

Variable Selection for Portfolio Choice

YACINE AÏT-SAHALIA and MICHAEL W. BRANDT*

ABSTRACT

We study asset allocation when the conditional moments of returns are partly predictable. Rather than first model the return distribution and subsequently characterize the portfolio choice, we determine directly the dependence of the optimal portfolio weights on the predictive variables. We combine the predictors into a single index that best captures time variations in investment opportunities. This index helps investors determine which economic variables they should track and, more importantly, in what combination. We consider investors with both expected utility (mean variance and CRRA) and nonexpected utility (ambiguity aversion and prospect theory) objectives and characterize their market timing, horizon effects, and hedging demands.

THERE IS BY NOW AMPLE EVIDENCE in the literature that the means, variances, covariances, and higher order moments of stock and bond returns are time-varying and predictable. However, it has proven difficult to translate this evidence of predictability into practical portfolio advice because the different moments of returns, which in turn determine the optimal portfolio weights, are typically predicted by different sets of economic variables. Perhaps because of this difficulty with modeling the conditional return distribution, most professional investment advice is given solely on the basis of variables that forecast expected returns, such as the dividend yield or the slope of the term structure.¹

Looking beyond expected returns, it is difficult to decide which selection or combination of predictive variables the investor should focus on. (As a result, the empirical literature on portfolio choice has relied on a predeter-

* Ait-Sahalia is with the Department of Economics, Princeton University, and NBER. Brandt is with the Wharton School, University of Pennsylvania, and NBER. We thank John Campbell, George Constantinides, David Chapman, Frank Diebold, Han Hong, Paul Pfleiderer, Jessica Wachter, and especially Anthony Lynch for their comments and suggestions. We also thank seminar participants at the 2001 Annual Meeting of the American Finance Association, Columbia University, Harvard University, MIT, NYU, Princeton University, Stanford University, and the University of Chicago. This research was conducted during the first author's tenure as an Alfred P. Sloan Research Fellow. Financial support from the NSF under Grant SBR-9996023, the Bendheim Center for Finance at Princeton University, and the Rodney White Center at the Wharton School is gratefully acknowledged.

¹ It is quite natural, of course, to focus on the first moment of the return distribution when making a conditional portfolio choice. First, expected returns are the most intuitive and arguably the most important input to the investor's objective function; second, the dependence of the optimal portfolio choice on the first moment of the return distribution is monotonic for most preferences, unlike the dependence on higher order moments; and third, even with relatively simple preferences, the dependence of the optimal portfolio choice on the whole return distribution is so complex that it can usually only be solved numerically.

mined choice of one or at most two concurrent state variables. In a single-period context, Avramov (1999) uses the dividend yield, book-to-market ratio, earnings yield, Treasury bill yield, and term spread; and Kandel and Stambaugh (1996) use the dividend yield. In a multiperiod setting, Balduzzi and Lynch (1999), Barberis (2000), and Campbell and Viceira (1999) use the dividend yield; Brandt (1999) uses the dividend yield, default spread, term spread, and lagged return; Brennan, Schwartz, and Lagnado (1997) use the dividend yield, bond yield, and Treasury bill yield; Campbell and Viceira (1998) use the Treasury bill yield; Chacko and Viceira (1999) use the observed returns variance; and Lynch (2000) uses the dividend yield and term spread.) This is true even in the few special cases where we have an explicit asset allocation formula, such as for mean-variance utility where the optimal allocation is proportional to the ratio of the conditional mean to the conditional variance of returns. In this mean-variance case, it is clear that we want to find variables that best predict the *ratio* of the first two conditional moments. Choosing variables that best predict the mean and variance separately is likely to be counterproductive. Indeed, what should we do if a variable has a positive effect on both means (which the investor likes) and variances (which are detrimental to the investor)? What should we do if this variable is highly significant for one of the moments but less so for the other? How do we capture the relative importance that the investor's preferences place on the different moments? These questions all suggest that in a portfolio choice context, we should select variables to directly predict optimal portfolio weights, rather than *first* select variables to predict separate features of the return distribution and *then* explore later their implications for asset allocation.

It is also intuitively clear that different objective functions place different emphases on the various features of the conditional return distribution. For example, a mean-variance investor wants to predict means and variances, while a loss-averse investor may be more concerned about forecasting the size of the left tail of the return distribution.² Since, again, the means, variances, and size of the tails are not always predicted by the same variables, these two investors may choose different predictors for their conditional portfolio choice. Furthermore, investors may also disagree about the variable selection because, at the optimal choice, they are holding different portfolios of risky securities.³

² If returns are normally distributed, a loss-averse investor cares effectively about the ratio of the mean to the standard deviation of returns, which measures the size of the tail of a Gaussian density, rather than the ratio of the mean to the variance that a mean-variance investor cares about.

³ Consider two investors in the same class of preferences, one who is relatively risk averse and holds mostly bonds and another who is less risk averse and holds primarily stocks. Since different variables help predict the moments of bond and stock returns, these two investors may also choose different predictors for their conditional portfolio choice. Naturally, these effects are further compounded when we compare investors with different objective functions and different portfolio holdings, such as a mean-variance investor who holds stocks and a loss-averse investor who holds bonds.

In this paper, we show how to select and combine variables to best predict an investor's optimal portfolio weights, both in single-period and multi-period contexts. Rather than first model the various features of the conditional return distribution and subsequently characterize the portfolio choice, we focus directly on the dependence of the portfolio weights on the predictors. We do so by solving sample analogues of the conditional Euler equations that characterize the portfolio choice, as originally suggested by Brandt (1999). However, unlike the existing literature, we determine endogenously, for a given set of utility preferences, which of the candidate predictors are important for the optimal portfolio weights (rather than important for separate moments of the return distribution).

The advantage of focusing directly on the optimal portfolio weights is that we bypass the estimation of the conditional return distribution. This intermediate estimation step is the Achilles' heel of conditional portfolio choice because although the moments of returns are predictable, this predictability is for some moments quite tenuous. In particular, in the literature on predicting returns, an R^2 of 10 percent is hailed, rightly so, as a great success. Our approach is based on the hope that the relationship between the portfolio weights, which are complicated functions of the return distribution, and the predictors is less noisy than the relationship between the individual moments and the predictors. Even if it is not, we avoid introducing additional noise and potential misspecifications through the intermediate, but unnecessary, estimation of the return distribution.

We form a linear combination or index of the conditioning variables that best predicts the investor's optimal portfolio weights and then judge the importance of each individual variable by the role it plays in this index. We make no further assumptions about the relationship between the optimal portfolio weights and the predictors for two reasons. First, the dependence of the portfolio choice on the predictors can be highly nonlinear, even when the conditional moments are approximately linear; and second, the particular form of the nonlinearities not only varies greatly with the investor's preferences but also cannot generally be determined explicitly. This leads us to a semiparametric approach, where the optimal portfolio weights depend nonparametrically on a parametric index of the predictors.

We study investors with both expected utility (mean variance and constant relative risk aversion (CRRA)) and nonexpected utility (ambiguity aversion and prospect theory) objectives in order to see how the optimal index composition depends on the characteristics of the investor's preferences. From a normative perspective, our results can help investors with any one of these preferences determine which economic variables they should track and, more importantly, in what single combination. Our index is a parsimonious way to describe the current state of the investor's investment opportunities, just as in different economic contexts indices summarize high-dimensional state vectors (the index of leading economic indicators, the business cycle index, the consumer confidence index, etc.). Macroeconomic indices are country specific, since different countries have different characteristics, and analo-

gously our investment opportunities index is investor specific because different investors have different preferences.

For the purpose of giving portfolio advice, one advantage of our index approach is that it helps investors understand their conditional asset allocation in a more intuitive manner. For instance, it delivers simple rules like “if the index increases, the allocation to stocks should increase”. By contrast, it is generally difficult to represent graphically variables in more than two dimensions, let alone develop economic intuition about their interactions.

We characterize the market timing, horizon effects, and hedging demands of different types of investors, where we disentangle the effects due to the time-varying return distribution from those due to the specific preference structure. Specifically, we show how the portfolio choice of both expected and nonexpected utility investors varies as a function of the predictors, investment horizon, and rebalancing frequency. We explain also how the source of these variations differs across the preference specifications.

The remainder of the paper is organized as follows. We motivate the variable selection problem in Section I with individual moment regressions. Section II explains our econometric approach of predicting optimal portfolio weights with an index that captures the current state of investment opportunities. In Section III we discuss four parameterizations of the investor's preferences, which we then use in Section IV for our empirical work. There we characterize the optimal index composition and portfolio rules for different types of investors, different horizons, and different rebalancing frequencies. We conclude in Section V.

I. Predicting Individual Moments

A. Data

We collect monthly, quarterly, semiannual, and annual returns on the Standard and Poors (S&P) 500 index, an equal-weighted portfolio of non-callable government bonds with more than 10 years to maturity, and a maturity-matched Treasury bill from CRSP. The returns are sampled monthly from January 1954 through December 1997. The sample consists of 528 observations.

An ever-growing set of economic variables has been shown to partly predict the means, variances, and covariances of returns. (The following is a partial list of academic papers that document various degrees of mean predictability and the variables they use: Campbell (1987), term spread; Campbell and Shiller (1988a, 1988b), dividend yield; Cochrane (1991), investment-to-capital ratio; Fama and Schwert (1977), Treasury bill yield; Fama and French (1988, 1989), default spread, dividend yield, term spread; Ferson and Harvey (1991), default spread, dividend yield, lagged returns, term spread, Treasury bill yield; Keim and Stambaugh (1986), default spread, trend; Lamont (1998), dividend-to-earnings ratio; Lettau and Ludvigson (2000),

consumption-to-wealth ratio; and Pontiff and Shall (1998), book-to-market ratio. Studies on variance predictability include: Bollerslev (1986), lagged squared return, lagged variance; Campbell (1987), term spread; Engle (1982), lagged squared return; French, Schwert, and Stambaugh (1987), lagged squared return, lagged variance; Harvey (1991), default spread, dividend yield, lagged squared return, lagged variance, term spread, Treasury bill yield; Schwert (1989), debt-to-equity ratio, default spread, lagged variance, volume; and Whitelaw (1994), default spread, lagged variance, paper spread, Treasury bill yield. Finally, representative papers on predicting covariances are: Bollerslev, Engle, and Wooldridge (1988), lagged covariances, lagged cross-products of returns; Campbell (1987), term spread; and Harvey (1989), default spread, dividend yield, term spread.) We collect monthly data on four popular predictors: the default spread, the log dividend-to-price ratio of the S&P index, the term spread, and an S&P index trend (or momentum) variable. The default spread is the yield difference between Moody's Baa- and Aaa-rated corporate bonds. The dividend yield is the sum of dividends paid on the S&P index over the past 12 months divided by the current level of the index. The term spread is the yield difference between the 10- and 1-year government bonds. The trend is the difference between the log of the current S&P index level and the log of the average index level over the previous 12 months. Fama and French (1988, 1989) show that the first three predictors capture cyclical time variations in excess stock and bond returns, and Keim and Stambaugh (1986) use a variable very similar to our trend to predict returns. The data for the predictors is from the DRI/Citibase database.

Table I and Figure 1 describe the data. Panel A of Table I presents univariate descriptive statistics for the monthly returns, annual returns, and predictors. We omit the quarterly and semiannual returns to preserve space. Panel B shows pairwise correlations of the predictors with each other and with excess stock and bond returns, their squares, and their cross-products. Figure 1 plots the time-series and autocorrelations of the predictors.

B. Predictive Regressions

We first verify that the variables we identified as potential predictors indeed capture time variations in at least the first and second moments of excess bond and stock returns. For that purpose, we set up the following regressions:

$$\mathbf{E}_t \begin{bmatrix} r_{t+1}^b \\ r_{t+1}^s \end{bmatrix} = \begin{bmatrix} \mathbf{Z}'_t \gamma_b \\ \mathbf{Z}'_t \gamma_s \end{bmatrix} \quad \text{and} \quad \text{Var}_t \begin{bmatrix} r_{t+1}^b \\ r_{t+1}^s \end{bmatrix} = \begin{bmatrix} \mathbf{Z}'_t \delta_{bb} & \mathbf{Z}'_t \delta_{bs} \\ & \mathbf{Z}'_t \delta_{ss} \end{bmatrix}, \quad (1)$$

where r_{t+1}^b and r_{t+1}^s denote bond and stock returns in excess of the Treasury bill return and the vector \mathbf{Z}_t contains subsets of the four predictors. We demean and standardize the data to eliminate the intercepts and then estimate γ and δ using Hansen's (1982) generalized method of moments (GMM)

Table I
Description of Returns and Predictors

Panel A of this table shows annualized descriptive statistics of monthly and annual returns on the S&P 500 index, a portfolio of long-term government bonds, and a maturity-matched Treasury bill. The panel also shows descriptive statistics of four predictors: the default spread Def, the log dividend-to-price ratio of the S&P index LnDP, the term spread Term, and the S&P index momentum variable Trend. Panel B shows correlations of the predictors and their first differences with the predictors, with excess stock returns r^s , excess bond returns r^b , their squares, and cross-products. The data is sampled monthly from January 1954 through December 1997. There are 528 observations.

Panel A: Descriptive Statistics											
	Mean	Median	StdDev	Skew	Kurtosis	Min	Max	Autocorrelations			
								ρ_1	ρ_3	ρ_6	ρ_{12}
One-Month Horizon											
T-Bill	0.053	0.049	0.008	1.043	4.437	0.000	0.013	0.972	0.920	0.874	0.784
Bonds	0.069	0.040	0.085	0.784	6.796	-0.079	0.141	0.111	-0.085	0.035	0.007
Stocks	0.128	0.161	0.144	-0.403	5.338	-0.225	0.166	0.056	-0.002	-0.062	0.030
One-Year Horizon											
T-Bill	0.053	0.050	0.026	0.895	3.728	0.008	0.141	0.998	0.987	0.954	0.852
Bonds	0.069	0.050	0.095	1.060	4.793	-0.153	0.454	0.933	0.785	0.566	0.048
Stocks	0.126	0.143	0.147	-0.487	3.466	-0.505	0.532	0.921	0.736	0.422	-0.230
Predictors											
Def	0.961	0.810	0.442	1.332	4.689	0.320	2.690	0.973	0.910	0.832	0.687
LnDP	1.275	1.247	0.251	-0.157	3.335	0.476	1.844	0.990	0.958	0.904	0.792
Term	0.562	0.630	1.286	-0.185	3.539	-3.950	3.960	0.968	0.867	0.752	0.616
Trend	0.018	0.022	0.039	-0.680	3.837	-0.164	0.106	0.886	0.657	0.306	-0.116

Panel B: Correlations														
	Def	LnDP	Term	Trend	One-Month Horizon					One-Year Horizon				
					r^b	r^s	$(r^b)^2$	$(r^s)^2$	$r^b r^s$	r^b	r^s	$(r^b)^2$	$(r^s)^2$	$r^b r^s$
Def	1.00	0.54	-0.07	-0.08	0.14	0.08	0.37	0.15	0.23	0.20	0.11	0.38	0.17	0.27
LnDP		1.00	-0.35	-0.32	-0.00	-0.03	0.22	0.13	0.12	-0.07	0.16	0.28	0.17	-0.02
Term			1.00	0.26	0.14	0.18	-0.07	-0.09	-0.03	0.29	0.29	0.10	-0.10	0.12
Trend				1.00	-0.12	0.01	-0.01	-0.19	-0.14	-0.13	-0.12	-0.05	-0.24	-0.13
Δ Def	-0.11	0.11	-0.21	-0.22	0.12	-0.16	0.09	0.20	0.04	0.12	0.06	0.10	0.04	0.07
Δ LnDP	-0.13	-0.03	-0.18	0.01	-0.29	-0.53	-0.02	0.08	-0.10	-0.05	-0.36	-0.07	-0.15	-0.05
Δ Term	0.15	0.11	-0.13	-0.17	0.31	-0.04	0.14	0.09	0.12	0.15	0.12	-0.00	0.14	0.02
Δ Trend	0.12	0.06	0.09	-0.24	0.29	0.95	0.05	-0.05	0.25	0.03	0.31	0.05	0.19	0.06

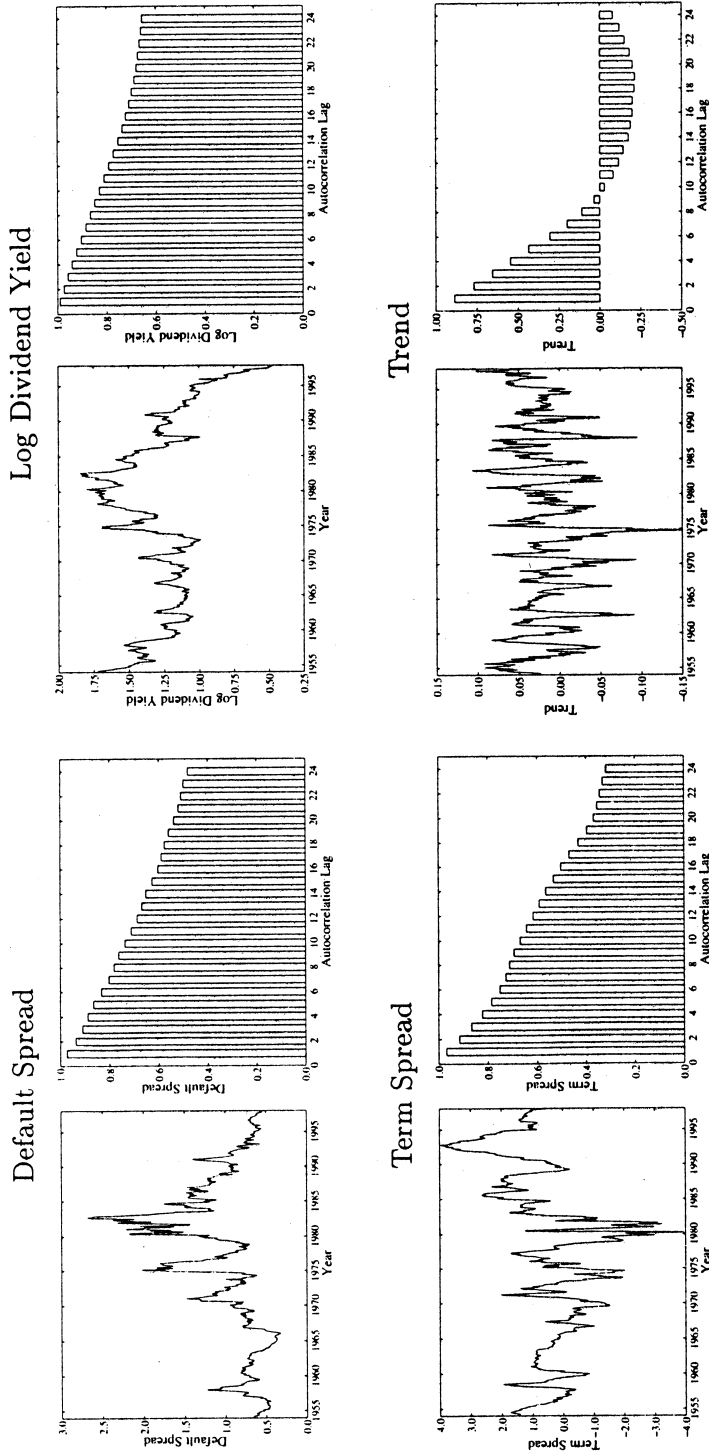


Figure 1. Predictors. This figure shows time-series plots and autocorrelation plots of four predictors: the default spread, the log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index trend variable. The data is sampled monthly from January 1954 through December 1997. There are 528 observations.

for all one-, two-, three-, and four-dimensional subsets of the four predictors. Whenever we use overlapping returns we compute autocorrelation-adjusted asymptotic standard errors using the procedure of Hodrick (1992).

Panel A of Table II presents the regression results. For each security (bonds and stocks), moment (mean, variance, and covariance), and return horizon (monthly and annual), the panel presents the best one-, two-, and three-variable partial regressions. It also shows the full regression with all four predictors. The best partial regressions are chosen according to the Akaike information criterion (AIC).⁴

The following facts emerge from Panel A of Table II:

- The default spread relates positively to the variances and covariance of both monthly and annual stock and bond returns. The regression coefficients are both statistically and economically significant (recall the data is demeaned and standardized, so the magnitude of the coefficients is meaningful), except for the stock return variance at the annual horizons. The default premium also relates positively, but not always significantly, to expected bond and stock returns.
- The log dividend-to-price ratio relates positively and significantly (at the 10 percent level) to expected stock returns at the annual horizon but not at the monthly horizon.⁵ The log dividend yield also relates positively to the variance of bond returns and negatively to the covariance between stock and bond returns at the annual horizon.
- The term spread is by far the most important predictor for expected returns. It relates positively to expected bond and stock returns at both horizons.⁶ The coefficients are statistically and economically significant, especially for stock returns. In addition, the term spread relates negatively to stock and bond variances.
- The trend variable relates negatively and significantly to expected bond returns at both horizons. It also relates negatively to the variance of stock returns, with regression coefficients that increase in magnitude and statistical significance with the horizon.
- Except for the covariance, the adjusted R^2 s of the regressions increase with the horizon, meaning that returns and squared returns are more predictable at long horizons than at short horizons. This pattern is due

⁴ Since the data is demeaned and standardized, the AIC criterion and the adjusted R^2 of the regressions are virtually identical. To save space, we only report the adjusted R^2 of the regressions in Panel A of Table II.

⁵ Because the dividend-to-price ratio does not appear to predict dividend growth, it must, on the basis of the present value formula, forecast returns. In fact, as pointed out by Cochrane (1999), “price divided by anything sensible” has this forecasting power. Since most price ratios are variables with very slow mean reversion (see Figure 1) they forecast long-term returns better than short-term returns.

⁶ Fama and French (1989) document that the slope of the yield curve moves in tandem with the business cycle. The yield curve is inverted at the peak of the cycle, where expected returns are low, and upward sloping when a recession turns into a recovery, where expected returns are high.

Table II
Individual Moment Predictability

Panel A of this table presents predictive regressions for expected excess bond returns $E_t[r_{t+1}^b]$, expected excess stock returns $E_t[r_{t+1}^s]$, the variance of excess bond returns $Var_t[r_{t+1}^b]$, the variance of excess stock returns $Var_t[r_{t+1}^s]$, and the covariance between excess bond and stock returns $Cov_t[r_{t+1}^b, r_{t+1}^s]$. The predictors are the default spread Def, the log dividend-to-price ratio of the S&P 500 index LnDP, the term spread Term, and the S&P index momentum variable Trend. The predictors are demeaned and standardized. The return horizon is one month or one year. For each moment, the table presents the best one-, two-, and three-variable partial regressions, selected based on the Akaike information criterion, and the full regression with all four predictors. Panel B summarizes the choice of the best one or two predictors for each moment, based on the Akaike information criterion, at the 1-, 3-, 6-, and 12-month horizon.

Panel A: Predictive Regressions									
Def	One-Month Horizon				R^2	One-Year Horizon			
	LnDP	Term	Trend	Def		LnDP	Term	Trend	R^2
$E_t[r_{t+1}^b]$									
		0.142*		0.02			0.286**		0.08
		0.185**	-0.166***	0.04	0.218		0.301**		0.13
0.137		0.192**	-0.158**	0.06	0.206		0.353**	-0.208**	0.17
0.189	-0.104	0.165**	-0.180**	0.07	0.317**	-0.217	0.296**	-0.254***	0.19
$E_t[r_{t+1}^s]$									
		0.177**		0.03			0.292**		0.09
0.092*		0.183***		0.04		0.305*	0.400***		0.16
0.090*		0.191***	-0.029	0.04		0.268*	0.423***	-0.142	0.18
0.110*	-0.040	0.180***	-0.038	0.04	0.022	0.281*	0.426***	-0.140	0.18
$Var_t[r_{t+1}^b]$									
0.370***				0.14	0.366***				0.13
0.366***		-0.057		0.14	0.276**	0.167*			0.15
0.368***		-0.068	0.045	0.14	0.248**	0.239*	0.159		0.17
0.359***	0.019	-0.063	0.050	0.14	0.245**	0.247*	0.155	0.025	0.17
$Var_t[r_{t+1}^s]$									
			-0.190*	0.04				-0.272**	0.07
0.120**			-0.181*	0.05			-0.097	-0.247**	0.08
0.117**		-0.066	-0.164*	0.05	0.146	-0.122		-0.305**	0.09
0.121*	-0.008	-0.068	-0.166*	0.05	0.160	-0.163	-0.127	-0.286**	0.10
$Cov_t[r_{t+1}^b, r_{t+1}^s]$									
0.220***				0.05	0.220**				0.05
0.212**			-0.110	0.06	0.327***	-0.198**			0.07
0.238***	-0.050		-0.124*	0.06	0.345***	-0.248***		-0.127	0.09
0.243***	-0.062	-0.032	-0.120*	0.06	0.357***	-0.280***	-0.080	-0.116	0.09
Panel B: Variable Selection									
	One-Month Horizon	Three-Month Horizon	Six-Month Horizon	One-Year Horizon					
$E_t[r_{t+1}^b]$	Term, Trend	Term, Trend	Term, Trend	Term, Trend					
$E_t[r_{t+1}^s]$	Term	LnDP, Term	LnDP, Term	LnDP, Term					
$Var_t[r_{t+1}^b]$	Def	Def	Def	Def					
$Var_t[r_{t+1}^s]$	Def, Trend	Def, Trend	Trend	Trend					
$Cov_t[r_{t+1}^b, r_{t+1}^s]$	Def	Def, Trend	Def, Trend	Def, LnDP					

***, **, and * denote statistical significance at the 1, 5, and 10 percent levels, respectively.

to the slowly mean-reverting nature of the predictors (see Figure 1) and is more pronounced for expected returns than for return variances. Stock and bond returns are about equally predictable at both horizons. Squared bond returns, however, are substantially more predictable than squared stock returns.

The most important finding for motivating our approach is the fact that if we had to restrict attention to only one or two predictors, the variable selection would depend on the conditional moments of returns that we most wanted to predict (bonds vs. stocks and first vs. second moments) as well as on the return horizon. This result is best illustrated in Panel B of Table II, which lists for all four return horizons the best one or two predictors for each moment. The problem with selecting variables in the portfolio choice context lies in the fact that the moments of returns (or functions of them) that we want to predict are endogenous to the investor's preferences. For example, an investor who is very risk averse and holds mostly bonds may want to focus on predicting the variance of bond returns with the default spread, while another investor who is less risk averse and holds mostly stocks may want to focus on predicting expected stock returns with the log dividend-to-price ratio and term spread. It is this endogeneity of the variable selection in the portfolio choice context that motivates our emphasis on predicting optimal portfolio weights, rather than individual moments.

II. Predicting Optimal Portfolio Weights

A. Investor's Problem

Consider a single-period investor who maximizes the conditional expectation of an objective function $v(W_{t+1})$ of next period's wealth W_{t+1} .⁷ The expectation is conditional on a vector of state variables Z_t . The maximization is over the portfolio weights α_t under the budget constraint $W_{t+1} = W_t(\alpha_t'R_{t+1})$, where R_{t+1} is a vector of gross returns on the securities the investor can buy and sell. Formally, the portfolio choice problem is:

$$\max_{\alpha_t} E[v(W_t(\alpha_t'R_{t+1}))|Z_t], \quad (2)$$

subject to the adding-up constraint $\alpha_t'\iota = 1$, where ι denotes a vector of ones. Although the portfolio weights sum to one, not all wealth must be invested in risky securities because one of the securities may be riskless. Furthermore, the portfolio choice may be subject to a set of constraints $a \leq c(\alpha_t) \leq b$, such as short-sale or borrowing constraints.

⁷ We describe our econometric approach for a single-period portfolio choice to keep the notation simple. However, our estimator extends readily to a multiperiod portfolio choice by replacing the single-period objective function and its derivative with a multiperiod "value function" and its derivative (see Brandt (1999)). In fact, we apply our approach to the multiperiod portfolio choice in Section IV.C.3.

The solution to the investor's problem is the mapping from the state vector Z_t to the portfolio weights α_t . Assuming that this mapping is time invariant, we denote it⁸:

$$\alpha_t \equiv \alpha(Z_t), \quad (3)$$

and refer to it as the investor's portfolio choice, policy, weight, or rule.

The relation between the portfolio policy and the predictability of individual moments of the returns R_{t+1} given the predictors Z_t obviously depends on the specification of the objective function $v(W_{t+1})$. To illustrate this point, consider an investor with standard mean-variance preferences. The investor's objective function is:

$$\mathbb{E}[v(W_{t+1})|Z_t] = \mathbb{E}[W_{t+1}|Z_t] - \frac{\gamma}{2} \text{Var}[W_{t+1}^2|Z_t], \quad (4)$$

with the coefficient of absolute risk aversion $\gamma \geq 0$. The investor's portfolio policy is:

$$\alpha_t = \Sigma_t^{-1} \iota \frac{\gamma W_t - \iota' \Sigma_t^{-1} \mu_t}{\gamma W_t \iota' \Sigma_t^{-1} \iota} + \frac{\Sigma_t^{-1} \mu_t}{\gamma W_t}, \quad (5)$$

where $\Sigma_t = \text{Var}[R_{t+1}|Z_t]$ and $\mu_t = \mathbb{E}[R_{t+1}|Z_t]$. This analytical expression for the portfolio weight illustrates two facts. First, even with this most simple preference specification, there is no straightforward link between predictability in first and second moments and the portfolio policy that allows the investor to identify which predictors are important for the portfolio choice. Second, even if the conditional moments are approximately linear in the state variables, the ratio form of the portfolio policy implies that it can be a highly nonlinear and even nonmonotonic function.

If the portfolio choice includes a risk-free rate, it simplifies to allocating a fraction α_t^{tgc} of wealth to the mean-variance efficient tangency portfolio, with excess return r_{t+1}^{tgc} , and to invest the remainder in the risk-free asset, where:

$$\alpha_t^{tgc} = \frac{1}{\gamma W_t} \frac{\mathbb{E}[r_{t+1}^{tgc}|Z_t]}{\text{Var}[r_{t+1}^{tgc}]}. \quad (6)$$

In other words, the mean-variance portfolio choice is proportional to the conditional mean-variance ratio of the tangency portfolio. Therefore, in selecting variables for the mean-variance portfolio choice, which is a normative issue, we equivalently select variables for predicting the mean-variance ratio of the tangency portfolio, which is a generic descriptive statistic of the conditional return distribution.

⁸ With time-invariant objective function $v(W_{t+1})$, this assumption only requires that the conditional distribution of the returns R_{t+1} given the predictors Z_t is time homogenous.

B. Indices for the Conditional Portfolio Choice

Brandt (1999) shows how to estimate the optimal portfolio policy $\alpha(Z_t)$ without further assumptions about the return dynamics or functional form of the decision rule by replacing the conditional expectation in the investor's problem (2) with a consistent estimator. For a given realization of the state vector Z_t , we define the fully nonparametric estimator $\hat{\alpha}(Z_t)$ of the true $\alpha(Z_t)$ as the portfolio weights that solve the investor's problem when the conditional expectation $E[\cdot|Z_t]$ is replaced with a consistent estimator $\hat{E}[\cdot|Z_t]$, such that

$$\hat{E}[v(W_t(\alpha'_t R_{t+1}))|Z_t] \rightarrow E[v(W_t(\alpha'_t R_{t+1}))|Z_t] \quad \text{as } T \rightarrow \infty, \quad (7)$$

for all portfolio weights α_t and state vector realizations Z_t . In particular, Brandt (1999) suggests estimating the conditional expectation with a standard nonparametric regression.

Unfortunately, this fully nonparametric approach does not allow us to address the issue of which predictors are important for the portfolio choice because nonparametric estimators typically cannot handle a large number of regressors. As the number of predictors increases, the convergence rate of most nonparametric estimators to their asymptotic distribution deteriorates exponentially. This feature of the estimators is commonly referred to as the "curse of dimensionality". Realistically, given our relatively short sample, we cannot reliably estimate the conditional expectation for more than two predictors.

To overcome this econometric problem, we adopt a semiparametric approach that explicitly recognizes the endogeneity of the variable selection. We assume that the investor's optimal portfolio weights depend on the predictors Z_t only through a single linear index or factor $Z'_t \beta$ with unknown parameters β .⁹ The dependence of the portfolio weights on this index, however, is left completely unrestricted. We classify our approach as semiparametric because the index is parametric but the portfolio policy is not.

Formally, we rewrite the investor's problem (2) as:

$$\max_{\alpha_t} E[v(W_t(\alpha'_t R_{t+1}))|Z'_t \beta], \quad (8)$$

which implies that the optimal α_t depend on Z_t only through the index $Z'_t \beta$.¹⁰

$$\alpha_t \equiv \alpha(Z'_t \beta; \beta). \quad (9)$$

⁹ We can relax the assumption of a linear index by introducing nonlinearities in Z_t . Our experimentations with nonlinear indices suggest, however, that in the portfolio choice context, the incremental expected utility loss from combining the predictors linearly, as opposed to nonlinearly, is minimal.

¹⁰ The optimal portfolio choice α_t depends on the index coefficients β not only through the index realization $Z'_t \beta$, but also through the functional form of the policy function $\alpha(\cdot)$. For example, consider two indices β and $\bar{\beta} = -\beta$. In this case, $\alpha(x; \beta) \neq \alpha(x; \bar{\beta})$, but instead $\alpha(x; \beta) = \alpha(-x; \bar{\beta})$ for all index realizations x .

From a statistical perspective, the index avoids the curse of dimensionality because it allows us to reduce the multivariate problem to one where we can implement the nonparametric approach described above in a univariate setting (since $Z_t'\beta$ is univariate). From an economic standpoint, the index offers a convenient univariate summary statistic that describes the current state of the time-varying investment opportunities.

The economic cost of collapsing the multidimensional information contained in Z_t into a linear index $Z_t'\beta$ is that the expected utility from the unconstrained problem (2) exceeds that from the constrained problem (8), unless our assumption about the index structure of the investor's problem is true (as opposed to just an approximation for econometric purposes). The magnitude of the expected utility loss due to the index depends on the application and can ultimately only be measured empirically. For our application, we present evidence in Section IV.C.1 that the expected utility loss is minor.

We estimate the index coefficients β through the conditional Euler equations of the investor's unrestricted problem. Specifically, we substitute the parametric restriction (9) into the first-order conditions of the problem (2) to obtain the following set of conditional moment conditions:

$$\mathbb{E}[m_{t+1}(\beta)|Z_t] \equiv \mathbb{E}[v'(W_t(\alpha(Z_t'\beta;\beta)'R_{t+1}))R_{t+1}|Z_t] = 0, \quad (10)$$

where $\alpha(Z_t'\beta;\beta)$ solves the investor's restricted problem (8). Multiplying these moment conditions by predetermined functions $g(Z_t)$ of the forecasting variables, taking unconditional expectations, and then applying the law of iterated expectations yields a standard GMM inference problem (see Hansen (1982)):¹¹

$$\min_{\beta} \mathbb{E}[m_{t+1}(\beta) \otimes g(Z_t)]' W \mathbb{E}[m_{t+1}(\beta) \otimes g(Z_t)] \quad (11)$$

with optimal weighting matrix $W = \text{Cov}[m_{t+1} \otimes g(Z_t)]^{-1}$. The final step in the construction of our estimator is to replace both the unconditional expectations and the optimal weighting matrix with consistent sample analogues.

The main difference between our estimator and standard GMM is that the portfolio weights in the conditional moments (10) are only defined implicitly through the investor's portfolio optimization. To evaluate the GMM criterion

¹¹ Notice that we need the instruments $g(Z_t)$ to identify β . The index coefficients are unconditionally not identified because $\alpha(Z_t'\beta;\beta)$ satisfies $\mathbb{E}[m_{t+1}(\beta)|Z_t'\beta] = 0$ and hence $\mathbb{E}[m_{t+1}(\beta)] = 0$ for any β .

This setup can be used not only for estimation, but also for testing whether a second set of predictors should be included in the index. With the instruments $g(Z_t, Y_t)$, the minimized GMM objective is an asymptotically χ^2 -distributed test for the hypothesis that the index has zero loadings on the predictors Y_t .

for a candidate β , we first estimate the sequence of optimal portfolio weights $\{\alpha(Z_t' \beta; \beta)\}_{t=1}^T$ by replacing the conditional expectation in the investor's problem (8) with nonparametric regressions, as in expression (7).

The estimator and its asymptotic distribution are described in more detail in a technical appendix that is available on request. The main results are that the estimator is consistent, asymptotically normal, and, although it has a nonparametric component, achieves the parametric convergence rate of \sqrt{T} irrespective of the number of predictors.¹²

The GMM estimator (11) treats the choice of β as an inference problem under the null that the unconstrained portfolio policy $\alpha(Z_t)$ has the index form $\alpha(Z_t' \beta; \beta)$. Under the alternative that the index form is suboptimal, a more natural way to choose β is through the unconditional utility maximization:

$$\max_{\beta} \mathbf{E} \left[\max_{\alpha_t} \mathbf{E}[v(W_t(\alpha_t' R_{t+1})) | Z_t' \beta] \right] = \max_{\beta} \mathbf{E}[v(W_t(\alpha(Z_t' \beta; \beta)' R_{t+1}))], \quad (12)$$

where the equality follows from the law of iterated expectations and the restriction (9). In words, the index defined by this maximization generates a sequence of conditional portfolio choices that is unconditionally optimal or, equivalently, that minimizes the unconditional expected utility loss from solving the constrained problem (8) as opposed to the unconstrained problem (2).

The expected utility maximization in equation (12) is nested by the GMM estimator (11) through an *optimal* (in an expected utility not statistical sense) set of instruments $g(Z_t)$. Specifically, the first-order conditions of the maximization:

$$\mathbf{E} \left[v^1(W_t(\alpha(Z_t' \beta; \beta)' R_{t+1})) R_{t+1} \frac{\partial \alpha(Z_t' \beta; \beta)}{\partial \beta} \right] = 0 \quad (13)$$

show that the maximization problem (12) is equivalent to the minimization problem (11) with instruments:

$$g(Z_t) = \frac{\partial \alpha(Z_t' \beta; \beta)}{\partial \beta} \quad (14)$$

and an arbitrary weighting matrix W (since the parameters are exactly identified).

¹² Because the index composition is estimated and the investor's portfolio choice depends on the estimated index, our subsequent estimates of the portfolio weights automatically incorporate the fact that ex ante predictability is uncertain. The standard errors of the estimated indices and portfolio policies are larger in a small sample but not asymptotically, since in sufficiently large samples, there is no uncertainty about predictability. Of course, the appropriate approach to fully address the issue of parameter uncertainty is a Bayesian framework (see Barberis (2000) and Kandel and Stambaugh (1996)).

However, just because these instruments are optimal in theory, they do not necessarily result in more reliable estimates of the index coefficients in practice for two reasons. First, to construct the instruments we need consistent estimates of the derivatives $\partial\alpha(Z'_t\beta;\beta)/\partial\beta$. A nonparametric estimator of these derivatives converges at a slower rate than a nonparametric estimator of the function $\alpha(Z'_t\beta;\beta)$ (see Härdle (1990)) and thus may introduce substantial noise. Second, since the functional form of the portfolio policy may depend on the index composition in a highly nonlinear and irregular way (even if the policy is well-behaved in the index for a given index composition), the derivatives $\partial\alpha(Z'_t\beta;\beta)/\partial\beta$ may cause the GMM objective function to be less well behaved. This makes the numerical minimization more difficult and increases the risk of ending up with a local instead of global minimum.

Given estimates of the optimal index composition, we judge the importance of each predictor in the conditional portfolio choice through the relative weight the predictor receives in the index. In other words, the relative weights the estimated index coefficients $\hat{\beta}$ place on the different variables, and their statistical significance, tell us which variables, and more importantly in what combination, are relevant for the investor's conditional portfolio choice.

III. Objective Functions

To see how the variable selection varies across investors with different preferences, we consider four parameterizations of the objective function $v(W_{t+1})$. The first two, mean-variance and constant relative risk aversion (CRRA) preferences, are standard expected utility objectives and result in fairly similar indices for the conditional portfolio choice. The second two, ambiguity aversion and prospect theory preferences, are generalized or non-expected utility objectives. They produce indices that are quite different from those of expected utility investors, demonstrating the endogeneity of the variable selection problem in the portfolio choice context. In effect, different investors focus on different aspects of the returns distribution, which different variables help predict.

A. Expected Utility

A.1. Mean-Variance Preferences

We already introduced the objective function of an investor with mean-variance preferences in equation (4), where γ measures the investor's absolute risk aversion $(\partial^2v(W)/\partial W^2)/(\partial v(W)/\partial W)$.¹³ An appealing feature of mean-variance preferences, and the reason we consider them here, is that the

¹³ For the empirical work we normalize $W_t = 1$, so relative risk aversion $[W_t(\partial^2v(W)/\partial W^2)]/\partial v(W)/\partial W$ also equals γ .

optimal portfolio weights depend exclusively and analytically on the first two moments of returns (see equation (5)). Thus, we can directly compare our semiparametric estimates of the portfolio policy to parametric estimates based on individual moment forecasts.

A.2. Constant Relative Risk Aversion Preferences

We also consider an investor with CRRA or power utility:

$$v(W_{t+1}) = \begin{cases} \frac{W_{t+1}^{1-\gamma}}{1-\gamma} & \text{if } \gamma > 1 \\ \ln W_{t+1} & \text{if } \gamma = 1 \end{cases}, \quad (15)$$

where γ now measures relative risk aversion $[W_t(\partial^2 v(W)/\partial W^2)]/(\partial v(W)/\partial W)$. CRRA preferences are by far the most popular objective function in the portfolio choice literature. This is largely because the investor's portfolio (and consumption) policy is proportional to wealth and the value function is homothetic in wealth. In a multiperiod setting, these features of CRRA preferences imply that wealth is not a state variable in the investor's problem.

B. Nonexpected Utility

B.1. Ambiguity Aversion

Expected utility theory assumes that the investor can compute expectations with respect to the return distribution, which requires that the agent knows the parametric structure of the return distribution and either knows its parameters or can form Bayesian beliefs about them. The investor is only exposed to the "risk" inherent in the returns and trades off this risk against expected rewards through the expected utility maximization. Knight (1921) and Ellsberg (1961) argue, however, that the investor may not have all of the information required to form such expectations. For example, an agent may not be able or willing to assign probabilities to a set of alternative parameterizations of the return distribution. Thus, the investor faces additional "ambiguity" that is not captured in the expected utility framework. Ambiguity aversion preferences formalize the idea that the investor dislikes not only risk but also this more vague uncertainty about the world (called Knightian uncertainty). There is an extensive experimental literature confirming that individuals indeed dislike ambiguity, both in gambling (e.g., Becker and Brownson (1964) or Curley and Yates (1985, 1989)) and in financial markets (e.g., Camerer and Kunreuther (1989) or Sarin and Weber (1993)).

Consider again an investor with CRRA preferences, except that now the agent is uncertain about whether the return distribution is \bar{p} (the empirical distribution, for example) or some other distribution $p \in P$ in the neighborhood of \bar{p} . The crucial difference between ambiguity aversion and expected

utility theory with model uncertainty is that with ambiguity aversion the investor cannot or does not want to assign probabilities to the set of alternative return distributions. Following Gilboa and Schmeidler (1989) and Dow and Werlang (1992), the investor's portfolio choice problem in this case is given by:

$$\max_{\alpha_t} \min_{p \in P} E_p[v(W_t(\alpha'_t R_{t+1})) | Z_t], \quad (16)$$

where $v(\cdot)$ is the CRRA utility function in equation (15). This maxmin criterion is quite intuitive. Given the complete ambiguity about the return distribution, the investor considers the worst case outcome (in the neighborhood of \bar{p}) through the interior minimization. The exterior maximization then achieves the usual risk versus expected reward trade-off.

To implement this form of ambiguity aversion, we need to characterize the set of possible return distributions P . We adopt the following ε -contamination parameterization:

$$P = \{(1 - \varepsilon)\bar{p} + \varepsilon p : p \in \mathcal{P}\}, \quad (17)$$

where \mathcal{P} denotes the σ -algebra generated by the support of the return distribution. The assumption of ε -contamination is informally introduced by Ellsberg (1961) and is now common in the literature on ambiguity aversion. It is used, for example, by Epstein and Wang (1994) and Liu (1998, 1999). The parameter ε reflects the investor's degree of ambiguity.¹⁴ With $\varepsilon = 0$, the investor's objective function reduces to that with standard CRRA preferences and return distribution \bar{p} .

The advantage of this parameterization is that the investor's problem simplifies to:

$$\max_{\alpha_t} (1 - \varepsilon) E_{\bar{p}}[v(W_t(\alpha'_t R_{t+1})) | Z_t] + \varepsilon \inf_{\mathcal{P}} v(W_t(\alpha'_t R_{t+1})), \quad (18)$$

where we assume that the support of the return distribution is independent of the predictors. Thus, to implement the notion of ambiguity aversion we only need to choose a value for the parameter ε and specify the support of the return distribution to evaluate the infimum.

Ambiguity aversion relates to the recent literature on robustness (e.g., Anderson, Hansen, and Sargent (1999), Hansen, Sargent, and Tallarini (1999), and Maenhout (1999)). Although the formalizations differ slightly, the behavioral motivation of the two theories is the same. Agents are uncertain

¹⁴ Alternatively, one can interpret the portfolio choice as the investor playing a two-stage game against nature. In the first stage, nature replaces with probability ε the return distribution \bar{p} with an arbitrary distribution $p \in \mathcal{P}$. In the second stage, nature then draws a set of returns from the return distribution.

about the true model and are unable or unwilling to assign probabilities to the set of alternative models. Not too surprisingly, robustness also results in maxmin policies. We focus on the ε -contamination version of ambiguity aversion because it captures the essence of the theory and is more tractable for our empirical application.

B.2. Prospect Theory and Loss Aversion

In another literature on decisions under uncertainty, Kahneman and Tversky (1979) argue that humans systematically violate the axioms of expected utility theory in two important ways. First, experimental subjects tend to overweight outcomes that are considered certain, relative to outcomes that are merely probable, which is referred to as the “certainty effect”. In financial markets, this certainty effect makes an investor risk averse in the case of gains, as a small certain gain is preferred to a probable risky gain, but risk seeking in the case of losses, as a probable risky loss is preferred to a small certain loss. In addition, subjects tend to simplify decisions by disregarding components common to the alternative choices and focusing on components that differentiate the choices, which is called the “isolation effect”. Since the decomposition of alternative choices into common and differentiating components is nonunique, however, the outcome of the decision problem depends on the investor’s perspective in the simplification process.

Motivated by this experimental evidence, Kahneman and Tversky (1979) formulate prospect theory, which consists of an editing stage, where alternatives are put into perspective, and a choice stage. (Finance applications of prospect theory preferences include Shefrin and Statman (1994), Benartzi and Thaler (1995), Shumway (1997), Barberis, Huang, and Santos (2001), and Barberis and Huang (2001).) Utility is defined over gains and losses relative to a reference point (such as the result of an all-cash investment strategy or last period’s wealth) rather than over the level of wealth as in expected utility theory. To capture the differential risk preferences over gains and losses generated by the certainty effect, Tversky and Kahneman (1992) propose the following objective function for the choice stage:

$$v(W_{t+1}) = \begin{cases} -l(\bar{W} - W_{t+1})^b & \text{if } W_{t+1} < \bar{W} \\ (W_{t+1} - \bar{W})^b & \text{otherwise} \end{cases}, \quad (19)$$

where \bar{W} is a reference wealth level determined in the editing stage. For example, \bar{W} could be the initial wealth W_t or its future value $R_t^{tb} W_t$, depending on the investor’s perspective. The parameter l measures the investor’s loss aversion and the parameter b captures the degree of risk seeking over losses and risk aversion over gains.¹⁵ The kink at the origin introduced by $l > 1$ makes losses (relatively) more painful than gains are pleasurable.

¹⁵ Tversky and Kahneman cite experimental evidence that suggests $b = 0.88$ and $l = 2.25$.

In addition, in Tversky and Kahneman's (1992) formalization of prospect theory the investor does not evaluate outcomes on the basis of true probabilities, but rather, as predicted by the certainty effect, on the basis of distorted probabilities. That is, instead of maximizing the true expectation of the objective function, the investor maximizes:

$$\mathbb{E} \left[v(W_{t+1}) \frac{\pi(p(W_{t+1}|Z_t))}{p(W_{t+1}|Z_t)} \middle| Z_t \right] = \int_{-\infty}^{+\infty} v(W_{t+1}) \pi(p(W_{t+1}|Z_t)) dW_{t+1},$$

where $\pi(\cdot)$ represents a subjective distortion of the objective probabilities $p(\cdot)$. Tversky and Kahneman (1992) suggest parameterizing this probability distortion as (see also Tversky and Wakker (1995)):

$$\pi(p(W_{t+1}|Z_t)) = \frac{p(W_{t+1}|Z_t)^c}{\{p(W_{t+1}|Z_t)^c + (1 - p(W_{t+1}|Z_t))^c\}^{1/c}}, \quad (20)$$

where c determines the degree of "irrationality". When $c = 1$, the decision weights $\pi(\cdot)$ reduce to the objective probabilities $p(\cdot)$. Notice also that when $0 < c < 1$, the weights are not a proper probability measure (hence they are called weights, not subjective probabilities) because they sum to less than one.

A special case of prospect theory is loss aversion, when $b = 1$, $c = 1$, and $l > 1$. In this case, the investor is risk neutral over gains and is risk neutral over losses, but realizes a greater incremental utility penalty for a loss than for an equally large gain. This results in unconditional risk aversion. Furthermore, since with $c = 1$ the decision weights reduce to the objective probabilities, this investor simply maximizes expected utility. Interestingly, Benartzi and Thaler (1995) find that the main aspect of prospect theory relevant for portfolio choice is loss aversion and that the concavity (convexity) of the value function on the upside (downside), as well as the subjectivity of the probability distortions, are only of second-order importance. Sharpe (1998) argues, however, that the local risk-neutrality property of loss aversion results in portfolio choices that are too extreme.¹⁶

A literature that is somewhat related to loss aversion involves Value-at-Risk (VaR) constraints that ensure that with probability of at least q , the investor's wealth next period (or some other target horizon) does not fall below some specified level.¹⁷ The extremes $q = 0$ and $q = 1$ correspond to an unconstrained investor and a portfolio insurer (e.g., Basak (1995) and Gross-

¹⁶ The problem with loss aversion is that the iso-expected utility curves are straight lines in mean versus standard deviation of returns space, which means that in a stylized portfolio choice between a single stock and cash the optimal portfolio choice is either 100 percent cash or 100 percent stock, depending on whether the iso-expected utility curves are more or less steep than the mean-variance frontier (see Sharpe (1998)).

¹⁷ VaR is the de facto standard measure of risk because of its simplicity and its popularity with regulators.

man and Zhou (1996)). Unfortunately, VaR preferences have two faults. First, Artzner et al. (1998) show that VaR measures have difficulties aggregating individual risks, even risks that are cross-sectionally independent, and sometimes discourage diversification. Second, Basak and Shapiro (1998) find that in a multiperiod setting, a VaR-constrained investor frequently chooses, quite paradoxically, a larger exposure to risky assets than an otherwise equivalent unconstrained investor.¹⁸ As a fix to this problem, they propose an alternative risk management constraint that incorporates the expected value of the loss. Similar in spirit to this extended VaR constraint, loss aversion preferences penalize for both the probability and magnitude of losses.

IV. Empirical Results

A. Unconditional Portfolio Choice

We begin our empirical work by characterizing the unconditional portfolio choice of investors with expected and nonexpected utility preferences. The unconditional portfolio choices are useful for understanding the optimal index compositions in Section IV.B and serve as a benchmark for the conditional asset allocations in Section IV.C.

A.1. Expected Utility Preferences

A.1.1. Mean-variance Investors. Panel A of Table III presents estimates of the unconditional portfolio choice of investors with mean-variances preferences and absolute risk aversion (and relative risk aversion since $W_t = 1$) of γ equal to 2, 5, 10, and 20.¹⁹ The investment horizon is one month or one year. The entries in the Without Risk-free Rate section of the panel are for a portfolio choice between stocks and bonds. The entries in the With Risk-free Rate section are for a portfolio choice between stocks, bonds, and Treasury bills. We assume that Treasury bills are risk free and fix the Treasury bill rate at its historical average. We impose the short-sale constraints $0 \leq x \leq 1$ to prohibit unrealistic leveraging and short selling. In brackets below each estimate are autocorrelation adjusted asymptotic standard errors.²⁰

Figure 2 helps visualize the mean-variance portfolio choice. The two graphs plot the expected return on wealth against the standard deviation of wealth for the estimated portfolio weights in Panel A of Table III. The two lines in

¹⁸ The intuition is that the VaR-constrained investor finds it optimal to insure against losses in states where insurance is relatively cheap (because losses are relatively small), but accepts the possibility of losses (up to probability $1 - q$) in states where insurance is expensive (because losses are potentially large).

¹⁹ We checked that our method of moments estimates are virtually identical to the results of plugging the unconditional moments from Table I into the analytic expression (5) for the optimal portfolio weights.

²⁰ Whenever the short-sale constraints are binding, we compute the asymptotic standard errors using the results of Moran (1971) and Andrews (1999), who derive the asymptotics of an extremum estimator of a parameter that is located at the boundary of a closed parameter space.

Table III
Unconditional Portfolio Choice with Expected Utility Preferences

This table shows estimates of the unconditional portfolio choice of investors with single-period objectives:

$$\text{Panel A: } \max_{\alpha} \left(E[W_{t+1}] - \frac{\gamma}{2} \text{Var}[W_{t+1}] \right) \quad \text{Panel B: } \max_{\alpha} E \left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma} \right],$$

where W_{t+1} is next period's wealth and $\alpha = [\alpha^s, \alpha^b, \alpha^{tb}]$ are the fractions of current wealth $W_t = 1$ invested in stocks, bonds, and Treasury bills, respectively. The investment horizon is one month or one year. The Without Risk-Free Rate estimates are for a portfolio choice between stocks and bonds (i.e., $\alpha^{tb} = 0$) and the With Risk-Free Rate estimates are for a portfolio choice between stocks, bonds, and Treasury bills. In both cases, the optimization is subject to the short-sales constraints $0 \leq \alpha \leq 1$. In brackets are autocorrelation adjusted asymptotic standard errors.

	Without Risk-Free Rate		With Risk-Free Rate				Without Risk-Free Rate		With Risk-Free Rate			
	α^s	α^b	α^s	α^b	α^{tb}	$\frac{\alpha^s}{\alpha^s + \alpha^b}$	α^s	α^b	α^s	α^b	α^{tb}	$\frac{\alpha^s}{\alpha^s + \alpha^b}$
	One-Month Horizon						One-Year Horizon					
Panel A: Mean-Variance Investors												
$\gamma = 2$	1.00 [0.45]	0.00 [0.45]	1.00 [0.73]	0.00 [0.74]	0.00 [0.37]	1.00 [0.73]	1.00 [0.36]	0.00 [0.36]	1.00 [0.55]	0.00 [0.44]	0.00 [0.31]	1.00 [0.44]
$\gamma = 5$	0.73 [0.23]	0.27 [0.23]	0.70 [0.37]	0.11 [0.38]	0.19 [0.24]	0.86 [0.42]	0.68 [0.22]	0.32 [0.22]	0.62 [0.31]	0.07 [0.24]	0.31 [0.24]	0.90 [0.32]
$\gamma = 10$	0.45 [0.12]	0.55 [0.12]	0.35 [0.19]	0.06 [0.19]	0.60 [0.12]	0.86 [0.42]	0.43 [0.14]	0.57 [0.14]	0.31 [0.15]	0.03 [0.12]	0.66 [0.12]	0.90 [0.32]
$\gamma = 20$	0.31 [0.06]	0.69 [0.06]	0.17 [0.09]	0.03 [0.10]	0.80 [0.06]	0.86 [0.42]	0.31 [0.10]	0.69 [0.10]	0.15 [0.08]	0.02 [0.06]	0.83 [0.06]	0.90 [0.32]
Panel B: Constant Relative Risk-Averse Investors												
$\gamma = 2$	1.00 [0.40]	0.00 [0.40]	1.00 [0.75]	0.00 [0.74]	0.00 [0.43]	1.00 [0.74]	1.00 [0.30]	0.00 [0.30]	1.00 [0.57]	0.00 [0.64]	0.00 [0.20]	1.00 [0.61]
$\gamma = 5$	0.69 [0.21]	0.31 [0.21]	0.66 [0.38]	0.16 [0.38]	0.18 [0.22]	0.81 [0.39]	0.55 [0.19]	0.45 [0.19]	0.54 [0.32]	0.08 [0.32]	0.37 [0.19]	0.87 [0.45]
$\gamma = 10$	0.43 [0.11]	0.57 [0.11]	0.33 [0.19]	0.08 [0.19]	0.59 [0.11]	0.81 [0.40]	0.31 [0.11]	0.69 [0.11]	0.28 [0.16]	0.04 [0.16]	0.68 [0.10]	0.87 [0.43]
$\gamma = 20$	0.30 [0.06]	0.70 [0.06]	0.17 [0.09]	0.04 [0.09]	0.79 [0.05]	0.81 [0.40]	0.22 [0.06]	0.78 [0.06]	0.14 [0.08]	0.02 [0.08]	0.84 [0.05]	0.87 [0.43]




Figure 2. Unconditional portfolio choice with mean-variance preferences. This figure plots the expected return on wealth against the standard deviation of wealth for the unconditional portfolio choice of investors with mean-variance preferences. The investment horizon is one month or one year. The straight line and hyperbola are the mean-variance frontiers with and without a risk-free rate, respectively. The stars (*) represent the portfolio choice with a risk-free rate. The circles (o) represent the portfolio choice without a risk-free rate. The cross (+) is a portfolio of 60 percent stocks, 20 percent bonds, and 20 percent Treasury bills.

each graph represent the mean-variance frontier with a risk-free rate (straight line) and without a risk-free rate (hyperbola). The stars and circles represent the corresponding optimal portfolios. As a reference point, we also plot a portfolio of 60 percent stocks, 20 percent bonds, and 20 percent Treasury bills, which happens to resemble the optimal portfolio of a mean-variance investor with $\gamma = 5$.

Several well-known but nevertheless interesting features of the mean-variance optimal portfolios emerge. Consider the portfolio choice with a risk-free rate. Except when $\gamma = 2$, in which case the short-sale constraints are binding, all mean-variance investors hold the same risky position of 86 or 90 percent stocks and 14 or 10 percent bonds, depending on the horizon but irrespective of risk aversion. Risk aversion only affects how much wealth the investor allocates to risky securities instead of to risk-free Treasury bills. This allocation ranges from 100 percent for $\gamma = 2$ to about 20 percent for $\gamma = 20$.

Graphically, the fact that all mean-variance investors hold the same risky position, which is the portfolio at the tangency of the two mean-variance frontiers, but allocate different fractions of wealth to it implies that the optimal portfolios are all arranged on a straight line in the expected return versus standard deviation space. Also, we notice from Figure 2 that the optimal standard deviation of wealth is inversely proportional to γ . This happens because the decision of how much wealth to invest in the risky tangency portfolio is inversely proportional to the investors' absolute risk aversion (see equation (6)).

These portfolio choice patterns are the direct consequence of two-fund separation and have important implications for the variable selection.²¹ Recall our example of two investors, one who is very risk averse, holds mostly bonds, and wants to predict variances, and another who is less risk averse, holds mostly stocks, and wants to predict means. We argued, based on this example, that the variable selection may differ across investors for two reasons: the investors' preferences for predicting the various moments of returns and their portfolio holdings. The implication of two-fund separation is that all mean-variance investors hold the same risky position, unless the borrowing constraints are binding or there is no risk-free rate, which means their variable selection can only differ due to different preferences for predicting means, variances, covariances, and higher order moments.

The conclusion that different investors hold the same risky position does not apply to the portfolio choice without a risk-free rate, although two-fund separation holds nevertheless (only now with two risky portfolios). Without a risk-free rate, the investors' stock holdings decrease and the bond holdings increase with the level of risk aversion. Also, the standard deviation of wealth decreases less than proportionally with γ , which means that relative to the portfolio choice with a risk-free rate the investors are taking on more risk.

A.1.2. Constant Relative Risk-averse Investors. Panel B of Table III presents estimates of the unconditional portfolio choice of investors with CRRA preferences and relative risk aversion of 2, 5, 10, and 20. The results are similar to those for mean-variance preferences, except that CRRA investors tend to hold less stocks (up to seven percent less) and more bonds (up to five percent more), relative to equally risk-averse mean-variance investors. These differences in the portfolio choices are attributed to the negative skewness in stock returns and the positive skewness in bond returns that we document in Panel A of Table I.²²

Given these differences in the optimal portfolio weights, it is interesting that the implications of two-fund separation for the mean-variance portfolio choice with a risk-free rate apply as well to CRRA preferences. This is an empirical result, not a theoretical one. In theory, the risky position of a CRRA investor can depend on relative risk aversion, since the investor's preferences for higher order moments, which differentiate a CRRA investor from

²¹ We use the notion of two-fund separation to refer to the fact that two mean-variance efficient portfolios span the mean-variance frontier. With a risk-free rate, these two portfolios are the Treasury bill and tangency portfolio. It is important, however, to realize that in the literature, two-fund separation is typically an equilibrium statement that refers to the "market portfolio" being mean-variance efficient.

²² We observe the greatest differences between the optimal portfolio weights of mean-variance and CRRA investors at the three-month horizon, where stock returns exhibit the most negative skewness. The three- and six-month estimates are not reported in Table III to conserve space.

an equally risk-averse mean-variance investor, are a function of relative risk aversion. In the data, however, the effect of the higher order moments is apparently not strong enough to be noticeable in the stock holdings of CRRA investors with different degrees of risk aversion, although it is clearly a factor in explaining the different stock holdings of equally risk-averse CRRA and mean-variance investors.²³

A.2. Nonexpected Utility Preferences

A.2.1. Ambiguity-averse Investors. The results for investors with ambiguity-aversion preferences are in Panel A of Table IV. We consider the cases $\gamma = \{5, 10\}$ and $\varepsilon = \{0.001, 0.005, 0.010\}$.²⁴ Recall that the case $\varepsilon = 0$ corresponds to the CRRA portfolio choices in Panel B of Table III. We parameterize the worst-case returns on stocks and bonds, which we need to evaluate the infimum in the objective function (18) as the empirical univariate minimums from Panel A of Table I.

Ambiguity aversion has two effects on the portfolio choice. First, an increase in ambiguity aversion leads investors to substitute Treasury bills for risky securities or bonds for stocks, depending on whether or not a risk-free security is available. This effect of an increase in ambiguity aversion parallels that of an increase in risk aversion, which is an observational equivalence formalized recently by Liu (1999) and Maenhout (1999). Second, in the case with a risk-free rate, the investor does not take a position (positive or negative) in bonds for $\varepsilon \geq 0.005$, even if we relax the short-sale constraints. This tendency to not hold a position in some securities, which happens because the expected returns are not sufficiently positive or negative to justify taking on the associated ambiguity, is the subject of Dow and Werlang's (1992) paper. It also motivates Liu's (1998) attempt to explain the limited participation of U.S. households in financial markets (see Mankiw and Zeldes (1991)) with ambiguity-aversion preferences. For the variable selection, the tendency of investors with ambiguity aversion to not hold a position in some securities has the same effect as two-fund separation. It causes different investors to hold similar portfolios.

²³ To be precise, the effects of higher order moments on the relative stock holdings of CRRA investors with different levels of risk aversion are not strong enough to be noticeable after rounding the estimates to the second decimal. They are actually noticeable, but, of course, not statistically significant at the third decimal.

Lynch (2000) reports analogous results for the multiperiod portfolio choice of CRRA and mean-variance investors under the assumption of joint log-normality of returns.

²⁴ The choice of ε is admittedly ad hoc. Camerer (1995) cites attempts to calibrate ambiguity aversion preferences to gambling experiments. Since it is unclear, however, how these experimental results relate to ambiguity aversion in financial markets, we present estimates for a relatively generous range of ε .

A.2.2. Prospect-theory Investors. Panel B of Table IV presents estimates of the unconditional portfolio choice of prospect theory investors.²⁵ We set $c = 1$ and, guided by the experimental calibrations of Tversky and Kahneman (1992), consider the parameter values $b = \{0.8, 0.9, 1.0\}$ and $l = \{2.0, 2.5, 3.0\}$, where $b = 1$ corresponds to pure loss aversion.²⁶ We set the wealth reference level \bar{W} , which is chosen in the editing stage of prospect theory, to the initial wealth $W_t = 1$.

By far the most striking feature of the prospect theory results are the strong horizon effects (which is why we report results for all four investment horizons here).²⁷ Consider the portfolio choice with a risk-free rate. At the one-month horizon the optimal allocation consists of more than 90 percent Treasury bills, while at the one-year horizon it is 100 percent stocks, irrespective of the preference parameters. This is in obvious contrast to the results for the other three sets of preferences, which exhibit only minimal horizon effects.

Benartzi and Thaler (1995) explain that the more often a loss-averse investor evaluates his or her portfolio, the less attractive are high expected return but high variance investments because losses of these investments are realized more often at short horizons than at long horizons. Loss aversion effectively causes short-term investors to be extremely risk averse, since the return distribution straddles the kink of the utility function, but long-term investors to be almost risk neutral, as the mass of the return distribution moves away from the kink. In contrast, Merton (1969) and Samuelson (1969) show that as long as returns form a martingale, the portfolio choices of mean-variance and CRRA investors are independent of the horizon.²⁸ Even if we relax the martingale assumption, the mean-variance and CRRA portfolio choices do not exhibit nearly the magnitude of horizon effects that we observe with loss aversion.

²⁵ To compute the standard errors, which requires first and second derivatives of the objective function with respect to α , we smooth the kink in the objective function at zero with a polynomial in-between the contact points -0.005 and $+0.005$ with continuous first and second derivatives at the contact points.

²⁶ Like Barberis, Huang, and Santos (2001), who use $c = 1$, $b = 1$, and $l = 2.25$, we abstract from the probability distortions. When $c = 1$, the weights $\pi(\cdot)$ reduce to the objective probabilities $p(\cdot)$. We checked, however, that our results are qualitatively robust to using Tversky and Kahneman's (1992) value of $c = 0.65$.

²⁷ In a single-period context, we define horizon effects as differences in the portfolio policies of investors with different buy-and-hold horizons. An alternative definition, and one that we employ in Section IV.C.3, refers to differences in the portfolio policies of multiperiod investors with the same rebalancing frequency but different numbers of rebalancing periods.

The behavioral finance literature distinguishes between the investor's rebalancing period when securities are traded and the "evaluation period" when gains and losses are realized mentally (as opposed to financially) (Benartzi and Thaler (1995) or Barberis and Huang (2001)). For simplicity, we assume that the rebalancing and evaluation periods are the same.

²⁸ This horizon irrelevance result breaks down if the returns are mean reverting or if the returns are contemporaneously correlated with innovations to the investment opportunity set (see Barberis (2000) and Campbell and Viceira (1999)).

Table IV
Unconditional Portfolio Choice with Nonexpected Utility Preferences

This table shows estimates of the unconditional portfolio choice of investors with single-period objectives:

$$\text{Panel A: } \max_{\alpha} \left((1 - \varepsilon) \mathbb{E} \left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma} \right] + \varepsilon \inf_R \frac{W_{t+1}^{1-\gamma}}{1-\gamma} \right) \quad \text{Panel B: } \max_{\alpha} \mathbb{E} \left[\begin{array}{ll} -l(\bar{W} - W_{t+1})^b & \text{if } W_{t+1} < \bar{W} \\ (W_{t+1} - \bar{W})^b & \text{if } W_{t+1} \geq \bar{W} \end{array} \right],$$

where W_{t+1} denotes next period's wealth, $\bar{W} = 1$ is a subjective wealth reference level, and $\alpha = [\alpha^s, \alpha^b, \alpha^{tb}]$ are the fractions of current wealth $W_t = 1$ invested in stocks, bonds, and Treasury bills, respectively. The investment horizon is 1, 3, 6, or 12 months. The Without Risk-Free Rate estimates are for a portfolio choice between stocks and bonds (i.e., $\alpha^{tb} = 0$) and the With Risk-Free Rate estimates are for a portfolio choice between stocks, bonds, and Treasury bills. In both cases, the optimization is subject to the short-sales constraints $0 \leq \alpha \leq 1$. In brackets are autocorrelation adjusted asymptotic standard errors.

		Without Risk-Free Rate		With Risk-Free Rate				Without Risk-Free Rate		With Risk-Free Rate			
		α^s	α^b	α^s	α^b	α^{tb}	$\frac{\alpha^s}{\alpha^s + \alpha^b}$	α^s	α^b	α^s	α^b	α^{tb}	$\frac{\alpha^s}{\alpha^s + \alpha^b}$
Panel A: Ambiguity-Averse Investors													
		One-Month Horizon						One-Year Horizon					
$\gamma = 5$	$\varepsilon = 0.1\%$	0.66 [0.20]	0.34 [0.20]	0.61 [0.36]	0.11 [0.38]	0.28 [0.21]	0.85 [0.46]	0.52 [0.17]	0.48 [0.17]	0.53 [0.29]	0.05 [0.31]	0.42 [0.18]	0.91 [0.49]
	$\varepsilon = 0.5\%$	0.52 [0.18]	0.48 [0.18]	0.48 [0.32]	0.00 [0.37]	0.52 [0.19]	1.00 [0.76]	0.43 [0.13]	0.57 [0.13]	0.47 [0.24]	0.00 [0.29]	0.53 [0.16]	1.00 [0.62]
	$\varepsilon = 1.0\%$	0.38 [0.16]	0.62 [0.16]	0.32 [0.29]	0.00 [0.36]	0.68 [0.17]	1.00 [1.12]	0.35 [0.11]	0.65 [0.11]	0.41 [0.21]	0.00 [0.29]	0.59 [0.15]	1.00 [0.70]
$\gamma = 10$	$\varepsilon = 0.1\%$	0.40 [0.10]	0.60 [0.10]	0.31 [0.18]	0.05 [0.19]	0.64 [0.11]	0.86 [0.46]	0.27 [0.08]	0.73 [0.08]	0.27 [0.15]	0.03 [0.15]	0.70 [0.09]	0.91 [0.48]
	$\varepsilon = 0.5\%$	0.29 [0.09]	0.71 [0.09]	0.24 [0.16]	0.00 [0.18]	0.76 [0.09]	1.00 [0.78]	0.17 [0.06]	0.83 [0.06]	0.24 [0.12]	0.00 [0.15]	0.76 [0.08]	1.00 [0.61]
	$\varepsilon = 1.0\%$	0.18 [0.07]	0.82 [0.07]	0.16 [0.15]	0.00 [0.18]	0.84 [0.09]	1.00 [1.12]	0.11 [0.05]	0.89 [0.05]	0.21 [0.11]	0.00 [0.14]	0.79 [0.08]	1.00 [0.68]

Panel B: Prospect-Theory Investors

		One-Month Horizon			Three-Month Horizon								
$b = 0.8$	$l = 2.00$	0.43 [0.20]	0.57 [0.20]	0.03 [0.19]	0.02 [0.19]	0.95 [0.11]	0.60 [4.27]	0.65 [0.26]	0.35 [0.26]	0.10 [0.11]	0.02 [0.10]	0.88 [0.07]	0.83 [1.53]
	$l = 2.50$	0.30 [0.14]	0.70 [0.14]	0.03 [0.15]	0.02 [0.15]	0.95 [0.09]	0.60 [3.87]	0.51 [0.20]	0.49 [0.20]	0.08 [0.09]	0.04 [0.09]	0.88 [0.06]	0.67 [1.36]
$b = 0.9$	$l = 2.00$	0.42 [0.19]	0.58 [0.19]	0.06 [0.18]	0.00 [0.19]	0.94 [0.11]	1.00 [3.01]	0.70 [0.35]	0.30 [0.35]	0.20 [0.17]	0.03 [0.14]	0.77 [0.13]	0.87 [1.49]
	$l = 2.50$	0.30 [0.13]	0.70 [0.13]	0.05 [0.14]	0.01 [0.15]	0.94 [0.09]	0.83 [3.32]	0.50 [0.22]	0.50 [0.22]	0.12 [0.10]	0.02 [0.10]	0.86 [0.07]	0.86 [1.33]
$b = 1.0$	$l = 2.00$	0.39 [0.17]	0.61 [0.17]	0.08 [0.18]	0.02 [0.18]	0.90 [0.11]	0.80 [2.35]	0.76 [0.39]	0.24 [0.39]	0.46 [0.50]	0.06 [0.23]	0.48 [0.45]	0.88 [1.58]
	$l = 2.50$	0.30 [0.11]	0.70 [0.11]	0.07 [0.14]	0.02 [0.14]	0.91 [0.08]	0.78 [2.07]	0.49 [0.24]	0.51 [0.24]	0.15 [0.13]	0.04 [0.12]	0.81 [0.09]	0.79 [1.41]
		Six-Month Horizon			One-Year Horizon								
$b = 0.8$	$l = 2.00$	1.00 [0.39]	0.00 [0.39]	0.22 [0.42]	0.13 [0.22]	0.65 [0.43]	0.63 [0.90]	1.00 [0.43]	0.00 [0.43]	1.00 [0.53]	0.00 [0.41]	0.00 [0.47]	1.00 [0.14]
	$l = 2.50$	0.53 [0.23]	0.47 [0.23]	0.16 [0.11]	0.07 [0.12]	0.77 [0.11]	0.70 [0.66]	1.00 [0.44]	0.00 [0.44]	1.00 [0.54]	0.00 [0.41]	0.00 [0.48]	1.00 [0.15]
$b = 0.9$	$l = 2.00$	1.00 [0.44]	0.00 [0.44]	1.00 [0.49]	0.00 [0.26]	0.00 [0.50]	1.00 [0.93]	1.00 [0.48]	0.00 [0.48]	1.00 [0.58]	0.00 [0.43]	0.00 [0.51]	1.00 [0.17]
	$l = 2.50$	0.60 [0.25]	0.40 [0.25]	0.24 [0.18]	0.14 [0.21]	0.62 [0.22]	0.63 [0.86]	1.00 [0.49]	0.00 [0.49]	1.00 [0.59]	0.00 [0.43]	0.00 [0.51]	1.00 [0.18]
$b = 1.0$	$l = 2.00$	1.00 [0.49]	0.00 [0.49]	1.00 [0.54]	0.00 [0.28]	0.00 [0.59]	1.00 [1.11]	1.00 [0.52]	0.00 [0.52]	1.00 [0.64]	0.00 [0.48]	0.00 [0.54]	1.00 [0.23]
	$l = 2.50$	0.65 [0.31]	0.35 [0.31]	0.49 [0.23]	0.18 [0.26]	0.33 [0.26]	0.73 [0.91]	1.00 [0.53]	0.00 [0.53]	1.00 [0.64]	0.00 [0.49]	0.00 [0.55]	1.00 [0.25]

Besides the horizon effects, the portfolio choice with loss aversion preferences is noteworthy because investors with different degrees of loss aversion hold substantially different positions in the risky securities (except at the annual horizon, where the short-sale constraints are binding). In particular, for the portfolio choice with a risk-free rate, the fraction of stocks in the risky part of the allocation ranges from 60 to 100 percent.

B. Optimal Index Composition

Turning now to the conditional portfolio choice, we estimate the optimal index composition as described in Section II.B. We use linear instruments $g(Z_t) = Z_t$, which may be suboptimal in the sense of problem (12), but result in numerically more reliable estimates.²⁹ We demean and standardize the variables Z_t to be able to interpret the magnitude of the index coefficients. We also normalize the index coefficients to sum to one in absolute values, meaning $|\beta|' \iota = 1$, since they are only identified up to scale. This normalization means that we can read the index coefficients as signed percentage loadings.

B.1. Expected Utility Preferences

B.1.1. Mean-variance Investors. Panel A of Table V presents estimates of the index coefficients for mean-variance investors with a one-month or one-year horizon. One, two, or three stars indicate that the coefficient is statistically significant at the 10, 5, or 1 percent levels, respectively. Table VI describes the estimated indices, where Panel A shows univariate descriptive statistics and Panel B shows pairwise correlations with the four predictors, with the returns on bonds, stocks, and wealth generated by the unconditional portfolio choice, and with the squared returns on bonds, stocks, and wealth.³⁰

²⁹ The “optimal” instruments in equation (14) tend to result in estimates that are sensitive to the starting values and are only locally optimal. The problem is that, due to the derivative of $\alpha(Z_t; \beta; \beta)$ with respect to its second argument, the optimal instruments cause the GMM objective function to be less well behaved than with linear instruments (see also the discussion in Section II.B).

We verified, however, that the estimates with linear instruments are virtually identical to the estimates that correspond to the global minimum of the GMM criterion with the optimal instruments. In particular, we solved for eight indices (two for each set of preferences) using a random search optimization algorithm that is robust to ill-behaved objective functions. In the *worst case* (ambiguity-aversion preferences with $\gamma = 10$ and $\varepsilon = 0.01$), the difference between the estimates with linear and optimal instruments implies an expected utility loss of 0.09 percent in certainty equivalent terms. The corresponding estimates of the index are [0.249, -0.124, 0.597, 0.030] with linear and [0.312, -0.119, 0.518, 0.051] with optimal instruments.

³⁰ We use the return on wealth generated by the unconditional portfolio choice as a meaningful way to collapse the bivariate return series (bonds and stocks) into a univariate variable. In particular, the return on wealth weights the bond and stock returns relative to their unconditional importance to the investor.

Table V
Index Composition with Expected Utility Preferences

This table shows estimates of the optimal index composition for the conditional portfolio choice of investors with single-period preferences:

$$\text{Panel A: } v(W_{t+1}) = W_{t+1} - \frac{\gamma}{2} (W_{t+1} - E[W_{t+1}])^2 \quad \text{Panel B: } v(W_{t+1}) = \frac{W_{t+1}^{1-\gamma}}{1-\gamma},$$

where W_{t+1} denotes next period's wealth. Each index is a linear combination of the default spread Def, the log dividend-to-price ratio of the S&P 500 index LnDP, the term spread Term, and the S&P index momentum variable Trend. The investment horizon is one month or one year. The Without Risk-Free Rate estimates are for a portfolio choice between stocks and bonds and the With Risk-Free Rate estimates are for a portfolio choice between stocks, bonds, and Treasury bills. In both cases, the optimization is subject to the short-sale constraints $0 \leq \alpha \leq 1$.

	Without Risk-Free Rate				With Risk-Free Rate			
	Def	LnDP	Term	Trend	Def	LnDP	Term	Trend
Panel A: Mean-Variance Investors								
One-Month Horizon								
$\gamma = 2$	-0.091	0.127	0.465**	0.317*	0.165	-0.134	0.549**	0.151
$\gamma = 5$	-0.075	0.122	0.444**	0.360*	0.164	-0.133	0.554**	0.149
$\gamma = 10$	0.002	0.168	0.401**	0.429**	0.158	-0.131	0.590**	0.121
$\gamma = 20$	0.049	0.230*	0.299*	0.422**	0.151	-0.129	0.599**	0.121
One-Year Horizon								
$\gamma = 2$	-0.235	0.418*	0.347	0.001	-0.107	0.291	0.559*	-0.043
$\gamma = 5$	-0.199	0.438*	0.327	0.036	-0.106	0.287	0.565*	-0.042
$\gamma = 10$	-0.159	0.459*	0.324	0.058	-0.105	0.285	0.570*	-0.041
$\gamma = 20$	-0.130	0.468*	0.318	0.083	-0.105	0.280	0.573*	-0.041
Panel B: Constant Relative Risk-Averse Investors								
One-Month Horizon								
$\gamma = 2$	-0.146	0.154	0.431**	0.268*	0.181*	-0.083	0.648**	0.089
$\gamma = 5$	-0.094	0.144	0.420**	0.342*	0.173*	-0.087	0.667**	0.073
$\gamma = 10$	-0.011	0.173	0.388*	0.428**	0.168	-0.091	0.683**	0.058
$\gamma = 20$	0.086	0.173	0.268*	0.472**	0.164	-0.091	0.687**	0.058
One-Year Horizon								
$\gamma = 2$	-0.136	0.441*	0.391	0.032	0.095	0.179	0.625*	-0.101
$\gamma = 5$	-0.124	0.462*	0.361	0.053	0.087	0.180	0.630*	-0.104
$\gamma = 10$	-0.102	0.470*	0.320	0.108	0.054	0.184	0.655*	-0.107
$\gamma = 20$	-0.055	0.504*	0.229	0.212	0.049	0.184	0.658*	-0.108

** and * denote statistical significance at the 5 and 10 percent levels, respectively.

Consider first the indices for the portfolio choice with a risk-free rate. The most striking feature of the results is that the index composition is remarkably insensitive to the level of risk aversion. This is again an implication of two-fund separation because with a risk-free rate and without short-sale

Table VI
Description of Indices with MV Preferences

Panel A of this table shows descriptive statistics of the optimal indices for the conditional portfolio choice of investors with single-period mean-variance preferences. Each index is a linear combination of the default spread Def, the log dividend-to-price ratio of the S&P 500 index LnDP, the term spread Term, and the S&P index momentum variable Trend. The index coefficients are shown in Panel A of Table V. The investment horizon is one month or one year. Panel B shows correlations of the indices with the predictors, excess bond returns r_{t+1}^b , excess stock returns r_{t+1}^s , and excess returns on wealth r_{t+1}^w generated by the unconditional portfolio choices from Panel A of Table III. It also shows correlations of the indices with squared excess returns on bonds, stocks, and wealth.

	One-Month Horizon					One-Year Horizon				
	Without Risk-Free Rate				With Risk-Free Rate	Without Risk-Free Rate				With Risk-Free Rate
	$\gamma = 2$	$\gamma = 5$	$\gamma = 10$	$\gamma = 20$	$\gamma = 5$	$\gamma = 2$	$\gamma = 5$	$\gamma = 10$	$\gamma = 20$	$\gamma = 5$
Panel A: Descriptive Statistics										
Median	0.05	0.04	0.04	0.00	0.05	-0.05	-0.05	-0.20	-0.02	-0.03
StdDev	0.59	0.60	0.61	0.53	0.66	0.39	0.39	0.00	0.41	0.52
Skew	-0.38	-0.42	-0.70	-0.23	-0.21	0.06	0.06	0.00	-0.01	0.03
Kurtosis	2.81	2.83	2.99	3.18	2.73	3.14	3.23	3.20	3.14	2.98
Min	-1.83	-2.01	-2.17	-1.90	-1.94	-1.11	-1.18	-1.26	-1.29	-1.56
Max	1.14	1.12	1.33	1.38	1.48	1.16	1.16	1.17	1.17	1.35
ρ_1	0.95	0.95	0.94	0.93	0.97	0.96	0.97	0.97	0.97	0.96
ρ_3	0.83	0.82	0.80	0.79	0.86	0.86	0.88	0.90	0.91	0.84
ρ_6	0.66	0.63	0.58	0.57	0.73	0.75	0.78	0.80	0.82	0.70
ρ_{12}	0.44	0.39	0.30	0.27	0.55	0.55	0.60	0.30	0.66	0.53
Panel B: Correlations										
Def	-0.13	-0.11	0.05	0.23	0.06	-0.08	0.04	0.16	0.23	0.03
LnDP	-0.32	-0.32	-0.18	0.03	-0.43	0.43	0.52	0.60	0.63	0.08
Term	0.86	0.83	0.74	0.61	0.95	0.55	0.50	0.47	0.44	0.89
Trend	0.68	0.73	0.79	0.79	0.48	-0.07	-0.01	0.02	0.06	0.04
r_{t+1}^b	0.03	0.02	0.01	-0.00	0.13	0.06	0.05	0.05	0.05	0.24
r_{t+1}^s	0.13	0.12	0.12	0.10	0.18	0.37	0.36	0.36	0.35	0.39
r_{t+1}^w	0.13	0.12	0.09	0.05	0.18	0.37	0.34	0.28	0.22	0.40
$(r_{t+1}^b)^2$	-0.07	-0.06	0.01	0.08	-0.01	0.17	0.21	0.25	0.27	0.20
$(r_{t+1}^s)^2$	-0.17	-0.18	-0.16	-0.13	-0.11	-0.04	-0.03	-0.01	0.00	-0.06
$(r_{t+1}^w)^2$	-0.17	-0.20	-0.14	-0.02	-0.11	-0.04	-0.03	0.02	0.06	-0.05

constraints, the optimal portfolio choice of a mean-variance investor is to allocate a fraction α_t^{tgc} (from equation (6)) of wealth to the tangency portfolio and to hold the remaining wealth in risk-free Treasury bills. Two-fund separation therefore implies that the portfolio policies of all mean-variance investors are proportional to each other and that, as a result, the optimal index of conditional variables must be the same for all investors (since the index is identified only up to scale). In line with this theoretical result, the estimated indices are in fact identical across different levels of risk aver-

sions if we relax the short-sale constraints. If we impose the short-sale constraints, as we do for Table V, the estimates are slightly different because the constraints are binding more often for an investor with $\gamma = 2$ than for an investor with $\gamma = 20$.

At the monthly horizon, the index loads positively on the term spread, positively on the default spread, negatively on the log dividend yield, and positively on the trend, with relative weights of about 55, 16, 13, and 15 percent, respectively. At the one-year horizon, it loads positively on the term spread and the dividend yield, with relative weights of 57 and 29 percent. The default spread and trend each account for less than 11 percent of the index.

To better understand the index composition, given that the predictors are correlated, consider the pairwise correlations in Panel B of Table VI. At the monthly horizon, the index is most correlated with the term spread and is about equally correlated with the log dividend yield and trend (correlations of 0.95, -0.43 , and 0.48 , respectively). Somewhat unexpectedly, it is nearly orthogonal to the default spread, although the loading on the default spread exceeds those on the log dividend yield and trend. At the annual horizon, the index is positively correlated with the term spread and dividend yield (correlations of 0.89 , 0.08 , respectively) and is virtually orthogonal to the default spread and trend.

At both horizons, the indices relate positively to the returns on wealth, with correlations of 0.18 and 0.40 , and negatively to the squared returns on wealth, with correlations of -0.11 and -0.05 .³¹ The signs and magnitudes of these correlations suggest the following:

- An increase in the index represents an unambiguous improvement in investment opportunities, since it increases the mean and decreases the variance of future wealth.
- The term spread receives the highest weighting in all indices.³² Tables I and II show that the term spread is the only variable that is strongly (in relative terms) positively correlated with the mean and negatively correlated with the variance of stock returns. Since mean-variance investors allocate unconditionally between 86 and 90 percent of their risky investments to stocks, they definitely cheer increases in the term spread. The log dividend-to-price ratio, the more traditional predictor of expected returns, receives a negative weight at the monthly horizon, since it is negatively correlated with the mean and positively correlated with the variance of stock returns, and a less positive (60 percent less) weight

³¹ At the monthly horizon, the positive correlation of the index with returns on wealth is attributed to the positive loading on the term premium. The negative correlation with squared returns comes from the positive loading on the trend. At the annual horizon both the positive correlation with returns and the negative correlation with squared returns on wealth are attributed to the term spread.

³² We checked that the dominance of the term spread in the indices is not attributed to a few influential data points, such as the period surrounding October 1987.

than the term spread at the annual horizon because it is positively correlated with the mean but also with the variance of stock returns. The less positive weight reflects the fact that an increase in the dividend yield is good news for expected returns but bad news for the variance of returns.

- The correlations of the indices with the returns and squared returns on wealth are both below their theoretical maximums implied by the individual moment regressions in Panel A of Table II. This implies that, in line with our intuition, the indices achieve a *compromise* between predicting means and variances.
- The balance between predicting means and variances depends on the relative degree of predictability of the two moments. At the monthly horizon, the mean and variance of stock returns are equally predictable (R^2 s of 0.04 and 0.05) and the indices are about equally correlated with returns and squared returns on wealth. At the annual horizon, however, the expected return of stocks is almost twice as predictable as the variance (R^2 s of 0.18 and 0.10) and, as a result, the indices focus almost exclusively on predicting mean returns.

This is exactly how we intend the indices to work. The only feature of our approach that is not well illustrated yet is that the index composition can be investor specific.

This feature of the indices is apparent from the results for the portfolio choice without a risk-free rate. Recall from Table III that as risk aversion increases, the unconditional portfolio choice without a risk-free rate varies from 100 percent stocks to the global minimum variance portfolio of 20 percent stocks and 80 percent bonds. The corresponding entries in Tables V and VI show that as the relative position in stocks versus bonds changes with the level of risk aversion, the index composition also changes. In particular, at the one-month horizon, the most important predictor is the term spread for $\gamma = \{2,5\}$ and the trend for $\gamma = \{10,20\}$.

B.1.2. Constant Relative Risk-averse Investors. The indices for CRRA preferences, in Panel B of Table V, resemble those for mean-variance preferences. Their likes and dislikes of higher order moments induce CRRA investors to focus somewhat more on the term premium and less on the dividend yield (except for the one-month portfolio choice without a risk-free rate). Despite the differences between the CRRA and mean-variance indices, the indices for CRRA investors with different levels of risk aversion are no more different than those for mean-variance investors with different levels of risk aversion. In fact, if we relax the short-sell constraints, the CRRA indices for the portfolio choice with a risk-free rate are identical across different levels of risk aversion, just as they are for mean-variance investors. This is true not only for γ ranging from 2 through 20, but also for γ as large as 100 (not shown in the table). We hence conclude that the preferences for higher

order moments, which depend on the level of risk aversion, are not strong enough to induce different CRRA investors to focus on different economic variables.

B.2. Nonexpected Utility Preferences

B.2.1. Ambiguity-averse Investors. Panel A of Table VII presents the estimated indices for investors with ambiguity aversion preferences. Overall, the results are quite similar to those for the expected utility preferences in Table V. For the portfolio choice with a risk-free rate, the term spread is by far the most important contributor to the index. The second most important predictor is the default spread at the monthly horizon and the log dividend yield at the annual horizon.

A more subtle feature of the results is that the index composition depends on the degree of ambiguity aversion. Consider, for example, the case $\gamma = 5$ for the one-month portfolio choice with a risk-free rate. With $\varepsilon = 0$ (in Panel B of Table V) the index coefficients are 0.173, -0.087 , 0.667, and 0.073, while with $\varepsilon = 0.01$ they are 0.184, -0.116 , 0.644, and -0.056 . Depending on whether the portfolio choice includes a risk-free rate or not, there are different reasons for why the index composition depends on the degree of ambiguity aversion. Without a risk-free rate, an increase in ambiguity aversion is observationally equivalent to an increase in risk aversion. Thus, the index composition changes in the same way as it does for the CRRA portfolio choice without a risk-free rate when risk aversion increases. With a risk-free rate, in contrast, the index changes because sufficiently ambiguity-averse investors avoid bonds. As the unconditionally optimal risky position shifts from 87 percent stocks and 13 percent bonds (in Panel B of Table III) to 100 percent stocks (in Panel A of Table IV), the index changes accordingly.

B.2.2. Prospect-theory Investors. The indices for prospect theory investors, shown in Panel B of Table VII, exhibit the most variation across different preference parameters and are the most different from the mean-variance benchmark. They are also the most challenging to interpret, due to the strong horizon effects that we documented in Panel B of Table IV.

To abstract from the horizon effects, we focus on the index composition for the portfolio choice with a risk-free rate at the one-year horizon, where all prospect theory investors hold only stocks. The indices are quite different from those for mean-variance preferences at the one-year horizon. In particular, when $b = 0.8$, the index loads positively and substantially on the default spread (correlation of 0.22). This is puzzling, at first sight, since the default spread is a strong positive predictor of the variance of stock returns and even relates negatively to the mean. However, the fact that prospect theory investors are *risk loving* in the case of losses explains the role of the default spread. When the default spread is high the investor anticipates relatively volatile losses, as opposed to certain losses, which is preferred in

Table VII
Index Composition with Nonexpected Utility Preferences

This table shows estimates of the optimal index composition for the conditional portfolio choice of investors with single-period preferences:

$$\text{Panel A: } v(W_{t+1}) = (1 - \varepsilon) \frac{W_{t+1}^{1-\gamma}}{1 - \gamma} + \varepsilon \inf_R \frac{W_{t+1}^{1-\gamma}}{1 - \gamma}$$

$$\text{Panel B: } v(W_{t+1}) = \begin{cases} -l(\bar{W} - W_{t+1})^b & \text{if } W_{t+1} < \bar{W} \\ (W_{t+1} - \bar{W})^b & \text{if } W_{t+1} \geq \bar{W} \end{cases},$$

where W_{t+1} denotes next period's wealth and $\bar{W} = 1$ is a subjective wealth reference level. Each index is a linear combination of the default spread Def, the log dividend-to-price ratio of the S&P 500 index LnDP, the term spread Term, and the S&P index momentum variable Trend. The investment horizon is 1, 3, 6, and 12 months. The Without Risk-Free Rate estimates are for a portfolio choice between stocks and bonds and the With Risk-Free Rate estimates are for a portfolio choice between stocks, bonds, and Treasury bills. In both cases, the optimization is subject to the short-sale constraints $0 \leq \alpha \leq 1$.

		Without Risk-Free Rate				With Risk-Free Rate			
		Def	LnDP	Term	Trend	Def	LnDP	Term	Trend
Panel A: Ambiguity-Averse Investors									
One-Month Horizon									
$\gamma = 5$	$\varepsilon = 0.1\%$	-0.024	0.169	0.415**	0.392*	0.178*	-0.097	0.659**	0.066
	$\varepsilon = 0.5\%$	0.008	0.175	0.384*	0.433**	0.183*	-0.115	0.645**	0.057
	$\varepsilon = 1.0\%$	0.014	0.183	0.327*	0.476**	0.184*	-0.116	0.644**	0.056
$\gamma = 10$	$\varepsilon = 0.1\%$	-0.008	0.186	0.346*	0.460**	0.196*	-0.099	0.669**	0.036
	$\varepsilon = 0.5\%$	0.013	0.195	0.307*	0.485**	0.235*	-0.121	0.612**	0.032
	$\varepsilon = 1.0\%$	0.027	0.221	0.255*	0.497**	0.249*	-0.124	0.597**	0.030
One-Year Horizon									
$\gamma = 5$	$\varepsilon = 0.1\%$	-0.080	0.483*	0.349	0.088	0.058	0.194	0.620*	-0.128
	$\varepsilon = 0.5\%$	0.012	0.512*	0.367	0.109	0.041	0.218	0.592*	-0.149
	$\varepsilon = 1.0\%$	0.045	0.535*	0.284	0.136	0.039	0.246	0.568*	-0.147
$\gamma = 10$	$\varepsilon = 0.1\%$	-0.063	0.491*	0.301	0.145	0.052	0.190	0.623*	-0.135
	$\varepsilon = 0.5\%$	0.024	0.521*	0.294	0.161	0.035	0.213	0.606*	-0.146
	$\varepsilon = 1.0\%$	0.053	0.543*	0.222	0.182	0.029	0.236	0.592*	-0.143

(continued)

prospect theory.³³ Indeed, as b increases and the investor becomes less risk loving over losses, the correlation of the index with the default spread drops sharply. The correlation increases for values of b less than 0.8 (not shown in the table).

³³ We checked that losses are indeed more volatile when the default spread is high than when it is low.

Table VII—Continued

		Without Risk-Free Rate				With Risk-Free Rate			
		Def	LnDP	Term	Trend	Def	LnDP	Term	Trend
Panel B: Prospect-Theory Investors									
One-Month Horizon									
$b = 0.8$	$l = 2.0$	0.069	-0.247	0.677**	-0.007	0.232*	-0.139	0.589**	-0.041
	$l = 2.5$	-0.032	-0.276*	0.587**	-0.105	0.229*	-0.138	0.584**	-0.049
$b = 0.9$	$l = 2.0$	0.101	-0.227	0.644**	-0.027	0.164	-0.135	0.590**	0.111
	$l = 2.5$	0.001	-0.276*	0.608**	-0.115	0.190	-0.150	0.624**	0.036
$b = 1.0$	$l = 2.0$	0.125	-0.210	0.603**	-0.062	0.204	-0.149	0.605**	0.042
	$l = 2.5$	0.032	-0.260*	0.582**	-0.126	0.189	-0.154	0.611**	0.046
Three-Month Horizon									
$b = 0.8$	$l = 2.0$	0.148	0.009	0.806**	0.037	0.235	0.085	0.674**	-0.006
	$l = 2.5$	0.109	-0.062	0.823**	-0.007	0.235	0.085	0.674**	-0.006
$b = 0.9$	$l = 2.0$	0.175	0.009	0.797**	0.019	0.156	0.146	0.636**	0.062
	$l = 2.5$	0.133	-0.052	0.800**	-0.015	0.185	0.127	0.654**	0.034
$b = 1.0$	$l = 2.0$	0.204	0.002	0.786**	-0.007	0.207	0.117	0.657**	0.018
	$l = 2.5$	0.154	-0.049	0.771**	-0.026	0.199	0.127	0.652**	0.022
Six-Month Horizon									
$b = 0.8$	$l = 2.0$	0.190	0.066	0.716*	0.028	0.205	0.116	0.642*	-0.037
	$l = 2.5$	0.198	-0.006	0.776*	-0.019	0.205	0.116	0.642*	-0.037
$b = 0.9$	$l = 2.0$	0.210	0.067	0.714*	0.008	0.140	0.186	0.653*	0.020
	$l = 2.5$	0.216	0.003	0.755*	-0.026	0.168	0.166	0.665*	-0.002
$b = 1.0$	$l = 2.0$	0.230	0.057	0.694*	-0.019	0.187	0.143	0.651*	-0.018
	$l = 2.5$	0.230	0.005	0.726*	-0.040	0.178	0.154	0.654*	-0.013
One-Year Horizon									
$b = 0.8$	$l = 2.0$	0.115	0.163	0.623*	-0.099	0.125	0.173	0.564*	-0.137
	$l = 2.5$	0.141	0.094	0.653*	-0.112	0.125	0.173	0.564*	-0.137
$b = 0.9$	$l = 2.0$	0.133	0.147	0.612*	-0.109	0.040	0.264	0.578*	-0.117
	$l = 2.5$	0.156	0.092	0.632*	-0.120	0.071	0.231	0.573*	-0.125
$b = 1.0$	$l = 2.0$	0.156	0.124	0.599*	-0.121	0.095	0.206	0.566*	-0.133
	$l = 2.5$	0.172	0.086	0.614*	-0.128	0.084	0.218	0.568*	-0.130

** and * denote statistical significance at the 5 and 10 percent levels, respectively.

Another interesting difference between the prospect theory and mean-variance indices is the negative loading on the trend. At the one-year horizon, the trend is a weak negative predictor of the mean, but a strong negative predictor of the variance of stock returns. Thus, mean-variance investors virtually ignore the trend, since the correlation with the variance is not strong enough to offset the correlation with the mean of the same sign, and

vice versa. More puzzling than the negative loading itself, however, is the fact that the role of the trend variable in the index does not depend much on the curvature of the utility function, which suggests that it is unrelated to our usual mean-variance intuition. Instead, there are two reasons to believe that loss-averse investors use the trend to predict positive skewness, or rather fewer and smaller losses. First, the importance of the trend depends on the loss aversion parameter l . In particular, when we consider $l < 2$, the loading on the trend decreases, while it increases when we consider $l > 2.5$ (both results are not shown in the table). Second, the correlation of the trend with cubed stock returns, a simple measure of skewness, is -0.21 .

C. Conditional Asset Allocation

Figures 3 and 4 plot the portfolio policies of investors with expected and nonexpected utility preferences. The black, gray, and white areas in each plot represent the estimated allocations of wealth to Treasury bills, bonds, and stocks as a function of the estimated indices from Table V (for Figure 3) and Table VII (for Figure 4).³⁴ We compute the conditional portfolio weights for 15 evenly spaced realizations of the index ranging from its mean minus two standard deviations to its mean plus two standard deviations, which roughly corresponds to the interior 95 percent of the empirical distribution of the index. We only plot the portfolio policies with a risk-free rate.

C.1. Market Timing

In our context, market timing is the response of the optimal asset allocation to changes in the index. Market timing is therefore visible in Figures 3 and 4 through changes in the bar charts from the left of each plot (where the index is equal to its mean minus two standard deviations) to the right (where the index is equal to its mean plus two standard deviations). In a nutshell, all of the investors we consider engage in market timing, unless they are very risk averse (with extremely high γ) or very ambiguity averse (with extremely high ε).

Otherwise, a few broad patterns emerge:

- More market timing takes place at long horizons than at short horizons. There are two reasons for this phenomenon. First, it reflects one of the basic findings from Section I, that long-horizon returns are more predictable than short-horizon returns, which is largely due to the slow mean reversion of the term premium and log dividend yield (see Figure 1). Second, short-horizon investors hold a larger percentage of their wealth in cash, which implies that the rewards from timing the market are more limited.

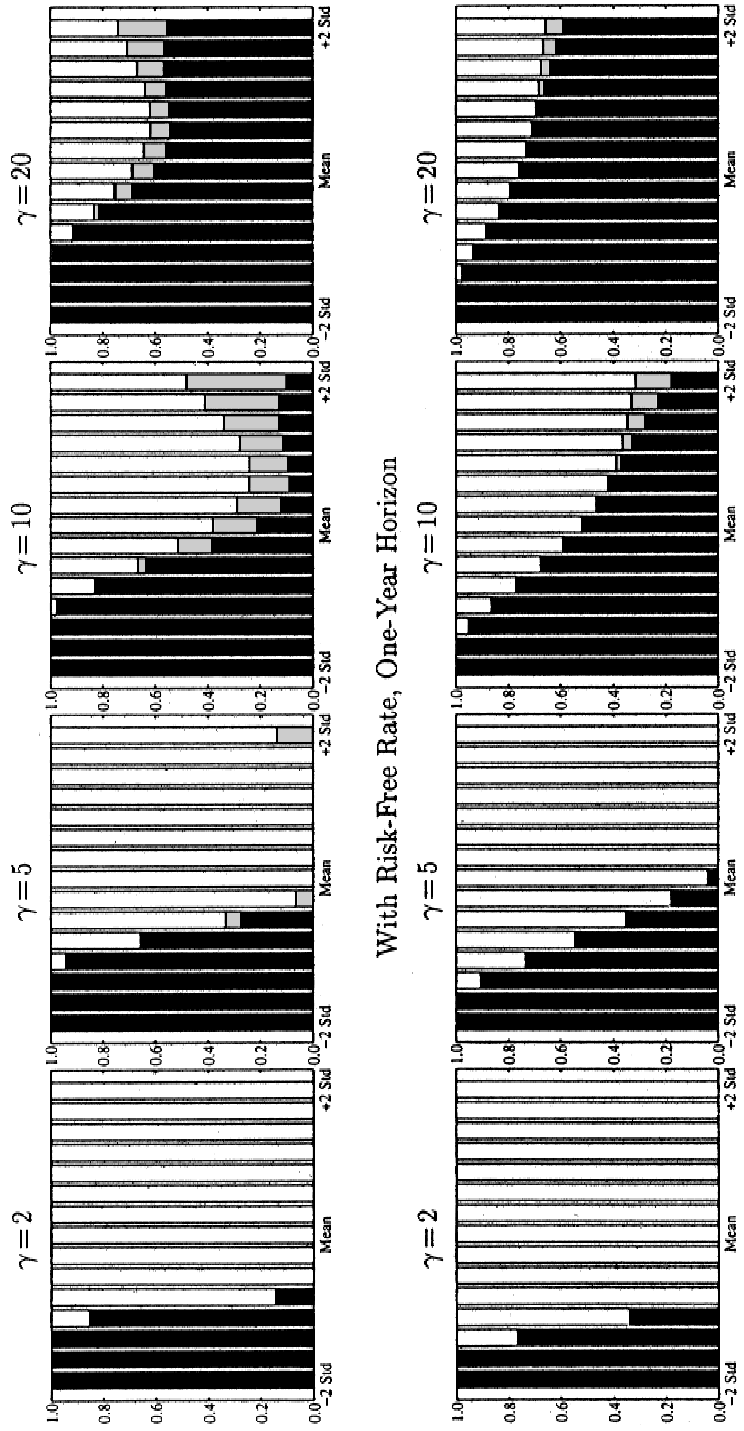
³⁴ The kernel bandwidths for the estimator of $\alpha(Z';\beta)$ are determined by leave-one-out cross-validation, which is a data-based procedure described in detail in Brandt (1999). We checked that the estimated portfolio rules are fairly insensitive to reasonable variations (± 25 percent) in the bandwidths.

- Less (more) risk-averse investors engage in more (less) market timing. It takes only a small increase in the index for an investor with $\gamma = 2$ to switch from 100 percent cash (or bonds in the case without a risk-free rate) to 100 percent stocks, and this switch occurs much before the index reaches its mean. More risk-averse investors react less abruptly to an increase in the index. Furthermore, there appears to be a limit to how far they are willing to switch from Treasury bills into stocks. This pattern is most pronounced at the one-month horizon, where a mean-variance investor with $\gamma = 20$, for example, reduces the cash holdings from 100 percent to 55 percent as the index moves from minus two standard deviations to its mean, but then keeps the 55 percent cash position even when the index rises to plus two standard deviations. Interestingly, the investor tends to substitute bonds for stocks as the index rises above its mean. As a result, the stock holdings peak around the mean of the index.³⁵
- Predictability induces even mean-variance investors with different levels of risk aversion to hold different relative positions in bonds and stocks for a given horizon. This creates a departure from the unconditional mean-variance allocation, where investors hold the same risky position, the tangency portfolio, irrespective of their risk aversion.
- The portfolio policies of mean-variance and CRRA investors with the same degree of risk aversion differ substantially, which can only be attributed to time variations in the higher order moments of the return distribution.³⁶
- Prospect theory investors are by far the most active market timers. They substitute stocks for Treasury bills more or less quickly, depending on their degree of loss aversion, as the index increases. Except at the monthly horizons, prospect-theory investors are invested 100 percent in stocks even before the index reaches its mean. At the one-month horizon and without a risk-free rate (not shown in the figure), their bond holdings reach a minimum of 25 to 40 percent around the mean of the index. Further increases in the index do not induce the investors to hold more stocks and fewer bonds because of the short-term loss potential of stocks. With a risk-free rate, prospect-theory investors introduce some bonds to their portfolio (up to 35 percent) when the index increases above its mean.

³⁵ This nonmonotonicity in the portfolio policy can be explained by the strong positive relation between the term spread and expected bond returns, as documented in Panel A of Table II, together with the fact that expected stock returns relate positively to the term spread only when the yield curve is inverted (see Boudoukh, Richardson, and Whitelaw (1997)).

³⁶ We explicitly control for the possibility of our econometric procedure inducing these differences in the policies by constraining the kernel bandwidths to be the same for mean-variance and CRRA preferences.

Panel A: Mean-Variance Investors
With Risk-Free Rate, One-Month Horizon



Panel B: Constant Relative Risk Averse Investors
With Risk-Free Rate, One-Month Horizon

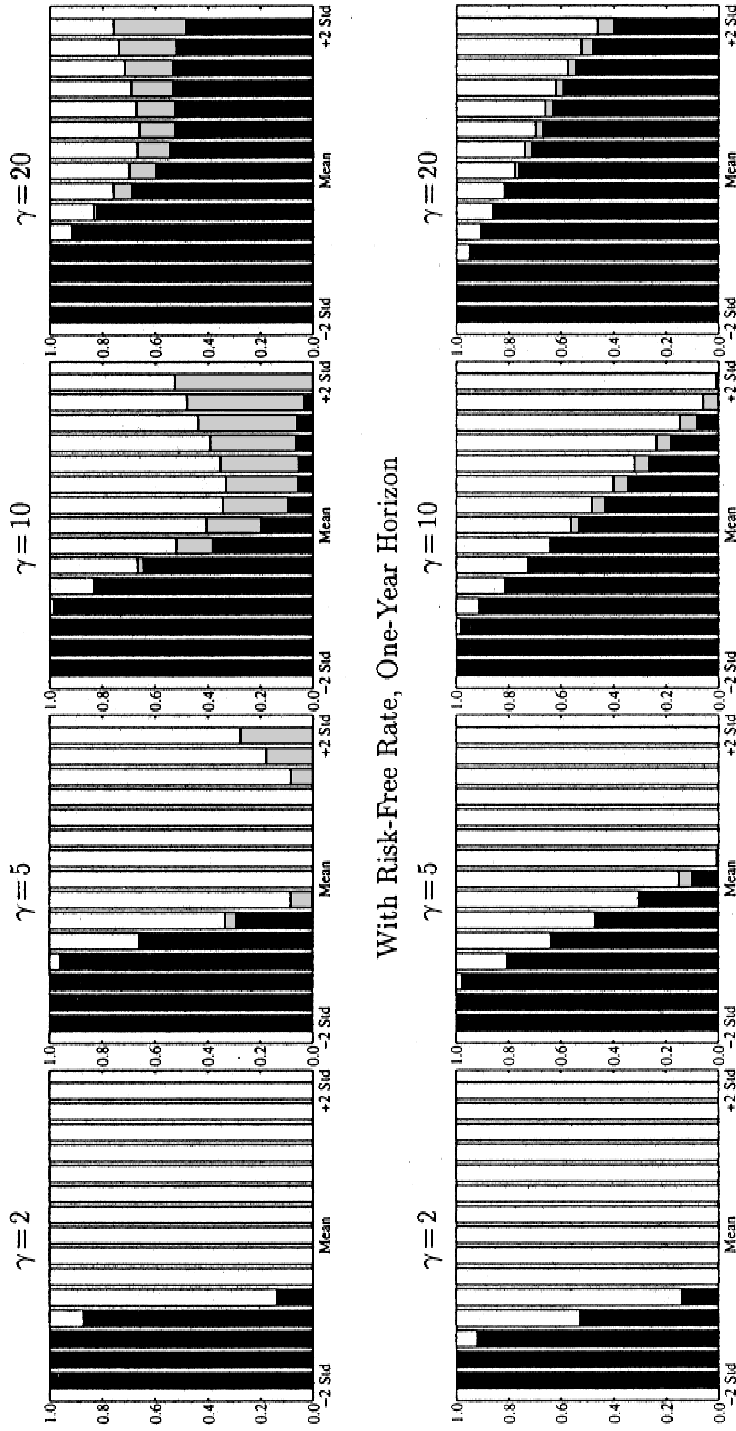
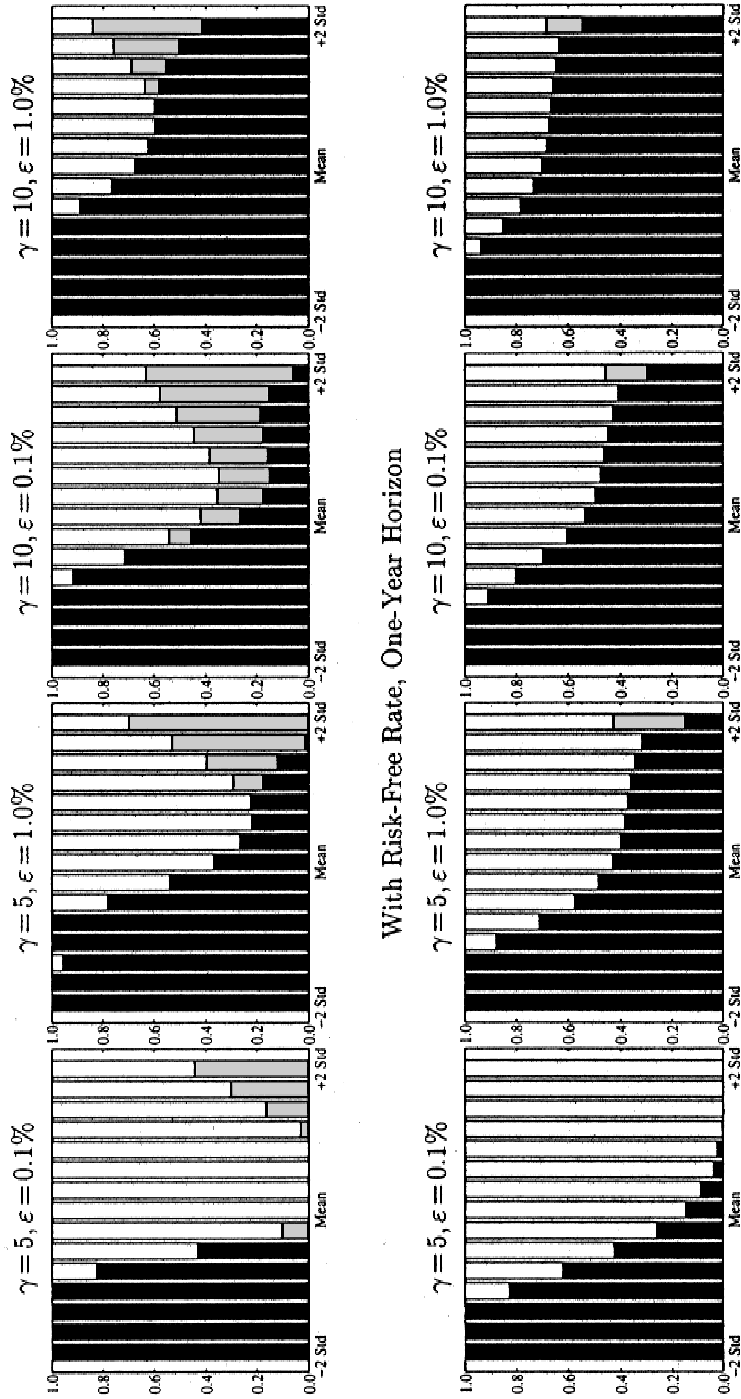



Figure 3. Conditional portfolio choice with expected utility preferences. This figure plots the conditional portfolio choice with a risk-free rate of investors with mean-variance preferences, in Panel A, or constant relative risk-aversion preferences, in Panel B. The investment horizon is one month or one year. The portfolio choice is conditional on an index of the default spread, log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index trend variable. The index coefficients are shown in Table V. The black, gray, and white areas represent the fractions of wealth allocated to Treasury bills, bonds, and stocks, respectively.

Panel A: Ambiguity Averse Investors

With Risk-Free Rate, One-Month Horizon



With Risk-Free Rate, One-Year Horizon






Figure 4. Conditional portfolio choice with nonexpected utility preferences. This figure plots the conditional portfolio choice with a risk-free rate of investors with ambiguity-aversion preferences, in Panel A, and prospect theory preferences, in Panel B. The investment horizon is 1, 3, 6, or 12 months. The portfolio choice is conditional on an index of the default spread, log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index momentum variable. The index coefficients are shown in Table VII. The black, gray, and white areas represent the fractions of wealth allocated to Treasury bills, bonds, and stocks, respectively.

Since mean-variance preferences yield an analytical expression for the optimal portfolio weights, we can compare our semiparametric estimates of the portfolio policies to plug-in policies based on individual estimates of the first and second moments of returns. This comparison allows us to measure the trade-off between the expected utility loss from combining the predictors into a single index as opposed to two indices (one for each moment) and the statistical gain (if any) from directly estimating the portfolio weights.

Table VIII compares our semiparametric estimates of the mean-variance portfolio policies in Panel A of Figure 3 to plug-in policies based on the regression estimates of the conditional moments in Panel A of Table II.³⁷ Since the linear moment structure imposed by the linear regressions makes this an unfair comparison, we also compare our estimates to nonlinear plug-in policies based on semiparametric index regressions, where each conditional moment depends nonparametrically on a *separate* linear index of the predictors (see Powell, Stock, and Stoker (1989)). We use index regressions, which still impose some parametric structure on the moments, because fully nonparametric plug-in policies, in which both moments are arbitrary functions of all four predictors, are infeasible due to the curse of dimensionality of nonparametric estimators.

Table VIII reports for each portfolio policy the mean and standard deviation of the return on wealth generated by the estimated asset allocations in our sample. It also shows the certainty equivalent rate of return ce_t , which we define as the risk-free rate of return on wealth that makes the investors indifferent between timing the market and the unconditional portfolio choice in Panel A of Table III. All entries in Table VIII are annualized.

The semiparametric policies clearly outperform the linear plug-in policies and perform about as well as the nonlinear plug-in policies, although they ignore the parametric structure of the optimal portfolio weights. The general pattern is that the semiparametric policies generate a mean return that is substantially larger than that of the linear plug-in policies but slightly smaller than that of the nonlinear plug-in policies. Furthermore, they generate a smaller standard deviation of wealth than both the linear and nonlinear plug-in policies. This implies that, according to the certainty equivalent rates of return, the semiparametric policies are dominated by the nonlinear plug-in policies for low levels of risk aversion, when investors focus on expected returns, but dominate the nonlinear plug-in policies for high levels of risk aversion, when investors are more sensitive to risk.

We draw two conclusions from Table VIII. First, comparing the linear plug-in policies to the nonlinear ones, it is clear that allowing for nonlinearities in the conditional moments is crucial. Second, comparing the semiparametric policies to the nonlinear plug-in ones, it is apparent that collapsing the conditioning information into a single index as opposed to two indices

³⁷ We use the full regressions, but the results with the best one- or two-variable regressions are very similar.

Table VIII

Semiparametric Versus Plug-in Estimates of the Conditional Portfolio Choice

This table shows the annualized mean and standard deviation of the return on wealth generated by semiparametric and plug-in estimates of the conditional portfolio choice of investors with single-period mean-variance preferences. The semiparametric estimates are based on indices of the default spread, the log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index trend variable. The index coefficients are shown in Panel A of Table V. The linear plug-in estimates are based on the individual moment regressions in Panel A of Table II. The nonlinear plug-in estimates are based on semiparametric index regressions for the individual moments. The table also shows the annualized certainty equivalent rate of return ce_t required to make the investors indifferent between the conditional portfolio choice and the unconditional portfolio choice in Panel A of Table III. The investment horizon is one month or one year.

	Without Risk-Free Rate									With Risk-Free Rate									
	Semiparametric Portfolio Weights			Plug-in Portfolio Weights						Semiparametric Portfolio Weights			Plug-in Portfolio Weights						
				Linear			Nonlinear						Linear			Nonlinear			
	Mean	Std	ce_t	Mean	Std	ce_t	Mean	Std	ce_t	Mean	Std	ce_t	Mean	Std	ce_t	Mean	Std	ce_t	
One-Month Horizon																			
$\gamma = 2$	0.141	0.129	0.017	0.137	0.119	0.014	0.143	0.131	0.014	0.149	0.117	0.028	0.143	0.104	0.025	0.150	0.118	0.029	
$\gamma = 5$	0.137	0.121	0.021	0.131	0.115	0.017	0.139	0.123	0.020	0.147	0.111	0.036	0.136	0.096	0.033	0.148	0.112	0.037	
$\gamma = 10$	0.121	0.098	0.018	0.117	0.101	0.011	0.122	0.100	0.015	0.127	0.083	0.027	0.120	0.082	0.020	0.130	0.086	0.027	
$\gamma = 20$	0.100	0.083	0.013	0.102	0.090	0.002	0.101	0.084	0.012	0.090	0.042	0.013	0.091	0.057	0.000	0.093	0.046	0.012	
One-Year Horizon																			
$\gamma = 2$	0.132	0.128	0.013	0.129	0.120	0.012	0.135	0.129	0.015	0.136	0.127	0.017	0.129	0.115	0.012	0.138	0.129	0.018	
$\gamma = 5$	0.129	0.117	0.020	0.118	0.106	0.016	0.131	0.119	0.021	0.130	0.115	0.022	0.116	0.098	0.016	0.132	0.116	0.023	
$\gamma = 10$	0.113	0.096	0.018	0.108	0.098	0.012	0.115	0.097	0.019	0.100	0.069	0.011	0.101	0.082	0.003	0.102	0.073	0.011	
$\gamma = 20$	0.100	0.086	0.022	0.100	0.094	0.007	0.102	0.087	0.021	0.076	0.035	0.006	0.088	0.068	-0.018	0.079	0.039	0.005	

(one for each moment) is inconsequential, or, if it matters, the expected utility loss is balanced by the statistical gain from directly estimating the portfolio policy.

As a more direct measure of the expected utility cost of collapsing the multidimensional information into a single index or, equivalently, of solving the constrained maximization problem (8) instead of the unconstrained maximization problem (2), we report in Table IX the differences in certainty equivalent rates of return for semiparametric policies with indices of two predictors and fully nonparametric policies (as in Brandt (1999)) with the same two predictors. Ideally, we would like to compare the semiparametric policies with all four predictors to their fully nonparametric counterparts but, as we explained in Section II.B, the nonparametric approach cannot be applied to more than two predictors due to the curse of dimensionality. However, we hope that the comparison with two predictors is representative of the results with all four predictors (especially since the index tends to be dominated by one or two predictors).

The most surprising result in Table IX is that most of the entries are positive, suggesting that collapsing the information into a single index generates an expected utility *gain*, not loss. Of course, the solution of the constrained maximization problem (8) cannot dominate that of the unconstrained maximization problem (2). Instead, the gains are due to our econometric approach. In particular, the reason is that the optimal bandwidths for the nonparametric estimator of the bivariate policy $\alpha(Z_{1,t}, Z_{2,t})$ are substantially larger than the optimal bandwidths for estimator of the univariate policy function $\alpha([Z_{1,t}, Z_{2,t}]' \beta; \beta)$.³⁸ This causes the estimates of the bivariate policies to be less responsive to the data. Since the theoretical expected utility difference cannot be positive, the difference due to the larger bandwidths can exceed 50 basis points per year (for low levels of risk aversion).

Focusing only on the negative entries in the table, where the expected utility loss is not erased by the statistical gain, it appears that without a risk-free rate there are significant interactions between the term spread and trend variable at the monthly horizon and between the dividend yield and term spread at the annual horizon. Similarly, with a risk-free rate, there appears to be significant interactions between the default spread and term spread at both the monthly and annual horizons and between the dividend yield and term spread at the annual horizon. In both cases, the expected utility loss due to the index can be as high as 50 basis points per year (for high levels of risk aversion).

To put these expected utility losses into perspective, we compare them to the expected utility gains from market timing based on the index (of all four predictors) instead of on each predictor individually. We show in Table X the

³⁸ The asymptotically optimal bandwidths of a nonparametric regression are proportional to $T^{-1/(K+4)}$, where K denotes the number of regressors (see Härdle (1990)). With $T = 528$ observations, the bandwidths for the bivariate estimator are therefore approximately 23 percent larger than those for the univariate estimator.

Table IX

Bivariate Indices Versus Pairs of Predictors for the Conditional Portfolio Choice

This tables shows the annualized certainty equivalent rate of return required to make an investor with single-period mean-variance preferences indifferent between semiparametric estimates of the conditional portfolio choice based on bivariate indices of the default spread, the log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index trend variable, and fully nonparametric estimates based on the same pairs of predictors. The index coefficients are shown in Panel A of Table V. The investment horizon is one month or one year.

	Without Risk-Free Rate						With Risk-Free Rate					
	Def LnDP	Def Term	Def Trend	LnDP Term	LnDP Trend	Term Trend	Def LnDP	Def Term	Def Trend	LnDP Term	LnDP Trend	Term Trend
	One-Month Horizon											
$\gamma = 2$	0.0042	0.0047	0.0032	0.0018	0.0007	-0.0011	0.0020	-0.0021	0.0044	0.0013	0.0025	0.0027
$\gamma = 5$	0.0034	0.0039	0.0024	0.0011	0.0001	-0.0015	0.0001	-0.0035	0.0031	0.0001	0.0015	0.0024
$\gamma = 10$	0.0021	0.0026	0.0019	0.0003	-0.0006	-0.0024	0.0006	-0.0048	0.0030	-0.0009	0.0011	0.0007
$\gamma = 20$	0.0019	0.0020	0.0013	-0.0009	-0.0017	-0.0036	0.0001	-0.0056	0.0024	-0.0014	0.0007	0.0001
	One-Year Horizon											
$\gamma = 2$	0.0034	0.0041	0.0016	-0.0033	0.0026	0.0030	0.0008	-0.0011	0.0033	-0.0020	0.0019	0.0025
$\gamma = 5$	0.0017	0.0027	0.0010	-0.0041	0.0010	0.0025	0.0003	-0.0016	0.0024	-0.0038	0.0013	0.0020
$\gamma = 10$	0.0012	0.0021	0.0010	-0.0049	0.0008	0.0017	-0.0010	-0.0031	0.0018	-0.0043	-0.0005	0.0011
$\gamma = 20$	0.0011	0.0010	0.0008	-0.0050	0.0001	0.0012	-0.0016	-0.0032	0.0009	-0.0051	-0.0010	0.0003

Table X
Indices Versus Individual Predictors for
the Conditional Portfolio Choice

This table shows the annualized certainty equivalent rate of return required to make an investor with single-period mean-variance preferences indifferent between estimates of the conditional portfolio choice based on indices of the default spread, the log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index trend variable, and estimates based on the individual predictors. The index coefficients are shown in Panel A of Table V. The investment horizon is one month or one year.

	Without Risk-Free Rate				With Risk-Free Rate			
	Def	LnDP	Term	Trend	Def	LnDP	Term	Trend
One-Month Horizon								
$\gamma = 2$	0.0178	0.0147	0.0007	0.0063	0.0287	0.0250	0.0001	0.0164
$\gamma = 5$	0.0160	0.0123	0.0017	0.0045	0.0297	0.0267	0.0023	0.0185
$\gamma = 10$	0.0160	0.0110	0.0012	0.0075	0.0229	0.0202	0.0022	0.0166
$\gamma = 20$	0.0114	0.0093	0.0016	0.0046	0.0115	0.0105	0.0026	0.0073
One-Year Horizon								
$\gamma = 2$	0.0138	0.0099	0.0072	0.0134	0.0171	0.0136	0.0068	0.0178
$\gamma = 5$	0.0195	0.0148	0.0109	0.0181	0.0202	0.0155	0.0055	0.0207
$\gamma = 10$	0.0174	0.0121	0.0111	0.0160	0.0101	0.0086	0.0011	0.0108
$\gamma = 20$	0.0199	0.0139	0.0153	0.0196	0.0056	0.0046	0.0014	0.0053

differences in certainty equivalent rates of return for the portfolio policies with indices and with individual predictors. Even when the individual predictor is the dominant variable in the index (without a risk-free rate: the dividend yield, term spread, or trend, depending on the level of risk aversion and horizon; with a risk-free rate: the term spread), the expected utility gain from the index can be as large as 1.4 percent per year. When the individual predictor is not the dominant variable in the index, the expected utility gain can be almost three percent.

Overall, Tables IX and X suggest that our index approach is sensible for general asset allocation problems where an analytical solution is not available. The expected utility gains from considering an index as opposed to individual predictors is substantial, while the expected utility loss from an index as opposed to multivariate predictors appears to be relatively small.

C.2. Horizon Effects

Horizon effects refer to differences in the portfolio choices of single-period investors with different buy-and-hold horizons (see also footnote 27). In a mean-variance setting, unconditional positive (negative) horizon effects are induced by an unconditional negative (positive) autocorrelation in returns, which makes

stocks more (less) attractive at long horizons than at short horizons. While the mean return grows linearly with the horizon, the variance of the returns also grows linearly when returns are uncorrelated, but grows less (more) than linearly when they are negatively (positively) autocorrelated. Furthermore, predictability can in some cases induce conditional mean reversion and thereby generate conditional horizon effects (e.g., Barberis (2000)).

The horizon effects for the conditional portfolio choice are most pronounced for prospect theory investors, who avoid stocks like the plague in the short-term but flock to them at longer horizons, just as in the unconditional case (see Table IV). However, these hedging demands are *not* due to unconditional or conditional mean reversion, but are rather attributed to the simple fact that losses on risky securities are realized more often at short horizons than at long horizons (see also Benartzi and Thaler (1995) and Barberis, Huang, and Santos (2001)). For prospect-theory investors, the mean reversion induced by predictability is a secondary source of horizon effects that reinforces the primary effect due to loss aversion.

For the expected utility investors, in contrast, mean reversion is the only source of horizon effects because it is the only reason for the Merton/Samuelson horizon-irrelevance result to fail in a single-period setting (recall footnote 28). The unconditional horizon effects in Table III are relatively small, reflecting the fact that returns are unconditionally not very autocorrelated. Both mean-variance and CRRA investors tend to allocate somewhat less wealth to stocks and bonds at the annual horizon than at the monthly horizon (notice the positive autocorrelations of returns in Panel A of Table I). Conditional on the indices, however, the horizon effects are much more pronounced. For example, consider the portfolio choice of a mean-variance investor with $\gamma = 10$. Although we cannot directly compare the portfolio policies across horizons, since the indices are not the same, it is immediately clear from the substantial bond positions at the monthly horizon that the portfolio choice depends dramatically on how soon utility is realized.

Unfortunately, the horizon effects we find for the expected utility investors are only partially consistent with the popular advice that investors should gradually shift from risky securities to cash as the end of their investment horizon (e.g., retirement) approaches. Consistent with this advice, sufficiently risk-averse investors hold more bonds at the monthly horizon than at the annual horizon. However, if we count bonds as risky securities, all investors allocate more wealth to risky securities at the monthly horizon. It is not clear, of course, that this pattern does not reverse at buy-and-hold horizons longer than a year or in a multiperiod context with intermediate rebalancing.³⁹ We therefore turn next to the multiperiod portfolio choice.

³⁹ Brandt (1999) also documents negative horizon effects (meaning less wealth allocated to stocks at long horizons relative to the monthly portfolio choice) for horizons between two months and one year. However, he finds positive horizon effects for buy-and-hold investments lasting between 1 and 10 years.

C.3. Hedging Demands

Hedging demands arise in a multiperiod (i.e., long-horizon with intermediate rebalancing) portfolio choice when investors can partially hedge against future changes in investment opportunities by deviating from their single-period portfolio choice. To measure the extent to which hedging demands arise in our setting, we solve for the multiperiod index coefficients β_n and portfolio weights $\alpha_n(Z'_t\beta_n; \beta_n)$ of investors with an annual rebalancing frequency and horizons n ranging from 1 to 10 years. We define hedging demands as $\alpha_n(Z'_t\beta_n; \beta_n) - \alpha_1(Z'_t\beta_1; \beta_1)$, where β_1 and $\alpha_1(\cdot)$ are the single-period index and portfolio weights, respectively.

Besides studying hedging demands, we are interested in seeing how the index composition β_n varies with the horizon n . Intuitively, the index serves two purposes in a multiperiod portfolio choice. First, the index forecasts the features of the return distribution that are important to the investor, just as in a single-period context. Second, it predicts next period's index realization, which in turn summarizes next period's investment opportunities, to help the investor hedge against future changes in investment opportunities. The first purpose of the index is independent of the horizon but the second one is not by definition.

We estimate β_n and $\alpha_n(Z'_t\beta_n; \beta_n)$ recursively, starting with the single-period problem at the end of the horizon and recursing backward, period by period, to the initial 10-period problem. Each period n we replace in our estimator the single-period objective function $v(W_{t+1})$ and its derivative $\partial v(W_{t+1})/\partial W_{t+1}$ with the multiperiod value function $J(W_{t+1}, Z'_{t+1}\beta_{n-1}, n-1)$ and its derivative $\partial J(W_{t+1}, Z'_{t+1}\beta_{n-1}, n-1)/\partial W_{t+1}$. The value function represents the expected utility of making optimal portfolio choices for the remaining $n-1$ periods, as a function of next period's wealth W_{t+1} and conditional on next period's index realization $Z'_{t+1}\beta_{n-1}$. We obtain the value function by evaluating the maximized objective function of the $(n-1)$ -period problem over a fine grid of initial wealth levels and index realizations.

We study only CRRA and ambiguity-aversion investors in this multiperiod context. The reason is that mean-variance utility is inherently a single-period objective that generates nonsensical multiperiod policies and prospect theory yields trivial results, since with annual rebalancing, the short-sale constraints are binding for all horizons and preference parameters.

Panels A and B of Table XI show estimates of the multiperiod index coefficients for CRRA and ambiguity-aversion investors with horizons of 1, 2, 5, and 10 years. Since the portfolio is rebalanced once a year, the one-year horizon results correspond to the single-period index coefficients in Panel B of Table V and Panel A of Table VII. To preserve space, we only present results for the cases $\gamma = \{5, 10\}$ and $\varepsilon = 0.005$. The results for different degrees of relative risk aversion and ambiguity aversion are quantitatively similar.

It is clear that the optimal index composition depends at least to some extent on the investor's horizon. Consider, for example, the portfolio choice with a risk-free rate of a CRRA investor with $\gamma = 10$. The single-period

Table XI
Index Composition for Multiperiod Portfolio Choice

This table shows estimates of the optimal index composition for the conditional portfolio choice of investors with multiperiod preferences:

$$\text{Panel A: } v(W_{t+n}) = \frac{W_{t+n}^{1-\gamma}}{1-\gamma} \quad \text{Panel B: } v(W_{t+n}) = (1-\varepsilon) \frac{W_{t+n}^{1-\gamma}}{1-\gamma} + \varepsilon \inf_R \frac{W_{t+n}^{1-\gamma}}{1-\gamma},$$

where W_{t+n} denotes the terminal wealth in n periods. Each index is a linear combination of the default spread Def, the log dividend-to-price ratio of the S&P 500 index LnDP, the term spread Term, and the S&P index momentum variable Trend. The portfolio is rebalanced annually and the investment horizon ranges from 1 to 10 years. The Without Risk-Free Rate estimates are for a portfolio choice between stocks and bonds and the With Risk-Free Rate estimates are for a portfolio choice between stocks, bonds, and Treasury bills. In both cases, the optimization is subject to the short-sale constraints $0 \leq \alpha \leq 1$.

	Without Risk-Free Rate				With Risk-Free Rate			
	Def	LnDP	Term	Trend	Def	LnDP	Term	Trend
Panel A: Constant Relative Risk-Averse Investors								
$\gamma = 5$								
$n = 1$	-0.124	0.462	0.361	0.053	0.087	0.180	0.630	-0.104
$n = 2$	-0.098	0.413	0.422	0.067	0.064	0.194	0.663	-0.079
$n = 5$	-0.089	0.395	0.430	0.086	0.049	0.197	0.679	-0.075
$n = 10$	-0.088	0.394	0.430	0.088	0.044	0.198	0.683	-0.075
$\gamma = 10$								
$n = 1$	-0.102	0.450	0.340	0.108	0.054	0.184	0.655	-0.107
$n = 2$	-0.088	0.441	0.360	0.111	0.041	0.192	0.674	-0.093
$n = 5$	-0.084	0.421	0.368	0.127	0.039	0.208	0.695	-0.058
$n = 10$	-0.084	0.420	0.368	0.128	0.022	0.209	0.721	-0.048
Panel B: Ambiguity-Averse Investors								
$\gamma = 5$ and $\varepsilon = 0.5\%$								
$n = 1$	0.012	0.512	0.367	0.109	0.041	0.218	0.592	-0.149
$n = 2$	0.015	0.461	0.405	0.119	0.028	0.229	0.608	-0.135
$n = 5$	0.016	0.442	0.419	0.123	0.020	0.232	0.617	-0.131
$n = 10$	0.016	0.438	0.420	0.126	0.016	0.234	0.622	-0.128
$\gamma = 10$ and $\varepsilon = 0.5\%$								
$n = 1$	0.024	0.521	0.294	0.161	0.035	0.213	0.606	-0.146
$n = 2$	0.036	0.509	0.289	0.166	0.022	0.222	0.625	-0.131
$n = 5$	0.042	0.487	0.298	0.173	0.016	0.229	0.627	-0.128
$n = 10$	0.043	0.483	0.300	0.174	0.011	0.231	0.633	-0.125

index has loadings of 0.054, 0.218, 0.592, and -0.149 on the default spread, dividend yield, term spread, and trend, respectively. For a 10-year portfolio choice, the corresponding loadings are 0.022, 0.209, 0.721, and -0.048 . The general pattern is that long-horizon investors focus even more on the term

spread than single-period investors. Also, when the portfolio choice includes a risk-free rate, long-term investors focus more on the dividend yield than single-period investors. When it does not, they instead pay more attention to the trend.

Panels A and B of Figure 5 plot the portfolio policies as a function of the horizon (along the horizontal axis) and the index (across plots). The three plots in each row present the portfolio choice conditional on the index being equal to its 25th percentile, median, and 75th percentile, respectively. The black, grey, and white areas in each plot represent the estimated allocations of wealth to Treasury bills, bonds, and stocks as a function of the horizon.

Relative to the recent and ever-growing literature on hedging demands (e.g., Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1998, 1999), Balduzzi and Lynch (1999), Brandt (1999), Chacko and Viceira (1999), Barberis (2000), and Lynch (2000)), we surprisingly find relatively small hedging demands. Even more striking, we find *negative* hedging demands, in the sense that long-term investors allocate less wealth to risky securities than short-term investors, instead of the consensus of positive hedging demands.

To understand our results, recall that for hedging demands to arise, the returns have to be contemporaneously correlated with changes in the investment opportunities. Looking at Panel B of Table I, we notice that at the one-year frequency, stock returns are (by construction) strongly negatively correlated with changes in the log dividend yield and strongly positively correlated with changes in the trend variable (-0.36 and 0.31 , respectively). They are weakly positively correlated with the term spread. Bond returns, in turn, are only weakly positively correlated with both the default and term spreads.

Most of the literature on hedging demands focuses on the dividend yield as the leading state variable. Since stock returns are strongly negatively correlated with innovations in the dividend yield, multiperiod investors in these settings allocate more wealth to stocks than single-period investors, since stock returns tend to be high when the dividend yield falls and the investment opportunities next period worsen. In our setting, however, the choice of the leading state variable is endogenously determined to be the term spread. Since returns are not as correlated with changes in the term spread as they are with changes in the dividend yield, our indices generate little hedging demands. In fact, the weak positive correlation of both stock and bond returns with changes in the term spread induces a slight reduction in risky investments at longer horizons.

Another noteworthy feature of our multiperiod policies is that at horizons of more than five years even CRRA investors do not invest in bonds. The reason is that the negative correlation of the bond returns with changes in the term spread (or the indices) leads investors to reduce their bond holdings as a hedge against changes in the investment opportunities. Therefore, our results confirm, at least qualitatively, the popular investment advice that short-term investors should hold a larger fraction of bonds relative to stocks than long-term investors (see also Canner, Mankiw, and Weil (1997)).

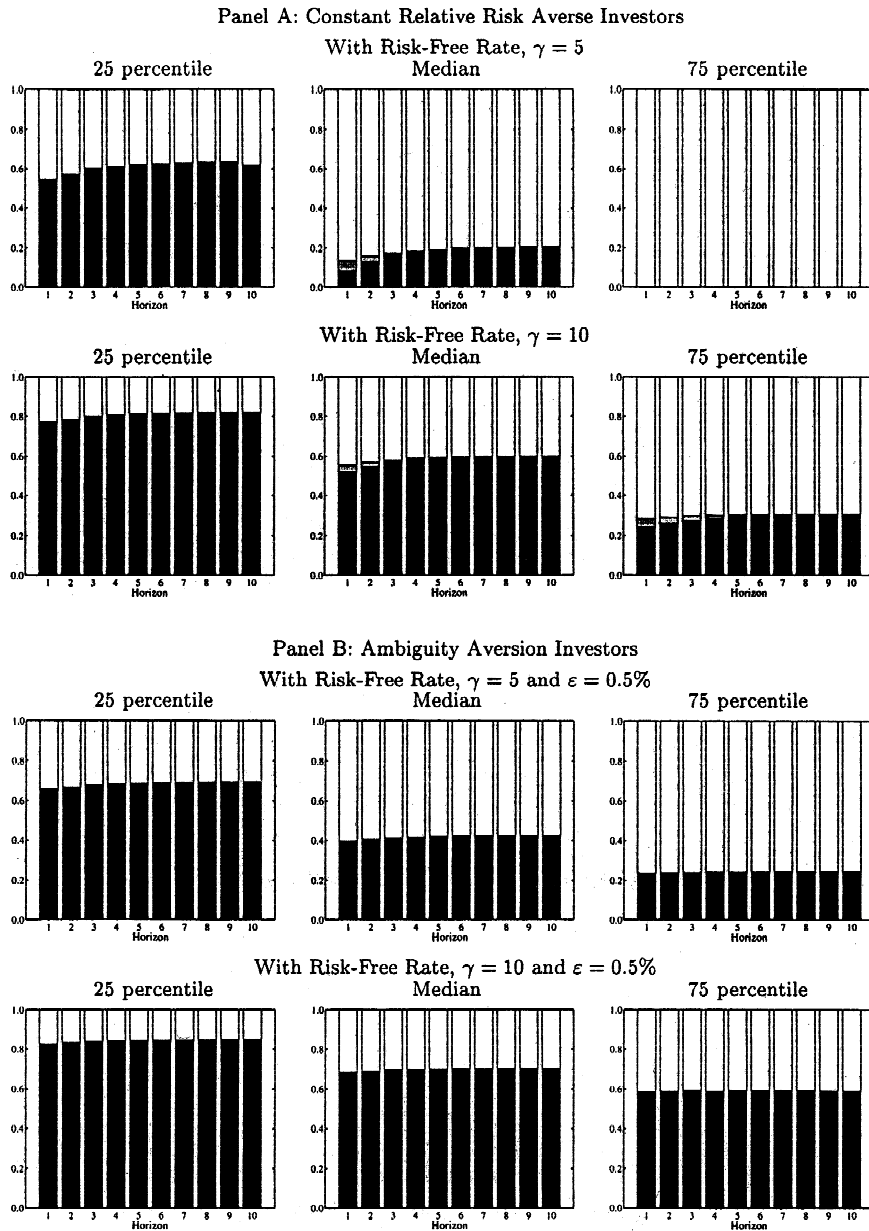


Figure 5. Conditional multiperiod portfolio choice. This figure plots the conditional multiperiod portfolio choices with a risk-free rate of investors with constant relative risk-aversion preferences, in Panel A, and ambiguity-aversion preferences, in Panel B. The horizon ranges from 1 year to 10 years with an annual rebalancing frequency. The portfolio choice is conditional on an index of the default spread, log dividend-to-price ratio of the S&P 500 index, the term spread, and the S&P index momentum variable. The index coefficients are shown in Table XI. The black, gray, and white areas represent the fractions of wealth allocated to Treasury bills, bonds, and stocks, respectively.

V. Conclusions

We showed how to use indices to select and combine predictive variables for the conditional portfolio choice of an investor with given preferences and investment horizons. We also characterized the optimal single-period and multiperiod asset allocations between asset classes for a variety of investor preferences as a function of their respective indices. This enables us to give investors sensible advice, we hope, regarding the economic variables to which they should pay attention and the sensitivity of their investments to the (sometimes mixed) signal these variables send. As Campbell and Viceira (1999) and others note, financial planners who provide advice to individual investors have received surprisingly little guidance from academics. We hope to have taken another step towards providing such advice.

Beyond the introduction of an index, an important feature of this paper is that we considered investors with a variety of preferences, including both expected and nonexpected utility objectives, which are viewed by psychologists as more realistic descriptions of human decision making under uncertainty. Along the way, we identified features of these preferences that lead investors to select different combinations of variables and react differently to them.

At least four stylized facts emerged from our empirical analysis:

- The term premium is a ubiquitous variable in our indices, appearing significantly across all preferences, investment horizons, and rebalancing frequencies. To a lesser extent, but fairly consistently, the S&P index momentum variable we termed “Trend” enters our indices at short horizons, whereas the dividend yield is the second most important variable at long horizons. The default risk premium generally records the lowest index loadings.
- All investors, when presented with their index of investment opportunities, find it optimal to engage in significant market timing.
- Horizon effects are most pronounced for prospect-theory investors, who find the likelihood of stock losses at short horizons to be prohibitively costly. For investors who are not subject to loss aversion, the relative lack of returns autocorrelation translates into relatively small horizon effects.
- Hedging demands are weak and negative because stocks do not provide a good hedge for innovations in the index (i.e., in the investment opportunities). However, the index coefficients vary with the horizons.

Finally, it should be noted that in the second half of the 1990s (which is only partially included in our sample), many of the patterns in expected returns previously identified have weakened: The dividend-to-price ratio stubbornly predicted negative returns for the stock market which never materialized; the size factor has all but disappeared; and value stocks earned substantially lower returns than predicted on the basis of previously estimated regressions. This should serve as a reminder that the magnitude of predictability in returns is small and subject to a tremendous amount of

noise, especially at short horizons. Our analysis incorporated this fact through the standard errors of the estimated indices and optimal portfolio rules. The size of these standard errors merely reflects the fact that predicting returns is difficult. Nevertheless, our estimates allow us to identify a few important stylized facts about the optimal selection and combination of predictors, as well as about market timing, horizon effects, and hedging demands.

REFERENCES

- Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent, 1999, Risk and robustness in equilibrium, Working paper, Stanford University.
- Andrews, Donald W. K., 1999, Estimation when a parameter is on a boundary, *Econometrica* 67, 1341–1383.
- Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath, 1998, Coherent measures of risk, Working paper, ETH Zurich.
- Avramov, Doron, 1999, Stock return predictability and model uncertainty, Working paper, University of Pennsylvania.
- Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Barberis, Nicholas, 2000, Investing for the long-run when returns are predictable, *Journal of Finance* 55, 225–264.
- Barberis, Nicholas, Ming Huang, and Jesus Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics*, 116, 1–53.
- Barberis, Nicholas, and Ming Huang, 2001, Mental accounting, loss aversion, and individual stock returns, *Journal of Finance* 56, 1247–1291.
- Basak, Suleyman, 1995, A general equilibrium model of portfolio insurance, *Review of Financial Studies* 8, 1059–1090.
- Basak, Suleyman, and Alex Shapiro, 1998, Value-at-risk based risk management: Optimal policies and asset prices, Working paper, University of Pennsylvania.
- Becker, Selwyn W., and Fred O. Brownson, 1964, What price ambiguity? Or the role of ambiguity in decision making, *Journal of Political Economy* 72, 62–73.
- Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics* 110, 73–92.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Bollerslev, Tim, Robert F. Engle, and Jeffrey M. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116–131.
- Boudoukh, Jacob, Mathew P. Richardson, and Robert F. Whitelaw, 1997, Nonlinearities in the relation between the equity risk premium and the term structure, *Management Science* 43, 371–385.
- Brandt, Michael W., 1999, Estimating portfolio and consumption choice: A conditional Euler equations approach, *Journal of Finance* 54, 1609–1646.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Camerer, Colin F., 1995, Individual decision making, in John H. Kagel and Alvin E. Roth, eds.: *The Handbook of Experimental Economics* (Princeton University Press, Princeton, NJ).
- Camerer, Colin F., and Howard C. Kunreuther, 1989, Experimental markets for insurance, *Journal of Risk and Uncertainty* 2, 265–300.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., and Robert J. Shiller, 1988a, The dividend price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y., and Robert J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

- Campbell, John Y., and Luis M. Viceira, 1998, Who should buy long-term bonds? Working paper, Harvard University.
- Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.
- Canner, Niko, N. Gregory Mankiw, and David N. Weil, 1997, An asset allocation puzzle, *American Economic Review* 87, 181–191.
- Chacko, George, and Luis M. Viceira, 1999, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, Working paper, Harvard University.
- Cochrane, John H., 1991, Production-based asset pricing and the link between stock returns and economic fluctuations, *Journal of Finance* 46, 209–235.
- Cochrane, John H., 1999, New facts in finance, *Economic Perspectives* 33, 36–58.
- Curley, Sean P., and J. Frank Yates, 1985, The center and range of the probability interval as factors affecting ambiguity preferences, *Organizational Behavior and Human Decision Processes* 36, 272–287.
- Curley, Sean P., and J. Frank Yates, 1989, An empirical evaluation of descriptive models of ambiguity reactions in choice situations, *Journal of Mathematical Psychology* 33, 397–427.
- Dow, James, and Sérgio Ribeiro Da Costa Werlang, 1992, Uncertainty aversion, risk aversion, and the optimal choice of portfolios, *Econometrica* 60, 197–204.
- Ellsberg, Daniel, 1961, Risk, ambiguity, and the Savage axioms, *Quarterly Journal of Economics* 75, 643–669.
- Engle, Robert F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation, *Econometrica* 22, 427–432.
- Epstein, Larry G., and Tan Wang, 1994, Asset pricing under Knightian uncertainty, *Econometrica* 62, 283–322.
- Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Fama, Eugene F., and G. William Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics* 5, 115–146.
- Ferson, Wayne E., and Campbell R. Harvey, 1991, The variation of economic risk premiums, *Journal of Political Economy* 99, 385–415.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Gilboa, Itzhak, and David Schmeidler, 1989, Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* 18, 141–153.
- Grossman, Sanford J., and Zhongquan Zhou, 1996, Equilibrium analysis of portfolio insurance, *Journal of Finance* 51, 1379–1403.
- Hansen, Lars Peter, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1053.
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini, 1999, Robust permanent income and pricing, *Review of Economic Studies* 66, 873–928.
- Härdle, Wolfgang, 1990, *Applied Nonparametric Regression* (Cambridge University Press, New York).
- Harvey, Campbell R., 1989, Time-varying conditional covariance in tests of asset pricing models, *Journal of Financial Economics* 24, 289–317.
- Harvey, Campbell R., 1991, The specification of conditional expectations, Working paper, Duke University.
- Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.
- Kandel, Shmuel, and Robert F. Stambaugh, 1996, On the predictability of stock returns: An asset allocation perspective, *Journal of Finance* 51, 385–424.
- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.

- Knight, Frank K., 1921, *Risk, Uncertainty, and Profit* (Houghton Mifflin, Boston, MA).
- Lamont, Owen, 1998, Earnings and expected returns, *Journal of Finance* 53, 1563–1588.
- Lettau, Martin, and Sydney Ludvigson, 2000, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance*, forthcoming.
- Liu, Wen-Fang, 1998, Heterogeneous agent economies with Knightian uncertainty, Working paper, University of Washington.
- Liu, Wen-Fang, 1999, Saving, portfolio, and investment decisions: Observable implications with Knightian uncertainty, Working paper, University of Washington.
- Lynch, Anthony W., 2000, Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability, *Journal of Financial Economics*, forthcoming.
- Maenhout, Pascal J., 1999, Robust portfolio rules and asset pricing, Working paper, Harvard University.
- Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The consumption of stockholders and non-stockholders, *Journal of Financial Economics* 29, 97–112.
- Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *Review of Economics and Statistics* 51, 247–257.
- Moran, Patrick A. P., 1971, Maximum-likelihood estimation in non-standard conditions, *Proceedings of the Cambridge Philosophical Society* 70, 441–450.
- Pontiff, Jeffrey, and Lawrence D. Schall, 1998, Book-to-market ratios as predictors of market returns, *Journal of Financial Economics* 49, 141–160.
- Powell, James L., James H. Stock, and Thomas M. Stoker, 1989, Semiparametric estimation of index coefficients, *Econometrica* 57, 1403–1430.
- Samuelson, Paul A., 1969, Lifetime portfolio selection by dynamic stochastic programming, *Review of Economics and Statistics* 51, 238–246.
- Sarin, K. Rakesh, and Martin Weber, 1993, The effect of ambiguity in a market setting, *Journal of Risk and Uncertainty* 5, 389–407.
- Schwert, G. William, 1989, Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Sharpe, William F., 1998, Morningstar's risk-adjusted ratings, *Financial Analysts Journal* 54, 21–33.
- Shefrin, Hershey, and Meir Statman, 1994, Behavioral portfolio theory, Working paper, Santa Clara University.
- Shumway, Tyler, 1997, Explaining returns with loss aversion, Working paper, University of Michigan.
- Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.
- Tversky, Amos, and Peter Wakker, 1995, Risk attitudes and decision weights, *Econometrica* 63, 1255–1280.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.

Discussion

JESSICA A. WACHTER*

THE CURRENT PAPER BY YACINE AIT-SAHALIA and Michael W. Brandt (henceforth AB) addresses two issues that are of central concern in portfolio choice: How can portfolio advice be made realistic while remaining tractable? How can

* Stern School of Business, New York University. I thank John Heaton and Anthony Lynch for helpful comments.