Cyclical β *

Paul Ehling† Costas Xiouros‡
BI BI

August 2023

Abstract

In a framework where the CAPM holds conditionally, we develop a model for the cross-section with fairly general dividend dynamics and a stochastic discount factor that accounts for standard asset pricing moments. Our model, which successfully captures unconditional betas across industries, features a cyclical component that induces significant variations in conditional betas. These variations exhibit non-linear patterns in relation to the business cycle. Using the time-series of the betas implied by our model, we find that the conditional CAPM explains the cross section of industry average quarterly returns over the period from 1927–2021.

Keywords: Conditional CAPM, cyclicality, cross-section of industry returns.

JEL Classification: E32, E37, G12.

*We would like to thank seminar participants at a BI Research Day and the 14th International Conference of the ERCIM WG on Computational and Methodological Statistics (CMStatistics 2021) at King's College London and Andrew Ang for helpful comments and suggestions. Paul Ehling and Costas Xiouros thank La Fondation du Risque through the Development of Quantitative Management Research Initiative and the Centre for Asset Pricing Research (CAPR) at BI for funding support. Lian Zhang provided excellent research assistance.

†Department of Finance, BI Norwegian Business School, paul.ehling@bi.no.

‡Department of Finance, BI Norwegian Business School, costas.xiouros@bi.no.

1 Introduction

We develop a conditional CAPM setting where the conditional market betas are endogenous in the cash-flow dynamics. In our framework, there are two determining factors of a conditional beta: The cyclicality of a stock, that is, a stock's cash-flows variability over the (business) cycle; and its cash-flow term-structure. In our reading, the literature has predominantly focused on the second component in an effort to explain certain CAPM anomalies, e.g., growth versus value. Our focus is on the cyclical component.

To this purpose, we find it natural to form portfolios with similar cyclical behavior and, therefore, we study the cross-section of industry portfolios. Using the 30 industry portfolio returns, which balances sample length with detail in industry classification, we re-confirm that industry portfolio average returns are quite similar and only weakly related to their unconditional market betas.[1] Specifically, in Figure 1 the black solid line shows the expected returns predicted by the static CAPM based on the sample unconditional betas; the red circles are the data and the red solid line is the linear fit to the data. The upshot is that the linear fit is considerably flatter than implied by the static CAPM.

In the figure, the blue circles are the average returns predicted by the conditional CAPM model where the conditional beta is driven by the cyclical component of the model and the blue solid line is its linear fit. The blue dashed line is the linear fit of the average returns predicted by the full model that also considers the second component, the term-structure of cash-flows. The model explains industry returns for two reasons: The first is based on the variability of the conditional betas over the business cycle. Specifically, according to the model, the conditional betas of pro-cyclical (counter-cyclical) stocks decrease (increase) during economic downturns, when the market risk premium increases. As a result, the average returns of pro-cyclical (counter-cyclical) stocks are lower (higher) than what their average conditional betas predict. The second reason is related to the average conditional betas. Specifically, according to the model, the average conditional betas of pro-cyclical (counter-cyclical) stocks are lower (higher)

[1]The early evidence that the security market line is flatter than what the static CAPM model predicts dates back to Friend and Blume (1970) and Black, Jensen, and Scholes (1972).

than their unconditional betas over the historical sample.

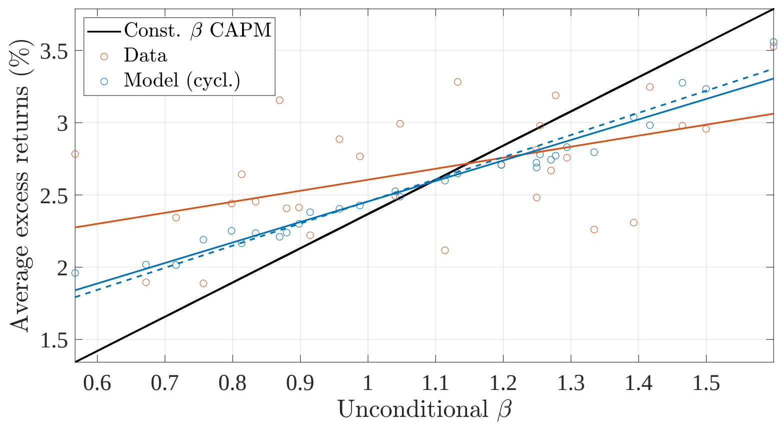

The figure shows the relation between average excess returns and corresponding (unconditional) market betas for the 30 industry portfolios. Data points are denoted by red circles, and the red solid line represents the linear fit. The black solid line represents the average excess returns projected by the static CAPM model. Blue circles depict the average excess returns projected by the cyclical component of the model (MC-I:$\hat{\beta}_{cyc}$), while the blue solid line represents the linear fit for these data. The blue dashed line represents the linear fit of the relation between average excess returns and unconditional betas predicted by the full model (MC-I:$\hat{\beta}_{all}$). The quarterly data are from 1932Q1 to 2020Q4.

Figure 1: Security Market Line – 30 industry portfolios.

The economic magnitude of these two effects is about the same. Together the two effects make the average returns predicted by the model jointly statistically indistinguishable from its empirical counterpart. Specifically, employing the model predicted betas, the conditional CAPM model passes the time-series asset pricing test using the Wald statistic and generates the lowest absolute unexplained mean returns (alpha's) among the empirical asset pricing models considered. In contrast, the static CAPM, the conditional CAPM with running window betas, the Fama-French three-factor (FF3) model, with and without the momentum factor, all fail the asset pricing test.[2]

We build on Hansen and Richard (1987) and Jagannathan and Wang (1996) as these works argue in favor of the view that the CAPM holds conditionally. Jagannathan and Wang (1996) find empirical support by using the corporate yield spread as conditioning variable for the

[2]The factor models considered are also rejected when tested against the 25 size and book-to-market portfolios, consistent with the results of Hou, Xue, and Zhang (2015), who show that these models as well as the $q-$factor model are all rejected by the GRS (Gibbons, Ross, and Shanken, 1989) test.

variations in the betas and the equity premium, and including a measure of returns on human capital as part of the returns on aggregate wealth. A number of other studies provide further empirical evidence in support for the conditional CAPM. For example, Ang, Chen, and Xing (2006) show that stocks whose systematic risk increases during market declines earn higher returns relative to what the static CAPM predicts.[3]

Lettau and Ludvigson (2001) find that the conditional versions of the CAPM and the consumption CAPM go a long way in explaining the size and value premia, where the conditioning variable is the log wealth-consumption ratio.[4] Yet, Lewellen and Nagel (2006) argue that the covariation between betas and the equity risk premium cannot be large enough to explain the value and momentum premia. Their empirical design uses series of short window OLS regressions to infer the variations in conditional betas, and they study the post-1963 data sample. Ang and Chen (2007) estimating instead a time-series model with time-varying betas, risk premium, and conditional volatilities, find that the conditional CAPM explains the value premium over the long data sample from 1926–2001.

More recently, Cosemans, Frehen, Schotman, and Bauer (2016) use a hybrid approach for estimating time-varying betas by shrinking rolling window estimates towards theoretically motivated conditional betas that depend on firm characteristics, the business cycle, and ex-ante heterogeneity. They show that their beta estimates improve portfolio construction and cross-sectional asset pricing tests. Another related study is Cederburg and O'Doherty (2016): Using an instrumental variable approach, they find that the conditional CAPM model explains the "betting against beta" anomaly.[5] Further, Buss and Vilkov (2012) provide support for a monotone relation between market betas and mean returns using option prices to infer forward-looking beta estimates. And Chang, Christoffersen, Jacobs, and Vainberg (2012) also find that option-implied betas are better predictors of future realized betas, compared to betas estimated

[3]In addition, there also is evidence indicating that both corporations (Graham and Harvey, 2001) and investors (Berk and Van Binsbergen, 2016, 2017) use the CAPM to make investment decisions and evaluate risk.

[4]Lustig and Van Nieuwerburgh (2005) provide further support that a conditional version of the CCAPM prices well the 25 size and book-to-market portfolios.

[5]Frazzini and Pedersen (2014) show that a "betting against beta" strategy had positive abnormal performance over the long sample.

using historical returns.

What sets us apart from these studies is mainly that we use a theoretical model, with particular focus on the cyclical behavior of cash-flows, to infer how the conditional betas vary over time. One distinctive feature of our model is that the behavior of the conditional betas is highly non-linear over the business cycle.

Menzly, Santos, and Veronesi (2004) present a comparable general equilibrium model to study the time-series predictability of excess returns and dividend growth. They employ a cash-flow model for the cross-section and the Campbell and Cochrane (1999) model for the stochastic discount factor to disentangle the offsetting effects of shocks to cash-flow expectations and risk-aversion. Using industry portfolios they find support for their model, specifically, how the sensitivity of cash-flow expectations to aggregate shocks affects the predictability relations. They do not, however, examine the cross-section of expected returns.

Generalizing the model of Menzly, Santos, and Veronesi (2004), Santos and Veronesi (2010) use it to examine the value premium. Their main conclusion is that the cyclicality (of what they refer to as cash-flow risk) required for value stocks to explain the value premium is counterfactually high. We abstain from making any preference assumption but instead adopt the approach of Brennan, Wang, and Xia (2004) for modelling the stochastic discount factor. We also use a distinct cash-flow model. We find that our model is unable to explain the average returns of the 25 Fama-French size and book-to-market portfolios. Of course, this is not surprising since we expect cyclicality to be a stronger driver of industry returns rather than affecting traditional stock characteristics. In other words, within a characteristic based portfolio, we expect to have both cyclical and counter-cyclical stocks. Our approach differs in that we use the unconditional beta of a portfolio to identify its cyclicality. Therefore, we find that the variation in conditional betas predicted by the model, given the unconditional betas of these portfolios, cannot explain their average returns. In this sense, our results complement and strengthen the conclusions reached in Santos and Veronesi (2010).

Lettau and Wachter (2007), on the other hand, with a similar approach to ours explain the value premium using a model for the cross-section where ex-ante identical firms differ in their growth state. The key distinctive feature is the downward sloping term structure of equity returns, which is consistent with the empirical evidence provided by Binsbergen, Brandt, and Koijen (2012) and Van Binsbergen and Koijen (2017). To produce this feature, Lettau and Wachter (2007) assume that discount rate shocks are not priced, while cash-flow shocks have a negative risk premium.[6] Consequently, growth firms exhibit longer durations, are more sensitive to cash-flow shocks and require a lower risk premium. In contrast, the model of Santos and Veronesi (2010) as well as our unconstrained model calibration imply instead an upward sloping term structure. Importantly, our model is flexible enough to accommodate an alternative calibration with a downward sloping term structure: Yet, even with the alternative calibration we still find that the conditional CAPM is unable to price the 25 size and book-to-market portfolios. Further, it is worth noting that the CAPM in Lettau and Wachter (2007) does not hold either unconditionally or conditionally by construction, since the market portfolio bears discount rate risk which is not priced. In our model the conditional CAPM holds approximately regardless of the slope of the term structure of equity returns.[7]

Another distinctive feature of our approach compared to the studies of Lettau and Wachter (2007) and Santos and Veronesi (2010) is that we use the theoretical model to generate the time-series of the conditional betas which we then use to test the conditional CAPM using historical returns. They instead compare the premia that their models predict to those found in the data. This distinction turns out to be important because we find that certain differences in average returns are due to the specific history of recessions we have observed.

Our study is also related to the long literature on production-based asset pricing, that has developed several theories in which the conditional CAPM holds, where firm characteristics are related to their conditional betas. Berk, Green, and Naik (1999) were among the first to relate a firm's systematic risk to real investment. Gomes, Kogan, and Zhang (2003) relate small firms to higher and riskier growth options, while growth firms to an expected decline

[6] The negative risk premium of cash-flow shocks in Lettau and Wachter (2007) refer to the assumption that negative systematic shocks increase the conditional mean of the aggregate dividend growth.

[7] The downward sloping term-structure in the alternative calibration requires that discount rate shocks have negative risk premium.

in productivity growth and shorter duration cash-flows. Carlson, Fisher, and Giammarino (2004) explain the size effect in a similar fashion, while they relate a firm's book-to-market to operating leverage. Zhang (2005) and Cooper (2006), on the other hand, propose that the value premium compensates investors for the higher risk of assets in place due to costly investment reversibility. Petkova and Zhang (2005) provide empirical support for the time-variation in the conditional betas of value and growth stocks, but question the economic significance. These theories typically consider firms to be ex-ante identical, however, our focus is on the implications of ex-ante heterogeneity in the cyclical behavior of cash-flows. For this reason we focus on industry rather than characteristics-based managed portfolios.

2 Explaining industry portfolio returns

Before we present the model in detail, we provide intuition for its mechanism and discuss how it performs relative to the data: In our conditional CAPM model the cash-flow dynamics determine the conditional beta of a stock, where the exogenous risk-free rate and price of risk drive the stochastic discount factor (SDF). A slow moving business cycle variable drives the price or risk and aggregate dividends. In our framework, shocks to the SDF exclusively drive aggregate quantities, including the risk-free rate, so that the conditional CAPM almost exactly describes expected returns.[8]

The cash-flow dynamics of a stock are specified by modeling its dividend share, which contains a cyclical and an idiosyncratic component. The latter is driven by a "short"-run and a "long"-run state. As a result, the conditional beta of a stock is a function of the state of the aggregate market and the individual state, which principally determines the term-structure of dividends.

For each industry, we estimate two regression models to predict its beta from observable quantities. The first captures the cyclical component through the relation with the risk-free

[8]Due to certain non-linearities, market returns do not perfectly correlate with the SDF. Nevertheless, deviations from perfect correlation are insignificant and, hence, the conditional CAPM model predicts almost exactly the model implied expected returns.

rate and the market log price-to-dividend ratio. The second captures the dependence on the individual state. We use these regressions to generate two predictions of the beta of each industry at quarterly frequency. We then predict the return of each industry at each point in the sample using the predicted betas and the realized market return. Finally, we average the conditional expected returns predicted by the model, to generate average returns for the industry portfolios over the sample.

We plot the model implied average industry returns against their sample market betas in Figure 1. The blue circles refer to the cyclical model. The solid blue line shows the linear fit, while the dashed line shows the linear fit of the average returns predicted by the full model. As we see, the model that considers only the cyclical component predicts a flatter relation between average industry returns and unconditional market betas compared to the static CAPM model, for the specific sample. The predicted relation is close to the empirical relation. For example, for a stock whose unconditional beta is around 0.6 (1.5) the model predicts a higher (lower) average return compared to the static CAPM, for this sample, of about 0.45% (0.35%).[9]

To highlight how the model generates these results and following Jagannathan and Wang (1996), we express the unconditional average return of a stock whose conditional market beta varies over time. According to the conditional CAPM model, an asset's conditional expected excess return is given by $E_t(R^i_{t+1} - R^f_t) = \beta^i_t \; E_t(R^m_{t+1} - R^f_t)$. After applying unconditional expectation and rearranging, we obtain

$$E(R^i_{t+1} - R^f_t) = \beta^i_{unc} \cdot E(R^m_{t+1} - R^f_t) + \left[E(\beta^i_t) - \beta^i_{unc}\right] \cdot E(R^m_{t+1} - R^f_t) + Cov(\beta^i_t, R^m_{t+1} - R^f_t).$$

Here, the first term corresponds to the expected return from the static CAPM predicted, where β^i_{unc} is the unconditional market beta. The second term measures the deviation from the static CAPM that arises when the average conditional beta differs from the unconditional beta. Covariations between the conditional beta and the expected market return generate the third

[9]The full model that also contains the term-structure of cash-flows component does not offer an improvement. We believe that this is due to secular industry trends in the sample that the model is unable to capture, since it assumes stationary dynamics. Specifically, the model assumes that there are no structural shocks and that the unconditional average dividend shares are equal to the sample averages.

term. We find that each of the two terms contribute about half to the improvement relative to the static CAPM.

We represent the sample estimates of these two terms in Figure 2 for the betas predicted by the cyclical model. The left panel plots the sample estimate of the third term, i.e., the covariance of $\hat{\beta}^i_{cyc}$ with the following period market excess returns against the unconditional betas. The right panel plots the sample estimate of the second term, i.e., the deviation of the average model predicted betas $E(\hat{\beta}^i_{cyc})$ from the unconditional betas, scaled by the average market excess return.

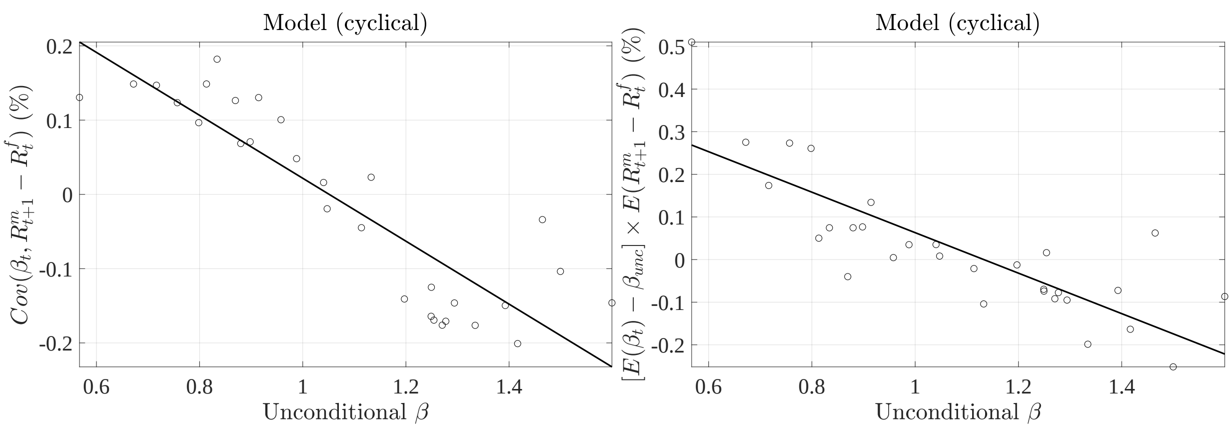

This figure decomposes the difference between the average excess returns of the 30 industry portfolios predicted by the static CAPM model and the average excess returns predicted by the cyclical component of the model (MC-I:$\hat{\beta}_{cyc}$). The left panel displays the covariance between the conditional beta and the market excess returns, while the right panel exhibits the difference between the average conditional beta and the unconditional beta, scaled by the average market excess return. Both graphs depict these measures in relation to the corresponding unconditional betas. The data, spanning from 1932Q1 to 2020Q4, are quarterly.

Figure 2: Decomposition of Security Market Line deviations – 30 industry portfolios.

As we see, for a stock whose unconditional beta is 0.6, the model predicts close to 0.20% higher average return for this sample, compared to the static CAPM, due to the positive covariation of its beta with the market excess returns. The other 0.25% of higher average return for this sample predicted by the model comes from the fact that the average conditional beta predicted by the cyclical model is higher than the unconditional beta. For a stock whose unconditional beta is 1.5 these numbers are about −0.18% and −0.17% that make up the 0.35% lower average return predicted by the model for this sample, compared to the static model.

3 Model for the aggregate stock market

We build a model of β that is consistent with the conditional CAPM in that one shock drives the stochastic discount factor (SDF), the riskless rate of return and aggregate dividends.

3.1 The SDF and aggregate dividends

In the model economy, a unique stochastic discount factor, M, prices all assets:

$$M_{t+1} = e^{-r_t^f - \frac{1}{2}(\lambda_t)^2 - \lambda_t \epsilon_{t+1}}, \tag{1}$$

where r^f stands for the one-period riskless rate of return (log), λ denotes the price of risk, ϵ is a one-dimensional standard Normal shock capturing short-run macro-movements in the economy and time runs from 0 to ∞. The riskless rate of return and the market price of risk are

$$\begin{aligned} r_{t+1}^f &= \phi_r r_t^f + (1 - \phi_r)\,\overline{r} + \sigma_r \epsilon_{t+1}, & (2)\\ \lambda_t &= \overline{\lambda} e^{-\sigma_\lambda x_t - q_\lambda x_t^2}, & (3) \end{aligned}$$

where x captures persistent business cycle movements in the economy through

$$x_{t+1} = \phi_x x_t + \epsilon_{t+1}. \tag{4}$$

Short-run macro-movements and persistent business cycle movements drive aggregate dividends, D, leading to an evolution according to

$$\ln\left(\frac{D_{t+1}}{D_t}\right) = \overline{\mu}_d + \sigma_{\mu_d} x_t + \sigma_d \epsilon_{t+1}. \tag{5}$$

Therefore, our model for the aggregate stock market includes the most common sources of stock price variations, namely fluctuations in the risk-free rate, the price of risk and cash-flow expectations.

Following Lettau and Wachter (2007), we analyze stock prices and their relation to conditional risk premia in terms of zero-coupon equities. Specifically, the stock market price-to-dividend ratio, which we denote with Q, is given by,

$$Q_t = \sum_{\tau=1}^{\infty} Q_t(\tau), \qquad Q_t(\tau) = E_t\left[\prod_{s=1}^{\tau} M_{t+s}\frac{D_{t+\tau}}{D_t}\right], \tag{6}$$

where $Q_t(\tau)$ is the price-to-dividend ratio of the zero-coupon equity that pays the aggregate dividend in τ periods.

3.2 Data and calibration

In our framework, the conditional CAPM holds almost exactly as there is only one macroeconomic shock.[10] Yet, as we show below, our model for the aggregate stock market matches the main moments of the stock market price-to-dividend ratio, the risk-free rate, the stock market returns, the predictability of stock market excess returns and the variance decomposition of the stock market price-to-dividend ratio.

We work with quarterly data that cover the period from 1926Q4 to 2020Q4. We impute the log price-to-dividend ratio $(p - d)$ and the dividends of the aggregate stock market from the CRSP stock returns with and without dividends. To minimize any effects of seasonality in dividends, we compute the $p - d$ ratio by normalizing the price of the stock market portfolio at the end of a quarter by the dividends of the last four quarters. Accordingly, the log dividend growth is computed as the quarterly growth rate in the running four quarter aggregate dividend. The stock market return (r^m) is taken to be the return of the CRSP stock index, and the risk-free rate is taken to be the yield to the 3-month Treasure bill, adjusted for expected inflation. The expected inflation is estimated by fitting an AR(1) process, and the inflation is computed as the growth rate in the CPI index.

Table 1 shows the ten parameters for two model calibrations (MC-I and MC-II). For MC-I,

[10]The conditional CAPM does not hold exactly because returns are not exactly normally distributed. However, since these deviations are negligible, there is an almost perfect linear relation between conditional expected returns and conditional betas.

we choose $\overline{r}$, ϕ_r and σ_r to fit the mean, volatility and first-lag autocorrelation of the risk-free rate in the data and set the unconditional mean of the log aggregate dividend growth ($\overline{\mu}_d$) to the corresponding sample mean in the data. To fit the autocorrelation of $p-d$, we employ the business cycle variable (ϕ_x). Further, to minimize the distance between the model and the data,[11] we set the two parameters of the aggregate dividend growth process (σ_d and σ_{μ_d}) and the three parameters that determine the price of risk ($\overline{\lambda}$, σ_λ and q_λ). In the estimations, we allow σ_d to be anywhere from 4% to 10%.

In the estimated model MC-I, the term-structure of expected excess returns is upward sloping. However, recent evidence, as in by Binsbergen, Brandt, and Koijen (2012) and Van Binsbergen and Koijen (2017), suggests that the contrary may be true. For this reason, we present calibration MC-II that produces a downward sloping term-structure of equity risk premia, through alternative values for σ_r, σ_d, σ_{μ_d}, $\overline{\lambda}$, σ_λ and q_λ. To do so and similar to Lettau and Wachter (2007), we force the price of risk to respond positively to aggregate shocks by allowing σ_λ to take only negative values. That is, we impose the price of risk in MC-II to be pro-cyclical whereas the unconstrained estimation yields a counter-cyclical price of risk in MC-I. Consequently, in MC-II longer duration cash-flows whose prices are more sensitive to discount rate shocks have lower risk premia because discount rate shocks have negative risk premia.

We gather the resulting asset pricing moments in Table 2, the predictive regressions in Table 3 and the Campbell-Shiller decomposition in Table 4, where we estimate model moments by running 1000 simulations of 400 quarters each, with a burn-in of 100 quarters. The t-statistics ($t-st.$) measure the distance between the model moments and the data statistics, which we compute using the standard errors of the data estimates and the standard deviation of the model moments across the 1000 simulations.

From Table 2, we see that both calibrations of the model match the main asset pricing

[11] The minimization is with respect to the following moments: the mean, volatility and skewness of $p-d$; the mean, volatility and first-lag autocorrelation of the stock market returns and excess returns; the stock market Sharpe ratio (SR); the predictability of the cumulative excess stock market returns by the $p-d$ ratio over horizons from 1 quarter up to 28 quarters and the Campbell-Shiller decomposition of $p-d$ into variations in future 60-quarter cumulative dividend growth, excess return, and risk-free rate and variations in $p-d$ 60 quarters ahead.

Table 1: Calibrated model parameters

Description	Parameter	MC-I	MC-II
Unconditional mean of aggregate log dividend growth	$\overline{\mu}_d$	0.435%	0.435%
Volatility of aggregate dividend growth shocks	σ_d	4.0%	10.0%
Volatility of shocks to aggregate conditional mean	σ_{μ_d}	0.019%	0.002%
Persistence of business cycle variable x	ϕ_x	0.985	0.985
Unconditional mean of one-period riskless rate of return	$\overline{r}$	0.125%	0.125%
Persistence of one-period riskless rate of return	ϕ_r	0.69	0.69
Exposure of one-period riskless rate of return to aggregate shocks	σ_r	-0.6%	0.6%
Price of risk at steady state	$\overline{\lambda}$	0.157	0.250
Exposure of log price of risk to $-x$	σ_λ	0.086	-0.067
Exposure of log price of risk to $-x^2$	q_λ	0.00192	0.00224

The table shows the parameters of the two model calibrations, MC-I and MC-II. The model is calibrated to quarterly data, where the sample is from 1926Q4 to 2020Q4.

moments, where the t-statistics for all targeted moments are below one, except for the stock market volatility that MC-II cannot match. Specifically, the volatility of the log price-to-dividend ratio for the stock market is too low compared to the data, which also results in the stock market return volatility to be counterfactually low. Based on these results, we cannot reject the null hypothesis that the data estimates are generated from model MC-I. Further, MC-II offers an important alternative calibration as it generates a high equity premium.

In Table 2, we also show statistics that were not targeted. One of them is the volatility of the annual stock market dividend growth. It creates certain tensions in the model: In MC-I the volatility of dividend growth is 2.02%, which is low relative to the data estimate of 3.31%. The opposite is true for MC-II with a volatility of 5.01%. A higher volatility in MC-I would generate even higher volatility for the stock market return, whereas a lower volatility in MC-II would result in a lower volatility for the stock market log price-to-dividend ratio. The table also shows the correlation between the one-period riskless rate of return and the stock market $p-d$ ratio, which in the data is close to zero. In MC-I the correlation is -0.473 and in MC-II -0.058, where for the latter we do not reject the null hypothesis that the data estimates are generated from model averages. Neither MC-I nor MC-II fit the correlation between aggregate dividend growth and changes in the $p-d$ ratio and the volatility and autocorrelation of the running four-quarter log aggregate dividend growth, although MC-II does get the sign of the

Table 2: Model versus data

	Data	MC-I	$t-st.$	MC-II	$t-st.$
			Targeted moments		
$\mu(p-d)$	3.419	3.370	(0.565)	3.459	(0.457)
$\sigma(p-d)$	0.447	0.407	(0.865)	0.179	(5.802)
$ac_1(p-d)$	0.965	0.957	(0.660)	0.971	(0.440)
$skew(p-d)$	0.143	0.095	(0.178)	0.001	(0.529)
$\mu(r^f)$	0.124	0.124	(0.000)	0.126	(0.020)
$\sigma(r^f)$	0.820	0.829	(0.100)	0.829	(0.100)
$ac_1(r^f)$	0.691	0.690	(0.026)	0.690	(0.026)
$\mu(r^m)$	1.644	1.368	(0.619)	1.243	(0.899)
$\sigma(r^m)$	10.521	12.582	(1.447)	5.074	(3.823)
$ac_1(r^m)$	-0.057	-0.049	(0.091)	0.164	(2.495)
$\mu(r^m-r^f)$	1.520	1.245	(0.577)	1.117	(0.844)
$\sigma(r^m-r^f)$	10.514	12.534	(1.352)	4.978	(3.706)
$ac_1(r^m-r^f)$	-0.044	-0.007	(0.457)	0.021	(0.810)
SR	0.145	0.099	(0.781)	0.224	(1.374)
$\mu(\Delta d)$	0.435	0.439	(0.014)	0.442	(0.027)
			Untargeted moments		
$\rho(\Delta d, \Delta p-d)$	-0.260	0.292	(6.973)	-0.427	(2.112)
$\rho(p-d, r^f)$	-0.003	-0.516	(3.320)	-0.052	(0.314)
$\sigma(\Delta d)$	3.306	2.018	(2.424)	5.011	(3.209)
$ac_1(\Delta d)$	0.543	0.755	(3.308)	0.747	(3.195)

The first column shows the variable name. The second column shows the data estimate. The third and fifth columns show the model estimates for model calibration 1 (MC-I) and model calibration 2 (MC-II). The fourth and sixth columns show the t-statistics ($t-st.$) of the hypotheses that the data estimates are generated from model averages. $p-d$ denotes the log price-to-dividend ratio, r^f stands for the one-period riskless rate of return, r^m is the return on the aggregate stock market, SR stands for the Sharpe ratio and Δd is the log aggregate dividend growth rate. The data from the model are time-aggregated. The data are quarterly: $1926Q4-2020Q4$.

correlation right. Overall we conclude here that both calibrations have their merits.

Table 3 shows that the model captures the predictability of stock market excess returns by the $p-d$ ratio. We inspect the predictability for horizons up to 28 quarters ahead and see that all the t-statistics measuring the difference between the model and the data for horizons up to 28 quarters are below 1.4. Therefore, the predictability regressions from both calibrations of the model relative to the data are statistically indistinguishable.

Table 4 shows the Campbell-Shiller decomposition of variations in the $p-d$ ratio. Variations in future excess returns are by far the most important driver of $p-d$, as indicated by the

Table 3: Predictive regressions: model versus data

Horizon	$Q=1$	$Q=2$	$Q=4$	$Q=8$	$Q=12$	$Q=16$	$Q=20$	$Q=24$	$Q=28$
Data	-0.082	-0.112	-0.172	-0.241	-0.272	-0.302	-0.350	-0.400	-0.434
MC-I	-0.066	-0.093	-0.131	-0.183	-0.221	-0.251	-0.276	-0.298	-0.318
$t-st.$	(0.289)	(0.243)	(0.415)	(0.490)	(0.392)	(0.360)	(0.517)	(0.772)	(0.918)
MC-II	-0.076	-0.106	-0.144	-0.189	-0.216	-0.234	-0.246	-0.254	-0.260
$t-st.$	(0.102)	(0.085)	(0.287)	(0.440)	(0.428)	(0.480)	(0.731)	(1.108)	(1.374)

The second (data), third (model calibration 1) and fifth (model calibration 2) rows show the regression coefficient from predictive regressions where excess returns are predicted by the log price-to-dividend ratio for horizons Q = 1, 2, 4, 8, 12, 16, 20, 24, 28. The forth and sixth rows show t-statistics ($t-st.$) of the hypotheses that the data coefficient estimates are generated from model averages. The data are quarterly: $1926Q4-2020Q4$.

Table 4: Campbell-Shiller decomposition: model versus data

	d	r^m	r^m-r^f	r^f	$(p-d)_{60}$
Data	0.154	0.605	0.657	-0.051	0.206
MC-I	0.064	0.701	0.676	0.025	0.207
$t-st.$	(0.850)	(0.689)	(0.095)	(0.712)	(0.006)
MC-II	0.095	0.637	0.650	-0.013	0.219
$t-st.$	(0.551)	(0.225)	(0.033)	(0.356)	(0.244)

The second (data), third (model calibration 1) and fifth (model calibration 2) rows show the regression coefficient from a Campbell-Shiller decomposition with variations in aggregate dividend cash flows (d), market returns (r^m), excess returns (r^m-r^f), riskless rate of return (r^f) and residual variation ($(p-d)_{60}$) with the log price-to-dividend ratio using 60 quarters. The forth and sixth rows show t-statistics ($t-st.$) of the hypotheses that the data coefficient estimates are generated from model averages. The data are quarterly: $1926Q4-2020Q4$.

results obtained from the data, where they account for about 65% of the volatility in $p-d$. The variations associated with future risk-free rate fluctuations in the data are slightly positive but insignificant (2%). In the model, the risk-free rate contributes small negative variations in $p-d$ (-7.6%) that are also insignificant, which supports our view that modeling more accurately the dynamics of the riskless rate of return will not alter our results for the cross-section. Thus, the variations in excess returns in the model account for about 72% of the variations in $p-d$, which is a bit higher than what we see in the data. Yet, the differences between the model and the data with respect to the variations in $p-d$ coming from r^m-r^f and r^f are statistically insignificant with t-statistics below one. Lastly, the model matches well the remaining variations in the

$p-d$ ratio that are associated with future cash-flows (g) and the $p-d$ ratio 60 quarters ahead ($p-d_{60}$), with t-statistics of 0.161 and 0.356, respectively.

4 Model for the cross-section

Armed with a model for the aggregate stock market, we analyze a cross-section of returns by modeling cash-flows. To understand persistent differences in β's across stocks, e.g. industries, the stocks are ex-ante heterogeneous through stationary dividend shares that neither vanish nor take over the aggregate dividends.

4.1 Dividend shares

There are I stocks. Individual stock dividends represent shares of aggregate dividends

$$\theta_t^i \coloneqq \frac{D_t^i}{D_t}, \qquad i = 1, \ldots, I, \tag{7}$$

where $\sum_{i=1}^{i=I} \theta^i = 1$. The dividend shares have a rich structure in that each dividend share has an idiosyncratic component, denoted with y_t^i, and a systematic component that depends on the business cycle variable x, as follows:

$$\theta_t^i = y_t^i \cdot \left[1 + (\eta_i - a_{1,t})\, f(x_t)\right]. \tag{8}$$

Consider for now that $a_{1,t}$ equals zero, which is a variable that ensures that dividend shares sum up to one. Here, the parameter η_i determines the cyclicality of a stock. We allow it to take any value between -1 and $+1$, where a positive (negative) value implies a pro-cyclical (counter-cyclical) stock. Consequently, $f(x)$, which is a monotonically increasing function, is also bounded by -1 and $+1$. We also require that $f(0) = 0$, and $f'(x) > 0$. A natural choice is $f(x) = \Phi(x; 0, \nu^2) - 1$, where $\Phi(\cdot; \mu, \sigma^2)$ is the normal cumulative distribution function with mean μ and standard deviation σ; and the parameter ν determines the sensitivity of a stock to the business cycle and how this varies over the business cycle.

The idiosyncratic components y_t^i sum up to one for each time period, that is, $\sum_{i=1}^I y_t^i = 1$, whle the cyclical components add up to zero. For this reason, we set $a_{1,t} = \frac{1}{I}\sum_{i=1}^I y_t^i \eta_i$, which equals zero at the steady state when $y_t^i = \overline{y}_0^i$ for all i. It is important that $a_{1,t}$ also has an expected value of zero at every other state of the idiosyncratic components of the cross section. Moreover, as the number of stocks approaches infinity, $a_{1,t}$ converges to zero almost surely.

When the aggregate economy is at its steady state ($x = 0$), the dividend shares are given by y_t^i. When the economy is growing at a faster rate than average such as during "booms," then a pro-cyclical ($\eta_i > 0$) stock's dividend share will be higher compared to the dividend share that the stock would have had if the economy were in its steady state. Similarly, counter-cyclical stocks "see" their dividend shares shrink during good times and expand during downturns. The formulation of the cyclical component of the dividend shares requires that $y_t^i \ll 0.5$, since the maximum number that the square bracket in (8) can take is 2.

The idiosyncratic component of a dividend share mean reverts around a time-varying long-run mean as follows:

$$y_{t+1}^i = \left[\phi_i y_t^i + (1-\phi_i)\,\overline{y}_t^i\right] a_{2,t} + \sigma_i y_t^i \cdot \left(l_y - y_t^i\right)\left(\epsilon_{t+1}^i - u_{t+1}\right), \tag{9}$$

where $\phi_i \in [0, 1)$ and each ϵ^i is an independent one-dimensional standard Normal shock. The parameter $l_y < 0.5$ denotes the upper limit of y. The variable $a_{2,t}$ ensures that the sum of the conditional means equals one every period; thus, $a_{2,t} = \left[1 + \sum_{i=1}^I \phi_i \cdot (y_t^i - \overline{y}_t^i)\right]^{-1}$. If we assume $\phi_i = \phi$ for all stocks, then $a_2 = 1$. Similarly, u is a weighted average of the individual shocks to ensure that the sum of the shocks equal zero; thus,

$$u_{t+1} = \sum_{i=1}^I \frac{\sigma_i y_t^i \cdot (l_y - y_t^i)}{\sum_{j=1}^I \sigma_j y_t^j \cdot (l_y - y_t^j)} \epsilon_{t+1}^i.$$

According to equation (9), the conditional volatility of y_{t+1}^i depends on the idiosyncratic dividend share y_t^i and tends to zero as y_t^i tends to 0 or l_y. This prevents any stock from vanishing, which is important for maintaining a sufficiently large cross-section.[12]

The long-run mean of a dividend share also fluctuates over time according to

$$\overline{y}^i_{t+1} = \varphi \overline{y}^i_t + (1-\varphi)\overline{y}^i_0 + \overline{\sigma}_i \overline{y}^i_t \cdot \left(l_y - \overline{y}^i_t\right)\left(\overline{\epsilon}^i_{t+1} - \overline{u}_{t+1}\right), \tag{10}$$

where $\varphi \in [0,1)$ and each $\overline{\epsilon}^i$ is an independent one-dimensional standard Normal shock. In the same fashion as u, $\overline{u}$ is defined as a weighted average of $\overline{\epsilon}^i$'s, so that the shocks to the long-run means sum up to zero, that is,

$$\overline{u}_{t+1} = \sum_{i=1}^{I} \frac{\overline{\sigma}_i \overline{y}^i_t \cdot (l_y - \overline{y}^i_t)}{\sum_{j=1}^{I} \overline{\sigma}_j \overline{y}^j_t \cdot (l_y - \overline{y}^j_t)} \overline{\epsilon}^i_{t+1}.$$

Our framework can accommodate long-run means $\overline{y}$ that follow a random walk ($\varphi = 1$), and in that case we would not need to specify $\overline{y}^i_0$. In our view, permanent shocks to the long-run means capture structural shocks to the economy, which could be a significant source of fluctuations in stock prices. However, then the boundaries would be absorbing states. This implies that we would need to model entry by new firms. For this reason, we restrict the analysis to cases where φ is close to but lower than one.

The aggregating variables $a_{1,t}$, $a_{2,t}$, u_t, and $\overline{u}_t$ always take their long-run values of either 0 or 1 when I tends to infinity. For our finite cross-section, we adopt an approach similar to Menzly, Santos, and Veronesi (2004) and assume that these variables always take their long-run values. Further, since the model is in discrete time, in simulations we impose the upper and lower bounds.

4.2 Prices and expected returns

Before we fit the model to the data, we analyze how a stock's price-to-dividend ratio, conditional beta, and expected return behave. Let $y^i_t(\tau) = E_t(y^i_{t+\tau})$ denote the expected dividend share τ

[12] The analysis could be adapted to include the possibility of a company being driven out of the market. Obviously, such an extension requires that new firms enter every period.

periods ahead. Given equations (9) and (10), we obtain

$$y_t^i(\tau) = \bar{y}_0^i + \phi_i^\tau \cdot (y_t^i - \bar{y}_0^i) + \frac{1-\phi_i}{\varphi - \phi_i}(\varphi^\tau - \phi_i^\tau)(\bar{y}_t^i - \bar{y}_0^i). \tag{11}$$

In the special case where $\phi_i = \varphi$, we have

$$y_t^i(\tau) = \bar{y}_0^i + \varphi^\tau \cdot (y_t^i - \bar{y}_0^i) + \tau\varphi^{\tau-1}(1-\varphi)(\bar{y}_t^i - \bar{y}_0^i). \tag{12}$$

Using the expression for $y_t^i(\tau)$, we derive an expression of a stock's price-to-dividend ratio Q_t^i, after expressing it as a sum of zero-coupon claims. This yields the following expression

$$Q_t^i = \sum_{\tau=1}^{\infty} \frac{y_t^i(\tau)}{\theta_t^i} \left[Q_t(\tau) + \eta_i Q_t^x(\tau) \right], \tag{13}$$

where $Q_t^x(\tau)$ represents the price-to-dividend ratio of the claim to the cash-flow τ periods ahead, that has unit exposure to the cyclical component $f(x)$. That is,

$$Q_t^x(\tau) = E_t \left[\prod_{s=1}^{\tau} M_{t+s} f(x_{t+\tau}) \right].$$

Similarly, we express a stock's conditional beta as a weighted average of the conditional betas of the zero-coupon equities as follows

$$\beta_t^i = \sum_{\tau=1}^{\infty} \frac{y_t^i(\tau)}{\theta_t^i} \cdot \frac{\beta_t(\tau) Q_t(\tau) + \eta_i \beta_t^x(\tau) Q_t^x(\tau)}{Q_t^i}. \tag{14}$$

The expressions in (13) and (14) serve to analyze expected returns in the cross-section and to expose their relation to prices.

The following proposition establishes a useful benchmark case.

Proposition 1. *When dividend shares of stocks follow a random walk, i.e., $y_t^i(\tau) = \theta_t^i, \forall\tau$, and exhibit no cyclicality, i.e., $\eta_i = 0$, then $\beta_t^i = 1$ and $Q_t^i = Q_t$.*

Proposition 1 implies that the conditional beta of a stock deviates from one because of two reasons. The first is cyclicality, which is determined by η_i. Since we assume η_i to be constant, the cyclicality of a stock implies that the average conditional beta will also differ from one. The second is the cash-flow term-structure. For example, deviations depend on whether the dividend share is expected to increase or decrease over time. We analyze these two features next.

4.3 Comparative static analysis of one stock

Naturally, a pro-cyclical (counter-cyclical) stock has an average conditional beta higher (lower) than one. However, the conditional beta depends on the state of the economy within the business cycle. To analyze how the conditional betas vary over the business cycle, we need to parametrize the function $f(x)$ that determines cash-flow cyclicality. In the analysis, we assume that $\nu = 3.5$ to produce a sufficiently large cross-section of unconditional betas.[13] Specifically, for this value the unconditional betas vary from slightly below 0.6 to close to 1.8 for model MC-I, which is about the same range as we see in the data for the 30 industry portfolios. The range for MC-II is from around 0.7 to around 2.0.

Figure 3 shows how equilibrium quantities vary over the business cycle in model MC-I. Low (high) values for x indicate economic contractions (expansions) and the gray area shows the distribution. In Panel A, we plot the log price-to-dividend ratio ($p-d$ ratio) and in Panel B the conditional beta for the market ($\eta = 0$), a pro-cyclical stock ($\eta = 0.75$) and a counter-cyclical stock ($\eta = -0.75$). In Panel A, we observe that the market's and the counter-cyclical stock's $p-d$ ratios monotonically increase with the business cycle variable x, that is, they increase when the state of the economy improves. The average $p-d$ ratio of the counter-cyclical stock is the highest because it offers some insurance to negative economic shocks. However, the $p-d$ ratio of the pro-cyclical stock does not vary monotonically with the business cycle. In an area close to the mean, the $p-d$ ratio decreases with x, despite of the decrease in the equity premium (see Panel C). The reason for this behavior is principally the way in which the conditional beta changes with economic shocks. At the local maximum in the area between $x = -10$ and $x = 0$,

[13] Figure 13 plots the function $f(x)$ for various values of ν along with the distribution of x.

the conditional beta is effectively constant. However, to the right of this local maximum, a positive economic shock leads to a decrease in the conditional beta, which implies a further increase in the stock price, and the opposite is true for a negative economic shock. Therefore, to the right of the local maximum the stock becomes more risky and the $p-d$ ratio decreases. To the right of $x = 0$, the $p-d$ ratio of the pro-cyclical stock increases again because the increase in the riskiness of the stock is lower while the equity premium continues to decrease.

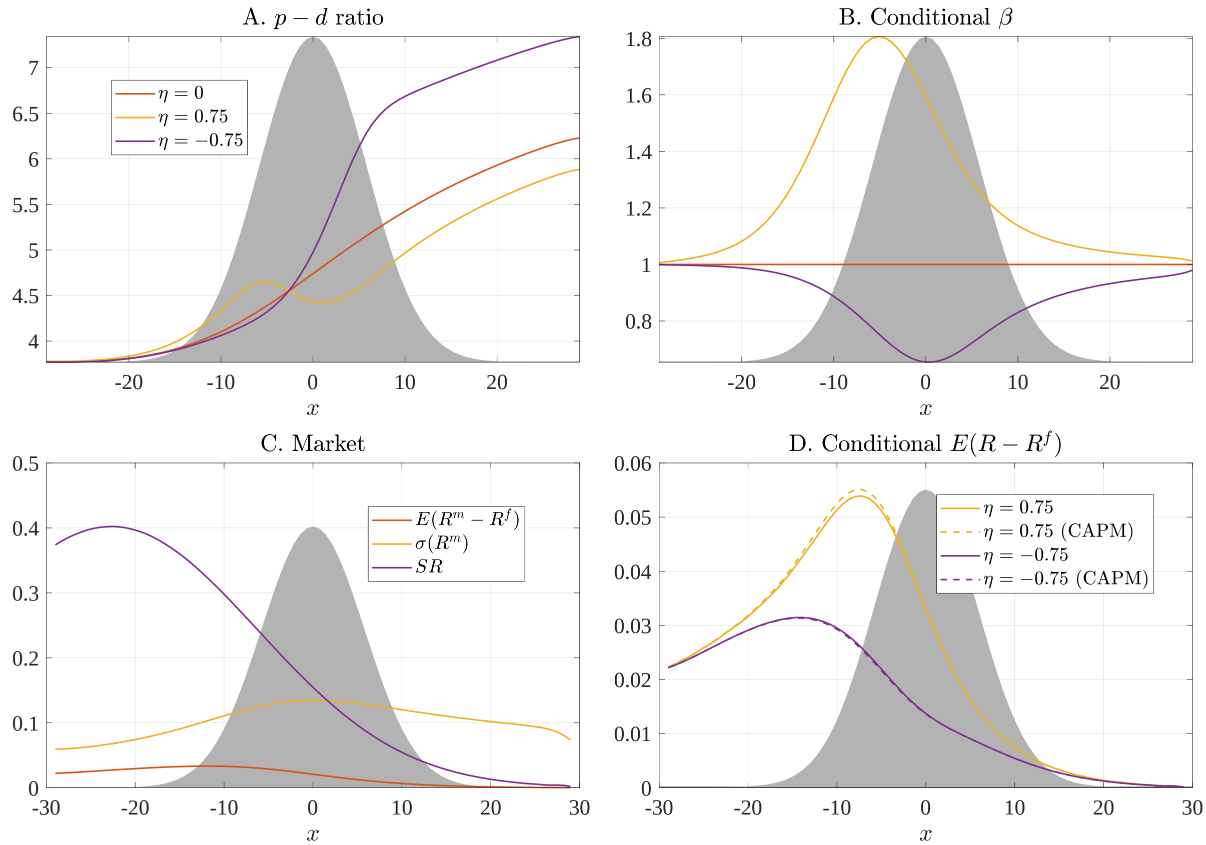

The figure contains various asset pricing quantities against the business cycle variable x. Panel A shows the log price-to-dividend ratio and Panel B the conditional β for the market ($\eta = 0$), for a pro-cyclical stock ($\eta = 0.75$), and for a counter-cyclical stock ($\eta = -0.75$). Panel C shows the equity premium, the conditional volatility of market returns and the market portfolio Sharpe ratio. In Panel D, the solid lines represents the expected excess returns of a pro-cyclical and counter-cyclical stock. The dashed lines show the expected excess returns given by the conditional CAPM model.

Figure 3: Asset price quantities over the business cycle (MC-I)

In Panel B, we see that the conditional betas of the cyclical stocks vary considerably over the business cycle, where the largest dispersion occurs in normal times around $x = 0$. This is principally due to the fact that the exposure to the business cycle, which is given by $f'(x)$ is highest around the mean. We also observe that the conditional betas are not monotonic, where the cross-section shrinks significantly both during contractions and during expansions.

In Panel C, we plot the equity premium, the market Sharpe ratio, and the conditional volatility of the market returns. The equity premium and the Sharpe ratio are largely counter-cyclical, while the conditional volatility shows only a small (and asynchronous) counter-cyclicality.

In Panel D, we plot the risk premia of the pro-cyclical and counter-cyclical stocks, which results from the behavior of the conditional betas and the equity premium. From it, we see that the pro-cyclical stock always demands a higher conditional risk premium; the spread however shrinks both during expansions and during contractions. Around $x = -5$, we observe the maximum spread, which is a bit lower than the unconditional mean. If we consider the $p - d$ ratio as a proxy for the market-to-book ratio, then the model predicts a value premium during normal times and during expansions, and a value discount during contractions. Specifically, for x higher than around -3, low $p - d$ ratio (high book-to-market) stocks require a higher expected return, while the opposite is true over the rest of the x domain. This is consistent with the empirical observation that the value strategy performs well during normal times and during economic expansions but under-performs during recessions.

In Figure 10 in Appendix A, we plot the same quantities as in Figure 3 with respect to the risk-free rate. From that figure we learn that the risk-free rate exerts minimal effects on the conditional betas. Moreover, in Figure 11 in Appendix A, we show the same plots as in Figure 3 for the model calibration MC-II. For this alternative model calibration, the $p - d$ ratios are mostly decreasing in the business cycle variable because the market Sharpe ratio and the equity premium are largely pro-cyclical. Otherwise, the conditional betas and the value premium behave in similar fashions as in model MC-I.

Next, we inspect how the individual state of a stock affects its conditional beta and $p - d$ ratio. We plot these quantities for model MC-I in Figure 4 against the short-run state y and the long-run state $\overline{y}$. When we vary the short-run state, we keep the long-run state at the unconditional mean, while when we vary the long-run state, we set the short-run state equal to

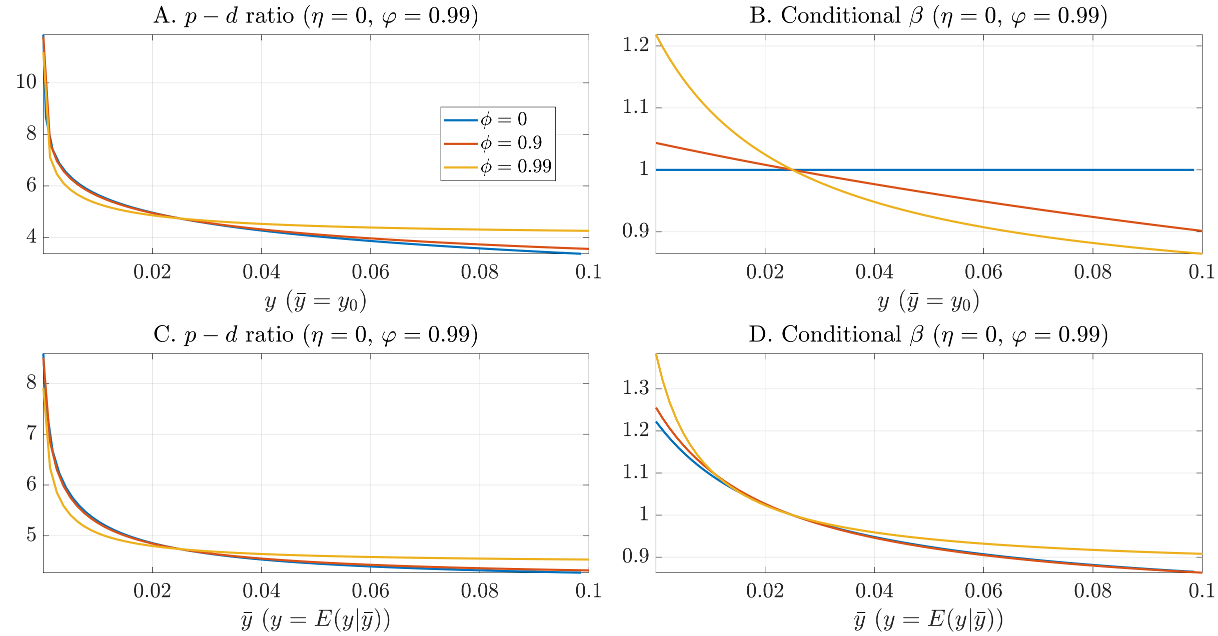

This figure shows for model MC-I the log price-to-dividend ratio ($p-d$) and the conditional β over the individual state $(y, \overline{y})$, for a non-cyclical stock ($\eta = 0$), for various values of φ. We set the aggregate state variables (x, r^f) at their steady state values. Panels A and B plot these quantities against y where we set $\overline{y}$ equal to the unconditional mean, and Panels C and D show the quantities against $\overline{y}$ where we set y equal to its conditional expectation.

Figure 4: Stock $p-d$ and market β in terms of y and $\overline{y}$ (MC-I)

its conditional mean. We show these plots for a non-cyclical stock, for $\varphi = 0.99$, and for three different values for the persistency of the short-run shocks ϕ: 0, 0.9, and 0.99. In Panels A and C, we plot the $p-d$ ratio and we observe that its behavior depends little on the persistency parameter. In contrast, the conditional beta in Panels B and D, are significantly affected by persistent shocks. Also, we see that the individual state generates a value discount in the sense that a relatively high $p-d$ ratio (low book-to-market) is associated with a high expected return. As in Santos and Veronesi (2010), the counter-cyclicality of the equity premium makes longer duration cash-flows more risky than shorter duration cash-flows. As a result, growth stocks are more risky than value stocks. In Figure 12 in Appendix A, we see that the model calibration MC-II generates a value premium due to the downward sloping term-structure of equity premia. This, however, comes at the cost of a pro-cyclical equity premium and too low stock market volatility.

5 Estimating the model for industries

5.1 Data and calibration

Our choice of industry portfolios is governed by several criteria. We need long enough data series to be able to estimate the model parameters for each portfolio. We need enough homogeneity within each portfolio so that the estimated model parameters accurately describe the cash-flow dynamics of its constituents. We also need a large enough cross-section to test the model. For our purpose, the best candidate is the set of 30 industry portfolios from Kenneth French's database.[14]

The data span the period from July 1926 to the January 2021. The original data is monthly from which we construct quarterly series to remove the monthly seasonality in the dividends and the fact that there are many months with zero dividends. We choose $\varphi = 0.99$ to match the levels of autocorrelation in the market capitalization shares as well as to generate enough idiosyncratic volatility in returns. Regarding the function $f(x)$, as already mentioned, we choose the parameter ν to obtain a cross-section of unconditional betas comparable to the data. To achieve this, we need high sensitivity of cash-flows around the steady state, which implies lower sensitivity further away from the steady state. A value of $\nu = 3.5$ allows to capture the range of unconditional betas across the 30 industries.

For each industry portfolio, we then estimate the dividend share process parameters as follows: We set $\overline{y}_0^i$ to the average dividend share in the sample. Next, we estimate the other four parameters $(\eta_i, \phi_i, \sigma_i, \overline{\sigma}_i)$ by targeting the following four moments: the unconditional market beta (β), the autocorrelation in the dividend share $(ac(\theta))$, the volatility of the dividend share $(\sigma(\theta))$, and the idiosyncratic return volatility (IV). Essentially, the unconditional beta allows us to identify the cyclicality parameter η_i. In this regard, Figure 5 shows that the unconditional beta is a good proxy for the cyclicality of an industry as measured by the correlations of the changes of dividend and capital shares with the recession indicator. Further, with ϕ_i and σ_i we target the dividend share moments, and with $\overline{\sigma}_i$ we target the idiosyncratic return volatility.

[14] Kenneth French data library.

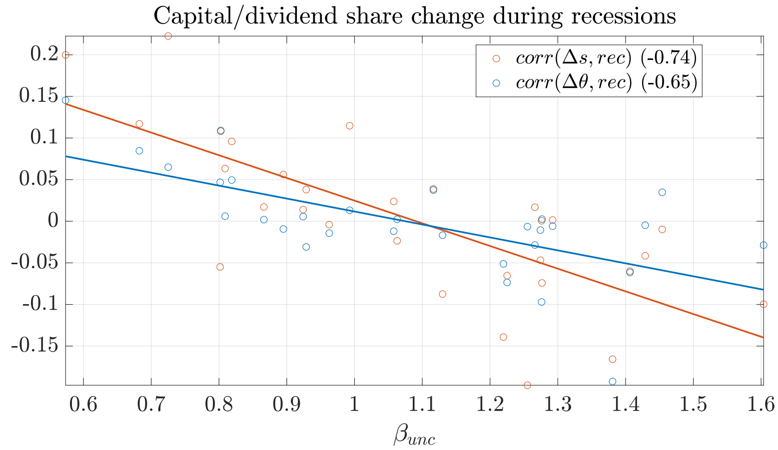

In the above figure, the x-axis represents the unconditional market betas of the 30 industry portfolios in the data. On the y-axis, we depict the correlations between the changes in capital shares (Δs) represented by red circles, and dividend shares ($\Delta\theta$) represented by blue circles, with the recession indicator (rec). The recession indicator takes the values of $[0, \frac{1}{3}, \frac{2}{3}, 1]$. The data covers the period from 1926Q4 to 2020Q4. The straight lines indicate the linear fits. The correlations between the quantities plotted on the y and x axes are shown in parentheses in the legend.

Figure 5: Recession shocks vs betas – 30 industry portfolios

In the estimation, we minimize the distance between the data and the model with respect to the above four statistics. In the model, these statistics are the simulation averages of $N = 100$ simulations of 400 quarters each, with a burn-in of 100 quarters.[15] The distance between the model and the data is given by the square root of the sum of the squared $t-$statistics. Specifically, the following formula shows the $t-$statistic with respect to a moment z:

$$t(z) \coloneqq \frac{|\mu_d(z) - \mu_m(z)|}{\sqrt{se(z)^2 + N^{-1}\sigma(z)^2}}, \tag{15}$$

where $\mu_d(z)$ corresponds to the data moment, $\mu_m(z)$ represents the model simulation mean, $se(z)$ is the data standard error, and $\sigma(z)$ represents the standard deviation of the moment across the N simulations. We compute GMM corrected standard errors with Newey and West (1987) weighting and 16 lags.

The estimation prioritizes achieving a better fit for both the unconditional beta and the idiosyncratic return volatility, placing more weight on these two statistics. Undoubtedly, the

[15]The 400 quarters match the length of the data sample.

unconditional beta holds primary significance. Simultaneously, the idiosyncratic volatility is essential to capture the variations of the term-structure of dividend shares.

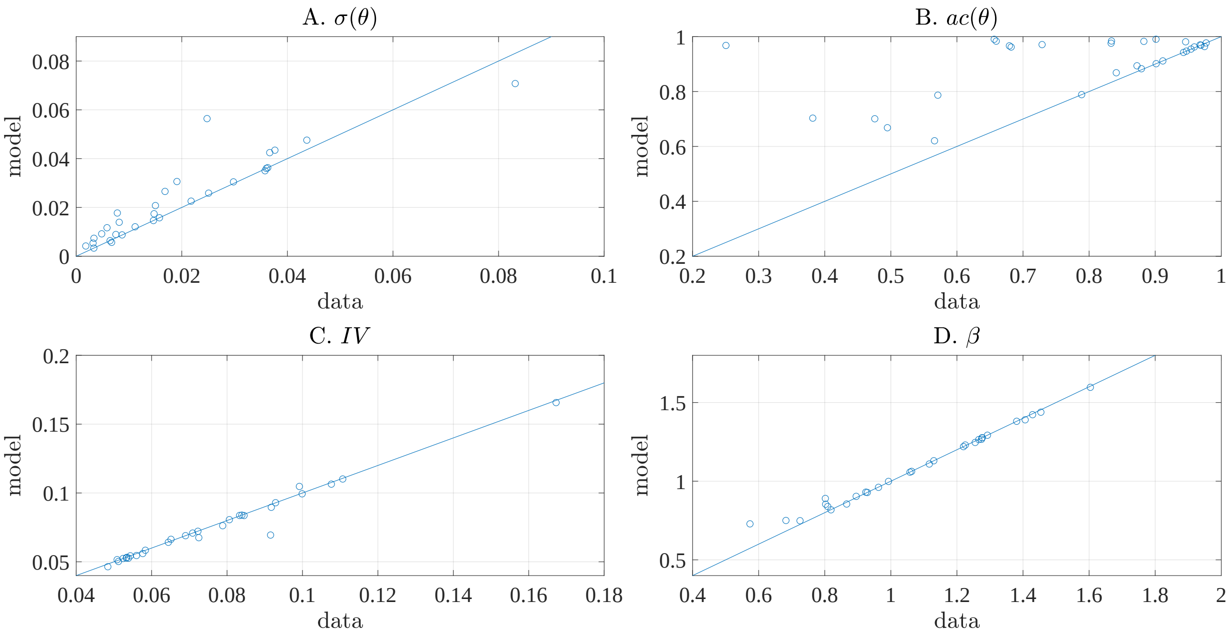

This figure shows the targeted moments regarding the estimation of the cash-flow model for each of the 30 industry portfolios, the volatility of dividend share (Panel A), the autocorrelation of the dividend share (Panel B), the static CAPM idiosyncratic return volatility (Panel C), and the unconditional β (Panel D). The $y-$axis corresponds to the model estimates and the $x-$axis to the data moments. The 45^o line is shown in blue.

Figure 6: Calibration fitted moments (MC-I) – 30 industry portfolios

In Figure 6, we compare the data estimates against those predicted by model MC-I. Panels C and D demonstrate a better fit for the idiosyncratic volatilities and unconditional betas, respectively, in comparison to the other two statistics presented in Panels A and B. This is confirmed by Panel A of Figure 14 where we show with box plots the $t-$statistics for the 30 industry portfolios. In addition, Panel B shows the box plots of the *model distance* ($md(z)$) for each statistic, where the model distance refers to the number of standard deviations that the data estimate is away from the model mean, that is,

$$md(z) \coloneqq \frac{|\mu_d(z) - \mu_m(z)|}{\sigma(z)}. \tag{16}$$

With the exception of the autocorrelation of the dividend shares, the data estimates of the remaining statistics are within 2 standard deviations from the model means. The maximum model distance for $\sigma(\theta)$, IV, and β are 1.17, 1.28, and 1.56, respectively. However, the model fails to generate low enough dividend share autocorrelations because the common autocorrelation of the long-run dividend share shocks φ is set to 0.99 and given that $\phi \geqslant 0$. Yet, low persistency shocks to the dividend shares have only small effects on conditional betas.[16]

Regarding the persistency parameter φ of the long-run shocks, Figure 15 shows that the chosen value generates autocorrelations of capital shares that are largely consistent with the data. Autocorrelations in the data are high, ranging from around 0.925 to 0.995. The model predicts for some industries relatively low autocorrelations, yet, md is slightly higher than 2 only for one industry portfolio. This implies that the data is consistent with the model.

Following Menzly, Santos, and Veronesi (2004), we assume that the aggregating variables, which ensure that the sum of dividend shares equals one, take their long-run values. Table 8 in the Appendix contains statistics of the aggregating variables across simulations. These show moderate deviations from their long-run values and, more importantly, the aggregating variables exhibit no correlations with the aggregate state variables and the aggregate shock. Consequently, the pricing implications derived from the simulations are consistent with the model.

5.2 Generating model implied conditional β's

Having estimated the dividend share process for each industry portfolio, we generate predictions for the conditional betas. There are, however, a number of challenges. First, the conditional beta is a function of latent variables, the business cycle variable x and the portfolio specific variables $(y^i, \overline{y}^i)$. Second, the function is non-linear and industry portfolio specific. Relating conditional betas to observable quantities is our approach to address these challenges.

Regarding the business cycle variable, the stock market $p - d$ ratio serves as a proxy, given that it has a monotonic, almost linear, relation with x, as shown by Panel A of Figure 3 for MC-I and Panel A of Figure 11 for MC-II. Also, the stock market $p-d$ ratio is largely insensitive

[16] Figures 16 and 17 show the corresponding plots for model MC-II.

to the other determinant, the risk-free rate, as shown by Panel A of Figure 10. To capture the industry portfolio specific state $(y^i, \overline{y}^i)$, we use the dividend share θ^i and the capital share s^i.

Regarding the functional form, it should strike a balance between flexibility in capturing the model's implied dependencies, while avoiding excessive complexity that could hinder parameter estimation. These considerations guide us to adopt the following functional form:

$$\begin{aligned}\hat{\beta}^i_{t,all} &= b^i_0 + b^i_1 \cdot \left[\theta^i_t - \mu(\theta^i)\right] + b^i_2 \cdot \left[s^i_t - \mu(s^i)\right] + b^i_3 \cdot \left[r^f_t - \mu(r^f)\right] \\ &\quad + b^i_4 \cdot [p - d_t - \mu(p - d_t)] + b^i_5 \cdot [p - d_t - \mu(p - d_t)]^2 + b^i_6 \cdot [p - d_t - \mu(p - d_t)]^3 .\end{aligned}$$

We estimate the parameters using regressions on simulated data and refer to the above as the full model, since we use both aggregate variables and industry portfolio variables to predict conditional betas.

However, structural shocks and secular industry trends over the sample period makes the reliance on the dividend and capital shares difficult when making predictions about the conditional betas. For example, it is questionable to assume that during the early period of the sample investors could predict either the growth or the decline of certain industries. For example, the market capitalization share of the transportation industry has declined from around 18% to around 2%, while the business equipment industry has grown from around 1.5% to close to 15%. Another example is the consumer goods industry which enjoyed a growth from around 0.6% to more than 10% to then decline again to levels a bit lower than 2%. Despite these secular trends, we expect the cyclical behavior of the industries to have not altered significantly over the sample period. Therefore, we also estimate a cyclical model where we predict conditional betas based on purely aggregate variables, as follows,

$$\begin{aligned}\hat{\beta}^i_{t,cyc} &= b^i_0 + b^i_1 \cdot \left[r^f_t - \mu(r^f)\right] + b^i_2 \cdot [p - d_t - \mu(p - d_t)] \\ &\quad + b^i_3 \cdot [p - d_t - \mu(p - d_t)]^2 + b^i_4 \cdot [p - d_t - \mu(p - d_t)]^3 .\end{aligned}$$

We next analyze the performance of these models and the importance of the two factors that drive the conditional betas according to the model.

5.3 Cyclicality vs term-structure

Before inferring the time-series of the conditional betas of the industry portfolios using $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$, we examine how well they capture the variations in the conditional betas in the model. Table 5 shows statistics regarding the R^2's from the estimations and variance contributions. For model MC-I the explanatory power of $\hat{\beta}_{all}$ ranges from 82.2% to 97.7% across the 30 industry portfolios with a mean of 90.9%. For model MC-II, the explanatory power ranges from 76.5% to 97.1% with a mean of 90.1%.

For $\hat{\beta}_{all}$, we also compute the contribution of the cyclical component that is determined by the aggregate state variables $p-d$ and r_f. These contributions are estimated as the covariance of the corresponding components with the true conditional betas divided by the variance of the conditional betas. The contribution of the risk-free rate is minimal, which confirms our earlier observation that the risk-free rate plays a minor role for the conditional betas. Further, the importance of the cyclical component in $\hat{\beta}_{all}$ varies considerably across industries. This is expected since the importance of the cyclical component depends on the cyclicality of a portfolio. The contribution of the cyclical component for model MC-I varies from 4.3% to 86.1% with a mean of 37.1%. Similarly, for model MC-II, the range is from 3.2% to 88.7% with a mean of 31.1%.

However, the true contributions of the cyclical component are larger than those implied by $\hat{\beta}_{all}$. This is because the portfolio specific state also depends on the business cycle variable. Since causality runs only from the business cycle variable to the portfolio specific state, the true importance is given by the explanatory power of $\hat{\beta}_{cyc}$. For MC-I (MC-II), the explanatory power ranges from 20.5% (13.8%) to 95.8% (92.3%) where the mean is 50.4% (42.4%). Yet, these estimates still slightly underestimate the true importance of the cyclical component, since $\hat{\beta}$ is only an approximation and does not account for all fluctuations. Nevertheless, on average cyclicality is at least as important for the fluctuations of the conditional beta as the term-

Table 5: R^2's and variance decomposition

	Avg.	St.dev.	Min	q25	q50	q75	Max
MC-I: $\hat{\beta}_{all}$	90.9	4.2	82.2	89.2	91.9	94.1	97.7
r_f, $p-d$	37.1	22.8	4.3	13.5	35.5	57.0	86.1
r_f	-0.7	0.6	-2.2	-1.2	-0.6	-0.2	0.0
MC-I: $\hat{\beta}_{cyc}$	50.4	22.9	20.5	28.1	48.4	73.2	95.8
MC-II: $\hat{\beta}_{all}$	90.1	5.4	76.5	87.3	90.8	94.8	97.1
r_f, $p-d$	31.1	28.3	3.2	11.8	17.5	46.1	88.7
r_f	-1.3	1.5	-4.3	-2.0	-0.6	-0.3	0.1
MC-II: $\hat{\beta}_{cyc}$	42.4	24.5	13.8	23.5	36.8	52.8	92.3

This table presents the explanatory power, in percentages, of $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$, along with individual components, in explaining the variation in the true conditional betas for the estimated cash-flow models of the 30 industry portfolios. Each line reports the average (Avg.), standard deviation (St.dev.), minimum (Min), maximum (Max), and the 25%, 50%, and 75% quantiles across these portfolios. The values presented are the averages obtained from 100 simulations. The components "r_f, $p-d$" and "r_f" correspond to the collection of terms of the full model $\hat{\beta}_{all}$ with their respective quantities. For $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$ we report the regression R^2's. For the components we report the covariances of those components with the true conditional β's divided by the variances of the true conditional β's.

structure, while in some cases cyclicality accounts for almost all fluctuations.

6 Asset pricing tests

With the models $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$ for each industry portfolio and for each of the calibrated models MC-I and MC-II, we can infer time-series of model implied conditional betas and test whether the conditional CAPM model explains returns.

We use standard time-series asset pricing test statistics. Suppose $\{\hat{\beta}^i_t, t = 1, \ldots, T\}$ is the estimated time-series of conditional betas of portfolio i. We compute the abnormal return every period,

$$\alpha^i_t = R^i_{t+1} - R^f_t - \hat{\beta}^i \cdot (R^m_{t+1} - R^f_t)$$

and then the average abnormal return (alpha) and the return residuals, via,

$$\alpha_i = \frac{1}{T}\sum_{t=1}^{T}\alpha_t^i, \qquad \varepsilon_t^i = \alpha_t^i - \alpha_i.$$

Under the null hypothesis the alpha is zero for all portfolios. The first test statistic we compute is the mean absolute alpha for the set of N test portfolios, i.e.,

$$\mu(|\alpha|) = \frac{1}{N}\sum_{i=1}^{N}|\alpha_i|. \tag{17}$$

For each portfolio, we perform a standard $t-$test to examine the significance of the estimated alpha, using GMM corrected standard errors with Newey and West (1987) weighting matrix. We then count the number of test portfolios for which the average abnormal return is significant at the 5% level,

$$\#p < 5\% = \sum_{i=1}^{N}\mathbb{1}\left[\mathcal{F}_t\left(\frac{|\alpha_i|}{se(\alpha_i)}, T-1\right) > 0.975\right], \tag{18}$$

where $\mathcal{F}_t(\tau, k)$ denotes the cumulative distribution function of the $t-$distribution with k degrees of freedom, $se(\alpha_i)$ is the GMM corrected standard error estimate of the alpha, and $\mathbb{1}(\cdot)$ denotes the indicator function.

Apart from the above two informal tests, we also perform the Gibbons, Ross, and Shanken (1989) (GRS) test and the more general and slightly stricter Wald test to examine whether the alphas are jointly zero. However, these tests are done with constant loadings. Let $\mathbf{f}_t \equiv (f_{1,t}, \ldots, f_{K,T})'$ denote the vector of K asset pricing factors with mean sample estimate denoted by $\overline{\mathbf{f}}$ and covariance matrix sample estimate denoted by $\mathbf{\Omega}$. Also let $\boldsymbol{\alpha} \equiv (\alpha_1, \ldots, \alpha_N)'$ denote the vector of alpha sample estimates and $\mathbf{\Sigma}$ the sample estimate of the covariance matrix of the return residuals. Then under the assumption of i.i.d. residuals that follow the Gaussian distribution, the statistic,

$$J_{GRS} \coloneqq \frac{T \cdot (T-N-K)}{N \cdot (T-K-1)} \left(1 + \overline{\mathbf{f}}' \boldsymbol{\Omega}^{-1} \overline{\mathbf{f}}\right)^{-1} \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha} \sim F_{N,T-N-K} \tag{19}$$

follows the F-distribution with degrees of freedom N and $T-N-K$ under the null hypothesis that the alphas are zero. The above statistic gives us the GRS finite sample test. Relaxing the assumption of Normality, the statistic,

$$J_{Wald} \coloneqq T \left(1 + \overline{\mathbf{f}}' \boldsymbol{\Omega}^{-1} \overline{\mathbf{f}}\right)^{-1} \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha} \sim \chi^2_N \tag{20}$$

follows asymptotically the χ^2 distribution with N degrees of freedom under the null.[17]

To transform the GRS and the Wald tests into tests with time-varying loadings, we assume that the conditional betas are observed without error. The term $\overline{\mathbf{f}}' \boldsymbol{\Omega}^{-1} \overline{\mathbf{f}}$ in the expressions of J_{GRS} and J_{Wald} account for the fact that the factor loadings are estimated with error. Removing this term, we obtain the following statistics,

$$J^*_{GRS} \coloneqq \frac{T-N}{N} \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha} \sim F_{N,T-N} \tag{21}$$

and

$$J^*_{Wald} \coloneqq T \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha} \sim \chi^2_N \tag{22}$$

to use in order to test the models where the betas vary over time. These are simply tests of whether the means of a set of variables jointly equal zero. We further note that J^*_{GRS} and J^*_{Wald} are stricter tests compared to J_{GRS} and J_{Wald}, respectively.

[17] For more details about the GRS and the Wald tests, see chapter 12 in Cochrane (2005).

6.1 The 30 industry portfolios

For each portfolio, both MC-I and MC-II models provide two sets of conditional betas: $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$. Table 6 presents the results of asset pricing tests conducted on various models and over different sample periods. Section A of the table presents the outcomes of testing the conditional CAPM model from 1932Q1 to 2020Q4, using six different estimates of the conditional betas. Apart from the four estimates generated by our models, we also evaluate the constant betas, as provided by the sample unconditional estimates β_{unc}, and the 5-year running window estimates β_{rw}. The results for the entire sample, spanning from 1927Q1 to 2020Q4, are presented in section B. We utilize the statistical measures J^*_{GRS} and J^*_{Wald}, that we refer to as GRS* and Wald* tests, respectively, to derive these results.

The table presents our primary finding that the conditional CAPM model using beta estimates from the MC-I model effectively explains the returns of the 30 industry portfolios based on the GRS* and Wald* tests. For the period from 1932Q1-2020Q4 (1927Q1-2020Q4) we find that the $p-$values for the GRS* and Wald* tests using MC-I:$\hat{\beta}_{all}$ are 11.7% (6.56%) and 5.0% (2.45%), respectively. These results indicate that, with the exception of the Wald test during the period 1927Q1-2020Q4, the null hypothesis cannot be rejected at the 5% confidence level. Additionally, none of the alpha estimates in either of the samples are statistically significant at the 5% confidence level.

The results are even more robust for MC-I:$\hat{\beta}_{cyc}$, as both tests fail to reject the null hypothesis at the 10% confidence level for either of the two sample periods. Specifically, the p-values for the GRS* and Wald* tests over the period 1932Q1-2020Q4 (1927Q1-2020Q4) are 28.7% (19.1%) and 17.0% (10.1%), respectively. These findings are consistent with our expectation, as we had previously noted that the capital and dividend shares, when compared to their respective sample means, may not accurately reflect the term-structure of dividend shares throughout the entire sample period.

The conditional CAPM with the MC-I:$\hat{\beta}_{cyc}$ beta estimates demonstrates superior performance also with respect to the mean absolute alpha metric, $\mu(|\alpha|)$. In this case, the mean

Table 6: Time-series AP tests – 30 industry portfolios

A. 1932Q1–2020Q4

			MC-I		MC-II	
	β_{unc}	β_{rw}	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$
$\mu(\lvert\alpha\rvert)$ (%)	0.454	0.399	0.339	0.317	0.714	0.666
GRS* p-val (%)	0.004	0.019	11.730	28.743	4.751	1.528
Wald* p-val (%)	0.012	0.289	5.031	16.950	1.489	0.323
$\#p < 5\%$	5	2	0	0	0	2

B. 1927Q1–2020Q4

		MC-I		MC-II	
	β_{unc}	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$
$\mu(\lvert\alpha\rvert)$ (%)	0.462	0.373	0.341	0.615	0.599
GRS* p-val (%)	0.085	6.563	19.112	4.738	0.854
Wald* p-val (%)	0.007	2.447	10.129	1.589	0.163
$\#p < 5\%$	4	0	0	1	2

	C. 1927Q1–2020Q4			D. 1932Q1–2020Q4		
	CAPM	FF3	FF3M	CAPM	FF3	FF3M
$\mu(\lvert\alpha\rvert)$ (%)	0.462	0.552	0.505	0.454	0.538	0.481
GRS p-val (%)	0.152	0.000	0.000	0.266	0.000	0.000
Wald p-val (%)	0.017	0.000	0.000	0.032	0.000	0.000
$\#p < 5\%$	4	10	10	4	9	10

	E. 1927M1–2021M1			F. 1963M7–2021M1			
	CAPM	FF3	FF3M	CAPM	FF3	FF3M	FF5
$\mu(\lvert\alpha\rvert)$ (%)	0.142	0.176	0.157	0.128	0.183	0.169	0.222
GRS p-val (%)	1.138	0.000	0.000	14.114	0.032	0.019	0.001
Wald p-val (%)	0.722	0.000	0.000	9.845	0.008	0.004	0.000
$\#p < 5\%$	5	7	9	2	5	8	11

The table presents the results of asset pricing tests for various models and periods, using the 30 industry portfolios as test assets. Panels A and B examine the conditional CAPM model with different estimates of conditional betas: β_{unc} represents constant betas equal to sample estimates, β_{rw} denotes betas estimated using 5−year running windows, while $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$ refer to model-implied betas. Panels C to F test static factor models: static CAPM, the Fama and French (1993) three-factor model (FF3), the FF3 model together with the Carhart (1997) momentum factor (FF3M), and the Fama and French (2015) five-factor model (FF5). The tests are defined by equations (17) through (22): $\mu(|\alpha|)$ represents the mean absolute α, where a portfolio's α is the mean unexplained return. GRS (GRS*) and Wald (Wald*) refer to the F and χ^2 specification tests for the static factor (conditional CAPM) models. The table reports the $p-$values. $\#p < 5\%$ indicates the number of portfolios with statistically significant α's at the 5% confidence level. All numbers are in percentages, except for $\#p < 5\%$. The title of each panel indicates the sample period and frequency, with Q (M) representing quarterly (monthly) data.

absolute alpha is 0.319% per quarter, as opposed to the 0.454% observed with constant betas. In contrast, the conditional CAPM model employing the beta estimates obtained from the

MC-II model performs poorly in this regard, despite improving performance in the GRS* and Wald* tests when compared to the constant beta and running window beta estimates.

When comparing the performance of the conditional CAPM model between MC-I and MC-II, several observations can be made. Firstly, the performance of MC-II is notably worse than that of MC-I, particularly with respect to the cyclical component of the conditional betas. This suggests that the counter-cyclical behavior of the stock market $p-d$ ratio and its relationship with the conditional betas is not consistent with the observed data. However, unlike MC-I, the term-structure component in MC-II does result in a slight improvement in performance. This could potentially indicate a downward sloping term-structure of equity premia, although it is difficult to draw any firm conclusions.

Notably, the conditional CAPM with the running window betas outperforms the unconditional CAPM (β_{unc}) across all the metrics considered, implying that betas are time-varying. However, these betas fail to capture the fluctuations implied by our model, particularly those related to the business cycle.

The empirical asset pricing literature emphasizes the importance of including multiple pricing factors to account for the cross-sectional variation in expected returns, as evidenced by studies such as Fama and French (2015) and Hou, Xue, and Zhang (2015). Hence, it is necessary to evaluate the performance of such models in explaining the returns of the 30 industry portfolios and compare them with the conditional CAPM model using our model's beta estimates. In our analysis, we consider several models, including the static CAPM, the Fama and French (1993) three-factor model (FF3), the FF3 model combined with Carhart (1997) momentum factor (FF3M), and the Fama and French (2015) five-factor model (FF5).[18] It is worth noting that data for the FF5 model is available only from July 1963. We conduct tests using both quarterly and monthly data for these models, and the results are presented in sections C

[18]The static CAPM model is the same as the conditional CAPM model when the betas are constant (β_{unc}). Hence, the performance of both models is identical with regard the mean absolute alpha and the number of significant alphas, as observed in sections B and C. However, the difference lies in the tests used to evaluate their performances. The β_{unc} model is tested using the more stringent GRS* and Wald* tests, whereas the static CAPM model is tested using the GRS and Wald tests. As a result, the $p-$values reported for β_{unc} are slightly higher than those reported for the static CAPM.

to F of Table 6.

The main observation is that none of the models in Panels C to F can account for the 30 industry portfolio returns, and surprisingly, the static CAPM model performs the best across all metrics. Regardless of the sample period or data frequency, the static CAPM consistently yields the lowest mean absolute alpha, the highest $p-$values, and the fewest statistically significant alphas. Moreover, during the sample period of 1963M7-2021M1, the static CAPM model is not rejected by either the GRS or Wald test, with $p-$values of 14.1% and 9.8%, respectively. Yet, the performance of the static CAPM model is slightly inferior compared to the conditional CAPM model with running window betas, and it is statistically rejected in most samples and data frequencies. Among all the models we consider, the conditional CAPM model with beta estimates provided by the MC-I model exhibits the best performance and it is not statistically rejected, which is mainly due to the cyclical behavior it predicts.

6.2 The Fama-French 25 size and book-to-market portfolios

While our model is successful in explaining the 30 industry portfolio returns, we acknowledge that it may not be sufficient to explain CAPM "anomalous" returns of managed portfolios, such as those based on size or the book-to-market ratio. There are several reasons for this limitation. Firstly, our model's ability to explain returns relies on its ability to identify the cyclicality of a portfolio. While this may be relatively constant and easily identifiable for the 30 industry portfolios, it is probably not the case for managed portfolios. Secondly, managed portfolio returns may be influenced by the term-structure of their cash-flows, which our model seems unable to capture. This could be due to the difficulty in identifying the true conditional cash-flow term-structure, or because the term-structure of equity returns is downward-sloping, while our model MC-I predicts instead an upward-sloping term-structure.

We examine this conjecture with the 25 Fama-French size and book-to-market portfolios. However, in order to apply our model and infer conditional betas, we have to make the strong assumption that the cash-flow parameters $(\overline{y}_i, \eta_i, \phi_i, \sigma_i, \overline{\sigma}_i)$ of each portfolio are constant over time. Applying the same procedure as for the 30 industry portfolios yields four sets of conditional betas, with two sets for each of the MC-I and MC-II models. We then perform the same tests and over the same sample periods as for the 30 industry portfolios. We present the results in Table 7.

As expected, none of the MC-I and MC-II models is able to explain the average returns of these portfolios whether over the full sample (section A) or the sample from 1932Q1-2020Q4 (section B). All $p-$values are effectively zero, the mean absolute alphas and the number of statistically significant alphas are higher than those produced by either the static CAPM or the conditional CAPM with running window betas.

However, from sections C to F of Table 7, we find that none of the other models we consider are able to explain the average returns of these portfolios either. All $p-$values are zero, and the multifactor models do not demonstrate any meaningful improvement in the mean absolute alpha when compared to the static CAPM model. Furthermore, these models score lower in terms of the number of statistically significant alphas. This observation is particularly evident in the quarterly data, as can be seen from sections C and D. While the multifactor models appear to offer an improvement in the mean absolute alpha in the post-war monthly sample from 1963M7 to 2021M1, all models are still strongly statistically rejected in this sample as well.

6.3 The mechanism and supporting evidence

The model MC-I is able to explain the average returns of the 30 industry portfolio through a key mechanism: the cross-section shrinks during recessions when the equity premium increases. As a result, the relationship between average returns and unconditional betas is flatter than what the static CAPM predicts, as illustrated in Figure 1. Too see why, consider the decomposition of the average return of a portfolio that we expressed in Section 2,

$$E(R^i_{t+1} - R^f_t) = \beta^i_{unc} \cdot E(\lambda^m_t) + \left[E(\beta^i_t) - \beta^i_{unc}\right] \cdot E(\lambda^m_t) + Cov(\beta^i_t, \lambda^m_t) \tag{23}$$

Table 7: Time-series AP tests: 25 Fama-French portfolios

	A. 1932Q1–2020Q4						
				MC-I		MC-II	
	β_{unc}	β_{rw}		$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$
$\mu(\lvert\alpha\rvert)$ (%)	0.433	0.636		0.698	0.726	1.901	1.682
GRS* p-val (%)	0.000	0.000		0.000	0.000	0.000	0.000
Wald* p-val (%)	0.000	0.000		0.000	0.000	0.000	0.000
$\#p < 5\%$	1	3		4	4	4	4
	B. 1927Q1–2020Q4						
				MC-I		MC-II	
	β_{unc}			$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$	$\hat{\beta}_{all}$	$\hat{\beta}_{cyc}$
$\mu(\lvert\alpha\rvert)$ (%)	0.427			0.671	0.683	1.623	1.486
GRS* p-val (%)	0.000			0.000	0.000	0.051	0.000
Wald* p-val (%)	0.000			0.000	0.000	0.006	0.000
$\#p < 5\%$	1			4	4	4	4
	C. 1927Q1–2020Q4			D. 1932Q1–2020Q4			
	CAPM	FF3	FF3M	CAPM	FF3	FF3M	
$\mu(\lvert\alpha\rvert)$ (%)	0.427	0.405	0.440	0.433	0.410	0.441	
GRS p-val (%)	0.000	0.000	0.000	0.000	0.000	0.000	
Wald p-val (%)	0.000	0.000	0.000	0.000	0.000	0.000	
$\#p < 5\%$	1	11	8	1	9	8	
	E. 1927M1–2021M1			F. 1963M7–2021M1			
	CAPM	FF3	FF3M	CAPM	FF3	FF3M	FF5
$\mu(\lvert\alpha\rvert)$ (%)	0.159	0.116	0.100	0.202	0.092	0.083	0.083
GRS p-val (%)	0.000	0.000	0.001	0.000	0.000	0.000	0.001
Wald p-val (%)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\#p < 5\%$	3	9	7	5	6	6	4

The table presents the results of asset pricing tests for various models and periods, using the 25 Fama-French size and book-to-market portfolios as test assets. Panels A and B examine the conditional CAPM model with different estimates of conditional betas: β_{unc} represents constant betas equal to sample estimates, β_{rw} denotes betas estimated using 5−year running windows, while $\hat{\beta}_{all}$ and $\hat{\beta}_{cyc}$ refer to model-implied betas. Panels C to F test static factor models: static CAPM, the Fama and French (1993) three-factor model (FF3), the FF3 model together with the Carhart (1997) momentum factor (FF3M), and the Fama and French (2015) five-factor model (FF5). The tests are defined by equations (17) through (22): $\mu(|\alpha|)$ represents the mean absolute α, where a portfolio's α is the mean unexplained return. GRS (GRS*) and Wald (Wald*) refer to the F and χ^2 specification tests for the static factor (conditional CAPM) models. The table reports the $p-$values. $\#p < 5\%$ indicates the number of portfolios with statistically significant α's at the 5% confidence level. All numbers are in percentages, except for $\#p < 5\%$. The title of each panel indicates the sample period and frequency, with Q (M) representing quarterly (monthly) data.

where $\lambda_t^m = E_t(R_{t+1}^m - R_t^f)$ denotes the conditional equity premium. Figure 2 plots the second and third terms of the above expression for the conditional betas MC-I:$\hat{\beta}_{cyc}$ of the 30 industry portfolios. The figure reveals that each term contributes about equally to explaining the average returns. During recessions, when the equity premium increases, the betas of pro-cyclical (counter-cyclical) stocks decrease (increase), leading to a negative (positive) value for the third term. The reason why the second term is also negative (positive) for pro-cyclical (counter-cyclical) stocks in our sample is related to the same reason. As a result of the specific pattern of recessions that we have observed, the sharp decline (increase) in the conditional beta of pro-cyclical (counter-cyclical) stocks during those episodes has resulted in average conditional betas that are lower (higher) than their unconditional estimates.

To examine further the behavior of the model in regards to the industry portfolio performances during recessions, we plot the average excess returns during recessions against the unconditional betas of the portfolios in Figure 7. We determine recessions in the model when the business cycle variable x falls below a certain threshold. In the data, about 17% of the time the economy is in recession, according to the NBER recession data. Therefore, recessions are considered to be those states for which x is below the 17%-percentile. We simulate 100 histories of 100 years each and compute the average excess returns for each portfolio during recessions for each history and then across all simulated histories. These are shown with the circles in the figure and the yellow line shows the linear fit with a linear correlation coefficient of -0.97. The red line shows the average excess returns predicted by the static CAPM. We notice that the red line is flatter than the yellow line and the two lines cross. Consequently, the counter-cyclical (pro-cyclical) portfolios, as indicated by their unconditional betas, perform on average better (worse) compared to their static CAPM predictions, according to the model.

We generate the same plot with the data in Figure 8. We observe that, similar to the model behavior depicted in Figure 7, the linear relationship between average excess returns during recessions and unconditional betas (represented by the yellow line) is flatter than what the static CAPM predicts (represented by the red line). However, we also observe that the two lines do not intersect. The yellow line lies above the red line, because the counter-cyclical portfolios outperformed the expectations set by the static CAPM, whereas the pro-cyclical portfolios performed more or less as expected. Furthermore, the linear correlation coefficient

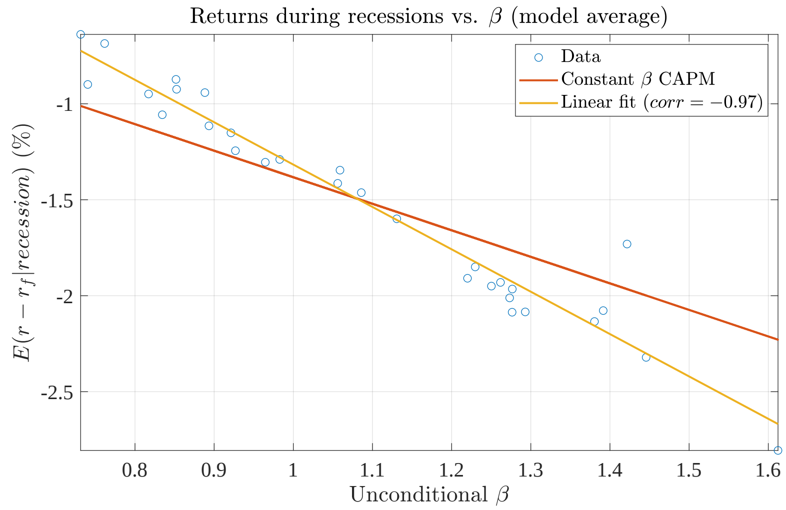

The figure plots with blue circles the average excess returns during recessions against the unconditional betas for the 30 industry portfolios predicted by the model MC-I. The data are obtained from 100 simulations of 400 quarters each, with a burn-in of 100 quarters. The yellow line is the linear fit to the simulated data, where the parenthesis in the legend reports the linear correlation coefficient. The red line represents the relation predicted by the static CAPM model.

Figure 7: Excess returns during recessions – model MC-I average.

between average excess returns and unconditional betas is calculated to be -0.73.

However, it's important to note that in the data, we only observe a single history, whereas in the model, we present the average across 100 simulated histories. When comparing the average excess return of the market (represented by the red line at the point where the unconditional beta is one), we find that it was approximately -4.5%, significantly lower than the model's prediction of around -1.4% across simulations. This implies that, according to the model, the recessions we have observed are more severe compared to their unconditional expectations. Consequently, this finding helps explain the previous observation concerning the second term in expression (23), where the average conditional betas of pro-cyclical (counter-cyclical) stocks were lower (higher) in the observed history than their unconditional estimates.

To enable a more appropriate comparison between the model and the data, we select the simulated history from the set of 100 paths that exhibits the closest average market excess return during recessions to that observed in the data. We then generate the same plot for this selected path, as shown in Figure 9. We observe a remarkably similar behavior to that seen

The figure plots with blue circles the average excess returns during recessions against the unconditional betas for the 30 industry portfolios in the data. The data is quarterly and the sample period is 1926Q4–2020Q4. The yellow line is the linear fit to the data, where the parenthesis in the legend reports the linear correlation coefficient. The red line represents the relation predicted by the static CAPM model.

Figure 8: Recessions - data.

in the data: the yellow line consistently lies above the red line, the slope of the yellow line closely matches the data, and the linear correlation between the average excess returns during recessions and unconditional betas is calculated to be -0.80, which is in close proximity to the data correlation of -0.73.[19]

Therefore, the resemblance between the data and the model supports the mechanism by which the model explains the returns of the 30 industry portfolios over the past century or so. Specifically, it suggests that, firstly, the cross-section of these portfolios in terms of conditional betas contracts during recessions. Secondly, the average conditional betas over the observed history deviate from their unconditional estimates in a manner that causes the relationship between average excess returns and the unconditional estimates to be flatter than what the static CAPM model predicts.

[19] In Figure 9, we observe that the yellow line intersects the vertical line at an unconditional β of 0.8 (1.5) at approximately -2.9% (-6.5%), mirroring the behavior observed in the data depicted in Figure 8.

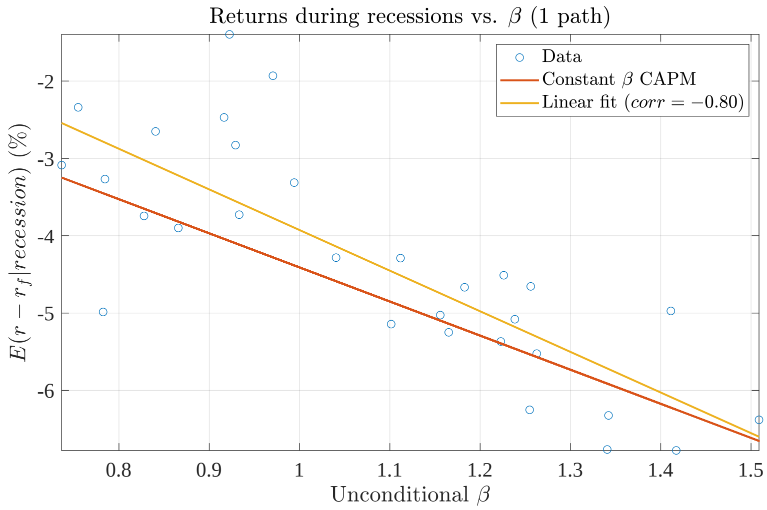

The figure plots with blue circles the average excess returns during recessions against the unconditional betas for the 30 industry portfolios obtained by one simulation path of the model MC-I. The path is chosen out of 100 simulations of 400 quarters each, with a burn-in of 100 quarters, that exhibits the closest average market excess return during recessions to that observed in the data. The yellow line is the linear fit to the simulated data, where the parenthesis in the legend reports the linear correlation coefficient. The red line represents the relation predicted by the static CAPM model.

Figure 9: Recessions - model 1 path.

7 Conclusions

In light of the shortcomings of the traditional static CAPM model in explaining average returns, extensive research has delved into assessing its conditional validity through diverse methodologies. In our study, we introduce a dynamic model with flexible cash-flow and stochastic discount rate dynamics, enabling the inference of asset or portfolio conditional betas. Applying this model to 30 industry portfolios, we find that the conditional CAPM effectively accounts for the cross-section of their average returns spanning approximately a century; an achievement that eludes several established factor models. A key determinant of our success lies in the covariance patterns between conditional betas and the market risk premium, alongside the distinctive impact of historical recessionary periods on observed returns.

Our model hinges on two drivers influencing the conditional beta of an asset. First, the cyclicality of the asset, reflecting variations in its dividend share across economic cycles, yields substantial nonlinear dynamics in conditional betas. Second, the term-structure of expected dividend share growth in a given period, also contributes significantly to conditional beta variations. Notably, our analysis of the 30 industry portfolios indicates that, on average, both components hold equal importance in generating conditional beta variations.

A notable challenge in implementing our approach arises from the necessity of observing a lengthy history of dividend shares characterized by stationary dynamics. This poses a specific issue for the industry portfolios in relation to the term-structure component. Despite this, the cyclicality component alone proves sufficient to explain the cross-section of average returns. However, for managed portfolios such as value or size strategies, this challenge becomes notably more formidable. Nevertheless, our findings strongly suggest that more effort to accurately map the true conditional market betas of assets and portfolios is well-justified.

References

Ang, Andrew, and Joseph Chen, 2007, Capm over the long run: 1926–2001, *Journal of Empirical Finance* 14, 1–40.

——— , and Yuhang Xing, 2006, Downside risk, *The review of financial studies* 19, 1191–1239.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.

Berk, Jonathan B., and Jules H. Van Binsbergen, 2016, Assessing asset pricing models using revealed preference, *Journal of Financial Economics* 119, 1–23.

——— , 2017, How do investors compute the discount rate? they use the capm, *Financial Analysts Journal* 73, 25–32.

Binsbergen, Jules van, Michael Brandt, and Ralph Koijen, 2012, On the timing and pricing of dividends, *American Economic Review* 102, 1596–1618.

Black, Fischer, Michael C. Jensen, and Myron Scholes, 1972, *The Capital Asset Pricing Model: Some Empirical Tests* (New York: Praeger).

Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2004, Estimation and test of a simple model of intertemporal capital asset pricing, *Journal of Finance* 59, 1743–1775.

Buss, Adrian, and Grigory Vilkov, 2012, Measuring equity risk with option-implied correlations, *The Review of Financial Studies* 25, 3113–3140.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Carlson, Murray, Adlai Fisher, and Ron Giammarino, 2004, Corporate investment and asset price dynamics: Implications for the cross-section of returns, *Journal of Finance* 59, 2577–2603.

Cederburg, Scott, and Michael S. O'Doherty, 2016, Does it pay to bet against beta? on the conditional performance of the beta anomaly, *The Journal of finance* 71, 737–774.

Chang, Bo-Young, Peter Christoffersen, Kris Jacobs, and Gregory Vainberg, 2012, Option-implied measures of equity risk, *Review of Finance* 16, 385–428.

Cochrane, John H., 2005, *Asset Pricing* (Princeton University Press: New Jersey) 2 edn.

Cooper, Ilan, 2006, Asset pricing implications of nonconvex adjustment costs and irreversibility of investment, *Journal of Finance* 61, 139–170.

Cosemans, Mathijs, Rik Frehen, Peter C Schotman, and Rob Bauer, 2016, Estimating security betas using prior information based on firm fundamentals, *The Review of Financial Studies* 29, 1072–1112.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3 – 56.

Fama, Eugene F, and Kenneth R French, 2015, A five-factor asset pricing model, *Journal of financial economics* 116, 1–22.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1 – 25.

Friend, Irwin, and Marshall Blume, 1970, Measurement of portfolio performance under uncertainty, *The American economic review* 60, 561–575.

Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.

Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross section of returns, *Journal of Political Economy* 111, 693–732.

Graham, John R., and Campbell R. Harvey, 2001, The theory and practice of corporate finance: Evidence from the field, *Journal of financial economics* 60, 187–243.

Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *The Review of Financial Studies* 28, 650–705.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional capm and the cross-section of expected returns, *Journal of Finance* 51, 3–53.

Lettau, Martin, and Sydney C. Ludvigson, 2001, Resurrecting the (c)capm: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.

Lettau, Martin, and Jessica A. Wachter, 2007, Why is long-horizon equity less risky? a duration-based explanation of the value premium, *Journal of Finance* 62, 55–92.

Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional capm does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Lustig, Hanno N, and Stijn G Van Nieuwerburgh, 2005, Housing collateral, consumption insurance, and risk premia: An empirical perspective, *The Journal of Finance* 60, 1167–1219.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1–47.

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

Petkova, Ralitsa, and Lu Zhang, 2005, Is value riskier than growth?, *Journal of Financial Economics* 78, 187–202.

Santos, Tano, and Pietro Veronesi, 2010, Habit formation, the cross section of stock returns and the cash flow risk puzzle, *Journal of Financial Economics* 98, 385–413.

Van Binsbergen, Jules H., and Ralph S.J. Koijen, 2017, The term structure of returns: Facts and theory, *Journal of Financial Economics* 124, 1–21.

Zhang, Lu, 2005, The value premium, *Journal of Finance* 60, 67–103.

A Further tables and figures

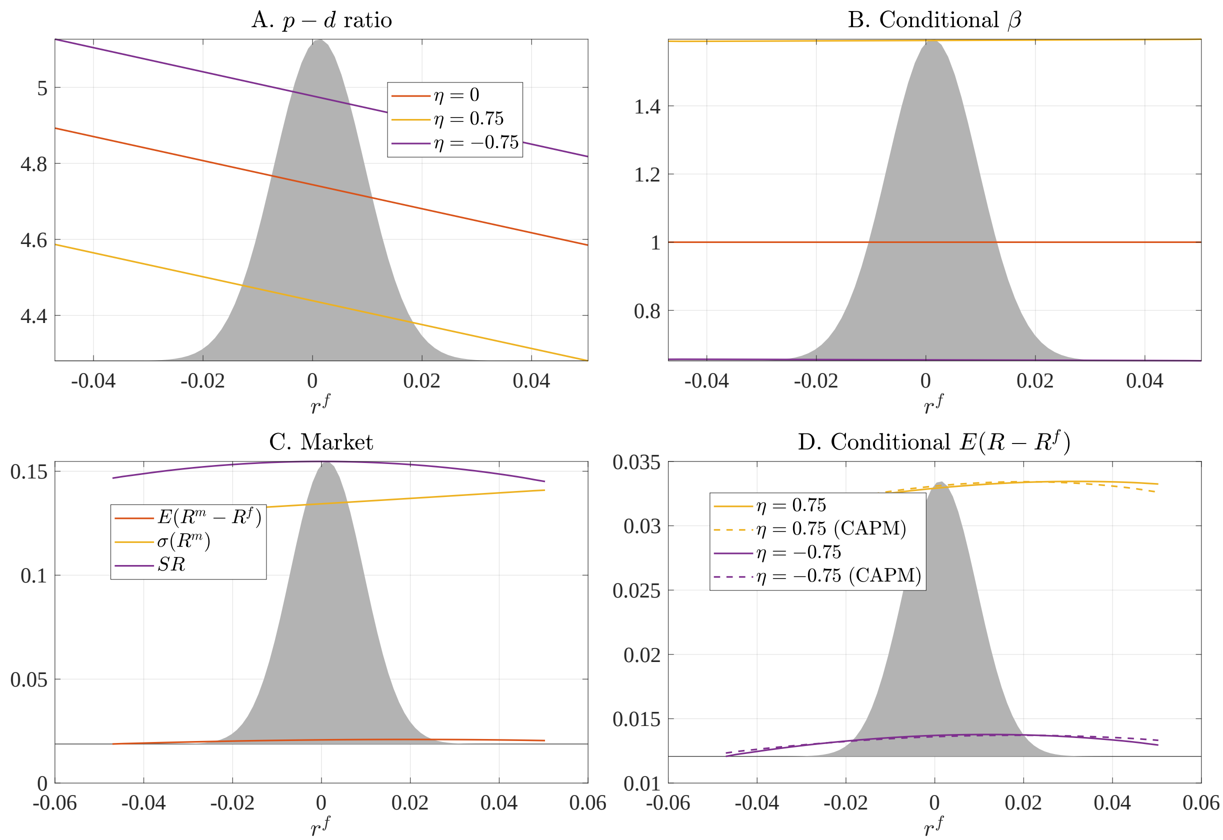

The graph plots various asset pricing quantities against the state of the risk-free interest rate r^f. Panel A plots the log price-to-dividend ratio and Panel B the conditional β for the market ($\eta = 0$), for a pro-cyclical stock ($\eta = 0.75$), and for a counter-cyclical stock ($\eta = -0.75$). Panel C plots the equity premium, the conditional volatility of market returns and the market portfolio Sharpe ratio. In Panel D, the straight lines show the expected excess returns of a pro-cyclical and counter-cyclical stock. The dashed lines show the expected excess returns given by the conditional CAPM model.

Figure 10: Asset price quantities against the interest rate (MC-I)

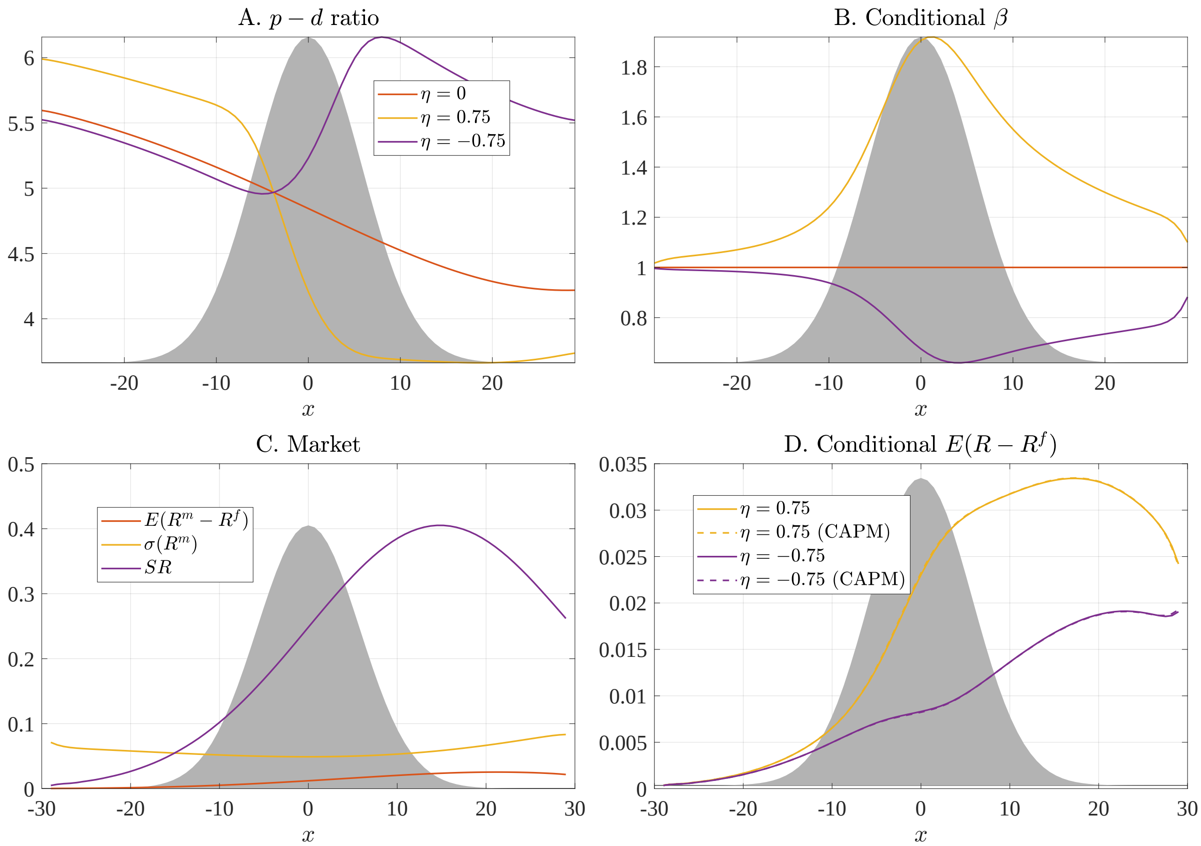

The graph plots various asset pricing quantities against the business cycle variable x. Panel A plots the log price-to-dividend ratio and Panel B the conditional β for the market ($\eta = 0$), for a pro-cyclical stock ($\eta = 0.75$), and for a counter-cyclical stock ($\eta = -0.75$). Panel C plots the equity premium, the conditional volatility of market returns and the market portfolio Sharpe ratio. In Panel D, the straight lines show the expected excess returns of a pro-cyclical and counter-cyclical stock. The dashed lines show the expected excess returns given by the conditional CAPM model.

Figure 11: Asset price quantities over the business cycle (MC-II)

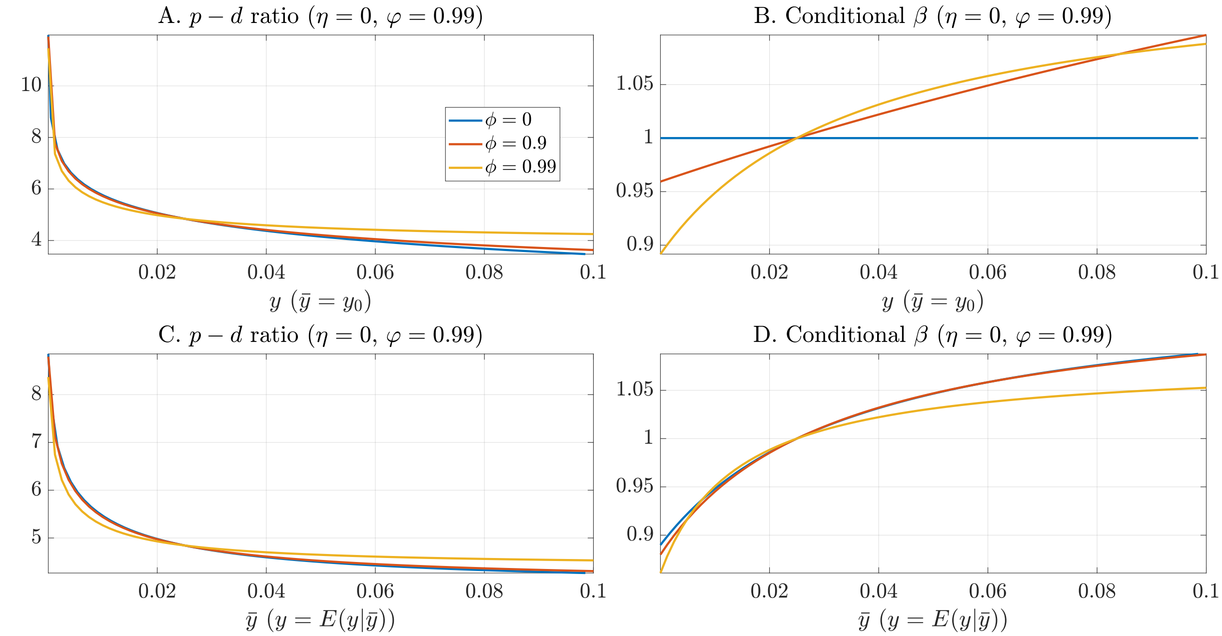

This figure shows for model MC-II the log price-to-dividend ratio ($p-d$) and the conditional β over the individual state $(y, \bar{y})$, for a non-cyclical stock ($\eta = 0$), for various values of φ. We set the aggregate state variables (x, r^f) at their steady state values. Panels A and B plot these quantities against y where we set $\bar{y}$ equal to the unconditional mean, and Panels C and D show the quantities against $\bar{y}$ where we set y equal to its conditional expectation.

Figure 12: Stock $p-d$ and market β in terms of y and $\bar{y}$

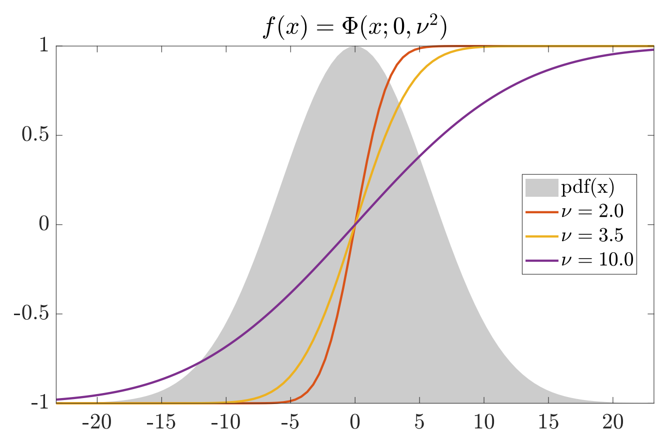

The figure plots the $f(x)$ function, as defined by $f(x) = \Phi(x; 0, \nu^2) - 1$, for various values of ν.

Figure 13: The function $f(x)$

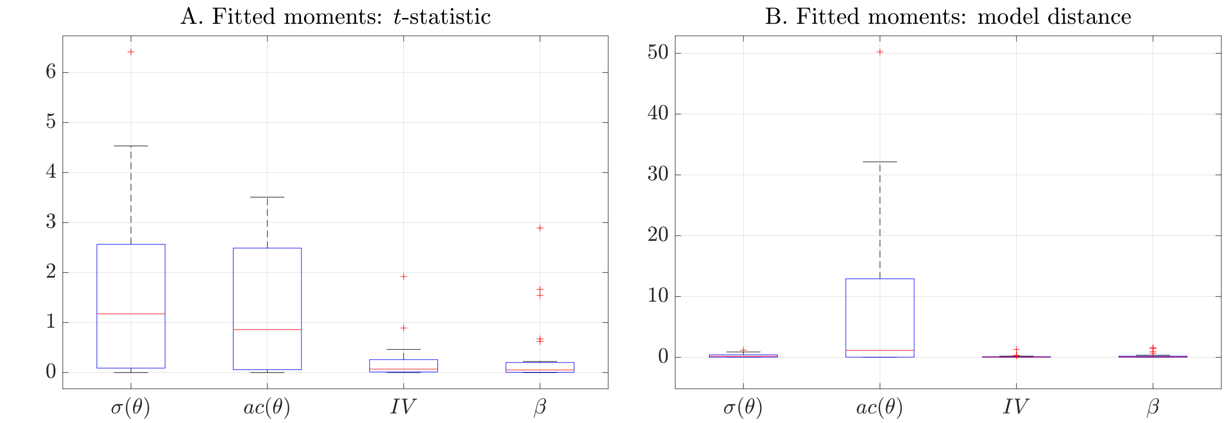

This figure displays box plots illustrating various statistics for the 30 industry portfolios. Panel A shows the t-statistics of the quantities employed in estimating the cash-flow model in model MC-I: $\sigma(\theta)$ and $ac(\theta)$ represent the volatility and autocorrelation of the dividend share, respectively. IV denotes the idiosyncratic return volatility, and β signifies the unconditional sample β. The t-statistics are estimated using the sample estimate standard errors and the standard deviations across 100 model simulations. In Panel B, the figure presents the model distance for each of these quantities. The model distance is calculated as the difference between the data estimate and the average of the model simulations, divided by the standard deviation across simulations.

Figure 14: Calibration metrics (MC-I) – 30 industry portfolios.

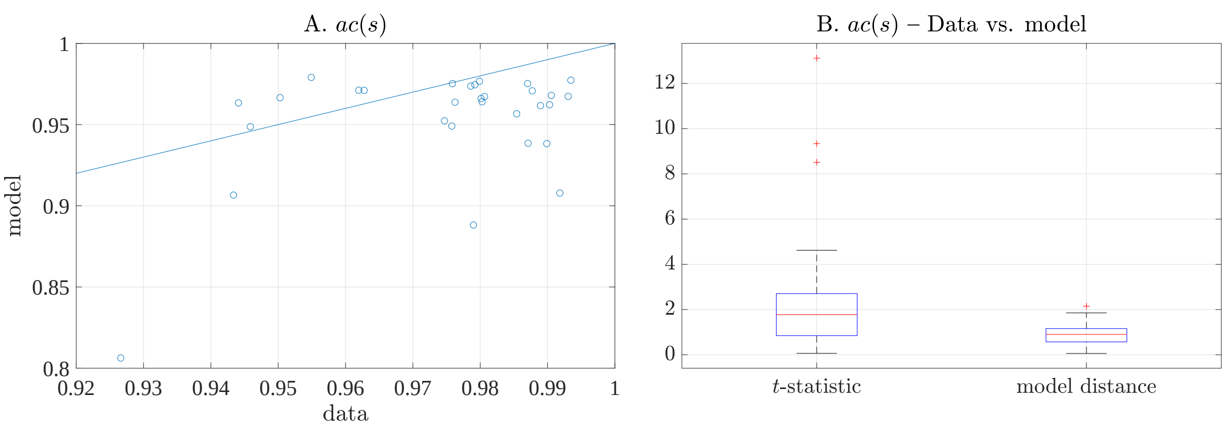

This figure relates the autocorrelations of the capital shares of the 30 industry portfolios implied by the estimated cash-flow model in model MC-I with the data estimates. Panel A plots model-implied values against the data estimates. Panel B shows box plots of the $t-$statistic and the model distance. The t-statistics are estimated using the sample estimate standard errors and the standard deviations across 100 model simulations. In Panel B, the figure presents the model distance for each of these quantities. The model distance is calculated as the difference between the data estimate and the average of the model simulations, divided by the standard deviation across simulations.

Figure 15: Serial correlation of capital shares (MC-I) – 30 industry portfolios.

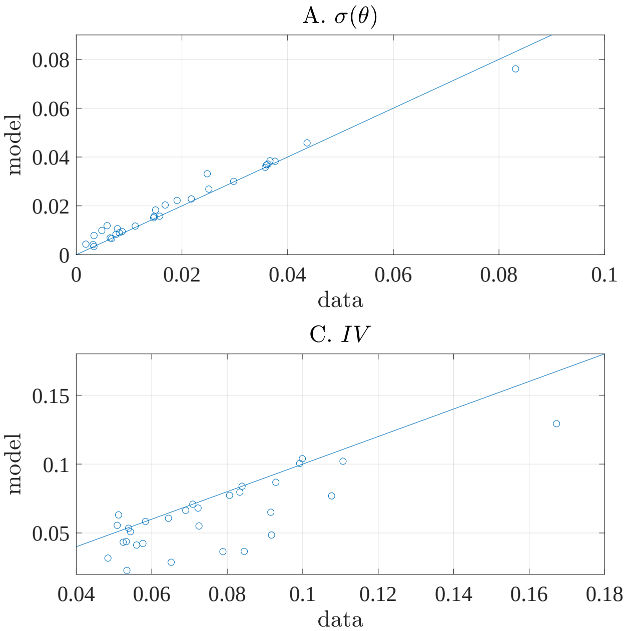

This figure shows the targeted moments regarding the estimation of the cash-flow model for each of the 30 industry portfolios, the volatility of dividend share (Panel A), the autocorrelation of the dividend share (Panel B), the static CAPM idiosyncratic return volatility (Panel C), and the unconditional β (Panel D). The $y-$axis corresponds to the model estimates and the $x-$axis to the data moments. The 45^o line is shown in blue.

Figure 16: Calibration fitted moments (MC-II) – 30 industry portfolios

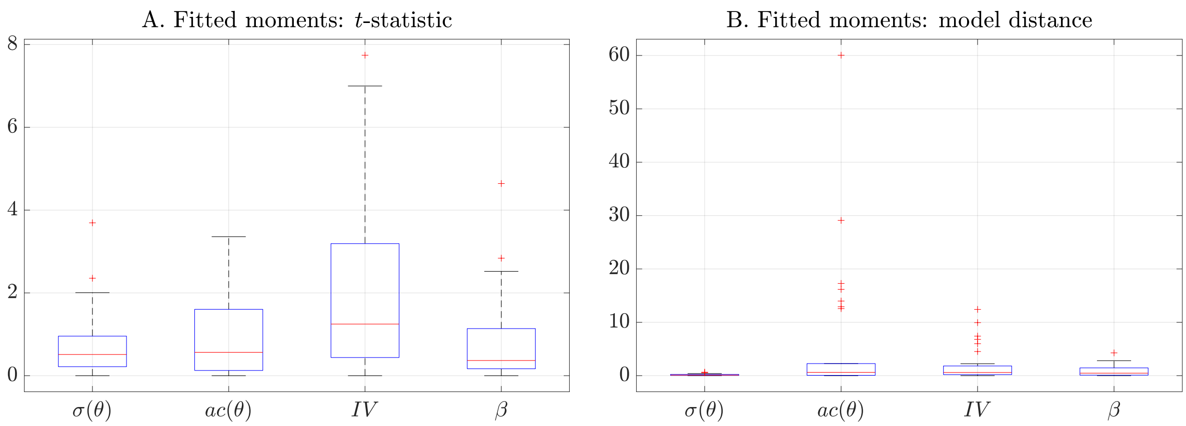

This figure displays box plots illustrating various statistics for the 30 industry portfolios. Panel A shows the t-statistics of the quantities employed in estimating the cash-flow model in model MC-II: $\sigma(\theta)$ and $ac(\theta)$ represent the volatility and autocorrelation of the dividend share, respectively. IV denotes the idiosyncratic return volatility, and β signifies the unconditional sample β. The t-statistics are estimated using the sample estimate standard errors and the standard deviations across 100 model simulations. In Panel B, the figure presents the model distance for each of these quantities. The model distance is calculated as the difference between the data estimate and the average of the model simulations, divided by the standard deviation across simulations.

Figure 17: Calibration metrics (MC-II) – 30 industry portfolios.

Table 8: Statistics of aggregating variables

	MC-I	MC-II		MC-I	MC-II
$\mu(a_1)$	0.0166	0.0148	$\mu(a_2)$	1.0006	1.0017
$\sigma(a_1)$	0.0706	0.0731	$\sigma(a_2)$	0.0185	0.0292
$ac_1(a_1)$	0.9646	0.9535	$ac_1(a_2)$	0.2750	-0.1325
$\rho(a_1, x)$	-0.0320	0.0066	$\rho(a_2, x)$	0.0053	-0.0011
$\rho(a_1, r^f)$	0.0120	0.0054	$\rho(a_2, r^f)$	0.0018	-0.0056
$\mu(u)$	0.0012	0.0011	$\mu(\bar{u})$	0.0057	0.0006
$\sigma(u)$	0.3196	0.3932	$\sigma(\bar{u})$	0.6304	0.3138
$ac_1(u)$	-0.0051	0.0076	$ac_1(\bar{u})$	0.0016	-0.0005
$\rho(u, \epsilon)$	0.0024	0.0029	$\rho(\bar{u}, \epsilon)$	0.0000	0.0001

The table shows several statistics of the aggregating variables from the simulated model economies. For each model we run 100 simulations of 400 quarters each (with a burn-in of 100 quarters) and the statistics reported are averages across simulations. We denote the mean with $\mu(\cdot)$, the standard deviation with $\sigma(\cdot)$, the first-lag autocorrelation with $ac_1(\cdot)$, and the correlation with $\rho(\cdot, \cdot)$.