Beta Ambiguity and Asymmetric Mispricing

Samuel Jaeger†

May 15, 2026

Abstract

After bad market news, the empirical Security Market Line (SML) is steep and closely aligned with the Capital Asset Pricing Model (CAPM); after good news, however, it slopes downward. This mispricing asymmetry is sharply at odds with canonical explanations for a flat SML. I show that beta ambiguity resolves the puzzle: ambiguity-averse investors act as if betas are inflated when news is bad and deflated when news is sufficiently good, generating the observed asymmetry. The mechanism further predicts that betting-against-beta (BAB) strategies crash after strongly negative news. A simple dynamic BAB strategy exploiting this prediction raises the Sharpe ratio by 0.48 on average across U.S. and international equity markets.

[1]I am particularly grateful to Julien Cujean and James Dow for valuable discussions and comments. I would also like to thank Philip Baumann, Clemens Boehlen, Lorenzo Bretscher, Lukas Jaeger, Norbert Trautmann and Philip Valta for valuable comments. All errors are my own.

†University of Bern, Institute for Financial Management, Engehaldenstrasse 4, 3012 Bern, Switzerland, samuel.jaeger@unibe.ch.

1 Introduction

The empirical Security Market Line (SML) is known to be flatter than predicted by the Capital Asset Pricing Model (CAPM), a failure commonly attributed to borrowing constraints (Black, 1972; Frazzini and Pedersen, 2014) or investor disagreement (Hong and Sraer, 2016; Andrei, Cujean, and Wilson, 2023). Consistent with these explanations, empirical evidence shows that the CAPM performs better in market states when constraints are loose and disagreement is low.[1]

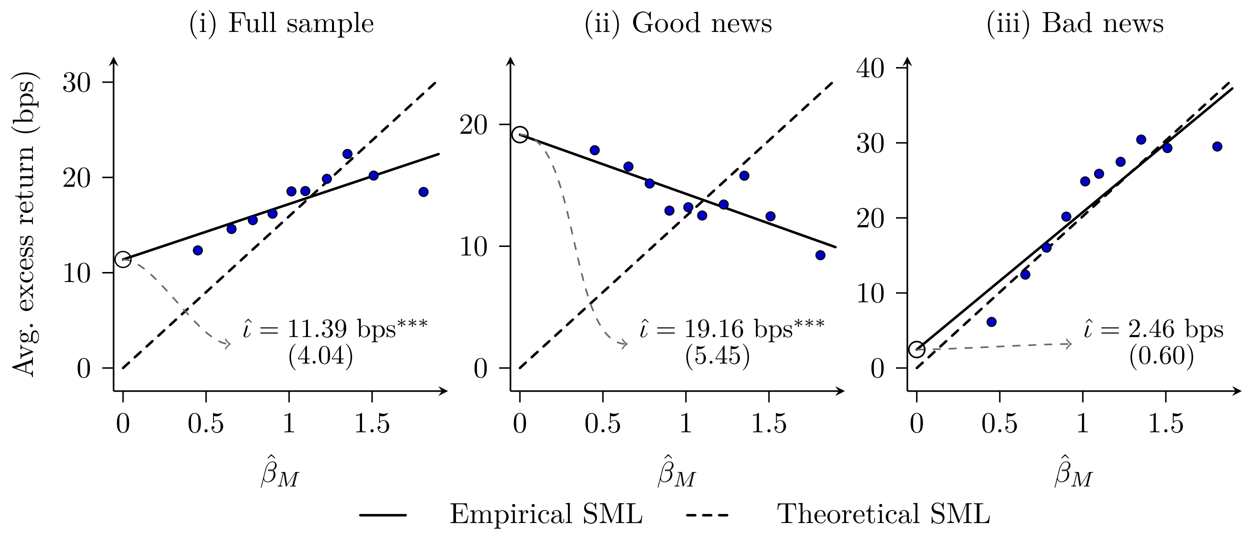

Figure 1: Security market line after positive and negative market returns, 1930–2024. This figure plots average value-weighted weekly excess returns against market betas ($\hat{\beta}_M$) for ten beta-sorted portfolios of U.S. common stocks listed on the NYSE, Amex, or NASDAQ. The dashed line denotes the CAPM-implied SML with zero intercept and slope equal to the average market excess return. The solid line shows the estimated SML from Fama and MacBeth (1973) cross-sectional regressions, with intercept $\hat{\iota}$ marked by an arrow. Market betas are estimated over the full sample using daily excess returns following Lewellen and Nagel (2006). Panel (i) shows the full-sample SML. Panels (ii) and (iii) condition on the lagged market excess return being above (good news) or below (bad news) its full-sample mean. t-statistics use Newey and West (1987) standard errors with nine lags. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

In this paper, I show that the flatness of the SML strongly varies with the *sign* of market news. Figure 1 illustrates this result for the U.S. stock market using 95 years of data from 1930 through 2024, with realized market excess returns as a proxy for news. Following bad news, the realized SML is steep and closely aligned with the CAPM, whereas following good news it becomes downward sloping with a large intercept.

[1] The CAPM performs better when leverage constraints are looser (Jylha, 2018; Boguth and Simutin, 2018; Lu and Qin, 2021). Greater analyst disagreement is associated with a flatter SML (Hong and Sraer, 2016), and Andrei et al. (2023) argue that public information crowds out heterogeneous private signals, rationalizing the steep SML on announcement days.

This mispricing asymmetry is hard to reconcile with the canonical theories for a flat SML mentioned above. For these theories to account for it, borrowing constraints would have to loosen and disagreement would have to fall on bad news. Yet empirical evidence suggests the opposite correlation: constraints tighten (Brunnermeier and Pedersen, 2008; Geanakoplos, 2010; Adrian and Shin, 2013) and disagreement rises (Cujean and Hasler, 2017) in bad times, implying the reverse pattern.

I show that the asymmetry arises when investors are averse to beta ambiguity. Because true betas are not known, investors face a set of plausible values when making investment decisions. Ambiguity-averse investors evaluate portfolios under worst-case betas, which generally differ from true betas and, crucially, vary with the sign of news. The resulting distortions in perceived betas translate into an asymmetric mispricing pattern across news states.

Consider a mean-variance investor in the spirit of Markowitz (1952). To choose her portfolio she needs the mean and covariance matrix of payoffs, both of which depend on assets' exposures to a common factor, their betas. Standard rational-expectations models assume the investor knows these betas because she knows the true data-generating process of payoffs. In practice, however, betas are not observable and must be estimated (e.g., Vasicek, 1973; Shanken, 1992; Handa and Linn, 1993; Adrian and Franzoni, 2009; Hollstein, Prokopczuk, and Wese Simen, 2020), and multiple defensible estimation procedures coexist: different methodologies, and different specification choices within a given methodology such as the estimation window, sampling frequency, or shrinkage procedure, can deliver materially different estimates for the same stock at the same point in time. Each plausible beta vector maps into a different mean and covariance matrix of payoffs, so the investor faces a family of plausible payoff distributions rather than a single one. Beta ambiguity therefore introduces a second layer of uncertainty beyond risk in the sense of Knight (1921): the investor is uncertain not only about the realization of payoffs, but about the mean and covariance matrix that govern their distribution.

Experimental evidence shows that ambiguity-averse investors behave as if they make decisions under worst-case scenarios (Ellsberg, 1961; Bossaerts, Ghirardato, Guarnaschelli, and Zame, 2010; Ahn, Choi, Gale, and Kariv, 2014). Applied to the case described above, this suggests the investor acts on worst-case betas among the set of plausible values. Two key implications follow. First, perceived betas are distorted relative to true betas, generating mispricing relative to the rational-expectations benchmark. Second, and central to the mechanism, which beta is worst depends on the *sign* of market news. When news is bad, inflated betas are worst because they amplify exposure to a factor with a negative expected realization. When news is sufficiently good, deflated betas are worst because they imply forgone gains. As a consequence, perceived betas of ambiguity-averse investors vary endogenously with news even though true betas are exogenous and fixed, generating news-dependent mispricing.

To derive the resulting asset pricing implications, I formalize this mechanism in a one-period model with multiple risky assets whose payoffs follow a common factor structure. Beta ambiguity is modeled as a range of plausible betas containing the true value. A representative mean-variance investor observes news about the factor and chooses a portfolio. She has max-min preferences in the sense of Gilboa and Schmeidler (1989), which axiomatize the ambiguity-averse behavior described above: she maximizes utility over the worst-case distribution. The mean-variance setting reveals why the response to bad and good news is not symmetric. Higher betas affect utility through two channels: they shift expected payoffs in the direction of news, and they raise variance regardless of news. When news is bad, both channels push the same way, and inflated betas are always worst. When news is good, the channels conflict, and inflated betas remain worst until news becomes sufficiently good for the mean channel to overturn the variance penalty. Beyond this threshold, the worst case switches to deflated betas. This threshold is the source of the asymmetry: bad news always tilts toward inflated betas, but good news tilts toward deflated betas only when it is sufficiently strong.

These distortions have observable consequences for cross-sectional mispricing. Realized payoffs are drawn from the true distribution, whereas equilibrium prices are formed under the worst-case distribution. As a result, realized returns carry an ambiguity premium. In the cross-section, this premium appears as a nonzero SML intercept. A positive intercept indicates an SML that is "too flat", with high-beta assets underperforming and low-beta assets outperforming the CAPM benchmark; a negative intercept indicates the reverse, an SML that is "too steep". The intercept decomposes into two components. A *premium bias* reflects misperceived factor risk and depends only on whether betas are inflated or deflated. An *expectation bias* reflects misperceived sensitivity of payoffs to news and scales with its magnitude. Together, these biases make the SML intercept a piecewise-linear function of news. As news improves from very bad to moderately good, the intercept rises. At the threshold where perceived betas switch from inflated to deflated, the intercept jumps upward. For stronger good news beyond this threshold, it declines. Because the switch occurs only on sufficiently good news, the model predicts a mispricing asymmetry: the realized SML intercept is higher after positive news than after negative news, a pattern in line with Figure 1.

I test the model's asset pricing predictions using weekly excess returns on beta-sorted portfolios of U.S. stocks from 1930 to 2024. I follow a large empirical literature in using realized market excess returns as a proxy for the sign and magnitude of news, because they aggregate and reveal newly arriving information (e.g., McQueen, Michael, and Thorley, 1996; Chordia and Swaminathan, 2000). I find strong support for the model's predictions. The weekly SML intercept is on average 15.45 basis points (8.03% annualized) higher following good news than following bad news and traces out the predicted piecewise relationship with news: it rises when lagged market returns are negative, peaks when they are positive, and declines for strongly positive realizations.

The findings are robust along several dimensions. They survive controls for the Carhart (1997) factors, an illiquidity factor, and proxies for borrowing constraints and investor disagreement, and they persist across alternative test assets, including industry, size, and book-to-market portfolios. Complementary to the baseline analysis, I conduct an event study around scheduled macroeconomic announcements following the selection procedure of Savor and Wilson (2013, 2014). The results confirm the baseline findings: the cumulative difference in the SML intercept following bad announcements relative to good or neutral ones declines over the subsequent eight trading days, consistent with the theoretical prediction.

The mechanism has direct implications for the betting-against-beta (BAB) strategy of Frazzini and Pedersen (2014). BAB exploits a flat SML by going long low-beta stocks and short high-beta stocks to construct a zero-beta portfolio. Since its expected return is directly linked to the SML intercept, the model's predictions translate naturally into predictions for BAB performance: it should be most profitable after moderately positive news and particularly prone to crashes when news turns sharply negative. To test this implication, I backtest a dynamic BAB strategy with a simple rule: hold a long BAB position when the lagged market return is at or above its long-run mean, switch to a short position when it falls at least 1.5 standard deviations below the mean, and otherwise maintain the current position. Across U.S., European, Pacific, and global equity markets, the dynamic strategy outperforms static BAB by 4.1 to 6.2 percentage points in annualized average excess returns and raises the Sharpe ratio by 0.48 on average. Because the rule responds only to strongly negative news, the strategy triggers only four to five additional trades per year, limiting turnover and transaction costs.

This paper contributes to the literature on ambiguity in asset pricing following the multiple-priors framework of Gilboa and Schmeidler (1989) and its dynamic extension in Epstein and Schneider (2003).[2] The closest papers are Epstein and Schneider (2008) and

[2] A large body of work studies how ambiguity affects asset prices and premia. Equilibrium implementations include Epstein and Wang (1994) and Chen and Epstein (2002). Related implications for market participation arise in Dow and Werlang (1992). Robust-control approaches to model uncertainty are developed in Hansen and Sargent (2001, 2007). Smooth ambiguity preferences are introduced by Klibanoff, Marinacci, and Mukerji (2005), with recursive extensions in Hayashi and Miao (2011) and asset-pricing applications in Ju and Miao (2012). Surveys include Epstein and Schneider (2010) and Guidolin and Rinaldi (2013). More recent empirical contributions on ambiguity and asset prices include Şahin and Danışoğlu

Illeditsch (2011), who study how ambiguity interacts with news and show that it can generate asymmetric price responses. My framework mainly differs along two dimensions. First, I study ambiguity about factor betas rather than signal precision. Second, I derive implications for cross-sectional mispricing, whereas their focus lies on return distribution properties such as volatility and skewness.[3]

This paper also relates to the literature on beta uncertainty and cross-sectional asset pricing. Hollstein et al. (2020) and Armstrong, Banerjee, and Corona (2013) show that stocks with high beta uncertainty underperform those with low beta uncertainty, which Armstrong et al. (2013) attribute to convexity in prices with respect to beta. In contrast, I study how beta uncertainty interacts with news to generate state-dependent mispricing, rather than how beta uncertainty affects average returns unconditionally.

Finally, the paper relates to the literature on factor timing. The closest contribution is Hedegaard (2018), who also times BAB using past market returns but links the predictability to time-varying leverage demand. While closely related empirically, the mechanism differs. In his framework, stronger leverage demand raises expected BAB returns; absent such demand, the CAPM benchmark is restored and expected BAB returns are zero. Thus, the channel naturally explains time variation in positive BAB premia, but speaks less directly to predictable BAB crashes. Under beta ambiguity, by contrast, bad news can turn expected BAB returns negative, naturally generating such crashes. More broadly, the dynamic BAB strategy studied here is related to Daniel and Moskowitz (2016) in using conditional exposure to mitigate factor crashes.

The remainder of the paper is organized as follows. Section 2 develops the model and derives the testable predictions. Section 3 presents the data and methodology. Section 4 presents the empirical results, and Section 5 concludes.

2 Model

2.1 Economy

I consider a one-period economy with two dates, t and $t+1$. Traded assets consist of N risky stocks, indexed by $i = 1, \ldots, N$, and a risk-free asset with gross return normalized to one. Each stock i is assumed to have an equal number of $1/N$ shares outstanding, so that

(2022).

[3] Ambiguity about signal precision in the multi-asset economy studied here could also generate a state-dependent SML intercept, but the peak would likely arise in the bad-news region rather than the good-news region, reversing the predicted asymmetry. The empirical evidence in this paper instead points to a peak in the good-news region, consistent with beta ambiguity.

aggregate supply across all stocks sums to one.[4] I define the market portfolio $\boldsymbol{M} = \frac{1}{N}\mathbf{1}$ as the vector whose i-th entry equals the shares outstanding of stock i.[5]

Let $\tilde{\boldsymbol{D}}_{t+1}$ denote the vector of payoffs per unit share realized at $t+1$, which follow a common factor structure

$$\tilde{\boldsymbol{D}}_{t+1} = \overline{D}\,\mathbf{1} + \boldsymbol{\beta}\,\tilde{F}_{t+1} + \tilde{\boldsymbol{\epsilon}}_{t+1}, \tag{1}$$

where $\overline{D} > 0$ is a constant. The common factor $\tilde{F}_{t+1}$ and the idiosyncratic component $\tilde{\boldsymbol{\epsilon}}_{t+1}$ are unobservable, independent, mean-zero normal random variables with distributions

$$\tilde{F}_{t+1} \sim \mathcal{N}(0, \tau_F^{-1}), \qquad \tilde{\boldsymbol{\epsilon}}_{t+1} \sim \mathcal{N}(\mathbf{0}, \tau_\epsilon^{-1}\mathbf{I}_N). \tag{2}$$

The only source of cross-sectional heterogeneity is the vector of factor exposures $\boldsymbol{\beta}$, which I refer to as *betas*.[6] The beta of the market portfolio is $\overline{\beta} := \boldsymbol{M}^\top \boldsymbol{\beta}$, which is normalized to one.

At time t, all agents observe the same public information about the factor in the form of a noisy signal,

$$s_t = \tilde{F}_{t+1} + \tilde{\xi}_t, \qquad \tilde{\xi}_t \sim \mathcal{N}(0, \tau_s^{-1}), \tag{3}$$

where τ_s denotes signal precision. Since the common factor drives payoffs of all stocks, s_t can be interpreted as macroeconomic news.

By standard Bayesian updating, the posterior distribution of the factor is normal with

$$\mathbb{E}_t[\tilde{F}_{t+1} \mid s_t] = \frac{\tau_s}{\tau}\, s_t = \rho\, s_t, \qquad \mathrm{Var}_t(\tilde{F}_{t+1} \mid s_t) = \tau^{-1}, \quad \tau = \tau_F + \tau_s, \tag{4}$$

where $\rho \equiv \tau_s/\tau$ denotes the weight placed on the signal in the posterior mean of the factor. The posterior expectations $\boldsymbol{\mu}_t$ of payoffs and conditional variance $\boldsymbol{\Sigma}_t$ are given by:

$$\begin{aligned} \boldsymbol{\mu}_t &= \mathbb{E}_t\Big[\tilde{\boldsymbol{D}}_{t+1} \mid s_t\Big] = \overline{D}\,\mathbf{1} + \boldsymbol{\beta}\,\rho\, s_t, \\ \boldsymbol{\Sigma}_t &= \mathrm{Var}_t\Big[\tilde{\boldsymbol{D}}_{t+1} \mid s_t\Big] = \tau^{-1}\,\boldsymbol{\beta}\boldsymbol{\beta}^\top + \tau_\epsilon^{-1}\mathbf{I}_N. \end{aligned} \tag{5}$$

Expected payoffs are linear in factor betas $\boldsymbol{\beta}$. The first term of the covariance matrix captures factor risk, which scales with $\boldsymbol{\beta}$, whereas the second term reflects idiosyncratic risk that is constant across stocks.

[4] The assumption of equal supply across stocks is made for simplicity and allows me to present the main results in the most transparent form. I discuss the implications of heterogeneous supply for the model's results in the Appendix 6.1.1.

[5] Throughout the paper, I place a tilde on random variables, write vectors and matrices in bold, and use plain font for scalars.

[6] Common factor betas should not be confused with endogenous *market betas* introduced in Section 2.3, which are affine transformations of exogenous betas.

2.2 Investor's portfolio choice under beta ambiguity

I consider a representative investor who observes the signal s_t at time t, predicts payoffs, chooses a portfolio, and then collects realized payoffs at $t+1$. The key assumption is that betas are ambiguous: the investor does not know the true betas but only that it lies within an ambiguity range, i.e. $\boldsymbol{\beta} \in [\boldsymbol{\beta}_L, \boldsymbol{\beta}_H]$. For expositional clarity, I assume that the true beta vector lies in the interior of this range; the exact specification is introduced in Section 2.3. I assume that the ambiguity range for the beta of the market portfolio is strictly positive, $0 < \overline{\beta}_L < \overline{\beta} = 1 < \overline{\beta}_H$, so that, even if some individual assets have negative betas (i.e., $\exists i : \beta_i < 0$), all admissible values of the market's factor beta remain positive.

Beta ambiguity relaxes the standard rational-expectations assumption that the investor knows the true data-generating process and thus the true betas $\boldsymbol{\beta}$. When betas are known, the investor can pin down the true posterior distribution in (5), so that uncertainty is limited to risk in the classical sense, captured by the covariance matrix. Under beta ambiguity, each plausible beta vector in the ambiguity range induces a different posterior mean and covariance matrix, such that the investor faces a *family* of plausible distributions, introducing a second layer of uncertainty in the sense of Knight (1921).

I assume that the investor is both risk and ambiguity averse, with constant absolute risk aversion (CARA) and max–min preferences in the sense of Gilboa and Schmeidler (1989). Under max–min preferences, the investor considers the worst-case scenario when making an investment decision.[7] Specifically, the investor maximizes the certainty equivalent under the worst-case beta vector in the ambiguity range,

$$CE_{\min} := \min_{\boldsymbol{\beta} \in [\boldsymbol{\beta}_L, \boldsymbol{\beta}_H]} \Big\{ \underbrace{\left(\overline{D}\,\mathbf{1} + \boldsymbol{\beta}\,\rho\, s_t\right)}_{:=\boldsymbol{\mu}_t} - \boldsymbol{P}_t \Big\}^{\top} \boldsymbol{\Theta} - \tfrac{\gamma}{2}\, \boldsymbol{\Theta}^{\top} \underbrace{\left(\tau^{-1}\,\boldsymbol{\beta}\boldsymbol{\beta}^{\top} + \tau_{\epsilon}^{-1}\mathbf{I}_N\right)}_{:=\boldsymbol{\Sigma}_t} \boldsymbol{\Theta}, \tag{6}$$

where $\gamma > 0$ is the risk aversion coefficient, and $\boldsymbol{P}_t$ and $\boldsymbol{\Theta}$ denote the given vectors of market prices and portfolio holdings. I denote the minimizing (worst-case) vector by $\hat{\boldsymbol{\beta}}$ and refer to it as the investor's *perceived betas*.

The minimization problem depends on the beta vector $\boldsymbol{\beta}$ only through the portfolio's beta, a scalar:

$$b := \boldsymbol{\Theta}^{\top}\boldsymbol{\beta}, \tag{7}$$

[7] This behavior is consistent with experimental evidence (Ellsberg, 1961; Bossaerts et al., 2010; Ahn et al., 2014).

such that the minimization problem (6) simplifies to the following concave objective,

$$CE_{\min} = C(\mathbf{\Theta}) \ + \min_{b \in [\,\underline{b}(\mathbf{\Theta}),\, \overline{b}(\mathbf{\Theta})\,]} \left\{ \rho\, s_t\, b \ - \ \frac{\gamma}{2}\, \tau^{-1}\, b^2 \right\}, \tag{8}$$

where $C(\mathbf{\Theta})$ collects all terms that do not depend on $\boldsymbol{\beta}$.[8]

The worst case for the investor is to hold a portfolio with high variance, low expected payoff, or both. Thus, the associated worst-case beta is jointly determined by the expected-payoff term, $\rho s_t b$, and the variance penalty, $\frac{\gamma}{2}\tau^{-1}b^2$, in (8). The variance penalty increases with b^2, so a larger absolute portfolio beta $|b|$ raises variance and lowers expected utility regardless of the signal s_t. In contrast, the effect of the portfolio beta on the expected payoff depends on the sign of the signal. When news is good ($s_t > 0$), low portfolio beta lowers expected utility. When news is bad ($s_t < 0$), high portfolio beta lowers expected utility.

To characterize the admissible range of the portfolio beta, note that the ambiguity set $[\boldsymbol{\beta}_L, \boldsymbol{\beta}_H]$ is rectangular, so that b lies in a compact interval $b \in [\,\underline{b}(\mathbf{\Theta}),\, \overline{b}(\mathbf{\Theta})\,]$. Decompose holdings into long and short components,

$$\mathbf{\Theta} = \mathbf{\Theta}^+ - \mathbf{\Theta}^-, \qquad \mathbf{\Theta}^+, \mathbf{\Theta}^- \geq \mathbf{0}. \tag{9}$$

The endpoints of the range are obtained by assigning long positions the lowest (highest) and short positions the highest (lowest) betas:

$$\begin{aligned} \underline{b}(\mathbf{\Theta}) &= (\mathbf{\Theta}^+)^\top \boldsymbol{\beta}_L - (\mathbf{\Theta}^-)^\top \boldsymbol{\beta}_H, \\ \overline{b}(\mathbf{\Theta}) &= (\mathbf{\Theta}^+)^\top \boldsymbol{\beta}_H - (\mathbf{\Theta}^-)^\top \boldsymbol{\beta}_L. \end{aligned} \tag{10}$$

The unique maximizer of the objective (8) is $b^{\max}(s_t) = \frac{\tau_s}{\gamma} s_t$, and by concavity, the minimum over the interval is attained at the endpoint farthest from this maximizer. Proposition 1 formalizes the resulting selection rule for the worst-case portfolio beta.

Proposition 1 (Selection of worst-case portfolio beta)**.** *Let the ambiguity range for the portfolio beta be* $[\underline{b}(\mathbf{\Theta}), \overline{b}(\mathbf{\Theta})]$ *with midpoint* $m(\mathbf{\Theta}) = \frac{1}{2}\big(\underline{b}(\mathbf{\Theta}) + \overline{b}(\mathbf{\Theta})\big)$*. Let* $b^{\max}(s_t) = \frac{\tau_s}{\gamma} s_t$ *denote the unconstrained maximizer of objective* (8)*. Then the perceived (worst-case) portfolio beta is*

$$\hat{b}(\mathbf{\Theta}, s_t) = \begin{cases} \overline{b}(\mathbf{\Theta}), & b^{\max}(s_t) < m(\mathbf{\Theta}), \\ \underline{b}(\mathbf{\Theta}), & b^{\max}(s_t) > m(\mathbf{\Theta}). \end{cases} \tag{11}$$

Figure 2 illustrates the selection rule in an example with positive news ($s_t > 0$), so the hypothetical exposure that maximizes utility is positive $b^{\max}(s_t) > 0$. Since $b^{\max}(s_t) > m(\mathbf{\Theta})$,

[8] $C(\mathbf{\Theta}) \coloneqq (\overline{D}\,\mathbf{1} - \boldsymbol{P}_t)^\top \mathbf{\Theta} - \frac{\gamma}{2}\, \tau_\epsilon^{-1}\, \|\mathbf{\Theta}\|^2$.

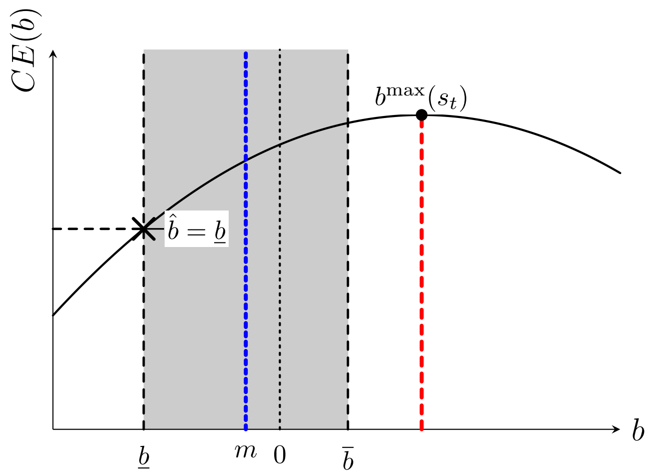

Figure 2: Choice of worst-case portfolio beta. The figure plots the certainty equivalent CE as a function of portfolio beta b for a positive signal $s_t > 0$. The shaded region represents the ambiguity range for the portfolio beta $b \in [\underline{b}(\mathbf{\Theta}), \overline{b}(\mathbf{\Theta})]$, with midpoint m. The parabola is concave, with maximizer $b^{\max}(s_t)$. Since $b^{\max}(s_t) > m$, the worst case is attained at the lower endpoint $\underline{b}(\mathbf{\Theta})$. The investor therefore assigns high betas to short positions and low betas to long positions, minimizing exposure.

the worst case portfolio beta lies at the lower bound $\underline{b}(\mathbf{\Theta})$. Intuitively, when news is strongly positive the investor would ideally want high factor exposure; the worst-case scenario is therefore to hold a portfolio with the lowest admissible exposure, forgoing expected gains while still incurring a variance penalty. As a result, the investor assigns the lowest admissible betas to long positions and the highest admissible betas to short positions. These assignments determine the perceived betas $\hat{\boldsymbol{\beta}}$ that the investor uses to make her investment decision.

Once the worst-case betas $\hat{\boldsymbol{\beta}}$ have been determined according to rule (11), the investor solves the outer maximization of the max–min problem in (6). The chosen portfolio is

$$\mathbf{\Theta}^{\star} = \frac{1}{\gamma}\hat{\mathbf{\Sigma}}_t^{-1}\big(\hat{\boldsymbol{\mu}}_t - \boldsymbol{P}_t\big), \tag{12}$$

where the hats indicate that the conditional mean $\hat{\boldsymbol{\mu}}_t$ and covariance matrix $\hat{\mathbf{\Sigma}}_t$ are computed at the perceived betas $\hat{\boldsymbol{\beta}}$. Portfolio demand therefore takes the familiar mean–variance form, with the key distinction that posterior moments reflect the worst-case rather than the true distribution.

2.2.1 Perceived betas in equilibrium

The preceding analysis characterizes the worst-case portfolio beta for a given candidate portfolio $\mathbf{\Theta}$. The result depends on the composition of that portfolio. In particular, long and short positions reverse the effect of news on expected payoffs: bad news lowers expected payoffs for long positions but raises them for short positions. As a consequence, the worst-case beta assignment differs across long and short positions. In equilibrium, however, the representative investor must hold the entire supply of assets, so the relevant portfolio is the market portfolio $\boldsymbol{M}$. Evaluating the max–min objective (6) at $\boldsymbol{\Theta} = \boldsymbol{M}$ yields two implications. First, because all assets are in positive net supply, the market portfolio contains no short positions. The worst-case selection rule in (11) therefore applies uniformly: all perceived betas are either inflated ($\boldsymbol{\beta}_H$) or deflated ($\boldsymbol{\beta}_L$) relative to the truth. Second, in the absence of short positions, the worst-case choice depends only on the signal s_t. Proposition 2 characterizes the resulting state-dependent selection rule.

Proposition 2 (Perceived betas in equilibrium depend on news)**.** *Let* $\overline{s} := \frac{\gamma}{2\tau_s}\left(\overline{\beta}_L + \overline{\beta}_H\right) > 0$*. Under the minimization problem* (6) *with* $\boldsymbol{\Theta} = \boldsymbol{M}$*, the perceived vector of betas satisfies*

$$\hat{\boldsymbol{\beta}}(s_t) = \begin{cases} \boldsymbol{\beta}_H, & s_t < \overline{s}, \\ \boldsymbol{\beta}_L, & s_t > \overline{s}. \end{cases} \tag{13}$$

Perceived betas are inflated on bad and moderately good news ($s_t < \overline{s}$) *and deflated on sufficiently good news* ($s_t > \overline{s}$)*.*

The selection rule follows from the mean–variance trade-off described above applied to the market portfolio. Increasing $\overline{\beta}$ scales both the conditional mean and the conditional variance of the market portfolio. While higher variance always lowers utility, the effect on the mean depends on the sign of the signal. When news is bad ($s_t < 0$), high beta amplifies expected losses, so inflated betas are systematically worst. When news is moderately good ($0 < s_t < \overline{s}$), high beta raises expected payoffs, but the variance penalty still dominates, and inflated betas remain worst. Only when the signal exceeds the threshold $\overline{s}$ do the higher expected payoffs from high beta outweigh the variance penalty, causing the worst case to switch to deflated betas.

Given the perceived betas from Proposition 2, equilibrium prices follow from substituting $\boldsymbol{\Theta}^\star = \boldsymbol{M}$ into the optimal demand (12) and solving for prices:

$$\boldsymbol{P}_t^\star(s_t) \;=\; \hat{\boldsymbol{\mu}}_t(s_t) \;-\; \gamma\, \hat{\boldsymbol{\Sigma}}_t(s_t)\, \boldsymbol{M}. \tag{14}$$

Equilibrium prices reflect perceived expected payoffs discounted by the risk adjustment $\gamma\hat{\boldsymbol{\Sigma}}(s_t)\boldsymbol{M}$, which compensates investors for bearing perceived risk.

2.3 Equilibrium (mis)pricing under beta ambiguity

The preceding section shows how beta ambiguity aversion distorts perceived betas away from their true values and makes this distortion depend on the sign and magnitude of news. I now derive the asset pricing implications of this mechanism. Because equilibrium prices are set under worst-case beliefs while payoffs are realized under the true distribution, beta ambiguity generates mispricing relative to the rational-expectations benchmark where the investor knows true betas and the CAPM holds. This mispricing manifests in the cross section as a nonzero SML intercept: positive when the SML is "too flat" and negative when it is "too steep." The main result of this section is that the sign and magnitude of the intercept vary systematically with the sign and magnitude of news.

Specification of the ambiguity range To obtain closed-form pricing implications, I impose a multiplicative structure on the ambiguity range by specifying

$$\boldsymbol{\beta}_H \equiv (1+\kappa)\boldsymbol{\beta}, \qquad \boldsymbol{\beta}_L \equiv (1-\kappa)\boldsymbol{\beta}, \tag{15}$$

where $\kappa > 0$ measures the degree of ambiguity. This structure is empirically motivated: Hollstein et al. (2020) document that beta uncertainty increases with the level of beta, and I provide additional supporting evidence in Section 3 and Section 6.2.1 in the Appendix.

Combined with the endogenous worst-case selection rule in Proposition 2, this structure implies that the perceived betas are proportional to true betas,

$$\hat{\boldsymbol{\beta}}(s_t) = a(s_t)\boldsymbol{\beta}, \tag{16}$$

where the state-dependent ambiguity multiplier is

$$a(s_t) = \begin{cases} 1+\kappa & s_t < \overline{s}, \\ 1-\kappa & s_t > \overline{s}. \end{cases} \tag{17}$$

Expected empirical returns. I now derive the model's empirical asset pricing implications from the perspective of an econometrician who observes realized payoffs $\tilde{\boldsymbol{D}}_{t+1}$ and equilibrium prices $\boldsymbol{P}_t^\star(s_t)$. The econometrician computes realized dollar returns,

$$\tilde{\boldsymbol{R}}_{t+1} = \tilde{\boldsymbol{D}}_{t+1} - \boldsymbol{P}_t^\star(s_t), \tag{18}$$

which I refer to as *empirical returns*. Since payoffs are normally distributed, dollar returns are the natural return measure in this environment. Because the risk-free rate is normalized to zero, all empirical returns are excess returns. Payoffs are drawn from the true distribution, whereas equilibrium prices are determined under the worst-case distribution as described in Section 2.2.

The conditional expected empirical returns are given by

$$\mathbb{E}_t\left[\tilde{\boldsymbol{R}}_{t+1} \mid s_t\right] = \boldsymbol{\mu}_t(s_t) - \boldsymbol{P}^\star_t(s_t) = \underbrace{\gamma\,\boldsymbol{\Sigma}_t\,\boldsymbol{M}}_{\text{RE premium}} + \underbrace{\gamma\big(\hat{\boldsymbol{\Sigma}}_t(s_t) - \boldsymbol{\Sigma}_t\big)\boldsymbol{M} + \big(\boldsymbol{\mu}_t - \hat{\boldsymbol{\mu}}_t(s_t)\big)}_{\text{ambiguity premium}}, \tag{19}$$

which represent the returns the econometrician is expected to observe given signal s_t. The first term is a rational-expectations (RE) risk premium, the compensation a mean–variance investor requires for bearing risk. In the absence of beta ambiguity, $a(s_t) = 1$, this is the only source of expected returns. The remaining terms reflect an *ambiguity premium* that arises from misperceptions of means and covariances. From a rational-expectations perspective, this premium constitutes a *mispricing*, as it compensates investors beyond what is justified by ”true” risk.[9]

Substituting the ambiguity multiplier $a(s_t)$ from (17) into the expression for expected returns (19) yields a decomposition of the ambiguity premium into two biases that are proportional to true beta. Corollary 1 states the result.

Corollary 1 (Expected empirical returns under beta ambiguity)**.** *Expected empirical returns satisfy*

$$\mathbb{E}_t\big[\tilde{\boldsymbol{R}}_{t+1} \mid s_t\big] = \underbrace{\gamma\Big(c_M\,\mathbf{1} + \tau^{-1}\,\boldsymbol{\beta}\Big)}_{\textit{RE premium}} + \underbrace{\Pi(s_t)\,\boldsymbol{\beta}}_{\textit{premium bias}} + \underbrace{\Xi(s_t)\,s_t\,\boldsymbol{\beta}}_{\textit{expectation bias}}, \tag{20}$$

where

$$\Pi(s_t) = \pi_H\,\mathbf{1}_{s_t<\bar{s}} - \pi_L\,\mathbf{1}_{s_t>\bar{s}}, \qquad \Xi(s_t) = -\eta\,\mathbf{1}_{s_t<\bar{s}} + \eta\,\mathbf{1}_{s_t>\bar{s}},$$

and π_H, π_L, $\eta > 0$ *are given by*

$$\pi_H = \gamma\,\kappa(2+\kappa)\,\tau^{-1}, \qquad \pi_L = \gamma\,\kappa(2-\kappa)\,\tau^{-1}, \qquad \eta = \kappa\,\tau_s.$$

The *premium bias* reflects misperceived factor risk. When perceived betas are inflated ($s_t < \bar{s}$), investors overstate factor risk and therefore require higher compensation than what can be justified under rational expectations, generating a positive premium bias. When perceived betas are deflated ($s_t > \bar{s}$), risk is understated and the premium bias is negative.

The *expectation bias* reflects distortions in the pricing sensitivity to news. When betas

[9] Other models with ambiguity premia have been proposed to explain the equity premium puzzle documented by Mehra and Prescott (1985); see, for example, Chen and Epstein (2002), Barillas, Hansen, and Sargent (2009), Gollier (2011), and Rieger and Wang (2012).

are inflated ($s_t < \bar{s}$), prices overreact to news; when betas are deflated ($s_t > \bar{s}$), prices underreact. These misreactions cause expected payoffs to deviate from their rational-expectations counterparts, with the gap widening in the magnitude of news.

The empirical SML. The premium and expectation bias in Corollary 1 translate into cross-sectional mispricing in the form of a nonzero SML intercept. To formalize these implications, suppose the econometrician runs a cross-sectional regression of expected returns on *market betas* $\boldsymbol{\beta}_{M,t}$ in the spirit of Fama and MacBeth (1973),

$$\mathbb{E}_t[\tilde{\boldsymbol{R}}_{t+1} \mid s_t] = \iota(s_t)\,\mathbf{1} + \phi(s_t)\,\boldsymbol{\beta}_{M,t}. \tag{21}$$

The coefficients $\iota(s_t)$ and $\phi(s_t)$ represent the intercept and slope of the empirical SML. The CAPM holds when the intercept is zero and the slope equals the market risk premium.

The econometrician is assumed to use the true market betas $\boldsymbol{\beta}_{M,t}$ when calculating the SML.[10] The market beta equals a stock's conditional covariance with the market portfolio divided by the conditional market variance $\sigma^2_{M,t}$,

$$\boldsymbol{\beta}_{M,t} = \frac{\boldsymbol{\Sigma}_t\,\boldsymbol{M}}{\boldsymbol{M}^\top \boldsymbol{\Sigma}_t \boldsymbol{M}} = \frac{\tau^{-1}\boldsymbol{\beta} + c_M \mathbf{1}}{\tau^{-1} + c_M} = \lambda\,\boldsymbol{\beta} + (1-\lambda)\,\mathbf{1}, \tag{22}$$

where $c_M = \frac{\tau_\epsilon^{-1}}{N}$ is the residual risk of the market portfolio and $\lambda = \frac{\tau^{-1}}{\tau^{-1}+c_M} \in (0,1)$. Since residual risk is identical across stocks, every stock receives the same additive term c_M in the numerator, making its market beta a convex combination of its true beta and one. This affine map compresses market betas toward one, so each stock's market beta lies between its true beta and one.

Projecting expected empirical returns in (20) onto the CAPM basis $\{\mathbf{1}, \boldsymbol{\beta}_{M,t}\}$ yields the intercept $\iota(s_t)$ and slope $\phi(s_t)$ of the empirical SML,

$$\begin{aligned}
\iota(s_t) &= c_M\,\gamma\left(1 - a(s_t)^2\right) \;-\; c_M\left(1 - a(s_t)\right)\tau_s\,s_t,\\
\phi(s_t) &= \gamma\,\sigma^2_{M,t} \;+\; \gamma\,\sigma^2_{M,t}\left(a(s_t)^2 - 1\right) \;+\; \sigma^2_{M,t}\left(1 - a(s_t)\right)\tau_s\,s_t.
\end{aligned} \tag{23}$$

Under rational expectations ($a(s_t) = 1$), the intercept is zero and the slope equals the market risk premium $\gamma\sigma^2_{M,t}$. Hence, as anticipated, the CAPM holds exactly. Under beta ambiguity ($a(s_t) \neq 1$), both coefficients deviate from these benchmark values, reflecting cross-sectional mispricing.

[10] The true beta is assumed to lie at the midpoint of the ambiguity range. This coincides with a natural point estimate of beta. The investor has access to the same estimate but, being ambiguity averse, evaluates her portfolio under worst-case betas.

The SML intercept provides a transparent summary measure of CAPM mispricing, on which I focus henceforth.[11] A nonzero intercept arises because the premium- and expectation bias in (20) are proportional to beta $\boldsymbol{\beta}$, whereas the SML is calculated using market betas. Inverting (22),

$$\boldsymbol{\beta} = \frac{1}{\lambda}\,\boldsymbol{\beta}_{M,t} - \frac{1-\lambda}{\lambda}\,\mathbf{1}, \tag{24}$$

terms linear in $\boldsymbol{\beta}$ become affine in $\boldsymbol{\beta}_{M,t}$: the compression of betas toward one introduces a constant that cannot be absorbed by the slope of the SML, which is reflected by a nonzero intercept. Any ambiguity-induced mispricing thus translates into *relative* mispricing of high-beta versus low-beta stocks. A positive intercept ($\iota(s_t) > 0$) implies a SML that is "too flat": high-beta stocks are overpriced relative to low-beta stocks. A negative intercept reverses the pattern, implying an SML that is "too steep".

I next characterize how the intercept depends on news. Substituting $a(s_t)$ from (17) into (23) yields the expression for the news-dependent SML intercept in Proposition 3, the main theoretical result of this paper.

Proposition 3 (SML intercept as a function of news)**.** *The empirical SML intercept satisfies*

$$\iota(s_t) = \underbrace{\bar{\Pi}(s_t)}_{\textit{premium bias}} + \underbrace{\bar{\Xi}(s_t)\,s_t}_{\textit{expectation bias}}, \tag{25}$$

where the state-dependent components are

$$\bar{\Pi}(s_t) = -\bar{\pi}_H\,\mathbf{1}_{s_t<\bar{s}} + \bar{\pi}_L\,\mathbf{1}_{s_t>\bar{s}}, \qquad \bar{\Xi}(s_t) = \bar{\eta}\,\mathbf{1}_{s_t<\bar{s}} - \bar{\eta}\,\mathbf{1}_{s_t>\bar{s}},$$

and $\bar{\pi}_H$, $\bar{\pi}_L$, $\bar{\eta} > 0$ are given by

$$\bar{\pi}_H = c_M\tau\,\pi_H, \qquad \bar{\pi}_L = c_M\tau\,\pi_L, \qquad \bar{\eta} = c_M\,\eta.$$

The intercept is piecewise linear in s_t, with slope $\bar{\eta}$ when $s_t < \bar{s}$ and $-\bar{\eta}$ when $s_t > \bar{s}$. At $s_t = \bar{s}$, the intercept exhibits an upward jump of size

$$\Delta\iota = \lim_{s_t\to\bar{s}^+}\iota(s_t) - \lim_{s_t\to\bar{s}^-}\iota(s_t) = \bar{\pi}_H + \bar{\pi}_L > 0. \tag{26}$$

[11] The intercept maps directly into asset-level CAPM alphas. The conditional alpha vector satisfies

$$\boldsymbol{\alpha}(s_t) \equiv \mathbb{E}_t[\tilde{\boldsymbol{R}}_{t+1} \mid s_t] - \boldsymbol{\beta}_{M,t}\,\mathbb{E}_t[\tilde{R}_{M,t+1} \mid s_t] = \iota(s_t)\,(\mathbf{1} - \boldsymbol{\beta}_{M,t}).$$

If $\iota(s_t) > 0$, high-beta stocks ($\beta_{M,i,t} > 1$) exhibit negative alphas and low-beta stocks positive alphas. If $\iota(s_t) < 0$, the pattern reverses.

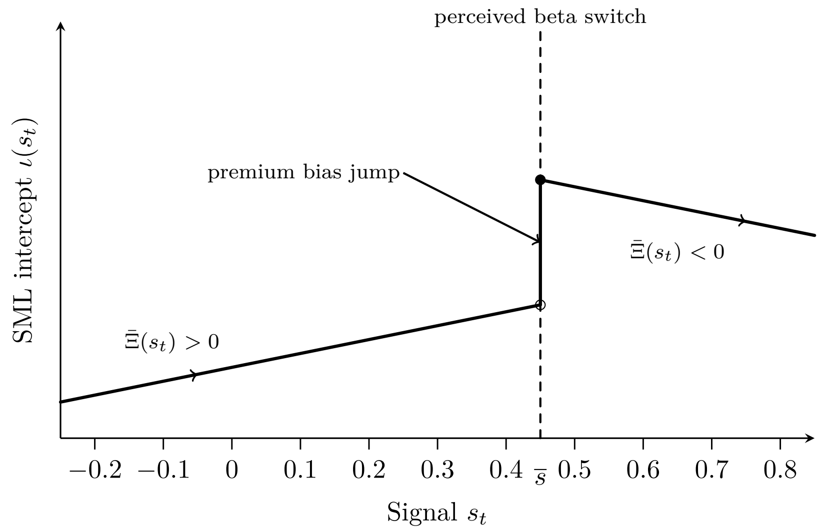

Figure 3: SML intercept as a function of news. This figure illustrates the model-implied SML intercept $\iota(s_t)$ as characterized in Proposition 3. Within each regime, the intercept varies linearly with the news realization s_t (expectation bias): the relationship is positive for $s_t < \bar{s}$ and negative for $s_t > \bar{s}$. At $s_t = \bar{s}$, perceived betas switch from inflated to deflated, inducing a discrete upward jump in the intercept (premium bias).

Figure 3 illustrates Proposition 3. Starting from the left, where news is bad and perceived betas are inflated, both biases push the intercept downward: the premium bias overstates factor risk, while the expectation bias causes prices to overreact to news. Prices fall too much, high-beta stocks become underpriced relative to low-beta stocks, and the SML steepens. As news improves toward zero, the expectation bias weakens and the intercept rises.

When news turns moderately positive but remains below the threshold ($0 < s_t < \bar{s}$), perceived betas are still inflated. The premium bias still lowers the intercept, but the expectation bias now works in the opposite direction: positive news, processed under inflated betas, overprices high-beta stocks and flattens the SML. At $s_t = \bar{s}$, perceived betas switch from inflated to deflated, generating a discrete upward jump of size $\bar{\pi}_H + \bar{\pi}_L$ through the premium bias. Beyond the threshold, positive news is processed under deflated betas, so the expectation bias turns negative and pulls the intercept back down.[12]

Proposition 3 implies a sharp asymmetry between good and bad news states, formalized in Corollary 2.

Corollary 2 (Asymmetry in $\iota(s_t)$ on good vs. bad news)**.** *For any $s > 0$,*

$$\iota(s) \ > \ \iota(-s), \qquad \textit{and} \qquad \mathbb{E}[\iota(s_t) \mid s_t > 0] \ > \ \mathbb{E}[\iota(s_t) \mid s_t < 0]\,. \tag{27}$$

[12] In the limiting case of risk neutrality, the premium bias disappears, the cutoff collapses to $s_t = 0$, and the discontinuity in the intercept vanishes. The intercept–news relation becomes continuous, with a kink at zero.

Equivalently, the empirical SML is expected to be flatter following good news than following bad news.

The asymmetry arises because the perceived-beta switch occurs when news is positive $s_t > \bar{s}$, so positive and negative news of equal magnitude do not generate symmetric distortions. Figure 4 illustrates the resulting pricing pattern, which is the theoretical counterpart of the empirical pattern documented in Figure 1 in the Introduction.

Relative to Figure 4, the empirical SMLs in Figure 1 are flatter overall across all states. This difference may reflect additional forces, such as borrowing constraints, that flatten the SML unconditionally. Yet such frictions shift the intercept upward in every state without changing the gap between good- and bad-news states. In other words, bad news partly offsets the unconditional flattening induced, for instance, by borrowing constraints, whereas good news amplifies it.[13] Existing theories therefore cannot eliminate the asymmetry implied by Corollary 2, nor—as discussed in the Introduction—can they explain it.

I next test the model's predictions empirically, evaluate their robustness to alternative explanations, and analyze their implications for betting-against-beta (BAB) strategies.

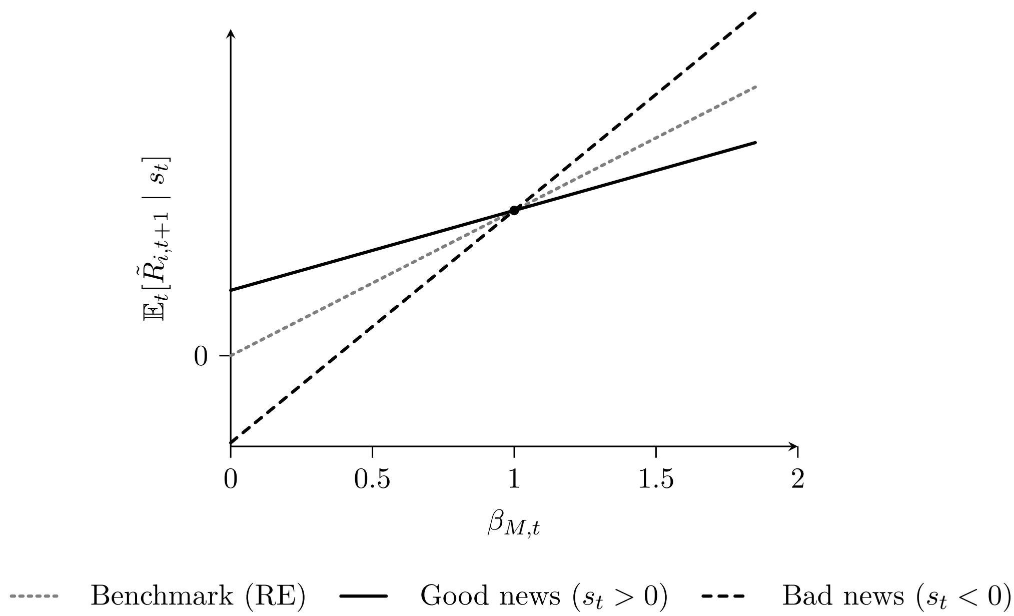

Figure 4: Model-implied SML on good and bad news. The solid (dashed) line depicts the expected SML following good (bad) news. Under beta ambiguity, the SML is flatter after good news ($s_t > 0$) and steeper after bad news ($s_t < 0$) relative to the rational-expectations benchmark (dotted).

[13]See the corresponding stylized Figure 10 in the Appendix.

3 Data and methodology

Data The baseline empirical analysis focuses on the U.S. equity market.[14] The sample spans 1930 to 2024 and includes all common stocks listed on the NYSE, NASDAQ, and Amex. I obtain stock-level data from the Center for Research in Security Prices (CRSP).[15] Stock returns are adjusted for delistings following Shumway (1997). Unless stated otherwise, all returns are expressed in excess of the risk-free rate.

I measure the market return using the CRSP value-weighted market portfolio and the risk-free rate using the one-month Treasury bill, both from Kenneth French's data library, which also provides the Carhart (1997) factor returns. I obtain an illiquidity factor based on the measure of Amihud (2002) from the database of Jensen, Kelly, and Pedersen (2023), daily U.S. and international betting-against-beta (BAB) factor returns from the AQR data library, and 25 size/book-to-market and 10 industry portfolio returns from Kenneth French's library.

I use additional data to construct control variables. Book-to-market ratios are obtained from Jensen et al. (2023). Aggregate disagreement is measured following Hong and Sraer (2016) using analyst forecasts from the Institutional Broker's Estimate System (I/B/E/S). Dividend yields of the S&P 500 are from Robert Shiller's database, while credit spreads, the adjusted National Financial Conditions Index (ANFCI), and the Consumer Price Index (CPI) are from the Federal Reserve Bank of St. Louis (FRED).[16]

Proxying for the sign and magnitude of news Tests of the theoretical predictions require a proxy for the sign and magnitude of macroeconomic news s_t. I follow a large empirical literature in using realized market excess returns (see, e.g., Abraham and Ikenberry (1994); McQueen et al. (1996); Chordia and Swaminathan (2000); Hou and Moskowitz (2005)). This proxy is well suited because market prices aggregate newly arriving information, and returns reveal its sign and magnitude without requiring direct observation of the underlying news s_t. I use a weekly sampling frequency to balance two competing concerns: at lower frequencies such as monthly, good and bad news within a period may offset each other, masking the alternating sign of underlying information arrival; at very high frequencies such as daily or intraday, price movements more likely reflect transitory noise than actual information.

[14] I extend the analysis to international equity markets when examining dynamic BAB strategies.

[15] The sample includes all securities in the CRSP universe with share codes (SHRCD) of 10 or 11 and exchange codes (EXCHCD) of 1, 2, 3, 31, 32, or 33.

[16] Data sources: Kenneth French's data library is available at `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`; the Jensen et al. (2023) factor database at `https://jkpfactors.com`; the AQR data library at `https://www.aqr.com`; Robert Shiller's database at `https://shillerdata.com`; and FRED at `https://fred.stlouisfed.org`.

I classify news states in two ways. Under the *all signals* classification, a period is labeled good (bad) whenever the market return exceeds (falls below) its long-run mean. Under the *strong signals* classification, the threshold is shifted by half a standard deviation, leaving an intermediate region unclassified; this restriction sharpens identification by excluding small price movements that may primarily reflect noise rather than meaningful information arrival. As a complementary analysis, I run an event study around scheduled macroeconomic announcements, which isolate periods in which the timing of information arrival is directly observable.

Market betas I estimate market betas using daily returns and account for nonsynchronous trading by including one week of lagged market returns in the regression, following Dimson (1979) and Lewellen and Nagel (2006). Let $R_{M,t}$ denote the market return. I estimate

$$R_{i,t} = \alpha_i + \beta_{i,M}^{(0)} R_{M,t} + \beta_{i,M}^{(1)} R_{M,t-1} + \beta_{i,M}^{(2)} \left(\frac{R_{M,t-2} + R_{M,t-3} + R_{M,t-4}}{3} \right) + \varepsilon_{i,t}, \tag{28}$$

defining the market beta as $\hat{\beta}_{i,M} = \beta_{i,M}^{(0)} + \beta_{i,M}^{(1)} + \beta_{i,M}^{(2)}$.

SML estimation I estimate the SML using weekly returns measured from the last trading day of week τ to the last trading day of week $\tau + 1$. Following Fama and MacBeth (1973), I run the cross-sectional regression

$$R_{i,\tau+1} = \iota_{\tau+1} + \phi_{\tau+1} \hat{\beta}_{i,M,\tau^-} + \varepsilon_{i,\tau+1}, \tag{29}$$

where $\hat{\beta}_{i,M,\tau^-}$ denotes the market beta estimated over a one-year rolling window ending at τ^-, the second-to-last trading day of week τ (typically Thursday). On that day I also compute the market excess return over the preceding week and classify the news state as described above. This timing is consistent with the model's two-date structure: mispricing arises at time τ when news is incorporated into prices and is reflected in realized returns over the subsequent period, $\tau + 1$. I introduce a one-day gap between news measurement and SML estimation for two reasons. First, it rules out look-ahead bias, which is particularly important for trading strategies that exploit ambiguity-induced mispricing (see Section 4.2): an investor observes the preceding week's market return at Thursday's close and can rebalance at Friday's close to trade on the predicted mispricing in week $\tau + 1$. Second, it reduces the possibility that sluggish price adjustment contaminates the results, a potential confound I discuss in more detail below.

Test assets In the baseline analysis, I use beta-sorted portfolios as test assets. At the end of each month, I estimate each stock's market beta over the preceding calendar year, sort stocks into beta deciles using NYSE breakpoints, and form value-weighted portfolios within each decile, requiring at least 120 daily return observations per stock. The ten portfolios are held for one month before being rebalanced again.

Table 1 reports summary statistics of characteristics across the beta-sorted portfolios, where P1 (P10) denotes the lowest (highest) beta decile. Both the average confidence-interval width $CI(\hat{\beta}_{M,\tau^-})$ and the beta-uncertainty measure $Unc(\hat{\beta}_{M,\tau^-})$ of Hollstein et al. (2020), which serve as proxies for beta ambiguity, increase monotonically from P1 to P10, with a P10–P1 spread that is large and significant at the 1% level.[17] This pattern confirms the assumed multiplicative ambiguity structure in equation (16) of the model. High-beta portfolios also differ systematically from low-beta portfolios along characteristics known to be priced in the cross-section (e.g., Fama and French, 1993; Carhart, 1997). The high-beta portfolio tends to have a lower average book-to-market ratio and smaller market capitalization, indicating a tilt toward small-growth firms.[18]

Control variables A potential concern is that autocorrelation and cross-correlation could mechanically generate a relation between lagged market returns and the SML intercept that is unrelated to beta ambiguity. Lo and MacKinlay (1990) document that returns on large stocks lead returns on small stocks, and McQueen et al. (1996) show that this pattern is asymmetric: small stocks adjust slowly to good market-wide news but quickly to bad news. This could mechanically affect the predicted mispricing asymmetry, given that high-beta portfolios tilt toward small stocks, as shown in Table 1. To ensure that these effects do not confound the results, I control for the risk factors of Carhart (1997) and for an illiquidity factor based on Amihud (2002).

[17]A detailed description of the construction of $Unc(\hat{\beta}_{M,\tau^-})$ is provided in the Appendix.

[18]Novy-Marx (2014) shows that returns on defensive equity strategies, such as low-beta portfolios, load significantly on standard risk factors.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P10-P1
$\hat{\beta}_{M,\tau^-}$	0.465	0.644	0.773	0.889	0.995	1.077	1.189	1.311	1.458	1.737	1.272***
											(82.79)
$CI(\hat{\beta}_{M,\tau^-})$	0.141	0.139	0.139	0.145	0.154	0.162	0.176	0.199	0.228	0.306	0.165***
											(33.39)
$Unc(\hat{\beta}_{M,\tau^-})$	0.279	0.284	0.292	0.321	0.346	0.369	0.409	0.455	0.519	0.651	0.372***
											(36.70)
$\overline{R}_{\tau+1}$	12.3	14.6	15.5	16.2	18.5	18.6	19.8	22.5	20.2	18.5	6.1
											(1.05)
$Size$	14.081	14.580	14.628	14.608	14.698	14.645	14.443	14.235	13.966	13.284	−0.796***
											(−5.22)
B/M	0.694	0.645	0.596	0.569	0.556	0.545	0.538	0.535	0.547	0.568	−0.123***
											(−2.97)
$Illiq$	0.354	0.171	0.141	0.134	0.124	0.140	0.173	0.210	0.252	0.416	0.062
											(1.17)
N	649.5	318.3	277.8	256.5	243.1	236.0	235.8	243.7	265.6	376.7	

Table 1: Summary statistics for beta-sorted portfolios. The table reports time-series averages of portfolio characteristics for ten beta-sorted portfolios. $\hat{\beta}_{M,\tau^-}$ denotes the portfolio's average market beta, with betas estimated weekly over a rolling one-year window. $CI(\hat{\beta}_{M,\tau^-})$ is the average width of the confidence band around the beta estimate, and $Unc(\hat{\beta}_{M,\tau^-})$ is the beta-uncertainty measure of Hollstein et al. (2020). $\overline{R}_{\tau+1}$ denotes the mean weekly excess return, in basis points. The variables $\hat{\beta}_M$, $CI(\hat{\beta}_M)$, $Unc(\hat{\beta}_M)$, and $\overline{R_{\tau+1}}$ are constructed at the portfolio level using weekly data. $Size$ is the average log market capitalization. The book-to-market ratio (B/M) follows Jensen et al. (2023). $Illiq$ is the logarithm of one plus the illiquidity measure of Amihud (2002), computed from data in the preceding calendar year. $Size$, B/M, and $Illiq$ are time-series averages of monthly value-weighted cross-sectional averages within each portfolio, calculated from stock-level data. N denotes the average number of stocks per portfolio. The final column reports the mean difference between portfolios P10 and P1. t-statistics for the P10–P1 spread, based on Newey–West standard errors with nine lags, are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

4 Empirical Results

The empirical analysis proceeds in three steps. Section 4.1 presents the paper's baseline results by examining the news-dependence of the SML intercept. Building on these findings, Section 4.2 provides a practical application by proposing a dynamic betting-against-beta (BAB) strategy. Section 4.3 then assesses the robustness of the findings along three dimensions: an event study around scheduled macroeconomic announcements; controls for alternative explanations; and an extended cross-section of test assets beyond beta-sorted portfolios.

4.1 Mispricing asymmetry across news states.

The model's central prediction is that the SML intercept is higher, on average, following good news than following bad news (Corollary 2). I evaluate this prediction by estimating weekly SMLs separately across good- and bad-news states. Table 2 reports the resulting estimates in basis points and shows strong support for the predicted asymmetry.

	Full	Good news		Bad news		Bad − Good	
		All	Strong	All	Strong	All	Strong
SML intercept ($\hat{\iota}_{\tau+1}$)	10.52***	17.68***	20.10***	2.23	−8.36	−15.45***	−28.46***
	(3.50)	(4.51)	(3.41)	(0.51)	(−1.18)	(−2.63)	(−3.08)
SML slope ($\hat{\phi}_{\tau+1}$)	5.90	−4.50	−6.98	17.68***	37.26***	22.18**	44.24***
	(1.30)	(−0.71)	(−0.75)	(2.79)	(3.47)	(2.47)	(3.11)
Market return ($\overline{R}_{M,\tau+1}$)	15.92***	12.46***	11.05	20.22***	29.15***	7.76	18.10
	(4.36)	(2.65)	(1.58)	(3.72)	(3.18)	(1.08)	(1.57)
Observations	4,905	2,635	1,273	2,268	1,184	4,903	2,457

Table 2: SML intercept, slope, and market return across news states. This table reports weekly means of the SML intercept ($\hat{\iota}_{\tau+1}$), SML slope ($\hat{\phi}_{\tau+1}$), and market excess return ($\overline{R}_{M,\tau+1}$), all in basis points, for the full sample and separately for good- and bad-news weeks. Under the all-signals classification, a week is classified as good (bad) if the lagged market excess return exceeds (falls below) its long-run mean. Under the strong-signals classification, the threshold is shifted by one-half of the long-run standard deviation, so that only weeks with sufficiently large positive or negative prior market returns are classified. The Bad − Good column reports the difference in conditional means, tested using a two-sample procedure that combines separate Newey–West long-run variance estimates for each group; negative values indicate a lower intercept following bad news. The sample runs from 1931 to 2024. t-statistics based on Newey–West standard errors with nine lags are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Under the all-signals classification, the SML intercept averages 17.68 basis points per week following good news versus 2.23 basis points following bad news, a differential of 15.45

basis points (approximately 8% annualized) that is significant at the 1% level. The SML slope mirrors this pattern. After good news, the slope is statistically indistinguishable from zero despite a positive market premium of 12.46 basis points; after bad news, it rises to 17.68 basis points, closely tracking the corresponding market premium of 20.22 basis points. These findings confirm the mispricing asymmetry previewed in Figure 1 and show that it persists when the SML is estimated with rolling-window rather than fixed in-sample betas.

	All signals		Strong signals	
	(1)	(2)	(3)	(4)
Constant	17.68***	21.60***	16.55***	18.20***
	(4.45)	(5.35)	(5.07)	(5.46)
$\mathbf{1}\{\text{bad news}_\tau\}$	−15.45***	−27.94***	−24.92***	−39.57***
	(−2.76)	(−4.96)	(−3.37)	(−5.35)
$R_{M,\tau+1}$		0.04		0.05*
		(1.62)		(1.69)
$\text{SMB}_{\tau+1}$		−0.28***		−0.29***
		(−2.69)		(−2.79)
$\text{HML}_{\tau+1}$		−0.02		−0.02
		(−0.32)		(−0.37)
$\text{Mom}_{\tau+1}$		0.17***		0.18***
		(4.18)		(4.19)
$\text{Illiq}_{\tau+1}$		0.02		0.03
		(0.24)		(0.30)
N	4,903	4,903	4,903	4,903
R^2	0.00	0.06	0.00	0.06

Table 3: News-dependent SML intercept with factor controls. This table reports estimates from time-series regressions of the weekly SML intercept $\hat{\iota}_{\tau+1}$ on a bad-news indicator $\mathbf{1}\{\text{bad news}_\tau\}$ and factor controls. The bad-news indicator equals one in weeks following a market excess return below the relevant threshold and zero otherwise; its coefficient therefore measures the average reduction in the SML intercept following bad news relative to good news. The constant captures the average intercept in good-news weeks. Under the all-signals specification (columns 1–2), the threshold is the long-run average market excess return. Under the strong-signals specification (columns 3–4), the threshold is shifted downward by one-half of the unconditional standard deviation, isolating weeks with particularly negative prior market performance. Columns (2) and (4) augment the baseline with contemporaneous excess returns on the market, SMB, HML, and momentum factors of Carhart (1997) and an illiquidity factor based on Amihud (2002). The sample runs from 1931 to 2024. t-statistics based on Newey–West standard errors with nine lags are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

Under the strong-signals classification, the asymmetry nearly doubles to 28.46 basis points, driven almost entirely by the bad-news subsample: restricting to strong bad news lowers the intercept by 10.59 to −8.36 basis points, while the good-news intercept rises only weakly by 2.42 to 20.10 basis points. The asymmetry is therefore not only in levels but also in sensitivity: the intercept falls as bad news intensifies, while responding far less to good news. This pattern is consistent with Proposition 3: the intercept declines monotonically with bad news, whereas its relation to good news is non-monotonic.

The full-sample result in Table 2 places the conditional results in perspective. Consistent with existing studies, the SML is too flat relative to the CAPM benchmark (e.g., Black, Jensen, and Scholes, 1972; Baker, Bradley, and Wurgler, 2011; Frazzini and Pedersen, 2014). The asymmetry emerges around this baseline flatness: bad news tends to offset it, steepening the SML toward the CAPM prediction, whereas good news reinforces it, flattening the SML further. As anticipated in Section 2.3, this pattern suggests that beta ambiguity adds a state-dependent component to mechanisms that flatten the SML unconditionally, such as borrowing constraints.

The mispricing asymmetry is robust to controlling for standard risk factors. As noted in Section 3, beta correlates with other priced characteristics, such as size and value, raising the possibility of confounding effects. To address this concern, I regress the weekly SML intercept on a bad-news indicator, with and without the four factors of Carhart (1997) and the illiquidity factor based on Amihud (2002). Table 3 reports the results. The coefficient on the bad-news indicator is negative across all specifications and, notably, increases in magnitude once these controls are included. Thus, the estimated asymmetry strengthens rather than weakens after absorbing exposure to other risk factors.

The mispricing asymmetry is persistent over time. Figure 5 plots cumulative SML intercepts by news state under the all-signals and strong-signals classifications, with and without factor controls. In all specifications, the cumulative paths diverge steadily between good- and bad-news states, although the pace of divergence weakens somewhat toward the end of the sample. The asymmetry is therefore persistent rather than concentrated in a few isolated episodes.

Taken together, the evidence supports the mispricing asymmetry predicted in Corollary 2: the SML intercept is significantly higher following good news than following bad news. The model, however, delivers a sharper prediction regarding the full shape of this relation (Proposition 3), which I now evaluate directly.

The SML intercept as a function of news. The piecewise relation predicted by Proposition 3 cannot be captured by a standard linear specification. I therefore estimate the conditional mean

$$m(r) = \mathbb{E}\big[\hat{\iota}_{\tau+1} \mid R_{M,\tau^-} = r\big] \tag{30}$$

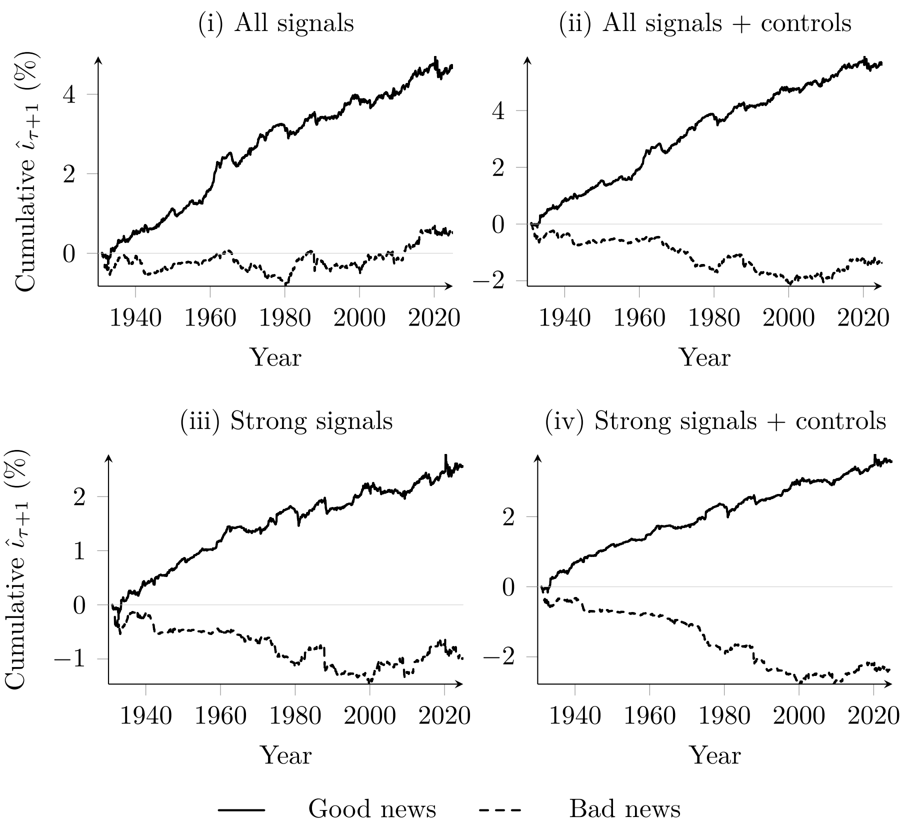

Figure 5: Cumulative SML intercepts by news state and signal strength. This figure plots cumulative sums of the weekly SML intercept from 1931 through 2024, computed separately for good- and bad-news weeks. Panels (i) and (ii) use the all-signals classification, under which a week is classified as good (bad) news if the lagged market excess return exceeds (falls below) the long-run average market excess return. Panels (iii) and (iv) apply the strong-signals classification, restricting good (bad) news weeks to those in which the lagged market excess return exceeds (falls below) the long-run mean by at least one-half of its unconditional standard deviation. Panels (i) and (iii) show raw cumulative intercepts. Panels (ii) and (iv) show cumulative intercepts orthogonalized with respect to the contemporaneous Carhart (1997) four-factor returns and an illiquidity factor based on Amihud (2002). All values are in percent. The solid (dashed) line corresponds to good- (bad-) news weeks.

nonparametrically using locally weighted scatterplot smoothing (LOWESS). This procedure fits local linear regressions around each value of the news proxy R_{M,τ^-}, assigning greater weight to nearby observations without imposing a global functional form. To account for potential confounding through priced risk factors, I also construct a control-adjusted fit by first regressing the SML intercept on the five factor controls and then applying LOWESS to the residuals.

Figure 6 plots the LOWESS estimates of the conditional mean of the weekly SML intercept as a function of news. The figure reveals a hump-shaped relation between the SML intercept and news, closely matching the pattern predicted by Proposition 3. The intercept rises from deeply negative news, peaks near $R_{M,\tau^-} \approx 5\%$, and declines for stronger positive

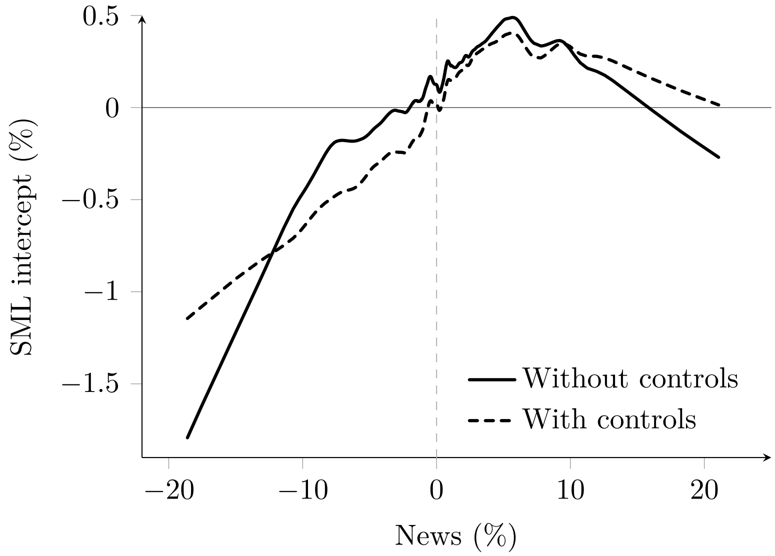

Figure 6: SML intercept as a function of news. This figure plots LOWESS estimates of the conditional mean of the weekly SML intercept $\hat{\iota}_{\tau+1}$ as a function of news, proxied by the lagged weekly market excess return R_{M,τ^-}. Both quantities are in percent. Both fits use a bandwidth of $h = 0.2$, so that each local regression uses the nearest 20% of observations. The solid line shows the raw fit; the dashed line shows the control-adjusted fit, obtained by applying LOWESS to the residuals from a regression of $\hat{\iota}_{\tau+1}$ on the contemporaneous Carhart (1997) four-factor returns and an illiquidity factor based on Amihud (2002). The sample runs from 1931 to 2024.

news. The pattern is robust to factor controls.

The only difference from the theoretical prediction is that the discontinuity is not visually apparent in Figure 6. A natural explanation is that the threshold $\bar{s}$ at which the discontinuity occurs depends on risk aversion and signal precision. While these parameters are fixed in the model, they likely vary over the nearly century-long sample period. As a result, the threshold shifts over time, so aggregating across periods smooths the discontinuity into a hump rather than a sharp jump.

Overall, the evidence shows that the SML intercept varies predictably with news, both across good- and bad-news states and over the full range of news realizations. This predictability can be exploited in trading strategies, an implication I examine next.

4.2 Application: A dynamic betting-against-beta strategy

The asset-pricing implications of beta ambiguity translate directly into predictions for betting-against-beta (BAB) returns. BAB exploits a flat SML by buying low-beta stocks, shorting high-beta stocks, and scaling the two legs to construct a market-neutral portfolio (Frazzini and Pedersen, 2014). Since BAB has zero market beta by construction, its expected return maps directly into the SML intercept. The strategy therefore performs well when the SML is too flat but crashes when the SML becomes too steep. Hence, the model predicts that BAB crashes are concentrated after strongly bad news, when the SML steepens sharply and the intercept reaches its trough.

Figure 7 confirms this prediction. It reports average weekly returns on the BAB strategy of Frazzini and Pedersen (2014), conditional on the lagged market return falling 0.5, 1.0, and 1.5 standard deviations below its long-run mean, alongside the unconditional full-sample mean. Results are shown for the U.S. market and for three international equity markets: Global, Europe, and Pacific.[19] In every market, BAB returns are negative after bad news and decline monotonically with news strength. In the most extreme states, weekly losses exceed 50 basis points (-26% annualized), consistent with the model's prediction that stronger bad news steepens the SML and triggers BAB crashes.

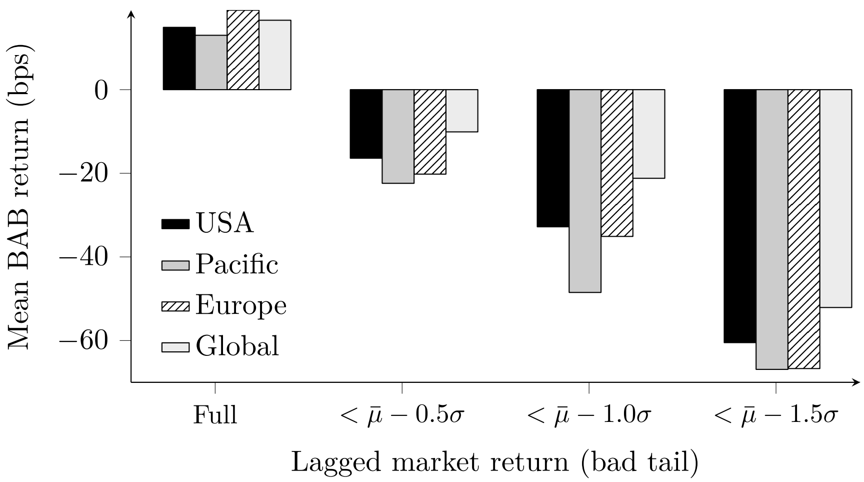

Figure 7: BAB crashes following bad news. This figure plots mean weekly returns of the betting-against-beta (BAB) factor in basis points, conditional on the strength of the prior week's news signal. For each region, the Full bar reports the unconditional mean BAB return. The remaining bars restrict the sample to weeks following lagged market excess returns that fall below the long-run mean ($\overline{\mu}$) by at least one-half, one, and one-and-a-half times the long-run standard deviation(σ), respectively. Weekly BAB returns are constructed from daily data following Frazzini and Pedersen (2014) and obtained from the AQR data library.

The predictability of BAB crashes motivates a dynamic strategy that conditions BAB exposure on lagged market returns. Although many implementations could exploit this predictability, I focus on a simple rule that keeps turnover low. The strategy is long BAB when the lagged market return is at or above its long-run mean and switches to short BAB when the lagged market return falls at least 1.5 standard deviations below the mean. In intermediate states, the strategy maintains its previous position. Thus, the strategy preserves the standard long-BAB exposure in normal times while exploiting crash-state predictability

[19]Sample periods for international equity markets are shorter due to limited data availability. The Global series begins in February 1987; the Pacific series in December 1988; and the European series in February 1989.

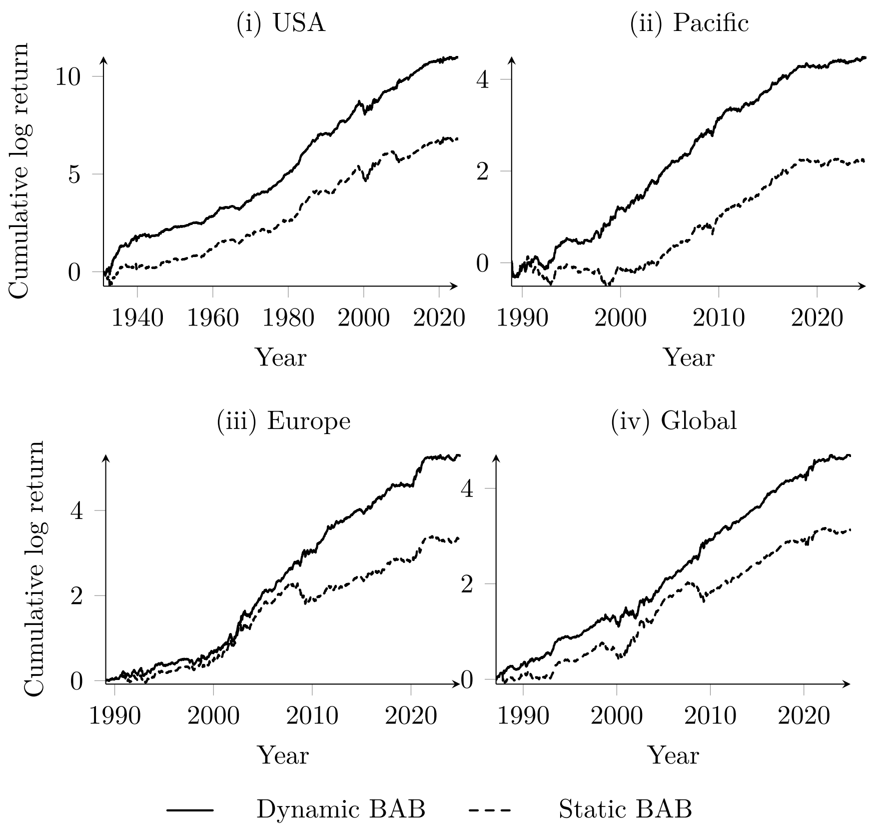

Figure 8: Cumulative performance of dynamic BAB versus static BAB. This figure plots cumulative weekly log returns of the dynamic BAB strategy (solid) and the standard betting-against-beta (BAB) factor (dashed) across four equity markets. The dynamic BAB strategy allocates to BAB based on the lagged weekly market excess return R_{M,τ^-} relative to its long-run distribution. It takes a long BAB position whenever R_{M,τ^-} is at or above the long-run mean, switches to a short BAB position whenever R_{M,τ^-} falls at or below the long-run mean minus one-and-a-half times the long-run standard deviation, and otherwise maintains its current position. Weekly BAB returns are constructed from daily data following Frazzini and Pedersen (2014) and obtained from the AQR data library.

by shorting BAB after extreme bad-news realizations.[20]

Figure 8 shows that the dynamic strategy outperforms static BAB in all four markets. The cumulative return paths diverge steadily, with the outperformance accumulating across repeated predicted crash states rather than a single episode.

Table 4 quantifies this outperformance. The dynamic strategy raises annualized returns by 4.1 to 6.2 percentage points relative to static BAB, with differences significant at the 1% level. Volatility is nearly unchanged because the strategy reverses BAB exposure in predicted crash states rather than moving to cash. The higher returns therefore translates into large Sharpe-ratio improvements: They rise from 0.60–0.97 for static BAB to 1.16–1.44 for the

[20] The logic is related to Daniel and Moskowitz (2016), who mitigate momentum crashes by dynamically adjusting momentum exposure to market conditions. Here, lagged market returns play an analogous role for BAB by identifying states that predict BAB crashes.

dynamic strategy. This outperformance requires limited additional trading, with fewer than five position switches per year on average.

	USA	Pacific	Europe	Global
Panel A: Returns				
Dynamic BAB mean (%, p.a.)	12.13***	13.00***	15.32***	12.70***
	(10.36)	(6.42)	(7.29)	(7.97)
Static BAB mean (%, p.a.)	7.70***	6.77***	9.88***	8.62***
	(6.31)	(3.19)	(4.77)	(5.00)
Dynamic − Static mean (%, p.a.)	4.43***	6.23***	5.44***	4.08***
	(4.79)	(4.04)	(3.59)	(2.83)
Panel B: Risk and Sharpe ratios				
Dynamic BAB volatility (%, p.a.)	9.90	11.24	11.20	8.83
Static BAB volatility (%, p.a.)	9.99	11.34	11.32	8.93
Dynamic BAB Sharpe ratio	1.22	1.16	1.37	1.44
Static BAB Sharpe ratio	0.77	0.60	0.87	0.97
Dynamic − Static Sharpe ratio	0.45	0.56	0.50	0.47
Panel C: Implementation				
Switches per year	3.96	5.16	4.72	5.07
Hit rate (%)	68.36	62.03	66.47	68.91
Observations	4,904	1,883	1,874	1,978

Table 4: Performance of dynamic BAB versus static BAB. This table compares the dynamic BAB strategy with the standard betting-against-beta (BAB) factor of Frazzini and Pedersen (2014) in U.S., Pacific, European, and global equity markets. The dynamic BAB strategy conditions on the lagged weekly market excess return in each region. It is long BAB when the lagged market return is at or above its long-run mean, short BAB when the lagged market return is at least 1.5 standard deviations below the mean, and otherwise maintains its previous position. Means and volatilities are annualized from weekly returns and reported in percent; Sharpe ratios are annualized using 52 weeks per year. t-statistics for mean returns, reported in parentheses, use Newey–West standard errors with nine lags. Switches per year is the average number of changes between long and short BAB positions. Hit rate is the fraction of switches after which the dynamic BAB strategy earns a positive cumulative return until the next switch. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

4.3 Additional Evidence and Robustness

I provide three sets of additional evidence. First, I use scheduled macroeconomic announcements as a complementary test in which the timing of information arrival is directly observable. The remaining two analyses serve as robustness checks: I show that time variation in borrowing constraints and investor disagreement, two leading explanations for a flat SML, does not account for the asymmetry across news states, and verify that the results extend beyond beta-sorted portfolios to a broader cross-section of test assets.

Event-study evidence from macroeconomic announcements. I substantiate the baseline findings with a complementary test based on scheduled macroeconomic announcements, which release public information at predetermined times. These announcements create short windows in which the timing of news arrival is directly observable and prices adjust to a common information shock, offering an additional setting to test the mispricing asymmetry.[21]

I consider inflation and unemployment releases and Federal Open Market Committee (FOMC) announcements, following the categorization of Savor and Wilson (2013, 2014). Inflation and unemployment release dates are available from 1958 and are obtained from the U.S. Bureau of Labor Statistics. FOMC announcement dates are from the Federal Reserve and are available from 1978.[22] The sample contains 1,988 announcement dates.

For each announcement, I define the event window as the release date and the preceding trading day, allowing for partial anticipation and early price incorporation. I use this interval as the empirical counterpart of the model's information-arrival period: prices adjust to a common news shock, and the post-announcement period captures the subsequent realization of the mispricing. I classify an announcement as bad if the cumulative market return over the event window is more than one-half standard deviation below its sample mean. Market betas are estimated over the one-year window preceding the announcement, and the SML is estimated over horizons $h \in \{0, \dots, 15\}$ trading days after the announcement, where $h = 0$ denotes the close of the announcement day.

Figure 9 plots the cumulative average difference in the SML intercept following bad versus non-bad announcements, $\hat{\iota}_h^{\text{bad}}$, over the three weeks after the announcement. The differential declines steadily and stabilizes around -40 basis points without controls and -60 basis points with controls. This differential reflects the asymmetry predicted by the model and unfolds over roughly two weeks following the announcement. This event-study evidence therefore corroborates the baseline results in a setting where the timing of information arrival is directly observed.

[21] Scheduled macroeconomic announcements, and FOMC announcements in particular, trigger large reactions in equity and bond prices, reveal information about both current policy and the future policy path, and command sizable announcement premia; see, among others, Bernanke and Kuttner (2005), Gürkaynak, Sack, and Swanson (2005), Savor and Wilson (2013), Lucca and Moench (2015), Hanson and Stein (2015), Cieslak, Morse, and Vissing-Jorgensen (2019), and Ai, Bansal, and Guo (2023).

[22] Inflation announcements use CPI releases prior to February 1972 and PPI releases thereafter, because PPI data are released a few days earlier, reducing the news content of subsequent CPI releases. Prior to February 1994, I assume FOMC decisions became public one trading day after the meeting, following Kuttner (2001).

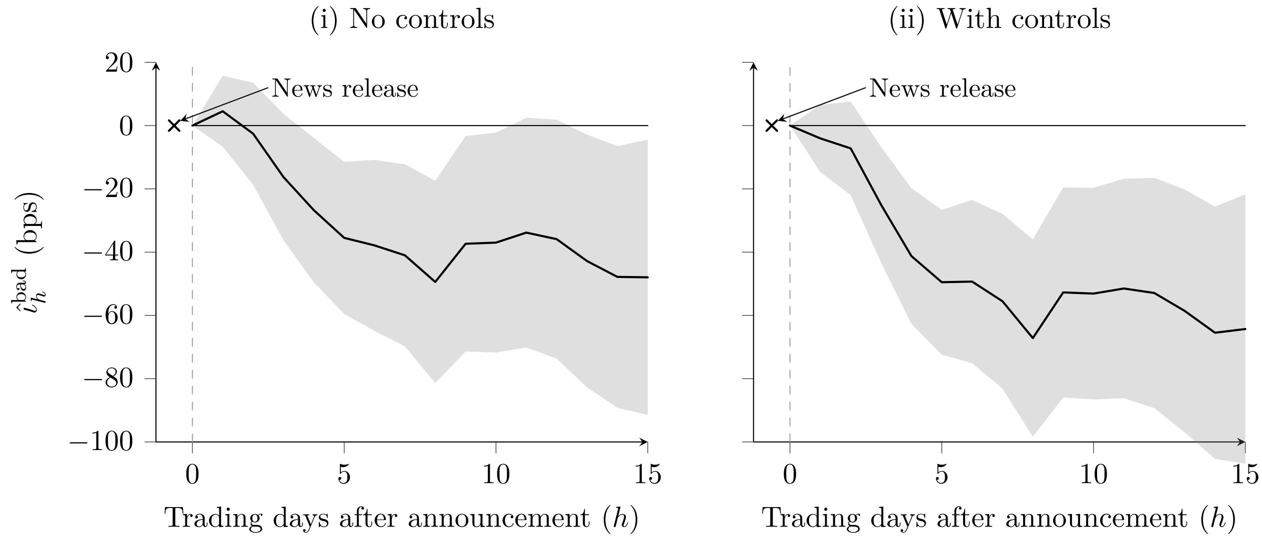

Figure 9: SML intercept following bad versus non-bad macroeconomic announcements. The solid line plots the cumulative average difference in the weekly SML intercept between bad and non-bad announcement weeks, $\hat{\iota}_h^{\text{bad}}$, over horizons $h \in \{0, \dots, 15\}$ trading days following the news release. An announcement is classified as bad if the cumulative market return over the announcement window falls more than one-half of a standard deviation below its historical mean. The shaded region denotes the 95% heteroskedasticity-robust confidence interval. Panel (i) reports results without controls. Panel (ii) controls for cumulative realizations of the Carhart (1997) four factors and an illiquidity factor based on Amihud (2002). The sample includes 1,988 announcements between 1958 and 2024.

Alternative explanations. I next examine whether the asymmetry reflects time variation in the two canonical mechanisms behind a flat SML: borrowing constraints and investor disagreement. As discussed in the Introduction, a borrowing-constraint explanation would require constraints to tighten on good news relative to bad news to produce the observed asymmetry. A disagreement explanation would require belief dispersion to rise on good news relative to bad news. Both requirements run counter to the usual cyclical behavior of these frictions: constraints tighten and disagreement rises in bad times, not good times (Brunnermeier and Pedersen, 2008; Geanakoplos, 2010; Adrian and Shin, 2013; Cujean and Hasler, 2017). The tests below confirm this directly: controlling for empirical proxies of these frictions do not affect the documented mispricing asymmetry.

I proxy for the tightness of borrowing constraints using the Adjusted National Financial Conditions Index (ANFCI).[23] I measure investor disagreement, following Hong and Sraer (2016), as the dispersion in analysts' long-term growth forecasts. Additional macro controls include inflation, the dividend yield, and the credit spread. All macro controls except ANFCI are standardized.[24] The sample is restricted to 1982–2024 because of limited data availability.

[23] Frazzini and Pedersen (2014) use the TED spread as a proxy for borrowing constraints. Following the LIBOR manipulation scandals, the TED spread has been phased out and is unavailable for the last part of my sample.

[24] ANFCI is not standardized because it is already constructed as a mean-reverting index.

Table 5 shows that the asymmetry is essentially unchanged after controlling for these mechanisms. The bad-news coefficient ranges from −36.5 to −38.9 basis points across specifications and remains significant at the 1% level throughout. Among the additional variables, only disagreement enters positively and is marginally significant, consistent with Hong and Sraer (2016) and Andrei et al. (2023): higher belief dispersion is associated with a flatter SML. I conclude that the documented mispricing asymmetry does not appear to be driven by time variation in borrowing constraints or investor disagreement.

Alternative test assets. As a final robustness check, I extend the set of test assets to include the 25 Fama–French size and book-to-market portfolios and 10 industry portfolios, as in Savor and Wilson (2014) and Hendershott, Livdan, and Rösch (2020). The results, reported in Appendix 6.2.2, confirm the baseline findings.

Tables 6 and 7 show that the SML intercept remains significantly higher following good news than bad news under both the all-signals and strong-signals classifications, with and without factor controls. The LOWESS estimate in Figure 14 displays the same hump-shaped pattern as in the baseline analysis: the intercept is lowest under deeply negative news, peaks under moderately positive news, and declines for strongly positive signals, consistent with Proposition 3. The control-adjusted fit preserves this shape, apart from a left-tail kink that likely reflects nonlinear interactions between factor returns and extreme negative news realizations.

Finally, the cumulative intercept paths in Figure 13 diverge persistently over the full sample horizon under all specifications. Overall, extending the set of test assets leaves the main results unchanged.

5 Conclusion

I show that beta ambiguity generates asymmetric mispricing across news states in the cross section of stock returns. Under ambiguity aversion, perceived factor betas vary with the sign and magnitude of factor news: bad news inflates perceived betas, whereas sufficiently good news deflates them relative to true betas. This mechanism produces a state-dependent SML intercept that is piecewise linear in news, with a peak in the good-news region and a trough following deeply negative news.

Empirical tests using weekly returns on beta-sorted portfolios of U.S. stocks over the period 1930–2024 strongly support these predictions. The asymmetry is robust to controls for borrowing constraints, investor disagreement, standard risk factors, and alternative test assets, and is corroborated by an event study around scheduled macroeconomic an-

	All signals		Strong signals	
	(1)	(2)	(3)	(4)
Constant	15.833***	16.058***	15.934***	15.516***
	(3.11)	(3.12)	(3.06)	(3.02)
$\mathbf{1}\{\text{bad news}_\tau\}$	−37.887***	−38.860***	−38.274***	−36.508***
	(−2.81)	(−2.86)	(−2.90)	(−2.79)
$R_{M,\tau+1}$	0.063	0.064	0.064	0.061
	(1.38)	(1.40)	(1.40)	(1.35)
$SMB_{\tau+1}$	−0.336**	−0.345**	−0.346**	−0.347**
	(−2.03)	(−2.08)	(−2.08)	(−2.06)
$HML_{\tau+1}$	0.078	0.074	0.074	0.076
	(0.69)	(0.66)	(0.65)	(0.66)
$Mom_{\tau+1}$	0.273***	0.274***	0.274***	0.277***
	(4.95)	(4.97)	(4.96)	(5.02)
$Illiq_{\tau+1}$	−0.096	−0.089	−0.088	−0.094
	(−0.62)	(−0.57)	(−0.57)	(−0.60)
$Disagreement_\tau$		6.118	5.969	10.628*
		(1.27)	(1.26)	(1.75)
$ANFCI_\tau$			−2.136	−15.100
			(−0.37)	(−1.55)
Dividend $Yield_\tau$				10.100
				(1.55)
Credit $spread_\tau$				11.512
				(1.24)
$Inflation_\tau$				−2.116
				(−0.44)
N	2,244	2,244	2,244	2,244
R^2	0.121	0.121	0.121	0.125

Table 5: SML intercept asymmetry controlling for borrowing constraints and disagreement. This table reports time-series regressions of the weekly SML intercept $\hat{\iota}_{\tau+1}$ on a bad-news indicator and control variables over the time period from 1982 through 2024. Under the all-signals classification, the bad-news indicator equals one when the lagged market excess return is below its long-run mean. Under the strong-signals classification, it equals one when the lagged market return is more than one-half standard deviation below its long-run mean; weeks above the corresponding good-news threshold form the omitted category. All columns include contemporaneous Carhart (1997) four-factor returns and an illiquidity factor based on Amihud (2002). Column (2) adds aggregate disagreement, measured as the dispersion in analysts' long-term growth forecasts following Hong and Sraer (2016). Column (3) adds the Adjusted National Financial Conditions Index (ANFCI) of the Federal Reserve Bank of Chicago. Column (4) adds the log dividend yield, the Moody's seasoned Baa–Aaa credit spread, and CPI inflation. Macro controls are standardized to zero mean and unit variance, except for ANFCI. t-statistics based on Newey–West standard errors with nine lags are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

nouncements. The model further implies that crashes in betting-against-beta strategies are predictable: a simple dynamic rule that conditions BAB exposure on lagged market returns raises the Sharpe ratio by 0.48 on average across U.S. and international equity markets.

The broader implication of the analysis is that relaxing the assumption that investors know the true betas, an assumption that is difficult to defend in practice, can account for patterns in the cross section of returns that existing theories of a flat SML cannot explain. The model is deliberately kept simple to isolate the key mechanism in its most transparent form. Extending the framework to multiple factors, heterogeneous beliefs, or a multi-period setting along the lines of Epstein and Schneider (2003) offers natural avenues for future research.

References

Abraham, A. and D. L. Ikenberry (1994). The individual investor and the weekend effect. *The Journal of Financial and Quantitative Analysis 29*(2), 263–277.

Adrian, T. and F. Franzoni (2009). Learning about beta: Time-varying factor loadings, expected returns, and the conditional capm. *Journal of Empirical Finance 16*(4), 537–556.

Adrian, T. and H. S. Shin (2013, 10). Procyclical leverage and value-at-risk. *The Review of Financial Studies 27*(2), 373–403.

Ahn, D., S. Choi, D. Gale, and S. Kariv (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics 5*(2), 195–223.

Ai, H., R. Bansal, and H. Guo (2023, November). Macroeconomic announcement premium. (31923).

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets 5*(1), 31–56.

Andrei, D., J. Cujean, and M. Wilson (2023, 02). The Lost Capital Asset Pricing Model. *The Review of Economic Studies 90*(6), 2703–2762.

Armstrong, C. S., S. Banerjee, and C. Corona (2013, 01). Factor-loading uncertainty and expected returns. *The Review of Financial Studies 26*(1), 158–207.

Baker, M., B. Bradley, and J. Wurgler (2011). Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal 67*(1), 40–54.

Barillas, F., L. P. Hansen, and T. J. Sargent (2009). Doubts or variability? *Journal of Economic Theory 144*(6), 2388–2418. Dynamic General Equilibrium.

Bernanke, B. S. and K. N. Kuttner (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of Finance 60*(3), 1221–1257.

Black, F. (1972). Capital market equilibrium with restricted borrowing. *The Journal of Business 45*(3), 444–455.

Black, F., M. Jensen, and M. Scholes (1972). The capital asset pricing model: Some empirical tests. *Studies in the Theory of Capital Markets.*

Boguth, O. and M. Simutin (2018). Leverage constraints and asset prices: Insights from mutual fund risk taking. *Journal of Financial Economics 127*(2), 325–341.

Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. R. Zame (2010, 01). Ambiguity in asset markets: Theory and experiment. *The Review of Financial Studies 23*(4), 1325–1359.

Brunnermeier, M. K. and L. H. Pedersen (2008, 11). Market liquidity and funding liquidity. *The Review of Financial Studies 22*(6), 2201–2238.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance 52*(1), 57–82.

Chen, Z. and L. Epstein (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica 70*(4), 1403–1443.

Chordia, T. and B. Swaminathan (2000). Trading volume and cross-autocorrelations in stock returns. *The Journal of Finance 55*(2), 913–935.

Cieslak, A., A. Morse, and A. Vissing-Jorgensen (2019). Stock returns over the fomc cycle. *The Journal of Finance 74*(5), 2201–2248.

Cujean, J. and M. Hasler (2017). Why does return predictability concentrate in bad times? *The Journal of Finance 72*(6), 2717–2758.

Daniel, K. and T. J. Moskowitz (2016). Momentum crashes. *Journal of Financial Economics 122*(2), 221–247.

Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics 7*(2), 197–226.

Dow, J. and S. R. d. C. Werlang (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica 60*(1), 197–204.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics 75*(4), 643–669.

Epstein, L. G. and M. Schneider (2003, November). Recursive multiple-priors. *Journal of Economic Theory 113*(1), 1–31.

Epstein, L. G. and M. Schneider (2008). Ambiguity, information quality, and asset pricing. *The Journal of Finance 63*(1), 197–228.

Epstein, L. G. and M. Schneider (2010). Ambiguity and asset markets. *Annual Review of Financial Economics 2*(1), 315–346.

Epstein, L. G. and T. Wang (1994). Intertemporal asset pricing under knightian uncertainty. *Econometrica 62*(2), 283–322.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fama, E. F. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy 81*(3), 607–636.

Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics 111*(1), 1–25.

Geanakoplos, J. (2010). The leverage cycle. *NBER Macroeconomics Annual 24*, 1–66.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Gollier, C. (2011). Portfolio choices and asset prices: The comparative statics of ambiguity aversion. *The Review of Economic Studies 78*(4), 1329–1344.

Guidolin, M. and F. Rinaldi (2013, February). Ambiguity in asset pricing and portfolio choice: a review of the literature. *Theory and Decision 74*(2), 183–217.

Gürkaynak, R. S., B. Sack, and E. Swanson (2005, March). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review 95*(1), 425–436.

Handa, P. and S. C. Linn (1993). Arbitrage pricing with estimation risk. *Journal of Financial and Quantitative Analysis 28*(1), 81–100.

Hansen, L. and T. J. Sargent (2001, May). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

Hansen, L. P. and T. J. Sargent (2007). *Robustness.* Princeton University Press.

Hanson, S. G. and J. C. Stein (2015). Monetary policy and long-term real rates. *Journal of Financial Economics 115*(3), 429–448.

Hayashi, T. and J. Miao (2011). Intertemporal substitution and recursive smooth ambiguity preferences. *Theoretical Economics 6*(3), 423–472.

Hedegaard, E. (2018). Time-varying leverage demand and predictability of betting-against-beta. Working paper, AQR Capital Management.

Hendershott, T., D. Livdan, and D. Rösch (2020). Asset pricing: A tale of night and day. *Journal of Financial Economics 138*(3), 635–662.

Hollstein, F., M. Prokopczuk, and C. Wese Simen (2020). Beta uncertainty. *Journal of Banking and Finance 116*, 105834.

Hong, H. and D. A. Sraer (2016). Speculative betas. *The Journal of Finance 71*(5), 2095–2144.

Hou, K. and T. J. Moskowitz (2005, 10). Market frictions, price delay, and the cross-section of expected returns. *The Review of Financial Studies 18*(3), 981–1020.

Illeditsch, P. K. (2011). Ambiguous information, portfolio inertia, and excess volatility. *The Journal of Finance 66*(6), 2213–2247.

Jensen, T. I., B. Kelly, and L. H. Pedersen (2023). Is there a replication crisis in finance? *The Journal of Finance 78*(5), 246–2518.

Ju, N. and J. Miao (2012). Ambiguity, learning, and asset returns. *Econometrica 80*(2), 559–591.

Jylha, P. (2018, June). Margin requirements and the security market line. *Journal of Finance 73*(3), 1281–1321.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Knight, F. H. (1921). *Risk, Uncertainty and Profit.* Boston and New York: Houghton Mifflin Company. Library of Congress Control Number 21012860.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics 47*(3), 523–544.

Lewellen, J. and S. Nagel (2006, November). The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics 82*(2), 289–314.

Lo, A. W. and A. C. MacKinlay (1990). When are contrarian profits due to stock market overreaction? *The Review of Financial Studies 3*(2), 175–205.

Lu, Z. and Z. Qin (2021). Leveraged funds and the shadow cost of leverage constraints. *The Journal of Finance 76*(3), 1295–1338.

Lucca, D. O. and E. Moench (2015). The pre-fomc announcement drift. *The Journal of Finance 70*(1), 329–371.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance 7*(1), 77–91.

McQueen, G., P. Michael, and S. Thorley (1996). Delayed reaction to good news and the cross-autocorrelation of portfolio returns. *The Journal of Finance 51*(3), 889–919.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*(2), 145–161.

Newey, W. K. and K. D. West (1987, May). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–08.

Novy-Marx, R. (2014, October). Understanding defensive equity. (20591).

Rieger, M. O. and M. Wang (2012). Can ambiguity aversion solve the equity premium puzzle? survey evidence from international data. *Finance Research Letters 9*(2), 63–72.

Savor, P. and M. Wilson (2013). How much do investors care about macroeconomic risk? evidence from scheduled economic announcements. *The Journal of Financial and Quantitative Analysis 48*(2), 343–375.

Savor, P. and M. Wilson (2014). Asset pricing: A tale of two days. *Journal of Financial Economics 113*(2), 171–201.

Shanken, J. (1992). On the estimation of beta-pricing models. *The review of financial studies 5*(1), 1–33.

Shumway, T. (1997). The delisting bias in crsp data. *The Journal of Finance 52*(1), 327–340.

Vasicek, O. A. (1973). A note on using cross-sectional information in bayesian estimation of security betas. *The Journal of Finance 28*(5), 1233–1239.

Şahin, B. C. and S. Danışoğlu (2022). Ambiguity and asset pricing: An empirical investigation for an emerging market. *International Review of Financial Analysis 84*, 102338.

6 Appendix

6.1 Theoretical Appendix

6.1.1 Heterogeneous Asset Supply

This appendix examines the robustness of the main results to heterogeneous asset supply. In the baseline model, each stock i has $1/N$ shares outstanding, so that the market portfolio is $\boldsymbol{M} = \frac{1}{N}\mathbf{1}$. I now replace this assumption with stock-specific supply $m_i > 0$, where $\sum_{i=1}^{N} m_i = 1$, so that $\boldsymbol{M} = (m_1, \ldots, m_N)^\top$. All other assumptions are maintained. In particular, the beta of the market portfolio is normalized to $\overline{\beta} = \boldsymbol{M}^\top \boldsymbol{\beta} = 1$.

Perceived betas and the ambiguity premium. The worst-case selection rule in Propositions 1 and 2 depends only on the portfolio beta $b = \boldsymbol{\Theta}^\top \boldsymbol{\beta}$ and on the requirement that the market portfolio $\boldsymbol{M}$ has strictly positive entries. Since both properties hold for any supply vector $\boldsymbol{M} > 0$ with $\overline{\beta} = 1$, the selection rule and the state-dependent perceived-beta switch at threshold $\overline{s}$ are unaffected by the supply structure.

Similarly, the decomposition of expected empirical returns in (19),

$$\mathbb{E}_t\Big[\tilde{\boldsymbol{R}}_{t+1} \mid s_t\Big] = \underbrace{\gamma\, \boldsymbol{\Sigma}_t\, \boldsymbol{M}}_{\text{RE premium}} + \underbrace{\gamma\big(\hat{\boldsymbol{\Sigma}}_t(s_t) - \boldsymbol{\Sigma}_t\big)\boldsymbol{M} + \big(\boldsymbol{\mu}_t - \hat{\boldsymbol{\mu}}_t(s_t)\big)}_{\text{ambiguity premium}},$$

remains valid. With the multiplicative ambiguity specification $\hat{\boldsymbol{\beta}}(s_t) = a(s_t)\, \boldsymbol{\beta}$ and $\overline{\beta} = 1$, the ambiguity premium for stock i is

$$\gamma\, \tau^{-1}\big(a(s_t)^2 - 1\big)\, \beta_i \; + \; \big(1 - a(s_t)\big)\, \rho\, s_t\, \beta_i, \tag{31}$$

which is proportional to β_i and independent of the supply vector $\boldsymbol{M}$. The decomposition into premium bias and expectation bias in Corollary 1 therefore continues to hold under heterogeneous supply, with the only modification being that the rational-expectations premium becomes $\gamma \boldsymbol{\Sigma}_t \boldsymbol{M}$, whose i-th entry is $\gamma(\tau^{-1}\beta_i + \tau_\epsilon^{-1} m_i)$ rather than $\gamma(\tau^{-1}\beta_i + c_M)$.

Effect on market betas and the empirical SML. Under equal supply, the covariance of stock i with the market is

$$[\boldsymbol{\Sigma}_t \boldsymbol{M}]_i = \tau^{-1}\beta_i + \frac{\tau_\epsilon^{-1}}{N},$$

so that the idiosyncratic-risk component $c_M = \tau_\epsilon^{-1}/N$ is common to all stocks. Market beta is therefore an affine function of factor beta, as stated in (22). This affine structure is the key property that allows the ambiguity premium to be absorbed by an intercept-and-slope decomposition when projected onto the CAPM basis $\{\mathbf{1}, \boldsymbol{\beta}_{M,t}\}$.

Under heterogeneous supply, the covariance becomes

$$[\boldsymbol{\Sigma}_t \boldsymbol{M}]_i = \tau^{-1}\beta_i + \tau_\epsilon^{-1}\, m_i,$$

and market beta is

$$\beta_{i,M,t} = \frac{\tau^{-1}\beta_i + \tau_\epsilon^{-1}\, m_i}{\sigma^2_{M,t}}, \qquad \sigma^2_{M,t} = \tau^{-1} + \tau_\epsilon^{-1}\sum_{j=1}^{N} m_j^2. \tag{32}$$

Since m_i varies across stocks, market beta is no longer a function of β_i alone. Inverting (32) gives

$$\beta_i = \tau\, \sigma^2_{M,t}\, \beta_{i,M,t} - \tau\, \tau_\epsilon^{-1}\, m_i,$$

so the ambiguity premium for stock i in (31) can be rewritten as

$$f(s_t)\, \beta_i = f(s_t)\, \tau\, \sigma^2_{M,t}\, \beta_{i,M,t} \;-\; f(s_t)\, \tau\, \tau_\epsilon^{-1}\, m_i,$$

where $f(s_t)$ collects the state-dependent scalars from (31). The first term is proportional to $\beta_{i,M,t}$ and is absorbed by the slope of the empirical SML. The second term, $-f(s_t)\, \tau\, \tau_\epsilon^{-1}\, m_i$, is stock-specific and cannot be absorbed by a constant intercept unless all m_i are equal.

Consequences for the SML intercept. Projecting expected empirical returns onto the CAPM basis $\{\mathbf{1}, \boldsymbol{\beta}_{M,t}\}$ via cross-sectional OLS yields an intercept

$$\iota(s_t) = \iota^{\text{equal}}(s_t) \;+\; R(s_t,\, \boldsymbol{M}),$$

where $\iota^{\text{equal}}(s_t)$ is the closed-form intercept from Proposition 3 and $R(s_t,\, \boldsymbol{M})$ is a residual that depends on the cross-sectional distribution of supply weights m_i relative to market betas $\beta_{i,M,t}$ and vanishes when supply is equal ($m_i = 1/N$ for all i). Two implications follow.

First, the closed-form intercept in Proposition 3 no longer obtains. The intercept now depends on the joint cross-sectional distribution of supply m_i and market beta $\beta_{i,M,t}$, not only on the model primitives $(\kappa, \gamma, \tau_s, c_M)$.

Second, the asymmetry in Corollary 2 may not hold for all supply configurations. Because the OLS intercept absorbs supply-driven terms that interact with the state-dependent scalar $f(s_t)$, the sign and magnitude of $\iota(s_t)$ depend on the cross-sectional covariance between m_i and $\beta_{i,M,t}$. For supply structures in which m_i is strongly correlated with $\beta_{i,M,t}$, the additional terms can dominate and potentially alter the asymmetry between good and bad news states.

To sum up, the equal-supply assumption is therefore not innocuous for the SML predictions. It ensures that market beta is a sufficient statistic for factor beta in the cross section, which delivers the closed-form intercept and the asymmetry result. The economic mechanism itself, namely that beta ambiguity aversion generates a state-dependent ambiguity premium proportional to factor beta, is entirely independent of the supply structure. The mapping from this mechanism to testable SML predictions, however, requires the additional structure that equal supply provides.

6.1.2 Asymmetry with additional frictions

Figure 10 illustrates how the asymmetry in the SML survives when an additional friction, such as borrowing constraints, generates an unconditionally flat SML. The friction flattens the benchmark (dotted), but beta ambiguity still pivots the conditional SMLs around it, preserving the asymmetry between good and bad news states.

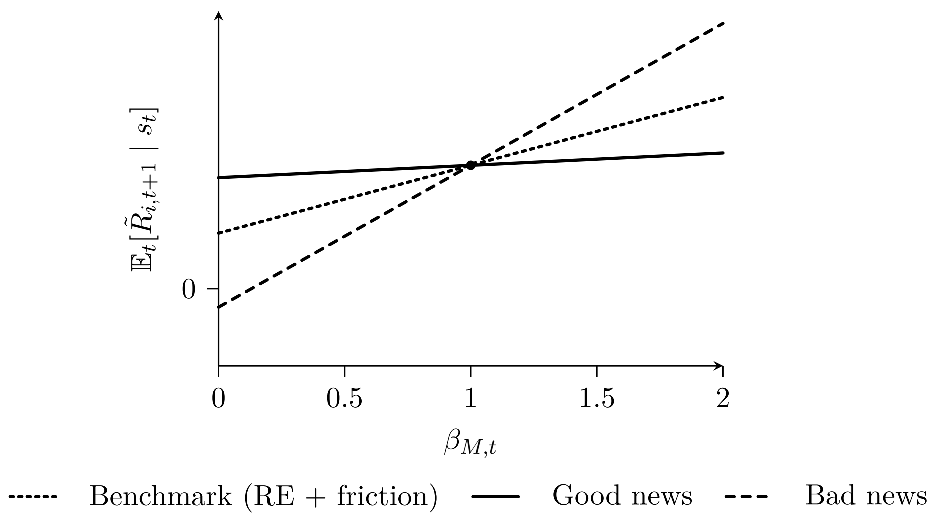

Figure 10: Model-implied SML with an additional friction. The dotted line is the benchmark SML under rational expectations combined with a friction, such as borrowing constraints, that flattens the SML unconditionally. The solid (dashed) line depicts the expected SML following good (bad) news under beta ambiguity. The asymmetry between good and bad news states is preserved.

6.2 Empirical Appendix

6.2.1 Beta uncertainty increases with beta

The model in Section 2 conjectures that beta ambiguity scales with the level of beta. This section provides empirical evidence consistent with that assumption.

I measure beta uncertainty using two complementary proxies based on market beta estimates. The first proxy is the width of the 95% confidence interval of the beta estimate, $\hat{\beta}_{i,M,\tau^-}$, estimated following the procedure outlined in Section 3, computed as $3.92 \times SE(\hat{\beta}_{i,M,\tau^-})$. The second proxy follows Hollstein et al. (2020) and measures the width of the union of 95% confidence intervals across a menu of beta estimators (historical windows, EWMA, Dimson, and Scholes–Williams corrections).

Figure 11 shows that both proxies increase from the low-beta portfolio (L) to the high-beta portfolio (H). Figure 12 shows that this relation is stable over time: the difference between the high- and low-beta portfolios is positive for nearly the entire sample.

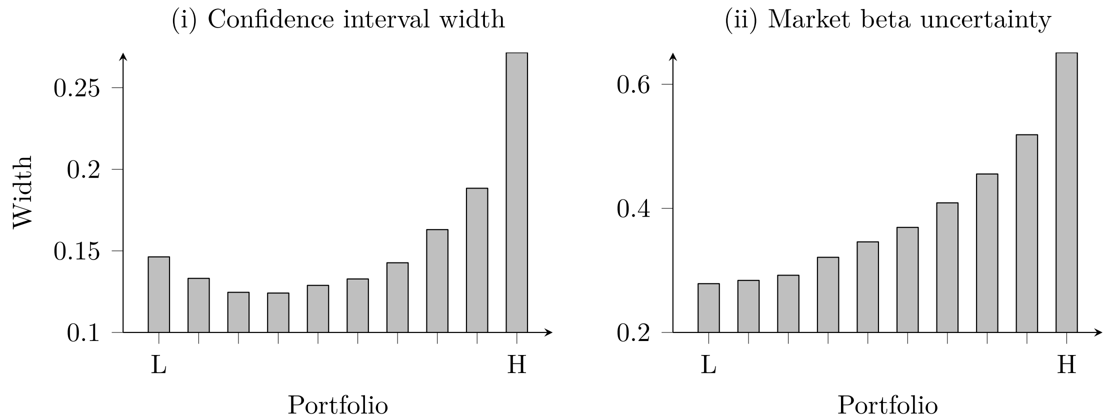

Figure 11: Market beta uncertainty across beta-sorted portfolios. Panel (i) shows the mean 95% confidence interval width for weekly market beta estimates, computed as $3.92 \times \mathrm{SE}(\hat{\beta}_{M,\tau^-})$. Panel (ii) shows mean market beta uncertainty following Hollstein et al. (2020), defined each week as the width of the union of 95% confidence intervals across a menu of estimators (historical windows, EWMA, Dimson, and Scholes–Williams corrections), then averaged over time. Portfolios are sorted from low (L) to high (H) market beta.

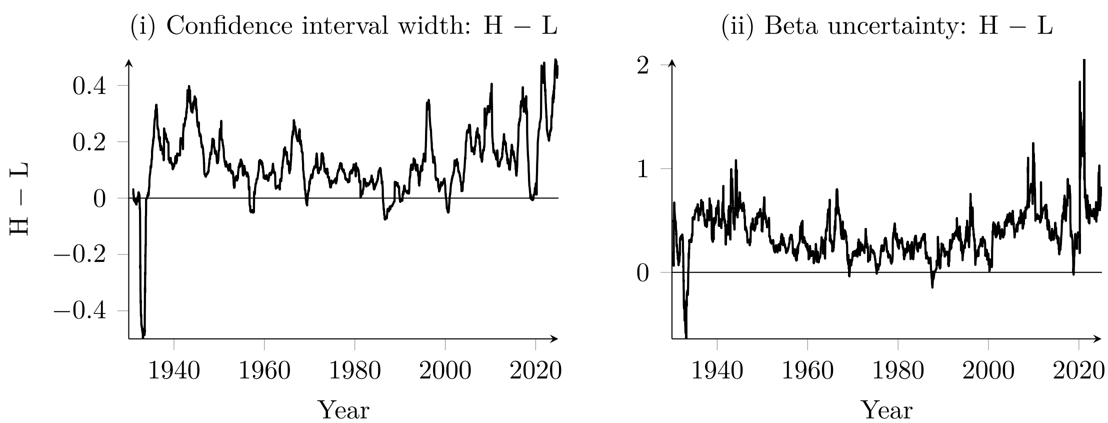

Figure 12: Time series of market beta uncertainty spread between the high- and low-beta portfolio. Each panel plots the weekly difference in market beta uncertainty between the high-beta (H) and low-beta (L) portfolios. Panel (i) uses the 95% confidence interval width $CI(\hat{\beta}_{M,\tau^-})$; panel (ii) uses the Hollstein et al. (2020) uncertainty measure $Unc(\hat{\beta}_{M,\tau^-})$. The thin line marks zero.

6.2.2 Alternative test assets

	Full	Good news		Bad news		Bad − Good	
		All	Strong	All	Strong	All	Strong
SML intercept ($\hat{\iota}_{\tau+1}$)	14.51***	20.70***	22.99***	7.32	1.18	−13.38**	−21.81**
	(4.48)	(5.11)	(3.47)	(1.59)	(0.16)	(−2.18)	(−2.22)
SML slope ($\hat{\phi}_{\tau+1}$)	3.70	−1.06	−0.46	9.22	20.87**	10.28	21.33
	(0.83)	(−0.18)	(−0.05)	(1.52)	(2.14)	(1.20)	(1.60)
Market return ($\overline{R}_{M,\tau+1}$)	15.92***	12.46***	11.05	20.48***	29.15***	8.02	18.10
	(4.36)	(2.65)	(1.58)	(3.77)	(3.18)	(1.12)	(1.57)
Observations	4,905	2,635	1,273	2,269	1,184	4,904	2,457

Table 6: SML intercept, slope, and market return across news states (FF25 + Beta-sorted + Industry). This table reports weekly means of the SML intercept ($\hat{\iota}_{\tau+1}$), the SML slope ($\hat{\phi}_{\tau+1}$), and the market excess return ($\overline{R}_{M,\tau+1}$), all measured in basis points. Means are reported for the full sample and separately for good and bad news weeks. A week is classified as a good (bad) news week if the lagged market excess return exceeds (falls below) a threshold based on the time-series distribution of weekly market excess returns. Under the all-signals classification, the threshold is the long-run average weekly market excess return. Under the strong-signals classification, the threshold is shifted by one-half of the long-run standard deviation of weekly market excess returns, so that only weeks with particularly positive or negative prior market performance are classified as good or bad news weeks, respectively. The Bad − Good column reports the difference in conditional means. The associated t-statistic is computed from a two-sample test that combines separate Newey–West long-run variance estimates for each group, and the observation count reflects the combined number of good and bad news weeks. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	All signals		Strong signals	
	(1)	(2)	(3)	(4)
Constant	20.700***	20.640***	18.753***	17.174***
	(5.02)	(5.10)	(5.32)	(4.98)
$\mathbf{1}\{\text{bad news}_\tau\}$	−13.378**	−22.348***	−17.572**	−28.591***
	(−2.41)	(−4.18)	(−2.35)	(−4.09)
$R_{M,\tau+1}$		0.128***		0.129***
		(4.68)		(4.73)
$\text{SMB}_{\tau+1}$		−0.192		−0.198
		(−1.47)		(−1.52)
$\text{HML}_{\tau+1}$		0.008		0.005
		(0.11)		(0.08)
$\text{Mom}_{\tau+1}$		0.182***		0.182***
		(4.38)		(4.38)
$\text{Illiq}_{\tau+1}$		0.061		0.065
		(0.50)		(0.54)
N	4,904	4,904	4,904	4,904
R^2	0.001	0.051	0.001	0.052

Table 7: News-dependent SML intercept with factor controls: expanded test assets. This table reports estimates from time-series regressions of the weekly SML intercept $\hat{\iota}_{\tau+1}$ on a bad-news indicator $\mathbf{1}\{\text{bad news}_\tau\}$ and factor controls, using an expanded cross-section of test assets consisting of the Fama–French 25 portfolios, ten beta-sorted portfolios, and ten industry portfolios. Specifications follow Table 3. The sample runs from 1931 to 2024. t-statistics based on Newey–West standard errors with nine lags are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels.

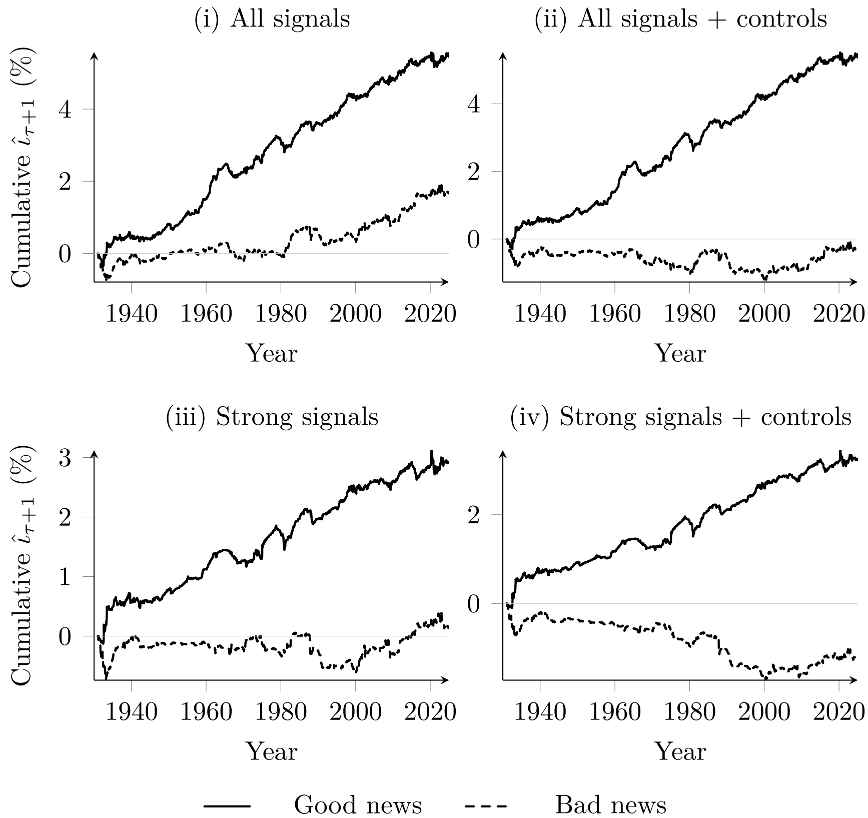

Figure 13: Cumulative SML intercepts by news state and signal strength: expanded test assets. This figure plots cumulative sums of the weekly SML intercept from 1931 through 2024, computed separately for good and bad news weeks, using an expanded cross-section of test assets consisting of the Fama–French 25 portfolios, ten beta-sorted portfolios, and ten industry portfolios. Panels (i) and (ii) use the all-signals classification; panels (iii) and (iv) use the strong-signals classification. Panels (i) and (iii) show raw cumulative intercepts. Panels (ii) and (iv) show cumulative intercepts orthogonalized with respect to the contemporaneous Carhart (1997) four-factor returns and an illiquidity factor based on Amihud (2002). All values are in percent. The solid (dashed) line corresponds to good (bad) news weeks.

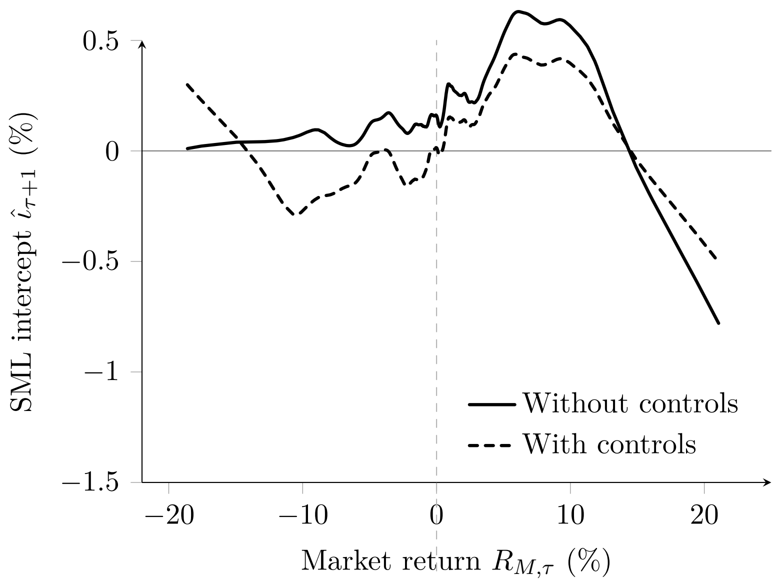

Figure 14: SML intercept as a function of lagged market return: extended test assets. This figure plots LOWESS estimates of the conditional mean of the weekly SML intercept $\hat{\iota}_{\tau+1}$ as a function of the lagged market excess return $R_{M,\tau}$, both in percent. The cross-section of test assets includes ten beta-sorted portfolios, 25 Fama–French size and book-to-market portfolios, and ten industry portfolios. Both fits use a bandwidth of $h = 0.2$, so that each local regression uses the nearest 20% of observations. The solid line shows the raw fit; the dashed line shows the control-adjusted fit, obtained by applying LOWESS to the residuals from a regression of $\hat{\iota}_{\tau+1}$ on the contemporaneous Carhart (1997) four-factor returns and an illiquidity factor based on Amihud (2002). The sample runs from 1931 to 2024.