 1945. The p-values for the full sample estimates are close to zero in both cases. The long-run estimates suggest that changes in expected returns account for about 40% of the variation of the dividend-to-price ratio in the full sample [around 45%, if we exclude the years before 1700 when identity (5) holds only approximately]. The remainder, or 55% of price variation, is driven by changes in expected dividend growth rates.

3.4. Long horizon predictability

The previous long-run estimates assume that the underlying Vector Auto Regression (VAR) model is correctly specified. In this subsection, we use the data to directly test whether dividend yields predict returns over longer horizons. Fama and French (1988) show that such an approach strengthens the evidence for return predictability. Boudoukh et al., (2008) argue that the improved fit of predictive model is largely mechanical and is not statistically significant.

Table 4
Long-horizon predictability.
This table reports ordinary least squares estimates of regressing sum of annual real returns ($\sum_{j=1}^{h} ret_{t+j}$) or dividend growth rates ($\sum_{j=1}^{h} dg_{t+j}$) on the dividend-to-price ratio. All variables are in logs. Horizon h is either one, three, or five years. Below the estimated coefficients (in parentheses) are Newey and West (1987) t-statistics with the number of lags equal to the length of the horizon. In brackets are t-statistics based on non-overlapping observations, calculated as the mean across alternative non-overlapping samples (e.g., in case of five-year predictions, we report the mean across five different non-overlapping samples starting in years one through five). The p-values (Sim., Direct) are based on Monte Carlo simulations. The implied coefficient for longer horizon predictions is based on the one-year coefficient and calculated using $\beta_{x,h} = \beta_x(1-\beta_{dp,1}^h)/(1-\beta_{dp,1})$. All regressions include a constant (not reported).

	1629–1945	1945–2015	1700–2015	1629–2015
	(1)	(2)	(3)	(4)
Dependent variable: 1-year *ret*				
dp_t	0.11	0.09	0.08	0.07
t-statistic	(3.22)	(2.03)	(2.47)	(2.81)
p-value (Sim., Direct)	.00	.24	.01	.00
R^2	0.04	0.06	0.03	0.03
Dependent variable: 3-year *ret*				
dp_t	0.29	0.25	0.21	0.20
t-statistic (Overlap.)	(4.40)	(2.84)	(3.12)	(3.41)
t-statistic (Non-overlap.)	[3.44]	[2.83]	[2.81]	[3.01]
p-value (Sim., Direct)	.00	.25	.01	.00
Implied coefficient	0.24	0.26	0.20	0.19
p-value (Sim., Direct)	.00	.20	.00	.00
R^2	0.10	0.15	0.08	0.07
Dependent variable: 5-year *ret*				
dp_t	0.44	0.44	0.34	0.32
t-statistic (Overlap.)	(4.51)	(5.25)	(4.02)	(4.29)
t-statistic (Non-overlap.)	[3.46]	[3.70]	[2.99]	[3.17]
p-value (Sim., Direct)	.00	.20	.00	.00
Implied coefficient	0.31	0.40	0.29	0.28
p-value (Sim., Direct)	.00	.15	.00	.00
R^2	0.15	0.28	0.13	0.12
Dependent variable: 1-year *dg*				
dp_t	−0.20	−0.01	−0.10	−0.10
t-statistic	(−5.30)	(−0.40)	(−3.45)	(−4.09)
p-value (Sim., Direct)	.00	.40	.00	.00
R^2	0.14	0.00	0.07	0.07
Dependent variable: 3-year *dg*				
dp_t	−0.33	0.01	−0.15	−0.18
t-statistic (Overlap.)	(−4.82)	(0.09)	(−2.59)	(−3.29)
t-statistic (Non-overlap.)	[−3.79]	[0.10]	[−2.19]	[−2.75]
p-value (Sim., Direct)	.00	.62	.01	.00
Implied coefficient	−0.44	−0.03	−0.25	−0.26
p-value (Sim., Direct)	.00	.39	.00	.00
R^2	0.15	0.00	0.06	0.08
Dependent variable: 5-year *dg*				
dp_t	−0.35	0.00	−0.15	−0.20
t-statistic (Overlap.)	(−3.58)	(−0.04)	(−2.48)	(−3.12)
t-statistic (Non-overlap.)	[−2.60]	[−0.04]	[−1.63]	[−2.23]
p-value (Sim., Direct)	.00	.57	.08	.02
Implied coefficient	−0.57	−0.04	−0.37	−0.38
p-value (Sim., Direct)	.00	.38	.00	.00
R^2	0.12	0.00	0.05	0.08

In Table 4, we directly estimate the predictability of three- and five-year real log returns and dividend growth rates. For comparison, we also report the one-year results. We report t-statistics based on overlapping and non-overlapping data. In the non-overlapping case, we run three or five different regressions using alternative samples of non-overlapping observations and report the average t-statistic across those samples. In addition, we report

Monte Carlo p-values (Sim., Direct). For comparison with the direct estimates, we report the three- and five year coefficients implied from the annual VAR, using $\beta_{x,h} = \beta_x(1-\beta_{dp,1}^h)/(1-\beta_{dp,1})$, where h is the length of the horizon and x stands for returns or dividend growth rates.

The evidence for dividend yields predicting returns is strong across the different horizons. When we increase the horizon from one to five years, the estimated coefficients increase monotonically and the full-sample R-squared increases from 3% to 12%. The t-statistics based on non-overlapping observations increase from 2.81 to 3.17. Simulated p-values are always close to zero. Also, the three- and five-year coefficients are close to those implied by the VAR. All in all, the evidence for return predictability is largely aligned across horizons. When we break the sample in 1945, the estimates for 1629–1945 and 1945–2015 are quantitatively similar.

The dividend growth results are somewhat different. The coefficients increase when we move from a one- to three-year horizon, but relatively less so than for returns. Moving from three to five years leaves the coefficients virtually unchanged. The full-sample R-squared is approximately the same across all horizons. All coefficients are statistically significant. Simulated p-values remain close to zero, but the t-statistics based on non-overlapping observations decrease with the increase of the horizon, from -4.09 to-2.23. The long horizon coefficients are also substantially smaller (in absolute value) than those implied from the annual VAR. This suggests that expected dividend growth rates contain some mean-reverting component that the VAR does not capture. It also means that the implied long-run coefficient β_{dg}^{lr} from Table 3 likely overstates the importance of expected dividends for price movements. The full-sample five-year coefficient on dividend growth implies an annual coefficient that is half the size of the actual coefficient (-0.05, rather than -0.10). It also suggests that around 60% (rather than 40%) of the variation in the dividend-to-price ratio can be attributed to discount rate news. Consistent with earlier results, we find no evidence that dividend yields predict dividend growth rates after 1945.

Results are similar if we discount future returns and dividend growth rates by the linearization constant ρ (Online Appendix, Table OA.5).

3.5. Out-of-sample predictability

Goyal and Welch (2003, 2008) show that the dividend-to-price ratio is a poor out-of-sample predictor of returns compared with simply using the historical average return. The out-of-sample R-squared (ROOS) is typically negative. This does not necessarily mean that expected returns do not vary over time. Simulations in Cochrane (2008) indicate that even when all variation in the dividend yield comes from changes in expected returns, out-of-sample R-squareds are seldom positive. The critique does suggest, however, that dividend yields perhaps do not help predicting returns in real time.

In this subsection, we present estimates for out-of-sample predictability in our data. Following Goyal and Welch (2008), we calculate the out-of-sample R-squared as

$$\text{ROOS} = 1 - \frac{\sum_{\tau=1}^{T}(y_\tau - \hat{y}_\tau)^2}{\sum_{\tau=1}^{T}(y_\tau - \bar{y}_\tau)^2}, \tag{11}$$

where y_τ is the actual return, $\hat{y}_\tau$ is the return predicted by the dividend-to-price ratio using a coefficient estimated on the sample up to $\tau-1$, and $\bar{r}_\tau$ is the mean return up to $\tau-1$. In our main run, the first training sample is from 1629 through 1700, with the first out-of-sample prediction in 1701. The final prediction is based on the coefficient estimated on data from 1629 to 2014. Besides one year returns, we calculate ROOS for predicting three- and five-year returns. We use real returns and express all the variables in logs. We assess the statistical significance of our results with a Clark and West (2007) t-statistic. We also report the fraction of ROOSs from Monte Carlos simulations that are larger than the empirical estimate of ROOS (p-value, Sim.).

Table 5 reports the results. The ROOS is always positive and increases with the return horizon, between 2% and 10%. This compares with 3% and 13% for the in-sample R-squareds for this period (Table 4). According to the statistical tests, the out-of-sample predictions based on the dividend-to-price ratio present a significant improvement over the historical mean.

In Fig. 3, analogous to Fig. 2, Panel B, we plot increases or declines in ROOS. The patterns in Fig. 3 are similar to the in-sample statistics reported in Fig. 2, Panel B. Overall, strong evidence exists for return predictability at both the annual and multi-annual horizons, with the possible exception of the UK period (1813–1870) for longer horizons and the most recent years of the sample.

We also consider two adjustments proposed by Campbell and Thompson (2008). In the second column of Table 5, we set the predictive coefficient on the dividend-to-price ratio equal to zero if the regression generates a negative estimate. ROOS is identical to the first column, indicating that this scenario never materializes. In the third column, we set return forecast to zero if the model predicts it to be negative. This happens occasionally in the recent US period, when the dividend-to-price ratio is very low. As a result, ROOS increases from 2% to 3% for one-year returns and from 10% to 11% for five-year returns.

Finally, we address the concern that out-of-sample evidence could depend on the exact sample split (Rossi and Inoue, 2012; Kolev and Karapandza, 2017). Instead of letting the first training sample end in 1700, we consider all the possible splits with at least 20% and at most 80% of observations used for the initial training period. In Table 5, we report the average ROOS across those sample splits. The out-of-sample R-squared decreases somewhat. This is consistent with Fig. 3 that suggests that the out-of-sample evidence is somewhat weaker in the more recent data. However, ROOS remains positive and statistically significant, at least at the 10% level.

4. Business cycle variation

The previous results indicate that dividend yields predict returns and, therefore, that discount rates vary over

Table 5
Out-of-sample return predictability.
This table reports out-of-sample R-squared (ROOS) and Clark and West (2007) t-statistics for out-of-sample predictions of one-, three-, and five- year real returns. All variables are in logs. The p-values (Sim.) are from Monte Carlo simulations. All predictions are based on an expanding window. For the sample split in 1700, the initial training period is 1629–1700, and the rest of the sample period 1701–2015 is used for the calculation of the out-of-sample statistics. For the many sample splits, we report average statistics across all the sample splits with at least 20% and at most 80% of observations used in the estimation of initial parameters. The unconstrained ROOS compares the forecast errors of the historical mean against the forecasts from the dividend-to-price ratio predictive regression. Positive coefficient imposes a restriction that the coefficient on the dividend-to-price ratio is positive, or else the historical mean is used as a forecast. Positive forecast requires that the forecast is positive, or else zero is used as a forecast.

	Sample split in 1700			Many sample splits (20%–80%)		
	Unconstrained	Positive coefficient	Positive forecast	Unconstrained	Positive coefficient	Positive forecast
	(1)	(2)	(3)	(4)	(5)	(6)
One-year predictions						
ROOS	0.02	0.02	0.03	0.01	0.01	0.02
t-statistic	(2.37)			(1.72)		
p-value (Sim.)	.00					
Three-year predictions						
ROOS	0.06	0.06	0.06	0.03	0.03	0.04
t-statistic	(2.57)			(1.89)		
p-value (Sim.)	.00					
Five-year predictions						
ROOS	0.10	0.10	0.11	0.07	0.07	0.08
t-statistic	(2.97)			(2.28)		
p-value (Sim.)	.00					

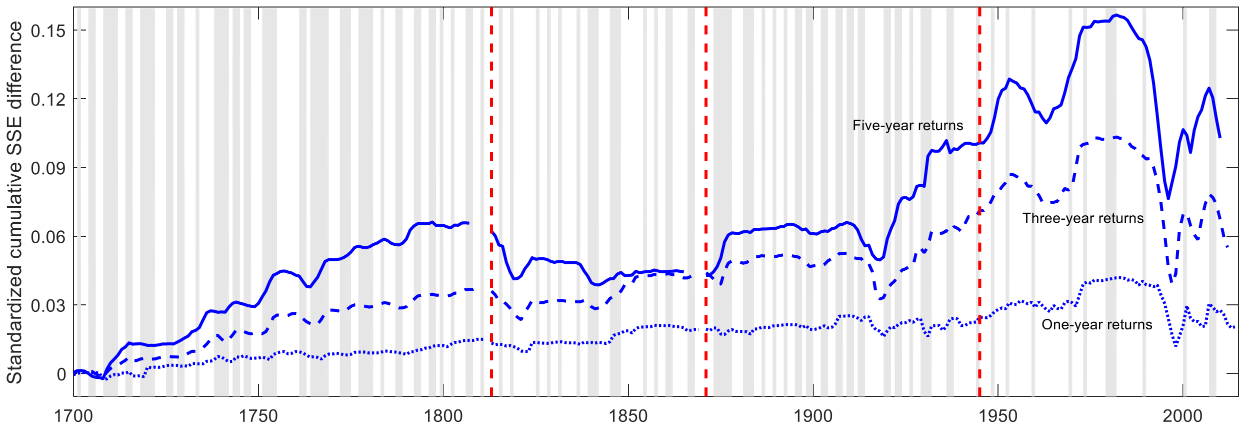

Fig. 3. Out-of-sample return predictive regressions**.** This figure is based on out-of-sample predictive regressions for one-,three-, and five-year returns. Out-of-sample predictions are based on an expanding window. The initial training period is 1629 – 1700 and the first out-of-sample prediction is in 1701. Each line plots $(SST_1^t - SSE_1^t)/SST_1^T$, where $SST_1^t = \sum_{\tau=1}^{t} (y_\tau - \bar{y}_\tau)^2$, $SSE_1^t = \sum_{\tau=1}^{t} (y_\tau - \hat{y}_\tau)^2$ and $SST_1^T = \sum_{\tau=1}^{T} (y_\tau - \bar{y}_\tau)^2$, with y_τ the one-, three-, or five-year-ahead return, $\hat{y}_\tau$ the return predicted by a constant and dp_τ using coefficients estimated on the sample up to $\tau - 1$, and $\bar{y}_\tau$ the mean return up to $\tau - 1$. The last observation (T) corresponds to the out-of-sample R-squared. Shaded areas denote recessions. Dashed vertical lines denote the time periods: Netherlands and UK (1629–1812), UK (1813–1870), US early (1871–1945), and US recent (1945–2015).

time. What drives these fluctuations in expected returns? One explanation is that the market has less risk-bearing capacity in downturns, giving rise to countercyclical discount rates (Fama and French, 1989). In this section, we study to what degree return predictability and business cycle fluctuations are related.

Table 6 provides summary statistics related to the business cycle. Because recession dates are available only from 1700 onward, results are based on the period 1700–2015. We also look at the subperiods 1700–1945 and 1945–2015 separately. A total of 62 individual recession events are in the data, ten of which take place after 1945. Realized returns are lower in recessions than in expansions. The same holds for dividend growth rates, but the difference is statistically significant only for returns. A visual inspection of Fig. 2, Panels A and C, suggests that the dividend-to-price ratio tends to peak in recessions. Table 6 confirms this pattern. The average dividend yield is 4.86% in downturns and 4.16% in expansions. The difference is statistically significant with a t-statistic of 4.58. Looking at 1700–1945 and 1945–2015 separately yields broadly similar results. In line with Fama and French (1989), these results suggest that the predictability of returns from dividend yields stems, at least in part, from business cycle fluctuations. During recessions, prices fall and expected returns go up.

To explore this further, we include a recession dummy as an additional predictor variable in the forecasting regressions:

Table 6
Business cycle variation: summary statistics.
This table reports summary statistics for real returns, real dividend growth rates, and the dividend-to-price ratio over the business cycle. From 1870 through 2015, recession and expansion dates come from the National Bureau of Economic Research (NBER). From 1700 until 1870, these dates are based on secondary sources that use NBER criteria to determine business cycle peaks and troughs in the UK (details are in the Online Appendix). In brackets are paired sample t-tests for the difference in mean.

	1700–1945			1945–2015			1700–2015		
	Recession		Expansion	Recession		Expansion	Recession		Expansion
	(1)		(2)	(3)		(4)	(5)		(6)
RET (percent)	4.88		8.78	−2.88		10.41	4.10		9.27
Standard deviation (percent)	14.92		13.47	19.74		15.99	15.55		14.25
t-statistic		[−2.12]			[−2.19]			[−2.95]	
DG (percent)	1.85		2.24	−0.47		3.21	1.62		2.53
Standard deviation (percent)	14.76		12.73	10.90		5.85	14.40		11.11
t-statistic		[−0.21]			[−1.14]			[−0.59]	
DP (percent)	4.94		4.56	4.14		3.22	4.86		4.16
Standard deviation (percent)	1.28		1.05	1.60		1.34	1.33		1.29
t-statistic		[2.48]			[1.85]			[4.58]	
N	108		135	12		58	120		193

$$r_{t+1} \simeq \alpha_r + \beta_r dp_t + \delta_r Recession_t + \varepsilon^r_{t+1}, \tag{12}$$

$$dg_{t+1} \simeq \alpha_{dg} + \beta_{dg} dp_t + \delta_{dg} Recession_t + \varepsilon^{dg}_{t+1}, \tag{13}$$

and

$$dp_{t+1} \simeq \alpha_{dp} + \beta_{dp} dp_t + \delta_{dp} Recession_t + \varepsilon^{dp}_{t+1}. \tag{14}$$

Similar to before, the coefficients in Eqs. (12)–(14) are linked by the approximate identities $\beta_r - \beta_{dg} + \rho\beta_{dp} \simeq 1$ and $\delta_r - \delta_{dg} + \rho\delta_{dp} \simeq 0$. We are particularly interested in how β_r, the coefficient on the dividend-to-price ratio in the return regression, changes after conditioning on the business cycle. If it falls, it would suggest that the business cycle is one of the channels through which dividend yields predict returns. A caveat here is that the dating of recessions happens ex post, relying on turning points in important macroeconomic series. Information about the start or end of a recession is therefore not necessarily available in real time. Also, a simple classification of contractions versus expansions does not capture the intensity of business cycles and misses long-term changes in business conditions (Fama and French, 1989).

Results are in Table 7. The approximate identities hold well in the data. Recessions are associated with higher expected returns. When we condition on the business cycle, β_r decreases. In the full sample it changes from 0.08 to 0.06. This confirms that dividend yields predict returns because, in part, prices are lower in recessions. This explains around 25% of the annual return predictability from dividend yields. Recessions are also associated with a lower dividend-to-price ratio next period, consistent with expected returns falling when the economy recovers. Controlling for the business cycle makes dividend yields more persistent. β_{dp} increases from 0.86 to 0.88. Finally, recessions appear unrelated to next period annual dividend growth rates, and coefficient β_{dg} in the dividend growth regression is unchanged at −0.10. Findings are similar for the subperiod 1700–1945. For 1945–2015, conditioning on the business cycle has only a limited effect on β_r.

Next, we explore whether the inclusion of a recession dummy in the predictive regressions changes the long-run coefficients on dividend yields. Following Cochrane (2011), we complete the VAR system by estimating a predictive regression for recessions:

$$Recession_{t+1} \simeq \alpha_{Rec} + \beta_{Rec} dp_t + \delta_{Rec} Recession_t + \varepsilon^{Rec}_{t+1}, \tag{15}$$

and we calculate the long-run coefficients from Eqs. (12)–(15) as

$$\begin{bmatrix} \beta^{lr}_r & \delta^{lr}_r \\ \beta^{lr}_{dg} & \delta^{lr}_{dg} \end{bmatrix} = \frac{B}{(I - \rho\phi)}, \quad \text{where } B = \begin{bmatrix} \beta_r & \delta_r \\ \beta_{dg} & \delta_{dg} \end{bmatrix} \text{ and }$$
$$\phi = \begin{bmatrix} \beta_{dp} & \delta_{dp} \\ \beta_{Rec} & \delta_{Rec} \end{bmatrix}. \tag{16}$$

Long-run coefficients are linked approximately by $\beta^{lr}_r - \beta^{lr}_{dg} \simeq 1$ and $\delta^{lr}_r - \delta^{lr}_{dg} \simeq 0$. Results are in Table 7.

Despite the ability of recessions to predict annual returns, the associated long-run coefficients δ^{lr}_r and δ^{lr}_{dg} are small. In addition, the inclusion of a recession dummy leads only to a small change in the long-run coefficients on the dividend-to-price ratio β^{lr}_r. This is largely due to the combination of two facts. Recessions or expansions have little persistence. They are also not a strong predictor of future dividend yields. Business cycles therefore mainly affect the short end of discount rates' term structure (Cochrane, 2011). Dividend yields are much more persistent and capture long-term variation in business conditions that goes beyond a simple classification of contractions versus expansions (Fama and French, 1989).

5. Summary and discussion

The results in this paper indicate that return predictability from dividend yields has been a robust characteristic of financial markets over the last four centuries. Our findings are robust to a number of statistical tests proposed in the literature, including Monte Carlo simulations that take the Stambaugh (1999) bias into account. The implied variation in expected returns lines up well with the business cycle, with on average high returns following downturns. This is true for both early and more recent data. Finally, we find that roughly until the mid-20th

Table 7
Business cycle variation: predictability results.
This table reports ordinary least squares estimates of regressing annual real returns, dividend growth rates, the dividend-to-price ratio, and the recession dummy on the lagged dividend-to-price ratio and the lagged recession dummy. The main variables are in logs. Recession takes a value one for recessions, and zero otherwise. Below the estimated coefficients (in parentheses) are Newey and West (1987) t-statistics with one lag. All regressions include a constant (not reported). The long-run coefficients are implied from the short-run coefficients using $B^{lr} = B/(I - \rho\Phi_{dp})$, where B is a matrix of predictive coefficients from the return and dividend growth regressions and Φ is a matrix of predictive coefficients from the dividend-to-price ratio and the recession regressions.

	1700–1945		1945–2015		1700–2015	
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: ret_{t+1}						
dp_t	0.12	0.10	0.09	0.09	0.08	0.06
t-statistic	(2.73)	(2.24)	(2.03)	(1.85)	(2.47)	(1.75)
$Recession_t$		0.06		0.03		0.06
t-statistic		(3.64)		(0.59)		(3.17)
R^2	0.04	0.09	0.06	0.06	0.03	0.06
Dependent variable: dg_{t+1}						
dp_t	−0.24	−0.24	−0.01	0.00	−0.10	−0.10
t-statistic	(−5.18)	(−5.15)	(−0.40)	(−0.04)	(−3.45)	(−3.24)
$Recession_t$		0.00		−0.05		0.00
t-statistic		(0.19)		(−1.73)		(−0.30)
R^2	0.16	0.16	0.00	0.07	0.07	0.07
Dependent variable: dp_{t+1}						
dp_t	0.67	0.69	0.93	0.94	0.86	0.88
t-statistic	(11.76)	(11.99)	(20.23)	(19.52)	(22.24)	(22.37)
$Recession_t$		−0.06		−0.08		−0.06
t-statistic		(−2.97)		(−1.28)		(−2.90)
R^2	0.46	0.48	0.85	0.86	0.73	0.74
Dependent variable: $Recession_{t+1}$						
dp_t		0.13		0.10		0.23
t-statistic		(0.86)		(0.94)		(2.91)
$Recession_t$	0.12	0.11	0.07	0.05	0.16	0.12
t-statistic	(2.34)	(2.06)	(0.56)	(0.38)	(3.25)	(2.35)
R^2	0.02	0.02	0.01	0.02	0.03	0.05
Long-run coefficient: ret_{t+1}						
dp_t	0.34	0.31	0.90	0.94	0.43	0.41
$Recession_t$		0.05		−0.04		0.03
Long-run coefficient: dg_{t+1}						
dp_t	−0.67	−0.70	−0.10	−0.06	−0.55	−0.57
$Recession_t$		0.05		−0.04		0.03

century the dividend-to-price ratio also predicts dividend growth rates, especially at a shorter (annual) horizon.

These results are supportive of modern asset pricing models that incorporate time-varying expected returns. At the same time, the failure of the dividend-to-price ratio to predict dividend growth rates in the recent period poses a puzzle (Cochrane, 2011). What changed after World War II so that movements in dividend yields stopped reflecting expected dividends?

We leave this for future research. A possible explanation lies in the observation that, as a fraction of earnings, firms have substantially reduced their dividends in the recent period (Fama and French, 2001). To illustrate this, Fig. 4 tracks the development of the dividend-to-earnings ratio for the Netherlands and UK period (1651–1812), and the US period (1871–2015) for which the data are available. We sum dividends and earnings over 20-year (trailing) periods and take the ratio. During the 17th and 18th centuries, firms paid out close to a 100% of their earnings to shareholders. In 1945, this number was still around 80%. By 1982, however, the dividend-to-earnings ratio had fallen to approximately 45%. (The payout ratio fell even more after 1982, but this is partly related to the increased use of repurchases.) This pattern is consistent with Fig. 1 and Table 1 that show that the dividend-to-price ration has decreased in the recent period and that returns to investors have increasingly come from capital appreciation in lieu of dividends.

The growing disconnect between earnings and dividends can reduce the ability of dividend yields to predict dividend growth rates in at least two ways. First, it enables firms to smooth dividends more, and, as argued by Chen et al., (2012), this can make changes in dividend yields less informative about future growth rates. Fig. 1, Panel B shows that dividends were least volatile in the recent US period. Second, a lower dividend-to-earnings ratio implies that firms push the eventual payouts to shareholders into the future. Given that expected returns appear to be more persistent than expected dividend growth rates, postponing dividend payments increases the relative sensitivity of dividend yields to shocks in expected returns and reduces the ability of dividend yields to predict dividend growth rates.[7]

[7] For evidence that expected returns are more persistent than dividend growth rates, see Fama and French (1988), Binsbergen and Koijen (2010), Koijen and Van Nieuwerburgh (2011), and Golez (2014). Our results in Section 3.4 are supportive of this finding.

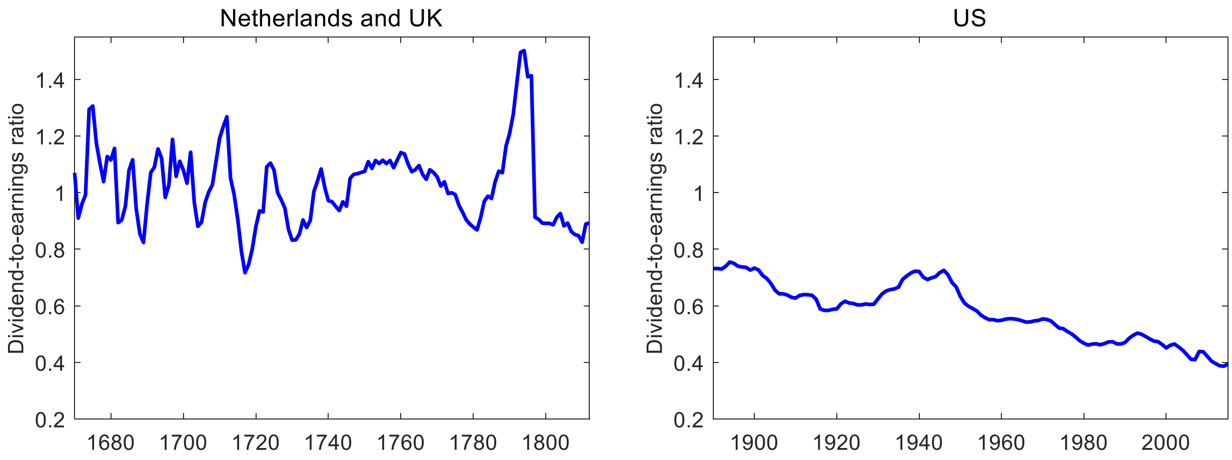

Fig. 4. Dividend-to-earnings ratio**.** This figure plots the ratio of 20-year trailing sum of dividends to 20-year trailing sum of earnings for the Netherlands and UK period (1651–1812) and the combined US period (1871–2015). The first observations are in 1670 and 1890, respectively.

References

Acheson, G., Hickson, C.R., Turner, J.D., Ye, Q., 2009. Rule Britannia! British stock market returns, 1825−1870. J. Econ. Hist. 69, 1107–1137.

Albuquerque, R., Eichenbaum, M., Luo, V.X., Rebelo, S., 2016. Valuation risk and asset pricing. J. Financ. 71, 2861–2906.

Ashton, T.S., 1959. Economic Fluctuations in England. Clarendon Press, Oxford, UK.

Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. J. Financ. 59, 1481–1509.

Binsbergen, J.H. van, Koijen, R.S.J., 2010. Predictive regressions: a present-value approach. J. Financ. 65, 1439–1471.

Boudoukh, J., Richardson, M., Whitelaw, R.F., 2008. The myth of long-horizon predictability. Rev. Financ. Stud. 21, 1577–1605.

Bowen, H.V., 1989. Investment and empire in the later 18th century: East India Company stockholding, 1756−1791. Econ. Hist. Rev. 42, 186–206.

Bris, D.le, Goetzmann, W.N., Pouget, S. The present-value relation over six centuries: the case of the Bazacle Company. J. Financ. Econ. forthcoming.

Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of stock market behavior. J. Polit. Econ. 107, 205–251.

Campbell, J.Y., Shiller, R.J., 1988a. Stock prices, earnings, and expected dividends. J. Financ. 43, 661–676.

Campbell, J.Y., Shiller, R.J., 1988b. The dividend-price ratio and expectations of future dividends and discount factors. Rev. Financ. Stud. 1, 195–227.

Campbell, J.Y., Shiller, R.J., 2005. Valuation ratios and the long-run stock market outlook: an update. In: Thaler, R.H. (Ed.), Advances in Behavioral Finance II. Princeton University Press, Princeton, NJ.

Campbell, J.Y., Thompson, S.B., 2008. Predicting excess stock returns out of sample: can anything beat the historical average? Rev. Financ. Stud. 21, 1509–1531.

Chen, L., 2009. On the reversal of return and dividend growth predictability: a tale of two periods. J. Financ. Econ. 92, 128–151.

Chen, L., Da, Z., Priestley, R., 2012. Dividend smoothing and predictability. Manag. Sci. 58, 1834–1853.

Clark, G., 2015. What were the British earnings and prices then? (New Series). MeasuringWorth http://www.measuringworth.com/ukearncpi/.

Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. J. Econom. 138, 291–311.

Cochrane, J.H., 1992. Explaining the variance of price-dividend ratios. Rev. Financ. Stud. 5, 243–280.

Cochrane, J.H., 2008. The dog that did not bark: a defense of return predictability. Rev. Financ. Stud. 21, 1533–1575.

Cochrane, J.H., 2011. Presidential address: discount rates. J. Financ. 66, 1047–1108.

Cowles, A., 1939. Common stock indices, 1871−1937. Cowles Commission for Research in Economics. Monograph no. 3. Principia Press, Bloomington, IN.

Fama, E.F., French, K.R., 1988. Dividend yields and expected stock returns. J. Financ. Econ. 22, 3–25.

Fama, E.F., French, K.R., 1989. Business conditions and expected returns on stocks and bonds. J. Financ. Econ. 25, 23–49.

Fama, E.F., French, K.R., 2001. Disappearing dividends: changing firm characteristics or lower propensity to pay? J. Financ. Econ. 60, 3–43.

Gayer, A.D., Rostow, W.W., Schwartz, A.J., 1953. The Growth and Fluctuation of the British Economy, 1790–1850, Volume 1. Oxford University Press, Oxford, UK.

Goetzmann, W.N., Ibbotson, R.G., Peng, L., 2001. A new historical database for the NYSE 1815 to 1925: performance and predictability. J. Financ. Mark. 4, 1–32.

Golez, B., 2014. Expected returns and dividend growth rates implied by derivative markets. Rev. Financ. Stud. 27, 790–822.

Goyal, A., Welch, I., 2003. Predicting the equity premium with dividend ratios. Manag. Sci. 49, 639–654.

Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of the equity premium prediction. Rev. Financ. Stud. 21, 1455–1508.

Grossman, R.S., 2002. New indices of British equity prices, 1870−1913. J. Econ. Hist. 62, 121–146.

Harrison, P., 1998. Similarities in the distribution of stock market price changes between the 18th and 20th centuries. J. Bus. 71, 55–79.

Hickson, C.R., Turner, J.D., Ye, Q., 2011. The rate of return on equity across industrial sectors on the British stock market, 1825−1870. Econ. Hist. Rev. 64, 1218–1241.

Koijen, R.S.J., Nieuwerburgh, S. van, 2011. Predictability of returns and cash flows. Ann. Rev. Financ. Econ. 3, 467–491.

Kolev, G.I., Karapandza, R., 2017. Out-of-sample equity premium predictability and sample split-invariant inference. J. Bank. Financ. 84, 188–201.

Koudijs, P., 2016. The boats that did not sail: asset price volatility in a natural experiment. J. Financ. 71, 1185–1226.

Lent, W. van, Sgourev, S.V., 2013. Local elites versus dominant shareholders: dividend smoothing at the Dutch East India Company. Acad. Manag. Proc. 1, 15592.

LeRoy, S.F., Porter, R.D., 1981. The present value relation: test based on implied variance bounds. Econometrica 49, 555–574.

Neal, L., 1990. The rise of financial capitalism. Cambridge University Press, Cambridge, UK.

Nelson, C.R., Kim, M.J., 1993. Predictable stock returns: the role of small sample bias. J. Financ. 48, 641–661.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroscedasticity, and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Rajan, R.,G., Zingales, L., 2003. The great reversals: the politics of financial development in the 20th century. J. Financ. Econ. 69, 5–50.

Rangvid, J., Schmeling, M., Schrimpf, A., 2014. Dividend predictability around the world. J. Financ. Quant. Anal. 49, 1255–1277.

Rossi, B., Inoue, A., 2012. Out-of-sample forecast tests robust to the choice of window size. J. Bus. Econ. Stat. 30, 432–453.

Rostow, W.W., 1972. Cycles in the British economy: 1790 – 1914. In: Aldcroft, D.H., Fearon, P. (Eds.), British Economic Fluctuations 1790–1913. Macmillan. Basingstoke, UK, pp. 74–96.

Santa-Clara, P., Maio, P., 2015. Dividend yields, dividend growth, and return predictability in the cross section of stocks. J. Financ. Quant. Anal. 50, 33–60.
Schwert, G.W., 1990. Indexes of US stock prices from 1802 to 1987. J. Bus. 63, 399–426.
Schwert, G.W., 2003. Anomalies and market efficiency. In: Constantinides, G., Harris, M., Stulz, R. (Eds.), Handbook of the Economics of Finance. North Holland, Amsterdam, pp. 937–972.
Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? Am. Econ. Rev. 71, 421–436.
Stambaugh, R.F., 1999. Predictive regressions. J. Financ. Econ. 54, 375–421.
Wright, J.F., 1997. The contribution of overseas savings to the funded national debt of Great Britain, 1750−1815. Econ. Hist. Rev. 50, 657–674.
Zanden, J.L. van, Leeuwen, B. van, 2012. Persistent but not consistent: the growth of national income in Holland, 1347−1807. Explor. Econ. Hist. 49, 119–130.