Market Ambiguity Attitude Restores the Risk-Return Tradeoff *

Soroush Ghazi Mark Schneider Jack Strauss

December 17, 2023

Abstract

The risk-return tradeoff, while viewed as a fundamental law of finance, has been challenging to find empirically in the aggregate stock market. We consider a representative agent asset pricing model with Knightian uncertainty and demonstrate that the positive relationship between the conditional equity premium and market volatility depends on the agent's ambiguity attitude (reflecting the agent's degree of optimism or pessimism). The model predicts the conditional equity premium is increasing in market volatility, but its slope flattens as market optimism rises. We develop a methodology to extract the representative agent's ambiguity attitude from our asset pricing model. Results validate our model predictions and document the significant in-sample and out-of-sample explanatory power of ambiguity attitude in explaining the risk-return tradeoff. In our sample, market volatility is not significant in forecasting returns. However, including the market ambiguity attitude leads to a significant positive relationship between volatility and future returns. Hence, our model and results can explain why the literature has found it difficult to empirically validate the risk-return tradeoff.

JEL Classification: D8, D81, G12, G40, G41

Keywords: Ambiguity Attitude; Risk-Return Tradeoff; Equity Premium; Knightian Uncertainty

*We thank Aurelien Baillon, Turan Bali, Alain Chateauneuf, Lauren Cohen, Stephen Dimmock, Sebastian Ebert, Simon Grant, Tigran Melkonyan, Julian Thimme, Alexander Zimper, and seminar participants at the University of Connecticut, the 2021 North American Summer Meeting of the Econometric Society, the 2021 China Meeting of the Econometric Society, and the 2023 International Symposium on Forecasting for insightful comments related to this research. We thank Amit Goyal, Yehuda Izhakian, and Travis Johnson for sharing their data with us. Affiliation of authors: Soroush Ghazi, University of Alabama (sghazi@cba.ua.edu); Mark Schneider, University of Alabama (maschneider4@cba.ua.edu); Jack Strauss, University of Denver (jack.strauss@du.edu).

1. Introduction

A positive relationship between the conditional mean and the conditional variance of aggregate stock returns (the risk-return tradeoff) is one of the central empirical implications of equilibrium asset pricing theory. Rational risk-averse investors require higher compensation in equilibrium for holding stocks during riskier periods (*i.e.*, periods with higher market volatility) (Merton, 1973, 1980). Ghysels et al. (2005) comment that "This risk-return trade-off is so fundamental in financial economics that it could be described as the "first fundamental law of finance." Unfortunately, the trade-off has been hard to find in the data. Previous estimates of the relation between risk and return often have been insignificant and sometimes even negative" (p.510). Recent work continues to find the absence of a risk-return relation for the aggregate stock market (Moreira and Muir, 2017; DeMiguel et al., 2021; Barroso and Maio, 2023) and theoretical asset pricing models have struggled to offer an explanation.

In this paper, we build on the literature related to Knightian uncertainty or ambiguity in which probabilities of events are unknown and investors have varying degrees of optimism and pessimism (ambiguity attitudes) towards this uncertainty. Prior research has found that ambiguity and ambiguity aversion help explain the equity premium puzzle (*e.g.*, Chen and Epstein, 2002; Ju and Miao, 2012), the stock market non-participation puzzle (*e.g.*, Dow and da Costa Werlang, 1992; Easley and O'Hara, 2009; Dimmock et al., 2016), and the cross-section of expected stock returns (*e.g.*, Thimme and Völkert, 2015; Bali and Zhou, 2016). We provide a theoretical asset pricing model that shows market ambiguity attitude plays a critical role in explaining the risk-return tradeoff. Our paper then empirically documents that a significant, positive mean-variance relationship only exists when we include market ambiguity attitude. This significant relationship is robust to inclusion of standard equity premium predictors, and it holds across different forecast horizons, sub-samples, and in both in-sample and out-of-sample tests.

We consider a representative agent asset-pricing model in a setting of Knightian uncertainty with an agent that has a strong decision theoretic foundation. We obtain the theoretical prediction that the agent's degree of optimism or pessimism (the market ambiguity attitude) systematically moderates the risk-return tradeoff. That is, the conditional equity premium is increasing in market volatility, and the slope of this relationship flattens as market optimism increases. To test this theoretical prediction, we develop a methodology for measuring market ambiguity attitude directly from the asset pricing model.

The expected utility model, which serves as the foundation for axiomatic decision theory, game theory, information economics, and asset pricing theory, has difficulty reconciling both ambiguity-averse and ambiguity-seeking behavior (Baillon and Bleichrodt, 2015; Kocher et al., 2018). This limitation is shared by the standard ambiguity models applied to market settings which focus only on ambiguity aversion (Gilboa and Schmeidler, 1989; Hansen and Sargent, 2001; Klibanoff et al., 2005). We consider a representative agent characterized by the NEO-EU (non-extreme outcome expected utility) model of choice under ambiguity as in Chateauneuf et al. (2007) and Zimper (2012), which permits a full spectrum of ambiguity attitudes ranging from purely pessimistic to purely optimistic.[1]

We introduce an interpretation of the NEO-EU pricing formula that provides an intuitive explanation for systematic deviations from market efficiency. The equilibrium price can be expressed as a weighted average of an information component (the asset's discounted expected value) and a noise component (depending on market ambiguity attitude). In the absence of ambiguity, prices fully reflect the information component. As ambiguity rises, prices increasingly reflect the market's ambiguity attitude. Hence, when ambiguity is high, prices may deviate from the standard implication of market efficiency (that prices fully reflect available information) as the price partially reflects information and partially reflects the market's ambiguity attitude.

From a broader perspective, the NEO-EU representative agent model provides a means of unifying different sources of time-varying expected returns that have traditionally been studied in isolation. In particular, the simple formula for the equity premium derived from a NEO-EU CAPM model includes a role for market optimism, Knightian uncertainty, positive skewness, and disaster risk, thereby unifying four strands of the asset pricing literature. The formula for the equity premium can be decomposed into a speculation premium, which is increasing in market optimism, in market positive skewness, and in market uncertainty, and an uncertainty premium, which is increasing in market ambiguity aversion, in market disaster risk, and in market uncertainty.

Chateauneuf et al. (2007) observe, "On an aggregate level, business cycles and stock market

[1] The NEO-EU model satisfies the axioms of both the α-maxmin multiple priors model (Gilboa and Schmeidler, 1989; Ghirardato et al., 2004), and Choquet expected utility theory (Schmeidler, 1989), two of the primary frameworks for modeling decisions under ambiguity in which objective probabilities of events are unknown.

fluctuations have been attributed to 'irrational' optimism and pessimism. Economic theory, however, finds it difficult to see in such moods a major factor determining economic behavior." By demonstrating that market ambiguity attitude explains time variation in the risk-return tradeoff and that it predicts market crashes and NBER recessions, our study indicates that market ambiguity attitude provides an economic rationale into business cycles and stock market fluctuations.

This paper makes three contributions. First, we contribute to the literature on empirical applications of ambiguity models to market settings. We add to the ambiguity literature empirically by (i) introducing a methodology for measuring the aggregate market ambiguity attitude (i.e., the ambiguity attitude of the representative agent) and aggregate market ambiguity; (ii) demonstrating that market ambiguity attitude generates time variation in the risk-return tradeoff; (iii) showing a strong empirical link between market ambiguity attitude and the skewness of the market risk-neutral distribution; (iv) documenting that market ambiguity attitude predicts stock market crashes, consistent with the theoretical implication that a high level of market ambiguity attitude reflects an over-valued market relative to the expected utility benchmark; (v) showing empirically that market ambiguity attitude predicts recessions, and thereby provides a new empirical link between the stock market and the real economy; (vi) demonstrating that the resulting expected return forecasts generated by the interaction between market volatility and market ambiguity attitude outperform the passive strategy that buys and holds the market portfolio and earns significant factor-adjusted returns relative to leading factor models in an out-of-sample investment application; and (vii) revealing that market ambiguity attitude is not far from the ambiguity attitude measured in lab experiments (Dimmock et al., 2015; Baillon et al., 2018), but that market ambiguity is smaller than ambiguity measured in experiments, consistent with the market representative agent being closer to the expected utility benchmark than individual decision makers.

Second, our paper contributes to the theoretical literature on ambiguity models in market settings by (i) introducing a representative agent asset pricing model with ambiguity, building on Chateauneuf et al. (2007), and Zimper (2012); (ii) deriving a theoretical relationship between a behavioral measure of ambiguity and a market-based measure of economic uncertainty (the variance risk premium); (iii) deriving a theoretical relationship between market ambiguity attitude and the skewness of the risk-neutral distribution; and most importantly, (iv) demonstrating theoretically that ambiguity attitude flattens the slope of the risk-return tradeoff. We obtain this prediction from a simple asset pricing formula that provides the theoretical foundation for our empirical analysis.

Our third contribution is to the equity premium literature and the risk-return tradeoff. The presence of this tradeoff for the aggregate stock market (whether market volatility positively predicts excess returns) has been investigated and debated in many empirical studies. However, the empirical evidence has been mixed. French et al. (1987), Baillie and DeGennaro (1990), Campbell and Hentschel (1992), Ghysels et al. (2005), Lundblad (2007), Guo and Whitelaw (2006), Brandt and Wang (2007), and Pástor et al. (2008) find a positive risk-return tradeoff. However, Campbell (1987), Nelson (1991), Whitelaw (1994), Brandt and Kang (2004), and Lettau and Ludvigson (2010) find a negative risk-return relation. As noted by Yu and Yuan (2011), "numerous studies over the past three decades find rather mixed empirical evidence of such a relation" (p.367). Yu and Yuan (2011) reconcile the evidence by finding that the slope of the mean-variance relation is flat (steep) at times of high (low) sentiment. However, we find that including sentiment does not produce a significant positive coefficient on market volatility over our more recent sample period. Instead, we show that including market ambiguity attitude produces a positive risk-return tradeoff, both theoretically and empirically, using both in-sample and out-of-sample tests.[2]

A preview of our results shows that across our sample period, 1990 – 2022, market volatility fails to positively predict the equity premium. However, when including the interaction between market ambiguity attitude and market volatility, the coefficient on market volatility is positive and significant while the interaction term is negative and significant, as predicted by the theory in Section 2. In out-of-sample tests, the out-of-sample R-squared (R^2_{OS}) at the one-month forecast horizon increases from -0.16% to 4.06% when the interaction term is added to the regression.

Results then differentiate our method from the Baker and Wurgler sentiment index. We find that ambiguity attitude and market volatility subsumes the BW predictive power for the risk-return trade-off. Out-of-sample results further document that our ambiguity attitude series has significant predictive power in forecasting returns. A real-time investment strategy based on our model leads to a near doubling of the Sharpe ratio; e.g., it increases from 0.48 for the historical average to

[2] More recently, Zhang and Zhou (2023) use the economic policy uncertainty (EPU) index of Baker et al. (2016) as a proxy for un-spanned uncertainties and finds that it helps explain the breakdown of the risk-return tradeoff. In unreported regressions, we find that our results are robust to including the lagged EPU index and that both market volatility and the interaction between market volatility and market ambiguity attitude remain significant. Further recent work by Liu et al. (2021) and Yang (2022) finds respectively that different degrees of extrapolative weighting or infrequent large volatility shocks can explain the risk-return tradeoff.

0.80 and certainty equivalence rises from 2.0 to 8.1. Ambiguity attitudes also generate a significant annualized Fama-French six-factor alpha of 6.5%. Goyal Welch graphs further illustrate that the market ambiguity attitude consistently outperforms the historical average benchmark over the out-of-sample period. Lastly, we investigate the sources of the predictive power of ambiguity attitudes. We document that a high ambiguity attitude predicts market crashes and NBER recessions, highlighting a new link between returns, bear markets and the real economy.

The relationship between the risk-return tradeoff and market ambiguity attitude is not subsumed by standard equity premium predictors. In a kitchen sink regression for predicting the log equity premium with lagged market volatility and 11 prominent lagged equity premium predictors in the recent literature from Cederburg et al. (2023), the adjusted R^2 is 4.1%. Adding the interaction term between lagged market volatility and lagged market ambiguity attitude to that regression roughly doubles the adjusted R^2 to 8.3%. The coefficient on market volatility changes from being negative and insignificant in the former regression to positive and significant in the latter, while the interaction term is negative and significant. In a regression with lagged market volatility and 25 lagged equity premium predictors including the 14 widely used monthly equity premium predictors from Welch and Goyal (2008) and 11 equity premium predictors from Cederburg et al. (2023), the adjusted R^2 is 7.9%. Adding the interaction term to that regression increases the adjusted R^2 to 11.1%, and the coefficients on market volatility and the interaction term remain significant. These results show that market ambiguity attitude adds substantial gains in predictive power beyond that accounted for by the standard equity premium predictors.

Market ambiguity attitude also improves the stability of the estimated coefficients for market volatility. For instance, at the three-month forecast horizon, the estimated volatility coefficient changes sign from -0.18 to 0.53 over the two halves of the sample period. In contrast, when the interaction between market ambiguity attitude and market volatility is included, the coefficient estimates on market volatility are 1.31 and 1.34 in the two halves of the sample period and they are not significantly different.[3]

The $\alpha-$maxmin multiple priors framework from Ghirardato et al. (2004), of which the NEO-EU

[3]Goyal et al. (2021) provide the following criterion as an important requirement that successful equity premium predictors should satisfy: "The model should be reasonably stable, i.e., a variable should not have statistically significantly different IS coefficients and/or a sign change in predicting the equity premium in our sample's first half and second half—for us, at least at the 5% level. If this fails, there is little reason to proceed." (IS denotes 'in-sample'.)

model of Chateauneuf et al. (2007) is a notable special case, is a prominent axiomatic framework from decision theory in which aversion and affinity to ambiguity coexist. The $\alpha-$maxmin multiple priors framework and the NEO-EU model have been studied in laboratory experiments at the level of individual behavior (Baillon and Bleichrodt, 2015; Baillon et al., 2018; Dimmock et al., 2015; Kocher et al., 2018; König-Kersting et al., 2023) and in market experiments (Bossaerts et al., 2010). The α-maxmin model has also been applied to asset pricing theory (Chateauneuf et al., 2007; Zimper, 2012; Anthropelos and Schneider, 2022). However, it has not yet been applied to the risk-return tradeoff, and it has not been investigated empirically using stock market data.

2. The NEO-EU CAPM

2.1. The Representative Agent

The α-maxmin multiple priors model (Ghirardato et al., 2004) is one of the primary frameworks for studying attitudes toward ambiguity in the decision theory literature. Let $\mathcal{S}$ represent a set of possible future states of nature, $\mathcal{C}$ a set of consumption levels, and $\mathcal{F}$ a set of acts where an act, $f : \mathcal{S} \to \mathcal{C}$ assigns a consumption level to each state. One state $s \in \mathcal{S}$ will be realized but that true state is presently unknown. Subsets of $\mathcal{S}$ are referred to as *events*. Let Ω denote the set of all possible events. Let $\Delta(\mathcal{S})$ denote the set of all probability distributions on $\mathcal{S}$.

It is typical to write the α-maxmin value function with α as the weight on the worst-case expected utility. However, the original NEO-EU model formulation includes α as the weight on the best-case expected utility and in the present context that formulation is more intuitive to describe how α flattens the slope of the risk-return relation.

DEFINITION 1: *An α-maxmin agent has the following value function for an act f:*

$$V(f) = \alpha \max_{\mathcal{P}\in M} E_{\mathcal{P}}\, u\big(C(s)\big) + (1-\alpha) \min_{\mathcal{P}\in M} E_{\mathcal{P}}\, u\big(C(s)\big), \tag{1}$$

where $M \subseteq \Delta(\mathcal{S})$ is a closed convex set of prior distributions that the agent deems plausible given the agent's information. In (1), α represents the agent's attitude (degree of optimism) toward ambiguity, and $E_{\mathcal{P}}\, u\big(C(s)\big)$ is the agent's expected utility with respect to prior distribution $\mathcal{P} \in M$.

In empirical applications, it is often useful to assume a parameterized set of prior distributions.

A common specification of M is the following (Chateauneuf et al., 2007; Dimmock et al., 2015):

$$M_\gamma = \{\mathcal{P} \in \Delta(\mathcal{S}) : \mathcal{P}(E) \geq (1-\gamma)\pi(E)\}, \tag{2}$$

for all $E \in \Omega$, where $\gamma \in [0,1]$. In (2), the agent has a reference prior distribution, π, and a degree of confidence in that reference prior of $1-\gamma$. As summarized in Dimmock et al. (2015), the set of priors M_γ implies the following restrictions on the probability distributions $\mathcal{P} \in M_\gamma$:

$$0 \leq (1-\gamma)\pi(E) \leq \mathcal{P}(E) \leq (1-\gamma)\pi(E) + \gamma \leq 1.$$

for all $E \in \Omega$. Dimmock et al. (2015) note that the set of priors M_γ "allows the probability $\mathcal{P}(E)$ to vary in an interval of length γ around the reference probability $\pi(E)$." In this model, γ is interpreted as the level of perceived ambiguity and the model reduces to the standard subjective expected utility model when the agent perceives no ambiguity (corresponding to $\gamma = 0$).

Chateauneuf et al. (2007) show that the α-maxmin model in (1) combined with the set of priors in (2) is equivalent to the NEO-EU representation of preferences in Equation (3) for which they provide an axiomatic foundation. A NEO-EU (non-extreme outcome expected utility) agent maximizes a weighted average of the expected utility of an uncertain prospect and the Hurwicz value of the prospect which takes a convex combination of the best and worst-case utilities.[4]

DEFINITION 2: *A NEO-EU agent has the following value function for an act f:*

$$V(f) = (1-\gamma)E_\pi u\big(C(s)\big) + \gamma\big(\alpha\, u(\overline{C}) + (1-\alpha)u(\underline{C})\big). \tag{3}$$

In (1), $V(f)$ is the valuation of act f for the NEO-EU agent, $E_\pi u\big(C(s)\big)$ is the agent's expected utility (EU) from consumption under act f with respect to her subjective probability distribution, π, while $u(\overline{C})$ and $u(\underline{C})$ are, respectively, the utility from the best-case and worst-case consumption levels across states under f. These preferences separate the agent's beliefs, ambiguity attitude, α, and perceived level of Knightian uncertainty, γ. The agent's ambiguity attitude can range from pure ambiguity aversion or pure pessimism ($\alpha = 0$) to pure ambiguity seeking or pure optimism ($\alpha = 1$). The agent's perceived level of ambiguity, γ, ranges from no ambiguity ($\gamma = 0$), in which case the

[4]We restrict our attention to acts that are simple functions as in Lemma 3.1 of Chateauneuf et al. (2007).

agent maximizes expected utility with respect to her subjective prior distribution, π, to complete uncertainty ($\gamma = 1$), where the agent places no confidence in her prior and relies on the Hurwicz criterion for robust decision making which is robust to all prior distributions over the same support. The NEO-EU model nests expected utility preferences ($\gamma = 0$), and the ϵ-contamination model of ambiguity aversion ($\gamma \in$(0,1], $\alpha = 0$), two prominent theoretical benchmarks in the literature (Dow and da Costa Werlang, 1992).

The NEO-EU model accommodates aversion toward left-tail ambiguity and a preference for speculating on right-tail ambiguity. By overweighting the extreme outcomes, Chateauneuf et al. (2007) show that the NEO-EU model explains the behavior of a consumer who purchases both lottery tickets and insurance, which has been a challenge for EU since its inception (Friedman and Savage, 1948; Ebert and Karehnke, 2021). More generally, the NEO-EU model generates a preference for ambiguity over low-likelihood events and an aversion to ambiguity over high-likelihood events, consistent with the experiments in Baillon and Bleichrodt (2015) and Kocher et al. (2018). Since the focus of our empirical strategy is to capture the low-frequency movements in the risk-neutral probability of tail events, NEO-EU is a natural choice among ambiguity models for our application. In contrast this overreaction to both positive and negative tail events is not captured by popular ambiguity models that permit only uniform ambiguity attitudes, such as the smooth model of ambiguity aversion (Klibanoff et al., 2005), the maxmin multiple priors model, (Gilboa and Schmeidler, 1989), robust control preferences (Hansen and Sargent, 2001), and the ϵ-contamination model (Dow and da Costa Werlang, 1992).

The NEO-EU model can also be viewed as characterizing an agent with prospect theory probability weighting who overweights the tails of the distribution. That is, the same functional representation of preferences can be interpreted as a model of choice under risk as a special case of rank-dependent utility (Quiggin, 1982) in which the agent knows the distribution but is systematically biased and overweights the tails relative to an unbiased expected utility agent (Dimmock et al., 2021). The rank-dependent utility formulation of the NEO-EU model can be parameterized to reflect the agent's degree of optimism and the agent's degree of insensitivity to objectively known probabilities (referred to as likelihood insensitivity). In this formulation, α reflects the agent's degree of optimism toward risk ranging from pure pessimism ($\alpha = 0$) to pure optimism ($\alpha = 1$). Wakker (2010) notes that the probability weighting function embedded in (3), is among the most promising families of weighting functions in the literature and "the interpretation of its parameters is clearer and more convincing than with other families."

2.2. Equilibrium

Motivated by Chateauneuf et al. (2007) and Zimper (2012), we consider an asset pricing model with a *NEO-EU* representative agent. Our goal is not to develop a full-fledged dynamic general equilibrium model, but rather to develop a transparent model that highlights the mechanism and underlying intuition and that serves as a guide for our empirical analysis. As in Chateauneuf et al. (2007), we present our analysis in a simple two-period model in which the economy has one risky asset representing the aggregate stock market and a risk-free bond in zero net supply. The risky asset's price in period t is P_t, and its stochastic payoff in state s in period $t+1$ is $X_{t+1}(s)$. The risk-free bond's price in period t is P_t^b, and its payoff is one unit of consumption with certainty. We consider the simplest case of the NEO-EU model in which the agent has linear utility. This specification thus focuses exclusively on the effects of ambiguity and ambiguity attitude. The agent's discounted utility in period t of a given consumption level in period $t+1$ is $\delta C_{t+1}(s)$, where $\delta \in (0,1)$ is the agent's time discount factor. To simplify notation, in our subsequent analyses, for any variable θ, we define $\theta_{t+1} \coloneqq \theta_{t+1}(s)$, and we denote the corresponding conditional expectation by $E_t\theta_{t+1} \coloneqq E_{\pi,t}\theta_{t+1}(s)$. At time t, the agent chooses its level of consumption and investment to maximize:

$$\max_{\{C_t,B_t,S_t\}} C_t + (1-\gamma_t)\delta E_t C_{t+1} + \gamma_t\delta[\alpha_t\overline{C}_{t+1} + (1-\alpha_t)\underline{C}_{t+1}], \tag{4}$$

where E_tC_{t+1} is the time t expected utility of consumption in period $t+1$, and $\overline{C}_{t+1}$ and $\underline{C}_{t+1}$ are utilities from the perceived best (optimistic) and worst (pessimistic) case consumption levels in period $t+1$. Note that the conditional expected utility, E_tC_{t+1}, and the conditional maximum and minimum consumption levels, $\overline{C}_{t+1}$ and $\underline{C}_{t+1}$ are known to the agent at time t. Moreover, γ_t and α_t represent the agent's perceived *ambiguity* and *ambiguity attitude* at time t. The budget constraints at time t and $t+1$ are $\Omega_t = C_t + P_t^bB_t + P_tS_t$, and $C_{t+1} = B_t + S_tX_{t+1} + \Omega_{t+1}$, where S_t and B_t are the agent's position in the risky and risk-free assets in time t, and Ω_t is the agent's endowment at time t.

Under the linear utility specification, buying one unit of stock at time t has a marginal (utility) cost of P_t, and its payoff X_{t+1} has a marginal utility gain of $(1-\gamma_t)\delta E_t X_{t+1} + \gamma_t\delta(\alpha_t\overline{X}_{t+1} + (1-\alpha_t)\underline{X}_{t+1})$. Thus, the equilibrium price P_t adjusts to equate the current marginal utility cost to the discounted marginal utility gains, that is:

$$P_t = (1-\gamma_t)\delta E_t[X_{t+1}] + \gamma_t\delta(\alpha_t\overline{X}_{t+1} + (1-\alpha_t)\underline{X}_{t+1}). \tag{5}$$

The price of a risky asset is a weighted average of a fundamental component, $\delta E_t[X_{t+1}]$, i.e., the asset's discounted expected value, that reflects the agent's information, and a noise component, $\delta(\alpha_t\overline{X}_{t+1} + (1-\alpha_t)\underline{X}_{t+1})$, that is a function of the agent's optimism and pessimism toward uncertainty, reflecting the agent's ambiguity preferences. The relative strength of these two components depends on the agent's perceived level of uncertainty in the market, γ_t, with the agent relying less on its information at times of high uncertainty.

As indicated in the introduction, the pricing formula provides an intuitive explanation why deviations from market efficiency are expected to occur in periods of high ambiguity. In the absence of ambiguity ($\gamma_t = 0$), the price fully reflects the information component. As ambiguity increases (i.e., for higher values of γ_t), prices increasingly reflect the market's ambiguity attitude, α_t, which is impounded in prices along with the information component. In this case, prices may deviate from the textbook notion of efficiency (that prices fully reflect available information), as they partially reflect information and partially reflect the market's ambiguity attitude.

Given that return in state s is the payoff in state s, divided by the price, we have $R_{t+1} = X_{t+1}/P_t$ and hence, we express (5) as the following Euler equation:

$$(1-\gamma_t)E_t\delta[R_{t+1}] + \delta\gamma_t\left(\alpha_t\overline{R}_{t+1} + (1-\alpha_t)\underline{R}_{t+1}\right) = 1. \tag{6}$$

Similar reasoning for the risk-free bond gives us $P_t^b = \delta$ and thus:

$$\delta R_{f,t} = 1. \tag{7}$$

2.3. The Equity Premium

Subtracting (7) from (6) and rearranging terms yields the equity premium:

$$E_t R_{t+1} - R_{f,t} = \underbrace{\frac{\left[(R_{f,t} - \overline{R}_{t+1})\right] \alpha_t \gamma_t}{(1-\gamma_t)}}_{\text{Speculation Premium}} + \underbrace{\frac{\left[(R_{f,t} - \underline{R}_{t+1})\right] (1-\alpha_t)\gamma_t}{(1-\gamma_t)}}_{\text{Uncertainty Premium}}. \tag{8}$$

Equation (8) decomposes the equity premium into two terms. We refer to the first term as a *speculation premium*, and it is negative, reflecting that investors pay to hold stocks that are more exposed to market optimism (or a market boom). The speculation premium becomes larger in magnitude with higher market optimism, α_t, market positive skewness, $\overline{R}_{t+1}$, or market uncertainty, γ_t. The second term is an *uncertainty premium* that becomes larger in magnitude with higher market ambiguity aversion, $(1-\alpha_t)$, market disaster risk (lower $\underline{R}_{t+1}$), or market uncertainty, γ_t.

Representation (8) includes a role for market optimism (α_t), Knightian uncertainty (γ_t), positive skewness ($\overline{R}_{t+1}$), and disaster risk ($\underline{R}_{t+1}$), thereby unifying these strands of the asset pricing literature. Since (8) permits deviations from rational expectations consistent with the NEO-EU model, we refer to (8) as the *NEO-EU CAPM*.

2.4. Best and Worst States

To apply the model empirically, we parameterize the best and worst-case returns perceived by the agent. To do so, let $\mu_t := E_t R_{t+1}$ and $q_t := \sigma_t(R_{t+1})$:

ASSUMPTION 1: *The agent's perceived return in state $s \in \mathcal{S}$ in period $t+1$ is $R_{t+1}(s) = \mu_t + \xi_s q_t$.*

Under Assumption 1, the perceived highest and lowest returns across states are then $\overline{R}_{t+1} = \mu_t + \overline{\xi} q_t$ and $\underline{R}_{t+1} = \mu_t - \underline{\xi} q_t$, where $\overline{\xi} := \max_{s\in\mathcal{S}} \xi_s$, and $\underline{\xi} := |\min_{s\in\mathcal{S}} \xi_s|$. Assumption 1 specifies returns to be within an interval of the asset's perceived mean return, and the size of this interval increases with market volatility, q_t. In our empirical analysis, we consider the simplest case in which the endpoints of the interval are symmetric around the mean (i.e., in which $\overline{\xi} := \xi > 0$ and $\underline{\xi} := \xi$). In this case, the agent perceives returns to be in a "confidence interval" that is a fixed number (ξ) of standard deviations above or below the mean.

The model implied equity premium in (8) together with Assumption 1 yield a simple expression for the equity premium stated in the following proposition.

PROPOSITION 1: *(Market ambiguity attitude and the risk-return tradeoff) The equity premium with a NEO-EU representative agent and linear utility who perceives the state space according to Assumption 1 is:*

$$E_t R_{t+1} - R_{f,t} = \underline{\xi}\gamma_t q_t - (\underline{\xi} + \overline{\xi})\alpha_t \gamma_t q_t. \tag{9}$$

Importantly, ambiguity attitude α_t moderates the slope of the relation between expected excess return and risk (measured as market volatility q_t). In the symmetric case where $\overline{\xi} = \underline{\xi} = \xi$:

$$E_t R_{t+1} - R_{f,t} = \xi q_t \gamma_t (1 - 2\alpha_t). \tag{10}$$

Proof. Replace $\overline{R}_{t+1} = \mu_t + \overline{\xi} q_t$ and $\underline{R}_{t+1} = \mu_t - \underline{\xi} q_t$ (obtained from Assumption 1) in Equation (8) and rearrange the terms. □

We expect an ambiguity attitude α_t less than 0.5, reflecting a bias towards ambiguity aversion rather than optimism. Empirical estimates of α_t measured in lab experiments (Dimmock et al., 2016; Baillon et al., 2018) also report $\alpha_t < 0.5$, which implies a positive relationship between risk and expected returns. Moreover, a higher level of optimism α_t means a shallower slope, i.e., a lower reward for accepting risk.

2.5. The Variance Risk Premium

The variance risk premium (VRP), the difference between the risk-neutral and physical market variance is commonly interpreted as a measure of economic uncertainty (Zhou, 2018; Bali and Zhou, 2016). Formally, the VRP is defined as $\text{VRP}_t \equiv \text{Var}_t^Q R_{t+1} - \text{Var}_t^P R_{t+1}$ (Zhou, 2018), where Var_t^Q and Var_t^P are the conditional variance under the risk-neutral and physical measures. For the NEO-EU agent with linear utility, the VRP is the following:

$$VRP_t = (1-\gamma_t) E_t (R_{t+1} - \tilde{E}_t R_{t+1})^2 + \gamma_t \left[\alpha_t (\overline{R}_{t+1} - \tilde{E}_t R_{t+1})^2 + (1-\alpha_t)(\underline{R}_{t+1} - \tilde{E}_t R_{t+1})^2\right] - q_t^2. \tag{11}$$

PROPOSITION 2: *Under Assumption 1, with $\overline{\xi} = \underline{\xi} = \xi$, the variance risk premium for a NEO-EU* *agent with linear utility can be approximated as the following:*

$$VRP_t \approx \gamma_t q_t^2(\xi^2 - 1). \tag{12}$$

Proof. See the Appendix. □

In Proposition 2, the VRP approximation is independent of the market ambiguity attitude, α_t. The propositions thus provide a partial separation between ambiguity and ambiguity attitude. Further, Proposition 2 provides a theoretical link between a market-based measure of Knightian uncertainty (VRP) and a behavioral measure of Knightian uncertainty, γ_t. Moreover, γ_t (scaled by the constant $\xi^2 - 1$) serves as a wedge between VRP_t and q_t^2.

2.6. Taking the Model to Data

In this subsection, we use the model expressions from the previous subsections to extract a measure of ambiguity attitude α_t. To that end, we first estimate q_t from a GARCH model and use it to construct a measure of γ_t. Then, α_t can be estimated using a Markov switching model. The details are provided below.

First, we measure q_t via a simple GARCH(1,1) model for the log returns, where pd_t is the price-dividend ratio for the S&P 500 index and is used to construct a measure of the expected equity premium:[5]

$$\begin{aligned} \log(R_{t+1}) &= \theta_0 + \theta_1 pd_t + u_{t+1}, \\ u_{t+1} &= q_{t+1} z_{t+1}, \quad \text{with} \quad z_{t+1} \sim \mathcal{N}(0,1) \\ q_{t+1}^2 &= \omega_0 + \lambda_1 u_t^2 + \beta_1 q_t^2. \end{aligned}$$

In this specification, the log market return is linear in the price-dividend ratio (pd_t), and the error is assumed to follow a GARCH(1,1) process. Lemma 1 in the Appendix shows how the log return is approximately linear in the price-dividend ratio when we replace the payoff with dividends. The estimation not only provides us with an estimate of the conditional market volatility $\hat{q}_t$, but

[5]The dividend price ratio is also used as a proxy for the equity premium empirically in Pástor and Veronesi (2020) and is closely related to the equity premium in traditional macro-finance models.

also yields a conditional market equity premium based on the conditional expected return and the market risk free rate $E_t R_{t+1} - R_{f,t} \approx \exp(\hat{\theta}_0 + \hat{\theta}_1 pd_t) - R_{f,t}$.[6] For notational convenience, we denote the conditional equity premium by $\text{EP}_t \equiv E_t R_{t+1} - R_{f,t}$.

Having established a theoretical link between market ambiguity attitude and the risk-return tradeoff, our next objective is to provide an estimate of the time series of α_t using the structural equations from the model under Assumption 1 and linear utility.

Second, following Zhou (2018) and Bekaert and Hoerova (2014), we use the square of the VIX index as a proxy for the risk-neutral variance, $\text{Var}_t^Q R_{t+1}$. Then using formula (12) we find γ_t to be:

$$\hat{\gamma}_t \approx \frac{1}{\xi^2 - 1}\Big(\frac{\text{VIX}_t^2}{\hat{q}_t^2} - 1\Big). \tag{13}$$

Finally, we use the relationship, $\text{EP}_t = \xi(1 - 2\alpha_t) q_t \gamma_t$ from Equation (10) to estimate α_t. In line with the intuition that α_t has persistent dynamics, we let α_t follow a Markov-switching structure with two states. Note that the relationship implies $\xi(1 - 2\alpha_t) = \dfrac{\text{EP}_t}{q_t \gamma_t}$. Thus, if α_t follows a Markov-switching model, so does the ratio $\dfrac{\text{EP}_t}{q_t \gamma_t}$. To estimate α_t, we estimate the following Markov-switching dynamic regression model

$$\frac{\hat{\text{EP}}_t}{\hat{q}_t \hat{\gamma}_t} = \mu_{m_t} + \epsilon_t, \tag{14}$$

where ϵ_t is a white noise and μ_{m_t} switches between two regimes according to a probability matrix. The estimated model gives us a predicted value of $\hat{\mu}_{m_t}$ using the information up to and including time t. We then find our estimate of α_t according to

$$\hat{\alpha}_t = \frac{1}{2}\Big(1 - \frac{\hat{\mu}_{m_t}}{\xi}\Big). \tag{15}$$

Importantly, both q_t and α_t are estimated dynamically so that *only* the information up to period t is used in the estimation of $\hat{q}_t$ and $\hat{\alpha}_t$ to avoid look-ahead bias. To elaborate, q_t is estimated from a GARCH model that uses data for the first half of the sample period (through June, 2006). Starting with the beginning of the out-of-sample period (July, 2006), we then run the GARCH model every period with an expanding window to estimate q_t using only information up through

[6] Allowing for the second order (Jensen) term virtually makes no difference in the final estimates.

period t. Similarly, the parameters of the Markov-switching model are estimated for the first half of the sample period. For the out-of-sample period, the Markov-switching model is re-estimated every period t using information only up through period t, and we obtain a new estimate of α_t each period that uses information only up through period t.

The quantity $\frac{\text{EP}_t}{q_t \gamma_t}$ is a measure like a conditional Sharpe ratio but which includes a role for market ambiguity, γ_t. In the Markov-Switching model there are two regimes: (i) a bear market regime with relatively low prices and high expected future returns per unit of risk, and (ii) a bull market regime with relatively high prices and low expected future returns per unit of risk. Market optimism, α_t, is then increasing in the probability of the bull market regime.

As market ambiguity is directly linked to the VRP by the model, and as Knightian uncertainty or ambiguity by itself has received much attention in prior literature, we focus on the new aspect of our approach which is the time series of market ambiguity attitude. This focus also reflects the motivation of the paper which is to investigate if market ambiguity attitude moderates the risk-return tradeoff which is predicted by the theory studied here. Assuming γ_t has little or no predictive content due to its low auto-correlation, q_t and α_t contain all of the information about the conditional equity premium in Equation (10).[7]

3. Properties of Market Ambiguity Attitude

3.1. The Level of Market Ambiguity and Ambiguity Attitude

The mean value of α is reasonably close to laboratory estimates from studies on individual decision making. In a representative sample, Dimmock et al. (2015) estimates $\alpha = 0.44$. In a lab experiment with a stock market setting, Baillon et al. (2018) estimates $\alpha = 0.45$. For candidate values of ξ (truncating the return distribution at 2, 3, or 4.77 standard deviations), the estimated α ranges between 0.42 and 0.27. The market-based estimates provided by our study indicate that market ambiguity attitude is not far from the ambiguity attitudes of individual agents. Our market γ is considerably smaller than laboratory estimates. Dimmock et al. (2015) estimates γ to be 0.40. Baillon et al. (2018) estimates γ to be 0.52. The aggregate market is less biased relative to

[7] We estimate that α has a one-month auto-correlation of 0.97, q has a one-month auto-correlation of 0.94, and γ has a one-month auto-correlation of 0.46.

the expected utility benchmark (revealed by a lower γ) and hence perceives less uncertainty than individual agents. This information is summarized in Table 1.

Table 1. Estimates of Ambiguity Attitude (α) and Ambiguity Perception (γ)

	Laboratory Estimates		Market Estimates		
	BBKHL	DKMP	$\xi = 2$	$\xi = 3$	$\xi = 4.77$
α	0.45	0.44	0.42	0.36	0.27
γ	0.52	0.40	0.31	0.13	0.05

Notes: The two leftmost columns display the estimated ambiguity attitude (degree of optimism), α, and ambiguity perception, γ, in the NEO-EU framework estimated from the individual choice experiments of Baillon et al. (2018) (BBKHL) and Dimmock et al. (2015) (DKMP). The remaining three columns display the mean α and γ estimated from market data for different values of the truncation parameter, ξ.

We use the specification $\xi = 4.77$, since it implies a best-case return and worst-case return that are roughly consistent with common definitions of a bull market and a bear market in the media and on Wall Street (Kurov, 2010). The specification does not rely on information in the out-of-sample period. Over the training sample period, the monthly mean market volatility, q, is 4.05% and the monthly mean expected market return from the GARCH model is 0.74%. Computing the maximum return under Assumption 1 with these values yields $\overline{R} = 0.0074 + 4.77(0.0405) \approx 0.20$. This is consistent with the threshold for a bull market (a return of 20% from a market's recent low). The corresponding worst-case return is approximately -0.19, and similar to the threshold for a bear market (a return of -20% from a market's recent high). Since ξ is constant, ξ affects the level but not the time variation of α. Consequently, our results are robust to different values of ξ. Of particular note, our main metrics for forecast evaluation (the in-sample and out-of-sample R^2 and the difference in cumulative sum of squared forecast errors presented in Section 4) are identical for all $\xi \in [2, 4.77]$.[8]

3.2. Time Variation in Market Ambiguity Attitude and the Risk-Neutral Distribution

This section motivates and illustrates what α captures empirically. Intuitively, a smaller α_t, consistent with a more pessimistic NEO-EU representative agent, leads to greater negative skewness

[8] A positive feature of a specification with $\xi \leq 4.77$ is that $\alpha \in (0, 1)$ for all periods in our sample spanning more than 30 years of monthly data. Under specifications with $\xi > 5$, the estimated α becomes negative in some periods. Truncating α at zero in those periods will slightly affect the R^2.

of the risk-neutral probability density. That is, one might expect the skewness of the risk-neutral density to increase in α_t. To investigate this relationship, we first provide a formal proposition in which market risk-neutral skewness (RNS) is increasing in α_t. The proposition provides an upper bound on γ_t necessary for this increasing relationship between RNS and α_t. Next, we test if α_t correlates positively and significantly with RNS. To push the proposition further, we test if the correlation between α_t and RNS is higher when the difference between the upper bound provided by the proposition and γ_t is high and so the condition necessary for the increasing relationship between RNS and α_t is more easily met. We also plot α_t and RNS to visualize the relationship, and we conduct Granger causality tests to infer potential dependencies between α_t and RNS.[9]

We document a positive and significant correlation of 0.40 between α_t and RNS across the OOS period in Table 12 in the Online Appendix. Remarkably, the correlation more than doubles when comparing α_t and the product of RNS and an indicator variable that equals one when the difference between the upper bound from Proposition 3 and γ_t is above its median value. When graphing these relationships, shown in Figure 1, it is apparent that α_t looks like a smooth version of RNS. These observations indicate that market ambiguity attitude, α_t, contains information about low-frequency movements (and hence the more predictable variation) in RNS.

Building on Proposition 1, the risk-neutral distribution has most of its mass in the middle (from the physical distribution), and on the two extreme events (the best and worst case returns). The NEO-EU model implies that the weight of the risk-neutral measure in the middle of the distribution is $1-\gamma_t$, on the best case return is $\gamma_t\alpha_t$, and on the worst case return is $(1-\gamma_t)\alpha_t$. To focus on the effect of α_t on the skewness of the risk-neutral distribution, we find that it is (empirically) enough to approximate the risk-neutral distribution with three points: the worst case, the middle, and the best case returns, giving us the following proposition.

PROPOSITION 3: *Approximating the risk-neutral distribution implied by the NEO-EU representative agent as a three-point distribution, the skewness of the risk-neutral distribution is increasing in* α_t *if* $\gamma_t < \frac{1}{1+8\alpha_t(1-\alpha_t)}$.

[9] Market risk-neutral skewness is measured from the SKEW index of the CBOE. RNS $= E[(\frac{R-\mu}{\sigma})^3]$, where R is the 30-day log-return on the S&P 500, μ and σ are respectively the mean and standard deviation of R, $x := (\frac{R-\mu}{\sigma})^3$ and RNS $=$ E$[x]$. RNS is obtained from the SKEW index by the relation RNS $= \frac{100-SKEW}{10}$. (See `https://cdn.cboe.com/resources/indices/documents/SKEWwhitepaperjan2011.pdf`). The series is converted from the daily frequency to the monthly frequency using the last observation in each month as the RNS value for that month.

Proof. See the Appendix. □

Supporting Proposition 3 empirically, we find that α_t is positively and significantly correlated with market risk-neutral skewness (RNS). Motivated by the proposition, we test if the correlation between α_t and market risk-neutral skewness is higher when $\frac{1}{1+8\alpha_t(1-\alpha_t)} - \gamma_t$ is high so that the inequality relating difference between the upper bound on γ_t necessary for the proposition to hold and γ_t, has more slack. To do so, we construct an indicator variable that equals one if $\frac{1}{1+8\alpha_t(1-\alpha_t)} - \gamma_t$ is above its median value over the OOS period (0.386) and zero otherwise, and we compute the product of this indicator variable and the market risk-neutral skewness. We find that the correlation between α_t and the product of RNS and the indicator variable is 0.85 as shown in Table 12 in the Appendix. Both relationships are plotted in Figure 1. Clearly, when the value $\frac{1}{1+8\alpha_t(1-\alpha_t)} - \gamma_t$ is farther from zero, the correlation between RNS and α_t is also higher.

From Figure 1, it appears that α_t looks like a smoothed version of RNS. To further investigate the relationship between α_t and RNS, we conduct Granger causality tests using the optimal lag lengths according to the Bayesian Information Criterion (one period) and according to the Akaike Information Criterion (two periods). As we show in Table 13 in the Appendix, in both tests, α_t significantly Granger causes RNS at the 5% significance level, whereas RNS does not Granger cause α_t (with p-values above 0.80 in both cases).

Recall that in theory, variation in α measures variation in the market risk-neutral probabilities of tail events (for both the right and left tail). We have seen in this section that empirically, α reflects low-frequency movements in the SKEW index of the CBOE. The SKEW index was introduced by the CBOE as an "indicator that measures perceived tail risk" and the CBOE notes that it is intended to be complementary to the VIX index (CBOE, 2011). The fact that α, whose construction does not depend on SKEW, recovers the low frequency movements in SKEW consistent with theory, shows the potential relevance of the NEO-EU model for empirical finance applications.

Figure 1. NEO-EU Optimism (α) and Market Risk-Neutral Skewness (RNS)

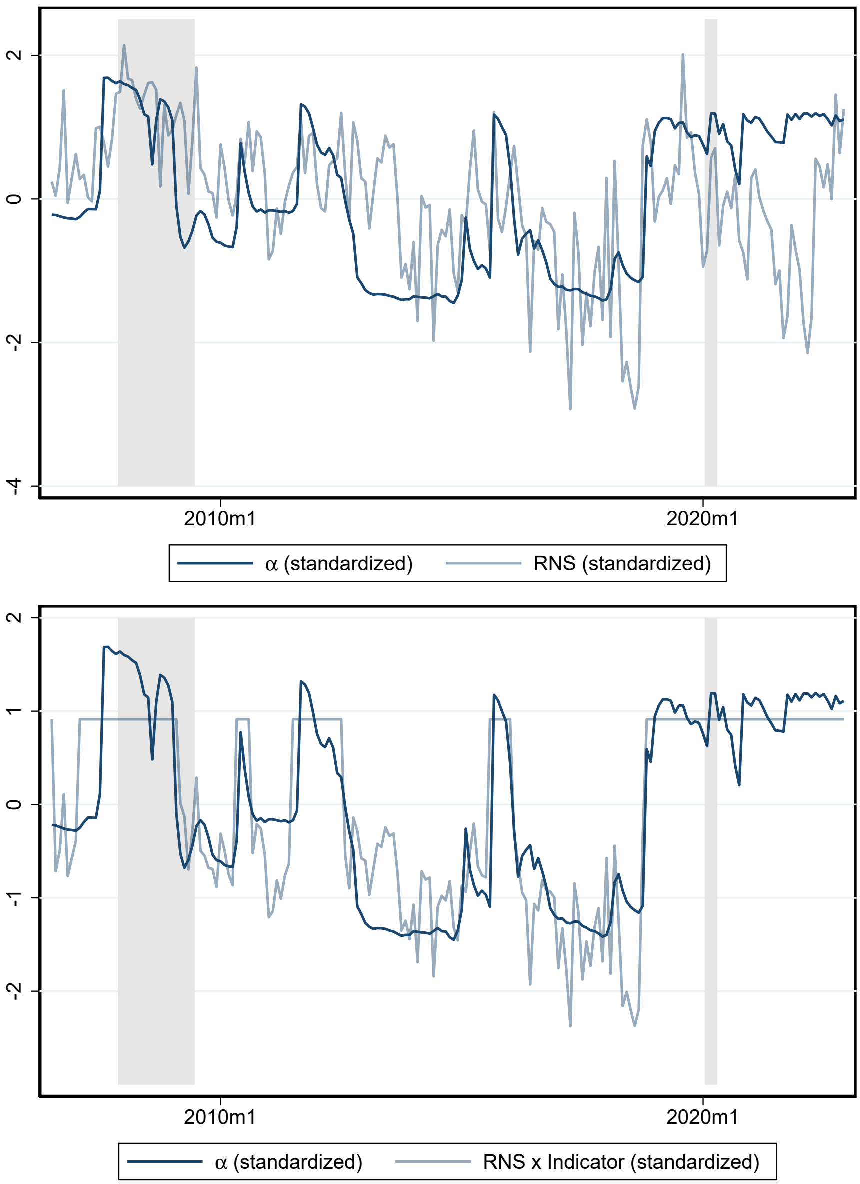

Notes: The top figure displays the time series of the market risk-neutral skewness (RNS) from the Chicago Board of Options Exchange, and the recursively updated market ambiguity attitude series (α). The bottom figure displays α and the product of market risk-neutral skewness and an indicator variable that equals one if the difference between the upper bound on γ in Proposition 3 and γ is above its median value (0.386) over the out-of-sample period (RNS $\times$ Indicator). Both figures span the out-of-sample period for α (2006:07 through 2022:12). Over this period, the correlation between the two series in the top figure is 0.40, and the correlation between the two series in the bottom figure is 0.85. For the purpose of comparison, all series have been standardized to have a mean of zero and a standard deviation of one over this period. Shaded areas show the NBER recession periods.

4. The Risk-Return Tradeoff

Under Proposition 1, there is a positive relationship between market volatility, q_t, and the expected equity premium, but the slope of this relationship flattens as market optimism, α_t, increases.[10] Our first analysis is motivated by three basic questions: First, does market ambiguity attitude moderate the risk-return tradeoff as predicted by Proposition 1? Second, if so, is the predictive power of market volatility, q_t, and the interaction between market ambiguity attitude and market volatility, $\alpha_t q_t$ subsumed by standard equity premium predictors? Third, what is the incremental increase in predictive power generated by including $\alpha_t q_t$ in the predictive regressions, beyond that delivered by the standard equity premium predictors? To probe these questions, we consider 25 equity premium predictors consisting of the 14 predictors in Welch and Goyal (2008) available at the monthly frequency and the 11 newer predictors used by Cederburg et al. (2023) for which data is available beginning in January, 1990.[11] For q and α, our data begins in 1990 (the first year available for the VIX which is needed for the construction of α) as it is a common measure of the risk-neutral variance. For details about the construction of the 25 predictors, see Welch and Goyal (2008), Cederburg et al. (2023).

Table 2 reports regressions following Equation (16). The full specification is a regression of the (realized) log equity premium, denoted R_t^e, in period t against market volatility (q_{t-h}), the product of market ambiguity attitude and market volatility ($\alpha_{t-h}q_{t-h}$), and a set of k alternative equity premium predictors. All predictors are lagged h periods. The table considers the simplest case (with lagged market volatility as the only regressor), as well as the case that includes $\alpha_{t-h}q_{t-h}$ and additional specifications with various sets of control variables. We present the results using monthly data for both the one-month forecast horizon, corresponding to $h = 1$ and the three-month forecast horizon corresponding to $h = 3$ in the regressions. We refer to the former as the monthly forecast horizon, and the latter as the quarterly forecast horizon.

[10]Our analysis is guided by the model prediction that α_t moderates the risk-return tradeoff (i.e., the relevant variables are q_t and $\alpha_t q_t$). For this reason and due to multicollinearity, we omit α_t from these regressions.

[11]This criterion enables us to include predictors that span the period 1990 - 2021, and omits only the left-tail jump variation (LJV) predictor from Cederburg et al. (2023) for which available data begins in 1996. Including that variable does not affect the results.

$$R_t^e = \beta_0 + \beta_1 q_{t-h} + \beta_2 \alpha_{t-h} q_{t-h} + \Sigma_{i=3}^k \beta_i x_{i,t-h} + \epsilon_t. \quad (16)$$

In Equation (16), k denotes the total number of predictor variables in the regression. Odd-numbered regressions in Table 2 do not include $\alpha_{t-h}q_{t-h}$, while this term is included in even-numbered regressions. Regressions summarized in columns (1), (2), (9), and (10) show our baseline results without controls. All data is updated from the original studies to span from 1990:01 - 2021:12. In columns (1) through (8) all predictor variables are lagged by one month (corresponding to the monthly forecast horizon). In columns (9) through (16) all predictor variables are lagged by three months (corresponding to the quarterly forecast horizon).

In Table 2, for ease of interpreting the estimated coefficients, the predictors q and αq are divided by their full-sample standard deviation. Regression specification (2) in the top panel of Table 2 which includes both q and αq reveals that a one standard deviation increase in the conditional market volatility leads to an increase in the future realized equity premium of 1.35%. In contrast, a one standard deviation increase in the product of market volatility and market optimism, αq, leads to a decrease in the future equity premium of -1.36%. Both coefficients have t-statistics above three, and hence are both statistically significant and economically large. Similar results hold at the quarterly forecast horizon as shown in regression specification (2) in Panel B of the table.

Table 2 answers our three questions. Regarding the first question, the table shows in Column (1) of Panels A and B that the relationship between R_t^e and lagged market volatility is not significant at the monthly or quarterly horizon. These regressions confirm the findings of Campbell (1987), Nelson (1991), Whitelaw (1994), Brandt and Kang (2004), Lettau and Ludvigson (2010), and Barroso and Maio (2023) that the risk-return relationship is neither strong nor robust in the data. The regressisons reveal that the absence of a clear risk-return tradeoff for the aggregate stock market continues to be a puzzle even when using more recent data than prior studies. In contrast, adding the interaction between market ambiguity attitude and market volatility to the regression yields a significant positive relation between R_t^e and lagged market volatility, q, and a significant negative relationship between R_t^e and lagged αq as noted in the preceding paragraph. These findings support the theoretical prediction that market ambiguity attitude moderates the risk-return tradeoff.

Table 2. Market Ambiguity Attitude and the Risk-Return Tradeoff controlling for 25 Predictors

	Monthly Forecast Horizon							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-1}	0.30	1.35***	1.49***	2.73***	-0.10	1.71***	1.10**	2.86***
	(1.18)	(5.44)	(4.07)	(4.45)	(-0.28)	(2.98)	(2.34)	(3.56)
$\alpha_{t-1}q_{t-1}$		-1.36***		-1.90***		-2.10***		-2.34***
		(-3.87)		(-2.70)		(-4.40)		(-3.17)
GW_{t-h}	NO	NO	YES	YES	NO	NO	YES	YES
CJO_{t-h}	NO	NO	NO	NO	YES	YES	YES	YES
k	1	2	15	16	12	13	26	27
adj. R^2	0.002	0.037	0.050	0.078	0.041	0.083	0.080	0.113
	Quarterly Forecast Horizon							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel B	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-3}	0.33	1.25***	1.01***	1.70***	-0.37	0.93**	0.36	1.23**
	(1.51)	(4.87)	(3.01)	(4.19)	(-1.14)	(2.36)	(0.77)	(2.32)
$\alpha_{t-3}q_{t-3}$		-1.20***		-1.05***		-1.50***		-1.15***
		(-3.48)		(-2.83)		(-4.70)		(-3.32)
GW_{t-h}	NO	NO	YES	YES	NO	NO	YES	YES
CJO_{t-h}	NO	NO	NO	NO	YES	YES	YES	YES
k	1	2	15	16	12	13	26	27
adj. R^2	0.003	0.029	0.058	0.065	0.057	0.078	0.081	0.087

Newey-West t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays regressions of the log equity premium, R_t^e, (in percent) against market volatility, q, and the set of 14 monthly equity premium predictors in Welch and Goyal (2008), the set of 11 newer equity premium predictors used by Cederburg et al. (2023) for which data are available beginning in January, 1990, and both sets of predictors. Even-numbered regressions also include αq. In the regression specifications in Panel A, all predictor variables are lagged by one month (monthly forecast horizon). In the regression specifications in Panel B all predictors are lagged three months (quarterly forecast horizon). The GW row indicates whether the 14 monthly Welch and Goyal (2008) predictors are included as controls. The CJO row indicates whether the 11 Cederburg et al. (2023) predictors available starting in January, 1990, are included as controls. k denotes the number of predictor variables in the regression including q and the control variables, and αq where applicable. For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation. The sample period spans monthly data from 1990:01 through 2021:12.

Regarding our second question, Table 2 shows that neither the Welch and Goyal (2008) predictors nor the Cederburg et al. (2023) predictors collectively subsume the risk-return tradeoff results. In all six specifications where q and αq are both included in the regressions, the coefficient on q remains positive and significant while the coefficient on αq remains negative and significant, even in the presence of 25 standard equity premium predictors as controls. In the absence of αq, the coefficient on q is significant in the presence of the Welch and Goyal (2008) predictors but insignificant using the more recent Cederburg et al. (2023) predictors.

Regarding the third question, Table 2 shows that including αq in the kitchen sink regressions substantially improves the predictive power. For instance, the adjusted R^2 for regression (5) in the table which includes the 11 recent predictors in Cederburg et al. (2023) is 4.1%. Adding αq to that regression roughly doubles the adjusted R^2 to 8.3%. This is a remarkable increase in predictive power relative to a regression that already contains 11 strong equity premium predictors from the recent literature. As a related example, the adjusted R^2 in Regression specification (7) which includes q in addition to 25 established equity premium predictors is 8%. Adding αq to that regression increases the adjusted R^2 by over 3% to 11.3%.

4.1. Market Sentiment and the Risk-Return Tradeoff

As noted in the introduction, Yu and Yuan (2011) have found that market sentiment, proxied by the Baker and Wurgler (2006) market sentiment index (BW), moderates the risk-return tradeoff. We test here if BW moderates the risk-return relationship for our sample period and whether that relationship explains the moderating effect of market ambiguity attitude. To investigate this, we run versions of regression Equation (17) which includes market volatility, q_t, the product of the BW index and market volatility ($bw_t q_t$), the BW index by itself, and the product $\alpha_t q_t$.

$$R_t^e = \beta_0 + \beta_1 q_{t-h} + \beta_2 bw_{t-h} q_{t-h} + \beta_3 bw_{t-h} + \beta_4 \alpha_{t-h} q_{t-h} + \epsilon_t. \tag{17}$$

The results are displayed in Table 3. Regression specification (1) in the table, which includes only q_t and $bw_t q_t$ as predictors, shows that a one standard deviation increase in the product of the BW index and market volatility leads to a decrease in the future log equity premium of 0.56% on average. This significant effect is reduced to an insignificant decrease of 0.13%, when the interaction between market ambiguity attitude and market volatility is included in regression specification (3). From the table, we see that BW does not subsume the predictive power of $\alpha_t q_t$. Instead, the interaction between market ambiguity attitude and market volatility subsumes the predictive power of BW for the risk-return tradeoff.

Table 3. Market Sentiment and the Risk-Return Tradeoff

	Monthly Horizon					Quarterly Horizon			
	(1)	(2)	(3)	(4)		(5)	(6)	(7)	(8)
Predictor	R_t^e	R_t^e	R_t^e	R_t^e	Predictor	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-1}	0.31	0.31	1.29***	1.32***	q_{t-3}	0.25	0.22	0.98***	0.94***
	(1.20)	(1.17)	(4.24)	(4.07)		(1.34)	(1.12)	(3.03)	(2.71)
$\alpha_{t-1}q_{t-1}$			-1.23**	-1.24**	$\alpha_{t-3}q_{t-3}$			-0.91**	-0.89**
			(-2.58)	(-2.53)				(-2.26)	(-2.13)
$bw_{t-1}q_{t-1}$	-0.56***	-0.62	-0.13	0.19	$bw_{t-3}q_{t-3}$	-0.67***	-1.39**	-0.36	-0.78
	(-3.05)	(-0.87)	(-0.46)	(0.24)		(-3.79)	(-2.08)	(-1.51)	(-1.02)
bw_{t-1}		0.07		-0.32	bw_{t-3}		0.74		(0.42)
		(0.09)		(-0.46)			(1.02)		(0.57)
R^2	0.022	0.022	0.044	0.045		0.027	0.028	0.040	0.040

Newey-West t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays regressions of the log equity premium (R_t^e) in percent, against the conditional stock market volatility (q_{t-h}) and the product of market volatility and the product of the Baker and Wurgler (2006) market sentiment index (bw_{t-h}) and market volatility in regression specifications (1) and (5). Regression specifications (2) and (6) include (bw_{t-h}). Regression specifications (1) through (4) use a forecast horizon of $h = 1$ month. Regression specifications (5) through (8) use a forecast horizon of $h = 3$ months. The sample period is 1990:01 - 2022:06 (June 2022 is the last month available for the Baker and Wurgler (2006) sentiment series). For ease of interpreting the coefficients, q_t, $\alpha_t q_t$, and $bw_t q_t$ and bw_t are divided by their unconditional standard deviation.

We also tested whether the ambiguity measure from Brenner and Izhakian (2018) or the time varying U.S. disaster probabilities from Barro and Liao (2021) subsume the relationship between market ambiguity attitude and the risk-return tradeoff. It should be noted that the measure in Brennan et al. (2004) is an index of ambiguity whereas α_t in the present paper is a measure of ambiguity attitude. In contrast to our approach, the ambiguity attitude in Brennan et al. (2004) is not time-varying. We find that the coefficient on volatility remains positive and significant and the coefficient on the interaction term $\alpha_t q_t$ remains negative and significant, but the ambiguity measure from Brenner and Izhakian (2018) is not significant at both the monthly and quarterly forecast horizons when added to the regressions. We similarly find that the coefficients on q_t and

$\alpha_t q_t$ remain significant while the coefficient on time-varying disaster probabilities is not significant when added to the regression with q_t and $\alpha_t q_t$.

4.2. Out-of-Sample Regressions

Following the classic work of Welch and Goyal (2008), it is increasingly common to test if evidence of return predictability from in-sample regressions also holds out-of-sample. Consequently, we investigate the risk-return tradeoff and the predicted moderating effect of α in out-of-sample regressions. We use three standard metrics to evaluate out-of-sample (OOS) predictability: (1) the R^2_{OS} statistic of Campbell and Thompson (2008); (2) the MSPE-adjusted statistic of Clark and West (2007); and (3) the difference in cumulative sum of squared errors between the historical average equity premium forecast and the forecast based on a set of predictor variables (Welch and Goyal, 2008). The Campbell-Thompson out-of-sample R^2 statistic is defined as follows:

$$R^2_{OS} = 1 - \frac{\sum_{t=1}^{T}(r_t - \hat{r}_t)^2}{\sum_{t=1}^{T}(r_t - \overline{r}_t)^2}. \tag{18}$$

In (18), $\hat{r}_t$ is the forecast (fitted value) from the predictive regression using information through period $t-h$, while $\overline{r}_t$ is the mean equity premium using historical values through period $t-h$. A positive R^2_{OS} statistic implies that the predictive regression has lower average mean-squared prediction error than the historical mean equity premium. Campbell and Thompson (2008) show that even an R^2_{OS} of 1% can result in substantial utility gains for a mean-variance investor.

The second metric is the mean squared prediction error (MSPE)-adjusted statistic of Clark and West (2007). This metric, henceforth CW, compares the predictive performance of two nested models accounting for the higher noise introduced in the larger (nesting) model. The null hypothesis is that the two models have equal mean-squared prediction errors. The alternative hypothesis is that the larger model has a smaller MSPE. Rejecting the null hypothesis means the predictive model outperforms the forecast based on the historical average.

Our third metric is the difference between cumulative sum of squared errors (CSSE) from the historical mean equity premium and the CSSE generated a set of predictor variables (Welch and Goyal, 2008). The out-of-sample difference in cumulative sum of squared errors ($\Delta CSSE_{OOS}$) is:

$$\Delta CSSE_{OOS} = \sum_{t=1}^{s}(r_t - \overline{r}_t)^2 - \sum_{t=1}^{s}(r_t - \hat{r}_t)^2. \tag{19}$$

We compute (19) over time for $s = s_0, ..., T-1$, which enables us to display graphically the evolution and robustness of the equity premium prediction and its performance. A value $\Delta CSSE_{OOS} > 0$ implies a greater cumulative sum of squared error for the equity premium forecast generated by the historical mean equity premium than for the equity premium forecast generated by a set of predictor variables. Thus, a positive slope indicates that the predictor forecasts are outperforming the benchmark. As an empirical benchmark, only one of the 15 equity premium forecasts plotted in Jondeau et al. (2019) using common predictors in the literature consistently exceeds a $\Delta CSSE_{OOS}$ value of 1.5% and only three of those forecasts exceed a value of 1% for a sustained period of time.

Table 4 displays the R^2_{OS} statistics for the out-of-sample predictive regressions based on the following sets of predictor variables: (q), $(q, \alpha q)$, and $(q, bw \cdot q)$. From Table 4, we see that market volatility does not have out-of-sample predictive power, with a negative R^2_{OS} statistic. In contrast, and supporting the predicted moderating effect from Proposition 1, the pair of predictors consisting of market volatility and the product of market volatility and market ambiguity attitude produce a positive R^2_{OS} statistic above 4%.

Table 4. R^2_{OS} for the Risk-Return Tradeoff

	Monthly Horizon		Quarterly Horizon	
Predictors	R^2_{OS}	CW	R^2_{OS}	CW
q_t	-0.16	0.16	-0.28	0.13
q_t, $\alpha_t q_t$	4.06	2.77***	4.08	2.93***
q_t, $bw_t q_t$	0.64	1.09	1.53	1.54*

Notes: The Table displays the Campbell and Thompson (2008) R^2_{OS} statistic (in percent) for three sets of predictor variables at the monthly forecast horizon (one month ahead) of the log equity premium. The sets of predictors are market volatility (q_t); market volatility and the product of volatility and ambiguity attitude (q_t, $\alpha_t q_t$); and market volatility and the product of volatility and the sentiment index (q_t, $bw_t q_t$). CW is the Clark and West (2007) MSPE-adjusted statistic. *, **, and *** denotes significance at the 10%, 5%, and 1% levels, respectively. The out-of-sample period spans the second half of our sample, from 2006:07 - 2022:12.

Figure 2 plots the evolution of the forecast performance of market volatility (left panel) and of

Figure 2. Market Ambiguity Attitude and the Risk-Return Tradeoff Out-of-Sample

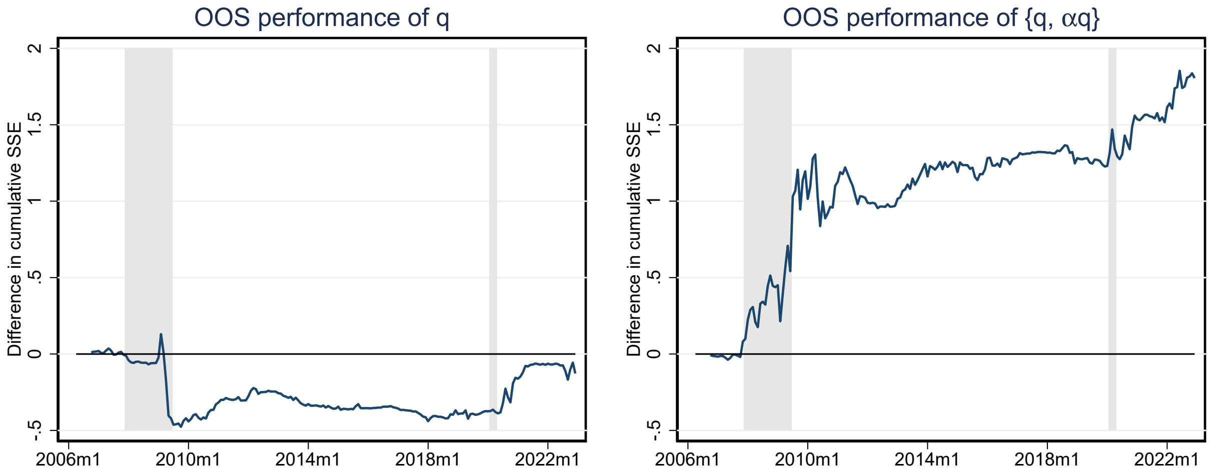

Notes: This figure displays the difference in cumulative sum of squared errors, $\Delta CSSE_{OOS}$ (in percent), between the three-month-ahead forecast of the log equity premium based on the historical average and the three-month-ahead forecast based on the conditional market volatility, q, from a GARCH(1,1) model (from Section 2.6) in the left panel. The right panel displays the $\Delta CSSE_{OOS}$ between the forecast based on the historical average and the forecast based on the pair of predictors consisting of the conditional market volatility, q, and the product of q and the conditional market ambiguity attitude, α. The out-of-sample period is 2006:07 - 2022:12. Shaded periods are NBER recessions.

market volatility and the product of ambiguity attitude and volatility (right panel) over the second half of the sample period (July, 2006, through December, 2022) which constitutes our out-of-sample period. The difference between the two panels is striking. Market volatility has greater cumulative sum of squared forecast errors than the historical average equity premium forecast throughout the sample period. In contrast, the forecast in the right panel has a consistent upward trend, indicating our model forecasts are outperforming the benchmark. Similar results hold for the one-month-ahead (monthly) forecast horizon as shown in the Internet Appendix in Figure 3.

5. Sources of Ambiguity Attitude Predictability

A natural question given the preceding results is what is the source of the ability of market ambiguity attitude to predict future returns? In this section we briefly explore two channels through which α_t might predict future returns - through predicting large market declines (such as market crashes) and through predicting recessions which could reflect time variation in macroeconomic risk related to the real economy.

5.1. Ambiguity Attitude and Market Crashes

Under the model in Section 2, a high α_t reflects an over-valued market relative to the benchmark of an expected utility representative agent. In the Markov-switching model, a high α_t also coincides with a regime in which $\frac{\mathrm{EP}_t}{q_t\gamma_t}$ is low, which corresponds on average to low expected market excess returns and high conditional market volatility. Given the persistent dynamics of α_t, it should then predict future large market declines which are natural consequences of a regime with low expected returns and high volatility.

Table 5 provides evidence that a high level of α_t systematically precedes large market declines. The table shows the frequency of large market declines (one-month market returns below -10% (top row) and below -5% (bottom row) in three sample cases. For the full sample, there were six market crashes of at least 10%, that occurred in roughly 1.5% of the periods, and 39 market crashes of at least 5% that occurred in approximately 10% of the periods. The second case shown in Table 5 is the frequency of crashes that occurred in periods in which α_t was in the top 33% of α_t values within the preceding three months (across the full sample of α_t values). The table shows that a high level of market ambiguity attitude in the three months prior to a given period t increases the frequency of 10% crashes in period t to above 4%, more than double the unconditional average. The frequency of 5% crashes also roughly doubles to nearly 20%. In contrast, none of the 10% market declines and less than five percent of the 5% market declines occurred in periods in which α_t was not in the top 33% of α_t values in the preceding three months.

The second and third column in Table 5, described in the previous paragraph, display the frequency of market crashes in period t, given a high or low level of α_t in the preceding three months. In contrast, the fourth column in Table 5 presents the frequency of a high level of α_t in the preceding three months, given a crash occurred in period t. From the table, we see that all six market crashes of at least 10% occurred in periods in which α_t was in the top 33% of α_t values in the previous three months. The bottom row shows that roughly 72% of all 5% crashes occurred in periods in which α_t was in the top 33% of all α_t values in the previous three months. Our approach complements Gormsen and Jensen (2022), who find that large market declines are preceded by periods of low market volatility, by demonstrating that they are also preceded by periods of high market ambiguity attitude.

Table 5. Frequency of Market Crashes

	Frequency of Market Crashes			Crashes Predicted
	Unconditional	α (Top 33%)	α (Bottom 67%)	α (Top 33%)
10% Market Declines	1.55% (6)	4.14% (6)	0.00% (0)	100.00% (6)
5% Market Declines	9.95% (39)	19.31% (28)	4.45% (11)	71.79% (28)

Notes: The table displays the frequency of large market declines in percent (with the total number in parentheses) across the (i) full sample period (Unconditional); (ii) across periods in which α surpassed the top 33% of (full-sample) α values within the preceding three months; (iii) across periods in which α did not surpass the top 33% of α values within the preceding three months. The fourth column displays the proportion of realized crashes that occurred in a period in which α surpassed the top 33% of α values within the preceding three months. The first row displays the results for one-month declines in the market exceeding 10%. The second row displays the results for one-month market declines exceeding 5%. The data covers the full sample period from 1990:01 - 2022:12.

Table 6 uses logistic regressions to test whether market ambiguity attitudes predicts 5% or 10% market crashes. We run logistic regressions in which the left-hand-side variable is an indicator of either a 10% market crash or a 5% market crash. In our baseline specification summarized in column (1) (for a 10% crash) and column (6) (for a 5% crash) we include only α on the right-hand-side (lagged three months). In the remaining columns, we include three control variables (each lagged three months) which are the variables that were used in the construction of α and which are each plausible predictors of a market crash: the conditional market volatility, q, the VIX index, and the market price-dividend ratio.

In Panel A of Table 6, we see that α is a significant predictor of both 10% and 5% market declines at the quarterly forecast horizon and that it subsumes any predictability present in q, VIX, or pd for market crashes. Panel A uses the full sample period (1990 - 2022). Panel B shows that similar results obtain when the logistic regression uses only data from the out-of-sample period (July, 2006 - December, 2022).

5.2. Ambiguity Attitude and NBER Recessions

The predictive ability of α might also reflect time-varying macroeconomic risk that is linked to the real economy. Indeed, Rapach et al. (2010) argue that a promising explanation for successful forecasts of the equity premium is that they reflect time-variation in macroeconomic risk linked to the real economy. In the case of the market ambiguity attitude, if α follows a mean-reverting

Table 6. Predicting Market Crashes with Market Ambiguity Attitude

Logistic Regressions for Predicting Market Crashes (Full Sample Period)

Panel A	(1) -10%	(2) -10%	(3) -10%	(4) -10%	(5) -10%	(6) -5%	(7) -5%	(8) -5%	(9) -5%	(10) -5%
α_{t-3}	1.65***	1.58***	1.58***	2.38***	2.66***	0.77***	0.81***	0.81***	0.83***	0.93***
	(4.06)	(3.47)	(3.15)	(3.60)	(5.23)	(4.19)	(3.89)	(3.60)	(3.14)	(2.89)
q_{t-3}		0.31			0.29		-0.08			-0.06
		(0.46)			(1.06)		(-0.38)			(-0.29)
VIX_{t-3}			0.17		-0.32			-0.07		-0.07
			(0.27)		(-0.79)			(-0.26)		(-0.27)
pd_{t-3}				-0.69	-0.81				-0.08	-0.12
				(-1.15)	(-1.57)				(-0.35)	(-0.51)
Pseudo R^2	0.164	0.171	0.167	0.199	0.206	0.079	0.080	0.080	0.080	0.081

Logistic Regressions for Predicting Market Crashes (Out-of-Sample Period)

Panel B	(1) -10%	(2) -10%	(3) -10%	(4) -10%	(5) -10%	(6) -5%	(7) -5%	(8) -5%	(9) -5%	(10) -5%
α_{t-3}	1.90***	2.11**	1.68***	1.85***	2.17***	0.76***	0.81***	0.80***	0.77***	0.86***
	(3.49)	(2.39)	(5.25)	(4.01)	(3.94)	(2.94)	(3.14)	(2.98)	(2.86)	(3.00)
q_{t-3}		0.54			0.12		-0.14			-0.17
		(1.20)			(0.31)		(-0.73)			(-0.83)
VIX_{t-3}			0.35		-0.39			-0.08		-0.03
			(0.82)		(-0.78)			(-0.31)		(-0.12)
pd_{t-3}				-1.60*	-2.21				-0.04	-0.17
				(-1.84)	(-1.37)				(-0.14)	(-0.45)
Pseudo R^2	0.159	0.214	0.188	0.271	0.284	0.072	0.076	0.073	0.072	0.077

Robust Z statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays the slope coefficients from logistic regressions. The left-hand-side variable equals one in period t if a market return less than -10% occurred in period t, and zero otherwise. In columns (6) - (10), the left-hand-side variable equals one in period t if a market return less than -5% occurred in period t, and zero otherwise. The right-hand-side variables (each lagged three months) are the market ambiguity attitude. α, the conditional market volatility, q, the VIX index of the Chicago Board of Options Exchange, and the price-dividend ratio, pd, of the S&P 500 index. Panel A displays the results for the full sample period (1990:01 - 2022:12). Panel B displays the results for the second half of the sample (the out-of-sample period, spanning 2006:07 - 2022:12). For convenience in interpreting the coefficients, each right-hand-side variables is divided by its full-sample standard deviation.

stochastic process, high α today predicts a lower α in the future. If the economy slows down as α declines (i.e., if a decline in optimism reduces consumption expenditures by consumers and investment by firms), then high α could positively forecast recessions. To explore this possibility, we consider whether α systematically predicts NBER recessions.

Table 7 reveals that roughly 9% of periods in our sample are classified as recession periods. In contrast, nearly 18% of periods in which α_t entered the top 33% of α values in the three months before period t are recession periods, roughly double the unconditional average. Only 4% of periods in which α_t did not surpass the top 33% of α values in the preceding three months are recession periods, less than half the unconditional average. The table also shows that 72% of NBER recession periods occurred within three months after α reached the top 33% of α values.

Table 7. Frequency of NBER Recession Periods

	Frequency of NBER Recession Periods			Recessions Predicted
	Unconditional	α (Top 33%)	α (Bottom 67%)	α (Top 33%)
NBER Recessions	9.18% (36)	17.93% (26)	4.05% (10)	72.22% (26)

Notes: The table displays the frequency of NBER recession periods in percent (with the total number in parentheses) across the (i) full sample period (Unconditional); (ii) across periods in which α surpassed the top 33% of (full-sample) α values within the preceding three months; (iii) across periods in which α did not surpass the top 33% of α values within the preceding three months. The fourth column displays the proportion of recession periods that occurred in a period in which α surpassed the top 33% of α values within the preceding three months. The data covers the full sample period from 1990:01 - 2022:12.

Table 8 summarizes logistic regressions for the full sample period (in Panel A) and the out-of-simple period (in Panel B), with α as a predictor variable for NBER recessions at the three-month horizon.[12] Liu and Moench (2016) investigate which variables best predict recessions and identify the term spread and the aggregate stock market return as the two strongest recession predictors at short horizons including the three-month forecast horizon. Guha and Hiris (2002) find that credit spreads are also useful predictors of recessions. We consider three related variables as candidate NBER recession predictors: (i) the term spread, TMS, defined as the difference between the long-term yield on U.S. government bonds and the U.S. treasury bill; (ii) the aggregate stock market return, R_m, from Kenneth French's data library; and (iii) the default yield spread (DFY), defined as the difference between BAA and AAA-rated corporate bond yields. Each of these variables has

[12]Similar results are obtained using probit regressions.

significant predictive power for NBER recessions over our sample period. As additional control variables we include the price dividend ratio of the S&P 500 index (PD), the VIX index from the CBOE, and q from the GARCH(1,1) model in Section 2.6 (which are the variables used in the construction of α). Finally, we include as controls the Baker and Wurgler (2006) market sentiment index and the lagged NBER recession indicator.

Table 8 shows that α significantly predicts NBER recessions across each set of control variables, for both the full-sample period and the out-of-sample period. Across the full sample period, the gain in Pseudo R^2 relative to not including α ranges from 4% to 13.4%. For regression specification (6) in Panel A with all eight control variables included, adding α to the regression increases the Pseudo R^2 by 11.2%. The effect is stronger for the out-of-sample period. For regression specification (6) in Panel B, adding α to the eight control variables which already include strong recession predictors, the Pseudo R^2 increases by nearly 15%. The results in this section provide initial evidence that market ambiguity attitude is an important determinant of stock market fluctuations and business cycle fluctuations.

6. Out-of-Sample Investment Performance

We next consider the investment performance of a dynamically optimized portfolio that uses the equity premium forecasts generated by q_t and $\alpha_t q_t$. As discussed by Campbell and Thompson (2008), Jondeau et al. (2019), and Giglio et al. (2021), the Markowitz optimal weight on the market portfolio is given by:

$$w_t = \frac{1}{\lambda}\left[\frac{E_t[R_{t+1}] - R_{f,t}}{\mathrm{var}_t[R_{t+1}]}\right]. \tag{20}$$

In equation (20), $E_t[R_{t+1}] - R_{f,t}$ is the time t expected market return in excess of the risk-free rate, $\mathrm{var}_t[R_{t+1}]$ is the conditional variance of the market return, λ is the investor's coefficient of relative risk aversion. Giglio et al. (2021) note that experimental estimates of this coefficient are between 3 and 10 and they consider 4 to be a realistic value of λ. Following Jondeau et al. (2019) we add the realistic portfolio constraint that $w_t \in [0, 2]$ which excludes short-selling and permits at

Table 8. Predicting Recessions with Market Ambiguity Attitude

Logistic Regressions for Predicting Recessions (Full Sample Period)

Panel A	(1) REC	(2) REC	(3) REC	(4) REC	(5) REC	(6) REC
α_{t-3}	0.85***	0.79***	0.77***	1.95***	0.79***	3.26***
	(4.68)	(3.18)	(3.00)	(6.06)	(4.46)	(4.00)
Lagged Recession Predictors	NO	YES	NO	NO	NO	YES
Lagged Recession Indicator	NO	NO	YES	NO	NO	YES
Lagged α Ingredients	NO	NO	NO	YES	NO	YES
Lagged Sentiment Index	NO	NO	NO	NO	YES	YES
Pseudo R^2	0.093	0.278	0.448	0.288	0.099	0.682
Δ(Pseudo R^2)	0.093	0.052	0.040	0.134	0.066	0.112

Logistic Regressions for Predicting Recessions (Out-of-Sample Period)

Panel B	(1) REC	(2) REC	(3) REC	(4) REC	(5) REC	(6) REC
α_{t-3}	2.27***	3.57***	2.10***	3.00***	2.77***	3.25***
	(3.02)	(3.55)	(3.14)	(4.44)	(3.94)	(3.26)
Lagged Recession Predictors	NO	YES	NO	NO	NO	YES
Lagged Recession Indicator	NO	NO	YES	NO	NO	YES
Lagged α Ingredients	NO	NO	NO	YES	NO	YES
Lagged Sentiment Index	NO	NO	NO	NO	YES	YES
Pseudo R^2	0.290	0.593	0.617	0.612	0.379	0.747
Δ(Pseudo R^2)	0.290	0.275	0.142	0.311	0.375	0.147

Robust z statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table displays the slope coefficients from logistic regressions. The left-hand-side variable is the NBER recession indicator (REC) from the St. Louis Federal Reserve database. REC is equal to one in period t if there was a recession in period t and is equal to zero otherwise. The right-hand-side variables (each lagged three months) include the market ambiguity attitude (α) and eight control variables: (i) the term spread (TMS) (the difference between the long-term U.S. government bond yield and the U.S. treasury bill) from Welch and Goyal (2008) (ii) the default yield spread (DFY) (the difference between BAA and AAA-rated corporate bond yields) from Welch and Goyal (2008) (iii) the aggregate market return (R_m) from Kenneth French's data library (The variables TMS, DFY, and R_m are our 'recession predictor' variables as these are known to have forecasting power for NBER recessions); (iv) the lagged NBER recession indicator; (v) the price-dividend ratio of the S&P 500 index from Robert Shiller's website (PD); (vi) the VIX index of the Chicago Board of Options Exchange; (vii) the conditional market volatility, q, from the GARCH model in Section 2.6 (The variables PD, VIX, and q are the 'α ingredients' as these variables were used in the construction of α); and (viii) the Baker and Wurgler (2006) market sentiment index. Panel A uses the full sample of available data from 1990:01 - 2022:12, except in regression specification (2) which ends in 2021:12, specification (5) which ends in 2022:06, and specification (6) which ends in 2021:12 due to data availability. Panel B displays the results for the out-of-sample period from 2006:07 through 2022:12, except in regression specifications (2), (5), and (6) due to the same data availability restriction as in Panel A. Δ(Pseudo R^2) denotes the change in Pseudo R^2 from including α in the regression relative to an otherwise identical regression that excludes α. For convenience in interpreting the coefficients, α is divided by its full-sample standard deviation.

most 100% leverage. The ex post portfolio excess return, $R^e_{p,t+1}$, at the end of month $t+1$ is then:

$$R^e_{p,t+1} = w_t \, R^e_{m,t+1}, \tag{21}$$

where $R^e_{m,t+1}$ denotes the market excess return in period $t+1$. As noted by Jondeau et al. (2019), repeating this process for each period from the first out-of-sample period through the end of the sample period, yields a time series of ex post excess returns for each optimal portfolio. We will evaluate the performance of each portfolio according to the portfolio's realized Sharpe ratio during the out-of-sample period: $SR_p = \frac{\overline{r}_p}{\sigma_p}$ where $\overline{r}_p$ is the sample mean and σ^2_p is the sample variance of the portfolio return. A second metric we compute is the certainty equivalent return on portfolio p, defined as: $CER = \overline{r}_p - (\lambda/2)\sigma^2_p$. This quantity is the risk-free return that would make a mean-variance investor with risk aversion λ indifferent between that return and investing in portfolio p. We also test if the investment strategies earn significant risk-adjusted returns relative to the Fama and French (2018) six factor model and the Hou et al. (2021) five-factor q-factor model.

We consider three equity premium forecasts, each at the one month forecast horizon. The sets of predictors used to generate the forecasts are: (i) q and $\alpha_t q_t$; (ii) q_t; and (iii) the historical average forecast. As a benchmark, we also consider a fourth investment strategy, the buy-and-hold strategy of passively holding the market portfolio. To compute the conditional variance in the optimal portfolio weight, we use q^2_t, the conditional market variance from the GARCH(1,1) model in Section 2.6. Table 9 summarizes the investment performance with strategies ranked by their realized Sharpe ratio. For each strategy, the table displays the average weight on the market portfolio ($\bar{w}$), the average monthly return (Ret), the average monthly volatility (Vol) of the portfolio return, the annualized monthly Sharpe ratio (SR), the annualized certainty equivalent return on the strategy in percent, and the annualized risk-adjusted returns relative to the Fama-French six factor model ($\hat{\alpha}^{FF6}$), and the Hou et al. (2021) five-factor q-factor model ($\hat{\alpha}^{q5}$). The abnormal returns relative to the factor models are estimated as the intercept in simple time series regressions.[13]

The investment strategy based on q_t and $\alpha_t q_t$ is the only one of the strategies to earn significant

[13] The Fama and French (2018) six factor model extends the CAPM (consisting of the market excess return as the only factor) to include five additional factors. These factors are long-short portfolios sorted on firm size, firm book-to-market ratio, firm profitability, firm investment, and firm past returns. The Hou et al. (2021) five-factor q model consists of a market factor, and factors formed from the following firm characteristics: firm size, investment-to-assets, return-on-equity, and expected growth.

abnormal returns relative to the Fama-French six factor model and the q-factor model. Neither the investment strategy based on market volatility, q_t, alone, or the strategy based on the historical average outperforms the passive buy-and-hold strategy in terms of the portfolio Sharpe ratio or certainty equivalent return. In contrast, the strategy that combines α_t with q_t generates a Sharpe ratio that is 43% higher than that of the passive buy-and-hold strategy and a CER that is more than double that of the buy-and-hold strategy.

Table 9. Out-of-Sample Investment Performance

Predictors	$\bar{w}$	Ret	Vol	SR	CER	α^{FF6}	α^{q5}
q, αq	1.08	1.15	4.98	0.80	8.14	6.55**	6.55*
Buy-and-hold	1.00	0.76	4.69	0.56	3.91	0.00	0.00
q	0.90	0.49	3.86	0.44	2.33	-1.07	-0.24
Historical avg.	0.93	0.48	3.95	0.42	2.03	-0.72	-0.36

Notes: The table displays the out-of-sample investment performance of investment strategies that dynamically optimize the weights on the market portfolio based on forecasts of the equity premium at the one month forecast horizon. The weight on the market portfolio in each period is the one-month-ahead equity premium forecast divided by the product of the coefficient of relative risk aversion (λ) and the conditional market variance (based on the GARCH(1,1) model from Section 2.6 estimated one-month ahead). For all strategies λ is set equal to 4 as suggested by Giglio et al. (2021). The investment strategies correspond to forecasts based on (i) market volatility, q, and the product of market optimism and market volatility, αq; (ii) q; (iii) the historical average forecast; and (iv) the passive strategy that buys and holds the market portfolio. For each strategy, the table displays the average weight on the market portfolio ($\bar{w}$), the average monthly return (Ret), the average monthly volatility (Vol) of the portfolio return, the annualized monthly Sharpe ratio (SR), the annualized certainty equivalent return, and the annualized risk-adjusted returns relative to the Fama and French (2018) six factor model (α^{FF6}), and the Hou et al. (2021) five-factor q-factor model (α^{q5}). Returns are in percent. The data spans the out-of-sample period from 2006:07, through 2022:12.

7. Robustness Checks

We perform various robustness checks to evaluate the robustness of our results. In particular, (i) We test if the regression coefficients are stable across the two halves of the sample period; (ii) We test whether α moderates the risk-return tradeoff in both halves of our sample period; and (iii) we test if the results hold using alternative GARCH volatility models.

7.1. Stability of Regression Coefficients

Motivated by Welch and Goyal (2008) and Goyal et al. (2021), We investigate if the regression coefficients are stable across the two halves of our sample (January, 1990 - June, 2006, and July, 2006 - December, 2022). Table 10 displays the in-sample regression coefficients from regressions of the form in (22) and (23). The first regression in Table 10 uses (22) where $y_{t+h} \equiv R^e_{t,t+h}$, $x_t \equiv q_t$ and D is a dummy variable that equals 0 in the first half of the sample period and 1 in the second half of the sample. In the table, $\beta_1 \equiv \beta$ and $\beta_2 \equiv \beta + \beta_{Dx}$. The second regression in Table 10 uses (23) and adds $z_t \equiv \alpha_t q_t$. The third regression in the table is similar to the second except that now $x_t \equiv \alpha_t q_t$ and $z_t \equiv q_t$.

$$y_{t+h} = a + \beta\, x_t + \beta_D\, D + \beta_{Dx}\, D\, x_t + \epsilon_{t+h} \tag{22}$$

$$y_{t+h} = a + \beta\, x_t + \beta_D\, D + \beta_{Dx}\, D\, x_t + \beta_z z_t + \epsilon_{t+h} \tag{23}$$

Similar to Goyal et al. (2021), sign changes between the estimated coefficients for the two halves of the sample period, β_1 and β_2, are marked with Δ. Statistically significant differences between β_1 and β_2 at the 10% significance level are marked with an asterisk if present.

Two observations emerge from Table 10. First, the regression with only q_t is unstable, as it changes sign from negative to positive across the two halves of the sample period. Second, the estimated coefficients for q_t are noticeably more stable when the interaction term $\alpha_t q_t$ is included in the regression. Further, the coefficients for $\alpha_t q_t$ are similar and not significantly different across the two halves of the sample period. For example, at the quarterly forecast horizon, the coefficient estimates on market volatility are -0.18 and 0.53 in the two halves of the sample period when only q is included in the regression. Including the interaction term αq in the regression yields estimated coefficients for q of 1.31 and 1.34 in the two halves of the sample period and they are not significantly different. The coefficients for αq are -1.28 and -1.31 across the two halves of the sample period and are also not significantly different.

Table 10. Stability of Coefficients for the Risk-Return Tradeoff

		Monthly			Quarterly		
x_t	z_t	β_1	β_2	Δ	β_1	β_2	Δ
q_t		-0.03	0.43	Δ	-0.18	0.53	Δ
q_t	$\alpha_t q_t$	1.78	1.43		1.31	1.34	
$\alpha_t q_t$	q_t	-1.35	-1.69		-1.28	-1.31	

Notes: The table displays coefficients from full-sample predictability regressions for the first and second half of the sample period for the monthly and quarterly forecast horizons. The first sub-sample spans monthly data from 1990:01 - 2006:06. The second sub-sample spans from 2006:07 - 2022:12. β_1 and β_2 denote the estimated coefficients for the first and second halves of the sample. They are estimated from the regression $y_{t+h} = a + \beta\, x_t + \beta_D\, D + \beta_{Dx}\, D\, x_t + \beta_z z_t + \epsilon_{t+h}$ where D is a dummy variable that equals 0 in the first half of the sample period and 1 in the second half of the sample, $\beta_1 \equiv \beta$ and $\beta_2 \equiv \beta + \beta_{Dx}$. The predictor variables include market ambiguity attitude, α_t, conditional market volatility, q_t, measured from a GARCH(1,1) model, and the product $\alpha_t q_t$. Δ denotes a sign change across the two halves of the sample (β_1 and β_2 are of opposite sign). For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

7.2. Performance Across Subsamples

We test whether market ambiguity attitude moderates the risk-return tradeoff in each half of our sample period in Table 11. The regression results presented in columns (6) and (9) in both Panels A and B of Table 11 reveal that α significantly moderates the risk-return tradeoff for both halves of the sample period (the training period from 1990:01 - 2006:06 and the out-of-sample period from 2006:07 - 2022:12) at both the monthly and quarterly forecast horizons. In each case, the significantly positive risk-return relation is recovered when αq is included in the regression with q. Further, note that including both q_t and $\alpha_t q_t$ at least doubles the R^2, relative to just including q_t for each of regressions (3), (6), and (9) at both the monthly and quarterly forecast horizons.

7.3. Alternative GARCH Volatility Models

We test if the risk-return tradeoff results also hold if q is constructed as a standard simple GARCH(1,1) model (without including the price-dividend ratio) or as a GJR GARCH model. As with our main specification (the GARCH(1,1) model from Section 2.6), both the simple GARCH(1,1) model and the GJR GARCH model are recursively estimated and are free of look-ahead bias.

Table 15 in the Online Appendix shows the risk-return tradeoff results for the case where q is a standard simple GARCH(1,1) model (similar to that in Section 2.6 but constructed with a constant

Table 11. Market Ambiguity Attitude and the Risk-Return Tradeoff Across Subsamples

Monthly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t
q_{t-1}	0.30		1.35***	0.43		1.51***	-0.02		1.56**
	(1.18)		(5.44)	(1.38)		(5.45)	(-0.04)		(2.29)
$\alpha_{t-1} q_{t-1}$		-0.32	-1.36***		-0.28	-1.69**		-0.35	-1.37**
		(-1.07)	(-3.87)		(-0.55)	(-2.59)		(-1.02)	(-2.35)
R^2	0.005	0.005	0.042	0.011	0.003	0.061	0.000	0.007	0.027

Quarterly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel B	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t	R^e_t
q_{t-3}	0.33		1.25***	0.54**		1.39***	-0.19		1.19**
	(1.51)		(4.87)	(2.48)		(4.34)	(-0.45)		(2.11)
$\alpha_{t-3} q_{t-3}$		-0.24	-1.20***		-0.05	-1.34**		-0.42	-1.19**
		(-0.96)	(-3.48)		(-0.12)	(-2.44)		(-1.27)	(-2.23)
R^2	0.005	0.003	0.034	0.018	0.000	0.049	0.001	0.010	0.021

Newey-West t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays regressions of the log equity premium, R^e_t, in percent, against the conditional stock market volatility (q_{t-h}), estimated from a GARCH(1,1) model (from Section 2.6), in regression specifications (1), (4), and (7); regressions of the log equity premium against the product of market ambiguity attitude and the conditional market volatility ($\alpha_{t-h} q_{t-h}$) in regression specifications (2), (5), and (8); and regressions of the log equity premium against both variables in regression specifications (3), (6), and (9). Formally, we run versions of the following regression that include one or both of the right-hand-side variables:

$$R^e_t = \beta_0 + \beta_1 q_{t-h} + \beta_2 \alpha_{t-h} q_{t-h} + \epsilon_t, \tag{24}$$

Regressions (1), (2), and (3) span monthly data from the full sample period (1990:01 - 2022:12). Regressions (4), (5), and (6) use data from the second half of this sample (2006:07 - 2022:12) which is the period for which q_t and α_t are recursively estimated using only information available to investors in real time. Regressions (7), (8), and (9) use data from the first half of the sample period (1990:01 - 2006:06) which served as the training period in which q_t and α_t were estimated using all data in the first half of the sample. Regressions in Panel A are over a forecast horizon of $h = 1$ month (monthly horizon). Regressions in Panel B are over a forecast horizon of $h = 3$ months (quarterly horizon). For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

mean instead of a time-varying mean based on the price-dividend ratio). The table shows the results for both monthly and quarterly forecast horizons and for the full sample and each subsample.

The results for the standard simple GARCH model are similar to our baseline results: By itself, q does not predict the equity premium, but including both q and αq yields a positive and significant coefficient on q and a negative and significant coefficient on αq at both the monthly and quarterly forecast horizon. At the monthly horizon, the R^2 jumps from 0.4% with only q to 4.2% with both q and αq. At the quarterly horizon, the R^2 jumps from 0.4% with only q to 3.3% with both q and αq. These results for the full sample are stronger for the out-of-sample period. For the training period, the coefficients for both the monthly and quarterly horizon forecasts are also significant at the 10% level.

Table 17 in the Online Appendix reveals that the coefficient estimates under the simple GARCH model are similar across the two halves of the sample period when both q and αq are included in the regression. For example, for q, the estimated coefficient is 1.18 for the first half and 1.31 for the second half at the quarterlyy horizon. For αq, the estimated coefficient is -1.28 for the first half and also -1.28 for the second half of the sample at the quarterly horizon.

Figure 4 displays the predictive performance over time (the $\Delta CSSE$ plots) for out-of-sample regressions with q (left panel) and both q and αq (right panel) for the one-month forecast horizon (top panel) and the one-quarter forecast horizon (bottom panel) where q is the standard GARCH volatility. The figure shows that while by itself, q under-performs the benchmark, the forecast with both q and αq consistently outperforms the benchmark with a $\Delta CSSE$ above 1% that increases across the out-of-sample period and is close to 2% by the end of the sample period.

In addition to the simple GARCH(1,1) model, we apply a GJR GARCH model. Tables 16 and 18 in the Online Appendix summarize the results and indicate the GJR specification yields similar results to the simple GARCH model. The out-of-sample results for the GJR GARCH model are shown in the bottom panel of Table 19 in the Online Appendix. Figure 5 shows the out-of-sample performance of the GJR GARCH model over time which is similar to that of the simple GARCH model shown in Figure 4.

8. Conclusion

This paper studies the effect of market ambiguity attitude on the risk-return tradeoff. We consider a representative agent asset pricing model in which equilibrium prices depend on an information component (reflecting the asset's discounted expected value) and a noise component (reflecting the market's ambiguity attitude). Market efficiency (in which prices fully reflect available information) deteriorates in periods of high uncertainty (at which time prices partially reflect information and partially reflect the market's ambiguity attitude). The equilibrium equity premium depends on market optimism, Knightian uncertainty, positive skewness, and disaster risk, linking these strands of the asset pricing literature. We derive the theoretical implication that the equity premium is increasing in market volatility and the slope of this relationship flattens as market ambiguity attitude increases. We develop a theory-based measure of the market's ambiguity attitude and tested the theoretical predictions that it explains time variation in the market risk-return tradeoff and that it predicts market crashes.

We find that the predicted positive relationship between the equity premium and the conditional market volatility is observed only after accounting for the market ambiguity attitude. This finding holds both in-sample and out-of-sample, at both the monthly and quarterly forecast horizons, and it is not explained by market sentiment or established equity premium predictors. We also document that market ambiguity attitude predicts market crashes, consistent with high levels of optimism reflecting an over-valued market relative to an expected utility representative agent. Further, market ambiguity attitude predicts NBER recessions, thereby providing a new link between the aggregate stock market and the real economy. Our results indicate that market ambiguity attitude is an important state variable in driving time-varying expected returns, and might help to bridge the gap between irrational exuberance in the stock market and equilibrium asset pricing theory.

Appendix

Proof of Proposition 2: Let us rewrite the definition of VRP in (11), and use the fact that under the risk-neutral measure, the expected market return equals the risk-free rate. Thus,

$$VRP_t = (1-\gamma_t)E_t(R_{t+1}-R_{f,t})^2 + \gamma_t\left[\alpha_t(\overline{R}_{t+1}-R_{f,t})^2 + (1-\alpha_t)(\underline{R}_{t+1}-R_{f,t})^2\right] - q_t^2.$$

Using $E_t(R_{t+1} - R_{f,t})^2 = q_t^2$, and applying Assumption 1 with $\overline{\xi} = \underline{\xi} = \xi$ gives us:

$$VRP_t = (1-\gamma_t)q_t^2 + \gamma_t\big[\alpha_t(EP_t + \xi q_t)^2 + (1-\alpha_t)(EP_t - \xi q_t)^2\big] - q_t^2,$$

where EP_t is the equity premium. Next, replacing EP_t from Equation (10), gives us:

$$VRP_t = (1-\gamma_t)q_t^2 + \gamma_t\Big[\alpha_t\big(\xi q_t\gamma_t(1-2\alpha_t) + \xi q_t\big)^2 + (1-\alpha_t)\big(\xi q_t\gamma_t(1-2\alpha_t) - \xi q_t\big)^2\Big] - q_t^2.$$

After simplification, rearranging the terms, and factoring $\xi^2 q_t^2$, we have

$$VRP_t = \big[1 + (\gamma_t^2 - 2\gamma_t)(1-2\alpha_t)^2\big]\xi^2 q_t^2\gamma_t - \gamma_t q_t^2.$$

On average, the term $(\gamma_t^2 - 2\gamma_t)(1-2\alpha_t)^2$ is about two orders of magnitude smaller than 1 (this is because both γ_t and $(1-2\alpha_t)$ are on average small values), and hence, we can drop that term in the bracket and approximate the bracket with one, which yields the result in Equation (12).[14] $\square$

Proof of Proposition 3: First, notice that we can approximate the risk-neutral probability density for the NEO-EU agent with three points: the mean of the distribution (for the EU part) and the two extreme points for the best and worst case returns. Without loss of generality, we can shift the three-point distribution such that the risk-neutral weights are placed on the three points $-\xi$, 0, and ξ. Since the probabilities of the points $(-\xi, 0, \xi)$ are $\big(\gamma(1-\alpha), 1-\gamma, \gamma\alpha\big)$, the skewness of this distribution is

$$\frac{(2\alpha-1)\gamma\left((1-\gamma)(2(1-\gamma)-1) - 8\alpha(1-\alpha)\gamma^2\right)}{(\gamma(4\alpha(1-\alpha)\gamma + 1-\gamma))^{\frac{3}{2}}}.$$

The derivative of this expression with respect to α is

$$\frac{2\left(1-\gamma-8\alpha(1-\alpha)\gamma\right)}{\left(\gamma\left(4\alpha(1-\alpha)\gamma + 1-\gamma\right)^5\right)^{\frac{1}{2}}}.$$

[14] For our baseline specification of $\xi = 4.77$, plugging the unconditional values of $\gamma = 0.05$ and $\alpha = 0.27$ from Table 1 yields an expression for the term in the bracket of 0.98. Plugging in the estimated values of γ_t and α_t to construct a time series of the term in the bracket yields an unconditional value for the bracketed term of 0.97. These observations confirm that the term in the bracket can reasonably be approximated by one.

Thus, $\frac{\partial \text{ skewness}}{\partial \alpha} > 0$ if

$$\gamma < \frac{1}{1 + 8\alpha(1-\alpha)}. \tag{25}$$

For the sake of numerical comparison, the right hand side is smallest at $\alpha = 0.5$, for which the upper bound on γ is $\frac{1}{3}$. As α approaches 0 or 1, the upper bound on γ approaches one. $\square$

LEMMA 1: *If the payoff X_{t+1} in the model is replaced with the dividend D_{t+1}, then replacing Assumption 1 with $\overline{D}_{t+1} = (1+\xi q_t)D_t$, $\underline{D}_{t+1} = (1-\xi q_t)D_t$, and $E_t[D_{t+1}] = D_t$, gives approximately the same equation for the equity premium in (10). Moreover, the expected log return is approximately linear in the price-dividend ratio.*

Proof. Given the best and worst case scenarios for future dividend, i.e., $\overline{D}_{t+1} = (1 + \xi q_t)D_t$ and $\underline{D}_{t+1} = (1 - \xi q_t)D_t$, the price Equation (5) becomes

$$P_t = (1-\gamma_t)\delta E_t[D_{t+1}] + \gamma_t\delta\big(\alpha_t\overline{D}_{t+1} + (1-\alpha_t)\underline{D}_{t+1}\big)$$
$$P_t = (1-\gamma_t)\delta D_t + \gamma_t\delta\Big(\alpha_t(1+\xi q_t)D_t + (1-\alpha_t)(1-\xi q_t)D_t\Big),$$

and hence, we have the price-dividend ratio

$$pd_t = \delta\big(1 - \xi\gamma_t q_t(1-2\alpha_t)\big).$$

Thus, $log(pd_t) = log(\delta) + log\big(1 - \xi\gamma_t q_t(1-2\alpha_t)\big)$, and using the approximation $log(1+x) \approx x$, we find that as is the case with pd_t, the log price-dividend ratio $log(pd_t)$ is (approximately) linear in $\xi\gamma_t q_t(1-2\alpha_t)$. Note that the approximation is accurate if $\gamma_t q_t$ is small, and in our monthly data, both γ_t (on average 0.05) and q_t (on average 0.04) are small.

As for the expected return, we have $E_t R_{t+1} = \frac{E_t[D_{t+1}]}{P_t} = \frac{1}{pd_t}$, so the log expected return is linear in $log(pd_t)$, which is linear $\xi\gamma_t q_t(1-2\alpha_t)$. Thus, we showed that approximately, both pd_t and the log expected return are linear in $\xi\gamma_t q_t(1-2\alpha_t)$, and hence, the log expected return is also linear in pd_t. Finally, note that the log expected return and expected log return are off by a Jensen's term. Thus, as long as this Jensen's term is negligible, the expected log return is approximately linear in the price-dividend ratio. The last condition on the negligibility of the Jensen's term can be checked in the data, and we find that in our monthly data, log expected returns and expected log returns have an almost perfect correlation.

□

REFERENCES

Anthropelos, Michail, and Paul Schneider, 2022, Optimal investment and equilibrium pricing under ambiguity, *arXiv preprint* arXiv:2206.10489.

Azimi, Mehran, Soroush Ghazi, and Mark Schneider, 2023, The market sharpe ratio, equity premium prediction, and investment under knightian uncertainty: The role of optimism. Manuscript.

Baillie, Richard T, and Ramon P DeGennaro, 1990, Stock returns and volatility, *Journal of Financial and Quantitative Analysis* 25, 203–214.

Baillon, Aurélien, and Han Bleichrodt, 2015, Testing ambiguity models through the measurement of probabilities for gains and losses, *American Economic Journal: Microeconomics* 7, 77–100.

Baillon, Aurélien, Han Bleichrodt, Umut Keskin, Olivier l'Haridon, and Chen Li, 2018, The effect of learning on ambiguity attitudes, *Management Science* 64, 2181–2198.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *The Journal of Finance* 61, 1645–1680.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, 2016, Measuring economic policy uncertainty, *The Quarterly Journal of Economics* 131, 1593–1636.

Bali, Turan G, and Hao Zhou, 2016, Risk, uncertainty, and expected returns, *Journal of Financial and Quantitative Analysis* 51, 707–735.

Barro, Robert J, and Gordon Y Liao, 2021, Rare disaster probability and options pricing, *Journal of Financial Economics* 139, 750–769.

Barroso, Pedro, and Paulo F Maio, 2023, The risk-return tradeoff among equity factors, *Available at SSRN* 2909085.

Bekaert, Geert, and Marie Hoerova, 2014, The vix, the variance premium and stock market volatility, *Journal of Econometrics* 183, 181–192.

Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of econometrics* 31, 307–327.

Bossaerts, Peter, Paolo Ghirardato, Serena Guarnaschelli, and William R Zame, 2010, Ambiguity in asset markets: Theory and experiment, *The Review of Financial Studies* 23, 1325–1359.

Brandt, Michael W, and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent var approach, *Journal of Financial Economics* 72, 217–257.

Brandt, Michael W, and Leping Wang, 2007, Measuring the time-varying risk-return relation from the cross-section of equity returns. Manuscript.

Brennan, Michael J, Ashley W Wang, and Yihong Xia, 2004, Estimation and test of a simple model of intertemporal capital asset pricing, *The Journal of Finance* 59, 1743–1776.

Brenner, Menachem, and Yehuda Izhakian, 2018, Asset pricing and ambiguity: Empirical evidence, *Journal of Financial Economics* 130, 503–531.

Campbell, John Y, 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.

Campbell, John Y, and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–318.

Campbell, John Y, and Samuel B Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *The Review of Financial Studies* 21, 1509–1531.

CBOE, 2011, The CBOE SKEW Index - SKEW, `https://cdn.cboe.com/resources/indices/documents/SKEWwhitepaperjan2011.pdf`, [Online; accessed 11-November-2023].

Cederburg, Scott, Travis L Johnson, and Michael S O'Doherty, 2023, On the economic significance of stock return predictability, *Review of Finance* 27, 619–657.

Chateauneuf, Alain, Jürgen Eichberger, and Simon Grant, 2007, Choice under uncertainty with the best and worst in mind: Neo-additive capacities, *Journal of Economic Theory* 137, 538–567.

Chen, Long, and Lu Zhang, 2011, Do time-varying risk premiums explain labor market performance?, *Journal of Financial Economics* 99, 385–399.

Chen, Zengjing, and Larry Epstein, 2002, Ambiguity, risk, and asset returns in continuous time, *Econometrica* 70, 1403–1443.

Clark, Todd E, and Kenneth D West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Cooper, Ilan, and Richard Priestley, 2009, Time-varying risk premiums and the output gap, *The Review of Financial Studies* 22, 2801–2833.

DeMiguel, Victor, Alberto Martin-Utrera, and Raman Uppal, 2021, A multifactor perspective on volatility-managed portfolios, *Available at SSRN 3982504* .

Dimmock, Stephen G, Roy Kouwenberg, Olivia S Mitchell, and Kim Peijnenburg, 2015, Estimating ambiguity preferences and perceptions in multiple prior models: Evidence from the field, *Journal of Risk and Uncertainty* 51, 219–244.

Dimmock, Stephen G, Roy Kouwenberg, Olivia S Mitchell, and Kim Peijnenburg, 2016, Ambiguity aversion and household portfolio choice puzzles: Empirical evidence, *Journal of Financial Economics* 119, 559–577.

Dimmock, Stephen G, Roy Kouwenberg, Olivia S Mitchell, and Kim Peijnenburg, 2021, Household portfolio underdiversification and probability weighting: Evidence from the field, *The Review of Financial Studies* 34, 4524–4563.

Dow, James, and Sérgio Ribeiro da Costa Werlang, 1992, Uncertainty aversion, risk aversion, and the optimal choice of portfolio, *Econometrica* 197–204.

Driesprong, Gerben, Ben Jacobsen, and Benjamin Maat, 2008, Striking oil: another puzzle?, *Journal of Financial Economics* 89, 307–327.

Easley, David, and Maureen O'Hara, 2009, Ambiguity and nonparticipation: The role of regulation, *The Review of Financial Studies* 22, 1817–1843.

Ebert, Sebastian, and Paul Karehnke, 2021, Skewness preferences in choice under risk, *Available at* SSRN.

Fama, Eugene F, and Kenneth R French, 2018, Choosing factors, *Journal of financial economics* 128, 234–252.

French, Kenneth R, G William Schwert, and Robert F Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Friedman, Milton, and Leonard J Savage, 1948, The utility analysis of choices involving risk, *Journal of Political Economy* 56, 279–304.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci, 2004, Differentiating ambiguity and ambiguity attitude, *Journal of Economic Theory* 118, 133–173.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509–548.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522.

Gilboa, Itzhak, and David Schmeidler, 1989, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* 18, 141–153.

Gormsen, Niels Joachim, and Christian Skov Jensen, 2022, Higher-moment risk, *Available at SSRN* 3069617.

Goyal, Amit, Ivo Welch, and Athanasse Zafirov, 2021, A comprehensive look at the empirical performance of equity premium prediction 2, *Available at SSRN* 3929119.

Guha, Debashis, and Lorene Hiris, 2002, The aggregate credit spread and the business cycle, *International Review of Financial Analysis* 11, 219–227.

Guo, Hui, and Robert F Whitelaw, 2006, Uncovering the risk–return relation in the stock market, *The Journal of Finance* 61, 1433–1463.

Hansen, Lars Peter, and Thomas J Sargent, 2001, Robust control and model uncertainty, *American Economic Review* 91, 60–66.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2021, An augmented q-factor model with expected growth, *Review of Finance* 25, 1–41.

Huang, Darien, and Mete Kilic, 2019, Gold, platinum, and expected stock returns, *Journal of Financial Economics* 132, 50–75.

Jondeau, Eric, Qunzi Zhang, and Xiaoneng Zhu, 2019, Average skewness matters, *Journal of Financial Economics* 134, 29–47.

Jones, Christopher S, and Selale Tuzel, 2013, New orders and asset prices, *The Review of Financial Studies* 26, 115–157.

Ju, Nengjiu, and Jianjun Miao, 2012, Ambiguity, learning, and asset returns, *Econometrica* 80, 559–591.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *The Review of Financial Studies* 27, 2841–2871.

Kelly, Bryan, and Seth Pruitt, 2013, Market expectations in the cross-section of present values, *The Journal of Finance* 68, 1721–1756.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.

Kocher, Martin G, Amrei Marie Lahno, and Stefan T Trautmann, 2018, Ambiguity aversion is not universal, *European Economic Review* 101, 268–283.

König-Kersting, Christian, Christopher Kops, and Stefan T Trautmann, 2023, A test of (weak) certainty independence, *Journal of Economic Theory* 209, 105623.

Kurov, Alexander, 2010, Investor sentiment and the stock market's reaction to monetary policy, *Journal of Banking and Finance* 34, 139–149.

Lettau, Martin, and Sydney C Ludvigson, 2010, Measuring and modeling variation in the risk-return trade-off, *Handbook of Financial Econometrics: Tools and Techniques,* 617–690.

Li, Jun, and Jianfeng Yu, 2012, Investor attention, psychological anchors, and stock return predictability, *Journal of Financial Economics* 104, 401–419.

Liu, Qi, Zhiwei Su, Huijun Wang, and Jianfeng Yu, 2021, Extrapolation and risk-return trade-offs, *Available at SSRN 4039296* .

Liu, Weiling, and Emanuel Moench, 2016, What predicts us recessions?, *International Journal of Forecasting* 32, 1138–1150.

Lundblad, Christian, 2007, The risk return tradeoff in the long run: 1836–2003, *Journal of Financial Economics* 85, 123–150.

Merton, Robert C, 1973, An intertemporal capital asset pricing model, *Econometrica* 867–887.

Merton, Robert C, 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *The Journal of Finance* 72, 1611–1644.

Nelson, Daniel B, 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica,* 347–370.

Pástor, L'uboš, Meenakshi Sinha, and Bhaskaran Swaminathan, 2008, Estimating the intertemporal risk–return tradeoff using the implied cost of capital, *The Journal of Finance* 63, 2859–2897.

Pástor, L'uboš, and Pietro Veronesi, 2020, Political cycles and stock returns, *Journal of Political Economy* 128, 4011–4045.

Pollet, Joshua M, and Mungo Wilson, 2010, Average correlation and stock market returns, *Journal of Financial Economics* 96, 364–380.

Quiggin, John, 1982, A theory of anticipated utility, *Journal of Economic Behavior & Organization* 3, 323–343.

Rapach, David E, Matthew C Ringgenberg, and Guofu Zhou, 2016, Short interest and aggregate stock returns, *Journal of Financial Economics* 121, 46–65.

Rapach, David E, Jack K Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *The Review of Financial Studies* 23, 821–862.

Schmeidler, David, 1989, Subjective probability and expected utility without additivity, *Econometrica* 571–587.

Thimme, Julian, and Clemens Völkert, 2015, Ambiguity in the cross-section of expected returns: An empirical assessment, *Journal of Business & Economic Statistics* 33, 418–429.

Wakker, Peter P, 2010, *Prospect theory: For risk and ambiguity* (Cambridge University Press).

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *The Review of Financial Studies* 21, 1455–1508.

Whitelaw, Robert F, 1994, Time variations and covariations in the expectation and volatility of stock market returns, *The Journal of Finance* 49, 515–541.

Yang, Aoxiang, 2022, Understanding negative risk-return trade-offs, *Available at SSRN* .

Yu, Jianfeng, and Yu Yuan, 2011, Investor sentiment and the mean–variance relation, *Journal of Financial Economics* 100, 367–381.

Zhang, Huacheng, and Guofu Zhou, 2023, Unspanned risk and risk-return tradeoff, *Available at SSRN 4344929* .

Zhou, Hao, 2018, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, *Annual Review of Financial Economics* 10, 481–497.

Zimper, Alexander, 2012, Asset pricing in a lucas fruit-tree economy with the best and worst in mind, *Journal of Economic Dynamics and Control* 36, 610–628.

Online Appendix

8.1. Data Appendix

This appendix contains the sources of data used in the paper.

1. **Market Excess Return:** The market excess return (Rm-Rf), market return (Rm), and risk-free rate (Rf) are from Kenneth French's data library: `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

2. **Baker-Wurgler Sentiment Index (BW):** The Baker and Wurgler (2006) market sentiment index (*bw*) is from Jeffrey Wurgler's website: `https://pages.stern.nyu.edu/~jwurgler/`.

3. **Barro-Liao U.S. Disaster Probabilities:** The Barro and Liao (2021) U.S. disaster probability data series is available from Gordon Liao's website at: `https://gliao.xyz/research/`.

4. **Cederburg, Johnson, & O'Doherty Equity Premium Predictors:** The eleven predictors used from Cederburg et al. (2023) were shared with us by the authors of that paper. Their data extends through December, 2017. We were able to have data updated through 2021 for all eleven of the predictors in their paper that have data available at the start of our sample period (January, 1990). The eleven predictors are: West Texas Intermediate oil price changes (Driesprong et al., 2008), the variance risk premium (Bollerslev, 1986), the output gap (Cooper and Priestley, 2009), average correlation (Pollet and Wilson, 2010), nearness to the DOW all-time high (Li and Yu, 2012), new orders-to-shipments of durable goods (Jones and Tuzel, 2013), the tail-risk measure of Kelly and Jiang (2014), the PLS book-to-market factor (Kelly and Pruitt, 2013), short interest (Rapach et al., 2016), employment growth (Chen and Zhang, 2011), and the gold-to-platinum ratio (Huang and Kilic, 2019). Data extended through 2021 for the out-of-sample short interest index is available from Guofu Zhou's website at `http://apps.olin.wustl.edu/faculty/zhou/zpublications.html`. Data extended through 2021 for eight other predictors (West Texas Intermediate oil price changes, the variance risk premium, the output gap, average correlation, nearness to the DOW-all time high, new orders-to-shipments of durable goods, the tail-risk measure of Kelly and Jiang (2014), and the PLS book-to-market factor) were provided to us by Amit Goyal. The remaining two

series (employment growth and the gold-to-platinum ratio) were extended through 2021 in Azimi et al. (2023) using publicly available data according to the procedures described in the original papers (Chen and Zhang, 2011; Huang and Kilic, 2019).

5. **Goyal-Welch Equity Premium Predictors:** The 14 Goyal-Welch equity premium predictors at the monthly frequency are available from Amit Goyal's website: `https://sites.google.com/view/agoyal145`.

6. **Ambiguity Index:** The ambiguity index from Brenner and Izhakian (2018) was provided to us directly by Yehuda Izhakian.

7. **NBER Recession Indicator:** The NBER recession indicator is from the St. Louis Federal Reserve Website (FRED), series USREC and is available at: `https://fred.stlouisfed.org/series/USREC`.

8. **Price Dividend Ratio (PD):** The price-dividend ratio (*pd*) of the S&P 500 index is computed as S&P composite price, P, divided by dividend D from Robert Shiller's website: `http://www.econ.yale.edu/~shiller/data.htm`.

9. **VIX and RNS:** The monthly VIX index and the marker Risk Neutral Skewness (RNS) are from the Chicago Board of Options Exchange (CBOE). Both are converted from daily to monthly series using the last index value for each month as the monthly value for that month. The daily VIX data is available at `https://www.cboe.com/tradable_products/vix/vix_historical_data/`. The daily SKEW index is available at `https://www.cboe.com/us/indices/dashboard/skew/`. RNS is constructed from the SKEW index of the CBOE according to the relation: RNS = (100 - SKEW)/10. (see the CBOE white paper on the SKEW index, page 5, at: `https://cdn.cboe.com/resources/indices/documents/SKEWwhitepaperjan2011.pdf`).

8.2. Supplementary Tables and Figures

Table 12. Correlations between recursively estimated α_t and Market Risk-Neutral Skewness

	RNS	RNS× Indicator
α_t	0.40^{***}	0.85^{***}

Notes: The table displays the correlations between the market ambiguity attitude, α_t, and (i) the aggregate market risk-neutral skewness (RNS) and (ii) the product of RNS and an indicator variable that equals 1 when $\frac{1}{1+8\alpha_t(1-\alpha_t)} - \gamma_t$ is above its median value and which equals zero otherwise. The correlations are for the out-of-sample period (during which α_t is recursively estimated to be free of look-ahead bias), which spans 2006:07 through 2022:12.

Table 13. Granger causality test between α_t and Market Risk-Neutral Skewness

	BIC	AIC
$\alpha \rightarrow$ RNS	$(0.010)^{**}$	$(0.035)^{**}$
RNS $\rightarrow \alpha$	(0.848)	(0.807)

Notes: The table reports Granger-causality test results for market ambiguity attitude, α, and market risk-neutral skewness (RNS). The tests are conducted for both the optimal lag length (one period) under the Bayesian Information Criterion (BIC) and for the optimal lag length (two periods) under the Akaike Information Criterion (AIC). The p-values of the tests are reported. ** denotes the 5% level of statistical significance. The tests are for the out-of-sample period (during which α is recursively estimated to be free of look-ahead bias), which spans 2006:07 through 2022:12.

Table 14. GARCH(1,1) specifications of market return for up to 3 lags of the *pd* ratio.

	Full Sample (1990 - 2022)			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)
pd_{t-1}	-0.037**			-0.033**		
	(0.015)			(0.017)		
pd_{t-2}		-0.037**			-0.032*	
		(0.015)			(0.017)	
pd_{t-3}			-0.036**			-0.033*
			(0.015)			(0.017)
Constant	2.838***	2.812***	2.744***	2.504***	2.45***	2.497***
	(0.827)	(0.834)	(0.831)	(0.916)	(0.930)	(0.935)
ARCH						
$ARCH_{t-1}$	0.192***	0.192***	0.195***	0.110	0.112	0.113
	(0.044)	(0.044)	(0.045)	(0.081)	(0.082)	(0.083)
$GARCH_{t-1}$	0.774***	0.774***	0.771***	0.864***	0.862***	0.860***
	(0.052)	(0.052)	(0.050)	(0.092)	(0.092)	(0.093)
Constant	0.985**	0.979**	0.998**	0.442	0.442	0.448
	(0.476)	(0.471)	(0.477)	(0.506)	(0.502)	(0.499)
N	395	394	393	197	196	195

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The table displays the statistics for the GARCH(1,1) model from Section 2.6 of the main text for the training period (1990:01 - 2006:06) and the for the full-sample period (1990:01 - 2022:12)). As highlighted in the main text, the GARCH model is recursively estimated each period after the training period to be free from look-ahead bias for the period from July, 2006, through December, 2022. The full sample and training sample results shown here provide a snapshot of the performance of the GARCH model at two points in time and demonstrate that the estimated coefficients are relatively stable. *pd* is the price-dividend ratio on the S&P 500 index from Robert Shiller's website. Returns are in percent. Standard errors are in parentheses.

Table 15. Market Ambiguity Attitude and the Risk-Return Tradeoff (Simple GARCH Volatility)

Monthly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-1}	0.28		1.34^{***}	0.42		1.53^{***}	-0.04		1.34^{*}
	(1.09)		(5.26)	(1.35)		(5.22)	(-0.09)		(1.93)
$\alpha_{t-1}q_{t-1}$		-0.32	-1.36^{***}		-0.28	-1.71^{**}		-0.35	-1.24^{**}
		(-1.07)	(-3.73)		(-0.55)	(-2.59)		(-1.01)	(-2.06)
R^2	0.004	0.005	0.040	0.011	0.003	0.060	0.000	0.007	0.023

Quarterly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel B	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-3}	0.30		1.23^{***}	0.54^{**}		1.41^{***}	-0.22		0.97^{*}
	(1.33)		(4.86)	(2.40)		(4.34)	(-0.55)		(1.71)
$\alpha_{t-3}q_{t-3}$		-0.25	-1.20^{***}		-0.06	-1.37^{**}		-0.42	-1.06^{*}
		(-0.99)	(-3.43)		(-0.15)	(-2.45)		(-1.29)	(-1.97)
R^2	0.004	0.003	0.033	0.017	0.000	0.048	0.002	0.010	0.019

Newey-West t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays regressions of the log equity premium, R_t^e, in percent, against the conditional stock market volatility (q_{t-h}), estimated from a standard simple GARCH(1,1) model, in regression specifications (1), (4), and (7); regressions of the log equity premium against the product of market ambiguity attitude and the conditional market volatility ($\alpha_{t-h}q_{t-h}$) in regression specifications (2), (5), and (8); and regressions of the log equity premium against both variables in regression specifications (3), (6), and (9). Formally, we run versions of the following regression that include one or both of the right-hand-side variables:

$$R_t^e = \beta_0 + \beta_1 q_{t-h} + \beta_2 \alpha_{t-h} q_{t-h} + \epsilon_t, \tag{26}$$

Regressions (1), (2), and (3) span monthly data from the full sample period (1990:01 - 2022:12). Regressions (4), (5), and (6) use data from the second half of this sample (2006:07 - 2022:12) which is the period for which q_t and α_t are recursively estimated using only information available to investors in real time. Regressions (7), (8), and (9) use data from the first half of the sample period (1990:01 - 2006:06) which served as the training period in which q_t and α_t were estimated using all data in the first half of the sample. Regressions in Panel A are over a forecast horizon of $h = 1$ month (monthly horizon). Regressions in Panel B are over a forecast horizon of $h = 3$ months (quarterly horizon). For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

Table 16. Market Ambiguity Attitude and the Risk-Return Tradeoff (GJR GARCH Volatility)

Monthly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-1}	0.22		1.32***	0.33		1.55***	-0.09		1.28*
	(0.79)		(4.75)	(0.94)		(5.11)	(-0.20)		(1.78)
$\alpha_{t-1}q_{t-1}$		-0.35	-1.40***		-0.31	-1.81***		-0.38	-1.19**
		(-1.13)	(-3.76)		(-0.61)	(-2.68)		(-1.09)	(-1.98)
R^2	0.002	0.006	0.039	0.007	0.004	0.059	0.000	0.008	0.021

Quarterly	Full Sample (1990 - 2022)			Out-of-Sample Period			Training Sample Period		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel B	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e	R_t^e
q_{t-3}	0.24		1.20***	0.41*		1.37***	-0.27		0.93
	(1.03)		(4.90)	(1.72)		(4.38)	(-0.61)		(1.59)
$\alpha_{t-3}q_{t-3}$		-0.28	-1.22***		-0.11	-1.44**		-0.45	-1.03*
		(-1.11)	(-3.55)		(-0.31)	(-2.60)		(-1.32)	(-1.91)
R^2	0.003	0.004	0.030	0.011	0.001	0.044	0.002	0.011	0.018

Newey-West t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: The table displays regressions of the log equity premium, R_t^e, in percent, against the conditional stock market volatility (q_{t-h}), estimated from a standard GJR GARCH model, in regression specifications (1), (4), and (7); regressions of the log equity premium against the product of market ambiguity attitude and the conditional market volatility ($\alpha_{t-h}q_{t-h}$) in regression specifications (2), (5), and (8); and regressions of the log equity premium against both variables in regression specifications (3), (6), and (9). Formally, we run versions of the following regression that include one or both of the right-hand-side variables:

$$R_t^e = \beta_0 + \beta_1 q_{t-h} + \beta_2 \alpha_{t-h} q_{t-h} + \epsilon_t, \tag{27}$$

Regressions (1), (2), and (3) span monthly data from the full sample period (1990:01 - 2022:12). Regressions (4), (5), and (6) use data from the second half of this sample (2006:07 - 2022:12) which is the period for which q_t and α_t are recursively estimated using only information available to investors in real time. Regressions (7), (8), and (9) use data from the first half of the sample period (1990:01 - 2022:06) which served as the training period in which q_t and α_t were estimated using all data in the first half of the sample. Regressions in Panel A are over a forecast horizon of $h = 1$ month (monthly horizon). Regressions in Panel B are over a forecast horizon of $h = 3$ months (quarterly horizon). For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

Table 17. Stability of Coefficients in Predictability Regressions (Simple GARCH)

		Monthly			Quarterly		
x_t	z_t	β_1	β_2	Δ	β_1	β_2	Δ
q_t		-0.05	0.43	Δ	-0.22	0.52	Δ
q_t	$\alpha_t q_t$	1.64	1.41		1.18	1.31	
$\alpha_t q_t$	q_t	-1.34	-1.67		-1.28	-1.28	

Notes: The table displays coefficients from full-sample predictability regressions for the first and second half of the sample period. The first sub-sample spans monthly data from 1990:01 - 2006:06. The second sub-sample spans from 2006:07 - 2022:12. β_1 and β_2 denote the estimated coefficients for the first and second halves of the sample. They are estimated from the regression $y_{t+h} = a + \beta\, x_t + \beta_D\, D + \beta_{Dx}\, D\, x_t + \beta_z z_t + \epsilon_{t+h}$ where D is a dummy variable that equals 0 in the first half of the sample period and 1 in the second half of the sample, $\beta_1 \equiv \beta$ and $\beta_2 \equiv \beta + \beta_{Dx}$. The predictor variables include market ambiguity attitude, α_t, conditional market volatility, q_t, measured from a simple GARCH(1,1) model, and the product $\alpha_t q_t$. The forecasts use a horizon of $h = 1$ month (monthly forecast horizon). Δ denotes a sign change across the two halves of the sample (β_1 and β_2 are of opposite sign). Significant differences at the 0.10 level using Newey-West standard errors with a lag length of 12 are denoted *. For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

Table 18. Stability of Coefficients in Predictability Regressions (GJR GARCH)

		Monthly			Quarterly		
x_t	z_t	β_1	β_2	Δ	β_1	β_2	Δ
q_t		-0.10	0.33	Δ	-0.26	0.40	Δ
q_t	$\alpha_t q_t$	1.68	1.37		1.22	1.25	
$\alpha_t q_t$	q_t	-1.32	-1.75		-1.24	-1.35	

Notes: The table displays coefficients from full-sample predictability regressions for the first and second half of the sample period. The first sub-sample spans monthly data from 1990:01 - 2006:06. The second sub-sample spans from 2006:07 - 2022:12. β_1 and β_2 denote the estimated coefficients for the first and second halves of the sample. They are estimated from the regression $y_{t+h} = a + \beta\, x_t + \beta_D\, D + \beta_{Dx}\, D\, x_t + \beta_z z_t + \epsilon_{t+h}$ where D is a dummy variable that equals 0 in the first half of the sample period and 1 in the second half of the sample, $\beta_1 \equiv \beta$ and $\beta_2 \equiv \beta + \beta_{Dx}$. The predictor variables include market ambiguity attitude, α_t, conditional market volatility, q_t, measured from a GJR GARCH model, and the product $\alpha_t q_t$. The forecasts use a horizon of $h = 1$ month (monthly forecast horizon). Δ denotes a sign change across the two halves of the sample (β_1 and β_2 are of opposite sign). Significant differences at the 0.10 level using Newey-West standard errors with a lag length of 12 are denoted *. For ease of interpreting the coefficients, q and αq are divided by their (full sample) standard deviation.

Table 19. R^2_{OS} (percent) for Log Equity Premium Forecasts

	Simple GARCH Volatility			
	Monthly		Quarterly	
Predictors	R^2_{OS}	CW	R^2_{OS}	CW
q_t	-0.16	0.14	-0.36	0.00
q_t, $\alpha_t q_t$	4.21	2.85***	4.03	2.99***
	GJR GARCH Volatility			
	Monthly		Quarterly	
Predictors	R^2_{OS}	CW	R^2_{OS}	CW
q_t	-0.73	-0.22	-1.13	-0.49
q_t, $\alpha_t q_t$	3.83	2.89***	3.48	2.78***

Notes: The Table displays the Campbell and Thompson (2008) out-of-sample R^2 statistic (in percent) for three sets of predictor variables at the monthly forecast horizon (one month ahead) and the quarterly forecast horizon (three months ahead) of the log equity premium. The sets of predictors are (i) market volatility (q_t); (ii) the product of market volatility and market ambiguity attitude ($\alpha_t q_t$); and (iii) market volatility and the product of volatility and ambiguity attitude (q_t, $\alpha_t q_t$). The top panel shows the results for which volatility q is generated by a simple GARCH(1,1) model. The bottom panel shows the results for which q is generated by a GJR GARCH model. CW is the Clark and West (2007) MSPE-adjusted statistic; *** denotes significance at the 1% level. The out-of-sample period spans monthly data from 2006:07 - 2022:12.

Figure 3. OOS Equity Premium Prediction with Volatility and Optimism (One-Month Forecast)

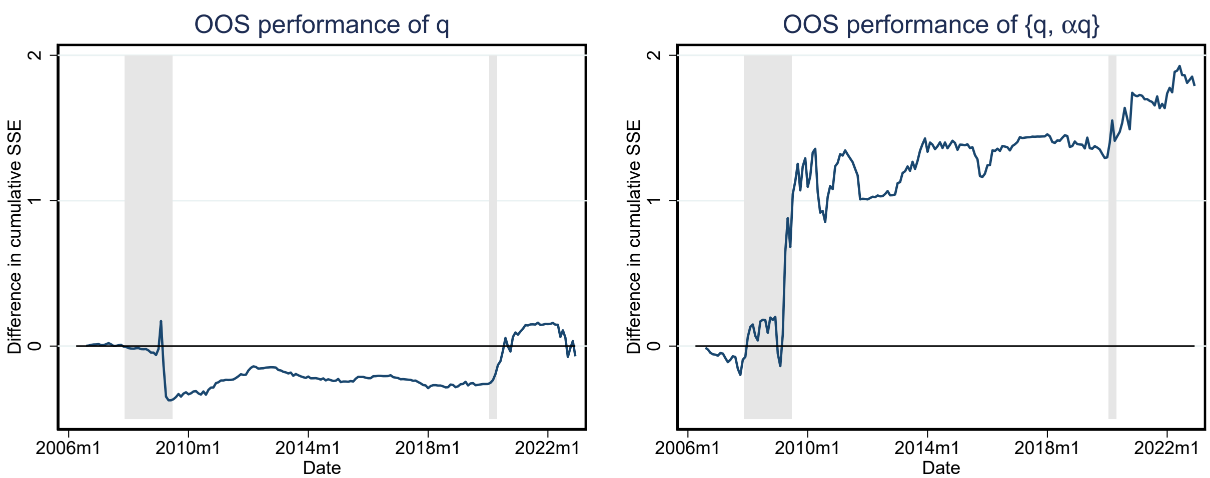

Notes: This figure displays the difference in the cumulative sum of squared errors, $\Delta CSSE_{OOS}$, between the one-month-ahead forecast of the log equity premium based on the historical average and the one-month-ahead forecast based on the conditional market volatility from a GARCH(1,1) model (from Section 2.6) in the left panel. The right panel displays the $\Delta CSSE_{OOS}$ between the forecast based on the historical average and the forecast based on the pair of predictors consisting of the conditional market volatility and the product of the conditional market volatility and the conditional market ambiguity attitude. The out-of-sample period spans from 2006:07 - 2022:12. Shaded periods are NBER recessions.

Figure 4. The Risk-Return Tradeoff Out-of-Sample with α and Simple GARCH Volatility

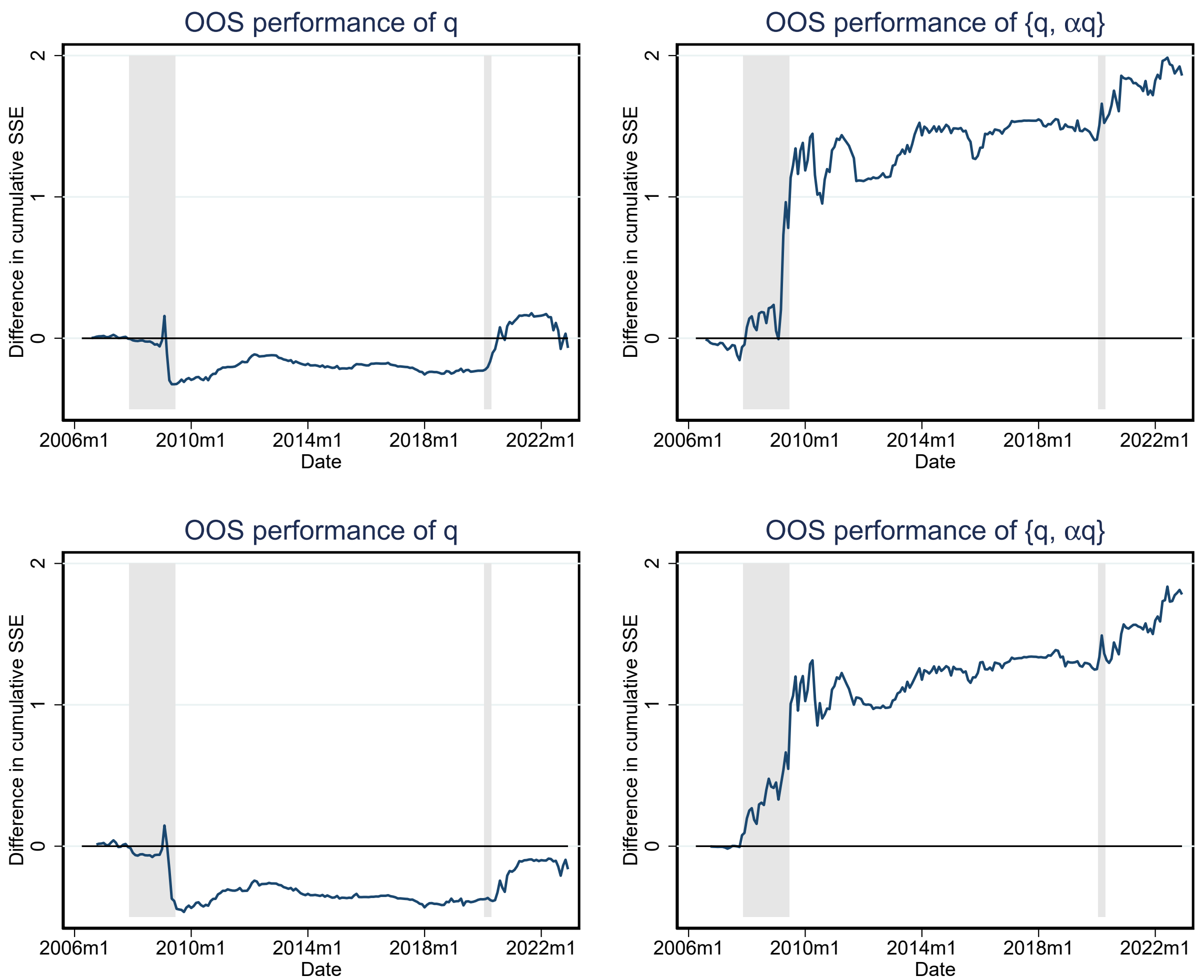

Notes: This figure displays the difference in the cumulative sum of squared errors, $\Delta CSSE_{OOS}$, between the forecast of the log equity premium based on the historical average and the forecast based on the conditional market volatility from a standard simple GARCH(1,1) model in the left panel. The right panel displays the $\Delta CSSE_{OOS}$ between the forecast based on the historical average and the forecast based on the pair of predictors consisting of the conditional market volatility and the product of the conditional market volatility and the conditional market ambiguity attitude. The top panel displays one-month-ahead forecasts (monthly forecast horizon). The bottom panel displays three-month-ahead forecasts (quarterly forecast horizon). The out-of-sample period spans from 2006:07 - 2022:12. Shaded periods are NBER recessions.

Figure 5. The Risk-Return Tradeoff Out-of-Sample with α and GJR GARCH Volatility

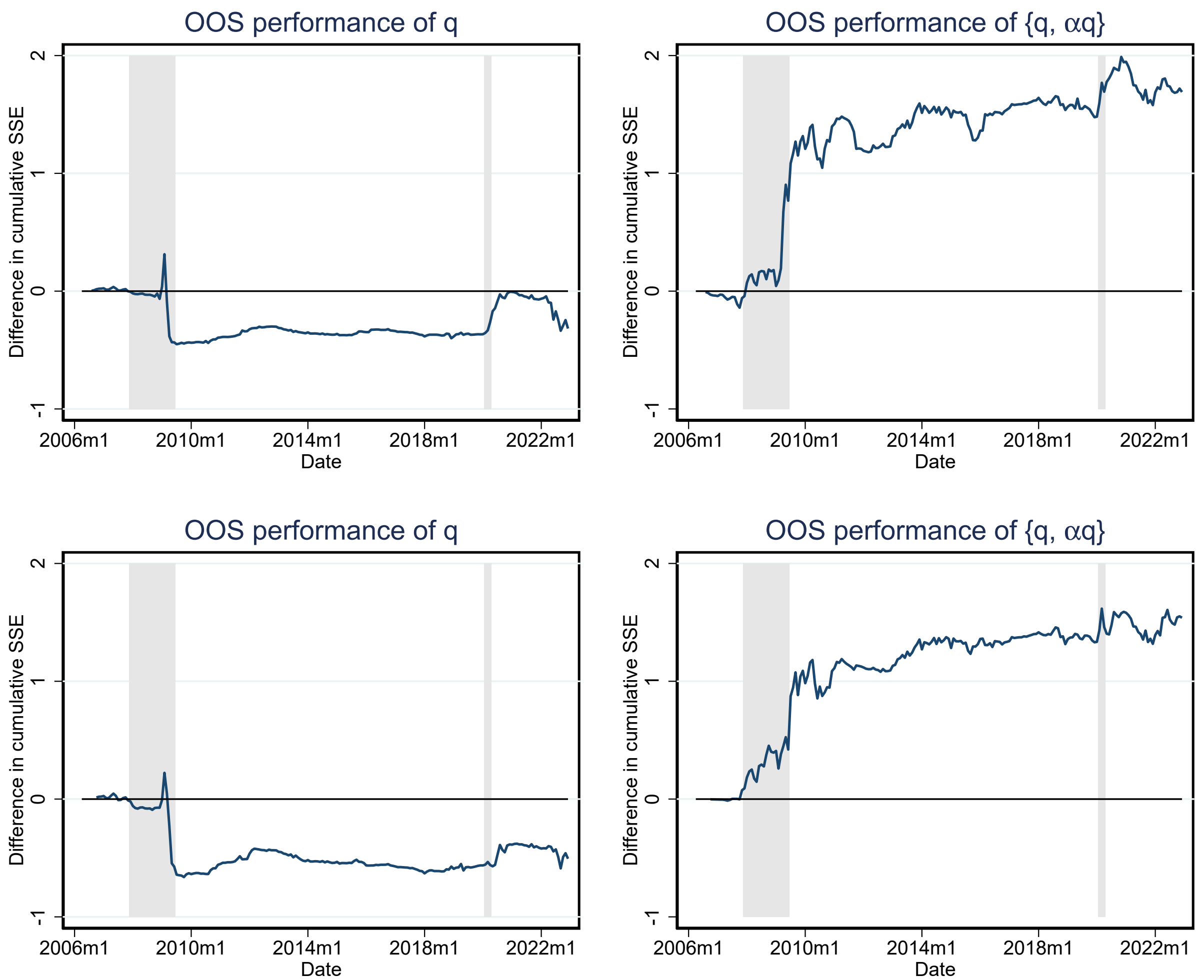

Notes: This figure displays the difference in the cumulative sum of squared errors, $\Delta CSSE_{OOS}$, between the forecast of the log equity premium based on the historical average and the forecast based on the conditional market volatility from a standard GJR GARCH model in the left panel. The right panel displays the $\Delta CSSE_{OOS}$ between the forecast based on the historical average and the forecast based on the pair of predictors consisting of the conditional market volatility and the product of the conditional market volatility and the conditional market ambiguity attitude. The top panel displays one-month-ahead forecasts (monthly forecast horizon). The bottom panel displays three-month-ahead forecasts (quarterly forecast horizon). The out-of-sample period spans from 2006:07 - 2022:12. Shaded periods are NBER recessions.