Excessive equity premium: the curious case of some low-beta portfolios

Miloš Božović*

University of Belgrade, Faculty of Economics and Business

Abstract

Consumption-based asset pricing can reconcile the high equity premium with the observed correlation between stock returns and consumption growth only if the investors' risk aversion is extremely high. In this paper, we focus on equity premia of portfolios obtained as univariate sorts of stocks based on their dynamic correlation with the market return. We show that the premia decrease with the size and stability of these dynamic correlations, whereas the return volatility and the correlation with consumption growth increase. As a result, the implied risk aversion parameters for portfolios with relatively low and unstable correlations with the market are greater than those obtained for the remaining stocks by a factor of three. These portfolios also have relatively low and unstable betas compared to the rest of the market. Thus, the highest equity premia seem to originate from portfolios more loosely coupled to systemic risk, which appears counterintuitive. A factor capturing the beta uncertainty premium captures the pattern and produces more reasonable values of the risk aversion parameters.

Keywords: asset pricing; risk premia; dynamic correlations; beta uncertainty

JEL classification: G12

*Email: milos.bozovic@ekof.bg.ac.rs

1 Introduction

Stocks pay a very high premium over bonds. The average return on the market portfolio of US common stocks between 1927 and 2022 was 11.84 percent. During the same period, a one-month US Treasury bill had an average annualized yield of 3.28 percent. This large spread of 8.55 percentage points is at the core of the equity premium puzzle of Mehra and Prescott (1985): stocks correlate with the economy, but the observed consumption volatility is too low to explain the risk associated with equity investments.[1] The considerable equity premium is persistent over more extended periods (Siegel, 2014) and is not unique to the United States but instead holds across all developed capital markets (Mehra and Prescott, 2003; Mehra, 2003). The long-term nature of the equity premium is of the essence for investment decisions, cost of capital estimation and numerous other practical and academic applications.

Expectedly, the last four decades saw many attempts to provide theoretical explanations behind the equity premium puzzle. Kocherlakota (1996), Koshy and Cochrane (1998), Mehra and Prescott (2003) and Mehra and Prescott (2006) are some of the standard entry points to the literature. Most can be regarded as modifications of the simple assumptions behind the consumption-based equilibrium model à la Lucas (1978) used in the original paper of Mehra and Prescott (1985). A considerable number of these modifications represent alternative forms of the utility function, mostly relying on time non-separability and similar adaptations of agents' preferences (see, for example, Abel, 1990; Constantinides, 1990; Ferson and Constantinides, 1991; Heaton, 1995; Campbell and Cochrane, 1999). Some apply modifications related to goods (Jagannathan and Wang, 1996; Lettau and Ludvigson, 2001) or states of nature (Epstein and Zin, 1989; Weil, 1989; Kandel and Stambaugh,

[1] The inflation-adjusted returns give similar results. The average annual *real* return on the US stock market between 1960 and 2021 was 8.30 percent, while the real return on US T-bills was 0.72 percent, resulting in an equity premium of 7.57 percentage points.

1991; Campbell, 1996; Tallarini, 2000). Other approaches analyze implications of market imperfections, such as financial constraints or market frictions (Bansal and Coleman, 1996; Palomino, 1996; Holmström and Tirole, 1998; McGrattan and Prescott, 2001). Alternative efforts rely on assumptions that go beyond the usual equilibrium approach, including the modification of state probabilities by introducing rare events having an extreme impact on the value of investment portfolio (Barro, 2005), or behavioral attempts such as those based on loss aversion (Benartzi and Thaler, 1995).

In this paper, we reexamine the US data searching for stocks that may have the most substantial impact on the high equity premium. We find the potential candidates by focusing on assets that produce a significant tilt in excess returns. The tilt originates from the discrepancies in the rolling-window correlations of these assets with the market excess return. The portfolios are constructed by a simple univariate sorting of stocks based on such dynamic correlations.

We demonstrate that the equity premia and the Sharpe ratios are substantially higher for portfolios exhibiting lower and less stable rolling-window correlations with the market. These portfolios also have a very low correlation with consumption growth. Consequently, their implied coefficients of relative risk aversion are extremely high, almost three times as high as for the market portfolio. These portfolios also have relatively low and unstable betas but much higher alphas than the rest of the market. We use the Intertemporal CAPM (ICAPM) framework to introduce a factor capturing the beta uncertainty premium that helps alleviate the problem.

The remainder of this paper is organized as follows. In Section 2, we provide the theoretical background and derive the equity premium from the consumption-based equilibrium model. In Section 3, we describe the construction of the sample portfolios and present the key results. Section 4 concludes.

2 Consumption-based model

The equity premium puzzle and the related paradoxes (for instance, Weil, 1989; Cochrane and Hansen, 1992) can have only one of the two mutually exclusive explanations: either the data are not representative enough, or the consumption-based model is an inadequate description of the reality. If the former is true, it would be reasonable to expect data accumulation to lead to progressively declining risk premia. However, this is not the case: historical averages were mainly increasing (Mehra and Prescott, 1985; Fama and French, 2002; Mehra and Prescott, 2003; Dimson et al., 2008).

We revisit the issues with the consumption-based model by starting from the Fundamental Equation of Asset Pricing in continuous time:

$$\mathbb{E}_t\left[\frac{\mathrm{d}(\Lambda_t V_t)}{\Lambda_t V_t}\right] = 0, \tag{1}$$

where $\Lambda_t = e^{-\delta t}u'(C_t)$ is the pricing kernel, with δ being the subjective discount rate and $u'(C_t)$ the marginal utility of a representative agent driven by the aggregate per-capita consumption C_t. The cumulative value process V_t captures the combined effect of capital gains and dividends. The infinitesimal relative change of the pricing kernel plays the role of an instantaneous stochastic discount factor, such that:

$$\mathbb{E}_t\left(\frac{\mathrm{d}\Lambda_t}{\Lambda_t}\right) = -r_t^f\,\mathrm{d}t, \tag{2}$$

where r_t^f represents the instantaneous risk-free rate. Through a simple manipulation of Equation (1), using Equation (2), we obtain:

$$\mathbb{E}_t\left(\frac{\mathrm{d}V_t}{V_t}\right) = r_t^f\,\mathrm{d}t - \mathrm{cov}_t\left(\frac{\mathrm{d}\Lambda_t}{\Lambda_t}, \frac{\mathrm{d}V_t}{V_t}\right). \tag{3}$$

Application of Itô's Lemma to the pricing kernel gives:

$$\frac{\mathrm{d}\Lambda_t}{\Lambda_t} = -\delta \mathrm{d}t + \frac{C_t u''(C_t)}{u'(C_t)} \frac{\mathrm{d}C_t}{C_t} + \frac{1}{2} \frac{C_t^2 u'''(C_t)}{u'(C_t)} \left(\frac{\mathrm{d}C_t}{C_t} \right)^2 . \tag{4}$$

Substituting this result into Equation (3) and casting it into discrete time, we find:

$$\mathbb{E}_t \left(R^e_{t+1} \right) = \eta_t \; \mathrm{cov}_t \left(\Delta c_{t+1}, R^e_{t+1} \right). \tag{5}$$

Here, R^e_{t+1} is the next-period excess return, Δc_{t+1} is the first difference of log consumption $c_{t+1} \equiv \ln C_t$ (i.e., approximately the growth rate of the aggregate per-capita consumption), while

$$\eta_t \equiv -\frac{C_t u''(C_t)}{u'(C_t)} \tag{6}$$

is the local curvature of the utility function.[2]

The intuition behind the result in Equation (5) is simple but profound: the risk premium, measured by the term on the left-hand side, originates from the interaction of asset returns with the investors' consumption. It scales with the local curvature of the utility function that commonly captures the relative risk aversion. Investors ultimately care about their consumption, not their portfolio returns or the performance of individual assets. Covariance measures how a slight change in return affects consumption patterns. Therefore, the average risk premium should be higher for assets that are generally highly correlated with consumption and should scale with the overall risk aversion of investors. However, the premium in Equation (5) is also time-varying, so we should expect it to be exceptionally high when the correlation and risk aversion increase.

[2] Note that the quadratic term in Equation (4) vanishes from the covariance in Equation (3) as it produces a variation of order higher than two.

An alternative way to write Equation (5) is

$$\mathbb{S}_t\left(R^e_{t+1}\right) = \eta_t \rho_t \sigma_t\left(\Delta c_{t+1}\right), \tag{7}$$

where

$$\mathbb{S}_t\left(R^e_{t+1}\right) \equiv \frac{\mathbb{E}_t\left(R^e_{t+1}\right)}{\sigma_t\left(R^e_{t+1}\right)}$$

denotes the conditional Sharpe ratio,

$$\rho_t \equiv \text{corr}_t\left(\Delta c_{t+1}, R^e_{t+1}\right)$$

is the conditional correlation between consumption growth and excess return and $\sigma_t\left(\Delta c_{t+1}\right)$ is the conditional volatility of consumption growth. For a power utility, $u'(C_t) = C_t^{-\gamma}$, the Equation (7) simplifies to

$$\mathbb{S}_t\left(R^e_{t+1}\right) = \gamma \rho_t \sigma_t\left(\Delta c_{t+1}\right). \tag{8}$$

The puzzle is that, historically, the annualized monthly Sharpe ratio for the market portfolio was around 0.44 while the annualized volatility of the monthly consumption growth was around 3 percent. These values impose the lower bound of $\gamma \gtrsim 14.5$ on the relative risk aversion coefficient. If we include the monthly correlation between the market excess return and consumption growth, which is around 0.18, we obtain the relative risk aversion coefficient of around 77. This value is, obviously, extremely high.[3] With the same observed covariance between the market excess return and consumption growth, the log utility ($\gamma = 1$) would imply the annualized equity premium of about ten basis points.

In the next section, we show that the equity premia and the Sharpe ratios are even higher for portfolios exhibiting lower and less stable dynamic correlations with the market.

[3] This version of the equity premium puzzle is often called the correlation puzzle according to Cochrane and Hansen (1992).

This effect originates from their low covariance with consumption growth.

3 Results

3.1 Data

The Arbitrage Pricing Theory implies that, under some reasonable assumptions, assets with higher Sharpe ratios will have a higher correlation with the market portfolio. Božović (2022a) provides evidence of a premium for holding stocks with smaller and less stable correlations with the market. Here, we exploit whether this tilt has some implications for the equity premium over a slightly shorter sample.

Our dataset combines consumption growth rates and stock returns. The US aggregate per-capita consumption levels were constructed by dividing monthly personal consumption expenditures (expressed in billions of US Dollars as a seasonally adjusted annual rate) by the total US population. Since consumption expenditure data is available with monthly frequency, whereas population data is annual, we used the same total population for each month in the respective year. Both series are available from the FRED Database of the Federal Reserve Bank of Saint Louis. We used Kenneth French's Database for the market portfolio return and the risk-free rate. The dataset covers the period between January 1959 and December 2020.

Our sample portfolios are created as univariate sorts on the five-year rolling window correlations with the market excess return. We use monthly returns on all NYSE, AMEX and NASDAQ stocks with share codes 10 or 11, traded between January 1954 and December 2020, available from CRSP to obtain these sorts. For each stock traded in months $t-60$ and $t-1$, we correlate its returns over the one-month US Treasury bill rate with the market excess return. Then, in each month t from January 1959 until December 2020, we sort all stocks based on these correlations and form equally-weighted portfolios by grouping the stocks into quintiles. This way, we obtain five series spanning the same period as consumption growth and the excess return of the market portfolio.

3.2 The excessive premium pattern

Table 1 shows key characteristics of monthly excess returns on portfolios sorted by five-year rolling-window correlation with the market excess return. Portfolios formed from stocks with low and less stable correlations with the market tend to have relatively higher average excess returns $\mathbb{E}\left(R^e_{i,t}\right)$ and lower volatility $\sigma\left(R^e_{i,t}\right)$ than stocks having higher and more stable correlations. They also have higher CAPM alpha, lower beta and higher (unconditional) Sharpe ratios than high-correlation stocks and the market portfolio (MKT). The alphas in the first and the second quintile are highly significant. Therefore, this pattern is very strong and strictly monotone. There is a 52 bp monthly premium when investing in a portfolio long on stocks with a low correlation with the market and short on stocks with a high correlation with the market. This long-short strategy carries a highly significant monthly alpha of about 1.09 percent.

The quintile portfolios obtained as univariate dynamic correlation sorts exhibit a substantial equity premium tilt. The lowest quintile is also one with the loosest coupling with the economy: it has a correlation coefficient of 0.13 with the growth rate of aggregate per-capita consumption Δc_t, compared to 0.19–0.22 for the other four quintiles or 0.18 for the entire market. The corresponding coefficient of relative risk aversion γ implied by the consumption-based model with a power utility function, Equation (8), is around 230, three times greater than the same parameter for the market portfolio. The implied coefficients also follow the monotone pattern across the quintiles, declining to 57 (a factor of four) for the stocks with high and stable correlations with the market.

Table 1: Equity premia and implied risk aversion for portfolios formed on dynamic correlations.
This table shows key characteristics of monthly excess returns on portfolios sorted by five-year rolling-window correlation with the market excess return. We provide the results for 1959/01–2020/12. For each quintile portfolio i, we report the average and the standard deviation of excess returns for all months, CAPM alpha and beta, Sharpe ratio for the excess returns, the unconditional correlations with consumption growth, and the implied risk aversion parameter. The last row shows the market excess return (MKT) results over the same observation period.
The asterisks indicate the usual significance levels: *** for significance at 1%; ** for significance at 5%; * for significance at 10%.

Portfolio i	$\mathbb{E}\left(R_{i,t}^e\right)$	$\sigma\left(R_{i,t}^e\right)$	α_i	β_i	$\mathbb{S}\left(R_{i,t}^e\right)$	$\text{corr}\left(\Delta c_t, R_{i,t}^e\right)$	γ_i
Low	1.23	4.69	0.95***	0.50***	0.91	0.13	229.7
2	0.99	5.40	0.44***	0.97***	0.64	0.20	99.2
3	0.85	5.87	0.19*	1.17***	0.50	0.22	73.5
4	0.81	6.35	0.06	1.33***	0.44	0.20	70.7
High	0.71	7.16	−0.15*	1.53***	0.34	0.19	57.0
MKT	0.57	4.43	0.00	1.00***	0.44	0.18	77.4

Table 2 shows the characteristics of rolling-window alphas and betas for monthly excess returns on quintile portfolios sorted by dynamic correlation with the market excess return. For each portfolio, we report the average and the standard deviation of CAPM alpha and beta estimated over the five-year rolling window. The average alpha, $A(\alpha_i)$, is relatively high for portfolios with low and less stable correlations with the market and equals almost one percent per month. It declines as we move through the dynamic correlation quintiles, reaching -9 basis points in the highest quintile. The alphas are also more stable for the higher quintiles, which we can see from their moving-window standard deviations $\sigma(\alpha_i)$ shown in the second column of Table 2.

On the other hand, the average beta, $A(\beta_i)$, increases with dynamic correlation with the market portfolio return, taking values between 0.49 and 1.49. Thus, it follows a pattern

Table 2: Stability of alphas and betas for portfolios formed on dynamic correlations. This table shows the characteristics of rolling-window alphas and betas for monthly excess returns on portfolios sorted by five-year rolling-window correlation with the market excess return. We provide the results for 1959/01–2020/12. For each portfolio, we report the average and the standard deviation of CAPM alpha and beta estimated over the five-year rolling window.

Portfolio i	$A(\alpha_i)$	$\sigma(\alpha_i)$	$A(\beta_i)$	$\sigma(\beta_i)$
Low	0.96	0.62	0.49	0.28
2	0.46	0.60	0.94	0.18
3	0.21	0.58	1.14	0.16
4	0.09	0.54	1.30	0.15
High	−0.09	0.39	1.49	0.23

similar to that presented in Table 1. The lowest quintile also has the least stable betas, with a moving-window standard deviation $\sigma(\beta_i)$ of 0.28, notably higher than the other four quintiles. The fact that the average return decreases while beta increases across dynamic correlation quintiles is puzzling and counterintuitive (see Figure 2). One would expect that beta should be able to capture at least the first-order effect in expected returns. Here, we find the opposite, resulting in a negative slope of the Security Market Line of these five portfolios. This result hints that the high premium of lower quintiles originates from the instability of correlations—or, equivalently, from uncertainty in betas—rather than their *levels*. It is also consistent with some earlier studies: Fama and French (1997) argue that uncertainty about genuine risk factors and relatively large variations of the factor loadings over time lead to substantial standard errors in betas, while Pastor and Stambaugh (1999) find that beta uncertainty is more critical for pricing than model uncertainty.

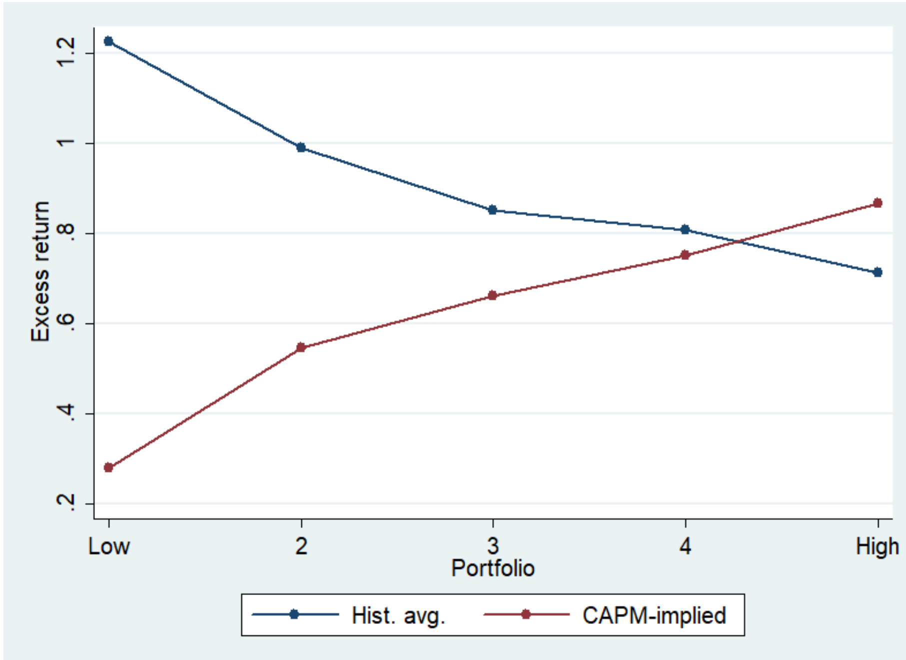

Figure 1: Excess returns for quintile portfolios obtained as univariate sorts based on five-year rolling-window correlations with the market portfolio: true averages compared to CAPM predictions.

3.3 Capturing beta uncertainty

To account for beta uncertainty, we introduce a state variable X_t that increases with the good news about future consumption driven by a decrease in the instability of correlations with the market portfolio. Using the ICAPM logic of Merton (1973), we define the value function as:

$$J(W_t, X_t) = \max_{\{C_t\}} \mathbb{E}_t \left[\int_{s=0}^{\infty} e^{-\delta s} u\left(C_{t+s}\right) \, \mathrm{d}s \right], \tag{9}$$

subject to the budget constraints on wealth W_t and given X_t. The value function represents the achieved level of happiness and is thus increasing in both W_t and X_t.

The envelope condition gives

$$\frac{\partial J(W_t, X_t)}{\partial W_t} = u'(C_t). \tag{10}$$

It sets equal the investor's marginal effects of a dollar spent and consumed. Applying it to the pricing kernel, using Itô's Lemma and discretizing Equation (3), we can generalize Equation (5) to:

$$\mathbb{E}_t\left(R^e_{t+1}\right) = \left(-\frac{W_t J_{WW}}{J_W}\right)\text{cov}_t\left(R^e_{t+1}, R^W_{t+1}\right) + \left(-\frac{X_t J_{WX}}{J_W}\right)\text{cov}_t\left(R^e_{t+1}, \Delta x_{t+1}\right), \tag{11}$$

where the subscripts to J denote the partial derivatives of the value function, R^W_{t+1} is the return on the representative investor's "wealth portfolio" and Δx_{t+1} is the change in the log state variable, $x_{t+1} \equiv \ln X_{t+1}$. The coefficients multiplying the covariance terms in Equation (11) measure the investor's relative risk aversion and the relative aversion to changes in the state variable. This second coefficient reacts to any change in the stability of correlations and betas that may impact the investor's wealth.

The beta representation of Equation (11) is:

$$\mathbb{E}_t\left(R^e_{t+1}\right) = \beta_W \lambda_{W,t} + \beta_X \lambda_{X,t}. \tag{12}$$

Therefore, Equation (12) is a cross-sectional form of a two-factor model, with return on the wealth portfolio and the correlation instability variable as risk factors. The corresponding factor loadings are β_W and β_X, while the prices of risk are given by:

$$\lambda_{W,t} = \left(-\frac{W_t J_{WW}}{J_W}\right)\text{var}_t\left(R^W_{t+1}\right) \tag{13}$$

$$\lambda_{X,t} = \left(-\frac{X_t J_{WX}}{J_W}\right)\text{var}_t\left(\Delta x_{t+1}\right). \tag{14}$$

We create proxies for the factors to test the predictions of Equation (11). For the wealth portfolio return, we use the market portfolio excess return. For the correlation instability variable, we use the return on a portfolio long on the first and short on the fifth quintile of the univariate correlation sort.

Table 3 shows the results of time-series regressions for each of the quintiles. The regression produces insignificant alphas, unlike CAPM (cf. Table 1). The correlation instability proxy X is highly significant and lifts the strong decreasing pattern of CAPM market betas. The regressions also explain a high fraction of variations in portfolio returns. The model performance is better in higher quintiles (81–90 percent in quintiles 3–5 and 77 percent in quintiles 1 and 2), which is expected given that the quintile portfolios were constructed to capture progressively stronger dynamic correlation with one of the regressional factors.

We can visualize the regression results from Table 3 in Figure 2. The near-constant pattern of market betas is apparent. Its combined contribution with the second factor—the beta uncertainty proxy—helps capture the cross-sectional average excess returns relatively well.

The mild variation in market betas also implies that the relative risk aversion coefficients will be balanced across the quintiles. The strong tilt implied by the standard consumption-based model observed in Table 1 is now picked up by variations in the state variable X. We run a cross-sectional regression of average excess returns on beta estimates from the time-series regression to assess the risk aversion parameters. We report these results in Table 4. The unconditional prices of risk are now obtained as regressional coefficients. We find $\lambda_M = 0.53$ and $\lambda_X = 0.51$, both highly significant. The adjusted R^2 of 0.995 indicates a near-perfect fit, consistent with insignificant intercepts in the time-series regressions.

Table 3: Time-series regression for a two-factor model.
The table summarizes the results of regressions of excess returns on five dynamic correlation portfolios on the market excess return and the correlation instability factor. For each quintile i obtained from the univariate sort of stocks based on their five-year rolling window correlation with the market, we run the following time-series regression:

$$R^e_{i,t} = \alpha_i + \beta_{iM} R^e_{M,t} + \beta_{iX} R^e_{X,t} + \varepsilon_{i,t},$$

where $R^e_{i,t}$ is the excess return on portfolio i; $R^e_{M,t}$ is the excess return on the market portfolio, acting as a proxy for the investor's wealth portfolio; $R^e_{X,t}$ is the excess return on a portfolio long on the first quintile ($i = 1$) and short on the fifth quintile ($i = 5$), acting as a proxy for the correlation instability variable; $\varepsilon_{i,t}$ is the usual regression residual. The results show the estimates of regressional coefficients and the adjusted R^2.
The asterisks indicate the usual significance levels: *** for significance at 1%; ** for significance at 5%; * for significance at 10%.

Portfolio i	α_i	β_{iM}	β_{iX}	R^2_{adj}
Low	0.01	1.38***	0.85***	0.77
2	−0.10	1.48***	0.50***	0.77
3	−0.10	1.45***	0.27***	0.81
4	−0.05	1.42***	0.09***	0.86
High	0.01	1.38***	−0.15***	0.90

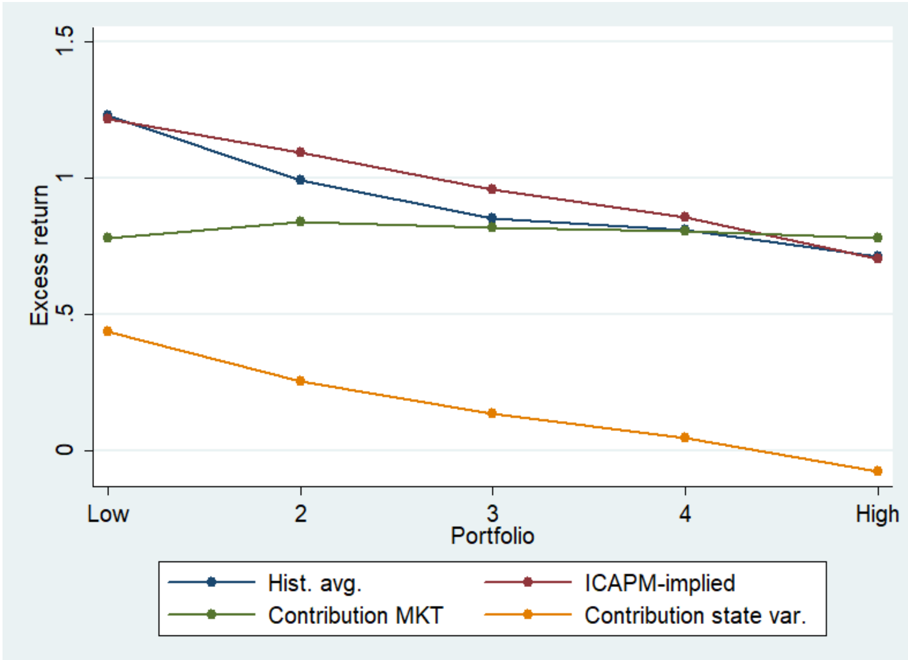

Figure 2: Excess returns for quintile portfolios obtained as univariate sorts based on five-year rolling-window correlations with the market portfolio: true averages compared to two-factor ICAPM predictions. The figure also shows the contributions of the market portfolio and the state variable that captures the uncertainty in betas.

Since both factors represent tradable portfolios, equalities

$$\mathbb{E}\left(R^e_{M,t}\right) = \lambda_M + \lambda_X \beta_{MX} \tag{15}$$

$$\mathbb{E}\left(R^e_{X,t}\right) = \lambda_M \beta_{MX} + \lambda_X. \tag{16}$$

must hold. Eliminating β_{MX} from Equation (15) and (16), we find

$$\lambda_M \left[\mathbb{E}\left(R^e_{M,t}\right) - \lambda_M\right] = \lambda_X \left[\mathbb{E}\left(R^e_{X,t}\right) - \lambda_X\right]. \tag{17}$$

The values of the left-hand-side and the right-hand side of Equation (17), obtained from the results of the cross-sectional regressions shown in Table 4 and the average factor returns, are very close (0.02 and 0.00, respectively), with well-overlapping 95-percent confidence intervals.

The implied risk aversion parameters are obtained by analogy with the Equations (13) and (14), dividing the unconditional prices of risk λ_M and λ_X by the unconditional variance of their corresponding factors, i.e.:

$$\eta_M = \frac{\lambda_M}{\mathrm{var}\left(R^e_{M,t}\right)} \tag{18}$$

$$\eta_X = \frac{\lambda_X}{\mathrm{var}\left(R^e_{X,t}\right)}. \tag{19}$$

The standard errors for the relative risk aversion parameters are obtained from the standard errors of prices of risk via simple error propagation rules. The 95-percent confidence intervals for the risk aversion parameters implied by these results are $\eta_M \in [2.45, 2.99]$ and $\eta_X \in [0.92, 1.79]$, which seems plausible.

Table 4: Cross-sectional regression for a two-factor model.
The table summarizes the results of regressions of average excess returns on the estimates of factor loadings obtained from the time series regression, Table 3. We run:

$$\overline{R^e_{i,\cdot}} = \lambda_M \widehat{\beta}_{iM} + \lambda_X \widehat{\beta}_{iX} + a_i,$$

where $\overline{R^e_{i,\cdot}}$ are the time-averaged excess returns for each quintile i; $\widehat{\beta}_{iM}$ is the estimate of the market portfolio factor loading; $\widehat{\beta}_{iX}$ is the estimate of the correlation instability factor loading; a_i is the cross-sectional regression residual. The results show the estimates of prices of risk λ_k, here obtained as regressional coefficients, their standard errors and the adjusted R^2. We also report the values $\lambda_k \left[\mathbb{E}\left(R^e_{k,t}\right) - \lambda_k\right]$ implied by the average factor returns. These values should be the same for $k = M$ and $k = X$, Equation (17).
The implied relative risk aversion parameters η_k are obtained by dividing the prices of risk by the variance of the corresponding factor. The standard errors for the risk aversion parameters are obtained via simple error propagation rules.
The asterisks indicate the usual significance levels: *** for significance at 1%; ** for significance at 5%; * for significance at 10%.

Parameter	Market portfolio, M	Corr. instability proxy, X
Price of risk, λ_k	0.53***	0.51***
Standard error for λ_k	0.03	0.08
R^2_{adj}	0.995	
$\lambda_k \left[\mathbb{E}\left(R^e_{k,t}\right) - \lambda_k\right]$	0.02	0.00
Rel. risk aversion, η_k	2.72	1.35
Standard error for η_k	0.14	0.22

4 Conclusion

We show that the premia of portfolios obtained as univariate sorts of stocks on rolling-window correlations with the market return sharply decrease along the sorting dimension while their beta increases. This anomaly produces a strong decreasing tilt in Sharpe ratios and an increasing pattern in the correlation of returns with consumption growth. Consequently, the implied coefficients of relative risk aversion differ between the first and the last dynamic correlation quintile four times, further exacerbating the equity premium puzzle. The stocks in the lowest quintile also exhibit relatively low and unstable betas compared to the rest of the market. This counterintuitive finding implies that the highest equity premia originate from assets with the weakest coupling to systemic risk.

Following the intuition based on findings of Fama and French (1997) and Pastor and Stambaugh (1999), we attribute this excessive premium to the beta uncertainty aversion. To this end, we introduce a state variable capturing the instability of correlations with the economy in the general ICAPM framework. Empirically, this translates to another "risk factor," for which we construct a proxy as a tradable portfolio long on stocks with low and unstable correlations with the economy and short on stocks with high and stable correlations. This approach removes the tilt as the additional factor picks up the decreasing premium while market betas become almost constant across the correlation quintiles. It also generates reasonable risk aversion parameters.

There are at least two limitations to this approach. First, it is a *deus-ex-machina* solution, where the factor is constructed exploiting a spread it tries to "explain." However, a similar factor seems to capture a shared pattern that emerges across the major asset pricing anomalies (Božović, 2022a) and characterize the returns of industry portfolios (Božović, 2022b). Second, a more general description within the ICAPM framework should include investors' outside income or other state variables for investment opportunities. From a purely empirical perspective, such an approach would only increase the number of factors in the regressions. Our goal here was to isolate the main effect and remove the tilt produced by the premium for beta uncertainty.

Acknowledgments

The author acknowledges the financial support of the Ministry of Education, Science and Technological Development of the Republic of Serbia. The usual disclaimer applies.

References

Abel, A., 1990. Asset prices under habit formation and catching up with the joneses. American Economic Review 80, 38–42.

Bansal, R., Coleman, W., 1996. A monetary explanation of the equity premium, term premium, and risk-free rate puzzles. Journal of Political Economy 104, 1135–1171.

Barro, R., 2005. Rare events and the equity premium. NBER Working Paper No. 11310.

Benartzi, S., Thaler, R., 1995. Myopic loss aversion and the Equity Premium Puzzle. Quarterly Journal of Economics 110, 73–92.

Božović, M., 2022a. A common pattern across asset pricing anomalies. Finance Research Letters 48, 103004.

Božović, M., 2022b. Can a dynamic correlation factor improve the pricing of industry portfolios? Finance Research Letters 53, 103626.

Campbell, J., 1996. Understanding risk and return. Journal of Political Economy 104, 298–345.

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

Cochrane, J., Hansen, L., 1992. Asset pricing explorations for macroeconomics. NBER Macroeconomics Annuals 7, 115–165.

Constantinides, G., 1990. Habit formation: A resolution of the equity premium puzzle. Journal of Political Economy 98, 519–543.

Dimson, E., Marsh, P., Staunton, M., 2008. The worldwide equity premium: A smaller puzzle, in: Mehra, R. (Ed.), Handbook of the Equity Risk Premium. Elsevier, Amsterdam.

Epstein, L., Zin, S., 1989. Substitution, risk aversion and the temporal behavior of asset returns. Journal of Political Economy 99, 263–286.

Fama, E., French, K., 1997. Industry costs of equity. Journal of Financial Economics 43, 153–193.

Fama, E., French, K., 2002. The equity premium. Journal of Finance 57, 637–659.

Ferson, W., Constantinides, G., 1991. Habit persistence and durability in aggregate consumption: Empirical tests. Journal of Financial Economics 29, 199–240.

Heaton, J., 1995. An empirical investigation of asset pricing with temporally dependent preference specifications. Econometrica 63, 681–717.

Holmström, B., Tirole, J., 1998. Private and public supply of liquidity. Journal of Political Economy 106, 1–40.

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

Kandel, S., Stambaugh, R., 1991. Asset returns and intertemporal preferences. Journal of Monetary Economics 27, 39–71.

Kocherlakota, N., 1996. The equity premium: It's still a puzzle. Journal of Economic Literature 34, 42–71.

Koshy, H., Cochrane, J., 1998. Where is the market going? uncertain facts and novel theories. Available at SSRN: https://ssrn.com/abstract=75371 or http://dx.doi.org/10.2139/ssrn.75371.

Lettau, M., Ludvigson, S., 2001. Consumption, aggregate wealth and expected stock returns. Journal of Finance 56, 815–849.

Lucas, R., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.

McGrattan, E., Prescott, E., 2001. Taxes, regulations, and asset prices. NBER Working Paper No. 8623.

Mehra, R., 2003. The equity premium: Why is it a puzzle? (corrected). Financial Analysts Journal 59, 54–69.

Mehra, R., Prescott, E., 1985. The equity premium: A puzzle. Journal of Monetary Economics 15, 145–161.

Mehra, R., Prescott, E., 2003. The equity premium in retrospect, in: Constantinides, G., Harris, M., Stulz, R. (Eds.), Handbook of the Economics of Finance: Volume 1B – Financial Markets and Asset Pricing. Elsevier, Amsterdam.

Mehra, R., Prescott, E., 2006. The equity premium: What have we learned in 20 years?, in: Mehra, R. (Ed.), Handbook of Investments: Equity Risk Premium in the Handbook of the Economics of Finance series. Elsevier, Amsterdam.

Merton, R., 1973. An intertemporal capital asset pricing model. Econometrica 41, 867–887.

Palomino, R., 1996. Noise trading in small markets. Journal of Finance 51, 1537–1550.

Pastor, L., Stambaugh, R., 1999. Costs of equity capital and model mispricing. Journal of Finance 54, 67–121.

Siegel, J., 2014. Stocks for the Long Run: The Definitive Guide to Financial Market Returns and Long-Term Investment Strategies, 5th ed. McGraw Hill.

Tallarini, T., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45, 507–532.

Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. Journal of Monetary Economics 24, 401–421.