Journal of Monetary Economics 131 (2022) 61–77

Contents lists available at ScienceDirect

Journal of Monetary Economics

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

Recessions and the stock market☆

Tim A. Kroencke

University of Neuchâtel, Rue A.-L. Breguet 2, CH-2000, Neuchâtel, Switzerland

ARTICLE INFO

Article history:
Received 24 June 2021
Revised 5 July 2022
Accepted 5 July 2022
Available online 7 July 2022

Keywords:
Business cycles and the stock market
Predictability of the price-dividend ratio
Empirical evaluation of macro-finance models

ABSTRACT

An event study approach is adopted to investigate the drivers of the stock market around recessions. First, stock prices and dividends drop contemporaneously when accounting for different timing conventions. Accordingly, stock prices do not anticipate recessions due to an economic mechanism (cash flow news). Second, the variance of price changes increases at least as much as the variance of dividend growth during recessions. This result suggests that changes in the price of risk (discount rate news) play an essential role. Implications and opportunities for standard asset pricing theories and recently proposed alternatives are also discussed.

© 2022 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)

1. Introduction

Using quarterly stock market data covering 42 recessions in 14 countries since 1951, I study the behavior and the potential factors that move the stock market around recessions. I find that real dividends fall by 13% on average within the first four quarters after the beginning of recessions. At the same time, stock prices drop even more sharply by 30%. This excess drop in prices implies considerable variation in the price-dividend ratio. If investors respect present-value relationships, either expected returns or expected future dividend growth must vary considerably during recessions (Shiller, 1981; Campbell and Shiller, 1988; Campbell, 1991). In this paper, I build on Muir (2017) and utilize an event study approach to shed novel light on the question of what drives the stock market down during recessions.

The Timing of the Price-Dividend Ratio. First, I investigate the possibility that stock prices (at least in part) anticipate the low future cash flows (and economic growth) observed during recessions. A long literature shows that the price-dividend ratio does not predict dividend growth (e.g., Campbell and Shiller, 1988; Cochrane, 2008). This finding suggests that price-dividend ratio movements around recessions plausibly reflect changes in expected returns and not changes in expected future dividend growth (Muir, 2017). That said, another strand of the literature shows that monthly/quarterly stock returns can be used to forecast future business conditions (e.g., Fama and French, 1989).[1] Consistent with the latter finding, Cieslak and Vissing-Jorgensen (2021) provide recent empirical evidence that U.S. central bankers are "asset price dependent" (Warsh, 2016) when they revise expectations about future growth.

☆ University of Neuchâtel, Switzerland. For helpful suggestions I am grateful to Urban Jermann (the Editor), Carolin Pflueger (the Associate Editor), an anonymous Referee, Pedro Barroso, Christoph Merkle, Stig Vinther Møller, Jesper Rangvid, Maik Schmeling, Michael Ungeheuer, Christian Wagner, Josef Zechner and the seminar participants of the Aalto Business School, Aarhus University, Central Bank of Ireland, University of Fribourg, University of St. Gallen, University of Strathclyde Glasgow, Vienna University of Economics and Business as well as the participants in conferences where the paper was presented. First draft: May 2018.

E-mail address: tim.kroencke@unine.ch

[1] See also, Fama (1990), Chauvet (1999), Vassalou (2003), Campbell and Diebold (2009), Backus et al. (2010), Maio and Cooper (2019), among others.

https://doi.org/10.1016/j.jmoneco.2022.07.004
0304-3932/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/)

I find that the price-dividend ratio falls significantly by 5.6% in the two quarters *before* recessions start and dividends drop. However, stock prices are commonly measured at the end of a period (e.g., the last trading day of December), while dividends are measured as a trailing 12-month sum (e.g., from January to December). This property gives rise to the so-called time-aggregation bias (Working, 1960; Taio, 1972), which implies that dividends *mechanically* lag stock market prices within the horizon over which the data are time-aggregated. While the mechanical lag is not a problem in a purely econometric sense, it can lead to erroneous conclusions when asking about the economic mechanism why stock prices anticipate recessions. Following Cochrane (1996), and Kroencke (2017), I consider stock prices calculated as a trailing 12-month mean to obtain an adjusted price-dividend ratio.[2] After aligning the timing of dividends and prices, I find that stock prices and dividends start falling contemporaneously with the beginning of a recession. This result speaks against the idea that the stock market can anticipate the low cash flows observed during recessions due to an economically rooted mechanism (e.g., long run risk, or time-varying rare disaster risk).

A recent literature provides evidence that subjective expectations about cash flows can explain the variation in price-dividend ratios (De La O and Myers, 2021; Bordalo et al., 2020). During recessions, it might be possible that investors (incorrectly) expect depressed future dividend growth, which would push the price-dividend ratio further down than realized dividends. While I recognize that the available survey data around recessions is limited, I provide evidence that speaks against this possibility. I study the forward term structure of expected U.S. real GDP growth from the Survey of Professional Forecasters (SPF) around the available U.S. subset of my events. During recessions, the nowcast and future one quarter expectations are most heavily revised. Revisions in the three quarter forecast are small and to a large extent insignificant.[3] Longer horizons (e.g., future ten years) are not available, and I cannot rule out that subjective expectations at such horizons affect stock prices around recessions. However, it would require a U-shaped relationship where short and long-horizon expectations vary while the medium horizon remains flat. I conclude that price-dividend ratios more plausibly fall during recessions due to elevated expected returns, in line with earlier evidence (Lustig and Verdelhan, 2012).

The Recession Variance Ratio. Second, I dissect the possible drivers of elevated expected returns during recessions. The first possibility is that cash flows are more uncertain during recessions. A higher amount of risk should then translate to higher expected returns and lower price levels. An alternative driver is that investors get more risk-averse and command a higher price of risk. The first channel affects the variance of cash flows and stock prices, while the latter channel only affects stock prices (e.g., Jurado et al., 2015). Therefore, comparing the change in the variance of prices versus cash flows reveals the importance of these two channels. To the best of my knowledge, this fact has not yet been exploited to evaluate discount rate news channels.

The recession variance ratio, defined as the ratio of the recession variance over the pre-recession variance, is as large as 2.1 for price changes and 1.7 for dividend growth (both significantly different from one). In the U.S. sample, the recession variance ratio of price changes is even as large as 4.8. In the absence of important long-run cash flow news, this result suggests that the changes in the perception of risk must be *at least* equally important to changes in the amount of risk, and therefore plays a key role during recessions.

Implications for Asset Pricing Theories. To compare my empirical results to the theoretical counterfactual, I simulate "recessions" in the long-run risk model of Bansal and Yaron (2004), the habit model of Campbell and Cochrane (1999), and in the rare disaster risk model by Wachter (2013).[4] I find that the classic models struggle to explain the facts. In the long-run risk model, revisions about expectations of future cash flows take a central role. The average recession is also a period of predictable low growth. Because of the predictable cash flow component, the (time-aggregated) price-dividend ratio starts dropping about one year ahead of recessions and is flat during recessions. This drop is about one year too early compared to the empirical data. Similarly, the model with time-varying rare disaster risk predicts that stock prices start dropping in anticipation of recessions as the "disaster" probability tends to increase ahead of the actual recession. The long-run risk as well as the time-varying rare disaster model (in the absence of realized disasters) do not push up the variance of price changes in a significant way. They struggle to do so because there is no feedback mechanism such that a recession can give rise to significantly elevated expected returns.

In the Campbell and Cochrane (1999) habit model, cash flows are unpredictable and homoscedastic. All changes in the price-dividend ratio come from changes in the price of risk. Indeed, stock prices fall contemporaneously with cash flows during recessions, similarly as in the data. However, the stock price variance increases only by a factor of 1.1 during recessions. This number is not close to 2.1 as in the broad sample (or 4.8 in the U.S. sample). The model, therefore, has problems explaining quantitatively the stock market around recessions.

My interpretation of these results is that the large recession variance ratio for price changes requires an amplifier mechanism (e.g., frictions, disagreement, accounting for corporate policies). I provide a discussion of recent examples in the literature that might help to bring the models closer to the empirical data. An alternative take on the results is that large

[2] While the time-aggregation bias is frequently accounted for in theoretical work (e.g., Campbell and Cochrane, 1999; Bansal and Yaron, 2004), it is often not taken into account in empirical work.

[3] I find variation regarding expected growth in the data at the medium horizon (future three quarters). But these shifts in expectations are not systematically related to recessions.

[4] The original rare disaster model studied by Rietz (1988) and Barro (2006) has a constant disaster risk probability and features a constant price-dividend ratio (as in the classic model). Gabaix (2012) and Gourio (2008) also present models that incorporate time-varying rare disasters, or time-varying sensitivity of the stock market to rare disaster risk.

recession variance ratios are driven by subjective beliefs about long horizon economic growth, which are difficult to measure and are not captured by the SPF revisions of expected real GDP growth.

The event study methodology applied in this paper is a non-parametric method that comes with a minimum of assumptions and data requirements. It also comes with a disadvantage. By construction, my results are silent on what drives the market outside of recession periods. Models that do a poor job during recessions might be useful in explaining longer-horizon movements in the stock market.

Related literature. I complement earlier work by Muir (2017). He compares the decline of asset prices and the change in fundamentals around financial crises and "normal" recessions. Asset price declines are larger during financial crises compared to normal recessions, even though fundamentals move similarly. For that reason, he concludes that standard consumption-based asset pricing theories must have difficulties in explaining the even larger drop of asset prices during financial crises. Since his analysis uses a "difference-in-difference" approach (comparing financial crises with normal recessions), his paper is silent about what kind of economic mechanisms drive stock prices around recessions in general. The goal of my paper is to fill this gap in the literature. I extend his analysis by an "absolute" approach and study the potential economic mechanisms that drive the stock market around "normal" recessions.

A few papers provide empirical tests of the predictions of different equilibrium asset pricing models. Zviadadze (2021) provides such a comprehensive test using structural estimation techniques and exploiting the term structure of risk. She finds that consumption uncertainty and short-run consumption shocks are empirically linked and thus present an important feature of the data. Bekaert et al. (2022) disentangle changes in the price of risk and cash flow uncertainty using a structural model. They conclude that changes in the price of risk are the main driver of conditional stock price variance. Campbell et al. (2013) apply the classic VAR decomposition to compare cash-flow news and discount rate news over time and across recessions. Other comprehensive comparisons of the predictions of multiple asset pricing models with a focus on the pricing of options-based payoffs can be found in Dew-Becker et al. (2017), Martin (2017), Beason and Schreindorfer (2022). Work with a focus on the term structure of expected returns is presented by van Binsbergen et al. (2012), Backus et al. (2014), and Belo et al. (2015). My findings largely corroborate this literature.

An important difference to these papers is that I apply a non-parametric method that comes with a minimum of assumptions and data requirements, e.g., the availability of derivative market data is not required. This allows me to study a longer sample period compared to many other contributions. My results are also valid if markets are segmented, or differ in liquidity.

Recently, Kozak and Santosh (2020) provide evidence that discount rates are high in periods of high marginal utility of wealth. These results are qualitatively in line with the predictions of the consumption-based asset pricing theories discussed in this paper. However, their results are silent on whether the models also quantitatively match the data. I show that these models have difficulty explaining the facts quantitatively.

My paper is also connected to the return predictability literature. Several papers have documented that expected returns are countercyclical, e.g., Fama and French (1989), Ferson and Harvey (1991), Harrison and Zhang (1999), Lettau and Ludvigson (2009), Møller and Rangvid (2018), Golez and Koudijs (2018). Lustig and Verdelhan (2012) show that realized future returns, as a proxy for expected returns, are higher during NBER recessions compared to normal times. Rapach et al. (2010), Henkel et al. (2011), Dangl and Halling (2012), find that returns are more predictable during recessions. Cooper and Priestley (2009) show that the output gap predicts future returns; Atanasov et al. (2020) document predictability by cyclical consumption. Both findings are in line with the idea that expected returns are high during recessions, when the output gap is negative and cyclical consumption is low. These papers document important empirical facts. I add to this literature by studying how these findings align with theoretical predictions of different models.

The problem of the time-aggregation bias in the price-dividend ratio can jointly explain non-predictability of dividends by the price-dividend ratio at low frequency (annual data, e.g., Cochrane, 2008), and why other parts of the literature find predictive power of business conditions at higher frequency (monthly/quarterly data, e.g., Fama and French, 1989). I am not aware of an alternative explanation that jointly explains both empirical findings presented in the literature. However, my paper is silent on why other authors find dividend predictability at the annual frequency when considering alternative methods or samples.[5] One possibility is that these papers capture a more structural component of dividend predictability unrelated to the business cycle, which is outside the scope of this paper.

What is required in a model to explain the stock market around recessions is some sort of amplification mechanism or asymmetric response to sharp declines in economic activity that makes risky assets more vulnerable in "bad times". My empirical results support recent theoretical work in this area. Some recent papers propose theoretical mechanisms where short-run and long-run consumption risks are combined (e.g., Branger et al., 2016), or that agents learn (or disagree) about the business cycle (e.g., Andrei et al., 2019; Cujean and Hasler, 2017). Others consider that investors might be disappointment-averse (Schreindorfer, 2020), underline the importance of institutional frictions during recessions (e.g., Adrian and Shin, 2014), or model the dividend dynamics in a more realistic way (Belo et al., 2015; Bekaert and Engstrom, 2017). The recession variance ratio proposed in this paper is appealing for this literature, as it is an easy to implement metric that can be used to benchmark future models.

[5] For example, Larrain and Yogo (2008), Chen (2009), Binsbergen and Koijen (2010), Bansal et al. (2012), Rangvid et al. (2014), Jank (2015), Golez and Koudijs (2018), Jagannathan and Liu (2019).

Outline. Section 2 presents the data employed, the determination of recession events, and the empirical methodology. Section 3 studies the empirical timing of stock prices around recessions. Section 4 shows that recession variance ratios give clues on the drivers of discount rate news. Further results are provided in Section 5. Section 6 provides a discussion of the implications for asset pricing theories; followed by the conclusion.

2. Data, recessions & econometric methodology

I first describe the source of the stock market data and the construction of the list of recession events. I then illustrate the event time methodology applied in this paper using real GDP growth. This exercise confirms that the event list successfully detects business cycle peaks.

Stock Market and Real GDP Data. I collect stock market prices, dividend yields, and real GDP from the Global Financial Database (GFD) for 14 countries (Australia, Canada, Denmark, France, Germany, Italy, Japan, Netherlands, Norway, Spain, Sweden, Switzerland, United Kingdom, United States).[6] I convert the foreign currency stock market prices to U.S. dollars using the according GFD exchange rate. I deflate U.S. dollar stock price indices using the GFD U.S. consumer price index. A measure of real dividends is obtained by multiplying dividend yields with the real prices. This database allows me study the stock market around recessions at the quarterly frequency in the sample period from 1951 to 2019 (main results), and at the annual frequency from 1872 to 2019 (robustness) for a large cross-section of countries. A list containing the variable names is provided in Table OA.1 of the Online Appendix.

Recessions. For the U.S., I use annual and quarterly business cycle peaks as reported by the NBER business cycle dating committee.[7] For the 13 additional countries, I collect annual business cycle peaks from Jordá et al. (2011), Table 7. To find the business cycle peak at the quarterly frequency, I use the closest monthly peak provided by Fushing et al. (2010). Muir (2017) finds that stock markets in recessions associated with a financial crisis behave different to "normal" recessions. To provide a complementary analysis, I want to avoid confounding events and exclude financial crises using the list provided by Muir (2017). In the annual sample, I also exclude observations +/- three years (the event window) around World War I and II, and the hyperinflation period of the 1920s in Germany (as suggested by Anarkulova et al., 2022).

In total, I have 43 "normal" recessions in the quarterly sample and 135 events in the annual sample when considering all countries. Complete stock market data (prices and dividend yields) are available for 42 (75) events in the quarterly (annual) sample. For the U.S., stock market data are available for all 9 (16) events in the quarterly (annual) sample. The complete event list is provided in Table OA.2 of the Online Appendix.

I acknowledge that there are alternative approaches available to define business cycle peaks. The selection of recessions in this paper ensures that I look at the same events as in the previous literature (Lustig and Verdelhan, 2012; Muir, 2017). Second, the definition of recessions is comparable across countries. Third, using an existing definition of business cycle peaks in contrast to using an own methodology mitigates sample selection bias concerns.

Econometric methodology and real GDP in recession event time. To explain the event study methodology employed in this paper and to illustrate that the business cycle peaks are reasonably determined, I report real GDP growth in recession event time in Fig. 1. The data are quarterly sampled from 1951 to 2019, and real GDP data are available for 33 out of the 43 events (i.e., for less events than for the stock market). To construct the figure, I first de-mean log real GDP growth by the country specific mean (i.e., I account for country fixed effects). Afterwards, real GDP growth is re-ordered in recession event time, where τ=0 corresponds to the dates from my event list, τ=+1 is one quarter later, τ=-1 is one quarter earlier and so on. I then take the mean across all event observations and cumulate them over the event window from τ=-12 to τ=+12 quarters. For statistical inference, I conduct 10,000 bootstrap samples of the events. Reported confidence intervals correspond to the cumulative estimate from starting at the event in τ=0 (i.e. from the middle of the figure) and are the 5% and 95% percentiles of the bootstrapped distribution. Because I re-sample from the events, the bootstrap accounts for event induced increases in the variance of the variable of interest.

In Fig. 1, I report real GDP growth in recession event time for all 14 countries (black, dots), and only for the U.S. (red, squares). I find that real GDP growth has a clear business-cycle pattern, which illustrates that the definition of business cycle peaks is reasonable. To make the timing of the drop in real GDP growth easier to see, I shift real GDP growth in the figure such that the maximum within one year of the events aligns with the y-axis at zero. On average, real GDP falls within five quarters by about 4%. The results are comparable in the "all countries" sample with the "U.S." sample. The definition of business cycle peaks could correspond to the beginning of a period as well as the end of a period. In the data, I find that the peak of (de-meaned) GDP growth is in quarter τ=-1 (i.e. closer to the beginning of the period), and for this reason, I label this quarter as the business cycle peak in the figure.[8] This is a cosmetic label to indicate in the following event time figures where real GDP peaks. For example, if stock prices drop before τ=-1 (the business cycle peak), the stock market predicts real GDP growth in recession event time.

Alternative Methods. The related literature often relies on panel dummy regressions to study a variable of interest in event time (see, e.g., Muir, 2017). This method is equivalent to the event study approach (Campbell et al., 1997, Chapter

[6] The Global Financial Database (GFD) has been extensively used in previous research. Anarkulova et al. (2022) provide further references and a comparison to other annually sampled databases.

[7] http://www.nber.org/cycles/recessions.html.

[8] In the U.S., the point estimates indicate already a drop from τ=-2 to τ=-1, but this drop is economically small and statistically insignificant.

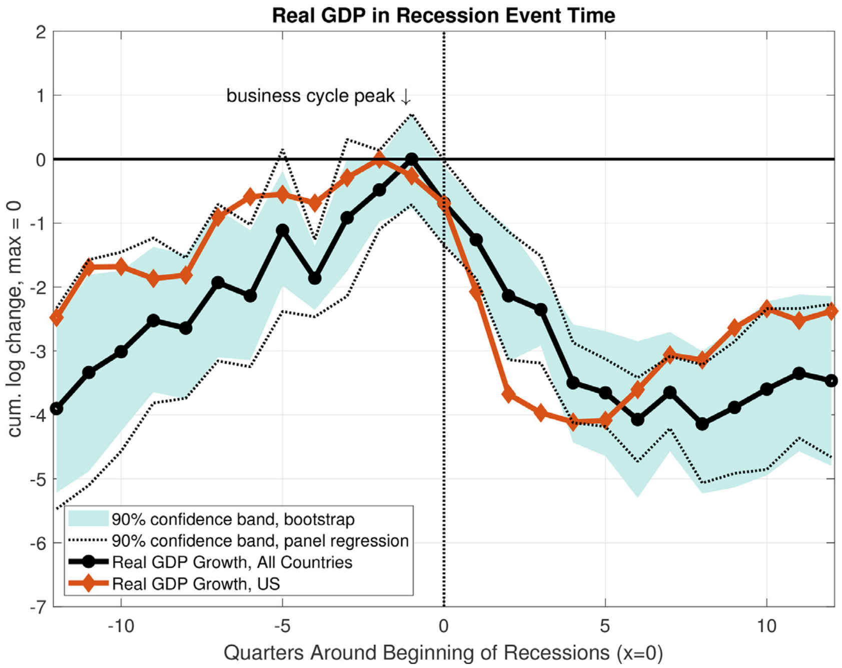

Fig. 1. Real GDP growth in recession event time (1951–2019, # Recessions: 33).

4).[9] I find that panel regressions provide the (almost) identical point estimates, and very similar confidence intervals when relying on heteroscedasticity robust standard errors; illustrated by the dotted lines in Fig. 1. The bootstrap procedure allows me to draw inferences about the recession variance ratio, which i s not straightforward with panel regressions.

3. The timing of the stock market around recessions

3.1. Accounting for the time-aggregation bias in dividends

In the empirical data, dividends are measured as a trailing 12-month sum, i.e. they are time-aggregated.[10] The motivation behind this convention is that dividends are highly cyclical with strong annual and quarterly patterns that reflect the accounting cycle. Because stock prices are (usually) measured at the end of a period, the different timing conventions give rise to a purely mechanical lead-lag relationship between dividends and prices (also known as the time-aggregation bias, see Working, 1960; Taio, 1972). Following Cochrane (1996) and Kroencke (2017), I consider time-aggregated stock prices by computing the trailing 12-month mean of stock prices. I then use the time aggregated stock price to obtain a modified time-aggregated price-dividend ratio, which aligns the timing of prices with the timing of dividends.

The alignment of the timing of dividends and prices matters when studying the timing of the price-dividend ratio around recessions. This is not a second order effect. For illustration, I simulate the price-dividend ratio in a simple model where the true ratio is constant. In Figure OA.4 (Online Appendix), the price-dividend ratio using 12-month trailing dividends makes large swings in recession event time even when the true price-dividend ratio is constant. In contrast, the price-dividend ratio based on time-aggregated stocks prices remains constant and allows for correct inference on the timing of prices relative to dividends.

3.2. The timing of stock prices and cash flows

Figure 2 shows the cumulative drop of stock prices and dividends around 42 recessions in 14 countries since 1951.[11] The upper figure shows the cumulative change of changes in the log price-dividend ratio and the lower figure shows the components of the ratio individually. I shift all variables by their own local peak (y-axis = 0) to make it easier to see the timing and the cumulative drop around recessions across variables. Table 1 reports the numerical results for the price-dividend ratio and provides more detailed statistical inference.

[9] Also the local linear projections of Jordá (2005) are equivalent when regressing on event dummies.

[10] Similarly, quarterly real GDP is measured as the sum of real GDP over three months. Macroeconomic variables are usually also aggregated over time, albeit to a possibly lesser extent.

[11] There are 43 “normal” recessions (Table OA.2, Online Appendix). Due to data availability, Denmark in Q3/1961 is missing.

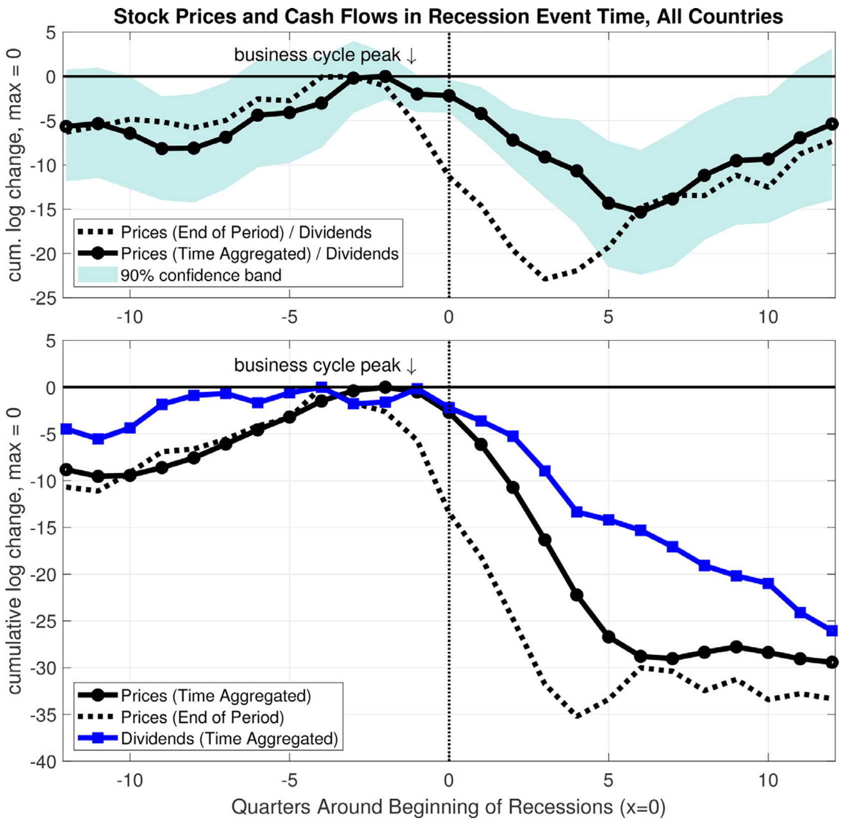

Fig. 2. The stock market in recession event time (All countries, 1951–2019, # Recessions: 42).

Table 1
The stock market in recession event time. This table shows cumulative log changes of the price-dividend ratio ($\triangle pd$) around the beginning of recessions from 1951 to 2019. In Panel A, estimates are based on 42 recessions in 14 countries; in Panel B, estimates are based on 9 recessions in the U.S. (the event list is provided in Table OA.2, Online Appendix). The data are sampled quarterly, and the event window ranges from τ=-12 to τ=+12 quarters around the beginning of a recession (τ=0). In the price-dividend ratio, dividends are always trailing 12-month sums, i.e., they are time-aggregated. Prices are measured at the end of a quarter (end of period, E.o.P), or they are trailing 12-month means such that they have the same timing as dividends and are also time-aggregated (T.A.). T-statistics (t) are based on a bootstrap resampling of the data in the event time.

	←pre-Recession					↓	after Beginning of Recession→				
Quarter	−12:-1	−8:-1	−4:-1	−2:-1	−1	0	0:+1	0:+2	0:+4	0:+8	0:+12
	Panel A: All Countries (# Recessions: 42)										
$\triangle pd_t$, T.A.	1.36	6.17	2.11	−1.80	−1.99	−0.18	−2.20	−5.20	−8.68	−9.18	−3.38
t	*0.35*	*1.63*	*0.67*	*−1.12*	*−1.57*	*−0.16*	*−1.25*	*−2.55*	*−2.43*	*−2.09*	*−0.65*
$\triangle pd_t$, E.o.P.	−0.39	−0.46	−2.85	−5.60	−4.50	−5.72	−9.00	−14.02	−16.35	−7.86	−1.75
t	*−0.09*	*−0.10*	*−0.83*	*−2.80*	*−2.52*	*−3.63*	*−3.76*	*−4.08*	*−3.24*	*−1.51*	*−0.32*
	Panel B: U.S. Sample (# Recessions: 9)										
$\triangle pd_t$, T.A.	6.64	8.01	3.18	−0.51	−1.26	−2.39	−6.01	−9.69	−12.42	−6.85	−6.94
t	*0.83*	*2.06*	*2.64*	*−0.41*	*−1.62*	*−3.64*	*−4.95*	*−5.50*	*−2.72*	*−0.98*	*−0.96*
$\triangle pd_t$, E.o.P.	1.53	1.64	−4.56	−5.08	−3.73	−9.59	−10.09	−12.85	−6.28	−5.45	0.85
t	*0.21*	*0.48*	*−1.98*	*−2.97*	*−1.90*	*−4.68*	*−4.52*	*−2.96*	*−0.91*	*−0.69*	*0.13*

I find that the time-aggregated price-dividend ratio (upper figure) starts dropping significantly only at the beginning of recessions (τ=0). Before the start of recessions (τ=-1), the time-aggregated price-dividend ratio drops only insignificantly by -2%. The cumulative drop from the beginning of recessions (τ=0) to four quarters later (τ=+4) is -9% (t-stat: -2.4). In contrast, the unadjusted price-dividend ratio starts to decline well in advance of recessions. For example, from two quarters to one quarter before the beginning of recessions, it is down by -5.6% (t-stat: -2.8). Therefore, the traditional price-dividend ratio using end of period prices and time-aggregated dividends can “predict” recessions. However, once I adjust for the time-aggregation bias and impose an equal timing of the measurement of prices and dividends, the price-dividend ratio does not anticipate recessions.

Table 2
The stock market in recession event time: Components. This table complements Table 1 and shows cumulative log changes of the components of the price-dividend ratio (prices $\triangle p$ and dividends $\triangle d$) around the beginning of recessions from 1951 to 2019. Prices are measured at the end of a quarter (end of period, E.o.P), or they are trailing 12-month means such that they have the same timing as dividends and are also time-aggregated (T.A.). T-statistics (t) are based on a bootstrap resampling of the data in the event time.

	←pre-Recession					↓	after Beginning of Recession→				
Quarter	−12:-1	−8:-1	−4:-1	−2:-1	−1	0	0:+1	0:+2	0:+4	0:+8	0:+12
	Panel A: All Countries (# Recessions: 42)										
$\triangle p_t$, T.A.	7.11	8.07	2.67	−0.14	−0.52	−2.18	−5.60	−10.20	−21.71	−27.83	−28.90
t	*1.57*	*2.01*	*1.03*	*−0.11*	*−0.86*	*−3.60*	*−4.98*	*−6.12*	*−7.37*	*−7.21*	*−6.42*
$\triangle p_t$, E.o.P.	5.01	1.21	−2.40	−3.99	−3.06	−7.75	−12.45	−19.10	−29.52	−26.77	−27.64
t	*1.00*	*0.29*	*−0.89*	*−2.03*	*−2.26*	*−5.38*	*−6.47*	*−6.27*	*−6.90*	*−5.83*	*−5.31*
$\triangle d_t$, T.A.	5.40	1.68	0.45	1.61	1.44	−2.03	−3.45	−5.09	−13.17	−18.91	−25.89
t	*1.18*	*0.45*	*0.18*	*1.17*	*1.12*	*−1.61*	*−1.74*	*−2.13*	*−3.63*	*−5.26*	*−5.56*
	Panel B: U.S. Sample (# Recessions: 9)										
$\triangle p_t$, T.A.	2.81	3.53	0.56	−1.87	−1.32	−3.51	−8.01	−12.45	−17.68	−13.52	−11.89
t	*0.35*	*0.93*	*0.29*	*−1.82*	*−2.37*	*−6.43*	*−7.40*	*−6.11*	*−3.50*	*−1.57*	*−1.33*
$\triangle p_t$, E.o.P.	−2.30	−2.84	−7.18	−6.45	−3.80	−10.70	−12.09	−15.61	−11.54	−12.11	−4.10
t	*−0.34*	*−0.73*	*−2.83*	*−4.01*	*−2.25*	*−5.28*	*−6.58*	*−3.35*	*−1.59*	*−1.31*	*−0.52*
$\triangle d_t$, T.A.	−3.83	−4.48	−2.62	−1.36	−0.07	−1.12	−2.00	−2.76	−5.27	−6.66	−4.95
t	*−1.10*	*−1.86*	*−1.61*	*−1.35*	*−0.09*	*−3.02*	*−2.75*	*−3.91*	*−5.92*	*−3.11*	*−1.56*

A look at the components in Table 2 shows that dividends drop contemporaneously with real GDP and up to -13% during recessions (from τ=0 to τ=+8; t-stat: -5.3). At the same time, stock prices drop -30% (t-stat: -7.2). Stock prices drop roughly 50% more during recessions, indicating a potential role for increased expected returns during recessions (in line with Lustig and Verdelhan, 2012).

In Panel B of Table 1 and Table 2, I report results when the sample is constrained to the U.S. (9 recessions). Figure OA.1 (Online Appendix) provides a visual presentation of the U.S. experience. Stock prices measured at the end of the period anticipate recessions, while stock prices aggregated over time do not. Prices drop about two time as much as dividends, indicating also a prominent role for increased expected returns during recessions in the U.S. sample.

In summary, I find that stock prices measured using the end of period timing convention anticipate dividends and recessions about two quarters ahead. Time-aggregated dividends decline along with real GDP at the onset of recessions. Once I measure stock prices as time-aggregated, stock prices drop simultaneously as dividends and no longer anticipate recessions. This result speaks against the idea that cash flows have an economically rooted predictable component such that stock prices anticipate recessions. Stock prices anticipate recessions because they are based on market prices, which are measured more timely and can therefore reflect the real economy more quickly.

3.3. The timing of growth expectations

Recessions might *lead* investors to revisions in future expected growth.[12] In Fig. 1, I find that real GDP growth increases with average speed in the "all countries" sample after recessions (τ=+5 to τ=+12), as indicated by a flat line.[13] For the U.S., real GDP growth is even upward trending. Dividends (Fig. 2) are mildly downward trending in the all countries sample, and flat in the U.S. sample. However, dividends are also measured rather inaccurately with increasing distance from τ=0 (Table 2). Furthermore, De La O and Myers (2021) provide evidence that subjective expectations about future cash flows can explain to a large extent the variation in the levels of price-dividend and price-earnings ratios (in contrast to realized cash flows, Campbell and Shiller, 1988).[14] Even if realized growth does not fall after recessions (Fig. 1), investors may mistakenly believe that it does.

To explore this possibility, I analyze changes in the *forward* term structure of U.S. real GDP growth forecasts from the Survey of Professional Forecasters (SPF) around 6 U.S. recessions from 1969 to 2019.[15] I compare revisions in short-term and longer term forward growth expectations of real GDP. For example, in Q4 of 2015, I first compute expected growth for Q1 2016, Q2 2016, Q3 2016, and Q4 2016. Second, in Q1 2016, I update expected growth for Q1 2016 (which becomes a nowcast), Q2 2016, Q3 2016, and Q4 2016. Finally, I compute the difference between expectations in Q1 2016 and Q4 2015

[12] Note that there is no such mechanism in traditional versions of the long-run risk model (e.g., Bansal and Yaron, 2004; Bansal et al., 2012). Similarly, such a link between short-run consumption risk and the rare disaster probability is not present in the Wachter (2013)-version of the rare disaster risk model.

[13] Country individual means are re-moved from the data, i.e. a flat line indicates average real growth. See Section 2 for further details.

[14] See also Bordalo et al. (2020).

[15] The data come from the website of the Philadelphia Fed: https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/data-files/rgdp. Real GDP Forecasts are available since Q4 1968 and are provided for a nowcast and for up to four quarters into the future.

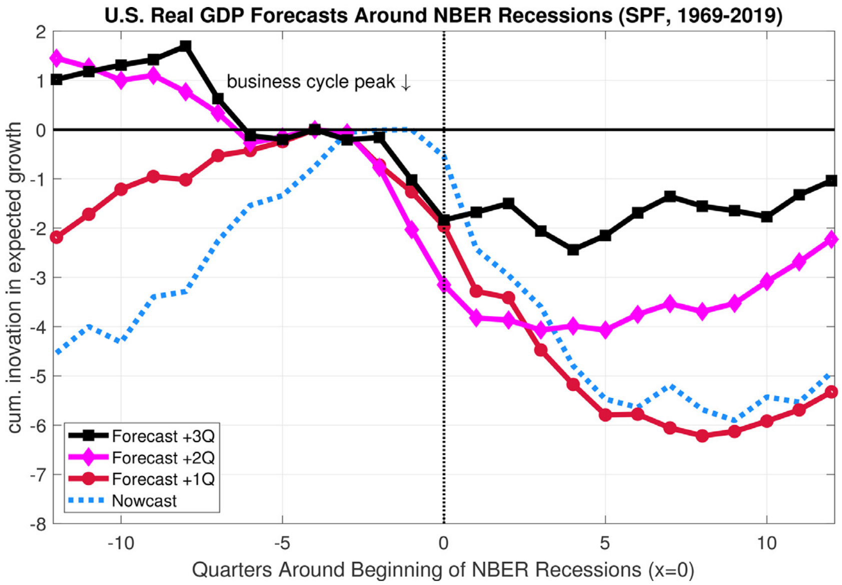

Fig. 3. Term structure of real GDP growth forecast revisions in recession event time (U.S. Sample, Q4/1969-2019, # Recessions: 6).

for the four quarterly horizons (nowcast, +1Q, +2Q, +3Q). Note that all four forecasts span a growth rate of one quarter and are directly comparable to each other.

Thus, this analysis allows me to compare short-term revisions in growth expectations versus longer term revisions in forward growth expectations. However, to derive a meaningful interpretation, I have to assume i) that survey expectations of professional forecasters are close to the expectations that are reflected by stock prices, and ii) that real GDP growth follows similar dynamics as real consumption growth and corporate cash flows, at least around the event window.[16]

The point estimates of the cumulative change in expectations are provided in Fig. 3. Numerical results and more detailed statistical inference are reported in Table OA.3 of the Online Appendix. First, focusing on the expected forward growth rate three quarters into the future (black line; the longest horizon available), I find that +3Q expectations drop -0.9% (t-stat: -1.7) at the business cycle peak (τ=-1). They drop another -0.8% (t-stat: -1.4) with the beginning of recessions (τ=0) and remain virtually flat during the remainder of the recession event time. In contrast, the +1Q and +2Q forecasts start dropping around the same time, but they drop steeper and significantly during the recessions. Finally, the nowcast drops with τ=0 and has the deepest cumulative drop during recessions (e.g., -4.8% within the first four quarters, t-stat: -2.2). These results suggest that cash flow news are mainly of contemporaneous nature around recessions.

In summary, I do not find support for the idea that revisions in longer horizon expectations of economic growth play a prominent role in explaining stock prices around recessions. The magnitude of revisions is quickly decreasing from the nowcast to the three quarter ahead forecast. Revisions in expectations beyond one year are not available. However, the term structure would have to be U-shaped to allow large revisions in long-term growth expectations. Such a pattern seems unusual.[17]

In the Online Appendix, Figure OA.2, I provide further results on the forward term structure of expected real GDP growth. These results show that there is variation in longer horizon forward growth rates (as suggested by long-run risk models). However, these variations are mainly unrelated to recessions and occur at a relatively low frequency.

4. Recession variance ratios

4.1. Motivation

It is well known from the literature that the stock return variance increases during recessions (e.g., Schwert, 1989; Hamilton and Lin, 1996; Nyberg, 2012; Jurado et al., 2015, among many others). However, stock returns have a price and cash

[16] The SPF also covers (nominal) corporate profit forecasts and real consumption forecasts as alternative measures of expected cash flows. I find that nominal corporate profit forecasts deflated by the GDP deflater behave very similarly to real GDP in event time. I prefer real GDP against real consumption forecasts (as, e.g., Andrei et al., 2019), because the consumption time-series does not start before 1981 and is covered by less survey participants.

[17] The weak one-quarter lead of short-term forecasts and the rather flat three-quarter ahead forecast are in line with the output forecast literature. For example, Chauvet and Potter (2013) report that the SPF has almost no forecast power beyond the one quarter ahead forecast. It is also plausible with the observation that end of period stock prices start declining a few months before the beginning of recessions. Naturally, survey expectations can faster incorporate new information than measured and quarterly time-aggregated real GDP.

flow component, and the two depend differently on different economic determinants. For example, as noted by Jurado et al. (2015, p.1178), the variance of stock price changes can depend on variation in risk aversion, while the variance of cash flow growth does not.

To see this, extend the definition of the log price change by dividend price changes in t+1 and t and re-arrange terms:

$$p_{t+1} - p_t = p_{t+1} - p_t + d_{t+1} - d_{t+1} + d_t - d_t = \underbrace{(p_{t+1} - d_{t+1}) - (p_t - d_t)}_{\textit{change in pd}} + \underbrace{(d_{t+1} - d_t)}_{\textit{change in d}} \tag{1}$$

Accordingly, the variance of stock prices can be decomposed as

$$\sigma^2_{p,t} = \sigma^2_{pd,t} + \sigma^2_{d,t} + 2\sigma_{d,pd}. \tag{2}$$

In words, the variance of price changes is equal to the sum of the variance of dividend growth, the variance of changes in the price-dividend ratio, and two times the covariance term.[18] Eq. (2) shows that an increase in uncertainty about dividend growth should impact the variance of dividend growth and the variance of prices, holding all else equal. On the other hand, an increase in the variance of changes in the price-dividend ratio affects only the variance of stock price changes and not the variance of dividend growth. Conceptually, the variance of changes in the price-dividend ratio is linked to innovations in expected returns or long-run cash flow news.[19] This must hold in any asset pricing model, as the relationship is purely based on an accounting identity.

My results in the previous section show that recessions are not anticipated by the price-dividend ratio (after accounting for the time-aggregation bias), mainly affect short-horizon expectations (the nowcast and one-quarter ahead forecast), and have little impact on longer horizons (future three quarter ahead forecast). For this reason, I conclude that objective as well as subjective long-run cash flow news are unlikely to play a prominent role around recessions.

Motivated by these considerations, I define the recession variance ratio of prices and dividend growth as:

$$\textit{recession variance ratio}(x) = \frac{Var(x \mid \textit{recession})}{Var(x \mid \textit{pre} - \textit{recession})}. \tag{3}$$

Comparing the recession variance ratio between prices and cash flows then allows me to infer the *relative importance* of changes in uncertainty (quantity of risk) and the changes in the perception of risk (price of risk). For example, suppose the variance of price changes increases by the same amount as the variance of dividend growth during recessions. In the absence of long-run cash flow news, the increase in the variance of price changes would then require the variance of expected return innovations to increase by a similar amount.

For practical purposes, I measure the pre-recession period over the quarterly event window τ=-5 to τ=-2 (i.e., one year) before the start of a recession (τ=0). I measure the recession variance over the quarterly event window τ=+2 to τ=+5 (i.e., one year after the start of a recession). This definition mitigates the use of overlapping data in the case of time-aggregated data (e.g., dividends). For statistical inference, I follow the methodology for changes of the price-dividend ratio in recession event time and bootstrap the data in event time 10,000 times (see Section 2).

Interpretation of recession-time conditional variances. The law of total variance states that the unconditional variance is equal to the expected conditional variance plus the variance of conditional means:

$$Var(x) = E[Var(x \mid \tau)] + Var[E(x \mid \tau)]. \tag{4}$$

The recession variance ratio is the ratio of the variance conditional on the time being pre-recession (τ=-5:τ=-2) or within a recession (τ=+2:τ=+5). The empirical counterparts can be viewed as two components of the full sample variance, where the variance outside of the recession event window and the variance of conditional means are the remaining parts. The recession variance ratio attempts to provide model-free evidence on the time-variation of variances.[20]

4.2. Empirical results

Panel A of Table 3 provides recession variance ratios and the components in the all countries sample. Panel B provides recession variance ratios for the U.S. sample.

In the all countries sample, dividend growth has a recession variance ratio of 1.7, which is significant different from 1.0. The variance of end of period stock price changes increases by a factor of 2.1 (90% confidence interval: {1.5, 2.9}) during recessions. The increase of the variance of time-aggregated stock price changes is closer to dividend growth (1.5, 90% c.i.: {1.09, 1.98}). The correlation coefficient between dividend growth and changes in the price-dividend ratio remains almost unchanged before and during recessions (-0.38 vs -0.33). In the U.S. sample, the recession variance ratio for dividend growth

[18] The same decomposition holds for point in time as well as for time-aggregated data, as it is purely based on an accounting identity. However, numerical results can be expected to differ as time-aggregated variances are downward biased, as shown by Working (1960).

[19] For example, in Campbell and Cochrane (1999), changes in the price-dividend ratio reflect changes in risk aversion. In Bansal and Yaron (2004) and Wachter (2006), changes in the price-dividend ratio reflect long-run cash flow news, or changes in the disaster risk probability. A detailed discussion including model simulations is provided in Section 6 of this paper.

[20] An alternative approach would be to rely on an explicit model of conditioning information. Duffee (2018, p.2144) provides a more detailed discussion in the context of inflation variance ratios.

Table 3
Recession variance ratios. This table reports the annualized variance (%) of log changes of stock prices ($\sigma^2_{p,t}$), the price-dividend ratio ($\sigma^2_{pd,t}$), and dividends ($\sigma^2_{d,t}$) around the beginning of recessions from 1951 to 2019. Below ($2 \times \sigma_{d,pd}$) is two times the covariance term between the price-dividend ratio and dividend growth. The variance of price changes can be decomposed as $\sigma^2_{p,t} = \sigma^2_{pd,t} + \sigma^2_{d,t} + 2\sigma_{d,pd}$. In Panel A, estimates are based on 42 recessions in 14 countries; in Panel B, estimates are based on 9 recessions in the U.S. (the event list is provided in Table OA.2, Online Appendix). The pre-recession variance is measured four quarters before the beginning of a recession (τ=-5 to τ=-2) and the recession variance is measured the four quarters after the beginning of a recession (τ=+2 to τ=+5). The reported recession variance ratio is the ratio of the recession variance and the pre-recession variance. 90% confidence intervals are based on a bootstrap resampling of the data in event time.

	End of Period Prices				Time-Aggregated		
	$\sigma^2_{p,t}$	$\sigma^2_{pd,t}$	$\sigma^2_{d,t}$	$2 \times \sigma_{d,pd}$	$\sigma^2_{p,t}$	$\sigma^2_{pd,t}$	$2 \times \sigma_{d,pd}$
	Panel A: All Countries (# Recessions: 42)						
	pre-recession variance, % p.a. (−5:−2):						
variance	3.57	3.66	1.98	−2.07	0.93	2.45	−3.49
c.i., 90%	[2.59, 4.47]	[2.58, 4.68]	[1.46, 2.44]		[0.71, 1.12]	[1.77, 3.11]	
	recession variance, % p.a. (+2:+5):						
variance	7.33	7.22	3.34	−3.22	1.36	3.64	−5.63
c.i., 90%	[5.76, 8.66]	[5.56, 8.72]	[2.39, 4.18]		[1.09, 1.58]	[2.74, 4.39]	
	recession variance ratio (recession/pre-recession):						
ratio	2.06	1.97	1.69		1.46	1.49	
c.i., 90%	[1.48, 2.92]	[1.38, 2.92]	[1.15, 2.44]		[1.09, 1.98]	[1.03, 2.14]	
	Panel B: U.S. Sample (# Recessions: 9)						
	pre-recession variance, % p.a. (−5:−2):						
variance	1.05	1.20	0.10	−0.25	0.20	0.19	−0.09
c.i., 90%	[0.60, 1.31]	[0.67, 1.49]	[0.04, 0.15]		[0.12, 0.24]	[0.11, 0.23]	
	recession variance, % p.a. (+2:+5):						
variance	5.00	4.99	0.06	−0.04	0.86	0.83	−0.03
c.i., 90%	[2.23, 6.87]	[2.25, 6.88]	[0.04, 0.08]		[0.47, 1.07]	[0.46, 1.02]	
	recession variance ratio (recession/pre-recession):						
ratio	4.76	4.16	0.63		4.32	4.42	
c.i., 90%	[2.24, 8.89]	[1.92, 7.79]	[0.33, 1.47]		[2.47, 7.21]	[2.53, 7.59]	

is not significantly different from 1.0, while the recession variance of end of period (time-aggregated) stock price changes increases by a factor of 4.8 (4.3).

Looking at the levels, in the all countries sample, the stock price variance is 3.57% (volatility of 18.89%) before recessions and as large as 7.33% (volatility of 27.07%) during recessions. In the U.S. sample, the respective variances (volatilities) are 1.05% (10.24%) vs 5.00% (22.36%).[21] Fig. 4 shows the variance estimates in event time to visualize the results in the table.[22] The figure suggests that the time-aggregated stock price variance in the sample for all countries responds with some lag and the recession variance ratio could be considered a conservative estimate.

To sum up, in the all countries sample, the variance of stock price changes increases at least by the same magnitude as the variance of dividend growth. In the U.S. sample, the variance of stock price changes increases even substantially more than the variance of dividend growth. In the absence of (subjective) long run cash flow news, this result implies that innovations to expected returns (discount rate news) must play an important role around recessions. In the following section, I provide a more detailed discussion of this argument.

4.3. Implications of recession variance ratios: A reduced form model

I first illustrate the implications of large stock price recession variance ratios using a simple reduced form model. This allows me to illustrate the importance of variation in expected returns without going into the details of fully specified models. In the absence of long-run cash flow news, a recession variance ratio of two (or larger) requires a large increase in the variance of expected returns. To illustrate this point, I adapt a reduced form model as presented in Cochrane (2008) and Cochrane (2005, Chapter 20). I assume that the economy is described by the following three equations:

$$r_t = z_t + \sigma_{r,t}\varepsilon_{r,t}, \tag{5}$$

$$z_t = bz_{t-1} + \sigma_{\delta,t}\delta_t, \tag{6}$$

$$\triangle d_t = \sigma_{d,t}\varepsilon_{d,t}, \tag{7}$$

[21] The relative large increase in the variance of U.S. stock price changes based on the event time methodology is similar to GARCH estimates that can be found in the literature. For example, the results reported in Nyberg (2012, Fig. 3) suggest a similarly large increase in the variance of stock returns around U.S. recessions.

[22] To reduce noise in the figure, I smooth estimates around the four neighboring observations using Gaussian weighting.

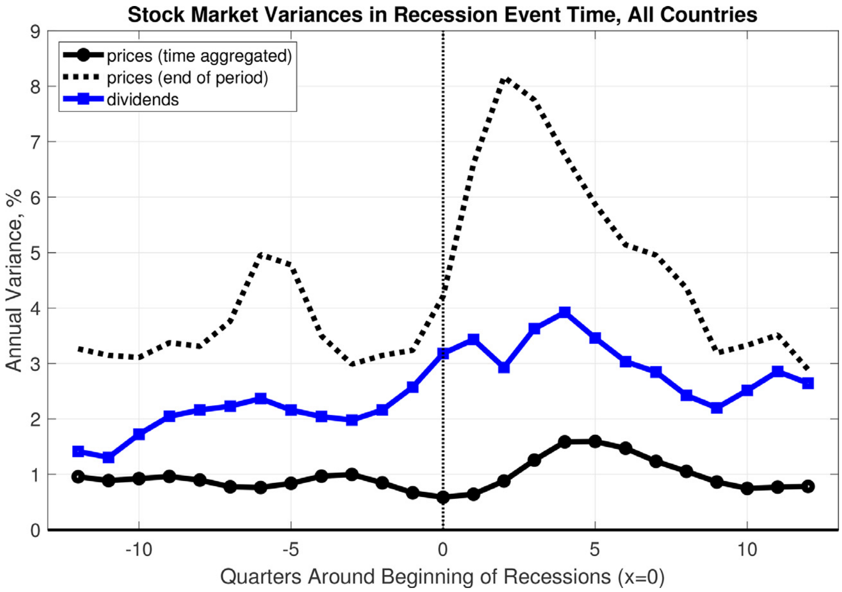

Fig. 4. Stock market variance in recession event time (All countries, 1951–2019, # Recessions: 42).

where z_t captures time-varying expected returns, r_t is the realized return, and $\triangle d_t$ is dividend growth. All variables are de-meaned and in logs; δ_t, $\varepsilon_{d,t}$ are standard normal shocks. As a result, the price-dividend ratio can only change when there are changes in expected returns (z_t). Return innovations, $\varepsilon_{r,t}$, are implied by the present-value relationship. Cochrane (2005) shows that:

$$d_{t+1} - p_{t+1} = b(d_t - p_t) + \frac{\sigma_{\delta,t+1}\delta_{t+1}}{1-\rho b}, \tag{8}$$

$$r_{t+1} = (1-\rho b)(d_t - p_t) + \left(\varepsilon_{d,t} - \frac{\rho}{1-\rho b}\sigma_{\delta,t+1}\delta_{t+1}\right), \tag{9}$$

$$\triangle p_{t+1} = (1-b)(d_t - p_t) + \left(\varepsilon_{d,t} - \frac{1}{1-\rho b}\sigma_{\delta,t+1}\delta_{t+1}\right), \tag{10}$$

which implies that the stock price variance can be decomposed as:

$$\sigma^2_{p,t} = \sigma^2_{d,t} + \sigma^2_{dp,t} + 2\rho_{d,dp}\sigma_{d,t}\sigma_{dp,t} \tag{11}$$

$$\sigma^2_{dp,t} = \frac{\sigma^2_{\delta,t}}{(1-\rho b)^2}. \tag{12}$$

In words, the stock price variance reflects the variance of innovations in dividends, plus the variance of innovations in expected returns, plus a covariance term. To calibrate the model, I use the same model parameters as suggested by Cochrane (2005) and set $b = 0.9$, $\rho = 0.96$. The other parameters are calibrated to match with the sample variance of dividends and stock prices $\sigma_d = 0.12$, $\sigma_\delta = 0.015$, $\rho_{d,dp} = 0.38$.

In the previous section of this paper (Table 3, Fig. 4), I find that the variance of price changes increases by a factor of 2.1 (U.S.: 4.8) during recessions. The variance of dividend growth increases by a factor of 1.7 (U.S.: not significantly different from one). How can it be explained that stock price volatility increases by at least the same amount or even more compared to cash flows?

Eq. 11 suggests a simple answer. As shown in Table 4, if the variance of dividend growth increases by 1.7 the variance of price changes should only increase by a factor of 1.36, holding all else equal. To push the recession variance ratio of stock prices up, discount rate shocks ($\sigma^2_{\delta,t}$) must be more volatile during recessions, or the correlation term in the covariance must go up. I find that the correlation does not change much in the empirical data. The more plausible route is to increase the variance of discount rate shocks $\sigma^2_{\delta,t}$. Increasing the variance of discount rates by a factor of 1.7 gives a variance ratio for stock price changes of 1.7; increasing by a large factor of 6 (!) gives variance ratio for stock price changes of 3.5.

My interpretation of these findings is that linking investor preferences and expectations such that discount rates are highly volatile during recessions is key to generate a large stock price variance during recessions. There is not sufficient variance in (realized) cash flows to explain the large variance of stock price changes.

Table 4
Recession variance ratios in a reduced-form model. This table shows the recession variance ratio for dividend growth and stock prices changes in a reduced form-model following Cochrane (2005). The variance of price changes can be decomposed as:

$$\sigma_{p,t}^2 = \sigma_{d,t}^2 + \sigma_{dp,t}^2 + 2\rho_{d,dp}\sigma_{d,t}\sigma_{dp,t},$$

where, $\sigma_{d,t}^2$ is the variance of dividends, $\sigma_{dp,t}^2 = \sigma_{\delta,t}^2/(1-\rho b)^2$is the variance of the dividend-price ratio with $\sigma_{\delta,t}^2$being the variance of innovations in expected returns, and $\rho_{d,dp}$is the correlation between dividend growth and the dividend-price ratio. The model parameters $\rho = 0.96$ and $b = 0.90$ are set to the same values as in Cochrane (2005). The other parameters are calibrated to match with the sample variance of dividends and stock prices (all countries sample). The recession variance ratio is the variance during recessions dividend by the variance before recessions.

		Cash Flows	Prices		
before recessions	σ_d	0.12	0.12	0.12	0.12
	$\rho_{d,dp}$		0.38	0.38	0.38
	σ_δ		0.015	0.015	0.015
during recessions	σ_d	$0.12 \times \sqrt{1.7}$	$0.12 \times \sqrt{1.7}$	$0.12 \times \sqrt{1.7}$	$0.12 \times \sqrt{1.7}$
	$\rho_{d,dp}$		0.38	0.38	0.38
	σ_δ		0.015	$0.015 \times \sqrt{1.7}$	$0.015 \times \sqrt{6}$
recession variance ratio		1.70	1.36	1.70	3.54

4.4. Implications of recession variance ratios: Simulations of standard asset pricing models

I simulate artificial data in the habit model (Campbell and Cochrane, 1999), long run risk model (Bansal and Yaron, 2004), and the rare disaster risk model (Wachter, 2013). As described in detail in the Online Appendix, I use calibrations of the models as close as possible to the original publications. I then search for recessions (i.e., large drops of realized consumption) in the simulated data and re-calculate the recession variance ratio for stock price changes as close as possible to the empirical data. The results are tabulated in Table OA.7 of the Online Appendix.

Overall, I find that the standard models do not generate a large recession variance ratio for stock price changes.[23] The model with the largest recession variance ratio for stock price changes is the habit model. This is intuitive, as this model is a prime example of a model where expected returns increase due to a sudden decline in consumption and the habit. However, the stock price variance increases only by a factor of 1.1. This illustrates that the required variance in expected returns is difficult to achieve for leading asset pricing models. I discuss the implications in more detail in Section 6.

5. Further empirical results

The baseline results focus on the period 1951 - 2019 (all countries: 42 recessions, U.S. 9 recessions), because this allows me to study the data sampled at the quarterly frequency. Quarterly data are helpful to pinning down the timing of prices and cash flows. For a longer sample period, 1872 - 2019, I also analyze an annually sampled dataset that covers a total of 75 recessions (U.S.: 16).[24]

Fig. 5 shows the results for the longer sample (numerical results can be found in Table OA.5, Online Appendix). The figure re-samples the baseline results using quarterly data, although the picture is indeed less granular. I find that end of period stock prices drop with the beginning of recessions, while dividends drop with a lag of one period. In the all countries sample, the maximum drop of stock prices is -17% (t-stat: -4.9), while the maximum drop in dividends is -11% (t-stat: -3.7). As in the quarterly data, stock prices drop more than dividends, indicating an important role for time-varying expected returns. In the Online Appendix, I report the recession variance ratio for the annual data (Table OA.6). The results are similar to the baseline results using quarterly data. In the all countries sample, I find that the recession variance ratio of stock price changes is 1.9 and for dividend growth 2.2, both are significant different from 1.0. For the U.S., also in the annual data the stock price variance increases substantially more compared to dividends (recession variance ratios of 6.6 vs 1.4). In summary, the annual sample gives very similar results when compared to the quarterly sample.

[23] For the rare disaster risk model, I provide results conditional on samples with a realized rare disaster and without. Given that I focus on "normal" recessions in this paper, it is plausible to compare with the model conditional on that a rare disaster has not been observed. This is common practice in this literature when comparing the model to empirical observations in the post war period (e.g., Wachter, 2013)

[24] There is a trade-off involved when choosing between quarterly and annual data in event studies (see Morse, 1984). A recession can occur as early as in January or as late as in December. In the annual dataset, it is assumed that the recession always happens at the same point in time within the year, which results in observations less precisely measured. On the other hand, the number of recessions is increased, which means that more (but less precisely measured) observations are available. Another disadvantage of the annual data is that I cannot compute time-aggregated stock prices for all countries, due to a lack of higher frequent observations.

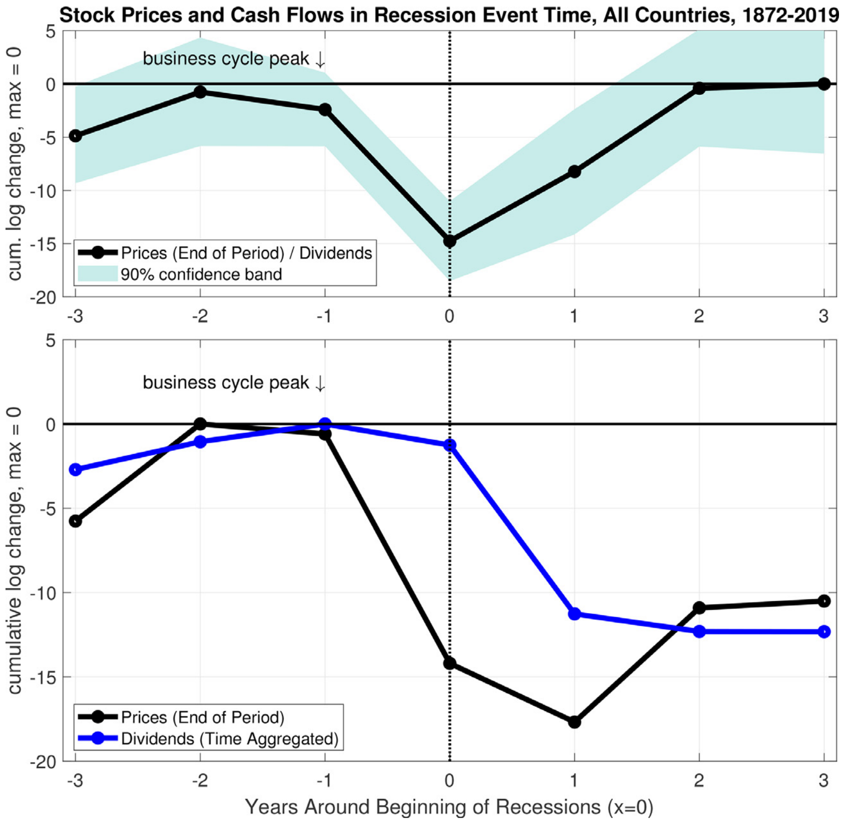

Fig. 5. The stock market in recession event time: Annual sample, 1872–2019 (All Countries, # Recessions: 75).

6. What are the implications for asset pricing theories?

In this section, I first discuss in more detail why standard asset pricing theories (long risk, habits, rare disasters) fail to predict high stock price variances around recessions. Second, I provide a broader discussion of modifications and alternative channels that have been proposed in the recent literature and how they might help to explain the empirical facts presented in this paper. Although my discussion is not necessarily conclusive, I argue that the timing of the price-dividend ratio around recessions and the recession variance ratio proposed in this paper are appealing metrics that can be used to benchmark new models.

6.1. Long-run risk

In the long-run risk model by Bansal and Yaron (2004), cash flows have three ingredients. Traditional one period consumption shocks ("short-run risk"), a persistent component in consumption growth ("long-run risk") and time-varying volatility. Interestingly, dividend growth only shares the persistent component (the long-run risk) of consumption risk and is, thus, to some degree predictable.

In the Online Appendix, Table OA.7 and Figure OA.5, I report results from model simulations to see how the stock market should behave around recessions in the long-run risk model.[25] I find that the time-aggregated price-dividend ratio starts dropping about one year ahead of recessions. Intuitively, by searching for business cycle peaks, one identifies periods when long-run growth and short run growth happen to be low at the same time, as both feed into consumption. The forward-looking investor, in turn, recognizes that future dividends will be lower going forward (due to the long-run component) and as a consequence stock prices decline several quarters before the recession actually starts. The figure also shows that the price-dividend ratio does not move much with the beginning of recessions.

The long-run risk model also features time-varying volatility. However, in the long-run risk model, changes in volatility are, by construction, uncorrelated to short-run or long-run consumption risk. Volatility risk is modeled as an "additive"

[25] Further simulation details are provided in the Online Appendix. I use the same model parameters as in Bansal and Yaron (2004), to make my results comparable to the literature. The re-calibration of the long-run risk model provided by Bansal et al. (2012) model leads to similar results; see Table OA.7.

channel. As a result, there is no systematic relationship between volatility and recessions (a fall in consumption).[26] The long-run risk model might be better able to explain the stock market around recessions by linking short-run consumption shocks to long-run shocks and / or the volatility process. My results are therefore supportive for such variations of the model.[27]

6.2. Habits

In the habit model by Campbell and Cochrane (1999), consumption is an unpredictable standard i.i.d. process. However, the price of risk and expected returns are time-varying. Consumption is evaluated relative to a habit, which is a reference level of consumption that can be thought of as a moving average of past realizations. As a result, expected returns are low after observing a "good run" of consumption shocks. This is by construction around the peak of the business cycle. In contrast, expected returns rise when consumption comes closer and closer to the reference level, i.e., during "significant declines" in economic activity. The habit model by Campbell and Cochrane (1999) is a prime example of a theory that features a link of short-term cash flow news to time-varying expected returns.

I report results from simulated recessions in the habit model in Table OA.7 and Figure OA.6 of the Online Appendix. It is easy to see that the price-dividend ratio drops contemporaneously with the beginning of recessions (and dividends). This reflects the fact that expected future cash flow growth is constant, but stock prices fall more than cash flows because expected returns go up. End of period stock prices — which are frequently used in empirical research — lead dividends and consumption for many quarters. The time-aggregation bias makes dividends and consumption predictable, even in the habit model. However, my simulation of recessions in the habit model shows that the model only increases stock price volatility by a factor of 1.1. When I condition on the largest 20% of the simulated recessions, I find a recession variance ratio of 1.5.[28] But then, the drop in consumption (and dividends) must be very large. The habit model therefore requires a stronger amplifier.[29]

6.3. Rare disasters

In the Wachter (2013) version of the rare disaster risk model, consumption has an unpredictable standard i.i.d. component similar to the habit model. In addition, with a small probability (on average 3.6% in annual terms), consumption is subject to a large decline — the disaster event. In this model, the price-dividend ratio is a decreasing function of the rare disaster probability; prices are depressed when investors believe that a disaster is around the corner.

Table OA.7 and Figure OA.7 of the Online Appendix show stock returns around simulated recessions according to the model. Following Wachter (2013), I present results based on all simulated recessions as well as results when rare disasters are conditioned out. The rare disasters literature (e.g., Barro, 2006; Barro and Ursua, 2008; Wachter, 2013) treats the U.S. postwar sample as a period in which no rare disaster took place and the "conditional" results are usually used to compare with the post-war sample. Around "normal" recessions, the price-dividend ratio is expected to be rather flat. The reason is that there is no structural link between the standard i.i.d. consumption risk and the time-varying rare disaster risk probability (and in turn the price-dividend ratio).[30] Similar to the long-run risk model, a potentially missing ingredient is a mechanism where a series of bad short-term consumption shocks (recessions) spill over to the disaster risk probability.

6.4. Intermediary-based models

Muir (2017) argues that asset pricing theories where asset prices are related to the health of the financial system are promising in explaining stock markets around financial crises (e.g., He and Krishnamurthy, 2013; Adrian and Shin, 2014; Brunnermeier and Sannikov, 2014). If the health of the financial system is also affected in "normal" recessions, at least on average, such theories might also help to explain high recession variance ratios of stock prices.

In Fig. 6, I report the estimated growth of bid-ask spreads of individual U.S. stocks averaged over three size groups (numerical results are reported in Table OA.4 of the Online Appendix).[31] The bid-ask spread is a central measure of trans-

[26] In Figure OA.8 of the Online Appendix, I provide results for the reverse experiment when the event time is specified by the trough in volatility. Average realized dividends are barely affected by increasing volatility.

[27] For example, Branger et al. (2016) propose a model where short-run and long-run consumption risks are linked with each other. An initial drop in consumption can unfold into a disaster. Tédongap (2014) presents a long run-risk model where consumption volatility is modeled as a GARCH process, i.e. short-run risk and volatility risk are linked.

[28] Figure OA.9 of the Online Appendix.

[29] Belo et al. (2015) show that changing the dividend (and consumption) dynamics can substantially change the properties of stock returns in the classic models, including the habit model. Bekaert and Engstrom (2017) propose a bad-environment good-environment process for consumption growth to generate an asymmetric relationship between option implied volatility and declines in consumption growth. Schreindorfer (2020) considers that investors might be disappointment-averse. As a result, asset prices respond also asymmetric to recessions.

[30] Actually, in the figure the price-dividend ratio slightly declines with the beginning of recessions. However, this result comes from the fact that I condition out rare disaster events (as in Wachter, 2013), which means that I systematically remove observations with high disaster probability around the event τ=0. However, this effect is rather small in magnitude.

[31] I rely on a novel estimator of bid-ask spreads based on transaction data as proposed by Ardia et al. (2021). This estimator is an improved and considerably more accurate version of the popular Roll (1984) estimator.

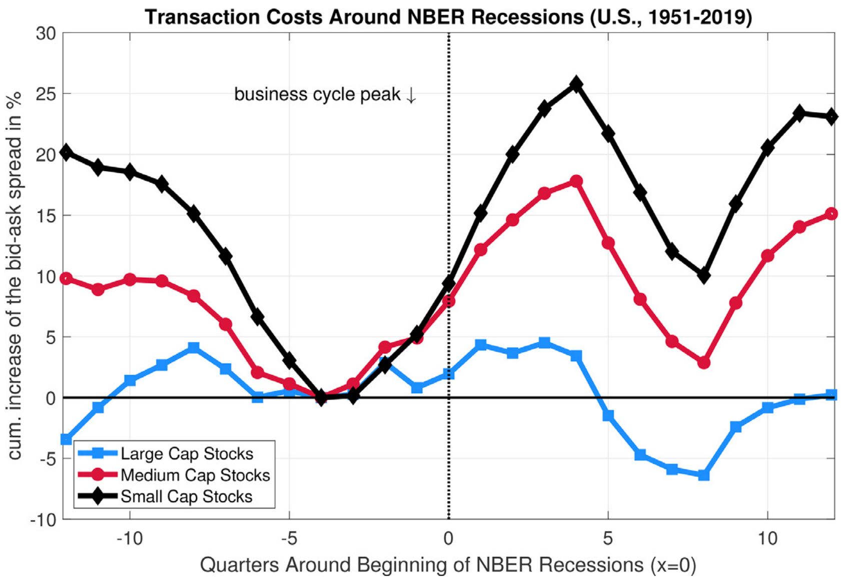

Fig. 6. Transaction costs in recession event time (U.S. Sample, # Recessions: 9).

action costs and can be more broadly regarded as a proxy of the "health" of the financial system in a wide sense.[32] While I find that transaction costs significantly increase for small cap stocks by 21% after four quarters (t-stat: 2.2), and somewhat weaker for mid cap stocks by 13% (t-stat: 1.7), there is only an increase of 3% (t-stat: 0.4) in transaction costs for large cap stocks. The increase for small cap stocks is substantial but not catastrophic from the investor's point of view either. While an in-depth analysis is beyond the scope of the paper, financial market conditions as measured by transaction costs are unlikely to deteriorate sufficiently during "normal" recessions to explain the facts.

6.5. Subjective beliefs

Another possibility is that stock market investors have incorrect cash flow beliefs around recessions. For example, they expect that the drop in economic activity is larger than it actually is. Indeed, De La O and Myers (2021), and Bordalo et al. (2020), provide evidence that the level of subjective cash flow beliefs explain a large part of the variation in the level of stock price-fundamental ratios. While De La O and Myers (2021) stress the importance of short-term expectations (one-year and two-year horizons), Bordalo et al. (2020) argue that longer horizon expectations also play an important role.

My results based on survey forecasts reported in Section 3 show that *revisions* in the very short-term expectations (nowcast, up to two quarters) vary a lot around recessions, while revisions already at the horizon of just three quarters remain almost flat. This result indicates that the very short-term horizon is important around recessions, not longer horizons (three quarters). It might be possible that the high recession variance ratio of stock price changes is to some extend explained by high subjective uncertainty about*the current state* of the economy. It would be therefore interesting to break down the one-year horizon in De La O and Myers (2021) into shorter intervals and to study in more detail which type of events trigger revisions in expectations. My finding that the very short-term revisions of expectations (nowcast, one quarter forecast) vary the most is therefore also supportive to the literature that stresses learning and the importance of disagreement about the state of the current business cycle (e.g., Andrei et al., 2019; Cujean and Hasler, 2017).

Due to data limitations, I do not examine subjective expectations over very long horizons in recession event time, such as future 10-year growth. Therefore, I cannot rule out the possibility that they play a role. However, it would require that revisions in expectations show a U-shaped pattern where the short and very long-term expectations are affected during recessions, but not the medium-term. I argue that this would be a rather unusual pattern. Yet, an interesting implication would be that investors confuse "normal" business cycle patterns with more structural changes (financial crises and war episodes are excluded from my analysis). Alternatively, equity investors' subjective expectations could be decoupled from the participants in the survey of professional forecasters. In any case, the recession variance ratio is also an informative metric for this strand of the literature. The recession variance ratios provide a quantitative indication of how large these revisions in expectations would have to be.

[32] The literature shows that the bid-ask spread is in theory determined by, e.g., inventory risk and risk capacity of intermediaries (Stoll, 1978), asymmetric information (Bagehot, 1971), and search frictions (Duffie et al., 2005). If the financial system deteriorates in a significant way, it is plausible to expect increasing transaction costs (see, e.g., Acharya and Pedersen, 2005; Acharya and Pedersen, 2019).

7. Conclusion

Stock market prices reflect the economy at a particular point in time, while dividends and variables commonly used to measure economic growth (real GDP) reflect the economy over a period of time. This feature leads to the time-aggregation bias and allows stock prices to anticipate economic developments within the time-aggregation period. Once I account for the time-aggregation bias, I find that stock prices drop almost contemporaneously with dividends and the business cycle. I interpret this finding as evidence against the idea that cash flows (and economic growth) have an economically rooted predictable component at the business cycle frequency.

Furthermore, I find that stock prices drop substantially more than dividends around recessions; the variance of stock price changes increases at least by the same amount as the variance of dividend growth. This result indicates that innovations in expected returns are highly volatile during recessions. I illustrate that these facts are difficult to explain with standard asset pricing theories and provide a discussion of potential channels that might be able to account for the observed asymmetries.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2022.07.004.

References

Acharya, V.V., Pedersen, L.H., 2005. Asset pricing with liquidity risk. J Financ Econ 77 (2), 375–410.
Acharya, V.V., Pedersen, L.H., 2019. Economics with market liquidity risk. Critical Finance Review 8 (1–2), 111–125.
Adrian, T., Shin, H.S., 2014. Procyclical leverage and value-at-risk. Review of Financial Studies 2014 (27), 373–403.
Anarkulova, A., Cederburg, S., ODoherty, M.S., 2022. Stocks for the long run? evidence from a broad sample of developed markets. J financ econ 143 (1), 409–433.
Andrei, D., Hasler, M., Jeanneret, A., 2019. Asset pricing with persistence risk. Review of Financial Studies 32 (7), 2809–2849.
Ardia, D., Guidotti, E., Kroencke, T.A., 2021. Efficient estimation of bid-ask spreads from open, high, low, and close prices. Working Paper.
Atanasov, V., Møller, S.V., Priestley, R., 2020. Consumption fluctuations and expected returns. Journal of Finance 75, 1677–1713.
Backus, D., Chernov, M., Zin, S., 2014. Sources of entropy in representative agent models. Journal of Finance 69 (1), 51–99.
Backus, D.K., Routledge, B.R., Zin, S.E., 2010. The cyclical component of us asset returns. Working Paper.
Bagehot, W., 1971. The only game in town. Financial Analysts Journal 27 (2), 12–14.
Bansal, R., Kiku, D., Yaron, A., 2012. An empirical evaluation of the long-run risks model for asset prices. Critical Finance Review 1, 183–221.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. Journal of Finance 59, 1481–1509.
Barro, R.J., 2006. Rare disasters and asset markets in the twentieth century. Quarterly Journal of Economics 121 (3), 823–866.
Barro, R.J., Ursua, J.F., 2008. Macroeconomic crises since 1870. Brookings Pap Econ Act 1 (Spring), 255–350.
Beason, T., Schreindorfer, D., 2022. Dissecting the equity premium. Journal of Political Economy. Forthcoming
Bekaert, G., Engstrom, E., 2017. Asset return dynamics under habits and bad environment-good environment fundamentals. Journal of Political Economy 125 (3), 713–760.
Bekaert, G., Engstrom, E.C., Xu, N.R., 2022. The time variation in risk appetite and uncertainty. Manage Sci 68 (6), 3975–4753.
Belo, F., Collin-Dufresne, P., Goldstein, R.S., 2015. Dividend dynamics and the term structure of dividend strips. Journal of Finance 70 (3), 1115–1160.
van Binsbergen, J., Brandt, M., Koijen, R., 2012. On the timing and pricing of dividends. American Economic Review 102 (4), 1596–1618.
Binsbergen, J.V., Koijen, R.S., 2010. Predictive regressions: a present-value approach. Journal of Finance 65 (4), 1439–1471.
Bordalo, P., Gennaioli, N., Porta, R.L., Shleifer, A., 2020. Expectations of fundamentals and stock market puzzles. Working Paper.
Branger, N., Kraft, H., Meinerding, C., 2016. The dynamics of crises and the equity premium. Review of Financial Studies 29 (1), 232–270.
Brunnermeier, M.K., Sannikov, Y., 2014. A macroeconomic model with a financial sector. American Economic Review 104, 379–421.
Campbell, J.Y., 1991. A variance decomposition for stock returns. Economic Journal 101, 157–179.
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107 (2), 205–251.
Campbell, J.Y., Giglio, S., Polk, C., 2013. Hard times. Review of Asset Pricing Studies 3, 95–132.
Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, New Jersey.
Campbell, J.Y., Shiller, R.J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1, 195–228.
Campbell, S.D., Diebold, F.X., 2009. Stock returns and expected business conditions: half a century of direct evidence. Journal of Business & Economic Statistics 27, 266–278.
Chauvet, M., 1999. Stock market fluctuations and the business cycle. J Econ Soc Meas 25 (3–4), 235–257.
Chauvet, M., Potter, S., 2013. Forecasting Output, Vol. 2. in Handbook of Economic Forecasting, Elsevier.
Chen, L., 2009. On the reversal of return and dividend growth predictability: a tale of two periods. J Financ Econ 92, 128–151.
Cieslak, A., Vissing-Jorgensen, A., 2021. The economics of the fed put. Review of Financial Studies 34 (9), 4045–4089.
Cochrane, J.H., 1996. A cross-sectional test of an investment-based asset pricing model. Journal of Political Economy 104, 572–621.
Cochrane, J.H., 2005. Asset Pricing. Princeton University Press, New Jersey.
Cochrane, J.H., 2008. The dog that did not bark: a defense of return predictability. Review of Financial Studies 21, 1533–1575.
Cooper, I., Priestley, R., 2009. Time-varying risk premiums and the output gap. Review of Financial Studies 22 (7), 2801–2833.
Cujean, J., Hasler, M., 2017. Why does return predictability concentrate in bad times? Journal of Finance 72 (6), 2717–2758.
Dangl, T., Halling, M., 2012. Predictive regressions with time-varying coefficients. J Financ Econ 106 (1), 157–181.
De La O, R., Myers, S., 2021. Subjective cash flow and discount rate expectations. Journal of Finance 76 (3), 1339–1387.
Dew-Becker, I., Giglio, S., Le, A., Rodriguez, M., 2017. The price of variance risk. J Financ Econ 123 (2), 225–250.
Duffee, G.R., 2018. Expected inflation and other determinants of treasury yields. Journal of Finance 73 (5), 2139–2180.
Duffie, D., Gârleanu, N., Pedersen, L.H., 2005. Over-the-counter markets. Econometrica 73 (6), 1815–1847.
Fama, E.F., 1990. Stock returns, expected returns, and real activity. Journal of Finance 45, 1089–1108.
Fama, E.F., French, K.R., 1989. Business conditions and expected returns on stocks and bonds. J Financ Econ 25, 23–49.
Ferson, W.E., Harvey, C.R., 1991. The variation of economic risk premiums. Journal of Political Economy 99, 385–415.
Fushing, H., Chen, S.-C., Berge, T.J., Óscar Jordá, 2010. A chronology of international business cycles through nonparametric decoding. Federal Reserve Bank ofSan Francisco working Paper Series.
Gabaix, X., 2012. An exactly solved framework for ten puzzles in macro-nance. Quarterly Journal of Economics 127 (2), 645–700.
Golez, B., Koudijs, P., 2018. Four centuries of return predictability. J Financ Econ 127, 248–263.

Gourio, F., 2008. Time-series predictability in the disaster model. Finance Research Letters 5 (4), 191–203.
Hamilton, J.D., Lin, G., 1996. Stock market volatility and the business cycle. Journal of Applied Econometrics 11 (5), 573–593.
Harrison, P., Zhang, H.H., 1999. An investigation of the risk and return relation at long horizons. Review of Economics and Statistics 81, 399–408.
He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. AmericanEconomic Review 103, 732–770.
Henkel, S.J., Martin, J.S., Nardari, F., 2011. Time-varying short-horizon predictability. J Financ Econ 99 (3), 560–580.
Jagannathan, R., Liu, B., 2019. Dividend dynamics, learning, and expected stock index returns. Journal of Finance 74 (1), 401–448.
Jank, S., 2015. Changes in the composition of publicly traded firms: implications for the dividend-price ratio and return predictability. Manage Sci 61, 1362–1377.
Jordá, O., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95 (1), 161–182.
Jordá, O., Schularick, M., Taylor, A.M., 2011. Financial crises, credit booms, and external imbalances: 140 years of lessons. IMF Economic Review 59 (2), 340–378.
Jurado, K., Ludvigson, S.C., Ng, S., 2015. Measuring uncertainty. American Economic Review 105, 1177–1216.
Kozak, S., Santosh, S., 2020. Why do discount rates vary? J financ econ 137 (3), 740–751.
Kroencke, T.A., 2017. Asset pricing without garbage. Journal of Finance 72, 47–98.
Larrain, B., Yogo, M., 2008. Does firm value move too much to be justified by subsequent changes in cash flow? J financ econ 87, 200–226.
Lettau, M., Ludvigson, S.C., 2009. Measuring and modeling variation in the risk return trade-off. Handbooks in Finance, Vol. Volume 1, Chapter 11. Handbook of Financial Econometrics: Tools and Techniques.
Lustig, H., Verdelhan, A., 2012. Business cycle variation in the risk-return trade-off. J Monet Econ 59, 35–49.
Maio, P.F., Cooper, I., 2019. Asset growth, profitability, and investment opportunities. Manage Sci 65 (9), 3949–4450.
Martin, I., 2017. What is the expected return on the market? Quarterly Journal of Economics 132, 367–433.
Morse, D., 1984. An econometric analysis of the choice of daily versus monthly returns in tests of information content. Journal of Accounting Research 22, 605–623.
Muir, T., 2017. Financial crises and risk premia. Quarterly Journal of Economics 132, 765–809.
Møller, S.V., Rangvid, J., 2018. Global economic growth and expected returns around the world: the end-of-the-year effect. Manage Sci 64 (2), 495–981.
Nyberg, H., 2012. Risk-return tradeoff in u.s. stock returns over the business cycle. Journal of Financial and Quantitative Analysis 47 (1), 137–158.
Rangvid, J., Schmeling, M., Schrimpf, A., 2014. Dividend predictability around the world. Journal of Financial and Quantitiative Analysis 49, 1255–1277.
Rapach, D.E., Strauss, J.K., Zhou, G., 2010. Out-of-sample equity premium prediction: combination forecasts and links to the real economy. Review of Financial Studies 23 (2), 821–862.
Rietz, T.A., 1988. The equity risk premium: a solution. J Monet Econ 22 (1), 117–131.
Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39 (4), 1127–1139.
Schreindorfer, D., 2020. Macroeconomic tail risks and asset prices. Review of Financial Studies 33 (8), 3541–3582.
Schwert, G.W., 1989. Why does stock market volatility change over time? Journal of Finance 44, 1115–1153.
Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? Amercian Economic Review 71, 421–436.
Stoll, H., 1978. The supply of dealer services in securities markets. Journal of Finance 33 (4), 1133–1151.
Taio, G.C., 1972. Asymptotic behaviour of temporal aggregates of time series. Biometrika 59, 525–531.
Tédongap, R., 2014. Consumption volatility and the cross-section of stock returns. Rev Financ 19 (4), 1–39.
Vassalou, M., 2003. News related to future GDP growth as a risk factor in equity returns. J Financ Econ 68, 47–73.
Wachter, J.A., 2006. A consumption-based model of the term structure of interest rates. J Financ Econ 79, 365–399.
Wachter, J.A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? Journal of Finance 68 (3), 987–1035.
Warsh, K., 2016. Fed policy running against capital investment. Squawk Box.
Working, H., 1960. Note on the correlation of first differences of averages in a random chain. Econometrica 28, 916–918.
Zviadadze, I., 2021. Term structure of risk in expected returns. Review of Financial Studies 34 (12), 6032–6086.