

How Income Risks Affect Financial Markets

An Empirical Analysis of Three Transmission Channels

Dissertation
zur Erlangung des Doktorgrades
der Wirtschafts- und Sozialwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen

vorgelegt von
Tobias Langen
aus Steinfurt

Tübingen
2015

Tag der mündlichen Prüfung: 21.10.2015

Dekan: Professor Dr. rer. soc. Josef Schmid

1. Gutachter: Professor Dr. rer. pol. Joachim Grammig

2. Gutachter: Professor Dr. rer. pol. Martin Biewen

Contents

Lists of Tables and Figures	V
1 Introduction	1
2 Empirical Asset Pricing with Large Idiosyncratic Income Shocks	5
2.1 Idiosyncratic Risk in Asset Pricing	6
2.2 Measuring Idiosyncratic Income Risk	9
2.2.1 Krebs' (2004) Idiosyncratic Income Risk Specification	9
2.2.2 The IR-Factor as Measure of Personal Disaster Risk	11
2.2.3 The IR-Model Stochastic Discount Factor	14
2.3 Data Description	16
2.4 Results and Discussion	18
2.4.1 IR-Factor Series Estimation	18
2.4.2 IR-Model Estimation and Model Comparison	24
2.5 Concluding Remarks	31
Appendix A	33
3 Empirical Asset Pricing with Reference-Dependent Heterogeneous Agents	45
3.1 Introduction	46

3.2	Model Setup and Methodology	49
3.2.1	Reference Grouping	49
3.2.2	Estimation Methodology	53
3.3	Data Description	56
3.4	Results and Discussion	60
3.4.1	GMM Estimation Results	60
3.4.2	Robustness	66
3.5	Concluding Remarks	68
4	A Cross-Country Analysis of Unemployment and Bonds with Long-Memory Relations	71
4.1	Introduction	72
4.2	Methodology	76
4.2.1	Nowcasting with Long-Run Relations	76
4.2.2	Unemployment and Bonds – Heterogeneous VAR	78
4.3	Data Description	78
4.3.1	Unemployment Rates and Bond Yields	78
4.3.2	Google Search Volume	82
4.4	Results and Discussion	87
4.4.1	Nowcasting Unemployment – M-HAR Estimation	87
4.4.2	Unemployment and Bonds – Heterogeneous VAR Estimation	96
4.5	Concluding Remarks	102
	Appendix B	103
5	Conclusion	105
	Bibliography	107

List of Tables

2.1	Descriptive Statistics – Personal Income Shock Series	17
2.2	Time Series and Cross-Sectional Least Squares Estimates	25
2.3	Model Comparison – GMM Estimates	29
2.4	Model Comparison – GMM Estimates (Full)	40
2.5	Model Comparison – GMM Estimates (Equally-Weighted Portfolios) .	41
2.6	GMM Estimates for Different Types of Central Moments	42
3.1	Descriptive Statistics – Income Growth and Factors	58
3.2	Monthly Average Fama-French Portfolio Excess Returns	59
3.3	GMM Estimation Results	61
3.4	GMM Estimation Results – Moment Condition Comparison	63
4.1	Descriptive Statistics – Unemployment Changes and Bond Yields . .	81
4.2	List of Search Terms and Descriptive Statistics	84
4.3	Nowcasting Unemployment – Estimation Results	92
4.3	Nowcasting Unemployment – Estimation Results (cont.)	93
4.4	HVAR – Estimation Results	103
4.4	HVAR – Estimation Results (cont.)	104

List of Figures

2.1	Time Periods Covered by the Nine SIPP Panels	16
2.2	Tail Approximation – Goodness of Fit	19
2.3	Structural Breaks in the Estimated IR-Factor Series (0.96)	21
2.4	IR-Factor Series (0.96) Corrected for Structural Breaks	23
2.5	Model Comparison – Goodness of Fit	27
2.6	Example of Peaks over Threshold Selection	36
2.7	Model Comparison – Goodness of Fit (Robustness Comparisons) . . .	43
3.1	Reference Grouping: Constructing Representative Agents A and B . . .	51
3.2	Time Structure of the SIPP Panles	56
3.3	Income Index Over Last Two Recessions	57
3.4	Model Comparison – Goodness of Fit (Value-Weighted Portfolios) . .	64
3.5	Model Comparison – Goodness of Fit (Equally-Weighted Portfolios) .	65
3.6	Robustness – Estimates are Stable for Different Group Quantiles . . .	67
4.1	Unemployment Rates and Ten-Year Bond Yields	80
4.2	Changes in Search Volume for "Jobs" and "Unemployment Benefits" .	86
4.3	Correlation of Principal Components and Keyword Categories	88
4.4	Cross-Correlation Peak Size and Timing	89
4.5	Cross-Correlation – Unemployment Changes and Principal Component	90
4.6	Scatter Plots of MSE and R^2	94
4.7	AR and M-HAR Nowcasts	96
4.8	Lag Structure	97
4.9	Impulse-Response Functions of Bond Yields	99
4.10	Scatter Plot of Impulse-Response Peak and Timing	100
4.11	Impulse-Response Functions of Unemployment	101

Chapter 1

Introduction

New job opportunities have the potential to drastically increase available income – perhaps more so than a good year at the stock market. In the US, 62% of personal income is derived from wages, as opposed to a mere 4.8% from capital gains and 3.7% from interest and dividends.¹ In the same way that a job promotion can positively affect available income, being laid off can immediately erode it. It seems likely that job stability plays an important role in personal investment decisions. An unstable job market may influence the willingness of investors to bear additional risks at the stock market, while a stable job market could alleviate concerns. Likewise, personal income determines a large fraction of a country's income. 81% of US tax revenue is generated from personal income taxes. Thus, overall unemployment is likely to be an important factor when investors assess the solvency of a country. This evaluation influences the risk compensation that investors demand.

This thesis analyzes how financial markets are affected by income risk and unemployment. I focus on two sectors of financial markets and evaluate three potential transmission channels. To model how asset returns are affected by individual income risk, two possible approaches are developed in chapters 2 and 3, while chapter 4 examines how government bond yields are affected by aggregate income risk in the form of unemployment changes. The first and more traditional approach models

¹Statistics are taken from the Federal Reserve's Survey of Consumer Finances and the Office of Management and Budget's Fiscal Year Report. Reported values are from 2013.

individual income risk as a factor in a linear asset pricing model. With this model I examine how income risk influences the willingness of investors to take risks and the premium different types of portfolios have to pay in order to compensate for the portfolio's co-movement with the risk factor. The second approach considers psychological biases that influence how investors perceive income changes. Specifically, I test for the presence of behavioral effects that may arise from the exposure to individual income risk. A third study shifts the focus to a macroeconomic level. Here the relationship between government bond yields and unemployment is analyzed in a high-frequency heterogeneous vector-autoregression.

Large Idiosyncratic Income Shocks

In chapter 2, the relevance of large idiosyncratic income shocks for asset prices is evaluated in a linear factor model. The channel through which income risk is linked to asset returns is described by Constantinides and Duffie (1996): exposure to income risk renders the situation more risky for investors, especially during a recession. In contrast to other empirical studies, I focus on the tails of the income shock distribution because events like job loss, divorce, or an offer from a prospective employer typically lead to large income changes. How dangerous are the negative events? How advantageous are the positive ones? To determine their importance, I develop a tail risk measure that identifies the factor by which cross-sectional income changes exceed a high quantile q on average at each point in time. The factor can be computed for large positive and large negative income changes. I use the tail income risk factors in a linearized Consumption-Based Model estimation on a cross-section of 25 portfolios sorted by size and book-to-market ratio. The results show that the introduction of large idiosyncratic income risk helps explain cross-sectional return differences. Both the upper tail and the lower tail idiosyncratic risk component are priced. To compensate for times of market turmoil, when the lower tail risk is high, large-value stocks have to offer an additional expected return of 8% p.a. in order to remain attractive for investors that are hit by large negative income shocks. Conversely, the expected annual return of large-value stocks is reduced by 4.5%, as compensation for prosperous market conditions, when the chance for an upper tail

shock is high and investors are willing to pay for the opportunity to smooth their consumption over time.

Reference-Dependent Heterogeneous Agents

In chapter 3 I examine psychological effects that are observed when individuals are exposed to income risk. These behavioral effects represent another link between income risk and financial markets. The approach draws on research that assesses limited rationality in decisions under risk by Kahneman and Tversky (1979). The authors devise a framework under which individuals display different attitudes towards risk. More precisely, individuals make decisions with respect to a reference level, where potential relative gains are devalued while potential relative losses are amplified. Ultimately, the location of a potential outcome in relation to the reference level decides whether the agent behaves risk-seeking or risk-averse. Experimental evidence for such psychological biases that influence individuals' decisions is plentiful. The study in chapter 3 presents empirical evidence that psychological biases have an effect on asset returns in a non-experimental setting. I use US panel data to sort individuals into a group that is above or a group that is below their reference level, depending on their recent income development. Representative agents are formed within these groups by capturing the average income development of investors above or below their reference level. According to prospect theory, these representative agents should display different attitudes towards risk. Indeed, the estimation results indicate that investors located below their reference level are risk-seeking while investors above their reference level are risk-averse. Furthermore, by allowing for risk-seeking behavior, the cross-sectional variation in returns of portfolios sorted by size and book-to-market value can be explained with reasonably low risk aversion coefficients, and a large fraction of the high average premium awarded for holding risky assets can be explained.

Unemployment Impacting Government Bond Yields

Chapter 4 investigates the relationship between unemployment and government bond yields. Economic theory suggests that a combination of a high level of government debt and low tax revenue due to a high unemployment rate may increase the default risk of a country. This should be reflected in government bond yields and constitutes yet another potential link between income risk and financial markets. The chapter is based on joint work with Dr. Thomas Dimpfl. We describe a general method to increase the frequency of unemployment data from monthly to weekly using a mixed frequency heterogeneous autoregressive model. Our method combines methodologies from the literature on mixed-frequency nowcasting (Marcellino and Schumacher, 2010) and heterogeneous autoregressions (Corsi, 2009). The approach exploits additionally available high-frequency data that are related to the original series. We show that to this end Google search query data can successfully be employed to nowcast unemployment changes. The resulting weekly unemployment time series is then used in a heterogeneous vector-autoregression of unemployment changes and bond yields. For a sample of seven European countries we consistently find that bond yields react positively to a rise in unemployment, while for the United States and Australia this effect is negative. Shocks to bond prices barely have any impact on unemployment.

Chapter 5 summarizes the key findings of the three approaches to analyze the relationship between unemployment, income risk and financial markets, outlines the main contributions of each chapter, and draws a conclusion.

Chapter 2

Empirical Asset Pricing with Large Idiosyncratic Income Shocks*

Abstract

In this chapter, I present evidence that large individual income changes can help explain the size and value premium in a cross-section of portfolio returns. I develop a tail risk measure, the tail income risk factor and estimate it based on US income data. In an augmented Consumption-Based Asset Pricing Model, the tail income risk factor emerges as a priced factor when explaining a cross-section of returns of 25 portfolios sorted by size and book-to-market ratio. Large-value stocks compensate for lower tail idiosyncratic income risk exposure by offering an additional expected return of 8% p.a. Conversely, large-value stocks compensate for upper tail idiosyncratic income risk exposure with an expected return reduction of 4.5% p.a. My findings support Krebs' (2004) critique of the Constantinides and Duffie (1996) idiosyncratic risk asset pricing model that central moments of the cross-sectional distribution of income cannot be used to test the implications of the Constantinides and Duffie (1996) model. The results back the notion of a fat-tail-generating personal disaster process affecting asset prices.

*Chapter 2 is based on the paper “Tracing Tails – Large Idiosyncratic Income Shocks in a Heterogeneous Agent Asset Pricing Model” by Langen (2013). I thank S. Bryzgalova, T. Dimpfl, J. Grammig, L. Huergo, S. Jank, R. Jung, F. Peter, and W. Pohlmeier, as well as participants of the Annual Meeting of the German Statistical Society (Berlin) and the European Meeting of the Econometric Society (Toulouse) for helpful comments and suggestions. Financial support from the German Research Foundation (DFG) is gratefully acknowledged.

2.1 Idiosyncratic Risk in Asset Pricing

Developing and testing asset pricing models that provide a link between asset prices and the real economy is a major concern of the financial and econometric literature. In the canonical framework of neoclassical rational models, such as the Consumption-Based Model, risk-averse investors demand a premium for holding risky assets. However, the empirical performance of these models is largely disappointing (Hansen and Singleton, 1982; Constantinides and Ferson, 1991; Fama and French, 1992).

A possible reason may be that these models do not sufficiently account for the risks that investors have to bear. One of these omitted risks is the focus of this chapter. Concretely, I examine the impact of idiosyncratic income shocks on how investors evaluate their position with respect to the state of the economy and thus, the risk premium they demand for different kinds of securities. Following Constantinides and Duffie (1996), I allow for an investor heterogeneity that results from the exposure to idiosyncratic income shocks. These idiosyncratic shocks can affect the investment decisions of individuals. When an economy is hit by a recession, investors do not only face the risk of potential losses in the stock market, there is also the additional risk of considerable income losses. An asset pricing model that neglects the effects of idiosyncratic income risk is likely to underestimate the severity of the situation that investors find themselves in. If investors are to hold their financial assets throughout the recession, they will demand compensation in the form of higher expected returns. In contrast to other empirical studies, I focus on the tails of the cross-sectional income shock distribution because the events that cause idiosyncratic income shocks (job losses, accidents, sicknesses, divorces, but also promotions, and new job opportunities) typically lead to large changes in available income. Focusing on the tails also acknowledges Krebs' (2004) critique who shows that central moments of the cross-sectional distribution cannot be used to test the implications of idiosyncratic risk models. I develop a tail risk measure that enables the measurement of the personal disaster risk component as it is specified by Krebs (2004) and is also suitable to deal with more general cases.

Previous studies suggest that accounting for idiosyncratic income risk helps explain the high equity premium that investors demand for holding risky assets and the cross-sectional return variation of size and book-to-market sorted portfolios. Constantinides and Duffie (1996) demonstrate that idiosyncratic income risk has an impact on asset prices. They derive a closed-form solution for the stochastic discount factor, assuming log-normally distributed income shocks. In this special case, the variance of the shocks enters the stochastic discount factor as a measure for the idiosyncratic income risk. Brav et al. (2002), Balduzzi and Yao (2007), as well as Grishchenko and Rossi (2012) use household micro-data from the US Consumer Expenditure Survey to evaluate the prediction of Constantinides and Duffie (1996) that the variance of log-consumption plays a relevant role in the stochastic discount factor. All present evidence that the Constantinides and Duffie (1996) model is able to explain the equity premium with economically plausible parameters. Jacobs and Wang (2004) use a linearized version of the Constantinides and Duffie (1996) discount factor to test the Constantinides and Duffie (1996) model's ability to price the Fama-French portfolios. They conclude that the cross-sectional variance of log-consumption growth is a priced factor.

Although the empirical results seem encouraging, some major drawbacks call for reconsideration. First of all, measurement errors in the Consumer Expenditure Survey are a common problem (Altonji, 1986; Altonji and Siow, 1987; Zeldes, 1989). Measurement errors affect the variance of log-consumption growth and bias the results. Vissing-Jørgensen (2002) and Balduzzi and Yao (2007) demonstrate that a larger measurement error skews risk aversion estimates downwards. Secondly, assuming log-normal income growth does not allow for enough short-run income variability to produce an economically significant contribution to pricing the equity premium (Cochrane, 2008). Plausible income processes should exhibit small changes for extended periods, but jump abruptly to a new level following a shock. However, log-normality is essential in all mentioned empirical applications, as it justifies the use of the variance of log-consumption growth as a risk factor. Thirdly, Krebs (2004) shows that central moments of the cross-sectional distribution cannot be used to test the implications of the Constantinides and Duffie (1996) model in an economy where

investors face a personal disaster risk. He demonstrates that the observed central moments can be arbitrary once the possibility of large income shocks is introduced. This is a consequential result as the events that Constantinides and Duffie (1996) invoke to motivate their model – job loss or accidents – should predominantly entail large income changes. Krebs' (2004) income specification that allows for large, drastic shifts seems to be more realistic. However, it is incompatible with the use of cross-sectional variance as a risk factor. In the light of these observations, it is unclear how to evaluate the empirical evidence shown in the above mentioned studies that rely on Constantinides and Duffie's (1996) log-normal individual income process.

In this study, I propose a measure for idiosyncratic income risk that is motivated by the specification of the income process in Krebs (2004) and accounts for the importance of the personal disaster risk component. I suggest a partition of the observed cross-section of income that enables the measurement of the personal disaster component as the factor by which shocks exceed a high quantile q on average. I show that this tail income risk factor (IR-Factor) can quantify changes of the personal disaster risk variables in the Krebs (2004) income growth specification. The IR-Factor is also suitable for more general income growth processes, as it does not hinge on restrictive assumptions. I use the IR-Factor in an augmented, linearized Consumption-Based Model estimation on a cross-section of 25 portfolios sorted by size and book-to-market ratio. The results show that the introduction of a fat-tail-generating idiosyncratic risk process helps explain cross-sectional return differences. Both an upper tail idiosyncratic risk component as well as a lower tail idiosyncratic risk component are priced. To compensate for turbulent times, when the lower tail risk is high, large-value stocks have to offer an additional expected return of 8% p.a. in order to remain attractive for investors that are hit by large negative income shocks. Conversely, large-value stocks offer an expected annual return that is reduced by 4.5%, as compensation for prosperous times, when the upper tail chance is high and people hit by big positive income shocks want to smooth their consumption.

The remainder of the chapter is structured as follows. Section 2.2 introduces and discusses the IR-Factor in relation to the personal disaster risk process proposed by Krebs (2004) as well as possible generalizations and derives the stochastic discount factor of the IR-Model. Section 2.3 describes the personal income data from the U.S. Survey of Income and Program Participation as well as the test portfolios, section 2.4 discusses the estimation results, and section 2.5 concludes.

2.2 Measuring Idiosyncratic Income Risk

2.2.1 Krebs' (2004) Idiosyncratic Income Risk Specification

Krebs (2004) builds on Constantinides and Duffie (1996) who construct a full information economy in equilibrium, in which symmetrical investors with von Neumann-Morgenstern preferences are hit by permanent and uninsurable idiosyncratic income shocks and show that these idiosyncratic shocks have asset pricing implications. Income risk is quantified by the variance of log-consumption growth in a closed-form solution for the stochastic discount factor. The derivation hinges on the income innovation specification. Constantinides and Duffie (1996) assume that innovations are log-normally distributed. Krebs (2004) and Cochrane (2008) criticize this assumption, arguing that it cannot generate enough short-run variability. Krebs (2004) suggests a generalization of the restrictive income process specification, allowing for disastrous income shocks to happen with low probability. The consequence of this change in specification is that central moments of the cross-sectional distributions of income and consumption no longer provide testable restrictions.

The original individual income growth process specified by Constantinides and Duffie (1996), has the following form:

$$\frac{y_{i,t+1}}{y_{i,t}} = (1 + g_{1i,t+1})(1 + G_{t+1}), \quad (2.1)$$

where

$$\ln(1 + g_{1i,t+1}) = s_{1i,t+1}Z_{t+1} - \frac{Z_{t+1}^2}{2} \quad \text{and} \quad s_{1i,t+1} \sim \mathcal{N}(0, 1).$$

$y_{i,t+1}$ is the income of individual i at time $t + 1$, G_{t+1} is the average income growth rate and $g_{1i,t+1}$ is the individual deviation from the average income growth rate. The process $\{g_{1i}\}_{t=0}^{\infty}$ generates what Krebs (2004) calls observed cross-sectional moments. $s_{1i,t+1}$ is the individual income shock and Z_{t+1}^2 is the cross-sectional variance of individual income shocks. Both $Z_{t+1}^2(S_{t+1})$ and $G_{t+1}(S_{t+1})$ are assumed to be functions of an unobserved state variable S_{t+1} . Constantinides and Duffie's (1996) stochastic discount factor then depends on aggregate growth and the cross-sectional variance as sources of risk. If any other moments of the cross-sectional distribution of income growth were a function of the state variable S_{t+1} , they would influence the stochastic discount factor as well.

Krebs (2004) generalizes the income process by introducing a personal disaster risk component

$$\frac{y_{i,t+1}}{y_{i,t}} = (1 + g_{1i,t+1})(1 + g_{2i,t+1})(1 + G_{t+1}), \quad (2.2)$$

where $g_{1i,t+1}$ can be the process in equation (2.1), or any other observed moment generating process and $g_{2i,t+1}$ is a process that allows for rare extreme events to occur – events that cause substantial individual income changes. Krebs (2004) defines $g_{2i,t+1}$ as a random variable that can have two outcomes $-\eta(S_{t+1})$ and $\frac{p(S_{t+1})}{1-p(S_{t+1})}\eta(S_{t+1})$. $-\eta(S_{t+1})$ occurs with probability $p(S_{t+1})$. Both, size and probability of an extreme negative shock depend on the state variable. Krebs (2004) shows that the overall income process can have arbitrary cross-sectional moments and still fulfill the Euler equation, given some extreme disaster process. The proof relies on the two characteristics of the large income shock process, η and p . First, Krebs (2004) shows that the stochastic discount factor growth can be arbitrarily adjusted by varying $\{\eta\}_{t=0}^{\infty}$, given any fixed choice of central moments of income or consumption (determined by $\{g_{1i}\}_{t=0}^{\infty}$ and $\{G\}_{t=0}^{\infty}$).¹ Then, Krebs (2004) demonstrates that the central moments of the generalized income process in equation (2.2) can get arbitrarily close to the observed moment process for small $\{p\}_{t=0}^{\infty}$. Although $\{\eta\}_{t=0}^{\infty}$ entirely determines the stochastic discount factor, it cannot be detected in the observed moment process. Consequently, central moments of the cross-sectional distribution of income shocks

¹The adjustment can also be caused through any combination of changes in η and p . Changes in η , however, are sufficient.

cannot be used to test the model. All pricing relevant dynamics are driven by the large income shock component.

2.2.2 The IR-Factor as Measure of Personal Disaster Risk

Now, consider the following partitioned individual income process specification:

$$y_{i,t+1} = (1 + g_{1i,t+1})(1 + G_{t+1})y_{i,t} + g_{i,t+1}^+ \Delta y_{t+1}[q] + g_{i,t+1}^- \Delta y_{t+1}[1 - q], \quad (2.3)$$

where $\Delta y_{t+1}[q]$ is the upper q quantile of the cross-sectional income shock distribution in levels, while $\Delta y_{t+1}[1 - q]$ is the lower $1 - q$ quantile. The random variables $g_{i,t+1}^+$ and $g_{i,t+1}^-$ define the tails of the income shock distribution. Their properties are of particular interest. Subtracting the individual's current income $y_{i,t}$ from equation (2.3) yields the following absolute income shock specification:

$$\begin{aligned} \Delta y_{i,t+1} &= [(1 + g_{1i,t+1})(1 + G_{t+1}) - 1]y_{i,t} + g_{i,t+1}^+ \Delta y_{t+1}[q] + g_{i,t+1}^- \Delta y_{t+1}[1 - q] \\ &= g_{i,t+1}^\circ y_{i,t} + g_{i,t+1}^+ \Delta y_{t+1}[q] + g_{i,t+1}^- \Delta y_{t+1}[1 - q], \end{aligned} \quad (2.4)$$

where $g_{i,t+1}^\circ$ is a random variable that generates economically non-extreme income fluctuations – changes that are below the q quantile. The tail generating random variables $g_{i,t+1}^+$ and $g_{i,t+1}^-$ are zero with probability q and follow any distribution > 1 with probability $1 - q$. The exact functional form is left unspecified. A distribution with a lot of mass allocated close to one will create a light-tailed income shock distribution. If more mass is allocated away from one, the distribution will have heavy tails. The IR-Factors are estimates for the expected values of $g_{i,t+1}^+$ and $g_{i,t+1}^-$. Taking expectations across i yields

$$\mathbb{E}[\Delta y_{i,t+1}] = \mathbb{E}[g_{i,t+1}^\circ y_{i,t}] + \mathbb{E}[g_{i,t+1}^+] \Delta y_{t+1}[q] + \mathbb{E}[g_{i,t+1}^-] \Delta y_{t+1}[1 - q]. \quad (2.5)$$

The last two terms can be re-written as

$$\begin{aligned}\mathbb{E}[g_{i,t+1}^+] \Delta y_{t+1}[q] &= \mathbb{E}[\Delta y_{i,t+1} | \Delta y_{i,t+1} > \Delta y_{t+1}[q]] \\ \mathbb{E}[g_{i,t+1}^-] \Delta y_{t+1}[1-q] &= \mathbb{E}[\Delta y_{i,t+1} | \Delta y_{i,t+1} < \Delta y_{t+1}[1-q]].\end{aligned}\quad (2.6)$$

Rearranging leads to a specification for the expected value of the large tail generating random variables:

$$\begin{aligned}\mathbb{E}[g_{i,t+1}^+] &= \frac{\mathbb{E}[\Delta y_{i,t+1} | \Delta y_{i,t+1} > \Delta y_{t+1}[q]]}{\Delta y_{t+1}[q]} \\ \mathbb{E}[g_{i,t+1}^-] &= \frac{\mathbb{E}[\Delta y_{i,t+1} | \Delta y_{i,t+1} < \Delta y_{t+1}[1-q]]}{\Delta y_{t+1}[1-q]}.\end{aligned}\quad (2.7)$$

A convenient definition of these expectations can be derived by recognizing the numerator as the expected shortfall. The resulting expectations are defined as the IR-Factors:

$$\begin{aligned}IR_{t+1}^+ &\equiv \mathbb{E}[g_{i,t+1}^+] = \frac{ES_q^+[\Delta y_{i,t+1}]}{\Delta y_{t+1}[q]} \\ IR_{t+1}^- &\equiv \mathbb{E}[g_{i,t+1}^-] = \frac{ES_q^-[\Delta y_{i,t+1}]}{\Delta y_{t+1}[1-q]}.\end{aligned}\quad (2.8)$$

The IR-Factors consist of an expected shortfall component in the numerator and a quantile in the denominator. If the tails are light, observations will lie close to the quantile and thus the IR-Factor will be small. If the tails are heavy, many observations will be far above the quantile values, so the IR-Factors will be large.

In the special case of Krebs' (2004) income process specification, the lower tail IR-Factor simplifies to the following expression, as $1 - q$ approaches p

$$IR_{t+1}^- = \frac{p}{1-q} \frac{(1-\eta)\mathbb{E}[y_{i,t+1}] - \mathbb{E}[y_{i,t}]}{(\tilde{y}_{i,t+1} - y_{i,t})[1-q]},\quad (2.9)$$

where $(1-\eta)\mathbb{E}[y_{i,t+1}] - \mathbb{E}[y_{i,t}]$ is the expected size of the observed large shocks, and $(\tilde{y}_{i,t+1} - y_{i,t})[1-q]$ is the lower tail q -quantile of the shocks, generated by the non-extreme income change component, with $\tilde{y}_{i,t+1} = (1 + g_{i,t+1}^o)y_{i,t}$.² The small

²A derivation of equation (2.9) can be found in the appendix.

probability p with which the large shock $-\eta$ occurs makes it hard to detect these shocks in the observed central moments. However, this probability is inflated, when the quantile is pushed further into the tail of the distribution. As $1 - q$ approaches p , $\frac{p}{1-q}$ approaches one. When the quantile is large enough, the IR-Factor measures how large the extreme events are in relation to the q quantile of the non-extreme income changes.

The partition of the observed cross-section of income into extreme and non-extreme components (2.4) also enables the measurement of the personal disaster risk component for more general personal income process specifications. Krebs' (2004) specification only models one tail of the distribution. Nevertheless, if large negative income shocks can influence the stochastic discount factor, the same is true for large positive income shocks. This possibility is accounted for by including an additional upper tail risk component in the personal income specification. Furthermore, Krebs' (2004) specification allows for one size of large income shocks only ($-\eta$), which is overly restrictive. The partitioned income specification in (2.4) relaxes this assumption. Lastly, Krebs (2004) assumes that the distribution of large income shocks is the same across individuals and does not depend on current income. This is not necessarily true. Consider a case where investors face the risk of losing a major part of their income. A symmetric exposure to this risk would generate an even distribution of small income shocks, as small income individuals are hit and high income shocks, as high income individuals are hit. However, the economic situation might instead be such that low income individuals have a much higher probability of being hit by a shock. This asymmetric exposure would generate more small income changes than large ones. Conversely, if high income individuals are more exposed to income risk, a higher number of large income changes will be observed. The estimation of a risk component measure in the symmetric case could be performed using all income changes. In the asymmetric case, the exposure to the income risk and even its size can vary with the individuals' income. Consequently, the income size and the resulting shock size have to be taken into account. Partitioning the observed income process allows for a distinction between those times when high income individuals lose big portions of their income and those times when low income individuals are

more exposed, while the risk for high earners is relatively low. Both cases certainly have different economic implications.

The scarcity of large income shocks makes the estimation of the IR-Factor challenging. While the quantile of the distribution may be estimated by the empirical quantile with adequate precision, an approximation of the expected shortfall component as the sample mean above the quantile is biased downward. The bias is especially prominent for heavy tailed distributions. Extreme Value Theory can help avoid this problem (compare Embrechts et al. (1999), McNeil and Frey (2000)). It justifies fitting a Generalized Pareto distribution to the peaks over a fixed threshold (Davidson and Smith, 1990). The expected shortfall component can be consistently estimated from the parameters of a fitted Generalized Pareto distribution that traces the available tail observations. A detailed discussion can be found in the appendix.

2.2.3 The IR-Model Stochastic Discount Factor

The stochastic discount factor of the IR-Model is a modification of a linearized Constantinides and Duffie (1996) stochastic discount factor. To motivate this, consider the following specification

$$m_{t+1}^{CD} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} \exp \left[\frac{\gamma(\gamma + 1)}{2} Z_{t+1}^2 \right]. \quad (2.10)$$

This factor is derived by Constantinides and Duffie (1996) for log-normal individual income growth as specified in equation (2.1). Here, β is the time preference of the investor and γ the risk aversion coefficient. The cross-sectional variance of log-consumption growth Z_{t+1}^2 augments the Consumption-Based Model stochastic discount factor (Lucas, 1987; Breeden, 1979; Grossman and Shiller, 1981) as a measure for the idiosyncratic income risk that investors are exposed to. Following Cogley (2002) and Jacobs and Wang (2004), the heterogeneous agent stochastic discount factor is linearized by a Taylor approximation. Although this introduces an approximation error, there are two convincing advantages: no assumptions need to be made about the functional form of the investors' utility function and the estimation results of a linearized model can be compared straightforwardly to other linear

benchmark models, such as the Capital Asset Pricing Model (Sharpe, 1964; Lintner, 1965; Mossin, 1966) and the Fama-French Three-Factor Model (Fama and French, 1995). The linearized stochastic discount factor has the following form:

$$m_{t+1}^{L-CD} = b_0 + b_1 \Delta c_{t+1} + b_2 \text{Var} [\Delta c_{i,t+1}], \quad (2.11)$$

where Δc_{t+1} is consumption growth, as before, and $\text{Var} [\Delta c_{i,t+1}]$ is the cross-sectional variance of consumption growth. Since the goal of this study is to trace the influence of large cross-sectional income changes on the stochastic discount factor, consumption is approximated by income. Jacobs and Wang (2004) do the same to avoid mixing income and consumption variables

$$m_{t+1}^{L-JW} = b_0 + b_1 \Delta y_{t+1} + b_2 \text{Var} [\Delta y_{i,t+1}]. \quad (2.12)$$

Finally, the generalization of Krebs (2004) and the defining impact of the personal disaster process on the stochastic discount factor is accounted for by replacing the variance with the IR-Factors

$$m_{t+1}^{IR} = b_0 + b_1 \Delta y_{t+1} + b_2 IR_{q,t+1}^+ + b_3 IR_{q,t+1}^-, \quad (2.13)$$

where Δy_{t+1} is the overall income growth and $IR_{q,t+1}^+$ and $IR_{q,t+1}^-$ are the upper tail chance and the lower tail risk of the cross-sectional distribution of income shocks.³ The IR-Model is evaluated in comparison to the CAPM and the Fama-French Model. The unknown parameters are estimated via GMM (Hansen, 1982), using the unconditional moment restrictions implied by the basic pricing equation. The corresponding return-beta representations are evaluated via two-pass least squares regressions.

³The same stochastic discount factor can be motivated in the ICAPM framework of Merton (1973). However, the theory-driven approach outlined here seems more compelling, albeit its reliance on several approximations.

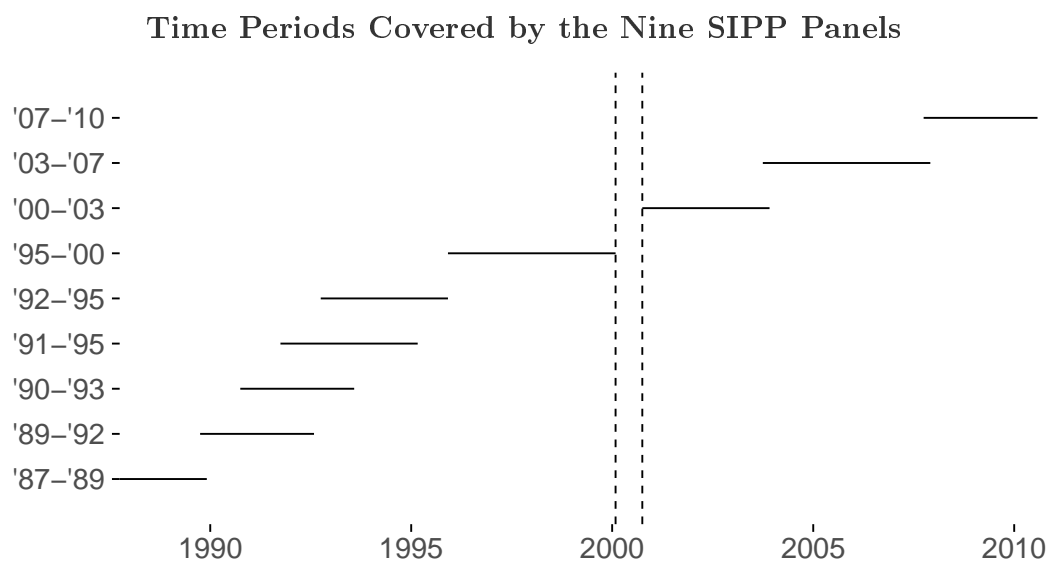


Figure 2.1: Time periods covered by each of the nine SIPP panels. The overlapping ends of the panel enable income shock computations at each point in time for at least one panel. Observations from overlapping panels are pooled. The time period between March and September 2000 is not covered by the SIPP and is excluded from the estimation.

2.3 Data Description

The Survey of Income and Program Participation⁴ (SIPP) is chosen in this study in favor of the Panel Study of Income Dynamics or the Current Population Survey because it allows for the computation of individual income data on a quarterly basis. The SIPP is conducted by the U.S. Census Bureau. It is a sequence of nine national representative panels, which record the labor income of households and of the adult individuals within each household. The panels cover a time period of almost 23 years, from October 1987 to August 2010. In the years before the 1996 panel, several smaller panels were in operation simultaneously. In a consolidation effort starting in 1996, only one larger panel was conducted at a time, with slight overlaps at the beginning and end of panels. Due to lack of funding, there is an observation gap between March and September 2000. Figure 2.1 depicts the time structure of the SIPP panels. Each panel covers the income of 30,000 to 100,000 adults. Using the panel structure, I generate monthly cross-sections of income shocks. Table 2.1 provides an overview of the income shock data. Each of the nine panels is composed

⁴www.census.gov/sipp

Descriptive Statistics – Personal Income Shock Series

Table 2.1: This table provides condensed descriptive statistics of the income shock series, generated from the nine SIPP panels. Reported personal labor income changes are in US Dollar. Each row provides the time period and length of the panel. All subsequent values refer to a median month. $\overline{\Delta y^-}$ and $\overline{\Delta y^+}$ list the cross-sectional average positive and negative income shocks, $\sigma_{\Delta y}$ is the cross-sectional variance of income shocks and $\Delta y[0.05]$ and $\Delta y[0.01]$ provide the cross-sectional, empirical 0.05 and 0.01 quantiles of income shocks. Finally, the number of cases in which no income change was reported are listed as well as the number of observations.

	Months	$\overline{\Delta y^-}$	$\overline{\Delta y^+}$	$\sigma_{\Delta y}$	$\Delta y[0.05]$	$\Delta y[0.01]$	Zeros	Obs.
Nov '87 - Dec '89	26	-842	784	1095	-1063	-3042	7048	15408
Nov '89 - Aug '92	34	-609	581	623	-500	-1512	33567	47538
Nov '90 - Aug '93	34	-641	623	670	-500	-1510	24607	35001
Nov '91 - Mar '95	41	-660	629	652	-500	-1580	34116	48223
Nov '92 - Dec '95	38	-664	638	683	-495	-1555	34577	48238
Jan '96 - Feb '00	50	-892	842	1164	-481	-2004	55709	76344
Nov '00 - Dec '03	38	-1058	1019	1197	-536	-2400	50300	67886
Nov '03 - Dec '07	50	-2592	2515	2752	-2755	-7053	41298	63668
Nov '07 - Aug '10	34	-2972	2907	2816	-3072	-8000	41746	61576

of different individuals, so it is not possible to compute income changes between panels. However, since the panels overlap, income shocks can be computed for each month for at least one of the panels. Observations covering the same months in different panels are pooled. Possible structural breaks, especially after the reorganization in 1996 and the funding break in 2000, are discussed and accounted for during the IR-Factor estimation.

The IR-Model is evaluated on 25 value-weighted quarterly portfolio returns. The portfolios are composed of assets that are sorted by market capitalization and book-to-market ratio. These range from small market value, low book-to-market ratio (small size, growth) to large market value, high book-to-market ratio (large size, value). The data can be downloaded from the homepage of Kenneth R. French, which also holds series for the Fama-French Factors (R^m , SMB and HML), as well as the one-month treasury bill data used to compute excess returns.⁵ Aggregate con-

⁵www.dartmouth.edu/~kfrench

sumption and income growth series are obtained from the homepage of the Federal Reserve Bank of St. Louis.⁶

2.4 Results and Discussion

2.4.1 IR-Factor Series Estimation

Unbiased IR-Factor estimates are obtained using Extreme Value Theory. Its application to the income shock data requires the determination of a threshold above which income shocks are considered extreme. If chosen correctly, the peaks above the threshold qualify as being approximately Generalized Pareto distributed. The threshold choice has to strike a balance between being sufficiently high, such that the observations can be considered extreme, and being sufficiently low, such that enough observations are available to estimate the parameters of the Generalized Pareto distribution with high precision. A suitable threshold region is identified with the aid of several goodness of fit measures. The first graph in Figure 2.2 shows the average root mean squared error (RMSE) of the upper and lower tail approximations for different thresholds. The values are computed by selecting different empirical income shock quantiles as a threshold (0.85 to 0.99). Income shocks above the respective thresholds are considered extreme and the peaks above the given threshold are fitted to a Generalized Pareto distribution. The RMSE of the approximation is computed in each quarter. The average RMSEs across all quarters are depicted, for each threshold quantile. One can see that the approximation of the tail of the empirical income shock distribution by the Generalized Pareto distribution improves, as the threshold moves closer towards the end of the empirical distribution. The approximation starts deteriorating again around the 0.96 quantile. A threshold in the region of the 0.95 quantile seems to be a reasonable choice. It is about as high as the threshold can be pushed into the tail of the distribution, while still maintaining a good model fit. The graph below depicts the RMSEs in each quarter for the threshold quantiles 0.85, 0.95 and 0.99. Choosing the threshold at the 0.95

⁶www.stlouisfed.org

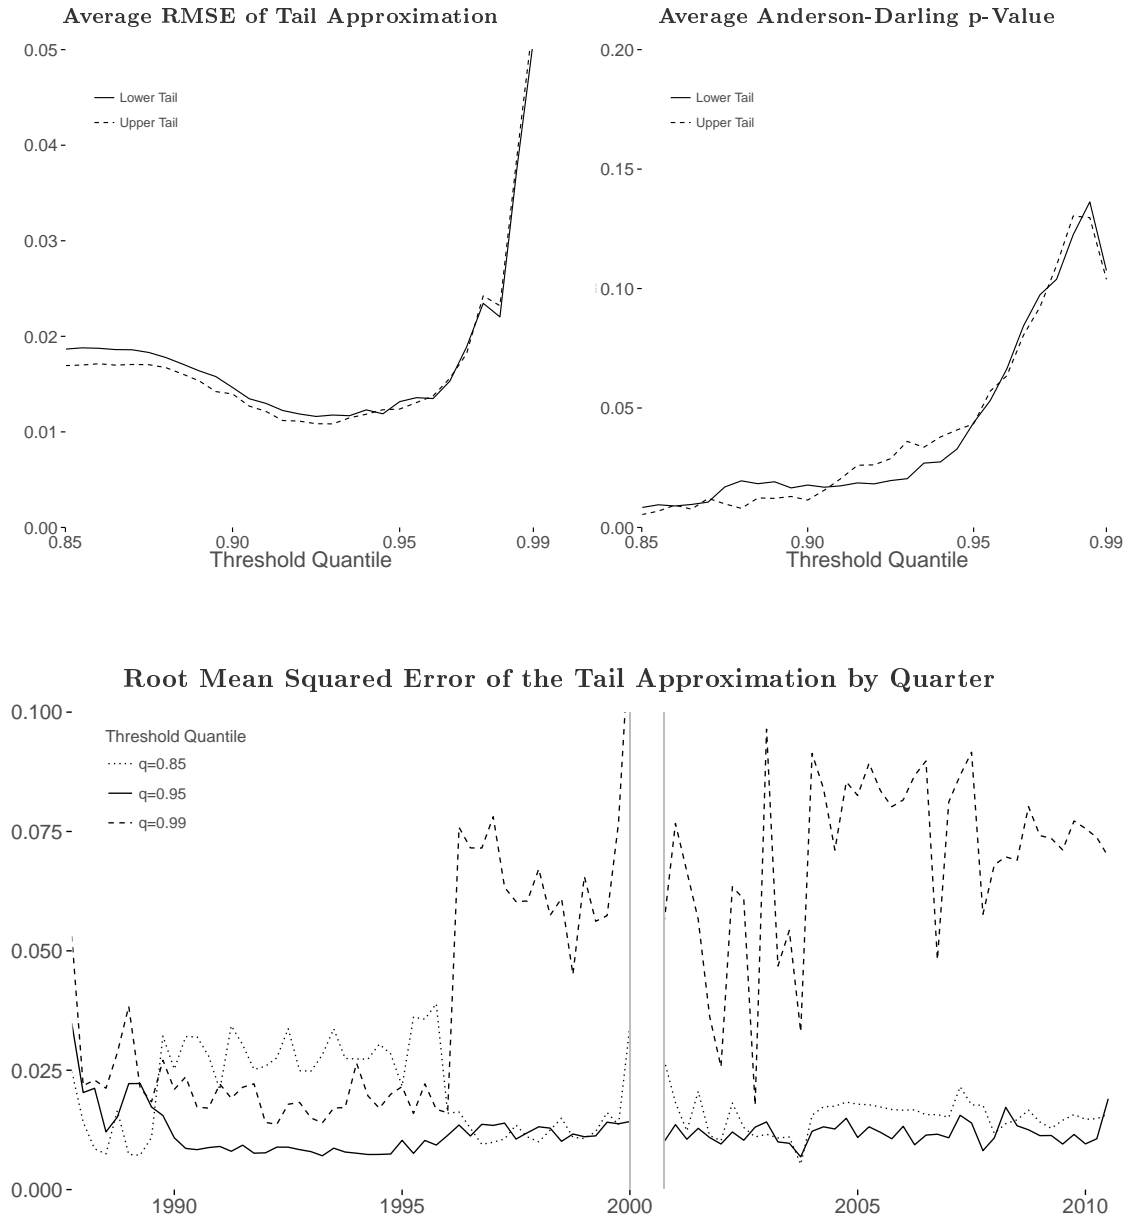


Figure 2.2: These plots represent several ways of evaluating the goodness of fit of the income shock tail approximation by the Generalized Pareto distribution. Ideally a threshold quantile area should be identified for which the fit is good and the approximation is valid. The upper left graph depicts the time-average root mean squared error of the upper and lower tail approximation for a range of threshold quantiles. The top right graph depicts the average Anderson-Darling p-value for a range of threshold quantiles. The bottom graph shows the time series of root mean squared errors for three threshold quantiles.

quantile provides a low RMSE throughout the whole observation period. The last graph in Figure 2.2 depicts the average p-value of an Anderson-Darling test with the null hypothesis being that the observed values are generated by a Generalized Pareto distribution. The further the threshold is pushed into the tail, the higher the p-values get. If feasible, a threshold should be chosen for which the null is not rejected in any quarter.⁷

These observations motivate the following estimation strategy: IR-Factor series are estimated for three quantiles (0.95, 0.96 and 0.97). To estimate these factor series, I use the Peaks over Threshold method, where the threshold is selected according to a quantile that is 0.01 below the estimated IR-Factor quantile. Since the threshold quantile is always below the designated IR-Factor quantile, each IR-Factor estimation is feasible. At the same time the threshold quantile is pushed as high as possible. From an economic perspective, IR-Factor quantiles around 0.95 seem to be a meaningful choice. Estimating the IR-Factors for a range of quantiles serves as a robustness check. If a threshold quantile is too low and the selected income shocks do not qualify as extreme events, the estimation results will not be robust to the threshold selection. If the values do qualify as extreme events, then varying the threshold should only influence efficiency. The higher-than-necessary threshold simply reduces the number of observations from the underlying Generalized Pareto distribution that can be used for estimation. Furthermore, if idiosyncratic income risk is priced, then it would be surprising if an effect only appeared at a specific quantile. Looking at a range of IR-Factor quantiles ensures not to be misled by an odd draw. All three IR-Model specifications should deliver sensible and comparable results.

IR-Factor series are constructed using the Generalized Pareto distribution parameter estimates. Figure 2.3 shows the upper and lower tail IR-Factor series for the 0.96 quantile alongside its 0.95 confidence interval. Two structural breaks in the series' average are apparent. The shifts coincide with the beginning of new panels

⁷It should be noted that as the threshold rises, the number of peaks naturally decreases, causing the Anderson-Darling test to lose power and reject less frequently. However, the number of peaks at the 0.96 quantile is still rather large, with an average of 7611 observations for the lower tail and 7628 for the upper tail.

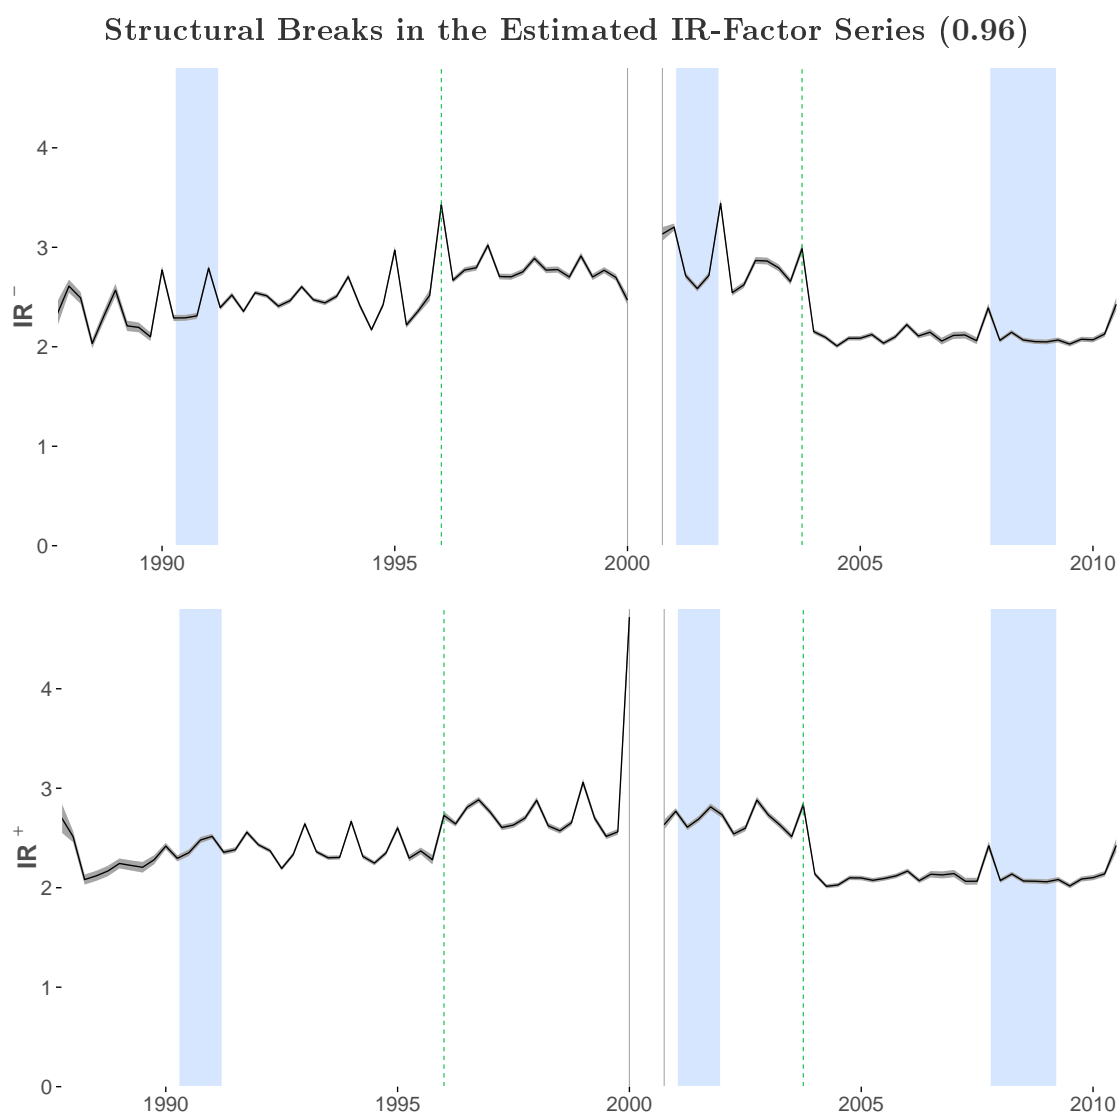


Figure 2.3: Estimated 0.96 IR-Factor time series for the lower tail (IR^-) and the upper tail (IR^+) along its 0.95 confidence band. There is a two quarter gap (2000 Q2, 2000 Q3). The blue shaded areas mark recession periods. Dashed lines indicate the start of the 1996 and the end of the 2000 panel.

and organizational as well as procedural changes within the Census Bureau's organization of the SIPP. The first clear break occurs at the beginning of the sixth panel (Q1 1996) and the second at the beginning of the eighth panel (Q4 2003). The 1996 panel is the first panel that does not have major overlaps with other running panels and features a sample size that is nearly twice as large as all the previous panels. It is the outcome of a general move to better pool and target available resources. The 2003 panel is the first full-fledged panel after the 2000 funding failure. The structural changes in the IR-Factor series might be caused by differences in the ability to avoid non-responses after a big income change. Individuals that experience a large negative income shock might be inclined to drop out of the panel in order to focus on readjusting after the shock. Individuals that experience a large positive income shock might have to move to a new job location and are possibly lost. Both these effects lead to an underestimation of the IR-Factor. This downward bias can be weakened or amplified by the financial and organizational ability of the Census Bureau to track down those cases. I account for the structural breaks by demeaning the IR-Factor series. The resulting series are depicted in Figure 2.4.

The lower tail IR-Factor series seems to display a higher volatility in the first two recession periods (depicted as blue shaded areas). In the last recession there is only one higher peak at the beginning of the period. The upper tail IR-Factor does not seem to be especially volatile in the first two recession periods, but is notably similar to the lower tail IR-Factor in the last recession period. The recessions appear to be structurally different. The first two recessions feature several waves of increases in the lower tail risk, while the upper tail risk remains stable. The last recession features just one increase in both upper and lower tail risk. In the first quarter of 2001, just before the peak of the dot-com bubble on March 10th, the upper tail risk rises considerably, while at the same time the lower tail risk falls. Unfortunately, the following two quarters are missing but in the subsequent recession, there is a strong reaction by the lower tail IR-Factor. It increases considerably right after the unobserved period and in the quarters thereafter. The dot-com bubble and its recession appear to be linked more closely to cross-sectional income distribution changes than the economy-wide recession caused by the sub-prime crisis. This ability




Figure 2.4: Demeaned 0.96 IR-Factor time series for the lower tail (IR^-) and the upper tail (IR^+) along its 0.95 confidence band. There is a two quarter gap (2000 Q2, 2000 Q3). The blue shaded areas mark recession periods. The large peak in the (IR^+) series coincides with the peak of the dot-com bubble. The following recession features increased lower tail risk fluctuations. This can also be observed for the first recession period. The recent recession seems structurally different.

to distinguish between symmetric and asymmetric recessions where big earners lose over-proportionally may be one of the characteristics that allow the IR-Factors to price the cross-section of portfolio returns that is considered in this study.

2.4.2 IR-Model Estimation and Model Comparison

How well can the IR-Model explain differences in a cross-section of portfolio returns? This question is investigated employing two-pass least squares regressions and GMM. The two-pass regression is performed in two steps. In a time series regression, excess returns are regressed on income growth and the IR-Factors to estimate a portfolio's exposure to the factors

$$R_{i,t}^e = a_i + \beta_i' Factors_t + \varepsilon_{i,t}, \quad (2.14)$$

where $R_{i,t}^e$ is the excess return of the portfolio i at time t , a_i is a portfolio specific constant, $Factors_t$ is a vector containing the income growth and IR-Factor series for the lower and upper tail, and β_i' is a vector containing the exposure factors $\beta_{i,\Delta y}$, $\beta_{i,IR-}$, and $\beta_{i,IR+}$. The estimate $\hat{\beta}_i$ is used in a cross-sectional regression of average excess returns to estimate a variable that is related to the factors' price of risk

$$\mathbb{E}_T [R_i^e] = \lambda' \hat{\beta}_i + \alpha_i, \quad (2.15)$$

where $\mathbb{E}_T [R_i^e]$ is the time series average excess return of portfolio i , λ' is a vector containing the price of risk variables, and $\hat{\beta}_i$ are the factor exposures, estimated in the time series regression. The results are summarized in Table 2.2.

The IR-Factors can explain differences in portfolio returns related to differences in size and book-to-market ratios. The lower tail exposure estimates from the time series regression $\hat{\beta}_{i,IR-}$ are significant – especially for larger stocks. The upper tail exposure estimates $\hat{\beta}_{i,IR+}$ are significant – especially for smaller stocks. The cross-sectional estimates $\hat{\lambda}_{IR-}$ suggest that the lower tail IR-Factor is priced. The upper tail price of risk estimates $\hat{\lambda}_{IR+}$ are weakly significant for the second and third IR-Factor quantiles. The overall size of the estimates is economically meaningful. In

Time Series and Cross-Sectional Least Squares Estimates

Table 2.2: List of two-pass least squares estimates for the Idiosyncratic Risk Model. The three panels refer to the different IR-Factor quantiles. The time series estimates $\hat{\beta}_{i,IR+/-}$ are reported in 5x5 tables. The estimates are sorted by portfolios, ranging from low to high book-to-market ratio, and small to large size. Heteroscedasticity corrected p-values are reported in parentheses. The cross-sectional estimates $\hat{\lambda}_{IR+/-}$ can be found below each 5x5 table. Shanken (1992) corrected p-values are reported in parentheses. Estimation results for the income growth factor are omitted as they are insignificant. The sampling frequency is quarterly, returns are denoted in % per month.

Panel A: Two-Pass Estimates for $q = 0.95$													
$\hat{\beta}_{i,IR-}$	Small					Large	$\hat{\beta}_{i,IR+}$	Small					Large
Low	-1.65	-0.55	-0.55	-0.95	0.12		Low	5.43	3.32	2.09	2.04	-0.66	
	(0.44)	(0.77)	(0.74)	(0.54)	(0.92)			(0.00)	(0.00)	(0.00)	(0.00)	(0.11)	
	-0.11	1.29	1.62	2.57	2.04			4.46	1.11	0.39	-0.26	-1.17	
	(0.94)	(0.32)	(0.15)	(0.01)	(0.02)			(0.00)	(0.09)	(0.39)	(0.46)	(0.00)	
	0.54	1.50	1.97	2.60	2.30		2.82	0.52	-0.17	-0.67	-0.89		
	(0.67)	(0.18)	(0.05)	(0.01)	(0.01)		(0.00)	(0.35)	(0.64)	(0.07)	(0.02)		
	0.65	1.86	2.71	2.37	2.72		2.35	0.85	-0.76	0.17	-1.78		
	(0.59)	(0.09)	(0.01)	(0.02)	(0.00)		(0.00)	(0.14)	(0.07)	(0.65)	(0.00)		
High	1.84	3.03	3.68	2.51	1.68		High	1.79	0.56	-0.61	-1.77	-0.92	
	(0.18)	(0.03)	(0.00)	(0.03)	(0.22)		(0.01)	(0.40)	(0.12)	(0.00)	(0.05)		
$\hat{\lambda}_{IR-}$	0.30						$\hat{\lambda}_{IR+}$	0.18					
	(0.04)							(0.18)					
Panel B: Two-Pass Estimates for $q = 0.96$													
$\hat{\beta}_{i,IR-}$	Small					Large	$\hat{\beta}_{i,IR+}$	Small					Large
Low	-4.20	-2.59	-2.40	-2.96	-0.64		Low	9.27	5.81	3.81	3.80	-0.84	
	(0.17)	(0.31)	(0.32)	(0.19)	(0.73)			(0.00)	(0.00)	(0.00)	(0.00)	(0.37)	
	-1.63	1.03	1.52	2.88	2.26			7.54	1.97	0.96	-0.17	-1.85	
	(0.45)	(0.54)	(0.31)	(0.03)	(0.06)			(0.00)	(0.05)	(0.17)	(0.81)	(0.01)	
	-0.11	1.54	2.38	3.21	2.67		4.76	0.89	-0.10	-0.97	-1.35		
	(0.95)	(0.30)	(0.07)	(0.03)	(0.02)		(0.00)	(0.32)	(0.88)	(0.16)	(0.12)		
	0.12	1.78	3.31	2.56	3.47		3.96	1.48	-1.29	0.38	-2.89		
	(0.94)	(0.22)	(0.03)	(0.07)	(0.01)		(0.00)	(0.10)	(0.07)	(0.54)	(0.00)		
High	1.82	3.48	4.43	3.09	2.28		High	2.92	0.85	-1.05	-2.91	-1.36	
	(0.33)	(0.06)	(0.01)	(0.05)	(0.17)		(0.02)	(0.42)	(0.13)	(0.00)	(0.19)		
$\hat{\lambda}_{IR-}$	0.24						$\hat{\lambda}_{IR+}$	0.17					
	(0.03)							(0.09)					
Panel C: Two-Pass Estimates for $q = 0.97$													
$\hat{\beta}_{i,IR-}$	Small					Large	$\hat{\beta}_{i,IR+}$	Small					Large
Low	-4.27	-3.09	-3.31	-3.83	-1.07		Low	10.60	7.43	5.19	5.15	-1.27	
	(0.31)	(0.35)	(0.28)	(0.20)	(0.65)			(0.04)	(0.00)	(0.00)	(0.00)	(0.33)	
	-1.97	0.88	1.21	2.73	2.21			9.05	3.00	1.93	0.41	-2.66	
	(0.53)	(0.67)	(0.52)	(0.12)	(0.16)			(0.03)	(0.04)	(0.08)	(0.76)	(0.03)	
	-0.25	1.13	2.07	3.07	2.93		5.79	1.66	0.52	-0.76	-1.95		
	(0.92)	(0.53)	(0.22)	(0.09)	(0.04)		(0.05)	(0.14)	(0.65)	(0.56)	(0.14)		
	-0.17	1.35	2.92	2.08	3.59		4.80	2.20	-1.35	0.58	-3.92		
	(0.94)	(0.44)	(0.12)	(0.24)	(0.04)		(0.08)	(0.08)	(0.25)	(0.55)	(0.04)		
High	1.81	3.55	4.27	2.88	2.85		High	3.38	1.31	-1.03	-3.11	-1.81	
	(0.43)	(0.12)	(0.04)	(0.20)	(0.17)		(0.14)	(0.35)	(0.43)	(0.09)	(0.25)		
$\hat{\lambda}_{IR-}$	0.22						$\hat{\lambda}_{IR+}$	0.15					
	(0.04)							(0.08)					

both regressions, excess returns are sampled quarterly and denoted in percentage points. Multiplying the price of risk factor with each portfolio's exposure factor yields an explainable return span of 4.8% per quarter for the lower tail 0.95 IR-Factor. The upper tail 0.95 IR-Factor covers a similarly large span of 3.6% per quarter. The exposure estimates $\hat{\beta}_{i,IR-}$ and $\hat{\beta}_{i,IR+}$ display an interesting pattern. Small-growth stocks have a negative exposure to the lower tail IR-Factor, while large-value stocks have a positive exposure to the lower tail IR-Factor. Exposures to the upper tail IR-Factor are just reversed. An economic intuition for the portfolios' risk exposures to the IR-Factor could be the following. When the lower tail IR-Factor is large, high income individuals are hit by large unexpected income shocks. These large income changes cause their previous portfolio allocation to be sub-optimal. In order to smooth their consumption path and to adjust it to the lower income, individuals liquidate long-term investments in secure "blue chips", most commonly large-value stocks. This allows for higher consumption at that moment at the cost of lower consumption in the following periods. In order to keep their investors, these large-value stocks have to offer a higher return. These expected return compensations can be detected in the two-pass regression. Conversely, when the upper tail IR-Factor is large, investors experience major unexpected income rises. In the process of readjusting their portfolio to realize a new smooth consumption path, they want to invest some of the additional income in secure long-term investments to safely transfer it into the future. As a result, large-value stocks can get away with paying a lower return.⁸ Overall, the direction of the effects and their sizes remain comparable across all three IR-Factor quantiles, which underlines that the results are robust with regard to threshold and IR-quantile selection. The overall significance of the first stage estimates decreases as the threshold is pushed further into the tail. This is an expected phenomenon. If the threshold is chosen to be higher than necessary, fewer observations are available, which translates into higher standard errors. To keep the discussion organized, the remainder of this section will focus on the results for the 0.95 IR-Factor. Detailed results for all three quantiles can be found in the appendix.

⁸Note that this intuition serves as a plausibility check. The actual mechanism through which the economy reaches equilibrium is not within the scope of this study.

Model Comparison – Goodness of Fit (Benchmark Models and IR-Model)

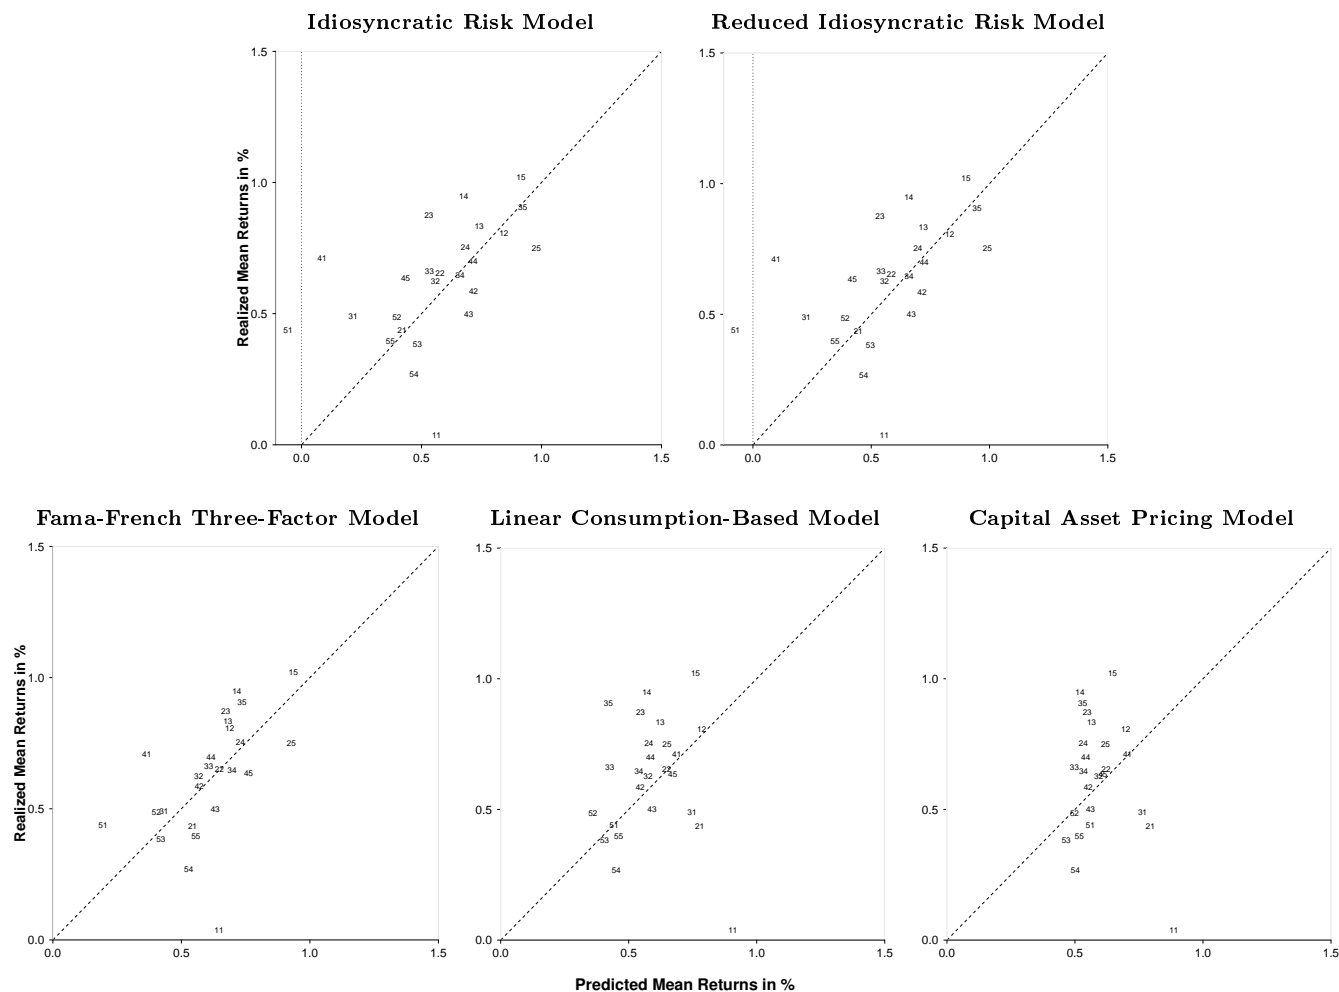


Figure 2.5: Realized and predicted mean excess returns. These graphs visualize the goodness of fit provided by the different stochastic discount factor specifications for the 25 Fama-French portfolio returns. The points are labeled according to their associated portfolio. The first digit refers to the book-to-market ratio (low to high), the second digit refers to the size (small to large). The bottom row plots the benchmark models' predicted mean returns against the observed values. The top row does the same for the (reduced) Idiosyncratic Risk Model. The closer the points align around the 45-degree line, the better is the model's ability to explain the observed mean returns. The sampling frequency is quarterly, returns are denoted in % per month.

Figure 2.5 shows a comparison of average return prediction quality between the Fama-French Model, the Linear CBM, the CAPM and the two IR-Model specifications. The better the portfolios align around the 45-degree line, the better is the model's ability to explain the portfolios' average returns. The variation that is explained by the IR-Model is comparable to that of the ad-hoc Fama-French Model. The returns align more clearly along the 45-degree line for the IR-Model than for the theory-driven alternatives, the Linear CBM and the CAPM. The return of the small-growth portfolio [11] is gravely overestimated in each model. This is a typical weakness of many asset pricing models and the IR-Model is no exception, although the overestimation is slightly reduced. On the other hand, the IR-Model struggles when pricing the larger growth portfolios [41], [51], even more so than the Fama-French Model.

Alternatively, all the above models can directly be estimated in their stochastic discount factor representation via GMM. First stage estimates use the identity matrix as a weighting matrix. Each portfolio receives the same weight in the objective function. Second stage GMM employs an optimal weighting matrix estimate, delivering more efficient estimates. To enable identification, the constant b_0 is set to one and factors are demeaned. The results are reported in Table 2.3. First stage estimates for the three benchmark models are not significant. This may be due to the length of the time period. None of the models can be rejected by the J-test; again, this is likely due to the length of the sample period. The signs of the coefficients are as expected. For the CAPM, a higher market return implies a lower stochastic discount factor, and thus a lower marginal rate of substitution. Investors are not willing to replace consumption in the next period with consumption now. The same is true for the Linear CBM. Higher consumption growth indicates a lower stochastic discount factor. In the Fama-French Model, the coefficient of the size factor SMB is not significant in both the first and the second stage. This is in line with the literature that finds the size effect to gradually disappear in recent decades.⁹ Second stage GMM estimates retain their sign and are significant. Cross-sectional root mean squared errors and adjusted R^2 s can provide a basis for comparing the models'

⁹Cochrane (2008) provides a comprehensive overview.

Model Comparison – GMM Estimates

Table 2.3: Comparison of the Generalized Method of Moments estimates for different stochastic discount factor specifications. Reported are first and second stage GMM estimates for the Capital Asset Pricing Model, the Linear Consumption-Based Model, the Fama-French Three-Factor Model and the the Idiosyncratic Risk Model. In the reduced IR-Model estimation (r), the income growth factor is excluded as a robustness check. p-values are given in parentheses. Furthermore, J-tests as well as adjusted R^2 s and root mean squared errors from the cross-sectional regression are provided below the estimates. The adjusted R^2 values refer to a cross-sectional regression that includes a constant to enable an interpretation as explained variance. None of the reported R^2 values are significantly different from R^2 s computed for randomly simulated factors.

	CAPM		L-CBM		FF		IR		IR (r)	
	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd
R^m	-0.07 (0.18)	-0.09 (0.02)			-0.06 (0.33)	-0.09 (0.04)				
Δc			-1.78 (0.37)	-2.04 (0.00)						
Δy							-0.17 (0.81)	-0.65 (0.22)		
SMB					-0.06 (0.38)	-0.04 (0.50)				
HML					-0.09 (0.17)	-0.12 (0.00)				
IR^-							-4.79 (0.31)	-2.09 (0.00)	-4.79 (0.29)	-2.30 (0.00)
IR^+							-1.12 (0.39)	-0.99 (0.00)	-1.20 (0.37)	-0.90 (0.00)
J-test	24.41	24.17	17.68	17.52	25.18	24.21	21.73	17.11	22.18	18.61
RMSE	0.27		0.26		0.19		0.23		0.24	
adj. R^2	0.03		-0.04		0.47		0.15		0.19	

goodness of fit, much like the graphs in Figure 2.5. Out of the three benchmark models, the Fama-French Model features the lowest average prediction error, as well as the highest explained variance. Both the CAPM and the Linear CBM perform worse, especially with respect to explained variance. The numbers reflect what can also be seen in Figure 2.5.

First stage estimates for the IR-Model are insignificant, just like the benchmark models. The IR-Factor coefficients retain the correct sign, consistent with the two-pass regression results. Second stage GMM estimates are significant. An economic interpretation for the upper tail coefficient is straightforward. If the IR-Factor is high, people experience large rises in income, resulting in an overall lower hunger. When the lower tail IR-Factor is high, high income individuals seek to liquidate some of their assets. To remain attractive, these assets have to offer higher returns. This return pressure will alleviate everyone else's hunger, resulting in an overall decrease in the stochastic discount factor. The RMSE is below that of the CAPM and the Linear CBM and comes close to the RMSE of the Fama-French Model. Explained variance for the IR-Model is considerably larger than for the Linear CBM or the CAPM but below that of the Fama-French Model. Adjusted R^2 are then simulated for 5000 draws of independent standard normal factors to evaluate the significance of the reported values (see Lewellen et al., 2010). None of the reported R^2 s are larger than the 0.95 quantile of the simulated R^2 values. One should avoid drawing conclusions based on these R^2 values alone.

The income growth factor coefficients feature the expected sign but are insignificant. Krebs (2004) argues that central moments can be arbitrary and do not provide any testable restrictions. The dynamics of the stochastic discount factor can be determined solely by the large income shock process. The estimation results support this notion. To further investigate the importance of the central moment factor to the estimation results, the aggregate income growth factor is excluded in a reduced version of the IR-Model. The influence on coefficient estimates is negligible. Figure 2.5 shows a comparison of the complete and reduced models' average return predictions. Here too, the difference is small. The reduced models' predictions of the big growth portfolio returns ($|51|$ and $|41|$) are even better. This is also re-

flected in the RMSE and adjusted R^2 values. The average prediction error remains nearly the same, whether income growth is included or not. Explained variance rises noticeably, which gives reason to believe that aggregate income growth does not contribute notably to the explained variance. Excluding an unnecessary factor leads to a smaller adjustment of the R^2 and thus to a higher adjusted R^2 . The importance of central moments in a stochastic discount factor that includes the IR-Factors is further explored in the appendix, where aggregate consumption growth is used in the stochastic discount factor instead of income growth. The implications are similar. Overall, the GMM results confirm the findings of the two-pass least squares regression.

2.5 Concluding Remarks

This study introduces the tail income risk Factor (IR-Factor) as a measure for the individual income risk that investors are exposed to. The IR-Factor is estimated on US income data and evaluated on returns of a cross-section of 25 portfolios sorted by size and book-to-market ratio. Two-pass regression and GMM results provide empirical evidence that large idiosyncratic income risk contributes to the explanation of this cross-section of portfolio returns. Its contribution to explaining the variation in portfolio returns is both economically and statistically significant. The direction and size of the effect varies across portfolios and reflects the impact of investors' portfolio adjustments to idiosyncratic income shocks. Aggregate income growth plays no significant role in the stochastic discount factor. This supports the theoretical claim of Krebs (2004), who argues that large individual income changes have a dominating effect on the stochastic discount factor while central moments can be arbitrary.

An interesting avenue for future research may be the development of a complete model with a closed-form solution for a stochastic discount factor that incorporates this personal disaster risk. In combination with a suitable utility function, this could enable estimations of risk aversion and time preference parameters and thus allow for a better economic assessment of the model. It might entail the departure from

the elegant representative agent paradigm, but the possible gain is a better and more complete understanding of the real economic processes that drive asset prices.

Appendix A

A.1 The IR-Factor for the Krebs (2004) Income Process Specification

To derive equation (2.9), recall the income growth specification of Krebs (2004):

$$\begin{aligned}\Delta y_{i,t+1} &= ((1 + g_{1i,t+1})(1 + g_{2i,t+1})(1 + G_{t+1}) - 1) y_{i,t} \\ &= ((1 + g_{i,t+1}^\circ)(1 + g_{2i,t+1}) - 1) y_{i,t}.\end{aligned}\tag{2.16}$$

Applying the IR-Factor (equation (2.7)) to $\Delta y_{i,t+1}$ yields

$$IR_{q,t+1}^- = \frac{\mathbb{E}[\Delta y_{i,t+1} | \Delta y_{i,t+1} < \Delta y_{t+1}[1 - q]]}{\Delta y_{t+1}[1 - q]}.\tag{2.17}$$

Substituting equation (2.16) into the denominator of (2.17) returns

$$\Delta y_{t+1}[1 - q] = (((1 + g_{i,t+1}^\circ)(1 + g_{2i,t+1}) - 1)y_{i,t}) [1 - q].\tag{2.18}$$

The definition simplifies for the case that $(1 - q) > p$ and η are sufficiently large, such that for all $\Delta y_{i,t+1}[u]$, where $u < q$, it holds that $g_{2i,t+1} = 0$. If individuals are hit by the shock η , these shocks are large enough to accumulate beyond the q quantile. Since we always want to choose a case where $(1 - q) > p$ and since Krebs (2004) claims that the large shock implicates a personal disaster, this is a mild assumption.¹⁰

$$\begin{aligned}\Delta y_{t+1}[1 - q] &= (((1 + g_{i,t+1}^\circ) - 1)y_{i,t}) [1 - q] \\ &= ((\tilde{y}_{i,t+1} - y_{i,t})) [1 - q].\end{aligned}\tag{2.19}$$

¹⁰The assumption is also inconsequential, as a similar equation to (2.9) can be derived without it. Instead, one can rely on shocks from the personal disaster risk component to be over-proportionally present in the tails. The resulting formulas are more complex, but the implications remain the same. I stay with the more straightforward case.

Substituting equation (2.16) into the numerator of (2.17) yields

$$\begin{aligned} & \mathbb{E} [\Delta y_{i,t+1} | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]] = \\ & \mathbb{E} [((1 + g_{i,t+1}^\circ)(1 + g_{2i,t+1}) - 1) y_{i,t} | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]]. \end{aligned} \quad (2.20)$$

There are two possibilities how the condition can be fulfilled for a given $y_{i,t}$. Either the corresponding $\Delta y_{i,t+1}$ belongs to the $1 - q - p$ income changes that are not the result of a personal disaster shock, but are still large enough, such that they lie above the q quantile or the condition is fulfilled for the p values that are a result of a personal disaster. With this in mind, equation (2.20) can be divided into

$$\begin{aligned} & \mathbb{E} [\Delta y_{i,t+1} | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]] = \\ & \frac{1 - q - p}{1 - q} \mathbb{E} [((1 + g_{i,t+1}^\circ) - 1) y_{i,t} | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]] + \\ & \frac{p}{1 - q} \mathbb{E} [((1 + g_{i,t+1}^\circ)(1 - \eta) - 1) y_{i,t}]. \end{aligned} \quad (2.21)$$

With the assumption that income shocks are independent of current income, the terms simplify further:

$$\begin{aligned} & \frac{1 - q - p}{1 - q} \mathbb{E} [((1 + g_{i,t+1}^\circ) - 1) y_{i,t} | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]] + \\ & \frac{p}{1 - q} - ((1 + G)(1 - \eta) - 1) \mathbb{E} [y_{i,t}] \\ & = \frac{1 - q - p}{1 - q} \mathbb{E} [(\tilde{y}_{i,t+1} - y_{i,t}) | \Delta y_{i,t+1} < \Delta y_{t+1} [1 - q]] + \\ & \frac{p}{1 - q} ((1 - \eta) \mathbb{E} [y_{i,t+1}] - \mathbb{E} [y_{i,t}]). \end{aligned} \quad (2.22)$$

Two effects drive the first part of the sum towards zero as q increases. First, the mean-of-excess function of a light tailed distribution is decreasing in q , so the expectation term approaches zero. Second, the fraction $\frac{1-q-p}{1-q}$ also approaches zero as $1 - q$ approaches p . As the threshold q is pushed further into the tail of the income shock distribution, the numerator of the IR-Factor converges to the second part of

the sum. Resubstituting equations (2.19) and (2.22) back into equation (2.17) yields the result obtained in equation (2.9)

$$IR_{q,t+1}^- = \frac{p}{1-q} \frac{(1-\eta)\mathbb{E}[y_{i,t+1}] - \mathbb{E}[y_{i,t}]}{(\tilde{y}_{i,t+1} - y_{i,t})[1-q]}. \quad (2.23)$$

A.2 Unbiased Estimation of the Expected Shortfall

The Peak over Threshold method rests on two limit theorems. The first theorem (Fisher and Tippett, 1928) describes the distribution of the maximum of a series of random variables. Let $\{X\}_n$ be iid random variables, b_n a constant > 0 , a_n a constant $\in \mathbb{R}$, \mathbb{V} a not degenerated distribution, and M_n the maximum value of the $\{X\}_n$ series such that

$$\frac{M_n - a_n}{b_n} \xrightarrow{d} \mathbb{V}, \quad (2.24)$$

then \mathbb{V} belongs to the class of Generalized Extreme Value distributions F_ξ . ξ is a shape parameter that determines the type of extreme value distribution – light, medium or heavy tailed. The distribution that has generated $\{X\}_n$ is said to be in the maximum domain of attraction of the Generalized Extreme Value distribution F_ξ . This convergence theorem can be utilized to model extreme events within certain predefined blocks, for example a month or a city. The distribution of the maxima within each block converges in distribution to the Generalized Extreme Value distribution F_ξ , if it is not degenerated.

In a cross-section of income changes, no natural blocks are present. When blocks are formed arbitrarily, multiple extreme events might be selected into the same block while other blocks might not contain any extreme events at all. This is likely to cause inefficiencies. In such a case, the Peak over Threshold method provides a more efficient solution.

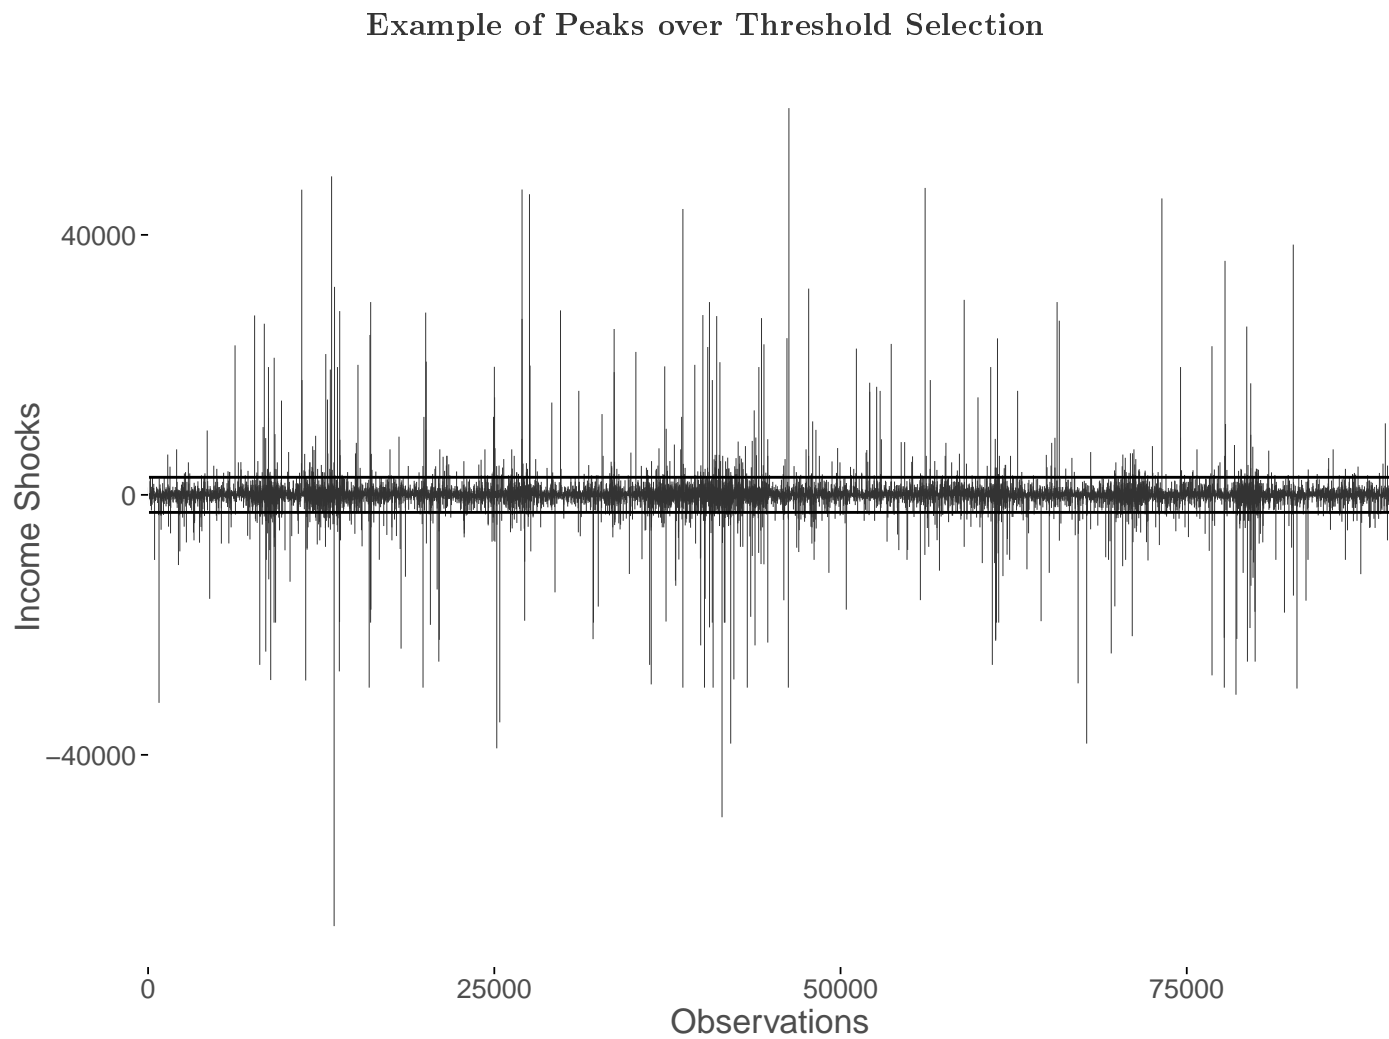


Figure 2.6: Income changes in US Dollar for individuals from the US Survey of Income and Program Participation in August 1996. Only the values above/below a threshold are used in estimating the parameters of the Generalized Pareto distribution.

Figure 2.6 illustrates the procedure. By selecting excesses above a certain threshold, and not maxima within a block, all large events can be taken into account. The theorem by Pickands (1975), Balkema and de Haan (1974) describes the distribution of the excesses F_X^u and relates it to the Fisher and Tippett (1928) theorem. Let X be a random variable with a distribution F_X and an excess distribution F_X^u . There exists a Generalized Pareto distribution $\tilde{F}_\xi(x, 0, \sigma(u))$ such that

$$\lim_{u \uparrow x_F} \left[\sup_{0 \leq x < x_F - u} \left(|F_X^u(x) - \tilde{F}_\xi(x; 0, \sigma(u))| \right) \right] = 0, \quad (2.25)$$

if and only if F_X is in the maximum domain of attraction of an extreme value distribution F_ξ with identical shape parameter ξ . Here, u is the selected threshold and x_F is the endpoint of the distribution ($x_F = \sup\{x \in \mathbb{R} | F_X(x) < 1\}$). The excess function F_X^u converges uniformly to a Generalized Pareto distribution \tilde{F}_ξ , if and only if the underlying value generating distribution F_X is in the maximum domain of attraction of an extreme value distribution F_ξ .

The Generalized Pareto distribution has the following form:

$$\tilde{F}_\xi(x, \mu, \sigma) = \begin{cases} 1 - (1 + \xi \frac{x-\mu}{\sigma})^{-\frac{1}{\xi}} & \text{for } \xi \neq 0 \\ 1 - \exp(-\frac{x-\mu}{\sigma}) & \text{for } \xi = 0. \end{cases} \quad (2.26)$$

The unknown parameters ξ and σ can be estimated via maximum likelihood, μ is known to be zero. These estimates can be used to compute a maximum likelihood estimate for the expected shortfall component

$$\widehat{\text{ES}}_q(\Delta y) = u + \frac{\hat{\sigma}(s-1)}{\hat{\xi}} + \frac{\hat{\sigma}s}{1-\hat{\xi}} \quad (2.27)$$

with: $s = \left(\frac{n}{N_u}(1-q) \right)^{-\hat{\xi}}$,

where n are the numbers of observations in the whole set and N_u are the number of peaks over the threshold. Dividing this estimate by the empirical q -quantile of the income shock distribution results in an unbiased estimate for the IR-Factor.

A.3 Robustness Checks

GMM Results for Higher Quantiles

Detailed GMM results for all three IR-Factor quantiles can be found in Table 2.4. The IR-Factor estimation results remain comparable. The sharp rise in RMSE for the higher IR-Factor quantiles is mostly due to the model's increasing problems in pricing big growth stocks ($|51|$ and $|41|$). The income growth factor coefficients display an interesting pattern. For the 0.95 quantile, they feature the expected sign but are insignificant. For the 0.96 and 0.97 quantile, the sign switches and second stage estimates are significant. This is at odds with the intuition that larger aggregate income growth should decrease hunger. However, this result is *not* at odds with the findings of Krebs (2004). He argues that central moments can be arbitrary and do not provide any testable restrictions. The dynamics of the stochastic discount factor can be determined solely by the large income shock process. Viewed in that light, the seemingly erratic and contradictory estimation results support the argumentation of Krebs (2004).

Alternative Test Portfolios

I check the robustness of the results by evaluating all models on a different set of portfolios. Each model specification is re-estimated using the equally-weighted Fama-French portfolios. The estimation results are expected to be similar. However, it is likely that small portfolios are prone to measurement errors, due to liquidity constraints. Table 2.5 summarizes the results. The sign, size and significance of all coefficients closely resemble the results for the value-weighted portfolios. The most notable difference are higher adjusted R^2 values, especially for the IR-Model, reflecting an overall better explanation of the observed average return variance.

Alternative Central Moments

While the erratic results obtained for the aggregate income growth factor support the dominating importance of the tail risk, it is possible that aggregate income growth is just a poor factor in general. To check the robustness of the above findings, alternative central moment factors are explored. In alternative ad-hoc IR-Model specifications, aggregate income growth is replaced by aggregate consumption growth

(C-IR) and the market excess return (R-IR). The estimation results are reported in Table 2.6. The inclusion of IR-Factors makes a notable difference in the explained average return variance for the aggregate consumption growth factor, as well as the market excess return factor. The adjusted R^2 rises from -0.04 to 0.17 when the IR-Factors are included alongside consumption growth and from 0.03 to 0.27 when included alongside the market excess return. The adjusted R^2 values for the Consumption IR-Model are nearly the same as for the Income IR-Model, and smaller than for the reduced IR-Model. This indicates that the contribution of the aggregate consumption growth factor to the explained variance of average returns is negligible. The inclusion of the market excess return instead of the aggregate consumption or income growth leads to a notable increase in the adjusted R^2 . Coefficient estimates for the aggregate consumption growth factor and the market excess return factor have the correct signs, but are insignificant at both GMM stages. The exception is the second stage estimate for the C-IR97 specification, with a p-value of 0.10 . Coefficient estimates for the lower tail IR-Factor are comparable to the original models' estimates. Coefficient estimates for the upper tail IR-Factor are notably smaller compared to the original IR-Model. Overall, the conclusions reached for the IR-Model can be maintained. The inclusion of aggregate consumption growth instead of aggregate income growth has no impact on explained variance. The inclusion of the market excess return does improve the explanatory power, but removes the model from the context of a Consumption-Based Model. Additionally, the market excess return factor becomes insignificant when the IR-Factors are included.

Figure 2.7 gives an overview of all robustness model specifications and their ability to predict average excess return for the 0.95 IR-Factor quantile.

Model Comparison – GMM Estimates (Full)

Table 2.4: Comparison of the Generalized Method of Moments estimates for different stochastic discount factor specifications. Reported are first and second stage GMM estimates for the Capital Asset Pricing Model, the Linear Consumption-Based Model, the Fama-French Three-Factor Model and the three Idiosyncratic Risk Model specifications for different IR-Factor quantiles (0.95, 0.96 and 0.97). In the reduced IR-Model estimation (r), the income growth factor is excluded as a robustness check. p-values are given in parentheses. Furthermore, J-tests as well as adjusted R^2 s and root mean squared errors from the cross-sectional regression are provided below the estimates. The adjusted R^2 values refer to a cross-sectional regression that includes a constant to enable an interpretation as explained variance.

	CAPM		L-CBM		FF		IR95		IR96		IR97		IR95 (r)		IR96 (r)		IR97 (r)		
	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	
R^m	-0.07 (0.18)	-0.09 (0.02)			-0.06 (0.33)	-0.09 (0.04)													
Δc			-1.78 (0.37)	-2.04 (0.00)															
Δy							-0.17 (0.81)	-0.65 (0.22)	0.34 (0.77)	0.29 (0.03)	0.53 (0.67)	0.70 (0.04)							
SMB					-0.06 (0.38)	-0.04 (0.50)													
HML					-0.09 (0.17)	-0.12 (0.00)													
IR^-							-4.79 (0.31)	-2.09 (0.00)	-6.75 (0.30)	-2.24 (0.03)	-8.99 (0.35)	-2.58 (0.08)	-4.79 (0.29)	-2.30 (0.00)	-6.78 (0.32)	-2.51 (0.01)	-9.18 (0.38)	-2.77 (0.02)	
IR^+							-1.12 (0.39)	-0.99 (0.00)	-2.48 (0.45)	-1.53 (0.00)	-3.79 (0.52)	-2.07 (0.00)	-1.20 (0.37)	-0.90 (0.00)	-2.25 (0.43)	-1.25 (0.00)	-3.14 (0.50)	-1.79 (0.00)	
J-test	24.41	24.17	17.68	17.52	25.18	24.21	21.73	17.11	21.14	15.06	20.40	13.31	22.18	18.61	20.76	16.26	20.33	14.52	
RMSE	0.27		0.26		0.19		0.23		0.29		0.32		0.24		0.29		0.33		
adj. R^2	0.03		-0.04		0.47		0.15		0.16		0.12		0.19		0.20		0.15		

Model Comparison – GMM Estimates (Equally-Weighted Portfolios)

Table 2.5: Comparison of the Generalized Method of Moments estimates for different stochastic discount factor specifications on the equally-weighted Fama-French portfolios. Reported are first and second stage GMM estimates for the Capital Asset Pricing Model, the Linear Consumption-Based Model, the Fama-French Three-Factor Model and the three Idiosyncratic Risk Model specifications for different IR-Factor quantiles (0.95, 0.96 and 0.97). In the reduced IR-Model estimation (r), the income growth factor is excluded as a robustness check. p-values are given in parentheses. Furthermore, J-tests as well as adjusted R^2 s and root mean squared errors from the cross-sectional regression are provided below the estimates. The adjusted R^2 values refer to a cross-sectional regression that includes a constant to enable an interpretation as explained variance.

	CAPM		L-CBM		FF		IR95		IR96		IR97		IR95 (r)		IR96 (r)		IR97 (r)		
	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	
R^m	-0.07 (0.17)	-0.08 (0.03)			-0.05 (0.39)	-0.06 (0.11)													
Δc			-1.91 (0.38)	-2.20 (0.00)															
Δy							0.25 (0.87)	0.42 (0.05)	0.78 (0.74)	0.53 (0.01)	0.78 (0.74)	0.66 (0.00)							
SMB					-0.08 (0.18)	-0.09 (0.02)													
HML					-0.11 (0.09)	-0.14 (0.00)													
IR^-							-4.56 (0.24)	-2.33 (0.00)	-6.30 (0.25)	-2.45 (0.00)	-8.72 (0.28)	-3.83 (0.00)	-4.61 (0.28)	-2.45 (0.00)	-6.60 (0.30)	-2.98 (0.00)	-8.80 (0.34)	-5.15 (0.00)	
IR^+							-1.03 (0.55)	-0.76 (0.00)	-2.22 (0.59)	-1.01 (0.05)	-2.63 (0.69)	-1.83 (0.04)	-0.90 (0.39)	-0.66 (0.00)	-1.51 (0.45)	-1.02 (0.01)	-1.41 (0.68)	-1.46 (0.02)	
J-test	24.37	24.35	17.62	17.20	26.01	25.48	24.90	20.46	24.74	18.43	21.74	17.42	24.82	20.93	24.87	20.34	21.58	18.98	
RMSE	0.29		0.27		0.19		0.25		0.29		0.30		0.25		0.31		0.33		
adj. R^2	0.05		-0.03		0.56		0.23		0.26		0.22		0.24		0.28		0.24		

GMM Estimates for Different Types of Central Moments

Table 2.6: Comparison of the Generalized Method of Moments estimates for different stochastic discount factor specifications. Reported are first and second stage GMM estimates for the three Idiosyncratic Risk Model specifications with different IR-Factor quantiles (0.95, 0.96 and 0.97) and two ad-hoc IR-Models with alternative central moments. C-IR includes consumption growth, R-IR includes the market excess return instead of income growth. p-values are given in parentheses. Furthermore, J-tests as well as adjusted R^2 s and root mean squared errors from the cross-sectional regression are provided below the estimates. The adjusted R^2 values refer to a cross-sectional regression that includes a constant to enable an interpretation as explained variance.

	IR95		IR96		IR97		R-IR95		R-IR96		R-IR97		C-IR95		C-IR96		C-IR97		
	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	
Δy	-0.17 (0.81)	-0.65 (0.22)	0.34 (0.77)	0.29 (0.03)	0.53 (0.67)	0.70 (0.04)													
R^m							-0.04 (0.58)	-0.02 (0.66)	-0.05 (0.45)	-0.04 (0.39)	-0.05 (0.42)	-0.05 (0.34)							
Δc													-0.99 (0.61)	-0.75 (0.43)	-1.29 (0.54)	-1.27 (0.13)	-1.43 (0.53)	-1.39 (0.10)	
IR^-	-4.79 (0.31)	-2.09 (0.00)	-6.75 (0.30)	-2.24 (0.03)	-8.99 (0.35)	-2.58 (0.08)	-2.86 (0.13)	-2.17 (0.00)	-3.38 (0.14)	-3.04 (0.00)	-4.23 (0.17)	-3.77 (0.00)	-2.60 (0.13)	-2.14 (0.00)	-3.01 (0.15)	-2.75 (0.00)	-3.73 (0.19)	-3.57 (0.00)	
IR^+	-1.12 (0.39)	-0.99 (0.00)	-2.48 (0.45)	-1.53 (0.00)	-3.79 (0.52)	-2.07 (0.00)	-0.59 (0.34)	-0.61 (0.00)	-0.74 (0.48)	-0.82 (0.05)	-0.73 (0.68)	-0.66 (0.43)	-0.46 (0.50)	-0.50 (0.01)	-0.49 (0.69)	-0.48 (0.19)	-0.34 (0.87)	-0.20 (0.79)	
J-test	21.73	17.11	21.14	15.06	20.40	13.31	25.11	24.30	25.05	24.88	24.24	24.11	21.08	20.28	19.26	19.06	19.24	19.17	
RMSE	0.23		0.29		0.32		0.23		0.23		0.24		0.22		0.22		0.23		
adj. R^2	0.15		0.16		0.12		0.27		0.25		0.26		0.16		0.17		0.13		

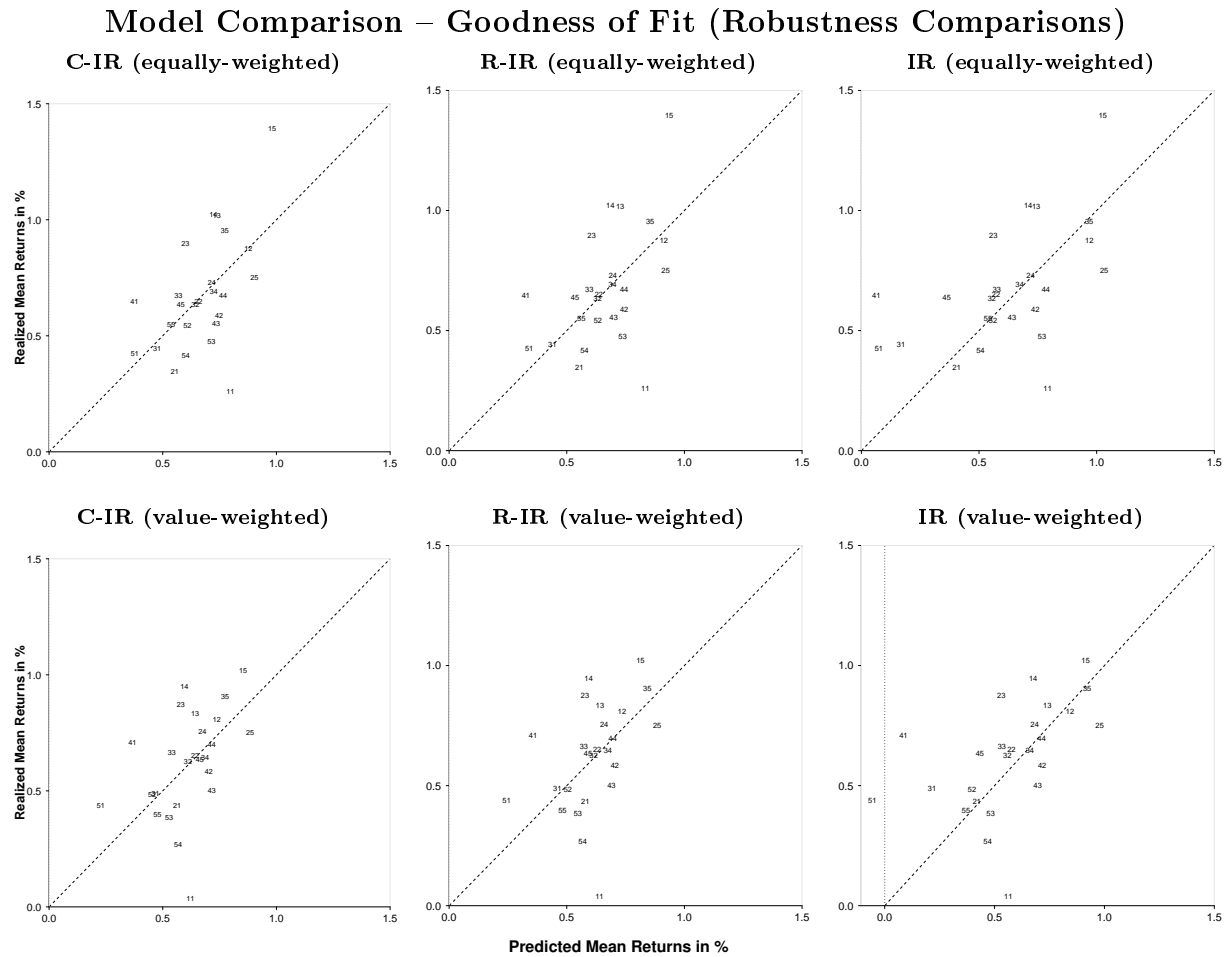


Figure 2.7: Realized and predicted mean excess returns. These graphs visualize the goodness of fit provided by different stochastic discount factor specifications for the 25 Fama-French portfolio returns. The points are labeled according to their associated portfolio. The first digit refers to the book-to-market ratio (low to high), the second digit refers to the size (small to large). In the C-IR version of the IR-Model, income growth is replaced by consumption growth. In the R-IR version, the excess return of the market portfolio replaces income growth. The top row plots the models' predicted mean returns against the observed values for the equally-weighted test portfolios. The bottom row does the same for the value-weighted test portfolios. The closer the points align around the 45-degree line, the better is the model's ability to explain the observed mean returns. The sampling frequency is quarterly, returns are denoted in % per month.

Chapter 3

Empirical Asset Pricing with Reference-Dependent Heterogeneous Agents*

Abstract

In this chapter I propose a strategy for the empirical evaluation of prospect theory that links concepts from behavioral finance to the literature on asset pricing with heterogeneous agents. I develop an asset pricing model in which two representative agents maximize their utility by investing in risky assets. One agent represents the behavior of investors above their reference level, one below. Using US income panel data, investors are sorted into groups depending on recent income development. In line with prospect theory, estimation results indicate that investors below their reference level act risk-seeking. The cross-sectional variation in returns of portfolios sorted by size and book-to-market value can be explained with a plausible risk aversion coefficient of ten while the unexplained equity premium is drastically reduced.

*Chapter 3 is based on the paper “Empirical Asset Pricing with Reference-Dependent Heterogeneous Agents” by Langen (2014). I thank T. Dimpfl, J. Grammig, R. Jung, J. Lahaye, F. Peter, W. Pohlmeier, E. Schaub, P. Sercu, and J. Sönksen, as well participants of the Congress of the European Economic Association (Toulouse) and the Journal of Banking and Finance’s Special Conference on Recent Developments in Financial Econometrics (Geelong) for helpful comments and suggestions. Financial support from the German Research Foundation (DFG) is gratefully acknowledged.

3.1 Introduction

Prospect theory (Kahneman and Tversky, 1979, 1992) is well established in describing choice under risk in experimental settings. So far few studies have implemented and evaluated prospect theory predictions in an empirical asset pricing context. To a certain degree, this is surprising. The psychological effects are well documented and it is hard to deny that a more accurate understanding of financial markets is needed. But the conversion of prospect theory to non-experimental settings entails numerous theoretical and practical difficulties so the translation is not straightforward.

This study tackles this challenge by linking concepts from behavioral finance to the literature on asset pricing with heterogeneous agents. In particular, I propose an asset pricing model in which two reference-dependent representative agents maximize their expected utility across time and states by investing in risky assets. One representative agent captures the behavior of investors above their reference level, the other represents the behavior of investors below their reference level. Prospect theory predicts that due to biases investors below their reference level act risk-seeking. The estimation results suggest that this is the case. By allowing for risk-seeking behavior, the cross-sectional variation in returns of portfolios sorted by size and book-to-market value can be explained with reasonably low risk aversion coefficients and the average unexplained equity premium (Mehra and Prescott, 1985) is reduced considerably.

My study contributes to the behavioral finance literature that analyzes the relation between returns and investor misevaluation. This relation exists, when psychological biases are shared systematically among groups of investors. These groups' systematic overbuying or overselling impacts prices and returns (Hirshleifer, 2001). Even in the presence of purely rational investors, equilibrium prices are affected by a systematic misevaluation (Figlewski, 1978; Campbell and Kyle, 1993; Shefrin and Statman, 1994). The psychological dispositions that are most relevant in the context of asset pricing can be subsumed by the notions of self-deception and self-control (Hirshleifer, 2001). The relevance of these two sources is evinced by a large body of experimental evidence (Kachelmeier and Shehata, 1992; Odean, 1998; Post

et al., 2008). When investors lack self-control, it manifests itself in time-inconsistent discounting and a high preference for immediate consumption. When investors are subject to self-deception, they misjudge their current situation and exhibit overconfidence in private information signals and their ability. Self-deception can also elicit biased self-attribution, where positive outcomes are associated with personal actions while negative outcomes are dismissed as bad luck. All behaviors subsumed by the notions of self-deception and self-control contribute to the misevaluation of securities. Self-deception is essential for the distinction of the curvatures in different parts of the utility function (risk-seeking or risk-averse behavior), because the behavioral impact of biases depends on the location of the investors in relation to their reference level. Self-deception creates the risk-seeking behavior below the reference level. An investor who is affected by self-deception may, for example, dismiss bad outcomes as bad luck and not adjust accordingly. This will seem risk seeking.

Since there is no obvious way to account for behavioral effects in an asset pricing context, recent empirical studies have proposed different approaches and have tested the predictions of prospect theory according to various interpretations.¹ Dimmock and Roy (2010) provide evidence that loss aversion can explain household non-participation in the stock market. Barberis and Huang (2008) show that if prospect theory holds, the skewness of returns is priced. Boyer et al. (2010), Bali et al. (2011) and Conrad et al. (2013) provide further empirical evidence. De Giorgi and Hens (2006) examine the size and value premium and are able to rationalize both with the calibration of a convex-concave exponential utility function. Finally, Barberis et al. (2001), Andries (2012) and Pagel (2012) build on Benartzi and Thalers' (1995) model to explain the high equity premium, and present additional evidence.

My study links these behavioral finance aspects to the literature on asset pricing with heterogeneous agents. In a seminal paper, Constantinides and Duffie (1996) demonstrate that expected returns are related to the cross-sectional variation of investors' individual income. Under restrictive assumptions, Constantinides and Duffie derive a closed-form solution for the stochastic discount factor. Brav et al.

¹Barberis (2013) provides a general assessment of recent research on prospect theory. Booij et al. (2010) give an extensive overview of both recent and older experimental studies.

(2002), Balduzzi and Yao (2007), and Grishchenko and Rossi (2012) account for investor heterogeneity by constructing an economy-wide stochastic discount factor as the average of the individual stochastic discount factors. In equilibrium, all investors' stochastic discount factors are valid, so a linear combination of the individual stochastic discount factors is also valid.

This study adopts the averaging of stochastic discount factors for the estimation of an asset pricing model with reference-dependent heterogeneous agents. I use US panel data and sort individuals into a group that is either above or below their reference level, depending on their recent income development. Representative agents are formed within these groups by capturing the average income development of investors above (below) their reference level. According to prospect theory, both representative agents may display different attitudes towards risk. The estimation results indicate that investors below their reference level are risk-seeking. The curvature of the utility function in the convex (loss) part is about twice as high as in the concave (gain) part which is in line with many empirical findings from experimental studies. The estimation results further suggest that forming groups by investors' disposition to their reference level helps explain the value premium. The cross-sectional variation in returns of size and book-to-market sorted portfolios can be explained with reasonable relative risk aversion coefficients of around ten. The average unexplained equity premium is the smallest among all benchmark models – with a reduction of 75% over the benchmark Fama-French (1995) Three-Factor Model. Compared to purely rational alternatives, return prediction accuracy is drastically increased and parameter estimates are more plausible.

The remainder of the chapter is structured as follows. Section 3.2 introduces the method of reference grouping and discusses the estimation methodology. Section 3.3 describes the data used for the estimation. Section 3.4 discusses the empirical results and section 3.5 draws a conclusion.

3.2 Model Setup and Methodology

3.2.1 Reference Grouping

The assumption that investors evaluate their income in relation to a reference level causes some theoretical and practical difficulties. One particularly troublesome question is how exactly do people determine their reference level? Literature does not provide a straightforward answer in this matter. One obvious reason for this is that a personal reference level is not directly observable. A major strand of studies assumes the reference level to be set externally (Campbell and Cochrane, 1999). Other studies assume that the reference level reflects expectations about the future (Koszegi and Rabin, 2006, 2007, 2009). For some investors the formation of these expectations may include notions of (un-)deservedness. Expectation forming could also take into account the overweighting of rare events. A big practical problem is the type of data needed to compute these potentially complex individual-specific reference levels. Reference grouping avoids these problems, without having to find a direct answer to them on an individual investor level.

Assume investor i derives utility in period $t+1$ from a reference-dependent power utility function

$$u_{i,t+1}(y_{i,t+1}|y_i^{ref}) = \begin{cases} \frac{1}{1-\gamma} (\Delta^* y_{i,t+1})^{1-\gamma} & \text{if } \Delta^* y_{i,t+1} > 1 \text{ (gain)} \\ \frac{1}{1-k\gamma} (\Delta^* y_{i,t+1})^{1-k\gamma} & \text{if } \Delta^* y_{i,t+1} < 1 \text{ (loss)}, \end{cases} \quad (3.1)$$

where γ is the relative risk aversion in the upper (gain) part of the utility function, and k is a loss aversion parameter. Utility is derived from changes of personal income $y_{i,t+1}$ in relation to an unobserved personal reference level y_i^{ref}

$$\Delta^* y_{i,t+1} = \frac{y_{i,t+1}}{y_i^{ref}}. \quad (3.2)$$

Note that the personal reference level y_i^{ref} is assumed to be sticky. Within a certain time frame, the investors have an idea of what they desire, or what they are used

to, and do not change this assessment quickly from one period to the next. The stickiness is what allows me to circumvent measuring the reference level directly. Also note that investors are assumed to derive utility directly from income, not consumption. Many studies suggest that investors are slow to adjust their consumption upon receiving bad income news as a result of loss-aversion (Shea, 1995; Bowman et al., 1999). To avoid this effect, I focus on income directly.

Utility-maximizing investors choose their portfolio allocations, such that marginal changes to their portfolio compositions today do not lead to an increase in expected, discounted future utility. Assuming time separable utility, this optimal allocation implies the stochastic discount factor

$$m_{i,t+1} = \beta \frac{u'(y_{i,t+1}|y_i^{ref})}{u'(y_{i,t}|y_i^{ref})}, \quad (3.3)$$

where β is the time preference parameter. If I had the knowledge, that the income of investor i is above its reference level, equation (3.3) in combination with equations (3.1) and (3.2) would simplify to

$$\begin{aligned} m_{i,t+1} &= \beta \left(\frac{y_{i,t+1}/y_i^{ref}}{y_{i,t}/y_i^{ref}} \right)^{-\gamma} \\ &= \beta \left(\frac{y_{i,t+1}}{y_{i,t}} \right)^{-\gamma}. \end{aligned} \quad (3.4)$$

Similarly, if I knew that the income of investor i is below its reference level, equation (3.3) would simplify to

$$m_{i,t+1} = \beta \left(\frac{y_{i,t+1}}{y_{i,t}} \right)^{-k\gamma}. \quad (3.5)$$

These simplifications make a direct measurement of an individual's reference level obsolete. To identify where the income of an investor is located in relation to his reference level, I take a look at his recent income path. Assume all investors are sorted by their recent income development. If investor i belongs to the upper q quantile of people that has recently experienced considerable positive income changes, it is reasonable to assume that he is currently above (A) his reference level. Conversely,

Reference Grouping: Constructing Representative Agents A and B

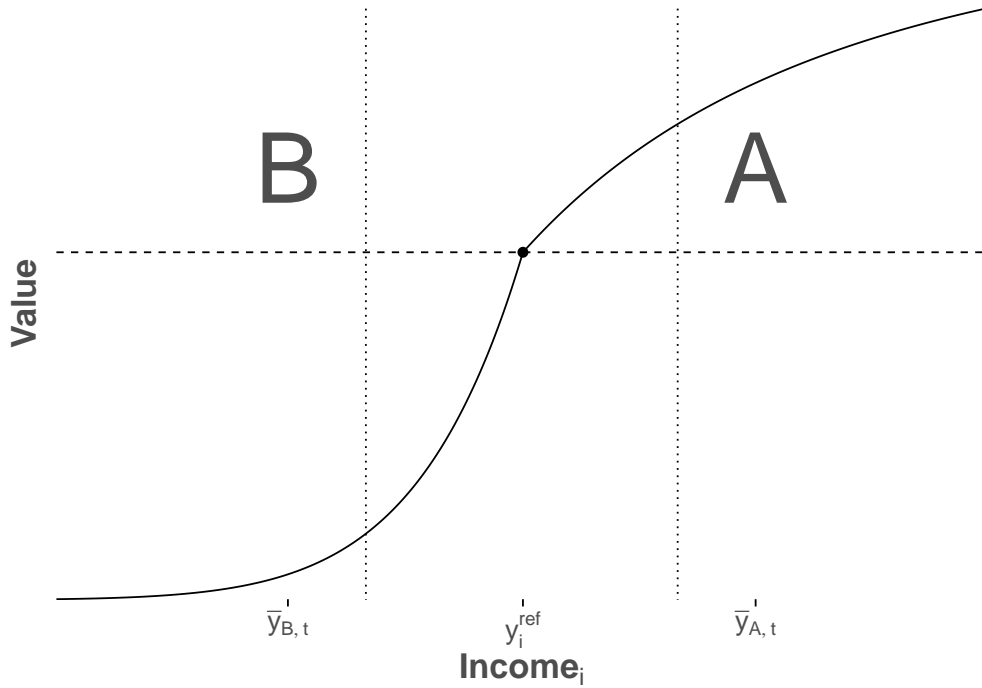


Figure 3.1: This graph visualizes the procedure of reference grouping. Investors are sorted into groups that are above (A) and below (B) their reference level. Representative investors are created in both groups. The income changes for representative agent A are constructed using only income information of agents well above their reference level. Since the reference level is sticky, only the upper part of the utility function is relevant for investor A. Likewise, the income changes for representative agent B are constructed using only income information of agents well below their reference level. So only the lower part of the utility function is relevant for investor B.

if investor i belongs to the lower q quantile of people that has recently experienced considerable negative income changes, it is reasonable to assume that he is currently below (B) his reference level. Figure 3.1 illustrates this approach. To capture the general behavior of investors above and below their reference level I form two representative agents within these groups. Representative agent A (above) has the average income development of the q investors with the most advantageous recent income developments. These investors are above their respective reference level and are assumed to share a common risk aversion parameter γ . All of the gains and losses in income are evaluated within the risk-averse part of the utility function because these q investors have recently surpassed their reference level by a large margin. Representative agent B (below) has the average income development of the

q investors with the least advantageous recent income developments. All of the gains and losses in income are evaluated within the lower part of the utility function. If k is negative, these agents display risk-seeking behavior. If k is positive but unequal to one, they just have a different risk aversion. This provides a test for the presence of behavioral effects.

Bhamra and Uppal (2014) prove the existence of an equilibrium in such an economy. As long as agents share the same time preference β , an equilibrium for i heterogeneous power utility agents exists without restricting the risk aversion coefficients γ_i . Translating this study's model to the model of Bhamra and Uppal (2014) results in a two agent economy where both agents have the same time preference β while $\gamma_A = \gamma$ and $\gamma_B = k\gamma$.

Expected returns are related to investor misevaluation when groups of investors systematically share the same biases. Through reference grouping I form representative agents among investors that potentially share certain psychological biases due to their disposition in relation to their reference level. Self-deception is the psychological mechanism that protects individuals from the pain of realizing that they are not where they want to be. Self-deception may lead investors below their reference level to appear risk-seeking. It may even lead them to be overconfident in their ability to invest and thus continue to hold high-risk securities. It may also lead them to have a biased-self attribution where gains are attributed to their ability while losses are attributed to external circumstances. This may then cause systematic overbuying or overselling and thus contribute to a misevaluation of assets. This potential effect materializes itself in the value of k , where negative values imply a risk-seeking behavior below the reference level. To ensure the existence of an equilibrium all agents share the same time preference β . If investors above and below the reference level lack self-control and engage in time-inconsistent discounting, this should manifest itself in a low time preference parameter.

3.2.2 Estimation Methodology

The model's capacity in pricing a cross-section of returns as well as the average equity premium can be evaluated using a set of moment restrictions in the spirit of Parker and Julliard (2005). These restrictions allow for the estimation of an average unexplained excess-return term α . In equilibrium the stochastic discount factors of both agents should be valid. This implies that they can price a cross-section of excess returns

$$\begin{aligned} 0 &= \mathbb{E} [m_{A,t+1} R_{t+1}^e | \mathcal{F}_t] \\ 0 &= \mathbb{E} [m_{B,t+1} R_{t+1}^e | \mathcal{F}_t], \end{aligned} \tag{3.6}$$

where R_{t+1}^e is a vector of future excess returns and \mathcal{F}_t is the information set available at time t . Using these individual moment conditions and the law of total expectations, an economy-wide unconditional moment restriction can be derived

$$0 = \mathbb{E} \left[\sum_{i=A}^B m_{i,t+1} R_{t+1}^e \right]. \tag{3.7}$$

The type of moment restriction in equation (3.7) is commonly used in GMM (Hansen and Singleton, 1982) estimations. An additional moment restriction can be included to identify the time preference β

$$0 = \mathbb{E} \left[\sum_{i=A}^B m_{i,t+1} R_{t+1}^f - 1 \right], \tag{3.8}$$

where R_{t+1}^f is the risk-free rate. Equation (3.7) can be rearranged using the definition of covariance to get an expression for the expected excess return

$$\mathbb{E} [R_{t+1}^e] = - \frac{\text{Cov} \left[\sum_{i=A}^B m_{i,t+1}, R_{t+1}^e \right]}{\mathbb{E} \left[\sum_{i=A}^B m_{i,t+1} \right]}. \tag{3.9}$$

If the model holds, the difference between expected excess returns on the left, and predicted excess returns on the right should be zero. If the model does not hold, part of the average level of excess returns may remain unexplained:

$$0 = \mathbb{E} [R_{t+1}^e] - \alpha + \frac{\text{Cov} \left[\sum_{i=A}^B m_{i,t+1}, R_{t+1}^e \right]}{\mathbb{E} \left[\sum_{i=A}^B m_{i,t+1} \right]}. \quad (3.10)$$

The additional parameter α measures the average unexplained excess return of the model. It is thus an estimate for the unexplained equity premium implied by the model. Equation (3.10) is used as a moment restriction in the GMM estimations, where equation (3.8) is employed as identifying condition for β . The set of restrictions (3.7) and (3.8) is used in a robustness check. In all estimations, first stage estimates rely on the identity matrix as a weighting matrix where each portfolio receives the same weight in the objective function. Second stage GMM employs an optimal weighting matrix estimate, which delivers more efficient estimates.

The performance of the Reference Group Model (RGM) above is evaluated in comparison to several benchmark models. The Consumption Based Model (CBM) (Lucas, 1987; Breeden, 1979; Grossman and Shiller, 1981) is estimated with

$$m_{t+1}^{CBM} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma}, \quad (3.11)$$

where c_{t+1}/c_t is consumption growth. The Income-Consumption Based Model (I-CBM) is estimated with

$$m_{t+1}^{I-CBM} = \beta \left(\frac{y_{t+1}}{y_t} \right)^{-\gamma}, \quad (3.12)$$

where y_{t+1}/y_t is income growth. The Capital Asset Pricing Model (CAPM) (Sharpe, 1964; Lintner, 1965; Mossin, 1966) is estimated with

$$m_{t+1}^{CAPM} = b_0 + b_1 R_{t+1}^m, \quad (3.13)$$

where R_{t+1}^m is the market excess return. The Fama-French (1995) Three-Factor Model (FF) is estimated with

$$m_{t+1}^{FF} = b_0 + b_1 R_{t+1}^m + b_2 SMB_{t+1} + b_3 HML_{t+1}, \quad (3.14)$$

where SMB_{t+1} is the excess return of small size portfolios over big size portfolios and HML_{t+1} is the excess return of high book-to-market portfolios over low book-to-market portfolios. These two factors are constructed using the returns on the basis of which the model is then evaluated, giving the Fama-French Model a natural advantage and the status of the benchmark model.

The different models are compared using several standard model performance and goodness of fit measures. The root mean squared error is calculated as

$$RMSE = \sqrt{\frac{1}{J}(\hat{R}_j^e - \mathbb{E}_T[R_j^e])^2}, \quad (3.15)$$

where

$$\hat{R}_j^e = -\frac{\text{Cov}(\hat{m}_{t+1}, R_{j,t}^e)}{\mathbb{E}_T[\hat{m}_{t+1}]} + \hat{\alpha} \quad (3.16)$$

and

$$\mathbb{E}_T[R_j^e] = \frac{1}{T} \sum_{t=1}^T R_{jt}^e, \quad (3.17)$$

where R_j^e is the excess return of the j th test portfolio. The RMSE provides a measure for the size of the model errors when explaining the different portfolios' average observed excess returns. A second goodness of fit measure relates the variation of the pricing errors to the overall observed variation

$$R^2 = 1 - \frac{\text{Var}(\mathbb{E}_T[R_j^e] - \hat{R}_j^e)}{\text{Var}(\mathbb{E}_T[R_j^e])} \quad (3.18)$$

The smaller the fraction, the higher the R^2 value, the better the ability of the model to explain the observed excess return variation.²

²As Lewellen et al. (2010) point out, one should avoid drawing conclusions based on these R^2 values alone. By simulating many unrelated factor series I find that it is likely to receive comparable and larger R^2 values by chance for all models.




Figure 3.2: This graph depicts the time structure of the nine SIPP panels. Panels are used as natural reference grouping blocks. Within each block income changes are calculated for both representative investors. Observations for the same time period from different panels are then averaged. The observation gap between March and September 2000 is excluded from the estimation.

3.3 Data Description

Individual income data is taken from the Survey of Income and Program Participation³ (SIPP) of the U.S. Census Bureau. It is chosen over other popular data sets (Panel Study of Income Dynamics or the Current Population Survey) as it provides income data on a monthly frequency. The SIPP consists of nine sequential, nationally representative panels. Each panel records the labor income of households and adult individuals within each household. The time series cover a period of nearly 23 years. The first panel starts in October 1987 and the last observations are from August 2010. There is a short time period of six months where no data is available between March and September 2000, due to a lack of funds. Figure 3.2 again depicts the time structure of the SIPP panels.

Cross-sectionally, the panels vary in size and contain the income of around 30,000 to 100,000 adults. The sample of households changes between the different panels which makes the linking of individual income series between panels impossible. The panels are used as natural blocks within which individuals are sorted by their recent

³www.census.gov/sipp



Figure 3.3: The income index shows the hypothetical income developments of two investors that start with an income of 1000 following the observation gap in September 2000. The changes of their income are equal to the calculated representative agents' income changes for $q = 0.12$. Recession periods are indicated as blue shaded areas.

income development. Individuals that are on an upwards path are considered for the group of investors who are above their reference level, individuals on a downward path are considered for the group of investors who are below their reference level. For each group the highest ranking q quantile is chosen. The higher the quantile, the further away the investors presumably are from their reference level, but at the same time the fewer observations are available to form the representative investor. The formation creates a trend so the income series are de-trended. All panels feature overlaps with others, so that there are no observation gaps between panels. Income changes can be computed for both groups of investors for each month. The resulting income growth series is depicted as an income index in Figure 3.3. The index shows the hypothetical income developments of two investors that start with an income of 1000 following the observation gap in September 2000. The changes of their income are equal to the calculated representative agents' income changes for a ranking quantile of $q = 0.12$. Descriptive statistics for the income growth series can be found in the left panel of Table 3.1.

The Reference Group Model is evaluated on the returns of 25 equally-weighted and value-weighted portfolios. The portfolios are constructed by sorting assets ac-

Descriptive Statistics – Income Growth and Factors

Table 3.1: The left panel of this table presents descriptive statistics for the generated monthly income growth series where y_A is the time series of investor above their reference level and y_B is the time series for investors below their reference level. Investor groups are formed with respect to different quantiles q . The right panel presents descriptive statistics for the factors of the benchmark models. The lower panel displays the correlations between all time series with $q = 0.12$. The sample period is November 1987 to August 2010 excluding a six month gap from April to September 2000.

<u>Time Series Means and Variances</u>							
q	$\overline{\Delta y_A}$	$\sigma_{\Delta y_A}$	$\overline{\Delta y_B}$	$\sigma_{\Delta y_B}$	Factors x	\bar{x}	σ_x
0.10	0.9996	0.0129	1.0023	0.0011	SMB	0.2718	9.7013
0.12	0.9992	0.0107	1.0023	0.0011	HML	0.1918	9.3665
0.14	0.9991	0.0093	1.0022	0.0011	R^m	0.5341	19.4243
0.16	0.9987	0.0082	1.0022	0.0012	Δc	1.0043	0.0000
0.18	0.9987	0.0073	1.0022	0.0012	Δy	1.0021	0.0000
0.20	0.9989	0.0066	1.0023	0.0012			

<u>Time Series Correlations</u>							
	SMB	HML	R^m	Δc	Δy	Δy_A	
HML	-0.23						
R^m	0.21	-0.24					
Δc	0.18	-0.06	0.14				
Δy	0.04	-0.06	0.11	0.11			
Δy_A	0.07	-0.02	0.09	0.10	0.01		
Δy_B	0.08	-0.09	0.08	0.14	0.10	0.47	

Monthly Average Fama-French Portfolio Excess Returns

Table 3.2: This table presents average observed monthly excess returns for the 25 Fama-French portfolios between November 1987 and August 2010 excluding the six month gap from April to September 2000. Assets are sorted into portfolios according to their size and book-to-market value. This sorting algorithm creates the large return variation displayed below. The left panel lists the average monthly excess returns for the value-weighted portfolios, the right panel lists the average monthly excess returns for the equally-weighted portfolios. The last row displays the average monthly equity premium that an investor who invested equally into all portfolios in 1987 would have earned by 2010. All values are in percent per month.

Value-Weighted Excess Returns						Equally-Weighted Excess Returns						
		<u>Book Market</u>						<u>Book Market</u>				
		Low	High	Low	High							
Size	Small	0.18	0.88	0.89	1.00	1.10	Small	0.45	1.01	1.10	1.10	1.48
		0.56	0.73	0.93	0.81	0.85	Size	0.53	0.76	0.99	0.81	0.88
		0.62	0.70	0.73	0.70	0.99		0.64	0.72	0.76	0.77	1.07
		0.80	0.65	0.58	0.74	0.72		0.77	0.67	0.64	0.74	0.74
	Large	0.55	0.57	0.47	0.35	0.52	Large	0.55	0.64	0.56	0.51	0.67
	Average	0.71					Average	0.78				

according to their market capitalization (size) and their book-to-market ratio (value). Table 3.2 provides an overview of the large dispersion in returns that is created by this sorting procedure. The average excess-return an investor would earn over the risk-free rate by investing in these portfolios is a monthly 0.71% and 0.78% respectively. The homepage of Kenneth R. French provides time series of these portfolios for the US stock market, alongside time series for the Fama-French Factors (R^m , SMB and HML) and the one-month treasury bill data used to compute excess returns.⁴ The data for aggregate consumption and income growth can be downloaded from the homepage of the Federal Reserve Bank of St. Louis.⁵ Descriptive statistics for all factors can be found in the right panel of Table 3.1.

⁴www.dartmouth.edu/~kfrench

⁵www.stlouisfed.org

3.4 Results and Discussion

3.4.1 GMM Estimation Results

GMM estimation results based on moment conditions (3.10) and (3.8) are reported in Table 3.3. Listed are first and second stage GMM estimates for the Reference Group Model (RGM) as well as several benchmark models. p -values are reported in parenthesis. For the formation of the representative agents, investors with $q = 0.12$ best (worst) income developments are used. As long as a q of 0.2 or lower is chosen, estimation results are stable. A more detailed discussion of the choice of q can be found in section 3.4.2.

The Reference Group Model (RGM) delivers the lowest unexplained equity premium estimate $\hat{\alpha}$ of all considered models. It is a less than half of the monthly average excess return implied by the test assets with an average unexplained excess return of 0.35% per month. The risk aversion parameter estimate $\hat{\gamma}$ of around ten borders on realistic values. Time preference parameter estimates $\hat{\beta}$ of 0.5 to 0.6 are extremely low. However, these values are in line with prospect theory when investors lack self-control and engage in time-inconsistent discounting. The loss-aversion parameter estimate \hat{k} is negative with values of around -2. The negativity implies that investors below their reference level act risk-seeking. The value of < -1 implies loss aversion. Both are in line with prospect theory. Estimates stay similar in the second stage estimation. Comparing the estimates of the value-weighted and the equally-weighted test portfolios also reveals no striking differences.

The estimates of the alternative models are greatly affected by the inclusion of the unexplained equity premium parameter α . For the Capital Asset Pricing Model (CAPM) and the Fama-French Model (FF), the estimates for the market excess return R^m have the opposite sign to what is expected. This seems to be offset by an overestimated average equity premium $\hat{\alpha}$. The values of $\hat{\alpha} = 1.06\%$ and 1.46% per month respectively are far above the observed average equity premium of 0.71%. The Consumption Based Model (CBM) and the Income CBM (I-CBM) estimates for the unexplained equity premium are more reasonable. Additionally, the parameter

GMM Estimation Results

Table 3.3: This table presents GMM estimation results for the five different stochastic discount factor specifications. The sample period is 1987-2010. Test assets are the 25 value-weighted Fama-French portfolios in the first panel and the 25 equally-weighted Fama-French portfolios in the second panel. Equations (3.10) and (3.8) are used as moment conditions. First stage estimates use an identity matrix as weighting matrix. Second stage estimates are computed using an estimate for the optimal weighting matrix to deliver more efficient parameter estimates. p -values are reported in parentheses. J -test statistics are based on second stage estimates.

Model	Stage	Value-Weighted Portfolios							$\hat{\alpha}$	J	RMSE	R^2
		$\hat{\beta}$	$\hat{\gamma}$	\hat{k}	R^m	SMB	HML					
RGM	1st	0.62 (0.02)	10.31 (0.04)	-1.91 (0.01)					0.35 (0.22)	26.78 (0.06)	0.178	0.268
	2nd	0.53 (0.00)	12.84 (0.00)	-1.96 (0.00)					0.44 (0.03)			
I-CBM	1st	0.96 (0.00)	-13.64 (0.91)						0.75 (0.06)	31.94 (0.10)	0.207	0.009
	2nd	0.92 (0.00)	-29.24 (0.41)						0.98 (0.00)			
CBM	1st	1.12 (0.00)	28.91 (0.71)						0.59 (0.10)	31.96 (0.10)	0.204	0.039
	2nd	1.21 (0.00)	48.05 (0.15)						0.91 (0.020)			
CAPM	1st	1.00 (0.00)			0.02 (0.66)				1.06 (0.14)	31.01 (0.12)	0.203	0.050
	2nd	0.99 (0.00)			0.03 (0.05)				1.34 (0.00)			
FF	1st	1.00 (0.00)			0.05 (0.04)	-0.04 (0.09)	-0.01 (0.62)		1.46 (0.00)	49.30 (0.00)	0.150	0.476
	2nd	0.99 (0.00)			0.08 (0.00)	-0.05 (0.01)	0.02 (0.38)		1.88 (0.00)			
Model	Stage	Equally-Weighted Portfolios							$\hat{\alpha}$	J	RMSE	R^2
		$\hat{\beta}$	$\hat{\gamma}$	\hat{k}	R^m	SMB	HML					
RGM	1st	0.58 (0.00)	9.97 (0.02)	-2.68 (0.01)					0.66 (0.18)	29.73 (0.13)	0.177	0.404
	2nd	0.49 (0.00)	12.17 (0.00)	-2.74 (0.00)					0.67 (0.00)			
I-CBM	1st	0.96 (0.00)	-16.88 (0.87)						0.84 (0.03)	36.06 (0.04)	0.227	0.020
	2nd	0.88 (0.00)	-43.06 (0.11)						0.92 (0.00)			
CBM	1st	1.23 (0.00)	57.48 (0.47)						0.50 (0.29)	36.32 (0.04)	0.222	0.070
	2nd	1.36 (0.00)	86.11 (0.01)						0.77 (0.00)			
CAPM	1st	1.00 (0.00)			0.03 (0.28)				1.47 (0.02)	36.96 (0.03)	0.211	0.157
	2nd	1.00 (0.00)			0.03 (0.03)				1.43 (0.00)			
FF	1st	1.00 (0.00)			0.05 (0.04)	-0.04 (0.11)	-0.01 (0.60)		1.62 (0.00)	76.16 (0.00)	0.148	0.585
	2nd	1.00 (0.00)			0.05 (0.00)	-0.03 (0.14)	0.03 (0.17)		1.57 (0.00)			

estimates for time preference $\hat{\beta}$ and risk aversion $\hat{\gamma}$ for the CBM are closer to sensible values, when α is included in the estimation. In contrast, the results for the I-CBM are defeating.

Table 3.4 provides a comparison between estimates based on the moment conditions in equations (3.10) and (3.8) that include the unexplained equity premium parameter α and the more conventional moment conditions in equations (3.7) and (3.8). The upper panel refers to the first set of conditions (Moment Conditions A), the lower panel to the second one (Moment Conditions B). The inclusion of α has virtually no impact on the Reference Group Model while all other models have coefficients that change drastically. In panel B, where no unexplained excess return α is estimated, the market excess return R^m parameter estimates have the expected sign, the Consumption Based Model features the typical high risk aversion parameter estimate $\hat{\gamma}$ of around 110 and the I-CBM risk aversion parameter is positive in the first stage estimation.

In both tables several model performance evaluation statistics are reported. In most of them, the Reference Group Model performs almost as well as the benchmark Fama-French Model and better than the others. The root mean squared error (RMSE) is always a little higher than that of the Fama-French Model. Estimated on the basis of value-weighted portfolios, the Reference Group Model's R^2 is about half as high as the R^2 of the Fama-French Model, while it is much closer for the equally-weighted portfolios. All this is also reflected in Figures 3.4 and 3.5. These figures depict plots of predicted mean excess returns against those that have been realized for both sets of test portfolios. The more the portfolio returns align around the 45-degree line, the better is the models' ability in pricing these portfolio returns.

GMM Estimation Results – Moment Condition Comparison

Table 3.4: This table presents GMM estimation results for the five different stochastic discount factor specifications. The sample period is 1987-2010. Test assets are the 25 value-weighted Fama-French portfolios. In the first panel equations (3.10) and (3.8) are used as moment conditions, in the second panel equations (3.7) and (3.8) are used as moment conditions. First stage estimates use an identity matrix as weighting matrix. Second stage estimates are computed using an estimate for the optimal weighting matrix to deliver more efficient parameter estimates. p -values are reported in parentheses. J-test statistics are based on second stage estimates.

		<u>Moment Conditions A</u>									
<u>Model</u>	<u>Stage</u>	$\hat{\beta}$	$\hat{\gamma}$	\hat{k}	R^m	SMB	HML	$\hat{\alpha}$	J	$RMSE$	R^2
RGM	1st	0.62 (0.02)	10.31 (0.04)	-1.91 (0.01)				0.35 (0.22)	26.78 (0.06)	0.178	0.268
	2nd	0.53 (0.00)	12.84 (0.00)	-1.96 (0.00)				0.44 (0.03)			
I-CBM	1st	0.96 (0.00)	-13.64 (0.91)					0.75 (0.06)	31.94 (0.10)	0.207	0.009
	2nd	0.92 (0.00)	-29.24 (0.41)					0.98 (0.00)			
CBM	1st	1.12 (0.00)	28.91 (0.71)					0.59 (0.10)	31.96 (0.10)	0.204	0.039
	2nd	1.21 (0.00)	48.05 (0.15)					0.91 (0.020)			
CAPM	1st	1.00 (0.00)			0.02 (0.66)			1.06 (0.14)	31.01 (0.12)	0.203	0.050
	2nd	0.99 (0.00)			0.03 (0.05)			1.34 (0.00)			
FF	1st	1.00 (0.00)			0.05 (0.04)	-0.04 (0.09)	-0.01 (0.62)	1.46 (0.00)	49.30 (0.00)	0.150	0.476
	2nd	0.99 (0.00)			0.08 (0.00)	-0.05 (0.01)	0.02 (0.38)	1.88 (0.00)			
		<u>Moment Conditions B</u>									
<u>Model</u>	<u>Stage</u>	$\hat{\beta}$	$\hat{\gamma}$	\hat{k}	R^m	SMB	HML	J	$RMSE$	R^2	
RGM	1st	0.53 (0.01)	12.38 (0.00)	-1.47 (0.00)				36.22 (0.02)	0.196	0.119	
	2nd	0.55 (0.00)	11.96 (0.00)	-1.56 (0.00)							
I-CBM	1st	1.04 (0.00)	49.60 (0.13)					55.41 (0.00)	0.643	-0.394	
	2nd	0.95 (0.00)	-18.76 (0.60)								
CBM	1st	1.39 (0.00)	106.15 (0.10)					37.21 (0.04)	0.233	-0.251	
	2nd	1.43 (0.00)	113.69 (0.00)								
CAPM	1st	1.00 (0.00)			-0.03 (0.00)			42.45 (0.01)	0.245	-0.388	
	2nd	1.00 (0.00)			-0.04 (0.00)						
FF	1st	1.00 (0.00)			-0.03 (0.16)	-0.03 (0.33)	-0.05 (0.15)	35.36 (0.05)	0.183	0.225	
	2nd	0.99 (0.00)			-0.03 (0.00)	-0.01 (0.22)	-0.03 (0.00)				

Model Comparison – Goodness of Fit (Value-Weighted Portfolios)




Figure 3.4: This figure depicts realized and predicted mean excess returns for the 25 value-weighted Fama-French portfolios in % per month. The plots allow for an assessment of the goodness of fit for each of the different model specifications. The portfolio points are labeled according to their position in the book-to-market sorting (first digit, low to high) and their size (second digit, small to big). The closer the portfolios align around the 45-degree line, the better are the average excess return predictions of that model.

Model Comparison – Goodness of Fit (Equally-Weighted Portfolios)

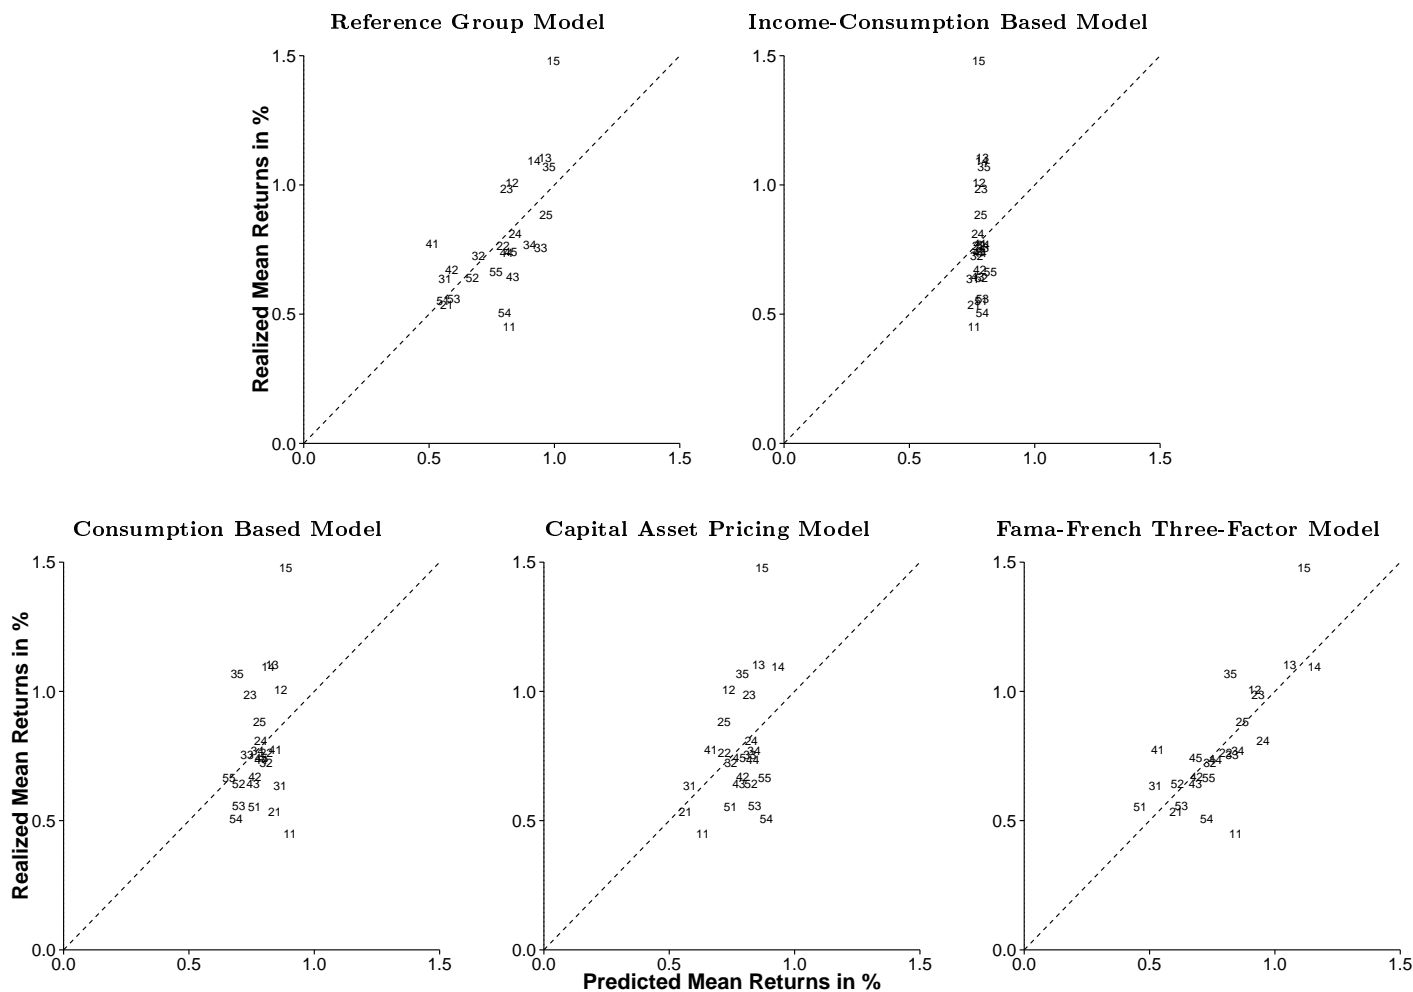


Figure 3.5: This figure depicts realized and predicted mean excess returns for the 25 equally-weighted Fama-French portfolios in % per month. The plots allow for an assessment of the goodness of fit for each of the different model specifications. The portfolio points are labeled according to their position in the book-to-market sorting (first digit, low to high) and their size (second digit, small to big). The closer the portfolios align around the 45-degree line, the better are the average excess return predictions of that model.

Only the Reference Group Model and the Fama-French Model can generate some of the necessary return variation, while the Capital Asset Pricing Model and the (Income-) Consumption Based Model predict nearly the same return for each portfolio. The premium that is earned by value portfolios cannot be explained by these traditional models. Allowing for risk-seeking behavior generates enough variation to explain a larger fraction of the dispersion in returns than the other theory-driven alternatives. An economic intuition for this is that firms with a small book value, relative to their market value, might be more difficult to evaluate, as fewer tangible assets are available. Thus, they may be more prone to misvaluation by investors that overestimate their personal ability or misinterpret private information signals.

3.4.2 Robustness

All above Reference Group Model estimates are produced with an investor group cutoff quantile of $q = 0.12$. Only investors that are among the q investors with the steepest income growth paths are used in the construction of representative agent A (above). A high group cutoff quantile q ensures that the assumption of this group of people to be representative for individuals above their reference level can be justified. Similarly, only those investors among the q investors with the steepest descending income paths are used in the construction of representative agent B (below). Figure 3.6 displays the impact, that the choice of the group cutoff quantile q has on the different model parameter estimates. The upper two graphs display the estimated values for the loss-aversion parameter \hat{k} for different quantiles q alongside the 0.95 confidence interval. The left column displays the estimates for the value-weighted test portfolios, the right column displays the estimates for the equally-weighted test portfolios. The middle and lower rows depict the same for the unexplained equity premium estimates $\hat{\alpha}$ and the risk-aversion parameter estimates $\hat{\gamma}$. The results indicate that estimates are relatively robust to the quantile selection. Risk-aversion and loss-aversion estimates stay significant throughout the different group size specifications. The unexplained equity premium $\hat{\alpha}$ stays insignificant. The level of all estimates also changes only slightly. In the data, many individuals do not report any income changes throughout a panel. So when using a total of 40% of

Robustness – Estimates are Stable for Different Group Quantiles

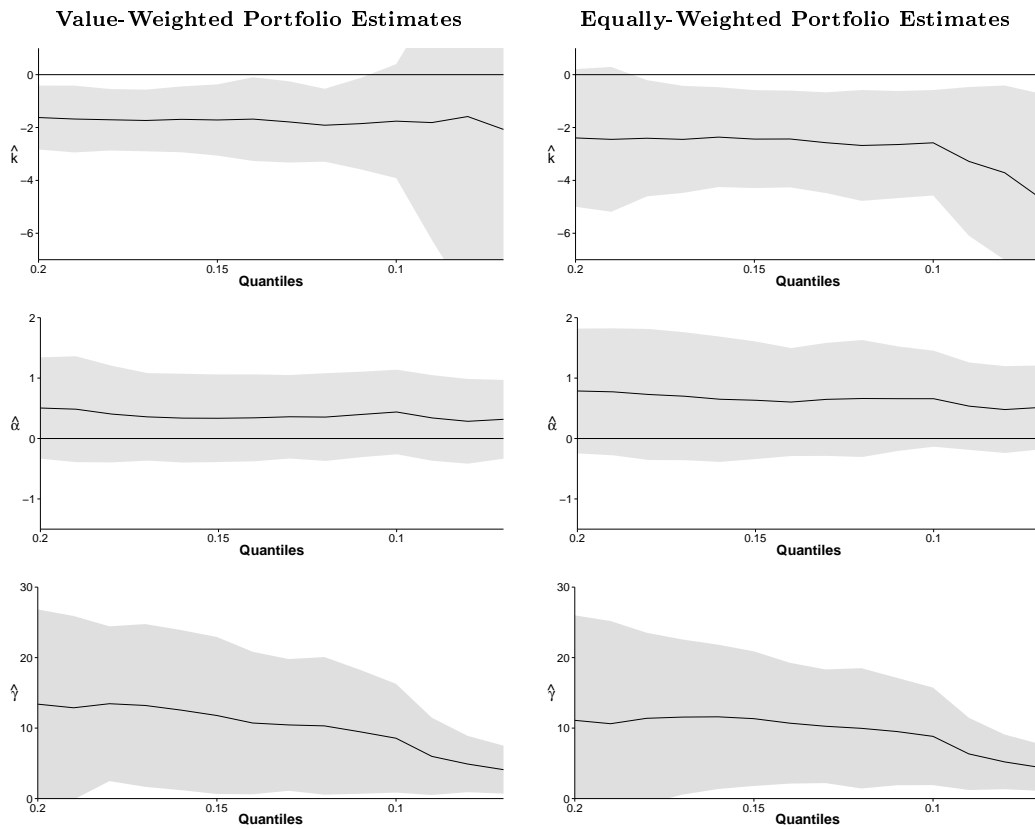


Figure 3.6: These graphs display GMM estimates for the key model parameters: the loss-aversion coefficient k , the unexplained average equity premium α and the risk aversion coefficient γ . Estimates are repeatedly computed for different values of the group cutoff quantile q . The left row depicts results for the 25 value-weighted Fama-French portfolios, the right row depicts results for the 25 equally-weighted Fama-French portfolios. 95% confidence intervals are indicated as gray areas.

all individuals' income developments in the representative agent construction (when $q = 0.20$) most of the income change information available in the data is exploited.

3.5 Concluding Remarks

Being biased can be subtly manipulative. So subtle that one might not even notice the influence on one's decisions. Where individuals are located in relation to a reference level can determine the biases that influence their behavior. This can sometimes be seen even more clearly when looking at examples that are far removed from questions of personal finance. Imagine that an individual intends to loose weight. In such a case people will have a target "reference" weight, a weight that they feel they should be able to reach – a weight they will be happy with. A commonly observed pattern is that individuals still tend to indulge in things that are wonderful right now but will have to be paid for later. In the case of loosing weight, this means eating that nice meal or watching TV instead of working out. Self-control is difficult. And it is closely tied to self-deception: People tend to believe, that they are going to change self-indulgent behavior in the future. But if past behavior is any indication, this may just be a case of overconfidence in future ability to resist temptations. To the individual, the behavior seems perfectly rational and it can see itself reaching its goal in the future. The individual is truly convinced of its perceived future ability to exercise self-control. To an outside observer, the same behavior appears to be extremely risk-seeking. The individual relies on an unlikely future commitment in order to make progress.

In experimental settings, evidence for psychological biases that influence individuals' decisions is plentiful. This study presents empirical evidence for the effect that psychological biases have on asset prices in a non-experimental setting. I propose a simple identification strategy based on individuals' income development that enables me to estimate the curvature of the investors' utility function in relation to their reference level. This allows for a test of the hypothesis that investors below their reference level act risk-seeking. Estimation results suggest that this is the case. Theory implies that risk-seeking investors overestimate their ability and overem-

phasize private information signals. Because the behavior is shared by a group of agents, asset prices are influenced. In the cross-sections of returns that I consider, assets are asymmetrically affected. When objective measures are less available (low book-to-market value), mispricings have a more pronounced effect. Allowing for psychological biases to affect asset prices, I am able to explain the cross-sectional variation in returns with a plausible risk aversion coefficient of ten. This is a vast improvement over values implied by purely rational models. Additionally, the model in this study can explain over 50% of the average equity premium – the largest value amongst all considered models. Besides providing evidence for risk-seeking behavior below the reference level, model parameter estimates are overall more plausible and model predictions more accurate than those of the purely rational alternatives.

Prospect theory is geared towards experimental research. For non-experimental empirical applications it is nonetheless a very useful generalization of many different psychological factors that influence decisions under risk. It facilitates a systematic categorization of psychological dispositions and resulting biases. And it provides predictions about the influence of these biases on certain decisions. Since these dispositions are so well documented, many recent empirical studies have adapted existing models to accommodate for those biases that are relevant to these models' application. Rather than trying to fit their model into a prospect theory framework, prospect theory tools are instead used to enhance existing theories and methodologies. It is still too early to draw any definite conclusions, but this strategy promises to be fruitful.

Chapter 4

A Cross-Country Analysis of Unemployment and Bonds with Long-Memory Relations*

Abstract

We analyze the relationship between unemployment rate changes and government bond yields during and after the most recent financial crisis across nine industrialized countries. The study is conducted on a weekly basis and we therefore nowcast unemployment data, which are only available once a month, on a weekly frequency using Google search query data. In order to account for the time series' long-memory components during the first-stage nowcasting and the second-stage modeling, we draw on Corsi's (2009) heterogeneous autoregressive time series model. In particular, we adapt this idea to a setting of mixed-frequency nowcasting. Our results indicate that Google searches greatly increase the nowcasting accuracy of unemployment rate changes. The impact of an idiosyncratic rise in unemployment on bond yields turns out to be positive for European countries while it is negative for the United States and Australia. The speed of the response also varies. Not unexpectedly, bond yields do not have an impact on unemployment. Our findings have interesting implications for the way shocks are absorbed in economic systems that differ, in particular, with respect to the central bank's core tasks.

*Chapter 4 is based on the paper "A Cross-Country Analysis of Unemployment and Bonds with Long-Memory Relations" by Dimpfl and Langen (2015). Financial support from the German Research Foundation (DFG) is gratefully acknowledged.

4.1 Introduction

Government bonds are by no means risk-free investments. In the aftermath of the financial crisis investors demanded considerable compensation for taking the risk of owning government bonds from nearly insolvent countries like Greece, Spain or Portugal. Economic theory suggests that a combination of a high level of government debt and low tax revenue due to high unemployment rates may increase the default risk of a country. Still, government bonds are affected by monetary policy which could halt or even reverse the effect. Earlier studies on the relationship between government debt and unemployment rates did not find a clear-cut pattern of interdependency. In particular, empirical studies are hampered by the fact that unemployment rates and government bond yields are observed on different frequencies. Bonds are continuously traded while unemployment rates are only announced once a month. The general solution is to align the data on a monthly frequency. Decreasing the frequency of the bond data obviously entails a substantial loss of information.

In this study, we examine the relationship between unemployment and government bond yields on a weekly basis, that is on a frequency which is higher than the publishing frequency of unemployment figures. To this end, we describe a general method to increase the resolution of a lower frequency time series when additional, related high-frequency data are available. The method implies a nowcasting of the low-frequency time series that feeds the supplementary data as a cascade of frequencies into the nowcast. The cascade structure accounts for long-run components of the higher frequency time series in a parsimonious way. We refer to this model as M-HAR model as it combines methodologies from the literature on mixed-frequency nowcasting (M; see, *inter alia*, Marcellino and Schumacher, 2010) and heterogeneous autoregressions (HAR; in particular following Corsi, 2009). We use Google search query data to increase the frequency of unemployment data from a monthly to a weekly level. For that purpose, the principal component of various search queries for information about unemployment is fed into the prediction model. In the process of nowcasting unemployment based on weekly Google data and monthly unemploy-

ment information, the M-HAR model can process both short-term and long-term relations between these time series. This is important because an individual's behavior that is indicative of becoming unemployed in the future can be observed far ahead of the date where the person actually becomes officially unemployed. Internet searches that were conducted one or multiple weeks before are used for the current unemployment nowcast. Therefore, our model can account for a general rise in searches that took place already during last three months, while current search intensity might be going down. We show that including Google search query data greatly increases the nowcasting accuracy compared to a pure autoregressive model. The resulting weekly unemployment time series is then used in a heterogeneous vector-autoregression (HVAR) of unemployment changes and bond yields. For a sample of seven European countries we consistently find that bond yields react positively to a rise in unemployment while for the United States and Australia this effect is negative. In contrast, there is virtually no impact of shocks in bond prices on unemployment.

Currently, the literature that directly analyzes the relationship between unemployment and bond prices is scarce. A notable exception is Bayoumi et al. (1995) who find that debt financing costs rise by nine basis points if unemployment rises by one percentage point. Still, there is quite a number of event studies that investigate how the announcement of unemployment figures impacts on government bond rates. Balduzzi et al. (2001) find that short-term (three-month T-bill, two-year note) and long-term bonds (ten-year note) in the US react positively to a surprise rise in jobless claims. Similarly, Fleming and Remolona (1999) report that employment announcements are the macroeconomic announcements that have the greatest effect on bond prices. In line with these findings, Afonso et al. (2011) document that the structural level of unemployment has a negative long-run impact on the credit rating of a country which may ultimately lead to rising refinancing costs, i.e. higher bond prices.

A possible explanation for why this kind of analysis is rare is a data issue. Unemployment figures and bond prices are not available at an identical frequency and one faces the question how to suitably aggregate bond data to monthly levels.

Recently, mixed data sampling (MIDAS) regression models (in particular Ghysels et al., 2004) have become popular as a means to alleviate this problem. For example, Ghysels et al. (2006) predict daily stock price volatility combining different intraday frequencies of returns. In this spirit, Marcellino and Schumacher (2010) combine the virtue of MIDAS regressions with factor models to reduce the dimensionality in macroeconomic models and to overcome the sampling frequency issue when now- and forecasting macroeconomic variables.

Another challenge in modeling government bond data is the long-memory property of this financial time series. Currently, there are two major ways to model stationary time series with long-memory properties. First, there is the class of ARFIMA (autoregressive fractionally integrated moving average) models. Sibbertsen et al. (2014) for example document that the bonds of France, Germany, Italy, and Spain are highly persistent and are close to unit-root behavior during the financial crisis. A second way to account for this high persistence is the cascading structure of heterogeneous autoregressive models as suggested by Corsi (2009). As the model is readily implemented it is now frequently and successfully applied (see, amongst others, Bauer and Vorkink, 2011; Dimpfl and Jung, 2012; Tseng et al., 2009). Chiriac and Voev (2011) compare the two methodologies and find a similar performance of the ARFIMA and the HAR model.

The forecast of a low-frequency time series at a higher frequency is termed nowcasting. This is of particular interest in macroeconomic modeling where GDP and other important economic indicators are only available on a quarterly or even yearly basis and are generally provided with a substantial time lag. Factors that influence the variable of interest are, however, available on a higher frequency and can be used to provide a nowcast of the still unobserved variable. Giannone et al. (2008) and Kuzin et al. (2011), for example, provide intra-quarter nowcasts of current GDP growth rates. On an even longer horizon Navicke et al. (2014) nowcast the at-risk-of-poverty rate in the EU which is only published with a delay of 2-3 years. This publication lag is obviously too long for policy recommendations.

In order to nowcast unemployment rates in France, Germany, Italy and Portugal, Barreira et al. (2013) draw on Google search query data. Google data have recently

gained particular attention as they allow to forecast a broad range of (economic) data. The ground-breaking work in this context is Ginsberg et al. (2009) who use search query data to find early warning signs of an upcoming influenza epidemic. The basic idea behind using Google search queries is to find a measure for the relative importance of a subject for the individuals that are directly concerned. Da et al. (2011) use this relation to predict stock price returns while Bank et al. (2011) explain time-varying liquidity supply in the German stock market. Choi and Varian (2012) were the first to use Google search query data in predicting unemployment figures. They find the addition of Google data particularly useful to identify turning points in unemployment.

This study exploits both the advances in the literature on modeling long-memory components and the literature on mixed-frequency nowcasting. Combining these methodologies we can conduct an analysis of the relationship between unemployment and government bond yields at a high-frequency. Changes in unemployment are preceded by changes in relevant search activity several months earlier. Using the information available in search activity data significantly increases the nowcasting accuracy of unemployment for each of the nine considered countries. The reaction of government bond yields to a rise in unemployment is country-specific. In Europe, government bond yields increase in reaction to a shock in unemployment. This result is in line with economic theory which suggests that countries with higher unemployment face increased insolvency risk. Government bond yields of Australia and the United States decrease in reaction to a shock in unemployment. This might be explained by a policy of more aggressive monetary interventions, as pursued in particular by the US federal reserve system.

The remainder of this chapter is structured as follows. Section 4.2 introduces the M-HAR method to nowcast unemployment and presents the heterogeneous VAR that is used to investigate the relationship of bonds and unemployment. Section 4.3 describes the data and extensively discusses the intricacies of working with Google's search volume data. Section 4.4 presents the empirical results and section 4.5 concludes.

4.2 Methodology

4.2.1 Nowcasting with Long-Run Relations

The M-HAR model can produce nowcasts of a low-frequency series by drawing on information from a related long-memory high-frequency series, where the timing of important informative events is not obvious ex-ante. Consider the case of Google searches and unemployment. How long does it take for an increased interest in unemployment benefits due to economic distress to translate into actual unemployment figures? How far in advance do individuals know they might be laid off? The relevant time frame could span a single week (e.g. in a hire-and-fire state like Louisiana in the US) or months (e.g. in Germany where the majority of workers have a three months notice period). To capture the entire, possibly important time structure in a vector-autoregression, a large number of lags is needed. Estimating such a model, however, suffers from the curse of dimensionality. In the M-HAR model this long lag structure is replaced by a much more parsimonious heterogeneous frequency cascade.

Let y_{t+1} be the low-frequency time series for which a nowcast is needed. Assume that y_{t+1} follows its own autoregressive structure of order p such that

$$\begin{aligned} y_{t+1} &= \phi_1^{(y)} y_t + \phi_2^{(y)} y_{t-1} + \dots + \phi_p^{(y)} y_{t-p+1} + \varepsilon_t \\ &= \sum_{j=1}^p \phi_j^{(y)} y_{t-j+1} + \varepsilon_t, \end{aligned} \quad (4.1)$$

where ε_t follows some distribution with finite first and second moments. Now let $x_{t,i}$ be a related higher frequency time series that can be used to produce nowcasts of y_{t+1} . Introducing a suitable structure of x into equation (4.1) results in

$$y_{t+1} = \sum_{j=1}^q \sum_{i=1}^{\tau} \phi_{j,i}^{(x)} x_{t-j+1,i} + \nu_{t,i} + \sum_{j=1}^p \phi_j^{(y)} y_{t-j+1} + \varepsilon_t, \quad (4.2)$$

where τ is the number of partitions of the interval $[t, t + 1]$ and q is the number of low-frequency lags of $x_{t,i}$. In principle, the parameters in equation (4.2) could be directly estimated. However, equation (4.2) potentially includes a large num-

ber of lags (namely $q \cdot \tau$) of the higher frequency series x , which might ultimately render estimation infeasible. Instead let us replace the homogeneous autoregressive structure of x with a parsimonious heterogeneous frequency cascade $\mathbb{C}(S)$ that captures the original autoregressive properties. S is a set of frequencies (e.g. weekly, monthly, quarterly, ...) and \mathbb{C} is the corresponding operator that aggregates x to the respective frequencies in S . Replacing the first term on the right hand side of equation (4.2) by this frequency cascade leads to

$$y_{t+1} = \mathbb{C}(S)x_{t,i} + \nu_{t,i} + \sum_{j=1}^p \phi_j^{(y)} y_{t-j+1} + \varepsilon_t \quad \forall i \in [1, \tau]. \quad (4.3)$$

In our application, unemployment data is available on a monthly basis while Google search volume is available on a weekly basis. In order to nowcast the monthly unemployment u_{t+1} we propose a frequency cascade structure of weekly data supplied with monthly, quarterly and yearly aggregates to cover the long-memory property of search query data:

$$u_{t+1} = c_w g_{t,i}^{(w)} + c_m g_{t,i}^{(m)} + c_q g_{t,i}^{(q)} + c_y g_{t,i}^{(y)} + \nu_{t,i} + \sum_{j=1}^p \phi_j^{(u)} u_{t-j+1} + \varepsilon_t \quad \forall i \in [1, 4], \quad (4.4)$$

where $g_{t,i}^{(w)}$ is the weekly search volume series as provided by Google, and $g_{t,i}^{(m)}$, $g_{t,i}^{(q)}$ and $g_{t,i}^{(y)}$ are series aggregated to monthly, quarterly and yearly averages, respectively. If we wanted to cover a similar time span without using the frequency cascade, we would have to include $12 \times 4 = 48$ lags of $g_{t,i}^{(w)}$. The higher frequency Google series $g_{t,i}^{(w)}$ has an additional amount of $\tau = 4$ observations for each time period t , allowing for four consecutive nowcasts before observing u_{t+1} :

$$\begin{aligned} \hat{u}_{t+1,1} &= c_w g_{t,1}^{(w)} + c_m g_{t,1}^{(m)} + c_q g_{t,1}^{(q)} + c_y g_{t,1}^{(y)} + \sum_{j=1}^p \phi_j^{(u)} u_{t-j+1} \\ \hat{u}_{t+1,2} &= c_w g_{t,2}^{(w)} + c_m g_{t,2}^{(m)} + c_q g_{t,2}^{(q)} + c_y g_{t,2}^{(y)} + \sum_{j=1}^p \phi_j^{(u)} u_{t-j+1} \\ \hat{u}_{t+1,3} &= c_w g_{t,3}^{(w)} + c_m g_{t,3}^{(m)} + c_q g_{t,3}^{(q)} + c_y g_{t,3}^{(y)} + \sum_{j=1}^p \phi_j^{(u)} u_{t-j+1} \\ \hat{u}_{t+1,4} &= c_w g_{t,4}^{(w)} + c_m g_{t,4}^{(m)} + c_q g_{t,4}^{(q)} + c_y g_{t,4}^{(y)} + \sum_{j=1}^p \phi_j^{(u)} u_{t-j+1}. \end{aligned} \quad (4.5)$$

The parameters c_i and $\phi_j^{(u)}$ in equation system (4.5) are estimated using quasi maximum likelihood (QML).

4.2.2 Unemployment and Bonds – Heterogeneous VAR

In order to analyze the relationship between unemployment and government bonds we draw on the well-established framework of vector autoregressive (VAR) models. The VAR model is implemented on a weekly basis. The nowcasts of unemployment based on the Google search volume data are matched with the quarterly observed data to construct a weekly time series. We then implement a heterogeneous VAR which accounts for past observations up to three years. Denote the weekly unemployment changes by Δu_t and the weekly bond yields by b_t . The set S includes the frequencies of one week, one, three and six months, and one, two and three years. The model reads as follows:

$$\begin{aligned} b_t &= a_{10} + \mathbb{C}_{11}(S)b_{t-1} + \mathbb{C}_{12}(S)\Delta u_{t-1} + \epsilon_{1,t} \\ \Delta u_t &= a_{20} + \mathbb{C}_{21}(S)b_{t-1} + \mathbb{C}_{22}(S)\Delta u_{t-1} + \epsilon_{2,t}. \end{aligned} \quad (4.6)$$

As we only have a reduced-form structure, the model can be estimated linewise with OLS.

4.3 Data Description

4.3.1 Unemployment Rates and Bond Yields

We analyze the relationship between unemployment rate changes and bond yields for a panel of nine countries: seven European countries (Austria, France, Germany, Portugal, Spain, Switzerland, and the UK) and two non-European countries (Australia and the US). All monthly unemployment series are obtained from Eurostat,¹ except for Australia and Switzerland for which the series are available on the Aus-

¹ec.europa.eu/eurostat

tralian Bureau of Statistics² and the Swiss Amstat³ websites, respectively. The time period considered is January 2004 (when Google search volume data became available) until April 2014. Daily government bond series are obtained from Datastream. Ten-year government bond yields are computed by subtracting the respective short-term bond rates. For the European countries for which short-term bond data are not available the EURIBOR is used instead.

Figure 4.1 depicts time series plots of the unemployment rates (upper graph) and bond yields (lower graph) for all countries in the panel. At first glance we can already observe some co-movement between rises in unemployment and bond yields. Furthermore, the effect of the European debt crisis is clearly visible for Spain and Portugal: both countries experienced a sharp increase in unemployment, starting during the financial crisis in 2008, which is accompanied by a similar rise in long-term government bond yields.

For the analysis, all unemployment series are seasonally adjusted and transformed into first differences. Subsequently, they are tested for stationarity using augmented Dickey-Fuller (DF) tests and are all found to be $I(0)$. Following Campbell et al. (1997) we calculate the yield spread using the theoretical cointegrating vector $(1, -1)$ and perform a cointegration analysis. We generally find p -values that reject the cointegration hypothesis. This finding shows that the yield still has a long-memory property which we will account for using the cascading lag structure as outlined in section 4.2. The fact that the DF test fails to reject the null hypothesis of a unit root might well be due to the lacking power of this test. Furthermore, the literature on ARFIMA models documents that the DF test fails to distinguish unit root behavior (with $d = 1$) from near unit root behavior ($d < 1$). We therefore also calculate the fractional difference parameter d using the method of Geweke and Porter-Hudak (1983) and find it to be close to one, depending on the bandwidth selection criteria. Again no clear answer can be provided by this test. Ultimately, we rely on the theoretical result in Campbell et al. (1997) and assume stationarity

²www.abs.gov.au

³www.amstat.ch

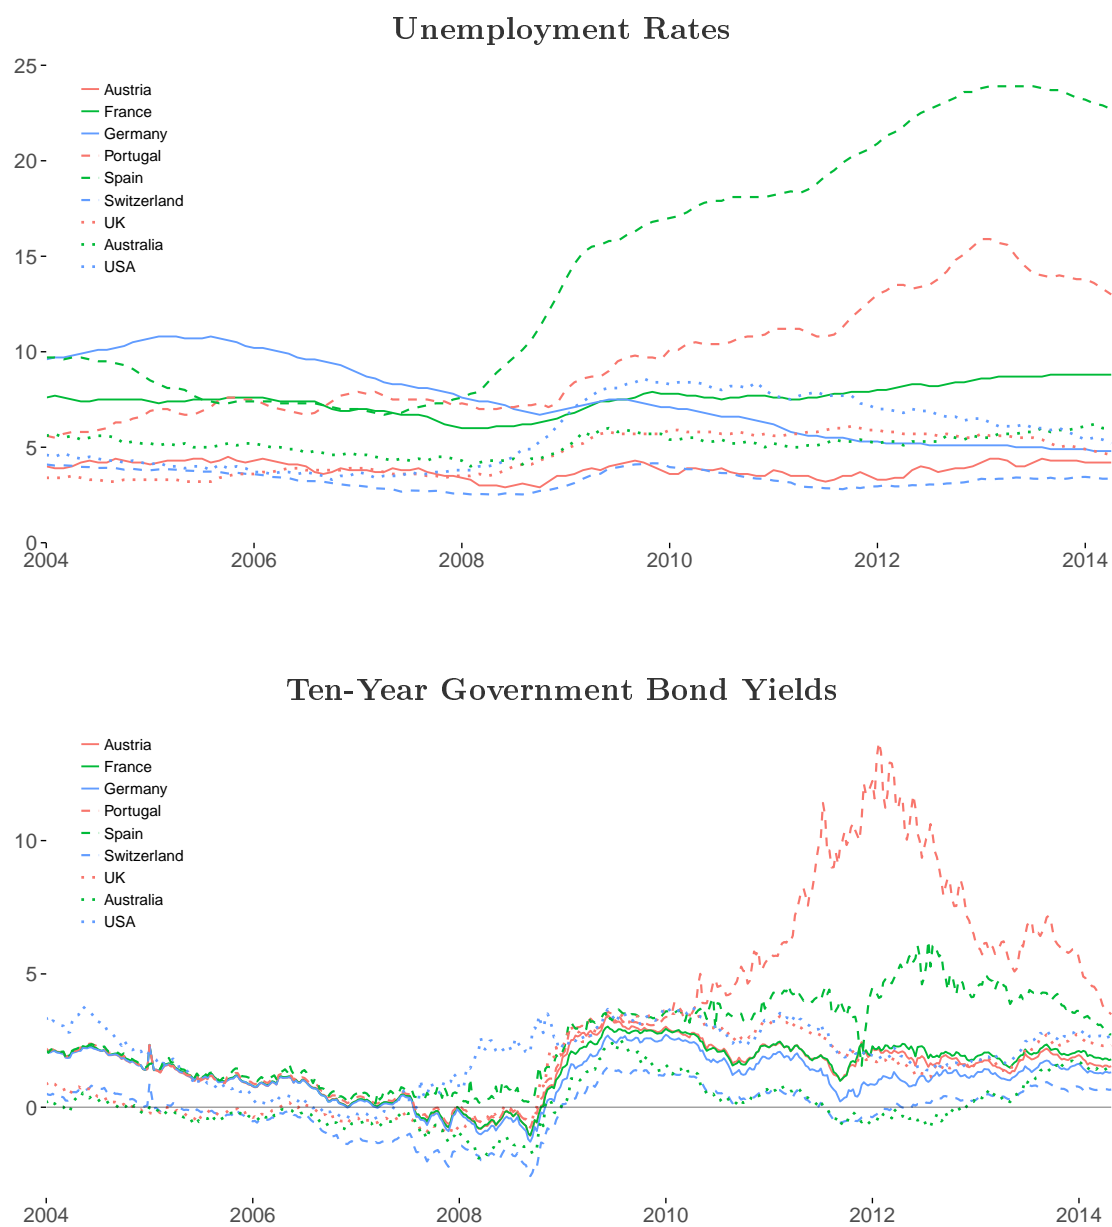


Figure 4.1: The figure presents time series plots of monthly unemployment rates per country (upper graph) and daily ten-year government bond yields (lower graph). Both series are depicted in percent, from January 2004 to April 2014.

Descriptive Statistics – Unemployment Changes and Bond Yields

Table 4.1: The table provides the mean (denoted by \bar{x}), the standard deviation σ , and the skewness coefficient ς of monthly unemployment changes Δu and weekly ten-year government bond yields b for the nine countries listed in the first column.

	$\overline{\Delta u}$	$\sigma_{\Delta u}$	$\varsigma_{\Delta u}$	\overline{b}	σ_b	ς_b
Austria	0.00	0.14	0.26	1.43	1.00	-0.50
France	0.01	0.08	-0.10	1.46	0.99	-0.69
Germany	-0.04	0.09	0.18	1.11	0.90	-0.50
Portugal	0.06	0.18	0.14	3.69	3.50	0.86
Spain	0.11	0.24	1.04	2.56	1.58	0.07
Switzerland	-0.01	0.07	0.78	-0.07	0.92	-0.64
UK	0.01	0.10	0.59	1.23	1.44	0.10
Australia	0.00	0.14	0.29	0.11	0.91	0.45
USA	0.00	0.18	0.57	1.96	1.15	-0.46

for the yield spreads while being aware that there is still a long-memory property which we have to account for.

Table 4.1 provides descriptive statistics. On average, only Germany and Switzerland could reduce the unemployment rate between 2004 and 2014. In Austria, Australia and the US, the unemployment rate remained roughly stable while it grew slightly in France and in the UK. In Spain and Portugal the sharp rise in unemployment after the financial crisis is reflected by the higher monthly growth rates: for Spain, this rate is ten times higher (on average) than for France or the UK. Similarly, these two countries exhibit the highest weekly bond yields, followed by the US and the UK. The lowest bond yields are observed for Australia and Switzerland; in case of the latter it is even negative. We also find that unemployment changes are generally left-skewed while the skewness of bond yields is not clearly drawn to left or right.

4.3.2 Google Search Volume

Ideally we would like to know how many people engage in preparations to file for unemployment and how far in advance of becoming unemployed they do so. With this information we could easily nowcast unemployment figures. This would also solve the drawback that even month-end data, which correspond to the reporting frequency, are not immediately available on the first day of the following month, but only with a lag. Of course we do not have this exact data, but we can try to approximate how many people prepare facing unemployment with the help of Google search volume data (see Choi and Varian, 2012; Barreira et al., 2013).

The data is obtained from Google Trends,⁴ where normalized weekly search volume data is provided for any search query that surpasses a certain volume threshold. These series are available starting in January 2004. The raw search volume data is corrected for multiple searches from one IP address and normalized by the total search volume at that time:

$$v_{s,t} = \frac{V_{s,t}}{V_t}, \quad (4.7)$$

where $v_{s,t}$ is the normalized search volume at time t for search query s , $V_{s,t}$ is the actual search volume, and V_t the total search volume. To better be able to compare changes in search terms with different volume levels, the normalized volumes $v_{s,t}$ are rescaled by their maximum historical value. The normalized search volume index $g_{s,t}$ is thus given as

$$g_{s,t} = \frac{v_{s,t}}{\max_t \{v_{s,t}\}} \cdot 100. \quad (4.8)$$

This preparation and pre-selection of the data by Google causes several complications. Firstly, the volume threshold criterion restricts the set of keywords we can use for our analysis. This is problematic especially for smaller countries. Secondly, the rescaling of the normalized volume series $v_{s,t}$ by its historical maximum causes the whole series to change when a new maximum occurs.

When working with Google's search data, it is vital to keep in mind what is measured: The number of times a keyword is entered into the search engine. To

⁴www.google.com/trends

make assumptions about the intentions that an individual had when she entered a query, the choice of keywords is crucial. They need to be specific to the intention and at best unique. However, most of them will also appear in other possibly related or – worse – competing contexts. The famous prediction of influenza epidemics with search query data by Ginsberg et al. (2009) is well-known to suffer from this keyword ambiguity. Some symptoms like fever or headache are not specific to having the flu, but occur in the course of numerous illnesses. Entering the keyword “fever” is, thus, not necessarily related to the illness of interest. When a fever causing epidemic unrelated to the flu spreads, the flu index rises erroneously.

Besides the choice of keywords, the context in which search volume data is used to draw conclusions about intentions can favor or hinder the application. Search queries are usually performed to acquire information on a specific topic. When the topic is only of interest to individuals when they are in a certain situation, then we can infer that changes in search volume indicate a higher amount of individuals in that situation (being sick, being unemployed). Still, this only works if the situation occurs rarely, so that individuals are not yet informed from previous experience and do not need to search for information anymore. When an individual suffers from flu symptoms in the flu season, he might not search for information on it anymore. For rare symptoms no such learning can occur.

These insights are very important for the application to unemployment. The context is favorable to the application because when an individual faces unemployment it is likely a rare event, unlike getting the flu. The situation is likely to be unfamiliar and new information needs to be retrieved. To make the keywords we use as situationally specific as possible, we focus on the process of filing for unemployment and receiving unemployment benefits. We sort all keywords into four categories, each linked to the search for information on the process of filing for unemployment. These categories guide the choice of keywords in the different languages and countries and ensure comparability. The category “Benefits” relates to queries about unemployment benefits while “Process or Agency” subsumes queries on how and where to file for unemployment. These two categories are very specific. The categories “Unemployed” and “Unemployment” are more general and as such more prone to be used

List of Search Terms and Descriptive Statistics

Table 4.2: The table presents the keywords used to extract Google search volume data and descriptive statistics thereof. Keywords marked with a * account for possible misspelling. v_{rel} is the relative average search volume normalized by the series denoted with $v_{rel} = 100$. For example $v_{rel} = 10$ means that the average search volume of that series is ten times lower than for the base series. σ is the standard deviation of the respective search volume over time.

Country	Benefits	Keyword Categories											
		v_{rel}	$\sigma_{\Delta v}$	Process or Agency		v_{rel}	$\sigma_{\Delta v}$	Unemployed		v_{rel}	$\sigma_{\Delta v}$	Unemployment	
Austria				ams		100	0.08	arbeitslos*		1	0.17		
France	allocation chômage*	14	0.28					au chômage*		14	0.21	chômage*	
	calcul chômage*	7	0.23									100 0.20	
Germany	arbeitslosengeld	61	0.10	hartz iv		89	0.21	arbeitslos*		29	0.11	arbeitslosigkeit*	
	arbeitslosengeld*	96	0.09	hartz iv*		100	0.20					39 0.10	
	arbeitslosengeld berechnung*	11	0.26										
Portugal	subsídio desemprego	35	0.17					desemprego		100	0.26		
	subsídio desemprego*	43	0.19										
Spain	subsídio desempleo*	8	0.13	inem*		100	0.21	en paro		8	0.22	paro	
												33 0.18	
Switzerland	arbeitslosengeld*	6	0.33	arbeitslosenkasse*		14	0.19	arbeitslos*		9	0.19	arbeitslosigkeit*	
				rav		100	0.16					11 0.31	
UK	jobseeker's allowance*	100	0.15					unemployed*		34	0.19	unemployment*	
	jobseeker's allowance rate*	2	0.34									62 0.15	
Australia	unemployment benefits*	38	0.10					unemployed*		27	0.16	unemployment*	
												100 0.15	
USA	unemployment benefits	6	0.23	file for unemployment*		2	0.24	unemployed*		4	0.25	unemployment	
	unemployment benefits*	13	0.27									unemployment*	
	unemployment benefits rate*	2	0.28									100	

by individuals not faced with unemployment. Still they are very specific to the context of unemployment. Out of these two categories, “Unemployed” is more likely to be part of individuals’ sentence queries like “I am unemployed”, while the category “Unemployment” more likely captures general attention. Table 4.2 provides an overview of the keywords used and the categories they represent. Queries marked with a star (*) are combined queries of several similar words and include different (mis-)spellings. Not all categories are represented for each country due to limited availability as a result of low search volume. We intentionally restrict ourselves to a very precise set of queries closely related to unemployment. Other studies like Askitas and Zimmermann (2009) and D’Amuri and Marcucci (2010) include different search terms that individuals who face unemployment might enter. The intention behind these queries is not to seek information about receiving unemployment benefits, but finding employment. Such queries range from searches for “jobs” or names of specific job search engines (like “indeed” or “job24”). We find that the volume of these queries behaves differently in every country.

An example is presented in Figure 4.2. The graphs depict the first principal component of changes in normalized search volume for queries related to jobs and unemployment benefits. For the United States the sharp increase in search volume of both series occurs almost at the same time during the crisis. For the United Kingdom however this does not hold. Here, the sharpest increase in search volume for jobs is observed years after the crisis. Searches for specific job websites have the additional problem that fluctuations in popularity lead to changes in search volume which are completely unrelated to unemployment. Moreover, the keywords would have to be frequently adjusted as new websites emerge. In our sample, only unemployment-related searches consistently occurred ahead of changes in unemployment for all nine countries. A possible explanation is that job search related queries are not necessarily specific to individuals in the situation of facing unemployment. They can just as well be a sign of a growing job market. The series just happen to be very similar for the US, which may be due to specifics of the social security system coupled with a high pressure to find a new job quickly.

Changes in Search Volume for "Jobs" and "Unemployment Benefits"

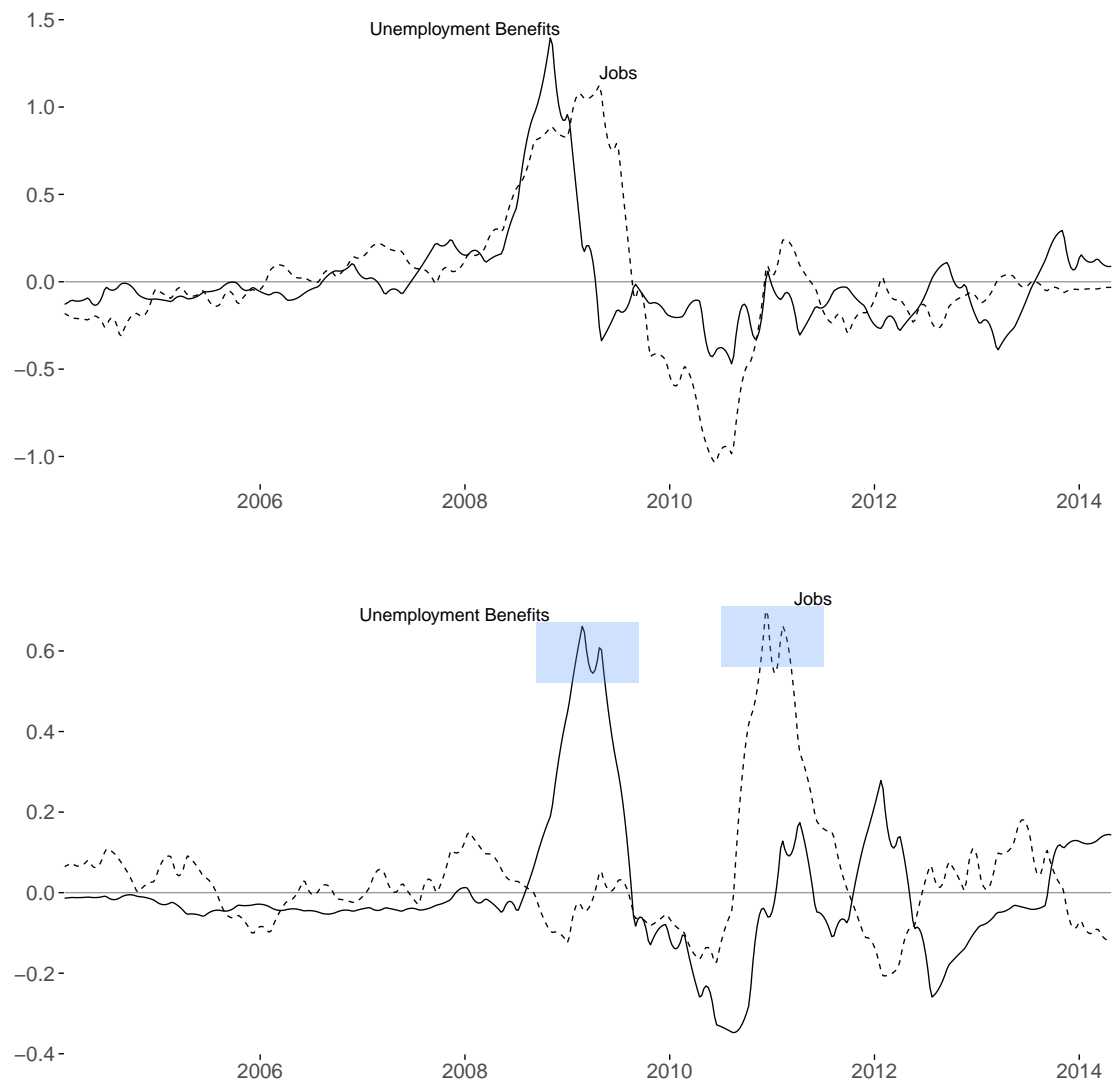


Figure 4.2: The figure presents time series plots illustrating the lagged co-movement of search volume for the keywords “jobs” and “unemployment benefits”. The upper graph presents data for searches conducted in the United States, the lower graph is UK data.

All search queries listed in Table 4.2 are different measures of the same underlying quantity, namely impending unemployment. We perform a principal component analysis to reduce the dimensionality of search volume data and produce a single indicator time series. Before computing the principal components, each normalized search volume series is smoothed and first differences are formed. These steps are necessary to ensure that the variance-maximizing principal component analysis does not simply pick the noisiest series as the most important variance contributor. To evaluate the fit of the first identified principal component with the search volume data of the different categories we compute the average correlation per category. Overall, the principal component represents all categories well, i.e. a high correlation is observed for all nine countries. The weakest relationship is found for the “Process” category and the principal component for Austria where the correlation is still 0.457. Figure 4.3 illustrates the results.

4.4 Results and Discussion

The analysis of the relation between changes in unemployment and government bond yields is separated into two steps. In a first step we use weekly Google search volume data to nowcast monthly unemployment changes on a weekly frequency employing the M-HAR method. In a second step we use the resulting higher frequency unemployment series to analyze the relationship with government bond yields in a heterogeneous VAR framework.

4.4.1 Nowcasting Unemployment – M-HAR Estimation

The M-HAR model requires the identification of a suitable cascading structure. As outlined in Subsection 4.2.1, the number of relevant lags might differ across countries. Therefore, we use cross-correlations between unemployment changes and the principal component of normalized search volume changes to determine an appropriate structure. Figure 4.5 depicts cross-correlations at different lags of search volume changes for France and the United States. The search volume changes feature a high




Figure 4.3: The figure illustrates the correlation of the extracted principal component with the search volume in the four categories “Benefits”, “Process”, “Unemployed” and “Unemployment” for the different countries.




Figure 4.4: The figure shows how many months h ahead of a change in unemployment, the correlation with past search query data is highest, i.e. a scatter plot of h against $\max_h \{\text{cor}(\Delta u_t, \Delta g_{t-h})\}$.

persistence as correlations change slowly between lags. One can observe that the correlation peaks at different lags: -2 months for the US, -5 months for France. The strength of the correlation also varies. This may be yet another indicator that the timing when people search for information about unemployment is country-specific.

Figure 4.4 provides an overview of the two key characteristics – time and strength – of the cross-correlation functions for all nine countries in the panel. The scatter plot depicts the maximum correlations between unemployment changes and search volume changes alongside the lag at which these occur. The blue points mark the two non-European countries, crisis countries are labeled by red squares, and the remaining European countries are tagged by green triangles. Again, we observe that there are large differences in the timing of the maximum correlation peak and the size of the correlation. There also seems to be a pattern that the closer the timing of changes in either time series, the higher is the correlation between them. This should be intuitively clear since a larger distance in time allows for more unrelated innovations to dilute the relation.

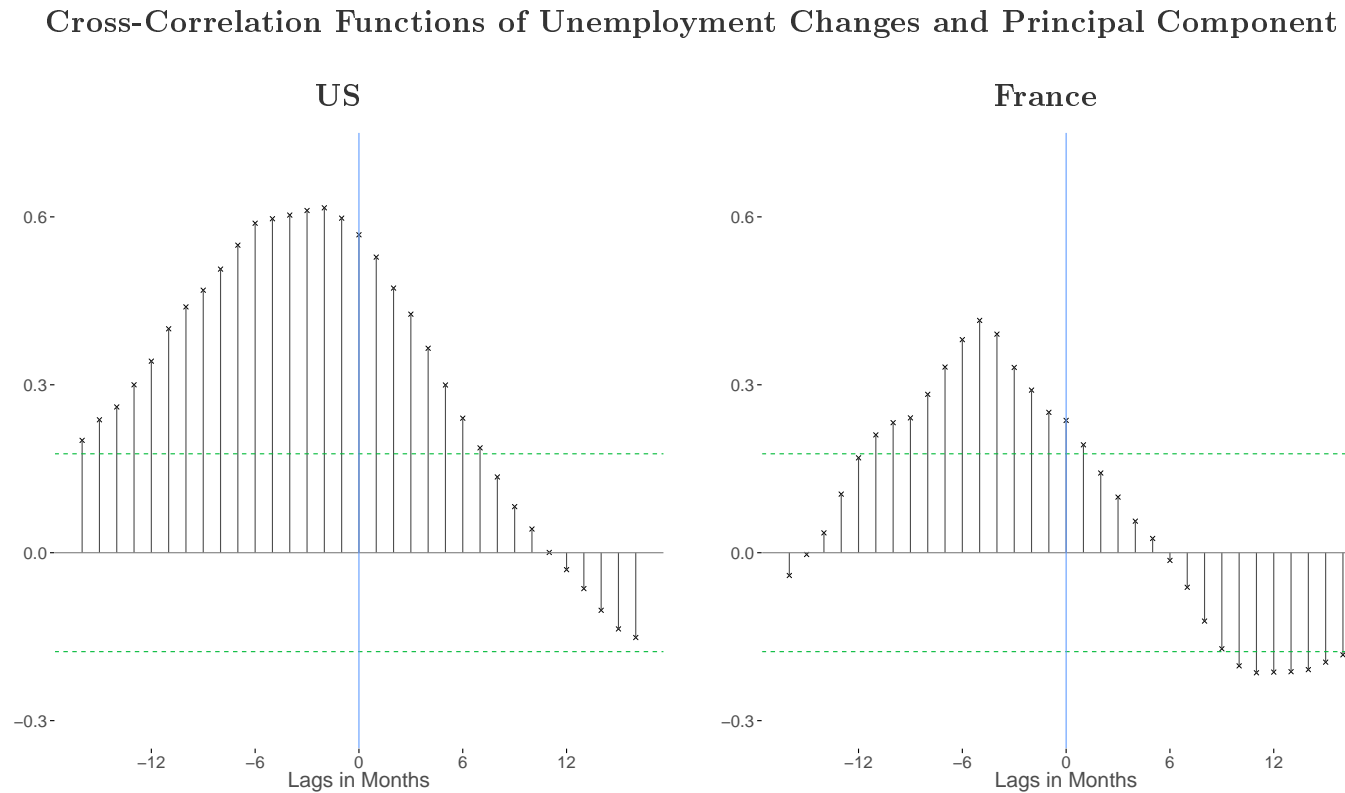


Figure 4.5: The figure presents the cross-correlation function of the current change in unemployment rates with different leads and lags of the principal component extracted from the search volume changes. The left graph depicts the situation for US data while the right graph presents data from France. The blue vertical line denotes lag 0. The green dotted horizontal lines represent 95% confidence intervals.

We use the M-HAR model based on equation (4.5) to produce four weekly nowcasts of unemployment each month. Figure 4.4 shows that the peak of the cross-correlation for Austria occurs at a lag of eight months. This is the longest distance observed, so a frequency cascade structure that covers one year should be able to capture a sufficient amount of lags. The 48 weekly lags are represented in the cascade by five frequency factors (week, month, three months, six months, year). The benefit of including the search volume frequency cascade in the estimation is evaluated by comparing the nowcasts to a forecast based on a pure autoregressive (AR) model of unemployment. The number of monthly lags of the autoregressive component for unemployment in both models is chosen in a way to guarantee the best model fit of the pure AR model. Selection criteria are forecasting performance, lag significance and BIC. The maximum of the optimal lag length p turns out to be three. Lag length is determined for each country individually. Parameters are estimated via QML. Table 4.3 summarizes the results.

The magnitude of the cascade coefficient estimates, as well as their significance, closely mirrors the time at which the respective cross-correlation functions peak in Figure 4.4. This result highlights that the cascade coefficients pick up the information occurring at these lags. Countries where the peak in the cross-correlation function occurs less than four months ahead have insignificant year coefficients in the M-HAR estimation. For countries with a peak more than four months ahead the year coefficients are significant. The pattern is more pronounced when the size of the coefficients is considered as well. The two exceptions are Portugal and the United Kingdom. For Portugal the overall high variance of the parameter estimates might indicate a data quality issue. The overall search volume for unemployment-related topics is so low that many queries that are used for other countries are below the Google-imposed threshold and thus unavailable. For the United Kingdom most coefficients have relatively low p -values but the largest absolute values are still observed before the four-months mark. The coefficient estimates for Spain are particularly large due to the big changes in unemployment during the financial and debt crises.

Mean squared forecasting error (MSE) and adjusted R^2 are reported in the last two columns of Table 4.3 to measure the benefit of including the Google search vol-

Nowcasting Unemployment – Estimation Results

Table 4.3: The model presents the coefficient estimates for the model used to nowcast unemployment. p -values are given in parentheses. Model fit is evaluated by the mean squared error (MSE) and the adjusted R^2 in the last two columns. The first row of each country is the pure autoregressive model, the third row is the M-HAR model. * (**) indicates that the reduction of the MSE is significant on the 5% (1%) level.

Country	$\phi_1^{(u)}$	$\phi_2^{(u)}$	$\phi_3^{(u)}$	c_w	c_m	c_q	c_h	c_y	σ_ε	MSE	R^2
Austria	0.28 (0.00)	-0.31 (0.00)							0.13 (0.00)	1.74	13.9
	0.21 (0.00)	-0.37 (0.00)		0.03 (0.96)	0.00 (1.00)	-0.03 (0.96)	-0.66 (0.11)	0.75 (0.00)	0.13 (0.00)	1.57**	22.1
France	0.39 (0.00)	0.17 (0.00)							0.07 (0.00)	0.55	23.9
	0.25 (0.00)	0.11 (0.01)		0.06 (0.70)	0.11 (0.64)	-0.43 (0.00)	0.44 (0.00)	-0.10 (0.02)	0.07 (0.00)	0.47**	38.4
Germany	0.53 (0.00)	0.24 (0.00)							0.07 (0.00)	0.43	32.9
	0.47 (0.00)	0.21 (0.00)		0.30 (0.05)	-0.26 (0.21)	0.02 (0.87)	-0.03 (0.77)	0.05 (0.16)	0.06 (0.00)	0.41*	35.4
Portugal	0.73 (0.00)								0.13 (0.00)	1.82	48.8
	0.61 (0.00)			-0.31 (0.26)	0.33 (0.38)	-0.17 (0.45)	0.26 (0.14)	0.05 (0.58)	0.13 (0.00)	1.67**	53.6

Table 4.3: Nowcasting Unemployment – Estimation Results (cont.)

Country	$\phi_1^{(u)}$	$\phi_2^{(u)}$	$\phi_3^{(u)}$	c_w	c_m	c_q	c_h	c_y	σ_ε	MSE	R^2
Spain	0.73 (0.00)	0.25 (0.00)	-0.07 (0.13)						0.12 (0.00)	1.36	75.0
	0.66 (0.00)	0.25 (0.00)	-0.07 (0.17)	1.38 (0.00)	-2.18 (0.00)	1.48 (0.00)	-0.71 (0.03)	0.13 (0.35)	0.11 (0.00)	1.25**	77.5
Switzerland	0.16 (0.00)	0.29 (0.00)	0.31 (0.00)						0.06 (0.00)	0.35	36.9
	-0.02 (0.62)	0.16 (0.00)	0.25 (0.00)	0.19 (0.04)	-0.29 (0.02)	0.28 (0.00)	-0.11 (0.08)	0.00 (0.95)	0.05 (0.00)	0.28**	48.9
UK	0.26 (0.00)	0.10 (0.03)	0.26 (0.00)						0.09 (0.00)	0.80	20.8
	0.17 (0.00)	0.05 (0.33)	0.18 (0.00)	-0.19 (0.38)	0.65 (0.03)	-0.31 (0.11)	-0.21 (0.14)	0.16 (0.01)	0.09 (0.00)	0.73**	28.6
Australia	-0.16 (0.00)								0.14 (0.00)	2.00	2.5
	-0.26 (0.00)			-0.50 (0.15)	0.69 (0.15)	-0.49 (0.07)	0.51 (0.01)	0.04 (0.67)	0.13 (0.00)	1.74**	15.4
USA	0.15 (0.00)	0.34 (0.00)	0.15 (0.00)						0.16 (0.00)	2.49	23.7
	-0.13 (0.01)	0.04 (0.42)	-0.07 (0.14)	-0.24 (0.31)	0.36 (0.27)	-0.06 (0.78)	0.09 (0.54)	0.06 (0.33)	0.14 (0.00)	1.90**	41.9




Figure 4.6: The figure compares the model fit of the pure autoregressive forecast model (Naive) and the M-HAR model based on mean squared error (left graph) and adjusted R^2 (right graph). The dashed 45-degree line indicates the reference line for equal performance of the two models.

ume frequency cascade in addition to a pure autoregressive model of unemployment. Overall, we find a significant reduction in MSE and a higher adjusted R^2 for all countries.⁵ The degree of improvement, however, varies greatly. Figure 4.6 depicts the changes as a scatter plot of MSE and R^2 for the pure AR and the M-HAR model. Interestingly, the largest improvements in both MSE and R^2 can be observed for the United States, even though none of the frequency cascade estimates are statistically significant. Large improvements in both criteria can also be observed for Australia. France and Switzerland, which already feature a low MSE in the pure model, display a large increase in R^2 . For the solvent European countries the unemployment forecasts based on the AR model result already in a comparatively low MSE. Therefore, the inclusion of the search query data has no large impact. The R^2 values indicate that while the search volume data does not greatly improve the average forecasting error in this case, it nevertheless helps explaining a larger fraction of the total variation in unemployment. For the non-European countries and the European countries hit by the solvency crisis the inclusion of search data distinctly improves the prediction of unemployment changes by both decreasing the initially high forecast error as well as explaining a larger part of total variation.

The plots in Figure 4.7 present a comparison of the M-HAR model and the autoregressive model based forecasts during the crisis for the United States and France. The vertical lines represent the observed changes in the official unemployment figures, while the solid line represents the nowcasts of the M-HAR model and the dashed line is the AR-based nowcast. Due to the additional search volume data, the M-HAR forecast can anticipate movements in unemployment much better than the autoregressive model. For the United States the large changes during the crisis cannot be inferred from its past changes alone, so the pure autoregressive model forecasts deviate by a large margin from actual data, which results in a high MSE for this model. The M-HAR forecast utilizes the increase in unemployment-related search activity that takes place one month ahead (cp. Figure 4.5), improving the forecast noticeably. But even for France, where the autoregressive model already

⁵Differences in MSE are tested for significance using the Diebold-Mariano test (Diebold and Mariano, 1995).

AR and M-HAR Nowcasts




Figure 4.7: The figure shows a comparison of the unemployment change nowcasts based on the pure autoregressive model (dotted line) to the one based on the M-HAR model (solid line). The horizontal lines represent the actual changes. The left graph is for the US, the right graph for France.

performs well, the inclusion of unemployment related search activity improves the forecasts.

4.4.2 Unemployment and Bonds – Heterogeneous VAR Estimation

To estimate the relationship between unemployment changes and bond yields we employ a heterogeneous VAR as described in equation (4.6) using the higher frequency unemployment series provided by the M-HAR model. The frequency cascades used in the heterogeneous VAR have to cover a longer time span than in the M-HAR model in the previous subsection to account for the extreme long-memory feature of the government bond yield time series (see also Diebold and Li, 2006). We find that a period of three years captures a sufficiently long lag structure, such that the impact of innovations is not permanent but slowly decreases over time. To cover this lag structure we use seven heterogeneous AR-terms, i.e. two frequency terms more than in the previously used cascade. Like in the standard VAR context, increasing the number of parameters in the HVAR leads to larger variances of the estimates,




Figure 4.8: The figure illustrates the impact of the lag structure of unemployment rate changes and bond yields on government bond yields. Significance of the heterogeneous autoregressive parameters is illustrated through fading blue where dark blue is a p -value < 0.001 and white is an insignificant parameter with a p -value > 0.2 .

which would in turn call for a parsimonious modeling strategy. Instead of using the same frequency cascade for each country, we could use cascades that are tailored to fit the individual countries best, i.e. leaving out insignificant frequencies or adding additional ones where necessary. This decreases the variance of the remaining estimates but it also leaves ample room for ambiguity. Instead, we opt for keeping the estimation as general as possible and comparable across countries, thereby accepting the fact that parameter estimates suffer from higher variance.

Figure 4.8 visualizes the estimation results of the bond equation, only displaying significant frequency estimates.⁶ The shading indicates the p -value of the parameter estimates on a scale from < 0.001 (blue) to > 0.2 (white). For Switzerland, Spain, and the US we find a rather timely response of bond yields to unemployment changes. For the remaining countries, it takes longer for the bond yields to react. However, the impact seems stronger in the latter case where the largest values of the estimates can be observed between the frequencies of six month and two years. The order of magnitude of the majority of the insignificant estimates is considerably lower than the reported values. Bond yield estimates are around one for the one

⁶Detailed estimation results are reported in the appendix.

week frequency which is due to the high persistence of the time series. This effect is countered by subsequent negative estimates on lower frequencies. The insignificant estimates of the bond yields cascade structure are again small in magnitude, similar to the insignificant unemployment cascade estimates. While for the impact of unemployment a more country-specific pattern can be observed, the autoregressive structure of the bond yields is rather similar across all countries.

Figure 4.9 provides impulse-response functions of bond yields to an idiosyncratic shock of one standard deviation in unemployment for all countries. Contemporaneous effects are identified via Cholesky decomposition. The ordering is not crucial in the present context as the important influences are not contemporaneous but lagged. The information contained in the impulse-response functions is condensed in Figure 4.10. The scatter plot depicts the peaks of the impulse-response functions and the lag at which they occur. The graph illustrates an interesting pattern. Government bond yields of all European countries react positively to idiosyncratic shocks in unemployment. This matches the intuition that a higher unemployment rate increases future debt of the country and decreases tax income revenue. This negatively impacts future solvency and increases the risk involved in holding long-term government bonds. For the European crisis countries the impact peaks at a much later stage, implying that the impact is much more persistent. This may be due to differences in the local economic systems and the way shocks to unemployment can be absorbed quickly. The German labor market was arguably the most flexible one during the crisis, which is reflected in this graph as well. The bond markets in the United States and Australia react differently to idiosyncratic shocks in unemployment. Here, shocks to unemployment decrease government bond yields. In the United States the shock is absorbed relatively quickly while the Australian system adjusts more slowly. Differences in labor market efficiency may serve as an explanation.

The fact that the response of bond yields is negative may be an indication of a different monetary policy in the United States and Australia compared to that of the European Union. The most likely explanation is that the European Central Bank's core task "shall be to maintain price stability" (EU, 2012). Therefore, rising

Impulse-Response Functions of Bond Yields to a Shock in Unemployment

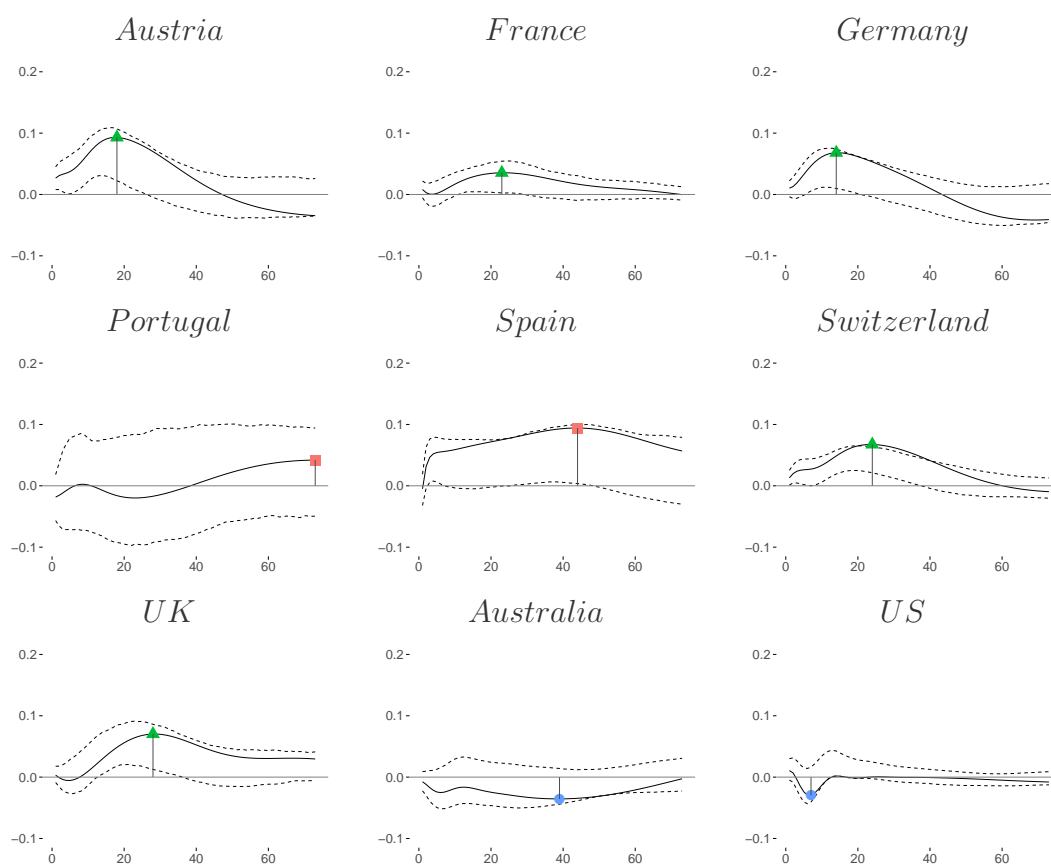


Figure 4.9: The figure presents the impulse-response functions of bond yields to a one standard deviation shock in unemployment (solid line). 95% confidence bounds (dotted lines) are based on a resampling bootstrap.

unemployment could be interpreted solely as a deterioration of the country's future solvency which might lead to a rising bond yield as in particular long-term refinancing costs should rise. In contrast, the central mission of the Federal Reserve Board in the United States entails the "pursuit of maximum employment, stable prices, and moderate long-term interest rates".⁷ Therefore the Fed might react immediately to rising unemployment rates by lowering interest rates to stimulate public and private investment. The excess supply of money might then lead to lower bond yields. The same holds true for Australia where the Reserve Bank of Australia's core task is also

⁷www.federalreserve.gov/aboutthefed/mission.htm

Scatter Plot of Impulse-Response Peak and Timing




Figure 4.10: The figure illustrates the timing and magnitude of the impulse-response of government bond yields to an idiosyncratic shock in unemployment changes.

to “contribute to the stability of the currency, full employment, and the economic prosperity”.⁸

The impact of idiosyncratic shocks in bond yields on unemployment changes is negligible. Figure 4.11 presents the corresponding impulse-response functions. As can be seen, the impact graphs lie close to zero for all nine countries. This result is in line with economic reasoning. Interest rates might have an impact on both public and private investment, which in turn determines the number of available jobs in the long run. A direct impact has, to the best of our knowledge, never been derived.

⁸www.rba.gov.au/about-rba/index.html

Impulse-Response Functions of Unemployment to a Shock in Bond Yields




Figure 4.11: The figure presents the impulse-response functions of unemployment to a one standard deviation shock in bond yields (solid line). 95% confidence bounds (dotted lines) are based on a resampling bootstrap.

4.5 Concluding Remarks

With this work we set out to investigate the relationship between government bond yields and unemployment. Government-issued bonds are priced on the assumption that the country as a whole will be able to pay the mortgage back in the future. This ability is entirely determined by the labor force and unemployment is therefore a great risk factor in this bet. Besides the short-run negative effects of unemployment, we show that at least in Europe, rising unemployment leads to an increase in government bond yields, indicating that financing costs for the country rise. In a worst-case scenario one might imagine that the country fails to refinance mature bonds because investors deem the risk as abundant, and require extremely high risk compensation (as happened in Greece or Spain) or refrain from investing at all, leading to a state bankruptcy. The current European sovereign debt crisis is such a situation. And even if it seems obvious, the reported result has an important policy implication: unemployment reduction has to be a major goal of government activity. Reducing unemployment not only raises the tax basis and, thus, alleviates the debt burden, but in the long run also leads to lower interest rates as the risk compensation requested by investors is reduced. A notable exception are the United States for which the trust of investors in the country's ability to meet its debt obligations is seemingly endless.

The investigation of the relationship between unemployment and bond yields requires a suitable preparation of the time series because bond and unemployment data are only available on different frequencies. Our methodological contribution consists of a proposition how to provide weekly data for unemployment using a mixed-frequency heterogeneous autoregressive model. We show that to this end Google search queries can successfully be used to nowcast unemployment changes. The application of the M-HAR model is of course not limited to unemployment nowcasting but can generally be applied to any context where a nowcast that employs a highly persistent time series is needed.

Appendix B

HVAR – Estimation Results

Table 4.4: The table reports coefficient estimates for the HVAR model. DV marks the dependent variable bond yields b or unemployment changes u . p -values are in parentheses.

Country	DV	c_{uw}	c_{um}	c_{u3m}	c_{u6m}	c_{uy}	c_{u2y}	c_{u3y}	c_{bw}	c_{bm}	c_{b3m}	c_{b6m}	c_{by}	c_{b2y}	c_{b3y}
Austria	b	0.04 (0.83)	-0.25 (0.28)	0.01 (0.98)	0.79 (0.12)	1.20 (0.14)	0.85 (0.49)	0.15 (0.87)	1.14 (0.00)	-0.18 (0.05)	-0.02 (0.77)	-0.03 (0.60)	0.03 (0.54)	-0.12 (0.03)	0.16 (0.04)
	u	0.91 (0.00)	-0.27 (0.00)	0.01 (0.94)	0.34 (0.06)	-0.94 (0.00)	-0.87 (0.05)	0.42 (0.21)	0.02 (0.29)	-0.01 (0.77)	-0.01 (0.78)	0.01 (0.51)	0.01 (0.63)	0.05 (0.02)	-0.09 (0.00)
France	b	-0.86 (0.21)	0.91 (0.31)	-0.91 (0.22)	0.00 (1.00)	4.17 (0.02)	2.82 (0.02)	2.37 (0.15)	1.06 (0.00)	-0.16 (0.06)	0.03 (0.72)	-0.17 (0.01)	-0.08 (0.34)	0.20 (0.08)	-0.03 (0.69)
	u	1.02 (0.00)	-0.10 (0.22)	-0.17 (0.02)	-0.01 (0.90)	0.46 (0.01)	0.06 (0.60)	-0.10 (0.52)	0.00 (0.95)	-0.01 (0.27)	0.01 (0.14)	-0.01 (0.04)	-0.02 (0.03)	0.03 (0.00)	-0.01 (0.06)
Germany	b	-0.05 (0.91)	0.70 (0.28)	-0.67 (0.25)	0.19 (0.76)	1.68 (0.08)	1.25 (0.14)	1.20 (0.53)	1.10 (0.00)	-0.15 (0.09)	-0.01 (0.90)	-0.06 (0.38)	-0.04 (0.57)	0.10 (0.16)	-0.08 (0.05)
	u	0.89 (0.00)	0.03 (0.71)	-0.07 (0.37)	-0.03 (0.75)	-0.49 (0.00)	0.18 (0.12)	0.40 (0.13)	0.01 (0.20)	0.00 (0.82)	0.00 (0.66)	0.02 (0.03)	-0.02 (0.07)	-0.04 (0.00)	0.02 (0.00)
Portugal	b	0.03 (0.94)	0.11 (0.88)	-0.11 (0.87)	-0.38 (0.59)	-0.25 (0.76)	-2.07 (0.26)	4.42 (0.13)	1.00 (0.00)	-0.20 (0.03)	0.28 (0.02)	-0.07 (0.59)	-0.08 (0.42)	0.21 (0.18)	-0.21 (0.10)
	u	0.92 (0.00)	-0.08 (0.37)	0.05 (0.58)	-0.01 (0.93)	0.13 (0.24)	-0.37 (0.13)	-1.02 (0.01)	0.01 (0.08)	-0.01 (0.35)	0.01 (0.74)	0.00 (0.78)	0.01 (0.56)	0.03 (0.20)	-0.03 (0.06)

Table 4.4: HVAR – Estimation Results (cont.)

Country	DV	c_{uw}	c_{um}	c_{u3m}	c_{u6m}	c_{uy}	c_{u2y}	c_{u3y}	c_{bw}	c_{bm}	c_{b3m}	c_{b6m}	c_{by}	c_{b2y}	c_{b3y}
Spain	<i>b</i>	0.97 (0.02)	-0.74 (0.20)	-0.12 (0.81)	-0.27 (0.57)	1.31 (0.01)	0.44 (0.37)	0.21 (0.81)	0.74 (0.00)	0.02 (0.85)	0.23 (0.07)	-0.26 (0.03)	-0.03 (0.81)	0.50 (0.03)	-0.23 (0.21)
	<i>u</i>	0.88 (0.00)	0.15 (0.13)	-0.17 (0.05)	0.14 (0.09)	-0.03 (0.66)	-0.24 (0.00)	0.04 (0.78)	0.03 (0.03)	-0.03 (0.12)	0.01 (0.62)	-0.01 (0.51)	0.04 (0.04)	-0.04 (0.25)	0.00 (0.91)
Switzerland	<i>b</i>	0.76 (0.15)	-1.13 (0.11)	0.69 (0.40)	1.09 (0.29)	0.61 (0.60)	2.89 (0.00)	-1.44 (0.11)	1.11 (0.00)	-0.18 (0.03)	-0.01 (0.86)	-0.10 (0.14)	0.04 (0.55)	0.17 (0.15)	0.02 (0.78)
	<i>u</i>	1.03 (0.00)	-0.26 (0.00)	0.22 (0.02)	0.04 (0.77)	-0.32 (0.02)	-0.13 (0.26)	0.11 (0.30)	0.01 (0.36)	0.00 (0.90)	-0.01 (0.42)	0.01 (0.23)	0.00 (0.52)	-0.03 (0.03)	0.01 (0.60)
UK	<i>b</i>	-0.44 (0.40)	0.25 (0.72)	-0.10 (0.88)	1.65 (0.04)	-0.18 (0.84)	0.20 (0.81)	-3.59 (0.10)	1.01 (0.00)	0.01 (0.88)	-0.11 (0.15)	0.05 (0.52)	0.03 (0.67)	0.13 (0.26)	-0.12 (0.20)
	<i>u</i>	1.18 (0.00)	-0.36 (0.00)	0.04 (0.54)	0.00 (1.00)	-0.12 (0.22)	0.07 (0.42)	0.47 (0.05)	0.00 (0.91)	0.00 (0.90)	0.00 (0.96)	0.01 (0.16)	-0.02 (0.00)	-0.01 (0.60)	0.01 (0.34)
Australia	<i>b</i>	-0.11 (0.69)	-0.22 (0.58)	0.94 (0.17)	-0.50 (0.63)	-2.60 (0.15)	3.71 (0.08)	-0.58 (0.50)	1.06 (0.00)	-0.12 (0.17)	0.07 (0.41)	0.04 (0.64)	-0.11 (0.14)	-0.09 (0.38)	0.13 (0.27)
	<i>u</i>	0.95 (0.00)	-0.45 (0.00)	0.22 (0.12)	-0.07 (0.76)	-0.41 (0.27)	0.23 (0.61)	-0.10 (0.56)	0.01 (0.60)	-0.01 (0.69)	0.01 (0.49)	0.00 (0.85)	-0.01 (0.50)	-0.03 (0.11)	0.01 (0.72)
US	<i>b</i>	0.13 (0.77)	-1.57 (0.02)	1.25 (0.23)	-0.08 (0.95)	0.23 (0.80)	0.40 (0.49)	1.08 (0.18)	0.97 (0.00)	-0.02 (0.85)	-0.02 (0.80)	0.00 (0.97)	0.05 (0.55)	-0.38 (0.14)	0.34 (0.23)
	<i>u</i>	1.03 (0.00)	-0.60 (0.00)	0.21 (0.11)	0.53 (0.00)	-0.40 (0.00)	-0.08 (0.29)	-0.02 (0.82)	-0.01 (0.16)	0.02 (0.12)	-0.01 (0.51)	0.00 (0.71)	0.02 (0.09)	0.01 (0.87)	-0.05 (0.19)

Chapter 5

Conclusion

How do income risks and unemployment affect financial markets? The study answers this question by looking at two sectors of financial markets. To that end, I analyze two models that link personal income risk to asset returns and one model that describes the impact of changes in unemployment on government bond yields.

Investigating how asset returns are affected, I examine how a purely rational asset pricing model changes with the introduction of an income risk factor and I describe how a behavioral model can be estimated using income risk as a basis for reference level discrimination. The evidence presented in chapter 2 suggests that income risk plays an important role in determining the risk premium different classes of portfolios have to pay in order to compensate for their exposure to the income risk factor, i.e. the co-movement of the returns with the severity of the income risk. The main contribution of chapter 2 is the development of a risk factor that captures large idiosyncratic income risk. The risk factor builds on Constantinides and Duffie (1996) who formulate a rational model for the effect of idiosyncratic income risk on asset returns. In the construction of the risk factor, I account for the Krebs (2004) critique by not relying on central moments of the income distribution for testable restrictions. Instead, the factor measures by how much a large quantile is exceeded on average. Finally, the income risk factor successfully prices a cross-section of portfolio returns.

Income risk also affects the investors' attitude towards risk. In the estimation of the behavioral model in chapter 3, I find that investors above a certain reference level feature plausible risk aversion, while investors below a certain reference level act risk-seeking. The study contributes to the existing literature on behavioral finance by introducing methods from the literature on asset pricing with heterogeneous agents. In that, it provides a test for the presence of behavioral effects using non-experimental data. The strategy relies on sorting individuals by their recent income development. If investors belong to an upper q quantile of individuals that has recently experienced considerable positive income changes, the model permits that they have a different risk aversion than investors that belong to the lower q quantile.

Chapter 4 examines how the market for government bonds is affected by changes in unemployment. We present evidence that bond yields respond to shocks in unemployment in a country-specific manner. Concretely, bond yields respond positively to rises in unemployment for European countries where the central bank may only engage in monetary policies and negatively for the US and Australia, where the central banks' core tasks also include economic goals. Varying response speeds hint at differences in the effectiveness with which economies can absorb these changes. Chapter 4 outlines a methodological contribution by describing how the frequency cascade of heterogeneous autoregressive models can be used in a nowcasting context to account for long lag structures in a parsimonious fashion. In the application of our method we provide a detailed discussion of the use of Google search query data in nowcasting.

All chapters conclude that income risks affect financial markets. The influence of income risk on asset returns is exerted through a change in investor behavior either as a result of additional risk exposure or through psychological biases. The impact of unemployment changes on government bond yields is country-specific and depends on the role of the central bank. Of course this list is incomplete and there is likely more to be discovered about how income risks affect financial markets. My exploration of three of these channels not only shows their importance but also demonstrates their diversity. This certainly makes a completion of the list an interesting agenda for future research.

Bibliography

- AFONSO, A., P. GOMES, AND P. ROTHER (2011): “Short- and Long-run Determinants of Sovereign Debt Credit Ratings,” *International Journal of Finance & Economics*, 16, 1–15.
- ALTONJI, J. G. (1986): “Intertemporal Substitution in Labor Supply: Evidence from Micro Data,” *Journal of Political Economy*, 94, 176–215.
- ALTONJI, J. G. AND A. SIOW (1987): “Testing the Response of Consumption to Income Changes with (Noisy) Panel Data,” *Quarterly Journal of Economics*, 102, 293–328.
- ANDRIES, M. (2012): “Consumption-based Asset Pricing with Loss Aversion,” Working Paper.
- ASKITAS, N. AND K. F. ZIMMERMANN (2009): “Google Econometrics and Unemployment Forecasting,” *Applied Economics Quarterly*, 55, 107–120.
- BALDUZZI, P., E. J. ELTON, AND T. C. GREEN (2001): “Economic News and Bond Prices: Evidence from the U.S. Treasury Market,” *Journal of Financial and Quantitative Analysis*, 36, 523–543.
- BALDUZZI, P. AND T. YAO (2007): “Testing Heterogeneous-Agent Models: An Alternative Aggregation Approach,” *Journal of Monetary Economics*, 54, 470–412.
- BALI, T., N. CAKICI, AND R. WHITELAW (2011): “Maxing Out: Stocks as Lotteries and the Cross-section of Expected Returns,” *Journal of Financial Economics*, 99, 427–446.

- BALKEMA, A. AND L. DE HAAN (1974): "Residual Life Time at Great Age," *Annals of Probability*, 2, 792–804.
- BANK, M., M. LARCH, AND G. PETER (2011): "Google Search Volume and its Influence on Liquidity and Returns of German Stocks," *Financial Markets and Portfolio Management*, 25, 239–264.
- BARBERIS, N. (2013): "Thirty Years of Prospect Theory in Economics: A Review and Assessment," *Journal of Economic Perspectives*, 27, 173–196.
- BARBERIS, N. AND M. HUANG (2008): "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review*, 98, 2066–2100.
- BARBERIS, N., M. HUANG, AND T. SNATOS (2001): "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116, 1–53.
- BARREIRA, N., P. GODINHO, AND P. MELO (2013): "Nowcasting Unemployment Rate and New Car Sales in South-Western Europe with Google Trends," *NET-NOMICS: Economic Research and Electronic Networking*, 14, 129–165.
- BAUER, G. H. AND K. VORKINK (2011): "Forecasting Multivariate Realized Stock Market Volatility," *Journal of Econometrics*, 160, 93–101.
- BAYOUMI, T., M. GOLDSTEIN, AND G. WOGLOM (1995): "Do Credit Markets Discipline Sovereign Borrowers? Evidence from U.S. States," *Journal of Money, Credit and Banking*, 27, 1046–1059.
- BENARTZI, S. AND R. THALER (1995): "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics*, 110, 73–92.
- BHAMRA, H. S. AND R. UPPAL (2014): "Asset Prices with Heterogeneity in Preferences and Beliefs," *Review of Financial Studies*, 27, 519–580.
- BOOIJ, A. S., B. M. S. VAN PRAAG, AND G. VAN DE KUILEN (2010): "A Parametric Analysis of Prospect Theory's Functionals For the General Population," *Theory and Decision*, 68, 115–148.

- BOWMAN, D., D. MINEHART, AND M. RABIN (1999): "Loss Aversion in a Consumption-Savings Model," *Journal of Economic Behavior & Organization*, 38, 155–178.
- BOYER, B., T. MITON, AND K. VORKINK (2010): "Expected Idiosyncratic Skewness," *Review of Financial Studies*, 32, 169–202.
- BRAV, A., G. M. CONSTANTINIDES, AND C. C. GECZY (2002): "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy*, 110, 793–824.
- BREEDEN, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.
- CAMPBELL, J. Y. AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- CAMPBELL, J. Y. AND A. S. KYLE (1993): "Smart Money, Noise Trading and Stock Price Behaviour," *Review of Economic Studies*, 60, 1–34.
- CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLEY (1997): *The Econometrics of Financial Markets*, Princeton, New Jersey: Princeton University Press.
- CHIRIAC, R. AND V. VOEV (2011): "Modelling and Forecasting Multivariate Realized Volatility," *Journal of Applied Econometrics*, 26, 922–947.
- CHOI, H. AND H. VARIAN (2012): "Predicting the Present with Google Trends," *Economic Record*, 88, 2–9.
- COCHRANE, J. H. (2008): "Financial Markets and the Real Economy," in *Handbook of the Equity Risk Premium*, Elsevier.
- COGLEY, T. (2002): "Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey," *Journal of Monetary Economics*, 49, 309–334.

- CONRAD, J., R. DITTMAR, AND E. GHYSELS (2013): “Ex Ante Skewness and Expected Stock Returns,” *The Journal of Finance*, 68, 85–124.
- CONSTANTINIDES, G. M. AND D. DUFFIE (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104, 219–240.
- CONSTANTINIDES, G. M. AND W. E. FERSON (1991): “Habit Persistence and Durability in Aggregate Consumption: Empirical Tests,” *Journal of Financial Economics*, 29, 199–240.
- CORSI, F. (2009): “A Simple Approximate Long-Memory Model of Realized Volatility,” *Journal of Financial Econometrics*, 7, 174–196.
- DA, Z., J. ENGELBERG, AND P. GAO (2011): “In Search of Attention,” *The Journal of Finance*, 66, 1461–1499.
- D’AMURI, F. AND J. MARCUCCI (2010): “‘Google It!’ Forecasting the US Unemployment Rate with A Google Job Search Index,” FEEM Working Paper No. 31.
- DAVIDSON, A. C. AND R. L. SMITH (1990): “Models of Exceedances Over High Thresholds,” *Journal of the Royal Statistical Society*, 52, 393–442.
- DE GIORGI, E. AND T. HENS (2006): “Making Prospect Theory Fit for Finance,” *Financial Markets and Portfolio Management*, 20, 339–360.
- DIEBOLD, F. AND R. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13, 253–263.
- DIEBOLD, F. X. AND C. LI (2006): “Forecasting the Term Structure of Government Bond Yields,” *Journal of Econometrics*, 130, 337–364.
- DIMMOCK, S. AND K. ROY (2010): “Loss-aversion and Household Portfolio Choice,” *Journal of Empirical Finance*, 17, 441–459.
- DIMPFL, T. AND R. C. JUNG (2012): “Financial Market Spillovers Around the Globe,” *Applied Financial Economics*, 22, 45–57.

- DIMPFL, T. AND T. LANGEN (2015): “A Cross-Country Analysis of Unemployment and Bonds with Long-Memory Relations,” Working Paper.
- EMBRECHTS, P., S. I. RESNICK, AND G. SAMORODNITSKYSS (1999): “Extreme Value Theory as a Risk Management Tool,” *North American Actuarial Journal*, 3, 30–41.
- EU (2012): “Consolidated Versions of the Treaty on European Union and the Treaty on the Functioning of the European Union,” *Official Journal of the European Union*, 55, 1–390.
- FAMA, E. F. AND K. R. FRENCH (1992): “The Cross-Section of Expected Stock Returns,” *The Journal of Finance*, 47, 427–465.
- (1995): “Size and Book-to-Market Factors in Earnings and Returns.” *The Journal of Finance*, 50, 131–155.
- FIGLEWSKI, S. (1978): “Market ‘Efficiency’ in a Market With Heterogeneous Information,” *Journal of Political Economy*, 86, 581–597.
- FISHER, R. A. AND L. H. C. TIPPETT (1928): “Limiting Forms of the Frequency Distribution of the Largest and Smallest Member of a Sample,” *Mathematical Proceedings of the Cambridge Philosophical Society*, 24, 180–190.
- FLEMING, M. J. . AND E. M. . REMOLONA (1999): “What Moves Bond Prices?” *The Journal of Portfolio Management*, 25, 28–38.
- GEWEKE, J. AND S. PORTER-HUDAK (1983): “The Estimation and Application of Long Memory Time Series Models,” *Journal of Time Series Analysis*, 4, 221–238.
- GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2004): “The MIDAS Touch: Mixed Data Sampling Regression Models,” Anderson Graduate School of Management – Finance Working Paper Series.
- (2006): “Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies,” *Journal of Econometrics*, 1-2, 59–95.

- GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): “Nowcasting: The Real-Time Informational Content of Macroeconomic Data,” *Journal of Monetary Economics*, 55, 665–676.
- GINSBERG, J., M. H. MOHEBBI, R. S. PATEL, L. BRAMMER, M. S. SMOLINSKI, AND L. BRILLIANT (2009): “Detecting Influenza Epidemics Using Search Engine Query Data,” *Nature*, 457, 1012–1014.
- GRISHCHENKO, O. V. AND M. ROSSI (2012): “The Role of Heterogeneity in Asset Pricing: The Effect of a Clustering Approach,” *Journal of Business & Economic Statistics*, 30, 297–311.
- GROSSMAN, S. J. AND R. J. SHILLER (1981): “The Determinants of the Variability of Stock Market Prices,” *The American Economic Review*, 71, 222–227.
- HANSEN, L. P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50, 1029–1054.
- HANSEN, L. P. AND K. J. SINGLETON (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, 50, 1269–1286.
- HIRSHLEIFER, D. (2001): “Investor Psychology and Asset Pricing,” *The Journal of Finance*, 56, 1533–1597.
- JACOBS, K. AND K. Q. WANG (2004): “Idiosyncratic Consumption Risk and the Cross Section of Asset Returns,” *The Journal of Finance*, 59, 2211–2252.
- KACHELMEIER, S. AND M. SHEHATA (1992): “Examining Risk Preferences Under High Monetary Incentives: Experimental Evidence from the People’s Republic of China,” *American Economic Review*, 82, 1120–1141.
- KAHNEMAN, D. AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica*, 47, 263–291.
- (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 5, 297–323.

- KOSZEGI, B. AND M. RABIN (2006): “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121, 1133–1165.
- (2007): “Reference-Dependent Risk Attitudes,” *American Economic Review*, 97, 1047–1073.
- (2009): “Reference-Dependent Consumption Plans,” *American Economic Review*, 99, 909–936.
- KREBS, T. (2004): “Testable Implications of Consumption-Based Asset Pricing Models with Incomplete Markets,” *Journal of Mathematical Economics*, 40, 191–206.
- KUZIN, V., M. MARCELLINO, AND C. SCHUMACHER (2011): “MIDAS vs. Mixed-Frequency VAR: Nowcasting GDP in the Euro Area,” *International Journal of Forecasting*, 27, 529–542.
- LANGEN, T. (2013): “Tracing Tails - Large Idiosyncratic Income Shocks in a Heterogeneous Agent Asset Pricing Model,” Working Paper.
- (2014): “Empirical Asset Pricing with Reference-Dependent Heterogeneous Agents,” Working Paper.
- LEWELLEN, J., S. NAGEL, AND J. SHANKEN (2010): “A Skeptical Appraisal of Asset Pricing Tests,” *Journal of Financial Economics*, 96, 175–194.
- LINTNER, J. (1965): “The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets,” *Review of Economics and Statistics*, 47, 13–37.
- LUCAS, R. E. (1987): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 1429–1445.
- MARCELLINO, M. AND C. SCHUMACHER (2010): “Factor MIDAS for Nowcasting and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP,” *Oxford Bulletin of Economics and Statistics*, 72, 518–550.

- MCNEIL, A. J. AND R. FREY (2000): "Estimation of Tail-Related Risk Measures for Heteroscedastic Financial Time Series: an Extreme Value Approach." *Journal of Empirical Finance*, 7, 271–300.
- MEHRA, R. AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- MERTON, R. C. (1973): "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.
- MOSSIN, J. (1966): "Equilibrium in a Capital Asset Market," *Econometrica*, 34, 768–783.
- NAVICKE, J., O. RASTRIGINA, AND H. SUTHERLAND (2014): "Nowcasting Indicators of Poverty Risk in the European Union: A Microsimulation Approach," *Social Indicators Research*, 119, 101–119.
- ODEAN, T. (1998): "Volume, Volatility, Price and Profit When all Traders are Above Average," *The Journal of Finance*, 53, 1887–1934.
- PAGEL, M. (2012): "Expectations-based Reference-Dependence and Asset Pricing," Working Paper.
- PARKER, J. A. AND C. JULLIARD (2005): "Consumption Risk and the Cross Section of Expected Returns," *Journal of Political Economy*, 113, 185–222.
- PICKANDS, J. (1975): "Statistical Inference Using Extreme Order Statistics," *Annals of Statistics*, 3, 119–131.
- POST, T., G. B. VAN DEN ASSEM, AND R. THALER (2008): "Deal or No Deal? Decision-making Under Risk in a Large-Payoff Game Show," *American Economic Review*, 98, 38–71.
- SHANKEN, J. (1992): "On the Estimation of Beta-Pricing Models," *The Review of Financial Studies*, 5, 1–33.
- SHARPE, W. F. (1964): "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *The Journal of Finance*, 19, 425–442.

- SHEA, J. (1995): "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis," *American Economic Review*, 85, 186–200.
- SHEFRIN, H. AND M. STATMAN (1994): "Behavioral Capital Asset Pricing Theory," *Journal of Financial and Quantitative Analysis*, 29, 323–349.
- SIBBERTSEN, P., C. WEGENER, AND T. BASSE (2014): "Testing for a Break in the Persistence in Yield Spreads of EMU Government Bonds," *Journal of Banking & Finance*, 41, 109–118.
- TSENG, T.-C., H. CHUNG, AND C.-S. HUANG (2009): "Modeling Jump and Continuous Components in the Volatility of Oil Futures," *Studies in Nonlinear Dynamics & Econometrics*, 13, Article 5.
- VISSING-JØRGENSEN, A. (2002): "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110, 825–853.
- ZELDES, S. P. (1989): "Consumption and Liquidity Constraints," *Journal of Political Economy*, 97, 305–346.