The Missing Horizon in Asset Pricing

Valerii Baidin*

August 24, 2026

Abstract

I show that a mismatch between investors' true investment horizons and those assumed by econometricians can cause beliefs or risk preferences to violate Euler constraints, helping explain the equity premium puzzle. I derive a pricing condition—the Horizon-Adjusted Fundamental Pricing Equation—stating that the Euler equation holds at the investor's chosen horizon and no nearby horizon does better. Using Smoothed Empirical Likelihood to extract subjective beliefs for pairs of short and long horizons, and assuming these beliefs are time-consistent and rational in terms of transitivity, I identify representative investors' planning horizons at eight to fourteen quarters under normal conditions and five to eight quarters during recessions—substantially longer than conventionally assumed. Implied risk aversion falls sharply beyond seven quarters, confirming that short-horizon assumptions drive the puzzle. In portfolios with heterogeneous horizons, growth stocks align with shorter horizons and value stocks with longer horizons.

Keywords: Investment Horizons, Equity Premium Puzzle, Subjective Beliefs, Consumption-Based Asset Pricing, Heterogeneous Horizons
JEL Codes: G11, G12, E21, C14, D84

*McGill University, Desautels Faculty of Management, valerii.baidin@mcgill.ca

1 Introduction

Most motivations for investing in the stock market emphasize a long-term outlook rather than short-term investment horizons, owing to the higher unpredictability and relative volatility of the latter. The majority of investors allocate their capital to the stock market as a form of hedging and saving, positioning themselves as long-term investors[1]. The investment horizon is a critical piece of information that determines the period for which the asset pricing Euler equation should hold. Given that investors do not explicitly state their investment horizons and no alternative sources of this information are available, I aim to extract such data from current asset prices.

Historically, econometricians assumed that the investment horizon was a quarter or a year. The main reasons for this are the availability of data at that time and the ability to develop longer non-overlapping time series. The potential pitfalls lie in the possible mismatch between the investor's optimal investment horizon and the horizon that the econometrician assumes to be the investor's investment horizon. Therefore, if the investor's true investment horizon is longer, their beliefs and/or risk preferences may not satisfy the short-horizon Euler constraints typically assumed in the literature, and this mismatch may be a reason for the arising puzzles.

Jensen (1969) was the first to raise the 'horizon problem' for the evaluation of investment portfolios where investors have heterogeneous horizons. He shows that if the assumed investment horizon differs from the 'true' holding period, the linear relationship ('the market line') is distorted. Levy (1972) demonstrates that as long as the 'true' horizon does not coincide with the horizon assumed in empirical research, performance indexes contain a systematic bias.

Levhari and Levy (1977); Levy (1984) estimate CAPM with different horizons and find that the investment horizon for which data are collected plays a crucial role and has a great impact on the regression coefficients. They show that if the 'true' investment horizon is longer than assumed, then the estimation of systematic risk (beta) is biased closer to one, and the opposite is also true. Handa et al. (1989) provide evidence that the size effect becomes statistically insignificant when betas are estimated using annual returns.

Gressis et al. (1976) applied the mean-variance approach to multi-period investments and developed the relationship between the efficient frontiers of different investment horizons, showing that the investment horizon affects portfolio choice. Meanwhile, Lee et al. (1990) provide a generalized asset pricing model for the equilibrium risk-return relationship under

[1]Yet, some investors are speculators who do not generate excess profits (Barber and Odean, 2000) and have no long-term impact on the market.

heterogeneous investment horizons and suggest that "some empirical anomalies that are inconsistent with the traditional CAPM have resulted from the misspecification of the CAPM by ignoring the discrepancy between observed data periods and true investment horizons".

The literature suggests that many apparent asset pricing puzzles dissolve—or weaken considerably—once a longer investment horizon is assumed. Parker and Julliard (2005) show that a three-year horizon provides a significantly better explanation of portfolio variation than any smaller horizon. Daniel and Marshall (1997) conclude that the equity premium and risk-free rate puzzles are primarily problems for shorter-horizon returns. Brainard et al. (1991) demonstrate that the longer the horizon of the investor, the better the CCAPM performs relative to the CAPM.

While the term 'investment horizon' is widely used in the literature, it is mainly referenced as common knowledge without strict definition. More consequentially, it is never optimized over. Samuelson (1969) and Mossin (1968) solve for portfolio policy conditional on a terminal date T, showing that under i.i.d. returns and isoelastic utility the optimal share is independent of T; the horizon is a parameter of the problem, not a choice. No optimality condition for the horizon itself exists: the literature has not asked which horizon an investor would choose.

The paper begins by defining the investment horizon formally. I then establish a benchmark irrelevance result: when log consumption and log returns follow random walks, every horizon prices identically, and the choice of horizon is irrelevant—the pricing-side analog of the portfolio irrelevance of Samuelson (1969). Once these dynamics are relaxed—as the data require—the optimal horizon is characterized by the Horizon-Adjusted Fundamental Pricing Equation (HAFPE): at that horizon, the standard pricing equation holds exactly, and no adjacent horizon offers the investor an improvement. Risky-asset demand peaks at this horizon, rising before it and declining after. When the econometrician imposes a shorter horizon than the investor's, the Euler equation fails to hold under the investor's true preferences, and forcing it to be satisfied requires implausibly high risk aversion and/or a discount factor above one—the equity premium and risk-free rate puzzles.

Turning to the empirical side, I make the first attempt, to the best of my knowledge, to extract information about the representative agent's investment horizons from asset pricing, discover the parameters on which they depend, and examine differences across various time periods. These findings shed light on some puzzles and help deepen understanding of asset pricing, taking into account varying agent investment horizons. To do so, I start by extracting the representative investor's subjective beliefs and risk preferences that satisfy the Euler constraints for various investment horizons using the Smoothed Empirical Likelihood (SEL) methodology (Chen et al., 2020; Ghosh et al., 2021; Ghosh and Roussellet, 2023). At a

first approximation, we conjecture that if, for a particular horizon, the investor's beliefs and/or risk aversion levels are unrealistic, then it is likely not a "true" investor horizon. This approach gives us some insights but not exact numbers.

The central exercise is a "quasi" test in which the agent chooses between two horizon types: short and long. The short-type agent invests two consecutive times. The long-type agent, with a horizon twice as long as the short type, invests only once. The test is constructed such that both agents live for the same period of time and face the same expected path for all conditional variables, with the only difference being their horizon type. Finally, the agent chooses the type that yields the higher utility of investing. I conduct this test for various possible horizons and construct boundaries for the likely true investment horizons.

In the final part, I extract portfolio compositions for various horizons. As Gressis et al. (1976) and Lee et al. (1990) show, an agent's optimal portfolio depends on the horizon. I obtain this information in two ways. First, I extract the portfolio composition for the given horizon h. Second, I find all possible allocations between two portfolios, the shorter and the longer horizons, for the representative investor who holds the entire market portfolio. Additionally, in this section, I explore the possibility that heterogeneous horizons could explain the differences in the investment portfolios of individual and institutional investors (Gompers and Metrick, 2001; Ferreira and Matos, 2008; Barber and Odean, 2000; Choi and Robertson, 2020).

The empirical findings are as follows: the most likely investment horizon is between 8 and 14 quarters and between 5 and 8 quarters during a recession. Investors are unlikely to invest at one- and two-quarter horizons, regardless of risk aversion. More risk-averse agents are less likely to invest in these horizons, even when they have a shorter planning horizon. Additionally, the results highlight a significant reduction in implied risk aversion for investing beyond 7 quarters, providing insight that the equity premium puzzle is less pronounced for the longer-horizon investors. Auxiliary results from the analysis of portfolios with heterogeneous horizons support the main results and suggest that low B/M assets are located in shorter-horizon portfolios, while high B/M assets are in longer-horizon portfolios.

The paper is organized as follows: Section II (Literature Review) discusses how investment horizons appear in different areas of literature. Section III introduces the theoretical model of the investment horizon and its implications for asset pricing, including the Horizon-Adjusted Fundamental Pricing Equation. Section IV describes the methodology, including the model setup, identification approach, and method to identify subjective beliefs and risk preferences from current asset prices. Section V presents the main results, including an overview of investors' beliefs and risk preferences over various horizons, and the identification of the true range of investment horizons. Section VI explores investment portfolios with

heterogeneous investment horizons and different points of view on how horizons, agents, and firms could correspond to each other. Finally, Section VII concludes. In addition, the appendix contains robustness check results and parsimonious numerical examples of the theoretical part.

2 Overview of Investment Horizons in Various Areas of Literature

In the introduction section, we covered the specific literature to which this paper primarily contributes. In this section, we provide a general overview of how the investment horizon is reflected in different areas of finance and how these areas suggest long-term horizons.

At this moment, we assume a fairly broad definition of investment horizon as the planned time to buy and hold an asset, without any obligation to keep the asset for exactly that duration. The topic of investment horizon is widely covered in literature from various perspectives including behavioral factors, savings strategies, predictability, information acquisition, and equity term structure. To take a common example, if someone is planning to buy a house soon, he is unlikely to buy stocks over the next week, yet many use stocks as long-term savings. Shiller (2014) highlights that asset prices exhibit "instantaneous unpredictability"(Sims, 1984) when prediction time approaches zero, and "accords with long-standing advice that investors should be patient, that they cannot expect to see solid returns over short time intervals." Thus, we can naturally link the investment horizon to the predictability of long-horizon returns, assuming that agents have an investment horizon longer than just the unpredictable immediate future.

One of the first papers to mention return predictability for the long horizon was Fama and French (1988a), which showed that the power of dividend yields to forecast stock returns increases with the horizon. This result was followed by many other papers (Campbell and Shiller, 1988; Hodrick, 1992; Fama and French, 2018). Mark (1995) emphasizes the economic significance of exchange rate long-horizon predictability for 12-16 quarters, against a random walk dominated in the short-horizon over 2-4 quarters. Cochrane (1999) states "Returns are predictable. In particular: Variables including the dividend price ratio and term premium can predict substantial amounts of stock return variation... Daily, weekly, and monthly stock returns are still close to unpredictable." Therefore, the long-horizon predictability could be one of the signs that the investment horizon should be longer than one month, and investors might even refrain from trading for some time.

Barberis (2000) demonstrates that "In the buy-and-hold case, we find that predictability

in asset returns leads to strong horizon effects: an investor with a horizon of 10 years allocates significantly more to stocks than someone with a one-year horizon." Balduzzi and Lynch (1999) establish that a long-horizon investor, who faces fixed and proportional transaction costs, may reduce the frequency of optimal rebalancing, leading to sizable no-trade regions. Thus, transaction costs are another reason for investors having various horizons.

It is worth mentioning that some long-horizon predictability is doubted, as Kothari and Warner (1997) show that tests for long-horizon abnormal security returns are severely misspecified and conclude that long-horizon studies require extreme caution. Similarly, Pástor and Stambaugh (2012) find that from an investor's perspective, which recognizes that parameters are uncertain, stocks are substantially more volatile over long horizons. At the same time, Barberis (2000) states that "even after incorporating parameter uncertainty, there is enough predictability in returns to make investors allocate substantially more to stocks, the longer their horizon."

The investment horizon might depend on the agent's ability to acquire information and learn from it. Hirshleifer et al. (2011) demonstrate that a subset of investors fails to fully incorporate the information contained in the most recent earnings surprises. Coibion and Gorodnichenko (2015) also show that individuals underreact to objective economic information. However, Bordalo et al. (2020) find that individual forecasters typically overreact to news, while consensus forecasters underreact relative to full-information rational expectations. Results are mixed, and both knowing the length of the investment horizon and understanding how information is learned will enhance each other. It is worth mentioning that price informativeness differs in good and bad times, and this may affect the length of the investment horizon depending on the state of the economy, as we will find below.

This paper naturally contributes to the large range of literature on multiperiod portfolio selection (Li and Ng, 2000; Basak and Chabakauri, 2010; Cochrane, 2022). In the classic i.i.d. world, portfolio allocation is independent of investment horizon, yet Kandel and Stambaugh (1996) are among the first to show how return predictability alters the optimal portfolio. Campbell and Viceira (1999) consider consumption and portfolio allocation and show that the long-horizon investors are free to rebalance their portfolio every period, even though the paper does not take into account transaction costs and portfolio restrictions. The same is also observed from real data; Ameriks and Zeldes (2004) report that in their sample data[2], individuals seldom change their portfolio allocation over time. Further, Viceira (2001) demonstrates, using the example of employed agents investing for retirement purposes, how horizon affects portfolio allocations and shows that longer horizons result in larger investments in stocks. Finally, in a recent working paper, Gabaix et al. (2024) explore

[2]Sample data is from the Survey of Consumer Finances and from TIAA-CREF.

the factors affecting the rebalancing of portfolios and found that the main one is responsible for the long-term equity risk premium.

Investment horizon is connected to the equity term structure literature, which has been growing lately (Gonçalves, 2021; Gormsen, 2021; Bansal et al., 2021). Binsbergen et al. (2012) show that long-term dividends have a lower risk premium than short-term dividends. Cassella et al. (2023) find that negative term premiums are associated with periods of above-average horizon bias, whereas periods of below-average horizon bias are associated with positive term premiums.

One of the behavioral factors affecting investment horizon is anxiety in the face of risk, or horizon-dependent risk aversion (HDRA), where agents tend to be more risk-averse when a risk is temporally close than when it is distant (Eisenbach and Schmalz, 2016). This preference generates a decreasing term structure of risk premia (Andries et al., 2014), and might be responsible for longer investment horizons. In addition, Daniel et al. (2020) provide other behavioral factors or psychological biases that are likely to affect asset prices, such as overconfidence and limited attention, which are responsible for longer and shorter horizon mispricing.

Finally, this paper contributes to the broad and growing literature on investors' subjective beliefs (Bianchi et al., 2022; Greenwood and Shleifer, 2014; Giglio et al., 2021; Chen et al., 2020; Ghosh and Roussellet, 2023). Survey data doesn't provide exact information about planned investment horizons, yet we can extract valuable insights from it. For instance, Giglio et al. (2021) show that even respondents who expect negative returns over the next year anticipate positive long-run returns. Therefore, we can make an assumption that an agent can make an investment decision based on long-horizon expectations, assuming they keep the stock for a longer period.

3 Investment horizon

This section develops the theory of the investment horizon and its implications for asset pricing. The first subsection defines the investment horizon and its corresponding Euler equation, and introduces a horizon function that incorporates the horizon directly into the model. The next subsection shows how neglecting the investment horizon can generate estimation bias and thereby help explain the equity premium puzzle. I then relate the presence of an investment horizon to the random-walk benchmark for consumption and returns. Finally, a model of the optimal investment horizon yields the Horizon-Adjusted Fundamental Pricing Equation.

3.1 Definition

This subsection presents a formal definition of the investment horizon, a term widely used in the literature but rarely defined precisely. I assume that investors have long investment horizons[3], as shown in the previous section and suggested by instantaneous unpredictability (Sims, 1984), long-horizon return predictability (Fama and French, 1988a, 2018), transaction costs (Balduzzi and Lynch, 1999), behavioral factors (Eisenbach and Schmalz, 2016; Daniel et al., 2020), and subjective beliefs (Giglio et al., 2021).

I define the investment horizon h as the planned duration for holding an asset, though investors do not need to hold the asset for exactly this period. For this specific horizon h, the Euler Equation 1 is satisfied, where S denotes subjective beliefs. The horizon h likely falls within a certain interval $\mathcal{H}$ (Jensen, 1969) (e.g., approximately three years, plus or minus a quarter). Therefore, Equation 1 must hold for at least one $h \in \mathcal{H}$.

$$\mathbb{E}_t^S[M_{t,t+h}R_{t,t+h}] = \mathbf{1}, \text{ for some } h \in \mathcal{H}, \tag{1}$$

where $M_{t,t+h} \equiv \beta^h U'(c_{t+h})/U'(c_t)$ is the stochastic discount factor (SDF) between periods t and $t+h$, and $R_{t,t+h}$ is the gross return over the same interval.[4] All expectations are under the investor's subjective measure; to simplify notation, I drop the superscript S in what follows.

At each time t, the investor chooses a horizon $h \in \mathcal{H}$ and makes decisions so that the Euler Equation 1 is satisfied over that horizon. Doing so requires forming expectations about consumption and returns over the full interval from t to $t+h$, accounting for both routine financial decisions and unpredictable events such as geopolitical shocks or career changes. When the next period $t+1$ arrives, the investor re-optimizes: the new horizon may either reset to the same length h (a fresh plan) or count down to $h-1$ (the same target date, now one period closer), provided the resulting horizon remains in $\mathcal{H}$.

Lee et al. (1990) assume investors hold their portfolio for the entire duration of their investment horizon without rebalancing or switching assets. In contrast, my approach does not impose such tight constraints and allows periodic rebalancing to be an integral part of the investment strategy. Conceptually, the investor primarily evaluates expected performance over their horizon (e.g., several years ahead) and reevaluates it in the next period.

The investment horizon framework aligns with time-inconsistent models, such as models

[3]Longer than one period of the model.

[4]I use the double-index notation $M_{t,t+h}$ rather than M_{t+h} to emphasize that the SDF is defined over the *interval* from t to $t+h$, as we will later refer to the SDF between two periods in the future. This convention makes explicit both the base period and the evaluation period, which is important when the investment horizon h varies over time. The same convention applies to $R_{t,t+h}$.

employing quasi-hyperbolic preferences ($\beta\delta$-discounting) (Laibson, 1997; Krusell and Smith, 2003; Laibson et al., 2026). Following the terminology of O'donoghue and Rabin (1999), I consider a *naive* investor: at time t, she commits to a plan for the near future as though her future selves will carry it out, although when the next period arrives, she re-optimizes and makes new decisions about portfolio allocation without having anticipated this revision.[5]

The intuition parallels how a driver focuses on a specific range of the road ahead depending on speed[6]. Only what falls within that range primarily affects the driver's actions: for instance, if they were going to make a right turn and found that it was already very close, they would miss it, as they would not have time to slow down. Similarly, a turn five miles ahead also does not affect their current actions. In contrast, if the turn is within the right range, it would affect their actions; they would slow down and prepare to turn. In the same way, investors concentrate on their investment horizon—the time range over which their Euler equation holds, since the very short-term market is too volatile compared to relatively stable consumption (income) in the near future, and what lies in the far future is heavily discounted and does not affect their decisions. This does not mean that instantaneous market events have no effect at all; it only means that investors mainly focus on the horizon range where their FOC must hold, just as a driver constantly focuses on the range where the car is heading and only rarely resorts to emergency braking.

3.2 The Horizon Function

To incorporate the investment horizon directly into a model, I formalize it as a behavioral rule—a horizon function $H(\tau)$, defined in Equation 2, that the investor plans to follow[7]—giving the share k_τ of invested wealth the agent plans to use at $t+\tau$:

$$H(\tau) = k_\tau, \qquad \sum_{\tau=1}^{\infty} k_\tau = 1, \qquad k_\tau \geq 0 \quad \text{for all } \tau \in \{1, 2, \ldots\}. \tag{2}$$

For instance, an investor with an exact two-quarter horizon sets $H(2) = 1$, planning to leave the investment untouched next period and use it entirely two periods ahead. The horizon function can also be interpreted as a probability distribution over future periods, capturing the likelihood that the investment is accessed at each date. A long-term investor concentrates probability mass several months or years into the future, assigning low probability to

[5] A *sophisticated* investor would anticipate her future re-optimization and adjust her current plan accordingly. The naive assumption is the simpler benchmark and is consistent with the fresh plan the agent makes each period in Equation 4.

[6] Appendix B provides illustrations of driver-focused vision (Figure 33, Figure 34, Figure 35).

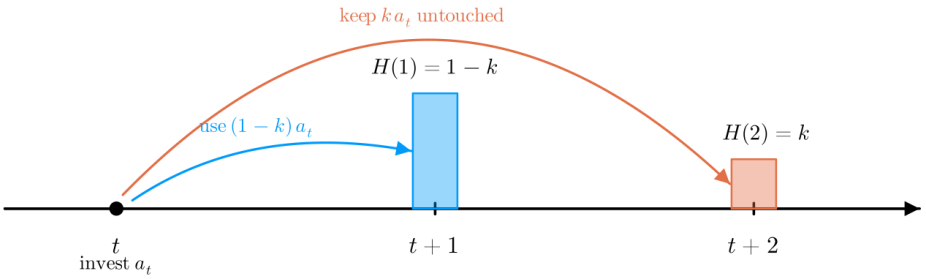
[7] Commonly recommended advice is "set it and forget it."

accessing the investment next period. An investor with $k_2 = 0.3$ plans to use 70% of the invested wealth next period and to leave 30% untouched for two periods; in the probabilistic reading, she expects to need the funds next period with probability 0.7.

For simplicity, let us assume the maximum horizon range is two periods. The horizon function then takes the form Equation 3, stating that in the next period the investor plans to use only $H(1)a_t = (1-k)a_t$ of the asset currently invested, and $H(2)a_t = ka_t$ will remain untouched until two periods ahead. Figure 1 illustrates this plan. Alternatively, the probabilistic approach states that in the next period, the investor will use the currently invested asset a_t with probability $1-k$ and will keep it until two periods ahead with probability k. In what follows, I adopt the deterministic-share approach.[8]

$$\begin{aligned} H(1) &= 1-k \\ H(2) &= k \\ k &\in [0,1]. \end{aligned} \tag{3}$$

Figure 1: Two-Period Horizon Function



Notes: The figure draws the two-period horizon function Equation 3 for $k = 0.3$. At time t the investor buys a_t and plans to use the share $H(1) = 1-k$ of the position at $t+1$ and to keep the remaining share $H(2) = k$ untouched until $t+2$; the weights sum to one. In the probabilistic reading, the bar heights are the probabilities of accessing the investment at each date.

Consider a parsimonious model: an agent receives stochastic income $\tilde{e}_t$ in each period, invests in a non-dividend-paying stock a_t with current price p_t and exogenous stochastic future prices $\tilde{p}_{t+1}$ and $\tilde{p}_{t+2}$—so the gross return is the price ratio $R_{t,t+h} = \tilde{p}_{t+h}/p_t$—and takes the horizon function in Equation 3 as given. To avoid excess notation, I suppress the risk-free asset and focus on the risky asset. This simplification does not affect the form of the risky-asset horizon conditions derived here. The investor retains access to a risk-free asset

[8]The probabilistic interpretation requires a different expected-utility formulation but delivers similar qualitative implications.

to finance intermediate consumption; its horizon-specific pricing conditions are introduced later. The agent maximizes lifetime utility in Equation 4.

At time t, let w_t denote cash on hand: current income plus assets that are actually available at t. Since each period the agent makes a fresh plan and has no obligation to follow the plan made earlier, the entire portfolio chosen in the previous period is available at t, so $w_t = e_t + a_{t-1}p_t$. Expected consumption next period depends on the stochastic income $\tilde{e}_{t+1}$, the asset price $\tilde{p}_{t+1}$, and the fraction $H(1)$ of currently invested assets the agent plans to use. The agent forms expectations about consumption two periods ahead in the same way, remaining consistent in applying the horizon function.

$$\begin{aligned} &\max_{\{a\}_t^\infty} \mathbb{E} \sum_{t=0}^{\infty} \beta^t U(c_t) \\ s.t.\ \ & c_t = w_t - a_t p_t \\ & \tilde{c}_{t+1} = \tilde{e}_{t+1} + H(1)a_t\tilde{p}_{t+1} - a_{t+1}\tilde{p}_{t+1} \\ & \tilde{c}_{t+2} = \tilde{e}_{t+2} + H(2)a_t\tilde{p}_{t+2} + H(1)a_{t+1}\tilde{p}_{t+2} - a_{t+2}\tilde{p}_{t+2}. \end{aligned} \tag{4}$$

Note that the budget constraint for c_t reflects the time-inconsistency of this framework: at time t, the agent uses the cash on hand w_t and is not bound by plans made in previous periods. That is, although at time $t-1$ the agent planned to keep a fraction $H(2) = k$ of a_{t-1} untouched until period $t+1$, upon reaching period t the agent re-optimizes and treats the available portfolio as liquid. This is the same form of time-inconsistency as in the quasi-hyperbolic preference models discussed above (Laibson, 1997; O'donoghue and Rabin, 1999). The constraints for c_{t+1} and c_{t+2}, in contrast, represent the agent's current forward-looking investment plan at time t.

The first-order condition Equation 5 depends on the horizon function constant k in Equation 3 and is a convex combination of the short- and long-horizon "Euler equations"[9]; I refer to each such term as a *horizon Euler term*, $F(h) \equiv \mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$. This parsimonious model serves only to present the Euler equation as a convex combination; no special assumptions are imposed on the income process or returns. The following discussion could apply to a broad range of models.

$$\begin{aligned} & U'(c_t)p_t = \mathbb{E}_t\left[(1-k)\beta U'(c_{t+1})p_{t+1} + k\beta^2 U'(c_{t+2})p_{t+2}\right] \\ \Rightarrow\ \ & 1 = (1-k)\mathbb{E}_t\left[M_{t,t+1}R_{t,t+1}\right] + k\mathbb{E}_t\left[M_{t,t+2}R_{t,t+2}\right]. \end{aligned} \tag{5}$$

[9] I use quotation marks because the individual terms $\mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$ do not themselves satisfy the standard Euler equation condition of equalling one.

3.3 Horizon Euler Terms and Estimation Bias

In Equation 5, either both terms exactly equal one, or one falls below one while the other exceeds one[10]. If an investor's true horizon is short, their one-period Euler equation holds ($\mathbb{E}_t[M_{t,t+1}R_{t,t+1}] = 1$) and, by the law of iterated expectations and the multiplicative structure of the SDF ($M_{t,t+2} = M_{t,t+1} \cdot M_{t+1,t+2}$), recursively extends to all longer horizons ($\mathbb{E}_t[M_{t,t+2}R_{t,t+2}] = 1$)[11]. In this scenario, Equation 5 holds trivially for any k assumed by the econometrician.

Conversely, if the investor places positive weight on the long horizon ($k > 0$), the FOC no longer *forces* the one-period Euler equation to hold, and $F(1)$ can lie either above or below one. Note that Equation 5 always holds as a whole; what varies is the value of each individual horizon Euler term.

Suppose first that $\mathbb{E}_t[M_{t,t+1}R_{t,t+1}] < 1$ and $\mathbb{E}_t[M_{t,t+2}R_{t,t+2}] > 1$. The short-horizon term being less than one means that, considered in isolation, the investor would prefer to sell (or short) the asset[12]. The long-horizon term being greater than one means that, considered in isolation, the investor would prefer to buy the asset and hold it for two periods. The investor thus balances opposing short- and long-horizon incentives to satisfy the FOC. This configuration is what the stock market offers: over a single period, the expected gain is too low to compensate for the high volatility, so a short-term investor would avoid it[13], yet the long-run expected return is attractive enough to justify holding.

Consider an econometrician who assumes a short investment horizon when the true horizon is longer. The econometrician would then observe Equation 6, which is less than one under the agent's true beliefs and risk preferences. Estimating the model under the constraint that the short-horizon Euler equation equals one yields an implausibly high β or risk-aversion γ, leading to the well-known equity premium and risk-free rate puzzles[14].

[10] This dichotomy requires $k \in (0, 1)$. At the endpoint $k = 1$, the FOC restricts only the two-period term, $F(2) = 1$, and leaves $F(1)$ unrestricted; symmetrically, at $k = 0$ only $F(1) = 1$ is restricted.

[11] Formally, the investor satisfies the one-period Euler equation at every date and returns are multiplicative, $R_{t,t+2} = R_{t,t+1}R_{t+1,t+2}$, so

$$\mathbb{E}_t[M_{t,t+2}R_{t,t+2}] = \mathbb{E}_t\big[M_{t,t+1}R_{t,t+1}\,\mathbb{E}_{t+1}[M_{t+1,t+2}R_{t+1,t+2}]\big] = \mathbb{E}_t[M_{t,t+1}R_{t,t+1}] = 1.$$

[12] $\mathbb{E}_t[M_{t,t+1}R_{t,t+1}] < 1 \Rightarrow U'(c_t) > \mathbb{E}_t[\beta U'(c_{t+1})R_{t,t+1}]$. The investor would want to increase current consumption and reduce investment to restore equality.

[13] Consider an example: if you are buying a house in one week, would you put your cash in the stock market for only that single week?

[14] To be more precise, the argument relates to excess returns. Higher γ increases the curvature of the SDF and raises the expectation toward one, so matching the data requires an implausibly high γ (Mehra and Prescott, 1985) or β (Weil, 1989). We turn on empirical examples of horizon Euler terms later in the paper.

$$\mathbb{E}_t\left[M_{t,t+1}R_{t,t+1}\right] < 1. \tag{6}$$

The opposite configuration, $\mathbb{E}_t[M_{t,t+1}R_{t,t+1}] > 1$ and $\mathbb{E}_t[M_{t,t+2}R_{t,t+2}] < 1$, reverses these incentives: the asset is highly attractive to hold for the immediate period, whereas its expected performance deteriorates over the longer horizon. A classic example is a momentum strategy: short-term expected returns are exceptionally high (Jegadeesh and Titman, 1993), yet the long-horizon Euler term falls below one. A related distortion arises in estimation: if an econometrician assumes only the short-horizon Euler equation while the true horizon is longer and $\mathbb{E}_t\left[M_{t,t+1}R_{t,t+1}\right] > 1$, the resulting estimate of risk aversion will be biased downward. Menkhoff (2011) provide survey evidence that momentum traders are characterized by short-term investment horizons and a lower degree of risk aversion.

Which configuration applies in practice depends on how the horizon Euler term $F(h)$ varies with h: if $F(h)$ increases with h, the first configuration prevails; if it decreases, the second prevails. In general, $F(h)$ need not be monotone: it may increase over shorter horizons and decrease over longer ones, or vice versa, so both configurations can apply to the same asset at different horizons. I next provide conditions under which $F(h)$ is increasing in h for moderate horizons, which I argue is the predominantly relevant empirical case. It is also worth recalling that at this stage we consider an investor with a given horizon and say nothing about its optimality; later we will show that $F(h)$ depends on the investor's k.

3.4 Random Walk Dynamics and Horizon Irrelevance

This subsection introduces a benchmark: if log consumption and log returns follow a random walk, the investor's horizon does not affect the pricing condition. Assume CRRA utility with risk aversion γ, and suppose log consumption and log returns have i.i.d. jointly normal increments. Then cumulative log consumption growth satisfies $\ln(c_{t+h}/c_t) \sim \mathcal{N}(gh,\, \sigma_c^2 h)$. Cumulative log returns satisfy $r_{t,t+h} \sim \mathcal{N}(\mu_r h,\, \sigma_r^2 h)$, with per-period covariance $\sigma_{c,r}$.[15] Under CRRA, $\ln(M_{t,t+h}R_{t,t+h}) = h\ln\beta - \gamma\ln(c_{t+h}/c_t) + r_{t,t+h}$ is normal because it is a linear combination of jointly normal variables. Applying the log-normal moment $\mathbb{E}_t[\exp(X)] = \exp(\mathbb{E}_t[X] + \frac{1}{2}\operatorname{Var}_t(X))$ then gives[16]:

$$F(h) = \mathbb{E}_t[M_{t,t+h}R_{t,t+h}] = \exp(\psi \cdot h), \quad \psi \equiv \ln\beta - \gamma g + \mu_r + \tfrac{1}{2}\gamma^2\sigma_c^2 + \tfrac{1}{2}\sigma_r^2 - \gamma\sigma_{c,r}. \tag{7}$$

[15] These are standard assumptions in the consumption-based asset pricing literature.

[16] Let $x_h = h\ln\beta - \gamma\ln(c_{t+h}/c_t) + r_{t,t+h}$. Then $\mathbb{E}_t[x_h] = h(\ln\beta - \gamma g + \mu_r)$ and $\operatorname{Var}_t(x_h) = h(\gamma^2\sigma_c^2 + \sigma_r^2 - 2\gamma\sigma_{c,r})$, using $\operatorname{Var}(aX + bY) = a^2\operatorname{Var}(X) + b^2\operatorname{Var}(Y) + 2ab\operatorname{Cov}(X,Y)$ and the fact that all variances and covariances scale linearly with h under the random walk.

The key implication of the random walk assumption is that both the conditional mean and the conditional variance of $\ln(M_{t,t+h}R_{t,t+h})$ scale linearly with h, so every horizon is governed by the same constant ψ. Before imposing equilibrium, we can see that $\psi > 0$ would make $F(h)$ increase with h, $\psi < 0$ would make it decrease, and $\psi = 0$ would make it flat. However, the pricing condition itself rules out the first two cases. Substituting $F(h) = \exp(\psi \cdot h)$ into Equation 5 yields

$$(1-k)\exp(\psi) + k\exp(2\psi) = 1. \tag{8}$$

Proposition 1 *(Horizon irrelevance under random-walk dynamics).* Assume CRRA utility, and suppose log consumption and log returns have i.i.d. jointly normal increments. Then the FOC Equation 5 implies $F(h) = 1$ for every horizon h and every $k \in [0, 1]$.

Proof. Set $y = \exp(\psi) > 0$. If $k \in (0, 1]$, Equation 8 becomes $ky^2 + (1-k)y - 1 = 0$, whose roots are $y = 1$ and $y = -1/k$; only $y = 1$ is admissible. If $k = 0$, the equation reduces directly to $y = 1$. Hence $\psi = 0$ in all cases, and Equation 7 gives $F(h) = \exp(\psi h) = 1$ for every horizon h. ■

In particular, $F(1) = F(2) = 1$ for every k[17]. The horizon function Equation 3 therefore has no effect, and the FOC Equation 5 collapses to the standard one-period Euler equation regardless of the investor's true horizon. Under the random walk benchmark, the investor's horizon is therefore irrelevant for pricing. This result parallels (Samuelson, 1969, pp. 244–245), who shows that, with i.i.d. returns and isoelastic utility, the optimal portfolio share chosen each period is independent of the terminal date T, and hence of the investor's remaining lifetime $T - t$; here, when $\ln(M_{t,t+h}R_{t,t+h})$ follows a random walk, the horizon is likewise irrelevant for the pricing condition in this model. Empirically, the random walk hypothesis is widely rejected for asset returns (Lo and MacKinlay, 1988). The benchmark therefore functions as an organizing device rather than a literal description of the data. Once uncertainty no longer accumulates linearly with time, horizon dependence arises even if each fixed-horizon marginal distribution remains log-normal.

3.5 Condition for Increasing and Decreasing Euler Terms

The preceding subsection showed that under the random walk benchmark, where both the conditional mean and the conditional variance of $\ln(M_{t,t+h}R_{t,t+h})$ scale linearly with h, the pricing condition implies that the horizon Euler term $F(h) = \mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$ is

[17] Crucially, the binding restriction stems from the random walk *dynamics*, not merely from log-normality of the marginal distribution at a given horizon: a process with richer temporal structure—such as persistent or mean-reverting consumption growth—can produce log-normal marginals at each horizon while still generating $F(h) \neq 1$.

constant in h, equal to one. To obtain a non-constant $F(h)$, we must relax the random walk assumption and allow richer dynamics in consumption growth or returns. I consider two cases: the variance of $\ln(M_{t,t+h}R_{t,t+h})$ grows faster than linearly in h (super-linear growth), and the variance grows slower than linearly in h (sub-linear growth). The former implies an increasing $F(h)$, whereas the latter implies a decreasing $F(h)$.

Let $X_h = \ln(M_{t,t+h}R_{t,t+h})$ denote the log of the product of the stochastic discount factor and the cumulative return over horizon h. I assume that X_h is conditionally normally distributed at each horizon h—a restriction on the marginal distribution at each horizon that does not constrain the temporal dynamics across horizons. In particular, the variance $V_h \equiv \mathrm{Var}_t(X_h)$ need not grow linearly with h. I refer to this as the *log-normal framework* throughout, to distinguish it from the full random walk in which both log-normality and linear variance growth hold simultaneously. Under this assumption, the horizon Euler term $F(h)$ is given by:

$$F(h) = \mathbb{E}_t[\exp(X_h)] = \exp\left(\mathbb{E}_t[X_h] + \frac{1}{2}\,\mathrm{Var}_t(X_h)\right). \tag{9}$$

Let $V_h \equiv \mathrm{Var}_t(X_h)$. To determine whether the Euler term is increasing or decreasing with the horizon, we differentiate $F(h)$ with respect to h:

$$\frac{dF(h)}{dh} = F(h)\left(\frac{d\mathbb{E}_t[X_h]}{dh} + \frac{1}{2}\frac{dV_h}{dh}\right). \tag{10}$$

Since $F(h)$ is strictly positive, the sign of $F'(h)$ is determined entirely by the sign of the term in the parentheses. Following standard asset pricing literature[18], we assume the expected value of the log product scales linearly with time such that $\mathbb{E}_t[X_h] = \mu h$, where $\mu < 0$ is the constant per-period drift[19]. Consequently, $\frac{d\mathbb{E}_t[X_h]}{dh} = \mu$.

Substituting this constant drift into Equation 10, we isolate the condition for an increasing Euler term ($F'(h) > 0$):

$$\begin{aligned} \mu + \frac{1}{2}\frac{dV_h}{dh} &> 0 \\ \frac{dV_h}{dh} &> -2\mu. \end{aligned}$$

[18]This linear-in-horizon drift specification is standard in log-normal asset-pricing environments with stationary increments; see, for example, (Cochrane, 2005, Ch. 1).

[19]The FOC itself implies $\mu < 0$: since $(1-k)\exp(\mu + \frac{1}{2}V_1) + k\exp(2\mu + \frac{1}{2}V_2) = 1$ is a weighted average of positive terms, at least one exponent is weakly negative, so $\mu \le -\frac{1}{2}V_1$ or $\mu \le -\frac{1}{4}V_2$; either way $\mu < 0$ whenever the variances are strictly positive. In the special case $F(1) = 1$ this sharpens to the familiar $\mu = -\frac{1}{2}V_1$.

We can link this inequality to the random walk benchmark. Under a pure random walk, X_h has i.i.d. increments, so $V_h = \sigma^2 h$ for some constant σ^2, and the condition $F'(h) = 0$ from the expression above becomes $\mu + \frac{1}{2}\sigma^2 = 0$, i.e. $\sigma^2 = -2\mu$. We define this required benchmark variance rate as $\sigma^2_{RW} \equiv -2\mu$. Equivalently, this result follows from the fundamental pricing equation: because $\mathbb{E}_t[\exp(X_h)] = 1$ in equilibrium[20], the log-normal moment condition requires $\mathbb{E}_t[X_h] + \frac{1}{2}\operatorname{Var}_t(X_h) = 0$, so that the variance per period must exactly offset twice the negative drift, confirming $\sigma^2_{RW} = -2\mu$.

- **Increasing Euler Term ($F'(h) > 0$):** Occurs at horizon h if and only if $\frac{dV_h}{dh} > \sigma^2_{RW}$. The variance of X_h grows faster than the random walk benchmark at that horizon (locally super-linear growth).

- **Decreasing Euler Term ($F'(h) < 0$):** Occurs at horizon h if and only if $\frac{dV_h}{dh} < \sigma^2_{RW}$. The variance of X_h grows slower than the random walk benchmark at that horizon (locally sub-linear growth).

Remark 1 *(Random-walk departures and the variance ratio).* Whether V_h grows faster or slower than the random-walk benchmark $h\,V$ is, in essence, a statement about serial dependence in the increments of X_h. Writing $X_h = \sum_{j=1}^{h} \Delta X_j$ with $\Delta X_j \equiv X_j - X_{j-1}$ and assuming stationary increments with per-period variance V, we obtain the standard variance decomposition

$$V_h \;=\; h\,V \;+\; 2\sum_{1\le i<j\le h} \operatorname{Cov}_t(\Delta X_i,\, \Delta X_j). \tag{11}$$

The first term is the random-walk benchmark itself; the second collects all departures from it. Under a pure random walk the increments are i.i.d., the cross-covariances vanish, and $V_h = h\,V$, recovering the constant Euler term of the previous subsection. $F(h)$ therefore exceeds the compounded one-period term $F(1)^h$ precisely when the cumulative cross-covariance is positive—the increments of the SDF–return product exhibit *persistence*—and falls short of it when the covariance is negative (*mean reversion*). Equivalently, the variance ratio $V_h/(h\,V)$, introduced by Lo and MacKinlay (1988) as a specification test of the random walk hypothesis, equals one under the random walk null, exceeds one in the persistence case, and falls below one in the mean-reverting case—though on the weekly return data they study, the variance ratio is above unity, providing no support for mean reversion at that frequency.

[20]Note that we are comparing against the random walk benchmark, under which, as shown above, all horizon Euler terms are exactly one.

Returning to the parsimonious two-period model, if the horizon Euler terms satisfy $F(1) < 1$ and $F(2) > 1$, then the variance of X_h must grow faster than the random-walk benchmark between $h = 1$ and $h = 2$. Moreover, this configuration is consistent with the FOC Equation 5 only for an investor with $k \in (0, 1)$. Figure 2 illustrates the increasing-Euler-term case in this setting, comparing the level dynamics and dispersion of $M_{t,t+h}R_{t,t+h}$ with the random-walk benchmark. An econometrician who incorrectly assumes this investor satisfies a one-period Euler equation would encounter the equity premium and risk-free rate puzzles, since forcing $F(1) = 1$ requires inflating the risk-aversion parameter and the discount factor. Note that at this stage I do not endogenize the choice of the horizon function, but rather solve the model conditional on a given rule.

Models that resolve the equity premium and risk-free rate puzzles relax the random walk assumption. (Parker and Julliard, 2005, Table 4, col. 7, and fn. 16) demonstrate that consumption growth does not follow a random walk: its variance grows faster than linearly up to two years and slower after that. They show that a three-year horizon ($S = 11$ quarters following the return) provides a significantly better cross-sectional explanation of portfolio return variation than any shorter horizon.

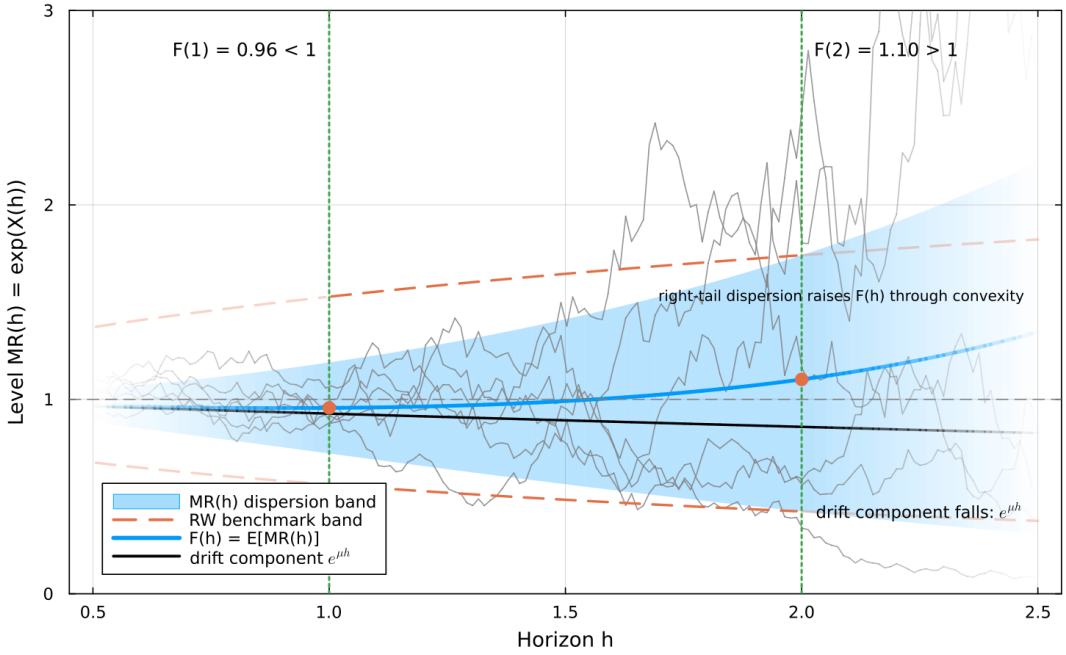
More broadly, these puzzles are typically resolved by modifying one of three ingredients: how investors evaluate consumption—as in the habit formation model of Campbell and Cochrane (1999); the dynamics of the consumption growth process—as in the long-run risk model of Bansal and Yaron (2004); or the distribution of shocks to consumption and returns—as in the rare disaster model of Barro (2006) and the rare-event framework of Julliard and Ghosh (2012). Using the habit formation model of Constantinides (1990), Daniel and Marshall (1997) conclude that the equity-premium and risk-free-rate puzzles are primarily problems for shorter-horizon returns.

3.6 Optimal Investment Horizons

Earlier in this section, we analyzed the model conditional on the exogenous horizon function Equation 3. I now relax this constraint and allow the optimal horizon function to be determined endogenously. Thus, investors choose not only how much to invest but also how long to hold their investments. For example, an investor in the stock market may know their minimum holding period will exceed one week, but must still decide whether to hold for one month, one year, or several years.

Consider the same parsimonious model as in Equation 4, but now the investor chooses the horizon function $H(\tau)$. The horizon parameter chosen at time t, denoted k_t, applies specifically to the asset a_t purchased at time t. In the next period, the investor buys a new

Figure 2: Horizon Euler Term under Negative Drift and Accelerating Variance



Notes: The figure shows the example of the horizon Euler term $F(h) = \mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$. Although the drift component $e^{\mu h}$ falls with the horizon because $\mu < 0$, the dispersion of $MR(h) = e^{X_h}$ increases sufficiently fast that convexity raises $F(h)$, so that $F(1) < 1$ while $F(2) > 1$. The blue band shows the dispersion of $MR(h)$, growing faster than a linear (random-walk) benchmark, and the simulated paths are illustrative. To build the figure: at $h = 2$ the random-walk process and the X process have the same variance, yet the variance of the X process grows faster, so at $h = 1$ it is lower than the random-walk benchmark. $X_h \sim \mathcal{N}(\mu h,\, V_h)$ with the accelerating variance schedule $V_h = \sigma^2 h^3$, $\sigma = 0.25$; the drift μ is the unique root of the two-period FOC Equation 5, $0.7\, e^{\mu + V_1/2} + 0.3\, e^{2\mu + V_2/2} = 1$ ($k = 0.3$), which gives $\mu = -0.076$, $F(1) = 0.96$, and $F(2) = 1.10$. The solid blue line is $F(h) = e^{\mu h + V_h/2}$, the black line is $e^{\mu h}$, and the two dots mark $F(1)$ and $F(2)$. The blue band is $\exp(\mu h \pm \sqrt{V_h})$. The dashed band is the random-walk benchmark, $\exp(\mu h \pm \sigma_{RW}\sqrt{h})$: the same drift μ but linear variance $V_h^{RW} = \sigma_{RW}^2 h$, with the rate anchored at the long horizon, $\sigma_{RW}^2 = V_2/2 = 0.25$. The two bands therefore coincide at $h = 2$, and the random-walk band is wider at every shorter horizon: the model's variance lies below the random-walk ray before the long horizon and catches up only at $h = 2$. Vertical dotted lines mark the two model horizons $h = 1, 2$; the region outside $[1, 2]$ is drawn faded.

asset a_{t+1} and chooses a new horizon parameter k_{t+1}. Thus, the investor does not choose one permanent k that governs all future investments. Instead, the investor chooses today's horizon for today's investment, and future investments are re-optimized when those future dates arrive. The first-order condition (FOC1) with respect to the asset choice a_t is

$$\begin{aligned} & U'(c_t)p_t = \mathbb{E}_t\left[(1-k_t)\beta U'(c_{t+1})p_{t+1} + k_t\beta^2 U'(c_{t+2})p_{t+2}\right] \\ \Rightarrow \quad & 1 = (1-k_t)\mathbb{E}_t\left[M_{t,t+1}R_{t,t+1}\right] + k_t\mathbb{E}_t\left[M_{t,t+2}R_{t,t+2}\right], \end{aligned} \tag{12}$$

and the first-order condition (FOC2) with respect to the horizon parameter k_t, for interior $k_t \in (0,1)$, is:

$$\begin{aligned} & a_t\,\mathbb{E}_t[-\beta U'(c_{t+1})p_{t+1} + \beta^2 U'(c_{t+2})p_{t+2}] = 0 \\ \Rightarrow \quad & \text{For an interior solution } a_t \neq 0, \text{ dividing by } U'(c_t)p_t \text{ yields:} \\ & \mathbb{E}_t[M_{t,t+2}R_{t,t+2}] - \mathbb{E}_t[M_{t,t+1}R_{t,t+1}] = 0. \end{aligned} \tag{13}$$

Because k_t is chosen specifically for the current investment, the optimality condition equates the expected marginal value of the asset's payoff across dates $t+1$ and $t+2$. One trivial case is if the investor has a short horizon: both horizon Euler terms equal one, so k_t is not identified. One may still interpret the short-horizon case as corresponding to $k_t = 0$.

For each candidate k_t, the investor solves for the optimal a_t so that FOC1 holds. Let $a_t^*(k)$ denote that optimal portfolio choice. Define the horizon Euler term

$$F_t(h;k) \equiv \mathbb{E}_t\left[M_{t,t+h}(a_t^*(k))\,R_{t,t+h}\right]. \tag{14}$$

Then Equation 12 can be written compactly as

$$(1-k_t)F_t(1;k_t) + k_t F_t(2;k_t) = 1. \tag{15}$$

Finally, define the local gap between two adjacent horizon Euler terms:

$$G_t(h;k_t) \equiv F_t(h+1;k_t) - F_t(h;k_t). \tag{16}$$

For simplicity, I focus on long positions, $a_t > 0$. If $a_t < 0$, the case is symmetric: the marginal condition is multiplied by a negative position, so the sign of the horizon adjustment is reversed. Theorem 1 at the end of this section collects the mirrored results.

The sign of the gap is the key object. Since the marginal value of increasing k_t is proportional to $a_t G_t(1;k_t)$, and since $a_t > 0$ in the case considered here, the sign of $G_t(1;k_t)$ tells us whether the investor would like to lengthen or shorten the horizon. If $G_t(1;k_t) > 0$,

then shifting the payoff of the date-t investment from the next period, $t+1$, to the following period, $t+2$, increases its marginal value. In this case, the longer horizon is locally more attractive. If $G_t(1;k_t) < 0$, the opposite is true: the payoff is more valuable at the shorter horizon. Finally, if $G_t(1;k_t) = 0$, the investor is exactly indifferent at the margin between the two adjacent horizons, so the horizon optimality condition is satisfied.

Implications of the First-Order Conditions

In this part, we interpret the first-order conditions and step through the investor's optimization problem to identify what determines the optimal investment horizon. We will see how the optimum relates to the random walk dynamics of consumption and returns, and provide a Horizon-Adjusted Fundamental Pricing Equation (HAFPE). To simplify notation, we suppress the time subscripts and write k, F, and G for k_t, F_t, and G_t when there is no ambiguity.

Because k is constrained to $[0,1]$, the equality in FOC2 characterizes only interior optima $k \in (0,1)$. We begin with the boundary cases. At $k=0$, FOC1 implies $F(1;0)=1$; by the recursive argument established above for the short-horizon case, this extends to $F(2;0)=1$, and more generally to $F(h;0)=1$ for all h. FOC2 therefore holds trivially. Both first-order conditions are therefore satisfied, so $k=0$ trivially appears optimal. However, this apparent optimality is illusory: as noted above, k is not identified in this case—any value satisfies FOC1—so $k=0$ is not truly optimal—the FOCs just happen to hold. To determine whether the investor actually benefits from lengthening the horizon, we must examine the gap from FOC2, $G(1;k) = F(2;k) - F(1;k)$, in a neighborhood of the boundary. We rearrange FOC1 to express $G(1;k)$ in two equivalent ways:

$$\begin{aligned} 1 &= F(1;k) + k[F(2;k) - F(1;k)] \\ \Rightarrow \quad [F(2;k) - F(1;k)] &= \frac{1 - F(1;k)}{k} \end{aligned} \tag{17}$$

$$G(1;k) = \frac{1 - F(1;k)}{k} = \frac{F(2;k) - 1}{1-k}. \tag{18}$$

At the boundaries, FOC1 reduces to $F(1;0)=1$ when $k=0$ and to $F(2;1)=1$ when $k=1$, so each representation takes the form $0/0$. Evaluating the one-sided limits:

$$\begin{aligned} \lim_{k\to 0^+} G(1;k) &= \lim_{k\to 0^+} \frac{1-F(1;k)}{k} = -\left.\frac{dF(1;k)}{dk}\right|_{k=0}, \\ \lim_{k\to 1^-} G(1;k) &= \lim_{k\to 1^-} \frac{F(2;k)-1}{1-k} = -\left.\frac{dF(2;k)}{dk}\right|_{k=1}. \end{aligned} \tag{19}$$

Suppose the investor's optimal horizon exceeds one period. We now show that $k = 0$ is not optimal, even though $F(1;0) = 1$ and both FOCs hold. To do so, we demonstrate that $\lim_{k\to 0^+} G(1;k) > 0$, which implies that FOC2 is positive, and it is optimal for the investor to increase k. As discussed above, having an optimal horizon $k^* > 0$ requires $F(1;k) < 1$ and $F(2;k) > 1$ for every $k \in (0, k^*)$.[21] Combined with the boundary value $F(1;0) = 1$, the inequality $F(1;k) < 1$ for $k > 0$ means $F(1;k)$ decreases as k rises from zero. Consequently, $dF(1;k)/dk|_{k=0} < 0$, which implies that the right-hand limit in Equation 19 is strictly positive. This confirms that the investor benefits from lengthening the horizon away from the short-horizon boundary, so $k = 0$ is not a true optimum. This is the only boundary conclusion needed for an interior optimum $k^* < 1$: the gap is strictly positive throughout the initial segment $k \in [0, k^*)$, with the value at $k = 0$ interpreted as the one-sided limit. If instead the optimum is the corner $k^* = 1$, the symmetric upper-boundary argument applies as well: $F(2;1) = 1$ together with $F(2;k) > 1$ for $k < 1$ gives $dF(2;k)/dk|_{k=1} < 0$, so the left-hand limit in Equation 19 is strictly positive and $G(1;k) > 0$ on the closed interval $[0, 1]$, where the endpoint values at $k = 0$ and $k = 1$ are understood as the corresponding one-sided limits.

Having ruled out the trivial optimum at $k = 0$, the next step is to describe the procedure that locates k^*. The procedure has two steps.

Step 1: Direction of optimization. Assume that the investor's optimal horizon is not short. Starting from $k = 0$, consider increasing the weight on the longer horizon to some $k' > 0$ and re-optimizing a_t so that FOC1 continues to hold. If FOC2 > 0, then increasing k' raises expected utility, so the optimum lies at a higher value of k. If FOC2 is eventually satisfied at some interior $k^* \in (0, 1)$, we obtain an interior optimum; if instead FOC2 remains positive for all $k \in (0, 1)$, the optimum is a corner solution.

Step 2: Change horizon function. In the previous step, the investor either finds an interior optimum k^* or reaches the corner $k^* = 1$. In the latter case, the investor shifts to a new horizon function that begins one period later. For example, starting from the initial segment $\{1, 2\}$ and reaching $k = 1$, the investor moves to the segment $\{2, 3\}$, where the horizon function takes the form

$$H(1) = 0, \quad H(2) = 1 - k, \quad H(3) = k, \quad k \in [0, 1]. \tag{20}$$

The same logic applies when the investor is choosing between any two adjacent horizons.

[21] Specifically, in the empirically relevant case for the stock market, the variance of X_h grows super-linearly across this segment at the re-optimized positions, so the horizon Euler term is increasing, $F(2;k) > F(1;k)$ (Subsection 3.5). In the opposite case, $F(1;k) > 1 > F(2;k)$, FOC2 is negative, so the investor would shorten the horizon.

For an arbitrary starting horizon $h \geq 1$, define the segment-specific horizon function

$$H(\tau) = 0 \quad \text{for } \tau < h, \qquad H(h) = 1 - k, \qquad H(h+1) = k, \qquad k \in [0, 1]. \tag{21}$$

The current investment a_t is now split between dates $t + h$ and $t + h + 1$. FOC1 becomes

$$1 = (1 - k)F(h; k) + kF(h + 1; k). \tag{22}$$

The direct effect of increasing k is to shift the payoff from $t + h$ to $t + h + 1$:

$$\begin{aligned} & \mathbb{E}_t \left\{ a_t \left[-\beta^h U'(c_{t+h}) p_{t+h} + \beta^{h+1} U'(c_{t+h+1}) p_{t+h+1} \right] \right\} = 0 \\ \Rightarrow \qquad & F(h + 1; k) - F(h; k) = 0. \end{aligned} \tag{23}$$

Thus, on the segment $\{h, h + 1\}$, the two first-order conditions reduce to

$$\begin{aligned} &\text{FOC1:} \qquad (1 - k)F(h; k) + kF(h + 1; k) = 1, \\ &\text{FOC2:} \qquad G(h; k) = F(h + 1; k) - F(h; k) = 0. \end{aligned} \tag{24}$$

This is the clean discrete characterization of the optimal horizon. The investor starts with the segment $\{1, 2\}$. If an interior value of k solves the two equations, the optimal horizon has been found. If the solution reaches $k = 1$ and the gap still points toward a longer horizon, the investor moves to the next segment $\{2, 3\}$ and repeats the same argument. The procedure continues across adjacent segments until the local gap is zero at an interior k^*,

$$\begin{aligned} \text{FOC1:} \qquad & (1 - k^*)F(h^*; k^*) + k^* F(h^*{+}1; k^*) = 1, \\ \text{FOC2 (Gap condition)} \qquad & G(h^*; k^*) = F(h^*{+}1; k^*) - F(h^*; k^*) = 0, \end{aligned} \tag{25}$$

or until it reaches a grid point where the previous segment ends at its corner $k = 1$ with a positive gap, while the next segment $\{h^*, h^*{+}1\}$ opens at $k^* = 0$ with a nonpositive gap, so the investor moves all weight to h^* but no further,

$$\begin{aligned} \text{FOC1:} \qquad & F(h^*; 0) = 1, \\ \text{FOC2 (Gap condition)} \qquad & G(h^*; 0) = F(h^*{+}1; 0) - F(h^*; 0) \leq 0. \end{aligned} \tag{26}$$

3.7 Horizon-Adjusted Fundamental Pricing Equation

Proposition 2 *(Continuous-horizon limit).* Consider a horizon grid with spacing $\Delta h > 0$, so that on each segment the investor places weight $(1 - k)$ on horizon h and weight k

on horizon $h + \Delta h$. In the limit $\Delta h \to 0$, the discrete FOCs Equation 25 reduce to[22]

$$\begin{aligned} &\text{FOC1:} \quad F(h^*) \equiv \mathbb{E}_t \left[M_{t,t+h^*} R_{t,t+h^*} \right] = 1, \\ &\text{FOC2:} \quad F'(h^*) = 0. \end{aligned} \tag{27}$$

The two conditions play different roles: FOC1 pins down the level, $F(h^*) = 1$, and FOC2 sets the slope to zero, $F'(h^*) = 0$. For any given h, FOC1 can be satisfied by choosing the appropriate asset position a_t. If FOC2 does not also hold, however, $F(h)$ crosses the horizontal line at one rather than touching it. Thus, the pair provides a value-matching and smooth-pasting characterization of the optimal horizon: the Euler term meets its benchmark of one smoothly rather than crossing it, since $F(h) \leq 1$ near the optimum (Remark 4). Figure 3 illustrates the two cases.

Proof. On a segment with grid spacing Δh, the horizon function assigns weight $(1 - k)$ to horizon h and weight k to horizon $h + \Delta h$.

FOC1. The first-order condition with respect to a_t is

$$(1 - k)\, F(h; k) + k\, F(h + \Delta h;\, k) = 1.$$

At an interior optimum, FOC2 requires

$$F(h + \Delta h;\, k) - F(h;\, k) = 0.$$

At the interior optimum (h^*, k^*)[23], both horizon Euler terms equal one: $F(h^*; k^*) = F(h^*{+}\Delta h;\, k^*) = 1$. As $\Delta h \to 0$, both evaluation points h^* and $h^*{+}\Delta h$ converge to the same horizon h^*, and by continuity of F the condition reduces to $F(h^*) = 1$.

FOC2. The optimality condition with respect to k requires $G(h^*; k^*) = 0$, so the Euler term takes the same value at the two ends of the segment:

$$F(h^* + \Delta h;\, k^*) - F(h^*;\, k^*) = 0.$$

By the mean value theorem, the horizon derivative therefore vanishes at some interior point

[22] Throughout the paper, the horizon Euler term is twice continuously differentiable in the horizon, and its first two horizon derivatives are jointly continuous in the horizon and the weight k. I maintain this regularity for every process considered and do not restate it.

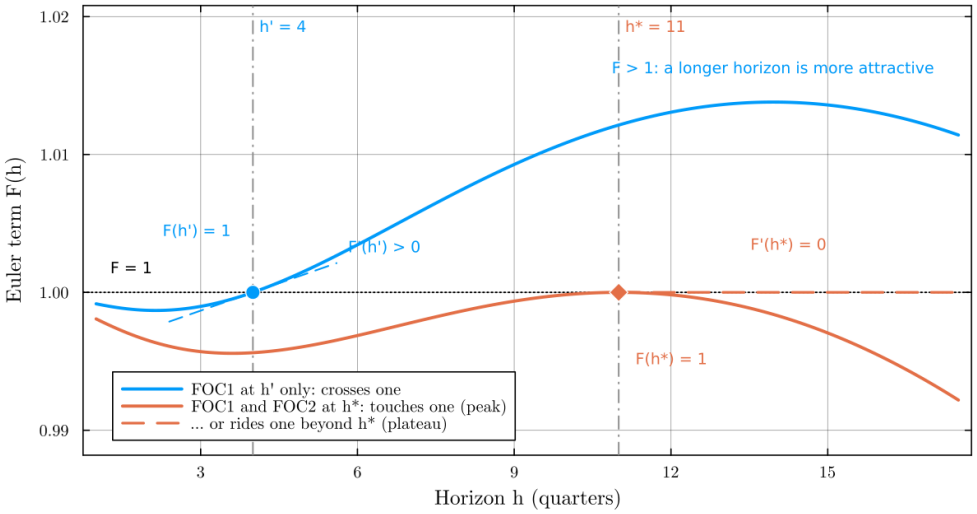
[23] The discrete optimum (h^*, k^*) depends on the grid spacing Δh. The continuous limit assumes that $(h^*(\Delta h),\, k^*(\Delta h)) \to (h^*,\, k^*)$ as $\Delta h \to 0$; the continuity of the derivatives of $F(h; k)$ near that limit is the regularity maintained throughout.

$\xi_{\Delta h} \in (h^*,\, h^* + \Delta h)$:

$$F_h(\xi_{\Delta h};\, k^*) = \frac{F(h^* + \Delta h;\, k^*) - F(h^*;\, k^*)}{\Delta h} = 0. \tag{28}$$

As $\Delta h \to 0$, the interval $(h^*,\, h^* + \Delta h)$ shrinks to its left endpoint, so $\xi_{\Delta h} \to h^*$, and continuity of F_h gives $F'(h^*) \equiv F_h(h^*; k^*) = 0$. ■

Figure 3: Value Matching and Smooth Pasting at the Optimal Horizon



Notes: The figure illustrates value matching and smooth pasting at the optimal horizon (illustrative calibration with $h' = 4$ and $h^* = 11$ quarters). Each curve plots the horizon Euler term $F(h)$, holding the asset position fixed at the value that satisfies FOC1 at the marked horizon. The level condition can hold at any candidate horizon: $F(h') = 1$. When FOC2 fails, however, the slope is positive, $F'(h') > 0$ (dashed tangent), so the Euler term crosses the horizontal line at one rather than touching it, and $F > 1$ beyond h' means a longer horizon is more attractive: h' is not optimal. At h^* both conditions hold, $F(h^*) = 1$ and $F'(h^*) = 0$, and the Euler term meets its benchmark of one smoothly. Beyond h^* the curve either falls back below one (peak, solid) or stays at one (plateau, dashed), the two cases of Figure 4.

Remark 2. In the Horizon-Adjusted Fundamental Pricing Equation (HAFPE) condition, $F'(h)$ denotes the horizon derivative holding the current asset position fixed. That is, if $a^*(h)$ satisfies FOC1,

$$F(a^*(h), h) = 1,$$

then

$$F'(h) \equiv F_h(a^*(h), h) = \left.\frac{\partial F(a, h)}{\partial h}\right|_{a = a^*(h)}.$$

It is not the total derivative $\frac{d}{dh}F(a^*(h), h)$, which equals zero by construction along the re-optimized FOC1 path. Instead, it is the partial horizon derivative evaluated at the

asset position chosen for the current segment. This follows directly from the segmented algorithm. On the segment, FOC1 holds for the given weight k^* and pins down the asset position $a_t^*(k^*)$; FOC2 then holds at this given $a_t^*(k^*)$. Both terms of the finite difference in Equation 28 are therefore evaluated at the same position, so its limit is the partial derivative in h.

Remark 3 *(Segmented Optimization and Discrete Horizon Dynamics).* It is worth noting that the optimization problem is solved within a discrete-time framework, which many find easier to follow and analytically more tractable. By defining the investment horizon as a convex combination of two adjacent periods—weighted by the parameter $k \in [0, 1]$ on successive segments $\{h, h+1\}$—the agent's decision-making process is reduced to a smooth optimization over a continuous parameter within a discrete grid.

This approach offers two distinct advantages over a direct optimization of a discrete horizon h:

1. **Analytical Tractability:** Optimizing over k allows for the application of standard calculus and first-order conditions without the technical difficulties of taking derivatives of discrete-valued functions or employing Dirac delta functions to represent point-mass horizons in continuous time.

2. **Structural Intuition:** The framework provides a more intuitive path to the continuous-horizon limit. As the grid spacing $\Delta h \to 0$, the discrete first-order conditions converge cleanly to $F(h^*) = 1$ and $F'(h^*) = 0$.

Consequently, this segmented approach clarifies the underlying economic mechanism while ensuring the model remains easy to optimize and interpret.

The optimal choice is also stable over time. The next proposition shows that if no news arrives, the investor who holds the position a^* with horizon h^* has nothing to change next period: he does not trade, and the horizon rolls forward with him, just as the driver's field of vision travels with the car and the wheel stays still while the road looks as expected.

Proposition 3 *(No trade without news).* Let (a^*, h^*) satisfy FOC1 and FOC2 at date t, and let $F_{t+1}(a, h) \equiv \mathbb{E}_{t+1}[M_{t+1,t+1+h} R_{t+1,t+1+h}]$ denote the date-$(t+1)$ Euler term. Suppose no news arrives between t and $t+1$: realized wealth and income equal their date-t expected values, and the Euler term is unchanged,

$$F_{t+1}(a, h) = F_t(a, h) \qquad \text{for every position } a \text{ and horizon } h.$$

Then (a^*, h^*) satisfies FOC1 and FOC2 at date $t+1$: the investor does not trade, keeps the horizon h^*, and the terminal date rolls from $t + h^*$ to $t + 1 + h^*$.

Proof. Because wealth and income realize as expected, the consumption entering the marginal utilities in F_{t+1} at the position a^* is the consumption chosen at t. FOC1 at $t+1$ then reads $F_{t+1}(a^*, h^*) = F_t(a^*, h^*) = 1$ by hypothesis. The two Euler terms are equal as functions of (a, h), so their horizon derivatives at the fixed position a^* are also equal, and FOC2 reads $F'_{t+1}(h^*) = F'_t(h^*) = 0$. The pair (a^*, h^*) thus satisfies at $t+1$ the same first-order conditions that selected it at t, so the position and the horizon are unchanged and no trade occurs. ■

The proposition places the horizon inside the standard view of trading. In frictionless models trade requires new information (Milgrom and Stokey, 1982), and in stationary economies equilibrium trade is limited (Judd et al., 2003). The investment horizon might seem to imply the opposite—that the investor plans to hold the position for a fixed period and then rebalance. The proposition shows there is no conflict: absent news, the same position a^* and the same horizon h^* satisfy the first-order conditions again next period, so trades respond to news, not to the passage of time. The result also means the horizon choice is time consistent: re-optimizing every period does not undo it, so no commitment is needed. Turnover tied to news is likewise consistent with the infrequent trading of households (Duffie and Sun, 1990; Abel et al., 2007).

Remark 4 *(Variance curvature at the optimal horizon).* The Euler term has a local maximum at the optimal horizon[24]: $F(h^*) = 1$, $F'(h^*) = 0$, and $F(h) \leq 1$ for nearby $h < h^*$, so $F''(h^*) \leq 0$ and hence $V''_{h^*} \leq 0$.[25] The weak inequality allows two cases, and the model does not choose between them: either curvature is strictly negative ($V''_{h^*} < 0$), so the Euler term has a peak at h^*; or variance grows at the random-walk rate from h^* on, so F stays at one and h^* starts a plateau. Figure 4 draws both cases in a local window around h^*: the cumulative-variance path approaches the random-walk ray from below, returns to it tangentially at h^*, then either falls back below it (the peak) or rides it (the plateau).

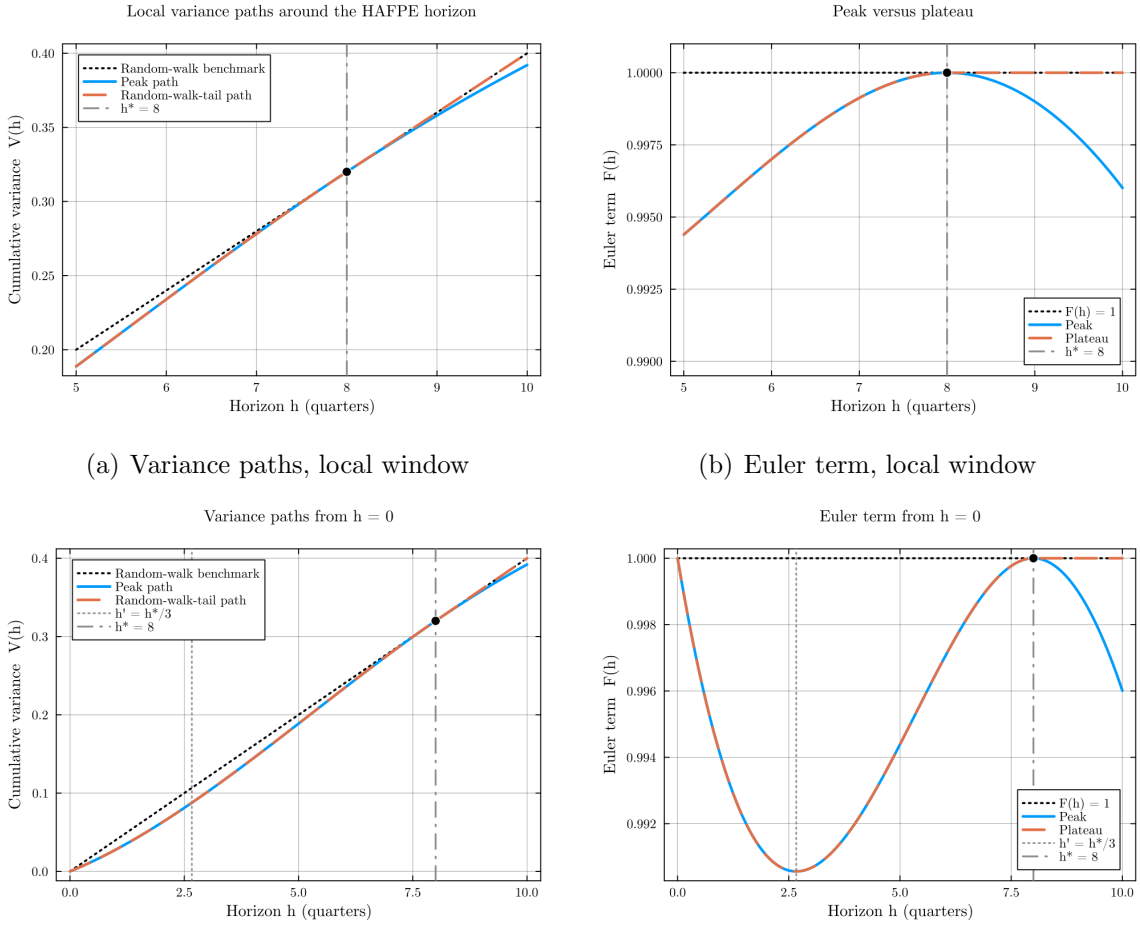
The preceding analysis characterizes the optimal investment horizon h^* but is silent about the size of the associated risky-asset position. I next show that risky-asset demand peaks at h^*: it rises before h^* and weakly declines thereafter. To analyze this relationship, I define $a^*(h)$ as the risky-asset demand that satisfies the pricing condition at horizon h. This formulation links the theory to observed portfolio positions.[26]

[24] If $F(h') > 1$ for some $h' < h^*$, take the two-point segment $\{h', h^*\}$; its gap $F(h^*) - F(h') = 1 - F(h')$ is negative, so the investor would shift weight toward h' — contradicting the optimality of h^*.

[25] Differentiating $F(h) = \exp(\mu h + \frac{1}{2}V_h)$ twice gives $F'' = F\left[(\mu + \frac{1}{2}V'_h)^2 + \frac{1}{2}V''_h\right]$. At h^*, $F(h^*) = 1$ and $F'(h^*) = 0$ imply $\mu + \frac{1}{2}V'_{h^*} = 0$, so the squared term drops out and $F''(h^*) = \frac{1}{2}V''_{h^*}$.

[26] Section 6 examines this relationship cross-sectionally: the portfolio weights that price each asset class

Figure 4: Tangency Geometry of the Optimal Horizon



(a) Variance paths, local window

(b) Euler term, local window

(c) Variance paths from $h = 0$

(d) Euler term from $h = 0$

Notes: The figure shows the tangency geometry of the optimal horizon ($h^* = 8$). The top row is a local window around h^*. Panel (a) shows the cumulative-variance path near the optimal horizon: before h^* it lies below the random-walk ray $\sigma^2 h$ and returns to it tangentially at h^*; beyond h^* the peak path falls back below the ray while the random-walk-tail path rides it. Panel (b) shows the corresponding Euler term: it rises toward one from below, reaches one exactly at h^*, and then either falls back (the peak) or stays at one (the plateau). The bottom row plots the same paths from $h = 0$. Panel (c): the variance path starts at zero on the ray, falls below it, and returns to it only at h^*; the ray is the RW line from the origin, tangent to the path at h^*. Panel (d): the Euler term starts at one, falls to a single minimum at $h' = h^*/3$, where the path's slope equals the ray's, and rises back to one at h^*—the shape in Equation 31.

The intuition is straightforward. Holding demand fixed, a longer horizon may raise or lower the Euler term in FOC1.[27] If a longer horizon raises the Euler term, the investor must increase risky-asset demand to restore FOC1; if it lowers the Euler term, demand must contract.

Proposition 4 *(Asset demand at the optimal horizon).* Consider a long position and let $a^*(h)$ denote the interior risky-asset demand that satisfies FOC1 at horizon h. Then risky-asset demand rises before h^* and weakly declines afterward: the maximizer is unique when the horizon derivative F' is strictly negative beyond h^*, and h^* is the smallest maximizer under the random-walk plateau of Remark 4.

The proof appears in the appendix to this section.

Remark 5 *(Shape of the Euler term).* The HAFPE conditions determine the global shape of the Euler term below the optimal horizon. Assume the asset position is $a^* \equiv a^*(h^*)$ and view $F(a^*, h)$ as a function over evaluation horizons $h \in [0, h^*]$. Both endpoints equal one: at $h = 0$, $M_{t,t} = R_{t,t} = 1$ by definition, so $F(a^*, 0) = 1$, while at h^* the HAFPE gives $F(a^*, h^*) = 1$. Between the endpoints the curve cannot exceed one: if $F(a^*, h) > 1$ at some $h < h^*$, the gap $F(a^*, h^*) - F(a^*, h) = 1 - F(a^*, h)$ is negative, so FOC2 favors shifting weight toward h, contradicting the optimality of h^*.

Remark 5 shows that $F(a^*, h) \leq 1$ on $(0, h^*)$ but not how many minima the curve has. The next proposition shows there is exactly one, at every step of the search. The idea is best seen by following the segmented algorithm step by step (Figure 6).

The moving random-walk benchmark. Coming back to the segmented algorithm, the investor does not know the random-walk benchmark in advance. The algorithm is set in discrete time, but from here on I work with its continuous-time version from Proposition 2 and keep the name. At each step h the position adjusts so that FOC1 holds at the current horizon, $F(a_h, h) = 1$ with $a_h \equiv a^*(h)$. In the log-normal framework, with the asset holding now included as an argument,[28] the log Euler term is the drift plus half the variance,

$$\log F(a, h) = \mu(a)\, h + \tfrac{1}{2} V_h(a), \tag{29}$$

vary systematically with the horizon, with low book-to-market assets appearing at short horizons and high book-to-market assets at long ones.

[27] In the segmented discrete formulation, FOC1 is $1 = (1-k)\mathbb{E}_t[M_{t,t+h}(a)R_{t,t+h}] + k\mathbb{E}_t[M_{t,t+h+\Delta h}(a)R_{t,t+h+\Delta h}]$, with $k \in [0,1]$. When the second Euler term exceeds the first, increasing k raises the right-hand side.

[28] The position enters through the SDF. The return $R_{t,t+h}$ is the same for every a: the investor takes prices as given. The SDF $M_{t,t+h} = \beta^h U'(c_{t+h})/U'(c_t)$ changes with a: a larger position lowers c_t and raises c_{t+h}, so $U'(c_{t+h})$ falls, $M_{t,t+h}$ falls in every state, and $\mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$ falls. Hence $\mu(a)$ depends on a through $\mathbb{E}_t \ln M_{t,t+h}$, and $V_h(a)$ through $\mathrm{Var}_t(\ln M_{t,t+h})$ and $\mathrm{Cov}_t(\ln M_{t,t+h}, \ln R_{t,t+h})$.

where $\mu(a)$ is the per-period drift of $\ln(M_{t,t+h}R_{t,t+h})$ and $V_h(a) \equiv \mathrm{Var}_t \ln(M_{t,t+h}R_{t,t+h})$ is its variance.

At step h of the segmented algorithm, FOC1 implies $\mu(a_h)\,h + \frac{1}{2}V_h(a_h) = 0$. We can construct a random-walk benchmark that matches the distribution at the horizon-h slice. The RW benchmark has a linear variance path. To have variance $V_h(a_h)$ at step h, its variance rate s_h must satisfy $s_h\,h = V_h(a_h)$. The benchmark has the same drift $\mu(a_h)$ and the constant variance rate s_h. Call it the *RW line*.

The log Euler term of the RW line is $\log F^{\mathrm{RW}}(a_h, u) = \mu(a_h)\,u + \frac{1}{2}s_h\,u$ (I write u for the horizon at which the curve is evaluated and keep h for the current step). To see how far the Euler term is from the benchmark at each horizon u, look at the log-Euler gap

$$\begin{aligned}\log F(a_h, u) - \log F^{\mathrm{RW}}(a_h, u) &= \big[\mu(a_h)\,u + \tfrac{1}{2}V_u(a_h)\big] - \big[\mu(a_h)\,u + \tfrac{1}{2}s_h\,u\big] \\ &= \tfrac{1}{2}\big[V_u(a_h) - s_h\,u\big].\end{aligned} \tag{30}$$

FOC1 gives $\mu(a_h) = -\frac{1}{2}s_h$,[29] so $\log F^{\mathrm{RW}}(a_h, u) = 0$ and the gap equals $\log F(a_h, u)$; hence Equation 30 can be seen as the log Euler term $\log F(a_h, u) = \frac{1}{2}\big[V_u(a_h) - s_h\,u\big]$, which allows us to see the log Euler term as the difference between the variance of the investor's Euler term and the variance of the possible RW benchmark.

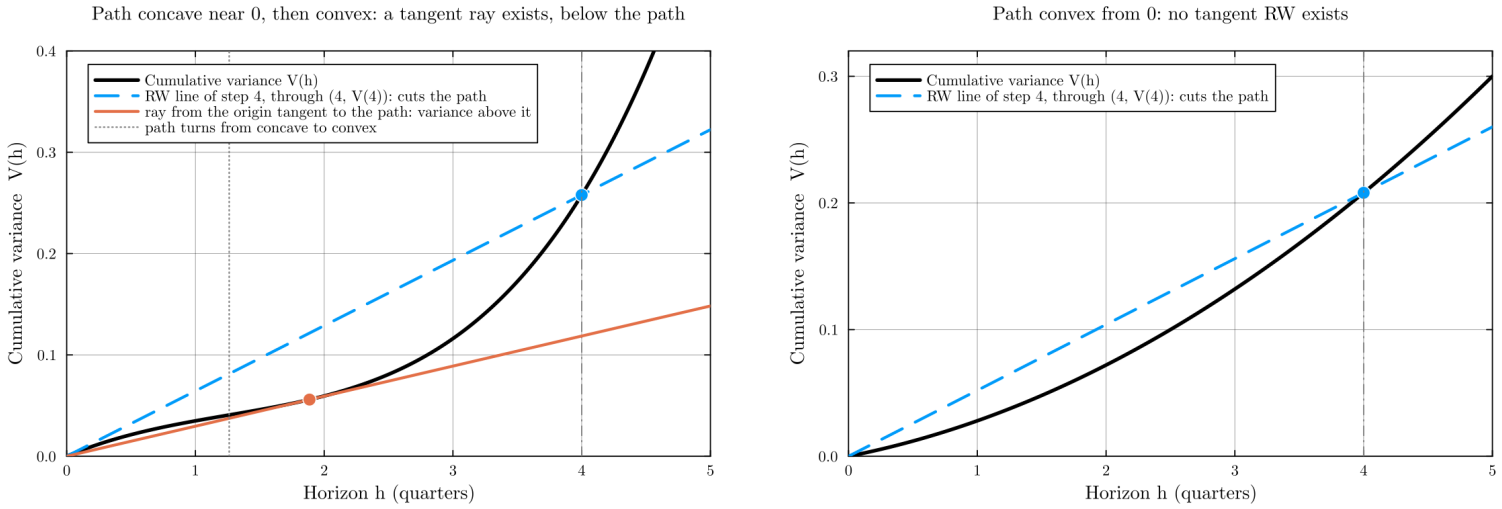
Differentiating this form of F with respect to u at the current step gives FOC2 of the HAFPE Equation 27, $\frac{1}{2}\big[V_h'(a_h) - s_h\big]$. Alternatively, Equation 30 is the gap between the Euler term and the RW benchmark, so FOC2 shows whether moving the horizon can change this gap. The HAFPE holds when the maximum possible RW benchmark — the steepest RW line — is reached.[30]

Consider the example in Figure 5, where the investor has a candidate horizon $h = 4$. The cumulative variance V_u, drawn as a function of the horizon, is the *variance path*; the RW line is the ray from the origin through $(4, V_4)$. The path crosses the line at $h = 4$: the variance lies above the line just beyond 4, so the path is steeper at the crossing, $V_4' > s_4$, and FOC2 is positive; the Euler term rises through one at $h = 4$. The search does not stop at $h = 4$: the investor lengthens the horizon to steepen the RW line. So the search for the optimal horizon can be seen as reaching for the steepest possible RW benchmark.

[29] At $u = h$, $s_h\,h = V_h(a_h)$, so $\mu(a_h)\,h + \frac{1}{2}s_h\,h = \log F(a_h, h) = 0$ by FOC1; dividing by h gives $\mu(a_h) = -\frac{1}{2}s_h$.

[30] Each step of the search is an investment alternative set: a position and a horizon (a_h, h) priced by FOC1. The investor moves to a longer alternative whenever that raises expected utility, FOC2 > 0, and every such move steepens the RW line. So the two rankings coincide: among all alternatives, the one with the steepest RW line gives the investor the highest utility, and the search stops there.

Figure 5: Step $h = 4$: No Tangent RW With the Variance Below It



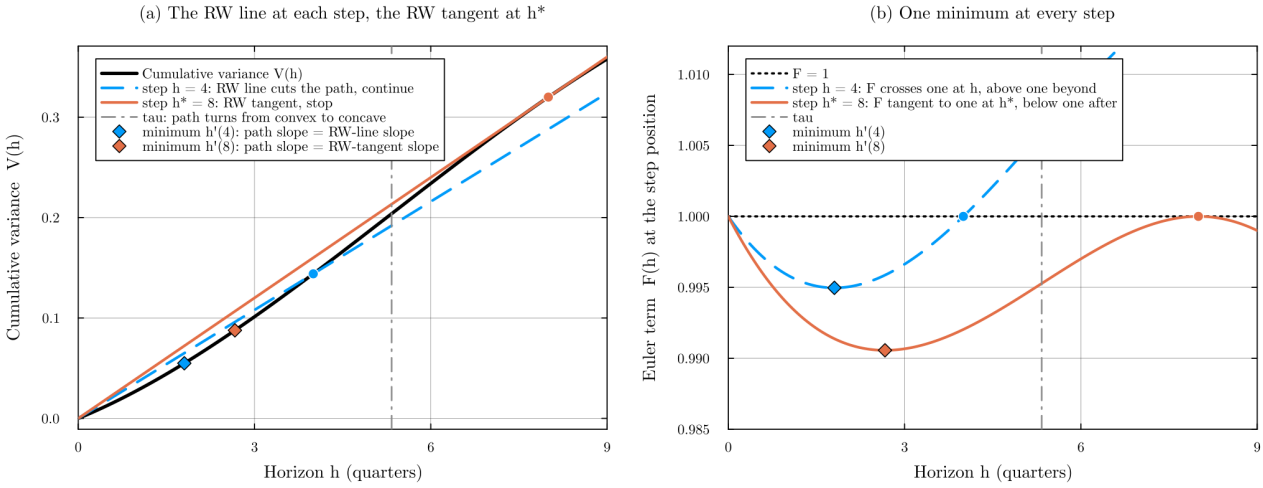
(a) (b)

Notes: Step $h = 4$ of the segmented search. All straight lines are rays from the origin, so all are RW lines. In both panels the RW line of the step goes through $(4, V_4)$ and cuts the path: beyond 4 the variance is above it. Panel (a): a path concave near zero and convex after. A ray from the origin is tangent to it, but the variance lies above that ray everywhere. The tangent ray touches the path at $u_0 < 4$, a horizon the search has already passed; it is the RW line of no step, and the Euler term of the step is below one there, since the path lies below the RW line of the step on all of $(0, 4)$. Such a path can arise only if its shape changes with the position: for a fixed path the search would not have started, because near zero a concave path is flatter than its own RW line. Panel (b): a path convex from zero. No ray from the origin is tangent to a convex path that starts at the origin, so no tangent RW exists. In both cases FOC2 is positive at $h = 4$ and the investor lengthens the horizon.

The search stops at the RW tangent. The RW line cannot steepen forever. As long as the path crosses the RW line from below, the horizon lengthens; the search stops when the line is tangent to the path—the slope of the RW line equals the slope of the path, $s_{h^*} = V'_{h^*}$, so $F_u(a^*, h^*) = 0$: the *RW tangent*. It is the steepest RW benchmark ray from the origin the path can touch: the investor reaches the highest slope the variance path allows and stops there (panel (a), $h^* = 8$).

It is worth noting that a tangent can exist before the current step h. In panel (a) of Figure 5, a lower ray touches the path at $u_0 < 4$, a horizon the search has already passed. At this point the path is convex, so it lies above its tangent ray: against that benchmark, lengthening horizon raises F—the tangent does not stop the search. A tangent stops the search only when the path lies below the RW line, so that no horizon offers an Euler term above one. In panel (b), the path is convex from zero and no lower tangent exists. Only at the stop is the benchmark known: its slope is $\sigma^2_{h^*} = s_{h^*}$.

Figure 6: The Segmented Algorithm Step by Step



Notes: Same variance path as Figure 4 ($h^* = 8$): convex on $(0, \tau)$, concave on (τ, h^*), $\tau = 2h^*/3$; the peak case beyond h^*. Panel (a): at each step h the benchmark is the RW line, the straight line from the origin to (h, V_h). At $h = 4$ it cuts the path from below at h, so the search continues and the line steepens; at $h^* = 8$ it is the RW tangent and the search stops. Diamonds mark the minimum $h'(h)$ of the Euler term at each step, where the path's slope equals the line's, always in the convex part. Panel (b): the Euler-term curves at the same steps, $\log F(a_h, u) = \frac{1}{2}[V_u - s_h u]$. Each starts at one, falls to one minimum, and returns to one at its step; at $h = 4$ it crosses one and lies above it beyond, at $h^* = 8$ it is tangent to one and lies below it after.

Remark 6 *(Random-walk benchmark and the variance ratio).* In the previous subsection, when we defined the investment horizon, Remark 1 compared the condition for the existence of a long horizon with departures from the random walk and the variance ratio.

In this section, we define the optimal horizon, Remark 5 provides the shape of the Euler term, and the discussion above shows the optimal horizon as a result of the maximum random-walk slope and its connection to the variance of log Euler terms.

The variance path $V_u(a^*)$ is the variance of $\ln(M_{t,t+u}R_{t,t+u})$ at horizon u, so $V_u(a^*)/u$ is the average variance rate up to u. The RW line through (h^*, V_{h^*}) has slope $s_{h^*} = V_{h^*}(a^*)/h^*$, the average rate up to the optimal horizon. The variance ratio of Subsection 3.5 (Lo and MacKinlay, 1988), which I denote $\mathrm{VR}(u)$, also compares the variance with a random-walk variance, yet since the optimal horizon is unknown, it cannot use the slope s_{h^*} as the random-walk rate; it takes instead the variance of the first period, $V \equiv V_1(a^*)$.

$$\mathrm{VR}(u) \equiv \frac{V_u(a^*)}{u\,V}, \qquad \mathrm{VR}(1) = 1.$$

The two random-walk rates are linked through the variance ratio at the optimal horizon, $s_{h^*} = \frac{V_{h^*}(a^*)}{h^*} = V\left[\frac{V_{h^*}(a^*)}{h^* V}\right] = V\,\mathrm{VR}(h^*)$, so Equation 30 at the optimal horizon is

$$\log F(a^*, u) = \tfrac{1}{2}\big[V_u(a^*) - s_{h^*}u\big] = \tfrac{1}{2}\,u\,V\big[\mathrm{VR}(u) - \mathrm{VR}(h^*)\big].$$

It would allow us to use the Lo–MacKinlay variance ratio at different h and see how it relates to the optimal horizon.

The log Euler term at horizon u is the variance ratio at u minus the variance ratio at h^*, times $\frac{1}{2}uV > 0$; the shape says two things. First, $F(a^*, u) < 1$ for $0 < u < h^*$ means $\mathrm{VR}(u) < \mathrm{VR}(h^*)$: the variance ratio is largest at h^*, and it is above one there because $\mathrm{VR}(1) = 1$ and $1 < h^*$. Second, FOC2, $F_u(a^*, h^*) = 0$, means $\mathrm{VR}'(h^*) = 0$: the slope of the variance ratio is zero at h^*.[31] So the optimal horizon is the horizon where the variance ratio is largest. This is the earlier statement in new words: the steepest RW line is the largest variance ratio.

Literature finds a variance ratio with this shape. Poterba and Summers (1988) report positive serial correlation in index returns over horizons of less than a year and negative serial correlation over horizons of three to eight years,[32] so the variance ratio of returns rises above one and later falls below one; it is largest at about two years.[33] Serial correlation and the variance ratio are connected (Equation 11). Positive serial correlation means the ratio rises. Negative means it falls. So the ratio of returns rises up to about one year and

[31] Differentiating, $F_u(a^*, u)/F(a^*, u) = \frac{1}{2}V\big[\mathrm{VR}(u) - \mathrm{VR}(h^*) + u\,\mathrm{VR}'(u)\big]$, which at $u = h^*$ equals $\frac{1}{2}Vh^*\,\mathrm{VR}'(h^*)$.

[32] Poterba and Summers (1988, Table 2)

[33] Rescaling the same ratios to a one-month base, the value-weighted ratio is 1.28 at 12 months, 1.31 at 24, 1.24 at 36, and 1.13 at 48 months; the equal-weighted ratio is 1.25 at 12 and 24 months and 1.14 at 36. The grid has no points between 12 and 24 months.

falls after two years. Its maximum lies between. Moskowitz et al. (2012) find the same rise-then-fall pattern in 58 liquid futures and forward instruments: returns persist for one to twelve months and partly reverse over longer horizons. Fama and French (1988b) find the reversal strongest for three- to five-year returns and weaker beyond, so at long horizons the ratio stops falling. In the model, the horizon where the ratio is largest is the optimal horizon. The HAFPE predicts this pattern: the ratio rises above one, is largest at h^*, and does not rise after h^* (Remark 4). The HAFPE does not say whether the ratio falls below one later. These studies measure returns. The theory is about the pricing product $M_{t,t+u}R_{t,t+u}$. The two have the same peak when the pricing kernel follows a random walk, because then the kernel adds only a linear term to the variance. So I use the studies as motivation only; **??** tests the shape directly on the Euler term.

Remark 5 and the HAFPE give the shape of the Euler term at the optimal position:

$$F(a^*, 0) = 1, \qquad F(a^*, h) < 1 \ \text{ for } 0 < h < h^*, \qquad F(a^*, h^*) = 1. \tag{31}$$

The curve equals one at $h = 0$ by definition and at the plan horizon by the HAFPE, lies below one in between, and by FOC2 touches one at h^* rather than crossing it. This shape is testable. An estimated Euler-term curve should be below one at every horizon before the plan horizon and return to one there. The appendix to this section presents empirical curves of the unconditional Euler term for the market portfolio. The curves follow this pattern except at very low levels of risk aversion.

Remark 7 *(Instantaneous unpredictability).* Consider the stock market as an econometrician observes it, either in the aggregate or at the level of an individual risky asset: the marginal investor holds a^* with a positive horizon h^*. Then, for any short holding period $0 < \Delta h < h^*$,

$$F(a^*, \Delta h) - F(a^*, 0) < 0,$$

because $F(a^*, 0) = 1$ by definition and $F(a^*, h) < 1$ for $h < h^*$ (Equation 31). An econometrician who evaluates the investment over the short period Δh therefore finds that its Euler term is below one: in marginal-utility units, a dollar invested for Δh returns less than a dollar. At the position that prices the asset, holding for Δh is worse than not holding at all—no one with a short horizon would take this position, and the investor who does take it enters only because they plan to stay. In the log-normal notation, the inequality reads $\mu\Delta h + \frac{1}{2}V_{\Delta h} < 0$—the drift of the pricing product is too low for its variance—and as $\Delta h \to 0$, $\mu \leq -\frac{1}{2}\sigma_0^2$ with $\sigma_0^2 \equiv V'_{0+}$ the instantaneous variance rate. The result echoes the "instantaneous unpredictability" that Shiller (2014) highlights, following

Sims (1984): investors "cannot expect to see solid returns over short time intervals." Here we see it from another angle: the drift μ is too low to compensate for the variance $V_{\Delta h}$ of the short-horizon pricing product.

Earlier in this section, we assumed a long position in the asset, so that it is optimal for the investor to hold the asset rather than to short it. This is a reasonable assumption, especially since we will later analyze a representative agent holding the market. Yet we do not provide any additional condition on the long position; a short position can also be optimal if HAFPE holds. Below we summarize the results for the case in which the investor holds a short position.

Corollary 1 (Short positions)**.** *Consider a short position, $a_t < 0$, and let (a^*, h^*) denote the optimum of the segmented search. Then:*

(i) The HAFPE is unchanged: $F(h^*) = 1$ and $F'(h^*) = 0$.

(ii) The Euler term lies above one before the optimal horizon: $F(a^*, h) \geq 1$ for $0 < h < h^*$, with $F(a^*, 0) = F(a^*, h^*) = 1$, so the curve touches one from above at h^*; the mirror image of **??** *gives a single peak on $(0, h^*)$.*

(iii) The size of the short position, $|a^*(h)|$, is maximal at h^*.

(iv) $F''(h^*) \geq 0$, so $V''_{h^*} \geq 0$. Relative to the random-walk benchmark of Remark 4, the geometry in Figure 4 is reversed: cumulative variance lies above the random-walk line $\sigma^2_{RW} h$ before h^* and touches it tangentially at h^*; beyond h^* it either moves above the ray again (a strict minimum of F) or rides it (a plateau).

The proof reapats the long-position case: for $a_t < 0$, FOC2 $= a_t G$ is positive when $G < 0$, so every inequality of the long-position arguments reverses and the proofs of Proposition 2, Proposition 4 (whose Step 3 does not use the sign of a), and **??** go through unchanged.

Risk aversion and the optimal horizon. Suppose we know only that returns are persistent, so that an interior optimal horizon exists, and that utility is CRRA, $U'(c) = c^{-\gamma}$. At the optimal horizon the HAFPE holds and the RW line exists. Under CRRA, with $\Delta c_u \equiv \ln c_{t+u} - \ln c_t$ the log consumption growth over the u periods and $r_u \equiv \ln R_{t,t+u}$ the log return, the variance rate at horizon u splits into three parts,

$$\frac{V_u}{u} = \gamma^2 v_c(u) + v_r(u) - 2\gamma\, v_{cr}(u), \qquad v_c(u) \equiv \frac{\mathrm{Var}_t(\Delta c_u)}{u},\ v_r(u) \equiv \frac{\mathrm{Var}_t(r_u)}{u},\ v_{cr}(u) \equiv \frac{\mathrm{Cov}_t(\Delta c_u, r_u)}{u}, \tag{32}$$

the same three risk terms as in ψ of Equation 7, with the random-walk constants replaced by their values at horizon u. A change in γ therefore reshapes the variance path, through $\gamma^2 v_c$ and $-2\gamma v_{cr}$; the tangency point moves, and with it the optimal horizon. To track the move we also need the mean, because the position responds through it. Write g for the mean of consumption growth per period, so that $\mathbb{E}_t \Delta c_h = g\,h$; the log Euler term written out is

$$\ln F(a,h) = h \ln \beta - \gamma g\, h + \mathbb{E}_t r_h + \tfrac{1}{2} h \big[\gamma^2 v_c(h) + v_r(h) - 2\gamma\, v_{cr}(h)\big]. \tag{33}$$

Proposition 5 *(Risk aversion and the optimal horizon).* Let h^* be an interior optimal horizon with the log Euler term Equation 33, and let an asset position raise mean consumption growth: $g = g(a)$, $g'(a) > 0$. Suppose, as in the data, that mean consumption growth dominates the risk rates near h^* and that the covariance of consumption with returns builds up with the horizon faster than the consumption variance at h^*.[34] Then a more risk-averse investor holds a smaller position and a shorter horizon.

The proof appears in the appendix to this section, in four short steps. A rise in γ lowers the Euler term at every horizon, more at longer ones. Since $F(h^*) < 1$, restoring FOC1 requires the investor to sell assets until $F(h^*) = 1$ again; the cut also acts in proportion to the horizon. The common per-period effect therefore cancels. The remaining gap depends on how the risk rates change with the horizon. When the covariance rate grows faster than γ times the variance rate, the Euler term at the longer end falls below one, so FOC2 is negative and the horizon shortens.

Without the covariance condition, the general direction is

$$\operatorname{sign} \frac{dh^*}{d\gamma} = \operatorname{sign} \big[\gamma\, v_c'(h^*) - v_{cr}'(h^*)\big]. \tag{34}$$

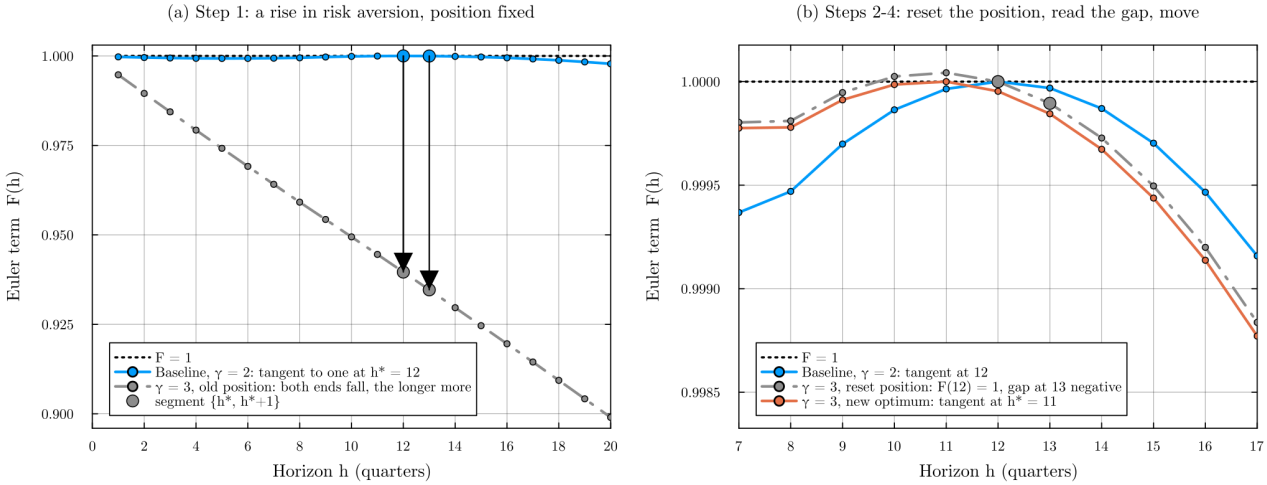
Return persistence alone does not decide it; consumption does. If consumption growth is a random walk, $v_c'(h^*) = 0$ and the covariance alone decides: a higher γ shortens h^*.

Figure 7 illustrates the proposition with data-like quarterly moments for a rise of γ from 2 to 3. Panel (a) shows the shock at the old position: at $\gamma = 2$ the Euler term is tangent to one at $h^* = 12$; the rise in γ lowers the log Euler term at every horizon, by about 6 percent on the segment, more at the longer end. Panel (b) shows the response: the cut in the position raises the log Euler term in proportion to the horizon and brings $F(12)$ back to one; the curve now crosses one from above, so the gap at 13 is negative, the investor shifts weight to the shorter end, and the new optimum is tangent at $h^* = 11$. In these moments the covariance rate rises at h^* faster than γv_c, so Equation 34 gives the one-quarter shortening

[34] Precisely: $g > \gamma\, v_c(h) - v_{cr}(h)$ for h near h^*, and $v_{cr}'(h^*) > \gamma\, v_c'(h^*)$.

seen in the figure.

Figure 7: A Rise in Risk Aversion on the Segment $\{h^*, h^*+1\}$



Notes: Quarterly moments from a moving-average model with data-like values: return standard deviation 8.5%, variance ratio 1.06 at 4 quarters, 1.14 at 8, 0.98 at 40; consumption growth mean $g = 0.5\%$, standard deviation 0.5%, first-order autocorrelation 0.2, and a slow response to return shocks, so that the correlation with returns rises from 0.21 at 1 quarter to 0.43 at 8. The per-period log Euler term at a FOC1 position is $q(h) = \frac{1}{2}[\gamma^2 v_c(h) + v_r(h) - 2\gamma v_{cr}(h)]$ and $\log F(a_h, u) = u[q(u) - q(h)]$. Panel (a), Step 1: at $\gamma = 2$ the Euler term is tangent to one at $h^* = 12$; raising γ to 3 at the old position lowers $\log F(h)$ by $hA(h)$, $A(h) = -g + \frac{1}{2}(3^2 - 2^2)v_c(h) - v_{cr}(h) \approx -0.005$, so both ends of the segment fall, the longer one more (arrows). Panel (b), Steps 2–4: the reset raises $\log F(h)$ by $-hA(12)$, so $F(12) = 1$ again; the curve now crosses one from above, the gap at 13 is negative, and the new optimum is tangent at $h^* = 11$.

Summary. This section developed the theory of the investment horizon: investors price assets over the horizon they plan to hold them, and this plan enters the standard model through the horizon function. The analysis delivers four main results. First, if the econometrician imposes a horizon different from the investors'—in particular, the one-period horizon of the data—the resulting estimates are biased: matching the data then requires implausibly high risk aversion or discounting, producing the well-known equity-premium and risk-free-rate puzzles. Second, the random walk is the benchmark against which horizons matter: under random-walk consumption and returns, every horizon prices identically, and the investment horizon becomes relevant only when variance departs from linear growth. Third, when the investor chooses the horizon optimally, the Euler term meets its benchmark of one smoothly: $F(h^*) = 1$ and $F'(h^*) = 0$. This pair is the Horizon-Adjusted Fundamental Pricing Equation. Fourth, risky-asset demand rises before the optimal horizon and weakly declines after it, so the position peaks at h^*. Finally, the section gives two testable implications of a long investment horizon: the horizon Euler terms, the expected pricing product

at each horizon, lie below one before the optimal horizon and return to one there; and the variance ratio is largest at the optimal horizon. So the theory predicts the variance-ratio pattern documented in empirical work.

The appendix below collects the proofs of this section's propositions. In the next section, Section 4, I present a methodology for extracting the investment horizon from asset-pricing data.

Appendix to this section

Proof of Proposition 4.

Step 1: Asset demand at a given horizon. The decision date t is fixed and the position is long, $a_t > 0$. To simplify the notation, we denote the investor's date-t risky-asset position by a and its re-optimized value at horizon h by $a^*(h)$. Time subscripts are suppressed unless they are necessary to indicate timing, as in c_t, c_{t+h}, $R_{t,t+h}$, and the conditional expectation operator $\mathbb{E}_t$.

Recall the horizon Euler term,

$$F(a,h) \equiv \mathbb{E}_t[M_{t,t+h}(a)R_{t,t+h}], \qquad M_{t,t+h}(a) \equiv \beta^h \frac{U'(c_{t+h})}{U'(c_t)}, \tag{35}$$

where the budget constraints

$$c_t = w_t - ap_t, \qquad c_{t+h} = w_{t+h} + ap_{t+h}, \qquad p_{t+h} = p_t R_{t,t+h},$$

make consumption at both dates depend on a. For a given horizon h, re-optimized asset demand solves FOC1,

$$F(a^*(h),h) = 1. \tag{36}$$

We use the notation

$$F_a(a,h) \equiv \frac{\partial F(a,h)}{\partial a}, \qquad F_h(a,h) \equiv \frac{\partial F(a,h)}{\partial h}.$$

In this notation, the HAFPE horizon derivative of the previous subsection is

$$F'(h) \equiv F_h(a^*(h),h), \tag{37}$$

where the derivative with respect to h is taken while holding the current optimized asset position fixed.

Step 2: What we need—the demand response. We seek the derivative $da^*(h)/dh$,

which measures how a change in the investor's horizon h affects optimal asset holdings. Because Equation 36 holds identically in h, we differentiate it with respect to h:

$$\begin{aligned} 0 &= \frac{d}{dh} F(a^*(h), h) \\ &= F_a(a^*(h), h) \frac{da^*(h)}{dh} + F_h(a^*(h), h). \end{aligned} \tag{38}$$

Therefore,

$$\frac{da^*(h)}{dh} = -\frac{F_h(a^*(h), h)}{F_a(a^*(h), h)}. \tag{39}$$

To sign the demand response we need only the sign of F_a, computed in Step 3; the sign of F_h is the HAFPE horizon derivative Equation 37 of the preceding subsections, restated in Step 4.

Step 3: The denominator—asset demand and the Euler term. We now compute $F_a(a, h)$. Define

$$X_{t,h}(a) \equiv M_{t,t+h}(a) R_{t,t+h} > 0,$$

so that $F(a, h) = \mathbb{E}_t[X_{t,h}(a)]$. Holding h fixed and treating the individual investor as a price taker, prices and returns do not depend on the individual's choice of a. We can therefore pass the derivative through the conditional expectation and, because $X_{t,h}(a) > 0$, express the derivative in log-derivative form:

$$F_a(a, h) = \mathbb{E}_t\left[\frac{\partial X_{t,h}(a)}{\partial a}\right] = \mathbb{E}_t\left[X_{t,h}(a) \frac{\partial}{\partial a} \log X_{t,h}(a)\right]. \tag{40}$$

Since

$$\log X_{t,h}(a) = h \log \beta + \log U'(c_{t+h}) - \log U'(c_t) + \log R_{t,t+h},$$

only the two marginal-utility terms depend on a. Differentiating them and substituting $\partial c_t / \partial a = -p_t$ and $\partial c_{t+h} / \partial a = p_{t+h} = p_t R_{t,t+h}$ from the budget constraints,

$$\frac{\partial}{\partial a} \log X_{t,h}(a) = \frac{U''(c_{t+h})}{U'(c_{t+h})} \frac{\partial c_{t+h}}{\partial a} - \frac{U''(c_t)}{U'(c_t)} \frac{\partial c_t}{\partial a} = p_t \left(R_{t,t+h} \frac{U''(c_{t+h})}{U'(c_{t+h})} + \frac{U''(c_t)}{U'(c_t)} \right). \tag{41}$$

Combining Equation 40 with Equation 41, we obtain

$$F_a(a, h) = p_t \mathbb{E}_t\left[X_{t,h}(a) \left(\frac{U''(c_t)}{U'(c_t)} + R_{t,t+h} \frac{U''(c_{t+h})}{U'(c_{t+h})} \right)\right]. \tag{42}$$

Under the standard assumptions $U' > 0$, $U'' < 0$, $p_t > 0$, and $R_{t,t+h} > 0$ state by state, each

term inside the parentheses is negative. Since $X_{t,h}(a) > 0$, it follows that

$$F_a(a, h) < 0. \tag{43}$$

Economically, this is the ordinary downward-sloping demand curve: increasing the risky-asset position shifts consumption from date t toward date $t + h$, and with concave utility this lowers the Euler term, because the marginal value of the additional future payoff falls relative to its marginal utility cost at date t. The derivation uses only $U' > 0$, $U'' < 0$, and positive prices and returns—never the sign of a—so Equation 43 holds for short positions as well.

Step 4: The numerator—the horizon and the Euler term. Recall the setup: the investor invests at a given horizon h, and the position $a^*(h)$ adjusts so that FOC1 holds at h; but h itself need not be optimal, so FOC2 need not hold—below the optimum it is positive, as shown above. The FOC2 expression is exactly the slope of the Euler term in the horizon: the object F_h. More precisely: On the grid the index h is discrete, so no derivative can be taken with respect to it; the horizon moves continuously only through the weight, as $h+k$. The analog of F_h is therefore the derivative with respect to k, and at a fixed position it is exactly the gap, because the segment Euler term $(1-k)F(a,h) + kF(a,h+1)$ is linear in k. The sign of the gap is already established: below h^* the investor is still lengthening the horizon, so the gap is positive; at the interior optimum it vanishes (FOC2); beyond h^*, $F \leq 1$ (Remark 4) keeps it nonpositive, since the term held equals one by FOC1. In the notation of this subsection,

$$F_h(a^*(h), h) > 0 \ \text{ for } h < h^*, \qquad F_h(a^*(h^*), h^*) = 0, \qquad F_h(a^*(h), h) \leq 0 \ \text{ for } h > h^*. \tag{44}$$

The last inequality is weak by design: Remark 4 does not rule out the random-walk plateau, in which F stays at one beyond h^*.

Step 5: The answer. We can now determine the direction of the asset-demand response. From Equation 39 and Equation 43, the negative denominator leaves the numerator's sign intact:

$$\operatorname{sgn}\left(\frac{da^*(h)}{dh}\right) = \operatorname{sgn}(F_h(a^*(h), h)) \,. \tag{45}$$

Thus, the asset-demand response has the same sign as the horizon derivative of the Euler term. Combining Equation 45 with Equation 44,

$$\frac{da^*(h)}{dh} > 0 \quad \text{for } h < h^*, \qquad \frac{da^*(h)}{dh} = 0 \quad \text{at } h = h^*, \qquad \frac{da^*(h)}{dh} \leq 0 \quad \text{for } h > h^*. \tag{46}$$

Risky-asset demand therefore rises up to the optimal horizon h^*. Before h^*, lengthening the horizon raises the Euler term, so the investor must raise asset demand to restore FOC1. After h^*, lengthening no longer raises the Euler term, so demand no longer rises: it falls or, on the random-walk plateau of Remark 4, stays unchanged. ■

Proof of Proposition 5. Follow the two-point segment $\{h^*, h^*{+}1\}$ of the segmented algorithm. Before the change in γ both conditions hold: $F(a^*, h^*) = F(a^*, h^*{+}1) = 1$.

Step 1: The shock, position fixed. Hold a fixed and differentiate Equation 33 in γ. Only three terms contain γ: $-\gamma g h$, $\frac{1}{2} h \gamma^2 v_c(h)$, and $-h\gamma v_{cr}(h)$. Their derivatives are

$$\frac{\partial \ln F(a,h)}{\partial \gamma} = -g\,h + h\,\gamma\,v_c(h) - h\,v_{cr}(h) = h\,A(h), \qquad A(h) \equiv -g + \gamma\,v_c(h) - v_{cr}(h).$$

By hypothesis $g > \gamma v_c(h) - v_{cr}(h)$ near h^*, so $A(h) < 0$; with standard numbers $g \approx 0.5\%$ per quarter, while γv_c and v_{cr} are of order 10^{-4}. So a rise in γ lowers both Euler terms of the segment below one: $\ln F(a^*, h^*)$ falls by $h^* A(h^*)$ and $\ln F(a^*, h^*{+}1)$ by $(h^*{+}1)A(h^*{+}1)$. The longer one falls more. The difference between the two drops is

$$(h^*{+}1)A(h^*{+}1) - h^* A(h^*) = A(h^*{+}1) + h^*\big[A(h^*{+}1) - A(h^*)\big] < 0\,:$$

the first term is one more period of the shock, about $-g$; the second is h^* times the change in the risk rates over one period, of order 10^{-4}, too small to change the sign.

Step 2: FOC1 is restored. Hold the horizon at h^*. FOC1, $F(a, h^*) = 1$, must hold at any optimum. Before the shock it held, $\ln F(a^*, h^*) = 0$. Step 1 gives the change, not the level: for a rise $\delta\gamma$ in risk aversion, $\ln F(a^*, h^*)$ moves by $h^* A(h^*)\,\delta\gamma$. Take $\delta\gamma = 1$ as in Figure 7. After the shock, then, $\ln F(a^*, h^*) = 0 + h^* A(h^*) < 0$: FOC1 fails. Only the position can move it back. By hypothesis a affects the mean of consumption growth, $g = g(a)$ with $g'(a) > 0$ (a larger position raises expected consumption), but not its variance or its covariance with returns. Differentiate Equation 33 in a: the only term with a is $-\gamma g(a) h$, so

$$\frac{\partial \ln F(a,h)}{\partial a} = -\gamma g'(a)\,h < 0.$$

Hence $\ln F(a, h^*)$ falls when a rises and rises when a falls. After the shock $\ln F(a^*, h^*) < 0$, so to bring it back to zero the investor must *cut* the position: $\delta a < 0$. For a change δa, the log Euler term at horizon h moves by

$$\frac{\partial \ln F(a,h)}{\partial a}\,\delta a = -\gamma g'(a)\,\delta a \cdot h \equiv h\,\delta\mu, \qquad \delta\mu \equiv -\gamma g'(a)\,\delta a.$$

With $g'(a) > 0$ and $\delta a < 0$, $\delta\mu > 0$: the cut raises the drift of $\ln(M_{t,t+h}R_{t,t+h})$ by $\delta\mu$ per period, so it raises $\ln F$ by $h\,\delta\mu$ at every horizon, again more at the longer end. FOC1 at h^* fixes the size of the cut. The new level at h^* is the pre-shock level, plus the Step 1 drop, plus the Step 2 rise, and it must be zero:

$$0 + h^*A(h^*) + h^*\delta\mu = 0 \quad \Longrightarrow \quad \delta\mu = -A(h^*) > 0.$$

Step 3: The gap. Now compare the two ends of the segment, one step at a time. Before the shock both Euler terms are one: $\ln F(a^*, h^*) = \ln F(a^*, h^*{+}1) = 0$.

At h^*: Step 1 lowers $\ln F$ by $h^*A(h^*)$; Step 2 raises it by $h^*\delta\mu$ with $\delta\mu = -A(h^*)$. The new level is

$$\begin{aligned}\ln F(a, h^*) &= 0 + h^*A(h^*) + h^*\delta\mu \\ &= h^*A(h^*) - h^*A(h^*) = 0.\end{aligned}$$

The two moves cancel by construction; FOC1 holds again.

At $h^*{+}1$: Step 1 lowers $\ln F$ by $(h^*{+}1)A(h^*{+}1)$; Step 2 raises it by $(h^*{+}1)\delta\mu$, with the same $\delta\mu = -A(h^*)$. The new level is

$$\begin{aligned}\ln F(a, h^*{+}1) &= 0 + (h^*{+}1)A(h^*{+}1) + (h^*{+}1)\delta\mu \\ &= (h^*{+}1)A(h^*{+}1) - (h^*{+}1)A(h^*) \\ &= (h^*{+}1)\big[A(h^*{+}1) - A(h^*)\big].\end{aligned}$$

The moves do not cancel because the shock at the longer horizon, $A(h^*{+}1)$, differs from the shock at h^* that the position was set to offset. This is the gap.

Now substitute the definition $A(h) = -g + \gamma v_c(h) - v_{cr}(h)$ from Step 1:

$$\begin{aligned}A(h^*{+}1) - A(h^*) &= \big[-g + \gamma v_c(h^*{+}1) - v_{cr}(h^*{+}1)\big] - \big[-g + \gamma v_c(h^*) - v_{cr}(h^*)\big] \\ &= \gamma\big[v_c(h^*{+}1) - v_c(h^*)\big] - \big[v_{cr}(h^*{+}1) - v_{cr}(h^*)\big].\end{aligned}$$

The mean term $-g$ is the same at both horizons and drops out. So the part of the shock that is the same per period at both horizons is removed by the position; only the change of the risk rates across the segment is left. This is the sign of FOC2, $a\,G$.

Step 4: The horizon. If the covariance rate v_{cr} rises faster than γv_c across the segment, the gap is negative, the investor lowers k, the weight on the longer horizon, and h^* falls; in the opposite case h^* rises. As $\Delta h \to 0$ the gap per unit of Δh becomes $h^*\big[\gamma v_c'(h^*) - v_{cr}'(h^*)\big]$,

which gives Equation 34. By hypothesis $v'_{cr}(h^*) > \gamma\, v'_c(h^*)$, so the gap is negative and the horizon shortens; the position was cut in Step 2. ■

Sample Horizon Euler Terms

Figure 8 and Figure 9 plot the unconditional sample mean of the horizon Euler term for the stock market. The pricing kernel of a CRRA investor over h quarters is $M_{t,t+h} = \beta(h)\, c_{t,t+h}^{-\gamma}$, where $\beta(h)$ is the discount factor applied to a payoff h quarters ahead and $c_{t,t+h}$ is gross consumption growth over the same h quarters. The sample analogue of $F(h) = \mathbb{E}_t[M_{t,t+h}R_{t,t+h}]$ is

$$\widehat{F}_h^s(\gamma) = \frac{1}{T_h}\sum_{t=1}^{T_h} \beta(h)\, c_{t,t+h}^{-\gamma} R_{t,t+h}^s, \qquad h = 1, \ldots, 40 \text{ quarters.} \tag{47}$$

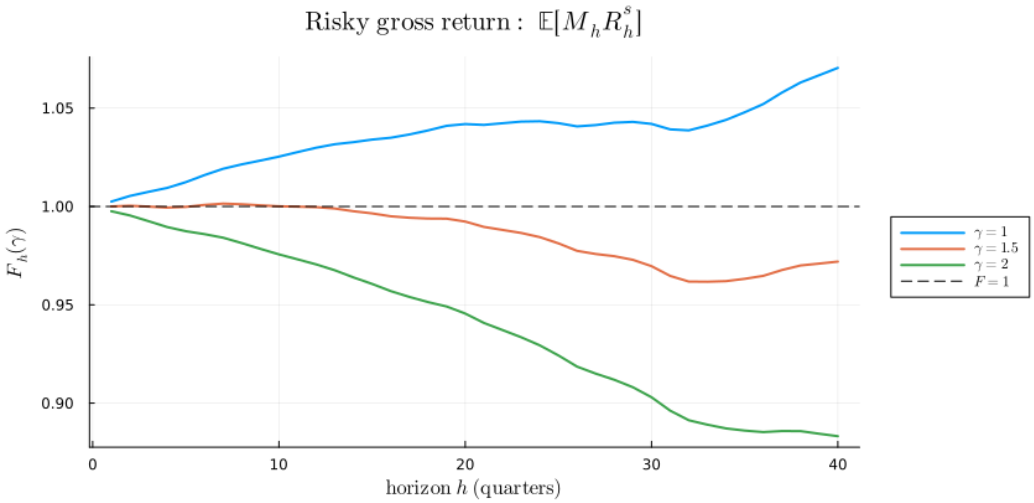
The data are quarterly, 1961Q3–2016Q4: $c_{t,t+h} = C_{t+h}/C_t$ is the growth of real personal consumption expenditure on nondurable goods and services, and $R_{t,t+h}^s$ is the real buy-and-hold return of the CRSP value-weighted market portfolio with dividends over the same h quarters; windows overlap, so $T_h = 222 - h$. The two figures differ only in the discount function $\beta(h)$: exponential, $\beta(h) = \beta^h$ with $\beta = 0.99$, in Figure 8; generalized hyperbolic (Loewenstein and Prelec, 1992), $\beta(h) = (1 + \alpha h)^{-\rho/\alpha}$ with $\rho = -\ln 0.99$, in Figure 9 ($\alpha \to 0$ returns β^h).[35]

The purpose of these figures is to check, on the market portfolio and for reasonable preference parameters, whether the horizon Euler term has the shape that Equation 31 predicts: below one before the optimal horizon, back to one at h^*, and not above one after. I ask three questions. (i) Is $F < 1$ at short horizons, the sign the theory requires before h^*? (ii) Does F return to one, and at what horizon? (iii) How do the answers depend on the discount function—exponential, or hyperbolic with curvature taken from experiments on dated money rather than from returns—and on how the expectation is formed: the sample mean, one realization of the data, or the mean under the estimated transition matrix P^o, which averages over the estimated law and can be conditioned on the initial state? The figures are descriptive. No parameter is estimated: at a single horizon β and γ are not

[35] Exponential discounting applies the same rate to every quarter: the discount between years two and three is the same as between years ten and eleven. Experimental evidence favors rates that fall with the distance to the payoff, so that a year of delay costs more between years two and three than between years ten and eleven (Loewenstein and Prelec, 1992; Laibson, 1997). The hyperbolic function has this property: its per-quarter rate is $\rho/(1 + \alpha h)$, equal to ρ at $h = 0$ and falling in h. For example, with $\beta = 0.99$ exponential discounting takes $1 - 0.99^4 = 3.9\%$ off a payoff for each year of delay, whichever year it is; the hyperbolic function with $\alpha = 0.1$ takes 2.0% for the year between years two and three and 0.8% for the year between years ten and eleven, and with $\alpha = 0.5$ it takes 0.7% and 0.2%.

separately identified ($\gamma = 1.8$ at $\beta = 0.99$ gives $F(1) = 1$, as does a higher γ at a higher β), so γ is set by the level condition given $\beta(h)$, and what is checked is the shape. Subsection 4.2 treats the parameters jointly, and **??** estimates the Euler term itself.

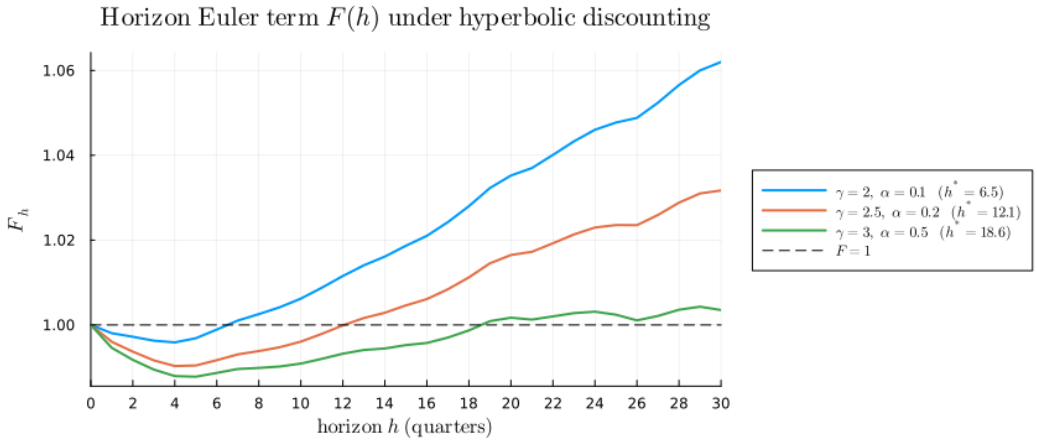
Figure 8: Sample Horizon Euler Term under Exponential Discounting



Notes: $\widehat{F}_h^s(\gamma)$ of Equation 47 with $\beta(h) = 0.99^h$ for $\gamma \in \{1, 1.5, 2\}$ and $h = 1, \ldots, 40$ quarters. The dashed line marks $F = 1$.

Figure 10 repeats the exercise on simulated paths from the objective transition matrix P^o. The left column uses exponential discounting, $\beta(h) = 0.99^h$. At $\gamma = 1.8$ the Euler term equals one at every horizon. The reason is simple. At $\gamma = 1.8$ the one-quarter Euler equation holds, $F(1) = 1$: a dollar invested for one quarter is worth a dollar. On these paths holding for h quarters is close to the one-quarter investment repeated h times, $F(h) \approx F(1)^h = 1$. No horizon is better than another, and the investor has no reason to choose one. This is the random-walk case of Subsection 3.4. Which way the other curves move is set by γ and β through $F(1)$: the same paths give a rising curve at $\gamma = 1$ and a falling one at $\gamma = 3$, and $F(h) \approx F(1)^h$ compounds up or down without turning back. The exponential panel therefore says one thing: under CRRA with $\beta(h) = 0.99^h$ and γ between the value that prices one quarter and 3, the data give $F < 1$ at every horizon up to forty quarters, the sign the theory requires before h^*. It does not locate h^*. A reader may object that it shows no horizon. It does not: $F < 1$ throughout is consistent with an h^* beyond the plotted range and with no horizon under these preferences, and the figure cannot separate the two. What it rules out is a short horizon under these preferences. The existence of h^* is a joint statement about preferences and data. The right column uses the same paths with hyperbolic discounting. Now the discount rate falls with the distance to the payoff, so $F(2) > F(1)$ even when

Figure 9: Sample Horizon Euler Term under Hyperbolic Discounting



Notes: $\widehat{F}^s_h(\gamma)$ of Equation 47 with $\beta(h) = (1+\alpha h)^{-\rho/\alpha}$, $\rho = -\ln 0.99$, for three pairs (γ, α) and $h = 0, \ldots, 30$ quarters; at $h = 0$ the term equals one by definition. h^* in the legend is the horizon at which the term crosses one from below (linear interpolation between quarters, on the grid $h = 1, \ldots, 40$). The largest value, $\alpha = 0.5$ per quarter, is 2.0 per year. For the same curvature parameter, denoted γ in their paper, Abdellaoui et al. (2010, Table 5, p. 859) report a median of 0.22 per year and an interquartile range up to 1.96 per year across individuals (gains), so $\alpha = 0.5$ lies just above their upper-quartile endpoint and $\alpha = 0.1$ lies within their interquartile range. Their estimates come from an experiment on choices between dated monetary outcomes, not from financial-market returns. The dashed line marks $F = 1$.

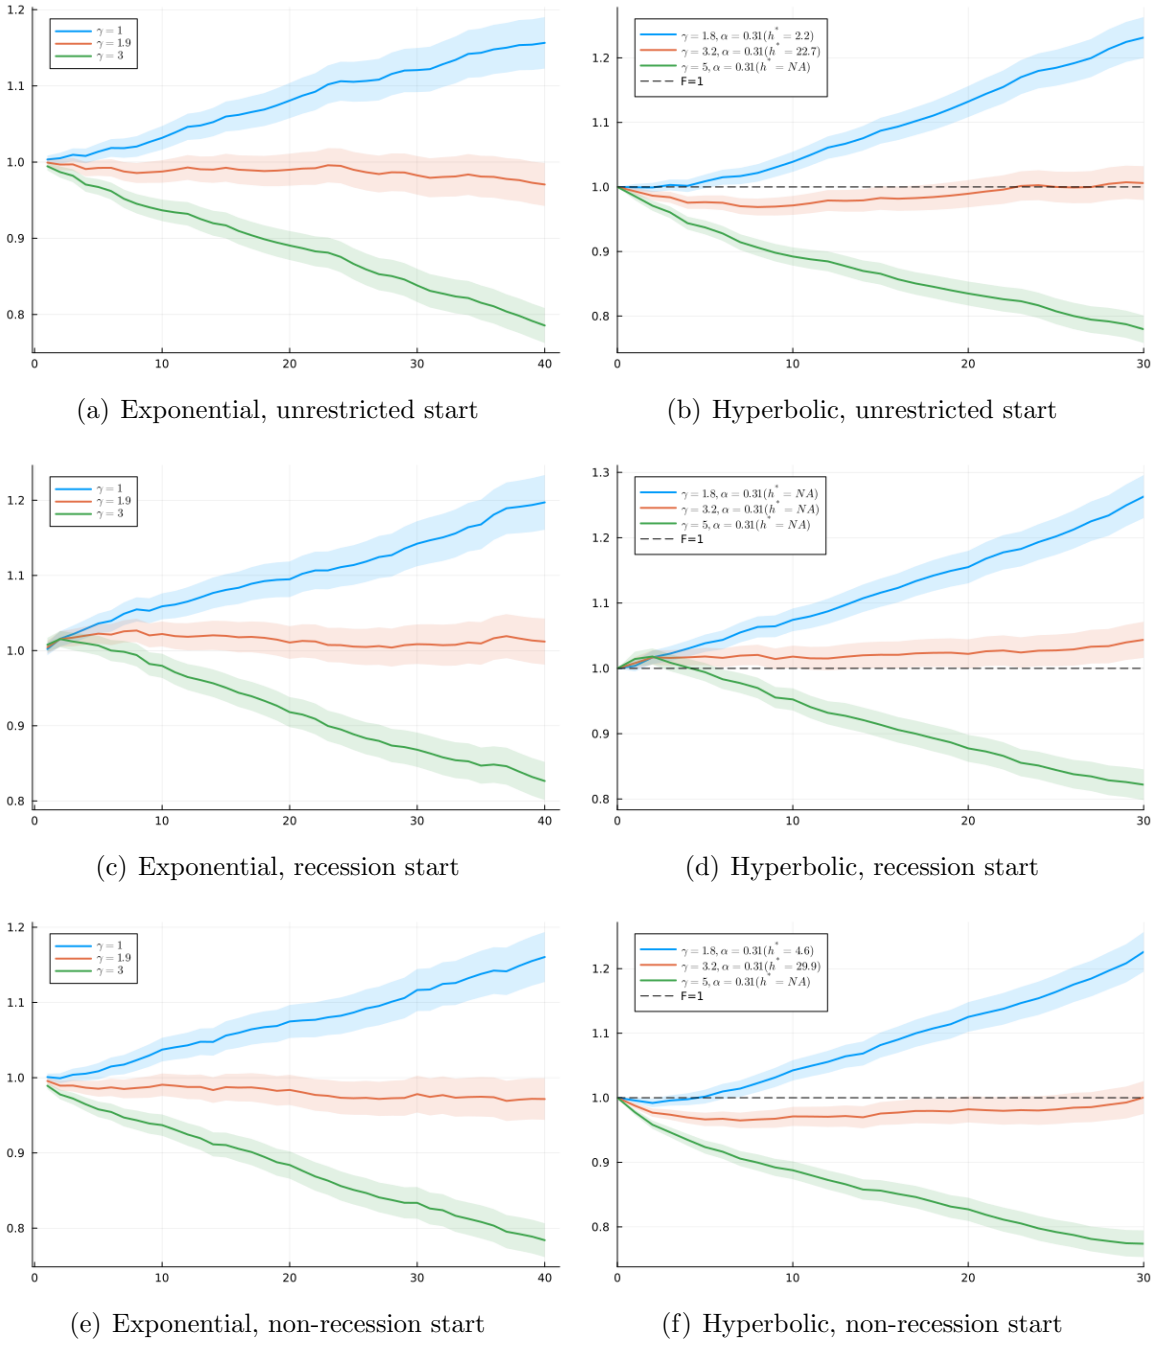
$F(1) = 1$, and at $\gamma = 3.2$ the curve falls below one, returns to one near 22 quarters, and stays there: the shape of Equation 31. At higher γ, **??** returns to one under exponential discounting.

The middle and bottom rows of the left column can be read with the shape result of Remark 5. We do not know γ, so the panels try several. The investor holds the market and cannot short, and at an optimum the Euler term is at or below one at every horizon up to the plan horizon and equal to one there. So $F < 1$ at some (γ, h) is consistent with the theory: it is the shape of the curve before h^*. $F > 1$ at (γ, h) is not: an investor with this γ, holding this position, would buy more, so he is not at an optimum with a plan of length h or longer — either γ is not his, or the position has not yet been adjusted to a shock. From a normal start, $\gamma = 1$ is ruled out and $\gamma = 1.9$ and 3 are admissible at every horizon. From a recession start, $F > 1$ in the first quarters at $\gamma = 1.9$ and 3, and at one to three quarters no $\gamma \geq 0$ brings the Euler term to one at this position (**??**). The rise is the recession's transitory effect, mainly low consumption growth,[36] which is collected in the first quarters and then

[36]With CRRA, $M_{t,t+h} = \beta^h (c_{t+h}/c_t)^{-\gamma}$. Consumption growth is negative in the first recession quarter and below normal for about two years, so the discount factor is high over these horizons and $\mathbb{E}_t[M_{t,t+h} R_{t,t+h}]$ rises above one at the unchanged position. The higher return variance adds a smaller part; the mean real return is lower in the first quarter and higher after. **??** shows the three pieces.

disappears. Two readings give the same conclusion. If the plan is long, the short-horizon $F > 1$ is the shock passing through a position set for a longer plan. If the plan were short, one to three quarters, no γ would price it. Either way the investor is not investing at one to three quarters in a recession, which is what the Euler-equation-error test of Section 5 finds, and the plan horizon is longer than the first quarters of the recession, where $F > 1$.

Figure 10: Objective-Simulation Horizon Euler Terms



(a) Exponential, unrestricted start

(b) Hyperbolic, unrestricted start

(c) Exponential, recession start

(d) Hyperbolic, recession start

(e) Exponential, non-recession start

(f) Hyperbolic, non-recession start

Notes: The panels report simulations from the full-sample objective transition matrix P^o. Each line is the mean across 1,000 paths; the shaded bands are pointwise 95% Student-t confidence intervals for the mean. The middle and bottom rows restrict only the initial state to recession and non-recession quarters, respectively; subsequent states follow P^o. The left column uses exponential discounting, $\beta(h) = 0.99^h$; the right column uses generalized hyperbolic discounting, $\beta(h) = (1 + \alpha h)^{-\rho/\alpha}$, with $\rho = 0.012$.

4 Methodology

The previous section introduced the investment horizon and the Horizon-Adjusted Fundamental Pricing Equation. This section focuses on extracting information about how investors perceive their horizons from observed data. Existing methods can recover subjective beliefs from asset prices by requiring only that the standard Euler equation holds (Kitamura et al., 2004; Ghosh and Roussellet, 2023). However, these methods cannot be applied directly to the Horizon-Adjusted Fundamental Pricing Equation, which relies on a first derivative with respect to the horizon and therefore requires knowledge of subjective beliefs not only at a single future date but also at neighboring dates, creating practical difficulties. I propose an alternative, quasi-testing approach that compares two candidate horizons and evaluates which is more likely to correspond to the investor's true horizon. The section proceeds in three steps. First, it introduces the representative-investor model. Second, it shows how to identify a range of plausible investment horizons when subjective beliefs are observed. Third, it explains how these beliefs can be recovered from the data.

4.1 Model Setup

We start with the class of consumption-based asset pricing models with a representative agent, who makes investment and consumption decisions at each time period. A representative agent framework is widely utilized because it provides analytical Euler equations for the joint behavior of asset returns and aggregate consumption. Even though the representative agent paradigm is sometimes criticized, it is still used due to its ability to explain the aggregate time series (Eichenbaum et al., 1988; Campbell and Cochrane, 1999).

We could state the problem from the individual agent's point of view, as many papers do (Weil, 1989; Campbell and Cochrane, 1999), and later aggregate the results, assuming that the representative agent's preferences could take the same form as the individual agents' preferences; or they could result from the aggregation of heterogeneous consumers with other preferences (Constantinides and Duffie, 1996). By assuming that the market is complete and hence households can insure their idiosyncratic shocks, the agent's SDF will depend on aggregate consumption growth. Brav et al. (2002) state, "Full consumption insurance implies that heterogeneous consumers are able to equalize, state by state, their marginal rate of substitution. Therefore, the equilibrium in a heterogeneous-consumer, full-information economy is isomorphic in its pricing implications to the equilibrium in a representative-consumer." Alternatively, to save space, we initially state the problem for a representative or single agent (Bansal and Yaron, 2004; Kocherlakota, 1990).

It is worth noting that at each time t, we assume that the market is in equilibrium in

such a way that the representative agent's Euler equation should be satisfied. The market reaches equilibrium through the adjustment of asset prices; thus, no direct trade is required in this framework. Typically, the production side is not considered, and it is assumed that the representative agent makes decisions about the allocation of consumption to firms and back.

4.2 Identification of Investment Horizon Range

In this subsection, we present an approach that is used to identify investment horizons. The approach is a sort of 'quasi' test, which compares the agent's decisions made under different hypothetical beliefs and identifies the most likely one that maximizes his utility and matches his behavior. Section 4.3 covers the method to recover investor beliefs as a non-parametric joint distribution of consumption, returns, and other conditional variables.

First, we examine the fundamentals of analyzing an agent with different beliefs. Specifically, we explore what can be inferred about the agent's decision-making when assuming alternative beliefs, while holding all other factors constant. We extract the beliefs in a manner that satisfies the first-order condition (48). The hat symbol denotes estimated beliefs.

$$U'(C_t) = \hat{\mathbb{E}}_t^S \left[\beta U'(C_{t+1}) R_{t,t+1}\right] \tag{48}$$

If we know the true subjective beliefs S and find that the first-order condition is not satisfied under these beliefs, it would mean that the agent, by solving his maximization problem, meets some constraints. If LHS of Equation 48 is less than RHS, the agent reaches the upper boundary solutions for the particular asset. He would like to move some of his current consumption to the future by buying this asset but could not. We do not know exactly the reasons for this; it could be explicit restrictions or any implicit behavior constrains[37]. Conversely, if LHS of Equation 48 is greater than RHS, the agent's solution touches the opposite, lower boundary. In this case, he would like to short more of the particular asset than he could, or he just doesn't invest in this asset at all. For simplicity, we define Euler Equation Error as Equation 49. We can state that if Euler equation error $EER(S) > 0$, the agent would like to increase his investment. By contrast, if $EER(S) < 0$, the agent would like to decrease his investment, or he just does not invest at all.

$$EER(S) = \hat{\mathbb{E}}_t^S \left[\beta U'(C_{t+1}) R_{t,t+1}\right] - U'(C_t) \tag{49}$$

We define the agent's beliefs S as a joint distribution of conditional variables such as

[37] For example, the first-order condition (FOC) for drug trading is not satisfied.

consumption, market returns, and others, which affect the agent's expectations and SDF. Assume that we have two possible candidates for the true agent beliefs, $S1$ and $S2$. For each of them, we calculate *EER()*, and for simplicity, let's assume that $EER(S1) > 0 > EER(S2)$. Since the FOC is not satisfied for either of them, we can't definitively identify the true one[38]. Yet, we can infer which is more probable by considering the agent's decisions. If the agent is investing, $S1$ would be more probable, and the agent would reach an upper bound solution, as demonstrated above. Conversely, if the agent does not invest, it is likely that they hold the beliefs $S2$; otherwise, if the agent invests while $EER(S2) < 0$, they would be incentivized to disinvest until $EER(S2) = 0$ or until they liquidate their stock position entirely (assuming short selling is prohibited). Therefore, our inference about the agent's beliefs depends on their actual behavior, not only on the sign of EER. It is noteworthy that EER() can be reformulated for excess return Equation 50 and interpreted similarly. In equilibrium, the agent should not be able to gain utility at the margin by selling bonds and investing in stocks. Otherwise, if $EER() > 0$, the agent would want to increase their investment.

$$EER(S) = \hat{\mathbb{E}}_t^S \left[\beta \frac{U'(C_{t+1})}{U'(C_t)} (R_{t,t+1} - R^{rf}_{t,t+1}) \right] \tag{50}$$

Short vs Long horizon test

In this part, we describe the approach that we use to extract information about investment horizons. We conduct the test for the given investment horizon $2h$ and as a result, we can conclude which investment horizon is more probable: $2h$ or half of them - h, which we refer to as the long and short horizon types respectively. The test is based on the assumption that agents' beliefs are time-consistent in terms of transitivity, so the expectation about two periods ahead equals the multiplication of two sequential expectations[39]. We construct the test by comparing two identical agents with the same lifetime and the same expected paths for all variables, differing only in the investment horizon type.

We define the agent who maximizes his lifetime utility with a lifetime $2h$. The agent could have one of two investment horizons: h vs $2h$. Thus his beliefs are either S (short) or L (long), corresponding to each type of horizon. If the agent is the short type, he invests with a horizon of h and reinvests in the same horizon in the next period. If he is the long horizon type, he invests only with a horizon of $2h$. It is important to note that both types live the

[38] At this moment, we consider the simplest model, without the possibility of having a different horizon, so we can state that $EER(S) = 0$ leads us to true beliefs.

[39] For simplicity, with only one possible state, if the agent expects 10% growth next month, they would expect 1.10^2 in two months. In the case of T possible states, the agent, two periods ahead, has T*T possible paths, and the beliefs transitional matrix remains the same.

same time period - $2h$ and have the same expected path of all conditional variables, the only difference is the investment horizon type. For simplicity, further, we assume h equals one quarter.

Consider, if the agent is the short type horizon investor with beliefs S, then his first order condition should be satisfied:

$$U'(C_t) = \hat{\mathbb{E}}_t^S \left[\beta U'(C_{t+1}) R_{t,t+1}\right] \tag{51}$$

Since beliefs are the joint distribution of consumption and returns for all possible states, by obtaining expectations about C_{t+1} for each state T, the next period first-order conditions should also be satisfied:

$$U'(C_{t+1}) = \hat{\mathbb{E}}_{t+1}^S \left[\beta U'(C_{t+2}) R_{t+1,t+2}\right] \tag{52}$$

The agent solved his problem backward, thus substituting Equation 52 into Equation 51 to obtain the two-period Euler Equations. It is noteworthy that the two-period Euler Equation 53 includes all possible inner paths for all states. In addition, we would like to highlight that short beliefs are formed for one time ahead, therefore, for expectations two periods ahead, two symbols of expectation are used (the first one is for the time t and second expectations for each state T in the next period).

$$U'(C_t) = \beta^2 \hat{\mathbb{E}}_t^S \left[R_{t,t+1} \hat{\mathbb{E}}_{t+1}^S \left[U'(C_{t+2}) R_{t+1,t+2}\right]\right] \tag{53}$$

Consider, if the agent is the long type, his belief $\mathbb{E}^L$ is extracted such that the first-order condition (FOC) is satisfied:

$$U'(C_t) = \hat{\mathbb{E}}_t^L \left[\beta^2 U'(C_{t+2}) R_{t,t+2}\right] \tag{54}$$

To sum up, we have an agent who lives time period $2h$. We have two candidates for his beliefs; he can be either long or short type horizon investor. Both candidate take into account all possible paths and states during this time. Thus, it is exactly the same agent; the only difference is in the beliefs he has used. We are going to test which beliefs are used by comparing the Euler Equation Error.

The Euler Equation Errors are given by Equation 55. If $EER(S) > EER(L)$ and if the agent actually invests, as we covered above, he likely has beliefs S, and we can conclude the agent is a short-horizon investor. On the other hand, If $EER(L) > EER(S)$, the agent is likely a long-horizon investor. Yet, we cannot state that it is the 'true' investment horizons,

as we are only comparing two candidates. By comparing the sequences of candidate h and $2h$, we would be able to determine the range of most likely investment horizons.

It is worth noting that, due to the constraint, EER(L) is zero, as shown in Equation 54. However, we cannot conclude that these are the true beliefs; we merely extract them, assuming they are true. Thus, we only compare the sign of EER() for two beliefs' candidates and guess which of them is more probable if the agent invests. In addition, as we see above (50), the EER() can easily be estimated by using excess returns.

$$\begin{aligned} EER(S) =& \beta^2 \hat{\mathbb{E}}_t^S \left[R_{t,t+1} \hat{\mathbb{E}}_{t+1}^S \left[U'(C_{t+2}) R_{t+1,t+2} \right] \right] - U'(C_t) \\ EER(L) =& \hat{\mathbb{E}}_t^L \left[\beta^2 U'(C_{t+2}) R_{t,t+2} \right] - U'(C_t) \end{aligned} \tag{55}$$

4.3 Recovery of Beliefs and Risk Preferences

This subsection outlines the extraction of subjective beliefs, following Ghosh and Roussellet (2023) and Kitamura et al. (2004). Because this procedure applies generically to any asset pricing model, the methodology is independent of the specific investment horizon. To simplify notation, we use single subscripts for one-period variables throughout this subsection (e.g., $R_{t+1} \equiv R_{t,t+1}$).

The absence of arbitrage opportunities implies the existence of a strictly positive pricing kernel M_{t+1} such that the Euler equation (56) is satisfied. Capital S indicates that expectations are based on the subjective agent's beliefs. The SDF of economic agents can be identified as a parametric function (57) of consumption growth $G_{t+1} = C_{t+1}/C_t$, set of risk factors Y_{t+1}, and θ_0 as the true value of the vector of parameters driving the SDF.

$$\mathbb{E}^{\mathcal{S}} \left[M_{t+1} \mathbf{R}_{t+1}^e \mid \mathcal{F}_t \right] = \mathbf{0}_k \tag{56}$$

$$M_{t+1} = M \left(G_{t+1}, Y_{t+1}; \theta_0 \right), \tag{57}$$

Our goal is to extract these subjective beliefs S in a way that minimally distorts the true Data Generating Process (DGP) and satisfies the conditional Euler equations (56). We define two matrices $P^{(0)}$ and $P^{(s)}$ for the objective and subjective joint distributions, such that $p_{i,j}^{(0)}$ and $p_{i,j}^{(s)}$ represent the objective and subjective conditional probabilities of transitioning from the state observed at time $t = i$ to that observed at $t' = j$. The number of states equals the number of observations T. Each state is defined by the vector x_t, which is a set of conditional variables (g_t,r_t,y_t,...) over which joint distributions are defined. The Euler equation (56) can be rewritten as follows:

$$\mathbb{E}^{\mathcal{S}}\left[M\left(G_{t+1};\theta_0\right)\mathbf{R}^e_{t+1}\mid\mathcal{F}_t\right]=\mathbf{0}_k\Rightarrow\sum_{j=1}^{T}p^{(s)}_{t,j}\cdot M\left(g_j,y_j;\theta_0\right)\mathbf{r}^e_j=\mathbf{0}_k \tag{58}$$

The objective distribution is estimated as a non-parametric kernel:

$$\forall i,j\in\{1,\ldots,T\},\quad p^{(0)}_{i,j}=\frac{\mathcal{K}\left(\frac{x_i-x_{(j-1)}}{b_T}\right)}{\sum_{t=2}^{T}\mathcal{K}\left(\frac{x_i-x_{(t-1)}}{b_T}\right)} \tag{59}$$

To obtain the result, we use the Smoothed Empirical Likelihood (SEL) estimator which minimizes the Kullback-Leibler Information Criterion divergence (KLIC, or relative entropy) between the objective and subjective measures:

$$\widehat{p}^{(s)}_{t,\cdot}=\underset{\left\{p^{(s)}_{t,j}\right\}_{j=1}^{T}\in\Delta}{\arg\min}\sum_{j=1}^{T}p^{(0)}_{t,j}\cdot\log\left(\frac{p^{(0)}_{t,j}}{p^{(s)}_{t,j}}\right)\ \text{s.t.}\ \sum_{j=1}^{T}p^{(s)}_{t,j}\cdot M\left(g_j,y_j;\theta_0\right)\mathbf{r}^e_j=\mathbf{0}_k \tag{60}$$

and, for a given $\forall t$ and $j\in\{1,\ldots,T\}$ the closed-form solution is as follows:

$$\widehat{p}^{(s)}_{t,j}\left(\theta_0\right)=\frac{1}{1+\widehat{\lambda}_t\left(\theta_0\right)'\left[M\left(g_j,y_j;\theta_0\right)\mathbf{r}^e_j\right]}\times p^{(0)}_{t,j},$$

where $\widehat{\lambda}_t\left(\theta_0\right)\in\mathbb{R}^k$ are the Lagrange multipliers associated with the conditional Euler equation constraints, and solve the following unconstrained maximization problem:

$$\widehat{\lambda}_t\left(\theta_0\right)=\underset{\lambda_t\in\mathbb{R}^k}{\arg\max}\sum_{j=1}^{T}p^{(0)}_{t,j}\log\left[1+\lambda_t'\left(M\left(g_j,y_j;\theta_0\right)\mathbf{r}^e_j\right)\right]$$

The recovered conditional probabilities $\left\{\widehat{p}_{t,j}\left(\theta_0\right)\right\}_{t,j=1}^{T}$ represent the beliefs of the representative investor that are consistent with observed asset prices at each time. It is worth noting that beliefs are recovered for the particular moment, a time t, and are independent of previous and future periods. Thus, the approach does not have drawbacks such as those which occur from overlapping periods in time series econometric analysis.

5 Results for the Homogeneous Investment Horizon

This section is dedicated to homogeneous investment horizons, as covered in the previous section. The first part presents an overview of subjective beliefs and risk preferences for a representative investor with various investment horizons, which provides insights into the

true investment horizon. The second part presents the results of identifying the investment horizon range, as described in Section 4.2.

Data

We utilize data spanning from 1961:3Q to 2016:4Q.[40] For the stock market, we use real market returns computed as the value-weighted index of a broad set of stocks provided by CRSP.[41] As for aggregate consumption, we use real personal expenditure on nondurable goods and services, and for risk-free rates, we use the deflated nominal yield curves (ranging from 1 year to 5 years). The set of state variables (or conditional variables) consists of exponentially weighted moving averages[42] of past real excess returns, consumption growth, three factors from PCA on nominal risk-free yield curves, and the first difference of the conditional variables. We use a Gaussian kernel as a baseline.

5.1 Overview of Subjective Beliefs and Risk Preferences over Varying Investment Horizons

In this subsection, we extract subjective beliefs for the CRRA utility agent over possible horizons h, ranging from 1 to 20 quarters.

Firstly, we extract implied risk aversion for different investment horizons by minimizing Smoothed Empirical Likelihood (SEL) (60) over γ, for a given h:

$$\gamma_h = \min_{\gamma} SEL(\gamma, h) \tag{61}$$

Figure (11) shows the results for various investment horizons. The long-run risk aversion is about $\gamma = 50$, which is six times lower than that of the 1Q horizon. It stabilizes for horizons longer than 8Q, which significantly reduces the severity of the equity premium puzzle. All of the above suggests that the risk aversion for the shorter horizon is unrealistically high. Since we can assume that risk aversion is a stable and constant trait of the investor, for the rest of the paper we will use $\gamma = 50$, unless otherwise specified.

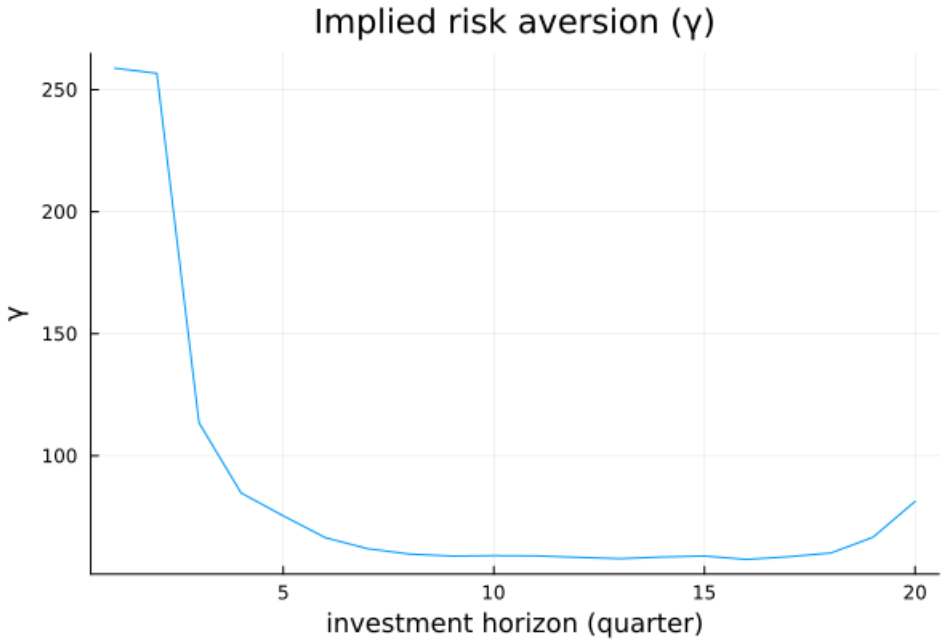
To support our findings, we extract risk aversion parameters for various setups. Figure (12)(a) presents results for gross returns on a five-asset portfolio and bonds. Figure (12)(b) shows the estimation of implied risk aversion preferences of the Epstein-Zin utility function

[40] The paper is currently being updated to cover a broader time span—1947:2Q to 2019:4Q. All results remain the same.

[41] In the fourth quarter of 2019, the index included 7,729 stocks.

[42] Our exponentially weighted variables are computed as $X_t^{(EW)} = \alpha X_t + (1-\alpha) X_{t-1}^{(EW)}$. We set $\alpha = 0.28$, whereby the past 13 quarters receive 99% of the weight. Our results are unaffected by the value of α.

Figure 11



Notes: The figure plots implied risk aversion γ for various investment horizons. Excess market returns are used in the Euler equation. The result is obtained by minimizing SEL over γ. The pricing kernels correspond to CRRA. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1947:2Q-2019:4Q.

on market returns and bonds. Finally, Figure (13) shows cross-sectional results, with estimations conducted separately for six non-overlapping portfolios. All these results demonstrate similar patterns of declining risk aversion values around the middle of investment horizons. Estimations, which include separate Euler equations for risk-free asset (gross returns Euler equations), show a slightly rising value at the longer tail. We do not delve deeply into this section; one hypothesis is that fixed income could have different investment horizons. We will speculate about this later.

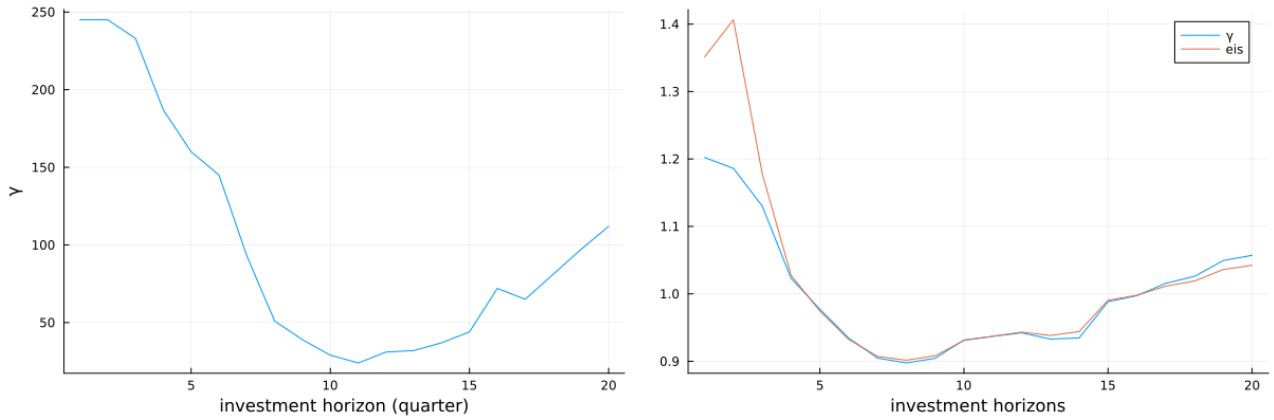
In the rest of this section, we will compare subjective and objective beliefs, exploring the differences and drawing insights about the 'true' investment horizon. The figure (14) presents the results of the divergence between objective and subjective excess returns for various horizons. Distortions between objective and subjective returns are much more pronounced at shorter horizons and subside by 8 quarters. Figure (15) shows annualized objective and subjective standard deviation for 1, 4, 8, and 10 quarters horizons. The first observation (which is not new) is that annualized standard deviation is declining over time. Objective and subjective standard deviations diverge up to 8Q, but the level isn't economically significant, so we can't conclude that it's not a realistic divergence. I want to highlight two additional

observations. The standard deviation is countercyclical after 10Q, suggesting that during a recession, agents are more certain that in the period two and a half years, some growth is likely expected, but they are not sure about the nearest future. Second, there are two time periods, before and after 1992, where 4Q subjective and objective standard deviations coincide before and diverges after. We will see the similar result later in the paper and discuss it there.

Figures (16 - 17) present the results of the divergence of objective and subjective Sharpe ratios for different investment horizons. Firstly, we observe that the Sharpe Ratio is growing over time, which is reasonable. However, as depicted in Figure 17, the Sharpe Ratio grows at a higher rate for up to 8 quarters and then stabilizes. Second, objective and subjective Sharpe Ratios significantly diverge, up to 5 times, for short horizons, which is not realistic and may be a sign that this horizon is not reasonable. Third, belief distortions are more pronounced at shorter horizons and less pronounced during a recession. This could be a sign that the horizon is shorter in 'bad' times. And finally, we observe the same pattern as for standard deviation, but much more pronounced; after 5Q horizons, there are several stages: before 1992, after 1992, and after 2009. The reasons for this are discussed below.

To conclude this section, we summarize all findings and highlight the main intuition. We examine subjective beliefs and attempt to identify whether they are plausible or unrealistic. The term 'unrealistic' is somewhat ambiguous and not strictly defined, but insights can be gained by comparing subjective and objective beliefs. Based on the findings in this section, we can make some inferences about the 'true' investment horizon. First, we observe that implied risk aversion declines rapidly and stabilizes around 8Q. Second, the subjective return and Sharpe ratio are approximately five times lower for 1Q, and only converge around 7-8Q. Finally, the Sharpe ratio exhibits two distinct patterns of growth, before and after 8Q. Something significant appears to happen around 8Q, leading us to infer that the true investment horizon may be close to this number.

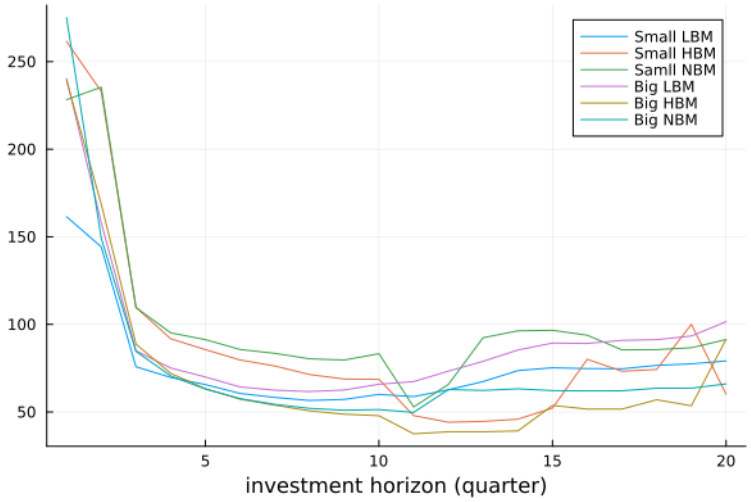
Figure 12: Implied risk aversion and EIS. Gross returns.



(a) (b)

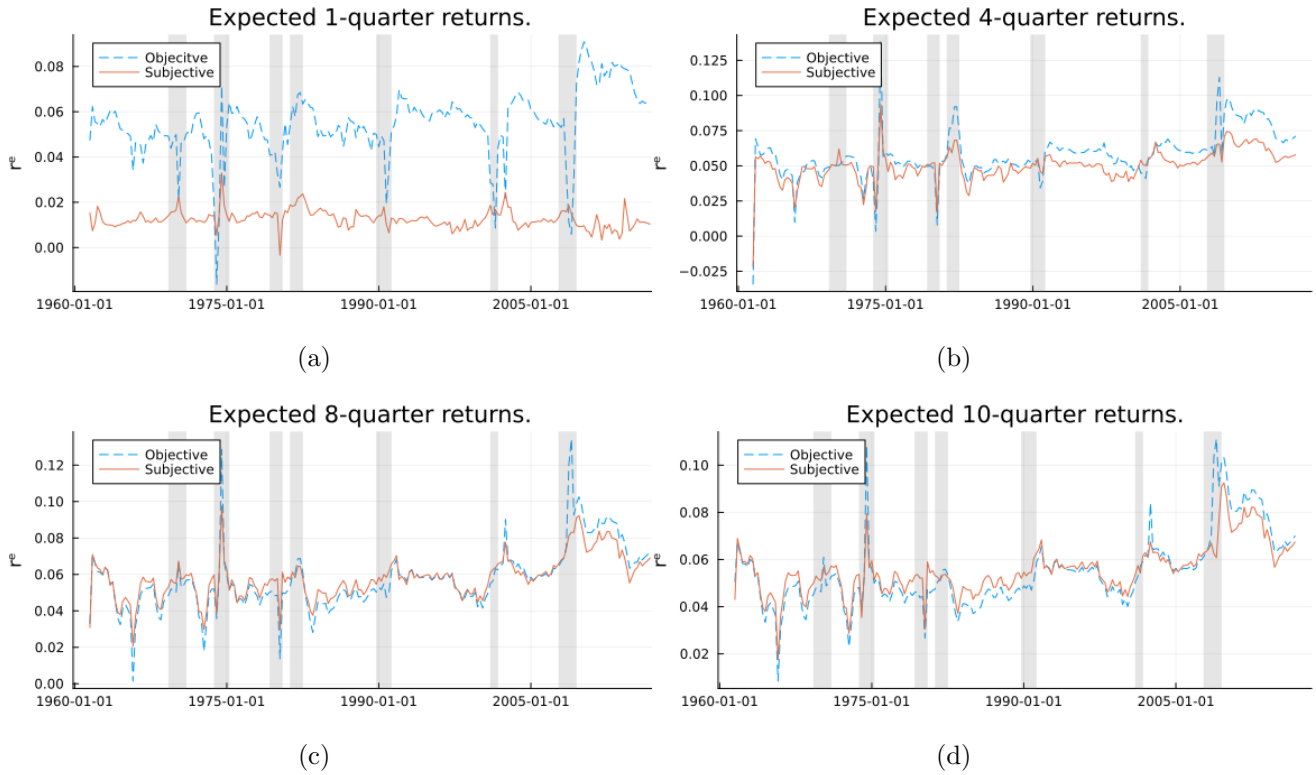
Notes: The figure plots implied risk aversion γ and EIS for various investment horizons. Assets: Market, Low and High B/M, Large and Small sizes. Panel (a) presents results for CRRA utility function and 5 assets and RF gross returns. Panel (b) presents results for Epstein-Zin utility function and gross market return and RF. The result is obtained by minimizing SEL. The pricing kernels correspond to CRRA. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 13: Implied risk aversion. Each asset class separately.



Notes: The figure plots implied risk aversion γ for various investment horizons. Excess returns on six, non overlapping portfolios are used in the Euler equation. The result is obtained by minimizing SEL over γ. The pricing kernels correspond to CRRA. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 14: Objective and Expected Excess Returns for Various Investment Horizons

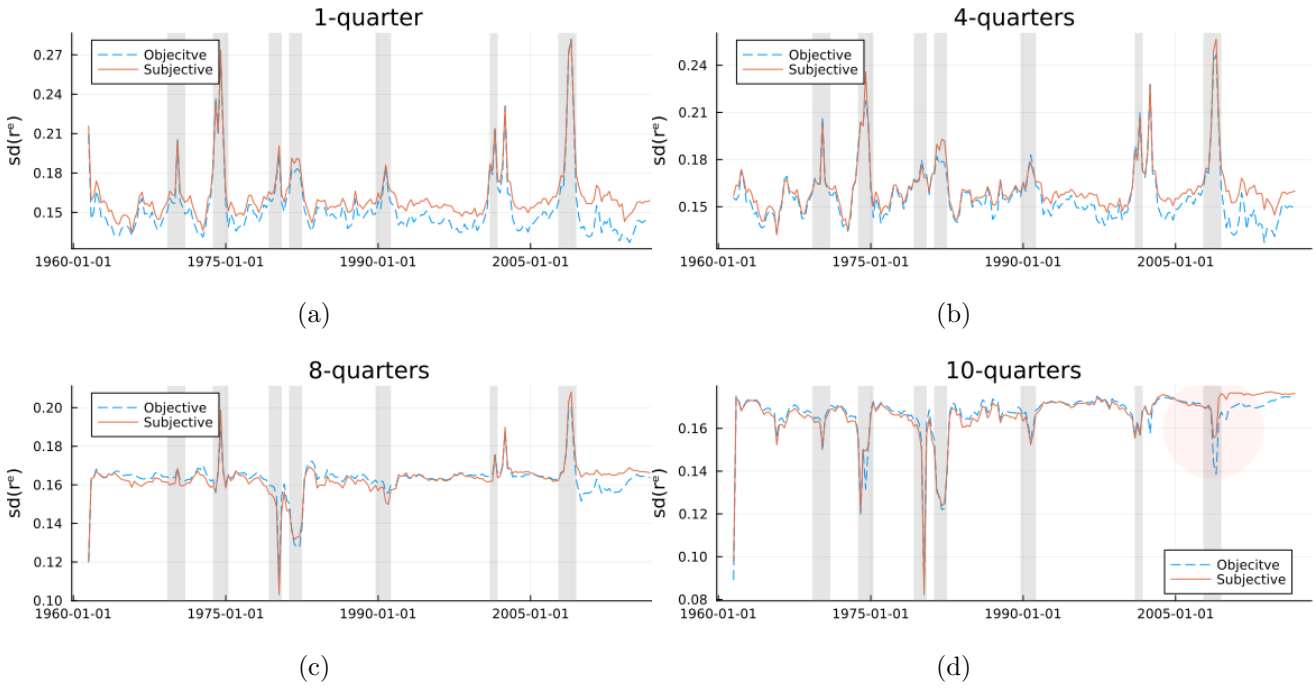


(a) (b)

(c) (d)

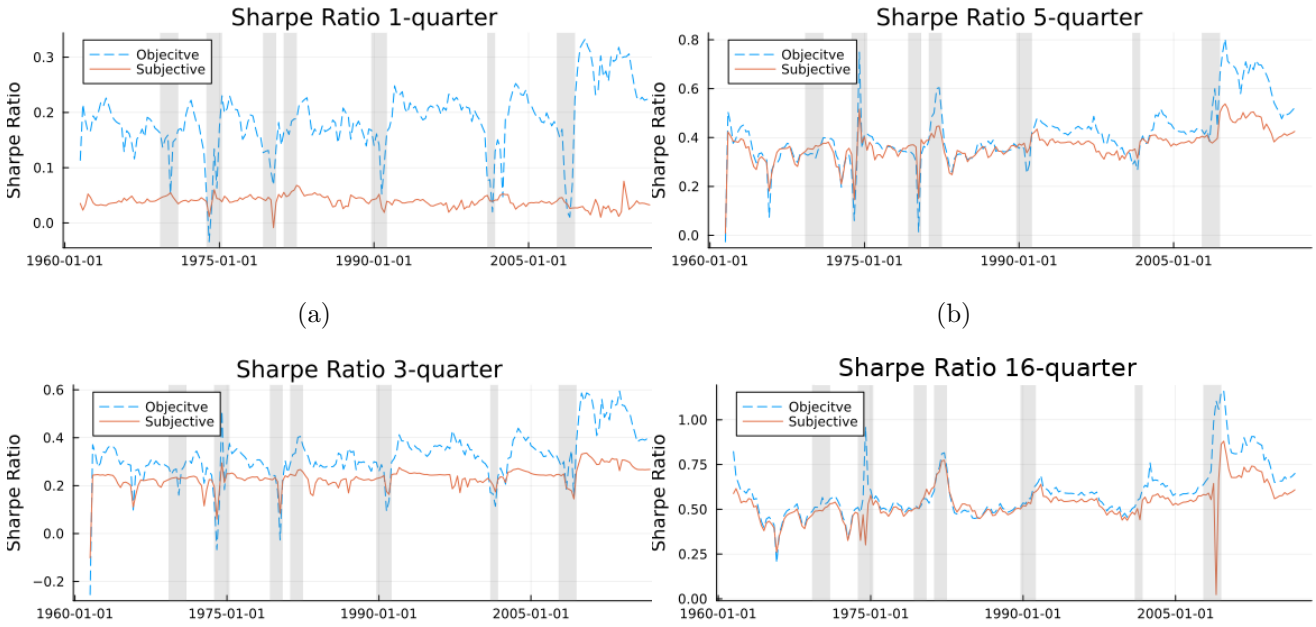
Notes: The figure plots annualized Objective and Subjective Expected Excess returns for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 15: Objective and Subjective Standard Deviation of Returns for Various Investment Horizons



Notes: The figure plots annualized Objective and Subjective Standard Deviations of returns for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 16: Objective and Subjective Sharpe Ratio for Various Investment Horizons

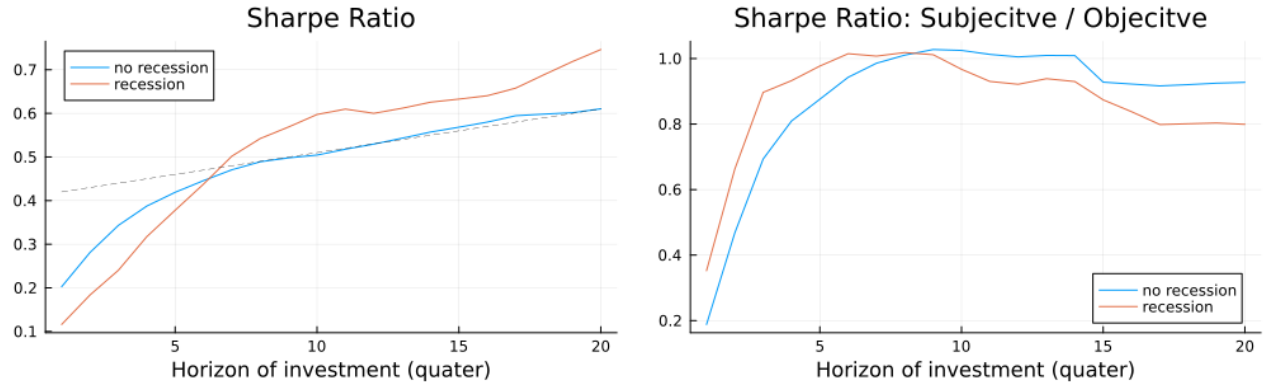


(a) (b)

(c) (d)

Notes: The figure plots Objective and Subjective Sharpe Ratio for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 17: Mean Sharpe Ratio for each Investment Horizon



(a) (b)

Notes: The left panel (a) shows objective Sharpe Ratio for various investment horizons. The right panel (b) shows the subjective Sharpe ratio relative to the objective one.

5.2 Identification Investment Horizon Range

In this subsection, we follow the methodology outlined in Section 4.2 and conduct a "quasi" test to identify the most likely investment horizon, h versus $2h$. We estimate $EER(S)$, as defined in equation (55), for various horizons h and present the results below, utilizing the Euler equation for excess market returns. If $EER(S) > 0$, the agent is more likely to have a short horizon h; otherwise, a long horizon $2h$ is indicated. This process is repeated for all possible values of h. For instance, as we will see below, when considering the pair 1Q versus 2Q, if the agent invests, they likely have a horizon of 2Q. Similarly, when comparing the pair 8Q versus 16Q, the agent likely has an investment horizon of 8Q.

Figure 18 presents a time series of $EER(S)$ for various values of h. Several observations can be made. Generally, during a recession, $EER(S)$ is elevated for horizons greater than 2Q, implying that the agent prefers a shorter horizon over a longer one during this unfavorable period. Conversely, $EER(S)$ is counter-cyclical for horizons up to 3Q. Thus, during a recession, the agent is less likely to invest in a short horizon of up to 3Q compared to a more favorable period. In summary, during a recession, the investment horizon tends to be shorter, yet farther from 1Q.

Figure 18



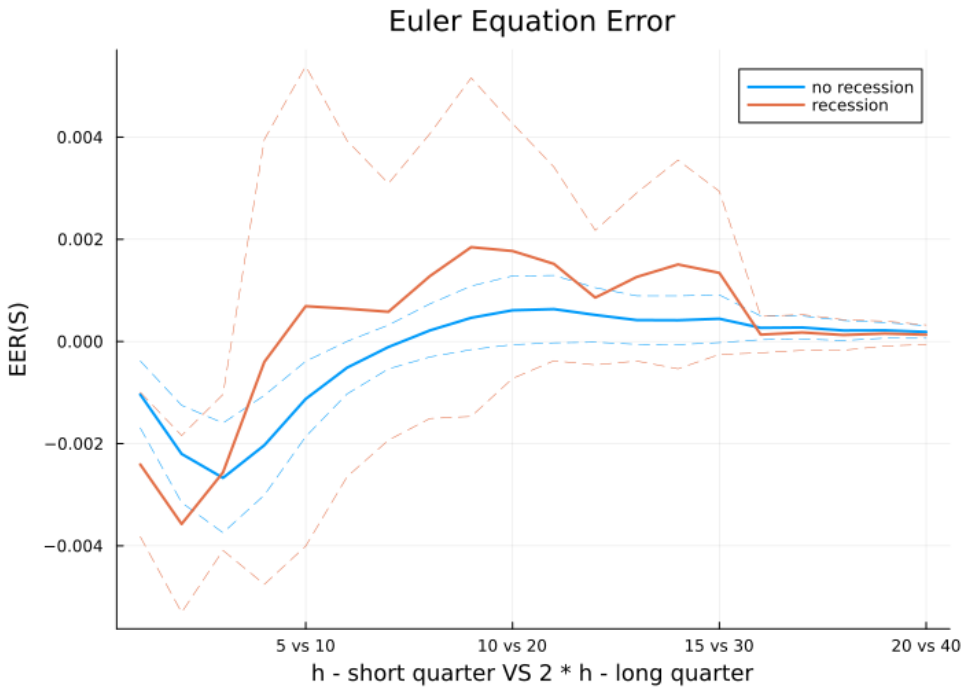
(a) (b)

Notes: The figure plots Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

The primary result can be derived from Figure 19, which illustrates the mean of EER for each possible horizon. When EER is below 0, it indicates that the agent is more likely to have a horizon of $2h$; conversely, if EER is positive, the agent is likely to have a short

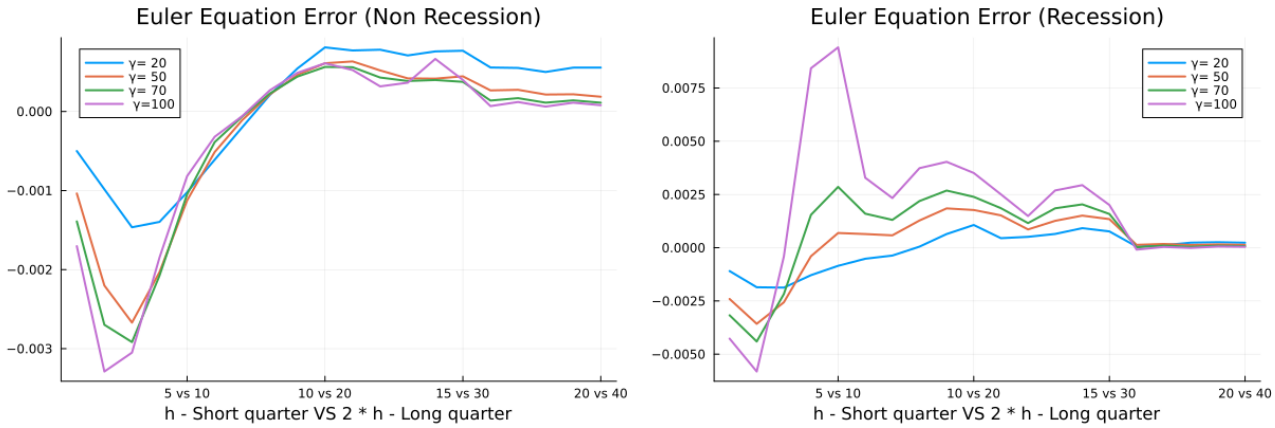
horizon. The key findings are as follows: (1) The most likely investment horizon ranges from 8 to 14 quarters. (2) During a recession, the most likely investment horizon narrows to between 5 and 8 quarters. (3) During a recession, the agent is highly unlikely to invest for a horizon of up to 3 quarters. To further support our findings, we conduct the same test using gross returns. Figure 22 presents results that show similar values. This raises important questions and hypotheses regarding the differing investment horizons between equity and fixed-income assets, which aligns with the concluding section of this paper and suggests the need for further research.

Figure 19



Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

Figure 20: EER for various Risk Aversion



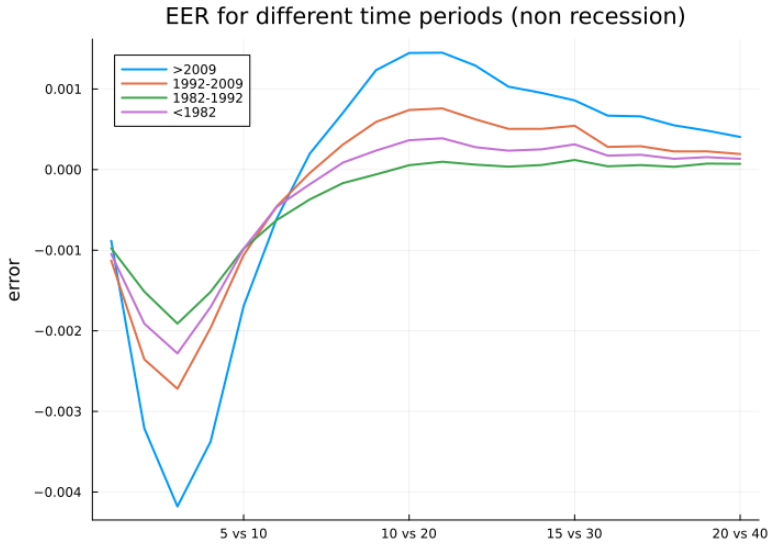
(a) (b)

Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment, for various levels of risk aversion γ. The left panel (a) shows results for non-recessionary times, while the right panel (b) shows results for recessionary times.

The next question we address is how the investment horizon depends on risk aversion. To explore this, we conduct the same test for various risk aversion parameters, with the results summarized in Figure 20. The key findings are as follows: (1) In non-recessionary periods, the investment horizon remains unchanged, regardless of the level of risk aversion. (2) More risk-averse agents exhibit a shorter investment horizon during recessions. (3) Additionally, more risk-averse agents are less likely to invest in a short horizon (1–2 quarters). The first finding is particularly noteworthy. It suggests that an agent's risk aversion influences the investment horizon in conjunction with market-wide risks, which are more pronounced during recessions but less so during favorable periods. Consequently, risk aversion does not impact the investment horizon in good times but becomes influential during recessions.

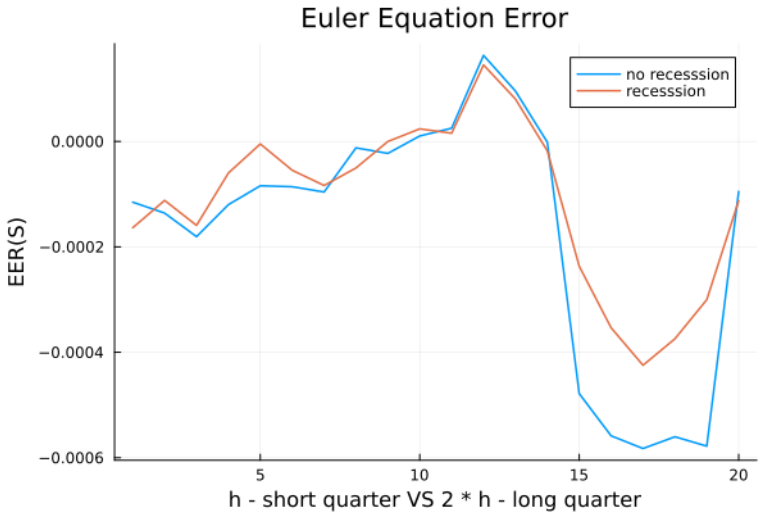
Finally, the results from Figure 21, along with the time series in Figure 18, demonstrate that the investment horizon shifts slightly across different time periods. For example, as shown in the figures, after the 2009 recession, agents are more likely to adopt a short-horizon approach compared to the period from 1982 to 1992. This trend, which began in 1992, can be attributed to easier access to the stock market for household investors, many of whom are (1) first-time participants in the stock market, focusing on shorter horizons, and (2) less experienced, leading to lower risk aversion, which makes them more inclined to favor short horizons of 1Q to 2Q.

Figure 21



Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment for different time periods. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

Figure 22: Short vs Long test for Gross Returns Euler Equation



Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. Gross returns for Market and RF. EER is sum of two. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

6 Heterogeneous Horizons Portfolio

The cross-sectional analysis of implied risk aversion reveals heterogeneity across asset classes. A natural question is whether portfolio composition also differs across horizons. Proposition 4 shows that asset holdings $a^*(h)$ peak at the optimal horizon h^*, so by decomposing the investor's holdings across horizons we can see which horizon is optimal for each asset class. This section examines that implication empirically.

In this section, we start by examining the portfolio composition for each given horizon separately. In a representative investor setup, this question is not particularly relevant since it is assumed that the investor holds the entire market. However, it can provide insight, reveal useful information, and prompt other possible research questions.

Additionally, from another perspective, we can consider an individual investor who is fully insured against idiosyncratic shocks and whose SDF depends on aggregate consumption growth. We can extract information about their possible portfolio composition for particular investment horizons. Gressis et al. (1976) show that agents with the same utility function but different investment horizons will choose different portfolios in the mean-variance approach. Lee et al. (1990) also state that each investor with an h-period horizon holds an identical portfolio, which should differ across different values of h, in the generalized CAPM.

In the second part, we examine a more restrictive model, where the representative investor holds the entire market, yet different asset classes in their portfolio have various investment horizons.

Overview of portfolio composition for a given investment horizon

Consider an agent's investment portfolio P as an asset; then, for a given asset/portfolio, the Euler equation Equation 62 should also be satisfied. Portfolio P consists of J assets with weight w_{jh}. For each investment horizon, portfolio P has a different composition.

$$\begin{aligned} \mathbb{E}^{\mathcal{S}}\left[M_{t+h}\mathbf{R}^{P}_{t+h} \mid \mathcal{F}_t\right] &= \mathbf{1} \\ \mathbf{R}^{P}_{t+h} &= \sum_{j=1}^{J} w_{jh} R^{j}_{t+h} \end{aligned} \tag{62}$$

We follow the same procedure as before and extract portfolio weights for different investment horizons by minimizing the Smoothed Empirical Likelihood (SEL) (60) over $w_{j\cdot}$, for a given h:

$$w_{jh} = \min_{\{w_{jh}\}} SEL(w_{\cdot h}, h) \tag{63}$$

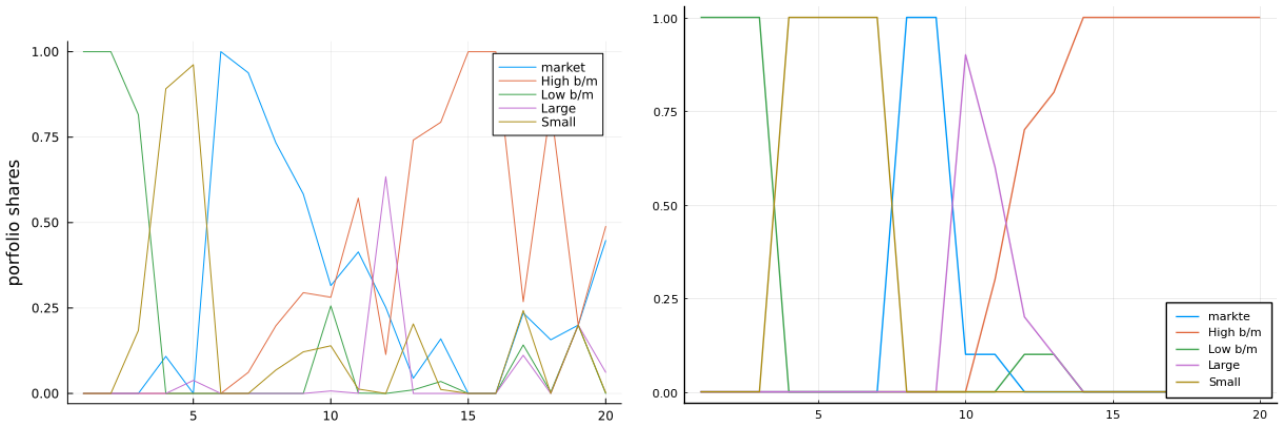
We use a five-asset investment portfolio, which consists of the following assets: total

market, high and low B/M, and large and small sizes. The model is estimated for different setups. First, we use excess returns and thus one Euler equation. Second, we use gross returns and two Euler equations: one for the portfolio and one for the risk-free asset.

It is worth noting that the asset portfolios are overlapping. This is done intentionally to easily identify where the investment portfolio coincides with the whole market; otherwise, we would see the composition of the market portfolio, and the results would be much noisier. In further work, we would use non-overlapping asset portfolios.

We expect that for the representative investor, there should be some horizons for which the investment portfolio coincides with the total market, as expected, since the representative agent holds the whole market. Second, since assets demonstrate different joint behavior with consumption growth, we expect that portfolio composition will differ over horizons.

Figure 23: Portfolio composition for a given investment horizon



(a) (b)

Notes: The figure shows asset weights in the investment portfolio if the agent invested for the given horizons (quarter, x-axis). Panel (a) represents results for excess returns, while Panel (b) represents results for gross returns.

Figure 23 shows weights for assets in the investment portfolio for various horizons. The left and right panels represent results for excess and gross returns, respectively[43]. First, we see that both results are quite similar and support each other. The main result is that, as expected, in the middle horizons, the investment portfolio coincides with the whole market. We can argue that this is highly likely to be the range for "true" investment horizons, and this result supports the previously obtained range. Third, we see that if the agent had shorter investment horizons, they would prefer low B/M assets, and if the agent had longer

[43]Results for gross returns are obtained on the grid, so they are less smooth.

horizons, they would invest in high B/M stocks. This seems quite reasonable and raises the hypothesis that an investor can have a portfolio with different horizons for various assets. We will explore this in more depth in the next part.

Investment portfolio with heterogeneous horizons

As we see above, the composition of the optimal portfolio depends on the investment horizon, as various asset classes exhibit different joint distributions with consumption over different time spans. That analysis was not scientifically rigorous and had the sole purpose of providing insight for further research. In this part, we assume that the representative agent holds a total market portfolio, as expected, and has different investment horizons for each class of assets. This setup can also be easily generalized to a heterogeneous-agents setup, where agents have different horizons. For instance, private and institutional investors could have different horizons and hold different portfolios.

The model setup is as follows: an agent has two possible investment horizons—short and long. The market consists of three asset classes: low book-to-market (B/M), neutral B/M, and high B/M (alternatively, six classes: separated into small and large sizes). The representative agent holds the total market and invests in each asset class for a specific investment horizon (an asset class cannot be divided between two horizons). The agent's FOC for each horizon should be satisfied:

$$\begin{aligned} \mathbb{E}\left[M_{t+S}(\mathbf{R}^{Short}_{t+S} - \mathbf{R}^{rf}_{t+S}) \mid \mathcal{F}_t\right] &= \mathbf{0} \\ \mathbb{E}\left[M_{t+L}(\mathbf{R}^{Long}_{t+L} - \mathbf{R}^{rf}_{t+L}) \mid \mathcal{F}_t\right] &= \mathbf{0} \end{aligned} \tag{64}$$

$$\begin{aligned} \mathbf{R}^{Short}_{t+S} &= \frac{w^h_S R^h_{t+S} + w^n_S R^n_{t+S} + w^l_S R^l_{t+S}}{w^h_S + w^n_S + w^l_S} \\ \mathbf{R}^{Long}_{t+L} &= \frac{w^h_L R^h_{t+L} + w^n_L R^n_{t+L} + w^l_L R^l_{t+L}}{w^h_L + w^n_L + w^l_L} \\ w^i_S + w^i_L &= 1, \quad \forall i \in \{h, n, l\} \\ w^i_S, w^i_L &\in \{0, 1\} \end{aligned} \tag{65}$$

where the superscripts denote the asset classes: h (high B/M), n (neutral B/M), and l (low B/M).

We do not know the exact "true" two horizons, short and long, so we estimate over all available pairs $\{S, L\}$. For given horizons S and L, we extract subjective beliefs for these horizons and portfolio weights w by minimizing the sum of $SEL^{S,L}$.

For each possible pair of S and L from 2 to 18 quarters, we extract portfolio compositions

and present the results as follows: a subplot labeled 'short' represents each pair where the S-short horizon is on the x-axis, and the y-axis shows the proportion of each particular asset in the short portfolios among all portfolios for this S as a short horizon. For instance, as Figure 24(a) shows, if the short horizon is three, all possible pairs $\{S, L\}$ are: $\{3, 4\}$, $\{3, 5\}, \ldots, \{3, 18\}$. Among all these pairs, low B/M assets are present in approximately half of the cases in short horizon portfolios, while high B/M assets are never present in short horizon portfolios.

The subplot labeled 'long' represents the opposite scenario, where each pair with the L-long horizon is on the x-axis. Figure 24(a) shows that if the long horizon is twelve, all possible pairs $\{S, L\}$ are: $\{2, 12\}$, $\{3, 12\}$, ..., $\{11, 12\}$. Among all these pairs, low B/M assets never appear in long horizon portfolios. Together, these two results suggest that if three quarters were a short horizon, it is very likely that for each possible pair, low B/M assets would always be in short portfolios.

Figure 24 shows results for three asset classes for normal and abnormal times, and Figure 25 shows results for six asset classes. First, we can state that during a recession, the investment horizon for all asset classes is shorter since the probability that an asset class is in a short-horizon portfolio is higher during a recession. This result is consistent with what we see in the previous sections.

Second, low B/M assets are located in short-horizon portfolios, regardless of their size, and high B/M assets are in long-horizon portfolios. In et al. (2011) support these results[44].

Third, considering sizes, large neutral and high B/M assets are more likely to be in short-horizon portfolios compared to small-sized assets, yet size has a quite subtle marginal effect. Therefore, we can conclude that more risky assets, such as small assets and undervalued assets (high B/M), are located in long-horizon portfolios, and that investors expect significant returns for currently undervalued assets and small assets in the long run. This can also be connected to momentum effects, since low B/M stocks are likely fast-growing and highly valued. Thus, they are located in short-horizon portfolios with the expectation of maintaining this momentum for a while.

We can consider this model from another perspective. Various asset classes have different horizons, and the short and long portfolios are held by different agents, such as individual investors and institutional investors. Both types of investors are fully insured against idiosyncratic shocks, so their SDF depends on the same aggregate consumption growth. This point of view can provide an additional reason why, in general, even in the case of many agents, some do not hold market portfolios: they may have various horizons.

[44]The paper uses the wavelet approach and adopts mean-variance frontiers with different investment horizons.

The differences in investment approaches between institutional and individual investors are widely covered in the literature. Ferreira and Matos (2008) and Gompers and Metrick (2001) show that U.S. institutions have a preference for value stocks (high B/M), whereas Barber and Odean (2000) show that households bias their investments toward high-beta stocks (low B/M can be considered as such stocks (Fama and French, 1993; Campbell and Vuolteenaho, 2004)). As we can see in our results, high B/M assets are in long portfolios and low B/M assets are in short ones. Many studies state that individual investors trade more often (Barber and Odean, 2000; Barber et al., 2009; Cohn et al., 1975). This aligns with our findings as short-horizon portfolios are likely to be traded more often than long-horizon portfolios. In addition, Basak and Pavlova (2013) find that institutional investors demand riskier stocks than retail investors. All these findings support our results, which suggest that individual investors could hold shorter-horizon portfolios, whereas institutional investors could hold longer-horizon portfolios.

Investment horizons as asset characteristics, not investors'

We would like to consider another point of view in a broad manner. We constructed an investment portfolio with different horizons for various assets and connected it to agent investment horizons, and even to heterogeneous agents. It is a plausible connection, yet we would like to return to the initial problem and propose that the horizon may depend not on the agent but on the asset characteristics. We see a strong separation between low and high B/M assets, where the former are associated with short-term portfolios and the latter are associated with long-term portfolios. Thus, some of the inherent characteristics of the asset may affect this allocation.

High B/M stocks are typically utilities, manufacturing, and financial services companies that constantly pay dividends. We can therefore characterize them as companies from stable sectors promising returns over long periods. Conversely, low B/M assets include technology companies, startups, and firms in rapidly expanding industries that rarely pay dividends. These companies are much less stable and promise returns in the relatively near future[45]. We can see how the distribution of expected returns over time is one of the key characteristics of these companies. Lettau and Wachter (2007) present firms as long-lived assets distinguished by the timing of cash flows and show that growth and value firms covary differently with the discount rate and cash flows, which explains their pricing. Along with this, Van Binsbergen (2020) presents an empirical paper on duration-based stock valuation, using returns from dividend strips. Recently, Gormsen and Lazarus (2023) introduces a duration factor and shows that cash-flow duration is an important determinant of the value premium.

[45]In this case, returns are typically generated by the fast-growing values of firms.

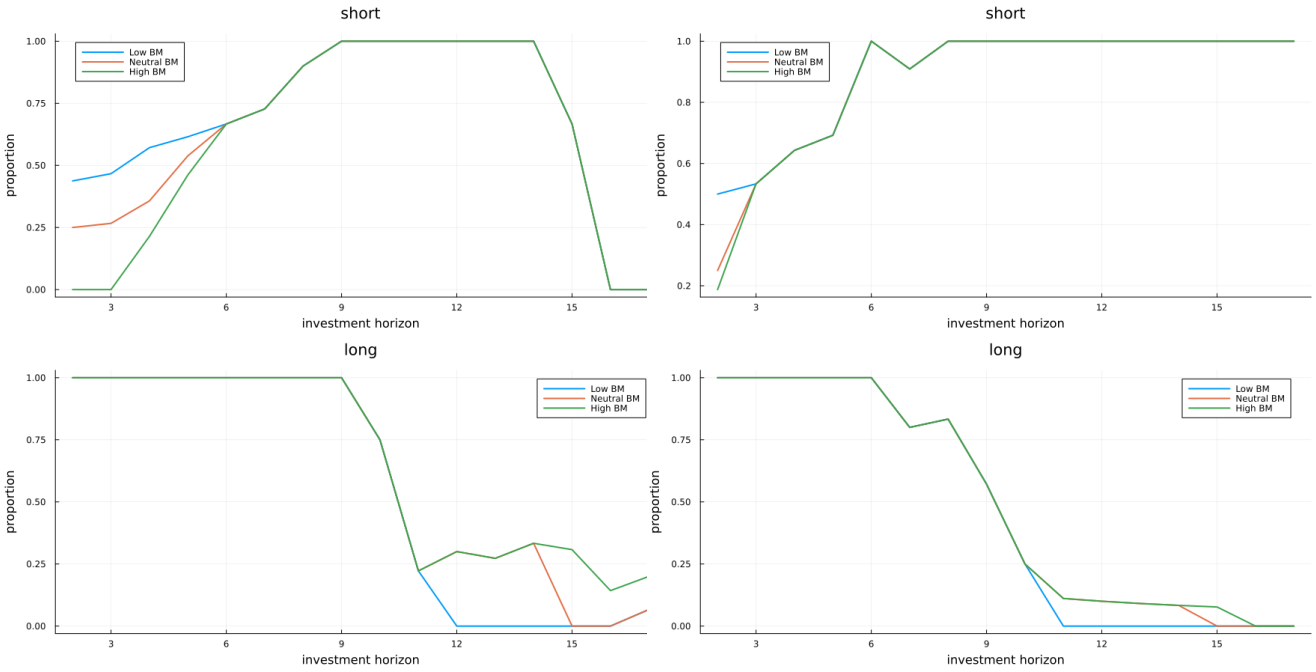
As we see above, 'implied equity duration'/'duration factor' is an inherent characteristic of an asset. Therefore, we can state that the investment horizons of the Euler equation[46] should match this expected duration of the asset. For simplicity, we will call these short-horizon and long-horizon assets.

This point of view, seeing investment horizons as an asset characteristic, can explain some of the features of book-to-market portfolios. High B/M stocks have a lower market beta than low B/M stocks; beta is typically computed on monthly data, and it is clear that current economic fluctuations[47] have a greater effect on short SDF than on long SDF. Thus, all else equal, long-horizon stocks would have a lower market beta. Campbell and Vuolteenaho (2004) present another point of view, the two-beta model, and show that value stocks have considerably higher cash-flow betas. This model is supported by and aligns with the view of value stocks as long-horizon stocks, since their cash flows are distributed in the distant future, and therefore their values are much more sensitive to these cash flows.

[46] We do not say 'agent's investment horizons' intentionally, since an agent can have many of them.

[47] Economic fluctuations include both market and consumption changes; thus, they can be applied to different models.

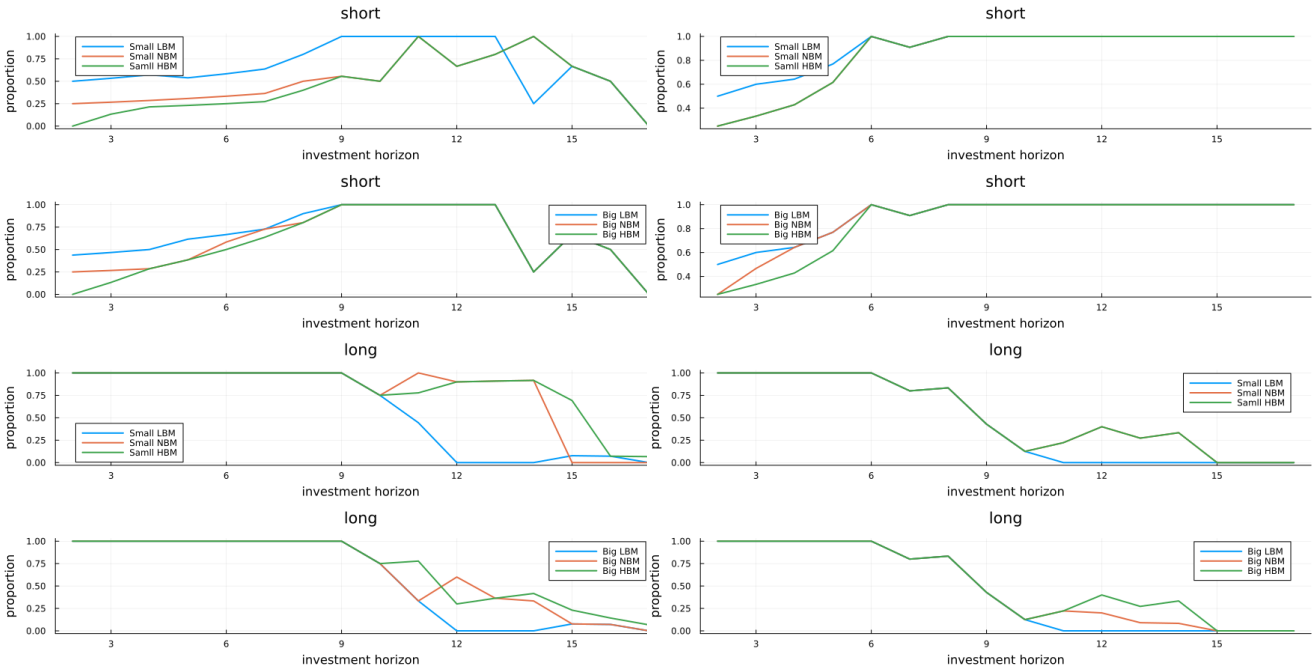
Figure 24: Short vs Long Portfolio. Three asset classes.



(a) (b)

Notes: The figures show the portfolio composition for all possible short and long horizon pairs. Panel (a) represents good times, while Panel (b) represents a recession. A subplot labeled 'short' represents each pair where the S-short horizon is on the x-axis, and the y-axis shows the proportion of each particular asset in the short portfolios among all portfolios for this S as a short horizon. The subplot labeled 'long' represents the opposite scenarios.

Figure 25: Short vs Long Portfolio. Six asset classes.



(a) (b)

Notes: The figures show the portfolio composition for all possible short and long horizon pairs. Panel (a) represents good times, while Panel (b) represents a recession. A subplot labeled 'short' represents each pair where the S-short horizon is on the x-axis, and the y-axis shows the proportion of each particular asset in the short portfolios among all portfolios for this S as a short horizon. The subplot labeled 'long' represents the opposite scenarios.

7 Conclusion

The investment horizon is an essential characteristic of investors that certainly needs to be taken into account, yet it has not received sufficient attention. Every investment decision specifies not only how much to invest but also for how long, and asset-pricing models are rarely explicit about the second choice. The main contribution of this paper is to raise this question—theoretically and empirically—and to provide a framework in which it can be asked and answered.

The starting point of the theory is the horizon function $H(\tau)$, which embeds investment horizons directly into financial models. Specifying the share of invested wealth the investor plans to use at each future date, $H(\tau)$ yields a first-order condition that is a convex combination of short- and long-horizon Euler terms rather than a single one-period equation. A benchmark result is horizon irrelevance: under random-walk dynamics for log consump-

tion and log returns, all horizon Euler terms equal one at every horizon, so the investment horizon has no effect on asset pricing. Conversely, when these random-walk assumptions are relaxed—as the empirical evidence requires—explicitly incorporating investor horizons into the model helps explain the equity premium puzzle.

The main theoretical result is the Horizon-Adjusted Fundamental Pricing Equation, which shows that the optimal investment horizon is characterized by the standard Fundamental Pricing Equation together with the requirement that the Euler term is locally maximized there. The Euler term is below one at every horizon except the optimal one—a property that can serve as a test of asset-pricing models, both theoretically and empirically. When econometricians force the one-period Euler equation to equal one, the moment condition is misspecified and the resulting parameter estimates are biased. The framework also delivers a demand-side implication: risky-asset demand, viewed as a function of the horizon, peaks at the optimal horizon, so positions are largest exactly where the pricing equation holds.

On the empirical side, this paper provides one of the first estimates of the investment horizon of the representative investor. The most likely investment horizon falls between 8 and 14 quarters and between 5 and 8 quarters during a recession; additionally, investors are unlikely to invest for one- and two-quarter horizons, regardless of risk aversion. These findings imply that the horizon ranges most commonly used in empirical finance are substantially shorter than true investment horizons—exactly the misspecification described above.

The estimated investment horizon of the representative investor has also been gradually shortening. This decline—likely attributable to the entry of less experienced household investors—suggests that financial market development and investor literacy are important determinants of the aggregate investment horizon, raising the question of how financial development and literacy programs can shape it.

The last part of the paper addresses heterogeneous investment horizon portfolios. This perspective, while novel, aligns with common sense: investors purchase different asset classes for their portfolios with varying intended holding periods. Specifically, our findings indicate that growth stocks are associated with shorter horizons, while value stocks correspond to longer horizons. Our results support the hypothesis that the investment horizon can also be a characteristic of assets (for instance, 'duration factor' or 'implied equity duration'). In a multi-investor environment, this asset-side characteristic should be matched with the investor whose horizon it fits.

From a policy perspective, investment horizons shorten during recessions and depend on financial access.[48] This means that policies encouraging retail investors to participate in

[48]The post-1992 shortening of the estimated horizon is attributable to easier stock market access for

markets—such as financial education programs or default savings plans—could shift aggregate horizons and affect asset prices. These findings open broad directions for future research, both theoretical and empirical, and may inform policy design aimed at stock markets and investor behavior.

household investors, many of whom were first-time participants with shorter horizons and lower risk aversion (see Figure 21 and Figure 18).

References

Abdellaoui, M., A. E. Attema, and H. Bleichrodt (2010). Intertemporal tradeoffs for gains and losses: An experimental measurement of discounted utility. *The Economic Journal 120*(545), 845–866.

Abel, A. B., J. C. Eberly, and S. Panageas (2007). Optimal inattention to the stock market. *American Economic Review 97*(2), 244–249.

Ameriks, J. and S. P. Zeldes (2004). *How do household portfolio shares vary with age?*

Andries, M., T. M. Eisenbach, and M. C. Schmalz (2014). Asset pricing with horizon-dependent risk aversion. Technical report, Staff Report.

Balduzzi, P. and A. W. Lynch (1999). Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics 52*(1), 47–78.

Bansal, R., S. Miller, D. Song, and A. Yaron (2021). The term structure of equity risk premia. *Journal of Financial Economics 142*(3), 1209–1228.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The journal of Finance 59*(4), 1481–1509.

Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean (2009). Just how much do individual investors lose by trading? *The Review of Financial Studies 22*(2), 609–632.

Barber, B. M. and T. Odean (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance 55*(2), 773–806.

Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance 55*(1), 225–264.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics 121*(3), 823–866.

Basak, S. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *The Review of Financial Studies 23*(8), 2970–3016.

Basak, S. and A. Pavlova (2013). Asset prices and institutional investors. *American Economic Review 103*(5), 1728–1758.

Bianchi, F., S. C. Ludvigson, and S. Ma (2022). Belief distortions and macroeconomic fluctuations. *American Economic Review 112*(7), 2269–2315.

Binsbergen, J. v., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Brainard, W. C., W. R. Nelson, and M. D. Shapiro (1991). The consumption beta explains expected returns at long horizons. *Manuscript, Economics Department, Yale University*.

Brav, A., G. M. Constantinides, and C. C. Geczy (2002). Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy 110*(4), 793–824.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of political Economy 107*(2), 205–251.

Campbell, J. Y. and R. J. Shiller (1988). Stock prices, earnings, and expected dividends. *the Journal of Finance 43*(3), 661–676.

Campbell, J. Y. and L. M. Viceira (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics 114*(2), 433–495.

Campbell, J. Y. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review 94*(5), 1249–1275.

Cassella, S., B. Golez, H. Gulen, and P. Kelly (2023). Horizon bias and the term structure of equity returns. *The Review of Financial Studies 36*(3), 1253–1288.

Chen, X., L. P. Hansen, and P. G. Hansen (2020). Robust identification of investor beliefs. *Proceedings of the National Academy of Sciences 117*(52), 33130–33140.

Choi, J. J. and A. Z. Robertson (2020). What matters to individual investors? evidence from the horse's mouth. *The Journal of Finance 75*(4), 1965–2020.

Cochrane, J. H. (1999). New facts in finance.

Cochrane, J. H. (2005). *Asset Pricing* (Revised ed.). Princeton, NJ: Princeton University Press.

Cochrane, J. H. (2022). Portfolios for long-term investors. *Review of Finance 26*(1), 1–42.

Cohn, R. A., W. G. Lewellen, R. C. Lease, and G. G. Schlarbaum (1975). Individual investor risk aversion and investment portfolio composition. *The Journal of Finance 30*(2), 605–620.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–2678.

Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of political Economy 98*(3), 519–543.

Constantinides, G. M. and D. Duffie (1996). Asset pricing with heterogeneous consumers. *Journal of Political economy 104*(2), 219–240.

Daniel, K., D. Hirshleifer, and L. Sun (2020). Short-and long-horizon behavioral factors. *The review of financial studies 33*(4), 1673–1736.

Daniel, K. and D. Marshall (1997). Equity-premium and risk-free-rate puzzles at long horizons. *Macroeconomic Dynamics 1*(2), 452–484.

Duffie, D. and T.-s. Sun (1990). Transactions costs and portfolio choice in a dynamic setting. *Journal of Economic Dynamics and Control 14*(1), 35–51.

Eichenbaum, M. S., L. P. Hansen, and K. J. Singleton (1988). A time series analysis of representative agent models of consumption and leisure choice under uncertainty. *The Quarterly Journal of Economics 103*(1), 51–78.

Eisenbach, T. M. and M. C. Schmalz (2016). Anxiety in the face of risk. *Journal of Financial Economics 121*(2), 414–426.

Fama, E. F. and K. R. French (1988a). Dividend yields and expected stock returns. *Journal of financial economics 22*(1), 3–25.

Fama, E. F. and K. R. French (1988b). Permanent and temporary components of stock prices. *Journal of Political Economy 96*(2), 246–273.

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

Fama, E. F. and K. R. French (2018). Long-horizon returns. *The Review of Asset Pricing Studies 8*(2), 232–252.

Ferreira, M. A. and P. Matos (2008). The colors of investors' money: The role of institutional investors around the world. *Journal of financial economics 88*(3), 499–533.

Gabaix, X., R. S. Koijen, F. Mainardi, S. Oh, and M. Yogo (2024). Asset demand of us households. Technical report, National Bureau of Economic Research.

Ghosh, A., T. Otsu, and G. Roussellet (2021). Properties of subjective beliefs estimators. *Available at SSRN 3784757*.

Ghosh, A. and G. Roussellet (2023). Identifying beliefs from asset prices. In *Proceedings of Paris December 2019 Finance Meeting EUROFIDAI-ESSEC*.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

Gompers, P. A. and A. Metrick (2001). Institutional investors and equity prices. *The quarterly journal of Economics 116*(1), 229–259.

Gonçalves, A. S. (2021). Reinvestment risk and the equity term structure. *The Journal of Finance 76*(5), 2153–2197.

Gormsen, N. J. (2021). Time variation of the equity term structure. *The Journal of Finance 76*(4), 1959–1999.

Gormsen, N. J. and E. Lazarus (2023). Duration-driven returns. *The Journal of Finance 78*(3), 1393–1447.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *The Review of Financial Studies 27*(3), 714–746.

Gressis, N., G. C. Philippatos, and J. Hayya (1976). Multiperiod portfolio analysis and the inefficiency of the market portfolio. *The Journal of Finance 31*(4), 1115–1126.

Handa, P., S. P. Kothari, and C. Wasley (1989). The relation between the return interval and betas: Implications for the size effect. *Journal of Financial Economics 23*(1), 79–100.

Hirshleifer, D., S. S. Lim, and S. H. Teoh (2011). Limited investor attention and stock market misreactions to accounting information. *The Review of Asset Pricing Studies 1*(1), 35–73.

Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *The Review of Financial Studies 5*(3), 357–386.

In, F., S. Kim, and R. Gençay (2011). Investment horizon effect on asset allocation between value and growth strategies. *Economic Modelling 28*(4), 1489–1497.

Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance 48*(1), 65–91.

Jensen, M. C. (1969). Risk, the pricing of capital assets, and the evaluation of investment portfolios. *The Journal of business 42*(2), 167–247.

Judd, K. L., F. Kubler, and K. Schmedders (2003). Asset trading volume with dynamically complete markets and heterogeneous agents. *The Journal of Finance 58*(5), 2203–2217.

Julliard, C. and A. Ghosh (2012). Can rare events explain the equity premium puzzle? *The Review of Financial Studies 25*(10), 3037–3076.

Kandel, S. and R. F. Stambaugh (1996). On the predictability of stock returns: an asset-allocation perspective. *The Journal of Finance 51*(2), 385–424.

Kitamura, Y., G. Tripathi, and H. Ahn (2004). Empirical likelihood-based inference in conditional moment restriction models. *Econometrica 72*(6), 1667–1714.

Kocherlakota, N. R. (1990). On tests of representative consumer asset pricing models. *Journal of Monetary Economics 26*(2), 285–304.

Kothari, S. and J. B. Warner (1997). Measuring long-horizon security price performance. *Journal of financial economics 43*(3), 301–339.

Krusell, P. and A. A. Smith (2003). Consumption-savings decisions with quasi-geometric discounting. *Econometrica 71*(1), 365–375.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Laibson, D., S. C. Lee, P. Maxted, A. Repetto, and J. Tobacman (2026). Estimating discount functions with consumption choices over the lifecycle. *The Review of Financial Studies 39*(6), 1580–1610.

Lee, C. F., C. Wu, and K. J. Wei (1990). The heterogeneous investment horizon and the capital asset pricing model: theory and implications. *Journal of Financial and Quantitative analysis 25*(3), 361–376.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *The journal of finance 62*(1), 55–92.

Levhari, D. and H. Levy (1977). The capital asset pricing model and the investment horizon. *The Review of Economics and Statistics*, 92–104.

Levy, H. (1972). Portfolio performance and the investment horizon. *Management Science 18*(12), B–645.

Levy, H. (1984). Measuring risk and performance over alternative investment horizons. *Financial Analysts Journal 40*(2), 61–68.

Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical finance 10*(3), 387–406.

Lo, A. W. and A. C. MacKinlay (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies 1*(1), 41–66.

Loewenstein, G. and D. Prelec (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics 107*(2), 573–597.

Mark, N. C. (1995). Exchange rates and fundamentals: Evidence on long-horizon predictability. *The American Economic Review*, 201–218.

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

Menkhoff, L. (2011). Are momentum traders different? implications for the momentum puzzle. *Applied Economics 43*(29), 4415–4430.

Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of Economic Theory 26*(1), 17–27.

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012). Time series momentum. *Journal of Financial Economics 104*(2), 228–250.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business 41*(2), 215–229.

O'donoghue, T. and M. Rabin (1999). Doing it now or later. *American economic review 89*(1), 103–124.

Parker, J. A. and C. Julliard (2005). Consumption risk and the cross section of expected returns. *Journal of Political Economy 113*(1), 185–222.

Pástor, L. and R. F. Stambaugh (2012). Are stocks really less volatile in the long run? *The Journal of Finance 67*(2), 431–478.

Poterba, J. M. and L. H. Summers (1988). Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics 22*(1), 27–59.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics 51*(3), 239–246.

Shiller, R. J. (2014). Speculative asset prices. *American Economic Review 104*(6), 1486–1517.

Sims, C. A. (1984). Martingale-like behavior of prices and interest rates.

Van Binsbergen, J. H. (2020). Duration-based stock valuation: Reassessing stock market performance and volatility. Technical report, National Bureau of Economic Research.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance 56*(2), 433–470.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of monetary economics 24*(3), 401–421.

A Appendix: Robustness Alternative Conditioning Sets and Test Assets

This appendix demonstrates the robustness of our results to the choice of conditioning set and test assets. We consider four conditioning sets:

1. The first conditioning set, which is used in the main part, includes 1-quarter real excess returns, 1-year nominal yield, 1-quarter consumption growth, and changes in consumption growth.

2. The second conditioning set contains 1-quarter consumption growth and 1-quarter gross real returns.

3. The third conditioning set contains 1-quarter consumption growth, 1-quarter gross real returns, and 1-quarter bond yield.

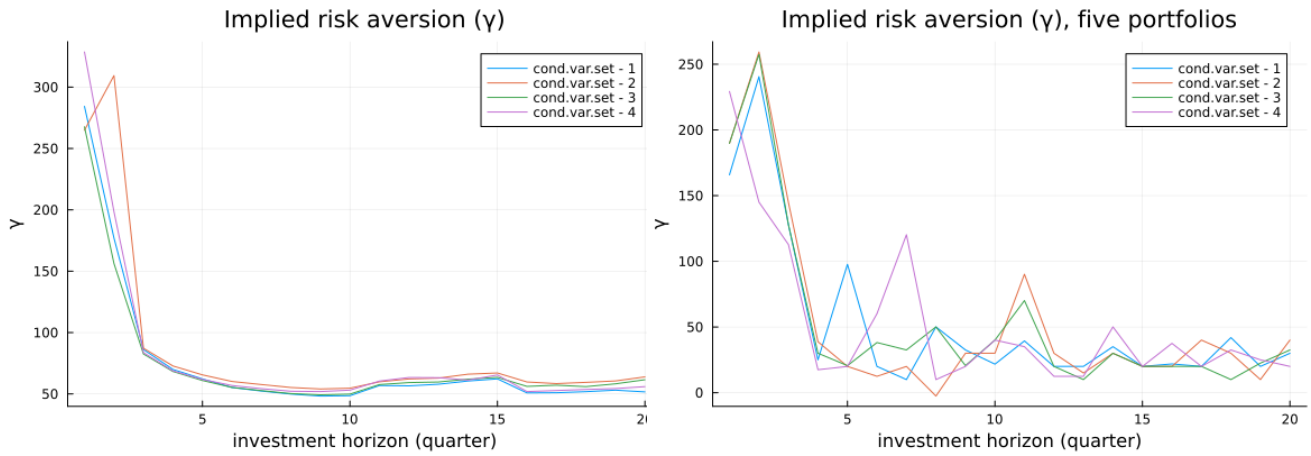
4. The fourth conditioning set contains 1-quarter consumption growth, 1-quarter real excess returns, and 1-quarter inflation.

Figure 26 shows the implied risk aversion γ for various investment horizons. Panel (a) shows results when only excess market returns are used in the Euler equation. Panel (b) presents results from five Euler equations: excess return for the total market, high and low B/M, and small and large size portfolios. The main pattern and results remain consistent regardless of the conditioning set and number of test assets.

Figures 27, 28, and 29 present comparisons of objective and subjective expected returns, standard deviations, and Sharpe ratios for various horizons. For all conditioning sets, results are similar, with conditioning sets 1 and 3 demonstrating subjective beliefs much closer to objective data.

Figures 30 and 31 show the Euler Equation error to compare whether short or long horizons are more likely. The general pattern remains consistent across all conditioning sets. Conditioning sets 2 and 4 demonstrate less volatility and do not exhibit a slight trend over the past 20 years. The reason for this is that these sets do not contain bond returns (although the 4th conditioning set includes excess returns, thus not explicitly incorporating the risk-free rate). Finally, Figure 32 presents a comparison of short and long horizon results for five portfolios separately. The results support the main findings from the second part of the paper.

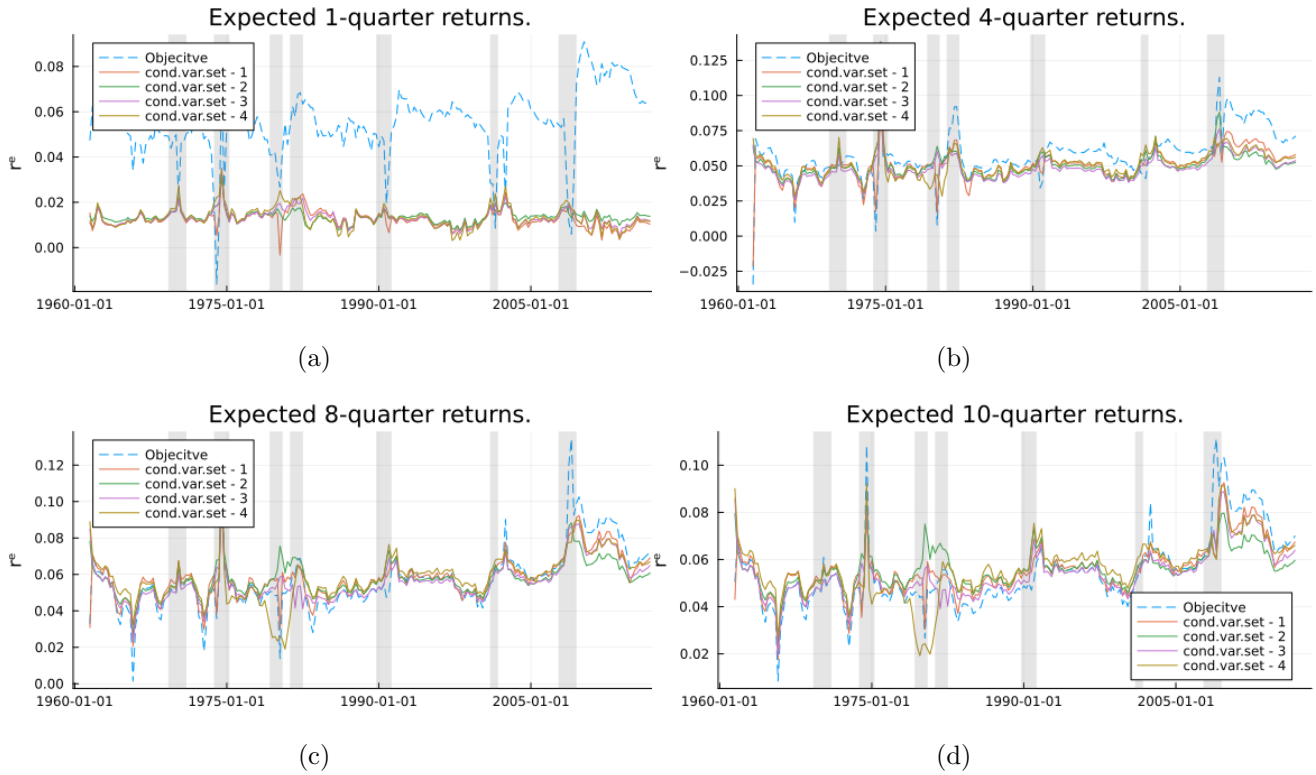
Figure 26



(a) (b)

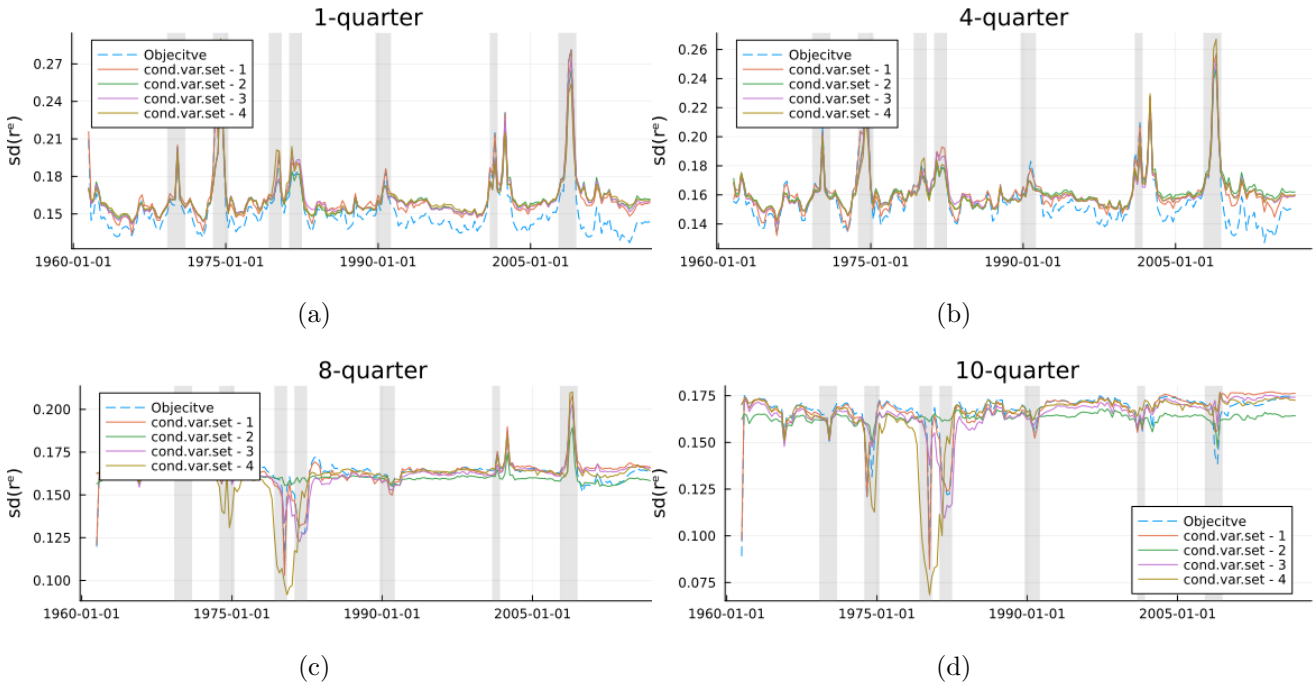
Notes: The figure plots implied risk aversion γ for various investment horizons. Panel (a) shows results when excess market returns are used in the Euler equation. Panel (b) presents results from five Euler equations: excess return for total market, high and low B/M, and small and large size portfolios. The results are obtained by minimizing SEL over γ. The pricing kernels correspond to CRRA utility. Real market returns are based on the S&P 500 index, consumption refers to real personal expenditure on nondurable goods and services, and risk-free rates are derived from deflated nominal yield curves for each quarter from 1 to 20. The sample period is 1961:3Q-2016:4Q.

Figure 27: Objective and Expected Excess Returns for Various Investment Horizons



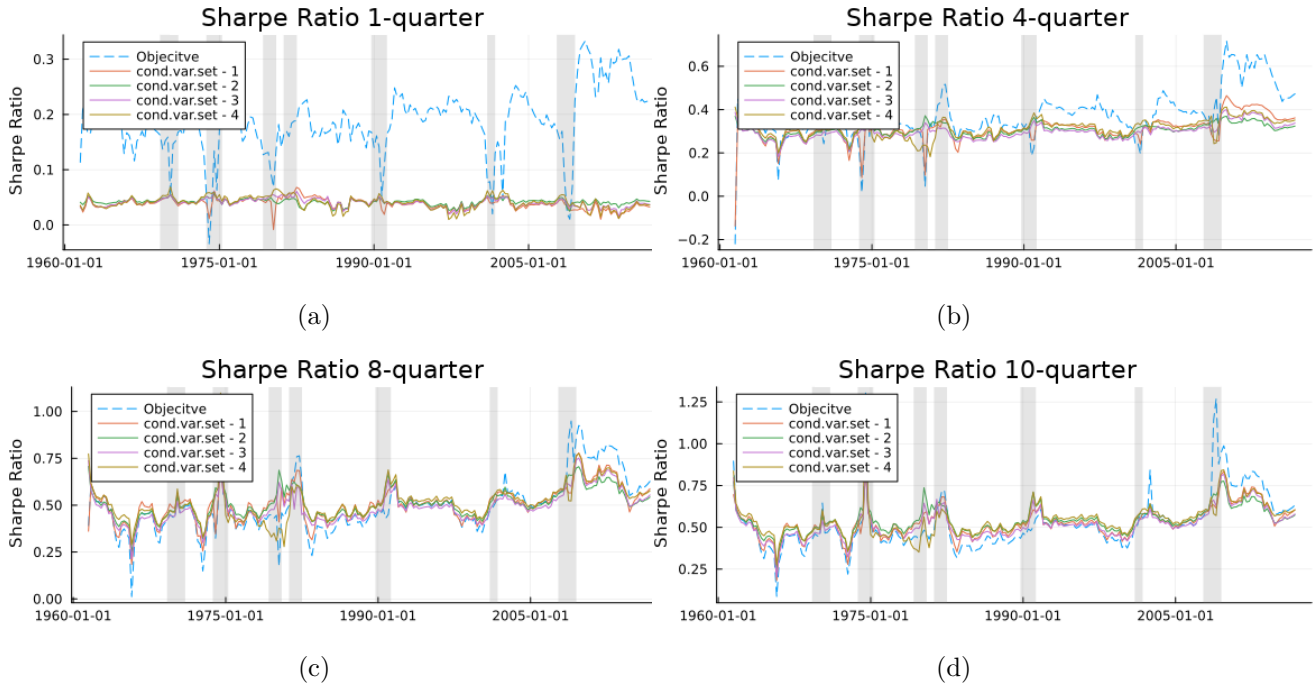
Notes: The figure plots annualized Objective and Subjective Expected Excess returns for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 28: Objective and Subjective Standard Deviation of Returns for Various Investment Horizons



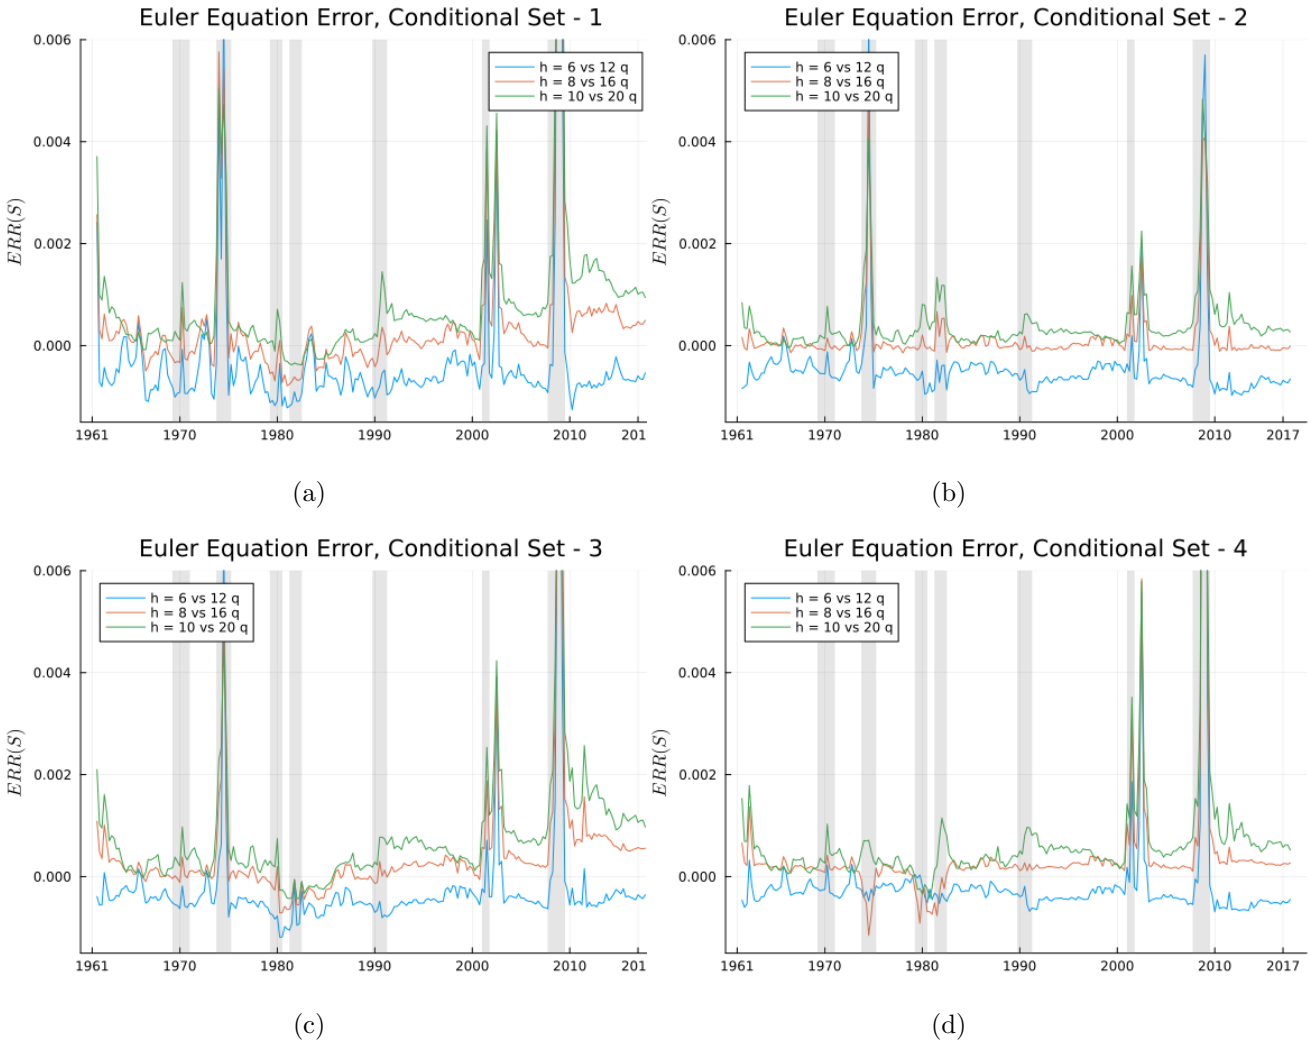
Notes: The figure plots annualized Objective and Subjective Standard Deviations of returns for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 29: Objective and Subjective Sharpe Ratio for Various Investment Horizons



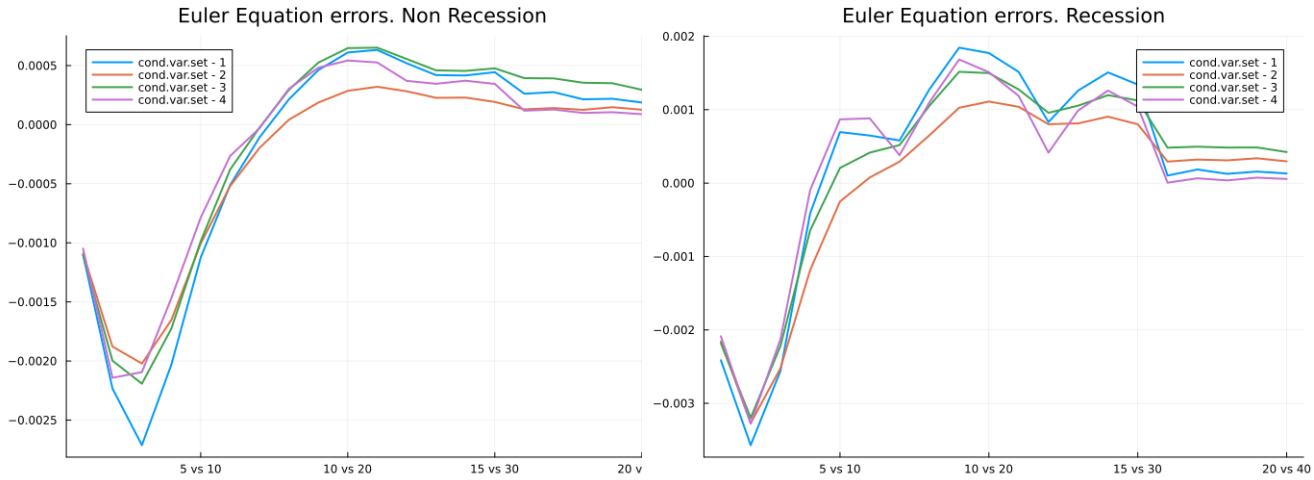
Notes: The figure plots Objective and Subjective Sharpe Ratio for various investment horizons. The pricing kernels correspond to CRRA, $\gamma = 50$. Real market returns correspond to S&P 500 index, consumption to the real personal expenditure on nondurable goods and services, and risk-free rates to the deflated nominal yield curves for each quarter from 1 to 20. Sample is 1961:3Q-2016:4Q.

Figure 30



Notes: The figure plots Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

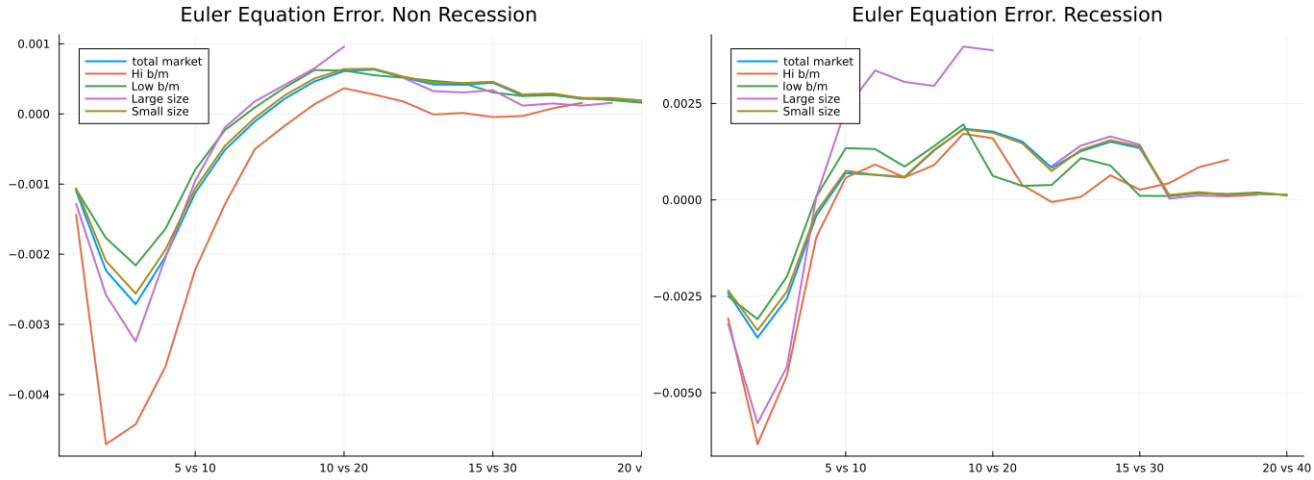
Figure 31



(a) (b)

Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

Figure 32



(a) (b)

Notes: The figure plots the mean Euler Equation errors (55) by applying short-horizon beliefs to twice the long-horizon investment. The value below zero means the long horizon is more likely to be the true investment horizon. A positive value means the opposite. The pricing kernels correspond to CRRA, $\gamma = 50$.

B Investor's investment horizon illustration figures

Figure 33

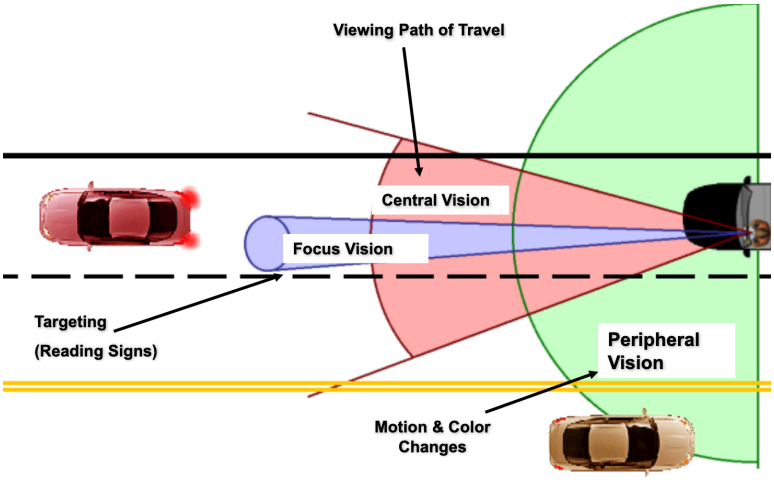
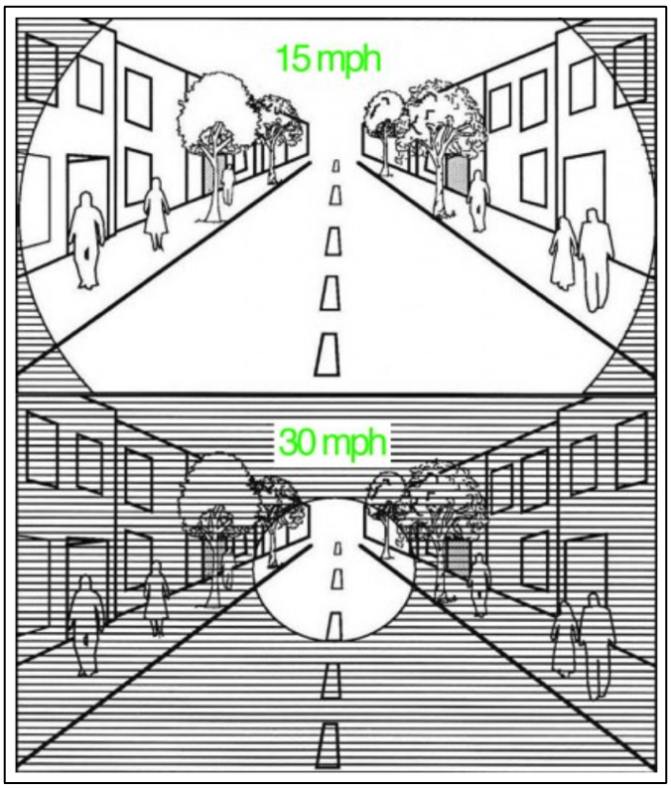
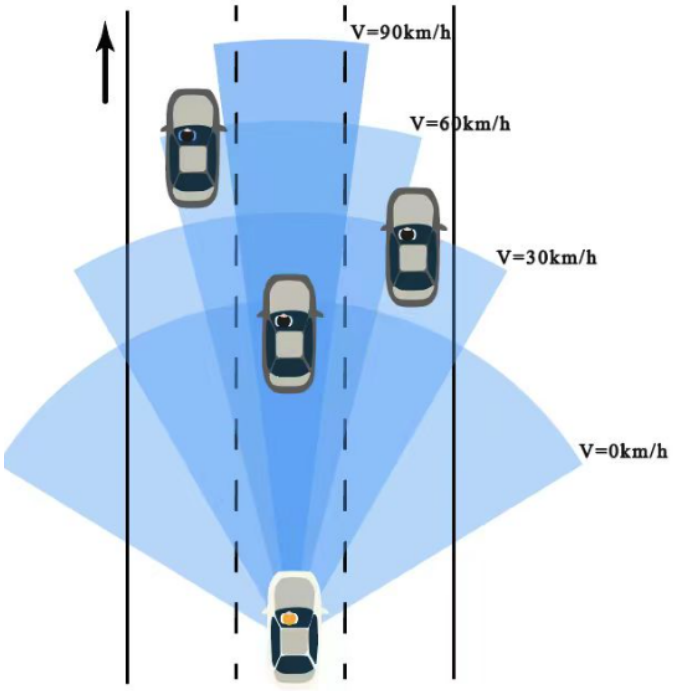


Figure 34



Notes: source "https://americawalks.org/part-2-reducing-vehicle-speeds-pedestrian-safety/"

Figure 35



Notes: source "https://arxiv.org/html/2402.04318v2"