2023

Adaptive Learning Gain in Asset Pricing

Sedealy Juste Lokossou

ADAPTIVE LEARNING GAIN IN ASSET PRICING

by

Juste Lokossou

A thesis submitted in partial fulfilment of the requirements for the

Master of Science

Major in Economics

South Dakota State University

2023

THESIS ACCEPTANCE PAGE

Juste Lokossou

This thesis is approved as a creditable and independent investigation by a candidate for the master's degree and is acceptable for meeting the thesis requirements for this degree. Acceptance of this does not imply that the conclusions reached by the candidate are necessarily the conclusions of the major department.

Zhiguang Wang
Advisor Date

Joseph Santos
Director Date

Nicole Lounsbery, PhD
Director, Graduate School Date

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor, Dr. Zhiguang Wang, for his unwavering support and guidance throughout my research journey in Finance. His passion and excellence have inspired me to tackle challenging and relevant research topics in asset pricing.

I am also grateful to my committee members, Dr. Nacasius Ujah, Dr. David Davis, and Dr. Gary Hatfield, for their invaluable insights and expertise, which have greatly enhanced the quality of this paper. Their guidance and feedback have been instrumental in shaping this research.

I am deeply appreciative to my mentors, Elizabeth Asiedu and Kagba Kousse, for their valuable insights and advice. I would like to extend my love and gratitude to my parents, Rene and Cica, whose unwavering support and encouragement have been instrumental in my academic pursuits.

I am eternally grateful for the extensive resources and funding provided by the Ness School of Management and Economics, including funding for the valuable data science program that have enhanced my research.

Lastly, I would like to express my appreciation to all the faculty and administration members who have contributed to my growth and learning. Your support and dedication will have a lasting impact on my academic and personal development.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

ADAPTIVE LEARNING GAIN IN ASSET PRICING

JUSTE LOKOSSOU

2023

This paper delves into the complexities of asset pricing, emphasizing the need to go beyond prevailing paradigms and constant learning gain assumptions. We examine the influence of personal experiences, adaptive learning processes, and subjective return expectations on asset pricing. By incorporating the concept of time-varying learning gain, we provide a more realistic portrayal of asset pricing.

Empirical analysis reveals a consistent negative correlation between experienced real payout growth and subsequent returns, indicating counter-cyclical behavior. Our findings also support the mean-reversion hypothesis in stock returns, although caution is needed due to some scenarios lacking statistical significance.

Theoretical exploration uncovers that higher uncertainty or variability compels investors to seek additional compensation, thus elevating the equity risk premium. Moreover, the information structure does not form a filtration, leading to no convergence to a specific value in the long run. Agents perceive future increments as negatively serially correlated but lack the memory to effectively exploit this correlation for forecasting. Consequently, the Law of Iterated Expectations does not hold. We propose the "resale" valuation method as ideal for agents with adaptive learning gains.

These findings contribute to an innovative asset pricing model with adaptive learning gains, enhancing our understanding of market dynamics. While this study does not provide calibration or validation, we outline the model's theoretical foundations and implications for

future research. Our work adds to the evolving landscape of asset pricing theory, highlighting the significance of adaptive learning in capturing complex dynamics.

Chapter 1

INTRODUCTION

The 1970s and 2008 economic crises have propelled behavioral approaches into the mainstream, driven by the lack of viable alternatives Malmendier and Wachter (2021). While it is a fact that early critiques of utility theory and rational expectations had little impact on the dominant status of these models, behavioral models gained traction due to empirical evidence challenging Bayesian updating and rational expectations.

The notion of rationality holds a crucial position in economics, particularly in asset pricing. It assumes that individuals, as rational agents, are mean-variance optimiser, possess complete knowledge of data-generating probabilities and have perfect information about all relevant parameters, including the utility function. However, this idealized assumption can be seen as delusional, as it overlooks cognitive biases, emotional factors (fear, greed, overconfidence, loss aversion, herding behavior, anchoring bias, regret aversion), and imperfect information that affect decision-making. Recognizing these deviations from rationality is vital for understanding the complexities of asset pricing and developing more accurate economic models. By acknowledging the limitations individuals face in acquiring and processing information, we gain a more nuanced understanding of decision-making behavior that goes beyond strict notions of rationality.

Although the diagnostic is effective, and novel studies such as those by Nagel and Xu (2019) and Wang (2021) have attempted to incorporate lifetime experience through learning gains in asset pricing models, to our knowledge, none of these models have yet considered an analytical formalism of a subjective model with adaptive learning gain. All the proposed models have used a constant learning gain, typically 0.018, as a proxy for the learning dynamics of individual experience. While this process aligns with the psychological concept of availability bias, which suggests that individuals are biased towards recent or readily available information, we consider this approach to add an additional layer of assumption to the problem at hand. In other words, we believe that utilizing the full potential of the learning gain series is likely to reveal the true dynamics of pricing an asset beyond the confines of assumed constraints.

A central concern is that some policy rules can lead to indeterminacy of equilibria, resulting in multiple rational expectations (RE) solutions, a notion supported by Bernanke and Woodford (1997), Woodford (1999), and Svensson and Woodford (2003). Another key problem is the questionable performance of these rules if private agents follow adaptive learning rules rather than rational expectations. Bullard and Mitra (2000) have noted that the stability of Taylor-type rules cannot be assumed if agents follow adaptive learning rules, extending Howitt (1992) warning about the potential instability of interest rate pegging and related rules under learning conditions.

The collective implications of the studies by Campbell and Cochrane (1999), Barberis et al. (2001) , Evans and Honkapohja (2001), Malmendier and Nagel (2016), Nagel and Xu (2019), and Wang (2021) have paved the way for this study to offer a more holistic approach to understanding asset pricing dynamics. This is achieved by considering the importance of individual's personal experiences, adaptive learning processes, and their subjective return expectations in influencing asset pricing.

Some studies such as Wang (2021) have primarily focused on a set of market participants such as analysts. This paper while using common shares and stocks traded on NYSE, AMEX, and NASDAQ is more inclusive, considering a broader set of market participants. This distinction provides a richer context to our study, enabling a more comprehensive understanding of asset pricing dynamics, accounting for the diversity of market participants.

Our research addresses a gap in the existing literature by incorporating the concept of adaptive learning gain into asset pricing dynamics. This variation adds an additional layer of complexity to the learning behavior of agents, thus providing a more realistic representation of asset pricing. Additionally, by considering a broader spectrum of market participants, our study allows for a deeper understanding of market dynamics.

This paper unfolds over seven chapters. In the first, we introduce the global objective of this research paper. In the second chapter, we survey the established literature on subjective beliefs, highlighting unexplored areas in the field. In the subsequent section, we explore the process of deriving the adaptive learning gains. Later, in chapter four, we delve into empirical hypotheses concerning asset pricing with adaptive learning gains.

In the fifth chapter, our focus shifts to the formulation of a new asset pricing model, which leverages adaptive learning gains. We articulate its theoretical basis and implications, while noting that this work doesn't extend to the model's calibration or validation - an essential task we reserve for a future study.

Moving forward, in chapter six, we present a discussion of our findings and their implications, underscoring how our adaptable learning gains model can contribute to the evolving landscape of asset pricing theory. The paper concludes with chapter seven where we present the appendix.

Chapter 2

RELATED LITERATURE

In the realm of finance, asset pricing stands as a dynamic and forward-thinking process, wielding significance for both individual investors and institutions alike. Its fundamental purpose lies in deciphering the drivers that propel asset prices. An amalgamation of factors, such as risk levels, anticipated returns, and the timing of their realization, interplay in the valuation of assets. Additionally, market conditions, available information, and regulatory influences exert their force, further shaping the price dynamics.

Moving forward, the perpetual discussion in asset pricing centers around equity risk premiums' cyclical behavior. Models like the Capital Asset Pricing Model (CAPM) and the Black-Scholes model typically assume investor rationality. However, Campbell and Cochrane (1999) introduced a time-varying 'habit' into the consumption-based asset pricing model, proposing that during economic downturns, as consumption nears this habit, the decrease in asset prices and increase in expected returns occur. Interestingly, although their approach doesn't fully embrace behavioral finance, it does serve as a transition from purely rational models to a more behavioral understanding of asset pricing.

Similarly, in their 2001 paper, Barberis, Huang, and Santos critique the traditional consumption-based model for explaining aggregate stock market behavior, stating it does not fully account for observed high market volatility, high average returns, and low correlation with consumption growth. Instead, they propose a model where investors derive utility from fluctuations in their financial wealth as well as from consumption. This model incorporates two crucial aspects: loss aversion, meaning investors are more sensitive to losses than gains, and a dependency of this loss aversion on prior investment performance.

Continuing on this theme, the introduction of variable risk aversion, influenced by past market movements, helps to generate high mean, high volatility, and low correlation of stock returns with consumption growth, while maintaining a low, stable riskless interest rate. It also explains the significant equity premium required to convince the loss-averse investor to hold stocks. This perspective, influenced by Kahneman and Tversky's prospect theory and psychological literature on risk-taking behavior is tightly linked to the position we defend in this paper as the authors stress that while loss aversion is vital in explaining the equity premium, the impact of prior outcomes on risk aversion also plays a critical role. Crucially, ignoring these prior outcomes would diminish an important source of stock price volatility, leading to less risk and a lower equity premium. We believe that agents update their future belief upon the formation of their priors.

On another front, Evans and Honkapohja (2001) challenge the oversimplification of rational expectations in monetary policy. They propose a rule that incorporates private sector expectations and economic structure, ensuring stability and convergence to rational expectations. Their findings support considering various indicators, given the potential destabilizing impact of deviations from rational expectations. The findings ofEvans and Honkapohja (2001) establish a direct connection with our study. Their emphasis on the role of adaptive learning in economic behavior and monetary policy formulation provides a robust theoretical foundation for our exploration of heterogeneity in the agents' expectations, supporting the notion that individuals' inflation expectations are significantly shaped by their personal experiences and adaptive learning processes.

Additionally, it is fair to say that our analysis of how agents perceive how asset pricing evolves heavily borrow from Malmendier and Nagel (2016) study that investigates how individuals form their expectations about future inflation, a subject critical to monetary policy and financial decisions. The authors propose a model where individuals' expectations are heavily influenced by personal experiences, particularly inflation rates they have encountered during their lifetime.

Following this line of thought, the model is built on adaptive learning algorithms, but with modifications to account for "learning from experience." That is, inflation experiences have more influence on people's forecasts than other historical data. Younger individuals react more strongly to unexpected inflation due to their shorter history of experiences, leading to differing inflation outlooks across generations. On the contrary to our position, the study posits that learning from experience provides a microfoundation for constant-gain learning and offers an alternative explanation for why older data is often disregarded - a phenomenon usually attributed to structural shifts and parameter drift. This process aligns with the psychological concept of availability bias, where individuals are biased towards information that is recent or readily available. However, as shown by Wang (2021) quoted by Malmendier and Wachter (2021) with the first model of memory in economics, “the long-term effects of past experiences are front and center of the model setup”.

Progressing further, building on Malmendier and Nagel (2011) and Malmendier and Nagel (2016), Nagel and Xu (2019) offer empirical findings, introducing novel insights about equity market returns and subjective stock return expectations. They propose a reduced-form framework incorporating constant-gain learning about dividend growth rates, leading to equity premium being counter-cyclical under objective expectations. The belief-updating rule's gain parameter significantly influences the volatility, persistence of the price-dividend ratio, and return predictability strength. In the final analysis, Malmendier and Nagel's research indicates a significant role for "learning from experience" in shaping expectations, providing a new perspective to understand the dispersion in inflation expectations documented in previous studies.

In contrast toNagel and Xu (2019), our study builds on the same theoretical framework but introduces a time-varying learning gain instead of a constant one. This significant departure from their approach potentially provides a more nuanced understanding of asset pricing dynamics. Our model, by allowing the learning gain to evolve over time, captures changes in agents' learning behavior more accurately, thus representing real-world asset pricing more realistically. This work enhances the existing literature by offering a more flexible approach to learning models in asset pricing.

Lastly, Wang (2021) examines subjective return expectations and their impact on investment decisions. The study reveals differences between objective and subjective expectations, with some investors being extrapolative and others contrarian. Extrapolative investors form their expectations about future asset prices based on recent price trends while contrarian investors tend to go against the prevailing market trends. By incorporating imperfect predictors and the Kalman Filter, the author explains the heterogeneity in return expectations. Understanding the role of subjective beliefs has implications for investor behavior and model refinement. Recognizing this heterogeneity has provided insight into different patterns of market behavior and help us refine our model. In the subsequent chapter, we delve into the intricacies of deriving the adaptive learning gains, shedding light on the process that underpins the core of our research.

Chapter 3

ADAPTIVE LEARNING GAINS

In the context of this research, adaptive learning refers to a dynamic model of economic behavior where agents form expectations based on a continually updated set of information, and consequently, revise their predictions or decisions over time in response to new data and outcomes. This approach contrasts with models assuming rational expectations, where agents are assumed to know the true model of the economy and make decisions accordingly. Here, adaptive learning acknowledges that investors do not possess perfect foresight and are constantly learning from their experiences and the evolving economic environment. They adjust their expectations and behavior based on the observed performance of assets, market trends, and economic indicators, thus creating a feedback loop between their expectations and actual market outcomes.

Adaptive learning thereby allows for a more realistic representation of investor behavior in asset pricing models, capturing the dynamism and uncertainty inherent in financial markets. Furthermore, it refers to the weight or significance investors assign to new information. It reflects how quickly or slowly they alter their expectations in response to changing market conditions, hence playing a critical role in the formation of asset prices.

Nagel and Xu (2019), presents a different approach to modeling learning behavior in financial markets when perfect rationality is not assumed. In their framework, the learning gain is assumed to be constant, implying that investors assign a fixed weight to new information regardless of market conditions or their past experiences. This model simplifies the learning process, allowing for more tractable computations and predictions. However, it presupposes that investors' learning adaptability is invariant over time, which may not always hold true in reality. In contrast, the "adaptive learning gain" model in this project recognizes that the weight investors assign to new information can fluctuate over time. This variation depends on numerous factors, including the nature of the information, its relevance to the investor's portfolio, the investor's confidence in their existing predictions, and the general economic environment.

As such, the adaptive learning gain model is more responsive and flexible, offering a more realistic depiction of investor behavior. It allows for variations in learning gain based on both the external economic environment and the individual's internal decision-making process, hence providing a more nuanced understanding of asset price dynamics.

Adaptive learning gains form the foundation for the rest of this master's thesis, making it a cornerstone endeavor. In Figure 3.1, we visualize the evolution of the adaptive learning gain over time, It contrasts actual average survey expectations with the predictions derived from three distinct forecasting approaches: Learning-from-Experience Forecasts, Constant-Gain Forecasts, and Sticky Information Forecasts. This comparative perspective illustrates the dynamic nature of adaptive learning gain and underscores its potential in contributing to more nuanced asset pricing models. Its examination reveals the pronounced volatility of the learning gain time series. This evidence points towards the complexity and dynamism inherent in the learning process. Consequently, the proposition by Nagel and Xu (2019) to collapse this extensive dynamism into a singular value of 0.018 appears overly simplistic and potentially detached from the observed reality. Thus, our task in this chapter is to explore a more nuanced approach that captures the inherent fluctuations and idiosyncrasies of adaptive learning gains.

In this research, our derivation of the learning gains is heavily influenced by the framework established in Malmendier and Nagel (2016). In their work, the authors model the perceived

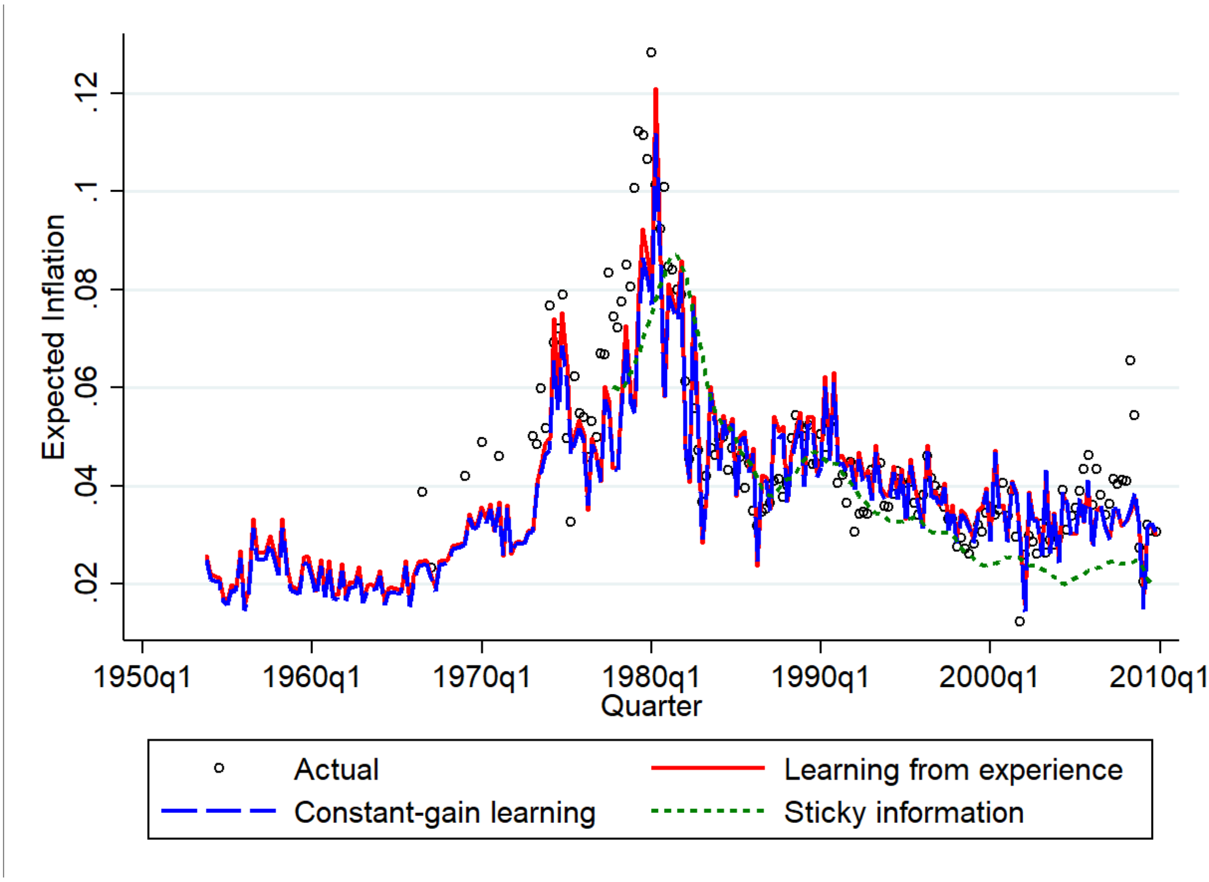

Figure 3.1: Learning Gain Over Time. This is reproduced from the data and algorithm shared by Nagel and Xu (2019).

law of motion that individuals are attempting to estimate as an autoregressive process of order one (AR(1)), given by:

$$\pi_{t+1} = \alpha + \phi\pi_t + \eta_{t+1}. \tag{3.1}$$

This process represents the inflation rate at time $t+1$ as a function of a constant term α, the previous period's inflation rate π_t weighted by a coefficient ϕ, and a random error term η_{t+1}.

In order to estimate the parameters $b \equiv (\alpha, \phi)'$, individuals use a recursive algorithm based on past data, according to the following rules:

$$b_{t,s} = b_{t-1,s} + \gamma_{t,s} R_{t,s}^{-1} x_{t-1}(\pi_t - b'_{t-1,s} x_{t-1}), \tag{3.2}$$

$$R_{t,s} = R_{t-1,s} + \gamma_{t,s}(x_{t-1}x'_{t-1} - R_{t-1,s}) \tag{3.3}$$

where $x_t \equiv (1, \pi_t)'$.

- R represents a covariance matrix used in the estimation process. It evolves over time based on past data and the learning gains (γ). The matrix R reflects the uncertainty or volatility in the parameter estimates.
- x_{t-1} is a vector containing the relevant variables used for estimation at time $t-1$. In this case, it is a 2-dimensional vector containing a constant term (1) and the inflation rate at time $t-1$ (π_{t-1}).
- $b_{t,s}$ is a vector containing the estimated parameters at time t for cohort s. In this case, it is a 2-dimensional vector containing the estimates for the constant term (α) and the coefficient for the previous period's inflation rate (ϕ). The subscript t indicates the time period, and the subscript s represents the cohort or group of individuals.

Equations 3.2 and 3.3 describe a recursive learning process where individuals use past data to update their parameter estimates for the autoregressive inflation process (3.1). The learning gains (γ) determine how much weight individuals assign to the current inflation surprise in updating their parameter estimates.

The recursion starts at some point in the distant past, and the sequence of gains $\gamma_{t,s}$ determines how much each cohort s updates their parameter estimates in response to an inflation surprise at time t. R represents the covariance matrix of the parameter estimates b, which is being updated over time as new data becomes available.

The learning gain $\gamma_{t,s}$, represents a departure from the standard adaptive learning framework, where it typically decreases over time as more data becomes available, implying that agents' expectations become more stable. However, in our model, following Malmendier and Nagel (2016), we let the gain γ depend on the age $t-s$ of the members of cohort s, leading to a decreasing-gain specification as follows:

$$\gamma_{t,s} = \begin{cases} \frac{\theta}{t-s} & \text{if } t-s \geq \theta \\ 1 & \text{if } t-s < \theta, \end{cases} \tag{3.4}$$

Here, θ is a parameter determining the rate of decrease of the learning gain. This specification captures the idea that older members of the cohort, having seen more inflation cycles, are less likely to significantly change their beliefs in response to new data. This approach provides a robust framework for analyzing the dynamics of learning and expectation formation in the context of inflation dynamics.

In our research, we propose to utilize a time-adaptable learning gain parameter in the adaptive learning framework, a modification that promises to shed more light on how learning behavior evolves over time and how it might adapt to the changing macroeconomic environment. The learning gain parameter in our model, denoted as ν_t, adapts with time according to the relation

$$\nu_t = \nu_1 + (\nu_2 - \nu_1)g(t), \tag{3.5}$$

where ν_1 and ν_2 are derived from a unique process.

In this equation, ν_1 and ν_2 represent the lower and upper bound of the learning gain, respectively. They are computed as the minimum and maximum of the estimated gain parameters (γ_1 and γ_2) obtained from an econometric analysis of historical inflation data. To ensure robustness and eliminate outliers, the gain parameters are estimated within rolling windows of data, with structural breaks identified through a breakpoint analysis. The parameter $g(t)$ in the equation is a time function, which may capture various temporal influences on the learning behavior, including demographic shifts, structural changes in the economy, or policy regime changes. Specifically, we use a linear regression model with breakpoints in order to detect structural breaks within our time series data. The analysis was focused on the learning gain time series, which was calculated using the learning gain algorithm from Nagel and Xu (2019), spanning from 1953 to 2009, resulting in 11250 observations. The Chow Test served as the technique underpinning our analysis.

In the linear regression model, the learning gain time series functioned as the response variable, with a constant term serving as the explanatory variable. Upon applying this technique, we arrived at values of $\nu_1 = 0.01858007$ and $\nu_2 = 0.08162585$. These values are used throughout the paper. Appendix F shows in detail the calculation of these values.

The introduction of this adaptive learning gain is an innovative addition to the adaptive learning framework. It provides a tool to capture possible non-linearity and time-dependent shifts in learning behavior, thereby enhancing the richness and the accuracy of the model's predictions. Furthermore, it allows for a more nuanced understanding of how different cohorts adapt their expectations to changing economic conditions. As such, our approach contributes to the current economic literature (Malmendier and Nagel (2011, 2016), Nagel and Xu (2019)) by incorporating more flexibility into the learning process, enhancing our understanding of expectation formation, and ultimately leading to more accurate macroeconomic predictions.

Chapter 4

EMPIRICAL REGRESSIONS

In this chapter, we turn our attention to a crucial empirical question—whether the proposed adaptive model indeed replicates the counter-cyclical behavior of the equity premium, a phenomenon consistently observed in the field of asset pricing. In addition, we consider other empirical regularities typically manifested in asset pricing. Undertaking this empirical assessment is of paramount importance, as it helps establish the veracity and applicability of our adaptive model. If it indeed demonstrates fidelity to these observed patterns, then our model could serve as a viable candidate in elucidating the mechanisms through which an economic agent prices assets. Consequently, the findings from this chapter will play an instrumental role in determining the robustness and relevance of the adaptive learning gain framework in the broader asset pricing landscape.

4.1 Return predictability

Let's assume that investors are learning about the mean growth rate μ_d of log real stock payouts, d:

$$\Delta d_t = \mu_d + \epsilon_t, \quad \text{epsilon is an IID shock} \tag{4.1}$$

As individuals in different birth cohorts learn from their life-time experience, their average belief is described by the two inflation points structure of constant gain learning rule. Thus we consider that the perceived growth rate $\tilde{\mu}_d$ evolves as:

$$\tilde{\mu}_{d,t+1} = \tilde{\mu}_{d,t} + [\nu_1 + (\nu_2 - \nu_1)g(t)](\Delta d_{t+1} - \tilde{\mu}_{d,t}) \tag{4.2}$$

where ν_1 and ν_2 represent respectively the two inflation points of the gain parameters. Just as in Nagel and Xu (2019), we can see that the perceived growth rate $\tilde{\mu}_d$ is updated every period with the observed surprise $\Delta d_{t+1} - \widetilde{\mu}_{d,t}$.

By applying the Campbell-Shiller Decomposition in Campbell and Shiller (1988) to the return innovations along with investors' subjective expectations, denoted $\widetilde{E}$ we obtain

$$\widetilde{E}_t r_{t+1} - r_f = \theta + [\frac{\rho}{1-\rho}(\nu_1 + (\nu_2 - \nu_1)g(t)) + 1](\Delta d_{t+1} - \tilde{\mu}_{d,t}) \tag{4.3}$$

as $g(t)$ is deterministic at every time t. where the term in parentheses times Delta $d_{t+1} - \tilde{\mu}_{d,t}$ represents the subjective growth-rate expectations revision that the econometrician anticipates, on average, in the next period, given her knowledge of $\mu_{d,t}$. The parameter ρ represents the coefficient of risk aversion in the investor's utility function. It quantifies how sensitive an investor is to risk and influences the impact of investor expectations on asset pricing dynamics, specifically in the context of the Campbell-Shiller decomposition. It appears clear that the returns is predictable by the observable surprise.

Having established in Section 4.1 the theoretical predictability of returns based on observable surprises, we now turn to Section 4.2 to put these theoretical underpinnings to the test. We aim to empirically assess the robustness of our adaptive learning gain asset pricing model, particularly in light of experienced payout growth, serving as a proxy for the broader concept of dividend surprises.

4.2 Model estimation

To verify whether the equation derived previously holds, we estimated equations 4.4 and 4.5. To be precise, we estimate a simple version of it, focusing only on experienced payout growth and not dividend surprise per se. By doing so, we empirically verify whether our adaptive learning gain asset pricing model is consistent with the counter-cyclical behavior of asset return or not.

$$\begin{aligned}\text{(Excess return)} = \alpha_0 + \alpha_1 \times \text{(Experienced payout growth)} \\ + \alpha_2 \times \text{(Inflation)} + \alpha_3 \times \text{(Price-dividend ratio)}\end{aligned} \tag{4.4}$$

$$\begin{aligned}\text{(Experienced returns)} = \beta_0 + \beta_1 \times \text{(Experienced payout growth)} \\ + \beta_2 \times \text{(Inflation)} + \beta_3 \times \text{(Price-dividend ratio)}\end{aligned} \tag{4.5}$$

Overall, equations 4.2 and 4.3 elucidate the role of adaptive learning in the dynamics of subjective dividends and the equity risk premium. Meanwhile, equations 4.4 and 4.5 assess the consistency of the adaptive learning model with well-established empirical paradigms in asset pricing, such as the counter-cyclical behavior of the equity premium and the mean reversion of asset returns.

4.3 Data and variable computations

In this project, we draw upon a rich dataset from the Center for Research in Security Prices (CRSP), featuring 98,253,046 daily observations (from 1926 to 2020) such as stock price, shares outstanding, returns with and without dividends, various adjustment factors, industry classification factors, and more. Additionally, we utilize the Consumer Price Index and 3-Month Treasury bill rate sourced from the Federal Reserve Bank Of St. Louis website.

However, the need for a more macroscopic lens prompts us to transition from these granular daily stock observations to a broader, quarterly level of analysis. This move allows us to capture the performance and behavior of stocks across distinct timeframes, thereby facilitating the identification of larger macroeconomic trends and market-wide patterns. Consequently, after cleaning the data, creating variables of interest, and aggregating the dataset, we are left with a final sample of 376 quarterly observations.

Moreover, it's worth noting that the presence of missing values in the 3-month Treasury bill time series has slightly compromised the completeness of our data. This has resulted in the final number of observations in subsequent regression analyses being relatively different, and somewhat reduced from one regression to the other. We have made every effort to ensure the robustness and validity of our analysis despite these limitations.

4.3.1 Definition of the variables

Table 4.1 provides an overview of the variable definitions used in the empirical part of this paper.

Table 4.1: Variable Definitions

Variable	Denomination	Explanation
Dependent variables	Excess return	Log return of the CRSP value-weighted index in quarter t+1, in excess of the return on a 3-month T-bill
	Experienced returns	An exponentially weighted average of quarterly log stock market index returns
Main predictor	Experienced payout growth	Exponentially weighted average of overlapping quarterly observations of four-quarter per-capita repurchase-adjusted real dividend growth rates up to and including quarter t
Control variables	Inflation	The average log CPI inflation rate during the four quarters $t-3$ to t
	Price-dividend ratio	Log price-dividend ratio of the CRSP value-weighted index at the end of quarter t

4.3.2 Adjusting for stock-buyback

Portfolio Dividends

From the CRSP dataset, we extract dividend payments from the value-weighted portfolio. Following Bansal et al. (2005), we denote the total return per dollar invested as

$$R_{t+1} = h_{t+1} + y_{t+1}, \tag{4.6}$$

where h_{t+1} denotes the price gain and y_{t+1} the dividend yield. To be precise, this latter refers to dividends at date $t+1$ per dollar invested at date t. In CRSP dataset, R_{t+1} refers to the the value-weighted return including dividends and h_{t+1} the price appreciation refers to the the value-weighted return excluding dividends. Therefore, $y_{t+1} = R_{t+1} - h_{t+1}$. From the dividend yield, we were able to compute the level of the dividends as follows:

$$D_{t+1} = y_{t+1} V_t, \tag{4.7}$$

where

$$V_{t+1} = h_{t+1} V_t, \tag{4.8}$$

with $V_0 = 1$. As recalled by Bansal et al. (2005), equation 4.7 suggests that the dividend series D_t corresponds to the total cash dividends given out by a mutual fund at t that extracts the dividends and reinvests the capital gains. The ex-dividends value of the mutual fund is V_t and the per dollar return available for the investors in the mutual fund is

$$R_{t+1} = \frac{V_{t+1} + D_{t+1}}{V_t} = h_{t+1} + y_{t+1}, \tag{4.9}$$

Equation (**??**) clearly suggests that V_t is the discounted value of the dividends that we use in this paper.

Dividends and Repurchases

Following Bansal et al. (2005) and Nagel and Xu (2019), we account for distortion from stocks repurchases that significantly affect the number of shares outstanding and thus dividend payout.

The inclusion of share buybacks in the calculation of dividends can create a skewed perception of a company's regular dividend policy. Since the 1980s, many public companies have increasingly used share buybacks as a means of returning capital to shareholders. While this is an important part of a company's capital distribution strategy, it represents a distinct approach from the regular distribution of dividends. Share buybacks often reflect a company's broader financial strategy and market conditions, rather than its ongoing operational profitability. Therefore, adjusting dividends to account for share buybacks allows us to better isolate and understand the company's regular dividend policy. This adjustment provides a more accurate picture of a company's recurring income distribution, separate from strategic capital actions such as buybacks.

The following describes the approach for estimating this alternative measure of payouts to equity shareholders. Let n_t denote the number of shares (after adjusting for stock dividends, splits, ets. using the CRSP share adjustment factor). For a given firm in the CRSP dataset, we construct the adjusted capital gain series as follows:

$$h^*_{t+1} = \left[\frac{P_{t+1}}{P_t}\right] min\left[\left(\frac{n_{t+1}}{n_t}\right), 1\right] \tag{4.10}$$

4.4 Econometric results

In this analysis, we assess the predictability of the log return of the CRSP value-weighted index in quarter t+1, in excess of the return on a 3-month T-bill, through two distinct lenses, namely: experienced real payout growth and experienced returns.

Panel A: Experienced Payout Growth

For the five periods under consideration, experienced payout growth consistently negatively correlates with subsequent returns. This suggests a possible counter-cyclical behavior in which periods of heightened growth are followed by reduced returns. This is an expected outcome given the pro-cyclical nature of dividends and the tendency of companies to distribute more dividends during prosperous times, which could result in diminished future returns. This effect holds true for all time periods, with the p-values indicating statistical significance in all but the 1927-2020 scenario in column (3).

Interestingly, the effect seems to be more pronounced in the more recent periods from 1980-2020, which could reflect the evolving market dynamics such as share buyback and the increasing importance of dividends as a predictor of returns. This finding is consistent with existing literature (Michaely et al. (1995), Jegadeesh and Titman (2001), Fama and French (2001), Baker and Wurgler (2004), and Baker and Wurgler (2007)) among others). Inflation and price-dividend ratio have been included as control variables, though they don't appear to exert a significant influence on the predicted returns in most cases, as evidenced by the higher p-values.

Panel B: Experienced Returns

Panel B considers experienced returns, defined as an exponentially weighted average of quarterly log stock market index returns. This variable also shows a consistent negative relationship with future returns, with the degree of the relationship seemingly more substantial than that of experienced payout growth.

This finding supports the mean-reversion hypothesis in stock market returns, which suggests that high past returns are often followed by lower future returns and vice versa (Fama and French (1988), Campbell and Shiller (1988), Jegadeesh and Titman (1993), Lo and MacKinlay (2011), and Conrad and Kaul (1998)). However, the p-values indicate a lack of statistical significance in some scenarios, calling for cautious interpretation of the results. The inflation and price-dividend ratio again serve as control variables. Their influence remains mostly non-significant, similar to the observations from Panel A.

Implications for Asset Pricing

These results have significant implications for asset pricing models. They provide empirical support for the use of adaptive learning measures, such as experienced payout growth and experienced returns, as predictive variables in asset pricing models.

The observed counter-cyclical behavior of returns following periods of experienced payout growth might indicate that investors adapt their expectations based on past dividend growth rates, a behavior which could be incorporated into asset pricing models to better forecast future returns.

The significant negative relationship between experienced returns and future returns could indicate mean-reversion in stock returns. This finding provides an empirical basis for incorporating mechanisms for capturing mean-reversion behavior in asset pricing models.

However, the inconsistency of the p-values across different time periods and predictor variables emphasizes the need for careful and context-specific application of these predictors in practice. This is why we develop in the next sections an asset pricing model with adaptable learning gain to pin down the dynamics of related factors.

Table 4.2: **Predicting Returns with Experienced Real Growth.**

Dependent variable is the log return of the CRSP value-weighted index in quarter $t+1$ in excess of the return on a 3-month T-bill. In Panel A, experienced payout growth denotes a long-run exponentially weighted average of overlapping quarterly observations of four-quarter per-capita repurchase-adjusted real dividend growth rates leading up to and including quarter t, constructed with weights implied by constant gain learning with adaptive learning gain $\nu_t = 0.01858007 + 0.06304578 \cdot g(t)$. In Panel B, experienced returns are constructed analogously as an exponentially weighted average of quarterly log stock market index returns. Inflation is measured as the average log CPI inflation rate during the four quarters $t-3$ to t; p-d refers to the log price-dividend ratio of the CRSP value-weighted index at the end of quarter t. The table shows slope coefficient estimates, with standard errors in brackets. Intercepts are not shown. p-values are shown in parentheses. The reported R^2 and number of observations are also shown.

	(1) 1927-2020	(2) 1927-2020	(3) 1927-2020	(4) 1980-2020	(5) 1980-2020
Panel A: Predicting returns with experienced real payout growth					
Experienced real payout growth	-0.003	-0.003	-0.002	-0.004	-0.004
	[0.001]	[0.001]	[0.001]	[0.001]	[0.002]
	(0.000)	(0.001)	(0.081)	(0.000)	(0.046)
Inflation		-0.161	-0.050	-0.200	-0.042
		[0.182]	[0.032]	[-3.774]	[0.033]
		(0.379)	(0.120)	(0.332)	(0.194)
p - d			0.017		0.005
			[0.017]		[0.022]
			(0.301)		(0.800)
Observations	201	196	199	159	162
Adjusted R-squared	0.054	0.051	0.060	0.072	0.073
Panel B: Predicting returns with experienced returns					
Experienced real returns	-0.217	-0.209	-0.155	-0.282	-0.292
	[0.061]	[0.062]	[0.100]	[0.080]	[0.179]
	(0.001)	(0.001)	(0.123)	(0.001)	(0.105)
Inflation		-0.156	-0.050	-0.155	-0.041
		[0.183]	[0.032]	0.200	[0.033]
		(0.393)	(0.125)	(0.440)	(0.216)
p - d			0.015		-0.002
			[0.200]		[0.028]
			0.451		(0.958)
Observations	201	196	199	159	162
Adjusted R-squared	0.054	0.049	0.057	0.063	0.065

Chapter 5

ASSET PRICING MODEL

This chapter sets up the theoretical foundation for the research. We offer a holistic perspective on the interplay between consumption growth, learning, and asset pricing. We provide the backdrop for the asset pricing model under study, based on the assumption that agents' consumption growth rates are stochastic. Later, we forecast future consumption growth rates based on historical data. This predictive distribution is crucial because it quantifies the uncertainty about future consumption growth rates, which is central to decision-making in financial markets. Moving forward, we develop a method to understand how learning affects agents' perception of shocks to the endowment process. This section forms the crux of the mathematical modelling part of this project and serves to establish the agent's learning mechanism under dynamic learning gains. Utilizing Kalman filtering method, we estimate unknown parameters (the expected consumption growth) from the observed data. In other words, we estimate the state of the dynamic system based on a series of measurements observed over time. We also explore how prior information affects the agent's learning and, in turn, asset pricing. Finally, by determining the stochastic discount factor (SDF), we analyze how the uncertainty about future consumption growth and learning impacts the SDF and, ultimately, asset prices.

By unraveling the relationships between consumption patterns and asset prices, this chapter contributes to the evolving landscape of asset pricing theory. It sheds light on the complexities inherent in decision-making processes, moving beyond traditional paradigms of rationality and constant learning gain.

5.1 Consumption in Asset Pricing

Let assume for simplicity that endowment growth follow a stochastic i.i.d. process

$$\Delta c_{t+1} = \mu + \sigma\epsilon_{t+1} \tag{5.1}$$

where μ and σ denote respectively the expected rate of growth and the volatility of growth. The agent is aware that Δc_{t+1} is i.i.d. , and she also knows σ. $\{\epsilon_t\}$ is a series of i.i.d. standard normal shocks.

The goal here is to approximate the unknown growth mean μ. To form an estimate of μ, the agent relies on the history of past endowment growth realizations, $H_t \equiv \{\Delta c_t, \Delta c_{t-1}, ...\}$, . Though we assume that the agent learns with fading memory, unlike Nagel and Xu (2019), we assume that she learns with adaptable gains $\nu_t = \nu_1 + (\nu_2 - \nu_1) * g(t)$. Moreover, unlike in most standard constant-gain learning models, we retain the modeling of the full posterior distribution—and hence the agent's subjective uncertainty—of the Bayesian approach. To do so, we use a weighted likelihood that has been used in the theoretical biology literature to model memory decay in organisms Mangel (1990). Nagel and Xu (2019) used almost the same approach with some nuance to what we did.

Given all these precedents, our representative agent, who has seen an infinite history of observations on Δc and possesses adaptable learning gains, forms a posterior as follows:

$$p(\mu \mid H_t) \propto p(\mu) \, \text{x} \, p(H_t \mid \mu) \tag{5.2}$$

with likelihood

$$p(H_t \mid \mu) \propto \left[\prod_{j=0}^{l} \left[\exp\left(-\frac{(\Delta c_{t+j} - \mu)^2}{2\sigma^2} \right) \right]^{(1-\nu_1)^j} \prod_{j=l+1}^{\infty} \left[\exp\left(-\frac{(\Delta c_{t+j} - \mu)^2}{2\sigma^2} \right) \right]^{(1-\nu_2)^j} \right] \tag{5.3}$$

and prior

$$p(\mu) = \frac{1}{\sqrt{2\pi\sigma_0^2}} exp\left(\frac{-(\mu - \mu_0)^2}{2\sigma_0^2} \right) \tag{5.4}$$

The finding about the variance of the posterior here is interesting with regard to finding when the agent have full memory from one hand, and when the memory decay occurs at a constant rate in the second hand. The variance of the posterior is the same as if the agent had observed, and retained with full memory with adaptive weights in the opposite of constant weight in the posterior a la Nagel and Xu (2019), $S \equiv 1/\nu$ realized growth rate observations. Though in our scenario the actual number of observed realizations is infinite, we come to the same conclusion with Nagel and Xu (2019) that the loss of memory implies that the effective sample size is finite and equal to S.

However, for simplicity, we work with an uninformative prior $(\sigma_0 \to \infty)$ that led to a relatively tractable posterior

$$\mu \mid H_t \sim \mathcal{N}\left(\frac{\sum_{j=0}^{\infty} \Delta c_{t+j} \left[(1-\nu_1)^j \mathbb{1}_{j \le l} + (1-\nu_2)^j \mathbb{1}_{j > l} \right]}{\sum_{j=0}^{\infty} \left[(1-\nu_1)^j \mathbb{1}_{j \le l} + (1-\nu_2)^j \mathbb{1}_{j > l} \right]}, \right.$$
$$\left. \frac{\sigma^2}{\sum_{j=0}^{\infty} \left[(1-\nu_1)^j \mathbb{1}_{j \le l} + (1-\nu_2)^j \mathbb{1}_{j > l} \right]} \right) \tag{5.5}$$

With the prior assumed uninformative, the posterior mean is updated recursively as

$$\tilde{\mu}_t = \tilde{\mu}_{t-1} + [\nu_1 + (\nu_2 - \nu_1) g(t)](\Delta c_t - \tilde{\mu}_{t-1}) \tag{5.6}$$

Starting from Equation (4.2), and replacing $\nu_t^* = \nu_1 + (\nu_2 - \nu_1) g(t)$ we can rearrange it as:

$$\tilde{\mu}_t = (1 - \nu_t^*) \tilde{\mu}_{t-1} + \nu_t^* \Delta c_t \tag{5.7}$$

Now we subtract $\tilde{\mu}_t$ on both sides and divide by $\sigma\sqrt{1 + \nu_t^*}$ to obtain:

$$\frac{\tilde{\mu}_t - \tilde{\mu}_{t-1}}{\sigma\sqrt{1 + \nu_t^*}} = \frac{\nu_t^*}{\sigma\sqrt{1 + \nu_t^*}} (\Delta c_t - \tilde{\mu}_{t-1}) \tag{5.8}$$

We recognize the left-hand side as a normalized increment in $\tilde{\mu}_t$, i.e., $\frac{\tilde{\mu}_t - \tilde{\mu}_{t-1}}{\sigma\sqrt{1+\nu_t^*}} = \tilde{\epsilon}_{t+1}$. Substituting this expression and multiplying both sides by $\sigma\sqrt{1 + \nu_t^*}$ yields:

$$\tilde{\mu}_{t+1} = \tilde{\mu}_t + \nu^*_{t+1}\sigma\sqrt{1+\nu^*_{t+1}}\tilde{\epsilon}_{t+1} \quad (5.9)$$

with $\nu^*_t = \nu_1 + (\nu_2 - \nu_1)g(t)$

$$\tilde{\mu}_{t+1} = \tilde{\mu}_t = (\nu_1 + (\nu_2 - \nu_1)g(t+1))\sigma\sqrt{1 + (\nu_1 + (\nu_2 - \nu_1)g(t+1))}\tilde{\epsilon}_{t+1} \quad (5.10)$$

Equation (5.10) shows that the posterior mean resulting from weighted-likelihood approach with uninformative prior is similar to the perceived endowment growth with adaptable gains scheme we saw in equation (4.1).

Moving forward, equation 5.11 represents the updating rule for the perceived mean consumption growth, denoted by $\tilde{\mu}_t$. This is just a new way to rewrite equation 4.2 by replacing the stock payout growth by the consumption growth. At each point in time, agents update their beliefs about the mean consumption growth based on the discrepancy between their current perceived mean consumption growth $\tilde{\mu}_t$ and the actual observed consumption growth Δc_{t+1}. The ν^*_{t+1} parameter represents the adaptive learning gains, which determines how heavily agents weigh new information relative to their prior beliefs. The term $(\Delta c_{t+1} - \tilde{\mu}_t)$ represents the prediction error - the difference between the actual consumption growth and the perceived mean consumption growth. In simple terms, if the actual consumption growth is higher than the expected one, the perceived mean consumption growth will be adjusted upwards, and vice versa. The speed and magnitude of the adjustment depend on the learning gains ν^*_{t+1}.

Equation 5.12 represents the same updating process as the first one, but it's reformulated in terms of standardized units. This equation essentially shows how the standardized perceived mean consumption growth changes in response to the standardized prediction error. It is useful for comparing the updating process across different scales of consumption growth and learning gains.

$$\tilde{\mu}_{t+1} = \tilde{\mu}_t + \nu^*_{t+1}(\Delta c_{t+1} - \tilde{\mu}_t) \quad (5.11)$$

$$\frac{\tilde{\mu}_{t+1} - \tilde{\mu}_t}{\sigma\sqrt{1+\nu^*_{t+1}}} = \frac{\nu^*_{t+1}}{\sigma\sqrt{1+\nu^*_{t+1}}}(\Delta c_{t+1} - \tilde{\mu}_t) \tag{5.12}$$

5.2 Predictive Distribution of Future Consumption Growth Given Current Information

To find the predictive distribution of $\Delta c_{t+j}|H_t$, we first recall the model equation:

$$\Delta c_{t+j} = \mu + \sigma\epsilon_{t+1}, \tag{5.13}$$

Since we have the posterior distribution for μ, we can now find the predictive distribution of $\Delta c_{t+j}|H_t$. We know that ϵ_{t+j} is normally distributed with mean 0 and variance 1, so the predictive distribution will be a convolution of the posterior distribution of μ and the distribution of $\sigma\epsilon_{t+j}$.

The posterior distribution of μ is a normal distribution with mean μ^* and variance σ^{*2}, as we found in the previous step. The distribution of $\sigma\epsilon_{t+j}$ is a normal distribution with mean 0 and variance σ^2.

The convolution of two normal distributions is also a normal distribution, with the mean being the sum of the means and the variance being the sum of the variances:

$$\Delta c_{t+j} \mid H_t \sim \mathcal{N}(\mu^*, \sigma^{*2} + \sigma^2), \text{ j = 1, 2, ...}, \tag{5.14}$$

So, the predictive distribution of $\Delta c_{t+j}|H_t$ is a normal distribution with mean μ^* and variance $\sigma^{*2} + \sigma^2$. The predictive distribution of $\Delta c_{t+j}|H_t$ is crucial in the sense that its variance can inform not only on the uncertainty facing future shocks ϵ_{t+j} but also the uncertainty about μ.

We denote expectations under the predictive distribution with $\tilde{\mathbb{E}}(.)$. In the following, we will rewrite the stochastic dynamics of the endowment process from the subjective viewpoint of the agent.

But first of all, let's define $\tilde{\epsilon}_{t+1}$ as the standardized unexpected endowment growth:

$$\tilde{\epsilon}_{t+1} = \frac{\Delta c_{t+1} - \tilde{\mu}_t}{\sigma\sqrt{1 + \nu_1 + (\nu_2 - \nu_1)g(t)}} \tag{5.15}$$

From the definition of $\tilde{\epsilon}_{t+1}$, we obtain:

$$\tilde{\epsilon}_{t+i} = \frac{\Delta c_{t+i} - \tilde{\mu}_{t+i-1}}{\sigma\sqrt{1 + \nu_1 + (\nu_2 - \nu_1)g(t+i-1)}} \tag{5.16}$$

And the updating scheme is as follows:

$$\tilde{\mu}_t = \tilde{\mu}_{t-1} + (\nu_1 + (\nu_2 - \nu_1)g(t))(\Delta c_t - \tilde{\mu}_{t-1}) \tag{5.17}$$

This definition essentially scales the unexpected endowment growth by dividing it by the standard deviation of the forecast error. This makes ϵ_{t+1} a dimensionless quantity and allows for easier comparison of forecast errors across different time periods or when endowment growth rates have different levels of volatility.

The factor $\sqrt{1 + (\nu_1 + (\nu_2 - \nu_1)g(t))}$ is used to adjust the scaling for the effect of the learning mechanism (i.e., constant gain learning) on the forecast error volatility. This adjustment is necessary because, as the agent learns from new information, the forecast error's volatility will be affected by the learning rate, which is represented by ν_t^* in this case.

Plugging $\Delta c_{t+1} - \tilde{\mu}_t$ in (5.10) yields $\tilde{\mu}_{t+1}$ as a function of ϵ_{t+1} is:

$$\tilde{\mu}_{t+1} = \tilde{\mu}_t + [\nu_1 + (\nu_2 - \nu_1)g(t)]\sigma\sqrt{1 + (\nu_1 + (\nu_2 - \nu_1)g(t))}\epsilon_{t+1}. \tag{5.18}$$

Now that we dispose of the analytical form of the standardized error (the standardized unexpected endowment growth), let's dive into deriving the properties of $\tilde{\epsilon}_{t+j}$, j= 1,2,... under the time-t predictive distribution.

5.3 Influence of Current Adjusted Expectations on Future Prediction Errors

In appendix C, we show that

$$\widetilde{cov}_t(\tilde{\epsilon}_{t+i}, \tilde{\epsilon}_{t+i+1}) = -\frac{(\nu_1 + (\nu_2 - \nu_1)g(t+i))}{\sigma\sqrt{1 + \nu_1 + (\nu_2 - \nu_1)g(t+i)}} < 0 \tag{5.19}$$

Thus, we have derived the expression for $c\tilde{o}v_t(\tilde{\epsilon}_{t+i}, \tilde{\epsilon}_{t+i+1})$ in terms of the learning parameters ν_1 and ν_2, as well as the smooth transition function $g(t)$.

Even if the cross-covariances between the forecast errors and the shocks are zero for all lags, the gain function can still generate persistence in the cross-covariances between consecutive forecast errors, as our demonstration shows. This is because the gain function depends on the past values of the forecast errors and can be non-zero even if the current forecast error is uncorrelated with the current shock.

In a scenario where Bayesian learning is subject to constant gain and fading memory, the agent's information structure deviates from the conventional filtration. In our model, the agent's posterior beliefs are formed based on a weighted likelihood, where more recent observations are given higher weights. Due to the nature of adaptable gains that result into fading memory, the posterior beliefs in future periods will be based on a different set of information that is not necessarily more informative about μ than the information available to the agent at time t. As a corollary, the information structure does not form a filtration, and there is no convergence to μ in the long run.

The time-t agent perceives future increments $\tilde{\epsilon}_{t+j}$ as negatively serially correlated, instead of martingale differences. However, the agent is unable to exploit this serial correlation to forecast $\tilde{\epsilon}_{t+1}$, as they lack the full memory to compare $\tilde{\mu}_t$ with $\tilde{\mu}_{t-1}$. This limitation arises from the adaptable-gain learning approach, which focuses on more recent data.

In our model, the predictive distribution of $\Delta c_{t+j}|H_t$ is given by $\Delta c_{t+j}|H_t \sim \mathcal{N}(\tilde{\mu}_t, \sigma^{*2} + \sigma^2)$. The agent's posterior mean, $\tilde{E}[\tilde{\mu}_{t+1}] = \tilde{\mu}_t$, is still the best forecast under the given learning process.

To various regards, this model illustrates the importance of considering the effects of fading memory and adaptable gain learning on the agent's information structure and the resulting implications for asset pricing. This approach provides a more realistic representation of learning in financial markets, as agents may not always possess full memory and may assign different weights to past information. Incorporating this paradigm could significantly reveals the way agents view and participate on the markets.

5.4 Asset Valuation

The nature of the predictive distribution of $\Delta c_{t+j}|H_t$ we derived in the previous section was a signal that the LIE could no longer be applied. A well-known valuation approach used in the asset pricing research is the "buy-and-hold" valuation, in which the agent values the asset based on the stochastically discounted payoff under the time-t predictive distribution

$$P_{H,t} = \tilde{\mathbb{E}}\left[M_{t+1|t} M_{t+2|t} C_{t+2}\right], \tag{5.20}$$

In this valuation approach, the agent bases her valuation on the expected future cash flows (e.g., dividends, interest, or other income) and the expected future price of the asset at the end of the holding period. The valuation takes into account the time value of money by discounting future cash flows to their present value using an appropriate discount rate. This is possible because she assumes the probabilities and expectations used in the calculations are based on the same information at each step of the process.

However, when factors such as learning, memory loss, or changing market conditions come into play, the information structure changes over time, and the consistency required for the Law of Iterated Expectations (LIE) to hold no longer exists.

It is fortunate that an alternative valuation does exist: the "resale" valuation. The core essence of resale valuation is to determine the value of an asset by considering the perspective of an investor who anticipates selling the asset in the future, and how that future investor (or the future self of the current investor) would value the asset based on their expectations at that time.

In the resale valuation approach, the agent at time t prices the asset based on the time-t predictive distribution of the stochastically discounted $t+1$ asset value:

$$P_{R,t} = \tilde{E}_t \left[M_{t+1|t} \tilde{E}_{t+1} \left[M_{t+2|t+1} C_{t+2} \right] \right], \tag{5.21}$$

where $M_{t+j|t}$ represents the one-period stochastic discount factor (SDF) from $t+j-1$ to $t+j$, given the agent's predictive distribution at t.

The difference in valuation between these approaches arises from the agent's perception of the statistical properties of the shock $\tilde{\epsilon}_{t+2}$ at times t and $t+1$. In the buy-and-hold valuation approach, the investor at time t takes into account the negative serial correlation between shocks $\tilde{\epsilon}_{t+1}$ and $\tilde{\epsilon}_{t+2}$. This means the investor acknowledges the relationship between these shocks when forming their expectations and valuation of the asset. As a result, this approach considers the predictive nature of these shocks when valuing the asset at time t. If the asset were priced at time t using the buy-and-hold valuation and the anticipation of a predictable $\tilde{\epsilon}_{t+2}$, the agent would find $\tilde{\epsilon}_{t+2}$ unpredictable at time $t+1$ due to memory loss. This time inconsistency is a major blow for this approach of valuation when the process does not have the properties of a martingale.

On the other hand, in the resale valuation approach, the investor at time t values the asset based on the expectation that the future investor at time $t+1$ (or the investor's future self) will perceive the shock $\tilde{\epsilon}_{t+2}$ as unpredictable. This is because, in this approach, the investor anticipates that the future investor will have a different set of beliefs and information, which may lead them to perceive the shock as unpredictable. As a result, the resale valuation approach does not incorporate the negative serial correlation between shocks $\tilde{\epsilon}_{t+1}$ and $\tilde{\epsilon}_{t+2}$

and the inconsistency raised in buy-and-hold does not occur with the resale valuation. In choosing the resale valuation approach, real asset resales take place as one generation is replaced by another.

5.5 Kalman Filtering

In this section, we describe an alternative interpretation of our model, where the agent perceives a stochastic trend and uses a Kalman filter with a diffuse prior to optimally track it.

The agent perceives the law of motion for the consumption growth rate, Δc_t, as described by equations 5.22 shown below. In this interpretation, the agent knows the variances of the shocks, σ_ξ^2 and σ_ζ^2, but not the actual trend growth rate, μ_t.

$$\Delta c_t = \mu_t + \xi_t, \quad \xi_t \sim \mathcal{N}(\prime, \sigma_\xi^{\in}), \tag{5.22}$$

$$\mu_{t+1} = \mu_t + \zeta_{t+1}, \quad \zeta_{t+1} \sim \mathcal{N}(\prime, \sigma_\zeta^{\in}) \tag{5.23}$$

We now describe a steady-state Kalman filter a la Hamilton (1994) that the agent uses to form predictions about the trend growth rate, μ_{t+1}, given an infinite history of observed data, H_t. The predictive distribution of μ_{t+1} is a normal distribution with mean $\tilde{\mu}_{t+1|t}$ and variance $\omega^2 + \sigma_\zeta^2$. Equation 5.24 illustrates the distribution of these predictions.

$$\mu_{t+1} \mid H_t \sim \mathcal{N}(\tilde{\mu}_{t+1|t}, \omega^2 + \sigma_\zeta^2), \tag{5.24}$$

The optimal forecast, $\hat{\mu}_{t+1|t} \equiv \hat{E}(\mu_{t+1} \mid H_t)$, evolves according to equation 5.25, with the Kalman gain, K, defined in equation K.5. Equation 5.27 relates the variance ω^2 to the Kalman gain and the variance of the shock ξ.

$$\hat{\mu}_{t+1|t} = \hat{\mu}_{t|t-1} + K(\Delta c_t - \hat{\mu}_{t|t-1}), \tag{5.25}$$

with

$$K = \frac{\omega^2 + \sigma_\zeta^2}{\omega^2 + \sigma_\zeta^2 + \sigma_\xi^2}, \tag{5.26}$$

and

$$\omega^2 = K\sigma_\xi^2. \tag{5.27}$$

The predictive distribution of the consumption growth rate at time $t+1$, Δc_{t+1}, is given by equation 5.28.

$$\Delta c_{t+1} \sim \mathcal{N}(\hat{\mu}_{t+1|t}, \omega^2 + \sigma_\zeta^2 + \sigma_\xi^2). \tag{5.28}$$

However, the time-t predictive distribution of Δc_{t+1} for $j > 1$ here is different from the adaptive learning gains setting because the agent perceives $\tilde{\mu}$ as a martingale (a stochastic process where the expected value of the next observation is equal to the current observation) and the predictive distribution inherits these martingale dynamics. In contrast, the adaptive memory setting has the predictive distribution converging to a stationary one at long horizons.

Despite this difference in perceived distribution for $j > 1$, the pricing is the same in both settings, as under resale valuation, pricing is based on a chain of valuations of one-period ahead payoffs from selling the asset.

In the case of the time-varying gain learning model, the learning gain is defined as $\nu(t) = \nu_1 + (\nu_2 - \nu_1)g(t)$. To match the time-t predictive distribution of Δc_{t+j} between the time-varying gain learning model and the equivalent full-memory model, we need to find expressions for $K(t)$, $\sigma_\xi^2(t)$, and $\sigma_\zeta^2(t)$ that depend on the time-varying learning gain $\nu(t)$.

We can start by modifying the expressions for K, σ_ξ^2, and σ_ζ^2 from the constant-gain learning model:

$$
\begin{aligned}
K(t) &= \nu(t) = \nu_1 + (\nu_2 - \nu_1)g(t), \\
\sigma_\xi^2 &= (1 - \nu(t)^2)\sigma^2 = (1 - [\nu_1 + (\nu_2 - \nu_1)g(t)]^2)\sigma^2, \\
\sigma_\zeta^2 &= (1 + \nu(t)^2)\sigma^2 = (1 + [\nu_1 + (\nu_2 - \nu_1)g(t)]^2)\sigma^2.
\end{aligned} \tag{5.29}
$$

5.6 Informative prior

Hamilton (1994) presents a full-memory model with an informative prior, in which the agent perceives the law of motion for consumption growth, Δc_t, and the latent trend growth rate, μ_t. The agent knows the parameters σ_ξ^2, σ_ζ^2, and h, $0 \leq h < 1$, and uses the Kalman filter to optimally track the latent trend.

The agent's perception of the law of motion for consumption growth is given by equations 5.30 and 5.31.

$$\Delta c_t = \mu_t + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2) \tag{5.30}$$

$$\mu_{t+1} = (1-h)\mu + h\mu_t + \zeta_{t+1}, \quad \zeta_{t+1} \sim \mathcal{N}(0, \sigma_\zeta^2) \tag{5.31}$$

The agent's optimal forecasts are derived using steady-state Kalman filter updating, as shown in equation 5.32.

$$\hat{\mu}_{t+1|t} = (1-h)\mu + h\hat{\mu}_{t|t-1} + K(\Delta c_t - \hat{\mu}_{t|t-1}), \tag{5.32}$$

where the parameters K and ω^2 are defined in equations 5.33 and 5.34.

$$K = h\frac{\sigma_\zeta^2 + h^2\omega^2}{\sigma_\zeta^2 + \sigma_\xi^2 + h^2\omega^2} \tag{5.33}$$

$$\omega^2 = K\sigma_\xi^2/h \tag{5.34}$$

Equation (K.13) below shows the iterative expression for the agent's optimal forecast of the latent trend growth rate

$$\hat{\mu}_{t+1|t} = \frac{1-h}{1-h+K}\mu + K\sum_{j=0}^{\infty}(h-K)^j \Delta c_{t-j}. \tag{5.35}$$

Equations 5.36 and 5.37 represent the agent's predictive distributions for μ_{t+1} and Δc_{t+1}, respectively, given the agent's information set H_t.

$$\mu_{t+1} \mid H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h^2\omega^2 + \sigma_\zeta^2), \tag{5.36}$$

$$\Delta c_{t+1} \mid H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h^2\omega^2 + \sigma_\zeta^2 + \sigma_\xi^2). \tag{5.37}$$

Finally, equations 5.38 show how to map the informative prior full-memory model into the time-varying learning memory setup. By choosing the parameters K, h, and σ_ξ^2 as shown, the subjective belief dynamics and asset prices are equivalent between the time-varying learning memory model with an informative prior and the full-memory model.

We can represent the new parameters as follows:

$$\begin{aligned} K(t) &= \phi\nu(t) \\ h(t) &= 1-\nu(t)+\phi\nu(t) \\ \sigma_\xi^2(t) &= \frac{1-\nu(t)}{1-\nu(t)+\phi\nu(t)}(1+\phi\nu(t))\sigma^2 \\ \sigma_\zeta^2(t) &= (1+\phi^2\nu(t))\sigma^2. \end{aligned} \tag{5.38}$$

With these modified parameters, the time-t predictive distribution of Δc_{t+j} in the time-varying gain learning model would match the predictive distribution in the equivalent full-memory model with an informative prior.

In summary, this section demonstrates how a full-memory model with an informative prior can be transformed into an equivalent time-varying learning memory model. This transformation helps in understanding the relationship between the two setups and highlights the importance of carefully choosing the parameters for the time-varying learning memory model to maintain the equivalence with the full-memory model.

5.7 Stochastic Discount Factor

Following Nagel and Xu (2019), we are considering a situation where the typical individual in our model uses a recursive utility framework, as first outlined by Epstein and Zin (1991), to determine the value of possible future outcomes. This allows us to disentangle risk aversion and the elasticity of intertemporal substitution.

$$V_t = \left[(1-\delta)C_t^{1-\frac{1}{\rho}} + \delta\widetilde{E}_t[V_{t+1}^{1-\gamma}]^{\frac{1-\frac{1}{\rho}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\rho}}} \tag{5.39}$$

where δ denotes the time discount factor, ρ the elasticity of intertemporal substitution (EIS), and γ represents the coefficient of relative risk aversion.

Following Hansen et al. (2008) Hansen et al. (2008), we can write the logarithm of the ratio of the continuation value to consumption as

$$v_t = \frac{1}{1-\rho}\text{log}\{(1-\delta) + \delta\,\text{exp}[(1-\rho)\Theta_t(v_{t+1} + c_{t+1} - c_t)]\} \tag{5.40}$$

where c_t denote the logarithm of the ratio of the continuation value (V_t) to consumption(C_t) and Θ_t is defined as

$$\Theta_t(v_{t+1}) = \frac{1}{1-\gamma}\text{log}\widetilde{\mathbb{E}}[exp[(1-\gamma)v_{t+1}]] \tag{5.41}$$

where $\widetilde{\mathbb{E}}$ denote the subjective expectation function.

Using the recursion scheme of Tallarini (1998) in a context of risk-sensitive business cycles and asset prices study, and solving for the specific case of $\rho = 1$, we can rewrite v_t as the following

$$v_t = \frac{\delta}{1-\gamma}\text{log}\widetilde{\mathbb{E}}_{\approx}[e^{(1-\gamma)(v_{t+1}+\Delta c_{t+1})}] \tag{5.42}$$

where δ denotes the time discount factor, and γ represents the coefficient of relative risk aversion.

Hansen et al. (2008) showed that this log-linear stochastic equation has a linear solution for the continuation value. Therefore, we postulate that v_t is linear in the state variable:

$$v_t = \mu_v + U_v\tilde{\mu}_t, \tag{5.43}$$

5.7.1 Model Solution for $\rho = 1$

In our scenario, SDF represents the factor by which the agent is willing to exchange future payoffs for current ones, under uncertainty. We show in appendix E that in the case of Epstein-Zin preferences, the SDF takes the following form:

$$m_{t+1|t} = \tilde{\mu}_{m,t} - \tilde{\mu}_t - \xi_t\sigma\tilde{\epsilon}_{t+1}, \tag{S1}$$

As g(t) is a deterministic and known function, the time-varying expressions of the component of $m_{t+1|t}$ are as the following:

$$\tilde{\mu}_{m,t} = log\delta - \frac{1}{2}(1-\gamma)^2(\nu_1 - (\nu_2-\nu_1)g(t)U_v + 1)^2(1+\nu_1-(\nu_2-\nu_1)g(t))\sigma^2 \tag{5.44}$$

$$\xi_t = [1-(1-\gamma)(\nu_1-(\nu_2-\nu_1)g(t)U_v+1)]\sqrt{\nu_1-(\nu_2-\nu_1)g(t)} \tag{5.45}$$

5.8 Evaluation of Consumption-Based Assets Claims

Let $\zeta \equiv W_t/C_t$ be the consumption-wealth ratio. The return on the consumption claim is

$$R_{W,t+1} \equiv \frac{W_{t+1}}{W_t - C_t} = \frac{C_{t+1}}{C_t}\frac{\zeta}{\zeta - 1}, \tag{5.46}$$

The corresponding log return is as follows:

$$
\begin{aligned}
r_{w,t+1} &= \Delta c_{t+1} + log(\zeta/(\zeta-1)) \\
&= \tilde{\mu}_t + \sigma\sqrt{1+\nu(t)}\tilde{\epsilon}_{t+1} + log(\zeta/(\zeta-1))
\end{aligned}
\tag{5.47}
$$

Plugging the return on consumption claim into the Euler equation and taking logs,

$$
\begin{aligned}
log(\zeta/(\zeta-1)) &= -\tilde{\mu}_m + \xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}(1+\nu)\sigma^2 - \frac{1}{2}\sigma^2\xi^2 \\
&= -log\ \delta
\end{aligned}
\tag{5.48}
$$

As shown by Nagel and Xu (2019), the wealth-consumption ratio

$$
\zeta = \frac{1}{1-\delta} \tag{5.49}
$$

In the context of Nagel and Xu's study, they begin by examining the log of the part of wealth at time t associated with the one-period ahead endowment flow, w_t^l. They relate this to the consumption at the same period, c_t. This relationship is described by:

$$
w_t^1 - c_t = log\widetilde{E}_t\left[M_{t+1|t}\frac{C_{t+1}}{C_t}\right] \tag{5.50}
$$

In our case, we're considering a scenario where the learning gain is time-varying, represented as $\nu_t = \nu_1 + (\nu_2 - \nu_1)g(t)$. This alters the dynamics of the equation above. Specifically, the time-varying learning gain changes the representation of the expected utility, yielding:

$$
w_t^1 - c_t = log\widetilde{E}_t[exp(\tilde{\mu}_m + (\sqrt{1+\nu_t} - \xi)\sigma\tilde{\epsilon}_{t+1})] \tag{5.51}
$$

We further simplify this expectation, which gives us:

$$
w_t^1 - c_t = \tilde{\mu}_m + \frac{1}{2}(\sqrt{1+\nu_t} - \xi)^2\sigma^2. \tag{5.52}
$$

The difference between the wealth and the consumption, $w_t^l - c_t$, is no longer constant over time due to the time-varying learning gain, ν_t. It results in a more intricate dynamics.

Proceeding with the valuation equation, we can infer the price of an n-period consumption strip, which evolves as:

$$w_t^n - c_t = n\tilde{\mu}_m + \frac{n}{2}(\sqrt{1+\nu_t} - \xi)^2\sigma^2. \quad (5.53)$$

Substituting the expressions for $\tilde{\mu}_m$ and ξ, we get:

$$w_t^n - c_t = nlog\ \delta + \frac{n}{2}(\sqrt{1+\nu_t} - \xi)^2\sigma^2. \quad (5.54)$$

This equation describes the price of an n-period consumption strip, adjusted for the time-adaptive learning gain. In this scenario, the price of the consumption strip is subject to changes over time following the time-varying learning gain.

In conclusion, the inclusion of a time-varying learning gain in the model leads to a more complex dynamics for the valuation of consumption strips. It enables the model to capture more nuanced aspects of adaptive learning, making it more suitable for applications in environments where learning parameters evolve over time.

5.9 Pricing Dynamics of Dividend Strips

We turn our attention to analyzing dividend strips, which are claims to single dividends to be paid in the future. These help us to transparently analyze the conditions necessary for a finite price. The price of the n-period dividend strip, denoted as P_t^n, is given by:

$$P_t^n \equiv \widetilde{E}_t[M_{t+1|t}\widetilde{E}_{t+1}[...\widetilde{E}_{t+n-1}[M_{t+1|t+n-1}D_{t+n}]]]. \quad (5.55)$$

here we make use of a time-adaptive learning gain ξ_t rather than a constant ξ, which introduces an additional layer of dynamism. Evaluating these expectations, we do so by iterating backwards from the payoff at $t+n$, evaluating one conditional expectation at a time.

Taking logs and evaluating, we now have:

$$p_t^n - d_t = [1-(1-\alpha)^n](c_t - d_t + \mu_{dc} + \frac{\lambda-1}{\alpha}\tilde{\mu}_t) + n\tilde{\mu}_m + \frac{1}{2}(A_{n,t}\sigma^2 + B_{n,t}\sigma_d^2), \quad (5.56)$$

where

$$A_{n,t} = \sum_{k=0}^{n-1}\{\sqrt{1+\nu_t}\left[\nu_t(\lambda-1)\frac{1-(1-\alpha)^k}{\alpha} + (\lambda-1)(1-\alpha)^k + 1\right] - \xi\}^2, \quad (5.57)$$

and

$$B_n = \frac{1-(1-\alpha)^{2n}}{1-(1-\alpha)^2} \quad (5.58)$$

For large n, approximately,

$$A_{n,t} \approx n\left[\sqrt{1+\nu_t}\left(1+\nu_t\frac{\lambda-1}{\alpha}\right) - \xi\right]^2, \quad (5.59)$$

And B_n, which does not grow with n, becomes very small relative to A_n. For the price to be well-defined, we need the terms that grow with n in equation 5.56 to be (weakly) negative. Using equation (34), this requires:

$$\tilde{\mu}_m + \frac{1}{2}\left[\sqrt{1+\nu_t}\left(1+\nu_t\frac{\lambda-1}{\alpha}\right) - \xi\right]^2\sigma^2 \leq 0. \quad (5.60)$$

This condition ensures that the price of consumption strips remains finite even under time-varying learning gain. The inequality provided is required to ensure that the price of the dividend strip, P_t^n, is well-defined and finite.

In this context, the inequality is designed to limit the growth of the terms in equation 5.60 that are proportional to n. If these terms are not (weakly) negative, then they could grow without bound as n increases, leading to an infinite price for the dividend strip, which is economically nonsensical. Specifically, the term $n\tilde{\mu}_m$ represents the part of the dividend strip price that grows with n, the number of periods. The term $\frac{1}{2}\left[\sqrt{1+\nu_t}\left(1+\nu_t\frac{\lambda-1}{\alpha}\right) - \xi\right]^2\sigma^2$ is the variance of the dividend strip price that also grows with n. Hence, to prevent the price from exploding as n gets large, the inequality is required.

This inequality is generally a condition for convergence in these types of valuation problems. It ensures that the infinite sum or product that typically arises when valuing these kinds of derivative securities converges to a finite value.

Equation 5.61 gives the return of a one-period claim, r^1_{t+1}.

$$r^1_{t+1} = \lambda \Delta c_{t+1} - (\lambda - 1)\tilde{\mu}_t - \tilde{\mu}_m - \frac{1}{2}(\sqrt{1+\nu_t}\lambda - \xi)^2\sigma^2 \tag{5.61}$$

Subtracting $r_{f,t} = -\tilde{\mu}_m + \tilde{\mu}_t - \frac{1}{2}\xi^2\sigma^2$ from the above equation, we get:

$$r^1_{t+1} - r_{f,t} = \lambda(\Delta c_{t+1} - \tilde{\mu}_t) - \tilde{\mu}_m - \frac{1}{2}(\sqrt{1+\nu_t}\lambda - \xi)^2\sigma^2 + \frac{1}{2}\xi^2\sigma^2 \tag{5.62}$$

The subjective conditional variance of r^1_{t+1} becomes $(1+\nu_t)\lambda^2\sigma^2$, so taking subjective expectations of 5.62, we get:

$$\log E_t[R^1_{t+1}] - r_{f,t} = \lambda\xi\sqrt{1+\nu_t}\sigma^2. \tag{5.63}$$

Now for the objective conditional variance, it will remain $\lambda^2\sigma^2$, and so taking objective expectations of 5.62 we get,

$$\log E_t[R^1_{t+1}] - r_{f,t} = \lambda\xi\sqrt{1+\nu_t}\sigma^2 - \frac{1}{2}\nu_t\lambda^2\sigma^2 + \lambda(\mu - \tilde{\mu}_t) \tag{5.64}$$

Moving on to the infinite-horizon claim. Starting from the general form, we have:

$$r^\infty_{t+1} = \Delta c_{t+1} + \frac{\lambda - 1}{\alpha_t}(\tilde{\mu}_{t+1} - \tilde{\mu}_t) - \tilde{\mu}_m - \frac{1}{2}\left[\sqrt{1+\nu_t}(1 + \nu_t\frac{\lambda - 1}{\alpha_t}) - \xi\right]^2\sigma^2. \tag{5.65}$$

After subtracting the risk-free rate, we obtain:

$$r_{t+1}^{\infty} - r_{f,t} = \Delta c_{t+1} + \frac{\lambda - 1}{\alpha_t}(\tilde{\mu}_{t+1} - \tilde{\mu}_t) - \tilde{\mu}_t - \frac{1}{2}\left[\sqrt{1+\nu_t}(1+\nu_t\frac{\lambda-1}{\alpha_t}) - \xi\right]^2 \sigma^2 + \frac{1}{2}\xi^2\sigma^2. \tag{5.66}$$

Next, we can calculate the subjective expectation and the objective expectation of r_{t+1}^{∞} using the same logic as before, adjusting for the time-adaptive learning gain, giving us:

$$\log \widetilde{E}_t[R_{t+1}^{\infty}] - r_{f,t} = \lambda\xi\sqrt{1+\nu_t}\sigma^2 - \frac{1}{2}\nu_t\left(1+\nu_t\frac{\lambda-1}{\alpha_t}\right)^2\sigma^2 \tag{5.67}$$

and

$$\log E_t[r_{t+1}^{\infty}] - r_{f,t} = \lambda\xi\sqrt{1+\nu_t}\sigma^2 - \frac{1}{2}\nu_t\left(1+\nu_t\frac{\lambda-1}{\alpha_t}\right)^2\sigma^2 + \lambda(\mu - \tilde{\mu}_t) + \frac{\lambda-1}{\alpha_t}(\mu - \tilde{\mu}_t). \tag{5.68}$$

Our model aims to numerically compute the price, P_t, of the equity claim to the entire stream of dividends, with adjustments for time-adaptive learning gain. In this context, P_t^n for $n > J$ and a sufficiently large J, is given by

$$P_t^n \approx C_t e^{\mu_{dc} + \frac{1}{2}\frac{1}{1-(1-\alpha)^2}\sigma_d^2 + \frac{\lambda-1}{\alpha}\tilde{\mu}_t} \exp(n\tilde{\mu}_m + \frac{1}{2}A_n\sigma^2), n > J, \tag{5.69}$$

where we approximate

$$A_n \approx A_j + (n-J)[\sqrt{1+\nu(t)}(\nu(t)\frac{\lambda-1}{\alpha} + 1) - \xi]^2, n > J. \tag{5.70}$$

In this context, $\nu(t) = \nu_1 + (\nu_2 - \nu_1)g(t)$ is our time-adaptive learning gain, and $g(t)$ is a smooth function of time.

Then, we can express the total price P_t as

$$P_t \approx \left(\sum_{n=1}^{J} P_t^n\right) + C_t V_J \exp\left(\mu_{dc} + \frac{1}{2}\frac{1}{1-(1-\alpha)^2}\sigma_d^2 + \frac{\lambda-1}{\alpha}\tilde{\mu}_t\right), \tag{5.71}$$

with

$$V_J = \frac{exp((J+1)\tilde{\mu}_m + \frac{1}{2}A_J\sigma^2 + \frac{1}{2}[\sqrt{1+\nu}(\nu\frac{\lambda-1}{\alpha}+1)-\xi]^2\sigma^2)}{1-\exp(\tilde{\mu}_m + \frac{1}{2}[\sqrt{1+\nu(t)}(\nu(t)\frac{\lambda-1}{\alpha}+1)-\xi]^2\sigma^2)} \tag{5.72}$$

We implement this by choosing a large enough J such that the value of P_t obtained is not sensitive to further changes in J.

For $\psi = 1$, the wealth-consumption ratio is constant

$$\log\frac{W_t - C_t}{C_t} = \log\frac{\delta}{1-\delta}, \tag{5.73}$$

and we only need to solve for the log price-dividend ratio. We express the $\log P/D$ ratio as a function of $\tilde{\mu}$ and $d_t - c_t$, i.e.,

$$\log\frac{P_t}{D_t} = H(\tilde{\mu}_t, d_t - c_t). \tag{5.74}$$

Because there are two state variables, we adopt basis functions, ψ_{ij}, in this form:

$$\psi_{ij}(\tilde{\mu}, d_t - c_t) \equiv \Lambda_i(\tilde{\mu})\Lambda_j(d_t - c_t) \tag{5.75}$$

where Λ_i denotes the Chebyshev polynomials. We approximate the log P/D ratio as:

$$\hat{H}(\tilde{\mu}, d_t - c_t; \beta_m) = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1}\beta_{m,ij}\psi_{ij}(\tilde{\mu}, d_t - c_t). \tag{5.76}$$

We rewrite the subjective Euler equation:

$$\tilde{E}_t[M_{t+1}R_{m,t+1}] = 1 \tag{5.77}$$

as

$$\begin{aligned} 0 &= I(\tilde{\mu}_t, d_t - c_t) \\ &\equiv \tilde{E}_t[e^{\tilde{\mu}_m - \tilde{\mu}_t - \xi\sigma\epsilon_{t+1} + \Delta d_{t+1}}\frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1}-c_{t+1})}+1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}] - 1 \\ &= e^{\tilde{\mu}_m + (\lambda-1)\tilde{\mu}_t - \alpha(d_t - c_t - \mu_{dc})}\tilde{E}_t[e^{(\lambda\sqrt{1+\nu(t)}-\xi)\sigma\tilde{\epsilon}_{t+1} + \sigma_d\eta_{t+1}}\frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1}-c_{t+1})}+1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}] - 1 \end{aligned} \tag{5.78}$$

In our investigation, we have developed a comprehensive model that accounts for the influence of time-adaptive learning gain in numerically computing the equity claim's price to the entire stream of dividends. Our equations give an effective approximation for large J values, thereby yielding a robust model for P_t. Additionally, our model effectively captures the effects of two state variables, further enhancing the model's predictive capacity. By numerically solving the given system of equations, we are able to draw detailed conclusions about equity prices and the equity premium in our specific time-adaptive learning gain scenario. This represents a significant advancement in our understanding of how learning mechanisms and information structures can influence financial market outcomes.

However, the development of this model also points towards several avenues for future work. The complexity of the equations, as well as the assumptions that underpin our model, calls for a further exploration. Specifically, we acknowledge that the model's calibration, an essential step for verifying its accuracy and applicability in different financial scenarios, has yet to be accomplished. Our future investigations will focus on model calibration to ensure the robustness of the model under a variety of conditions.

Chapter 6

CONCLUSION

In conclusion, this study sheds light on the complexities of asset pricing and challenges the prevailing notions of rationality and constant learning gain. The empirical analysis highlights the counter-cyclical behavior of returns following periods of experienced payout growth and provides support for the mean-reversion hypothesis in stock market returns. These findings underscore the significance of adaptive learning measures, such as experienced payout growth and experienced returns, in asset pricing models.However, the varying statistical significance across different scenarios and predictor variables calls for cautious interpretation and context-specific application of these predictors.

From the theoretical standpoint, within the adaptive learning gains model, the equity risk premium rises in line with increases in consumption change, learning gain, and the risk level as depicted by the standard deviation of shocks. Essentially, higher uncertainty or variability compels investors to seek additional compensation, thus elevating the equity risk premium.

Furthermore, the time-t agent perceives future increments of shocks $\tilde{\epsilon}_{t+j}$ as negatively serially correlated, instead of martingale differences. However, the agent is unable to exploit this serial correlation to forecast $\tilde{\epsilon}_{t+1}$, as they lack the full memory to compare $\tilde{\mu}_t$ with $\tilde{\mu}_{t-1}$. Recognizing the limitations of the traditional "buy-and-hold" valuation approach, this research introduces an innovative asset pricing model with adaptable learning gains. The theoretical analysis reveals the absence of convergence to a specific value and the breakdown of the Law of Iterated Expectations (LIE) under adaptable learning. It also emphasizes the importance of the "resale" valuation method, which considers the time-t predictive distribution of stochastically discounted future asset values.

The findings presented in this paper contribute to a deeper understanding of asset pricing dynamics by incorporating personal experiences, adaptive learning processes, and subjective return expectations. Moreover, the study expands on previous research by considering a broader set of market participants, providing a more comprehensive context for the analysis.

Moving forward, we aim in a future research to calibrate and validate the proposed asset pricing model with adaptable learning gains.

Overall, this research contributes to the evolving landscape of asset pricing theory by embracing the complexities of decision-making and highlighting the importance of adaptive learning in capturing the intricate dynamics of asset pricing. By incorporating these insights into economic models, we can gain a more accurate understanding of market behavior and improve predictions in the field of finance.

BIBLIOGRAPHY

Baker, M. and Wurgler, J. (2004). A catering theory of dividends. *The Journal of finance*, 59(3):1125–1165.

Baker, M. and Wurgler, J. (2007). Investor sentiment in the stock market. *Journal of economic perspectives*, 21(2):129–151.

Bansal, R., Dittmar, R. F., and Lundblad, C. T. (2005). Consumption, dividends, and the cross section of equity returns. *The Journal of Finance*, 60(4):1639–1672.

Barberis, N., Huang, M., and Santos, T. (2001). Prospect theory and asset prices. *The quarterly journal of economics*, 116(1):1–53.

Bernanke, B. S. and Woodford, M. (1997). Inflation forecasts and monetary policy.

Bullard, J. and Mitra, K. (2000). mdeterminacy, learnability, and monetary policy inertia, v federal reserve bank of st. Technical report, Louis Working Paper 2000% 030B.

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

Conrad, J. and Kaul, G. (1998). An anatomy of trading strategies. *The Review of Financial Studies*, 11(3):489–519.

Epstein, L. G. and Zin, S. E. (1991). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of political Economy*, 99(2):263–286.

Evans, G. W. and Honkapohja, S. (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

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

Fama, E. F. and French, K. R. (2001). Disappearing dividends: changing firm characteristics or lower propensity to pay? *Journal of Financial economics*, 60(1):3–43.

Hamilton, J. D. (1994). State-space models. *Handbook of econometrics*, 4:3039–3080.

Hansen, L. P., Heaton, J. C., and Li, N. (2008). Consumption strikes back? measuring long-run risk. *Journal of Political economy*, 116(2):260–302.

Howitt, P. (1992). Interest rate control and nonconvergence to rational expectations. *Journal of Political Economy*, 100(4):776–800.

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

Jegadeesh, N. and Titman, S. (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of finance*, 56(2):699–720.

Lo, A. W. and MacKinlay, A. C. (2011). *A non-random walk down Wall Street*. Princeton University Press.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The quarterly journal of economics*, 126(1):373–416.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Malmendier, U. and Wachter, J. A. (2021). Memory of past experiences and economic decisions. *Available at SSRN 4013583*.

Mangel, M. (1990). A dynamic habitat selection game. *Mathematical Biosciences*, 100(2):241–248.

Michaely, R., Thaler, R. H., and Womack, K. L. (1995). Price reactions to dividend initiations and omissions: Overreaction or drift? *the Journal of Finance*, 50(2):573–608.

Nagel, S. and Xu, Z. (2019). Asset pricing with fading memory (no. w26255).

Svensson, L. E. and Woodford, M. (2003). Indicator variables for optimal policy. *Journal of monetary economics*, 50(3):691–720.

Wang, R. (2021). *Subjective Beliefs and Asset Prices*. Columbia University.

Woodford, M. (1999). Optimal monetary policy inertia. *The Manchester School*, 67:1–35.

APPENDICES

A Deriving the innovation return

By applying the Campbell-Shiller Decomposition in Campbell-Shiller (1988) to the return innovations along with investors' subjective expectations, denoted $\widetilde{E}$ we obtain

$$r_{t+1} - \widetilde{E}_t r_{t+1} = (\widetilde{E}_{t+1} - \widetilde{E}_t)\sum_{s\geq 0} \rho^s \Delta_{t+1+s} - (\widetilde{E}_{t+1} - \widetilde{E}_t)\sum_{s\geq 1} r_{t+1+s} \tag{1}$$

In general, $\tilde{\mu}_d$ is an exponential-weighted average of past Δd observations.
Because investors with subjective beliefs do not revise they return expectations once they are set, return expectations stay fixed, $\widetilde{E}_t(r_{t+1}) = \theta + r_f$, with θ the constant risk premium and r_f the real constant risk-free rate.

$$\begin{aligned}
r_{t+1} - \widetilde{E}_t r_{t+1} &= (\widetilde{E}_{t+1} - \widetilde{E}_t)\sum_{s\geq 0} \rho^s \Delta_{t+1+s} \\
&= (\widetilde{E}_{t+1} - \widetilde{E}_t)[\Delta d_{t+1} + \sum_{s\geq 1} \rho^s \Delta_{t+1+s}] \\
&= \frac{\rho}{1-\rho}(\tilde{\mu}_{d,t+1} - \tilde{\mu}_{d,t}) + (\Delta d_{t+1} - \tilde{\mu}_{d,t}) \\
&\textit{By replacing} \quad (\tilde{\mu}_{d,t+1} - \tilde{\mu}_{d,t}) \quad \textit{by the law of motion of the perceived growth rate}, \\
&= \frac{\rho}{1-\rho}[\nu_1 + (\nu_2 - \nu_1)g(t)](\Delta d_{t+1} - \tilde{\mu}_{d,t}) + (\Delta d_{t+1} - \tilde{\mu}_{d,t}) \\
&= [\frac{\rho}{1-\rho}(\nu_1 + (\nu_2 - \nu_1)g(t)) + 1](\Delta d_{t+1} - \tilde{\mu}_{d,t}) \\
r_{t+1} &= [\frac{\rho}{1-\rho}(\nu_1 + (\nu_2 - \nu_1)g(t)) + 1](\Delta d_{t+1} - \tilde{\mu}_{d,t}) + \theta + r_f
\end{aligned} \tag{2}$$

If we further assume that we know the true growth rate μ_d when applying econometric techniques to a large sample of data, we can take expectations of the above final equation. Under these objective beliefs, we obtain:

$$\widetilde{E}_t r_{t+1} - r_f = \theta + [\frac{\rho}{1-\rho}(\nu_1 + (\nu_2 - \nu_1)\widetilde{E}_t g(t)) + 1](\Delta d_{t+1} - \tilde{\mu}_{d,t}) \tag{3}$$

where the term in parentheses times Delta $d_{t+1} - \tilde{\mu}_{d,t}$ represents the subjective growth-rate expectations revision that the econometrician anticipates, on average, in the next period, given her knowledge of $\mu_{d,t}$.

$$\widetilde{E}_t r_{t+1} - r_f = \theta + [\frac{\rho}{1-\rho}(\nu_1 + (\nu_2 - \nu_1)g(t)) + 1](\Delta d_{t+1} - \tilde{\mu}_{d,t}) \tag{4}$$

as $g(t)$ is deterministic at every time t.

B Derivation of the posterior

Combining the prior and the weighted likelihood:

$$p(\mu \mid H_t) \propto \exp\left(-\frac{(\mu-\mu_0)^2}{2\sigma_0^2}\right) \times \left[\prod_{j=0}^{l}\left[\exp\left(-\frac{(\Delta c_{t+j}-\mu)^2}{2\sigma^2}\right)\right]^{(1-\nu_1)^j} \times \prod_{j=l+1}^{\infty}\left[\exp\left(-\frac{(\Delta c_{t+j}-\mu)^2}{2\sigma^2}\right)\right]^{(1-\nu_2)^j}\right] \quad (5)$$

Now, let's combine the exponentials:

$$p(\mu \mid H_t) \propto exp\left(-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \sum_{j=0}^{l}(1-\nu_1)^j\frac{(\Delta c_{t+1}-\mu)^2}{2\sigma^2} - \sum_{j=l+1}^{\infty}(1-\nu_2)^j\frac{(\Delta c_{t+1}-\mu)^2}{2\sigma^2}\right) \quad (6)$$

$$p(\mu \mid H_t) \propto exp\left(-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \sum_{j=0}^{l}(1-\nu_1)^j\frac{(\Delta c_{t+1}-\mu)^2}{2\sigma^2} - \sum_{j=l+1}^{\infty}(1-\nu_2)^j\frac{(\Delta c_{t+1}-\mu)^2}{2\sigma^2}\right) \quad (7)$$

$$p(\mu \mid H_t) \propto exp\left[-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \frac{1}{2\sigma^2}\left(\sum_{j=0}^{l}(1-\nu_1)^j(\Delta c_{t+1}-\mu)^2 - \sum_{j=l+1}^{\infty}(1-\nu_2)^j(\Delta c_{t+1}-\mu)^2\right)\right] \quad (8)$$

$$p(\mu \mid H_t) \propto exp\left[-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \frac{1}{2\sigma^2}\left(\sum_{j=0}^{\infty}(\Delta c_{t+1}-\mu)^2\left[(1-\nu_1)^j\mathbb{1}_{j\le l} + (1-\nu_2)^j\mathbb{1}_{j>l}\right]\right)\right] \quad (9)$$

where $\mathbb{1}_{\text{condition}}$ is an indicator function, which is equal to 1 if the condition inside is true and 0 otherwise.
Now, we can rewrite the expression inside the exponential function as a quadratic function of μ:

$$-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \frac{1}{2\sigma^2}\sum_{j=0}^{\infty}(\Delta c_{t+1}-\mu)^2\left[(1-\nu_1)^j\mathbb{1}_{j\le l} + (1-\nu_2)^j\mathbb{1}_{j>l}\right] = -\frac{1}{2}A\mu^2 + B\mu - C \quad (10)$$

The coefficients A, B, and C will be functions of the data and the parameters. Since the posterior is proportional to the exponential of this quadratic function, it will have a normal distribution. To find the mean and variance of this distribution, we need to find the maximum of the quadratic function (i.e., the value of μ that maximizes the function) and the curvature of the function at the maximum.
To find the maximum of the quadratic function, we can differentiate it with respect to μ and set the derivative equal to zero:

$$\frac{d}{d\mu}\left(-\frac{1}{2}A\mu^2 + B\mu - C\right) = -A\mu + B = 0 \tag{11}$$

Solving for μ, we get the mean of the posterior distribution:

$$\mu* = \frac{B}{A} \tag{12}$$

To find the curvature, we can differentiate the quadratic function again with respect to μ:

$$\frac{d^2}{d\mu^2}\left(-\frac{1}{2}A\mu^2 + B\mu - C\right) = -A \tag{13}$$

Since the curvature is constant, the variance of the posterior distribution will be the reciprocal of the negative curvature:

$$\sigma*^2 = \frac{1}{A} \tag{14}$$

We just need now to derive the expressions for A, B, and C.
Let's recall the expression for the posterior:

$$p(\mu \mid H_t) \propto exp\left[-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \frac{1}{2\sigma^2}\left(\sum_{j=0}^{\infty}(\Delta c_{t+1}-\mu)^2\left[(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}\right]\right)\right] \tag{15}$$

Now, let's expand the terms in the exponential:

$$-\frac{(\mu-\mu_0)^2}{2\sigma_0^2} - \frac{1}{2\sigma^2}\sum_{j=0}^{\infty}(\Delta c_{t+1}-\mu)^2\left[(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}\right] = -\frac{1}{2}A\mu^2 + B\mu - C \tag{16}$$

Expanding the terms and combining similar terms, we have:

$$\begin{aligned} &-\frac{\mu^2}{2\sigma_0^2} + \frac{\mu\mu_0}{\sigma_0^2} - \frac{\mu_0^2}{2\sigma^2} \\ &-\frac{1}{2\sigma^2}\left[\sum_{j=0}^{\infty}((\Delta c_{t+j})^2 - 2\mu\Delta c_{t+j} + \mu^2)[(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}]\right] \\ &= -\frac{1}{2}A\mu^2 + B\mu - C \end{aligned} \tag{17}$$

Now, we can identify the coefficients A, B, and C:

$$A = \frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}\left[\sum_{j=0}^{\infty}(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}\right] \tag{18}$$

$$B = \frac{\mu_0}{\sigma_0^2} + \frac{1}{\sigma^2}\left[\sum_{j=0}^{\infty}\Delta c_{t+j}\left[(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}\right]\right] \tag{19}$$

$$C = \frac{\mu_0^2}{2\sigma_0^2} + \frac{1}{2\sigma^2}\left[\sum_{j=0}^{\infty}(\Delta c_{t+j})^2\left[(1-\nu_1)^j \mathbb{1}_{j\leq l} + (1-\nu_2)^j \mathbb{1}_{j>l}\right]\right] \tag{20}$$

Another method:

$$
\begin{aligned}
p(\mu \mid H_t) &\propto \exp\left[-\frac{(\mu-\mu_0)^2}{2\sigma_0^2}-\frac{1}{2\sigma^2}\left(\sum_{j=0}^{\infty}(\Delta c_{t+1}-\mu)^2\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]\right)\right]\\
&\propto \exp\left[-\frac{\mu^2}{2\sigma_0^2}+\frac{\mu\mu_0}{\sigma_0^2}-\frac{\mu_0^2}{2\sigma^2}-\frac{1}{2\sigma^2}\left[\sum_{j=0}^{\infty}((\Delta c_{t+j})^2-2\mu\Delta c_{t+j}+\mu^2)\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)\right.\right.\right.\\
&\propto \exp\left[-\frac{1}{2}A\mu^2+B\mu-C\right]\\
&\propto \exp\left[-\frac{1}{2}A\left(\mu-\frac{B}{A}\right)^2\right]\\
&\propto \exp\left[-\frac{1}{2}\frac{1}{1/A}\left(\mu-\frac{B}{A}\right)^2\right]
\end{aligned}
\tag{21}
$$

We made use of the fact that in the second line, the constant term does not depend on μ, and it will be absorbed into the proportionality constant for the posterior distribution. Starting with prior $\mu \sim \mathcal{N}(\mu_{\prime}, \sigma_{\prime})$ before seeing any data, the agent finally form a posterior as follows:

$$
\begin{aligned}
\mu \mid H_t \sim \mathcal{N}\Bigg(&\frac{\frac{\mu_0}{\sigma_0^2}+\frac{1}{\sigma^2}\left[\sum_{j=0}^{\infty}\Delta c_{t+j}\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>t}\right]\right]}{\frac{1}{\sigma_0^2}+\frac{1}{\sigma^2}\left[\sum_{j=0}^{\infty}(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]},\\
&\frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{\sigma^2}\left[\sum_{j=0}^{\infty}(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]}\Bigg)
\end{aligned}
\tag{22}
$$

As stated in the section 5.1 , for simplicity, we work with an uninformative prior $(\sigma_0 \to \infty)$ that led to a relatively tractable posterior

$$
\begin{aligned}
\mu \mid H_t \sim \mathcal{N}\Bigg(&\frac{\sum_{j=0}^{\infty}\Delta c_{t+j}\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]}{\sum_{j=0}^{\infty}\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]},\\
&\frac{\sigma^2}{\sum_{j=0}^{\infty}\left[(1-\nu_1)^j\mathbb{1}_{j\le l}+(1-\nu_2)^j\mathbb{1}_{j>l}\right]}\Bigg)
\end{aligned}
\tag{23}
$$

C Deriving the expression for the standardized unexpected endowment growth

Starting with equation (29) and using equation (30), we have:

$$\begin{aligned}
\tilde{\epsilon}_{t+i} &= \frac{\Delta c_{t+i} - \tilde{\mu}_{t+i-1}}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}} \\
&= \frac{\Delta c_{t+i} - (\tilde{\mu}_{t+i-2} + (\nu_1+(\nu_2-\nu_1)g(t))(\Delta c_{t+i-1} - \tilde{\mu}_{t+i-2}))}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}} \\
&= \frac{\Delta c_{t+i} - \tilde{\mu}_{t+i-2} - (\nu_1+(\nu_2-\nu_1)g(t))(\Delta c_{t+i-1} - \tilde{\mu}_{t+i-2})}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}} \\
&= \frac{\Delta c_{t+i} - \tilde{\mu}_{t+i-2}}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}} - \frac{(\nu_1+(\nu_2-\nu_1)g(t+i-1))(\Delta c_{t+i-1} - \tilde{\mu}_{t+i-2})}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}} \\
&= \tilde{\epsilon}_{t+i-1} - \frac{(\nu_1+(\nu_2-\nu_1)g(t+i-1))(\Delta c_{t+i-1} - \tilde{\mu}_{t+i-2})}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i-1)}}
\end{aligned} \tag{24}$$

so

$$\tilde{\epsilon}_{t+i+1} = \tilde{\epsilon}_{t+i} - \frac{(\nu_1+(\nu_2-\nu_1)g(t+i))(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i)}}$$

$$\begin{aligned}
\tilde{cov}_t(\tilde{\epsilon}_{t+i}, \tilde{\epsilon}_{t+i+1}) &= \widetilde{E}_t[\tilde{\epsilon}_{t+i}\tilde{\epsilon}_{t+i+1}] - \widetilde{E}_t[\tilde{\epsilon}_{t+i}]\widetilde{E}_t[\tilde{\epsilon}_{t+i+1}] \\
&= \widetilde{E}_t\left[\tilde{\epsilon}_{t+i}\left(\tilde{\epsilon}_{t+i} - \frac{(\nu_1+(\nu_2-\nu_1)g(t+i))(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i)}}\right)\right] - \widetilde{E}_t[\tilde{\epsilon}_{t+i}]^2 \\
&= \widetilde{E}_t[\epsilon^2_{t+i}] - \frac{(\nu_1+(\nu_2-\nu_1)g(t+i))\widetilde{E}_t[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+i}]}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i)}} - \widetilde{E}_t[\tilde{\epsilon}_{t+i}]^2 \\
&= -\frac{(\nu_1+(\nu_2-\nu_1)g(t+i))}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i)}}.\widetilde{E}_t[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+i}] \\
&= -\frac{(\nu_1+(\nu_2-\nu_1)g(t+i))}{\sigma\sqrt{1+\nu_1+(\nu_2-\nu_1)g(t+i)}} < 0
\end{aligned} \tag{25}$$

where we used the fact that $\widetilde{E}_t\left[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\epsilon_{t+i}\right] = \widetilde{E}_t[\Delta c_{t+i} - \tilde{\mu}_{t+i-1}]\widetilde{E}_t[\epsilon_{t+i}] = 0$ due to the orthogonality of the forecast error and the realized shock. The justification is that even if $\widetilde{E}_t\left[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+i}\right] = 0$ for all $i = 0$, the term $-\frac{\nu_1+(\nu_2-\nu_1)g(t+i)}{1+\nu_1+(\nu_2-\nu_1)g(t+i)}$ in the expression of $\tilde{cov}t(\tilde{\epsilon}t+i, \tilde{\epsilon}_{t+i+1})$ does not necessarily vanish. This is because this term is the product of the gains ν_1 and ν_2, which are constants, with the time-varying smooth transition function $g(t+i)$. As long as $g(t+i)$ is not constant and different from zero, this term will not collapse to zero even if the cross-covariances $\widetilde{E}_t\left[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+j}\right]$ are all equal to zero.

The reason why the entire $\tilde{cov}_t(\tilde{\epsilon}_{t+i}, \tilde{\epsilon}_{t+i+1})$ does not collapse to zero, even if $\widetilde{E}_t[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+i}] = 0$ for all i, is due to the persistence of the perception shocks. Even though the current and future shocks may not be correlated with the current perceived consumption growth, the shocks will still affect the future perceived consumption growth through their persistence.

In other words, the perception shocks have a persistent effect on the perceived consumption growth, and this persistence is captured by the auto-covariance structure of the standardized unexpected endowment growth. As a result, the shocks will affect the future perceived consumption growth even if they are uncorrelated with the current perceived consumption growth. This is why the entire auto-covariance structure of the standardized unexpected endowment growth does not collapse to zero, even if $\widetilde{E}_t[(\Delta c_{t+i} - \tilde{\mu}_{t+i-1})\tilde{\epsilon}_{t+i}] = 0$ for all i.

D Proving the analytical expressions of the shocks variances σ_ξ^2 and σ_ζ^2

Now, to confirm that these expressions result in a time-t predictive distribution of Δc_{t+j} in the time-varying gain learning model that matches the predictive distribution in the equivalent full-memory model, we need to substitute these expressions into the full-memory model's predictive distribution:
But we need to find the expression for ω^2. Using (K.5) and (K.6), we get:

$$\omega^2(t) = K(t)(\sigma_\zeta^2 + \sigma_\xi^2) - \sigma_\zeta^2. \tag{26}$$

Now we can substitute $K(t)$, $\sigma_\xi^2(t)$, and σ_ζ^2 into the equation for $\omega^2(t)$:

$$\omega^2(t) = [\nu_1 + (\nu_2 - \nu_1)g(t)]\sigma^2 - [\nu_1 + (\nu_2 - \nu_1)g(t)]^2\sigma^2. \tag{27}$$

Now, we can substitute our expressions for $\omega^2(t)$, σ_ζ^2, and $\sigma_\xi^2(t)$ into the predictive distribution (K.7):

$$\begin{aligned}\Delta c_{t+1} \sim \mathcal{N}\big(&\hat{\mu}_{t+1|t}, [\nu_1 + (\nu_2 - \nu_1)g(t)]\sigma^2 \\ &- [\nu_1 + (\nu_2 - \nu_1)g(t)]^2\sigma^2 + (1 + [\nu_1 + (\nu_2 - \nu_1)g(t)]^2)\sigma^2 \\ &+ (1 - [\nu_1 + (\nu_2 - \nu_1)g(t)]^2)\sigma^2\big).\end{aligned} \tag{28}$$

Now, we simplify the variance part:

$$\Delta c_{t+1} \sim \mathcal{N}(\hat{\mu}_{t+1|t}, \sigma^2). \tag{29}$$

So, the expressions for $K(t)$, $\sigma_\xi^2(t)$, and σ_ζ^2 result in a time-t predictive distribution of Δc_{t+j} in the time-varying gain learning model that matches the predictive distribution in the equivalent full-memory model.

E Proving that $K(t)$, $h(t)$, and $\sigma_\xi^2(t)$ result in a time-t predictive distribution of Δc_{t+j}

To confirm that the expressions for $K(t)$, $h(t)$, and $\sigma_\xi^2(t)$ result in a time-t predictive distribution of Δc_{t+j} in the time-varying gain learning model that matches the predictive distribution in the equivalent informative prior full-memory model, let's substitute these expressions into the full-memory model's predictive distribution.
Recall the predictive distribution for the informative prior full-memory model (K.15):

$$\Delta c_{t+1} \mid H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h^2\omega^2 + \sigma_\zeta^2 + \sigma_\xi^2).$$

Now, we can substitute our expressions for $h(t)$, $\omega^2(t)$, σ^2_ζ, and $\sigma^2_\xi(t)$ into the predictive distribution:

$$\Delta c_{t+1} \mid H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h(t)^2\omega(t)^2 + \sigma(t)^2_\zeta + \sigma(t)^2_\xi).$$

This distribution represents the time-t predictive distribution of Δc_{t+j} in the time-varying gain learning model. By substituting the expressions for the parameters, we have derived a distribution that matches the predictive distribution in the equivalent informative prior full-memory model, which confirms that the expressions for $K(t)$, $h(t)$, and $\sigma^2_\xi(t)$ result in the desired time-t predictive distribution.
Following Hansen, Heaton, and Li (2008), we start with value function iteration

$$v_t = \frac{\delta}{1-\gamma} log\widetilde{E}_t[e^{(1-\gamma)(v_{t+1}+\Delta c_{t+1})}], \tag{30}$$

where $v_t = \log(Vt/Ct)$ and V_t is the continuation value. We conjecture the solution to be linear in the state variable, i.e.

$$v_t = \mu_v + U_v\tilde{\mu}_t. \tag{31}$$

Plugging in the conjectured solution we get

$$\mu_v + U_v\tilde{\mu}_t = \frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(\mu_v+U_v\tilde{\mu}_{t+1}+\tilde{\mu}_t+\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}] \tag{32}$$

$$\begin{aligned}
\mu_v + U_v\tilde{\mu}_t &= \frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(\mu_v+U_v\tilde{\mu}_{t+1}+\tilde{\mu}_t+\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}] \\
&= \frac{\delta}{1-\gamma}\log\widetilde{E}_t\left[e^{(1-\gamma)(\mu_v+U_v(\tilde{\mu}_t+\nu_{t+1}\sigma\sqrt{1+\nu}\tilde{\epsilon}_{t+1})+\tilde{\mu}_t+\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1}}\right] \\
&= \frac{\delta}{1-\gamma}\log\widetilde{E}_t\left[e^{(1-\gamma)(\mu_v+(1+U_v)\tilde{\mu}_t+\nu_{t+1}\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1}U_v+\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}\right] \\
&= \frac{\delta}{1-\gamma}\log\widetilde{E}_t\left[e^{(1-\gamma)(\mu_v+(1+U_v)\tilde{\mu}_t+(1+\nu_{t+1}U_v)\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}\right] \\
&= \delta[\mu_v + (1+U_v)\tilde{\mu}_t] + \frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(1+\nu_{t+1}U_v)\sigma\sqrt{1+\nu_{t+1}}\epsilon_{t+1})}] \\
&= \delta\mu_v + \frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(1+\nu_{t+1}U_v)\sigma\sqrt{1+\nu_{t+1}}\epsilon_{t+1})}] + \delta(1+U_v)\tilde{\mu}_t
\end{aligned} \tag{33}$$

By identification,

$$\begin{gathered}
U_v = \delta(1+U_v) \\
\text{Thus, we get: } U_v = \frac{\delta}{1-\delta}
\end{gathered} \tag{34}$$

$$\begin{aligned}
\mu_v &= \delta\mu_v + \frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(1+\nu_{t+1}U_v)\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}] \\
\mu_v &= \frac{1}{1-\delta}\frac{\delta}{1-\gamma}\log\widetilde{E}_t[e^{(1-\gamma)(1+\nu_{t+1}U_v)\sigma\sqrt{1+\nu_{t+1}}\tilde{\epsilon}_{t+1})}] \\
\mu_v &= \frac{1}{2}(1-\gamma)U_v(\nu_{t+1}U_v+1)^2(1+\nu_{t+1})\sigma^2.
\end{aligned} \tag{35}$$

Now, let's turn to the stochastic discount factor (SDF). In our scenario, it represents the factor by which the agent is willing to exchange future payoffs for current ones, under uncertainty. In the case of Epstein-Zin preferences, the SDF takes the following form:

$$\begin{aligned} m_{t+1|t} &= \text{log}\left(\delta \frac{C_t}{C_{t+1}} \frac{V_{t+1}^{1-\gamma}}{\widetilde{E}_t[(V_{t+1})^{1-\gamma}]}\right) \\ &= \text{log}\delta - \Delta c_{t+1} + (1-\gamma)\text{log}(V_{t+1}) - \text{log}\widetilde{E}_t[(V_{t+1})^{1-\gamma}] \\ &= \text{log}\delta - \Delta c_{t+1} + (1-\gamma)(v_{t+1} + c_{t+1}) - \text{log}\widetilde{E}_t[e^{(1-\gamma)(v_{t+1}+c_{t+1})}] \\ &= \tilde{\mu}_{m,t} - \tilde{\mu}_t - \xi_t \sigma \tilde{\epsilon}_{t+1}, \end{aligned} \tag{36}$$

As g(t) is a deterministic and known function, the time-varying expressions of the component of $m_{t+1|t}$ are as the following:

$$\tilde{\mu}_{m,t} = log\delta - \frac{1}{2}(1-\gamma)^2(\nu_1 - (\nu_2 - \nu_1)g(t)U_v + 1)^2(1 + \nu_1 - (\nu_2 - \nu_1)g(t))\sigma^2 \tag{37}$$

$$\xi_t = [1 - (1-\gamma)(\nu_1 - (\nu_2 - \nu_1)g(t)U_v + 1)]\sqrt{\nu_1 - (\nu_2 - \nu_1)g(t)} \tag{38}$$

F Determination of Values ν_1 and ν_2

The objective of our analysis was to ascertain two specific values, ν_1 and ν_2, associated with inflation data by scrutinizing various data windows and pinpointing breakpoints. The methodology is delineated below:

1. Initial Setup:
   - To ensure reproducibility, a seed of 602 is fixed.
   - A window size of 40 observations is set, complemented by a minimum segment size of 10 for internal segments.
2. Preparation for Iterative Data Analysis:
   - Vectors are initialized to accumulate the computed gamma values for each data segment and their associated dates.
3. Window-Based Analysis:
   - The dataset is looped iteratively, taking 40 observations in each window.
   - Within each window, a simple linear regression model is employed:

     $$\text{Learning Gain} = \gamma_0 + \epsilon$$

   - As we are primarily interested in detecting structural breaks in the data, this simple regression with just a constant is adequate. The goal in this context is not to explain variations in the dependent variable but rather to identify points where the mean level of the series shifts. By focusing on the constant, we are targeting changes to the average value over different windows, making it easier to spot these shifts.

- Within this model, breakpoints are discerned, segmenting the data to reveal distinct behavioral patterns.
- If breakpoints are absent, have missing data, or are invalid, the iteration skips forward.
- If valid breakpoints are present, the gamma (coefficient) values for each segment are extracted and stored.

4. Consolidation of Findings :

   - Gamma values, paired with their relevant dates, are combined into a dataframe.
   - Entries with absent gamma values are purged from this dataframe.

5. Identification of Extreme Gamma Points:

   - The gamma values of utmost significance, which correspond to the highest and lowest values, are identified within the dataframe.
   - These distinct values represent our sought-after ν_1 and ν_2.

Upon completion of this methodology, we established that $\nu_1 = 0.01858007$ and $\nu_2 = 0.08162585$.