JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 55, No. 3, May 2020, pp. 783–827
COPYRIGHT 2018, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109018001412

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Capital Asset Pricing with a Stochastic Horizon

Michael J. Brennan and Yuzhao Zhang*

Abstract

In this paper we present empirical tests of an extended version of the capital asset pricing model (CAPM) that replaces the single-period horizon with a probability distribution over different horizons. Adopting a simple parameterization of the probability distribution of the length of the horizon, we estimate the parameters of the distribution as well as the parameters of the CAPM. We find that the extended model is not rejected for several different samples of common stocks, and for these samples it outperforms not only the standard CAPM but also the Fama–French (1993) 3-factor model. The probability distribution over horizon dates varies over time with the New York Stock Exchange (NYSE) turnover rate. We also find that returns satisfy the Euler equation of a representative financial institution that holds the market portfolio and has horizon probabilities estimated from 13F filings.

I. Introduction

The capital asset pricing model (CAPM) of Sharpe (1964), Lintner (1965), and Mossin (1966) is the oldest and perhaps the most elegant and intuitive of the asset pricing paradigms. It is this intuitive elegance that has allowed the model to survive newer and more sophisticated paradigms in the classroom, the courtroom and the boardroom,[1] despite its perceived lack of empirical success and general rejection among academics. After early acceptance (Black, Jensen, and Scholes (1972), Fama and MacBeth (1973)), the model has come to be regarded by many academics as a useful classroom example that fails to capture the complexity of the real world relation between risk and return. Indeed, rejection of the CAPM has become a touchstone for assessing the power of statistical tests of asset pricing models. For empirical purposes, the standard asset pricing paradigm has become

*Brennan, michaelbrennan830@gmail.com, University of California–Los Angeles Anderson School and Manchester University; Zhang (corresponding author), yzhang@business.rutgers.edu, Rutgers University Business School. This paper has benefited greatly from the comments and suggestions of Jennifer Conrad (the editor) and Joost Driessen (the referee). We also thank John Campbell, Jay Shanken, William Sharpe, Grigor Vilkov, Yan Xu, and participants in seminars at Tilburg University, Erasmus University, and the Stockholm School of Economics, as well as Jonathan Lewellen, our discussant at the Western Finance Association, for comments and suggestions on an earlier version of this paper.

[1]Graham and Harvey (2001) report that 73.5% of firms in their survey "always or almost always use" the CAPM, and that its use was much more common among large firms.

the Fama–French (FF) 3-factor model, which can be justified as a variant of either the intertemporal CAPM or Ross' arbitrage pricing theory (FF (1996)). In this paper, we show that the classical CAPM can be made empirically as well as theoretically relevant by a simple change in its assumptions. Instead of assuming a collection of mean-variance investors with a known single horizon or liquidation date, we assume a continuum of identical investors with mean-variance preferences but with a stochastic horizon or portfolio liquidation date. Each investor is assumed to trade only twice: when he enters the market, and when he realizes the value of his portfolio on the stochastic liquidation date. This assumption is intended to capture the fact that agents do not trade continuously as would be the case if transaction costs were 0, but instead when they purchase shares they have some trading horizon that is unknown both to them and to the econometrician.[2] The effect of the assumption is to ensure that investors are concerned with the distribution of wealth at several different future dates on which assets may be liquidated and returns realized. By parameterizing the probability distribution of the liquidation date we can estimate the model parameters using a standard generalized method of moments (GMM) approach, and the probability distribution of liquidation dates yields an estimate of the implied average horizon of investors.

Since the CAPM is a single-period model which does not specify the length of the time period or horizon for which the investor makes decisions, it is necessary to make some assumption about this for empirical purposes. Our assumption of a stochastic horizon or trading interval, which is critical to our analysis, contrasts with the usual assumption made in applying the single-period CAPM, that the horizon is fixed and exogenous. In the great majority of asset pricing tests the horizon is taken as 1 month and returns are measured over monthly intervals. This assumes that the 1-month beta coefficient is a sufficient statistic for the systematic risk of a security. Table 1 reports the results of multiple regressions of returns on selected FF (1993) portfolios on concurrent and lagged market excess returns, where the market return is proxied by the Center for Research in Security Prices (CRSP) value-weighted index return. Panel A reports results for the period 1926–2013, and Panel B for the more recent period 1970–2013 only. We report F-statistics for the null hypothesis that the coefficients of lagged market returns are jointly 0 and statistics that are significant at the 5% level are shown in bold. It is apparent, particularly in the full sample, that portfolio returns initially underreact to market returns and there is evidence of a longer-run overreaction in the significant negative coefficients at lag 4. These lagged cross-correlations between security and market returns mean that the measured 1-month beta coefficient is not a sufficient statistic for the beta coefficient computed for longer horizons.

One approach to dealing with these lagged cross-correlations is simply to ignore them for asset pricing purposes. An alternative approach is to calculate Dimson (1979) betas (see also Scholes and Williams (1977)) by summing the coefficients on the current and lagged market returns. But this assumes that the

[2]Infrequent trading by certain classes of investors has been invoked by Chien, Cole, and Lustig (2012) to explain time variation in the Sharpe ratio, and Brunnermeier and Nagel (2008) and Calvet, Campbell, and Sodini (2009) provide empirical evidence that individual investors revise their portfolios only infrequently.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

investor was able to trade at prices that were not actually observed, namely the prices that would have been observed if the prices had adjusted instantaneously to the market return. This seems unsatisfactory. Our approach assumes that the investor trades at the end of each month, but only with a given probability.

We are by no means the first to consider the effect of the horizon or time interval over which returns are measured on the fit of the CAPM. As long ago as 1977, Levhari and Levy complained that in various empirical tests of the CAPM, the investment horizon (return interval) was selected arbitrarily. They showed that if returns are serially independent and homoscedastic then the beta coefficient computed using n-period returns will differ systematically from that computed using

TABLE 1

Multiple Regression of Monthly Excess Returns on Selected FF Portfolios on Contemporaneous and Lagged Market Returns

Table 1 reports estimates of the equation:

$$R_{j,t} = \alpha_j + \beta_{j0}R_{m,t} + \beta_{j1}R_{m,t-1} + \beta_{j2}R_{m,t-2} + \beta_{j3}R_{m,t-3} + \beta_{j4}R_{m,t-4} + e^{\tau}_{j,t},$$

where $R_{j,t}$ and $R_{m,t}$ are the excess returns on portfolio j and the market portfolio in month t. The market portfolio is proxied by the CRSP value-weighted index. Test portfolios are 25 FF size- and BM-sorted portfolios. R^2_c denotes the R^2 of the constrained model in which $R_{m,t}$ is the only explanatory variable. We report the F-statistic for the test of the hypothesis that the coefficients on the lagged market returns are jointly 0, with p-values below the F-statistics. * indicates statistical significance at the 5% level. t-statistics are shown in parentheses. Panel A reports results for the full sample and Panel B for the period 1970–2013.

Panel A. 1926–2013

		β_{j0}	β_{j1}	β_{j2}	β_{j3}	β_{j4}	R^2	R^2_c	F
Small	Low	1.6	0.112	0.081	−0.074	−0.121	0.530	0.521	6.571*
		(33.37)	(2.32)	(1.68)	(−1.53)	(−2.52)			(0.00)
	3	1.365	0.142	0.085	0.005	−0.124	0.678	0.661	18.300*
		(45.40)	(4.69)	(2.82)	(0.17)	(−4.12)			(0.00)
	High	1.36	0.253	0.027	−0.025	−0.078	0.657	0.632	25.357*
		(42.52)	(7.86)	(0.84)	(−0.78)	(−2.45)			(0.00)
2	Low	1.228	0.099	0.072	0.013	−0.07	0.736	0.725	14.367*
		(52.63)	(4.23)	(3.08)	(0.56)	(−2.99)			(0.00)
	3	1.201	0.13	0.014	−0.027	−0.076	0.805	0.791	24.967*
		(63.62)	(6.87)	(0.72)	(−1.43)	(−4.02)			(0.00)
	High	1.37	0.123	0.008	−0.028	−0.091	0.728	0.718	12.928*
		(51.48)	(4.59)	(0.30)	(−1.05)	(−3.42)			(0.00)
3	Low	1.259	0.111	0.015	−0.04	−0.051	0.824	0.815	18.279*
		(67.97)	(5.96)	(0.83)	(−2.12)	(−2.78)			(0.00)
	3	1.128	0.059	−0.013	−0.026	−0.018	0.856	0.853	7.688*
		(77.19)	(4.03)	(−0.85)	(−1.76)	(−1.22)			(0.00)
	High	1.351	0.049	−0.005	−0.047	−0.045	0.749	0.746	4.188*
		(54.75)	(1.98)	(−0.19)	(−1.88)	(−1.83)			(0.01)
4	Low	1.084	−0.015	−0.004	0.011	−0.004	0.874	0.873	0.722
		(84.32)	(−1.15)	(−0.34)	(0.82)	(−0.28)			(0.54)
	3	1.089	0.036	0.005	−0.016	−0.021	0.878	0.876	4.769*
		(85.21)	(2.81)	(0.37)	(−1.27)	(−1.67)			(0.00)
	High	1.395	0.093	0.008	−0.08	−0.047	0.760	0.753	10.373*
		(56.02)	(3.71)	(0.33)	(−3.18)	(−1.90)			(0.00)
Big	Low	0.969	−0.011	−0.007	0.002	0.004	0.922	0.922	0.879
		(110.08)	(−1.20)	(−0.78)	(0.17)	(0.44)			(0.45)
	3	0.979	−0.009	−0.017	−0.006	−0.005	0.856	0.855	0.978
		(77.90)	(−0.68)	(−1.32)	(−0.44)	(−0.40)			(0.40)
	High	1.149	−0.035	0.034	−0.033	−0.012	0.693	0.692	1.845
		(48.11)	(−1.46)	(1.41)	(−1.38)	(−0.50)			(0.14)

(continued on next page)

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 1 (continued)
Multiple Regression of Monthly Excess Returns on Selected FF Portfolios on Contemporaneous and Lagged Market Returns

Panel B. 1970–2013									
Small	Low	1.387 (30.75)	0.237 (5.25)	0.033 (0.74)	0.007 (0.15)	−0.103 (−2.30)	0.659	0.638	11.245* (0.00)
	3	1.056 (32.50)	0.161 (4.97)	0.009 (0.27)	0.039 (1.22)	−0.069 (−2.16)	0.682	0.664	10.011* (0.00)
	High	1.035 (30.29)	0.263 (7.72)	0.002 (0.06)	0.052 (1.53)	−0.032 (−0.95)	0.663	0.623	20.652* (0.00)
2	Low	1.375 (39.93)	0.1 (2.91)	−0.016 (−0.48)	−0.026 (−0.77)	−0.075 (−2.20)	0.758	0.752	4.692* (0.00)
	3	1.028 (40.09)	0.103 (4.04)	−0.035 (−1.37)	−0.008 (−0.30)	−0.038 (−1.51)	0.762	0.753	6.482* (0.00)
	High	1.079 (32.99)	0.146 (4.46)	−0.034 (−1.03)	0.018 (0.55)	−0.01 (−0.29)	0.687	0.674	6.767* (0.00)
3	Low	1.311 (44.48)	0.045 (1.53)	−0.018 (−0.61)	−0.025 (−0.87)	−0.049 (−1.67)	0.794	0.791	2.115 (0.10)
	3	0.983 (44.40)	0.072 (3.24)	−0.039 (−1.77)	−0.011 (−0.49)	−0.022 (−1.03)	0.795	0.79	4.684* (0.00)
	High	0.996 (33.76)	0.079 (2.67)	−0.089 (−3.01)	0.018 (0.60)	−0.018 (−0.60)	0.694	0.685	4.969* (0.00)
4	Low	1.222 (54.21)	(−0.02)	−0.001 (−0.05)	−0.026 (−1.16)	−0.022 (−1.00)	0.850	0.849	0.854 (0.46)
	3	1.015 (47.64)	0.043 (2.01)	−0.046 (−2.17)	0.006 (0.27)	−0.016 (−0.74)	0.816	0.813	2.770* (0.04)
	High	1.021 (35.48)	0.052 (1.80)	−0.011 (−0.39)	0.009 (0.32)	0.025 (0.89)	0.711	0.709	1.432 (0.23)
Big	Low	0.991 (63.16)	−0.018 (−1.18)	0.001 (0.09)	0.003 (0.20)	−0.007 (−0.44)	0.885	0.884	0.555 (0.64)
	3	0.88 (44.97)	−0.028 (−1.42)	−0.018 (−0.91)	−0.003 (−0.17)	0.011 (0.57)	0.795	0.794	1.146 (0.33)
	High	0.871 (28.45)	0.041 (1.34)	0.024 (0.78)	−0.014 (−0.45)	0.002 (0.08)	0.612	0.609	0.939 (0.42)

m-period returns, which in turn affects tests of the CAPM.[3] This result was generalized by Longstaff (1989), who also showed that the continuous time CAPM becomes a multifactor model when discrete period returns are used if there is time variation in expected returns and risk. Lee (1976) considered a generalization of the CAPM to a single but unknown investment horizon under the assumption that periodic returns were serially independent, and showed how to estimate the horizon. Lee, Wu, and Wei (1990) develop a model in which returns are independent and identically distributed (IID), but investors have different horizons.

Turning to the literature on asset pricing and return intervals, Handa, Kothari, and Wasley (1993) show that while the standard CAPM is rejected using monthly returns on size-ranked portfolios, it is not rejected by a standard F-test when annual returns on the same portfolios are used (see also Handa, Kothari, and Wasley (1989)). Kothari, Shanken, and Sloan (1995) provide further evidence of a statistically significant risk premium associated with beta when annual data are used. Finally, Jagannathan and Wang (2007) show that when returns are measured on a calendar year basis, the consumption CAPM explains the returns on 25 FF size- and book-to-market (BM)-sorted portfolios as well as the FF 3-factor model. Their calendar year assumption is motivated by their observation that many decisions

[3] If $m > n$ and $\beta_n > 1 (< 1)$, then $\beta_m > (<) \beta_n$. See also Hawawini (1983).

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

are made at the end of the year because of Christmas, bonuses, and end-of-year capital gains tax considerations. Cohen, Polk, and Vuolteenaho ((2009), p. 2740) "adopt the perspective of a buy-and-hold investor" to "explain price levels instead of short-horizon expected returns" in order to avoid problems of measuring risk with high frequency return betas. This paper differs from theirs by allowing for a range of possible horizons for investors rather than specifying the horizon *a priori*; our model nests the single fixed horizon models.

The setup in this paper should be distinguished from that of Richard (1975) and others who analyze portfolio problems with uncertain lifetime but with continuous trading, since in these models investors trade continuously (or every period in discrete time models) to maintain constant risky portfolio weights. However, in our model investors trade only twice. More directly relevant is Martellini and Urosevic (2006), who study the optimal portfolio strategy of a mean-variance investor with an uncertain horizon and establish that, when stock prices follow a random walk, the portfolios on the efficient set are identical to those on the 1-period efficient set. This implies that a zero beta version of the 1-period CAPM would hold even when the representative investor has a stochastic horizon. In this paper, we do not impose the random walk assumption, and the empirical analysis reveals that the time-series properties of the returns play an important role in the empirical fit of the stochastic liquidation model.

Huang (2003) employs an overlapping generations model of investors with stochastic horizons that is similar to our setup but with only two assets, to study the implications of transaction costs for asset returns. Beber, Driessen, and Tuijp (2011) also study the effect of transaction costs in a model with IID returns in which investors have mean-variance utility functions and heterogeneous horizons. However, the focus of their paper is on the clientele effects that develop when investors have different horizons. Kamara, Korajczyk, Lou, and Sadka (2016) argue that there are long and short horizon risk factors arising from the different autocorrelation properties of the factors. They find evidence that the long horizon factors (which include the market return) are priced when betas are estimated using annual returns, but not when using monthly returns. Gilbert, Hrdlicka, Kalodimos, and Siegel (2014) argue that uncertainty about some firms' true systematic risk introduces a $\Delta\beta$ factor[4] into the CAPM which improves the fit of the model at a high (daily) frequency but makes little difference when quarterly returns are used.

In Section II we motivate our approach by showing how measured returns and betas vary systematically with the length of the return interval, and with the characteristics of the underlying firms. In Section III we show how the length of the return interval affects tests of the CAPM. In Section IV we present our simple version of the CAPM with stochastic liquidation. In Section V we report the empirical results of tests of the model and estimate the distribution of liquidation dates. Section VI considers the robustness of the results and reports results using different sets of test assets, as well as the results for a version of the model based on power utility. In Section VII we show that the probability distribution of liquidation dates is related to the turnover rate on the NYSE; when the turnover rate is

[4]This is the return on an arbitrage portfolio that is long high $\Delta\beta$ securities and short low $\Delta\beta$ securities where $\Delta\beta$ is the difference between betas estimated using quarterly and daily returns.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

low, pricing reflects the relatively long average horizons of investors. Section VIII compares the performance of the model with that of the FF factor models, using as test assets portfolios formed on firm size, investment, and profitability. We also show that returns satisfy the Euler equation of a representative financial institution that holds the market portfolio and has horizon probabilities that we estimate from institutional holdings data in 13F filings. Section IX concludes. For most of our results, we rely on the FF size- and BM-sorted portfolios as our test assets since these have become almost canonical in tests of asset pricing models. However, we also present results for the FF industry portfolios. This is for purposes of robustness, and because differences in costs of capital across industries are likely to be more important for corporate capital budgeting decisions than differences across size- and BM-sorted portfolios since small size and high BM are typically transient firm characteristics.

II. Returns, Betas, and the Measurement Interval

In this section, we show first how average returns and betas vary empirically with the length of the return interval, and then relate the variation of beta with the return interval to firm and portfolio characteristics.

A. Horizon Effects for the FF Portfolios

We motivate our analysis by showing first how both the mean returns and the betas on the FF size- and BM-sorted portfolios are affected by the choice of the time interval over which returns are measured. The τ-month return starting in month t, R_t^τ, is defined by compounding the 1-month returns, R_t:

$$R_t^\tau \equiv \Pi_{k=0}^{k=\tau-1}(1+R_{t+k})-1.$$

For each of the 25 size- and BM-sorted FF portfolios, τ-month betas, $\beta_{\tau j}$, $j=1,\ldots,25$, were estimated by ordinary least squares (OLS) for $\tau=3,6,9,\ldots,24$ months using overlapping monthly observations for the period July 1926–Dec. 2013. The estimation equation using τ-month returns is:

$$R_{j,t}^\tau = \alpha^\tau+\beta_{\tau j}R_{m,t}^\tau+e_{jt}^\tau,$$

where $R_{m,t}^\tau$ is the τ-month return on the market portfolio starting in month t, which is proxied by the CRSP value-weighted index return.

For brevity, we only report results for 15 of the 25 portfolios; results for the other portfolios follow a very similar pattern.[5] Panel A of Table 2 reports for each of the selected portfolios the 1-month betas and the ratio of the τ-month beta to the 1-month beta for each τ. If the betas were independent of the return interval, the entries after the 1-month beta would all be equal to unity. Instead, we see that there is a tendency for the betas to rise with the return interval.[6] Note, however, that this effect is not the same for all portfolios. For the 3- and 12- month return intervals for example, there is a clear tendency for the betas of high- BM-sorted

[5]Within each size quintile, we report numbers 1, 3, and 5 BM portfolios.

[6]Boguth, Carlson, Fisher, and Simutin (2016) show that when stocks adjust to new information slowly, average short-horizon portfolio returns are biased down.

TABLE 2

The Effect of the Return Interval on Estimates of Betas and Mean Returns for Selected FF Size- and BM-Sorted Portfolios

Panel A of Table 2 reports the ratio of βs estimated using τ-month returns to those estimated using 1-month returns, where the τ-month β, $\beta_{\tau j}$, is estimated using monthly overlapping observations from the equation:

$$R^{\tau}_{j,t} \quad = \quad \alpha^{\tau} + \beta_{\tau j} R^{\tau}_{m,t} + e^{\tau}_{j,t},$$

where $R^{\tau}_{j,t} = \Pi_{k=0}^{k=\tau-1}(1+R_{j,t+k})-1$, and the τ-month market return, $R^{\tau}_{m,t}$ is similarly defined. The market return is proxied by the CRSP value-weighted index return. Panel B reports the ratio of the mean annualized returns using a τ-month return interval to the mean annualized returns using a 1-month return interval. The annualized return using a τ-month return interval, ${}^{a}\bar{R}^{\tau}_{j}$, is defined by ${}^{a}\bar{R}^{\tau}_{j} \equiv (1+\bar{R}^{\tau}_{j})^{(12/\tau)}-1$, where the (nonannualized) average τ-month return, $\bar{R}^{\tau}_{j}$, is defined by $\bar{R}^{\tau}_{j} = \Sigma_{t=1}^{T-\tau}\Pi_{k=0}^{\tau-1}(1+R^{1}_{j,t+k})$. The sample period is July 1926–Dec. 2013.

Panel A. 1-Month βs, β_{1j}, and βs Ratios

			Ratio of τ-Month β to 1-Month β							
		β_{1j}	$\tau=3$	6	9	12	15	18	21	24
Small	Low	1.618	1.153	1.002	0.938	0.935	0.938	0.857	0.837	0.797
	3	1.379	1.240	1.197	1.078	1.149	1.155	1.085	1.080	1.020
	High	1.390	1.236	1.202	1.125	1.254	1.243	1.189	1.206	1.156
2	Low	1.235	1.069	1.094	1.061	1.043	1.026	1.003	0.996	0.984
	3	1.217	1.157	1.104	1.037	1.103	1.074	1.030	1.031	0.994
	High	1.385	1.160	1.074	0.984	1.064	1.065	0.998	1.022	0.979
3	Low	1.274	1.119	1.069	1.009	1.060	1.043	0.999	1.010	0.973
	3	1.137	1.089	1.032	0.986	1.032	1.021	1.006	1.025	1.009
	High	1.360	1.106	1.003	0.936	0.988	0.971	0.932	0.934	0.892
4	Low	1.081	0.995	1.021	1.019	1.007	0.998	0.983	0.961	0.943
	3	1.094	1.054	1.055	1.030	1.064	1.053	1.026	1.033	1.014
	High	1.413	1.138	1.025	0.944	1.010	0.986	0.913	0.909	0.855
Big	Low	0.968	0.977	0.988	0.995	0.988	0.995	1.012	1.026	1.041
	3	0.978	1.031	0.986	0.982	1.019	1.017	1.007	1.008	1.009
	High	1.149	1.052	1.026	0.995	1.024	1.044	1.021	1.028	1.002

Panel B. Annualized Average 1-Month Returns and Return Ratios

			Ratio of Annualized τ-Month Returns to Annualized 1-Month Returns							
		${}^{a}\bar{R}^{1}_{j}$	$\tau=3$	6	9	12	15	18	21	24
Small	Low	0.105	1.106	0.970	0.930	0.914	0.914	0.879	0.861	0.830
	3	0.175	1.091	1.068	1.027	1.044	1.047	1.030	1.017	0.989
	High	0.221	1.077	1.058	1.032	1.051	1.052	1.040	1.032	1.012
2	Low	0.112	1.065	1.057	1.047	1.060	1.052	1.042	1.033	1.018
	3	0.169	1.056	1.034	1.015	1.033	1.026	1.014	1.008	0.992
	High	0.202	1.060	1.028	0.992	0.999	0.997	0.980	0.975	0.957
3	Low	0.127	1.068	1.042	1.024	1.047	1.041	1.024	1.017	0.996
	3	0.161	1.037	1.016	1.003	1.016	1.011	1.003	0.999	0.985
	High	0.189	1.042	1.011	0.993	1.004	0.997	0.983	0.974	0.953
4	Low	0.126	1.023	1.018	1.018	1.024	1.016	1.005	0.991	0.973
	3	0.147	1.033	1.028	1.023	1.036	1.033	1.024	1.019	1.006
	High	0.173	1.060	1.020	0.989	0.998	0.987	0.964	0.952	0.929
Big	Low	0.114	1.016	1.011	1.013	1.018	1.015	1.010	1.008	0.998
	3	0.121	1.026	1.013	1.018	1.034	1.034	1.029	1.023	1.012
	Low	0.154	1.021	1.005	0.993	0.997	0.995	0.983	0.972	0.949

portfolios to be high relative to the 1-month betas. This effect is concentrated in the small firm portfolios, to a lesser extent in the large firm portfolios, and it is much less pronounced for the medium size firm portfolios.[7] In contrast, for the 24-month return interval there is a clear tendency for the betas of low-BM firm portfolios to be below their 1-month equivalents, and this is most pronounced for the small firm portfolios. In some cases these effects are very large. For example, the beta of the small low-BM firm portfolio declines by 7% (20%) relative to the

[7]Longstaff (1989) shows that the ranking of betas for different firms may reverse as a function of the interval over which returns are measured.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

1-month value as the return interval is increased to 12 (24) months. In contrast, the beta of the small high book-to-market firm portfolio increases by 25% (16%) as the return interval is increased from 1 month to 12 (24) months.

It follows from the definition of the τ-month return that, if 1-period returns are independent, then the expected 1 plus τ-period return is equal to the expected 1 plus 1-period return raised to the power τ:

$$\mathrm{E}[1+R^{\tau}] \quad = \quad \{\mathrm{E}[1+R^{1}]\}^{\tau}.$$

Raising both sides of the equation to the power $12/\tau$, this implies that the annualized expected τ-month return is given by:

$$\{\mathrm{E}[1+R^{\tau}]\}^{12/\tau} \quad = \quad \{\mathrm{E}[1+R^{1}]\}^{12}, \tag{1}$$

such that the annualized expected return is independent of the return interval, τ, and the ratio of the annualized expected τ-month return to the annualized expected 1-month return is equal to unity.

We consider the sample equivalent of relation (1) by comparing the annualized average τ-month return to the annualized average 1-month return. The average τ-month return is defined by:

$$\bar{R}^{\tau} \quad = \quad 1/(T-\tau)\Sigma_{t=1}^{T-\tau} R_t^{\tau}.$$

The *annualized* average τ-month return, ${}^{a}\bar{R}^{\tau}$, is defined by:

$${}^{a}\bar{R}^{\tau} \quad \equiv \quad (1+\bar{R}^{\tau})^{(12/\tau)} - 1.$$

Panel B of Table 2 reports the annualized average of the 1-month returns for the 25 portfolios, and the ratios of the annualized average τ-month returns to the annualized averages of the 1-month returns. For most portfolios, the ratios tend to be above unity for horizons of 3 months and longer, which is consistent with the positive short run autocorrelation in portfolio returns. Overall, the effect of the return interval is stronger on betas than on mean returns. This means that the relation between beta and mean return will depend on the horizon over which they are measured. While in frictionless markets it is natural to follow Merton (1973) in taking the appropriate horizon to be instantaneous, this approach is less defensible for markets in which transaction costs are significant, making the length of the horizon a significant empirical issue for asset pricing, as we show in Section III. The stochastic liquidation model of Section V allows for a probability distribution over horizon dates, and the parameters of this distribution are estimated along with the risk aversion parameter in the subsequent empirical tests. This contrasts with prior specifications in which the horizon is specified exogenously.

B. Determinants of Relative Betas for Different Horizons

There are at least 4 factors in addition to sampling variability that lead to differences between betas that are calculated for the same security using different return intervals. The first factor, which is purely mechanical, is that even if returns are independently and identically distributed over time, the multiplicative relation between n-period returns and 1 period returns implies a complex relation between

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

1-period betas, β_1, and n-period betas, β_n. Levhari and Levy (1977) show that if the CAPM holds over either 1 period or n periods, β_n will exceed (be less than) β_1 if and only if β_1 is greater (less) than unity.

The second factor, identified by Dimson (1979) and Scholes and Williams (1977), is purely empirical. Betas that are estimated using short horizon data may be biased due to thin trading, which causes the observed (closing) prices of the security and the market portfolio to be nonsynchronous. Since the (value-weighted) market portfolio consists mainly of the frequently traded stocks of large companies, the effect is typically to bias down the betas of small stocks that are infrequently traded.

The remaining 2 factors are related to the informational efficiency of markets and the way this differs across securities. First, there are systematic differences across firms in the timing of information release and in the speed of adjustment of stock prices to market-wide information. These differences give rise to complex patterns of cross-temporal covariances and thereby affect betas calculated over different return intervals. Lo and MacKinlay (1990) show that the returns on large firm stocks lead the returns on small firm stocks. Brennan, Jegadeesh, and Swaminathan (1993) show that the speed of adjustment to market-wide information also depends on the number of analysts following the stock holding firm size constant.[8] Finally, Cohen et al. ((2009), p 2740) point out that if markets are even slightly inefficient, mispricing that is systematically related to firm fundamentals may affect measures of risk. Gilbert et al. (2014) show that the difference between quarterly and daily beta estimates is related to measures of the opacity of a firm's accounting statements.

To substantiate the role of information in the relation between long and short return interval betas, we form 25 portfolios each year from 1984 to 2013 on the basis of firm size and analyst following and compare their 1- and 12-month betas. Specifically, we estimate each year from 1984 to 2013 for all common stocks (firms with share code 10 or 11 and stock price higher than 5) on the CRSP tape with at least 36 monthly returns over the past 5 years, a 1-month beta, and a 12-month beta, where the CRSP value-weighted index is taken as the market index. The 12-month beta is estimated using compounded overlapping returns over the past 60 months. At the end of each year (from 1984 to 2013) firms are sorted into size quintiles, and are further sorted into subquintiles by size within each size quintile. Within each size subquintile, firms are then sorted into quintiles by analyst coverage at the end of the year. The portfolio log beta ratio is obtained by averaging the log beta ratios on the 5 analyst quintile portfolios across the size subquintiles. For example, for the log beta ratios on the size quintile 3/analyst quintile 4 portfolio, we average the log beta ratios for size quintile 3 on stocks in analyst quintile 4 portfolio for size subquintiles 1–5. The second stage of sorting into size subquintiles reduces any relation between firm size and analyst following. The logarithms of the beta ratios on the portfolios are reported in

[8]Brennan and Copeland (1988) show that the betas of firms that split their stock increase by approximately 18% if daily returns are used, and Wiggins (1992) shows that the effect is attenuated if returns over longer time intervals are used to compute betas. This is consistent with stock splits increasing the speed of adjustment to information.

Table 3. We also report the difference in the log ratios between low and high coverage portfolios (LMH) within each size quintile. In Panel A the portfolio beta ratios are value-weighted and in Panel B equal-weighted. The standard errors of the log ratios are computed by assuming that the ratios are independent across firms. For both the equal- and value-weighted portfolios, the log beta ratio tends to decrease as analyst coverage increases across all firm size quintiles. The effect of analyst coverage is smallest for the value-weighted big firm portfolios, where the difference between the natural log beta ratios of the high and low analyst coverage portfolios is statistically insignificant. Overall, the results strongly support the hypothesis that analyst coverage increases the speed of price adjustment to market-wide information and thus tends to reduce the 12-month return beta relative to the 1-month return beta.

TABLE 3

Ratio of 12-Month Return β to 1-Month Return β for Portfolios Formed on Firm Size and Number of Analysts

Table 3 reports the log β ratios ($\ln(\beta_{12}/\beta_1)$) for 25 portfolios formed on the basis of firm size and the number of analysts following the firm. β_{12} (β_1) is the β of a security computed from 12-month (1-month) returns. At the end of each year, from 1984 to 2013, firms are sorted into 5 size quintiles, and within each size quintile are further sorted into subquintiles by size. Within each size subquintile firms are sorted into 5 quintiles by analyst coverage at the end of the year. The portfolio log β ratio for a given size quintile/analyst quintile is obtained by averaging the log β ratios on the 5 analyst quintile portfolios across the size subquintiles. The average log β ratios are valued-weighted in Panel A and equal weighted in Panel B. The standard errors of the log ratios are computed by assuming the ratios are independent across firms. We also report the difference in the ratios between low and high coverage portfolios (LMH) within each size quintile. *t*-statistics are shown in parentheses.

Panel A. Value-Weighted $\log(\frac{\beta_{12}}{\beta_1})$

	Small	2	3	4	Big
Low coverage	0.302	0.247	0.099	0.134	0.018
	(6.79)	(4.91)	(2.34)	(3.06)	(0.92)
2	0.258	0.163	0.127	0.068	−0.016
	(6.59)	(3.64)	(2.62)	(1.60)	(−0.52)
3	0.283	0.145	0.113	0.023	−0.046
	(7.97)	(2.90)	(2.76)	(0.56)	(−1.54)
4	0.183	0.169	0.115	0.065	−0.053
	(4.89)	(3.78)	(2.80)	(1.68)	(−2.09)
High coverage	0.190	0.109	0.051	−0.007	0.006
	(4.55)	(2.57)	(1.659)	(−0.26)	(0.18)
LMH coverage	0.112	0.138	0.047	0.140	0.012
	(3.75)	(4.00)	(1.69)	(3.08)	(0.32)

Panel B. Equal-Weighted $\log(\frac{\beta_{12}}{\beta_1})$

	Small	2	3	4	Big
Low coverage	0.309	0.248	0.102	0.136	0.053
	(6.97)	(5.15)	(2.35)	(3.23)	(1.98)
2	0.265	0.174	0.130	0.078	−0.016
	(7.44)	(3.98)	(2.77)	(1.79)	(−0.52)
3	0.288	0.152	0.118	0.029	−0.009
	(8.51)	(3.01)	(2.80)	(0.71)	(−0.32)
4	0.202	0.174	0.116	0.062	−0.021
	(6.08)	(4.03)	(2.80)	(1.64)	(−0.82)
High coverage	0.182	0.115	0.056	0.008	−0.009
	(4.77)	(2.67)	(1.82)	(0.31)	(−0.43)
LMH coverage	0.126	0.133	0.046	0.128	0.061
	(3.78)	(4.23)	(1.57)	(2.97)	(2.53)

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

To provide further evidence that the differences between long and short return interval betas are associated with variables that capture differences in delays and rate of information release across firms, we regress the ratios of long to short return interval betas on firm characteristics that are chosen to capture delays and rate of information release. As in the previous exercise, we estimate each year from 1969 to 2013 for all common stocks with at least 36 monthly returns over the prior 5 years, a 1-month beta, a 6-month beta, and a 12-month beta using compounded overlapping returns over the prior 60 months, where the CRSP value-weighted index is taken as the market index. For each firm, the ratio of the 2 long horizon betas to the 1-month beta is calculated and regressed on a set of firm characteristics. The first regressor is firm size, which is the market value of the equity. We expect that the long horizon beta ratios will be declining in firm size because large firms release more timely information, are followed by more analysts, and are more intensively traded.[9] The second regressor is the BOOK_TO_MARKET ratio. We hypothesize that the ratios of long to short return interval betas will be declining in the BOOK_TO_MARKET ratio because as Brennan and Wang (2010) have shown, the volatility of security mispricing is negatively related to this variable. Both size and BOOK_TO_MARKET ratio are measured at the end of the prior fiscal year. The third variable is the seasonality of earnings. We anticipate that firms whose earnings are more seasonal will have seasonal patterns in information release that will increase the ratio of long horizon to short horizon betas. Seasonality is measured by the average over the prior 5 years of the ratio of the highest quarterly earnings per share for each year to the average quarterly earnings per share for that year. Following Brennan et al. (1993), the fourth variable is the number of analysts following the firm, which is given by the number of analysts who had issued 1-year earnings forecasts in the prior year according to IBES. The fifth variable TO is the annual turnover rate. These variables are demeaned by the cross-sectional average of the same year to account for time trends, and then subjected to a logarithmic transformation. To take account of the "mechanical" tendency of 1-month betas to increase (decrease) with the return horizon if the 1-month is greater (smaller) than unity discussed by Levhari–Levy (1977), we include the estimated 1-month beta. This requires two adjustments. First, the Levhari–Levy effect is nonlinear since it predicts that long-period betas will exceed (fall short of) short-period betas if the beta is greater than unity. Therefore, in addition to including an estimate of β_1, we include a linear function of the estimated β_1 when this exceeds unity, $\mathbf{1}_{\beta_1} > 1 + \max[\beta_1 - 1, 0]$. The Levhari–Levy analysis implies that the coefficient of the 1-month beta will be negative, while the coefficient of the truncated β_1 variable will be positive. Secondly, it is necessary to take account of the fact that β_1 is measured with error since this will induce spurious correlations with the dependent variable, the ratio of β_n to β_1. Therefore we instrument the estimate of the β_1 regressor using firm size, BOOK_TO_MARKET ratio, book leverage, and monthly volatility, estimated over the prior 5 years.

[9]Handa et al. (1989) have previously shown that betas of small firms increase and betas of large firms decrease with the measurement interval.

For each year from 1969 to 2013,[10] the ratios of the estimated betas $\hat{\beta}_6/\hat{\beta}_1$ and $\hat{\beta}_{12}/\hat{\beta}_1$ are regressed on the independent variables described above. Following Fama and MacBeth (1973), the coefficients of the cross-sectional regressions are averaged over time, and standard errors are computed adjusting for serial correlation in the annual coefficient estimates. The results, which are reported in Table 4, show a strong negative size effect which is consistent with the prior

TABLE 4
Long Horizon Market β_s and Firm Characteristics

Panel A of Table 4 reports annual Fama–MacBeth regression estimates of $\ln(\beta_{6,i,t}/\beta_{1,i,t})$ on $\ln(\text{SIZE}_{i,t})$, $\ln(\text{BM}_{i,t})$, $\ln(\text{SEASONALITY}_{i,t})$, $\ln(1+\text{ANALYSTS}_{i,t})$, $\ln(\text{TO}_{i,t})$, $\hat{\beta}_1$, $\mathbf{1}_{\hat{\beta}_1>1}$, and $\hat{\beta}_1\mathbf{1}_{\hat{\beta}_1>1}$. $\beta_{1,i,t}$ ($\beta_{6,i,t}$) is estimated using monthly (6-month compounded) returns over the last 5 years (at least 36 observations required). The market index is proxied by the CRSP value-weighted index. $\text{SIZE}_{i,t}$, $\text{BM}_{i,t}$, $\text{ANALYSTS}_{i,t}$, and $\text{TO}_{i,t}$ are all normalized by their cross-sectional average at time t. $\text{SEASONALITY}_{i,t}$ is measured by the average over the past 5 years of the ratio of the highest quarterly EPS over the average EPS for the year. $\hat{\beta}_1$ is the predicted value of β_1 using size, BM, book leverage, and past 5-year return volatility. $\mathbf{1}_{\hat{\beta}_1>1}$ is a dummy variable which is equal to 1 if $\hat{\beta}_1$ is larger than 1, and 0 otherwise. In Panel B the dependent variable is $\ln(\beta_{12,i,t}/\beta_{1,i,t})$, where $\beta_{12,i,t}$ is estimated using 12-month compounded returns over the last 5 years (at least 36 observations required). All standard errors are adjusted for serial correlation using Newey–West with 6 lags. t-statistics are shown in parentheses. The sample period is 1984–2013 for regressions controlling for analyst coverage, and 1968–2013 for other specifications.

Panel A. 6-Month β over 1-Month β						
ln(SIZE)	−0.058	−0.060	−0.053	−0.045	−0.051	−0.048
	(−7.08)	(−7.26)	(−8.88)	(−6.98)	(−6.50)	(−6.10)
ln(BM)	−0.031	−0.030	−0.049	−0.046	−0.046	−0.046
	(−1.59)	(−1.53)	(−4.14)	(−3.01)	(−3.22)	(−3.01)
ln(SEASONALITY)		0.007	0.007	0.009	0.008	0.008
		(2.44)	(1.55)	(1.82)	(1.75)	(1.68)
ln(1 + ANALYSTS)			−0.047	−0.045	−0.033	−0.030
			(−8.36)	(−5.67)	(−4.27)	(−3.75)
$\mathbf{1}_{\hat{\beta}_1>1}$				−0.370		−0.328
				(−4.27)		(−3.90)
$\hat{\beta}_1$				−0.326	−0.044	−0.266
				(−3.42)	(−0.64)	(−2.76)
$\hat{\beta}_1\mathbf{1}_{\hat{\beta}_1>1}$				0.449	0.090	0.413
				(4.54)	(2.26)	(4.15)
ln(TO)					−0.034	−0.029
					(−4.16)	(−3.97)
Avg. Adj. R^2	0.038	0.039	0.040	0.052	0.051	0.054
Panel B. 12-Month β over 1-Month β						
ln(SIZE)	−0.055	−0.058	−0.047	−0.039	−0.046	−0.043
	(−5.51)	(−5.63)	(−5.43)	(−4.08)	(−4.16)	(−3.78)
ln(BM)	−0.036	−0.036	−0.058	−0.054	−0.056	−0.054
	(−1.71)	(−1.67)	(−4.20)	(−3.57)	(−3.80)	(−3.55)
ln(SEASONALITY)		0.019	0.022	0.021	0.020	0.020
		(4.95)	(4.45)	(4.86)	(5.01)	(4.71)
ln(1 + ANALYSTS)			−0.057	−0.049	−0.038	−0.034
			(−6.72)	(−5.13)	(−3.66)	(−3.14)
$\mathbf{1}_{\hat{\beta}_1>1}$				−0.462		−0.399
				(−2.80)		(−3.43)
$\hat{\beta}_1$				−0.451	−0.147	−0.394
				(−3.81)	(−1.16)	(−2.35)
$\hat{\beta}_1\mathbf{1}_{\hat{\beta}_1>1}$				0.572	0.144	0.527
				(4.16)	(2.63)	(3.78)
ln(TO)					−0.040	−0.032
					(−3.98)	(−3.67)
Avg. Adj. R^2	0.031	0.032	0.030	0.048	0.046	0.050

[10] When the number of analysts is included in the regression the sample starts in 1984.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

studies cited above. The betas of large firms tend to increase with the horizon at a slower rate than those of the betas of small firms. The BM effect is negative in all the regressions but is not significant until ANALYSTS is introduced. However, the introduction of this and other variables does not have a dramatic effect on the magnitude of the coefficient of BOOK_TO_MARKET. Considering these 2 variables alone, we see that they imply that big firms and value firms will have relatively less systematic risk when their risk is measured at long horizons. Thus, if the "true" horizon is, for example, 1 year, then measuring risk and return at the 1-month horizon will tend to understate the risk of both small firms and growth firms, since their 12-month betas tend to be greater than their 1-month betas. This is consistent with the apparent outperformance of small firms relative to standard CAPM benchmarks that measure risk at the 1-month horizon, but is inconsistent with the apparent outperformance of value firms.[11]

The coefficients of the other variables are consistent with the discussion above. The betas of firms with *seasonal* earnings tend to increase as the horizon is extended, and the effect for the 12-month beta ratio is approximately twice as large as for the 6-month beta ratio. We also confirm that the more analysts following the firm, the smaller is the change in beta as the return horizon is extended. These regression results confirm the importance for the difference between short and long return interval betas of variables that are associated with information dissemination.

Table 5 reports the results of similar regressions in which the dependent variable is the 12-month beta and the independent variables are the 1-month beta and firm characteristics, size, and BOOK_TO_MARKET and the betas with respect to the FF high-minus-low (HML) and small-minus-big (SMB) factors. Panel A reports the results for regressions using individual securities. The first regression shows that size and BOOK_TO_MARKET are significant determinants of the 12-month beta after allowing for the effect of the 1-month beta. The signs of the coefficients of these variables are consistent with the sign of the coefficients in Table 4; small and low BOOK_TO_MARKET firms have high 12-month betas given their 1-month betas. The second regression shows a modest increase in explanatory power when the SMB and HML betas are substituted for these firm characteristics, and the coefficients of these betas are consistent with the coefficients of the corresponding physical characteristics.[12]

Panel B of Table 5 reports for the 25 FF portfolios the results of regressing the 12-month beta on the FF HML and SMB betas, as well as β_1. Note that while the positive coefficient of β_{SMB} is consistent with the negative coefficient of size in Panel A and in Table 4 (small firms have higher 12-month betas), the positive coefficient of β_{HML} is inconsistent with the negative coefficient on BOOK_TO_MARKET in Panel A and in Table 4. The second and third columns show that the SMB and HML betas alone explain over 90% of the variance of the

[11] However, see the discussion of Panel B of Table 5 later in this section.
[12] Similar results are obtained if the 1-month beta is instrumented.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 5
Long Horizon Market βs and HML and SMB βs

Panel A of Table 5 reports Fama–MacBeth estimates of:

$$\beta_{12,i,t} = a + b_1\beta_{1,i,t} + b_2\mathbf{1}_{\hat{\beta}_{1,i,t}>1} + b_3\ln(\text{SIZE}_{i,t}) + b_4\ln(\text{BM}_{i,t}) + b_5\beta_{\text{SMB},i,t} + b_6\beta_{\text{HML},i,t} + \varepsilon_{i,t},$$

using firm level data for the period 1968–2013. $\beta_{1,i,t}$ ($\beta_{12,i,t}$) is estimated using monthly (12-month compounded) returns over the last 5 years (at least 36 observations required), where the CRSP value-weighted index is taken as the market return. $\text{SIZE}_{i,t}$ and $\text{BM}_{i,t}$ are normalized by their cross-sectional average at time t. $\mathbf{1}_{\hat{\beta}_{1,i,t}>1}$ is a dummy variable which is equal to 1 if β_1 is larger than 1, and 0 otherwise. Avg. N is the average number of stocks in the cross section. Panel B reports Fama–MacBeth estimates for the 25 FF size- and BM-sorted portfolios of:

$$\beta_{12,i} = a + b_1\beta_{1,i} + b_2\beta_{\text{RMRF},i} + b_3\beta_{\text{SMB},i} + b_4\beta_{\text{HML},i} + \varepsilon_i,$$

where β_{12} (β_1) is the 12-month (1-month) CAPM β, and β_{RMRF}, β_{SMB}, and β_{HML} are the (1-month) FF 3-factor βs. The sample period is 1926–2013 in Panel B. Standard errors are adjusted for serial correlation using Newey–West with 6 lags. t-statistics are shown in parentheses.

Panel A. Firm-Level Results

β_1	0.971	0.929	0.931	0.930
	(12.72)	(12.91)	(12.75)	(12.53)
$\mathbf{1}_{\hat{\beta}_1>1}$	0.102	0.107	0.096	0.093
	(2.61)	(2.83)	(2.97)	(2.99)
ln(SIZE)	−0.050		−0.021	−0.021
	(−2.96)		(−1.69)	(−1.67)
ln(BM)	−0.095		−0.076	−0.076
	(−3.57)		(−3.47)	(−3.44)
β_{SMB}		0.115	0.109	0.108
		(2.11)	(2.18)	(2.18)
β_{HML}		−0.093	−0.079	−0.076
		(−2.48)	(−2.22)	(−2.16)
Adj. R^2	0.25	0.29	0.301	0.30
Avg. N	3464	3464	3464	3464

Panel B. Portfolio Results

β_1	1.138			0.253
	(9.89)			(1.84)
β_{RMRF}			0.253	
			(1.84)	
β_{SMB}		0.335	0.325	0.274
		(15.68)	(15.36)	(7.01)
β_{HML}		0.211	0.209	0.173
		(7.21)	(7.49)	(4.97)
CONST	−0.101	1.014	0.755	0.755
	(0.71)	(5.63)	(5.33)	(5.33)
Adj. R^2	0.80	0.93	0.94	0.94

12-month betas and adding the market beta of the multiple regression (β_{RMRF}) does little to improve the explanatory power. The fourth column shows that the coefficient of β_1 is insignificant in the presence of the SMB and HML betas, so that the 12-month betas for these portfolios are close to being spanned by the HML and SMB betas, and inclusion of the 1-month beta does not improve the spanning. This implies that *if* the CAPM holds with a 12-month horizon, then we would expect to find that 12-month average returns would also be well explained by the FF model loadings on SMB and HML estimated from monthly data. Moreover, we would expect the coefficients of both SMB and HML betas to be positive as in fact they are. Whether the FF model loadings would then explain the cross section of 1-month average returns would depend on the relation between 1-month and 12-month average returns.

III. A Preliminary Analysis of the Effects of the Return Interval on Capital Asset Pricing

We have seen that the interval over which returns are measured has large effects on both the annualized return on a portfolio of stocks and on its measured beta coefficient, and that these effects vary across portfolios. In other words, the return interval affects both the left- and right-hand sides of standard CAPM regressions and therefore presumably affects the fit of the regression.

To illustrate this phenomenon and to assess its potential importance, Table 6 reports the results of OLS cross-sectional regressions of average portfolio excess returns on betas for the period 1926–2013 from the equation:

$$\overline{R_j - R_F} \quad = \quad a + \lambda\beta_j,$$

where betas are computed with respect to the CRSP value-weighted index. Figure 1 plots annualized monthly returns and annual returns against the predicted returns from regressions that use either annualized monthly returns or annual returns, and betas calculated from either monthly returns or annual returns.

In the first row of Table 6 both average returns and betas are calculated from 1-month returns and the R^2 is 12%. The regression slope and the mean absolute error from the regression are annualized by multiplying by 12. The annualized slope is 5.6%, which compares with the average market excess return of 7.8%, and the mean absolute error is 2.2% per annum (p.a.). The intercept is 4.3% p.a. These estimates are typical of CAPM estimates for similar sample periods.[13]

In the second row of Table 6, the dependent variable remains the average monthly return but the independent variable is now the annual beta, obtained by regressing overlapping 12-month compounded returns on the corresponding market return. Replacing the monthly beta by the annual beta raises the R^2 of the

TABLE 6

Cross-Sectional CAPM Regressions Using Monthly and Annual Data

Table 6 reports the results of ordinary least squares cross-sectional regressions of mean excess returns on βs with respect to the CRSP value-weighted index for the 25 FF size- and BM-sorted portfolios for the period 1926–2013.

$$\overline{R_j - R_F} \quad = \quad a + \lambda\beta_j.$$

t-statistics, which were calculated using the Shanken (1992) correction, are in parentheses. For regressions that use monthly returns, the returns are annualized by multiplying by 12. p-values in square brackets are computed from bootstrapped samples that ensure that returns are serially independent. MAE denotes the annualized mean absolute error.

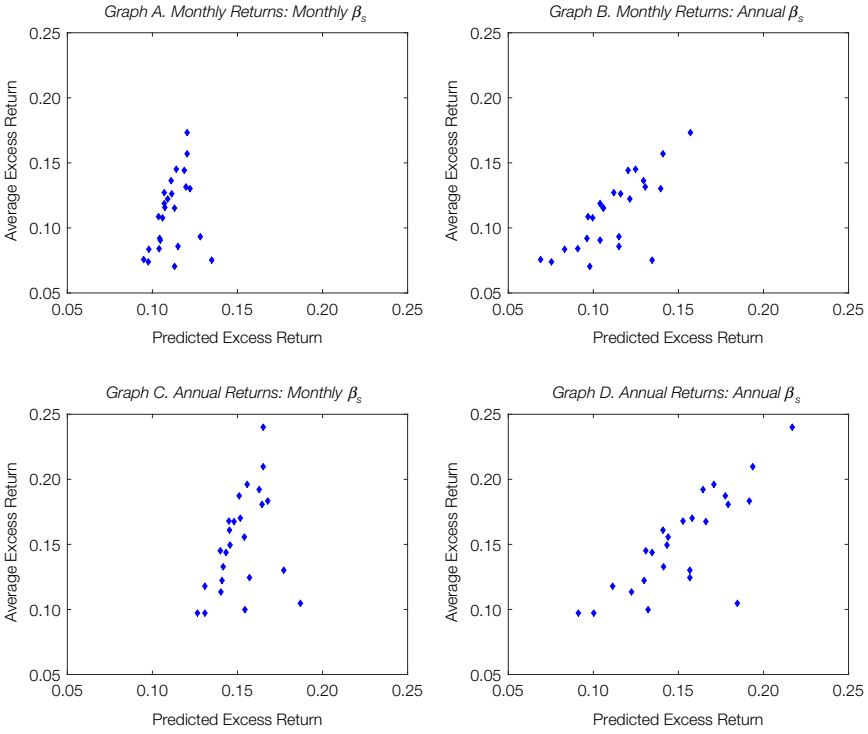
	a	λ	Adj. R^2	MAE
Monthly returns and β_s	0.043	0.056	0.12	0.022
	(0.11)	(1.26)		
		[0.47]	[0.47]	[0.33]
Monthly returns and annual β_s	0.011	0.069	0.56	0.014
	(0.10)	(1.75)		
		[0.09]	[0.00]	[0.00]
Annual returns and monthly β_s	0.049	0.084	0.15	0.028
	(0.12)	(1.73)		
		[0.01]	[0.01]	[0.21]
Annual returns and β_s	0.009	0.098	0.62	0.017
	(0.08)	(2.33)		
		[0.05]	[0.00]	[0.00]

[13] FF (1992), Kothari et al. (1995).

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

FIGURE 1
Actual and Predicted Returns

Figure 1 shows plots of different measures of average excess return against CAPM predicted returns for the 25 FF size- and BM-sorted portfolios for the period 1926–2013. The average excess returns are either average annual excess returns (annual returns) or annualized average monthly excess returns (monthly returns). The CAPM predicted return is provided by the CAPM cross-sectional regressions using annual or monthly returns for the period 1926–2013, whose parameters are reported in Table 5.



cross-sectional regression to 56% and the annualized mean absolute error falls to 1.4%. It is clear that the returns line up much more closely with the annual betas than they do with the monthly betas. Moreover, the slope of the line is now 6.9%, and the intercept is now only 1.1% p.a.[14]

In the third row of Table 6, average 12-month returns are regressed on the monthly betas. Now the R^2 is 15% and the mean absolute error is 2.8%. This is similar to the results we obtained with monthly returns and monthly betas, except that the slope of the line is steeper at 8.4%. Thus the R^2 of the monthly cross-sectional regression in the first row of the table is greatly improved by substituting annual betas for monthly betas. However, substituting annual returns for monthly returns makes relatively little difference to the R^2, although it does increase the slope of the line. In the fourth row of the table and Graph D of Figure 1, we report the results of regressing average annual returns on annual betas. Now the R^2 is

[14]Kothari et al. (1995) also regress monthly returns on annual betas and report significant slope coefficients.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

62% and the slope of the line is 9.8%. The intercept is only 0.9% p.a., and the mean absolute error is now only 1.7% p.a.

Figure 1 shows why the annual betas perform so much better than the monthly betas. In these plots the points will lie around the 45 degree line if realized returns line up with the model predicted returns. The slope of the regression line through the points is given by the product of the correlation between realized and predicted returns and the ratio of the standard deviations of realized and predicted returns. When monthly betas are used to predict returns in Graphs A and C, the points cluster around a line which is much too steep relative to the 45 degree line. This is because the standard deviation of predicted returns is too small due to insufficient variation across portfolios in the monthly betas. In contrast, in Graphs B and D, which are based on annual betas, the points are more closely aligned to the 45 degree line because the standard deviation of predicted returns (and betas) is higher. In fact, the standard deviation of the annual betas is 27% larger than that of the monthly betas.

We construct bootstrapped samples to assess the importance of the serial dependence structure of returns on the cross-sectional risk-return relation at different horizons. Specifically, we generate 10,000 data samples from the joint time series of monthly market and portfolio returns by sampling randomly from the vectors of joint returns. This approach maintains the means and the cross sectional covariance properties of the data while ensuring that there is no time-series dependence within the generated samples. For each generated sample we estimate the parameters (a, λ) and calculate the resulting R^2. The p-values which are reported in square brackets are the proportion of the generated samples in which the estimated parameter exceeds the reported point estimate. A low p-value suggests that the parameter estimates are inconsistent with the assumption of serially independent returns, given the means and the cross-sectional moments of the joint distribution of monthly returns. Not surprisingly the p-values for the first regression, which uses only monthly data, are 47%, since for this regression the temporal ordering of the sample data makes no difference to the parameter estimates. For the other regressions the temporal ordering of the monthly returns will be important if the returns are serially dependent. The p-values for the R^2 statistic are much lower than 0.47 for the other regressions which involve either annual returns or annual betas. When annual returns are regressed on monthly betas, the p-value for the R^2 statistic drops to 0.01, implying that it is highly unlikely that the observed relation between annual returns and monthly beta could have been observed if the returns were serially independent. Moreover, when the independent variable is the annual beta, the p-values for R^2 are 0.00, so that it is highly unlikely that the observed risk-return relations could have been generated by monthly returns with the same joint distribution if the returns were serially independent. Therefore the serial dependence of monthly returns must be taken into account in assessing risk-return relations, and this can only be done by considering returns measured over longer horizons than 1 month.

While these simple OLS regressions do not allow us to draw strong conclusions because the OLS assumptions are not satisfied, they do illustrate in dramatic fashion that the return interval for the purposes of both measuring mean returns and estimating betas makes a substantial difference to the relation between

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

realized average returns and measured betas. The stochastic liquidation model that we develop in Section IV allows the data to determine the optimal horizon over which to measure betas and returns.

IV. Capital Asset Pricing with Stochastic Liquidation

While the original CAPM was a model with a single period of unspecified length, it was soon extended by Merton (1973) to an infinite horizon setting in which investors could and would trade continuously in frictionless securities markets. Subsequent partial equilibrium models that considered transaction costs showed that an investor would trade optimally at discrete intervals. The time interval between trades is a constant in models such as Duffie and Sun (1990), in which there is a fixed cost of observing the portfolio and making decisions, while it is state dependent in models such as Constantinides (1986) and Davis and Norman (1990) in which there are proportional costs of transacting.[15] In contrast to these models we consider a setting in which portfolio adjustment takes place at random intervals that depend on investor-specific shocks. Such behavior can be motivated by considering the problem of an investor whose wealth is allocated between a transactions account invested in a liquid asset and a portfolio of investment assets when there is a cost of transferring funds between the investment portfolio and the liquid asset. If the investor withdraws cash from his investment portfolio to meet random consumption needs, such as those generated by shocks to labor income, medical expenses, the need to move house, and repairs and purchases of major appliances and automobiles, he will be forced to liquidate at least a part of his investment portfolio at a time that will depend on the magnitude and timing of these shocks. Even investors who are accumulating cash for investment to finance their future consumption may experience shocks to their rate of liquid asset accumulation which induce randomly timed reallocations between the liquid asset and the investment portfolio. Significant periods between portfolio adjustments seem to be consistent with the empirical evidence, at least for individual investors. Ameriks and Zeldes (2004) report that 44% of a sample of TIAA-CREF investors made no portfolio reallocations over a period of 10 years. Surveys cited by Chien et al. (2012) found that 57%–60% of U.S. households owning equities had made no trades in the prior year, and 45% of surveyed Italian households had made no more than 1 trade in the previous year.

We abstract from the allocation decision between the transaction and investment accounts and the issue of optimal consumption behavior to focus on the pricing of risky assets. Further, to avoid the difficulties introduced by optimal hedging behavior in a dynamic setting we assume that investors participate in the market for investments for only a single investment period of uncertain length. Specifically, consider an overlapping generations economy in which each generation consists of a continuum of identical investors. Investors are assumed to trade only twice: when they enter the market and purchase assets from previous generations, and when they leave the market. This overlapping generations economy

[15] Abel, Eberly, and Panageas (2013) develop a model which encompasses both time-dependent and state-dependent portfolio adjustment.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

is intended to capture the fact that, as a result of transaction costs which we do not model explicitly, investors do not trade continuously or even at fixed discrete intervals, and that when they purchase an asset they are typically uncertain about how long they will hold it. We assume for simplicity that if any of an investor's portfolio is to be liquidated, then the whole portfolio will be liquidated. We assume further that when an investor enters the market she is uncertain of the date at which she will liquidate her portfolio: π_τ, $\tau = 1,\ldots,\tau^*$ denotes the probability that she will liquidate after τ periods. At the beginning of any period t, assets will be held by members of τ^* generations, those who purchased at times $t-\tau^*$ to $t-1$. All trading (except for shares issued as part of a dividend reinvestment plan) takes place between the new generation of investors and members of past generations who are forced to liquidate in the current period.

We consider a steady state equilibrium in which a new generation of investors with constant mass enters the market at each period, and the same mass of investors who have entered the market in previous periods liquidates. In the steady state a constant fraction of the market portfolio will be liquidated each period and will be purchased by the new generation of investors. Denote by μ the mass of investors liquidating and entering the market each period, and let $\boldsymbol{\pi}$ denote the $(\tau^* \times 1)$ vector of holding period probabilities with typical element π_τ. Of the generation who enter at time t, a mass $\mu\pi_\tau$ will liquidate at time $t+\tau$, $\tau = 1,\ldots,\tau^*$. Although the liquidation date is uncertain for each individual investor, the total mass of investors from generation t who will liquidate at $t+\tau$ is known.

Denote by $\mathbf{x}_t$ the $(\tau^* \times 1)$ vector of the mass of investors in period t who will be liquidating τ periods in the future, $\tau = 1,\ldots,\tau^*$, then

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mu\boldsymbol{\pi}, \tag{2}$$

where the $(\tau^* \times \tau^*)$ time shift matrix $\mathbf{A}$ is given by:

$$A = \begin{bmatrix} 01000\cdots\cdots \\ 00100\cdots\cdots \\ \cdots\cdots \\ 00\cdots\cdots 1 \\ 00\cdots\cdots 0 \end{bmatrix}.$$

Equation (2) expresses the fact that the mass of investors who will be liquidating in τ periods is equal to the sum of the mass of investors who in the previous period were going to liquidate in $\tau+1$ periods, $x_{\tau+1}$, and the mass of new investors who will liquidate in τ periods, $\mu\pi_\tau$. In the steady state $\mathbf{x}_t = \mathbf{x}, \forall t$, where $\mathbf{x}$, the steady state vector of the mass of investors liquidating at each horizon, $\tau = 1,\ldots,\tau^*$, must satisfy:

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mu\boldsymbol{\pi}. \tag{3}$$

Then

$$\mathbf{x} = (\mathbf{I}-\mathbf{A})^{-1}\mu\boldsymbol{\pi}. \tag{4}$$

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Carrying out the inversion, the ith element of the vector $\mathbf{x}$ is given by $x_i = \mu \Sigma_{k=i}^{\tau^*} \pi_k$. The turnover per period, TO, is then given by the proportion of investors who liquidate in the current period, which is the ratio of $x_1 = \Sigma_{k=1}^{\tau^*} \pi_k = 1$ to $\mathbf{j}'\mathbf{x}$:

$$\text{TO} = \frac{1}{\Sigma_{i=1}^{\tau^*}[\Sigma_{k=i}^{\tau^*}\pi_k]}. \tag{5}$$

There are N risky assets in the economy in addition to 1-period riskless bonds which are in zero net supply. The τ-period riskless rate at time t is:

$$R_{F,t}^{\tau} \equiv \Pi_{s=1}^{\tau}(1 + r_{F,t+s}) - 1,$$

where $r_{F,t}$ is the 1-period riskless rate. Asset j has price $p_{j,t}$ and pays a dividend $d_{j,t}$ at time t. Investors who do not liquidate their position in a given period are assumed to invest their dividends in a dividend reinvestment plan (DRIP) run by the issuer. Investors who participate in the DRIP receive at time t, *in lieu of* a dividend, $d_{j,t}/p_{j,t}$ shares for every share held (including shares previously issued by the DRIP). When the investor liquidates his holding at time $t+\tau$, all shares held in the DRIP are redeemed by the issuer at the current price $p_{j,t+\tau}$. Define the 1-period return on stock j by:

$$r_{j,t} = (p_{j,t+1} + d_{j,t+1} - p_{j,t})/p_{j,t}.$$

Then the τ-period return on stock j for an investor who buys shares at time t and holds them until time $t+\tau$, reinvesting dividends in the DRIP, is R_{jt}^{τ}, where $R_{j,t}^{\tau} = \Pi_{s=1}^{\tau}(1+r_{j,t+s})-1$. The number of shares of asset j, excluding DRIP shares, is denoted by $\bar{n}_j$.

Each period, the new generation of investors enters the market with aggregate wealth:

$$W_{t,0} \equiv \int_i w_{t,0} di = \mu w_{t,0},$$

where $w_{t,0}$ is the wealth of an individual investor which is used to purchase shares from the liquidating investors. Let $n_j(t)$ denote the number of shares of asset j bought by each investor of the new generation at time t. The mass of investors who liquidate in period t is given by the sum of investors who bought in period $t-1$ and liquidate after 1 period, $\mu\pi_1$, plus the mass of investors who bought in period $t-2$ and liquidate after 2 periods, $\mu\pi_2$, etc., where μ is the mass of investors who enter the market each period. Therefore the total number of non-DRIP shares of asset j brought to the market by investors who liquidate in period t is $\mu \Sigma_{s=1}^{\tau^*} \pi_s n_j(t-s)$. We consider a stationary equilibrium in which the number of shares purchased each period is the same: $n_j(t) = n_j(t') = n_j, \forall j, t$. Thus, the supply of asset j at each period from liquidating investors is μn_j, where we have used the fact that $\Sigma_{s=1}^{\tau^*}\pi_s = 1$. In equilibrium this is equal to the number of shares demanded by the new generation of investors entering the market with mass μ.

Since the number of shares of asset j purchased at each period is μn_j, and the total outstanding supply of shares is held, we have:

$$\mu n_j = \lambda \bar{n}_j, \quad j = 1, \ldots, N, \tag{6}$$

where λ is the turnover rate given by equation (5).

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Since all investors of a given generation are identical, there is no borrowing or lending and market clearing requires that the aggregate value of shares bought by the new generation be equal to the endowment of the new generation:

$$\Sigma_{j=1}^{N} \mu n_j p_{j,t} = W_{t,0}, \tag{7}$$

where n_j, the total number of shares of asset j bought by each member of the new generation, is determined by the optimization problem of the investors, to which we now turn.

The final wealth of an individual investor who enters the market at time t and liquidates after τ periods is $w_{t,t+\tau}$, where:

$$w_{t,t+\tau} = w_{t,0}(1 + R^{\tau}_{F,t}) + \Sigma_{j=1}^{N} n_j p_j (R^{\tau}_{j,t} - R^{\tau}_{F,t}). \tag{8}$$

Let $U(w)$ be the utility of wealth of the investor if liquidation occurs τ periods after purchase. Then, taking into account the stochastic liquidation date, the objective function of the investor is:

$$\max_{n_j} \Sigma_{\tau=1}^{\tau^*} \pi_\tau \mathrm{E}\left\{U\left(w_{t,0}(1 + R^{\tau}_{F,t}) + \Sigma_{j=1}^{N} n_j p_j (R^{\tau}_{j,t} - R^{\tau}_{F,t})\right) | \Phi(F_t)\right\}. \tag{9}$$

In expression (9) F_t denotes the information available at time t and $\Phi(F_t)$ denotes the possibly coarser information set used by investors. If investors use all of the available information then $\Phi(F_t) = F_t$ and the market is said to be efficient. As we have noted above, there is evidence that the prices of some securities respond only slowly to new market-wide information which suggests that investors are not using the full information set to price securities. The first-order condition for a maximum in expression (9) is:

$$\Sigma_{\tau=1}^{\tau^*} \pi_\tau \mathrm{E}\left\{U'(w_{t,t+\tau})(R^{\tau}_{jt} - R^{\tau}_{F,t}) | \Phi(F_t)\right\} = 0, \quad j = 1, \ldots, N. \tag{10}$$

Taking into account the market clearing conditions, equation (6), and recalling that in equilibrium there is no borrowing or lending, each investor purchases only a share of the market portfolio, such that $w_{t,t+\tau} = w_{t,0}(1 + R^{\tau}_{M,t})$, where

$$R^{\tau}_{M,t} = \Sigma_{k=1}^{N} \bar{n}_k p_{k,t} R^{\tau}_{k,t} / \Sigma_{k=1}^{N} \bar{n}_k p_{k,t} - 1$$

is the τ-period return on the market portfolio starting from time t. We assume that the utility function takes the form $U(w) = 1/(1-b)w^{1-c}$. Then $U'(w) = w^{-c}$ and substituting for $w_{t,t+\tau}$ in (10), we obtain the market equilibrium condition:

$$\Sigma_{\tau=1}^{\tau^*} \pi_\tau \mathrm{E}\left\{(1 + R^{\tau}_{M,t})^{-c}(R^{\tau}_{jt} - R^{\tau}_{F,t}) | \Phi(F_t)\right\} = 0, \quad j = 1, \ldots, N. \tag{11}$$

Finally, taking unconditional expectations in (11) and dropping the subscript t, we have:

$$\Sigma_{\tau=1}^{\tau^*} \pi_\tau \mathrm{E}\left\{(1 + R^{\tau}_{M})^{-c}(R^{\tau}_{j} - R^{\tau}_{F})\right\} = 0, \quad j = 1, \ldots, N. \tag{12}$$

Equation (12) is the exact version of the asset pricing model with a stochastic horizon for power utility. However, given the attention that has been directed at

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

the CAPM, we shall focus most of our empirical analysis on an approximate version of equation (12), which corresponds more closely to the CAPM. By taking a Taylor series expansion of $(1+R_M^\tau)^{-c}$ around the mean of $\bar{R}_M^\tau$:

$$(1+R_M^\tau)^{-c} \quad \approx \quad (1+\bar{R}_M^\tau)^{-c}(1-c(1+\bar{R}_M^\tau)^{-1}(R_M^\tau-\bar{R}_M^\tau)),$$

the equilibrium condition (12) may be approximated by:

$$\Sigma_{\tau=1}^{\tau^*}\pi_\tau^*\mathrm{E}\left\{[1-c_\tau(R_M^\tau-\bar{R}_M^\tau)](R_j^\tau-R_F^\tau)\right\} \quad = \quad 0, \tag{13}$$

where $\pi_\tau^*=\pi_\tau(1+\bar{R}_M^\tau)^{-c}/\Sigma_{s=1}^{\tau^*}\pi_s(1+\bar{R}_M^s)^{-c}$ and $c_\tau=c(1+\bar{R}_M^\tau)^{-1}$. Since $(1+\bar{R}_M^\tau)^{-c}$ is the marginal utility of wealth in the event of liquidation at time τ when the market return realization is equal to its expectation, $\pi_\tau^*>0$ is the marginal utility of (expected) wealth weighted probability of liquidation at time τ.

Equation (13) nests as special cases discrete time versions of the CAPM with different trading intervals. For example, $\pi_{\hat{\tau}}^*=1,\pi_\tau^*=0$, for $\tau\neq\hat{\tau}$ implies that the CAPM holds with a $\hat{\tau}$-period trading interval. Equation (13) suggests that if investors trade only infrequently then the simple CAPM will hold for a long holding interval, τ. In general, equation (13) represents a weighted average of different versions of the CAPM, which differ only in the definition of the holding period. Of course it is possible that the distribution of holding periods will differ for different securities; for example, Amihud and Mendelson (1986) suggest that securities with higher transaction costs will be held by investors with longer trading horizons. However, the assumption of different liquidation probabilities for different securities would complicate our model considerably. For example, if the whole of the portfolio is not liquidated at time τ then the investor's marginal utility will depend, not just on total wealth as is the case here, but also on the composition of the wealth portfolio. For this reason we stick with the simple formulation (13).

Using the definition of covariance, equation (13) can be written as:

$$\Sigma_\tau^{\tau^*}\pi_\tau^*\left\{(\mathrm{E}[R_j^\tau]-R_F^\tau)-c_\tau\mathrm{cov}(R_M^\tau,R_j^\tau)\right\} \quad = \quad 0. \tag{14}$$

Then

$$\Sigma_\tau^{\tau^*}\pi_\tau^*\left\{(\mathrm{E}[R_j^\tau]-R_F^\tau)-\lambda_\tau\frac{\mathrm{cov}(R_M^\tau,R_j^\tau)}{\mathrm{var}(R_M^\tau)}\right\} \quad = \quad 0, \tag{15}$$

where $\lambda_\tau=c_\tau\mathrm{var}(R_M^\tau)$. Finally, defining $\beta_{\tau_j}\equiv\mathrm{cov}(R_j^\tau,R_M^\tau)/\mathrm{var}(R_M^\tau)$,

$$\Sigma_\tau^{\tau^*}\pi_\tau^*\left\{\mathrm{E}[R_j^\tau]-R_F^\tau\right\} \quad = \quad \Sigma_\tau\pi_\tau^*\lambda_\tau\beta_{\tau j}. \tag{16}$$

Thus, the CAPM with stochastic liquidation leads to an equilibrium condition that expresses the sum of expected excess returns for different horizons weighted by the marginal utility of wealth weighted probability of liquidation as the same probability weighted sum of betas times the market risk premium for each horizon. A more complete description of the model is given in Appendix A. In what follows, we shall assume that the marginal utility of (expected) wealth is the same for each possible liquidation date. Then $\pi_\tau^*\equiv\pi_\tau$.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

V. Estimation and Testing of an Asset Pricing Model with Stochastic Liquidation

We estimate the model parameters and compute test statistics using the moment condition corresponding to equation (13) which can be written as:

$$\text{(17)} \qquad \mathrm{E}\left\{\Sigma_{\tau=1}^{\tau^*}\phi_\tau(1/\tau)(1-c_\tau R_m^\tau)(R_j^\tau - R_F^\tau)\right\} \quad = \quad 0,$$

where $\phi_\tau = \tau\pi_\tau$ and τ^* is the longest possible liquidation horizon. We write the moment condition in the form (17) so that the returns are all expressed on a monthly basis after the division by τ. We refer to ϕ_τ as the *horizon* weight to distinguish it from the *probability* weight, π_τ. In order to specify the horizon weights, ϕ_τ, flexibly and parsimoniously, we follow Ghysels, Santa-Clara, and Valkanov (2006) and define:

$$\text{(18)} \qquad \phi(\tau,\boldsymbol{\gamma}) \quad \equiv \quad \frac{f\left(\frac{\tau}{\tau^*},\gamma_1;\gamma_2\right)}{\sum_{\tau=1}^{\tau^*} f\left(\frac{\tau}{\tau^*},\gamma_1;\gamma_2\right)},$$

for $f(z,a,b) \equiv z^{a-1}(1-z)^{b-1}/\beta(a,b)$ and $\beta(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$, where $\Gamma()$ is the Gamma function. This specification allows the weights, which sum to 1 by construction, to take various patterns. For example, gradually decaying over time or hump-shaped. At the same time, the weighting function requires only 2 free parameters $\boldsymbol{\gamma} = (\gamma_1,\gamma_2)$.

To limit the number of parameters to be estimated, we set $\tau^* = 24$, allowing for only 7 possible liquidation dates (1, 4, 8, 12, 16, 20, and 24 months), and set the market risk aversion parameters, c_τ, equal to a constant, c, for all horizons so that the model has only 3 free parameters, $\boldsymbol{\theta} \equiv (\gamma_1,\gamma_2,c)$.

Define $h(\mathbf{x_t},\boldsymbol{\theta})$ as an $N \times 1$ vector, with jth element:

$$\left\{\Sigma_{\tau=1}^{\tau^*}\phi(\tau,\boldsymbol{\gamma})(1/\tau)(1-cR_{m,t}^\tau)(R_{j,t}^\tau - R_{F,t}^\tau)\right\},$$

where $\mathbf{x}_t$ is the vector of market, portfolio, and risk free returns starting at time t: $(R_{j,t}^1, R_{j,t}^\tau, j=1,.N, R_{m,t}^1, R_{m,t}^\tau, R_{F,t}^1, R_{F,t}^\tau)$.

Define the $N \times 1$ vector:

$$g_T(\theta) \quad \equiv \quad \frac{1}{T-\tau}\Sigma_{t=1}^{T-\tau} h(\mathbf{x_t},\boldsymbol{\theta}).$$

The parameters $\boldsymbol{\theta}$ are chosen to minimize the weighted sum of squared pricing errors given by the quadratic form $g_T(\theta)'W g_T(\theta)$, where W is a $N \times N$ weighting matrix. The identity matrix in the first step, and inverse of the variance-covariance matrix in the second step.

Define $d = \partial g_T(\theta)/\partial\theta$. Then

$$d \quad = \quad \frac{1}{T-\tau}\begin{bmatrix} \Sigma_{\tau=1}^{\tau^*}\phi_{\gamma_1}(\tau,\boldsymbol{\gamma})(1/\tau)(1-cR_{m,t}^\tau)(R_{j,t}^\tau - R_{F,t}^\tau) \\ \Sigma_{\tau=1}^{\tau^*}\phi_{\gamma_2}(\tau,\boldsymbol{\gamma})(1/\tau)(1-cR_{m,t}^\tau)(R_{j,t}^\tau - R_{F,t}^\tau) \\ -\Sigma_{\tau=1}^{\tau^*}\phi(\tau,\boldsymbol{\gamma})(1/\tau)R_{m,t}^\tau(R_{j,t}^\tau - R_{F,t}^\tau) \end{bmatrix},$$

where each row of d is a $1 \times N$ vector. If the matrix d is of full rank, then

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

$$\sqrt{T}(\widehat{\theta}_T - \theta) \quad \sim \quad N(0, V), \tag{19}$$

where $V = (d W^{-1} d')^{-1}$.

The sample variance-covariance ($\widehat{S}$) is calculated following Newey–West to take account of autocorrelations in long horizon returns:

$$\widehat{S} \quad = \quad \sum_{j=-k}^{k} \left(\frac{k - |j|}{k} \right) \frac{1}{T} \sum_{t=1}^{T} (u_t u'_{t-j}), \tag{20}$$

where $u_t \equiv h(\mathbf{x_t}, \theta)$ is the vector of pricing errors. We set the number of lags, k, equal to 30.[16] To test the overidentifying restrictions, we compute the J-statistic,

$$J \quad = \quad T g_T(\hat{b})' \widehat{S}^{-1} g_T(\hat{b}), \tag{21}$$

which follows a χ^2 distribution with degrees of freedom equal to the number of overidentifying restrictions, $(N - 3)$, in which N is the number of test assets.

We first estimate the model using the 25 FF size and book-to-market portfolios for the period 1926–2013. The market return is proxied by the CRSP value-weighted index return, and the risk-free interest rate is taken as the 1-month Treasury bill rate. The parameter estimates are presented in Panel A of Table 7. Panel B reports for each liquidation date the horizon weight, ϕ_τ, that is implied by the parameters reported in Panel A. Panel C reports the implied average horizon and turnover rate calculated from equation (5), as well as the probability weights, π_τ, which are also shown graphically in Graph A of Figure 2. We note first that the model, which has only 3 free parameters, is not rejected either for the full sample period or for either of the 2 subperiods by the standard J-statistic, although the p-values for 2 halves of the sample period are around 10%. Secondly, a likelihood ratio test is unable to reject the hypothesis of equality of the risk aversion parameter, c_τ, across horizons for the whole sample period, and for the subsamples.[17] The risk aversion parameter, c, which is highly significant, increases from 1.69 in the first half of the sample period to 2.48 in the second half, and the full sample estimate is 1.67.[18]

Both parameters of the weighting function are highly significant for all sample periods. For the full sample and 2 half periods, the estimated probability of liquidation after 24 months is 0, providing assurance that our arbitrary 24-month cutoff is not binding. Most significantly, as seen in Panel C of Table 7, the estimated probability weighting system assigns *zero* probability to a liquidation at the 1-month horizon for either the whole sample or for the first half. In addition, the estimated probability of a liquidation at the 4-month horizon is less than 4% in the full sample. In the first half of the sample period the probability of liquidation is 0 for the 1-, 4-, and 8-month horizons. In the second half of the sample period the probability of a liquidation after 4 months rises to 32.1% and after 8 months to 48.4%. In the first half of the sample period the modal liquidation horizon is 16

[16] Similar results were obtained using 24–36 lags.

[17] We set the cutoff between the 2 halves at 1963 to facilitate comparisons with the extant literature.

[18] FF ((1992), p. 433) report that there is no "obvious relation" between β and average returns over the period 1963–1990.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

months; this falls to 8 months in the second half, and for the whole sample period the modal liquidation horizon is 12 months. The probability weighted number of months to liquidation is 16.7 in the first half of the sample, 7.4 in the second half and 12.3 for the whole period. The implied annual turnover rate rises from 72% in the first half of the sample period to 161% in the second half, and for the whole sample is 98%. This reduction in the expected holding period and increase in turnover rate is consistent with casual empiricism which shows that turnover rates have increased, particularly since the abandonment of fixed commissions in May 1975. We explore this further in Section VII.

TABLE 7

Asset Pricing with Stochastic Liquidation

Table 7 reports GMM estimates of the moment conditions from equation (17). The system is estimated assuming that the probability of liquidation takes a specific form, as explained, below using monthly returns on the 25 FF size- and BM-sorted portfolios and the CRSP value-weighted index for the period 1926–2013. The moment condition is

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau X_\tau\right] = 0, \quad \text{where } w_\tau = w(\tau,\gamma) = \frac{f\left(\frac{\tau}{K},\gamma_1;\gamma_2\right)}{\sum_{i=1}^{K} f\left(\frac{i}{K},\gamma_1;\gamma_2\right)}$$

for $\tau=1$ to K, where $f(z,a,b)=z^{a-1}(1-z)^{b-1}/\beta(a,b)$ and $\beta(a,b)$ is based on the Gamma function, or $\beta(a,b)=\Gamma(a)\Gamma(b)/\Gamma(a+b)$, and $X_\tau \equiv (1-cR_m^\tau)(R_j^\tau - R_f^\tau)$. A likelihood ratio test is performed against the unrestricted model where each horizon is allowed to have a different risk premium, c_τ. The likelihood ratio, LR, is reported and the p-values are based on a χ^2 distribution with 6 *d.f.* t-statistics are reported in parentheses; standard errors of the γ coefficients are reported in italics, and p-values for the J-statistic and likelihood ratio are in square brackets. Panel A reports the estimation results. Panel B reports the horizon weights, $\phi(\tau_i)\equiv w_i$ and Panel C reports the probability of liquidation at each horizon date, $\pi(\tau_i)=(\phi(\tau_i)/\tau_i)/\Sigma_{j=1}^{K}\phi(\tau_j)/\tau_j$, the estimated mean horizon, and the implied turnover rate. MAE (RMSE) is the annualized mean absolute (root mean square) pricing error.

Panel A. Parameters

	1926–2013	1926–1962	1963–2013	1926–1947	1947–1968	1968–1988	1988–2013
γ_1	6.11	37.11	5.33	11.21	19.31	14.18	6.35
	0.67	*0.76*	*0.90*	*1.12*	*1.77*	*2.74*	*1.21*
γ_2	5.29	16.88	10.00	5.36	24.69	29.76	49.89
	0.43	*0.75*	*0.53*	*1.11*	*1.22*	*0.90*	*1.54*
c	1.67	1.69	2.48	0.94	3.38	1.86	3.06
	(6.43)	(8.94)	(11.27)	(4.27)	(10.24)	(3.71)	(2.39)
J	28.82	30.48	29.62	13.72	18.42	15.58	10.06
	[0.17]	[0.11]	[0.12]	[0.88]	[0.62]	[0.79]	[0.98]
LR	6.99	10.58	10.18	4.93	7.28	7.30	9.53
	[0.32]	[0.10]	[0.12]	[0.55]	[0.30]	[0.29]	[0.15]
MAE	0.015	0.014	0.022	0.040	0.015	0.036	0.024
RMSE	0.021	0.022	0.028	0.054	0.018	0.045	0.033

Panel B. Horizon Weights (ϕ_τ)

	Horizon (months)						
	1	4	8	12	16	20	24
1926–2013	0.000	0.011	0.156	0.392	0.352	0.088	0.000
	(0.00)	(0.32)	(0.57)	(2.81)	(1.21)	(0.52)	(0.01)
1926–1962	0.000	0.000	0.000	0.008	0.775	0.216	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(6.14)	(2.41)	(0.00)
1963–2013	0.001	0.173	0.521	0.270	0.034	0.001	0.000
	(0.05)	(1.28)	(8.79)	(2.01)	(0.25)	(0.00)	(0.00)
1926–1947	0.000	0.000	0.007	0.130	0.492	0.370	0.002
	(0.00)	(0.00)	(0.10)	(0.89)	(3.13)	(2.58)	(0.01)
1947–1968	0.000	0.000	0.250	0.727	0.023	0.000	0.000
	(0.00)	(0.00)	(1.66)	(4.28)	(0.27)	(0.00)	(0.00)
1968–1988	0.051	0.179	0.562	0.052	0.051	0.051	0.051
	(0.36)	(1.28)	(3.56)	(0.31)	(0.30)	(0.22)	(0.22)
1989–2013	0.288	0.698	0.003	0.002	0.002	0.002	0.002
	(1.92)	(4.75)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)

(continued on next page)

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 7 (continued)
Asset Pricing with Stochastic Liquidation

Panel C. Probability Weights (π_τ), Mean Horizon, and Turnover Rate

	Horizon (months)							Mean Horizon (months)	Turnover
	1	4	8	12	16	20	24		
1926–2013	0.000	0.034	0.240	0.401	0.270	0.054	0.000	12.3	0.98
	(0.00)	(0.32)	(0.57)	(2.81)	(1.21)	(0.52)	(0.01)		
1926–1962	0.000	0.000	0.000	0.011	0.808	0.180	0.000	16.7	0.72
	(0.00)	(0.00)	(0.00)	(0.00)	(6.14)	(2.41)	(0.00)		
1963–2013	0.010	0.321	0.484	0.168	0.016	0.000	0.000	7.4	1.61
	(0.05)	(1.28)	(8.79)	(2.01)	(0.25)	(0.00)	(0.00)		
1926–1947	0.000	0.000	0.014	0.177	0.504	0.303	0.002	16.4	0.73
	(0.00)	(0.00)	(0.10)	(0.89)	(3.13)	(2.58)	(0.01)		
1947–1968	0.000	0.000	0.335	0.650	0.016	0.000	0.000	10.7	1.12
	(0.00)	(0.00)	(1.66)	(4.28)	(0.27)	(0.00)	(0.00)		
1968–1988	0.287	0.251	0.393	0.024	0.018	0.014	0.012	5.6	2.15
	(0.36)	(1.28)	(3.56)	(0.31)	(0.30)	(0.22)	(0.22)		
1989–2013	0.622	0.376	0.001	0.000	0.000	0.000	0.000	2.2	5.57
	(1.92)	(4.75)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)		

We also estimate the stochastic liquidation model for each of the 4 quarters of the sample. The results, also reported in Table 7, show that the average horizon is greater than 1 month in all subperiods and the average horizon declines monotonically over time, which is consistent with the increasing turnover rate observed in the data. Only in the final quarter of the sample period 1989–2013, does the modal horizon become 1 month, as is assumed in most tests of asset pricing models. Even in this final quarter the mean horizon is over 2 months. The estimated risk aversion coefficient, c, which is 1.67 for the full sample, ranges from 0.94 to 3.38 for the quarterly estimates.

Graph A of Figure 3 plots the estimated horizon weighted average excess returns,

$$\Sigma_{\tau=1}^{\tau^*}\phi_\tau(1/\tau)\overline{(R_j^\tau - R_F^\tau)},$$

against the estimated horizon weighted expected returns from the model,

$$\Sigma_{\tau=1}^{\tau^*}\phi_\tau(1/\tau)c\overline{R_M^\tau(R_j^\tau - R_F^\tau)}.$$

With the exception of the smallest and second smallest firm low-BM-sorted portfolios, the points cluster fairly tightly about the line implied by the theoretical relation. The correlation between the horizon weighted average excess returns and the horizon weighted predicted excess returns is 0.64. This rises to 0.81 when the small firm low-BM-sorted portfolio is excluded.[19]

The improved fit of the stochastic liquidation model may be understood by referring back to the "beta" version of the model, equation (16). We saw in Panel B of Table 2 that when appropriately annualized, mean returns are fairly insensitive to the horizon. This means that the left-hand side of equation (16) is also fairly insensitive to the probability weights. On the other hand, Panel A of Table 2

[19] As Campbell and Vuolteenaho (2004) remark, "this small growth portfolio is well known to present a particular challenge to asset pricing models, for example the 3-factor model of FF (1993) which does not fit this portfolio well."

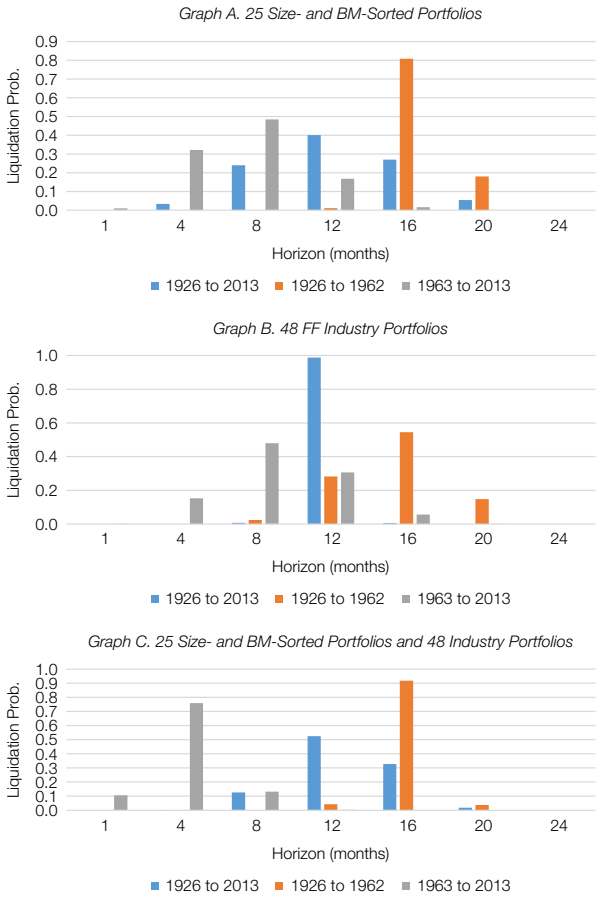
https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

FIGURE 2

Estimated Probability Distributions of Liquidation Dates for FF Size- and BM-Sorted Portfolios and Industry Portfolios

Figure 2 displays the estimated liquidation probabilities for the FF size- and BM-sorted portfolios and industry portfolios from 1926 to 2013 and subsamples of that period.



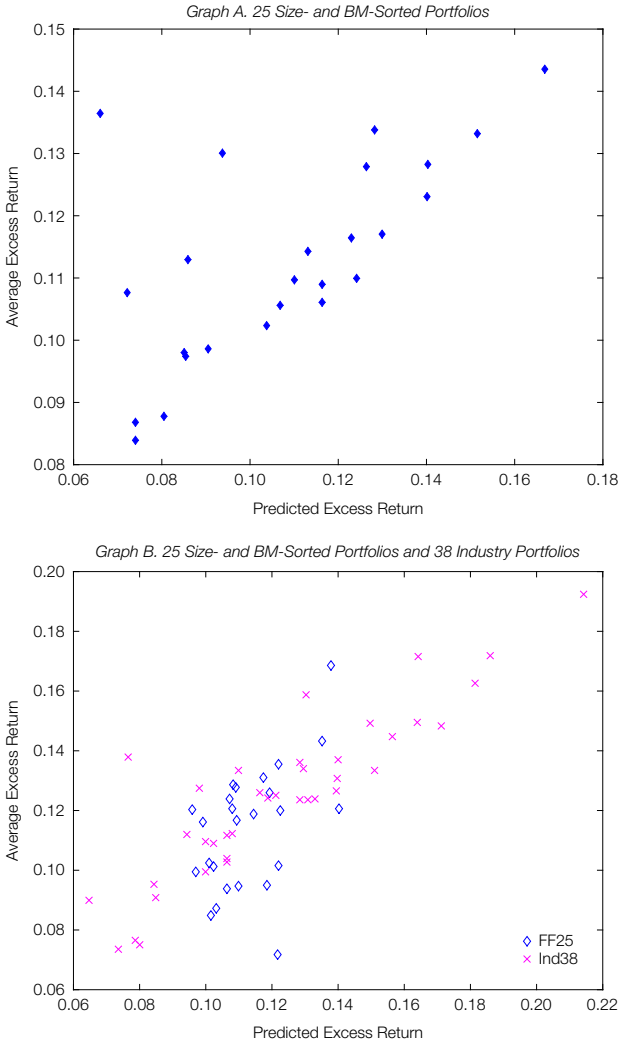
showed that the betas of some of the portfolios are very sensitive to the horizon over which they are measured. Thus, by choosing the probability weights, the stochastic liquidation model effectively chooses the horizon(s) whose betas line up best with the realized mean returns.

VI. Additional Tests of the Stochastic Liquidation Model

In this section we provide evidence on the power of our tests and report additional tests of the main results using different samples. To assess the power of the tests, we report the results of model tests on bootstrapped data, which preserve the co-moments of the data but impose serial independence on the returns. Specifically, we randomly sample 10,000 times from the vectors of market and

FIGURE 3
Actual and Stochastic Liquidation CAPM Predicted Returns

Figure 3 displays the horizon weighted annualized average excess returns and horizon weighted annualized predicted returns from the stochastic liquidation CAPM for the FF size- and BM-sorted portfolios and industry portfolios from 1926 to 2013.



portfolio returns so as to preserve the cross sectional dependence of returns and the mean values while eliminating any serial dependence. We then estimate the parameters and test statistics for each bootstrap sample that has the same number of observations as the full sample used in Table 7. Panel A of Table 8 reports the distributions of the estimated mean horizon, the estimated risk aversion coefficient, c, the J-statistic, and the associated p-value (from the χ^2 table) that are obtained from the bootstrap samples. For 97.5% of the bootstrapped samples, the

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 8

Bootstrap Statistics and Likelihood Ratio Tests

Panel A of Table 8 reports the distributions of the mean horizons, c estimates, J-statistics, and corresponding p-values estimated from 10 000 random samples of returns obtained by randomly sampling with replacement monthly vectors of the excess returns on the portfolios. Each random sample has the same number of observations as the full 1926–2013 sample reported in Table 7. Panel B reports likelihood ratio tests of a simple CAPM with a single horizon of τ^* months against the unrestricted stochastic liquidation model. The system is estimated using monthly returns on the 25 FF size- and BM-sorted portfolios and the CRSP value-weighted index for the period 1926–2013. The moment condition for the unrestricted model is

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau X_\tau\right] = 0, \quad \text{where } w_\tau = w(\tau,\gamma) = \frac{f\left(\frac{\tau}{K},\gamma_1;\gamma_2\right)}{\sum_{i=1}^{K} f\left(\frac{i}{K},\gamma_1;\gamma_2\right)}$$

for $\tau=1$ to K, where $f(z,a,b)=z^{a-1}(1-z)^{b-1}/\beta(a,b)$ and $\beta(a,b)$ is based on the Gamma function, or $\beta(a,b)=\Gamma(a)\Gamma(b)/\Gamma(a+b)$, and $X_\tau \equiv (1-cR_m^\tau)(R_j^\tau - R_F^\tau)$. For the τ^* fixed horizon model, $w_\tau=0$ for $\tau\neq\tau^*$, $w_\tau=1$ for $\tau=\tau^*$. The likelihood ratio is reported and the p-values based on χ^2 distribution with 2 *d.f.* are in parentheses. π_τ is the estimated probability of liquidation after τ months from the unrestricted model whose parameter estimates are shown in Table 7.

Panel A. Bootstrapped Horizon Weights and Test Statistics

	Mean Horizon (months)	c	J-Stat.	p-Value
2.5th percentile	1.07	3.17	29.73	0.125
5th percentile	1.09	3.01	30.46	0.108
25th percentile	1.09	2.45	33.12	0.060
50th percentile	1.12	2.01	34.53	0.043
75th percentile	1.94	1.68	35.83	0.032
95th percentile	2.80	1.21	37.77	0.019
97.5th percentile	7.96	1.08	38.10	0.018

Panel B. Likelihood Ratio Tests of Fixed Horizon Model for Different Horizons

	Fixed Horizon (months)						
	1	4	8	12	16	20	24
Log likelihood ratio	18.18	14.27	4.91	2.35	1.37	4.19	7.54
p-value	(0.00)	(0.00)	(0.09)	(0.31)	(0.50)	(0.12)	(0.02)
π_τ	0.00	0.01	0.16	0.39	0.35	0.09	0.00

J-statistic is greater than 29.7 and the p-value is less than 12.5%. For the actual data, the corresponding statistic was only 28.8 with a p-value of 17%. Only 0.8% of the bootstrap sample yield lower J-statistics than the statistic from the actual data. 75% of the bootstrap samples lead to a rejection of the model at the 6% significance level or better. These results suggest that our tests of the stochastic liquidation model do not lack power and show that the success of the model in pricing the 25 portfolios is highly dependent on the serial dependence structure of returns. This of course is not taken into account in standard asset pricing tests that use only monthly returns.

Similarly, the pattern of estimated horizon weights and resulting mean horizon reported in Table 7 is highly unlikely if the returns are serially independent. For example, 97.5% of the bootstrap samples yield a mean horizon shorter than 7.9 months and only 1.5% of the bootstrap sample yield a mean horizon longer than that estimated from the actual data, 12.3 months. It is clear that the parameter estimates and model fit that we obtain are highly unlikely if the time-series properties of the returns are ignored.

Panel B of Table 8 reports the results of likelihood ratio tests of a simple CAPM with a fixed τ^*-period horizon against the alternative of the stochastic liquidation model which allows for a range of possible liquidation horizons. The test assets are the 25 FF size- and BM-sorted portfolios for the period 1926–2013.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

The p-values show that the fixed horizon model is easily rejected for horizons of 1, 4, and 24 months. It is rejected at around the 10% level for horizons of 8 and 20 months, but is not rejected for horizons of 12 and 16 months. The p-values correspond closely to the probabilities of liquidation, π_τ, for the stochastic liquidation model shown in Table 7 and repeated here for convenience. However, it would be premature to conclude from this that the stochastic liquidation model is equivalent to a simple CAPM with a single fixed horizon. This is because, as seen in Panel C of Table 7 and Graph B of Figure 4, the liquidation probabilities and the mean horizon vary over time and, as we shall see in Section VII, this time-series variation in liquidation probabilities is related to the NYSE turnover rate.

A. Extending the Set of Test Assets

Lewellen, Nagel, and Shanken (2010) have drawn attention to the low power of tests of asset pricing models that use as test assets the Fama–French size- and BM-sorted portfolios because of the strong factor structure of the returns on these portfolios. Following their suggestion to expand the set of test assets to include other portfolios, we report the results of estimating the model on the 48 Fama–French industry portfolios, on the 25 Fama–French size and momentum portfolios, and then on both the 48 industry portfolios and the 25 size and BM-sorted portfolios. The results are shown in Tables 9 and 10.

TABLE 9
Pricing Industry and Momentum Portfolios

Table 9 reports GMM estimates of the moment conditions in equation (17). The system is estimated using monthly returns on first the 48 FF industry portfolios and the CRSP value-weighted index, and then the 25 FF size and momentum portfolios, for the period 1926–2013. Panel A reports the estimated parameters and test statistics and Panel B reports the probability weights for the Industry Portfolios. Panels C and D report the same statistics for the size and momentum portfolios. the same statistics for the size and momentum portfolios. t-statistics are reported in parentheses; standard errors of the γ coefficients are reported in italics, and p-values for the J-statistic are in square brackets. See Table 7 for more details.

Panel A. Industry Portfolios: Parameter Estimates

	γ_1	γ_2	c	J	MAE	RMSE
Full sample	46.72	49.97	2.05	36.41	0.014	0.018
	0.56	*0.66*	(7.84)	[0.36]		
1926–1962	11.96	7.00	1.27	20.88	0.017	0.023
	0.84	*0.81*	(5.35)	[0.93]		
1963–2013	6.00	8.71	1.89	19.10	0.014	0.020
	0.82	*0.99*	(6.83)	[0.98]		

Panel B. Industry Portfolios: Probability Weights (π_τ), Mean Horizon, and Turnover Rate

	Horizon (months)							Mean Horizon	
	1	4	8	12	16	20	24	(months)	Turnover
Full sample	0.000	0.000	0.007	0.988	0.006	0.000	0.000	12.0	1.00
	(0.00)	(0.00)	(0.15)	(5.27)	(0.21)	(0.00)	(0.00)		
1926–1962	0.000	0.000	0.025	0.282	0.545	0.148	0.000	15.3	0.79
	(0.00)	(0.00)	(0.52)	(1.64)	(9.64)	(1.21)	(0.00)		
1963–2013	0.002	0.153	0.480	0.307	0.057	0.001	0.000	9.1	1.32
	(0.00)	(0.01)	(2.12)	(2.03)	(0.01)	(0.00)	(0.00)		

(continued on next page)

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 9 (continued)
Pricing Industry and Momentum Portfolios

Panel C. Size and Momentum Portfolios: Parameter Estimates

	γ_1	γ_2	c	J	MAE	RMSE
Full sample	25.66	5.88	1.49	36.17	0.028	0.036
	0.19	*0.18*	(9.41)	[0.02]		
1926–1962	25.46	8.90	1.29	34.19	0.033	0.041
	0.04	*0.04*	(8.72)	[0.03]		
1963–2013	14.83	1.59	1.86	34.91	0.025	0.031
	0.02	*0.02*	(10.82)	[0.03]		

Panel D. Size and Momentum Portfolios: Probability Weights (π_τ), Mean Horizon, and Turnover Rate

	Horizon (months)							Mean Horizon	
	1	4	8	12	16	20	24	(months)	Turnover
Full sample	0.002	0.000	0.000	0.001	0.080	0.891	0.026	19.7	0.62
	(0.00)	(0.00)	(0.00)	(0.02)	(0.64)	(23.45)	(0.27)		
1926-1962	0.000	0.000	0.000	0.005	0.305	0.690	0.000	18.5	0.65
	(0.00)	(0.00)	(0.00)	(0.01)	(4.58)	(4.89)	(0.00)		
1963-2013	0.000	0.000	0.000	0.000	0.011	0.171	0.818	23.0	0.52
	(0.00)	(0.00)	(0.00)	(0.00)	(0.05)	(4.40)	(5.35)		

TABLE 10
Pricing 48 Industry Portfolios and 25 Size- and BM-Sorted Portfolios

Table 10 reports GMM estimates of the moment conditions in equation (17). The system is estimated using monthly returns on the 25 FF size- and BM-sorted portfolios and 48 industry portfolios, and the CRSP value-weighted index for the period 1926–2013. Panel A reports the stochastic liquidation model results. Panel B reports the probability weights. t-statistics are reported in parentheses, standard errors of the γ coefficients are reported in italics, and p-values for the J-statistic are in square brackets. See Table 7 for more details.

Panel A. Parameters

	γ_1	γ_2	c	J	MAE	RMSE
Full sample	11.00	9.33	1.92	36.21	0.013	0.017
	0.91	*0.87*	(6.87)	[0.99]		
1926–1962	36.15	19.42	1.29	27.35	0.015	0.023
	1.30	*1.61*	(4.03)	[1.00]		
1963–2013	5.234	19.90	2.37	21.55	0.021	0.027
	1.44	*1.45*	(9.76)	[1.00]		

Panel B. Probability Weights (π_τ), Mean Horizon, and Turnover Rate

	Horizon (months)							Mean Horizon	
	1	4	8	12	16	20	24	(months)	Turnover
Full sample	0.000	0.001	0.127	0.525	0.328	0.018	0.000	12.9	0.93
	(0.00)	(0.00)	(0.01)	(3.75)	(3.38)	(0.00)	(0.00)		
1926–1962	0.000	0.000	0.000	0.043	0.919	0.037	0.000	16.0	0.75
	(0.00)	(0.00)	(0.00)	(2.75)	(8.32)	(2.95)	(0.00)		
1963–2013	0.106	0.759	0.132	0.003	0.000	0.000	0.000	4.2	2.84
	(0.47)	(3.09)	(0.78)	(0.24)	(0.00)	(0.00)	(0.00)		

Panel A of Table 9 shows that when the test assets are the industry portfolios, the stochastic liquidation model is not rejected for either the whole sample period or for the 2 half-periods. Moreover the parameter estimates are similar to those reported in Table 7 for the size- and BM-sorted portfolio sample. Thus, for the whole sample period the point estimate of the risk aversion coefficient,

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

c, for the industry (size and BM-sorted) portfolios is 2.05 (1.67), and for the 2 half periods the point estimates are 1.27 (1.69) and 1.89 (2.48), respectively. The estimated patterns of probabilities of liquidation are also similar. Thus, the estimated mean horizons for the 3 sample periods are 12.0, 15.3, and 9.1 months for the industry portfolios, as compared with 12.3, 16.7, and 7.4 months for the size and book-to-market portfolios. Graph B of Figure 2 shows the implied probability distribution over the horizons.

When the 48 industry and 25 size and book-to-market portfolios are combined the results are similar, as seen in Table 10. A test for the equality of the coefficients across the 2 sets of portfolios does not reject the null of equality. Graph C of Figure 2 shows the estimated probability distribution of liquidation horizons. Graph B of Figure 3 shows the relation between (horizon weighted) actual and predicted returns. The industry portfolios have less dispersion in predicted returns than the size- and BM-sorted portfolios, reflecting their reduced dispersion in beta. Moreover, their returns exhibit greater dispersion about the predicted return, reflecting the lower level of diversification achieved by industry portfolios.

Panels C and D of Table 9 report the corresponding results for the 25 FF size and momentum portfolios. Consistent with prior models of rational asset pricing, the model cannot explain the momentum portfolios. The model is rejected at the 3% level or better for the whole sample period and for the 2 halves of the sample period, although the pricing errors of the stochastic horizon model are much smaller than those of the 1-month CAPM and FF 3-factor models. For the full sample, the mean absolute error for the 1-month CAPM (FF 3-factor model) is 0.045 (0.031) as compared to 0.028 for the stochastic liquidation model. We note that although momentum strategies are often referred to as high turnover strategies, the estimated turnover rate is below that reported for the 25 size- and BM-sorted portfolios in Table 7.

B. Power Utility

The stochastic liquidation model has been developed under the assumption of quadratic utility, which is of course quite restrictive, especially when returns are measured over many months. Therefore we also estimate a power utility version of the model, replacing the assumption of quadratic utility with the assumption that τ-period marginal utility is proportional to $(1+R_m^\tau)^{-c}$, where c is the coefficient of relative risk aversion.

The moment condition is then:

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau(\gamma_1,\gamma_2)(1+R_m^\tau)^{-c}(R_j^\tau - R_F^\tau)\right] \quad = \quad 0,$$

where $w_\tau(\gamma_1,\gamma_2)$ are the horizon weights.

The power utility version of the model was estimated using the FF 25 size- and BM-sorted portfolios, and the CRSP value-weighted index for the period 1926–2013. The results are reported in Table 11. While the model is not rejected, the p-value is lower than for the quadratic utility version reported in Table 7: 0.07 versus 0.17. The model predictions are also slightly worse, the mean absolute pricing error (MAE) is 0.020 as compared with 0.015 for the quadratic

TABLE 11
Power Utility

Table 11 reports GMM estimates of the stochastic liquidation model assuming power utility. The system is estimated assuming that the probability of liquidation takes a specific form, as explained below, using monthly returns on the 25 FF size- and BM-sorted portfolios, and the CRSP value-weighted index for the period 1926–2013. The moment condition is

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau X_\tau\right] = 0, \quad \text{where } w_\tau = w(\tau,\gamma) = \frac{f\left(\frac{\tau}{K},\gamma_1;\gamma_2\right)}{\sum_{i=1}^{K} f\left(\frac{i}{K},\gamma_1;\gamma_2\right)}$$

for $\tau=1$ to K, where $f(z,a,b)=z^{a-1}(1-z)^{b-1}/\beta(a,b)$ and $\beta(a,b)$ is based on the Gamma function, or $\beta(a,b)=\Gamma(a)\Gamma(b)/\Gamma(a+b)$, and $X_\tau \equiv (1+R_m^\tau)^{-c}(R_j^\tau - R_F^\tau)$. t-statistics are reported in parentheses, standard errors of the γ coefficients are reported in italics, and the p-value for the J-statistic is in square brackets. Panel A reports the estimation results. Panel B reports the probability of liquidation at each horizon date, $\pi(\tau_i)=(\phi(\tau_i)/\tau_i)/\Sigma_{j=1}^{K}\phi(\tau_j)/\tau_j$, where $\phi(\tau_i)=w_i$.

Panel A. Parameters

γ_1	γ_2	c	J	Mean Horizon	MAE	RMSE
9.99	7.13	1.96	31.09	14.0	0.020	0.027
0.58	*0.54*	(11.06)	[0.07]			

Panel B. Probability Weights

Horizon (Months)						
1	4	8	12	16	20	24
0.000	0.001	0.083	0.412	0.433	0.070	0.000
(0.00)	(0.01)	(0.21)	(2.25)	(3.67)	(0.15)	(0.00)

utility version. The estimated weight on the 1-month horizon continues to be 0 while the estimated mean horizon increases slightly from 12.3 to 14.0 months. In short, while the estimated probabilities for the different horizons are broadly similar, the quadratic utility version has a somewhat better fit than the power utility version of the model.

VII. A Generalized Stochastic Liquidation Model

Panel C of Table 7 shows a consistent reduction in the estimated mean horizon and a concomitant increase in the estimated turnover rate over the 4 quarters of the sample period. It is natural to ask therefore whether the probabilities of the different horizons are related in a systematic fashion to the actual turnover rate of stocks. To address this question, data on annual turnover on the NYSE were taken from the NYSE Web site, https://www.nyse.com. The 2 parameters of the probability distribution of liquidation dates, γ_1 and γ_2, were then modeled as functions of NYSE turnover:

$$\gamma_{1y} = \theta_1 + \theta_2 \mathrm{TO}_y, \quad \gamma_{2y} = \theta_3 + \theta_4 \mathrm{TO}_y, \tag{22}$$

where γ_{1y}, γ_{2y} are the parameters of the liquidation date probability function for investors entering the market in year y and TO_y is the NYSE turnover rate in year y.

The moment condition for the generalized model, corresponding to equation (13) for the basic model, is then derived as:

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau X_\tau\right] = 0, \tag{23}$$

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

where

$$w_{\tau,t} \equiv w(\tau,\gamma_t) = f\left(\frac{\tau}{K},\gamma_{1,t};\gamma_{2,t}\right)\Big/\sum_{i=1}^{K} f\left(\frac{i}{K},\gamma_{1,t};\gamma_{2,t}\right),$$

$f(z,a,b)$ is defined following equation (18) above, $\gamma_{1,t}$ and $\gamma_{2,t}$ are given by equation (22) for month t in year y, and $X_\tau \equiv (1-cR_m^\tau)(R_j^\tau - R_F^\tau)$.

The generalized model was estimated for the period 1926–2013, first using the 25 FF size- and BM-sorted portfolios, and then these 25 portfolios plus the 48 industry portfolios. The results are reported in Table 12. Panel A shows that the estimates of both θ_2 and θ_4 are highly significant, thus the probability weights do vary over time with the NYSE turnover rate. Comparing the results for the

TABLE 12
Time Varying Probabilities of Liquidation Horizon

Table 12 reports the parameter estimates for a generalized model of asset pricing with stochastic liquidation, in which the horizon probabilities depend on the turnover ratio for the New York Stock Exchange. The moment condition is

$$\mathrm{E}\left[\sum_{\tau=1}^{K} w_\tau X_\tau\right] = 0, \quad \text{where } w_{\tau,t} = w(\tau,\theta_t) = \frac{f\left(\frac{\tau}{K},\theta_{1,t};\theta_{2,t}\right)}{\sum_{i=1}^{K} f\left(\frac{i}{K},\theta_{1,t};\theta_{2,t}\right)}$$

for $\tau = 1$ to K, $f(z,a,b) = z^{a-1}(1-z)^{b-1}/\beta(a,b)$, $\gamma_{1,t} = \theta_1 + \theta_2 TO_t$, $\gamma_{2,t} = \theta_3 + \theta_4 TO_t$, and $\beta(a,b)$ is based on the Gamma function, or $\beta(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$, and $X_\tau \equiv (1-cR_m^\tau)(R_j^\tau - R_F^\tau)$. The sample period is 1926–2013. TO is the New York Stock Exchange Turnover Rate for the year. Implied TO is the model turnover rate for each year using the probabilities implied by the model estimates and the NYSE turnover rate for the year, averaged over all years. Mean horizon is the time-series average of the average horizon implied by the model each year using that year's NYSE turnover rate. Mean horizon (probability) weights are the time-series averages of the horizon (probability) weights implied by the model estimates. *t*-statistics are reported in parentheses, standard errors of the θ coefficients are reported in italics, and *p*-values for the *J*-statistic are in square brackets. See Table 7 for more details.

Panel A. Estimates

θ_1	θ_2	θ_3	θ_4	c	J	MAE	RMSE	Implied Turnover	Mean Horizon (months)
25 Size- and BM-Sorted Portfolios									
15.22	−1.45	−0.43	44.63	1.72	24.24	0.014	0.021	1.24	12.1
0.56	*0.14*	*0.03*	*1.34*	(−5.28)	[0.19]				
25 Size- and BM-Sorted Portfolios and 48 Industry Portfolios									
19.98	15.02	10.01	6.99	1.59	39.99	0.014	0.018	0.73	16.3
0.74	*0.51*	*0.43*	*0.17*	(−5.41)	[0.99]				

Panel B. Mean Horizon Weights

1	4	8	12	16	20	24
25 Size- and BM-Sorted Portfolios						
0.000	0.126	0.238	0.198	0.272	0.163	0.003
(0.36)	(4.22)	(7.06)	(8.86)	(9.70)	(6.60)	(1.72)
25 Size- and BM-Sorted Portfolios and 48 Industry Portfolios						
0.000	0.000	0.000	0.046	0.763	0.191	0.000
(0.07)	(0.12)	(0.72)	(1.45)	(11.68)	(5.55)	(0.42)

Panel C. Mean Probability Weights

1	4	8	12	16	20	24
25 Size- and BM-Sorted Portfolios						
0.000	0.150	0.237	0.202	0.267	0.141	0.002
(0.36)	(4.22)	(7.06)	(8.86)	(9.70)	(6.60)	(1.72)
25 Size- and BM-Sorted Portfolios and 48 Industry Portfolios						
0.000	0.000	0.000	0.062	0.781	0.156	0.000
(0.07)	(0.12)	(0.72)	(1.45)	(11.68)	(5.55)	(0.42)

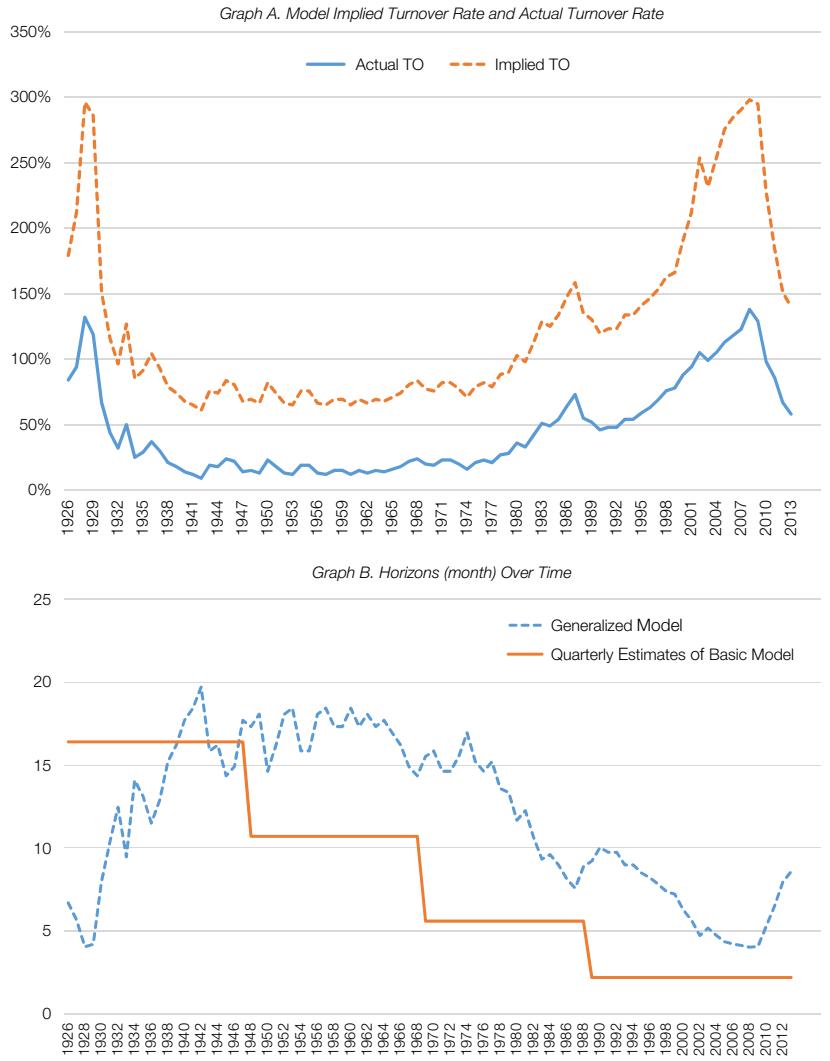
https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

25 size- and BM-sorted portfolios with those reported in Table 7 for the whole sample period, we see that the effect of allowing the horizon probabilities to depend on the turnover rate is to reduce the J-statistic from 28.82 to 24.24. Taking account of the 2 additional parameters to be estimated, the p-value improves from 0.17 to 0.19. The estimated risk aversion coefficient changes slightly from 1.67 to 1.72. Each year, the model implied turnover and mean horizon were calculated using the estimates of γ_{1t}, γ_{2t} implied by the estimates of the θ's and the

FIGURE 4

Model Implied Turnover and Horizons (1926–2013)

Graph A of Figure 4 plots of the model implied turnover rate against the actual turnover rate. Graph B plots the horizon (month) over time estimated with the baseline model using 4 quarters of the time-series data and estimated with the generalized model.



https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

NYSE turnover rate for that year. Graph A of Figure 4 plots the time series of the model implied turnover rate and the historical NYSE turnover rate. They track each other closely, with a correlation of 0.99. However, the mean model implied turnover rate of 124% (which compares with 98% for the basic model reported in Panel C of Table 7) is 2.7 times the historical average NYSE turnover rate of 46%. This difference in the means could be accounted for by an extension of the model to include a class of buy and hold investors who do not trade, and play no role in price-setting; we do not pursue this here. The mean horizon of 12.1 months for the generalized model compares with 12.3 months for the basic model. The pattern of mean horizon probabilities shown in Panel C is somewhat flatter than for the basic model. For both models the probabilities of horizons of 1 month or 24 months are essentially 0. Graph B of Figure 4 plots the time series of the mean horizon estimated from the generalized model for the whole sample period along with the estimates generated by repeated estimation of the basic model for the 4 quarters of the sample period. Both series capture the decline in the mean horizon in the second half of the sample period. However, the NYSE-turnover-driven generalized model does not capture the decline in the average horizon in the second quarter of the sample period that is apparent from the quarterly estimations of the basic model; the correlation between the two series is 0.51.

Table 12 also reports parameter estimates from the generalized model using the 25 Fama–French portfolios and the 48 industry portfolios as test assets. Again, both θ_2 and θ_4 are highly significant, highlighting the dependence on the turnover rate. Comparing with the basic model reported in Table 10, the J statistics and p values are almost identical. The estimated risk aversion coefficient changes from 1.92 to 1.59. The model implied turnover rate decreases from 93% to 73%, which is still much higher than the historical average turnover rate of 46%. The lower turnover rate is consistent with the increase in the average horizon from 12.9 months in the basic model to 16.3 months in the generalized model. Panel C clearly shows that the modal liquidation probability is at 16 months, compared with 12 months for the basic model.

VIII. A Comparison of Models

FF (2016) have proposed a new 5-factor model to price alternative sets of portfolios sorted by size, investment, and profitability. We use these 3 additional sets of portfolios as test assets to estimate the basic stochastic liquidation model and compare its performance with the CAPM and the FF 3- and 5-factor models. The sample period is 1963–2013 and the test assets are 25 size-profitability portfolios (MEOP), 25 size-investment portfolios (MEINV), and 25 profitability-investment portfolios (OPINV). Panel A of Table 13 reports the estimation results for these sets of portfolios along, with the original estimates for the size and book-to-market-sorted portfolios (MEBM) for the same sample period. The estimates of the risk aversion parameter, c, are highly significant and cluster in the interval from 2.11 to 2.57; the J-statistic tests fail to reject the model for any of the 4 sets of test assets. The estimated mean horizon ranges from 4.0 to 5.3 months for the new sets of test assets as compared with 7.4 months when the MEBM portfolios

are the test assets. Panel C shows that the estimated probability of liquidation in the first month is less than 2% for all 4 sets of test assets.

Panel B of Table 13 reports the results of estimating the CAPM, and the 3-factor and 5-factor models using GMM and the same 4 sets of test assets.

The moment condition for the 3-factor model, corresponding to equation (17) is:

$$\mathrm{E}\left\{(1 - b_{\mathrm{RMRF}}\mathrm{RMRF}^{\tau} - b_{\mathrm{SMB}}\mathrm{SMB}^{\tau} - b_{\mathrm{HML}}\mathrm{HML}^{\tau})(R_j^{\tau} - R_F^{\tau})\right\} \quad = \quad 0$$

and for the 5-factor model is:

$$\begin{aligned}\mathrm{E}\{&(1 - b_{\mathrm{RMRF}}\mathrm{RMRF}^{\tau} - b_{\mathrm{SMB}}\mathrm{SMB}^{\tau} - b_{\mathrm{HML}}\mathrm{HML}^{\tau} - b_{\mathrm{RMW}}\mathrm{RMW}^{\tau} \\ &-b_{\mathrm{CMA}}\mathrm{CMA}^{\tau})(R_j^{\tau} - R_F^{\tau})\} \quad = \quad 0,\end{aligned}$$

where RMRF, SMB, HML, RMW, and CMA are the 5 FF (2016) factors.

The CAPM is rejected for the MEBM and MEINV sets of portfolios and is close to rejection for the OPINV portfolios. The 3-factor model is rejected for the MEBM portfolios and is close to rejection for the MEOP and OPINV portfolios. Moreover, the 3-factor model yields inconsistent risk premium estimates for the size and value factor loadings across the different sets of test assets. Only the 5-factor model is not rejected for any of the 4 sets of portfolios. However, it has 5 parameters to be estimated in contrast to the 3 parameters of the stochastic liquidation model and, as Lewellen et al. (2010) point out, with 3 or 5 factors the asset pricing tests lack power. In terms of the mean absolute pricing error, the 5-factor model performs better for the MEBM portfolios, does about as well as the stochastic liquidation model for the MEOP portfolios, and roughly halves the MAE for the MEINV and OPINV portfolios. The stochastic liquidation model yields a smaller MAE than either the CAPM or 3-factor model for all 4 sets of portfolios. As discussed in Section V, the improvement of the stochastic liquidation model stems from its ability to select the horizon(s) over which betas are more closely aligned with average returns.

Overall, the stochastic liquidation model performs well relative to the other factor models, taking account of its parsimony and the fact that it was not designed to price any particular set of portfolios.

While the classical CAPM, and the stochastic liquidation extension presented in this paper, are based on the fiction of a representative investor, more recent asset pricing models have stressed the importance of intermediaries in asset pricing.[20] Prompted by these analyses while retaining the framework of the CAPM, we consider the performance of the stochastic liquidation model when the liquidation probabilities are exogenously set equal to the empirical liquidation probabilities of the institutional sector derived from 13F filings.[21] This is equivalent to testing the Euler equation restriction for a representative institution that holds the market portfolio and has the estimated liquidation probabilities. Panel A

[20] For theoretical models see He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014); for empirical evidence see Adrian, Etula, and Muir (2014), and He, Kelly, and Manela (2017).

[21] All institutional investment managers with more than $100 m of asset under management are required to file a quarterly 13F report detailing their asset holdings. The data are available as of 1980.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE 13
A Comparison of Models

Table 13 reports GMM estimates of the moment conditions for the stochastic liquidation model in equation (17), for the CAPM, and for the FF 3- and 5-factor models. The models are estimated using monthly returns on the FF size and book-to-market (MEBM), size and operating profitability (MEOP), size and investment (MEINV), and operating profitability and investment (OPINV) sorted portfolios for the period 1963–2013. Panel A reports the stochastic liquidation model results. Panel B reports results for the CAPM, FF 3-factor, and FF 5-factor models. Panel C reports the estimated probability weights, mean horizon, and turnover for the stochastic liquidation model. *t*-statistics are reported in parentheses, standard errors of the γ coefficients are reported in italics, and *p*-values for the *J*-statistic are in square brackets. See Table 7 for more details.

Panel A. Stochastic Liquidation Model

	γ_1	γ_2	c	J	Mean Horizon	MAE	RMSE	No. of Portfolios
MEBM	5.33	10.00	2.48	29.62	7.4	0.022	0.028	25
	0.90	*0.53*	(11.27)	[0.12]				
MEOP	11.10	36.12	2.35	21.16	5.0	0.017	0.022	25
	1.01	*1.13*	(6.66)	[0.45]				
MEINV	9.55	50.01	2.57	22.80	4.0	0.019	0.027	25
	0.94	*1.05*	(4.56)	[0.35]				
OPINV	11.10	34.42	2.11	23.58	5.3	0.020	0.027	25
	1.01	*1.27*	(5.21)	[0.31]				

Panel B. CAPM and FF Models

		b_{RMRF}	b_{SMB}	b_{HML}	b_{RMW}	b_{CMA}	J	MAE	RMSE
CAPM	MEBM	2.10					43.09	0.022	0.027
		(4.44)					[0.01]		
	MEOP	3.26					29.50	0.019	0.025
		(3.59)					[0.20]		
	MEINV	3.25					37.61	0.022	0.031
		(4.12)					[0.04]		
	OPINV	2.97					35.38	0.023	0.031
		(3.22)					[0.06]		
FF 3-factor	MEBM	1.17	0.62	4.59			43.58	0.018	0.027
		(1.96)	(0.74)	(4.97)			[0.00]		
	MEOP	3.37	−0.01	0.53			32.15	0.019	0.025
		(3.14)	(0.01)	(0.22)			[0.07]		
	MEINV	5.36	−1.89	14.11			26.74	0.024	0.028
		(6.31)	(1.51)	(8.75)			[0.22]		
	OPINV	2.38	3.76	1.84			33.91	0.023	0.032
		(2.23)	(2.39)	(0.92)			[0.05]		
FF 5-factor	MEBM	3.79	6.52	−2.36	15.28	15.93	20.33	0.010	0.013
		(3.87)	(5.90)	(−0.58)	(5.59)	(2.11)	[0.44]		
	MEOP	4.41	3.81	3.96	10.68	3.82	17.89	0.017	0.019
		(3.85)	(3.54)	(0.71)	(4.91)	(0.44)	[0.59]		
	MEINV	4.10	5.40	−2.63	13.63	15.78	19.99	0.010	0.014
		(5.79)	(4.60)	(0.85)	(4.24)	(4.34)	[0.46]		
	OPINV	3.55	8.99	−4.24	16.69	17.44	18.07	0.012	0.016
		(4.23)	(3.74)	(1.24)	(5.65)	(4.35)	[0.58]		

Panel C. Probability Weights (π_τ), Mean Horizon, and Turnover Rate

	Horizon (months)							Mean Horizon (months)	Turnover
	1	4	8	12	16	20	24		
MEBM	0.010	0.321	0.484	0.168	0.016	0.000	0.000	7.4	1.61
	(0.05)	(1.28)	(8.79)	(2.01)	(0.25)	(0.00)	(0.00)		
MEOP	0.000	0.752	0.247	0.001	0.000	0.000	0.000	5.0	2.40
	(0.00)	(3.77)	(2.82)	(0.05)	(0.00)	(0.00)	(0.00)		
MEINV	0.019	0.975	0.006	0.000	0.000	0.000	0.000	4.0	3.03
	(0.10)	(8.48)	(0.20)	(0.00)	(0.00)	(0.00)	(0.00)		
OPINV	0.000	0.679	0.319	0.002	0.000	0.000	0.000	5.3	2.27
	(0.00)	(3.67)	(3.45)	(0.08)	(0.00)	(0.00)	(0.00)		

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

of Table 14 reports the time-series averages of the estimated quarterly liquidation probabilities whose estimation is described in Appendix B. Our estimation horizon is truncated at 8 quarters and assets that are held for longer than 8 quarters are assumed to be sold at the end of the 8th quarter. After truncation, the mean liquidation horizon is approximately 5 quarters. Panel B reports the results of testing the Euler equation restriction by estimating the risk aversion coefficient, c, from the moment condition corresponding to condition (17) when the liquidation probabilities, π_τ^I, are given:

$$\mathrm{E}\left\{\Sigma_{\tau=1}^{8}\pi_\tau^I(1-cR_m^\tau)(R_j^\tau-R_F^\tau)\right\} \quad = \quad 0.$$

The estimation uses quarterly returns on the 25 FF size- and BM-sorted portfolios from 1980:Q1 to 2012:Q4. For the first estimation the liquidation probabilities vary over time, while in the second estimation they are held constant at their time-series means. It makes little difference whether or not the probabilities are held constant. In both cases the estimated risk aversion coefficient is close to 2.2, and the moment condition is not rejected. Panel C of Table 14 reports the result of estimating both the risk aversion coefficient and the liquidation probabilities of the stochastic liquidation model assuming that the probabilities follow a beta distribution as in Table 7. The estimate of c is the same as when the liquidation

TABLE 14

Asset Pricing with Stochastic Liquidation: Institutional Liquidation Probabilities

Table 14 reports the results of tests of the stochastic liquidation model for the period Jan. 1980–April 2012 using quarterly returns on the Fama–French size- and BM-sorted portfolios and *exogenous* quarterly liquidation probabilities corresponding to the liquidation probabilities of institutional portfolios constructed from 13F filings. The liquidation probabilities for each institution for a given quarter are constructed using the realized holding period for all purchases in that quarter and are value-weighted. The average institutional probabilities for a given quarter are then obtained by averaging across all reporting institutions using the value of purchases as weights. Panel A reports time-series means of the quarterly value-weighted average institutional liquidation probabilities. Panel B reports first GMM estimates of the risk aversion parameter c from the moment conditions:

$$\Sigma_{t=1}^{132}\left\{\Sigma_{\tau=1}^{8}\pi_{t,\tau}^I(1-cR_m^\tau)(R_j^\tau-R_F^\tau)\right\} \quad = \quad 0.$$

For the first equation, $\pi_{t,\tau}^I$ is the average institutional probability of liquidation in τ quarters for purchases made in quarter t. In the second equation the time varying probabilities $\pi_{t,\tau}^I$ are replaced by their time-series means, which are reported in Panel A. Panel C reports estimates for the stochastic liquidation model using quarterly returns for the same sample period when the liquidation probabilities are estimated, as in Table 7. t-statistics are reported in parentheses, standard errors of the γ coefficients are reported in italics, and p-values for the J-statistic are in square brackets. MAE (RMSE) is the annualized mean absolute (root mean square) pricing error.

Panel A. Mean Institutional Liquidation Probabilities from 13F Reports

Quarter								Mean Horizon Quarters
1	2	3	4	5	6	7	>7	
0.127	0.127	0.113	0.098	0.081	0.068	0.056	0.330	4.96

Panel B. Stochastic Liquidation Model with Institutional Liquidation Probabilities

	c	J	MAE	RMSE
Time-varying liquidation probabilities	2.20 (40.84)	19.93 [0.70]	0.018	0.024
Mean liquidation probabilities	2.16 (47.01)	19.47 [0.73]	0.018	0.024

Panel C. Stochastic Liquidation Model with Estimated Liquidation Probabilities

γ_1	γ_2	c	J	MAE	RMSE	Mean Horizon Quarters
2.03 *0.75*	2.00 *0.74*	2.16 (6.44)	16.26 [0.80]	0.020	0.028	4.25

probabilities are set equal to the average institutional liquidation probabilities, and the estimated mean horizon of 4.25 quarters is close to the average institutional horizon of 4.96 quarters (after truncation). In summary, the model is not rejected using the empirical institutional liquidation probabilities and the evidence is consistent with price-setting by a representative institutional investor who holds the market portfolio.

We also estimated a version of the Gilbert et al. (2014) opaque assets pricing model which assumes that investors learn with a delay about the sensitivity of opaque firm cash flows to an aggregate cash flow factor. Their model calibrations suggest an asset pricing model with an additional factor. This factor is constructed by going long on a portfolio of stocks that have a large difference between the beta measured using low frequency returns and the beta measured using high frequency returns and going short on a portfolio for which this difference is small. We implement the model using the difference between betas measured with 6-month returns and with 1-month returns. As shown in Appendix C, the model is rejected at the 1% level for the FF size- and BM-sorted portfolios for the sample periods 1926–2013 and 1963–2013. While our construction of the additional factor differs somewhat from that of Gilbert et al., our results are consistent with the results in their online appendix.

IX. Conclusion

In this paper we have shown that the monthly returns on the 25 Fama–French portfolios are not serially independent, and that this has major consequences for tests of the CAPM. Measured average returns and portfolio betas, and the relation between them, are all strongly affected by the length of the period over which returns are measured. We show that the ratio of long-horizon (12-month) betas to short-horizon (1-month) betas depends on such firm characteristics as the number of analysts following the firm and the seasonality of earnings, as well as firm size and book-to-market ratio. All of these variables, with the exception of seasonality, have been shown in previous papers to be associated with the efficiency of pricing of the firm's equity. We also find that for the FF size- and BM-sorted portfolios, the 12-month beta is significantly related to the SMB and HML betas, but is not significantly related to the 1-month market beta, given the other 2 betas. This raises the possibility that the empirical success of the FF 3-factor model is due to its ability to capture the long-horizon risk characteristics of the portfolios.

We develop an extended version of the CAPM with a continuum of identical investors who have a stochastic horizon or liquidation date. This gives rise to a relation between a probability weighted sum of expected returns over different horizons, and the same probability weighted sum of betas times market risk premia for the corresponding horizons. When we parameterize the probability weighting function, we find that the modified asset pricing model, which has only 3 parameters, is not rejected at conventional significance levels for either the whole sample period or for the 2 halves of the period. The correlation between the horizon weighted average excess returns and the horizon weighted predicted excess return is 0.64, and the correlation rises to 0.81 when the small firm low-BM-sorted portfolio is excluded.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

The model is also not rejected when it is fit to returns on 48 industry portfolios, or to the 73 portfolios that include both the size- and BM-sorted portfolios and industry portfolios, and we cannot reject the null hypothesis that the model parameters are the same across the 2 sets of portfolios. Moreover, the model fits relatively well for alternative portfolios sorted by investments and profitability. Finally the model is robust to the assumption of a power utility function.

Using returns on the 25 size- and BM-sorted portfolios, the estimated probability of the liquidation function implies that for the whole sample period, as well as for the first half of the sample, there is a 0 probability of liquidation after 1 month. This contrasts with the implicit assumption made in most conventional tests of asset pricing models that the probability of liquidation after 1 month is 100%. In the second half of the sample the probability of a liquidation after 1 month is 1%. The probability weighted number of months to liquidation decreases monotonically over the 4 quarters of the sample, from 16.4 months in the first quarter to 2.2 months in the last quarter. The reduction in the average liquidation horizon is consistent with the reduction in trading commissions and the increase in turnover that has followed the abandonment of fixed brokerage commissions in 1975 and the further development of competition in trading securities. To see whether the model captures the effects of increasing turnover rates on asset pricing, we generalize the model by allowing the 2 parameters of the distribution of liquidation dates to depend on the New York Stock Exchange (NYSE) turnover rate. We find that this improves the fit of the model without changing the estimated risk aversion coefficient significantly. The probability distribution of liquidation dates does depend significantly on the turnover rate. A higher NYSE turnover rate implies a shorter average horizon and a higher model implied turnover rate. The model implied turnover rate is almost perfectly correlated with the NYSE turnover rate but is over twice as high. This suggest that in addition to the price setting investors with high turnover rates that are captured by the model, there exists a class of buy and hold investors who do not trade and do not affect prices.

The model that we have developed and tested is extremely simple. It assumes that all the assets in the investor's portfolio are liquidated at the same time. A more realistic model would allow for the fact that, as Amihud and Mendelson (1986) have pointed out, investors with short horizons (high probabilities of liquidation) will be likely to hold more liquid securities, while investors with longer horizons will be likely to hold the more illiquid securities. It is beyond the scope of this paper to explore these interesting extensions which are left for future work. Some progress has been made in the paper by Beber et al. (2011), but at the expense of assuming that returns are IID. The focus of this paper has been on exploring the implications of non-IID returns.

Appendix A. Model

In this appendix, we provide a more complete description of the model underlying the CAPM with stochastic liquidation presented in Section IV. Formally, the model consists of:

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

TABLE A1
The Opaque Assets CAPM

Table A1 reports the results of GMM tests of the Gilbert et al. (2014) opaque assets 2-factor CAPM using monthly returns on the 25 FF size- and BM-sorted portfolios. The first factor is the CRSP value-weighted index, R_M, and the second factor is a $\Delta\beta$ factor, $R_{\Delta\beta}$, constructed by ranking firms on the difference between the 6-month beta and the 1-month beta with respect to the CRSP value-weighted index. The pricing kernel coefficients, b_β and $b_{\Delta\beta}$, are estimated for 1926–2013 and 1963–2013 using the moment condition:

$$\mathrm{E}\left\{(1 - b_\beta(R_M - R_F) - b_{\Delta\beta}R_{\Delta\beta})(R_j - R_F)\right\} \quad = \quad 0.$$

The J-statistic is a test of the overidentifying restrictions. t-statistics are reported in parentheses and p-values for the J-statistic are in square brackets.

	1926–2013		
	b_β	$b_{\Delta\beta}$	J
EW	2.41	−8.40	40.81
	(3.71)	(−2.98)	[0.01]
VW	3.82	−9.00	41.51
	(4.34)	(−3.65)	[0.01]
	1963–2013		
	b_β	$b_{\Delta\beta}$	J
EW	4.68	−11.28	43.35
	(4.43)	(−2.27)	[0.01]
VW	7.03	−11.31	45.81
	(4.53)	(−3.51)	[0.00]

i) A set of state variables $\mathbb{X}$ whose evolution under the investors' information sets is governed by the probability law $G(\mathbb{X}_t|\boldsymbol{\phi}(\mathbb{X}_{t-1}))$ where: $G(\mathbb{X}_t|\boldsymbol{\phi}(\mathbb{X}_{t-1})) = Pr(\mathbb{X} < \mathbb{X}_t|\boldsymbol{\phi}(\mathbb{X}_{t-1}))$, and $\boldsymbol{\phi}(\mathbb{X}_t)$ denotes the investors' information sets at time t.

ii) A set of N risky securities, $j = 1, \ldots, N$, paying dividend per share of $d_{j,t} = d_j(\mathbb{X}_t)$ at time t, and with prices $p_{j,t} = p_j(\mathbb{X}_t)$. Dividends are paid either in cash (to liquidating investors) or are reinvested in additional shares issued as part of a dividend reinvestment plan (DRIP) at the market price $p_{j,t}$. DRIP shares pay dividends like tradable shares and are redeemed by the issuer when the investor holding them liquidates. Define the 1-period return on security j by $r_{j,t} = (p_{j,t+1} + d_{j,t+1} - p_{j,t})/p_{j,t}$. The return on the security for an investor who purchases at time t and holds for τ periods, reinvesting dividends in the DRIP is then $R^\tau_{j,t} = \Pi^\tau_{s=1}(1 + r_{j,t+s}) - 1$. The number of shares of security j excluding DRIP shares is $\bar{n}_j, j = 1, \ldots, N$.

iii) A risk-free asset in zero net supply with interest rate in period t, $r_{F,t} = r_F(\mathbb{X}_t)$. The τ-period return on a rolled over riskless investment is $R^\tau_{F,t} \equiv \Pi^\tau_{s=1}(1 + r_{F,t+s}) - 1$.

iv) An infinite series of generations of investors, each consisting of a continuum of identical individuals with mass μ. Generation t enters the market at time t. Individual i has initial wealth $w_0(\mathbb{X}_t)$ and the aggregate initial wealth of the generation is $W_{t0} = W_0(\mathbb{X}_t) = \int_i w_0(\mathbb{X}_t)di = \mu w_o(\mathbb{X}_t)$. Each member of a generation invests for a single period of uncertain duration, at the end of which his portfolio is liquidated. DRIP shares are redeemed by the issuer and the remainder of the portfolio is sold to the new generation of investors. In particular, each investor who enters the market at time t liquidates at time $t + \tau$ with probability $\pi_\tau, \tau = 1, \ldots, \tau^*$. The mass of investors who liquidate in period t is given by $\mu\Sigma^{\tau^*}_{\tau=1}\pi_\tau = \mu$, that is, the sum of investors who bought in period $t - 1$ and liquidate after 1 period, $\mu\pi_1$, plus the mass of investors who bought in period $t - 1$ and liquidate after 2 periods, $\mu\pi_2$, etc. This is equal to the mass of the new generation of investors entering the market. Each investor of generation t maximizes a utility function of the form:

$$\max_{n_j} \Sigma^{\tau^*}_{\tau=1}\pi_\tau \mathrm{E}\left\{U\left(w_0(1 + R^\tau_{F,t}) + \Sigma^N_{j=1} n_j p_j (R^\tau_{j,t} - R^\tau_{F,t})\right)|\Phi(\mathbb{X}_t)\right\}, \tag{A-1}$$

where $\Phi(\mathbb{X}_t)$ denotes the information set of investors.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Thus, an equilibrium, if it exists, consists of: A set of asset demand functions, $n_{j,t} = n_j(\mathbb{X}_t)$, price functions, $p_{j,t} = p_j(\mathbb{X}_t)$, and an interest rate process, $r_{F,t} = r_F(\mathbb{X}_t)$, such that:

i) Asset demands satisfy the first-order conditions to the investor problem (A-1).

ii) Markets clear: $\int_i n_j(\mathbb{X}_t)di = \mu n_j(\mathbb{X}_t) = \bar{n}_j, j = 1, \ldots, N$.

Appendix B. Institutional Liquidation Probabilities

The institutional liquidation probabilities are calculated as follows. Let $\alpha_{i,j,t}$ denote the change in the fraction of ownership of stock j reported by institution i at the end of quarter t, calculated as the number of shares purchased/sold divided by the total number of shares outstanding. $\text{MV}_{j,t}$ is the market capitalization of stock j at the end of quarter t. All transactions are assumed to take place at the end of the quarter. Then at every quarter end, for every purchase ($\alpha_{i,j,t} > 0$), we look forward and identify when the institution cumulatively sells a fraction $\alpha_{i,j,t}$ of the same stock, and the number of quarters between the two transactions is denoted as the realized horizon, $\text{RLH}_{i,j,t}$. If the realized horizon is greater than 8 quarters it is set equal to 8 quarters. For each institution i and each quarter t, we compute the total transaction value, $\text{TV}_{i,t}(\text{rlh}) = \alpha_{i,j,t} \times \text{MV}_{j,t}$, for stocks that had a RLH equal to 1 to 8 quarters. The liquidation probability for horizon RLH is then $\text{Prob}_{i,t}(\text{RLH}) = \frac{\text{TV}_{i,t}(\text{RLH})}{\sum_{\text{RLH}=1}^{8} \text{TV}_{i,t}(\text{RLH})}$. Finally, we compute the value-weighted mean probabilities across all institutions to obtain the aggregate liquidation probabilities for all horizons from 1 to 8 quarters. The estimated average probabilities and derived mean horizons are presented in Table 14. The sample period is 1980–2012.

Appendix C. The Opaque Assets CAPM

We report asset pricing results for a 2-factor asset pricing model based on Gilbert et al. (2014). The βs are estimated with a 60-month rolling window. At the end of each year, from 1930 to 2013, we sort stocks into quintiles based on the difference between β's estimated with respect to the CRSP value-weighted index using 6-month returns and using 1-month returns[22] and construct the monthly returns on equal (EW) and value (VW) weighted portfolios. The $\Delta\beta$ factor is the difference between the return on the high and low $\Delta\beta$ portfolios. The kernel coefficients, b_β and $b_{\Delta\beta}$, are estimated for 1930–2013 and 1963–2013 using the moment condition:

$$\text{E}\left\{(1 - b_\beta(R_M - R_F) - b_{\Delta\beta}R_{\Delta\beta})(R_j - R_F)\right\} \quad = \quad 0.$$

The results, which are reported in Table A1, show that the overidentifying restrictions are rejected by the J-statistic at the 1% level or better.

References

Abel, A. B.; J. C. Eberly; and S. Panageas. "Optimal Inattention to the Stock Market with Information Costs and Transactions Costs." *Econometrica*, 81 (2013), 1455–1481.

Adrian, T.; E. Etula; and T. Muir. "Financial Intermediaries and the Cross-Section of Asset Returns." *Journal of Finance*, 69 (2014), 2555–2596.

Ameriks, J., and S. P. Zeldes. "How Do Household Portfolio Shares Vary with Age?" Working Paper, Columbia University (2004).

Amihud, Y., and H. Mendelson. "Asset Pricing and the Bid–Ask Spread." *Journal of Financial Economics*, 17 (1986), 223–249.

[22]Gilbert et al. (2014) formed the $\Delta\beta$ factor by sorting stocks based on the difference between betas calculated using quarterly data and daily data. We avoid the use of daily data for calculating the factor in order to prevent contamination by liquidity effects which are strong in daily data.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Beber, A.; J. Driessen; and P. Tuijp. "Pricing Liquidity Risk with Heterogeneous Investment Horizons." Working Paper, Tilburg University (2011).
Black, F.; M. C. Jensen; and M. Scholes. "The Capital Asset Pricing Model: Some Empirical Tests." In *Studies in the Theory of Capital Markets*, M. C. Jensen, ed. New York: Praeger (1972), 79–121.
Boguth, O.; M. Carlson; A. Fisher; and M. Simutin. "Horizon Effects in Average Returns: The Role of Slow Information Diffusion." *Review of Financial Studies*, 29 (2016), 2241–2281.
Brennan, M. J., and T. E. Copeland. "Beta Changes Around Stock Splits: A Note." *Journal of Finance*, 43 (1988), 1009–1014.
Brennan, M. J.; N. Jegadeesh; and B. Swaminathan. "Investment Analysis and the Adjustment of Prices to Common Information." *Review of Financial Studies*, 6 (1993), 799–824.
Brennan, M. J., and A. Wang. "The Mispricing Return Premium." *Review of Financial Studies*, 23 (2010), 3437–3468.
Brunnermeier, M. K., and S. Nagel. "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation." *American Economic Review*, 98 (2008), 713–736.
Brunnermeier, M. K., and Y. Sannikov. "A Macro-Economic Model with a Financial Sector." *American Economic Review*, 104 (2014), 379–421.
Calvet, L. E.; J. Y. Campbell; and P. Sodini. "Fight or Flight? Portfolio Rebalancing by Individual Investors." *Quarterly Journal of Economics*, 124 (2009), 301–346.
Campbell, J. Y., and T. Vuolteenaho. "Bad Beta, Good Beta." *American Economic Review*, 94 (2004), 1249–1275.
Chien, Y.; H. Cole; and H. Lustig. "Is the Volatility in the Market Price of Risk due to Intermittent Portfolio Rebalancing?" *American Economic Review*, 102 (2012), 2859–2896.
Cohen, R. B.; C. Polk; and R. Vuolteenaho. "The Price is (Almost) Right." *Journal of Finance*, 64 (2009), 2739–2782.
Constantinides, G. M. "Capital Market Equilibrium with Transaction Costs." *Journal of Political Economy*, 94 (1986), 842–862.
Davis, M. H. A., and A. R. Norman. "Portfolio Selection with Transaction Costs." *Mathematics of Operations Research*, 15 (1990), 676–713.
Dimson, E. "Risk Measurement When Shares are Subject to Infrequent Trading." *Journal of Financial Economics*, 7 (1979), 197–226.
Duffie, D., and T. S. Sun. "Transactions Costs and Portfolio Choice in a Discrete-Continuous Time Setting." *Journal of Economic Dynamics and Control*, 14 (1990), 33–51.
Fama, E. F., and K. R. French. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, 47 (1992), 427–465.
Fama, E. F., and K. R. French. "Common Risk Factors in the Returns on Stock and Bonds." *Journal of Financial Economics*, 33 (1993), 3–56.
Fama, E. F., and K. R. French. "Multifactor Portfolio Efficiency and Multifactor Asset Pricing." *Journal of Financial and Quantitative Analysis*, 31 (1996), 441–465.
Fama, E. F., and K. R. French. "Dissecting Anomalies with a Five-factor Model." *Review of Financial Studies*, 29 (2016), 69–103.
Fama, E. F., and J. D. MacBeth. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81 (1973), 607–636.
Ghysels, E.; P. Santa-Clara; and R. Valkanov. "Predicting Volatility: Getting the Most Out of Return Data Sampled at Different Frequencies." *Journal of Econometrics*, 131 (2006), 59–95.
Gilbert, T.; C. Hrdlicka; J. Kalodimos; and S. Siegel. "Daily Data is Bad for Beta: Opacity and Frequency-Dependent Betas." *Review of Asset Pricing Studies*, 4 (2014), 78–117.
Graham, J. R., and C. R. Harvey. "Theory and Practice of Corporate Finance: Evidence from the Field." *Journal of Financial Economics*, 60 (2001), 187–243.
Handa, P.; S. P. Kothari; and C. Wasley. "The Relation Between Return Interval and Betas: Implications for the Size Effect." *Journal of Financial Economics*, 23 (1989), 79–100.
Handa, P.; S. P. Kothari; and C. Wasley. "Sensitivity of Multivariate Tests of the Capital Asset-Pricing Model to the Return Measurement Interval." *Journal of Finance*, 48 (1993), 1543–1551.
Hawawini, H. "Why Beta Shifts as the Return Interval Changes." *Financial Analysts Journal*, 39 (1983), 73–77.
He, Z.; B. Kelly; and A. Manela. "Intermediary Asset Pricing: New Evidence from Many Asset Classes." *Journal of Financial Economics*, 126 (2017), 1–35.
He, Z., and A. Krishnamurthy. "Intermediary Asset Pricing." *American Economic Review*, 103 (2013), 732–770.
Huang, M. "Liquidity Shocks and Equilibrium Liquidity Premia." *Journal of Economic Theory*, 109 (2003), 104–129.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press

Jagannathan, R., and Y. Wang. "Lazy Investors, Discretionary Consumption, and the Cross-Section of Stock Returns." *Journal of Finance*, 62 (2007), 1623–1661.
Kamara, A.; R. A. Korajczyk; X. Lou; and R. Sadka. "Horizon Pricing." *Journal of Financial and Quantitative Analysis*, 51 (2016), 1769–1793.
Kothari, S. P.; J. Shanken; and R. G. Sloan. "Another Look at the Cross-Section of Expected Stock Returns." *Journal of Finance*, 50 (1995), 185–224.
Lee, C. F. "Investment Horizon and the Functional Form of the Capital Asset Pricing Model." *Review of Economics and Statistics*, 58 (1976), 356–363.
Lee, C. F.; C. Wu; and J. Wei. "The Heterogeneous Investment Horizon and the Capital Asset Pricing Model: Theory and Implications." *Journal of Financial and Quantitative Analysis*, 25 (1990), 361–376.
Levhari, D., and H. Levy. "The Capital Asset Pricing Model and the Investment Horizon." *Review of Economics and Statistics*, 59 (1977), 92–104.
Lewellen, J.; S. Nagel; and J. Shanken. "A Skeptical Appraisal of Asset Pricing Tests." *Journal of Financial Economics*, 96 (2010), 175–194.
Lintner, J. "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *The Review of Economics and Statistics*, 47 (1965), 13–37.
Lo, A., and C. MacKinlay. "When are Contrarian Profits due to Stock Market Overreaction?" *Review of Financial Studies*, 3 (1990), 175–205.
Longstaff, F. "Temporal Aggregation and the Continuous-Time Capital Asset Pricing Model." *Journal of Finance*, 44 (1989), 871–887.
Martellini, L., and B. Urosevic. "Static Mean-Variance Analysis with Uncertain Time Horizon." *Management Science*, 52 (2006), 955–964.
Merton, R. C. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, 41 (1973), 867–887.
Mossin, J. "Equilibrium in a Capital Asset Market." *Econometrica*, 34 (1966), 768–783.
Richard, S. F. "Optimal Consumption, Portfolio, and Life Insurance Rules for an Uncertain Lived Individual in a Continuous Time Model." *Journal of Financial Economics*, 2 (1975), 187–203.
Scholes, M. S., and J. Williams. "Estimating Betas from Non-Synchronous Data." *Journal of Financial Economics*, 5 (1977), 308–328.
Shanken, J. "On the Estimation of Beta-Pricing Models." *Review of Financial Studies*, 5 (1992), 1–33.
Sharpe, W. F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, 29 (1964), 425–442.
Wiggins, J. B. "Betas around Stock Splits Revisited." *Journal of Financial and Quantitative Analysis*, 27 (1992), 631–640.

https://doi.org/10.1017/S0022109018001412 Published online by Cambridge University Press