Recessions and the Stock Market

TIM A. KROENCKE*

This version: 31st January 2020

Abstract

Why do stock prices fall more sharply than dividends around recessions? This paper provides an assessment of alternative potential economic channels suggested in the literature. One possible explanation is that stock prices fall in anticipation of low future cash flows. I find that prices and cash flows drop contemporaneously, which speaks against such a channel. Alternatively, prices drop because expected returns are rising. I find that price volatility increases substantially more than cash flow volatility during recessions, which suggests that changes in the price of risk play an important role. These results allow for a fresh empirical assessment of competing asset pricing theories.

*University of Neuchâtel, Switzerland. Contact: tim.kroencke@unine.ch. I would like to thank Pedro Barroso, Jesper Rangvid and Maik Schmeling for detailed comments and the seminar participants of the Aalto Business School, University of St. Gallen, University of Fribourg and the participants in conferences where the paper was presented for their helpful suggestions. First draft: May 2018.

I. Introduction

Around recessions, the dividends paid out by stock companies fall on average by about 10%. Stock prices drop even more sharply by about 20%. These large declines happen fast; the average duration of recessions is less than one year. If investors respect present-value relationships, all movements in stock prices that are not reflected in changes of today's dividends must come from i), changes in expectations about *future* cash flows, or ii), changes in expected returns (Shiller, 1981; Campbell and Shiller, 1988; Campbell, 1991). But how much of the additional 10% drop in stock prices around recessions can be attributed to which of these two channels?

Competing asset pricing theories make fundamentally different predictions (e.g., Campbell and Cochrane, 1999; Bansal and Yaron, 2004; Wachter, 2013).[1] If prices fall around recessions because investors expect lower *future* cash flows, one should observe that prices fall *before* the economy is in a recession. Intuitively, if there is a predictable component in cash flows, forward-looking prices should reflect the bad times of tomorrow already today. Another explanation could be that expectations about future cash flows do not change a lot, but expected returns are elevated during recessions such that future dividends get more heavily discounted. As a result, prices fall *contemporaneously*, but *more* than dividends. For that reason, I argue that timing and deepness of the fall in prices relative to cash flows provides information on the channels that drive asset prices. The goal of this paper is to pin down the theoretical predictions of popular models in greater detail and to confront them with the empirical data. I go a step further and argue that the behaviour of the stock market around recessions restricts the channels that drive the market, thus enabling a new empirical evaluation of competing theories that has not yet been explored in detail by the literature.

An important empirical issue that arises in this exercise is that stock prices are market-based data which are easily observed at the end of a period. In contrast, earnings, dividends, and consumption are all flow variables that are usually measured over an entire period, typically

[1] See also, e.g., Cochrane (2017), Campbell (2017), or Ferson (2019) for comprehensive summaries.

a year (usually to avoid seasonality). These are "time-aggregated" data. However, comparing end of period data with time-aggregated data gives rise to a time aggregation bias (e.g., Working, 1960; Taio, 1972; Breeden, Gibbons, and Litzenberger, 1989). By construction, time-aggregated data are a weighted moving average of past observations and as such lag against end of period data.[2] Thus, to be sure that I compare apples with apples in my empirical analysis, I convert end of period stock prices to time-aggregated prices such that the timing of prices corresponds closely to the timing of cash flows.[3]

My empirical analysis focuses on quarterly U.S. data of stock prices, earnings, dividends, and consumption around recessions as defined by the NBER business cycle dating committee. The baseline results focus on the post-war period from 1950 to 2016. This sample allows me to study quarterly data, which is important to pin down the timing of the variables as precisely as possible. I find that stock prices, dividends, earnings, and consumption start falling contemporaneously with the beginning of a recession. As prices drop more compared to dividends, the price-dividend ratio declines with the beginning of recessions as well. I interpret these results as direct evidence against the idea that cash flows and consumption have a predictable component at the business cycle frequency such that stock prices anticipate recessions. I corroborate these findings based on market data by studying the forward term structure of expected economic growth as reported in the Survey of Professional Forecasters (SPF). I find that expected future real GDP growth does not anticipate recessions and these expectations are only mildly revised down. Instead, most revisions in expected growth take place at the short-term horizon.[4]

To compare my empirical results to the theoretical counterfactual, I simulate "recessions" in

[2]This bias is big. I simulate artificial data for the canonical consumption-based asset pricing model that features constant expected returns and unpredictable cash flows (Breeden, 1979). In this model, the true price-dividend ratio is constant, no matter how bad the recession is. However, I show that if prices are measured end of period and dividends are time-aggregated — as is usually done in empirical applications — the price-dividend ratio will appear to anticipate future dividends and will strongly predict recessions.

[3]Dividends and earnings have a strong seasonal pattern, so it is not possible to construct non-time-aggregated versions of these variables. Time-aggregated prices have been previously used by, e.g., Cochrane (1996) and Kroencke (2017) to reduce the time-aggregation bias.

[4]As reported in the Internet Appendix, I do find variation in longer horizon expected growth in the data. But these shifts in long horizon expectations are not systematically related to recessions.

the long-run risk model of Bansal and Yaron (2004). In this model, revisions about expectations of future cash flows take a central role. I show that the average recession is, by construction, also a period of predictable low growth. Because of the predictable component of cash flows, the (time-aggregated) price-dividend ratio starts dropping about one year ahead of recessions and is rather flat during recessions. This is about one year too early compared to the empirical data.[5]

I show that this result transfers to models that feature cash flow predictability at higher frequencies. Time-varying rare disasters as studied in Wachter (2013) were previously used to describe the stock market around recessions and periods of crises (e.g., Seo and Wachter, 2018).[6] However, they predict very similar to long run risk models that stock markets start dropping in anticipation of recessions as the probability of a "disaster" tends to increase ahead of the actual recession.

Next, I turn to the discount rate channel to explain why prices drop so much during recessions. On the outset, there are at least two sources that could give rise to expected returns, a), cash flows are more risky during recessions, or b), the price of risk increases as investors get more risk averse. My event study type approach allows me to estimate non-parametrically the volatility of stock prices, earnings, dividends, and consumption before and after the beginning of recessions. In line with a large literature (starting with Schwert, 1989; most recently Boguth and Kuehn, 2013; Bansal, Kiku, Shaliastovich, and Yaron, 2014; Tédongap, 2014; Jurado, Ludvigson, and Ng, 2015), I find that financial and macroeconomic uncertainty increases during recessions. Because I look at prices and cash flows separately, I can add to the literature by comparing the rise in volatilities between the different components of stock returns. The ratio of the recession variance over the pre-recession variance (which I dub the "the recession variance ratio" in the following) is as large as 4 for stock prices, 2.7 for earnings and below 2 for dividends

[5] I get similar results for the recent recalibration of the long-run risk model in Bansal, Kiku, and Yaron (2012).

[6] The original rare disaster model studied by Rietz (1988) and Barro (2006) have a constant disaster risk probability and feature 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.

and consumption. Stock prices get even more volatile compared to their own cash flows.

What does a model need to explain the data? I make some simple back of the envelope calculations to better understand the implications of the empirical results. Changes in stock prices are the sum of innovations in dividends and innovations in expected returns. Intuitively, to get a stock price variance that increases more than cash flow variance, I need the variance of expected return innovations to increase by a huge factor of 5. This is in line with the idea that the price of risk must be larger during recessions, so that expected returns become more sensitive to news. The prime example of a model that features such a mechanism is the habit model by Campbell and Cochrane (1999).

To further compare my empirical results to the theoretical counterfactual, I also simulate "recessions" in the habit model. Cash flows are unpredictable and homoscedastic. All changes in prices come from changes in the price of risk. Indeed, stock prices fall contemporaneously with cash flows during recessions, similar as in the data. However, stock price volatility increases only by a factor of 1.1 during recessions. This is not even close to 5.0.[7] The long-run risk model (Bansal and Yaron, 2004) does feature time-varying cash flow volatility. However, because times of high volatility are not systematically related to times of low cash flows, recession variance ratios are 1.0. The time-varying rare disaster model also struggles. Because there is no feedback mechanism such that a recession gives rise to a larger disaster probability, the price-dividend ratio also does not collapse during recessions and recession variance ratios are flat.

I provide several robustness checks. First, I show that my results do not just hold on average but also recession by recession for the period 1950-2016. I find that price-dividend ratios are in all cases lower after the beginning of a recession compared to before. Put differently, the stock market did not anticipate one single recession. Second, I document that my results are robust to an extended dataset of annual data that spans the period from 1871-2016. Stock prices also do not anticipate recessions in this sample. These results are, of course, less granular, but qualitatively and quantitatively similar to the baseline results that are based on quarterly

[7] The habit model is non-linear. But even looking at the 20% largest recessions does only increase the stock price variance by a factor of 1.5.

data. Taken together, these findings suggest that my results are not driven by a few extreme observations but rather reflect a common feature of recessions.

To sum up, I argue that recessions can be used for a fresh empirical assessment of consumption-based asset pricing theories. Many models do well in generating a large equity premium, stock market volatility or predictability of stock returns by the price-dividend ratio (see Table I for a horse race). They achieve these goals by relying on quite different mechanisms (e.g., Cochrane, 2017). But which of these mechanisms is important in which kind of situation? The literature is surprisingly silent on this question. Studying the behaviour of the stock market around recessions allows me to compare different channels proposed by competing theories in a parsimonious and tractable way. My results provide evidence that large changes in the price of risk play a prominent role for explaining the stock market around recessions. This channel requires some kind of amplifier mechanism such that movements in expected returns happen during recessions when we see large drops in economic activity.

Outline: The next section describes how my paper adds to the literature. Section III shows how stock prices and cash flows "should" respond to recessions according to well-known asset pricing theories. Section IV provides the empirical counterparts; followed by some back of the envelope calculations on how discount rates behave during recessions in Section V. Further results are provided in Section VI, followed by the conclusion.

II. Contribution to the Literature

Hall (2017) summarizes the current state of asset pricing as follows (p.327):

"This literature has reached the inescapable conclusion that the large movements in the value of the stock market arise mainly from changes in discount rates and only secondarily from changes in the dividend or profit flow capitalized in the stock market. The field is far from agreement on the reasons for the volatility of discount rates."

Also it is well understood that discount rates rise substantially during recessions, there is little empirical research that tries to discriminate between alternative theoretical explanations. The behaviour of the stock market around recessions that I document provides an opportunity to confront existing asset pricing theories with the data and helps me to shed novel light on the sources of large discount rate volatility.

My analysis complements earlier work by Muir (2017). He compares the decline of asset prices and the change in fundamentals around "normal" recessions and financial crises. He finds that asset price declines are larger during financial crises compared to normal recessions, even though fundamentals move similar. 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 actually drive stock prices around normal recessions. My paper fills 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. I add to previous findings by deriving quantitative predictions of leading asset pricing models on the timing of prices, cash flows and the rise in volatility. I then demonstrate that these model predictions help me to addresses a potential concern regarding the analysis of Muir (2017). He assumes that changes in dividend yields are to 100% due to revisions in future returns and do not capture revisions in cash flows. Even though traditional dividend yield regressions point towards this interpretation (e.g., Campbell and Shiller, 1988), it is well known that changes in the estimation method (e.g., Binsbergen and Koijen, 2010), variable definition (e.g., Larrain and Yogo, 2008; Jank, 2015), or the considered sample period (pre-1950: e.g., Chen, 2009; Golez and Koudijs, 2018; internationally: e.g., Rangvid, Schmeling, and Schrimpf, 2014; out-of-sample: Goyal and Welch, 2008) can lead to evidence that suggests that dividend yields also predict future dividends to some degree. I add to Muir (2017) by showing that it is possible to avoid the potential issues of dividend yield regressions by looking at changes of the dividend yield and its components before and after the beginning of a recession.

However, in order to study the timing of cash flows and stock prices with sufficient precision, one has to move to higher frequent data than annual, which brings me to the importance of the time aggregation bias. I show that end of period stock prices start dropping four quarters before time-aggregated stock prices, earnings, dividends, consumption and the NBER declares the beginning of a recession. This finding also explains the seemingly puzzling fact that a large literature finds that monthly/quarterly stock returns are a good proxy for *future* business conditions (e.g., Fama and French, 1989; Fama, 1990; Vassalou, 2003; Campbell and Diebold, 2009; Backus, Routledge, and Zin, 2010; Maio and Cooper, 2018), while at the same time another large literature finds that dividend yields do not predict future dividends (e.g., Campbell and Shiller, 1988; Cochrane, 2008; Muir, 2017). When looking at frequencies higher than one year, the time aggregation bias in cash flow measures will make market-based data appear to be leading. When looking at lower frequencies, as is typically done in the case of dividend yield predictability regressions, dividends might appear to be (close to) unforecastable. My results suggest that stock prices help to "nowcast" cash flows (and consumption, or recessions) within the horizon of about one year, simply because they are based on market prices and are time aggregation bias free. Thus, my results are also in line with the common newspaper narrative that the stock market "anticipates" recessions (e.g., Samuelson, 1966) — it simply does so because of the more timely measurement based on market data.[8]

[8]Nobel prize winner Paul A. Samuelson once commented in his *Newsweek* column on the conventional wisdom that the stock market predicts recessions (Samuelson, 1966). He famously joked that the stock market

My paper adds to the ongoing discussion whether there is a predictable component in consumption or not. The methods that are usually used to answer this question have a large margin of error (e.g., variance ratio tests or dividend yield regressions). As a result, with the available methods and data, it is not possible to definitely reject that consumption has a long-run predictable component that is sizeable in economic terms (see, for example, the discussions in Cochrane (1994); Hansen, Heaton, and Li, 2008; Marakani, 2009; Constantinides and Gosh, 2011; Beeler and Campbell, 2012; Bansal, Kiku, and Yaron, 2012; Dew-Becker, 2017). My event study-based tests show that dividends, earnings and consumption are not anticipated by stock prices around recessions. This holds on average but also recession by recession. After accounting for the time aggregation bias, the relationship between stock prices and cash flows around recessions is better described contemporaneously. I argue that these results constitute direct evidence that the business cycle is not well described by changes in expected future cash flows. However, my results are, by construction, silent on everything outside of the event window, e.g., whether there might be shifts in expected cash flows that drive stock prices at a considerably lower (or considerably higher) frequency.

My paper also adds to the return predictability literature. While 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), there are surprisingly few papers that focus on expected returns during specific kinds of "bad times" in the spirit of an event study. 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, Strauss, and Zhou (2010), Henkel, Martin, and Nardari (2011), Dangl and Halling (2012), find that returns are more predictable during recessions. These papers document interesting empirical facts. I add to these paper by studying how these findings connect to theoretical predictions of different models.

A rather recent literature finds that stock returns are predictable at surprisingly short ho-

has predicted nine of the past five recessions.

rizons. Bollerslev, Tauchen, and Zhou (2009) show that returns are predictable by the variance risk premium at horizons of less than one year. Bekaert, Engstrom, and Xu (2017) 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. Martin (2017) uses option prices to derive the lower bound on expected returns at a daily frequency. He argues that expected returns must be highly volatile, particularly during bad times. I come to a similar conclusion by studying stock prices and cash flows around recessions directly and without relying on derivative markets.

Finally, some recent papers propose theoretical mechanisms where short-run and long-run consumption risks are combined (e.g., Branger, Kraft, and Meinerding, 2016), agents learn (or disagree) about the business cycle (e.g., Andrei, Hasler, and Jeanneret, 2019; Cujean and Hasler, 2017), consider that investors might be disappointment-averse (Schreindorfer, 2019), or underline the importance of institutional frictions during recessions (e.g., Adrian and Shin, 2014). All these papers suggest some kind of amplifier mechanism which makes risky assets more "vulnerable" during "bad times". My empirical results are supportive for such amplifier mechanisms. I find that traditional models leave a surprisingly large part of stock price movements around recessions unexplained. A common shortcoming of traditional models is that the state-variable which does most of the work is not (strongly) linked to times of low economic activity.

III. Recessions in Asset Pricing Models

How "should" stock prices respond to recessions? To answer this question, I simulate 10,000 years of artificial monthly data for four consumption-based asset pricing models i), the canonical consumption-based asset pricing model (Breeden, 1979) ii), the long-run risk model as in Bansal and Yaron (2004) iii), the habit model as in Campbell and Cochrane (1999), and iv) the time-varying rare disaster risk model as in Wachter (2013). I then convert monthly data to the sum of 12 month dividends and 12 month consumption to get "time-aggregated" versions of these variables that are comparable to their frequently used empirical counterparts. Next, I compute overlapping *quarterly* log changes of these variables, as this is the highest frequency at which dividends and consumption are reported in my empirical dataset. NBER recessions are defined as:[9]

> *"A recession is a period between a peak and a trough, and an expansion is a period between a trough and a peak. During a recession, a significant decline in economic activity spreads across the economy and can last from a few months to more than a year."*

To mimick the NBER business cycle dating committee, I search for large local peaks in economic activity of the simulated data.[10] I proxy economic activity by consumption in these models. More specifically, I search for all peaks in the simulated data and then mark the 25% largest peaks as the observation just before the beginning of a recession, to get "significant declines".[11] As a result, between 2%-3% of my simulated observations are "beginning of recessions", which is comparable to the empirical data (3%). In the simulated data, similar to the empirical data, the duration of a "recession" varies. I then use local linear projections (Jorda, 2005) to estimate the cumulative average log change of the de-meaned price-dividend ratio, prices, dividends, and consumption in recession event time. The event window is set from 12 quarters before the beginning of recessions and ends 12 quarters after the beginning of recessions.

The Classic Consumption-based Model: In the classic model, future consumption is unpredictable and expected returns are constant. Realized returns are basically the realized dividends the investor receives in each period. As is common in the literature, dividends are a leveraged consumption claim and thus dividends and consumption are correlated — which

[9] `http://www.nber.org/cycles/recessions.html`, 02-14-2018.

[10] Importantly, I do not want to predict recessions in the simulated or in the empirical data later on. Instead, following the argument in Muir (2017), I want to study how stock prices behave given a large drop in economic activity.

[11] I use the Matlab built-in function findpeaks.m to detect local maxima in the level of consumption.

determines the equity premium in this model. I set consumption volatility to 2.5% and assume a high correlation between consumption growth and dividend growth.[12]

Figure 1 summarizes the simulation results for the simple model. The figure shows the cumulative change of the price-dividend ratio, prices, dividends and consumption around recessions. 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.[13] In this model, the price P_t is:

$$P_t = \frac{D_t \times (1 + E(\triangle D_{t+1}))}{E(R_{t+1}) - E(\triangle D_{t+1})}, \tag{1}$$

where D_t is the dividend, $E(R_{t+1})$ is the expected return and $E(\triangle D_t)$ is expected dividend growth. Because expected returns and expected dividend growth are constant, log prices move 1:1 with changes in log dividends, $p_t = d_t + constant$.[14] Thus, the (log) price-dividend ratio is constant. This, of course, only holds if prices and dividends are both measured exactly at the same point in time. The figure shows that if prices are measured at the end of a period (a quarter) and dividends are measured time-aggregated (the sum over the last year), the log price-dividend ratio makes wild swings even if the true price-dividend ratio is constant. This bias simply results from a timing mismatch; yet it is the commonly used measurement in empirical analysis. The figure also illustrates that using time-aggregated prices in the price-dividend ratio leads to a constant ratio and resolves the issue.

[12] The picked values for the consumption volatility and correlation are in line with recent empirical evidence provided by Savov (2011) and Kroencke (2017). However, the parameters are not essential for this paper, since the classical model serves only for illustrative purposes. The fact that the empirical log price-dividend ratio drops around recessions leads to rejection of the classic model for *any* parametrization. Further simulation details are provided in the Internet Appendix.

[13] The y-axis will also be the same in all figures such that one can easily compare results between all simulated models and the empirical data later on.

[14] Notice that p_t and d_t perfectly overlap in the figure and prices might be difficult to see.

Figure 1

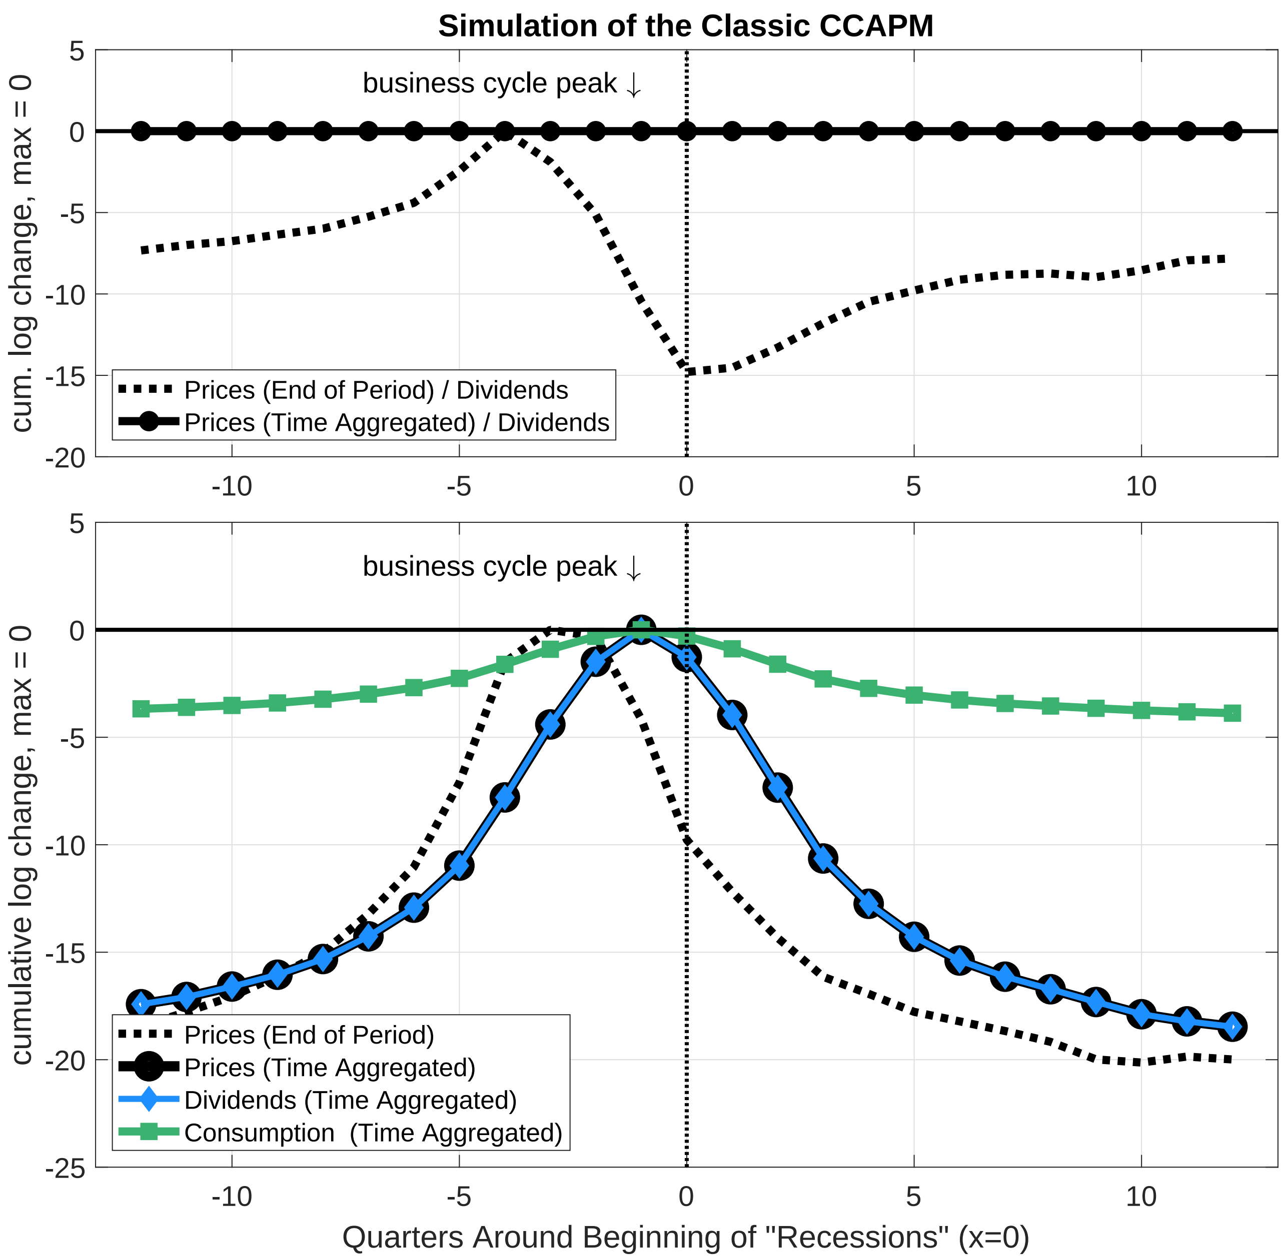

Simulation of the Classic CCAPM
business cycle peak ↓
cum. log change, max = 0
Prices (End of Period) / Dividends
Prices (Time Aggregated) / Dividends
5
0
-5
-10
-15
-20
-10
-5
0
5
10
business cycle peak ↓
cumulative log change, max = 0
Prices (End of Period)
Prices (Time Aggregated)
Dividends (Time Aggregated)
Consumption (Time Aggregated)
5
0
-5
-10
-15
-20
-25
-10
-5
0
5
10
Quarters Around Beginning of "Recessions" (x=0)

The Long-Run Risk Model: In the long-run risk model by Bansal and Yaron (2004), cash flows have three ingredients. Traditional one period consumption shocks ("short-run risk") as in the simple model before, and in addition, 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. I use the same model parameters as in Bansal and Yaron (2004), to make my results comparable to the literature.[15]

Figure 2 summarizes results for the long-run risk model. I find that the time-aggregated price-dividend ratio starts dropping about one year ahead of recessions. Intuitively, if one mimicks the NBER business cycle dating committee and searches for significant declines in economic activity (consumption), one identifies periods when long-run growth and short run growth happen to be low at the same time. The forward-looking investor, in turn, recognizes that future dividends will be lower going forward and as a consequence stock prices decline several quarters before the recession actually starts. The picture also shows that the price-dividend ratio does not move much with the beginning of recessions. It is also instructive to see that the end of period price-dividend ratio looks similar to its counterpart in the simple model with constant expected cash flows (Figure 1). The mismatch in the timing of end of period prices and time-aggregated dividends obscures the differences between both models.

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 modelled as an "additive" channel. As a result, there is no systematic relationship between volatility and recessions (a fall in consumption). I will provide the numbers and discuss this aspect of the model in more detail later in the paper.

[15] Further simulation details are provided in the Internet Appendix. The re-calibration of the long-run risk model provided by Bansal, Kiku, and Yaron (2012) model leads to similar results; see Table I. The timing of the drop of the price-dividend ratio is very similar. Because the role of long-run risks is toned down, the price-dividend ratio drops somewhat less.

Figure 2

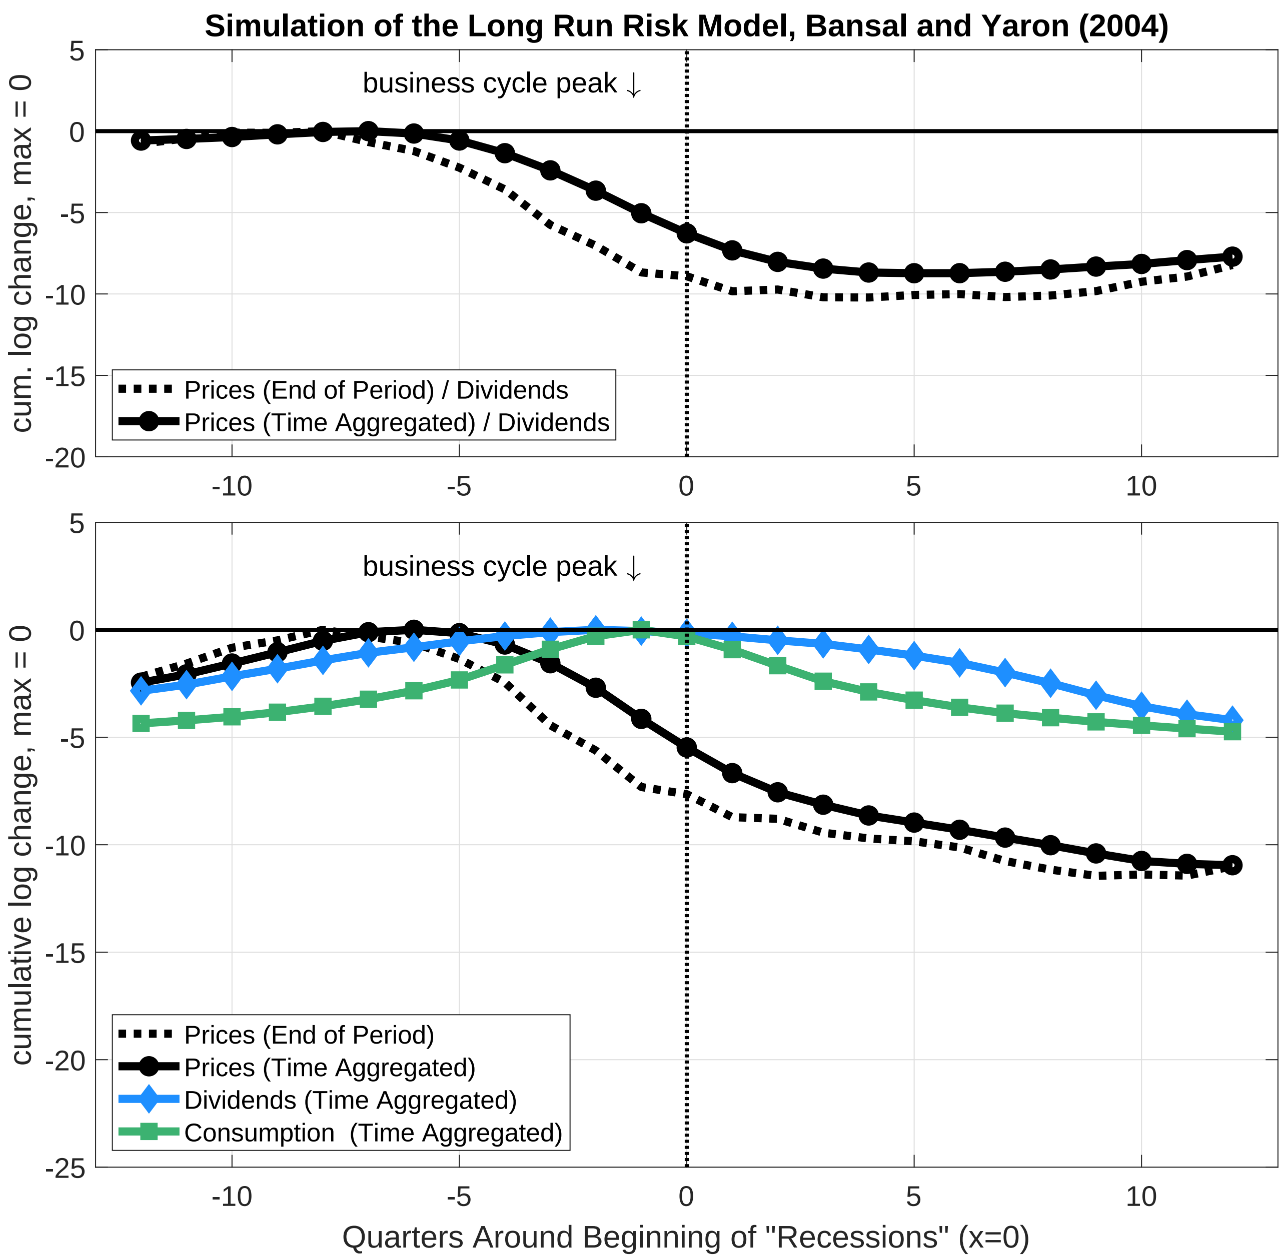
Simulation of the Long Run Risk Model, Bansal and Yaron (2004)
business cycle peak ↓
cum. log change, max = 0
Prices (End of Period) / Dividends
Prices (Time Aggregated) / Dividends
business cycle peak ↓
cumulative log change, max = 0
Prices (End of Period)
Prices (Time Aggregated)
Dividends (Time Aggregated)
Consumption (Time Aggregated)
Quarters Around Beginning of "Recessions" (x=0)

The Habit Model: Next, I simulate the habit model by Campbell and Cochrane (1999). As in the simple model, consumption is an unpredictable standard i.i.d. process. However, the price of risk and, thus, 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 (consumption). Again, I use the same model parameters as in Campbell and Cochrane (1999), to make my results comparable to the literature.[16]

Figure 3 summarizes the simulation results for the habit model. It is easy to see that the price-dividend ratio drops contemporaneously with the beginning of recessions (and dividends and consumption). 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. Getting the timing "wrong" makes dividends and consumption predictable. Even in the habit model.[17]

[16] Further simulation details are provided in the Internet Appendix. Notice that Campbell and Cochrane (1999) and Bansal and Yaron (2004) use quite different levels of consumption volatility. There is, unfortunately, no "benchmark" calibration available that sets both models on equal grounds, which implies that I can only compare the models qualitatively but not quantitatively.

[17] Of course, if one is interested in predicting recessions in real time (and not in the economic mechanisms that push stock prices down), using end of period stock prices would be the "right" way to go.

Figure 3

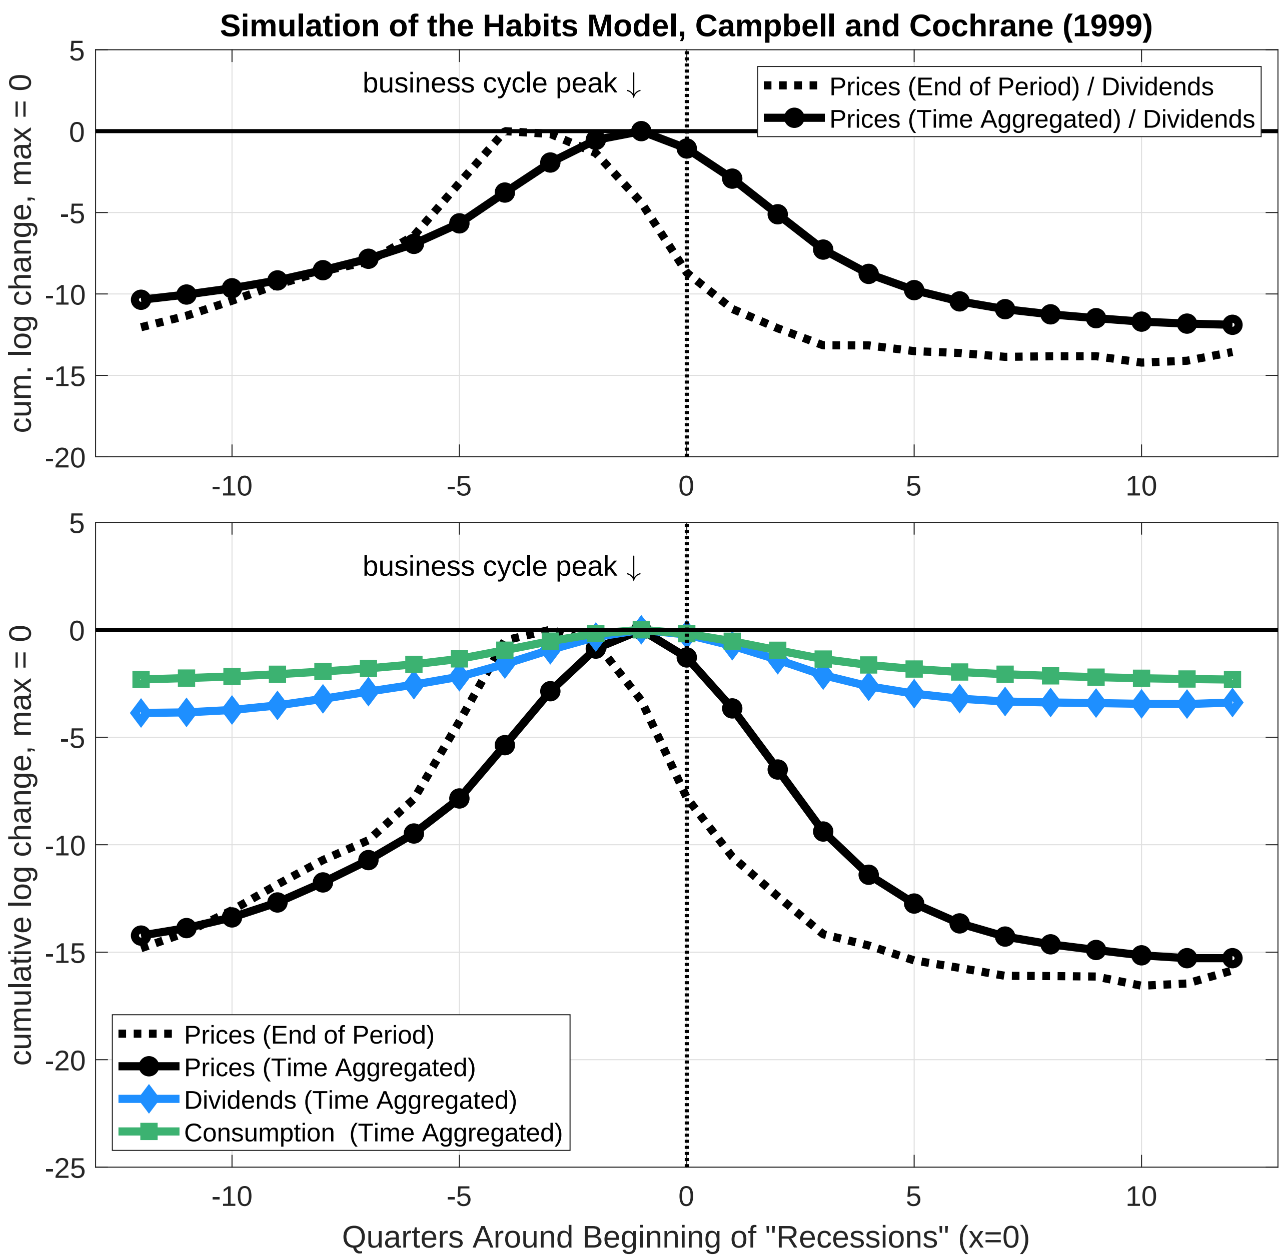
Simulation of the Habits Model, Campbell and Cochrane (1999)
business cycle peak ↓
Prices (End of Period) / Dividends
Prices (Time Aggregated) / Dividends
cum. log change, max = 0
5
0
-5
-10
-15
-20
-10
-5
0
5
10
business cycle peak ↓
cumulative log change, max = 0
Prices (End of Period)
Prices (Time Aggregated)
Dividends (Time Aggregated)
Consumption (Time Aggregated)
-25
Quarters Around Beginning of "Recessions" (x=0)

The Rare Disasters Model: Finally, I simulate the rare disasters model by Wachter (2013). Similar to the earlier models, consumption has an unpredictable standard i.i.d. component. 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. Again, I use the same model parameters as in Wachter (2013), to make my results comparable to the literature.[18]

Figure 4 summarizes the simulation results. 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. This convention also makes clear that the model features two types of recessions. "Normal" recessions coming from bad runs of the i.i.d. component of consumption (as in the previous models) and "large" recessions stemming from the disaster component. The black line shows that 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).[19] The red line is based on all simulated data and contains "normal" but also "large" recessions stemming from rare disasters. Because the "large" recessions are predictable, prices decline on average before recessions. This is similar to the long-run risk model, and is in line with the discussion in Wachter (2013, pp 1015)

[18] Further simulation details are provided in the Internet Appendix.

[19] Actually, in Figure 4 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 is Wachter, 2013), which means that I systematically remove observations with high disaster probability around the event t=0. However, this effect is rather small in magnitude.

Figure 4

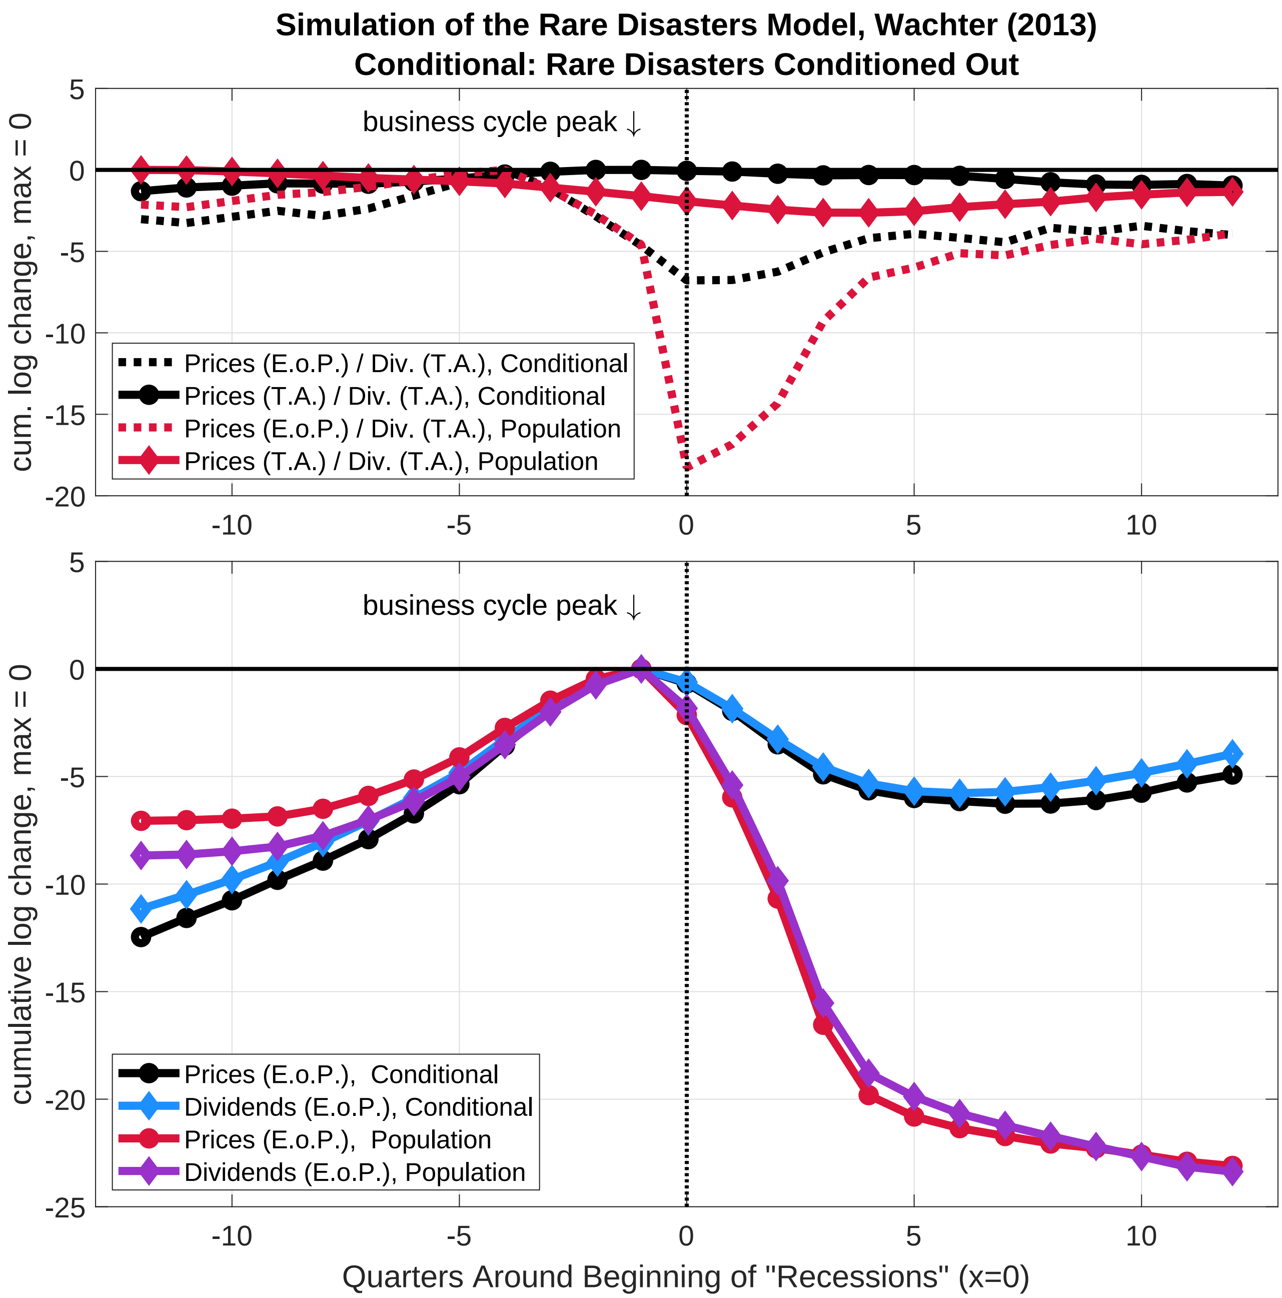

Simulation of the Rare Disasters Model, Wachter (2013)
Conditional: Rare Disasters Conditioned Out
business cycle peak ↓
cum. log change, max = 0
Prices (E.o.P.) / Div. (T.A.), Conditional
Prices (T.A.) / Div. (T.A.), Conditional
Prices (E.o.P.) / Div. (T.A.), Population
Prices (T.A.) / Div. (T.A.), Population
business cycle peak ↓
cumulative log change, max = 0
Prices (E.o.P.), Conditional
Dividends (E.o.P.), Conditional
Prices (E.o.P.), Population
Dividends (E.o.P.), Population
Quarters Around Beginning of "Recessions" (x=0)

Summary of the Simulation Experiment: Table I summarizes the results. Panel A confirms that all advanced models considered make progress in the sense that they can generate a large equity premium, high stock market volatility, and a price-dividend ratio that negatively predicts future returns — as what can be observed in the empirical data. However, as shown in Panel B, the models make rather different predictions on how the stock market should behave around recessions, e.g., *when* and *how* much the price-dividend ratio should drop around periods of low economic activity. The reason is that the models rely on fundamentally different channels that "drive" the market (Cochrane, 2017).

If recessions are periods of lower than usual growth such that there is a predictable component in cash flows, as in the long run risk model, the price-dividend ratio should start falling several quarters ahead of the recession (Figure 2 and the pre-drawdown in Table I). A similar logic applies to the recessions stemming from rare disasters (Figure 4). Finally, a drop in the price-dividend ratio contemporaneously with fundamentals (dividends and consumption) is in line with a change in discount rates as in the habit model (Figure 3 and the post-drawdown in Table I). Panel B of Table I reveals that advanced asset pricing model make fundamentally different predictions on how the stock market should behave around recessions. The timing of the drop in prices compared to cash flows around recessions is informative when it comes to disentangling different potential economic channels.

Table I Comparison of Asset Pricing Theories

This table compares stock market metrics as predicted by consumption-based asset pricing models with empirical counterparts. The traditional metrics are the annual equity premium ($E[R - rf]$, %), the annual stock market volatility ($\sigma[R - rf]$, %), and the slope (b_{pd}) and measure of fit (R^2, %) of a return predictive regression using the log price-dividend ratio. The lower panel shows the properties of the stock market around recessions. The drawdown, %, is the drop in the log price-dividend ratio (time aggregated) four quarters before the business cycle peaks (pre-drawdown), or four quarters afterwards (post-drawdown). The variance ratio is the stock price variance (end of period) measured four quarters after the business cycle peaks dividend by the stock price variance measured from the quarters four quarters before the business cycle peaks. The models are the classic C-CAPM (Breeden, 1979, parametrized to the results provided in Savov, 2011), the long-run risk model (the original version as in Bansal and Yaron, 2004, "BY", as well as the re-calibration as in Bansal, Kiku, and Yaron, 2012, "BKY"), the habit model (Campbell and Cochrane, 1999) and the time-varying rare disaster risk model (Wachter, 2013). For the rare disaster model, I provide results based on all simulated observations ("population") and a sample that conditions out realized rare disasters ("conditional"). All models are simulated 10,000 years; further details are provided in the main text and the Internet Appendix. The empirical data span the period 1950 to 2016.

	Prediction						
	CCAPM S2011	**LRR BY2004**	**LRR BKY2012**	**Habits CC1999**	**Disasters W2013 population**	**Disasters W2013 cond.**	**Data**
	Panel A: Traditional Metrics						
$E[R - rf]$, % p.a.	4.58	5.25	7.82	5.34	7.89		6.90
$\sigma[R - rf]$, % p.a.	19.48	16.34	23.25	15.84	20.25		16.78
	$r_{t+1} - r_{f,t+1} = a + b_{pd} \times (p_t - d_t) + \epsilon_{t+1}$						
b_{pd}	0.01	-0.01	-0.07	-0.15	-0.12		-0.09
R^2, %	0.00	0.00	0.76	2.41	4.77		4.30
	Panel B: Stock Market Around Recessions						
log price-dividend ratio (T.A:):							
pre-drawdown, % (1y)	0.00	-5.04	-2.84	7.84	-1.12	0.85	5.01
post-drawdown, % (1y)	0.00	-3.64	-1.92	-8.77	-1.02	-0.31	-10.80
log stock price changes (E.o.P):							
variance ratio: post/pre	1.00	1.04	1.07	1.12	3.27	1.01	3.96

IV. The Stock Market Around Recessions

A. Data & Method

My baseline results focus on quarterly sampled data for the period from 1950 to 2016. Time-aggregated stock prices, earnings, and dividends are trailing sums over the past 12 months. Earnings and dividends are highly seasonal and for this reason it is the common standard in the literature to look at trailing 12 month sums of these variables. I compute overlapping quarterly log changes to track the variables of interest around recessions. The quarterly sampling of the data allows me to pin down the timing of the variables as precisely as possible. The baseline sample spans ten NBER business cycle peaks.[20] Results for an annual dataset that spans the period from 1871 to 2018 and 29 recessions is provided as a robust test later in the paper. For the period from (1871) 1950 to 1974, I use the data provided by Robert J. Shiller on his website.[21] For the period from 1975 to 2016, I rely on prices, dividends, and earnings for the S&P 500 index as provided by the Thomson Reuters Datastream.[22] All stock market data are adjusted for inflation using the CPI.[23]

I use local linear projections (Jorda, 2005) to estimate the de-meaned log change of the variables of interest:

$$\triangle x_{t+h} = a_h + b_h \times D_{Beginning\,of\,Recession,t} + \zeta_{t+h},$$

where $\triangle x_{t+h}$ is the log change of price-dividend ratio, prices, dividend, earnings or consumption, $D_{Beginning\,of\,Recession,t}$ is a dummy that is one in the quarter the NBER declares the business cycle peak (and a recession begins), and h is the event window and ranges from

[20]Peak date (duration in months): Q2-1953(10), Q3-1957(8), Q2-1960(10), Q4-1969(11), Q4-1973(16), Q1-1980(6), Q3-1981(16), Q3-1990(8), Q1-2001(8), Q4-2007(18). My results are not sensitive with regard to the event Q3-1981, which closely follows the previous recession ("double dip"), or any other single observation as I show later in the paper.

[21]Shiller (2000).

[22]Series: S&PCOMP; datatypes: MV, DSDY, DSPE. A detailed comparison between the two datasets can be found in the Internet Appendix.

[23]Available from FRED (St. Louis Fed); series CPIAUCSL.

$h = -12$ to $h = +12$ quarters. The coefficient b is my estimate of the (de-meaned) log change. Standard errors are based on the full coefficient covariance matrix such that I account for correlation between bs. I then sum up bs to get the impulse response of my variables around recessions.

B. The Timing of Stock Prices and Cash Flows

Figure 5 shows the cumulative drop of prices relative to cash flows around recessions. The upper figure shows the log price-dividend ratio and the figure below 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. I find that the time-aggregated price-dividend ratio (upper figure) starts dropping with the beginning of recessions. The 90% confidence bands indicate that the drop in the price-dividend ratio is precisely measured. The lower figure reveals that both components, prices and dividends, fall very much contemporaneously. The decline in dividends cumulates to about -10%, and to about -20% for prices. Together, this results in a cumulative decline in the price-dividend ratio of -10% ($\triangle pd_t = \triangle p_t - \triangle d_t$).

Dividends might be smoothed due to dividend policy (e.g., Lintner, 1956, Fama and Babiak, 1968, Brav, Graham, Harvey, and Michaely, 2005, Leary and Michaely, 2011). For that reason, the figure also provides the cumulative response of earnings. I find that earnings fall very much together with prices and dividends. End of period stock prices, which are often used in empirical analysis, indeed lead recessions by several quarters. However, the comparison to time-aggregated stock prices reveals that this lead is spurious and is simply driven by the time aggregation bias — and not by an economically rooted mechanism.

Cash flows might be unpredictable before recessions but predictable afterwards. Recessions might *lead* investors to revisions in future expected consumption growth. To be clear, there no such mechanism in traditional versions of the long-run risk model (e.g., Bansal and Yaron, 2004, Bansal, Kiku, and Yaron, 2012).[24] The median duration of the covered recessions is 10 months (min 6 months; max 18months), i.e. "afterwards" starts after about one year. The time-aggregated data reach their trough after 6 quarters, the end of period prices earlier; in line with the rather short duration of recessions. Looking at the quarters after the end of recessions (t=+6 to t=+12), it is apparent that cash flows grow with average speed (or slightly above), as is the case before a recession starts (t=-12 to t=-2).[25] This suggests that the drop in stock prices also does not reflect expected cash flows lower than usual *after* a recession has occurred.

Table II shows cumulative log changes in the de-meaned variables of interest and provides detailed statistical inference. Coefficients at the left hand side are cumulated for all variables from quarter t=-1 up to quarter t=-12 before the beginning of recessions (t=0). The table reassures that the time-aggregated price-dividend ratio, or just time-aggregated prices, do not significantly decline ahead of recessions. Finally, the right hand side of the table reports the cumulative change starting with the beginning of recessions (t=0). The drop in all variables is highly significant for the first couple of quarters during recessions. As the start of the following expansion varies between recessions (between 0.5 - 1.5 years), longer horizon estimates get more and more imprecise, however.

In summary, these results speak against the idea that cash flows have a predictable component such that stock prices anticipate recessions.

[24] Similarly, such a link between short-run consumption risk and the rare disaster probability is not present in theWachter (2013)-version of the rare disaster risk model.

[25] All variables are de-meaned, hence, horizontal lines mean that the variable grows with average speed.

Figure 5. The Stock Market Around NBER Recessions, 1950-2016

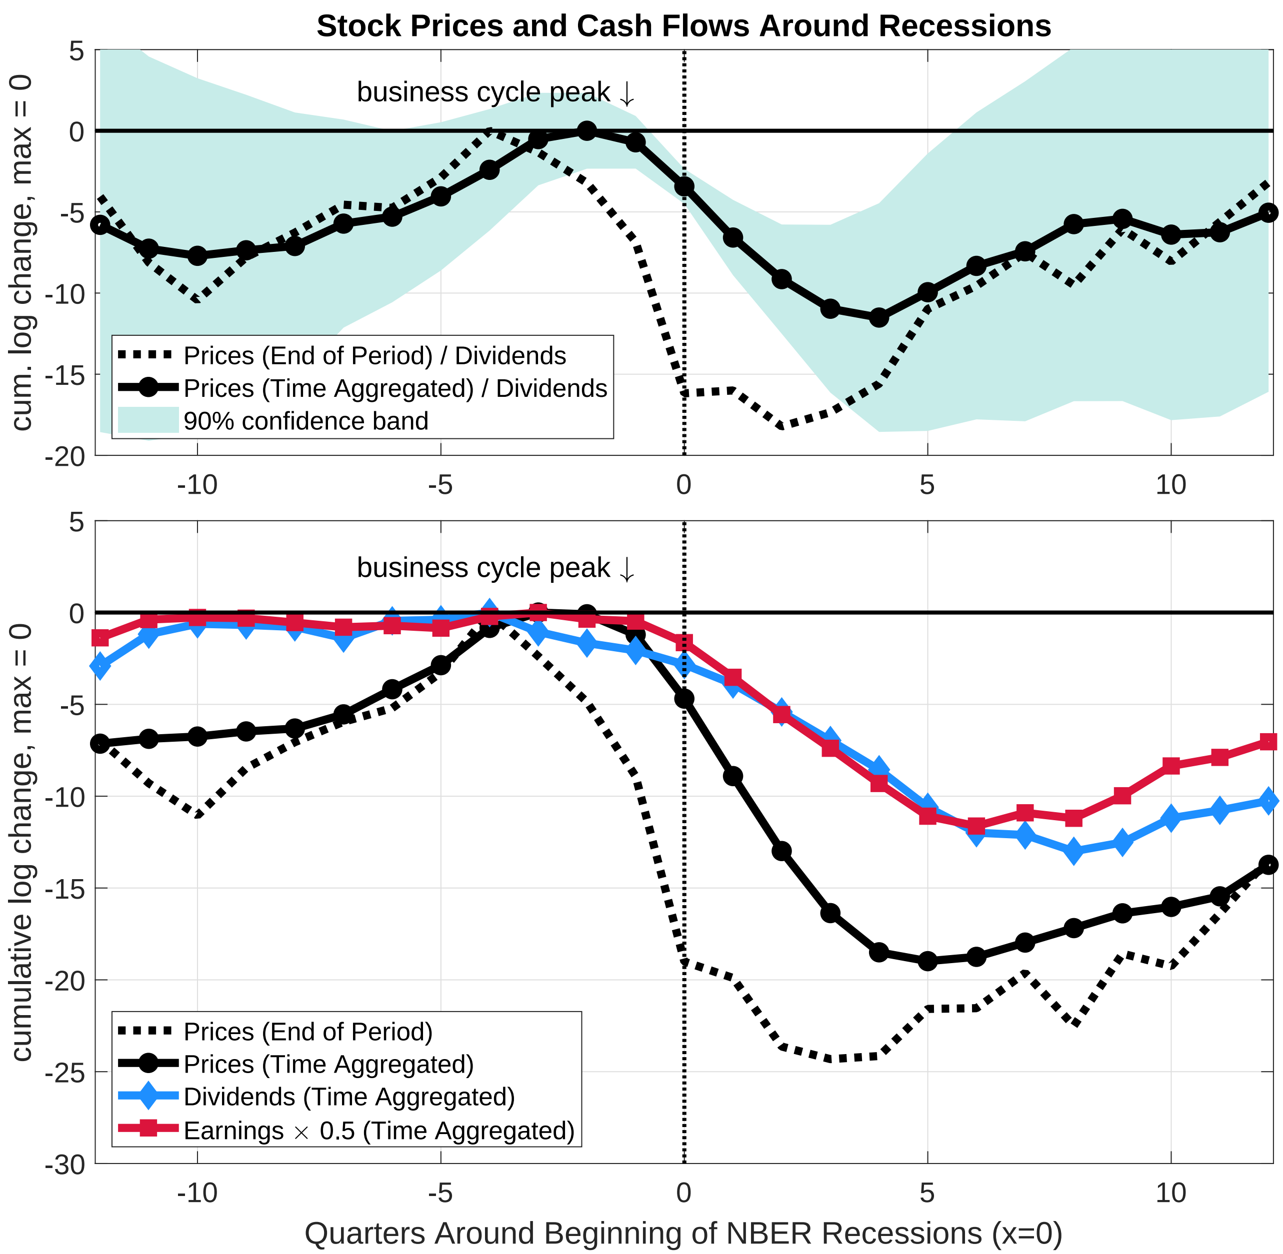

Table II Linear Projections: Stock Market

This table shows cumulative log changes of the price-dividend ratio ($\triangle pd$), prices ($\triangle p$), dividends ($\triangle d$), and earnings ($\triangle e$) around the beginning of NBER recessions from 1950 to 2016. The data are sampled quarterly and the event window ranges from h=-12 to h=+12 quarters around the beginning of a recession (h=0). Estimates are based on local linear projections,

$$\triangle x_{t+h} = a_h + b_h \times D_{Beginning\ of\ Recession,t} + \zeta_{t+h},$$

where $D_{Beginning\ of\ Recession,t}$ is one at the beginning of a recession and zero otherwise. Cumulative changes of the (de-meaned) variables are measured as sums of the coefficients b; t-statistics (t) on cumulated effects are based on the full b coefficient covariance matrix. Dividends and earnings are trailing 12 month sums, i.e., they are time-aggregated. Prices are as 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 earnings and dividends and are time-aggregated data (T.A.).

	←pre-Recession					**↓**	**after Beginning of Recession→**					
Quarter	**-12:-1**	**-8:-1**	**-4:-1**	**-2:-1**	**-1**	**0**	**0:+1**	**0:+2**	**0:+3**	**0:+4**	**0:+8**	**0:+12**
	Price-dividend ratio											
$\triangle pd_t$, E.o.P.	-0.46	0.82	-2.41	-4.02	-5.56	-3.68	-9.32	-9.10	-11.31	-8.72	-2.74	3.52
t	*-0.06*	*0.20*	*-0.79*	*-1.45*	-3.00	*-1.76*	*-4.53*	*-4.11*	*-2.89*	*-1.12*	*-0.39*	*0.53*
$\triangle pd_t$, T.A.	3.98	6.53	4.90	3.30	-0.23	-0.73	-2.72	-5.87	-8.43	-10.78	-5.10	-4.47
t	*0.52*	*1.32*	*1.54*	*1.46*	-0.16	*-0.74*	*-4.18*	*-4.18*	*-4.14*	*-2.52*	*-0.77*	*-0.67*
	Prices, dividends, and earnings											
$\triangle p_t$, E.o.P.	1.79	-0.51	-3.02	-5.68	-6.55	-4.06	-10.09	-10.93	-14.65	-15.21	-13.70	-4.72
t	*0.25*	*-0.11*	*-0.79*	*-1.77*	*-3.34*	*-2.19*	*-4.91*	*-5.16*	*-3.71*	*-1.92*	*-1.59*	*-0.59*
$\triangle p_t$, T.A.	6.23	5.20	4.29	1.64	-1.22	-1.12	-3.49	-7.70	-11.77	-17.27	-16.06	-12.72
t	*0.85*	*1.12*	*1.16*	*0.64*	*-0.88*	*-1.51*	*-5.26*	*-5.94*	*-5.87*	*-3.91*	*-1.96*	*-1.63*
$\triangle d_t$, T.A.	2.44	-0.20	0.75	0.79	-0.94	-0.22	-2.35	-6.13	-10.19	-17.69	-21.53	-13.28
t	*0.50*	*-0.05*	*0.19*	*0.26*	-0.56	*-0.24*	*-2.72*	*-5.07*	*-4.49*	*-4.98*	*-3.19*	*-2.56*
$\triangle e_t$, T.A.	2.25	-1.33	-0.60	-1.66	-0.99	-0.39	-0.77	-1.83	-3.34	-6.50	-10.95	-8.25
t	*0.75*	*-0.48*	*-0.32*	*-0.86*	-0.94	*-0.46*	*-2.01*	*-2.90*	*-3.93*	*-4.99*	*-3.20*	*-1.97*

C. *The Timing of Cash-Flows and Consumption*

In this section, I use consumption data as a further proxy for investor "cash flows" and also to connect my results more deeply to consumption-based asset pricing models. Consumption data are taken from the NIPA, available on the website of the Bureau of Economic Analysis.[26]

Reported consumption is not highly correlated with stock returns which gives rise to the equity premium puzzle (Grossman and Shiller, 1980, 1981; Hansen and Singleton, 1982, 1983; Mehra and Prescott, 1985). Savov (2011) argues that if there are measurement problems with officially reported consumption, the more simple measure of garbage might be a better approximation of true consumption. Indeed, he finds that garbage growth is highly correlated with stock returns. Kroencke (2017) provides a possible explanation of why garbage is more correlated with stock returns. If consumption is measured with an error, it is optimal for NIPA statisticians to filter reported consumption to smooth out measurement errors. However, filtering reduces the variance of reported consumption (reported consumption drops too little) and will introduce a lag compared to true but unobservable consumption (reported consumption lags market based data, e.g. stock prices).

Kroencke (2017) suggests that the true consumption estimation process might be approximated by a simple Kalman filter model which allows to compute the series of "unfiltered" NIPA consumption. Because the series of "unfiltered" NIPA consumption is likely to be better suited to assess the drop and the timing of consumption during recessions, I report results for NIPA consumption ("as reported") as well as unfiltered NIPA consumption.[27] Furthermore, I provide results for aggregate consumption, as measured by nondurables and services, as well as nondurables excluding services. There is evidence that nondurables are easier to measure than services and are accordingly less plagued by measurement problems (see, Wilcox, 1992; Kroencke, 2017).

[26]NIPA tables 2.3.4 and 2.3.5; lines "nondurable goods" and "services"; real per capita growth weighted by their nominal share. Capita numbers are from NIPA table 1.1.6.

[27]Because garbage is not available at the quarterly frequency, I complement results on reported NIPA consumption with unfiltered NIPA consumption sampled at the quarterly frequency. Further details on the unfilter procedure are provided in the Internet Appendix.

Figure 6 and Table III provide the results. All consumption measures start falling with the beginning of recessions. The figure also shows scaled dividends and earnings to allow for visual comparisons to stock market cash flows.[28] Unfiltered consumption, particularly nondurables, fall contemporaneously with firm cash flows. The figure also suggests that consumption grows with average speed (or slightly above) after the end of a recession (t=+6 to t=+12) as is the case before recessions start (t=-12 to t=-2).

In summary, I find that consumption falls very much together with firm cash flows around recessions, which corroborates the results from the previous section.

Figure 6. Consumption Around NBER Recessions, 1950-2016

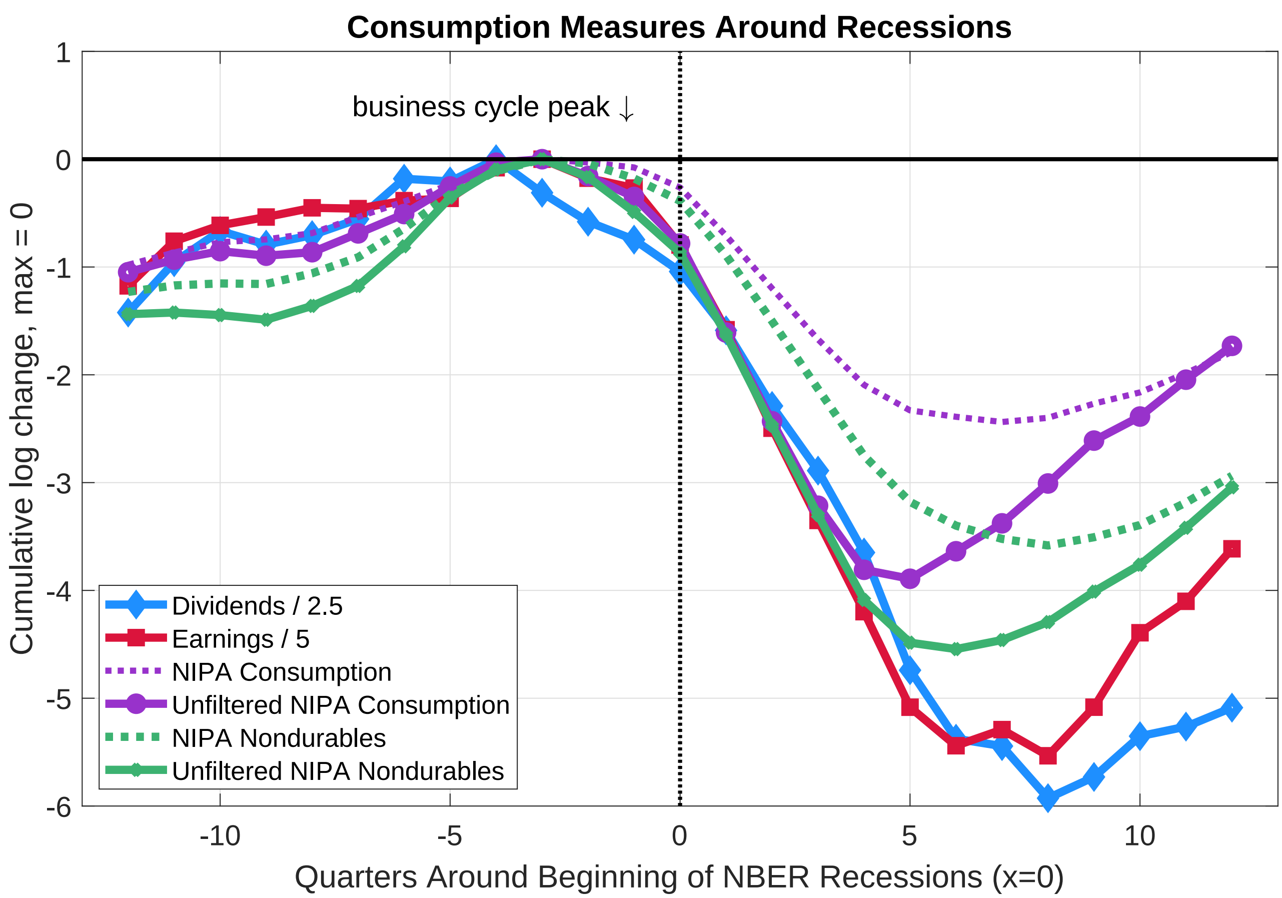

[28]The scaling of earnings and dividends can be interpreted as an adjustment for firm leverage, as is common in the asset pricing literature (e.g., Campbell and Cochrane, 1999; Bansal and Yaron, 2004); the larger scaling factor for earnings is in line with the average payout ratio.

Table III Linear Projections: Consumption Measures

This table shows cumulative log changes of aggregate consumption (services and nondurables) ($\triangle c$) and nondurable consumption ($\triangle ndr$) around the beginning of NBER recessions from 1950 to 2016. All consumption data are real per capita. The event window ranges from h=-12 to h=+12 quarters around the beginning of a recession (h=0). Estimates are based on local linear projections,

$$\triangle x_{t+h} = a_h + b_h \times D_{Beginning\,of\,Recession,t} + \zeta_{t+h},$$

where $D_{Beginning\,of\,Recession,t}$ is one at the beginning of a recession and zero otherwise. Cumulative changes of the (de-meaned) variables are measured as sums of the coefficients b; t-statistics (t) on cumulated effects are based on the full b coefficient covariance matrix. Consumption data are trailing 4 quarter sums, i.e. they are time-aggregated. Consumption data are as "reported" in the NIPA tables, or "unfiltered" as in Kroencke (2017). Details on the unfilter procedure are provided in the Internet Appendix.

	←pre-Recession					↓	after Beginning of Recession→					
Quarter	**-12:-1**	**-8:-1**	**-4:-1**	**-2:-1**	**-1**	**0**	**0:+1**	**0:+2**	**0:+3**	**0:+4**	**0:+8**	**0:+12**
	Aggregate consumption and nondurable consumption											
$\triangle c_t$, reported, T.A.	1.09	0.78	0.62	0.34	0.04	0.00	-0.16	-0.59	-1.06	-1.91	-2.04	-1.38
t	*1.47*	*1.20*	*1.04*	*0.76*	*0.21*	*0.05*	*-1.89*	*-3.35*	*-4.28*	*-4.73*	*-3.01*	*-1.60*
$\triangle c_t$, unfiltered, T.A.	0.95	0.80	0.70	0.26	-0.13	-0.12	-0.39	-1.22	-2.06	-3.43	-2.38	-1.14
t	*0.93*	*0.83*	*0.75*	*0.38*	*-0.40*	*-0.92*	*-2.86*	*-4.19*	*-4.94*	*-5.30*	*-2.50*	*-1.16*
$\triangle ndr_t$, reported, T.A.	1.04	1.00	0.82	0.30	-0.06	-0.09	-0.23	-0.74	-1.35	-2.57	-3.18	-2.48
t	*1.49*	*1.65*	*1.35*	*0.63*	-0.29	*-1.12*	*-2.66*	*-4.02*	*-5.19*	*-5.58*	*-4.19*	*-3.08*
$\triangle ndr_t$, unfiltered, T.A.	0.87	1.11	0.91	0.14	-0.29	-0.26	-0.41	-1.19	-2.06	-3.65	-3.56	-2.32
t	*1.01*	*1.40*	*1.05*	*0.21*	-1.06	*-2.31*	*-2.66*	*-3.52*	*-4.55*	*-5.27*	*-3.58*	*-2.59*

D. The Timing of Growth Expectations

In this section, I analyse changes in the *forward* term structure of real GDP growth forecasts from the Survey of Professional Forecasters (SPF).[29] More specifically, I compare changes 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 for the four quarterly horizons (nowcast, +1Q, +2Q, +3Q). To make results more comparable to the time-aggregated cash flow measures, I take the sum of the current quarter revision and the previous three revisions in expectations and then track the quarterly change of this "time-aggregated" variable.

Thus, this exercise 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 market prices, and ii) that real GDP growth follows similar dynamics as real consumption growth and corporate cash flows, at least around the event window.[30]

The point estimates of the cumulative change in expectations are provided in Figure 7. Details and statistical significance are reported in Table IV. First, focusing on the expected forward growth rate three quarters into the future (black line; the longest horizon available), I find that longer horizon expectations start to drop modestly and together with the business cycle peek. Put differently, professional forecasters do not see recessions coming and they

[29] 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; this time period spans seven U.S. recessions.

[30] 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, Hasler, and Jeanneret, 2019), because the consumption time-series does not start before 1981 and is covered by less survey participants.

update longer horizon expected growth only modestly. On the other hand, the nowcast of GDP growth is heavily revised downwards with the onset of a recession and drops three times as much as longer term forward expectations of GDP growth. These results suggest that cash flow news are mainly of contemporaneous nature around recessions.

Further Results: I provide further results on the forward term structure of expected growth in the Internet Appendix. 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.

Figure 7. Term Structure of Real GDP Growth Forecasts Around Recessions

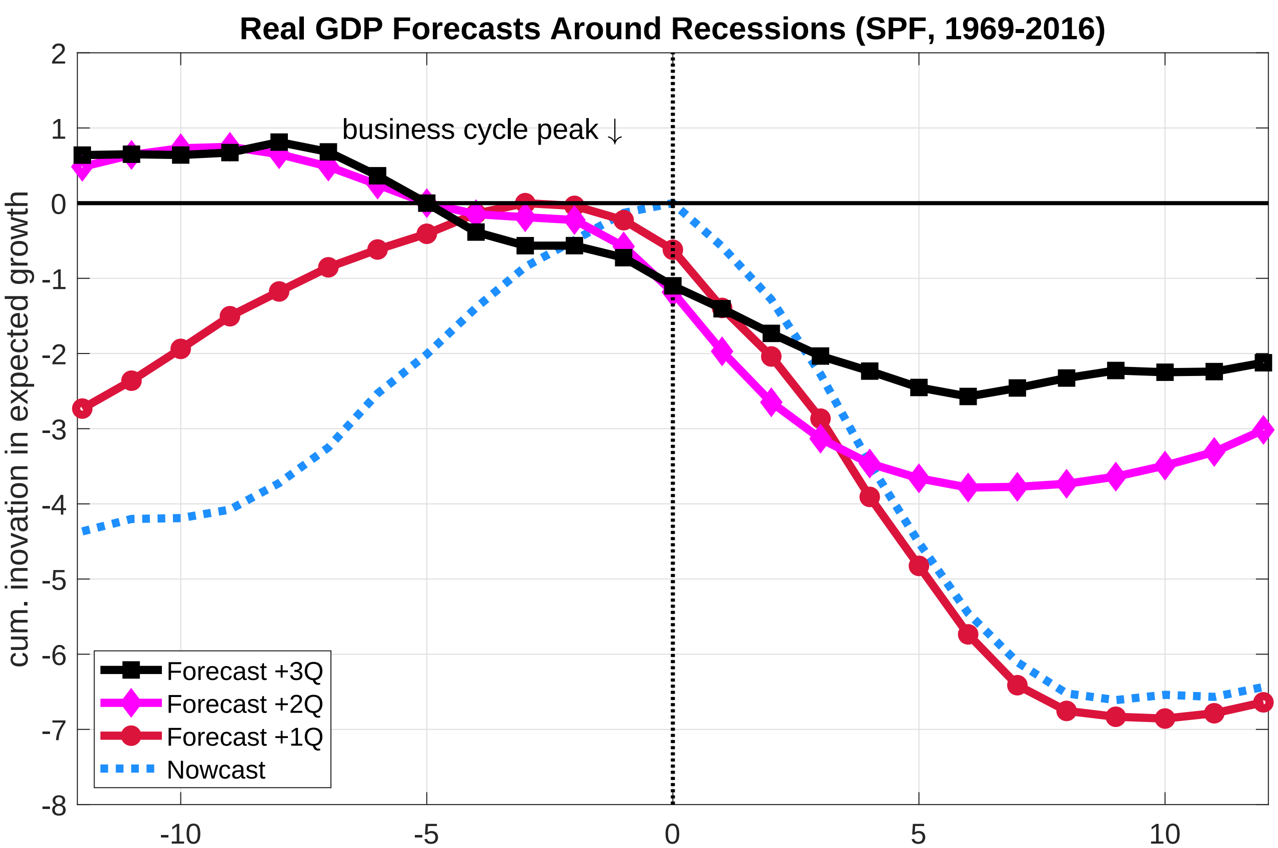

Table IV Linear Projections: Term Structure of Revisions in Real GDP Growth Forecasts

This table provides cumulative changes of revisions in expected real GDP log growth as reported by the Survey of Professional Forecasters from 1969 to 2016. Revisions in expected log growth rates are calculated as the difference of the forecast at time t and the forecast at time t-1 for the log mean forecast of the same quarter k; k refers to the current quarter at time t and up to three quarters into the future. All information are provided by the Survey of Professional Forecasters at the website of the Philadelphia Fed (www.philadelphiafed.org, series "mean RGDP forecasts"). The event window ranges from h=-12 to h=+12 quarters around the beginning of a recession (h=0). Estimates are based on local linear projections as described in the previous table captions (Table II and III).

	←pre-Recession					**↓**	**after Beginning of Recession→**					
Quarter	**-12:-1**	**-8:-1**	**-4:-1**	**-2:-1**	**-1**	**0**	**0:+1**	**0:+2**	**0:+3**	**0:+4**	**0:+8**	**0:+12**
	Changes in expected real GDP growth, % p.a.											
nowcast	4.31	3.95	3.12	1.88	0.72	0.36	0.13	-0.45	-1.15	-3.33	-6.39	-6.31
t	*3.61*	*4.61*	*4.19*	*3.09*	*2.16*	*1.86*	*0.67*	*-1.84*	*-2.65*	*-3.12*	*-4.67*	*-6.08*
forecast +1Q	2.84	1.28	0.63	0.18	-0.23	-0.19	-0.39	-1.17	-1.81	-3.68	-6.53	-6.41
t	*2.02*	*1.07*	*0.62*	*0.27*	*-0.53*	*-0.85*	*-2.16*	*-3.24*	*-3.84*	*-5.40*	*-5.25*	*-4.23*
forecast +2Q	-0.83	-1.32	-1.06	-0.58	-0.39	-0.35	-0.61	-1.40	-2.07	-2.89	-3.16	-2.44
t	*-0.45*	*-0.75*	*-0.75*	*-0.73*	*-0.93*	*-1.30*	*-1.96*	*-2.24*	*-2.61*	*-3.86*	*-3.96*	*-2.79*
forecast +3Q	-1.35	-1.40	-1.41	-0.73	-0.16	-0.16	-0.37	-0.68	-1.01	-1.51	-1.60	-1.40
t	*-0.75*	*-0.95*	*-1.08*	*-0.91*	*-0.71*	*-1.32*	*-2.11*	*-2.03*	*-1.99*	*-2.33*	*-2.28*	*-1.62*

E. Recession Variance Ratios

The previous literature has documented that price and cash flow volatility is higher than usual during recessions (starting with Schwert, 1989; most recently Boguth and Kuehn, 2013; Bansal, Kiku, Shaliastovich, and Yaron, 2014; Tédongap, 2014). My event study setting allows me to supplement these earlier results with non-parametric estimates that are free of any parametric assumption on how volatility can change. Importantly, I compare the increase of volatility between cash flows and stock prices around recessions. I argue that the increase of price volatility compared to cash flow volatility allows to make conclusions on what drives discount rates.[31] Stock price changes are the sum of cash flow innovations and expected return innovations. Intuitively, if price volatility increases more compared to cash flow volatility, it implies that expected return innovations become even more volatile compared to their own cash flows. Such a pattern is, for example, qualitatively in line with models where the price of risk can change (e.g., Campbell and Cochrane, 1999), such that expected returns become more sensitive to shocks.

Figure 8 and Table V provide the conditional volatility of stock prices, dividends, and earnings in recession event time. Estimates are based on the sample standard deviation conditional on the quarter around recessions; multiplied by 2 to provide annualized values. To reduce noise in the figure, I smooth estimates around the two neighbouring observations {-1, 0, +1} using the weights {0.25, 0.50, 0.25}.

Before recessions, end of period (time-aggregated) stock prices have a volatility of around 10% (5%).[32] Stock price volatility doubles during recessions to 20% (11%), which means that the ratio of the pre-recession variance (h=-5 to h=-2) to the post-recession variance (h=+1 to h=+4) quadruples (quintuples). Earnings volatility increases from 6% to 10%, which gives a considerably smaller recession variance ratio of 2.7. The conditional volatility of dividends does not respond quickly to recessions, and on average increases only slightly. I find that nondurable consumption volatility increases during recessions, however, the measured increase also does not result in variance ratios larger than 1.9. Aggregate consumption volatility shows very similar recession variance ratios of 1.5 (reported) and 1.6 (unfiltered). Evidence from the three measures of cash flows (earnings, dividends, and consumption) are somewhat inconclusive on the precise increase of volatility. Given the rise in earnings and consumption volatility, the flat volatility of dividends might be explained by a smoothing effect caused by the dividend payout policy at the firm level (Lintner, 1956, Fama and Babiak, 1968). Taken together, it seems to

[31] Jurado, Ludvigson, and Ng (2015) recently proposed a procedure to estimate macroeconomic uncertainty (almost) model free. However, their estimates rely on a large cross-section of variables and does not identify differences between the increase in the volatility of stock prices and cash flows, which is key for my analysis.

[32] It is not possible to map the exact relationship between the variances of end of period and time-aggregated data for an arbitrary process. Working (1960) shows that the variance of a time-aggregated i.i.d. process is 2/3 of the end of period counterpart; which is roughly in line with the figure.

be safe to say that cash flow volatility increases during recessions, in line with the literature, but to a lesser degree compared to stock price volatility (a result that is new to the literature).

Furthermore, the table provides results for the recession variance ratio from simulations of the long-run risk model, the habit model, and the rare disaster risk model as discussed in Section III. Even though the long-run risk model as in Bansal and Yaron (2004) features time-varying volatility in cash flows, the model is set up such that volatility does not systematically change around recessions when economic activity drops. As a result, all simulated recession variance ratios are around 1.[33] In the habit model by Campbell and Cochrane (1999), cash flows are homoscedastic. However, because the price of risk is linked to large drops in consumption, expected returns can get more sensitive to shocks during recessions. The table shows that the stock price variance indeed increases during recessions in the habit model, however, the effect is substantially smaller (recession variance ratio = 1.1) compared to the increase that can be measured in the empirical data (4). To see whether changes in my definition of simulated "recessions" change results, the table also provides results for the 20% largest recessions. In this case, the recession variance ratio increases to about 1.5.[34] Even looking at the tails of the distribution does not generate enough discount rate volatility. Turning to the rare disaster risk model as in Wachter (2013), I find that when rare disasters are conditioned out ("conditional"), recession variance ratios do not increase. In samples including rare disasters ("population"), I find that the stock price recession variance ratio is indeed very large (above 3), but cash flow recession variance ratios reflect the realized disasters and are dramatically larger.

In summary, I find that stock price volatility increases substantially more than cash flow volatility. For that reason, I conclude that an increase of the price of risk during recessions plays a key role in explaining the data. The habit model by Campbell and Cochrane (1999) qualitatively goes in the right direction, but does not quantitatively get close to the data.

[33] Further discussion on this point is provided in the Internet Appendix. I show that times of large changes in volatility are also not systematically related to consumption.

[34] The Internet Appendix provides figures for "large" recessions in the habit model that allow for further comparisons.

Figure 8. Stock Market Volatility Around NBER Recessions, 1950-2016

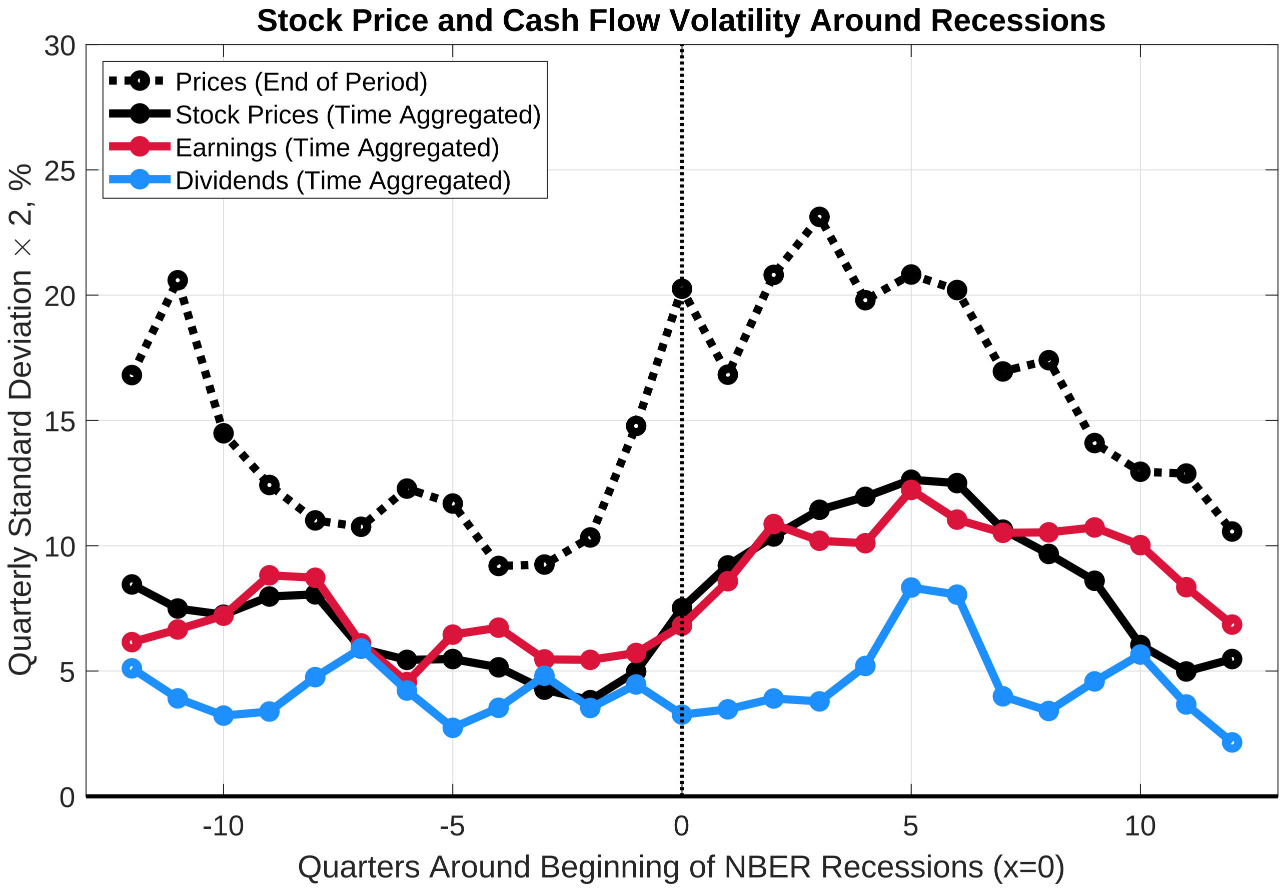

Table V Stock Market and Consumption Volatility Around Recessions

This table reports the annualized volatility (%) of log changes of stock prices ($\triangle p_t$), earnings ($\triangle e_t$), dividends ($\triangle d_t$), and nondurable consumption ($\triangle ndr_t$) the year before the beginning of a recession (h=-5 to h=-2) and the year after the beginning of a recession (h=+1 to h=+4). The reported recession variance ratio is the squared ratio of the recession volatility (+1:+4) over the pre-recession volatility (-5:-2). The last four rows report the recession variance ratio from simulations of the classic model (with homoscedastic stock prices and cash flows), the long-run risk model (Bansal and Yaron, 2004, LRR), the habit model (Campbell and Cochrane, 1999) and the rare disaster risk model (Wachter, 2013, RDR) as in Section III of the paper (Figures 1, 2 and 3). Habit "large rec." provides results based on 20% of the largest recessions in the habit model (Figures A.2). RDR "conditional" provided results from the rare disaster risk model when conditioning out realized rare disaster, while RDR "population" provides results based on all recessions observed in the simulation.

	Prices		**Cash flows**			
	$\triangle p_t$ E.o.P	$\triangle p_t$ T.A.	$\triangle e_t$ T.A.	$\triangle d_t$ T.A.	$\triangle ndr_t$ T.A. reported	$\triangle ndr_t$ T.A. unfiltered
	Empirical data, 1950 - 2016					
pre-recession vola. (-5:-2)	10.12	4.68	6.02	3.67	0.94	1.43
post-recession vola. (+1:+4)	20.15	10.73	9.93	4.18	1.29	1.82
variance ratio: post/pre	3.97	5.25	2.72	1.29	1.90	1.62
	Simulated models					
variance ratio, classic	0.99	1.06	1.06		1.04	
variance ratio, LRR	0.99	1.04	0.92		1.06	
variance ratio, Habit	1.12	1.06	0.94		1.03	
variance ratio, Habit "large rec."	1.44	1.55	0.99		0.97	
variance ratio, RDR "conditional"	1.01	0.94	0.78		0.77	
variance ratio, RDR "population"	3.27	6.63	32.81		30.68	

V. What Does a Model Need to Explain the Data?

Expected returns can rise if a), the amount of risk in the economy increases (cash flow/consumption uncertainty increases) or b), the price of risk changes (risk aversion goes up). In this section, I show that the recession variance ratio (variance during recessions divided by the variance just before recessions) of stock prices compared to cash flows is informative to learn about the two channels.

Model: I adapt a reduced form model as presented in Cochrane (2005, Chapter 20). I assume that the economy is described by the following three equations:

$$r_t = z_t + \sigma_{r,t}\varepsilon_{r,t}, \tag{2}$$

$$z_t = bz_{t-1} + \sigma_{\delta,t}\delta_t, \tag{3}$$

$$\triangle d_t = \sigma_{d,t}\varepsilon_{d,t}, \tag{4}$$

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:

$$d_{t+1} - p_{t+1} = b\left(d_t - p_t\right) + \frac{\sigma_{\delta,t+1}\delta_{t+1}}{1-\rho b}, \tag{5}$$

$$r_{t+1} = \left(1-\rho b\right)\left(d_t - p_t\right) + \left(\varepsilon_{d,t} - \frac{\rho}{1-\rho b}\sigma_{\delta,t+1}\delta_{t+1}\right), \tag{6}$$

$$\triangle p_{t+1} = \left(1-\rho b\right)\left(d_t - p_t\right) + \left(\varepsilon_{d,t} - \frac{1}{1-\rho b}\sigma_{\delta,t+1}\delta_{t+1}\right), \tag{7}$$

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

$$\sigma^2_{p,t} = \sigma^2_{d,t} + \sigma^2_{dp,t} + 2\rho_{d,dp}\sigma_{d,t}\sigma_{dp,t} \tag{8}$$

$$\sigma^2_{dp,t} = \frac{\sigma^2_{\delta,t}}{(1-\rho b)^2}. \tag{9}$$

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.

Parameters: I use almost the same model parameters as suggested by Cochrane (2005) and set $b = 0.9$, $\rho = 0.96$, $\sigma_d = 0.10$, $\sigma_\delta = 0.017$, $\rho_{r,dp} = -0.7$, $\rho_{d,dp} = 0.14$.[35] I scale these annual parameters to monthly counterparts such that I can compute the change in prices and dividends for 12-month "end of period" and "time-aggregated" data observed at a quarterly frequency (in the same way as described in Section III).

Shocks: To mimick the behaviour of the stock market during recessions, I assume that dividends are hit by a negative shock ($\sigma_{d,t}\varepsilon_{d,t}$) that totals to -10% and is evenly distributed over the four quarters from -2 to +1 around the beginning of a recession. Because of the time aggregation bias, it will appear as if dividends drop with the beginning of a recession. For all other observations, I set $\sigma_{d,t}\varepsilon_{d,t} = 0$ to focus on recessions. Similarly, I assume that expected returns increase by 1.36% such that the price-dividend ratio drops contemporaneously by -10% $(0.10 = 0.0136/(1 - b\rho))$.[36]

Results: Figure 9 illustrates the price-dividend ratio around recessions in this simple model. The parameters and shocks are set to get close to the average empirical recession (Figure 5).

In the previous section of this paper (Figure 8, Table V), I find that the variance of prices increases by a factor of four during recessions. The variance of cash flows, however, only increases by a factor between 2.7 (earnings) and 1.4 (dividends and consumption). How can it be explained that stock price volatility increases so dramatically?

Equation 8 suggests a simple answer. As shown in Table VI, if the variance of dividends increases by 2.5 (which is about the upper bound of my estimates) the variance of prices should

[35] These parameters are chosen to match with evidence from price-dividend ratio regressions.

[36] For all other observations, I set shocks to expected returns such that the price-dividend ratio is constant.

only increase by a factor of 1.6, 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 covariance term must go up ($\rho_{d,dp}\sigma_{d,t}\sigma_{dp,t}$). The covariance term is not promising. Numerically, the maximal possible covariance I can achieve is by setting $\rho_{d,dp} = 1$. However, this parameter would imply a low correlation between returns and the price-dividend ratio, which is counter-factual to the data (e.g., Cochrane, 2005). In any case, the hypothetical scenario of $\rho_{d,dp} = 1$ would result in a recession variance ratio that is 2.8, i.e. still not close to 4. The more plausible route is to increase the variance of discount rate shocks $\sigma^2_{\delta,t}$ by a large factor of 5(!).

The habit model by Campbell and Cochrane (1999) is a prime example of a theory that features such a mechanism. As consumption drops towards the habit, relative risk aversion increases and expected returns get more sensitive to shocks (in this model, a change in the consumption surplus ratio). As a result, stock price volatility increases, even though cash flow volatility is constant by assumption. However, my simulation of recessions in the habit model (Table V) shows that the model only increases stock price volatility by a factor of 1.1 (1.5, if one looks at very "large" simulated recessions). This suggests that the large degree of discount rate volatility required to explain the data is arguably difficult to generate by leading asset pricing theories, including the habit model.

My interpretation of these findings is that linking investor preferences and expectations such that discount rates are highly volatile during recessions is key to get asset pricing theories closer to the data. In addition, discount rate volatility must be linked to short-term consumption news such that the drop in prices is amplified during recessions.

Figure 9

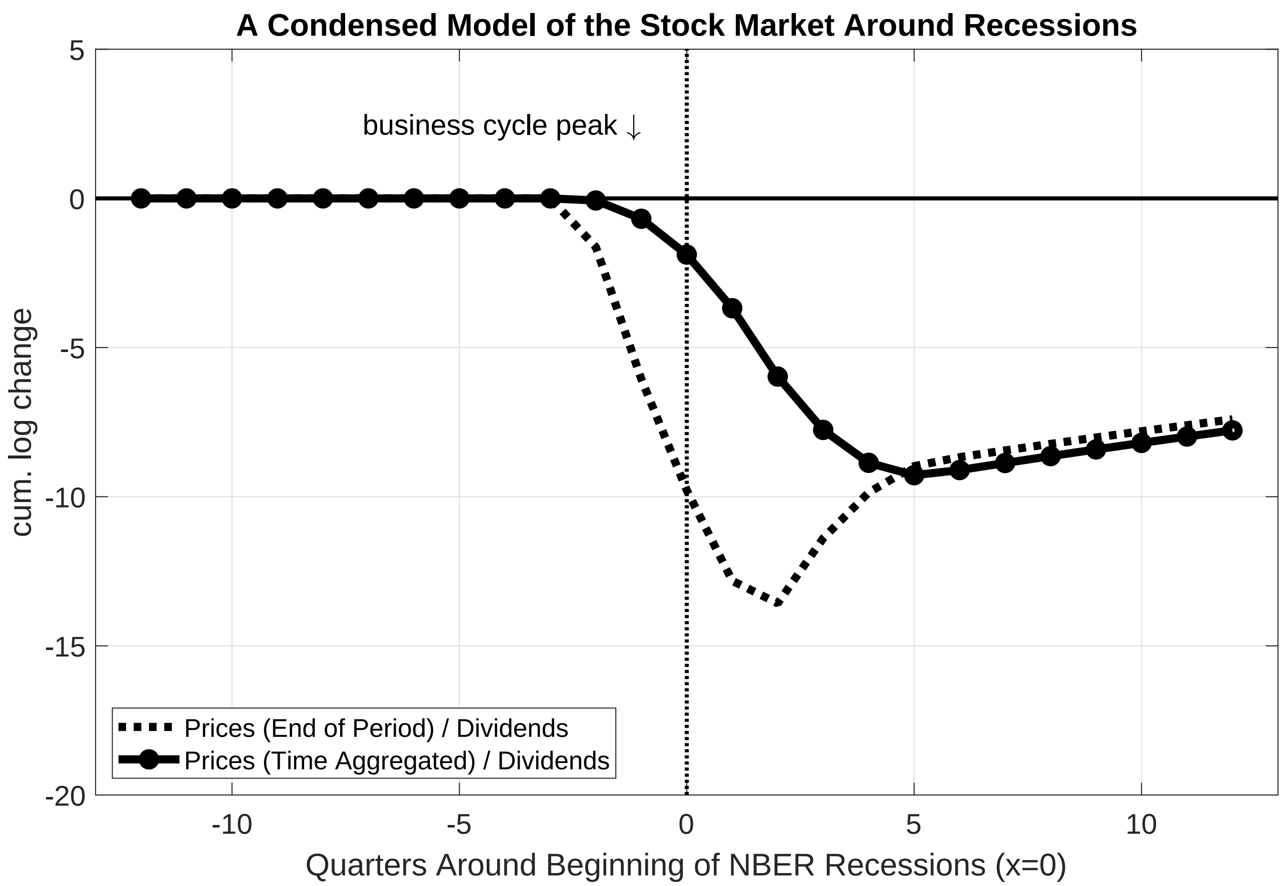
A Condensed Model of the Stock Market Around Recessions
business cycle peak ↓
5
0
-5
-10
-15
-20
cum. log change
-10
-5
0
5
10
Quarters Around Beginning of NBER Recessions (x=0)
Prices (End of Period) / Dividends
Prices (Time Aggregated) / Dividends

Table VI Back-of-the-Envelope Calculation of Stock Market Variances

This table shows the recession variance ratio for dividend and stock prices in a condensed model. Stock price volatility has three components:

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

where, $\sigma_{d,t}^2$ is the variance of dividends, $\sigma_{dp,t}^2 = \sigma_{\delta,t}^2 / (1-\rho b)^2$ is the scaled variance of innovations in expected returns, $\rho_{d,dp}$ is the correlation between dividend growth and the dividend yield. The recession variance ratio is the variance during recessions dividend by the variance before recessions. As is explained in the text, a large value for $\rho_{d,dp}$ is not plausible and is only provided as point of reference.

		Cash flows, $\triangle d_t$	**Prices,** $\triangle p_t$		
before recessions	σ_d	.10	.10	.10	.10
	$\rho_{d,dp}$	.14	.14	.14	.14
	σ_δ	.017	.017	.017	.017
during recessions	σ_d	$0.10 \times \sqrt{2.5}$	$0.10 \times \sqrt{2.5}$	$0.10 \times \sqrt{2.5}$	$0.10 \times \sqrt{2.5}$
	$\rho_{d,dp}$	.14	.14	1.00	.14
	σ_δ	.017	.017	.017	$.017 \times \sqrt{5}$
recession variance ratio		2.5	1.6	2.8	4.0

VI. Further Results

A. *Did Stock Prices Predict Any Recessions? (1950-2016)*

In Section IV I show that average stock prices do not drop ahead of recessions. They rather fall contemporaneously with cash flows and consumption. But do these results also hold for individual recessions? Figure 10, from left to right, shows the cumulative price-dividend ratio, dividends and earnings recession by recession. Because I compare one variable across different recessions, the y-axis now shows the level as measured from the peak, i.e. one quarter before the beginning of a recession. It is easy to see that stock prices are always higher, recession by recession, before compared to after the beginning of a recession. Stock prices drop contemporaneously with cash flows. This also shows that the baseline results are not driven by one or two extreme observations. There is, of course, large variation in the degree of the drop, but the timing of prices and cash flows is arguably quite similar.

Figure 10. The Stock Market Around NBER Recessions, 1950-2016

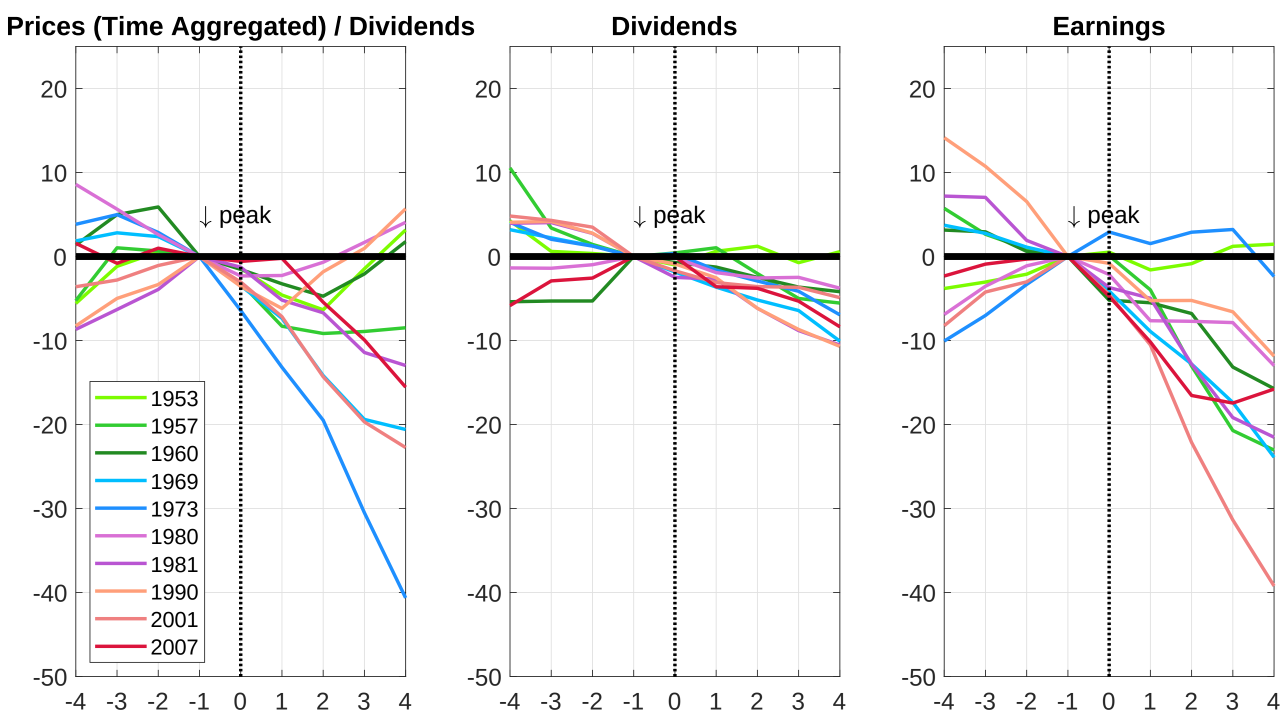

B. Evidence From Annual Data Since 1871

The baseline results focus on the period 1950 - 2016 (ten recessions), because this allows me to study the data sampled at a quarterly frequency. Quarterly data are helpful to pinning down the timing of prices and cash flows as precise as possible. For a longer sample period, 1871 - 2016, I also analyse an annually sampled dataset that covers a total of 29 NBER recessions. There is an important 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 basically 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.

Figure 11. The Stock Market Around NBER Recessions, 1871-2016

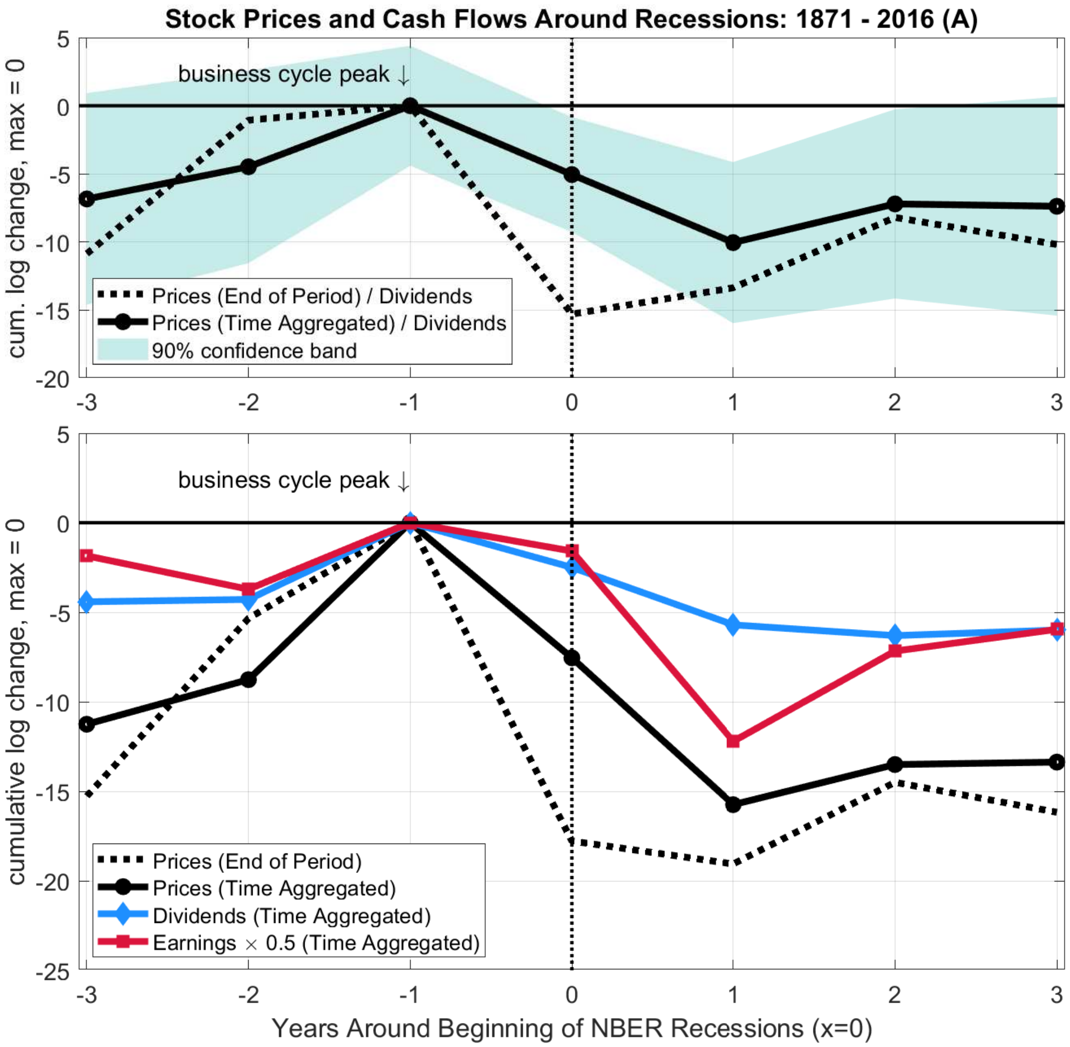

Figure 11 show the results for the longer sample. The figure closely resamples the baseline results, although the picture is indeed less granular. I find that time-aggregated stock prices and cash flows drop around recessions contemporaneously. As in the quarterly data, end of period stock prices lead cash flows, which I can attribute to the time aggregation bias. I find that the significance levels of the linear projections are comparable to the quarterly dataset (reflecting the trade-off mentioned above).

C. Further Results (Internet Appendix)

It might appear a bit surprising that the time-varying volatility channel of the long-run risk model does not show up during recessions. However, the volatility channel simply is not systematically linked to consumption growth. To further illustrate this point, the Internet Appendix also shows sorts by volatility troughs. During these times, price-dividend ratios fall contemporaneously when volatility starts to rise. But now cash flows and consumption remain unchanged.

Regarding the habit model, I find that the increase in stock price volatility is much larger compared to what the habit model predicts. Looking into the simulation details, I find that stock market volatility increases more visibly during large recessions. As I show in the Internet Appendix, these large recessions also have a larger and more persistent effect on stock prices compared to the "normal" recessions in the data. In short, although the habit model heads into the right direction, it falls short in explaining the large increase in expected return volatility.

The Internet Appendix also provides results for earnings based on alternative cash flow measures, in particular earnings as reported in the Shiller dataset. Finally, I show in the Internet Appendix that parametric estimates of the conditional volatility of stock prices and cash flows corroborate the conclusion based on non-parametric estimates. The timing of the increase in volatilities is somewhat different to the non-parametric estimates; the measured increases in conditional volatilities is fairly similar.

VII. Conclusion

Why do stock prices fall so much around recessions? I find that stock prices move almost contemporaneously with dividends, earnings, and consumption. I interpret this finding as direct evidence against the idea that cash flows and consumption have a predictable component at the business cycle frequency. This suggests that stock prices drop because expected returns rise. I find that the variance of stock prices increases more during recessions than the variance of cash flows. This result indicates that the price of risk substantially increases during recessions. I provide evidence that innovations in expected returns must be highly volatile during recessions to generate sufficient high stock price volatility. Theoretical asset pricing models have difficulties in matching the behaviour of the stock market around recessions. I conclude that finding a way of linking investor preferences and expectations such that discount rates are highly volatile during recessions is key to get asset pricing theories closer to the data.

It takes a "better" model to beat a model. The usual metrics (equity premium, volatility, return predictability) for the evaluation of asset pricing models are no longer helpful for model comparisons; they are matched by many models. I hope that the documented behaviour of the stock market around recessions provides a novel opportunity to develop "better" models.

References

Adrian, T., and H. S. Shin (2014): "Procyclical leverage and value-at-risk," *Review of Financial Studies*, 2014(27), 373–403.

Andrei, D., M. Hasler, and A. Jeanneret (2019): "Asset Pricing with Persistence Risk," *Review of Financial Studies*, 32(7), 2809–2849.

Backus, D. K., B. R. Routledge, and S. E. Zin (2010): "The Cyclical Component of US Asset Returns," *Working Paper*.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014): "Volatility, the Macroeconomy, and Asset Prices," *Journal of Finance*, 69, 2471–2511.

Bansal, R., D. Kiku, and A. Yaron (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221.

Bansal, R., and A. Yaron (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., and J. F. Ursua (2008): "Macroeconomic crises since 1870," *Brookings Papers on Economic Activity*, 1(Spring), 255–350.

Beeler, J., and J. Y. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bekaert, G., E. C. Engstrom, and N. R. Xu (2017): "The Time Variation in Risk Appetite and Uncertainty," *Working Paper*.

Binsbergen, J. V., and R. S. Koijen (2010): "Predictive Regressions: A Present-Value Approach," *Journal of Finance*, 65(4), 1439–1471.

Boguth, O., and L.-A. Kuehn (2013): "Consumption Volatility Risk," *Journal of Finance*, 68(6), 2589–2615.

Bollerslev, T., G. Tauchen, and H. Zhou (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22(11), 4463–4492.

Branger, N., H. Kraft, and C. Meinerding (2016): "The Dynamics of Crises and the Equity Premium," *Review of Financial Studies*, 29(1), 232–270.

Brav, A., J. R. Graham, C. R. Harvey, and R. Michaely (2005): "Payout policy in the 21st century," *Journal of Financial Economics*, 77, 483–527.

Breeden, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Breeden, D. T., M. Gibbons, and R. Litzenberger (1989): "Empirical Tests of the Consumption-oriented CAPM," *Journal of Finance*, 44, 231–262.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (2017): *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton Univers. Press.

Campbell, J. Y., and J. H. Cochrane (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., and R. J. Shiller (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–228.

Campbell, S. D., and F. X. Diebold (2009): "Stock Returns and Expected Business Conditions: Half a Century of Direct Evidence," *Journal of Business & Economic Statistics*, 27, 266–278.

Chen, L. (2009): "On the reversal of return and dividend growth predictability: A tale of two periods," *Journal of Financial Economics*, 92, 128–151.

Cochrane, J. (1994): "Permanent and Transitory Components of GNP and Stock Prices," *Quarterly Journal of Economics*, 109(1), 241–266.

Cochrane, J. H. (1996): "A Cross-Sectional Test of an Investment-Based Asset Pricing Model," *Journal of Political Economy*, 104, 572–621.

——— (2005): *Asset Pricing*. Princeton University Press, New Jersey.

——— (2008): "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies*, 21, 1533–1575.

——— (2017): "Macro-Finance," *Review of Finance*, 21(3), 945–985.

Constantinides, G. M., and A. Gosh (2011): "Asset Pricing Tests with Long-run Risk Consumption Growth," *Review of Asset Pricing Studies*, 1(1), 96–136.

Cujean, J., and M. Hasler (2017): "Why Does Return Predictability Concentrate in Bad Times?," *Journal of Finance*, 72(6), 2717–2758.

Dangl, T., and M. Halling (2012): "Predictive regressions with time-varying coefficients," *Journal of Financial Economics*, 106(1), 157–181.

Dew-Becker, I. (2017): "How risky is consumption in the long-run? Benchmark estimates from a robust estimator," *Review of Financial Studies*, 30, 631–666.

Fama, E. F. (1990): "Stock Returns, Expected Returns, and Real Activity," *Journal of Finance*, 45, 1089–1108.

Fama, E. F., and H. Babiak (1968): "Dividend Policy: An Empirical Analysis," *Journal of the American Statistical Association*, 63, 1132–1161.

FAMA, E. F., AND K. R. FRENCH (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

FERSON, W. (2019): *Empirical Asset Pricing: Models and Methods*. The MIT Press.

FERSON, W. E., AND C. R. HARVEY (1991): "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385–415.

GABAIX, X. (2012): "An exactly solved framework for ten puzzles in macro-nance," *Quarterly Journal of Economics*, 127(2), 645–700.

GOLEZ, B., AND P. KOUDIJS (2018): "Four Centuries of Return Predictability," *Journal of Financial Economics*, 127, 248–263.

GOURIO, F. (2008): "Time-series predictability in the disaster model," *Finance Research Letters*, 5(4), 191–203.

GOYAL, A., AND I. WELCH (2008): "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455–1508.

GROSSMAN, S. J., AND R. J. SHILLER (1980): "Preliminary Results on the Determinants of the Variability of Stock Market Prices," Working Paper, University of Pennsylvania.

——— (1981): "The Determinants of the Variability of Stock Market Prices," *American Economic Review: Papers & Proceedings*, 71, 222–227.

HALL, R. E. (2017): "High Discounts and High Unemployment," *American Economic Review*, 107(2), 305–330.

HANSEN, L. P., J. C. HEATON, AND N. LI (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116, 260–302.

HANSEN, L. P., AND K. J. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1286.

——— (1983): "Stochastic consumption, risk aversion, and the temporal behavior of asset returns," *Journal of Political Economy*, 91, 249–268.

Harrison, P., and H. H. Zhang (1999): "An Investigation of the Risk and Return Relation at Long Horizons," *Review of Economics and Statistics*, 81, 399–408.

Henkel, S. J., J. S. Martin, and F. Nardari (2011): "Time-varying short-horizon predictability," *Journal of Financial Economics*, 99(3), 560–580.

Jank, S. (2015): "Changes in the Composition of Publicly Traded Firms: Implications for the Dividend-Price Ratio and Return Predictability," *Management Science*, 61, 1362–1377.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Kroencke, T. A. (2017): "Asset Pricing without Garbage," *Journal of Finance*, 72, 47–98.

Larrain, B., and M. Yogo (2008): "Does firm value move too much to be justified by subsequent changes in cash flow?," *Journal of Financial Economics*, 87, 200–226.

Leary, M. T., and R. Michaely (2011): "Determinants of Dividend Smoothing: Empirical Evidence," *Review of Financial Studies*, 24, 3197–3249.

Lettau, M., and S. C. Ludvigson (2009): *Measuring and modeling variation in the risk return trade-off*, vol. Volume 1, Chapter 11 of *Handbooks in Finance*. Handbook of Financial Econometrics: Tools and Techniques.

Lintner, J. (1956): "Distribution of Income of Corporations Among Dividends, Retained Earnings and Taxes," *Amercian Economic Review*, 46, 37–113.

Lustig, H., and A. Verdelhan (2012): "Business Cycle Variation in the Risk-Return Trade-off," *Journal of Monetary Economics*, 59, 35–49.

Maio, P. F., and I. Cooper (2018): "Asset growth, profitability, and investment opportunities," *Management Science*, forthcoming.

Marakani, S. (2009): "Long-run Consumption Risks: Are They There?," Working Paper, City University of Hong Kong.

Martin, I. (2017): "What is the Expected Return on the Market?," *Quarterly Journal of Economics*, 132, 367–433.

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

Møller, S. V., and J. Rangvid (2018): "Global Economic Growth and Expected Returns Around the World: The End-of-the-Year Effect," *Management Science*, 64(2), 495–981.

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.

Rangvid, J., M. Schmeling, and A. Schrimpf (2014): "Dividend Predictability Around the World," *Journal of Financial and Quantitiative Analysis*, 49, 1255–1277.

Rapach, D. E., J. K. Strauss, and G. Zhou (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," *Journal of Monetary Economics*, 22(1), 117–131.

Samuelson, P. (1966): ""Science and Stocks," *Newsweek*, September 19, 66.

Savov, A. (2011): "Asset Pricing with Garbage," *Journal of Finance*, 66, 177–201.

Schreindorfer, D. (2019): "Macroeconomic Tail Risks and Asset Prices," *Review of Financial Studies*, forthcoming.

Schwert, G. W. (1989): "Why Does Stock Market Volatility Change Over Time?," *Journal of Finance*, 44, 1115–1153.

Seo, S. B., and J. A. Wachter (2018): "Do Rare Events Explain CDX Tranche Spreads?," *Journal of Finance*, 73(5), 2343–2383.

Shiller, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *Amercian Economic Review*, 71, 421–436.

——— (2000): *Irrational Exuberance*. Princeton University Press 2000, Broadway Books 2001, 2nd ed., 2005.

Siegel, J. J. (2016): "The Shiller CAPE Ratio: A New Look," *Financial Analysts Journal*, 72(3), 41–50.

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," *Review of Finance*, 19(4), 1–39.

Vassalou, M. (2003): "News Related to Future GDP Growth as a Risk Factor in Equity Returns," *Journal of Financial Economics*, 68, 47–73.

Wachter, J. A. (2005): "Solving Models with External Habit," *Finance Research Letters*, 2, 210–226.

Wachter, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *Journal of Finance*, 68(3), 987–1035.

Wilcox, D. W. (1992): "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work," *Amercian Economic Review*, 82, 922–941.

Working, H. (1960): "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica*, 28, 916–918.

For Online Publication

Internet Appendix

Recessions and the Stock Market

Details on the Simulated Asset Pricing Models

Classic Model: For the classic model, I assume Epstein-Zin preferences with a coefficient of relative risk aversion of 15, and an elasticity of inter-temporal substitution of 1.5. In this setting, Epstein-Zin preferences avoid the risk-free rate puzzle; with respect to the equity premium, one would get comparable results using CRRA preferences. Consumption follows an i.i.d. process. Following the evidence provides by Savov (2011) and Kroencke (2017), is set the (annual) volatility of consumption growth to 2.5% (see also Dew-Becker, 2017) and the correlation to stock dividend growth to 0.60. I then use the log-linearization technique described in Beeler and Campbell (2012), and Bansal, Kiku, and Yaron (2012), to solve for stock prices, dividends, and returns. To simulate the model, I use the MATLAB function BKY_generate_m.m from Beeler and Campbell (2012). I slightly modify the function to get prices.

Long-Run Risk Model: I use the model and parameters as in Bansal and Yaron (2004). The coefficient of relative risk aversion is 10 and the elasticity of inter-temporal substitution is 1.5; consumption growth volatility is set to 2.5%. The model is solved using the log-linearization technique described in Beeler and Campbell (2012), and Bansal, Kiku, and Yaron (2012). To simulate the model, I use the MATLAB function BKY_generate_m.m from Beeler and Campbell (2012), as for the classic model. I slightly modify the function to get prices.

Habit Model: I use the model and parameters as in Campbell and Cochrane (1999). To simulate the model, I convert the Gauss programs used in Campbell and Cochrane (1999) to MATLAB. The GAUSS files are provided by John Cochrane on his website. By the conversion to MATLAB, I have to make some adjustments. First, to solve the price-dividend ratio, I use the GaussLegendre.m function provided by Pavel Holoborodko on his website. This numerical integration routine is much faster than the build-in routine. Second, as recommended by Wachter (2005), I use a much finer grid for the consumption ratio to solve the mode. More specifically, I use an upper segment of 50 equally spaced points between zero and maximum consumption surplus ratio and a lower segment of 450 logarithmically spaced points between the lowest value of the upper segment and exp(-300). I slightly modify the function to get prices.

Rare Disaster Model: I use the model and parameters as in Wachter (2013). To simulate the model, I adapt the MATLAB programs provided to the paper Seo and Wachter (2018). These programs can be found on the website of the *Journal of Finance* to the paper. Again, I slightly modify the function to get prices.

Comparison of Simulation Results to the Original Papers: For the advanced asset pricing models, I use the same parameters as in the original papers and I use the same, or very similar, solution techniques. Table I shows the equity premium, stock market volatility, and return predictive regressions from my model simulations for a sample with a large number of observations (10,000 years). These values can be directly compared to the original papers:

- Habits: Campbell and Cochrane (1999), Tables 2 and 5; or Wachter (2005), Table 2, who proposes the finer grid, which leads to somewhat different results.
- Long-run risk: Bansal and Yaron (2004), Table 4; Bansal, Kiku, and Yaron (2012), Tables 2 and 5; or Beeler and Campbell (2012) Tables 2 and 4.
- Time-varying rare disaster risk: Wachter (2013), Table 2 and 3 .

Overall, I get very similar results to the original publications. There are some relatively small differences, which are likely to be explained by i) simulation noise, or ii) the fact that I provide results from one large sample where the literature sometimes reports results from a finite sample. An exception might be the return predictive regression in the Bansal and Yaron (2004) model, which leads to stronger results in the original paper. As discussed in Beeler and Campbell (2012), the difference comes from the fact that Bansal and Yaron (2004) use a different model simulation technique compared to Beeler and Campbell (2012) and Bansal, Kiku, and Yaron (2012), which is supposed to deliver less accurate results.

Volatility Shocks in the Long-Run Risk Model

Figure A.1 shows the impact of changes in volatility on stock prices according to the long-run risk model by Bansal and Yaron (2004). The y-axis is the same as in all other comparable figures to facilitate comparisons. Dividend growth is flat around large changes of conditional volatility. This is the mirror image of Figure 2 and further illustrates that recessions (low economic activity) and time-varying volatility are unrelated in the long-run risk model. The figure also illustrates that prices move approximately 1:1 with changes in volatility and generate drops in prices that are relative mild in magnitude.

Figure A.1

This figure shows results for the simulated long-run risk model of Bansal and Yaron (2004) around times of volatility troughs. The simulation is the same as described in Section III of the main paper. The top figure shows the price-dividend ratio, the lower figure prices, dividends, long-run consumption and dividend growth, and the conditional variance.

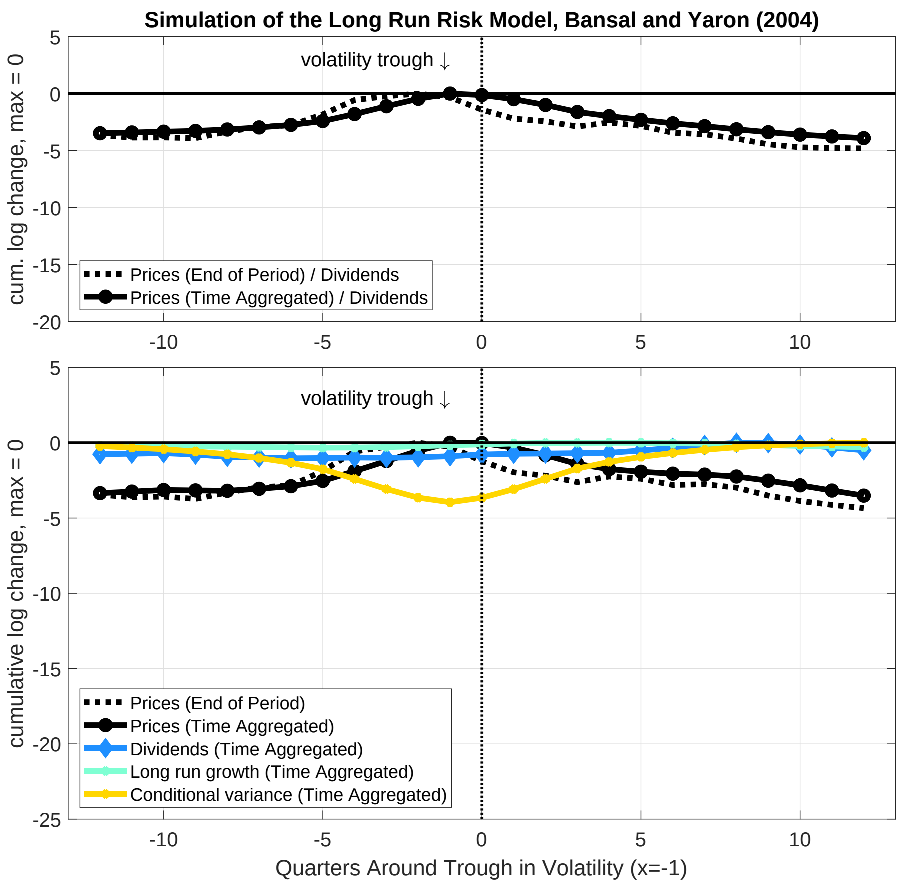

Large Recessions in the Habit Model

Figure A.2 shows results for "large" recessions in the habit model by Campbell and Cochrane (1999). The results in the main paper (black lines in this figure) are based on the 25% highest local peaks in annual consumption (this converts to about 2.5% of all simulated observations). I want to check how sensitive my results are with respect to my definition of "simulated recessions". For that purpose, I simply pick the 20% of the "largest" recessions (5% of the highest local peaks) and re-draw prices and the conditional volatility (red lines) in this figure. Looking at very large recessions reveals that prices drop deeper and the increase in the conditional volatility is more pronounced the stronger the recession is. The recession variance ratio (the variance one year before recessions dividend by the variance during recessions) is as large as 1.5 for the largest recessions (compared to 1.1 for all recessions). However, this number is still much less compared to the empirical data (4).

Figure A.2

This figure shows results for the simulated habit model of Campbell and Cochrane (1999). Black lines show the same results as described in Section III of the main paper. Red lines show results for 20% of the "largest" recessions.

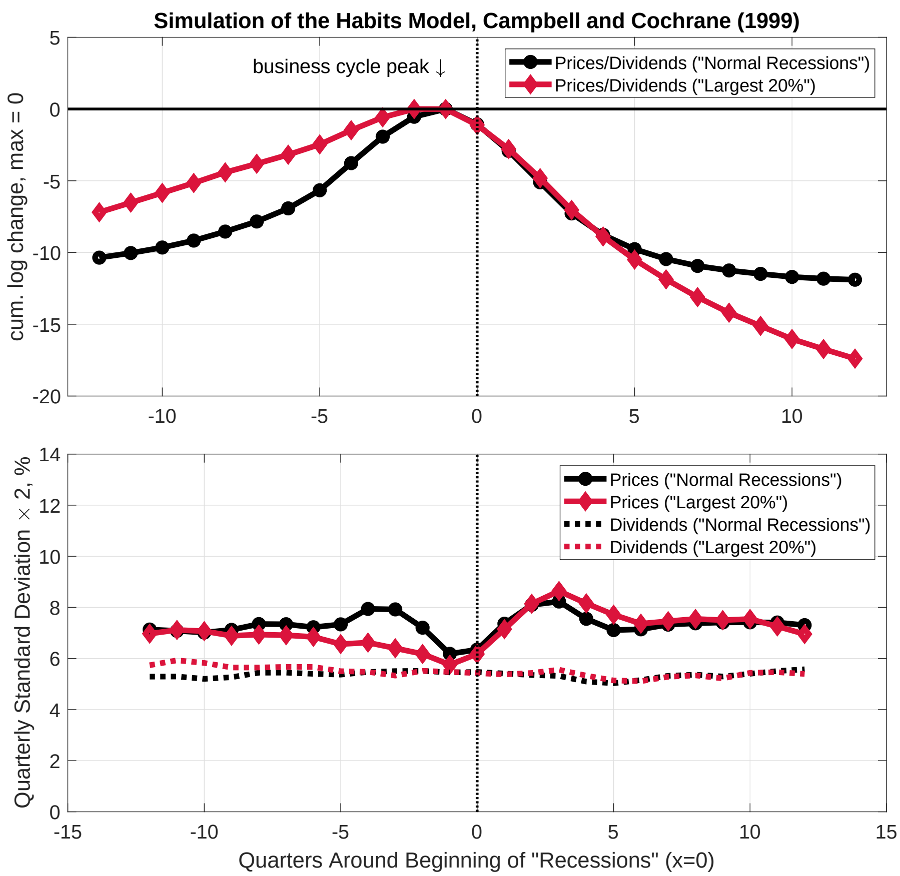

Unfiltered NIPA Consumption

True consumption is unobservable. NIPA statisticians estimate consumption based on proxies that can be thought of as true consumption plus measurement error. The quite complex estimation procedure of NIPA consumption potentially drives a wedge between the properties of reported consumption and true consumption, such that empirical inference can be substantially affected (e.g., Wilcox, 1992). Indeed, Savov (2011) finds that garbage growth as a more simple measure of consumption is highly correlated with stock returns and substantially reduces the equity premium puzzle.

Kroencke (2017) suggests that if observable consumption is measured with error, it is optimal for consumption statisticians to filter the observed proxies of consumption to pin down the level of true consumption as precise as possible. Even if filtering is optimal for the purpose of estimating the level of consumption, it can be hazardous for other applications where measurement error cancels out anyway, e.g., when computing consumption covariances in asset pricing applications, or when averaging consumption over many events as in this study. Kroencke (2017) argues that the true and complex estimation procedure might be approximated by a simple Kalman-filter model:

$$\hat{c}_t = \hat{c}_{t-1} + K\left(y_t - \hat{c}_{t-1}\right), \tag{10}$$

where $\hat{c}_t$ is the reported (or estimated) level of NIPA consumption, y_t is a noisy measure of the level of consumption (e.g., retail sales, or garbage) and K is the filter parameter. This equation can be reversed to:

$$\hat{y}_t = \frac{\hat{c}_t - (1-K)\,\hat{c}_{t-1}}{K}, \tag{11}$$

which can be used to derive a series of "unfiltered" NIPA consumption. The filter parameter K is a signal to noise ratio. It is close to 1 if there is almost no measurement error, i.e. in this case consumption statisticians fully update. If there is (a lot of) measurement error, K should be (a lot) below 1. Kroencke (2017) shows that the new measure of unfiltered NIPA consumption has a large correlation with stock returns and matches other properties of the garbage series of Savov (2011), such that the equity premium puzzle can be substantially mitigated. Importantly, unfiltered NIPA consumption can be used when garbage data are not available, for example at the quarterly frequency, as in this study.

Calibration: I use the same (constant) unfilter parameters as derived in Kroencke (2017), i.e. $K = 0.58$ for aggregate consumption (nondurables and services) and $K = 0.71$ for the easier to measure and thus less heavily filtered nondurable consumption series. I unfilter first, second, third, and fourth quarter consumption separately to avoid the time aggregation bias in the unfilter equation. In a second step, I then time aggregate all consumption data such that the timing is comparable to time-aggregated dividends and the other time-aggregated data.

Further Results on the Forward Term Structure of Expected Growth

Figure A.3 shows mean expected real GDP growth for the current quarter (nowcast) as well as the expected real GDP growth four quarters into the future ("forecast Q4") as reported by the Survey of Professional Forecasters. In line with the event-time figure provided in the main paper, I find that short-horizon expectations drop during recessions whereas longer horizon expectations do not show such a consistent business cycle behaviour. However, it is also clear that longer horizon expectations change over time, but apparently not at the frequency of the business cycle.

Figure A.3

This figure shows forward expected real GDP growth. The nowcast refers to the current quarter expected growth rate whereas the forecast refers to the expected growth rate four quarters into the future. The data come from the Survey of Professional Forecasters and are available at the website of the Philadelphia Fed; the sample period is from Q4/1968 to Q4/2016. Further details on the variable construction are provided in the caption to table **??**

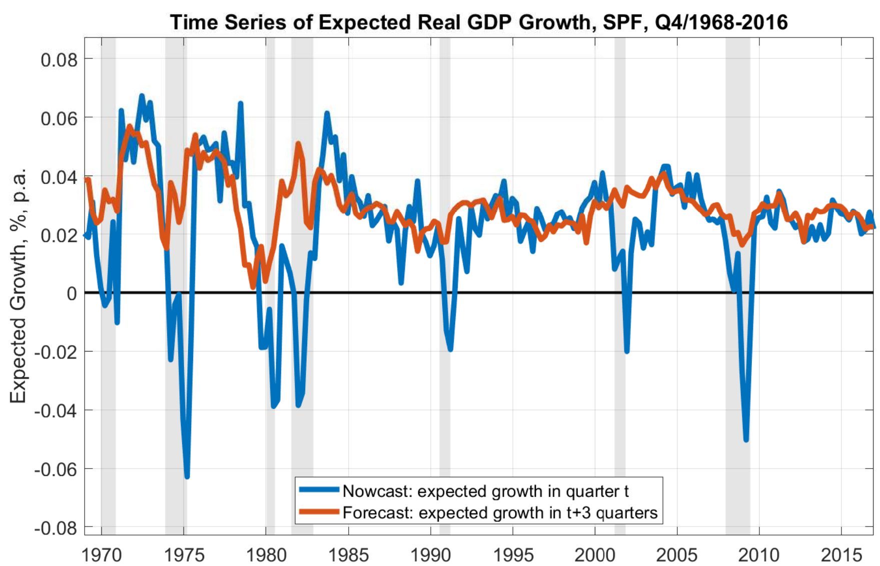

Alternative Cash Flow Measures

The baseline results rely on dividends and earnings coming from two sources. For the period from 1950 to 1974, I use the data provided by Robert J. Shiller on his website.For the period from 1975 to 2016, I rely on prices, dividends, and earnings for the S&P 500 index as provided by the Thomson Reuters Datastream (Series: S&PCOMP; datatypes: MV, DSDY, DSPE.). The main difference between the two datasets can be observed for the earnings series. Shiller's series is sourced from Standard and Poors and comes without any adjustments. The Datastream series is on a "continuing operations" basis. For example, earnings are adjusted for extraordinary profits/losses. As can be seen in Figure A.4, before 1990, there is a relative small difference between the two series. However, after 1990, corporations seem to make more and more use of extraordinary profits/losses which leads to a much more volatile unadjusted earnings series.[37] I use the adjusted Datastream series for the baseline results, as this series seems to be less affected by corporate policy and, thus, should be closer cash-flows as in the theoretical models.

Figure A.4. Alternative Cash Flow Measures, 1950-2016

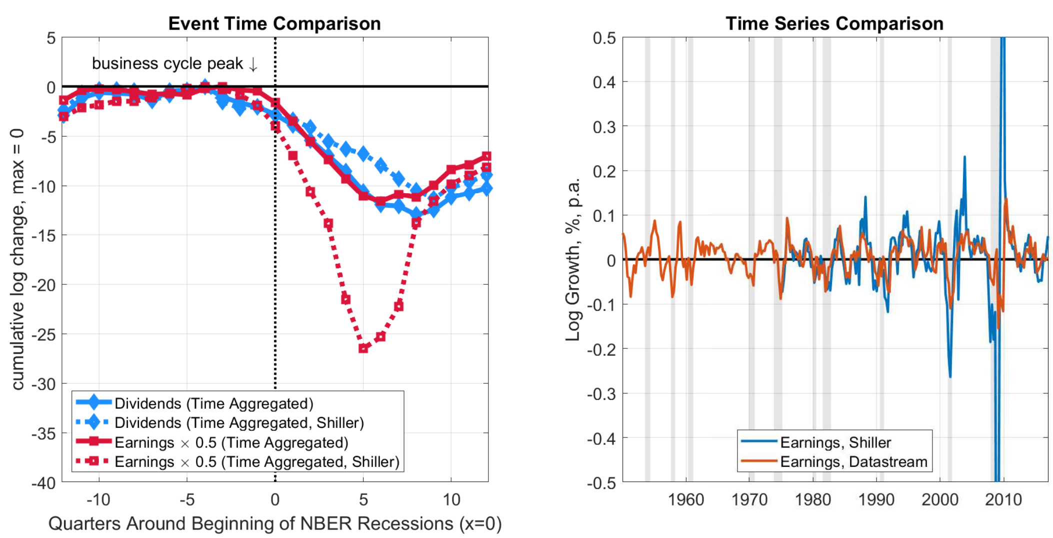

[37] See Siegel (2016) for a discussion.

Parametric Estimates of Conditional Volatility (1950-2016)

The baseline results on the conditional volatility of stock prices and cash flows are based on non-parametric estimates. Evidence on the change in the volatility for cash flows is inconclusive, in the sense that earnings volatility increases during recession while dividend volatility remains more or less constant. A potential concern could be that the non-parametric estimates are simply noisy and obscure the true pattern. Figure A.5 provides EGARCH(4,4,4) estimates of the conditional volatility. The figure suggests that the parametric method needs more time to catch the increase in volatility. In Figure 8, the non-parametric estimates of end of period stock prices jump immediately with the beginning of a recession. Otherwise, the pattern in the figure is very similar compared to the non-parametric estimates reported earlier.

Figure A.5. The Stock Market Around NBER Recessions, 1950-2016