



Should Long-Term Investors Time Volatility?☆

Alan Moreira^{a,*}, Tyler Muir^b^aSimon Graduate School of Business, Rochester, NY 1462, Box 270100, USA^bUCLA Anderson School of Management, 110 Westwood Plaza, Los Angeles, CA 90024, USA

ARTICLE INFO

Article history:

Received 28 August 2017

Revised 12 January 2018

Accepted 9 February 2018

Available online 28 September 2018

JEL classification:

G11

Keywords:

Volatility

Portfolio choice

Market timing

Mean reversion

ABSTRACT

A long-term investor who ignores variation in volatility gives up the equivalent of 2.4% of wealth per year. This result holds for a wide range of parameters that are consistent with US stock market data, and it is robust to estimation uncertainty. We propose and test a new channel, the volatility composition channel, for how investment horizon interacts with volatility timing. Investors respond substantially less to volatility variation if the amount of mean reversion in returns disproportionately increases with volatility and also if mean reversion happens quickly. We find that these conditions are unlikely to hold in the data.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

In October 2008, stock market volatility spiked to very high levels at the same time that a steep drop in the stock

market made valuation ratios appear attractive. This presented investors with a dilemma: was the potential increase in expected return implicit in the lower valuation ratios sufficient compensation for the unprecedented levels of volatility? Answering this question is relevant not only for extreme periods of volatility like those seen in 2008 but also much more generally as the negative correlation between volatility and realized returns is a salient feature of stock markets.

A common view among many practitioners and academics is that the low valuations provide more than adequate compensation for investors that are able to withstand the heightened short-term market volatility, i.e. investors with long investment horizons. For example, in response to the high volatility period in 2008 John Cochrane suggests: “If you . . . have a longer horizon than the average, it makes sense to buy,” a sentiment echoed by Warren Buffet around the same time.¹ This view has also translated into standard financial advice. For example, Vanguard—a leading mutual fund company—argue that

☆ University of Rochester Simon Graduate School of Business and UCLA Anderson School of Management. We thank Toni Whited (the editor) and Jessica Wachter (the referee) for many useful comments. We also thank John Campbell, John Cochrane, Ian Dew-Becker, William Goetzmann, Valentin Haddad, Ben Hebert, Jon Ingersoll, Ravi Jagannathan, Serhiy Kosak, Lars Lochstoer, Hanno Lustig, Justin Murfin, Stefan Nagel, Lubos Pastor, Myron Scholes, Ivan Shaliastovich, Ken Singleton, Tuomo Vuoltenahoo, Michael Weber, Lu Zhang, and participants at Yale SOM, University of Rochester Simon Business School, Washington University's Olin Business School, Duke University's Fuqua School of Business, Robert H. Smith School of Business at University of Maryland, Fisher College of Business at Ohio State University, UNC Kenan-Flagler Business School, UCSD Rady School of Management, USC Marshall School of Business, Imperial College School of Business, Questrom School of Business at Boston University, UCLA Anderson, Stanford GSB, the Chicago Booth Empirical Asset Pricing Conference, Arrowstreet Capital, and Norges Bank Investment Management for comments. We especially thank Nick Barberis for many useful discussions.

* Corresponding author.

E-mail addresses: alan.moreira@simon.rochester.edu (A. Moreira), tyler.muir@anderson.ucla.edu (T. Muir).¹ See Cochrane (2008a) and Buffett (2008).

long-term oriented investors are better off ignoring movements in volatility.² The argument is that, since volatility is typically associated with market downturns, and downturns are attractive buying opportunities for those with long horizons, it is not wise to sell when volatility spikes. Further, and most importantly, because of mean reversion in stock returns, investors with long horizons should not view increases in volatility as an increase in risk—the idea is that an increase in volatility makes stock prices more uncertain tomorrow but not more uncertain over long horizons that these investors care about. Thus, the argument is that periods of high volatility can be much more attractive to long-horizon investors relative to short-horizon investors. Empirically, there is direct evidence that investors seem to follow this advice: investors with longer horizons do in fact react much less to volatility changes than those with short horizons.³

The goal of our paper is to provide a structural quantitative framework for thinking through these facts and for evaluating the conventional wisdom. While an extensive literature on portfolio choice shows that long-horizon investors may indeed behave very differently than those with a short horizon, this paper is the first to comprehensively study the portfolio response to volatility when both expected returns and volatility are allowed to vary over time in a way that is consistent with the main empirical facts about the US market portfolio.

Specifically, we answer three questions: (1) how much volatility timing should long-term investors do, if any; (2) what are the utility benefits of volatility timing; and (3) what features of the return process are critical for understanding volatility timing? We study the portfolio problem of a long-lived investor that allocates her wealth between a risk-less and a risky asset in an environment where both volatility and expected returns are time varying and where the parameters governing these processes are estimated using US data from 1925 to 2016. We provide comprehensive, quantitative answers that show the effects of investor horizon as well as which parameters of the return dynamics are most relevant for our conclusions.

We begin our analysis by estimating a rich model for the dynamics of excess stock returns using simulated method of moments (SMM) and the last 90 years of stock return data. Our process for returns allows for (potentially independent) time variation in both volatility and expected returns. Allowing for both features is essential to capture the common argument that high volatility periods are “buying opportunities” for long-horizon investors. It also enables us to fit the most salient features of the US data, i.e., that both expected returns and volatility vary

significantly over time (Campbell and Shiller, 1988; Schwert, 1989) but are not strongly related to each other at short horizons, despite the fact that increases in volatility are associated with market downturns (Glosten, Jagannathan and Runkle, 1993). Finally, it allows us to compare the utility benefits from timing variation in volatility to the long literature on the utility benefits of timing expected returns (for example, Campbell and Viceira, 1999; Barberis, 2000; Wachter, 2002). We use both the parameter point estimates and the associated estimation uncertainty to consider a range of parameters governing the return process that are likely given the data.

In keeping with the portfolio choice literature, our analysis is partial equilibrium. That is, we leave open the question of what determines the joint dynamics of expected returns and volatility that we recover from the data, but given these dynamics, we study how an investor with standard preferences should invest. Specifically, we focus on the portfolio problem of an infinite-lived investor with recursive preferences (Epstein and Zin, 1989) with unit elasticity of intertemporal substitution (EIS). These preferences allow us to conveniently control the horizon of the investor, i.e., the timing of her consumption, while at the same time keeping the environment stationary. These preferences should capture individuals and institutions that target a constant expenditure share of their wealth (e.g., university endowments, sovereign wealth funds, or pension funds).

Our main finding is that the optimal response of a long-term investor strongly depends on the composition of volatility shocks. If the variance of expected and unexpected returns vary proportionally, as commonly assumed in the empirical literature (Campbell et al., 2016), long-horizon investors should substantially decrease their weight in the stock market after an increase in volatility. A strategy that ignores variation in volatility leads to large utility losses. This result fills an important gap in the literature. While previous work (Chacko and Viceira, 2005; Liu, 2007) shows that the volatility hedging demand is small relative to the myopic demand, their framework cannot speak directly to the argument that high volatility periods are “buying opportunities” for long-horizon investors. Specifically, because expected returns must be an affine function of stock market variance in their framework, their model will either imply too little time variation in expected returns or that the risk-return trade-off is too strong. Both possibilities are inconsistent with the data. Therefore, our paper is the first to study the portfolio of a long-term investor in framework that is flexible enough to fit the most important facts about the aggregate stock market.

We then allow for time variation in the composition of volatility shocks. Specifically, we show that if variation in volatility is driven by the volatility of expected returns, then the optimal portfolio can become much less sensitive to volatility variation than implied by variation in the risk-return trade-off. Intuitively, prices become more volatile only in the short term but not more volatile in the long term because the higher volatility of expected returns leads to an increased degree of mean reversion. Therefore when volatility is very low, returns are entirely

² For the Vanguard reference see “What to do during market volatility? Perhaps nothing.” See <https://personal.vanguard.com/us/insights/article/market-volatility-082015>. In the Internet Appendix, we reference similar advice from Fidelity, the New York Times, and many other sources. The Internet Appendix is available on the Journal of Financial Economics website.

³ Nagel et al. (2016) show that during the 2008–2009 period, older investors (those with a shorter investment horizon) sold much more heavily in response to increases in volatility compared to younger investors (those with a longer investment horizon), even when controlling for a host of other factors.

driven by permanent shocks, i.e., there is no mean reversion in returns and little return predictability. In that case, long-term and short-term investors will choose similar stock allocations. However, in high volatility times, stock returns become strongly mean reverting because the volatility of expected returns increases. In these periods, a short-term investor sees a reduction in the risk-return trade-off and wants to sell. The long-term investor weighs two effects: the myopic desire to sell due to the lower risk-return trade-off, but also the large hedging demand that now arises from mean reversion. Thus, the long-term investor will react less strongly to the increase in volatility in this case because it is accompanied by an increase in the degree of mean reversion in returns. This mechanism can be thought as the time-series counterpart of the cross-sectional “bad” beta in [Campbell and Vuolteenaho \(2004\)](#), and it is consistent with [Cochrane \(2008a\)](#) and [Buffett's \(2008\)](#) argument that the huge spike in volatility in the fall of 2008 was mostly about “short-term volatility,” i.e., increases in volatility that do not change the long-term distribution of asset prices.

Quantitatively, we show that this mechanism requires not only that volatility variation to be completely driven by variation in the volatility of mean-reverting shocks but also that mean reversion in returns happens fairly quickly. That is, shocks to expected returns are not too persistent. We show that absent these two conditions, long-term investors still find it optimal to time volatility.

We then use the data to put empirical bounds on the amount of mean reversion in returns coming from high versus low volatility periods. Specifically, we show empirically that the autocorrelation of stock market returns at a one-year horizon is not different in high versus low volatility periods. In the model, if there was substantially more mean reversion in returns during high volatility periods, and this mean reversion happened fairly quickly, then there would indeed be a large disparity in return autocorrelations, with a more negative return autocorrelations during high volatility periods. Importantly, our test only rejects the joint hypothesis that volatility is driven by the volatility of expected return shocks and expected returns have low persistence. This restriction on how much mean reversion and volatility interact is sufficient to determine that a long-term investor should still respond relatively aggressively to changes in volatility.

Our results are important for investors such as pension funds, endowments, sovereign wealth funds, individuals saving for retirement, or other long-term investors as they provide guidance in how to optimally respond to volatility. Specifically, we show that the optimal timing strategy is closely approximated by a strategy that is affine in the mean-variance portfolio $\frac{\mu_t}{\sigma_t^2}$. This simple strategy achieves nearly the same utility as the fully optimal strategy so that the optimal portfolio can be approximated fairly accurately by the myopic portfolio plus a constant weight investment in the buy-and-hold portfolio. This affine form for the portfolio strategy holds for a wide range of parameters that are likely given the data.

We next evaluate the utility benefits from volatility timing, where we define a volatility timing strategy as a

strategy that only uses conditional information on volatility but not expected returns. Specifically, we restrict ourselves to constant weight combinations of the buy-and-hold portfolio and the volatility managed portfolio from [Moreira and Muir \(2017\)](#), i.e., strategies that are affine in $\frac{\mu_t}{\sigma_t^2}$ where $\bar{\mu}$ sets the expected return to its unconditional mean. We compare the utility of this strategy to the fully optimal strategy, w^* , that conditions on both expected returns and volatility and also to the naive buy-and-hold strategy that does not do any timing. We find very large gains from volatility timing, which we measure using an annualized per period fee the investor is willing to pay to switch from a static buy-and-hold portfolio to a volatility timing portfolio. For the baseline estimates (no variation in the composition of volatility), the naive buy-and-hold investor would be willing to pay a 2.36% per period annualized fee to volatility time. In wealth equivalent terms this a 60% increase relative to the buy-and-hold portfolio. These gains are about 80% of the total gain of switching from the buy-and-hold strategy \bar{w} to the fully optimal strategy w^* (i.e., a strategy that also conditions on expected returns as well as volatility). Thus, ignoring variation in volatility is very costly, and the benefits to timing volatility are significantly larger than the benefits to timing expected returns.

We then show that these gains survive even after we take into account the parameter uncertainty implied by our estimation procedure for return dynamics,⁴ estimate the model at different frequencies to allow for the possibility that there are persistent movements in volatility, and use different proxies for expected returns.

What does the conventional view highlighted earlier miss? First, while it is true that valuation ratios tend to go down when volatility increases, signaling higher expected returns ahead, the conventional view misses that this increase in expected returns is much more persistent than the increase in volatility. Investors can avoid the short-term increase in volatility by first reducing their exposure to equities when volatility initially increases and capture the increase in expected returns by coming back to the market as volatility comes down. In other words, valuation ratios will remain attractive for an extended period of time while volatility is shorter lived, and the long-term investor can exploit this large difference in persistence. Second, time variation in expected returns implies that return volatility over short holding periods will overstate the amount of risk that an investor with a long-horizon faces. That is, a fraction of return variation is purely transitory and therefore is not risk for an investor with long enough investment horizons. The conventional view correctly uses this fact to argue that long-term investors should invest more in equities than short-term investors ([Campbell and Viceira, 1999](#)). However, this level effect on portfolios is often also interpreted to mean long-term investors should care less about volatility variation. This is a fallacy because long-term investors already have a higher exposure to equities due to mean reversion. We show that it is only if

⁴ This result contrasts to the gains from expected return timing that are much more sensitive to parameter uncertainty (see [Barberis, 2000](#); [Pástor and Stambaugh, 2012](#)).

there is relatively more mean reversion in returns when volatility increases that long-horizon investors will react differently. That is, the composition of permanent versus transitory shocks to returns must shift as volatility changes (we call this the volatility composition channel).

Our paper builds on the prolific literature on long-term asset allocation. Starting with the seminal work of Samuelson (1969) and Merton (1971), this literature has studied carefully the implications of mean reversion for portfolio choice. Campbell and Viceira (1999), Barberis (2000), and Wachter (2002) study the optimal portfolio problem in the presence of time-varying expected returns. The key result is that the presence of mean reversion in market returns imply investors with longer horizons should invest more in the stock market. In addition, this literature argues that market dips are good buying opportunities (Campbell and Viceira, 1999). An important caveat is that parameter uncertainty can attenuate these horizon effects (see Barberis, 2000; Xia, 2001).

Less studied, but we think equally important, is time variation in second moments. Chacko and Viceira (2005), Liu (2007), and Zhou and Zhu (2012) study variation in volatility, and Buraschi, Porchia and Trojani (2010) study variation in correlations. This literature finds only modest deviations from myopic behavior. The absence of large hedging demands in these papers suggest that volatility timing as in Moreira and Muir (2017) is desirable, and investment horizon effects are not first order. However, these papers abstract from (independent) variation in expected returns. Thus, they cannot speak directly to the conventional wisdom that volatility spikes are mostly “buying opportunities” or that return volatility is mostly due to transitory shocks that mean revert over the long run. It is precisely this gap that this paper fills. Consistent with the intuition behind the traditional view, we show that there is an important interaction between volatility and expected return variation through the volatility composition channel.

Related papers that account for both volatility and expected returns include Shanken and Tamayo (2012), Collin-Dufresne and Lochstoer (2017), and Johannes, Korteweg and Polson (2014). Collin-Dufresne and Lochstoer (2017) have a time-varying risk-return relationship in a general equilibrium setting and point out that long-term investors only want to buy at “low prices” if effective risk aversion, rather than risk itself, has increased to cause the fall in prices. Johannes et al. (2014) solve a Bayesian problem that accounts for time-varying volatility when forming out of sample expected return forecasts.

Finally, we build on the results in Fleming, Kirby and Ostdiek (2001), Fleming, Kirby and Ostdiek (2003), and Moreira and Muir (2017) who study volatility timing empirically in the context of a short-term mean-variance investor. We go well beyond the results in these papers by solving for the optimal portfolio for long investment horizons, estimating utility gains for long-horizon investors along with associated uncertainty surrounding such gains, considering parameter uncertainty in estimating expected returns and volatility (which we show is important), and studying the interaction of volatility and mean reversion in returns and showing a novel channel through which this

feature matters for an investors’ horizon, something argued to be important by academics and practitioners and something that appears empirically relevant in understanding who responds most to volatility (Nagel et al., 2016).

The paper proceeds as follows. Section 2 describes the process for returns and investor preferences. Section 3 analyzes the optimal portfolio and associated utility gains from volatility timing. Section 4 allows for time variation in the composition of volatility shocks. Section 5 studies robustness of our results to parameter uncertainty. Section 6 contains extensions to our main results. Section 7 concludes.

2. The portfolio problem

We study the problem of a long-horizon investor and investigate how much they should adjust their portfolio to changes in volatility.

2.1. Investment opportunity set

We assume there is a riskless bond that pays a constant interest rate r and a risky asset S_t , with dynamics given by

$$\frac{dS_t}{S_t} = (r + \mu_t)dt + \sigma_t dB_t^S, \quad (1)$$

where S_t is the value of a portfolio fully invested in the asset and reinvests all dividends. We model expected (excess) returns as an autoregressive process with stochastic volatility,

$$d\mu_t = \kappa_\mu(\bar{\mu} - \mu_t)dt + \sigma_\mu \sigma_t dB_t^\mu. \quad (2)$$

Notice that this means that the volatility of shocks to expected returns scale up and down proportionally with shocks to realized returns. In later sections, we consider cases in which we break this proportionality. We write log volatility $f(\sigma_t^2) = \ln(\sigma_t^2 - \underline{\sigma}^2)$ as an auto-regressive process with constant volatility,

$$df(\sigma_t^2) = \kappa_\sigma(\bar{f} - f(\sigma_t^2))dt + \nu_\sigma dB_t^\sigma, \quad (3)$$

where the parameter $\underline{\sigma}^2$ controls the lower bound of the volatility process. This lower bound is important in eliminating near arbitrage opportunities (i.e., infinite Sharpe ratios). Our assumption about a lognormal volatility process should not be seen as crucial, although it allows for easier solutions in our numerical exercise and a better fit for the unconditional distribution of volatility. Results using a square root process (Heston, 1993; Cox, Ingersoll and Ross, 1985) for volatility along the lines of Chacko and Viceira (2005) are similar.

Shocks to realized returns, expected returns, and volatility are thus captured by the Brownian motions dB_t^S , dB_t^μ , and dB_t^σ . We now specify the correlation of these shocks. First, we impose

$$E_t[dB_t^\mu dB_t^S] = -\frac{\sigma_\mu}{\kappa_\mu}. \quad (4)$$

Note that the correlation between expected returns and realized returns is not a free parameter. The correlation $-\frac{\sigma_\mu}{\kappa_\mu}$ implies that shocks to expected returns induce an immediate change in prices so that in the long

run, it exactly offsets expected return innovations, i.e., it imposes that expected return shocks have no effect on the long-run value of the asset. We make this choice to emphasize that we want to consider transitory shocks to returns that have no long-run impact; however, we also note that if one freely estimates this correlation in the data, one recovers roughly this value (Cochrane, 2008b). Hence, it is not an overly restrictive assumption. This correlation also defines the share of “discount rate shocks” that drive returns—that is, when the correlation is one, then all variation in returns is driven by discount rate shocks, and when it is zero, expected return shocks play no role. We label this correlation $E_t[dB_t^\mu dB_t^S] = -\alpha_\mu^{1/2}$. We thus write $\sigma_\mu = \alpha_\mu^{1/2} \kappa_\mu$ and focus on estimating α_μ and κ_μ in the data, as these parameters have direct economic interpretations as the share and persistence of discount rate shocks.

We next specify the remaining correlations

$$E_t[dB_t^\sigma dB_t^\mu] = \rho_{\sigma,\mu}, \tag{5}$$

$$E_t[dB_t^\sigma dB_t^{S-\mu}] = \rho_{\sigma,S}, \tag{6}$$

where $dB_t^{S-\mu}$ captures shocks to returns that are orthogonal to discount rate shocks, i.e., cash flow shocks.⁵ The correlation between volatility and expected and realized returns are free parameters that must satisfy $\rho_{\sigma,S}^2 + \rho_{\sigma,\mu}^2 \leq 1$ so that all Brownian motions are well defined.⁶ Finally, we set the unconditional mean of the log volatility process $f(\sigma_t^2)$ to $\bar{f} = \ln(\bar{\sigma}^2 - \sigma^2) - \frac{v_\sigma^2}{2\kappa_\sigma}$.

This parametrization leads to a natural interpretation of the parameters: $\bar{\mu}$ is the average expected excess return of the risky asset, $\bar{\sigma}^2$ is the average conditional variance of returns, v_σ^2 is the conditional variance of log variance, $\rho_{\sigma,\mu}$ controls the covariance between variance and discount rate shocks, and $\rho_{\sigma,S}$ controls the covariance between variance and cash flow shocks (return innovations uncorrelated to innovations in discount rates). Throughout, we adopt the language from the literature (Campbell and Shiller, 1988; Campbell, 1996; Campbell and Vuolteenaho, 2004), using “cash flow shocks” to denote permanent shocks to returns that are uncorrelated to shocks that affect expected returns. Next, α_μ denotes the discount rate share of return variation.

The stochastic environment described by Eqs. (1) and (2) allows for variation in volatility, variation in expected returns (i.e., mean reversion in returns), and flexible time-series relation between expected returns and volatility ($\rho_{\sigma,\mu}$). The latter governs the risk-return trade-off relationship between variance and the risk premium. In the Internet Appendix, we discuss even more sophisticated and flexible ways of modeling this relationship. In particular, we discuss allowing expected returns to more directly depend on volatility by having two frequencies for expected

returns: a shorter frequency component that is related to volatility and a slower moving component (specifically, we write $\mu_t = x_t + b\sigma_t^2$, where b governs the risk-return relation, and x and σ^2 are allowed to move at different frequencies). It turns out, however, that because the risk-return relation is empirically weak, we do not lose much by incorporating a less rich relationship between expected returns and variance. In fact, we will show in our estimation that the current model is able to capture the essential empirical moments relating risk and return in the time series, meaning our modeling of the risk-return trade-off is appropriate. Finally, in later sections we also allow for variation in the composition of volatility shocks (that is, we consider the case in which α_μ is not constant but varies over time). This will allow for variation in the share of return volatility due to discount rate shocks.

Together, these ingredients are novel and essential to study the optimal response to volatility variation. Earlier work on portfolio choice has studied expected return variation, volatility variation, or volatility variation with a constant risk-return trade-off. Examples of work that study volatility timing in a dynamic environment are Chacko and Viceira (2005) and Liu (2007). But these papers do not study the interaction of discount rate and volatility shocks that are the basis for the conventional view that long-horizon investors should ignore volatility variation.

2.2. Estimation of parameters

We estimate the model using SMM (Duffie and Singleton, 1993) and use the estimated parameters in Table 1 to discuss the model implications for portfolio choice. We use an optimal weight matrix as in Bazzdrusch et al. (2018) who find an SMM estimator with the optimal weight matrix has much better finite sample properties compared to an identity matrix. We separately calibrate the real riskless rate ($r = 1\%$) that reflects the US experience in the postwar sample. We also calibrate the volatility lower bound ($\underline{\sigma} = 7\%$) based on the data.⁷

Our goal is for the model to match the key dynamic properties of US stock returns shown in the empirical finance literature. With that in mind, we use the US market excess return from 1926–2015 from Ken French. We describe our estimation in more detail in the Internet Appendix. We briefly describe our choice of moments and how these moments identify the parameters of interest.

We use daily data to construct a monthly series of realized volatility, RV , which we use to match the properties of volatility in the model. Specifically, we will simulate the model at daily frequency and compute realized volatility in the same manner as in the data, thus the true volatility process is unobserved. We then aggregate to monthly frequency in the data and model to match all moments. Finally, we bring in additional monthly data on the US dividend price ratio from Robert Shiller to match moments

⁵ The Brownian motion representing the cash flow shocks can be constructed as $dB_t^{S-\mu} = \frac{dB_t^S + \alpha_\mu dB_t^\mu}{\sqrt{1-\alpha_\mu}}$.

⁶ We specify these correlations as constant. In particular, we don't consider time variation in the correlation between volatility and expected returns. See Collin-Dufresne and Lochstoer (2017) for a case in which this time variation plays a role in a general equilibrium model for long-term portfolio choice.

⁷ Here we use that the minimum of the Chicago Board Options Exchange Volatility Index (VIX) from 1990–2015 is 10%, so our 7% for the longer 90-year sample is reasonable. Note we use VIX to calibrate this number, rather than realized volatility, because realized volatility is noisy and hence would not properly measure a lower bound for true volatility.

Table 1

Estimates for the stochastic processes.

Panel A provides the calibrated parameters; Panel C provides the estimated parameters. Panel B provides the matched moments in the data and in the model. Parameters are estimated using SMM. We use monthly data on the market portfolio for 1926 to 2015. The standard errors uses the influence function approach described in [Erickson and Whited \(2002\)](#) to estimate the joint distribution of moments from the data. In Panel C, we also use the influence function approach to estimate the joint distribution of estimated parameters. Specifically, the parameter standard errors are constructed as follows: we sample 1,000 moment realizations from the estimated moment distribution, reestimate the model for each of the moment realizations, and then compute the parameter standard deviations across moment realizations. Note that $RV_t = \sum_{d=1/22}^1 (R_{t+d}^e - (\sum_{d=1/22}^1 R_{t+d}^e / 22))^2$ is the daily realized variance in month t , where R_{t+d}^e is the excess return on the market on date $t + d$. $R_{t,t+T}^e$ is the cumulative excess return on the market between dates t and T . The alpha of the volatility managed portfolio $\alpha(SR_t \rightarrow R_t)$ is the intercept of a regression of a volatility managed (excess) portfolio on the (excess) portfolio itself (see [Moreira and Muir, 2017](#), for details). See text for more details.

Panel A: Parameters				
Parameter	Description	Value		
r	Riskfree rate	0.01		
$\underline{\sigma}$	Volatility lower bound	0.07		
Panel B: Estimated moments				
Moment	Description	Data	Model	S.E.
$mean(R_{t,t+1})$	Equity premium (ann.)	0.078	0.074	0.017
$mean(RV_t)$	Avg. realized variance (ann.)	0.028	0.032	0.006
$var(\log(RV_t))$	Variance of realized variance	1.109	1.075	0.175
$\beta(\log(RV_{t+1}) \rightarrow \log(RV_t))$	Autocorr. of realized variance	0.722	0.734	0.046
$R_{sqrd}(R_{t,t+12} \rightarrow pd_t)$	Predict. regression R-sq (1 year)	0.082	0.052	0.042
$R_{sqrd}(R_{t,t+60} \rightarrow pd_t)$	Predict. regression R-sq (5 years)	0.234	0.169	0.079
$\beta(R_{t,t+1} \rightarrow \Delta \log(RV_{t+1}))$	Realized variance β of the market	-0.023	-0.023	0.002
$\alpha(SR_t \rightarrow R_t)$	Alpha of vol. managed portfolio (ann.)	0.048	0.051	0.017
$\beta(R_{t,t+1} \rightarrow RV_t)$	Conditional risk-return trade-off	-0.186	0.319	0.880
Panel C: Estimated parameters				
Parameter	Description	Estimate	S.E.	
$\bar{\sigma}^2$	Average variance	0.034	0.006	
κ_σ	Variance persistence	2.961	0.613	
ν_σ	Variance of variance	4.347	0.469	
$\bar{\mu}$	Equity premium	0.069	0.016	
α_μ	Share of discount rate shocks	0.449	0.143	
κ_μ	Discount rate shocks persistence	0.129	0.108	
$\rho_{\sigma,\mu}$	Corr vol and DR shocks	0.669	0.131	
$\rho_{\sigma,S-\mu}$	Corr vol and CF shocks	0.743	0.154	

related to time variation in expected returns and return predictability. We emphasize that we do not model dividends or firm price-dividend ratios in the model, but we use the predictive regressions in the data to infer the variability and dynamics of the conditional expected return process.

Our model for volatility is univariate and therefore is bound to miss the dynamics of volatility at some frequency. Because we are interested not in the level of hedging demands per se, but the overall gains from volatility timing, we choose to fit dynamics of volatility at the monthly frequency. In [Section 6.1](#) we discuss this choice in detail and show that our results are robust to targeting lower frequency movements in volatility.

2.2.1. Moments

We choose the vector of target moments to be informative about the parameters. Our target moments are (1) average excess return on the risky asset (equity premium); (2) average realized monthly variance; (3) the variance of the logarithm of realized monthly variance; (4)

the autocorrelation coefficient of the logarithm of realized monthly variance; (5–6) the R -squared of a predictability regression of one-year and five-year-ahead returns on the price-dividend ratio; (7) the regression coefficient of realized returns on changes in log realized variance (contemporaneous); and (8–9) the alpha of the volatility managed market portfolio on the market portfolio (see [Moreira and Muir, 2017](#)) as well as the risk-return trade-off estimated from regressions of future returns on realized variance. The alpha of the volatility managed portfolio is defined by the regression $\frac{c}{RV_t^2} R_{t+1} = \alpha + \beta R_{t+1} + \varepsilon_{t+1}$, where the alpha measures whether one can increase Sharpe ratios through volatility timing. [Moreira and Muir \(2017\)](#) show this alpha measures the strength of the risk-return trade-off over time but is a sharper measure than standard forecasting regressions.

While there is not an exact one-to-one mapping between moments and parameters, the link between parameters and moments is intuitive, and the moments are very informative about the parameters of interest. Average realized monthly variance identifies $\bar{\sigma}^2$ and the average

excess stock return identifies $\bar{\mu}$. The autocorrelation of volatility and the variance of log realized variance identify ν_σ and κ_σ . These moments imply that the estimated volatility process is highly volatile but not very persistent. The return predictability R -squares at one-month and five-year horizons identify α_μ , the discount rate share, and κ_μ , the volatility and persistence of discount rate shocks. Intuitively, the one-year R -square implies the share of discount rate shocks must be large and the fact that five-year R -squares are substantially larger implies that expected returns must be highly persistent. The covariance between realized returns and changes in log variance, the volatility managed alpha identify, and the conditional risk-return trade-off identify $\rho_{\sigma,\mu}$ and $\rho_{\sigma,S-\mu}$. In the data, the large negative correlation between volatility innovations and realized returns implies that $\rho_{\sigma,\mu}^2 + \rho_{\sigma,S-\mu}^2$ is close to one. The alpha of the volatility managed portfolio disciplines the extent to which this comovement is due to a correlation between discount rates and volatility shocks. In the data, a portfolio that takes less risk when volatility is high generates a large Sharpe ratio, implying that the comovement between volatility and discount rate shocks is not strong (see Moreira and Muir, 2017).

2.2.2. Estimation results

Table 1 reports targeted moments in the model and in the data. Overall the model matches the data extremely well and matches the key empirical facts on the dynamics of stock returns shown in the finance literature. Specifically, all model implied moments are within one standard deviation of the data moments. In particular, the estimated volatility process is highly volatile, so there is substantial time-variation in conditional volatility (Schwert, 1989). Expected returns are quite variable, i.e., discount rate volatility is an important component of stock market volatility (Campbell and Shiller, 1988), and these discount rate shocks are strongly correlated with volatility shocks (French, Schwert and Stambaugh, 1987). That is, increases in volatility are associated with low realized returns and increases in expected returns. However, this correlation does little to dampen variation in the risk-return trade-off because shocks to expected returns are much more persistent than shocks to volatility, and also the correlation between volatility and expected returns, while large, is not equal to one. Thus, the model is able to produce positive volatility managed alphas consistent with Moreira and Muir (2017) because the model, like the data, does not feature an overly strong risk-return trade-off. That is, consistent with a long literature, there is some risk-return trade-off in the data, but it appears to be fairly weak (French et al., 1987; Glosten et al., 1993; Lettau and Ludvigson, 2003). Further, the estimation procedure matches the volatility managed alpha more closely than the risk-return trade-off regression due to the large standard error associated with the risk-return trade-off regression coefficient (this means the optimal weight matrix will put less weight on this moment). Notice, the model features a stronger conditional risk-return trade-off compared to the data (which shows a negative relation and would suggest even higher benefits to timing volatility) though again, this is poorly estimated in the data. Thus, taken to-

gether, our process for returns matches the key empirical features about the properties of expected returns, conditional volatility, and realized returns shown by a long literature in asset pricing.

We also report standard errors for the estimated moments and parameters in Table 1. We follow the influence function technique from Erickson and Whited (2002) to estimate variance-covariance matrix of target moments, and we use the inverse of this variance-covariance matrix as the SMM weighting matrix. We report the moment standard errors in Panel C of Table 1. We note that the model matches the moments well such that all model implied moments are within one standard error of the empirical counterpart. The moment that is farthest from the data is the risk-return trade-off regression coefficient (the coefficient is positive in the model implying a positive risk-return trade-off, but the empirical point estimate is negative). However, we note that this moment is poorly estimated empirically. Instead, the volatility managed alpha provides a more precise measure of the risk-return trade-off and the model matches this moment closely (Moreira and Muir, 2017). Note that because we use the optimal weighting matrix, the estimation will pay less attention to moments that are less precisely estimated. Hence, the model is able to match the key properties of the risk-return trade-off in the data as well as match the key dynamics of both volatility and expected returns.

We report the parameter estimates and standard errors in Panel B of Table 1. Consistent with the large literature on market timing, the dynamics of expected returns are the least well estimated aspect of our model. This estimation uncertainty will play a role in later sections where we consider that the investor may not know the true parameters in making his portfolio decision. In contrast, almost all other parameters are well estimated in the model. The parameters imply a highly volatile and moderately persistent process for volatility, a volatile and very persistent process for expected returns with a large equity premium (although again, there is uncertainty about the degree of expected return variation), and a strong negative correlation between changes in volatility and realized stock returns.

2.3. Preferences and optimization problem

Investors preferences are described by Epstein and Zin (1989) utility, a generalization of the more standard constant relative risk aversion (CRRA) preferences that separates risk aversion from elasticity of intertemporal substitution. We adopt the Duffie and Epstein (1992) continuous time implementation and focus on the case of unit elasticity of substitution, though we relax this in later sections:

$$J_t = E_t \left[\int_t^\infty f(C_s, J_s) ds \right], \tag{7}$$

where $f(C_t, J_t)$ is an aggregator of current consumption and continuation utility that takes the form

$$f(C, J) = \beta(1 - \gamma)J \times \left[\log(C) - \frac{\log((1 - \gamma)J)}{1 - \gamma} \right], \tag{8}$$

where β is rate of time preference; γ is the coefficient of relative risk aversion. The unit elasticity of substitution is

convenient for our purposes because it allow us to directly vary the investor horizon in a way that is independent of the attractiveness of the investment opportunity set. Specifically, $1 - \exp(-\beta)$ is the share of investors wealth consumed within one year. Thus $1/\beta$ can be thought as the horizon of the investor. In Section 6.2 we consider alternative preference specifications.

Let W_t denote the investor wealth and w_t the allocation to the risky asset, then the budget constraint can be written as

$$\frac{dW_t}{W_t} = w_t \left(\frac{dS_t}{S_t} - rdt \right) + rdt - \frac{C_t}{W_t} dt. \quad (9)$$

The investor maximizes utility subject to his intertemporal budget constraint, Eq. (9), and the evolution of state variables, Eq. (2).

3. Analysis

Our aim is to quantify the optimal amount of volatility timing for a realistic portfolio problem in which an investor decides how much to invest in the market portfolio and in a riskless asset. We solve for the investor value function numerically and study how the optimal portfolio should respond to changes in volatility. Our analysis is quantitative in nature and it is therefore important that our model for returns described in Eq. (1)–(2) fit the dynamics of returns in the data.

In the baseline case we study the problem of an investor with a 20-year horizon ($\beta = 1/20$) and risk aversion of five, and we investigate the sensitivity of our results to these parameter choices.

3.1. Solution

The optimization problem has three state variables: the investor's wealth plus the investment opportunity set state variables μ_t and σ_t .

The Bellman equation for this problem is standard

$$\begin{aligned} 0 = & \sup_{w,C} \{ f(C_t, J_t) + [w_t \mu_t W_t + rW_t - C_t] J_W \\ & + \frac{1}{2} w_t^2 W_t^2 J_{WW} \sigma_t^2 + w_t W_t (- J_{W\mu} \sqrt{\alpha_\mu} \sigma_t^2 \\ & + J_{W\sigma} \nu_\sigma (\rho_{\sigma,S} \sqrt{1 - \alpha_\mu} - \rho_{\sigma,\mu} \sqrt{\alpha_\mu}) \sigma_t) \\ & + J_\mu \kappa_\mu (\bar{\mu} - \mu_t) + J_\sigma \kappa_\sigma (\bar{\sigma} - \sigma_t^2) \\ & + \frac{1}{2} J_{\mu\mu} \alpha_\mu \kappa_\mu^2 \sigma_t^2 + \frac{1}{2} J_{\sigma\sigma} \nu_\sigma^2 \\ & + J_{\mu\sigma} \rho_{\sigma,\mu} \nu_\sigma \alpha_\mu^{1/2} \kappa_\mu \sigma_t \}, \end{aligned} \quad (10)$$

where we omit the argument on $J_t = J(W_t, \mu_t, \sigma_t)$ for convenience. It is well known that the value function for this type of problem is of the form $J(W, \mu_t, \sigma_t) = \frac{W^{1-\gamma}}{1-\gamma} e^{(1-\gamma)g(\mu_t, \sigma_t)}$ for $\gamma > 1$. Plugging this form in Eq. (10) we obtain that the optimal consumption to wealth ratio is constant, $C_t = \beta W_t$ and the optimal portfolio weight satisfies

$$w^*(\mu_t, \sigma_t^2) = w^m(\mu_t, \sigma_t^2) + H^\mu(\mu_t, \sigma_t^2) + H^\sigma(\mu_t, \sigma_t^2), \quad (11)$$

where the first term in Eq. (11) is the myopic portfolio weight

$$w^m(\mu_t, \sigma_t^2) = \frac{1}{\gamma} \frac{\mu_t}{\sigma_t^2}. \quad (12)$$

It calls the investor to scale up his position on the risky asset according to the strength of the risk-return trade-off and her coefficient of relative risk aversion. This is also the optimal portfolio weight of a short-horizon mean-variance investor (or log investor). The additional terms in Eq. (11) are Mertonian hedging demands (Merton, 1971) that are given by

$$H^\mu(\mu_t, \sigma_t^2) = \frac{1-\gamma}{\gamma} g_\mu(\mu_t, \sigma_t^2) \sqrt{\alpha_\mu} \quad (13)$$

$$\begin{aligned} H^\sigma(\mu_t, \sigma_t^2) = & \frac{1-\gamma}{\gamma} g_\sigma(\mu_t, \sigma_t^2) \\ & \times \frac{(-\rho_{\sigma,S} \sqrt{1 - \alpha_\mu} + \rho_{\sigma,\mu} \sqrt{\alpha_\mu}) \nu_\sigma}{\sigma_t}. \end{aligned} \quad (14)$$

The hedging demands H arises because a long-horizon investor is concerned with the overall distribution of her consumption and not only the short-term dynamics of her wealth. Changes in the risky asset expected returns or volatility lead to changes in the distribution of the investor wealth, resulting in a demand for assets that hedge these changes. To the extent that the risky asset is correlated with changes in the opportunity set, this demand for hedging impacts the investor's position in the risky asset.

This hedging effect means that a long-horizon investor might behave very differently from a short-term oriented investor. An increase in volatility might generate an increase in the hedging demand that is enough to completely offset the reduction in exposure due to the myopic demand, i.e., it might be that long-horizon investors should just ignore time variation in volatility, in line with the argument articulated in Cochrane (2008a).⁸

The empirical fact that expected returns increase after low return realizations, $dB^\mu dB^\sigma < 0$ makes investment in the risky asset a natural investment hedge for changes in expected returns. This effect has been studied extensively in the literature (e.g. Campbell and Viceira, 1999; Barberis, 2000; Wachter, 2002), which has shown that when $\gamma > 1$, this hedging demand leads a long-horizon investor to have a larger average position in the risky asset.

A similar hedging demand arises due to changes in volatility, though with the opposite sign. The fact that increases in volatility tend to be associated with low return realizations also implies that the risky asset comoves with the investment opportunity set. Specifically, when $\gamma > 1$

⁸ "And what about volatility?... [I]f you were happy with a 50/50 portfolio with an expected return of 7% and 15% volatility, 50% volatility means you should hold only 4.5% of your portfolio in stocks!... [E]xpected returns would need to rise from 7% per year to 78% per year to justify a 50/50 allocation with 50% volatility... The answer to this paradox is that the standard formula is wrong... Stocks act a lot like long-term bonds... If bond prices go down more, bond yields and long-run returns will rise just enough that you face no long-run risk... [T]he same logic explains why you can ignore "short-run" volatility in stock markets." (Cochrane, 2008a)

the hedging demand due to volatility pushes investors to hold smaller positions in the risky asset (Chacko and Viceira, 2005; Buraschi et al., 2010).

The direction of these hedging demands follows from the interaction between changes in the Sharpe ratio and the coefficient of relative risk aversion. An investor that is more conservative than a log investor ($\gamma > 1$) wants to transfer resources from states where the opportunity set is better to states where it is worse. Because expected and realized returns are negatively correlated, a positive tilt toward the risky asset implies her wealth increases following reduction in the Sharpe ratio due to a reduction in expected returns. Symmetrically, because volatility and realized returns are negatively related, a negative tilt toward the risky asset implies her wealth increases following a reduction in the Sharpe ratio due to an increase in volatility.

Investment horizon, together with the persistence of the state variables ($\kappa_\mu, \kappa_\sigma$), shapes the strength of the hedging demand through the sensitivity of the value function to changes in the state variables (g_μ, g_σ). Intuitively, persistent changes to the state impact the investment opportunity set for longer, and this impact is larger for investors with a longer horizon, which are naturally more exposed to persistent changes in the opportunity set. As a result, the value function is typically more sensitive to the state variables, and the resulting hedging demands are larger for investors with longer horizons. Here the unit elasticity of intertemporal substitution (EIS) is particularly convenient, as the patience parameter β directly controls the effective horizon of the investor, i.e., the timing of their consumption.

3.2. Optimal portfolios

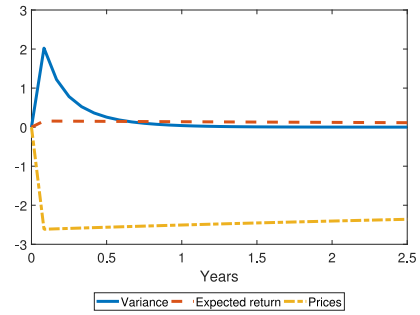
We use projection methods to solve for $g(\mu_t, \sigma_t)$ (see the Internet Appendix for details). Here we discuss the implications of $g(\mu_t, \sigma_t)$ for the optimal portfolio in terms of an impulse response function (Fig. 1) and then as a function of the state variables (2).

It is illuminating to discuss our results by contrasting the optimal choices of long- and short-term investors. In the top panels of Fig. 1, we start by showing the response of variance, expected returns, and prices to a one standard deviation shock to variance, and then show how long- and short-term investors respond.

Expected returns go up in response to a volatility shock, though quantitatively this increase is small. This is due to the high correlation between realized returns and volatility innovations present in the data. Thus, an innovation in volatility is correlated with innovations in expected returns. Nevertheless, the myopic and the optimal portfolio go down sharply and in parallel. This means that two investors with the same risk aversion, but different horizons, will reduce the fraction of their wealth allocated to stocks by exactly the same amount. In this sense, horizon has no impact on how investors should respond to changes in volatility.

There are, however, large level differences across portfolios. The long-term investor invests on average a much higher fraction of their wealth in stocks. Level differences across the optimal and the myopic portfolio are shown in

Variance, expected returns, and price response to a variance shock



Portfolio response to a variance shock

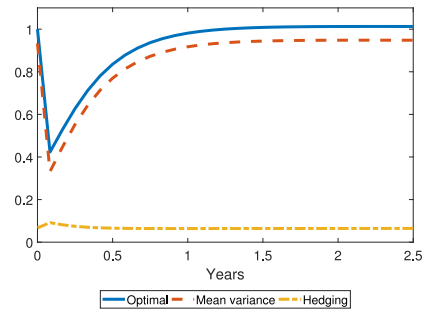


Fig. 1. Optimal portfolio response to a volatility shock. The top panel shows the behavior of the conditional expected return, conditional variance, and prices after a volatility shock. The bottom panel shows the optimal portfolio response to a volatility shock for both a mean variance (short horizon) and a long-horizon investor (labeled optimal). It also plots the conditional hedging demand term of the long-horizon investor. The x-axis is in years.

the level of the flat yellow line, which plots $w^*(\mu_t, \sigma_t^2) - w^m(\mu_t, \sigma_t^2)$ normalized by the steady state long-term portfolio $w^*(\bar{\mu}, \bar{\sigma}^2)$. The long-term investor has a risk exposure that is about 20% larger than the myopic investors in the steady state, but this difference, as a fraction of the risky portfolio share, grows large as volatility goes up and the myopic weight goes down.

The flat yellow line implies that the hedging demand is, at least locally, not strongly related to volatility. The hedging demand term drives a difference between short- and long-term investors, but this hedging demand term is roughly constant so that conditional responses to volatility variation are not substantially different. The response of a long-term investor to volatility is completely driven by the myopic component of her portfolio, i.e., variation due to the instantaneous risk-return trade-off.

In Fig. 2 we look at the global behavior of the optimal portfolio and show optimal policies as function of each state variable. Consistent with earlier work we find that the expected return hedging demand (H^μ) increases with expected returns.⁹ A new result unique to our analysis is that movements in the volatility hedging demand (H^σ) counteract this increase to some extent. Intuitively, investors are more exposed to volatility variation when their position in the risky asset is large. Therefore as

⁹ See for example Campbell and Viceira (1999).

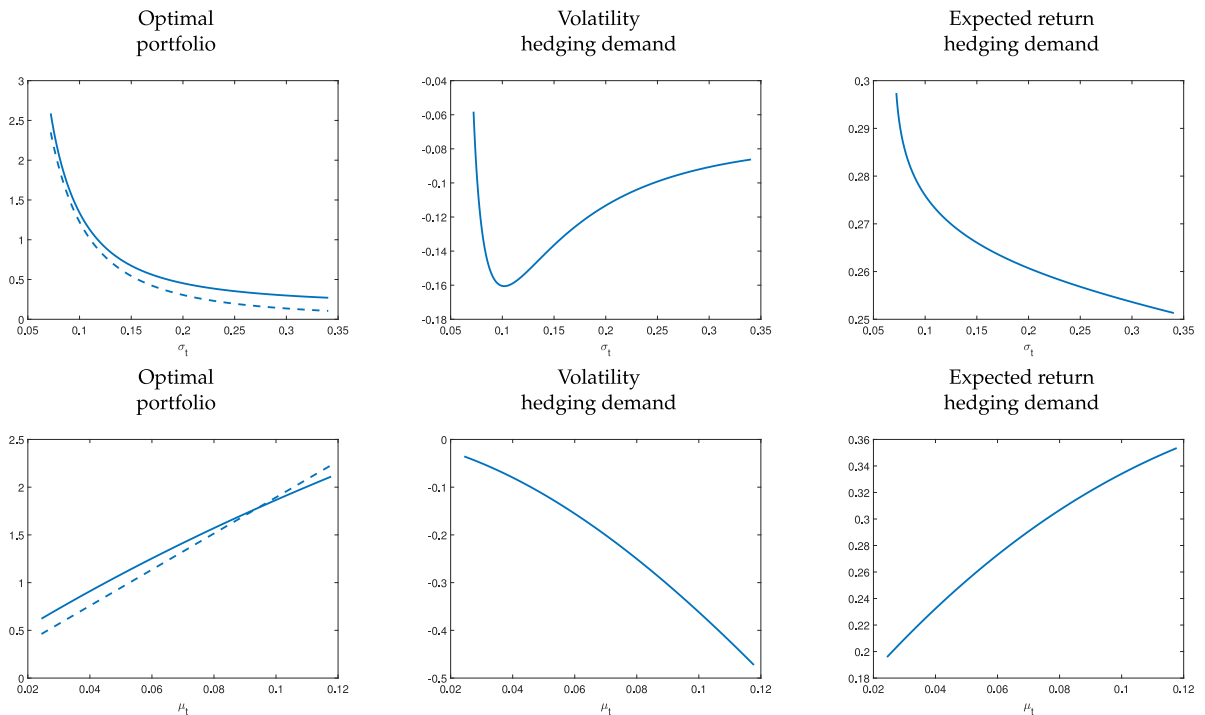


Fig. 2. Optimal portfolio policies. The top panels show how the optimal portfolio and its components vary as a function of the volatility state variable. Specifically, in the left panel we have the optimal portfolio (solid) and the myopic demand (dashed), and in the center and right panels we have the hedging demands due to volatility and expected return variation, respectively. The bottom panel shows how the same quantities vary with the expected return state variable.

expected returns increase, the volatility hedging demand becomes more negative. Increases in volatility reduce both volatility and expected return hedging demands. An increase in volatility reduces the expected return hedging demand because the investor has a smaller position on the risky asset when volatility is high and is therefore less exposed to variation in expected returns. The behavior of the volatility hedging demand is more nuanced, as it is U-shaped in volatility. As the volatility approaches the lower bound $\bar{\sigma}^2$, its volatility shrinks to zero. This effect makes the volatility hedging demand to become more negative as volatility increases from a low level. Eventually, the exposure channel dominates, and the hedging demand starts to become less negative as the investor reduces exposure to the risky asset as volatility increases further.¹⁰

Most importantly we see in the two right panels of Fig. 2 that the optimal portfolio tracks the myopic portfolio fairly closely as time variation in the myopic demand is an order of magnitude larger than time variation in the hedging demands. This motivates our analysis below.

3.2.1. The optimal portfolio is simple

Motivated by the result that movements in the hedging term are smaller compared to movements in the myopic demand (Fig. 1 and 2), we consider portfolio strategies that

invest in the myopic portfolio plus a constant position in the buy-and-hold portfolio,

$$\tilde{w}^*(\mu_t, \sigma_t^2) = a^* + b^* w^m(\mu_t, \sigma_t^2). \quad (15)$$

To find the loadings in \tilde{w}^* , we project w^* onto a constant and the managed portfolio w^m , i.e., $\tilde{w}^* = \text{proj}(w^* | 1, w^m)$. Given this portfolio rule \tilde{w}^* , we then solve for the investor's lifetime utility and compare this to the utility obtained under the fully optimal portfolio w^* . We find that the portfolio strategy \tilde{w}^* with $b^* = \text{cov}(w^m, w^*) \text{var}(w^m)^{-1}$ and $a^* = E[w_t^* - b^* w_t^m]$ attains the same lifetime utility as the optimal portfolio, i.e., the linear projection of the optimal portfolio on the myopic portfolio. The approximation \tilde{w}^* is not only a good local approximation for the optimal portfolio but also an excellent global approximation. We refer to $\tilde{w}^*(\mu_t, \sigma_t^2)$ as the optimal linear portfolio because its loading is a linear function of the myopic portfolio.¹¹

This result can be seen in Table 2 that shows the optimal policy loadings and the percentage lifetime expected utility loss from switching from the optimal portfolio to the affine approximation. A utility loss close to zero implies the investors forego almost no consumption if it adopts the simpler strategy. Thus \tilde{w}^* provides a good global approximation to w^* .¹²

¹⁰ Chacko and Viceira (2005) show that the exposure channel always dominates in a setting where the volatility process that does not feature a (strictly positive) lower bound. We confirm these findings in our setting.

¹¹ Formally, it is an affine function of w^m .

¹² A 1% utility loss is equivalent to decreasing the investor consumption by 1% state by state.

Table 2

Optimal portfolio as a function of investor horizon and risk aversion.

We first show ζ , the local elasticity of the optimal portfolio loading (w_t^*) to changes in variance $\zeta = -d\log(w_t^*)/d\log(\sigma_t^2)$ as we vary the investors' horizon (β^{-1}) from 10 to 100 years across columns and as we vary the investors' risk aversion (γ) from 3 to 10 across rows. We then show the approximation of the optimal portfolio that is affine in the myopic portfolio $\tilde{w}^*(\mu_t, \sigma_t^2) = a^* + b^* w^m(\mu_t, \sigma_t^2) = \frac{\mu_t}{\gamma \sigma_t^2}$.

We report loadings in the static buy-and-hold portfolio (a^*) and the myopic portfolio (b^*). The panel $\Delta\tilde{U}$ computes the utility losses from following the affine portfolio \tilde{w} compared to the fully optimal portfolio w^* , i.e., $\Delta\tilde{U} = U[\tilde{w}^*]/U[w^*] - 1$. It shows that our linear strategy \tilde{w} approximates the fully optimal portfolio well, as it results in small utility losses. Utility losses are in expected return units (e.g., $\Delta\tilde{U} = -1$ is equivalent to a 1% per year lower loss). Last, the panel ΔU^* shows the utility gains from switching from a static buy-and-hold portfolio (no timing of either state variable, with constant loading $\bar{w} = \bar{\mu}/(\gamma\bar{\sigma}^2)$) to the optimal timing portfolio ($\Delta U^* = U(w^*)/U(\bar{w}) - 1$).

ζ						$\Delta\tilde{U}$					
γ/β^{-1}	10	20	50	75	100	γ/β^{-1}	10	20	50	75	100
3.00	0.88	0.86	0.84	0.84	0.84	3.00	-0.03	-0.02	-0.02	-0.05	-0.05
5.00	0.88	0.85	0.83	0.83	0.83	5.00	-0.03	-0.02	-0.02	-0.02	-0.03
10.00	0.87	0.86	0.83	0.82	0.82	10.00	-0.02	-0.03	-0.01	-0.01	-0.01

a^*						ΔU^*					
γ/β^{-1}	10	20	50	75	100	γ/β^{-1}	10	20	50	75	100
3.00	0.16	0.21	0.25	0.25	0.25	3.00	5.00	4.71	4.40	4.20	4.15
5.00	0.11	0.14	0.17	0.18	0.17	5.00	3.14	2.84	2.59	2.52	2.42
10.00	0.06	0.07	0.09	0.09	0.09	10.00	1.67	1.41	1.28	1.24	1.22

b^*					
γ/β^{-1}	10	20	50	75	100
3.00	0.99	0.99	1.00	1.00	1.00
5.00	0.98	0.99	1.00	1.00	0.99
10.00	0.98	0.98	0.99	1.00	1.00

The result that the optimal linear portfolio \tilde{w}^* achieves the optimal utility is important because the numerical solution generates simple and implementable portfolio advice. Every investor can implement their strategy with two mutual funds: one that holds the market and one that times the risk-return trade-off. Tables 2 and 3 also show that the (approximate) optimality of the linear portfolio holds up across a wide range of parameters for the stochastic process, investment horizons, and risk aversion (specifically, we show how quantities change as we increase each parameter by one standard deviation). Note further that $\omega_1 \approx 1$ across a wide range of parameters. This tells us that the investment horizon plays a quantitatively important role only on the allocation to the buy-and-hold mutual fund. Investors, irrespective of their investment horizon, allocate the same fraction of their wealth to the timing mutual fund. Thus, investor portfolio response to volatility—as a fraction of their wealth—is always the same regardless of the investment horizon.

3.2.2. The optimal portfolio elasticity to changes in volatility

Another way to evaluate how responsive to volatility changes investors should be is in terms of an elasticity, i.e., the percentage change in the portfolio allocation resulting from a 1% increase in volatility, which is defined as

$$\zeta = -\frac{d\log(w^*(\mu_t, \sigma_t^2))}{d\log(\sigma_t^2)} \tag{16}$$

This perhaps provides a more direct measure of the importance of volatility driven changes for a particular in-

vestor. For a myopic investor $\zeta = 1 - d \ln(\mu_t) / d \ln(\sigma_t^2)$, which goes to one as the conditional risk-return trade-off goes to zero. Our estimates imply $\zeta^m = 0.96$, which reflects the small increase in expected return following a volatility shock we see in Fig. 1. An elasticity of one implies an investor reduce their exposure to stocks by 10% for a 10% increase in volatility.

The approximation (15) implies $\zeta \approx \frac{b^* w_t^m}{a^* + b^* w_t^m} \zeta^m$, with the long-term investor elasticity lower than the myopic as long $a^* > 0$. The elasticity also goes to zero as the myopic loading goes down due, for example to an increase in volatility. In Table 2 we focus on the elasticity of the optimal portfolio around the median value of the state variables to a one standard deviation increase in variance, i.e., the typical response to volatility. For the baseline parameters we find an elasticity of 0.85, which implies that as a share of her portfolio a long-horizon investor should respond less aggressively to variation in volatility. This happens because the long-horizon investor has a larger investment in the stock market to begin with (from the hedging demand term).

In Tables 2 and 3 we see that variation in ζ tracks variation in a^* , the optimal allocation to the buy-and-hold portfolio. For example, when the expected return is very volatile, high α_μ , the loading a^* is extremely high and the volatility elasticity is lower, around 0.68. To a smaller extent this also happens as we increase the investment horizon with the elasticity going down from 0.88 to 0.83 as the investment horizon increases from 10 to 50 years.

Table 3

Optimal portfolio and estimated parameters.

Here we show how key quantities of the portfolio policy change as we vary the model parameters. We use our estimation standard errors and report key quantities as we increase each parameter (parameter by parameter) by one standard deviation. We first show ζ , the local elasticity of the optimal portfolio to changes in variance $\zeta = -d\log(w_t^*)/d\log(\sigma_t^2)$. We then show the approximation of the optimal portfolio that is affine in the myopic portfolio $\tilde{w}^\sigma(\mu_t, \sigma_t^2) = a^* + b^* w^m(\mu_t, \sigma_t^2)$ where $w^m(\mu_t, \sigma_t^2) = \frac{\mu_t}{\gamma \sigma_t^2}$. We report loadings in the static buy-and-hold portfolio (a^*) and the myopic portfolio (b^*). The column $\Delta \tilde{U}$ computes the utility losses from following the affine portfolio \tilde{w} compared to the optimal portfolio, i.e., $\Delta \tilde{U} = U[\tilde{w}^*]/U[w^*] - 1$. It shows that our linear approximation captures the true optimal portfolio well in terms of resulting in small utility losses. Utility losses are in expected return units (e.g., $\Delta \tilde{U} = -1$ is equivalent to a 1% per year lower loss). Finally, the column ΔU^* shows the utility gains from switching from a static buy-and-hold portfolio (no timing of either state variable, with constant loading $w = \bar{\mu}/(\gamma \bar{\sigma}^2)$) to the optimal timing portfolio ($\Delta U^* = U(w^*)/U(\bar{w}) - 1$).

Parameter	Value	ζ	a^*	b^*	$\Delta \tilde{U}$	ΔU^*
Baseline		0.85	0.14	0.99	-0.02	2.84
$\bar{\sigma}$	3.97	0.83	0.13	0.98	-0.02	3.13
κ_σ	3.57	0.79	0.16	0.98	-0.01	2.44
ν_σ	4.82	0.90	0.14	1.00	-0.01	3.39
μ	8.51	0.89	0.13	0.99	-0.04	5.15
α_μ	0.59	0.78	0.23	1.00	-0.03	2.79
κ_μ	0.24	0.83	0.16	0.98	-0.05	2.73
$\rho_{\sigma,\mu}$	0.80	0.84	0.15	0.99	-0.02	2.76
$\rho_{\sigma,S-\mu}$	0.90	0.85	0.14	0.99	-0.02	3.06

In summary, the data on stock market returns, when looked at through the lens of the standard moments studied in the literature, strongly rejects the conjecture that investors should ignore movements in volatility. Investors with long investment horizons are somewhat less responsive to changes in volatility in terms of the percentage change in the size of their equity portfolio a given volatility movement calls for. However, as a percentage of their total wealth, both short- and long-term investors respond by identical amounts.

3.3. The (large) costs of ignoring variation in volatility

It is now clear that long-horizon investors should time volatility quite aggressively. Yet one could think that because volatility shocks are not very persistent, it might not be very costly to deviate from the optimal strategy. Here we evaluate the benefits of volatility timing by comparing increases in utility of only using information on conditional volatility, with the fully optimal policy that also uses information on conditional expected returns.

Specifically, we focus on a volatility timing strategies of the form

$$\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\bar{\mu}}{\gamma \sigma_t^2}, \tag{17}$$

where $\frac{\bar{\mu}}{\gamma \sigma_t^2}$ is the loading of a volatility managed portfolio from Moreira and Muir (2017). We again find the loadings a^σ and b^σ by projecting the fully optimal portfolio, w^* , onto a constant and the volatility managed portfolio so that $\tilde{w}^\sigma(\sigma_t^2) = \text{proj}(w^* | 1, \frac{\bar{\mu}}{\gamma \sigma_t^2})$. We refer to $\tilde{w}^\sigma(\sigma_t^2)$ as

the optimal volatility timing portfolio. We then compute the increase in utility from switching from the myopic buy-and-hold portfolio to $\tilde{w}^\sigma(\sigma_t^2)$, a portfolio that invests fixed amounts in the buy-and-hold portfolio and volatility-managed portfolio. We then compare these gains from volatility timing with the utility increase from switching from buy-and-hold to the fully optimal policy. We report these utility gains in terms of an annualized per period fee, i.e., the management fee that the naive buy-and-hold investor would be willing to pay to switch to one of the timing strategies (see Section C of the Internet Appendix for details).

Table 4 show results for our point estimates for horizons ranging from 10 to 100 years and coefficient of relative risk aversion ranging from three to ten. Following the myopic buy-and-hold strategy is very costly. For the baseline estimates, the naive buy-and-hold investor would be willing to pay 2.36% per period annualized fee to volatility time. In wealth equivalent terms this a 60% increase relative to buy-and-hold. A switch to the full optimal policy that also uses expected return information implies a fee of 2.84%. Thus the gains from volatility timing are large when compared to the total benefits of exploiting conditioning information. Specifically, Table 4 shows that investors can capture 80% of the total gains from timing by only timing volatility for an investment horizon of 20 years, and this share is larger than 70% for investors with horizons up to 100 years. Gains decay with the investment horizon but only very slowly.

Looking across rows in Table 5 we see that that this pattern holds up across a wide set of parameters. Volatility timing always increases utility relative to buy-and-hold, and always leads to increases that are more than 50% of the total gains from timing. All comparative statics are very intuitive. Gains make up a larger fraction of total gains when volatility is more variable and more persistent, expected returns are less volatile and more persistent, and investment horizons are shorter.

Lastly, it is worth highlighting the importance of the unconditional expected return. We see in Table 5 that the gains from volatility timing are strongly increasing in the unconditional expected return. This is consistent with the analysis in Moreira and Muir (2017) who find that when there is no conditional risk-return trade-off, the gains from volatility timing for a myopic investor increase proportionally to the unconditional risk-return trade-off.

Overall, these results show that ignoring volatility is likely to be very costly. These costs are large not only for our point estimates but also for a wide range of parameters that are consistent with the data.

4. The composition of volatility shocks

We have so far followed the empirical literature and assumed that the composition of volatility shocks is constant, equal to α_μ (see for example Campbell et al., 2016). Thus, when volatility changes, discount rate and cash flow volatility change proportionally (this is captured by the conditional volatility of $d\mu_t$ being proportional to σ_t). As a result, the amount of mean reversion in returns is constant. While this assumption is plausible a priori, it rules

Table 4

Utility gains from timing as a function of horizon and risk aversion.

We report loadings (a^σ , b^σ) in the volatility timing portfolio that is given by $\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\bar{\mu}}{\gamma\sigma_t^2}$. This portfolio is the approximation of the optimal portfolio that is an affine function of the volatility managed portfolio $\frac{\bar{\mu}}{\gamma\sigma_t^2}$. We then compare utility gains from several portfolio strategies. The first strategy is the optimal linear portfolio $\tilde{w}^*(\mu_t, \sigma_t^2) = a^* + b^* w^m(\mu_t, \sigma_t^2)$ (with associated utility $U[\tilde{w}^*]$), where $w^m(\mu_t, \sigma_t^2) = \frac{\mu_t}{\gamma\sigma_t^2}$ is the myopic portfolio loading. The second is the volatility timing portfolio. The third is the myopic buy-and-hold portfolio $\bar{w} = \bar{\mu}/(\gamma\bar{\sigma}^2)$. The panel ΔU^σ shows the utility gains from switching from buy-and-hold to the volatility timing portfolio ($\Delta U^\sigma = U(\tilde{w}[\sigma_t^2])/U(\bar{w}) - 1$). The last panel shows the fraction of the total utility gain that occurs from switching from the zero-timing buy-and-hold portfolio to the optimal portfolio can be achieved with just the volatility timing portfolio ($\Delta U^\sigma/\Delta U^*$). Utility gains are in expected return units (e.g., $\Delta U^\sigma = 1$ is equivalent to a 1% per year gain).

a^σ						ΔU^σ					
γ/β^{-1}	10	20	50	75	100	γ/β^{-1}	10	20	50	75	100
3.00	0.17	0.22	0.26	0.22	0.23	3.00	4.21	3.88	3.48	3.37	3.30
5.00	0.11	0.14	0.16	0.16	0.14	5.00	2.69	2.36	2.02	1.92	1.87
10.00	0.05	0.05	0.08	0.08	0.08	10.00	1.45	1.21	0.99	0.93	0.89
b^σ						$\Delta U^\sigma/\Delta U^*$					
γ/β^{-1}	10	20	50	75	100	γ/β^{-1}	10	20	50	75	100
3.00	0.88	0.89	0.89	0.90	0.90	3.00	84.19	82.38	79.11	80.24	79.39
5.00	0.88	0.89	0.90	0.90	0.91	5.00	85.69	82.92	78.12	76.16	77.24
10.00	0.89	0.90	0.91	0.91	0.91	10.00	87.03	85.44	77.35	74.84	73.32

Table 5

Utility gains from timing and estimated parameters.

Here we show how key quantities of the portfolio policy change as we vary the model parameters. We use our estimation standard errors and report key quantities as we increase each parameter (parameter by parameter) by one standard deviation. We report loadings (a^σ , b^σ) in the volatility timing portfolio that is given by $\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\bar{\mu}}{\gamma\sigma_t^2}$. This portfolio is the approximation of the optimal portfolio that is an affine function of the volatility managed portfolio $\frac{\bar{\mu}}{\gamma\sigma_t^2}$. We then compare utility gains from several portfolio strategies. The first strategy is the optimal linear portfolio $\tilde{w}^*(\mu_t, \sigma_t^2) = a^* + b^* w^m(\mu_t, \sigma_t^2)$ (with associated utility $U[\tilde{w}^*]$), where $w^m(\mu_t, \sigma_t^2) = \frac{\mu_t}{\gamma\sigma_t^2}$ is the myopic portfolio loading. The second is the volatility timing portfolio. The third is the myopic buy-and-hold portfolio $\bar{w} = \bar{\mu}/(\gamma\bar{\sigma}^2)$. The panel ΔU^σ shows the utility gains from switching from buy-and-hold to the volatility timing portfolio ($\Delta U^\sigma = U(\tilde{w}[\sigma_t^2])/U(\bar{w}) - 1$). The last panel shows the fraction of the total utility gain that occurs from switching from the zero-timing buy-and-hold portfolio to the optimal portfolio can be achieved with just the volatility timing portfolio ($\Delta U^\sigma/\Delta U^*$). Utility gains are in expected return units (e.g., $\Delta U^\sigma = 1$ is equivalent to a 1% per year gain).

Parameter	Value	a^σ	b^σ	ΔU^σ	$\Delta U^\sigma/\Delta U^*$
Baseline		0.14	0.89	2.36	82.92
$\bar{\sigma}$	3.97	0.13	0.87	2.61	83.14
κ_σ	3.57	0.17	0.87	1.95	79.87
ν_σ	4.82	0.11	0.92	2.86	84.42
μ	8.51	0.11	0.92	4.68	90.81
α_μ	0.59	0.22	0.88	2.13	76.34
κ_μ	0.24	0.17	0.85	1.94	70.85
$\rho_{\sigma,\mu}$	0.80	0.16	0.88	2.34	84.85
$\rho_{\sigma S-\mu}$	0.90	0.12	0.92	2.38	77.95

out the idea that movements in volatility are mostly due to “short-term volatility.” For example, [Cochrane \(2008a\)](#) argues that the huge spike in volatility in the fall of 2008 was fundamentally about an increase in the volatility of transitory shocks:

And what about volatility? ...[T]he standard formula is wrong... Stocks act a lot like long-term bonds ... If bond

prices go down more, bond yields and long-run returns will rise just enough that you face no long-run risk... [T]he same logic explains why you can ignore “short-run” volatility in stock markets. ([Cochrane, 2008a](#))

Empirically, there is no direct empirical evidence that confirms or refutes [Cochrane \(2008a\)](#) conjecture. In light of the fact that measuring the average share of discount rate shocks is already challenging ([Goyal and Welch, 2008](#)), measuring time variation in the share is even harder. Nevertheless, [Henkel, Martin and Nardari \(2011\)](#) and [Goez and Koudijs, 2018](#) provide some indirect evidence that suggests that there might be a positive correlation between volatility and the discount rate share. Using data that goes back to the seventeenth century, [Henkel et al. \(2011\)](#) and [Goez and Koudijs, 2018](#) show that most of the evidence for return predictability comes from periods identified as economic recessions. Together with the evidence that volatility tends to be high in recessions ([Schwert, 1989](#)), this evidence suggests that there might be a positive comovement between volatility and the degree of mean reversion in stock returns consistent with the idea that at least some of the movements in volatility are due to short-term volatility.

This comovement might matter a great deal to a long-term investor. If increases in volatility are entirely due to increases in discount rate volatility, the increase in return mean reversion will offset the increase in volatility, making the risky asset just as safe in the long run. Thus, the intuition is that investors with long investment horizons should not perceive periods of high discount rate volatility as much riskier than low volatility periods.

We model this comovement by focusing on the extreme case in which all variation in volatility is due to variation in discount rate volatility. Specifically, we decompose return innovations into transitory discount rate (dB_t^μ) and

permanent “cash flow” shocks¹³ (dB_t^c) as

$$dB_t^S = -\sqrt{\frac{\sigma_t^2 - \underline{\sigma}^2}{\sigma_t^2}} dB_t^\mu + \sqrt{\frac{\sigma_t^2}{\sigma_t^2}} dB_t^c, \quad (18)$$

which implies that the volatility of cash flow shocks is constant and only discount rate volatility varies. Consistent with Eq. (18) we set the volatility of expected returns in Eq. (2) to $\kappa_\mu \sqrt{\sigma_t^2 - \underline{\sigma}^2}$, which implies the discount rate share of return shocks is

$$\frac{\sigma_t^2 - \underline{\sigma}^2}{\sigma_t^2}. \quad (19)$$

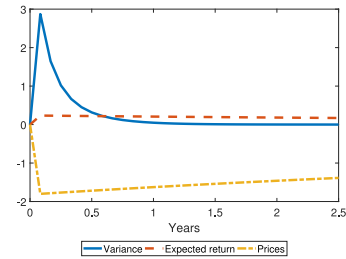
This share goes to 1 as volatility spikes to high levels and goes to zero as volatility drifts to the lower bound.¹⁴ Thus, when volatility is low returns are driven by permanent shocks and when volatility is high they are driven more by transitory shocks.

Fig. 3 show the impulse response functions (IRFs) for this extreme case in which all volatility variation is due to variation in the volatility of discount rate shocks. In the middle panel we see how the discount rate share spikes up with volatility and then slowly comes down. In the bottom panel we see that this results in an increase in hedging demand, which counteracts the decrease in exposure due to the myopic demand. It is still optimal to reduce the portfolio exposure after a volatility shock, but the response is less aggressive. Intuitively, stocks become relatively safer for a long-term investor than for a short-term investor when the share of discount rate shocks goes up. In response to a one standard deviation shock, the investor reduces his position in the risky asset by 30%, substantially less than the 45% in the constant discount rate share case.


The optimal portfolio can still be implemented with a constant position in the buy-and-hold and the myopic portfolio, but now the exposure to the myopic portfolio deviates from one. The elasticity to changes in volatility goes down from 0.83 in our baseline case to 0.39 in the discount rate volatility case, a large decline in the portfolio response to volatility variation. Nevertheless, the long-horizon investor still finds it optimal to time volatility.

Investors with a long investment horizon find it optimal to time the volatility of purely transitory shocks because our estimation implies these shocks take a long time to mean revert—the speed of mean reversion is very slow. Specifically, our SMM estimation interprets the return forecasting R -squares increase from 0.6% at the monthly horizon to 23% at the five-year horizon as evidence that the expected return process should be very persistent. This steep slope implies a mean reversion coefficient of $\kappa_\mu = 0.13$ for movement in expected returns, which translates into an autocorrelation of about 0.8 at the yearly frequency (a half-life of about six years). To put in perspective, an investor with a 20-year horizon (our baseline calibration) consumes about 30% of her wealth during this period.

Variance, expected returns, and price response to a variance shock



Share of discount rate shocks response to a variance shock



Portfolio response to a variance shock




Fig. 3. The composition of volatility shocks and the optimal portfolio response to a volatility shock. We repeat the exercise of plotting the optimal response to a volatility shock, but now we allow the composition of volatility shocks to change. Specifically, we now assume that the entire variation in volatility is due to variation in the volatility of discount rate shocks. Thus, the share of discount rate shocks in this case is increasing with volatility. To implement this, we assume the volatility of permanent shocks to returns (which we label “cash flow shocks,” consistent with the literature) is constant and equal to 7% and all time variation in volatility is due to discount rate volatility. This generates more mean reversion in returns when volatility spikes and hence increases the hedging demand term in high volatility periods as shown in the lower right figure. See Section 4 for a detailed description. The top plot shows the share of discount rate shocks in response to a one standard deviation volatility shock, the bottom left plot shows the response of volatility and expected return to a one standard deviation volatility shock, and the bottom right panel shows the optimal portfolio response.

Thus, investors respond to variation in the volatility of discount rate shocks because a substantial fraction of their consumption responds to variation in discount rates.

However, it is possible that return forecasting R -squares based on price-dividend ratios might lead us to overestimate the persistence of expected returns. Among others, Lettau and Van Nieuwerburgh (2008) and more recently Kelly and Pruitt (2013) find evidence that expected returns are less persistent than implied the return forecasting R -squares we use in our estimation. For example, Lettau and Van Nieuwerburgh (2008) show that when they allow for

¹³ We use the term “cash flow” shocks as is often used in the literature, but we simply mean permanent shocks to returns.

¹⁴ The correlation between cash flow and volatility shocks is simply $< dB_t^{S-\mu}, dB_t^c > = -\rho_{\sigma,S}$.

Table 6

Time varying composition of volatility shocks.

Here we show how time variation in the discount rate share impacts the optimal investment strategy. In the first columns (“baseline”), this share is constant (our baseline case), and in the next columns the discount rate share is increasing with volatility, which we label “positive comovement.” In short, while the degree of mean reversion in returns is constant in our baseline case, in the second case there is relatively more mean reversion in returns when volatility is high (see text for more details). Next, as we go from left to right we show what happens when we decrease the persistence of discount rates (increase in κ_μ) and show implied annual autocorrelations of expected returns. The idea is that lower discount rate persistence makes mean reversion in returns occur more quickly and thus makes discount rate shocks less risky for long-term investors. We show elasticities to volatility ζ and compare the utility from alternative portfolio strategies. The first, $\tilde{w}^*(\mu_t, \sigma_t^2)$ is the optimal linear portfolio (with associated utility $U[\tilde{w}^*]$). The second, $\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\mu}{\gamma \sigma_t^2}$ is the volatility timing portfolio. It is an approximation of the optimal portfolio that is affine in the volatility managed portfolio $\frac{\mu}{\gamma \sigma_t^2}$. The third is the myopic buy-and-hold portfolio $\tilde{w} = \bar{\mu} / (\gamma \bar{\sigma}^2)$. The row denoted ΔU^* shows the utility gain for an investor going from the myopic buy-and-hold portfolio to the optimal linear portfolio ($\Delta U^* = U[\tilde{w}^*] / U[\tilde{w}] - 1$). The row denoted ΔU^σ shows the utility gain for an investor going from the buy-and-hold portfolio to the volatility timing portfolio ($\Delta U^\sigma = U[\tilde{w}^\sigma(\sigma_t^2)] / U[\tilde{w}] - 1$). The last row shows the fraction of the total utility gain from the optimal portfolio can be achieved with the volatility timing portfolio ($\Delta U^\sigma / \Delta U^*$). Utility losses are in expected return units (e.g., $\Delta U^\sigma = 1$ is equivalent to a 1% per year gain). We report the data moments for comparison. The data column uses monthly market portfolio returns and its price-dividend ratio for the sample 1926 to 2015.

	Data	Baseline				Positive comovement					
κ_μ		0.06	0.15	0.25	0.40	0.60	0.06	0.15	0.25	0.40	0.60
$\rho(\mu_t, \mu_{t+12})$	0.91	0.90	0.83	0.75	0.65	0.54	0.90	0.83	0.75	0.65	0.53
R-sqrd 1 month	0.59	0.29	0.50	0.73	1.07	1.51	0.49	0.87	1.30	1.93	2.74
R-sqrd 1 year	8.18	3.32	5.52	7.66	10.20	12.56	5.55	9.75	13.90	19.14	24.52
R-sqrd 5 years	23.43	12.80	16.70	17.98	17.16	14.74	21.71	31.58	37.06	39.69	38.89
ζ		0.88	0.83	0.80	0.76	0.71	0.52	0.39	0.33	0.28	0.23
ΔU^*		2.69	2.64	2.70	2.88	3.21	2.73	3.35	3.92	4.62	5.40
ΔU^σ		2.44	2.10	1.86	1.59	1.31	1.65	1.62	1.48	0.03	-0.36
$\Delta U^\sigma / \Delta U^*$		90.64	79.40	68.71	55.28	40.69	60.16	48.20	37.82	0.58	-6.72

a structural break in the sample, the half-life of expected return shocks decrease to about three years.¹⁵ On the other hand, it is also possible that expected returns are more persistent than what we might estimate in these regressions (see [Stambaugh, 1999](#), for why persistence may be understated in shorter samples).¹⁶

Motivated by these findings, in [Table 6](#) we study how positive comovement between volatility and mean reversion interacts with the persistence of the expected return process. We focus the most on cases where expected returns are less persistent than implied by the price-dividend ratio because this implies mean reversion in returns would happen more quickly, amplifying the hedging demand for long-horizon investors. The return forecasting R-squares reported in the top rows of [Table 6](#) show that the expected return dynamics implied by our the various calibrations are comparable to what the empirical literature finds using various signals and methodologies. Specifically, they imply potentially much higher predictability at short horizons.

In the second column of Panel A we have our baseline estimation (expected return persistence of $\kappa_\mu u$ and constant share of discount rate shocks). In the second column of Panel B we have this same persistence but now with positive comovement between volatility and mean reversion. As expected returns become less persistent we see that the elasticity decays from 0.53 to 0.2 in the positive comovement case. The relative benefits of timing volatility

also decay, going down from 60% of the total gains from timing when expected returns are most persistent to zero when expected returns are the least persistent.

While measuring the exact degree of comovement between mean reversion and volatility is challenging, variation in return autocorrelations allow us to impose some restrictions on expected return dynamics. Intuitively, in the presence of positive comovement between volatility and mean reversion, the autocorrelation in returns should be more negative in high volatility months than in low volatility months. Specifically we compute,

$$\text{Corr}_{H-L}(R_{t,t+k}, R_{t-1,t}) = \text{Corr}(R_{t,t+k}, R_{t-1,t} | RV_{t-1}^2 > RV_{50\%}^2) - \text{Corr}(R_{t,t+k}, R_{t-1,t} | RV_{t-1}^2 < RV_{50\%}^2), \tag{20}$$

where $RV_{50\%}^2$ is the median of the realized variance in the sample.

[Table 7](#) shows these results. While in the data the difference in return autocorrelations between high and low volatility months is 0.45%, in the model it ranges from -2% to -14% depending of how persistent are expected returns and whether the discount rate share co-varies positively with volatility. While estimation uncertainty cannot rule out the possibility that there is positive comovement with very persistent expected returns, this moment is informative enough to rule out (with 4% p-value) positive comovement and expected returns that are equally or less persistent than implied by a expected return persistence of $\kappa_\mu = 0.4$ ($\text{corr}(\mu_{t+12}, \mu_t) \approx 0.6$). Put differently, if all high volatility episodes were associated with a lot more mean reversion in returns and this mean reversion happened very quickly, then returns would be much more negatively autocorrelated during high volatility periods. Thus, the data rule out this combination, and place a joint bound

¹⁵ See also evidence presented in [Drechsler and Yaron \(2011\)](#).

¹⁶ We thank a referee for this point. However, we note that if expected returns are more persistent than we estimate, this further limits the possibility of the volatility composition channel playing a role—hence it makes our conclusion that long-term investors should time volatility stronger.

Table 7

Time varying composition of volatility shocks in the data.

This table evaluates the time varying composition of volatility shocks in the data and the model. Specifically, we place joint bounds on the amount of comovement between volatility and mean reversion in returns by studying conditional return autocorrelations (at a one-year horizon) in high volatility versus low volatility episodes. We show that, as we increase the relationship between mean reversion in returns and volatility (“positive comovement”), this implies that stock returns should have a more negative autocorrelation in high volatility compared to low volatility periods. We show this by computing $corr_{H-L}$, which is the difference in return autocorrelations conditional on volatility being above (“High”) and volatility being below (“Low”) its median value. We then show that this difference is more pronounced when expected returns are less persistent (intuitively because this persistence controls the speed of mean reversion), and the columns in the table go from more persistent to less persistent expected returns as we go from left to right. The table shows that the data can reject the joint presence of positive co-moment between volatility and the discount rate share and an expected return process with low persistence. This is key because it is only in this case that volatility timing provides low benefits for long-horizon investors (see previous table). The p -value is computed by simulating the given process in the model 10,000 times and comparing the conditional return autocorrelations in these simulations with the empirical autocorrelation of returns in high and low volatility periods given in the top panel of the table. The data panel uses the market portfolio return from 1926–2015 and uses realized volatility in each month to condition on high and low volatility periods.

	Data				
$corr_{H-L}(R_{t,t+12}, R_{t-1,t}) \times 10^2$	0.452				
	Model				
κ	0.060	0.150	0.250	0.400	0.600
Baseline					
$corr_{H-L}(R_{t,t+12}, R_{t-1,t}) \times 10^2$	-2.011	-2.961	-3.716	-4.436	-4.903
p -value	0.359	0.315	0.276	0.243	0.218
Positive comovement					
$corr_{H-L}(R_{t,t+12}, R_{t-1,t}) \times 10^2$	-2.812	-6.012	-8.791	-11.808	-14.362
p -value	0.332	0.166	0.087	0.038	0.023

on the amount of comovement in volatility and mean reversion together with the speed of mean reversion. This rules out the range of parameters for which volatility timing is no longer beneficial.

Overall, this section shows that the composition of volatility shocks is a quantitatively important determinant of the optimal response to volatility. The comovement of volatility with the discount rate share determines whether it is optimal to respond more or less aggressively to changes in volatility. While the results in Table 7 leave open the possibility that there is positive comovement between volatility and mean reversion, it rules out the joint possibility that there is positive comovement and mean reversion which happens too quickly for long-term investors to care about transitory shocks. Therefore for the range of expected return dynamics that is likely given the data, it is always optimal to reduce the position in the risky asset when volatility goes up, and the benefits of such a strategy are sizable given likely parameters.¹⁷

5. Incorporating uncertainty

This section does two things. First, we assess the uncertainty surrounding our utility gains. Specifically, rather

than only reporting the point estimate for average utility gains, we study the full distribution of utility gains where we use the uncertainty from our estimation procedure about the parameters. We find the gains from volatility timing are extremely likely to be positive. Second, we then incorporate the fact that the parameters are unlikely to be known by the investor ex-ante. We incorporate parameter uncertainty by assuming an investor observes a signal for expected returns and volatility but does not know the true process for each and thus faces estimation risk, along the lines of Barberis (2000).

5.1. What is the probability that ignoring volatility variation is optimal?

The uncertainty surrounding our SMM parameter estimates indicates this probability is close to zero. We reach this conclusion by leveraging our SMM estimation to recover the uncertainty surrounding our parameter estimates and convert the uncertainty in this estimation to uncertainty about utility gains from volatility timing. We sample from the distribution of parameters recovered by the estimation (see Section 5) and solve for the optimal portfolio choice for each parameter realization. We then use the optimal portfolio solution to calculate the elasticity of portfolio with respect to changes in volatility and to compute the utility gains from switching from the myopic buy-and-hold strategy to a volatility managed strategy (and to the fully optimal strategy). This approach allow us to recover

¹⁷ In unreported results we study the case in which all volatility is driven by cash flow shocks. We find that in this case the optimal response to volatility variation and utility gains from volatility timing are larger than in our baseline case.

Table 8

Estimation uncertainty.

This table shows how estimation uncertainty impacts our results. We use our SMM estimation and optimal portfolio solution to recover the distribution of several features of the optimal portfolio policy. We show ζ , the local elasticity of the optimal portfolio to changes in variance $\zeta = -\partial \log(w^*) / \partial \log(\sigma^2)$, the optimal linear portfolio loadings in the static buy-and-hold portfolio (a^*) and the myopic portfolio (b^*), and compare utility gains from alternative investment strategies. We report utility gains in terms of the annual fee (as a share of wealth) the naive buy-and-hold investor would be willing to pay to switch to a timing strategy. ΔU^* is the annual fee to the fully optimal strategy $w^*(\mu_t, \sigma_t^2)$, and ΔU^σ is the annual fee to switch to the optimal volatility timing strategy $\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\mu_t}{\gamma \sigma_t^2}$. We report the 5%, 50%, and 95% percentiles of the distribution of each quantity implied by our estimation.

Quantity	5%	50%	95%
ζ	0.60	0.76	0.93
a^*	0.07	0.25	0.57
b^*	0.98	1.01	1.04
ΔU^*	1.61	3.01	6.25
a^σ	0.06	0.29	0.63
b^σ	0.64	0.84	0.95
ΔU^σ	0.60	2.04	5.53

the full distribution of the optimal portfolio elasticity and the economic gains of volatility timing.

The distribution of these quantities are shown in Table 8. The portfolio elasticity to volatility ranges from 0.6 to 0.9. This implies that the optimal response to a volatility increase is very likely to be a reduction in the weight allocated to the risky asset. The economic gains of volatility timing reflect this large elasticity. We find utility gains of volatility timing range from 0.6% to 5.5% in terms of an annual equivalent fee. Thus, ignoring volatility variation is extremely likely to be very costly.

5.2. Imperfect information

Our analysis so far endows investors' with perfect information with respect to variation in the investment opportunity set, i.e., we assume the state variables are perfectly observed by the investor. In practice investors have to form portfolios and trading strategies while facing uncertainty about their conditional estimates of the expected return and volatility.

We investigate how sensitive the utility gains of volatility timing and expected return timing are to forecasting uncertainty present in the data. This is important because, as many papers have shown, parameter uncertainty surrounding return predictability can have very large effects (Barberis, 2000; Goyal and Welch, 2008; Cochrane, 2008b; Pástor and Stambaugh, 2012). We confirm these results but show that uncertainty about conditional forecasts is not a crucial issue for volatility timing. Our approach follows closely the method described in Barberis (2000).

Specifically, we take the perspective of an investor who is given a 90-year sample for returns and who must estimate a rule to forecast returns and to forecast volatility using this 90-year sample. The investor then permanently adopts a portfolio timing rule based on this forecasting relationship. We then compute the expected utility for the investor going forward and compare this expected utility

to the case in which the investor knows the true process for returns (as computed earlier).

We begin by describing how we quantify uncertainty about expected return forecasting. In a given 90-year sample, the investor observes the variable μ_t that is the true conditional expected return for the risky asset, and uses this variable to forecast returns in a given sample.¹⁸ One can think of this as running predictive regressions with a candidate predictor such as the price-dividend ratio. The regression the investor runs in each sample is

$$r_{t+1} = \beta_{\mu,0} + \beta_{\mu,1} \mu_t + \varepsilon_{t+1}, \quad t = 1, \dots, T.$$

The true value for this regression is $\beta_{\mu,0} = 0, \beta_{\mu,1} = 1$, but the investor does not know this; he only sees μ_t as a candidate predictor of returns. In estimating this regression using a given sample, s , the investor estimates $\beta_{\mu,0}^s$ and $\beta_{\mu,1}^s$, where these are the coefficients recovered in a given sample. He then devises a trading strategy for what he believe is the expected return process going forward using these coefficients as the fitted value from this regression $\hat{\mu}_{T+t,s} = \beta_{\mu,0}^s + \beta_{\mu,1}^s \mu_{T+t}$ applied to the optimal portfolio rule $\tilde{w}(\mu_t, \sigma_t^2)$ described earlier in the paper. Notice that if, given 90 years of data, the investor always recovers the true coefficients $\beta_{\mu,0}$ and $\beta_{\mu,1}$, then this would be equivalent to the utility benefits of the full timing case studied earlier.

Next, we consider that the investor needs to make a forecast for volatility. As was the case before with the return forecasting regression, the investor is given a perfect signal, σ_t , about volatility, but must use this signal in the given sample to make a forecast about future volatility. He uses this signal to try to forecast the volatility next period. This is captured by $r_{t+1}^2 - \bar{r}_{t+1}^2 = \beta_{\sigma,0} + \beta_{\sigma,1} \sigma_t^2 + \varepsilon_{t+1}$ where we study the variation in these coefficients as before. Thus, the investor can poorly estimate the relationship between realized volatility and true volatility—analogue to the difficulty in predicting expected returns.

We then ask what is the expected utility associated with this rule given that these coefficients can vary from sample to sample, i.e., given that, even with 90 years of data, the investor may not know the true relationship between the predictor variable and future returns and future volatility. These results are given in Table 9 for three different sample sizes. We consider our baseline sample period in which the investor has 90 years of data, and we also consider the cases where the investor has 60 years or 180 years of data in which to estimate these relationships. Finally, we report the mean across many simulations as well as the lower percentiles of the distribution.

It turns out that the expected utility of the investor is highly sensitive to estimation uncertainty only when the investor tries to time expected returns (full timing column). The reason is that the estimated coefficients $\beta_{\mu,0}^s$ and $\beta_{\mu,1}^s$ vary dramatically from sample to sample, even with 90 years of data. This largely has to do with expected returns being very persistent making the relationship

¹⁸ Note that the estimation uncertainty we show here would be even larger if we were to consider imperfect predictors, that is, the investor only observes a noisy signal of μ_t , see Pástor and Stambaugh (2012).

Table 9

The costs of estimation uncertainty.

This table evaluates the robustness of adopting a volatility timing strategy and a full timing strategy with respect to parameter uncertainty. Specifically, we evaluate the utility costs when the investors must use a T -year sample to estimate a forecasting model for expected returns and volatility. We use sample sizes of 60, 90 (our sample size), and 180 years. Section 5.2 describes the calculation in detail. The last row shows utility gains when the investor faces no estimation uncertainty and knows the expected return and volatility signals (i.e., what the agent would realize with infinite data). Full timing means the agent times both expected returns and volatility, and volatility timing only times volatility. Both utility gain measures are given in percent per year relative to the naive buy-and-hold strategy that does not time either state variable. The first column reports the mean out of sample utility gain, and additional columns show the distribution of utility gains once estimation uncertainty is factored in, focusing on the lower end of this distribution.

Sample size (years)	Volatility timing (ΔU^v)				Full timing (ΔU^f)			
	Mean	5%	10%	25%	Mean	5%	10%	25%
60.00	1.95	1.49	1.69	1.89	-0.54	-13.33	-5.48	0.40
90.00	2.03	1.50	1.75	1.97	0.21	-8.93	-0.63	1.30
180.00	2.10	1.58	1.89	2.07	1.70	0.16	1.27	1.84
∞	2.36				2.84			

in the predictability regression difficult to estimate in a given sample. This point is well recognized by Goyal and Welch (2008), among others. When the investor only attempts to time volatility, estimation uncertainty turns out to be largely inconsequential, given 90 years of data, and the much lower persistence of volatility, the investor faces very low estimation risk. Table 9 contains these results and shows that the utility gains for volatility timing are essentially preserved when we take into account estimation risk.

In both the return and volatility forecasts we assumed the investor had perfect signals of the true process—in practice this is much more likely to be true of volatility, as investors observe realized volatility and have signals like the VIX that give near perfect signals of volatility in real time. In contrast, it is less likely that the investor would have a perfect signal for expected returns. Thus, our analysis here if anything understates the effects of parameter uncertainty on expected returns if one incorporates imperfect predictors.

In summary, because expected returns appear very persistent, predictive variables in a given sample can work poorly as forecasts for returns out of sample. This means that the benefits of timing expected returns are very sensitive to parameter uncertainty. We confirm this fact here, but this fact is well documented. However, this result is not true with volatility—this is essentially because volatility is easy to forecast both in and out of sample. Hence, the utility gains from volatility timing are far more robust to parameter uncertainty.

6. Extensions

We consider a number of extensions to our model including alternative preferences and outside income risk, and we briefly discuss how these extensions might interact with volatility timing.

6.1. The persistence of volatility shocks

Our model for volatility is univariate and therefore is bound to miss the dynamics of volatility at some

frequency. Previous work, most notably Chacko and Viceira (2005) and Campbell et al. (2016), has focused on low frequency aspects of the volatility dynamics (see also Zhou and Zhu, 2012). This makes sense if the goal is to study the level of hedging demands, as persistent variation in the opportunity set typically generates larger hedging demands. Our goal in this paper is different. We are interested instead in quantifying the utility gains of timing volatility. For this question, not only the persistence but also the volatility of the state variable is important. Because of this goal we choose to fit dynamics of volatility at the monthly frequency. As a result of this choice our model misses the dynamics of volatility at lower frequencies. For example, while our estimated volatility process has a 12-month autocorrelation close to zero, the point estimate in the data is approximately 0.4. In this section we show that our results are not very sensitive to this choice.

In Table 10 we present results when we target 3-month and 12-month autocorrelations of realized volatility instead on the 1-month we target in our baseline results. Gains from timing barely change; if anything, they increase slightly. Essentially, while focusing on lower frequency dynamics changes the average hedging demand, it does not change much the conditional hedging demands enough to change our basic results. Moreover, lower frequency movements in volatility do not have any effect on the myopic demand at all (recall that this demand does not depend at all on the persistence of volatility), and this drives most of the benefits of volatility timing. Thus, while we acknowledge that the model is misspecified in terms of capturing both high and low frequency aspects of volatility, we argue that this will not have much effect on our results. This is essentially confirmed in Table 10 that shows strong benefits of volatility timing even when we target lower frequency movements in volatility.

6.2. Alternative preferences

Thus far we have studied Epstein–Zin (EZ) preferences with unit elasticity of intertemporal substitution. These preference are not only standard in the portfolio choice

Table 10

Persistence of volatility shocks.

This table repeats our analysis but uses a volatility process targeted at lower frequency movements in volatility. That is, we specify volatility as an AR(1) in logs, as before, but we estimate the persistence parameter to match quarterly and annual autocorrelations in volatility rather than only monthly ones. Thus, this table considers how our main results can change if we focus on lower frequency movements in volatility. We report the data moments for comparison. The sample is monthly from 1926 to 2015. We compare the utility from alternative portfolio strategies. The first, $\tilde{w}^*(\mu_t, \sigma_t^2)$, is the optimal linear portfolio (with associated utility $U[\tilde{w}^*]$). The second, $\tilde{w}^\sigma(\sigma_t^2) = a^\sigma + b^\sigma \frac{\mu_t}{\gamma \sigma_t^2}$, is the volatility timing portfolio. It is an approximation of the optimal portfolio that is affine in the volatility managed portfolio $\frac{\mu_t}{\gamma \sigma_t^2}$. The third is the myopic buy-and-hold portfolio $\tilde{w} = \bar{\mu} / (\gamma \bar{\sigma}^2)$. The row denoted ΔU^* shows the utility gain for an investor going from the myopic buy-and-hold portfolio to the optimal linear portfolio ($\Delta U^* = U[\tilde{w}^*] / U[\tilde{w}] - 1$). The row denoted ΔU^σ shows the utility gain for an investor going from the buy-and-hold portfolio to the volatility timing portfolio ($\Delta U^\sigma = U[\tilde{w}^\sigma(\sigma_t^2)] / U[\tilde{w}] - 1$). The last row shows the fraction of the total utility gain from the optimal portfolio that is achieved with the volatility timing portfolio ($\Delta U^\sigma / \Delta U^*$). Utility losses are in expected return units (e.g., $\Delta U^* = 1$ is equivalent to a 1% per year gain).

Frequency targeted:	Data	1-month	3-month	12-month
$var(\log(RV_t))$	1.11	1.06	1.06	0.98
$corr(\log(RV_{t+1}), \log(RV_t))$	0.72	0.73	0.79	0.85
$corr(\log(RV_{t+3}), \log(RV_t))$	0.58	0.40	0.54	0.74
$corr(\log(RV_{t+12}), \log(RV_t))$	0.45	0.03	0.11	0.42
ζ		0.84	0.91	0.95
ΔU^*		2.64	2.66	2.45
ΔU^σ		2.16	2.18	2.28
$\Delta U^\sigma / \Delta U^*$		81.89	82.19	92.96

literature but also very convenient, as we can directly control the investor horizon by varying the impatience parameter β . Alternative preferences studied in the literature include (1) EZ preferences with nonunit EIS, (2) constant relative risk aversion preferences (i.e., EIS=1/RRA), and (3) preferences with habit formation.

Here we extend our analysis to (1) and (2). Fig. 4 show these results by comparing impulse responses across preferences. We normalize each portfolio weight by its steady state value so we can focus exclusively of the weight elasticity to a volatility shock. The top plot shows results for the baseline assumption that the composition of volatility shocks is constant; we see that all investors respond identically. In the bottom plot we see results for the case in which volatility and the amount of mean reversion comove. More specifically, when all volatility variation is due to variation in the volatility of transitory shocks to returns (discount rate volatility). We discuss this case at length in Section 4. In this case we see differences across investors responses as we vary the IES. Most interesting, we see that high IES investors tend to respond less to a volatility shock in this case. The reason behind this result is intuitive. Investors with higher IES tend to be less responsive to discount rate volatility because they optimally choose to save more when the investment opportunity set is very attractive, i.e., because they are more willing to postpone consumption their horizon is endogenously longer when the opportunity set is more attractive.

Case (3), habit, is substantially more complicated, as it requires adding a habit state variable. We haven't analyzed this case explicitly, but the analysis in Detemple and Zapatero (1992) and Gomes and Michaelides (2003)

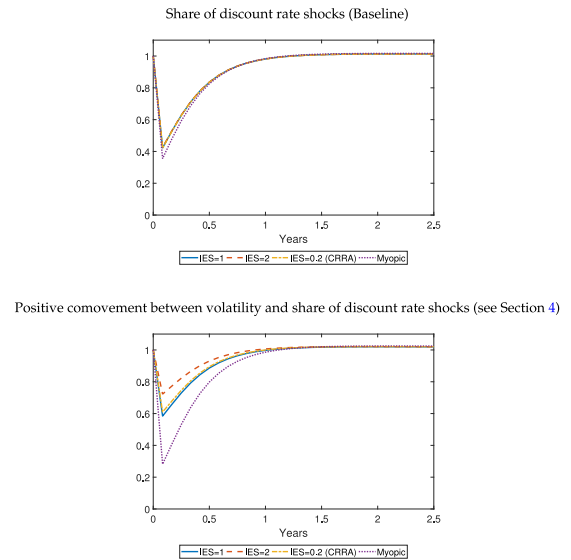


Fig. 4. Alternative preference parameters. We plot the optimal response to volatility shocks as we change the IES from 0.2 (the CRRA case) to 1 (our benchmark case) to 2. The top panel does this for our baseline model parameters, while the bottom panel does this for the case in which we allow changes in volatility to affect the discount rate share.

suggests that such preferences will lead us to similar results. For example, Detemple and Zapatero (1992) show that habit formation leads investors to first invest in a perfectly safe portfolio that finances habit consumption and then invest as a standard CRRA agent that had only

the residual wealth (the wealth minus the safe portfolio) would. Thus, they respond to a volatility shock as a CRRA agent with a similar allocation to the risky asset would. Their result suggests that while agents with habit forming preferences will invest much less in the market, their elasticity to a volatility shock is equal to a standard CRRA investor.

6.3. Nonfinancial income

Our baseline analysis is purposefully stark, as it relies on the assumption that the investor's only source of income is her financial wealth. A more realistic assumption is that the investor also earns wages or other sources of income. For example, Merton (1971), Viceira (2001), Cocco et al. (2005), and Polkovnichenko (2007) are examples of recent work that study how nonfinancial income shape portfolio decisions. For the baseline case in which outside income is riskless, these papers show that optimal portfolio is simply

$$w_t \frac{W_t}{W_t + PDV_t(E)} = \frac{\mu_t}{\gamma \sigma_t^2} + \text{hedging demand}, \quad (21)$$

where $PDV_t(E)$ is the present discounted value of the investor nonfinancial income, W_t is the investor financial wealth, and w_t is the share of financial wealth invested in the risky asset. The solution implies that the investor target the same share of total wealth allocated to the risky asset, what implies a much higher share of financial wealth, as $\frac{W_t}{W_t + PDV_t(E)} < 1$. In this simple riskless case, the solution is analogous to the investor having a lower risk aversion, $\tilde{\gamma} = \frac{W_t}{W_t + PDV_t(E)} \gamma$. Thus, all our our result will carry through to this case. We simply need to use $\tilde{\gamma}$ as the investor coefficient of relative risk aversion.

The impact of risk in the nonfinancial income stream can be understood by decomposing it in an idiosyncratic component that cannot be hedged or diversified and a component that covaries with the risky asset. Both components have the effect of reducing the present discounted value of the outside income. Intuitively, the higher the volatility, the higher the covariance between the income and the investor marginal utility. The end result is a higher discount rate. Again, our results will apply as in the riskless case after adjusting the outside income present discounted value. The aggregate component has a second effect because it not only impacts the value of the income stream, but it can also be hedged. The effect on the present value of the income stream is straightforward: a positive exposure increases the discount rate according to the risk premium earned in the risky asset. The covariance with the risky asset induces a new kind of hedging demand to emerge. Intuitively, the optimal portfolio choice adjusts for any exposure the investors income already has to the risky asset. A positive covariance thus induces a negative hedging demand, reducing the share of the investor financial wealth allocated to stocks.

While in practice, it is hard to find industries with wage income that is sufficiently strongly correlated with the stock market for these hedging demands to be large; more sophisticated modeling of labor income risk emphasizes a long-run relation between the stock market and

wages. For example, Benzoni et al. (2007) show that if labor income is cointegrated with dividends, the hedging demand can be large for empirically plausible parameters. Could this type of hedging demand overturn our results? As we have seen in Section 3.2, a constant negative hedging demand has the effect of increasing the elasticity of the portfolio weight to volatility. Thus, the “level” of the hedging demand will tend to amplify the optimal response to volatility. Our results can be overturned only if outside income hedging demand increases with volatility so that it pushes the portfolio toward stocks when volatility is high. The logic of cointegration is that all permanent shocks to stock prices, i.e., cash flow shocks, end up eventually impacting the labor income. Thus, variation in cash flow volatility should translate one-to-one to variation in the hedging demands, i.e., the hedging demand should become more negative in response to an increase in volatility. Variation in discount rate volatility, on the other hand, would not impact the hedging demand in this case. Thus, this co-integration channel would either increase or not impact the portfolio elasticity to volatility.

In order for the hedging demand to actually go up as volatility increases, the correlation between stock returns and wage income would have to go down enough to more than offset the increase in volatility. That is, a constant covariance between wage income and stock returns is not sufficient to overturn our results. In fact, this covariance would have to be strongly negatively related to volatility in order for hedging demand to increase with volatility. We are not aware of any empirical evidence pointing in this direction.

7. Conclusion

We study the portfolio problem of a long-lived investor that allocates her wealth between a riskless and a risky asset in an environment where both volatility and expected returns are time varying. We then comprehensively and quantitatively study how investors should respond to changes in volatility and what the utility costs to ignoring volatility variation are. We study how these results change with the investor's horizon and which features of the return dynamics we estimate are most important for our conclusions. Importantly, our analysis also takes into account that investors face parameter uncertainty regarding the dynamics of volatility and expected returns.

The main finding in this paper is that investors should substantially decrease risk exposure after an increase in volatility and that ignoring variation in volatility leads to large utility losses. The benefits of volatility timing are on the order of 2% of wealth per year for our preferred parameterization of an investor with risk aversion of 5- and a 20-year horizon. These benefits are significantly larger than those coming from expected return timing (i.e., from return predictability), particularly when parameter uncertainty is taken into account. We approximate the optimal volatility timing portfolio and find that its dependence on volatility is very simple: all investors, regardless of horizon, will choose fixed weights on a buy-and-hold portfolio that invests a constant amount in the risky asset, and a volatility managed portfolio that scales the risky asset

exposure by the inverse of expected variance $1/\sigma_t^2$. Further, we show that the weight on the volatility timing portfolio is independent of the investors' horizon in our baseline results.

We then show a novel channel through which long-horizon investors may differ in their response to volatility: they respond less aggressively to increases in volatility when only the volatility of mean-reverting shocks increases. Intuitively, this effect makes stock prices more volatile in the short run but doesn't change the distribution of long-run stock prices. Nevertheless, we provide empirical evidence that this channel is not strong enough to substantially decrease the gains from volatility timing.

References

- Barberis, N., 2000. Investing for the long run when returns are predictable. *J. Finance* 55 (1), pp. 225–264.
- Bazdresch, S., Kahn, R.J., Whited, T.M., 2018. Estimating and testing dynamic corporate finance models. *Rev. Financ. Stud.* 1, 322–361.
- Benzoni, L., Collin-Dufresne, P., Goldstein, R.S., 2007. Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *J. Finance* 62 (5), 2123–2167.
- Buffett, W. E., 2008. Buy American. I am. *The New York Times*.
- Buraschi, A., Porchia, P., Trojani, F., 2010. Correlation risk and optimal portfolio choice. *J. Finance* 65 (1), 393–420.
- Campbell, J.Y., 1996. Understanding risk and return. *J. Polit. Econ.* 104 (2), 572–621.
- Campbell, J. Y., Giglio, S., Polk, C., Turley, R., 2016. An intertemporal CAPM with stochastic volatility. Unpublished working paper. National Bureau of Economic Research.
- Campbell, J.Y., Shiller, R.J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Rev. Financ. Stud.* 1 (3), 195–228.
- Campbell, J.Y., Viceira, L.M., 1999. Consumption and portfolio decisions when expected returns are time varying. *Q. J. Econ.* 114 (2), 433–495.
- Campbell, J.Y., Vuolteenaho, T., 2004. Bad beta, good beta. *Am. Econ. Rev.* 94 (5), 1249–1275.
- Chacko, G., Viceira, L.M., 2005. Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. *Rev. Financ. Stud.* 18 (4), 1369–1402.
- Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. *Rev. Financ. Stud.* 18 (2), 491–533.
- Cochrane, J.H., 2008a. Is now the time to buy stocks? *Wall Street J.* Retrieved from: <https://www.wsj.com/articles/SB122645226692719401>.
- Cochrane, J.H., 2008b. The dog That did not bark: a defense of return predictability. *Rev. Financ. Stud.* 21 (4), 1533–1575.
- Collin-Dufresne, P., Lochstoer, L., 2017. Long-horizon investment, mean reversion, and structural breaks: a general equilibrium perspective. UCLA.
- Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. An intertemporal general equilibrium model of asset prices. *Econometrica* 53 (2), 363–384.
- Detemple, J.B., Zapatero, F., 1992. Optimal consumption–portfolio policies with habit formation. *Math. Finance* 2 (4), 251–274.
- Drechsler, I., Yaron, A., 2011. What's vol got to do with it. *Rev. Financ. Stud.* 24 (1), 1–45.
- Duffie, D., Epstein, L.G., 1992. Stochastic differential utility. *Econometrica* 60 (2), 353–394.
- Duffie, D., Singleton, K.J., 1993. Simulated moments estimation of Markov models of asset prices. *Econometrica* 61 (4), 929–952.
- Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57 (4), 937–969.
- Erickson, T., Whited, T.M., 2002. Two-step GMM estimation of the errors-in-variables model using high-order moments. *Econom. Theory* 18 (3), 776–799.
- Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *J. Finance* 56 (1), 329–352.
- Fleming, J., Kirby, C., Ostdiek, B., 2003. The economic value of volatility timing using “realized” volatility. *J. Financ. Econ.* 67 (3), 473–509.
- French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. *J. Financ. Econ.* 19 (1), 3–29.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Finance* 48 (5), 1779–1801.
- Golez, B., Koudijs, P., 2018. Four centuries of return predictability. *J. Financ. Econ.* 127 (2), 248–263.
- Gomes, F., Michaelides, A., 2003. Portfolio choice with internal habit formation: a life-cycle model with uninsurable labor income risk. *Rev. Econ. Dyn.* 6 (4), 729–766.
- Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium prediction. *Rev. Financ. Stud.* 21 (4), 1455–1508.
- Henkel, S.J., Martin, J.S., Nardari, F., 2011. Time-varying short-horizon predictability. *J. Financ. Econ.* 99 (3), 560–580.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* 6 (2), 327–343.
- Johannes, M., Korteweg, A., Polson, N., 2014. Sequential learning, predictability, and optimal portfolio returns. *J. Finance* 69 (2), 611–644.
- Kelly, B., Pruitt, S., 2013. Market expectations in the cross-section of present values. *J. Finance* 68 (5), 1721–1756.
- Lettau, M., Ludvigson, S., 2003. Measuring and modeling variation in the risk-return tradeoff. In: Ait-Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*. North Holland, Amsterdam, pp. 617–690.
- Lettau, M., Van Nieuwerburgh, S., 2008. Reconciling the return predictability evidence. *Rev. Financ. Stud.* 21 (4), 1607–1652.
- Liu, J., 2007. Portfolio selection in stochastic environments. *Rev. Financ. Stud.* 20 (1), 1–39.
- Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. *J. Econ. Theory* 3 (4), 373–413.
- Moreira, A., Muir, T., 2017. Volatility-managed portfolios. *J. Finance* 72, 1611–1644.
- Nagel, S., Reck, D., Hoopes, J., Langetieg, P., Slemrod, J., Stuart, B., 2016. Who sold during the crash of 2008–9? Evidence from tax return data on daily sales of stock. National Bureau of Economic Research.
- Pástor, L., Stambaugh, R.F., 2012. Are stocks really less volatile in the long run? *J. Finance* 67 (2), 431–478.
- Polkovnichenko, V., 2007. Life-cycle portfolio choice with additive habit formation preferences and uninsurable labor income risk. *Rev. Financ. Stud.* 20 (1), 83–124.
- Samuelson, P.A., 1969. Lifetime portfolio selection by dynamic stochastic programming. *Rev. Econ. Stat.* 51 (3), 239–246.
- Schwert, G.W., 1989. Why does stock market volatility change over time? *J. Finance* 44 (5), 1115–1153.
- Shanken, J., Tamayo, A., 2012. Payout yield, risk, and mispricing: a Bayesian analysis. *J. Financ. Econ.* 105 (1), 131–152. doi:10.1016/j.jfineco.2011.12.002.
- Stambaugh, R.F., 1999. Predictive regressions. *J. Financ. Econ.* 54 (3), 288–302.
- Viceira, L.M., 2001. Optimal portfolio choice for long-horizon investors with nontradable labor income. *J. Finance* 56 (2), 433–470. doi:10.1111/0022-1082.00333.
- Wachter, J.A., 2002. Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. *J. Financ. Quant. Anal.* 37 (01), 63–91.
- Xia, Y., 2001. Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation. *J. Finance* 56 (1), 205–246.
- Zhou, G., Zhu, Y., 2012. Volatility trading: what is the role of the long-run volatility component? *J. Financ. Quant. Anal.* 47 (2), 273–307.