Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

____________________________ ____________________

Mark Rachwalski Date

Essays on Asset Pricing

By

Mark Rachwalski
Doctor of Philosophy

Business

Jay Shanken
Advisor

Tarun Chordia
Committee Member

Ilia Dichev
Committee Member

T. Clifton Green
Committee Member

Narasimhan Jegadeesh
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Essays on Asset Pricing

By

Mark Rachwalski
B.S., Northwestern University, 1996
M.S., University of North Carolina at Charlotte, 2008

Advisor: Jay Shanken
B.S., State University of New York at Stony Brook, 1973
M.A., Cornell University, 1976
M.S., Carnegie Mellon University, 1980
Ph.D., Carnegie Mellon University, 1983

An abstract of
A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory Univeristy
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Business
2014

Abstract

Essays on Asset Pricing

By Mark Rachwalski

Abstract

This dissertation consists of three essays. The first essay, "Idiosyncratic Risk Innovations and the Idiosyncratic Risk-Return Relation" (with Quan Wen), examines the role of idiosyncratic risk in asset pricing. I find that stocks with increases in idiosyncratic risk tend to earn low subsequent returns for a few months. However, high idiosyncratic risk stocks eventually earn persistently high returns. These results are consistent with positively priced idiosyncratic risk and temporary underreaction to idiosyncratic risk innovations. Because risk levels and innovations are correlated, the relation between historical idiosyncratic risk and returns may reflect both risk premia and underreaction and yield misleading inference regarding the price of risk. The results reconcile previous work, which offers conflicting evidence on the price of idiosyncratic risk, and help to discriminate among explanations of the idiosyncratic risk-return relation. In the second essay, "Stock Wealth, Consumption, and Return Predictability", I construct a novel empirical model of expected stock returns. The stock wealth-consumption ratio reflects expected stock returns and consumption growth. Because consumption growth is mostly unpredictable, much of the variation of this ratio likely reflects changing expected stock returns. In contrast, isolating expected stock return information from other variables may be difficult (in addition to stock returns, the dividend yield may predict dividend growth, while the consumption-wealth ratio may predict non-stock wealth returns). Empirically, a detrended version of this ratio strongly predicts U.S. and international stock returns. In contrast to other predictive variables, predictability does not deteriorate after 1980 and out-of-sample performance is impressive. The third essay, "Bonds, Aggregate Wealth, and Stock Market Risk", examines the role of bond risk in determining stock prices. I find that bond returns predict consumption growth after controlling for equity returns, which suggests that bonds capture important information about aggregate wealth. Consistent with this, bond risk is priced in the cross section of stocks. Bond risk partially explains some well-known anomalies. The results suggest stock indices are an insufficient proxy for aggregate wealth and that bond risk is an important component of consumption risk.

Essays on Asset Pricing

By

Mark Rachwalski
B.S., Northwestern University, 1996
M.S., University of North Carolina at Charlotte, 2008

Advisor: Jay Shanken
B.S., State University of New York at Stony Brook, 1973
M.A., Cornell University, 1976
M.S., Carnegie Mellon University, 1980
Ph.D., Carnegie Mellon University, 1983

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory Univeristy
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Business
2014

Contents

List of Figures

List of Tables

List of Tables (continued)

Essay Three: Bonds, Aggregate Wealth, and Stock Market Risk

Idiosyncratic Risk Innovations and the Idiosyncratic Risk-Return Relation

Mark Rachwalski and Quan Wen

March 19, 2014

Abstract

Stocks with increases in idiosyncratic risk tend to earn low subsequent returns for a few months. However, high idiosyncratic risk stocks eventually earn persistently high returns. These results are consistent with positively priced idiosyncratic risk and temporary underreaction to idiosyncratic risk innovations. Because risk levels and innovations are correlated, the relation between historical idiosyncratic risk and returns may reflect both risk premia and underreaction and yield misleading inference regarding the price of risk. The results reconcile previous work, which offers conflicting evidence on the price of idiosyncratic risk, and help to discriminate among explanations of the idiosyncratic risk-return relation.

A large number of empirical studies have examined the cross-sectional price of idiosyncratic risk. Recent empirical work has focused on the negative relation between historical idiosyncratic risk and returns documented by Ang, Hodrick, Xing, and Zhang (2006, 2009), which suggests a negative price of idiosyncratic risk. This result is provocative because theory generally suggests that the price of idiosyncratic risk should be zero or positive (see Merton (1987)). Some researchers find support for the results of Ang, Hodrick, Xing, and Zhang (2006)[1], while others find the results to be fragile and possibly a manifestation of liquidity-related return patterns or skewness rather than reflective of an idiosyncratic risk-return relation[2]. Additionally, some studies have documented, under certain alternative specifications, a positive idiosyncratic risk-return relation (Lehmann (1990b), Spiegel and Wang (2005), Malkiel and Xu (2006), and Fu (2009)). Overall, the existing literature offers a confusing picture of idiosyncratic risk-related return patterns.

In this paper, we empirically examine the relation between idiosyncratic risk and returns. However, in contrast to prior work, we also consider the relation between idiosyncratic risk *innovations* and subsequent returns. This is important because risk levels and innovations are correlated, so disentangling the respective relations takes some care. We find strong and robust evidence supporting a negative relation between idiosyncratic risk innovations and returns, but far less compelling evidence for a negative relation between idiosyncratic risk levels and returns. In particular, we find that (1) some good proxies for idiosyncratic risk are positively related to returns, (2) in "horse race" regressions, changes in idiosyncratic risk are generally highly significant when explaining subsequent returns while idiosyncratic risk levels are often insignificant, and (3) the empirical relation between historical idiosyncratic risk levels and subsequent returns is strongest when using a relatively poor proxy for subsequent idiosyncratic risk (and often insignificant otherwise), while the relation between

[1]See Guo and Savickas (2010), Peterson and Smedema (2011), and George and Hwang (2011).

[2]See Bali and Cakici (2008), Huang, Liu, Rhee, and Zhang (2009), Fu (2009), Han and Lesmond (2011), Bali, Cakici, and Whitelaw (2011), and Boyer, Mitton, and Vorkink (2010).

innovations and returns is robust to the choice of proxy. This evidence helps us distinguish between various hypotheses about the processes driving these relations (e.g. a process that explains a negative relation between risk levels and returns may not easily explain a negative relation between risk innovations and returns).

We find that historical idiosyncratic risk can be positively or negatively related to subsequent returns, depending on the return measurement period. This can be easily seen in Figure 1, which plots the returns of idiosyncratic volatility-sorted hedge portfolios for up to ten years after portfolio formation. Equal-weighted hedge portfolio returns are negative for six months after portfolio formation, with the most negative return in month one. However, starting about six months after portfolio formation, returns are always positive (value-weighted returns turn positive later, but follow the same pattern). Clearly, month one returns, which are the focus of many empirical studies, are not sufficient to fully characterize the idiosyncratic risk-return relation. Indeed, a reasonable interpretation of Figure 1 is that the negative returns immediately after portfolio formation are likely attributable to a transitory friction, while the long-run (equilibrium) price of idiosyncratic risk is positive. Our empirical results are consistent with this interpretation. More generally, our results suggest that examining the timing and persistence of return patterns, as well as distinguishing between the effects of innovations and levels, can lead to a deeper understanding of what is really driving returns.

Our empirical results yield new insights into the return patterns associated with idiosyncratic risk, which appear to be richer than is commonly understood. In particular, our results suggest a positive price of idiosyncratic risk and a negative relation between idiosyncratic risk innovations and subsequent returns. In contrast, the recent literature generally focuses on the negative relation between historical idiosyncratic risk and returns documented by Ang, Hodrick, Xing, and Zhang (2006). Additionally, we find that idiosyncratic risk may explain a meaningful portion of average market portfolio excess returns (approximately 3% annually, although this estimate is not precise). This is an exciting result, as traditional beta does not appear to be priced in the cross section (see Fama and French (2004) and Lewellen and Nagel (2006)), which suggests that beta risk may not explain much of average market excess returns.

We find that the return patterns associated with idiosyncratic risk are more robust and economically important than previous studies suggest. Our results are robust to omitting low-priced or illiquid stocks from the sample, value-weighting returns in a way that focuses on large stocks, controlling for well-known liquidity-related return patterns, and extending the sample from 1966-2012 to 1929-2012. In contrast, some researchers find that the Ang, Hodrick, Xing, and Zhang (2006) anomaly may be driven by a small subset of small stocks (as little as 2% of aggregate market capitalization) and/or liquidity-related return patterns[3]. Our results also stand in contrast to many important and well-known return patterns (in particular, those associated with market-to-book, size, momentum, and liquidity), which are often meaningfully attenuated, and sometimes insignificant, when focusing on larger stocks.

We develop a simple model, featuring a positive price of idiosyncratic risk and temporary price underreaction to idiosyncratic risk innovations, to help understand our empirical results. We show that, under this model, the sign of the relation between historical idiosyncratic risk and subsequent returns is ambiguous, even if one assumes that the equilibrium price of idiosyncratic risk is positive. This occurs because idiosyncratic risk levels and innovations may be positively correlated, but have opposing relations with returns. A positive idiosyncratic risk premium implies that, ceteris paribus, stocks with high idiosyncratic risk will tend to earn high returns. In this case, temporary underreaction to idiosyncratic risk innovations implies that stocks with increases in idiosyncratic risk will tend to earn low subsequent returns for a while. This occurs because prices reflect risk news (which im-

[3]See Bali and Cakici (2008), Huang, Liu, Rhee, and Zhang (2009), Fu (2009), and Han and Lesmond (2011).

plies a higher discount rate and lower price) with a delay. Since idiosyncratic risk levels and innovations are positively correlated, stocks with high historical idiosyncratic risk will tend to earn high returns due to the idiosyncratic risk premium, but low returns due to underreaction. A priori, it is not clear which effect should dominate.

Price underreaction to idiosyncratic risk innovations could be caused by frictions and/or investor biases, and could be rational.[4] Prior studies find evidence of apparent underreaction in a wide variety of settings.[5] Also, there is evidence that investors underreact to volatility innovations when setting option prices (Poteshman (2001)). Then, it should not be surprising if investors underreact to idiosyncratic risk innovations.

In this paper, we do not directly test for investor underreaction (primarily because we do not observe investors' idiosyncratic risk estimates), so we do not determine whether price underreaction is caused by investor underreaction to risk innovations or some trading friction. Instead, we focus on fully documenting the return patterns associated with idiosyncratic risk and estimating the price of idiosyncratic risk in the presence of price underreaction. Of course, fully documenting these return patterns is a critical step in understanding the processes that drive the idiosyncratic risk premium and any related price underreaction. A proposed explanation of idiosyncratic risk-related return patterns should address, or at least be consistent with, all of the related return patterns. Our explanation (featuring positively priced idiosyncratic risk and price underreaction to risk innovations) appears to be fully consistent with the empirical results. In contrast, existing explanations generally focus on

[4]Such biases and frictions include biased investor beliefs (Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrahmanyam (1998)), slow information diffusion (Hong and Stein (1999), Hong, Torous, and Valaknov (2007)), information capacity constraints (Sims (2003)), non-trivial transactions or information gathering costs, and short-sale constraints.

[5]Prior studies suggest investors underreact to earnings announcements (Ball and Brown (1968), Bernard and Thomas (1989)), prior returns (Jegadeesh and Titman (1993)), dividend news (Michaely, Thaler, and Womack (1995)), share repurchases (Ikenberry, Lakonishok, and Vermaelen (1995)), seasoned equity offerings (Loughran and Ritter (1995), Spiess and Affleck-Graves (1995)), increased R&D expenditures (Eberhart, Maxwell, and Siddique (2004)), predictable demographic trends (DellaVigna and Pollet (2007)), industry returns (Hong, Torous, and Valaknov (2007)) and news about related firms (Cohen and Frazzini (2008), Menzly and Ozbas (2010)).

the Ang, Hodrick, Xing, and Zhang (2006) anomaly and are often inconsistent with some of our empirical results.

Our model implies that, in the presence of price underreaction, one should not infer that the equilibrium price of idiosyncratic risk is negative based on a negative empirical relation between historical idiosyncratic risk and returns. Instead, one should control for the effects of underreaction when estimating this price. Omitting such a control can result in misleading inference, including incorrectly estimating the sign of the price of risk. Therefore, the Ang, Hodrick, Xing, and Zhang (2006) anomaly, where stocks with high historical idiosyncratic risk tend to earn low returns, can be consistent with a positive price of idiosyncratic risk. The anomaly may largely reflect a negative relation between idiosyncratic risk innovations and subsequent returns. Consistent with this, we find little evidence of negatively priced idiosyncratic risk after controlling for the effects of idiosyncratic risk innovations.

Our paper contains three main contributions to the idiosyncratic volatility literature. First, we empirically document a long-run positive relation between idiosyncratic risk and returns. Documenting this relation is important because theory generally suggests that the equilibrium price of idiosyncratic risk should be non-negative. Fu (2009) estimates a positive relation, but this result has been challenged by Guo, Kassa, and Ferguson (2012) and Fink, Fink, and He (2012), who note that the estimator used by Fu is biased, and that an unbiased estimator yields no significant relation. Second, we document a short-lived negative relation between idiosyncratic risk *innovations* and subsequent returns that explains much of the anomaly documented by Ang, Hodrick, Xing, and Zhang (2006).

Our third contribution is to develop an underreaction framework that can reconcile these empirical results, as well as the non-negative price of idiosyncratic risk suggested by theory. The framework suggests why different idiosyncratic risk proxies can offer conflicting evidence on the sign of the idiosyncratic risk-return relation. Viewed from our framework, it is not surprising that empirical studies that focus on recent data (e.g. Ang, Hodrick, Xing, and Zhang (2006) use a trailing one-month window to calculate historical idiosyncratic volatility) often find a negative idiosyncratic risk-return relation, while studies that focus more on distant data (e.g. Lehmann (1990b) uses five years of monthly data) are more likely to find a positive, or insignificant, relation. This occurs because recent data may not be fully assimilated into prices (so that underreaction is relatively important), while distant data is more likely fully priced. Also, it is not surprising that the Ang, Hodrick, Xing, and Zhang (2006) anomaly is strong when using alphas but weak, and often insignificant, when using raw returns. Under our framework, a high less low idiosyncratic volatility-sorted hedge portfolio should earn high returns due to a positive idiosyncratic risk premium, but low returns due to underreaction to idiosyncratic risk innovations. However, we find that expected market returns are positively correlated with idiosyncratic risk premia. Then, controlling for the market factor should partially control for idiosyncratic risk premia, and the hedge portfolio alpha should be lower than the raw return. More generally, our framework, methodology, and empirical results highlight the potential importance of price underreaction in standard empirical procedures, and may be useful in a variety of settings.

The paper proceeds as follows. Section I describes our data and methodology and Section II presents our empirical results. Section III presents a simple model featuring price underreaction and revisits the results within the context of this model. Section IV concludes.

1 Data

This section describes the methods and data used in our empirical examination of the idiosyncratic risk-return relation. One of the primary results of our paper is that different proxies for expected idiosyncratic risk suggest different idiosyncratic risk-return relations.

In particular, we find that proxies constructed from recent data tend to suggest a negative idiosyncratic risk-return relation, while proxies constructed from more distant data tend to suggest a positive relation. In this section, we show that both of these proxies are informative about subsequent idiosyncratic risk and can be reasonably used to explore the idiosyncratic risk-return relation.

1.1 Stock Sample, Filters

We obtain data from CRSP and Compustat. Stocks with a lagged price less than one dollar are removed from the sample.[6] In the absence of such a filter, Ang, Hodrick, Xing, and Zhang (2006, 2009) find a negative cross-sectional relation between stock returns and idiosyncratic risk. However, Bali and Cakici (2008) show that this result is driven by small, illiquid stocks. This makes the Ang, Hodrick, Xing, and Zhang anomaly less economically interesting because the result may be driven by a small subset of the stock market, and may be difficult to exploit due to short selling constraints and low liquidity associated with this subset. Also, Bali and Cakici's finding raises concerns that the anomaly is driven by return reversals or other return patterns that are likely important for small, illiquid stocks. This appears to be a valid concern, as Fu (2009) and Huang, Liu, Rhee, and Zhang (2009) show that the anomaly seems to be related to short-term reversals of the type documented by Jegadeesh (1990) and Lehmann (1990a) (see also Han and Lesmond (2011)). Because our paper's primary objective is to understand economically important return patterns associated with idiosyncratic risk, rather than to explain the anomaly of Ang, Hodrick, Xing, and Zhang (2006), we impose a filter to remove small and illiquid stocks from the sample.

[6]Alternative filters, such as eliminating the most illiquid stocks, yield similar results.

1.2 "Recent" and "Distant" Idiosyncratic Volatility

Following Ang, Hodrick, Xing, and Zhang (2006, 2009) and others, idiosyncratic volatility is calculated as the standard deviation of the residuals from a time-series regression of individual stock returns on the contemporaneous factors of Fama and French (1996):

$$r_{i,t} = \alpha + \beta_{i,MKT} MKT_t + \beta_{i,HML} HML_t + \beta_{i,SMB} SMB_t + \epsilon_{i,t}. \tag{1}$$

We show later in the paper that our results are robust to using the market model (i.e. omitting the HML and SMB factors). Daily data is used in the factor regressions.

We distinguish between "recent" historical idiosyncratic volatility (IVR, calculated using data from day $-t$ to day -7) and "distant" historical idiosyncratic volatility (IVD, calculated using data from day $-t-365$ to $-t$). t is the threshold that partitions the historical data. Although we often focus on a six-month (183 day) IVR-IVD threshold, thresholds of one and twelve months (38[7] days and 365 days respectively) are also examined.[8] When calculating IVR, we exclude data from the most recent seven calendar days to alleviate the effects of short-term reversals. To ensure that IVR and IVD generate precise estimates of subsequent realized idiosyncratic volatility, we require $t/2-5$ daily observations to calculate IVR and 180 daily observations to calculate IVD.

1.3 Proxying for Expected Idiosyncratic Volatility

Idiosyncratic volatility calculated from a relatively long historical time series (six months or greater) of daily data is a good predictor, in the cross section, of subsequent realized idiosyncratic volatility. In particular, this measure of historical idiosyncratic volatility is

[7] The one-month threshold is set to 38 days because a week is skipped before measuring returns. Given this skipped week, a 31-day threshold corresponds to three weeks of data.

[8] Our choice of IVR-IVD thresholds is guided by previous research related to price underreaction, which we explore as an explanation of our results later in the paper. Bernard and Thomas and Jegadeesh and Titman show that (respectively) earnings announcement underreaction and momentum persist for about a year.

a better predictor than historical idiosyncratic volatility calculated using a shorter (e.g. one month) time series of daily data, historical idiosyncratic volatility calculated using monthly data, and predicted idiosyncratic volatility from a monthly EGARCH model (as in Spiegel and Wang (2005) and Fu (2009)). This can be seen in Table 1, where we report the time-series average R^2 of cross-sectional regressions of realized idiosyncratic volatility (calculated using daily data) on IVR, IVD, and other measures of historical and predicted idiosyncratic volatility. Individually, IVR and IVD offer the greatest explanatory power.[9] Adding other measures to a model that includes IVR and IVD yields essentially no increase in explanatory power. Also, one-month IVR is, individually, a suboptimal predictor of subsequent idiosyncratic volatility. Use of a longer historical sample (e.g. six-month IVR) or including IVD yields greater explanatory power.

We conclude that historical idiosyncratic volatility calculated from a long time series of daily data is an appealing instrument, in the cross section, for expected idiosyncratic volatility. This is important because, as stressed by Fu (2009), theory relates expected risk and expected returns, not historical risk and expected returns.[10] For this reason, we use IVR and IVD (with various thresholds) to investigate the idiosyncratic risk-return relation.

1.4 The Cross-Sectional Distribution of Idiosyncratic Volatility

Table 2 reports descriptive statistics of single- and sequentially-sorted stock portfolios[11]. We refer to IVR- and IVD-sorted portfolios as IVR1-IVR5 and IVD1-IVD5, respectively. We add an “S” to indicate a sequentially-sorted portfolio (e.g. the IVRS5 portfolio is formed by first sorting on IVD, then IVR). The columns labeled IVR and IVD report the idiosyn-

[9] These results are similar to those of Guo, Kassa, and Ferguson (2012), who find that one-month historical idiosyncratic volatility is a better predictor of subsequent idiosyncratic volatility than predicted volatility from a monthly EGARCH model.

[10] As noted by Fu (2009), historical idiosyncratic volatility need not be a good point estimate of expected idiosyncratic volatility. However, historical idiosyncratic volatility may still be useful when forming portfolios with dispersion in expected idiosyncratic volatility, or when generating a point estimate of expected idiosyncratic volatility (e.g. through an AR(1) model).

[11] These results use a six-month IVR-IVD threshold.

cratic volatility associated with each portfolio (e.g. the IVR5 portfolio contains stocks with an average daily idiosyncratic volatility of 5.44% over the recent historical period (IVR) and 4.93% over the distant historical period (IVD)). Table 2 demonstrates that a single sort on IVD is very similar to a sort on IVR; both IVD and IVR increase in a similar way from IVR1 to IVR5 and from IVD1 to IVD5. This occurs because idiosyncratic volatility, not surprisingly, exhibits positive autocorrelation.[12] Then, a wide variety of historical idiosyncratic volatility measures capture similar information about expected idiosyncratic volatility. Lagging idiosyncratic volatility, or calculating idiosyncratic volatility over different intervals, does not matter much when using historical idiosyncratic volatility as a proxy for expected idiosyncratic volatility. However, sequential sorts break the tight link between IVD and IVR. The IVDS portfolios exhibit substantial variation in IVD and little variation in IVR. Similarly, the IVRS portfolios exhibit substantial variation in IVR but little variation in IVD.

Table 2 also reports the mean change in idiosyncratic volatility (IVC, defined as IVR-IVD) for the portfolios. Both the IVRS- and IVDS-sorted portfolios exhibit substantial variation and are monotonic in IVC. This occurs because varying IVR while holding IVD constant is equivalent to varying IVC while holding IVD constant. Then, the IVRS and IVDS portfolios can be used to examine the relation between idiosyncratic risk innovations and subsequent returns.

1.5 Idiosyncratic Volatility and Firm Size

Table 2 contains additional portfolio descriptive statistics. On average, high idiosyncratic volatility stocks are small and illiquid. A difficulty encountered when interpreting the returns of hedge portfolios formed by single sorts on idiosyncratic volatility (e.g. IVR5

[12] Consistent with this, the cross-sectional IVR-IVD correlation is quite high. The time series average of the IVR-IVD cross-sectional Pearson correlation is 0.78. The time-series average of the Spearman correlation is 0.87.

less IVR1) is that such sorts turn out to be similar to a sort on size or liquidity (i.e. high idiosyncratic volatility portfolios tend to contain many small, illiquid stocks and the reverse). Then, any difference in mean returns across the portfolios could be driven by a subset of small, illiquid firms. This could be true even if portfolio returns are value weighted; because the average firm in the IVR5 and IVD5 portfolio is small, even small stocks could receive a large portfolio weight. For this reason return patterns associated with small stocks (e.g. bid-ask bounce, reversals, short-selling constraints) are a plausible explanation for anomalous returns associated with the IVR5 or IVD5 portfolios, or any portfolios formed from these portfolios (e.g. the IVR hedge portfolio). Also, nonzero mean returns of IVD or IVR hedge portfolios may have little economic importance because the returns may be driven by a small subset of small stocks.

Such concerns are alleviated when examining the returns of the IVRS and IVDS hedge portfolios because, by construction, the IVRS5 hedge portfolio must contain stocks from the IVD1 portfolio (which consists of many large stocks). Similarly, the IVDS5 portfolio must contain stocks from the IVR1 portfolio. Also, the IVRS portfolios can be interpreted as a sorting procedure that induces variation in (IVR-IVD) while controlling for IVD; this should alleviate microstructure concerns because idiosyncratic risk innovations are less obviously related to liquidity and size than idiosyncratic risk levels. Consistent with this, the IVRS and IVDS portfolios exhibit much more size balance than the IVR and IVD portfolios. For example, the IVR5 portfolio consists of, on average, 2% of total market capitalization. The IVRS5 portfolio consists of, on average, 12% of total market capitalization.

2 The Cross-Sectional Price of Idiosyncratic Volatility

In this section, we explore the idiosyncratic risk-return relation using different proxies for idiosyncratic risk. The previous section demonstrates that both IVR and IVD are good proxies for expected idiosyncratic risk. Then, given a negative idiosyncratic risk-return relation, we should observe a negative relation between both IVR and IVD and subsequent returns. We find that this is not the case. Our results suggest that some proxies for idiosyncratic risk are negatively related, and some positively related, to subsequent returns. Then, one should not conclude that the idiosyncratic risk-return relation is negative based on evidence from a single proxy.

2.1 Empirical Strategy

We examine the mean returns of idiosyncratic volatility-sorted hedge portfolios. We also consider Fama and MacBeth (1973) cross-sectional regressions of the following form:

$$r_{i,t+1} = \alpha + \beta_{t,IVD} IVD_{i,t} + \beta_{t,IVR} IVR_{i,t} + \gamma_t X_t + \epsilon_{i,t+1}, \tag{2}$$

where X is a vector of controls. This specification is used to test whether different proxies for idiosyncratic risk (IVD and IVR) have opposing relations with subsequent returns and should be informative about whether idiosyncratic risk levels or innovations are driving subsequent returns. For example, suppose there is a negative relation between idiosyncratic risk levels and subsequent returns. Then, one should expect $\beta_{t,IVD} \leq 0$ and $\beta_{t,IVR} \leq 0$. Importantly, there is no reason to expect either $\beta_{t,IVD} > 0$ or $\beta_{t,IVR} > 0$.[13] In contrast, a negative relation between idiosyncratic risk innovations and subsequent returns suggests $\beta_{t,IVR} < 0$ and $\beta_{t,IVD} > 0$, provided that $IVR - IVD$ is a good proxy for such innovations.[14]

IVR and $IVR - IVD$ are positively correlated (see Table 2). Then, the relation documented by Ang, Hodrick, Xing, and Zhang (2006) could mostly reflect a negative relation

[13] Both IVR and IVD are positively related to subsequent idiosyncratic volatility, even after controlling for the other proxy, so there is no reason to expect a partial negative relation.

[14] We assume that $IVR - IVD$ is correlated with the true idiosyncratic risk innovation, but do not require these quantities to be equal. In untabulated results, we find a high correlation between $IVR - IVD$ and the residual from a regression of IVR on IVD and other explanatory variables (e.g. past returns and additional lags of IVD).

between idiosyncratic risk innovations and subsequent returns, rather than a relation between idiosyncratic risk levels and subsequent returns (we discuss the potential causes of an idiosyncratic risk innovation-return relation later in the paper). Determining the relative importance of these relations is a critical step in understanding idiosyncratic risk-return patterns because the underlying causes of these relations are likely substantially different. Misinterpreting a negative relation between idiosyncratic risk innovations and returns as a relation between idiosyncratic risk levels and returns may lead a researcher to believe that certain mechanisms drive stock returns, when in fact they do not (as discussed later in the paper, the problem is quite severe, as this misinterpretation may lead a researcher to infer that the sign of the idiosyncratic risk-return relation is negative when it is really positive).

To examine the idiosyncratic risk-return relation, expunged of the idiosyncratic risk innovation-return relation, we also examine deferred returns. Because risk innovations and levels are correlated, and the way in which these quantities influence returns is unknown, it is difficult to fully control for either. Then, disentangling the effects of risk levels and innovations may be difficult. However, provided that the effects of innovations are temporary, one can estimate the underlying price of idiosyncratic risk by examining idiosyncratic volatility-sorted hedge portfolio returns for many months after portfolio formation. Although idiosyncratic risk innovations may influence the returns for some time after portfolio formation, eventually the effects of these innovations should dissipate, and the returns should reflect compensation for risk. If idiosyncratic risk is positively priced, we should eventually see a positive relation between idiosyncratic risk and returns, although this relation may not be apparent for some time after portfolio formation.

2.2 Risk Adjustment

A somewhat subtle issue that arises in this application is the appropriate use of a risk-adjustment model, which researchers often use to demonstrate that a certain portfolio exhibits anomalous returns. It is important to remember that returns are only anomalous relative to the risk-adjustment model. Importantly, a risk-adjustment model may not be useful when determining whether a characteristic is driving returns (which is the primary question of this paper).

For example, suppose for simplicity that the cross-section of stock returns is determined *solely* by idiosyncratic risk (although other variables could be included): $E[R_{i,t}] = \gamma_t IV_i$, where IV_i is stock i's idiosyncratic risk and γ_t is a risk aversion parameter which allows for time-variation in expected market returns (more generally, γ_t can be any time-varying state variable that maps idiosyncratic risk to returns). Then, an increase in risk aversion will be associated with negative returns for the market portfolio and the high less low idiosyncratic risk hedge portfolio (high IV stocks will have more negative returns than low IV stocks contemporaneous with the increase in risk aversion). Therefore, market portfolio returns will be contemporaneously positively correlated with idiosyncratic risk hedge portfolio returns, and the market factor should "explain" hedge portfolio returns. However, in this example, all cross sectional variation in returns is attributable to idiosyncratic risk. More generally, if the market premium is correlated with the idiosyncratic risk hedge portfolio premium, then the market factor will "explain" hedge portfolio returns, although this does not imply that idiosyncratic risk is not priced in the cross section.

This issue is particularly relevant when examining idiosyncratic risk, as researchers often examine the alphas of idiosyncratic risk hedge portfolios (the Ang, Hodrick, Xing, and Zhang (2006) anomaly is stronger under a risk adjustment model, but often weak and insignificant when using raw returns). Hedge portfolio alphas will reflect the price of idiosyncratic risk that is orthogonal to the risk-adjustment model. However, if idiosyncratic risk drives the risk adjustment model, then alphas may not be informative about the price of idiosyncratic risk.

2.3 Sorted Stock Portfolio Returns

Panel 1 of Table 3 reports mean returns of the hedge (high minus low idiosyncratic volatility) portfolio for the IVD, IVR, IVDS, and IVRS sorting procedures. We report raw portfolio returns and a four-factor alpha (using the three factors of Fama and French (1996) and a momentum portfolio formed using one-month prior return[15]). Equal- and value-weighted returns are reported for IVR-IVD thresholds of one, six, and twelve months.

Table 3 confirms the findings of Ang, Hodrick, Xing, and Zhang (2006). The one-month IVR hedge portfolio has a negative mean return, especially under value weighting or when using the risk-adjustment model. However, this relation is often not significant when examining raw or equal-weighted returns.

Mean IVD hedge portfolio returns are generally insignificant. This is an important non-result because IVD is a good proxy for idiosyncratic risk (see Table 1). This suggests that the negative IVR-return relation is not reflective of an idiosyncratic risk-return relation, but is instead reflective of something else (i.e. some variable correlated with IVR, but less correlated with IVD).

IVRS hedge portfolio mean returns are always negative and highly statistically significant. For example, using a six-month IVR-IVD threshold, the equal-weighted IVRS hedge portfolio has a mean return of -51.8 basis points per month, with a heteroskedasticity-robust t-statistic of -4.25. For each threshold and weighting scheme, the statistical evidence for a negative IVRS-return relation is stronger than the evidence for a negative IVR-return relation. Therefore, controlling for distant idiosyncratic volatility reveals a stronger relation between recent idiosyncratic volatility and returns. Then, it should not be surprising that the negative IVRS-return relation documented here is robust to controls and filters that authors have used to challenge the results of Ang, Hodrick, Xing, and Zhang (2006) (such

[15]We focus on one-month prior returns because Huang, Liu, Rhee, and Zhang (2009) show that return reversals may help to explain the returns of idiosyncratic volatility portfolios. We control for other prior returns later in the paper.

controls will be examined later in the paper). Additionally, the extreme IVRS portfolios have far more market capitalization balance than the extreme IVR hedge portfolios (the market capitalization shares of the extreme IVRS portfolios are 0.30 and 0.12, while the shares of the extreme IVR portfolios are 0.56 and 0.02, see Table 2). This suggests that the IVRS-return relation may be of greater economic importance than the IVR-return relation.

The IVDS hedge portfolio mean return is always positive and significant when using a twelve-month IVR-IVD threshold, always positive and generally statistically significant when using a six-month threshold, but generally insignificant when using a one-month threshold. Therefore, this analysis reveals, under certain conditions, a positive relation between idiosyncratic risk and subsequent returns. We attribute the insignificant one-month IVDS results to IVR-IVD serving as a poor proxy for priced idiosyncratic risk innovations when the IVR-IVD threshold is short (this is discussed in more detail later in the paper).

If idiosyncratic risk innovations are priced, then the mean returns of the IVRS and IVDS hedge portfolios should have different signs. In contrast, a negative price of idiosyncratic risk suggests that one or both of these portfolios should have negative returns, but does not suggest positive mean returns for either portfolio. Overall, our results are fully consistent with a negative relation between idiosyncratic risk *innovations* and subsequent returns, but sometimes inconsistent with a negative relation between idiosyncratic risk levels and returns.

Panel 4 of Table 3 reports mean returns of hedge portfolios formed using idiosyncratic risk proxies generated from monthly data. The IVM hedge portfolio is formed by sorting stocks on idiosyncratic volatility calculated from monthly data. The IVM hedge portfolio's mean raw returns are insignificant, although the alpha is negative under value weighting. We also report the mean returns of a hedge portfolio formed by sorting on predicted volatility from the best fitting EGARCH(i,j) model for each stock, where $i, j \leq 3$ (following Fu (2009)). There is little evidence of nonzero returns when examining the EGARCH(i,j) hedge portfolio. This result contrasts with that of Fu (2009), who finds that predicted idiosyncratic volatility from the best fitting EGARCH(i,j) model is positively related to subsequent returns. However, Guo, Kassa, and Ferguson (2012) and Fink, Fink, and He (2012) show that Fu's estimator is biased, and that use of an unbiased estimator results in an insignificant relation.[16] We confirm both results. Fu's predicted volatility is positively related to returns (not tabulated), although use of an unbiased (out-of-sample) estimator results in no significant positive relation.

2.4 Fama-MacBeth Regressions

In this section, the Fama and MacBeth (1973) procedure is used to estimate the relation between idiosyncratic volatility and the cross-section of returns. One advantage of the Fama-MacBeth procedure is that it is easy to simultaneously control for many characteristics. However, many of the standard controls used in cross-sectional regressions plausibly capture information about idiosyncratic volatility. Then, some of the control variables may serve to control for idiosyncratic risk, which is not desirable when examining the idiosyncratic risk-return relation.[17] For this reason, results are reported from a regression with only IVR and IVD and from a regression with IVR, IVD, and controls. We also examine a regression of returns on IVR and IVC (equal to IVR-IVD). This regression can be interpreted as a horse race between idiosyncratic risk levels and innovations as explanatory variables.

Cross-sectional regressions are run using both OLS and WLS (with weights equal to market capitalizations[18]). The OLS and WLS regressions correspond to an equal-weighted and

[16] This bias arises from the positive skewness of individual stock returns and the use of month t data when using an EGARCH model to predict time t volatility and time t returns. Excluding month t data eliminates the bias.

[17] Suppose idiosyncratic risk is priced, so that firms with high idiosyncratic volatility have high discount rates. Then market-to-book may capture information about idiosyncratic volatility; a high market-to-book value could indicate low discount rates and low idiosyncratic volatility. Small stocks may earn high returns because small stocks tend to have high idiosyncratic volatility. Illiquid stocks may earn high returns because illiquid stocks tend to have high idiosyncratic volatility.

[18] Under WLS, we minimize $\sum w_i e_i^2$, where w_i is market capitalization and e_i is the difference between

value-weighted approach (respectively). The WLS regressions impose a different weighting scheme than the value-weighted portfolios of Table 3, where each stock receives a weight equal to the stock's share of the sorted portfolio. For example, a small stock may receive a large weight in the IVR5 portfolio, which consists of many small stocks, but still have a small weight in the WLS regression.

Stock characteristics considered in the cross-sectional regressions are IVR, IVD, IVC, market capitalization, market-to-book ratio, prior return from month -6 to month -2, prior return over month -1, illiquidity[19], and maximum daily return over the prior month (see Bali, Cakici, and Whitelaw (2011)). We consider the latter because this characteristic appears to be informative about expected skewness. Results are robust to using other skewness measures, including expected skewness as constructed by Boyer, Mitton, and Vorkink (2010) and historical skewness. The sample of stocks used in this section is smaller than the sample of stocks used in the sorted portfolio section because turnover and book value data are required. This data requirement eliminates small, illiquid stocks from the sample and should further alleviate concerns that results are driven by this subset of stocks, but also reduces power by reducing the cross-sectional dispersion of idiosyncratic volatility and the number of observations.

Table 4 reports results using a one-, six-, and twelve-month IVR-IVD threshold. Consistent with the results of Ang, Hodrick, Xing, and Zhang (2006), as the sole explanatory variable, one-month IVR is negatively related to subsequent returns. However, this relation is fragile. As the sole explanatory variable, IVR is not significant in any of the WLS regressions. This suggests that the negative IVR-return relation is only important for smaller stocks. Also, the negative IVR-return relation is strongest for one-month IVR and weakest (insignificant) for twelve-month IVR, even when using OLS. This is troubling if one

the actual and fitted return. Under OLS, $w_i = 1$.

[19] Following Amihud (2002), illiquidity is calculated as the log of the trailing one-year average of daily $|R_{i,t}|/DVOL_{i,t}$, where $R_{i,t}$ is the return of stock i on day t and $DVOL$ is dollar volume.

interprets the negative IVR-return relation as evidence supporting a negative idiosyncratic risk-return relation because use of a stronger proxy for idiosyncratic risk (twelve-month IVR, see Table 1) yields a weak, and sometimes insignificant, idiosyncratic risk-return relation.

Regressions with both IVR and IVC suggest that there is a more robust relation between idiosyncratic risk innovations and subsequent returns than idiosyncratic risk levels and subsequent returns. In each of these regressions, the t-statistic associated with the IVC parameter is larger in magnitude than the t-statistic associated with the IVR parameter. Of these six regressions, IVR is marginally significant in one case, and insignificant in the others. IVC is always significant, usually at the 1% level. These results suggest that idiosyncratic risk innovations, rather than levels, drive returns, particularly when using better proxies for idiosyncratic risk (six- and twelve-month IVR).

Consistent with innovations, rather than levels, driving returns, the IVR-return relation is far stronger when controlling for IVD (varying IVR while controlling for IVD is equivalent to varying IVC while controlling for IVD). For example, in the twelve-month IVR-IVD threshold OLS regression, including IVD decreases the IVR parameter from -0.064 to -0.266, and the t-statistic from -0.95 to -4.58. Similarly, the IVR parameter is not significant in the univariate WLS regressions, although controlling for IVD reveals a significant IVR parameter.

In most regressions, the estimated IVR and IVD parameters are of a similar magnitude, but have opposing signs (the estimated IVR parameter is always negative while the estimated IVD parameter is always positive). This is also consistent with a negative relation between idiosyncratic risk innovations (IVR-IVD) and subsequent returns. In the regressions that contain both IVR and IVD, the IVR slope parameter can be interpreted as the relation between idiosyncratic risk innovations and returns (this can be seen by rewriting Equation 2 with (IVR-IVD) and IVD as regressors).

Many of the characteristics exhibit a weaker relation with returns when using WLS. For example, when using OLS and a six-month IVR-IVD threshold, illiquidity appears to be an important characteristic both economically and statistically, although this is no longer true when using WLS (the point estimate associated with illiquidity is 0.188 in the equal-weighted regressions (with a t-statistic of 3.91) and 0.028 in the value-weighted regressions (with a t-statistic of 0.63)). Similarly, the economic and statistical significance of market-to-book, one-month prior return, and lagged maximum return are attenuated in the value-weighted regressions. In contrast, the IVD and IVR parameters have similar magnitudes in the equal- and value-weighted regressions. This suggests that the IVR- and IVD-return relations are pervasive, not likely to be explained by return patterns associated with small stocks (e.g. bid-ask bounce or reversals), and likely relevant to the average investor. In contrast, in the WLS regression with only IVR, the slope parameter is never significant. Therefore, the results of this section suggest that the return patterns documented in this paper are of greater economic importance than previously documented results related to idiosyncratic risk.

Overall, the results of Table 4 confirm the sorted portfolio results. Controlling for IVD, IVR is negatively related to subsequent returns. Controlling for IVR, IVD is positively related to subsequent returns. In a horse race between IVR and IVC, IVC is generally highly significant while IVR is not. Finally, in the absence of controls, the IVR-return relation is strongest when using a relatively weak proxy for expected idiosyncratic risk (one-month IVR), but weak and often insignificant when using a stronger proxy for expected idiosyncratic risk (six- or twelve-month IVR). These results suggest that risk innovations, rather than risk levels, are driving the relation with subsequent returns.

2.5 Deferred Holding Periods

In the cross section, idiosyncratic risk innovations are negatively related to subsequent returns. This complicates estimation of the price of idiosyncratic risk because risk levels and innovations are correlated (see Table 2).[20] However, it seems sensible that the relation between idiosyncratic risk innovations should attenuate relatively quickly after measuring idiosyncratic risk, as risk innovations are not very persistent (i.e. an increase in idiosyncratic risk today does not suggest a similar increase in a month). In contrast, the cross-sectional distribution of idiosyncratic risk levels is quite persistent. Then, stocks that currently earn a high idiosyncratic risk premium should continue to do so for some time. Therefore, one way to separate the effects of idiosyncratic risk levels and innovations is to observe returns for many months after measuring idiosyncratic risk. Provided the effects of risk levels are more persistent than the effects of risk innovations, sufficiently deferred hedge portfolio returns should primarily reflect the price of idiosyncratic risk.

Importantly, examining deferred returns will only be useful if idiosyncratic volatility is persistent. The time series average of the cross-sectional correlation between six-month IVR and three-, five-, and ten-year subsequent six-month IVR is .56, .50, and .44 respectively (the Spearman correlations are 0.72, 0.67, and 0.58). This can also be seen in Figure 1, which presents the results of this section graphically. In particular, the dashed plot shows the evolution of the difference in IVR across the high and low idiosyncratic volatility portfolios. Even after ten years, the equal-weighted IVR difference of 0.028 (and value-weighted difference of 0.022) is large compared to the average cross-sectional mean and standard deviation of IVR (0.028 and 0.017, see Table 2). We conclude that IVR can be reasonably used to form portfolios with dispersion in expected idiosyncratic volatility long after portfolio formation.

Table 5 reports six-month equal-weighted returns of idiosyncratic volatility-sorted hedge portfolios for up to ten years after portfolio formation.[21] First, we note that the return

[20] A regression of returns on IVR and IVC does not fully solve this problem because (1) we do not know the "correct" IVR-IVD threshold, so IVR-IVD is likely an imperfect proxy for risk innovations, (2) the idiosyncratic risk premium and relation with innovations likely varies across firms, and (3) both IVR and IVD measure expected idiosyncratic risk with error.

[21] Some stocks drop out of the sample between the idiosyncratic volatility and return measurement periods. However, these portfolios do not suffer from look-ahead bias because investors are aware of the stocks that

patterns documented in the prior section are not very persistent. The negative returns of the equal-weighted IVR and IVRS hedge portfolio persist for only six months after portfolio formation. In fact, the six-month equal-weighted IVR hedge portfolio raw return is slightly positive. Although Table 5 does not report value-weighted results, Figure 1 indicates that the negative idiosyncratic risk-return relation is not very persistent even for value-weighted returns. In this figure, equal-weighted IVR hedge portfolio returns are negative for about five months after portfolio formation, while value-weighted returns are negative for about one year. In both cases the negative returns are largest in magnitude in the month immediately after portfolio formation, then quickly attenuate. After about eighteen months, the monthly returns of both equal- and value-weighted hedge portfolios are always positive.

In Table 5, the mean returns of every equal-weighted idiosyncratic volatility-sorted portfolio (IVR, IVD, IVRS, and IVDS), are always positive starting six months after portfolio formation (although these returns are not always significant). Value-weighted IVR returns are always positive starting eighteen months after portfolio formation. Starting about five years after portfolio formation, six-month returns of the IVR and IVD hedge portfolio are generally around 3% and always at least marginally significant. Combined with previously reported results, this provides evidence that a single measure of idiosyncratic risk (IVR) can be sometimes positively and sometimes negatively related to subsequent returns, depending on the return measurement period. The long-run positive relation is consistent with theory, which predicts a non-negative price of idiosyncratic risk. Indeed, compensation for risk seems a particularly appealing explanation for the long-run positive returns of Table 5 and Figure 1, as most types of mispricing are likely corrected within five years. In untabulated results, we find that the long-run positive returns of idiosyncratic volatility-sorted portfolios are robust to the use of different return aggregation periods and IVR-IVD thresholds (including use of one-month IVR), value weighting, and controlling for size and liquidity.[22]

have dropped out of the sample before the return measurement period.

[22] We control for size and illiquidity using sequentially-sorted portfolios, where stocks are first sorted into

Overall, this evidence suggests that the underlying price of idiosyncratic risk is positive, with the negative returns immediately after portfolio formation likely attributable to some sort of transitory friction.

2.6 Idiosyncratic Risk and Market Returns

Table 5 reports raw portfolio returns. In untabulated results, we find that the long-run alphas of the IVR and IVD portfolios are insignificant. For example, for six-month returns seven years after portfolio formation, raw equal-weighted idiosyncratic risk hedge portfolio returns are 3.03% (standard error of 1.29%), while the three-factor alpha is 1.43% (standard error of 1.01%).[23] This suggests that commonly-used factors can "explain" the returns of idiosyncratic volatility hedge portfolios. However, as discussed above, this does not imply that idiosyncratic risk is not priced, as market returns may reflect compensation for idiosyncratic risk. Reversing the analysis, mean excess market returns are 2.87% (standard error of 1.16%), while the intercept from a regression of market returns on idiosyncratic risk hedge portfolio returns is 1.33% (standard error of 1.02%). Therefore, our results suggest that idiosyncratic risk may explain much of market portfolio excess returns (about 3% annually). While it remains possible that traditional beta exposure may explain the positive price of idiosyncratic risk, the fact that idiosyncratic risk, but not beta, is priced in the cross section (see Fama and French (2004) and Lewellen and Nagel (2006) for evidence that beta is not priced) suggests that aggregate market returns may be more attributable to empirical measures of idiosyncratic risk than beta.

As an alternative approach, we (roughly) estimate the portion of the market premium that may be attributable to idiosyncratic risk by first estimating the price of idiosyncratic risk using the deferred returns of Figure 1, then applying this price to the idiosyncratic risk of

size (or liquidity) quintiles and then IVR or IVD quintiles. The IVR and IVD sequentially-sorted hedge portfolios exhibit statistically significant positive long-run returns. When we reverse the sort order, the long-run returns of the size and liquidity-sorted portfolios are not significantly different from zero.

[23] Results are generally similar for other sufficiently deferred return measurement periods.

the average stock. In Figure 1, deferred hedge portfolio returns are generally around 0.50% monthly, while the difference in idiosyncratic volatility across the high and low portfolios is approximately 2.9% and 2.5%, for equal- and value-weighted returns respectively. The time-series average of value-weighted idiosyncratic volatility for all stocks is 1.58% (the corresponding equal-weighted average, reported in Table 2, is 2.81%). Using value-weighted hedge portfolio returns, this suggests an aggregate idiosyncratic risk premium of 3.27% (equal to 1.58% * 12 * 0.50% / 2.90%). Use of equal-weighted hedge portfolio returns suggests an aggregate risk premium of 3.79%. Therefore, this approach also suggests that a substantial portion of aggregate market returns may be explained by idiosyncratic risk (and yields estimates of a similar magnitude).

2.7 Market Model and Alternative Samples

The primary sample consists of data from 1966-2012. This facilitates comparison to the work of Ang, Hodrick, Xing, and Zhang (2006), who examine a similar sample. In this section, we examine the robustness of our results to two changes. First, we compute idiosyncratic volatility using the market model rather than the three-factor model. Second, we consider an expanded sample (1929-2012).

Under the market model, idiosyncratic volatility is the standard deviation of the residuals from a time series regression of a stock's returns on the returns of the value-weighted stock index. Idiosyncratic volatility calculated under the market model is highly correlated with idiosyncratic volatility calculated under the three-factor model (pooled sample correlation in excess of 0.99). Therefore, idiosyncratic risk estimates appear to be insensitive to this change in empirical models.

Table 6 reports six-month mean returns of hedge portfolios for ten years subsequent to portfolio formation. Table 6 is similar to Table 5, although we use an expanded sample and the market model in Table 6. Results are generally similar. The IVRS hedge portfolio exhibits short-lived negative returns shortly after portfolio formation (although the six-month IVRS return is insignificant, the one-month IVRS hedge portfolio return is -0.45%, with a t-statistic of -4.71). Six months after portfolio formation, the mean returns of the IVR, IVD, IVRS, and IVRS portfolios are always positive and generally significant. Therefore, use of the market model and an expanded sample provides additional evidence that the underlying idiosyncratic risk-return relation is positive.

3 Interpreting the Evidence

Stocks with high idiosyncratic risk tend to earn high returns starting approximately six months after measuring idiosyncratic risk. However, in the short run, there is a negative relation between idiosyncratic risk innovations and and subsequent returns. In this section, we take a step toward understanding the economic mechanisms that drive these return patterns.

We start by presenting a simple model, featuring a positive price of idiosyncratic risk and price underreaction to idiosyncratic risk innovations, that can accommodate all of our empirical results. In contrast, existing explanations of idiosyncratic risk-related return patterns, which often focus on the anomaly of Ang, Hodrick, Xing, and Zhang (2006), are generally inconsistent with some of these results. Our model highlights how, in the presence of underreaction, the sign of the relation between historical idiosyncratic risk and returns is ambiguous, even if the price of idiosyncratic risk is assumed to be positive. Then, standard empirical procedures, which rely on this relation, may yield misleading inference regarding the price of idiosyncratic risk. Also, the model aids in interpreting our empirical results.

We reiterate that the primary objective of this paper is not to explain the anomaly documented by Ang, Hodrick, Xing, and Zhang (2006), but to understand *all* of the return patterns associated with idiosyncratic risk. While the Ang, Hodrick, Xing, and Zhang

(2006) anomaly is a piece of this puzzle, the prior section demonstrates that there are also other economically important return patterns associated with idiosyncratic risk.

3.1 A Simple Model of Investor Underreaction to Risk Innovations

Because we are focused on risk (and discount rates), rather than cash flows, we adopt a dividend discount model, where expected cash flows are held constant throughout. We assume that expected returns are determined solely by idiosyncratic risk[24] and, as described below, idiosyncratic risk innovations (i.e. other characteristics are not relevant). This assumption is clearly not realistic (e.g. we ignore size and value), but is made to simplify the exposition. Extending the model to account for other characteristics is straightforward.

We assume that the price of the stock follows

$$p_t = \frac{d}{r_t^*} = \frac{d}{\gamma IV_t^*}, \tag{3}$$

where p_t is the price, d is the dividend, IV_t^* is "priced" idiosyncratic volatility, and r_t^* is the commensurate expected return. γ is a positive risk aversion parameter that maps idiosyncratic risk to discount rates.

Under our model, priced idiosyncratic volatility (IV_t^*) may differ from true idiosyncratic volatility (IV_t). We discuss reasons why this may occur later in this section. We assume that the log of IV follows a random walk,

$$log(IV_{t+1}) = log(IV_t) + \epsilon_{t+1}. \tag{4}$$

[24]Canonical asset pricing models (e.g. the Sharpe-Lintner CAPM) say that, because investors are free to diversify, idiosyncratic risk is not priced. However, violations of the assumptions underlying these models can lead to a nonzero price of idiosyncratic risk. An idiosyncratic risk premium may be caused by constraints or frictions that limit investors' ability to diversify (see Levy (1978), Merton (1987), Hirshleifer (1988), Malkiel and Xu (2006)).

Priced idiosyncratic volatility (IV^*) evolves according to

$$IV^*_{t+1} = IV^*_t + \Theta(IV_t - IV^*_t), \tag{5}$$

with $0 < \Theta < 1$. Therefore, in the absence of additional shocks (ϵ) priced idiosyncratic volatility will eventually converge with true idiosyncratic volatility. When $IV_t = IV_t^*$, p_t is determined solely by true idiosyncratic risk (i.e. underreaction does not affect prices); this can be interpreted as the equilibrium stock price.

The effects of a positive shock to idiosyncratic volatility under this model are presented graphically in Figures 2 and 3. Examining these figures, it is clear that prices underreact to risk innovations. This results in predictable returns for a period of time as the shock is gradually incorporated into prices. However, in the long run, prices reflect true idiosyncratic volatility.

Under this model, the expected gross return,

$$E_t[\frac{p_{t+1} + d}{p_t}] = \frac{IV^*_t}{IV^*_t + \Theta(IV_t - IV^*_t)} + \gamma IV^*_t, \tag{6}$$

will depend on the idiosyncratic volatility error ($IV_t - IV_t^*$) and priced idiosyncratic volatility (IV_t^*). It is straightforward to show that if priced idiosyncratic volatility is too low (i.e. $IV_t > IV_t^*$), then next period's expected return will be low (relative to the case where $IV_t = IV_t^*$). This low expected return corresponds to an expected increase in priced idiosyncratic volatility and the discount rate, which reduces the price of the stock. Also, holding $IV_t - IV_t^*$ constant, higher priced idiosyncratic volatility will be associated with higher expected returns.

3.2 Empirical Implications

The model implies that, controlling for the idiosyncratic volatility error $(IV_t - IV_t^*)$, priced idiosyncratic volatility is positively related to subsequent returns. Also, controlling for the level of priced idiosyncratic volatility, recent innovations in idiosyncratic volatility are negatively related to subsequent returns. Alternatively, the empirical implications of the model can be stated as

$$\frac{\partial E_t[R_{t+1}]}{\partial IV_t} < 0 \tag{7}$$

and

$$\frac{\partial E_t[R_{t+1}]}{\partial IV_t^*} > 0, \tag{8}$$

where R_{t+1} is the time $t+1$ gross return.

Prior studies often focus on the relation between returns and a measure of the level of expected idiosyncratic volatility. Such a test can be expressed as

$$\frac{dE_t[R_{t+1}]}{dX_t} \neq 0, \tag{9}$$

where X is some proxy for expected idiosyncratic risk (possibly some measure of historical idiosyncratic volatility). Our framework does not offer a clear prediction of the sign of this relation because X_t could be positively correlated with both IV_t^* and $(IV_t - IV_t^*)$, which have opposing relations with subsequent returns.

3.3 Motivating the Idiosyncratic Risk Innovation-Return Relation

Should prices underreact to risk innovations? There is apparent evidence of underreaction in a wide variety of settings (see Footnote 5). Then, it should not be surprising to find underreaction to risk innovations. A simple explanation of price underreaction is investor underreaction (i.e. investors incorporate idiosyncratic risk news into prices with a delay).

This can be caused by behavioral biases (Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrahmanyam (1998))[25], slow information diffusion (Hong and Stein (1999), Hong, Torous, and Valaknov (2007)), or information capacity constraints (Sims (2003)). More generally, price underreaction may be attributable to investors' collective inability to behave as rational and computationally unconstrained agents (see Hirshleifer (2001) and Barberis and Thaler (2003)).

Is this sort of investor underreaction likely? Ultimately, we cannot offer a decisive answer to this question because we do not observe stock[26] investors' idiosyncratic risk estimates. However, prior research suggests that investor underreaction is particularly likely in this application, for a number of reasons.

First, idiosyncratic risk must be estimated from market data and other information sources. The relevant information set could easily be large, diverse, and continuously changing. Monitoring this information set in real time is likely a challenging task. Therefore, this would seem to be a setting where the assumption of rational and computationally unconstrained investors is particularly questionable. In contrast, earnings announcements are released in an easily processed form at a pre-scheduled time, which suggests that earnings announcements should be relatively easy for investors to process. However, investors appear to underreact to earnings announcements (Bernard and Thomas (1989, 1990)). If investors underreact to earnings announcements, and earnings announcements are easier to identify and interpret than idiosyncratic risk innovations, then investors are likely to underreact to idiosyncratic risk innovations. Empirical studies find that post-earnings announcement drift is exacerbated by investor distraction (see Hirshleifer, Lim, and Teoh (2009) and DellaVigna

[25]One explanation of momentum (Jegadeesh and Titman (1993)) is price underreaction caused by behavioral biases. Barberis, Shleifer, and Vishny develop a model where earnings follow a random walk, but the representative investor believes earnings either trend or mean revert, and show that such a model can lead to short-run underreaction (and momentum). Daniel, Hirshleifer, and Subrahmanyam show how overconfidence and biased self-attribution can also lead to these return patterns. Similar biases may explain price underreaction to risk innovations.

[26]Because the representative stock and option investor's idiosyncratic volatility estimates may differ, use of option implied volatilities also cannot decisively answer this question.

and Pollet (2009)). Because the task of continuously estimating future idiosyncratic risk for all stocks in real time would seem to be associated with relatively high levels of investor distraction, these results are consistent with investor underreaction to idiosyncratic risk innovations.

Second, idiosyncratic risk estimates may be imprecise; this could exacerbate the effects of behavioral biases and investor underreaction (see Zhang (2006)). Third, a particularly useful predictor of idiosyncratic risk is historical idiosyncratic risk. Historical idiosyncratic risk is often calculated as the standard deviation of the residuals from a time-series regression of returns on contemporaneous factors (e.g. market returns or the three factors of Fama and French (1996)). For much of the historical sample, many investors likely lacked the technical expertise and/or computing power required to calculate this measure for a large number of stocks in real time. In this case, these investors could not use all publicly-available information when forming an idiosyncratic risk estimate. Finally, there is evidence that investors underreact to volatility innovations when setting prices for S&P 500 index options (Poteshman (2001)). Because non-traded firm-level idiosyncratic volatility estimates almost certainly suffer from more severe underreaction than traded stock index options, this suggests that investors underreact to idiosyncratic volatility innovations.

Are there rational explanations for price underreaction? It seems sensible that transactions costs and short-sale constraints will, to some degree, limit arbitrage activity and may lead to price underreaction. Additionally, investors may face substantial information gathering costs. Properly accounting for such costs could make investor underreaction a rational outcome.

In this paper, we do not attempt to identify the source of price underreaction or classify this underreaction as rational or irrational. We view this as an exciting, albeit challenging topic for future research. The goal of this paper is to more fully characterize the return patterns associated with idiosyncratic risk and present a framework that can simultaneously accommodate all of these results. Of course, fully documenting these return patterns is a critical step toward identifying the underlying economic mechanisms. For example, the Ang, Hodrick, Xing, and Zhang (2006) anomaly suggests a negative idiosyncratic risk-return relation, and has led researchers to consider mechanisms that are consistent with such a relation. Our results suggest that such explanations do not fully address the return patterns associated with idiosyncratic risk.

3.4 The Importance of Idiosyncratic Risk Innovations

Price underreaction could easily influence the conclusions of empirical researchers. For example, in Figure 3, expected returns are negative for a period of time after the shock. If underreaction persists for several months, an empirical study that forms portfolios monthly may infer that the idiosyncratic risk-return relation is negative. However, Figure 3 was generated under the assumption of a positive price of idiosyncratic risk. It is interesting to note that, in the presence of price underreaction to risk innovations, an empirical study may possess more power detecting returns associated with underreaction than changing risk premia (or even the level of risk premia). This occurs because a persistent change in required returns has a large effect on current prices (see Campbell (1991), who shows that the contemporaneous return associated with a change in expected returns is equal to the sum of the discounted expected return changes).

For example, consider a stock that pays a dividend of one dollar annually in perpetuity. Immediately after a dividend, if the discount rate is 10%, the value of the stock is 10. If the discount rate increases by 1%, due to a change in risk, the new value of the stock is 9.09. If prices fully adjust to the risk innovation within one month, then the monthly return of the stock will be -8.3%[27]. A monthly return of -8.3% should be easier to empirically detect than the change in risk premia, equal to 1% annually or 0.08% monthly. This return should also

[27] The return is not -9.1% because, after one month, each dividend is one month closer.

be easier to detect than the level of the new risk premium, which will be less than 0.92% monthly (assuming the risk-free rate is greater than zero).

3.5 Revisiting the Results

Next, we revisit the empirical results within the context of our price underreaction framework. If our framework is correct, we should observe a negative relation between true idiosyncratic risk when controlling for priced idiosyncratic risk and a positive relation between priced idiosyncratic risk when controlling for true idiosyncratic risk, as in Equations 7 and 8 (alternatively, we could state these implications as a negative partial relation between idiosyncratic risk innovations and returns and a positive partial relation between priced idiosyncratic risk and returns). Also, when priced idiosyncratic risk equals true idiosyncratic risk, there should be a positive relation between idiosyncratic risk and returns (theory suggests the equilibrium price of idiosyncratic risk is non-negative). Finally, our framework allows for *temporary* underreaction, which suggests that the strength of certain return patterns will depend on when idiosyncratic risk and returns are measured (e.g. risk innovations are unlikely to affect five-year deferred returns).

Our framework can be easily mapped into our empirical results. Distant idiosyncratic volatility (IVD) is used to proxy for priced idiosyncratic risk, while recent idiosyncratic volatility (IVR) is used to proxy for true idiosyncratic risk.[28] This is sensible because, at some fixed point in time, distant information is more likely fully incorporated into prices than recent information. Then, our framework suggests a positive IVD-return relation and a negative IVR-return relation, after implementing the appropriate controls. This is exactly what we see in Tables 3 and 4. Additionally, our framework offers no clear prediction for the

[28]Under the simple model presented in this paper, idiosyncratic risk follows a random walk and IVR-IVD can be interpreted as the innovation in true idiosyncratic risk. Also, IVR-IVD is correlated with the with the effects of price underreaction related to this innovation. Under a different specification (e.g. an AR(1) process for true idiosyncratic risk), IVR-IVD may not equal the innovation in true idiosyncratic risk. However, in this case, IVR-IVD is likely correlated with both the innovation and the effects of price underreaction related to this innovation.

IVR- and IVD-return relations without the appropriate controls. Therefore, our framework is consistent with the insignificant IVR- and IVD-return relations in the absence of controls (see Tables 3 and 4), as well as the sometimes positive and sometimes negative IVR-return relation of Figure 1. Finally, there is a strong and robust relation between idiosyncratic risk innovations and subsequent returns that largely subsumes the negative relation between historical idiosyncratic risk and returns (see Table 4). Overall, these results are consistent with our framework.

Assuming our framework is correct, our empirical results can be used to estimate the duration of price underreaction. The IVR- and IVD-return relations will be consistent with our framework only under a suitable IVR-IVD threshold. For example, if the IVR-IVD threshold is very short, IVD may capture information that has not yet been fully assimilated into prices. In this case, the IVD parameter (and the IVDS hedge portfolio mean return) should be small and insignificant. If the IVR-IVD threshold is very long, then investors may have largely assimilated the information content of IVR, and the IVR parameter (and IVRS hedge portfolio mean return) should be small and insignificant. In Tables 3 and 4, we find weaker results when using a one-month threshold. This suggests that price underreaction meaningfully affect returns for longer than a month.

Our framework allows for temporary price underreaction to idiosyncratic risk innovations. However, in the long run, idiosyncratic risk is properly reflected by prices. This suggests that the relation between idiosyncratic risk innovations and returns may be important immediately after measuring idiosyncratic risk, but should decline in importance as the measurement interval becomes more distant. Therefore, our framework can accommodate the deferred return patterns of Figure 1 and Tables 5 and 6, where returns of high idiosyncratic risk stocks are negative immediately after measuring idiosyncratic risk, then increase (eventually becoming positive). Importantly, our framework predicts that the short-run innovation-return relation and the long-run level-return relation should have opposing signs, which is consistent with the data.

The short-lived nature of the negative relation (in Figure 1 the negative IVR-return relation is concentrated in the first month, and quickly attenuates) suggests that this relation is not driven by a persistent explanatory variable, which should be associated with a persistent negative relation. Then, idiosyncratic risk innovations (which are not persistent) are more likely to explain the short-lived negative relation than idiosyncratic risk levels (which are persistent). Therefore, the return patterns of Figure 1 and Tables 5 and 6, which indicate that the negative relation is not persistent, provide additional evidence that risk innovations, rather than levels, are driving this relation (as is also suggested by Table 4).

Overall, our framework can accommodate the empirical results of this paper. In contrast, other explanations are often inconsistent with some of the results. For example, one may infer that the idiosyncratic risk-return relation is negative based on the empirical results of Ang, Hodrick, Xing, and Zhang (2006). Our results present this interpretation with several difficulties. First, this explanation is difficult to reconcile with the high deferred returns of high idiosyncratic risk stocks. Second, this explanation does not address the strong and robust positive relation between IVD and returns when controlling for IVR. Third, the empirical evidence for a negative idiosyncratic risk-return relation is strongest when using a relatively weak proxy for subsequent idiosyncratic risk (one-month IVR), and often insignificant when using a stronger proxy, which is troubling. Finally, the return pattern of Figure 1 (in particular, the transitory nature and quick attenuation of the initial negative relation) suggests that the negative relation documented by Ang, Hodrick, Xing, and Zhang (2006) is not attributable to a persistent explanatory variable, such as idiosyncratic risk.

For these reasons, our results offer little support for general interpretations of existing explanations of the Ang, Hodrick, Xing, and Zhang (2006) anomaly. Authors have attempted to explain this relation by appealing to investor preference for skewness (Boyer, Mitton, and

Vorkink (2010), Bali, Cakici, and Whitelaw (2011)), heterogeneous investor beliefs combined with short sale constraints (Boehme, Danielsen, Kumar, and Sorescu (2009), George and Hwang (2011)), investor preference for high idiosyncratic volatility stocks (Han and Kumar (2008)), and risks omitted from the three-factor model (Chen and Petkova (2012), Barinov (2011)). Many of these explanations rely on persistent explanatory variables and do not address some of our empirical results. We do not rule out these explanations in specific settings (e.g. these explanations may explain why one-month IVR is negatively related to returns, even after controlling for risk innovations, in Table 4). However, these explanations do not appear to be more generally valid and may be of limited economic importance (the Ang, Hodrick, Xing, and Zhang (2006) anomaly appears to be concentrated in small stocks, while other return patterns associated with idiosyncratic risk appear to be more pervasive).

3.6 Alternative Explanations

An alternative explanation of our results is that the short- and long-run relations are distinct. Perhaps idiosyncratic risk can be thought of as a composite stochastic process, with a short-run and long-run component, where each component is priced differently. However, simultaneously asserting a positive risk premium for long-run idiosyncratic risk and a negative risk premium for short-run idiosyncratic risk does not seem sensible. It is possible that the short-run idiosyncratic risk-return relation is not driven by idiosyncratic risk, but some characteristic correlated with idiosyncratic risk (e.g. one of the explanations of the negative relation discussed above). However, as discussed above, the negative relation between idiosyncratic risk and returns is fragile, and seems to be driven by innovations in idiosyncratic risk. Still, the short-run relation could be caused by some characteristic correlated with idiosyncratic risk innovations. However, we have controlled for many return patterns that may plausibly explain this relation (e.g. short-term reversals, bid-ask bounce, liquidity). A remaining alternative explanation for the short-run relation is temporary misreaction to cash flow news. It is difficult to fully eliminate this possibility because cash flow news is difficult to measure (see Chen and Zhao (2009)). However, we control for the portion of prior cash flow news that is correlated with prior returns, which is likely substantial.

3.7 Cash Flows and the Contemporaneous Idiosyncratic Risk-Return Relation

Table 2 shows that, on average, high idiosyncratic volatility stocks tend to have high prior returns. The positive contemporaneous relation between returns and idiosyncratic volatility was documented by Duffee (1995), and is consistent with the positive skewness of the average stock's return distribution. However, partial underreaction to idiosyncratic risk innovations and a positive price of idiosyncratic risk suggests that idiosyncratic risk innovations should be negatively related to contemporaneous returns. This apparent contradiction can be explained by changing expected cash flows. High idiosyncratic risk firms, and firms with increases in idiosyncratic risk, tend to have positive cash flow news, as measured by the change in earnings contemporaneous with the IVR measurement period. This suggests that, given some innovation in idiosyncratic risk, cash flow and discount rate news tends to offset (i.e. a positive shock to idiosyncratic risk suggests lower prices due to increased risk, but tends to occur at the same time as good cash flow news, which suggests higher prices). Although it may be of interest to decompose these returns into a cash flow and discount rate component, this is a challenging task.[29]

[29] The contemporaneous change in earnings is almost certainly an insufficient measure of cash flow news. A present-value decomposition, as in Vuolteenaho (2002), may allow one to isolate the cash flow and discount rate news, although such decompositions appear to be unreliable because cash flow news is estimated as the residual after modeling firm-level expected returns, which are likely estimated with substantial error (see Chen and Zhao (2009)).

4 Conclusion

In this paper, we document a short-lived negative relation between idiosyncratic risk innovations and subsequent returns and a persistent positive relation between idiosyncratic risk levels and subsequent returns. These relations are consistent with a positive price of idiosyncratic risk and temporary price underreaction to idiosyncratic risk innovations. Because idiosyncratic risk levels and innovations are correlated, these relations tend to offset in standard empirical studies that examine the relation between historical idiosyncratic risk and subsequent returns. Importantly, we find the relations documented in this paper to be more robust and likely more economically important than previously documented return patterns related to idiosyncratic risk.

We develop a simple model to examine the idiosyncratic risk-return relation in the presence of underreaction. The model highlights how, in the presence of underreaction, standard empirical procedures may yield misleading inference regarding the cross-sectional price of idiosyncratic risk. We show that, even if the price of idiosyncratic risk is positive, in the presence of underreaction, the sign of the relation between historical idiosyncratic risk and subsequent returns can be negative. Therefore, our framework can simultaneously accommodate theory, which generally suggests a non-negative price of idiosyncratic risk, and a negative empirical relation between some measures of historical idiosyncratic risk and returns (as documented by Ang, Hodrick, Xing, and Zhang (2006)). Also, our framework reconciles empirical studies that may estimate different prices of idiosyncratic risk (including differing signs). In the presence of underreaction, studies that examine the relation between recent idiosyncratic risk and immediately subsequent returns should find a negative relation if recent information is not yet fully incorporated into prices. However, studies that use a longer window to calculate historical idiosyncratic risk, or examine deferred returns, are less focused on underreaction and more likely to estimate a positive relation.

We rule out many alternative explanations of our results. We find that return patterns associated with short-term reversals, momentum, and liquidity cannot explain our results. Still, it remains possible that some omitted stock characteristic, correlated with historical idiosyncratic volatility, could explain our findings. However, it is not easy to find an alternative explanation that predicts the return pattern observed in the data (short-lived negative returns, related to risk innovations, and persistent positive returns, related to the level of risk). Overall, we find a positive price of idiosyncratic risk and underreaction to idiosyncratic risk innovations to be a compelling explanation of the return patterns associated with idiosyncratic risk.

More generally, our methodology may be useful in a variety of settings, especially when empirical studies suggest a price of risk that runs counter to theory or intuition. For example, stocks with a high probability of default (high distress risk) tend to earn low returns (see Dichev (1998) and Campbell, Hilscher, and Szilagyi (2008)). This is a counterintuitive empirical result because distressed stocks are likely at least as risky as other stocks. In preliminary work, we apply the methodology of this paper to distress risk and find results consistent with a positive price of distress risk and underreaction to innovations in distress risk.

References

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5(1), 31–56.

Ang, A., R. J. Hodrick, Y. Xing, X. Zhang, 2006. The Cross-Section of Volatility and Expected Returns. Journal of Finance 61(1), 259–299.

——— , 2009. High idiosyncratic volatility and low returns: International and further U.S. evidence. Journal of Financial Economics 91(1), 1–23.

Bali, T. G., N. Cakici, 2008. Idioysncratic Volatility and the Cross Section of Expected Returns. Journal of Financial and Quantitative Analysis 43(1), 29–58.

Bali, T. G., N. Cakici, R. F. Whitelaw, 2011. Maxing Out: Stocks as lotterires and the cross-section of expected returns. Journal of Financial Economics 99(2), 427–446.

Ball, R., P. Brown, 1968. An Empirical Evaluation of Accounting Income Numbers. Journal of Accounting Research 6(2), 159–178.

Barberis, N., A. Shleifer, R. Vishny, 1998. A model of investor sentiment. Journal of Financial Economics 49(3), 307–343.

Barberis, N., R. Thaler, 2003. A Survey of Behavioral Finance. pp. 1051–1121.

Barinov, A., 2011. Idiosyncratic Volatility, Growth Options, and the Cross-Section of Returns. Working Paper.

Bernard, V. L., J. K. Thomas, 1989. Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?. Journal of Accounting Research 27, 1–36.

——— , 1990. Evidence that Stock Prices do not Fully Reflect the Implications of Current Earnings for Future Earnings. Journal of Accounting and Economics 13, 305–340.

Boehme, R. D., B. R. Danielsen, P. Kumar, S. M. Sorescu, 2009. Idiosyncratic Risk and the cross-section of stock returns: Merton (1987) meets Miller (1977). Journal of Financial Markets 12, 438–468.

Boyer, B., T. Mitton, K. Vorkink, 2010. Expected Idiosyncratic Skewness. Review of Financial Studies 23(1), 169–202.

Campbell, J. Y., 1991. A Variance Decomposition for Stock Returns. The Economic Journal 101(405), 157–179.

Campbell, J. Y., J. Hilscher, J. Szilagyi, 2008. In Search of Distress Risk. The Journal of Finance 63(6), 2899–2939.

Chen, L., X. Zhao, 2009. Return Decomposition. Review of Financial Studies 22(12), 5213–5249.

Chen, Z., R. Petkova, 2012. Does Idiosyncratic Volatility Proxy for Risk Exposure?. Review of Financial Studies 25(9), 2745–2787.

Cohen, L., A. Frazzini, 2008. Economic Links and Predictable Returns. Journal of Finance 63(4), 1977–2011.

Daniel, K., D. Hirshleifer, A. Subrahmanyam, 1998. Investor Psychology and Security Market under- and Overreactions. Journal of Finance 53(6), 1839–1885.

DellaVigna, S., J. M. Pollet, 2007. Demographics and Industry Returns. American Economic Review 97(5), 1667–1702.

——— , 2009. Investor Inattention and Friday Earnings Announcements. Journal of Finance 64(2), 709–749.

Dichev, I., 1998. Is the Risk of Bankruptcy a Systematic Risk?. Journal of Finance 53(3), 1131–1147.

Duffee, G. R., 1995. Stock Returns and Volatility a Firm Level Analysis. Journal of Financial Economics 37, 399–420.

Eberhart, A. C., W. F. Maxwell, A. R. Siddique, 2004. An Examination of Long-Term Abnormal Stock Returns and Operating Performance Following R&D Increases. Journal of Finance 59(2), 623–650.

Fama, E. F., K. R. French, 1996. Multifactor Explanations of Asset Pricing Anomalies. Journal of Finance 51(1), 55–84.

——— , 2004. The Capital Asset Pricing Model: Theory and Evidence. Journal of Economic Perspectives 18, 25–46.

Fama, E. F., J. D. MacBeth, 1973. Risk, Return, and Equilibrium: Empirical Tests. Journal of Political Economy 81(3), 607–636.

Fink, J. D., K. E. Fink, H. He, 2012. Expected Idiosyncratic Volatility Measures and Expected Returns. Financial Management 41(3), 519–553.

Fu, F., 2009. Idiosyncratic Risk and the Cross-Section of Expected Stock Returns. Journal of Financial Economics 91(1), 24–37.

George, T. J., C.-Y. Hwang, 2011. Analyst Coverage and the Cross Sectional Relation Between Returns and Volatility. Working Paper.

Guo, H., H. Kassa, M. F. Ferguson, 2012. On the Relation between EGARCH Idiosyncratic Volatility and Expected Stock Returns. Forthcoming, Journal of Financial and Quantitative Analysis.

Guo, H., R. Savickas, 2010. Relation between time-series and cross-sectional effects of idiosyncratic variance on stock returns. Journal of Banking & Finance 34(7), 1637–1649.

Han, B., A. Kumar, 2008. Retail Clienteles and the Idiosyncratic Volatility Puzzle. Working Paper.

Han, Y., D. Lesmond, 2011. Liquidity Biases and the Pricing of Cross-sectional Idiosyncratic Volatility. Review of Financial Studies 24(5), 24–37.

Hirshleifer, D., 1988. Residual Risk, Trading Costs, and Commodity Futures Risk Premia. Review of Financial Studies 1(2), 173–193.

——— , 2001. Investor Psychology and Asset Pricing. Journal of Finance 56(4), 1533–1597.

Hirshleifer, D., S. Lim, S. H. Teoh, 2009. Driven to Distraction: Extraneous Events and Underreaction to Earnings News. Journal of Finance 64(5), 2289–2325.

Hong, H., J. C. Stein, 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. Journal of Finance 54(6), 2143–2184.

Hong, H., W. Torous, R. Valaknov, 2007. Do Industries Lead Stock Markets?. Journal of Financial Economics 83, 367–396.

Huang, W., Q. Liu, S. G. Rhee, L. Zhang, 2009. Return Reversals, Idiosyncratic Risk, and Expected Returns. Review of Financial Studies 23(1), 147–168.

Ikenberry, D., J. Lakonishok, T. Vermaelen, 1995. Market underreaction to open market share repurchases. Journal of Financial Economics 39(2-3), 181–208.

Jegadeesh, N., 1990. Evidence of Predictable Behavior of Security Returns. Journal of Finance 45(3), 881–898.

Jegadeesh, N., S. Titman, 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. Journal of Finance 48(1), 65–91.

Lehmann, B., 1990a. Fads, Martingales, and Market Efficiency. Quarterly Journal of Economics 105(1), 1–28.

———, 1990b. Residual Risk Revisited. Journal of Econometrics 45(1-2), 71–97.

Levy, H., 1978. Equilibrium in an Imperfect Market: A Constraint on the Number of Securities in the Portfolio. American Economic Review 68(4), 643–658.

Lewellen, J., S. Nagel, 2006. The Conditional CAPM does not Explain Asset-Pricing Anomalies. Journal of Financial Economics 82(2), 289–314.

Loughran, T., J. R. Ritter, 1995. The New Issues Puzzle. Journal of Finance 50(1), 23–51.

Malkiel, B. G., Y. Xu, 2006. Idiosyncratic Risk and Security Returns. Working Paper.

Menzly, L., O. Ozbas, 2010. Market Segmentation and Cross-Predictability of Returns. Journal of Finance 65(4), 1555–1580.

Merton, R. C., 1987. A Simple Model of Capital Market Equilibrium with Incomplete Information. Journal of Finance 42(3), 483–510.

Michaely, R., R. H. Thaler, K. L. Womack, 1995. Price Reactions to Dividend Initiations and Omissions: Overreaction or Drift?. Journal of Finance 50(2), 573–608.

Peterson, D. R., A. R. Smedema, 2011. The return impact of realized and expected idiosyncratic volatility. Journal of Banking & Finance 35(10), 2547–2558.

Poteshman, A. M., 2001. Underreaction, Overreaction, and Increasing Misreaction to Information in the Options Market. Journal of Finance 56(3), 851–876.

Sims, C. A., 2003. Implications of Rational Inattention. Journal of Monetary Economics 50, 665–690.

Spiegel, M. I., X. Wang, 2005. Cross-sectional Variation in Stock Returns: Liquidity and Idiosyncratic Risk. Working Paper.

Spiess, D. K., J. Affleck-Graves, 1995. Underperformance in long-run stock returns following seasoned equity offerings. Journal of Financial Economics 38(3), 243–267.

Vuolteenaho, T., 2002. What Drives Firm-Level Stock Returns?. Journal of Finance 57(1), 233–264.

Zhang, X. F., 2006. Information Uncertainty and Stock Returns. Journal of Finance 61(1), 105–137.

Table 1: Predictive Cross Sectional Idiosyncratic Volatility Regressions, R^2

	One-Year IV			One-Month IV		
	IVR-IVD Threshold			IVR-IVD Threshold		
Explanatory Variables	1	6	12	1	6	12
IVR	0.55	0.78	0.75	0.37	0.47	0.47
IVD	0.72	0.59	0.52	0.45	0.40	0.35
IVM	0.45	0.45	0.45	0.30	0.30	0.30
PV(1,1)	0.28	0.28	0.28	0.18	0.18	0.18
PV(i,j)	0.12	0.12	0.12	0.08	0.08	0.08
IVR, IVD	0.76	0.79	0.75	0.49	0.49	0.48
IVR, IVD, IVM, PV(1,1), PV(i,j)	0.77	0.80	0.76	0.50	0.50	0.48

Notes - Table reports time series average R^2 from monthly cross-sectional regressions of realized idiosyncratic volatility on an intercept, IVR, IVD, 60-month trailing idiosyncratic volatility calculated using monthly data (IVM), out-of-sample conditional idiosyncratic volatility derived from an EGARCH (1, 1) model (PV(1,1)), and out-of-sample predicted idiosyncratic volatility derived from the best fitting EGARCH(i, j) model, where $i, j <= 3$ (PV(i,j)). Results are reported for one-year (365 day) realized idiosyncratic volatility (columns 2-4) and one-month (31 day) realized idiosyncratic volatility (columns 5-8). Data spans 1966-2012.

Table 2: Sorted Portfolio Descriptive Statistics

Sort	IVR	IVD	IVC	ME	MB	PRET6	PRET1	MAXRET	ILLIQ	Cap. Share	ΔEarn
All	2.81	2.82	-0.01	11.79	0.39	6.75	1.41	0.94	0.07		2.07
All SD	1.73	1.65	1.14	1.80	0.91	28.37	13.33	1.08	0.06		19.54
IVR1	1.15	1.32	-0.17	13.05	0.43	5.05	0.92	0.43	0.03	0.56	1.08
IVR2	1.76	1.91	-0.15	12.63	0.39	6.17	1.16	0.62	0.04	0.26	1.18
IVR3	2.40	2.55	-0.15	11.91	0.39	6.73	1.27	0.83	0.06	0.11	1.04
IVR4	3.27	3.38	-0.11	11.18	0.40	6.35	1.21	1.10	0.08	0.05	1.54
IVR5	5.44	4.93	0.52	10.19	0.36	9.48	2.50	1.72	0.13	0.02	5.80
IVD1	1.29	1.19	0.10	13.08	0.42	5.04	1.03	0.46	0.03	0.58	1.13
IVD2	1.88	1.81	0.07	12.63	0.38	5.84	1.18	0.65	0.04	0.25	1.19
IVD3	2.51	2.45	0.06	11.89	0.37	6.13	1.23	0.86	0.06	0.11	1.31
IVD4	3.34	3.30	0.04	11.16	0.38	6.13	1.29	1.12	0.08	0.05	1.85
IVD5	5.01	5.34	-0.34	10.20	0.43	10.64	2.32	1.61	0.12	0.02	4.93
IVRS1	1.82	2.55	-0.73	12.12	0.42	6.10	0.94	0.67	0.04	0.30	1.24
IVRS2	2.28	2.67	-0.39	12.01	0.42	6.55	1.12	0.79	0.05	0.22	1.35
IVRS3	2.63	2.78	-0.15	11.88	0.41	6.58	1.24	0.89	0.06	0.20	1.72
IVRS4	3.07	2.92	0.15	11.69	0.39	6.83	1.42	1.01	0.07	0.17	2.30
IVRS5	4.23	3.17	1.06	11.26	0.32	7.71	2.32	1.34	0.10	0.12	3.84
IVDS1	2.55	1.83	0.72	12.20	0.33	4.74	1.27	0.84	0.06	0.33	1.94
IVDS2	2.65	2.33	0.31	12.00	0.37	5.61	1.29	0.88	0.06	0.22	1.90
IVDS3	2.75	2.69	0.06	11.85	0.40	6.71	1.36	0.92	0.06	0.19	2.05
IVDS4	2.89	3.11	-0.22	11.66	0.42	7.63	1.49	0.98	0.07	0.16	2.14
IVDS5	3.18	4.12	-0.94	11.25	0.43	9.08	1.64	1.08	0.07	0.10	2.38

Notes - Table reports time series averages of sorted stock portfolio characteristics. For all stocks, characteristic means (All) and standard deviations (All SD) are reported. For sorted stock portfolios, only means are reported. Characteristics are distant idiosyncratic volatility (IVD), recent idiosyncratic volatility (IVR), the change in idiosyncratic volatility (IVC, defined as IVR-IVD), log of the market value of equity (ME), log of the market-to-book ratio (MB), 1- and 6-month prior returns (PRET1 and PRET6), log of the trailing one-year average of the absolute value of the daily return divided by dollar volume (ILLIQ, see Amihud (2002)), maximum daily return over the last month (MAXRET, see Bali et al. (2011)), market capitalization share, and earnings change contemporaneous with the IVR measurement period (defined as one-year earnings ending in the IVR measurement period less lagged one-year earnings, with the difference scaled by lagged market capitalization). Volatility and returns are reported as a percent. IVD and IVR quintile portfolios are formed by sorting stocks by IVD and IVR (respectively). IVDS quintile portfolios are formed by sequentially sorting stocks into IVR then IVD quintiles. Corresponding IVD quintile portfolios are then aggregated to form IVDS portfolios. IVRS quintile portfolios are formed similarly. Data spans 1966-2012.

Table 3: Idiosyncratic Volatility-Sorted Hedge Portfolio Returns

Panel 1: One-Month IVR-IVD Threshold

	EW Raw	EW Alpha	VW Raw	VW Alpha
IVR	-0.363	-0.480***	-0.683***	-0.844***
	0.278	0.178	0.284	0.174
IVD	-0.042	-0.259	-0.446	-0.737***
	0.306	0.191	0.334	0.195
IVRS	-0.549***	-0.504***	-0.683***	-0.668***
	0.086	0.074	0.121	0.105
IVDS	0.301*	0.096	-0.028	-0.276**
	0.161	0.104	0.186	0.130

Panel 2: Six-Month IVR-IVD Threshold

	EW Raw	EW Alpha	VW Raw	VW Alpha
IVR	-0.212	-0.405**	-0.680**	-0.964***
	0.303	0.190	0.339	0.193
IVD	0.134	-0.121	-0.237	-0.515***
	0.294	0.183	0.313	0.176
IVRS	-0.518***	-0.526***	-0.670***	-0.705***
	0.122	0.097	0.154	0.128
IVDS	0.522***	0.367***	0.280**	0.065
	0.102	0.081	0.127	0.114

Panel 3: Twelve-Month IVR-IVD Threshold

	EW Raw	EW Alpha	VW Raw	VW Alpha
IVR	-0.015	-0.266	-0.490	-0.846***
	0.305	0.187	0.333	0.193
IVD	0.230	-0.029	-0.153	-0.410**
	0.283	0.174	0.301	0.174
IVRS	-0.372**	-0.443***	-0.573***	-0.646***
	0.141	0.102	0.166	0.135
IVDS	0.448***	0.355***	0.338***	0.258**
	0.084	0.076	0.107	0.112

Panel 4: Monthly Measures

	EW Raw	EW Alpha	VW Raw	VW Alpha
EGARCH(i,j)	-0.030	-0.123	-0.010	-0.193
	0.223	0.136	0.241	0.135
IVM	-0.163	-0.250	-0.328	-0.446***
	0.290	0.164	0.296	0.156

Notes - Table reports monthly returns (in percent) of idiosyncratic volatility-sorted hedge portfolios. Portfolios are reformed monthly, from 1966-2012. Equal- and value-weighted raw returns and four-factor alphas (using the three factors of Fama and French (1996) and a one-month formation period momentum portfolio) are reported. Panels 1-3 report mean returns of hedge portfolios formed by sorts on recent historical idiosyncratic volatility (IVR, calculated over the last one, six, or twelve months), distant historical idiosyncratic volatility (IVD, calculated over the year prior to IVR), and sequential sorts on IVR then IVD (IVDS) and IVD then IVR (IVRS). Panel 4 reports mean returns of hedge portfolios formed by a sort on the last five years of idiosyncratic volatility using monthly data (IVM) and out-of-sample predicted volatility from the best fitting EGARCH(i,j) model, where $i, j \leq 3$ (see Fu (2009), Fink, Fink, and He (2012), and Guo, Kassa, and Ferguson (2012)). ***/**/* indicates significance at the 1%/5%/10% level, using a two-sided test.

Table 4: Fama-MacBeth Regressions

Panel 1: One-Month IVR-IVD Threshold

Weight	IVR	IVD	IVC	ME	MB	PRET1	PRET6	ILLIQ	MAXRET
OLS	-0.157***								
	0.045								
OLS	-0.118*		-0.124**						
	0.066		0.064						
OLS	-0.241***	0.124***							
	0.021	0.056							
OLS	-0.110***	0.005		0.071	-0.215***	-0.051***	0.008***	0.185***	-2.747***
	0.024	0.055		0.055	0.044	0.004	0.002	0.048	0.752
WLS	-0.079								
	0.072								
WLS	-0.068		-0.036						
	0.186		0.091						
WLS	-0.104**	0.036							
	0.040	0.091							
WLS	-0.248***	-0.115		-0.092**	-0.088	-0.031***	0.009***	0.027	3.170**
	0.044	0.089		0.055	0.069	0.006	0.003	0.045	1.570

Panel 2: Six-Month IVR-IVD Threshold

Weight	IVR	IVD	IVC	ME	MB	PRET1	PRET6	ILLIQ	MAXRET
OLS	-0.121*								
	0.063								
OLS	-0.066		-0.284***						
	0.068		0.041						
OLS	-0.350***	0.284***							
	0.043	0.042							
OLS	-0.323***	0.255***		0.082	-0.234***	-0.052***	0.009***	0.188***	-2.943***
	0.039	0.043		0.055	0.044	0.004	0.002	0.048	0.622
WLS	-0.089								
	0.098								
WLS	-0.024		-0.262***						
	0.108		0.078						
WLS	-0.286***	0.262***							
	0.079	0.078							
WLS	-0.406***	0.160**		-0.079	-0.098	-0.030***	0.009***	0.028	0.272
	0.069	0.078		0.055	0.069	0.006	0.003	0.044	1.316

Panel 3: Twelve-Month IVR-IVD Threshold

Weight	IVR	IVD	IVC	ME	MB	PRET1	PRET6	ILLIQ	MAXRET
OLS	-0.064								
	0.067								
OLS	-0.023		-0.242***						
	0.068		0.036						
OLS	-0.265***	0.242***							
	0.058	0.036							
OLS	-0.179***	0.202***		0.098*	-0.241***	-0.051***	0.008***	0.182***	-4.409***
	0.052	0.034		0.055	0.044	0.004	0.002	0.048	0.606
WLS	-0.056								
	0.104								
WLS	0.000		-0.269***						
	0.110		0.076						
WLS	-0.269***	0.269***							
	0.096	0.075							
WLS	-0.366***	0.181***		-0.063	-0.097	-0.030***	0.008***	0.036	-0.767
	0.086	0.071		0.055	0.069	0.006	0.003	0.044	1.280

Notes - Table reports Fama-MacBeth regressions of one-month stock returns on lagged stock characteristics. Characteristics are distant idiosyncratic volatility (IVD), recent historical idiosyncratic volatility (IVR), the log of the market value of equity (ME), the log of the market-to-book ratio (MB), 1- and 6-month prior returns (PRET1 and PRET6), a measure of illiquidity (ILLIQ) based on Amihud (2002), and the maximum daily return over the last month (MAXRET, see Bali et al. (2011)). Results are reported using an IVR-IVD threshold of one, six, and twelve months. Parameter point estimates are reported above standard errors. Results are reported for a standard cross-sectional regression (OLS) and a cross-sectional regression where each observation is weighted by market capitalization (WLS). Data is monthly and spans 1966-2011. ***, **, * indicates significance at the 1%, 5%, and 10% level, respectively.

Table 5: Six-Month Returns of Historical Idiosyncratic Volatility-Sorted Hedge Portfolios for Ten Years after Portfolio Formation

Return (Months)	IVR	IVD	IVRS	IVDS	IVR VW
$R_{1,6}$	0.17	1.36	-1.63**	2.38***	-2.36
	1.78	1.75	0.69	0.59	1.67
$R_{7,12}$	2.29	2.38	0.53	1.23**	-0.76
	1.78	1.77	0.61	0.59	1.57
$R_{13,18}$	2.74	2.65	0.89	0.99*	-0.11
	1.78	1.77	0.58	0.58	1.55
$R_{19,24}$	2.66	2.18	1.11*	0.35	0.48
	1.74	1.66	0.60	0.51	1.56
$R_{25,30}$	2.31	1.96	0.87	0.36	0.98
	1.68	1.59	0.54	0.47	1.67
$R_{31,36}$	2.11	1.82	1.04**	0.28	1.03
	1.57	1.53	0.52	0.48	1.51
$R_{37,42}$	1.93	2.01	0.62	0.56	1.45
	1.50	1.54	0.48	0.53	1.52
$R_{43,48}$	2.20	2.25	0.51	0.76	2.13
	1.50	1.55	0.43	0.52	1.63
$R_{49,54}$	2.69*	2.54*	0.73*	0.65	2.28
	1.51	1.54	0.44	0.48	1.60
$R_{55,60}$	2.84*	2.93**	0.73*	0.94*	2.03
	1.47	1.46	0.44	0.50	1.49
$R_{61,66}$	2.80*	2.83**	0.68	0.97**	2.30
	1.45	1.39	0.48	0.44	1.40
$R_{67,72}$	2.85**	3.11**	0.51	1.35***	2.71**
	1.37	1.36	0.42	0.47	1.33
$R_{73,78}$	2.76**	2.61**	0.69	0.97**	2.19*
	1.35	1.34	0.43	0.48	1.28
$R_{79,84}$	3.02**	2.82**	0.95**	0.73	3.00**
	1.32	1.35	0.41	0.48	1.34
$R_{85,90}$	3.03**	2.90**	1.19***	0.87*	3.10**
	1.29	1.31	0.44	0.46	1.39
$R_{91,96}$	3.10**	2.97**	1.12**	0.56	3.33**
	1.29	1.28	0.47	0.44	1.30
$R_{97,102}$	2.97**	3.06**	0.76*	0.67	3.17***
	1.24	1.30	0.40	0.52	1.22
$R_{103,108}$	2.96**	3.13**	0.70*	0.79	3.16**
	1.25	1.31	0.38	0.54	1.19
$R_{109,114}$	3.15**	3.06**	0.95**	0.66	2.90**
	1.26	1.29	0.45	0.55	1.31
$R_{115,120}$	3.06**	2.96**	0.83*	0.92*	2.59**
	1.27	1.25	0.49	0.49	1.30

Notes - Table reports returns (in percent) of equal-weight idiosyncratic volatility-sorted hedge portfolios for months 1-120 subsequent to portfolio formation, in six-month increments. Portfolios are reformed monthly, from 1966-2012. Hedge portfolios are formed by single sorts on recent historical idiosyncratic volatility (IVR, calculated over the last six months), distant historical idiosyncratic volatility (IVD, calculated over the year prior to IVR) and sequential sorts on IVR then IVD (IVDS) and IVD then IVR (IVRS). The IVR VW column reports value-weighted results, with weights determined at month zero. Newey-West standard errors are reported below each mean return. ***/**/* indicates significance at the 1%/5%/10% level, using a two-sided test.

Table 6: Six-Month Returns of Historical Idiosyncratic Volatility-Sorted Hedge Portfolios for Ten Years after Portfolio Formation, Market Model

Return (Months)	IVR	IVD	IVRS	IVDS
$R_{1,6}$	2.78*	3.42**	0.15	1.96***
	1.42	1.33	0.64	0.48
$R_{7,12}$	4.10***	3.65***	1.80***	0.89*
	1.38	1.29	0.62	0.47
$R_{13,18}$	3.95***	3.46***	1.89***	1.04**
	1.33	1.23	0.65	0.45
$R_{19,24}$	4.35***	3.76***	1.90***	0.76*
	1.40	1.24	0.71	0.42
$R_{25,30}$	4.36***	4.10***	1.55**	1.08**
	1.37	1.25	0.64	0.41
$R_{31,36}$	4.00***	3.92***	1.38**	1.26***
	1.27	1.18	0.54	0.42
$R_{37,42}$	3.79***	3.81***	1.00**	1.63***
	1.21	1.18	0.45	0.47
$R_{43,48}$	3.07***	3.08***	0.93**	1.33***
	1.08	1.06	0.42	0.42
$R_{49,54}$	3.25***	2.95***	1.09**	1.17***
	1.10	1.04	0.44	0.39
$R_{55,60}$	3.20***	3.09***	1.19***	1.11***
	1.06	1.02	0.44	0.41
$R_{61,66}$	3.31***	3.04***	1.03**	1.09***
	1.06	0.98	0.46	0.40
$R_{67,72}$	3.26***	3.29***	1.12**	1.27***
	1.04	0.97	0.48	0.41
$R_{73,78}$	3.07***	2.86***	1.01**	0.97**
	0.99	0.93	0.48	0.40
$R_{79,84}$	3.31***	3.15***	1.18**	0.89**
	0.99	0.94	0.51	0.41
$R_{85,90}$	3.10***	2.72***	1.37**	0.80**
	1.02	0.91	0.55	0.40
$R_{91,96}$	3.13***	2.78***	1.47***	0.42
	0.96	0.88	0.49	0.39
$R_{97,102}$	2.97***	2.79***	1.28***	0.49
	0.89	0.86	0.42	0.40
$R_{103,108}$	3.13***	2.93***	1.33***	0.60
	0.89	0.87	0.41	0.39
$R_{109,114}$	3.25***	2.74***	1.55***	0.31
	0.88	0.85	0.39	0.40
$R_{115,120}$	3.09***	2.70***	1.38***	0.23
	0.85	0.84	0.40	0.41

Notes - Table reports returns (in percent) of equal-weighted idiosyncratic volatility-sorted hedge portfolios for months 1-120 subsequent to portfolio formation, in six-month increments. Portfolios are reformed monthly, from 1929-2012. Hedge portfolios are formed by sorts on recent historical idiosyncratic volatility (IVR, calculated over the last six months), distant historical idiosyncratic volatility (IVD, calculated over the year prior to IVR), and sequential sorts on IVR then IVD (IVDS) and IVD then IVR (IVRS). Newey-West standard errors are reported below each mean return. ***/**/* indicates significance at the 1%/5%/10% level, using a two-sided test.

Figure 1: Idiosyncratic Volatility Hedge Portfolio Returns

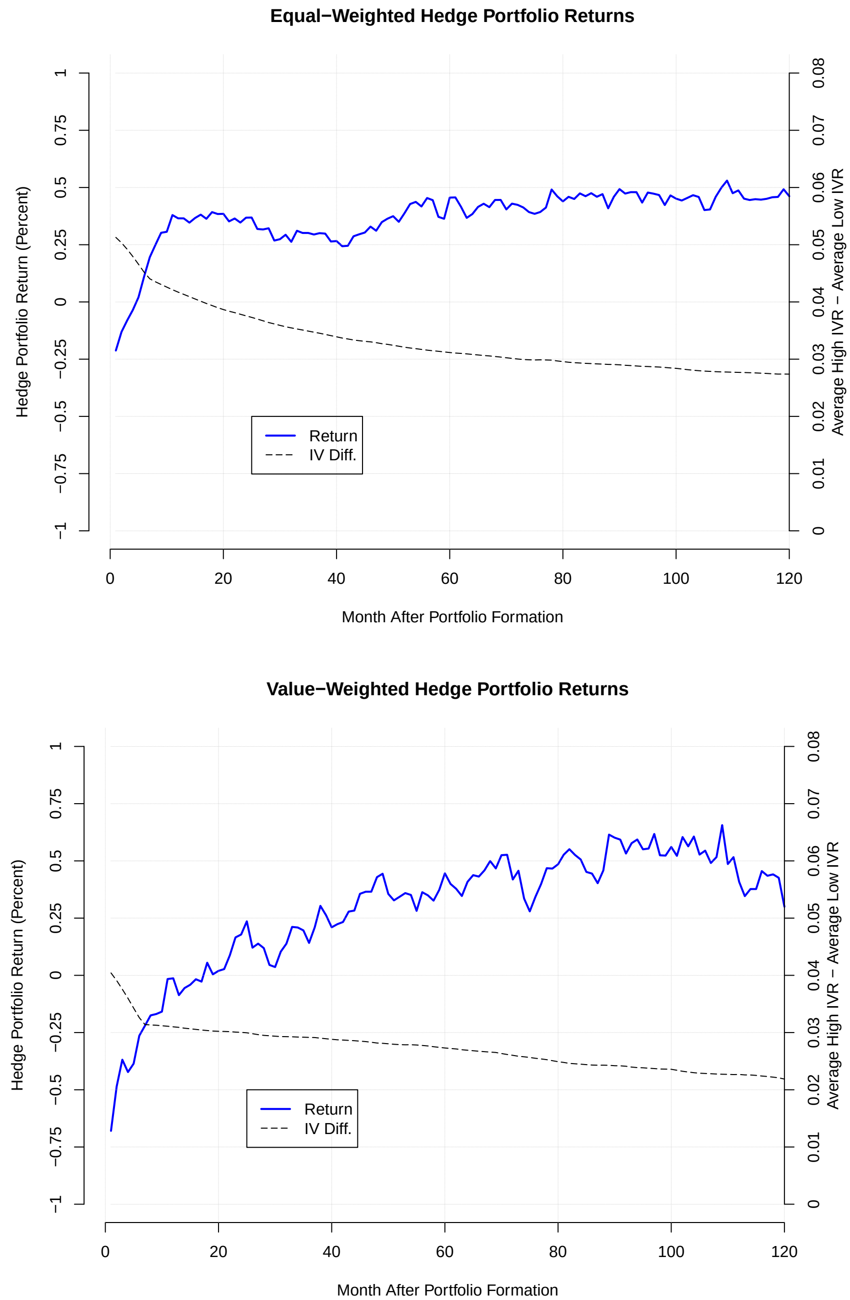

Notes - Figure displays mean returns and dispersion in idiosyncratic volatility of high-low idiosyncratic volatility hedge portfolios, by month after portfolio formation. Idiosyncratic volatility is calculated as six-month historical idiosyncratic volatility (IVR). Portfolios are reformed monthly, from 1966-2012. Value-weighted returns are calculated using weights at month zero. IV difference is the difference in return measurement period IVR across the extreme quintile portfolios.

Figure 2: Forecast and Actual IV, Positive IV Shock

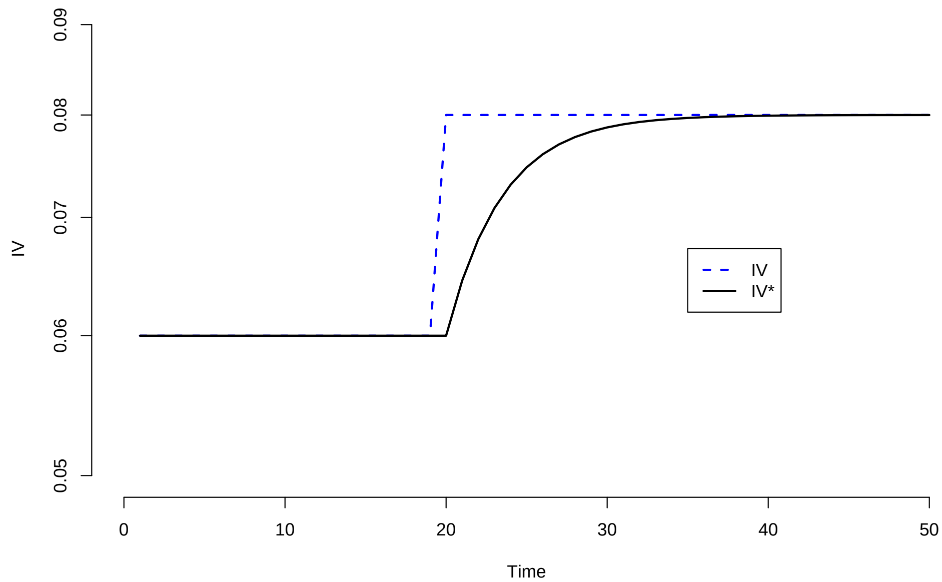

Figure 3: Prices under Forecast and Actual IV, Positive IV Shock

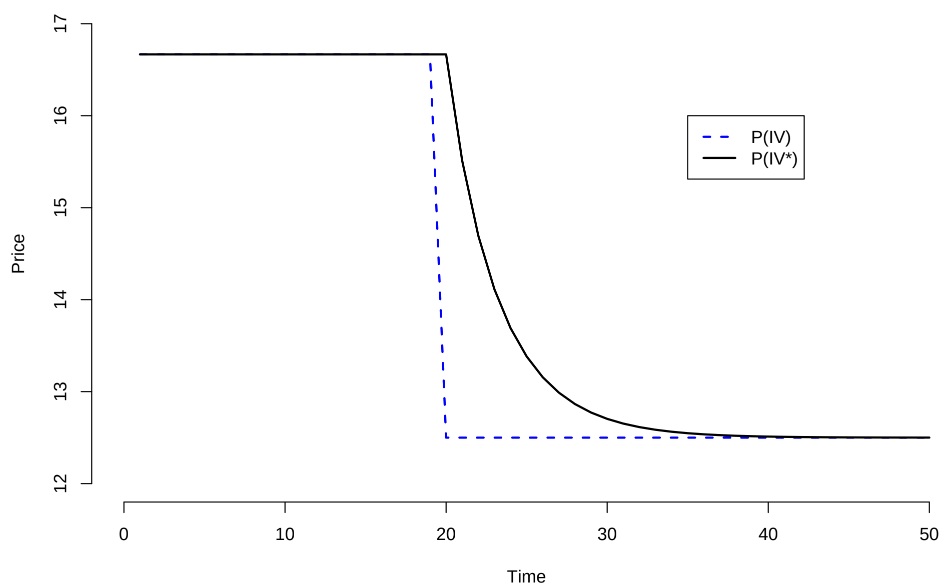

Notes - Figures display actual (IV) and perceived (IV*) idiosyncratic volatility, as well as the associated prices, P(IV) and P(IV*), under the price underreaction model. Θ is set to 0.25, initial idiosyncratic volatility is set to 0.06, and the shock, which occurs at time 20, is 0.02.

Stock Wealth, Consumption, and Return Predictability

Mark Rachwalski

March 19, 2014

Abstract

The stock wealth-consumption ratio reflects expected stock returns and consumption growth. Because consumption growth is mostly unpredictable, much of the variation of this ratio likely reflects changing expected stock returns. In contrast, isolating expected stock return information from other variables may be difficult (in addition to stock returns, the dividend yield may predict dividend growth, while the consumption-wealth ratio may predict non-stock wealth returns). Empirically, a detrended version of this ratio strongly predicts U.S. and international stock returns. In contrast to other predictive variables, predictability does not deteriorate after 1980 and out-of-sample performance is impressive.

1 Introduction

The question of whether and how stock returns are predictable has been at the forefront of financial research for decades. Researchers have identified several variables that likely capture information about expected stock returns, such as the dividend yield (see, for example, Fama and French (1988)) and the consumption-wealth ratio (Lettau and Ludvigson (2001a)). However, commonly-used empirical models of expected stock returns appear to be unstable over time (Paye and Timmermann (2006), Lettau and Van Nieuwerburgh (2008)) and often perform poorly out-of-sample and over the last thirty years (see Goyal and Welch (2008) and the evidence presented in this paper), which calls into question the predictive power and reliability of these models, especially in recent years. Then, such models may be of limited value when addressing important questions that require good estimates of expected stock returns (e.g. determining whether expected stock returns are time varying or how much of return variation can be attributed to variation in expected returns). This paper develops a novel empirical model of expected stock returns that may be useful when addressing such questions. Also, an alternative model of expected stock returns may serve to validate existing models (i.e. if two different models generate similar results, then we should have greater confidence in both models). This is important because expected returns are unobservable, so one cannot directly test expected return models.

For brevity, I refer to the fact that a predictive variable, such as the dividend yield, may predict stock returns and/or something else as "composite predictability". This is not a new concept, and it has long been understood that composite predictability may adversely affect return forecasts. Fama and French (1988) note that the dividend yield is a "noisy proxy for expected returns", as the dividend yield may predict dividend growth as well as stock returns (see also Cochrane (2008)). Composite predictability is also a concern for the consumption-wealth ratio of Lettau and Ludvigson (2001a,b), which may predict non-stock wealth returns, including human capital returns, as well as stock returns. Empirically, I find that the dividend yield predicts dividend growth[1] and that the consumption-wealth ratio predicts non-stock wealth returns. This suggests that a regression of stock returns on the lagged dividend yield or consumption-wealth ratio may serve as an unreliable model of expected stock returns, even if theory connects these variables to expected stock returns.

I develop a new predictive variable of stock returns that may be less influenced by composite predictability. Just as a high dividend yield should predict high stock returns, low dividend growth, or both, a low stock wealth-consumption ratio should predict high stock returns, low consumption growth, or both. However, consumption growth should be mostly unpredictable (see Hall (1978) and Lettau and Ludvigson (2004)).[2] Empirically, this appears to be a reasonable approximation.[3] Then, much of the variation in the stock wealth-consumption ratio should reflect variation in expected stock returns, and a predictive regression based on this ratio may yield a good empirical model of expected stock returns.

The stock wealth-consumption ratio can be viewed as a modification of the dividend yield, where stock wealth is scaled by consumption rather than dividends. Alternatively, the ratio can be interpreted as a stock wealth-specific modification of the consumption-wealth ratio of Lettau and Ludvigson (2001a,b). Intuitively, it should not be surprising that the stock wealth-consumption ratio may be better suited to predict stock returns while the consumption-wealth ratio may be better suited to predict aggregate wealth returns (indeed, the consumption-wealth ratio is a highly appealing predictor of aggregate wealth returns).

[1]Although Cochrane (2008) finds no evidence that the dividend yield predicts dividend growth, other authors find evidence of dividend growth predictability (see Menzly, Santos, and Veronesi (2004), Lettau and Ludvigson (2005), Bansal, Dittmar, and Kiku (2009), Larrain and Yogo (2008), and Chen (2009)). I find strong evidence of dividend growth predictability when using methods that account for structural changes in payout policy (see Fama and French (2001), Grullon and Michaely (2002), Boudoukh, Michaely, Richardson, and Roberts (2007)).

[2] Although there may be some predictability in consumption growth (due to adjustment constraints, habits (Campbell and Cochrane (1999)), or a persistent component of consumption growth (Bansal and Yaron (2004))), this should not have a large impact on return predictability provided that consumption is "close to" a random walk, as appears to be the case.

[3] See Cochrane (1994), Lettau and Ludvigson (2001a), Lettau and Ludvigson (2004) and Lettau and Ludvigson (2005).

Empirically, the observable proxy for the consumption-wealth ratio, CAY, is only modestly correlated with the stochastically detrended stock wealth-consumption ratio used to predict returns in this paper (correlation of -0.22). This suggests that the variables capture unique information about expected stock returns.

Although scaling by consumption, rather than another quantity, likely alleviates some of the adverse effects of composite predictability, the stock wealth-consumption ratio may still vary for reasons unrelated to expected stock returns. For example, changes in stock exchange listing requirements may encourage more firms to list on exchanges, increasing the stock wealth-consumption ratio, while expected returns may be unchanged. However, such changes are likely to affect the ratio at a low frequency (and may have permanent effects), while variation in expected stock returns is likely to operate at a somewhat higher frequency. Exploiting this logic, I adopt a parsimonious empirical framework to model transitory deviations in the ratio that are more likely related to expected stock returns. These deviations, referred to as SDEV, should be less influenced by composite predictability than the raw ratio and are used to predict stock returns. The empirical framework involves estimating a cointegrated relation, similar to that of Lettau and Ludvigson (2001a), and is shown to yield an appealing predictor of stock returns given (1) cointegration of aggregate wealth and consumption, (2) a stock wealth-aggregate wealth ratio bounded by zero and one, and (3) the likely importance of the stock wealth-aggregate wealth ratio when predicting stock returns.

Importantly, theory suggests that SDEV, the consumption-wealth ratio, and the dividend yield all reflect expected stock returns. However, theory offers no guarantee that the variation in any of these predictive variables can be solely attributed to variation in expected stock returns. For example, the cointegration framework of Lettau and Ludvigson (2001a) does not imply a one-to-one relation between expected stock returns and CAY, which is the specification used to predict stock returns. Instead, the framework relates aggregate wealth and consumption. Then, one cannot assert, based on theory alone, that the dividend yield, CAY, or SDEV should be preferred when constructing an empirical model of expected stock returns. While SDEV relies upon a long-run/transitory decomposition to empirically isolate expected stock returns, a standard predictive regression based on the dividend yield or CAY relies upon the assumption that controlling for expected dividend growth or expected non-stock wealth returns is not important. None of these assumptions are addressed by theory. Empirically, I find that the dividend yield predicts dividend growth and that CAY predicts non-stock wealth returns, which calls into question the validity of the latter assumptions.

Empirically, SDEV is an excellent predictor of U.S. stock returns. In-sample predictability associated with SDEV is at least as strong as other stock price-scaled variables, such as the dividend yield. In a post-war (1952-2012) sample, SDEV and CAY are both strong predictors of stock returns, although the predictability of the two variables appears to be largely complementary. SDEV continues to perform well in the more recent 1982-2012 subsample, where other variables, including CAY, exhibit deteriorating and/or no significant predictability. Also, in a sample of six non-U.S. countries (Australia, Canada, Japan, Germany, France, and the U.K.), there is widespread evidence that country-specific SDEV is a good predictor of national stock returns.

Out-of-sample tests indicate that SDEV could have been used by an investor with only ex ante information to predict U.S. stock returns. The increase in out-of-sample predictive accuracy associated with SDEV is often quite large (at 3-5 year horizons, 25% or more compared to forecasts based on historical mean returns). This sort of out-of-sample performance is unusual. Goyal and Welch (2008) show that, when using popular predictive variables and methods, stock returns are generally not predictable out-of-sample, even when in-sample methods suggest substantial predictability. Also, SDEV's out-of-sample predictability is consistent with in-sample predictability, in that the estimated in-sample model could easily have generated the observed out-of-sample predictability. In contrast, for several other predictive variables, including the dividend yield and CAY, large differences between in- and out-of-sample predictability are not likely explained by sampling error.

Recently, researchers have considered other approaches to addressing composite predictability. Cochrane (2008) and Binsbergen and Koijen (2010) jointly model the dividend yield, expected stock returns, and expected dividend growth. Kelly and Pruitt (2012) use cross-sectional book-to-market information that may allow one to isolate expected return and dividend growth information. These approaches attempt to decompose the information content of predictive variables that are likely affected by composite predictability, and require a specification that can correctly separate expected return and dividend growth information. In contrast, the approach taken in this paper is to construct an alternative predictive variable that should, for the most part, predict only stock returns.

This paper contains four main contributions. First, I highlight and empirically document the challenge posed by composite predictability. Importantly, theory used to motivate the dividend yield and CAY does *not* address the degree to which composite predictable influences the associated return forecasts. Therefore, theory offers no guarantee that these variables will yield good empirical models of expected stock returns (i.e. estimated expected stock returns may substantially differ from actual expected stock returns). The empirical evidence presented in this paper suggests that controlling for expected dividend growth or non-stock wealth returns may be a critical part of predicting stock returns with the dividend yield or CAY.

Second, I develop a predictive variable of stock returns, SDEV, that may be less influenced by composite predictability. Just like the dividend yield and CAY, SDEV will reflect expected stock returns. Because theory offers no guarantee that any of these variables will yield a good empirical model of expected stock returns, there is little reason, a priori, to prefer one of these variables over another. Ultimately, determining whether a model is accurately fitting expected stock returns, and can be used to forecast future returns, is an empirical question. In contrast to many other predictive variables, SDEV is an empirically potent and stable predictor of stock returns, and could have been used by investors to predict returns out-of-sample. This is important because the poor out-of-sample and recent in-sample performance of many predictive variables suggest flawed or unstable empirical models of expected stock returns (which may not perform well in the future even if they have performed well in the past) and unreliable in-sample expected return estimates.

Third, I find widespread evidence of international return predictability using the same methods as in the U.S. sample. This is important because most predictability research has focused on U.S. returns, so international predictability is less well understood.[4]

Finally, the results can be used to understand the dynamics of discount rates and the nature of predictability. In particular, the results suggest that (1) stock returns are highly predictable, especially over longer horizons; (2) expected stock returns are quite persistent; (3) stock and aggregate wealth discount rates exhibit meaningful independent variation that is especially apparent in the recent financial crisis; and (4) cross-country expected stock returns are generally highly correlated.

The paper proceeds as follows. Section 2 discusses composite predictability and develops the stock wealth-consumption ratio as a predictive variable. Section 3 examines predictive regressions of stock returns and related results. Section 4 examines out-of-sample forecasts of stock returns. Section 5 examines the international evidence. Section 6 concludes. An appendix examines predictability associated with the dividend yield and consumption-wealth ratio.

[4]However, see Paye and Timmermann (2006), Ang and Bekaert (2007) and Hjalmarsson (2010) for recent studies on international predictability.

2 Predicting Stock Returns with the Stock Wealth-Consumption Ratio

In this section, I discuss the challenge posed by composite predictability. Next, I develop an empirical model of expected stock returns based on the stock wealth-consumption ratio. This model relies on alternative assumptions to address composite predictability and identify variation in expected stock returns. I also discuss similarities and differences between the stock wealth-consumption ratio and consumption-wealth ratio of Lettau and Ludvigson (2001a,b). Intuitively, the variables differ because the stock wealth-consumption ratio is more like a traditional price-scaled variable (e.g. the dividend yield) and is focused on stock wealth, not aggregate wealth.

2.1 Composite Predictability and the Dividend Yield

The present value decomposition of Campbell and Shiller (1988), written here omitting the constant term,

$$d_t - p_t = E_t \sum_{i=1}^{\infty} \rho^j (r_{t+i} - \Delta d_{t+i}), \tag{1}$$

implies that variation in the dividend-price ratio must reflect variation in expected stock returns or expected dividend growth (or both). Many authors have used this decomposition to motivate the use of the dividend-price ratio in a standard predictive regression of stock returns:

$$r_{t+1} = \alpha + \beta(d_t - p_t) + \epsilon_{t+1}. \tag{2}$$

However, Equation 1 does not imply that Equation 2 will be a good empirical model of expected returns. Rather, this requires omitting the expected dividend growth term from Equation 1. If expected dividend growth is relevant, then Equation 2 is missing a potentially important control. This highlights the difference between asserting that a predictive variable captures information about expected stock returns and that a standard predictive regression, based on that variable, will serve as a good empirical model of expected returns. Importantly, the former does not imply the latter.

Cochrane (2008) stresses the relationship between the dividend yield, expected dividend growth, and expected returns. Cochrane notes that regressions of dividend growth on the lagged dividend yield provide no evidence of dividend growth predictability, which suggests that variation in the dividend yield should primarily reflect variation in expected stock returns. I revisit dividend growth predictability in the appendix and find that dividend growth is predictable based on the dividend yield, especially when using methods that account for structural changes in payout policy, such as a shift from traditional dividends to repurchases. This is consistent with recent research that finds that dividend growth is predictable (Menzly, Santos, and Veronesi (2004), Lettau and Ludvigson (2005), Bansal, Dittmar, and Kiku (2009), Chen (2009), and Larrain and Yogo (2008)). These results suggest that Equation 2 may yield a poor empirical model of expected returns even if the dividend yield captures important information about expected returns.

From a practical perspective, a standard predictive regression of stock returns on the dividend yield implies sharp historical declines in the equity premium. For example, in a regression of annual excess stock returns on the dividend yield, the average fitted value from 1930-1950 is 9.02%, while the average fitted value from 1992-2012 is 1.43%. Such a decline in the equity premium seems unlikely. More likely, the historical decline in the dividend yield is due to shifts in payout policy (e.g. a shift from dividends to repurchases). This suggests that the use of the traditional price-dividend ratio, as in Equation 2, yields a poor empirical model of expected stock returns. A potential solution is to use a more comprehensive measure of cash flows from firms to investors (e.g. the payout yield, as in Boudoukh, Michaely, Richardson, and Roberts (2007)). However, I find that predictability associated with the payout yield is unstable; there is no evidence of payout yield-based predictability in the post-war sample (discussed later in the paper). One potential explanation of this result is that Equation 2 serves as a poor empirical model of expected returns, even when using comprehensive measures of firm payouts, because there is no control for expected payout growth. Implementing such a control may be quite difficult because expected payout growth may be conditional and/or exhibit structural breaks. This paper explores an alternative predictive variable that requires different controls (which may be easier to implement).

2.2 Composite Predictability and the Consumption-Wealth Ratio

The consumption-wealth ratio of Lettau and Ludvigson (2001a,b) may also be influenced by composite predictability. Equation 3 of Lettau and Ludvigson (2001a) (which can be derived by manipulating the intertemporal budget constraint and ruling out asset bubbles),

$$c_t - w_t = E_t \sum_{i=1}^{\infty} \rho^i (r_{w,t+i} - \Delta c_{t+i}), \tag{3}$$

says that a high consumption-wealth ratio ($c - w$) indicates high expected returns to aggregate wealth (r_w) or low expected consumption growth (Δc). Because stocks are part of aggregate wealth, r_w is likely correlated with expected stock returns. In this way, $c - w$ may be useful when predicting stock returns. However, it remains possible that variation in the consumption-wealth ratio does not reflect variation in expected stock returns, but instead reflects variation in expected consumption growth and/or expected non-stock wealth returns (i.e. composite predictability may be important). Therefore, Equation 3 provides no guarantee that a standard predictive regression of stock returns on the lagged consumption-wealth ratio will serve as a good empirical model of expected stock returns.

In the appendix, I show that CAY, the observable proxy for the consumption-wealth ratio used by Lettau and Ludvigson (2001a), predicts non-stock wealth returns. This suggests that controlling for expected non-stock wealth returns may be an important part of predicting stock returns with CAY, in the same way that controlling for expected dividend growth may be an important part of predicting stock returns with the dividend yield. This is entirely consistent with the theory of Lettau and Ludvigson (2001a), as the intertemporal budget constraint (Equation 3) does not specify whether CAY should predict stock returns, non-stock wealth returns, or consumption growth. Additionally, the degree to which CAY predicts these quantities may vary over time. Then, theory offers little reason to expect that a regression of stock returns on CAY will serve as a stable model of expected stock returns.

2.3 The Stock Wealth-Consumption Ratio

Next, I develop a predictive variable of stock returns that may be less influenced by composite predictability. I start by decomposing aggregate wealth (W) into stock market wealth (S) and other wealth (O). Then,

$$W = S + O = E[\sum_t \frac{d_{s,t}}{(1+r_s)^t}] + E[\sum_t \frac{d_{o,t}}{(1+r_o)^t}] \tag{4}$$

where $d_{s,t}$ is the period t stock dividend, $d_{o,t}$ is the period t other wealth dividend, and r_s and r_o are the respective discount rates. It is straightforward to show that

$$\frac{\partial(S/W)}{\partial r_s} < 0. \tag{5}$$

Ceteris paribus, if the equity discount rate increases, then the equity share of aggregate wealth will decline. Therefore, the stock wealth-aggregate wealth ratio captures information about expected stock returns.[5] Holding other variables constant, a large stock wealth-aggregate wealth ratio indicates low expected stock returns.

[5]If r_s and r_o are correlated in a particular way, or equal, then S/W could be constant and not useful when predicting stock returns. However, there is little reason to believe, a priori, that S/W is constant.

Aggregate wealth is not observable. However, because consumption is observable and cointegrated with aggregate wealth[6], the stock wealth-consumption ratio is a plausible and potentially useful proxy for expected stock returns.[7] Because consumption growth is largely unpredictable (see Footnotes 2 and 3), little of the variation in the stock wealth-consumption ratio is likely attributable to variation in expected consumption growth. Therefore, the above analysis, combined with largely unpredictable consumption growth, suggests that much of the variation in the stock wealth-consumption ratio is attributable to variation in expected stock returns.

A similar analysis could be used to motivate other predictive variables of stock returns. For example, one could construct a proxy for aggregate wealth and the stock wealth-aggregate wealth ratio, and use this ratio to predict stock returns. Alternatively, one could predict returns with the stock wealth-GDP ratio (this is an intuitively appealing predictive variable, as it measures the stock market "share" of total output). However, transitory deviations in GDP and non-stock aggregate wealth returns appear to be predictable (see Cochrane (1994) and the evidence presented in this paper, respectively), while consumption appears to follow a random walk. This makes the stock wealth-consumption ratio a more appealing predictive variable of stock returns, as variation in this ratio is more likely to mostly reflect variation in expected stock returns. The other ratios may be substantially influenced by transitory deviations in GDP growth or expected non-stock wealth returns. For this reason, the remainder of the paper focuses on the stock wealth-consumption ratio as a predictor of stock returns. In untabulated results, I find that the stock wealth-GDP and a proxy for the stock wealth-aggregate wealth ratio can also be used to predict stock returns

[6]Models of intertemporal consumption allocation generally imply that consumption and aggregate wealth are cointegrated. Consumption will ultimately be limited by aggregate wealth. The assumption that consumers are never satiated ensures that all wealth is eventually consumed. Cochrane (1994) presents a formal treatment, showing that consumption and aggregate wealth are cointegrated provided that income growth is stationary. Also, see Lettau and Ludvigson (2001a).

[7]The stock wealth-aggregate wealth ratio is bounded by zero and one. Because consumption and aggregate wealth are cointegrated, the stock wealth-consumption ratio will also be nonexplosive.

(following the methods of this paper). However, as suggested by the above discussion, this predictability is not as strong as that of the stock wealth-consumption ratio.

Equation 5 indicates that the stock wealth-consumption ratio is likely correlated with expected stock returns, but does not imply that the stock wealth-consumption ratio will yield a good empirical model of expected stock returns. Similarly, Equations 1 and 3 do not imply that the dividend yield and consumption-wealth ratio will yield such a model. Ultimately, theory alone does not tell us which of these variables can be used to construct a good empirical model of expected stock returns. This is an empirical question.

2.4 Comparing the Consumption-Wealth and Stock Wealth-Consumption Ratios

Similarities and differences between the stock wealth-consumption and consumption-wealth ratios can be illustrated by adding s_t (log stock wealth) to both sides of Equation 3 and rearranging,

$$s_t - c_t = s_t - w_t - E_t \sum_{i=1}^{\infty} \rho^i (r_{w,t+i} - \Delta c_{t+i}). \tag{6}$$

The $s_t - w_t$ term on the right-hand side of this equation is likely correlated with expected stock returns (see Equations 4 and 5). Then, the stock wealth-consumption ratio will capture information about expected stock returns contained in $s_t - w_t$ and the infinite sum of aggregate wealth returns. In contrast, the consumption-wealth ratio does not directly reflect $s_t - w_t$ (see Equation 3). Therefore, the stock wealth-consumption ratio captures information about expected stock returns that is not directly captured by the consumption-wealth ratio.

One apparent advantage of the stock wealth-consumption ratio over CAY is that no aggregate wealth proxy is required. Equity wealth can be precisely estimated using CRSP data (provided equity wealth is defined as publicly-traded equity wealth). In contrast, aggregate wealth is unobservable. Lettau and Ludvigson (2001a) use an aggregate wealth proxy composed of labor income and financial wealth. Although this proxy is plausible, determining the quality of the proxy is ultimately quite difficult, if not impossible, because aggregate wealth is unobservable. For this reason, it is worthwhile to examine S/C as a predictor of stock returns even if one believes that S/W is constant. Empirically, CAY and stochastically detrended S/C (used to predict stock returns) have a correlation of -0.22. This modest correlation suggests that the variables capture unique information, and is consistent with substantial time variation in S/W, aggregate wealth measurement errors, or both.

2.5 Modeling Transitory Deviations in the Stock Wealth-Consumption Ratio

Although Equation 6 might appear to suggest that one should use the stock wealth-consumption ratio to predict returns, this assumes that other variation in the ratio is not important. For example, a change in technology or stock exchange listing requirements may cause some private firms to go public. Then, the stock wealth-consumption ratio will increase, although expected stock returns may remain unchanged. In Equation 6, this can be understood as a change in S/W that is not related to discount rates. When predicting stock returns, controls should be implemented for such variation in S/W.

To implement such a control, I start by recognizing that many such effects are likely to operate at a low frequency, and may be permanent. For example, changes in listing requirements are not likely to be important in explaining monthly variation in the ratio because only a small share of aggregate market capitalization lists or delists in a given month. Instead, these effects are likely to cause gradual changes in the ratio (although the cumulative effect of these changes over a long period of time could be quite large).

To capture this intuition, I adopt the following empirical framework to model the long-run stock wealth-consumption relation:

$$s_t = \alpha + \beta c_t + SDEV_t. \tag{7}$$

$SDEV$ is defined as the residual in this relation and can be interpreted as the transitory deviation in the stock wealth-consumption relation.[8] This framework allows one to separately focus on long-run variation in the stock wealth consumption relation (the predicted values from Equation 7) and the transitory component ($SDEV$). To the degree that variation in expected returns tends to be transitory in nature, $SDEV$ will serve as a less noisy expected return proxy than the raw stock wealth-consumption ratio.

Using annual data, I estimate[9] the stock wealth-consumption relation (with standard errors[10] below) as

$$\begin{array}{cccc} s_t = & 12.43+ & 1.53c_t & +SDEV_t. \\ & (0.35) & (0.13) & \end{array} \tag{8}$$

Not surprisingly, consumption and stock wealth are positively related over this time period. The estimated slope coefficient is 1.53. This implies that, over the historical sample (1929-2012), log stock wealth grows by more than one-to-one with log consumption and the mean of the stock wealth-consumption ratio varies with the level of consumption (which generally

[8]The stock wealth-aggregate wealth ratio is bounded by zero and one. Therefore, because consumption and aggregate wealth are cointegrated, the stock wealth-consumption ratio will be non-explosive. Provided Equation 7 models changes in the mean of the stock wealth-consumption ratio, $SDEV$ will be stationary.

[9]Because stock wealth and consumption are both integrated variables, this relation is estimated using dynamic least squares (Stock and Watson (1993)). In addition to contemporaneous levels, a DLS specification includes leads and lags of the differenced explanatory variable. One lead and lag are used, although results are not sensitive to the number of leads and lags.

[10]Although cointegration tests are inconclusive (narrowly missing rejecting no stock wealth-consumption cointegration), standard errors are estimated under the assumption that $SDEV_t$ is stationary, following the procedure of Li and Maddala (1997). This is consistent with a non-explosive stock wealth-consumption relation (implied by a bounded stock wealth-aggregate wealth ratio and cointegration of consumption and aggregate wealth) and analogous to assuming that scaled-price ratios, such as the dividend yield, are stationary. I estimate standard errors by assuming the data follows the cointegrated system $s_t = \beta c_t + \eta_t$, $c_{t+1} = c_t + v_t$, and $\eta_{t+1} = \rho\eta_t + \epsilon_t$. β and ρ are set equal to their sample estimates. I then resample (v_t, ϵ_t), and form an empirical distribution for the parameters of Equation 7. Use of the sieve bootstrap of Chang, Park, and Song (2006) yields similar results.

increases with time). This may be explained by changing technology, financial reporting, or listing requirements affecting the number of firms that choose to list on stock exchanges.

I reject the hypothesis that the consumption slope coefficient is equal to one (equivalently, I reject a cointegration vector of (-1, 1) for s and c). This is consistent with Equation 6, which allows for such variation in the stock wealth-aggregate wealth ratio. Provided expected returns are stationary (or can be assumed to be stationary for modeling purposes)[11], this suggests that a predictive regression of stock returns on the raw stock wealth-consumption ratio is not likely to serve as a good empirical model of expected returns. However, the transitory component of the ratio may still be useful when predicting the transitory component of expected stock returns.

Controlling for long-run variation in the stock wealth-consumption ratio appears to be an important part of predicting returns with the ratio. Similarly, use of the dividend yield and consumption-wealth ratio to predict returns may require adequately controlling for expected dividend growth and non-stock wealth returns (respectively). The empirical success of these predictive variables may hinge on the ability of researchers to implement adequate controls. Ultimately, determining the effectiveness of these controls is an empirical question. For this reason, it is difficult to say, a priori, that one of the variables should be preferred when constructing an empirical model of expected stock returns.

$SDEV$ measures transitory deviations in the stock wealth-consumption ratio, and likely captures information about expected stock returns. The standardized $SDEV$ time series is plotted in Figure 1. This plot largely conforms with reasonable intuition about historical expected stock returns (i.e. expected stock returns were low in 1929 and 1999, and high in the 1940s and 1980s). Figure 1 indicates that $SDEV$ is quite persistent ($SDEV$'s annual autocorrelation is 0.79). Provided variation in $SDEV$ primarily reflects variation in expected stock returns, this suggests that expected stock returns are also persistent.

[11] Expected stock returns may have trended higher or lower over the historical sample. The approach taken in this paper will not identify such variation, but may still be useful in modeling transitory variation.

Figure 1 suggests that $SDEV$ is approximately equally volatile from 1929-1969 and 1970-2012 (the respective standard deviations are both 0.99). This may be somewhat surprising as both consumption growth and stock returns were more volatile during the Great Depression. However, it is important to remember that $SDEV$ does not measure changes. Rather, $SDEV$ measures the degree to which stock prices are high or low relative to consumption. The most extreme $SDEV$ observation occurs late in the sample, at the height of the "tech bubble" of the 1990s.

2.6 Consumption Measurement Errors

Because $SDEV$ compares the level of stock wealth and contemporaneous consumption, and consumption is far less volatile than stock wealth, the level of stock wealth, not consumption, tends to drive variation in $SDEV$. This suggests that consumption measurement errors are likely unimportant. For example, suppose that consumption is mismeasured by 3% in a particular year. If consumption is accurately measured in the next year, annual consumption growth will be mismeasured by 3%. The standard deviation of annual real consumption growth from 1929-2012 is 3%, so a 3% measurement error corresponds to one standard deviation, a large effect. However, a 3% consumption measurement error corresponds to only 0.15 $SDEV$ standard deviations. Therefore, $SDEV$ is not likely to be meaningfully affected by consumption measurement errors. Also, trends in consumption measurement errors should be captured by the long-run relation, and are not likely to meaningfully affect $SDEV$ (i.e. such trends will be reflected in the predicted values from Equation 7, but are less likely to affect the residual).

3 Predictive Regressions of Stock Returns

3.1 Annual Predictive Regressions

Predictive regressions of annual excess stock index returns are presented in Table 1. Results are reported for three horizons (one, three, and five years). Results are reported for the full sample (1929-2012), as well as a post-war (1952-2012) and recent (1982-2012) subsample. I report Newey-West standard errors (Newey and West (1987)) and the associated p-value. However, because the Newey-West standard error is an asymptotic standard error and finite sample considerations may be important (Nelson and Kim (1993), Stambaugh (1999), Valkanov (2003)), p-values based on a parametric bootstrap are also reported.[12]

Results are reported for several price-scaled variables related to the dividend yield. The dividend-price ratio is the ratio of dividends to stock wealth. The payout-price ratio replaces dividends with payouts, defined as dividends less the change in aggregate market capitalization due to stock issuances and repurchases (see Boudoukh, Michaely, Richardson, and Roberts (2007)). The adjusted price-dividend ratio is the residual from a regression of stock wealth on contemporaneous dividends, similar to SDEV. This variable can be interpreted as a measure of transitory deviations in the dividend yield, expunged of long-run variation that may be related to firms' shifting payout policy. See the appendix for additional information about these variables. Figures 2 and 3 plot the historical dividend yield, payout yield, and the detrended dividend yield.

In Table 1, the dividend-price, payout-price, and adjusted price-dividend ratio all show some ability to predict stock returns. For example, the dividend yield is sometimes a statistically significant predictor of stock returns (e.g. one-year returns in the post-war

[12]The bootstrap employed depends on the predictive variable. When the bootstrap must account for estimation of a long-run relationship (i.e. the stock wealth-consumption and adjusted price-dividend ratio), I proceed under the assumption that the first-stage variables are cointegrated (see Footnote 10). In this case, the bootstrapped standard errors take into account the uncertainty associated with estimating the cointegrated relation. When there is no long-run relationship to estimate, an AR(1) parametric bootstrap similar to Nelson and Kim (1993) and Goyal and Welch (2008) is used.

sample), although not in other samples (one-year returns in the full sample, where the dividend-price ratio just misses marginal significance with a p-value of 0.129). Perhaps more importantly from the perspective of developing a practical and reliable model of expected returns, the point estimates associated with the dividend yield are often quite different across subsamples. For example, the five-year beta estimates across the subsamples are 0.207, 0.129, and 0.098 (see Boudoukh, Michaely, Richardson, and Roberts (2007) and Lettau and Van Nieuwerburgh (2008) for evidence that shifts in the dividend-price ratio affect return predictability). The fitted values associated with a standard dividend-price ratio predictive regression appear to trend lower over time (see Figure 4, which plots fitted values from the one-year regressions of Table 1).

Somewhat surprisingly, use of the payout-price ratio, rather than the dividend-price ratio, results in less convincing evidence of stock return predictability. Although the payout-price ratio performs well in the full sample, there is no evidence of predictability in the post-war (1952-2012) or recent (1982-2012) subsample. The strong payout-price ratio predictability in the full sample appears to be attributable to a few early data points. As can be seen in Figure 3, the payout-price ratio was much more volatile around 1930 than throughout the rest of the sample. Omitting these early data points results in no significant return predictability. Therefore, it is difficult to conclude that predictability associated with the payout-price ratio is robust across time, or that any past predictability is likely to hold in the future.

Overall, there is certainly some evidence that the dividend-price, payout-price, and adjusted price-dividend ratio predict returns. However, this predictability appears to be unstable, as the predictability evidence is quite strong in some subsamples but nonexistent in others. It seems difficult to conclude, despite some evidence of predictability, that any of these variables can be used to create a reliable empirical model of expected returns. One limitation of all of these predictive variables is that variation in these ratios may reflect variation in expected dividend (or payout) growth, rather than variation in expected returns. This is supported by evidence presented in the appendix, which demonstrates dividend and payout growth predictability. Then, an empirical model of expected stock returns should control for expected dividend or payout growth.

Table 1 also presents the results of predictive regressions of stock returns on SDEV.[13] In the full sample (1929-2012), SDEV is a statistically significant predictor of stock returns at all horizons when using either Newey-West standard errors or the bootstrap. The SDEV point estimates are economically large and sensible. For example, at the one-year horizon, a one standard deviation increase in SDEV (equivalent to a 32% log return) is associated with a decrease of 7.6% in expected excess log returns. At the five-year horizon, a one standard deviation increase in SDEV is associated with a decrease of 20.6% in expected excess log returns. The amount of predictability, as measured by $\bar{R}^2$, increases with the forecast horizon. This is consistent with persistent expected returns explaining more, on average, of total returns over longer forecast horizons. At the three- and five-year horizons, stock returns appear to be highly predictable ($\bar{R}^2$ of 28.2% and 37.6%, respectively).[14]

Statistical significance becomes more difficult to detect in the subsamples. This is to be expected because expected returns are likely a small part of total returns, so that predictive regressions often exhibit modest power. However, SDEV remains at least marginally significant in every regression in the 1952-2012 subsample. Power is further reduced in the 1982-2012 subsample, where SDEV is always significant when using the asymptotic standard errors although not when using the bootstrapped p-values (although the bootstrapped

[13] I use seasonally-adjusted consumption data, which may induce a look-ahead bias as the full sample may be used to generate the seasonal adjustments. However, this is likely to be of little importance as changes in stock wealth, rather than consumption, drive variation in SDEV (See Section 2.6). Similarly, use of revised, rather than real-time, consumption data is not likely to substantially alter the results.

[14] As noted by Boudoukh, Richardson, and Whitelaw (2008), sampling error and persistent regressors can lead to highly correlated parameter estimates across forecast horizons, as well as a pattern of increasing $\bar{R}^2$. Therefore, high $\bar{R}^2$ in a long-horizon regression does not, by itself, constitute evidence of return predictability. Rather, this requires properly accounting for the substantial biases in a long-horizon regression, and rejecting sampling error as an explanation.

p-values are at least close to marginal significance, with p-values of 0.129, 0.107, and 0.067 for one-, three-, and five-year returns respectively). To increase power, I examine quarterly data later in the paper and find that SDEV is a significant predictor of returns in the recent subsample.

Across all subsamples and forecast horizons, SDEV always offers more explanatory power, as measured by $\bar{R}^2$, than the other price-scaled variables (the dividend-price, payout-price, and adjusted price-dividend ratios). Also, the absolute value of the SDEV parameter estimate is, with one exception (five-year returns in the full sample), larger than that of the other scaled-price variables. Because all predictive variables are standardized, this suggests that SDEV is a relatively important predictor of stock returns. Finally, the SDEV parameter estimates seem to be reasonably stable across samples. For example, the one-year SDEV parameter estimates are -0.076, -0.058, and -0.062 across the full, post-war, and recent subsamples, respectively. Overall, these results suggest that SDEV captures important information about expected stock returns and that an SDEV-based predictive regression is an appealing way to construct an empirical model of expected stock returns.[15]

Table 1 also presents the results of univariate predictive regressions using the risk-free rate (see Fama and Schwert (1977) and Ang and Bekaert (2007)) and the consumption-wealth ratio (CAY, see Lettau and Ludvigson (2001a,b))[16]. The risk-free rate exhibits modest predictive power. Consistent with prior studies, CAY often exhibits very strong predictive power. The predictive ability of CAY and SDEV appears to be complementary. For example, in the 1952-2012 sample, CAY and SDEV offer similar explanatory power in univariate regressions ($\bar{R}^2$ of 8.1% and 9.6% respectively), which seems to be largely cumulative in the multiple regression (the $\bar{R}^2$ of the SDEV, CAY, and RF multiple regression

[15] In untabulated regressions, I find that the transitory deviation in the long-run stock wealth-GDP relation also predicts stock returns, although results are slightly weaker. This is consistent with cointegrated consumption and GDP and more predictable GDP growth than consumption growth. See Cochrane (1994).

[16] Following Lettau and Ludvigson (2001a), I assume that the cointegrating vector for CAY is known when calculating standard errors.

is 23.8%).[17] However, CAY's predictive power is reduced in the recent subsample, especially when using one-year returns. I focus on one-year returns because inference derived from overlapping returns from such a short subsample may be unreliable, and because overlapping returns underweight the most recent data (e.g. the 2012 return is part of one five-year return, while the 2000 return is part of five five-year returns). Using annual returns from 1982-2012, CAY is not a statistically significant predictor of stock returns (the empirical p-value is 0.216). One explanation for this reduced predictability is low power (i.e. sampling error is large in a small sample). However, this does not appear to be a complete explanation because SDEV exhibits reasonably strong evidence of predictability in this setting. The explanatory power of SDEV ($\bar{R}^2$ of 0.104) is substantially larger than that of CAY ($\bar{R}^2$ of 0.010) when examining annual returns over the recent subsample. A multiple regression with SDEV, the risk-free rate, and CAY has a $\bar{R}^2$ (0.092) lower than that of the univariate SDEV regression.

3.2 Quarterly Predictive Regressions

To increase power, I examine quarterly regressions in Table 2.[18] Results are similar to the annual regressions. SDEV is a statistically significant predictor of stock returns in both the post-war and recent subsamples. SDEV has higher explanatory power than the other price-scaled variables, as measured by $\bar{R}^2$. SDEV and CAY exhibit similar explanatory power in the post-war sample, and the multiple regression again suggests that this predictability is complementary. However, in the recent sample, SDEV continues to predict stock returns,

[17] One may be concerned that the $\bar{R}^2$ of some of the multiple regressions are "too high". Although it is difficult to know, a priori, how much of return variation should be attributable to expected return variation, a $\bar{R}^2$ of 23.8% at an annual horizon does seem quite high. This is one reason why it is important to search for stable models of expected returns; stability provides some assurance that a model is not just fitting unexpected return variation ex post. Also, $\bar{R}^2$ is a point estimate. The sampling distribution of $\bar{R}^2$ under the null that returns are modestly predictable can be quite wide, and consistent with the results reported here.

[18] Because quarterly consumption data is not available until the late 1940s, the bootstrap in this section uses annual data to estimate the long-run stock wealth-consumption relationship, but quarterly data to predict returns.

while CAY does not. In this subsample, CAY is not significant when using either the asymptotic or empirical p-values, and offers little explanatory power ($\bar{R}^2$ of -0.003). In contrast, SDEV significantly predicts stock returns (empirical p-value of 0.021) and offers reasonable quarterly explanatory power ($\bar{R}^2$ of 0.036).

3.3 Recent Divergence

It is interesting to note that the fitted values from the predictive regressions, plotted in Figure 4, suggest a recent divergence in expected returns. All of the fitted values exhibit a spike in 2009, suggesting that the poor stock market performance of 2008 was associated with higher future expected returns. However, the 2009 fitted value associated with CAY is approximately equal to the time series average, and is slightly below the fitted value from 2000, at the height of the “internet bubble”. In contrast, the SDEV, dividend-price ratio, and payout-price ratio fitted values suggest that expected returns were high in 2009, at least compared to recent (post-1990) fitted values. Also, in the last year of the sample (2012), the CAY fitted values is close to a minimum, while the other variables generally suggest average expected returns (the dividend-price ratio suggests low returns, but this is likely attributable to shifts in payout policy).

One interpretation of these results is that, while expected stock and aggregate wealth returns have historically been correlated, these expected returns have recently diverged. Then, after the stock market selloff of 2008, expected stock returns may have been high while expected returns to labor (and aggregate wealth) may have been low. Although one cannot infer much from a few years of returns, it is interesting to note that from 2009-2012 stock returns were indeed high while labor income growth was low. Over these three years, average quarterly real labor income growth was 0.48%, substantially lower than the 1948-2012 average of 0.71%. Over the same three years, the average quarterly excess stock return was 3.60%, substantially higher than the 1948-2012 average of 1.50%.

3.4 Consumption Growth Predictability

When building a SDEV-based empirical model of expected returns, an important consideration is SDEV-based consumption growth predictability (i.e. the composite predictability of SDEV). This is analogous to the joint predictability associated with the dividend yield discussed by Cochrane (2008). The dividend yield should predict dividend growth, stock returns, or both. An absence of dividend growth predictability suggests return predictability. Similarly, an absence of SDEV-based consumption growth predictability suggests that variation in SDEV primarily reflects variation in expected returns.

To examine consumption growth predictability, I estimate a regression of annual real consumption growth on lagged SDEV (using the full sample):

$$\Delta c_{t+1} = \quad 0.0295+ \quad 0.0049 SDEV_t, \tag{9}$$

with an $\bar{R}^2$ of 0.017.[19] The point estimate of the slope parameter has the anticipated (nonnegative) sign; a high stock-wealth consumption ratio suggests high subsequent consumption growth. However, the slope parameter is barely marginally significant when using a Newey-West standard error (p-value 0.099) and insignificant when using a parametric bootstrap (p-value 0.159). Additionally, the point estimate suggests that consumption does not react strongly to SDEV. A one-standard deviation increase in SDEV is associated with a 0.50% increase in expected consumption growth, while the mean of annual consumption growth is 3%. These results are consistent with some SDEV-based consumption growth predictability. However, the results are not conclusive, and consumption does not appear to react strongly to SDEV. This is consistent with prior studies, which have found consumption growth to be difficult to predict. Overall, this suggests that much of the variation in SDEV reflects variation in expected stock returns.

[19] Results are similar when lagged consumption growth is included as an additional regressor.

4 Out-of-Sample Predictability

This section examines out-of-sample predictability associated with SDEV and other predictive variables. Recent predictability research often focuses on out-of-sample predictability (e.g. Lettau and Van Nieuwerburgh (2008), Cochrane (2008)), as out-of-sample tests are an appealing way to replicate investors' evolving information set and real-time predictability is of obvious practical importance.[20]

Goyal and Welch (2008) show that few predictive variables are useful out-of-sample.[21] This may be attributable to the large sampling error associated with out-of-sample statistics (see Cochrane (2008)) or unstable empirical models. In this section, I conduct a test for statistically compatible in- and out-of-sample predictability and find that, for some predictive variables, poor out-of-sample performance is not likely explained by sampling error.

4.1 Out-of-Sample Stock Return Forecasts

Recursive out-of-sample stock return forecast results are presented in Table 3. Results are reported for multiple forecast horizons (1, 4, 12, and 20 quarters). The out-of-sample period starts in 1982. This date is chosen because the long-run stock wealth-consumption relation is not precisely estimated before this date. Figure 5 displays the evolution of the recursively estimated slope parameter from the long-run stock wealth-consumption relation (Equation

[20] Out-of-sample tests have generated some controversy. Inoue and Kilian (2005) argue that in-sample tests should be preferred due to higher power. Other authors view out-of-sample tests as a complement, rather than a competitor, to in-sample tests that are informative about stability (Stock and Watson (2003), Goyal and Welch (2008)). For additional perspective, see Lettau and Ludvigson (2003), Clark and McCracken (2005), Cochrane (2008), and Campbell and Thompson (2008).

[21] Ferreira and Santa-Clara (2011) find out-of-sample predictability contingent on certain ex ante investor beliefs about dividend growth, earnings growth, and price-earnings ratio predictability. Cooper and Preistley (2009) find out-of-sample predictability using detrended industrial production. Kelly and Pruitt (2012) find out-of-sample predictability using a latent factor constructed from the cross-sectional distribution of book-to-market. Rapach, Strauss, and Zhou (2009) and Dangl and Halling (2012) show that combination forecasts and time-varying parameters, respectively, can lead to out-of-sample predictability. This paper finds out-of-sample predictability using simple methods and a single intuitive and well-motivated predictive variable, related to both the dividend yield and consumption-wealth ratio, without making assumptions about investors' prior parameter beliefs.

7). Estimates of the slope parameter do not appear to converge until approximately 1980. This suggests that SDEV is not likely to have been useful to investors before this date, unless investors were able to estimate this relation using other data. This highlights an important limitation of SDEV as a predictive variable. Because calculating SDEV requires estimation of a long-run relation with a persistent error term as well as the predictive relation, SDEV-based stock return forecasts may be imprecise without a long historical time series. For this reason, a useful SDEV-based empirical model likely requires a longer historical time series than models based on other predictive variables.

For each predictive variable, the maximum available amount of historical data is used to generate out-of-sample forecasts. Forecast series are evaluated using mean squared forecast error (MSFE). Models with low MSFE generate accurate out-of-sample forecasts. Table 3 reports an out-of-sample R^2 statistic (following Campbell and Thompson (2008) and Goyal and Welch (2008)), equal to the percentage reduction in MSFE associated with including the predictive variable, as well as the associated empirical p-value. The p-value is determined using bootstraps similar to those used in the in-sample section.

The row labeled "MSFE null" reports results based on the historical mean model, where the out-of-sample forecast equals the average historical excess return. The SDEV model outperforms the historical mean model at every forecast horizon. In several cases, the out-of-sample gains associated with SDEV are quite large. For example, use of SDEV reduces MSFE by over 25% at both the three- and five-year horizons. SDEV out-of-sample R^2 are generally close to in-sample R^2. For example, at the one-year horizon, the SDEV out-of-sample R^2 is 0.12, similar to the in-sample $\bar{R}^2$ (equal to 0.135, 0.096, and 0.104 for the full, post-war, and recent samples respectively).

Compared to other predictive variables, SDEV performs exceptionally well out-of-sample. The other predictive variables generally exhibit no evidence of out-of-sample return predictability. CAY exhibits some evidence of out-of-sample predictability at longer (three-and five-year) forecast horizons, but performs poorly at the quarterly and annual forecast horizon. Also, most variables perform worse than the historical mean model (consistent with the results of Goyal and Welch (2008)).

Overall, the out-of-sample results indicate that SDEV could have been used by a real-time investor to predict stock returns over the last thirty years. As discussed above, this period is limited by investors' ability to estimate the long-run stock wealth-consumption relationship. One can argue that this is a short period to evaluate return predictability. However, the results are statistically significant, so sampling error is not a likely explanation.

4.2 Testing for Compatible In- and Out-of-Sample Predictability

Tables 3 and 1 indicate that many variables predict stock return in-sample, but not out of sample (see also Goyal and Welch (2008)). This raises important questions about predictability. Poor out-of-sample predictability may be a result of the high sampling error of out-of-sample statistics. However, differences between in- and out-of-sample predictability could also be caused by unstable models (i.e. a model may work in part, but not all, of the sample). The latter explanation raises important concerns when constructing an empirical model of expected stock returns, as one should prefer a model that works all of the time to one that works part of the time. This section explores this issue by determining whether in- and out-of-sample predictability is statistically compatible for a given predictive variable. If in- and out-of-sample predictability is compatible, then we should observe no statistical differences. However, statistically incompatible in- and out-of-sample predictability is not likely explained by sampling error, and suggests an unstable model of expected stock returns.

Table 3 reports the results of a test for compatible in- and out-of-sample predictability. For each regression, I form a sampling distribution for the out-of-sample R^2 statistic under the assumption that the (bias-adjusted[22]) estimated in-sample model is true. I then compare the actual out-of-sample R^2 to the estimated sampling distribution to determine whether out-of-sample predictability is statistically distinguishable from in-sample predictability. This procedure follows Cochrane (2008), who shows that low (or negative) out-of-sample R^2 need not be interpreted as evidence against predictability because the sampling error of out-of-sample statistics can be large.

For the dividend-price and payout-price ratio, in- and out-of-sample predictability is statistically different. This inference differs from that of Cochrane (2008), and can be attributed to the extended sample used in this paper (use of a shorter sample results in more statistically compatible in- and out-of-sample predictability for the dividend yield). In- and out-of-sample predictability of CAY is often significantly different, consistent with the previously discussed Wald test. For SDEV, in- and out-of-sample predictability is statistically compatible.

These results are consistent with Tables 1 and 2, where many predictive variables exhibit little evidence of predictability in the recent subsample. The statistical incompatibility of the in- and out-of-sample predictability of some variables suggests that, for these variables, the recent deterioration in predictability cannot be explained by sampling error. Instead, this incompatibility may be explained by unstable empirical models of expected stock returns. Although these variables may capture important information about expected stock returns, extracting this information and building a reliable empirical model of expected returns may still be a challenging task. This appears to be especially true when generating real-time forecasts, as few models are useful out-of-sample. However, an SDEV-based predictive regression appears to be useful even when generating real-time forecasts.

[22] I adjust the autocorrelations of the predictive variables by $(1 + 3 * \rho)/T$ (see Kendall (1954)). I adjust the estimated slope parameter of the predictive relation (β) using a bootstrap. The bootstrap forms a sampling distribution for β under estimated in-sample system. The difference between the average β, using this sampling distribution, and β from the in-sample system is the bias.

5 International Results

In this section, I examine stock return predictability in Canada, Japan, Germany, France, Australia, and the U.K.[23] These countries were chosen because they had a reasonably long historical time series of both consumption and return data. Because these time series are often shorter than the U.S. time series, only in-sample results are considered. International market capitalization data is often only available for a short time series. For this reason, a cumulative stock return index (following Cochrane (1994)) is used to proxy for market capitalization. The stock return index is calculated separately for each country and is equal to cumulated log returns of the national stock index. Substituting this index for market capitalization is reasonable because market capitalization changes are largely driven by returns. For example, the correlation between U.S. quarterly market capitalization percentage changes and both price and total returns is 0.99.

In this section, the first stage is a regression of the market index on log consumption, performed separately for each country.[24] Results from the first stage are not reported. The second stage is a predictive regression of stock returns on the lagged residual from the first stage. Results from the predictive regressions are reported in Table 4. Results are reported for both real and excess returns.[25]

Overall, there is strong evidence of predictability across many of the countries exam-

[23] International consumption and inflation data were obtained from either government websites or the St. Louis Federal Reserve website. National stock index return data were purchased from Global Financial Data.

[24] One difficulty with comparing national stock returns and country-specific consumption is that stock markets are integrated. For example, German firms may list in the U.S., which suggests that German consumption should be tied to U.S. stock returns. Such effects should not affect SDEV-based return predictability much if the amount of cross-listing is small, constant over time, or if the first stage regression adequately controls for time-variation in cross-country listings. Also, because the effects of cross-listing, as well as listing requirements and other determinants of stock wealth, likely differ across countries, there is little reason to expect similar parameters across countries in the first-stage regressions.

[25] Real equity returns are calculated by subtracting country-specific inflation from nominal equity returns, calculated using prices expressed in each country's currency. Excess returns are calculated by subtracting the U.S. risk-free asset's real return from real equity returns. Similar results are obtained when excess returns are calculated as nominal equity returns less the nominal country-specific treasury bill rate, although requiring country-specific treasury bill rate data results in a smaller sample.

ined. There is widespread evidence of predictability in the U.K., Japan, Australia, and Germany. There is some evidence of predictability in France (in the five-year forecasts), but no significant evidence of predictability in Canada. However, all estimated slope parameters are negative (which is consistent with theory), and are often of a similar magnitude. The amount of predictability is often similar across countries, and to U.S predictability. For example, the average $\bar{R}^2$ from the one-year excess return predictive regression across the six countries is 0.107, which is similar to the U.S. value of 0.096 (using the post-war sample). Additionally, $\bar{R}^2$ increases sharply with the forecast horizon, in a manner that is consistent across all countries (including the U.S.). Overall, these results indicate that SDEV-based stock predictability is not limited to the U.S.

Power tests suggest that sampling error can easily explain cross-country variation in the results, including the insignificant results. For example, when simulating the data under the (bias-adjusted) one-year U.K. results of Table 4, the null hypothesis of no predictability is rejected 86% of the time. When the simulated sample is shortened by 12 years (corresponding to a more typical international time series of 1960-2011), the null hypothesis of no predictability is rejected 81% of the time. Therefore, it should not be surprising to find some insignificant results, particularly in the countries with short time series.

Fitted values from the international predictability regressions are plotted in Figure 6. Figure 6 suggests that there is substantial common variation in cross-country expected stock returns. For example, most countries have low expected returns around 2000, the height of the "internet bubble." This is verified in Table 5, which reports cross-country expected return correlations. Excluding Japan, the correlations are quite high (always over 0.60). However, Japanese expected returns are only modestly correlated with other countries' expected returns.

6 Conclusion

This paper is concerned with building a good empirical model of expected returns. Commonly-used predictive variables, such as the dividend yield, almost certainly capture information about expected stock returns. However, because the dividend yield may also predict dividend growth, isolating expected stock return information, and constructing a useful empirical model of expected stock returns, may be difficult. I refer to the fact that a variable may predict stock returns and/or something else as "composite predictability".

I develop a new predictive variable of stock returns that may be less influenced by composite predictability. A high stock wealth-consumption ratio should predict low stock returns, high consumption growth, or both (and the reverse). However, because consumption growth is, theoretically and empirically, largely unpredictable, much of the transitory variation in the stock wealth-consumption ratio should reflect variation in expected stock returns. The stock wealth-consumption ratio can be viewed as a modification of either the dividend-price ratio or consumption-wealth ratio.

Theory suggests that the stock wealth-consumption ratio, consumption-wealth ratio, and dividend yield reflect expected stock returns. However, each of these variables is influenced by quantities other than expected stock returns, and the extent of this influence is not addressed by theory. Therefore, it is difficult to assert, based on theory alone, that one of these variables should be preferred to another when forecasting stock returns.

Empirically, a stochastically detrended stock wealth-consumption ratio (SDEV) is an excellent predictor of stock returns. This predictability extends to a recent (1982-2012) subsample, where many predictive variables exhibit weak and/or deteriorating predictability. SDEV could have been used by a real-time investor to predict stock returns out-of-sample. Also, there is widespread evidence of SDEV-based predictability in non-U.S. stock returns.

The empirical models of expected stock returns developed in this paper can be used to understand the dynamics of expected stock returns. The results suggest that (1) stock returns are highly predictable, especially over longer horizons; (2) expected stock returns are quite persistent; (3) stock wealth and aggregate wealth discount rates exhibit meaningful independent variation that is especially apparent in the recent financial crisis; and (4) cross-country expected stock returns are generally highly correlated.

Finally, the effects of collective data mining are always a concern in a predictability study. Although it is impossible to fully eliminate such concerns, data mining concerns should be alleviated by theory, which suggests that SDEV is a reasonable expected stock return proxy, just like the consumption-wealth or dividend-price ratio. Data mining concerns should also be alleviated by the strength of the U.S. evidence, compatibility of the in- and out-of-sample evidence, and widespread evidence of predictability in non-U.S. data (non-U.S. stock index returns have been examined much less thoroughly than U.S. stock returns, so the effects of collective data mining should be less severe).

A Predicting Dividend Growth

Cochrane (2008) finds that, empirically, the dividend yield does not predict dividend growth. This suggests that variation in the dividend yield largely reflects variation in expected stock returns. In this appendix, I find that, under certain specifications, dividend growth is predictable, which suggests that controlling for expected dividend growth may be important when using the dividend yield to forecast stock returns. These differing results are partly attributable to the way dividends are constructed and partly attributable to the use of an extended sample. Inference also appears to be sensitive to the variable used to predict and/or measure dividend growth. In particular, there is stronger evidence of dividend growth predictability under specifications that account for firms' shift from dividends to repurchases and merger-related payouts (see Grullon and Michaely (2002), Boudoukh, Michaely, Richardson, and Roberts (2007), and DeAngelo, DeAngelo, and Skinner (2008)). Note that the use of cash dividends, as often constructed from CRSP data, does not account for this shift. From a return predictability perspective, the predictability of traditional dividend growth is less relevant than predictability associated with more comprehensive measures of firm payouts.

From 1929-2011, dividend yields have declined (see Figure 2). One explanation for this decline is a shift from dividends to other forms of payouts (repurchases and cash payouts associated with mergers). For example, in the 1980s and 1990s (when this shift appears to be most intense) dividend growth may have been low even if total payout growth was high. Use of a non-comprehensive measure of payouts may lead to substantially mismeasured payout growth, and misleading inference about payout growth predictability.

In this appendix, I examine predictive regressions of dividend and payout growth[26]. Payouts are calculated entirely from CRSP data as dividends less the change in aggregate

[26]Because payouts can be negative, payout growth cannot be calculated as the change in log payouts. Payout growth is defined as the difference in annual real payouts divided by the lagged real market value.

market capitalization due to stock issuances and repurchases (see Boudoukh, Michaely, Richardson, and Roberts (2007))[27]. I also examine an "adjusted price-dividend ratio", which measures transitory variation in the price-dividend ratio. Such a measure may be useful if there is substantial long-run variation in the dividend yield that is related to changes in the composition of payouts (e.g. a shift from dividends to repurchases). Such variation could mask higher-frequency variation in expected payout growth. For example, suppose that, over the historical sample, firms increasingly substitute repurchases for dividends. Then, the dividend yield will decrease, even if expected payout growth has not decreased. In this case, much of the historical variation in dividend yield may be attributable to changing payout policy. However, there may still be substantial predictable transitory variation in expected dividend growth. The adjusted price-dividend ratio is defined as the residual in the following regression (variables are in logs):

$$p_t = \alpha + \beta d_t + pd_t^*. \tag{10}$$

This follows the procedure used elsewhere in the paper to model long-run variation in the stock wealth-consumption ratio.

I estimate[28] the long-run relation between prices and dividends to be (with standard errors below)

$$\begin{aligned} p_t = \quad & \underset{(0.36)}{-1.32} + \quad \underset{(0.13)}{1.36 d_t} \quad + pd_t^*. \end{aligned} \tag{11}$$

In this regression, p_t is log real stock wealth at time t. d_t is the log of the sum of twelve-month trailing real monthly dividends[29]. This result implies that, over the historical sample,

[27]In the months in which AMEX and NASDAQ stocks are added to the CRSP data (September 1962 and December 1972 respectively), I adjust repurchases by the market capitalization associated with the added stock exchange, so the added stocks are not counted as (negative) payouts.

[28]This relation (and the associated standard errors), is estimated in the same way as the stock wealth-consumption relation.

[29]Dividends are computed monthly, and are assumed to not be reinvested, see Chen (2009).

stock wealth has increased more than one-to-one with dividends. This is consistent with traditional dividends capturing a smaller share of total payouts later in the sample. Figure 2 plots the raw and adjusted dividend yield (the adjusted dividend yield is equal to $-pd_t^*$). While the raw dividend-price ratio appears to trend lower over the historical sample, the adjusted ratio does not.

Table 6 reports predictive regressions of dividend and payout growth. Predictive variables considered are the dividend-price ratio, log dividend-price ratio, payout-price ratio, and the adjusted price-dividend ratio. Dividend growth is often predictable. This inference differs from that of Cochrane (2008), who finds little evidence of return predictability. This can be explained by the way dividends are constructed; Cochrane assumes that dividends are reinvested while I do not (following Chen (2009)). If dividends are assumed to be reinvested, there is no evidence of predictability from a regression of dividend growth on the lagged dividend yield.[30] These results also differ somewhat from those of Chen (2009), who finds that dividend growth is not predictable over a similar sample. This can be explained by the choice of predictive variables. Chen uses the log of the dividend-price ratio, which exhibits no significant predictability (using two-sided tests) in Table 6. However, there is stronger evidence of dividend growth predictability using other predictive variables. In particular, there is strong evidence of both dividend and payout growth predictability when using predictive variables that account for changes in firms' payout policy (i.e. the price-payout and adjusted dividend-price ratio). This suggests that controlling for shifts in payout policy is important when examining dividend or payout growth predictability.

The results of Table 6 provide, in some cases, strong evidence of dividend and payout growth predictability. It is somewhat disconcerting that other specifications provide no

[30] As discussed by Chen (2009), this can be explained by the interaction of returns and dividends. When calculating annual dividends under reinvestment, a dividend received in January will be influenced by February through December returns. Assuming dividend reinvestment can lead to a potentially misleading measure of dividend growth. For example, firms may have increased dividends paid to investors, but sufficiently negative returns could lead to low dividend growth.

evidence of dividend growth predictability.[31] However, given this mixed evidence, and the strong predictability evidence associated with predictive variables that account for changes in payout policy, it would seem difficult to conclude that payout and dividend growth are unpredictable. Overall, these results suggest that variation in expected dividend or payout growth could play an important role when predicting stock returns with the dividend or payout yield.

[31] One explanation for mixed dividend growth predictability results is that such predictive regressions suffer from the same composite predictability problem as predictive regressions of stock returns. Just as an expected return regression should control for expected dividend growth, an expected dividend growth regression should control for expected returns.

B Dissecting CAY's Predictability

To better understand the changing predictability associated with CAY in the recent subsample, I examine the ability of CAY to predict consumption growth, asset growth, and labor income growth (the "C", "A", and "Y" of CAY, respectively). If consumption and the aggregate wealth proxy are cointegrated, then variation in CAY must predict changes in one of these variables. If CAY does not predict stock returns but CAY is nonconstant, then CAY must predict something else. I examine this in Table 7.

Over the post-war (1952-2011) subsample, CAY is, empirically, an excellent predictor of stock returns. However, this predictability appears to be largely attributable to the 1952-1981 subsample. In the recent subsample (1982-2011), CAY exhibits no significant stock return predictability. This is not just due to low power and sampling error in the recent subsample. Using a Wald test, I reject the null hypothesis that the slope parameter in a regression of excess stock returns on lagged CAY is equal in the 1952-1981 and 1982-2011 subsamples with an empirical p-value of 0.016 (the asymptotic p-value is less than 0.01). (There is no significant difference in predictability across subsamples for SDEV.) The changing predictability of CAY is consistent with Tables 1 and 2, where CAY often does not predict stock returns in the recent subsample.

CAY appears to predict asset returns in the recent subsample, just not stock returns. To examine this, I project asset growth (ΔA from CAY) onto contemporaneous stock returns (r_s), and define the residual as asset growth orthogonal to stock returns ($r_a \perp r_s$). Table 7 demonstrates that, in the recent subsample, CAY is a marginally significant predictor of $r_a \perp r_s$. However, CAY is not a significant predictor of r_s in this subsample. Therefore, in the recent subsample, CAY appears to predict asset growth, but only asset growth orthogonal to stock returns. Also, CAY significantly predicts labor income growth in both samples[32],

[32] This result differs from that of Lettau and Ludvigson (2005) using data up to 2001, and can be explained by the expanded sample considered in this paper. I obtain results similar to Lettau and Ludvigson when using their sample.

which suggests that CAY predicts returns to human capital. These results suggest that, while CAY almost certainly captures information about expected stock returns, CAY also captures information about expected non-stock wealth returns. Then, forecasting stock returns with CAY may be somewhat challenging because one should control for expected non-stock asset returns (i.e. composite predictability appears to be a concern with CAY).

This does not contradict the theory of Lettau and Ludvigson (2001a), where the intertemporal budget constraint relates consumption, asset returns, and wealth. Theory does not say that CAY must predict stock returns, as variation in CAY may largely reflect variation in expected non-stock wealth returns and/or consumption growth. Also, theory does not say that predictability must be stable over time (e.g. theory is consistent with CAY predicting only stock returns in one sample and only labor income growth in another). Determining what CAY predicts is, ultimately, an empirical question.

References

Ang, A., and G. Bekaert, 2007, "Stock Return Predictability: Is it There?," *Review of Financial Studies*, 20(3), 651–707.

Bansal, R., R. Dittmar, and D. Kiku, 2009, "Cointegration and Consumption Risks in Asset Returns," *Review of Financial Studies*, 22(3), 1343–1375.

Bansal, R., and A. Yaron, 2004, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.

Binsbergen, J. H. V., and R. S. J. Koijen, 2010, "Predictive Regressions: A Present-Value Approach," *Journal of Finance*, 65, 1439–1471.

Boudoukh, J., R. Michaely, M. Richardson, and M. R. Roberts, 2007, "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing," *Journal of Finance*, 62(2), 877–915.

Boudoukh, J., M. Richardson, and R. F. Whitelaw, 2008, "The Myth of Long-Horizon Predictability," *Review of Financial Studies*, 21(4), 1577–1605.

Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Campbell, J. Y., and R. J. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1(3), 195–228.

Campbell, J. Y., and S. B. Thompson, 2008, "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies*, 21(4), 1509–1531.

Chang, Y., J. Y. Park, and K. Song, 2006, "Bootstrapping Cointegrating Regressions," *Journal of Econometrics*, 133(2), 703–739.

Chen, L., 2009, "On the Reversal of Return and Dividend Growth Predictability: A Tale of Two Periods," *Journal of Financial Economics*, 92, 128–151.

Clark, T. E., and M. W. McCracken, 2005, "The Power of Tests of Predictive Ability in the Presence of Structural Breaks," *Journal of Econometrics*, 124, 131.

Cochrane, J. H., 1994, "Permanent and Transitory Components of GNP and Stock Prices," *Quarterly Journal of Economics*, 109(1), 241–265.

——— , 2008, "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies*, 21(4), 1533–1575.

Cooper, I., and R. Preistley, 2009, "Time-Varying Risk Premiums and the Output Gap," *Review of Financial Studies*, 22(7), 2801–2833.

Dangl, T., and M. Halling, 2012, "Predictive Regressions with Time-Varying Coefficients," *Journal of Financial Economics*, 106, 157–181.

DeAngelo, H., L. DeAngelo, and D. J. Skinner, 2008, "Corporate Payout Policy," *Foundations and Trends in Finance*, 3(2-3), 95–287.

Fama, E. F., and K. R. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22(1), 3–25.

——— , 2001, "Disappearing dividends: changing firm characteristics or lower propensity to pay?," *Journal of Financial Economics*, 60, 3–43.

Fama, E. F., and G. W. Schwert, 1977, "Asset Returns and Inflation," *Journal of Financial Economics*, 5(2), 115–146.

Ferreira, M. A., and P. Santa-Clara, 2011, "Forecasting Stock Market Returns: The Sum of the Parts is More than the Whole," *Journal of Financial Economics*, 100(3), 514–537.

Goyal, A., and I. Welch, 2008, "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455–1508.

Grullon, G., and R. Michaely, 2002, "Dividends, Share Repurchases, and the Substitution Hypothesis," *Journal of Finance*, 57(4), 1649–1684.

Hall, R. E., 1978, "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86(6), 971–987.

Hjalmarsson, E., 2010, "Predicting Global Stock Returns," *Journal of Financial and Quantitative Analysis*, 45(1), 49–80.

Inoue, A., and L. Kilian, 2005, “In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?,” *Econometric Reviews*, 23(4), 371–402.

Kelly, B., and S. Pruitt, 2012, “Market Expectations in the Cross Section of Present Values,” *Working Paper*.

Kendall, M. G., 1954, “Note on Bias in the Estimation of Autocorrelation,” *Biometrika*, 41(3-4), 403–404.

Larrain, B., and M. Yogo, 2008, “Does Firm Value Move Too Much to be Justified by Subsequent Changes in Cash Flow?,” *Journal of Financial Economics*, 87, 200–226.

Lettau, M., and S. C. Ludvigson, 2001a, “Consumption, Aggregate Wealth, and Expected Stock Returns,” *Journal of Finance*, 56(3), 815–849.

——— , 2001b, “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying,” *Journal of Political Economy*, 109(6), 1238–1287.

——— , 2003, “Measuring and Modeling Variation in the Risk-Return Tradeoff,” in *Handbook of Financial Econometrics*. pp. 617–690.

——— , 2004, “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption,” *American Economic Review*, 94(1), 276–299.

——— , 2005, “Expected Returns and Expected Dividend Growth,” *Journal of Financial Economics*, 76(3), 583–626.

Lettau, M., and S. Van Nieuwerburgh, 2008, “Reconciling the Return Predictability Evidence,” *Review of Financial Studies*, 21(4), 1607–1652.

Li, H., and G. Maddala, 1997, “Bootstrapping Cointegrating Regressions,” *Journal of Econometrics*, 80(2), 297–318.

Menzly, L., T. Santos, and P. Veronesi, 2004, “Understanding Predictability,” *Journal of Political Economy*, 112(1), 1–47.

Nelson, C. R., and M. J. Kim, 1993, “Predictable Stock Returns: The Role of Small Sample Bias,” *Journal of Finance*, 48(2), 641–661.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," vol. 55, pp. 703–708.

Paye, B. S., and A. Timmermann, 2006, "Instability of Return Prediction Models," *Journal of Empirical Finance*, 13, 274–315.

Rapach, D. E., J. K. Strauss, and G. Zhou, 2009, "Out-of-sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy," *Review of Financial Studies*, 23(2), 821–862.

Stambaugh, R. F., 1999, "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.

Stock, J. H., and M. W. Watson, 1993, "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61(4), 783–820.

——— , 2003, "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 41(3), 788–829.

Valkanov, R., 2003, "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics*, 68(2), 201–232.

Table 1: Predictive Regressions, Annual Value-Weighted Excess Stock Returns

	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
				Panel 1: 1929-2012					
			Univariate Regressions						
	DP	PP	PD*	SDEV	RF				
				One-Year Returns					
β	0.035	0.071	-0.035	-0.076	-0.026				
se	0.027	0.018	0.019	0.019	0.016				
p	0.100	0.000	0.036	0.000	0.045				
emp p	0.129	0.002	0.078	0.002	0.054				
$\bar{R}^2$	0.018	0.113	0.018	0.135	0.005				
				Three-Year Returns					
β	0.122	0.157	-0.100	-0.165	-0.071				
se	0.045	0.028	0.041	0.034	0.043				
p	0.004	0.000	0.006	0.000	0.048				
emp p	0.028	0.001	0.033	0.003	0.057				
$\bar{R}^2$	0.143	0.256	0.097	0.282	0.039				
				Five-Year Returns					
β	0.207	0.121	-0.136	-0.206	-0.105				
se	0.042	0.051	0.043	0.049	0.060				
p	0.000	0.008	0.001	0.000	0.040				
emp p	0.003	0.032	0.029	0.018	0.066				
$\bar{R}^2$	0.365	0.121	0.154	0.376	0.084				

Panel 2: 1952-2012

			Univariate Regressions					Multiple Regression	
	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
				One-Year Returns					
β	0.049	0.017	-0.044	-0.058	-0.024	0.055	-0.074	-0.062	0.051
se	0.022	0.023	0.020	0.019	0.014	0.023	0.021	0.015	0.022
p	0.012	0.229	0.013	0.001	0.045	0.008	0.000	0.000	0.010
emp p	0.039	0.237	0.038	0.009	0.063	0.051	0.063	0.003	0.029

	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
$\bar{R}^2$	0.062	-0.007	0.045	0.096	0.000	0.081		0.238	
					Three-Year Returns				
β	0.099	0.040	-0.089	-0.121	-0.039	0.114	-0.150	-0.121	0.092
se	0.064	0.047	0.052	0.044	0.021	0.042	0.041	0.030	0.046
p	0.063	0.200	0.043	0.003	0.027	0.003	0.000	0.000	0.023
emp p	0.152	0.268	0.128	0.040	0.041	0.037	0.042	0.009	0.064
$\bar{R}^2$	0.133	0.008	0.107	0.214	0.005	0.169		0.466	
					Five-Year Returns				
β	0.129	0.047	-0.113	-0.167	-0.033	0.158	-0.222	-0.167	0.101
se	0.040	0.073	0.031	0.049	0.037	0.045	0.033	0.051	0.035
p	0.000	0.249	0.000	0.000	0.190	0.000	0.000	0.001	0.002
emp p	0.049	0.339	0.031	0.058	0.185	0.023	0.009	0.043	0.025
$\bar{R}^2$	0.166	0.005	0.125	0.299	0.000	0.229		0.597	

Panel 3: 1982-2012

			Univariate Regressions					Multiple Regression	
	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
					One-Year Returns				
β	0.041	0.016	-0.043	-0.062	-0.001	0.036	-0.063	-0.047	0.044
se	0.027	0.035	0.030	0.033	0.017	0.026	0.043	0.025	0.036
p	0.064	0.329	0.073	0.028	0.480	0.083	0.073	0.032	0.111
emp p	0.190	0.363	0.168	0.129	0.503	0.216	0.484	0.127	0.195
$\bar{R}^2$	0.026	-0.027	0.029	0.104	-0.027	0.010		0.092	
					Three-Year Returns				
β	0.098	0.027	-0.108	-0.152	-0.012	0.133	-0.152	-0.150	0.148
se	0.078	0.085	0.072	0.054	0.034	0.071	0.060	0.035	0.072
p	0.103	0.373	0.066	0.002	0.361	0.029	0.006	0.000	0.020
emp p	0.291	0.454	0.232	0.107	0.409	0.162	0.305	0.026	0.114
$\bar{R}^2$	0.112	-0.027	0.142	0.330	-0.027	0.185		0.334	
					Five-Year Returns				
β	0.156	0.018	-0.161	-0.215	0.045	0.212	-0.207	-0.157	0.163
se	0.041	0.092	0.035	0.046	0.058	0.057	0.032	0.048	0.067
p	0.000	0.422	0.000	0.000	0.219	0.000	0.000	0.001	0.008

	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
emp p	0.086	0.534	0.045	0.067	0.310	0.061	0.071	0.121	0.139
$\bar{R}^2$	0.252	-0.039	0.245	0.504	-0.029	0.335		0.703	

Notes - Table presents predictive regressions of value-weighted excess stock returns. Data is annual. Results are reported for 3 forecast horizons (1, 3, and 5 years) and 7 models. The models are: univariate predictive regressions with the dividend-price ratio (DP), the payout-price ratio (PP), the adjusted price-dividend ratio (PD*), SDEV, the risk-free interest rate (RF), CAY, and a multiple regression with SDEV, RF, and CAY. I report the parameter point estimate (β), a p-value based on the Newey-West standard error (NW p), a bootstrapped p-value (emp p), and the regression $\bar{R}^2$. P-values are based on one-sided tests.

Table 2: Predictive Regressions, Quarterly Value-Weighted Excess Stock Returns

Panel 1: 1952-2012

	Univariate Regressions						Multiple Regression		
	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
β	0.012	0.009	-0.010	-0.014	-0.008	0.012	-0.020	-0.019	0.012
se	0.008	0.008	0.006	0.005	0.006	0.005	0.006	0.006	0.005
p	0.024	0.121	0.046	0.006	0.096	0.008	0.001	0.001	0.010
emp p	0.086	0.188	0.164	0.019	0.109	0.011	0.008	0.003	0.013
$\bar{R}^2$	0.015	0.008	0.011	0.022	0.004	0.017		0.055	

Panel 2: 1982-2012

	Univariate Regressions						Multiple Regression		
	DP	PP	PD*	SDEV	RF	CAY	SDEV	RF	CAY
β	0.014	0.012	-0.014	-0.018	0.002	0.007	-0.020	-0.007	0.004
se	0.007	0.012	0.008	0.007	0.009	0.007	0.086	0.009	0.007
p	0.018	0.155	0.045	0.005	0.395	0.177	0.011	0.206	0.267
emp p	0.091	0.262	0.173	0.021	0.413	0.212	0.082	0.268	0.257
$\bar{R}^2$	0.018	0.011	0.018	0.036	-0.008	-0.003		0.025	

Notes - Table presents predictive regressions of value-weighted excess stock returns. Data is quarterly. The models are: univariate predictive regressions with the dividend-price ratio (DP), the payout-price ratio (PP), the adjusted price-dividend ratio (PD*), SDEV, the risk-free interest rate (RF), CAY, and a multiple regression with SDEV, RF, and CAY. I report the parameter point estimate (β), a p-value based on the Newey-West standard error (p), a bootstrapped p-value (emp p), and the regression $\bar{R}^2$. P-values are based on one-sided tests.

Table 3: Out-of-Sample Value-Weighted Excess Stock Return Forecasts

	DP	PD*	PP	SDEV	RF	CAY
			Quarterly Forecasts			
MSFE null	0.774	0.774	0.774	0.774	0.774	0.776
MSFE	0.796	0.780	0.790	0.758	0.786	0.817
OOS R^2	-0.028	-0.008	-0.019	0.021	-0.016	-0.052
p-value, $OOS\ R^2 > 0$	0.951	0.650	0.931	0.021	0.894	0.985
p-value, IS > OOS	0.059	0.289	0.039	0.772	0.094	0.015
			One-Year Forecasts			
MSFE null	3.075	3.075	3.075	3.075	3.075	3.074
MSFE	3.336	3.141	3.846	2.707	3.310	3.383
OOS R^2	-0.095	-0.021	-0.251	0.120	-0.077	-0.100
p-value, $OOS\ R^2 > 0$	0.967	0.737	0.964	0.009	0.951	0.964
p-value, IS > OOS	0.001	0.011	0.001	0.154	0.026	0.001
			Three-Year Forecasts			
MSFE null	7.679	7.679	7.679	7.679	7.679	7.468
MSFE	9.022	8.236	11.051	5.547	8.380	6.518
OOS R^2	-0.175	-0.073	-0.439	0.287	-0.091	0.149
p-value, $OOS\ R^2 > 0$	0.970	0.786	1.000	0.017	0.865	0.065
p-value, IS > OOS	0.001	0.006	0.002	0.240	0.020	0.029
			Five-Year Forecasts			
MSFE null	9.825	9.825	9.825	9.825	9.825	9.860
MSFE	10.971	9.798	12.548	5.544	11.887	7.384
OOS R^2	-0.117	0.003	-0.277	0.435	-0.210	0.251
p-value, $OOS\ R^2 > 0$	0.867	0.445	0.973	0.013	0.916	0.042
p-value, IS > OOS	0.023	0.046	0.015	0.629	0.023	0.283

Notes - Table presents recursive out-of-sample value-weighted excess stock return forecast results. Predictive variables are the dividend-price ratio (DP), the adjusted price-dividend ratio (PD*), the payout-price ratio (PP), SDEV, RF, and CAY. For SDEV, the long-run stock wealth-consumption relationship is estimated recursively using annual data from 1929-2012. Out-of-sample forecasts start in 1982 and end in 2012. For each model, I report mean squared forecast error (MSFE), the historical mean model MSFE (MSFE null), out-of-sample R^2 (equal to 1 - MSFE / MSFE null), the p-value associated with a test of out-of-sample R^2 greater than zero, and the p-value associated with a test that out-of-sample predictability is less than the in-sample predictability. All MSFEs are multiplied by 100.

Table 4: International Stock Return Predictability

		U.K. (1948-)	Japan (1955-)	Australia (1960-)	Canada (1961-)	Germany (1970-)	France (1955-)
				One-Year Forecasts			
R	β	-0.077	-0.076	-0.089	-0.041	-0.109	-0.067
	emp p	0.050	0.126	0.041	0.299	0.063	0.325
	$\bar{R}^2$	0.104	0.083	0.157	0.043	0.195	0.064
R_x	β	-0.078	-0.082	-0.086	-0.039	-0.104	-0.064
	emp p	0.058	0.090	0.056	0.357	0.096	0.353
	$\bar{R}^2$	0.113	0.102	0.143	0.036	0.186	0.059
				Three-Year Forecasts			
R	β	-0.201	-0.218	-0.195	-0.106	-0.265	-0.201
	emp p	0.077	0.043	0.018	0.115	0.013	0.203
	$\bar{R}^2$	0.327	0.299	0.389	0.180	0.551	0.262
R_x	β	-0.193	-0.227	-0.180	-0.095	-0.236	-0.189
	emp p	0.067	0.038	0.020	0.212	0.029	0.195
	$\bar{R}^2$	0.340	0.342	0.336	0.136	0.497	0.256
				Five-Year Forecasts			
R	β	-0.276	-0.326	-0.279	-0.174	-0.331	-0.343
	emp p	0.069	0.014	0.088	0.183	0.001	0.045
	$\bar{R}^2$	0.468	0.438	0.534	0.360	0.674	0.515
R_x	β	-0.241	-0.322	-0.233	-0.129	-0.265	-0.296
	emp p	0.074	0.005	0.130	0.379	0.011	0.031
	$\bar{R}^2$	0.429	0.473	0.405	0.189	0.547	0.478

Notes - Table presents predictive regressions of value-weighted stock returns. Predictive variable is the lagged residual from a regression of the national return index (cumulative log returns) on national consumption. Results are presented for real returns (R) and real returns less the real U.S. risk-free rate (R_x). Data is annual. Last annual return is 2011. Results are reported for forecast horizons of 1, 3, and 5 years. For each regression, I report a parameter point estimate (β), bootstrapped p-value (emp p), and the regression $\bar{R}^2$. P-values assume one-sided tests.

Table 5: International Expected Stock Return Correlations

	Germany	Canada	Japan	U.K.	Australia	U.S.
France	0.611	0.826	0.245	0.733	0.739	0.769
Germany		0.629	0.287	0.714	0.662	0.654
Canada			-0.015	0.800	0.917	0.876
Japan				0.279	0.037	-0.002
U.K.					0.794	0.835
Australia						0.777

Notes - Table presents correlations of international expected stock returns. Expected stock returns are the fitted values from the one-year regression of Table 4 (non-U.S. data) and the one-year SDEV regression of Table 1.

Table 6: Predicting Dividend and Payout Growth

Panel 1: Predicting Dividend Growth

	dp	log dp	pp	pd*
β	-0.049**	-0.035	-0.034**	0.053***
se	0.022	0.022	0.015	0.020
$\bar{R}^2$	0.175	0.085	0.079	0.211

Panel 2: Predicting Payout Growth

	dp	log dp	pp	pd*
β	-0.075	-0.048	-0.647***	0.149**
se	0.087	0.086	0.107	0.069
$\bar{R}^2$	-0.007	-0.010	0.411	0.010

Notes - Table presents annual predictive regressions of dividend and payout growth. Data spans 1929-2011. Annual dividends are measured as the sum of real monthly dividends. Dividend growth is defined as the difference in log annual dividends. Payout growth is the difference in annual payout divided by lagged market value. Predictive variables are the dividend-price ratio (dp), the log of the dividend-price ratio, the payout-price ratio (pp), and the adjusted price-dividend ratio (pd*). Intercepts are included in all regressions, although intercept results are not reported. ***/**/* indicates significance at the 1%/5%/10% level.

Table 7: Regressing Consumption Growth, Asset Returns, and Labor Income Growth on CAY

	Δc	Δa	Δy	r_s	$r_a \perp r_s$
			1952-2011		
β	0.017	0.385***	0.122**	1.385***	0.107
se	0.029	0.140	0.054	0.477	0.079
$\bar{R}^2$	-0.003	0.026	0.014	0.022	0.005
			1982-2011		
β	0.047*	0.316*	0.124**	0.757	0.190*
se	0.028	0.174	0.059	0.543	0.101
$\bar{R}^2$	0.013	0.014	0.028	0.003	0.025

Notes - Table presents results of predictive regressions of consumption growth (Δc), asset growth (Δa), labor income growth (Δy), value-weighted stock returns (r_s), and the component of asset growth orthogonal to stock returns ($r_a \perp r_s$) on lagged standardized CAY (see Lettau and Ludvigson (2001a)). $r_a \perp r_s$ is the residual from a regression of Δa on contemporaneous r_s. Intercepts are included in all regressions, although intercept results are not reported. ***/**/* indicates significance at the 1%/5%/10% level.

Figure 1: Standardized Stock Wealth-Consumption Deviation (SDEV), 1929-2012

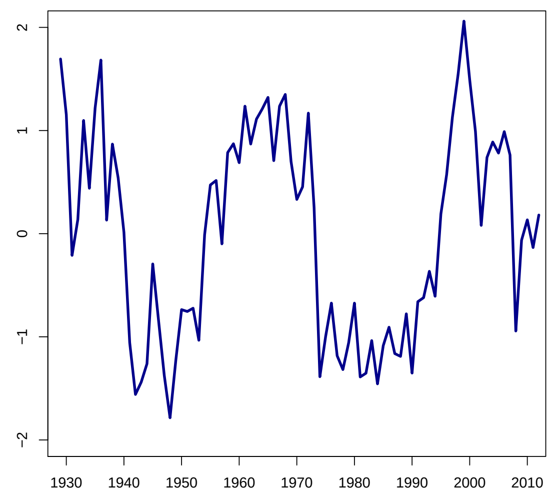

Figure 2: Standardized Dividend Yield, 1929-2012

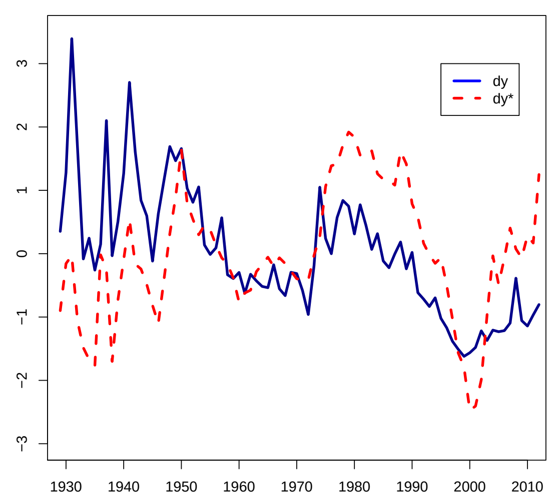

Notes - Figure plots the standardized dividend yield (dy) and adjusted dividend yield (dy*).

Figure 3: Standardized Payout Yield, 1929-2012

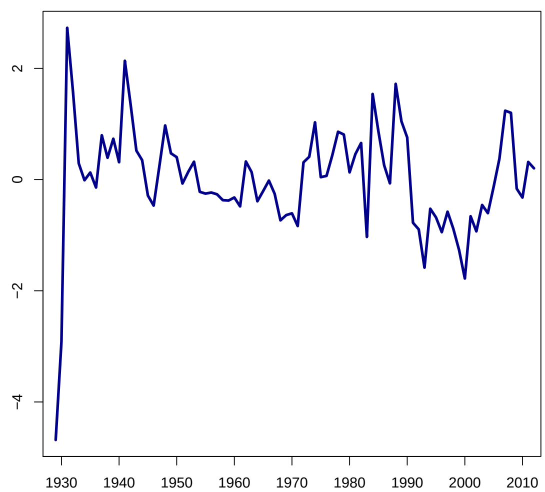

Figure 4: U.S. Fitted Excess Returns

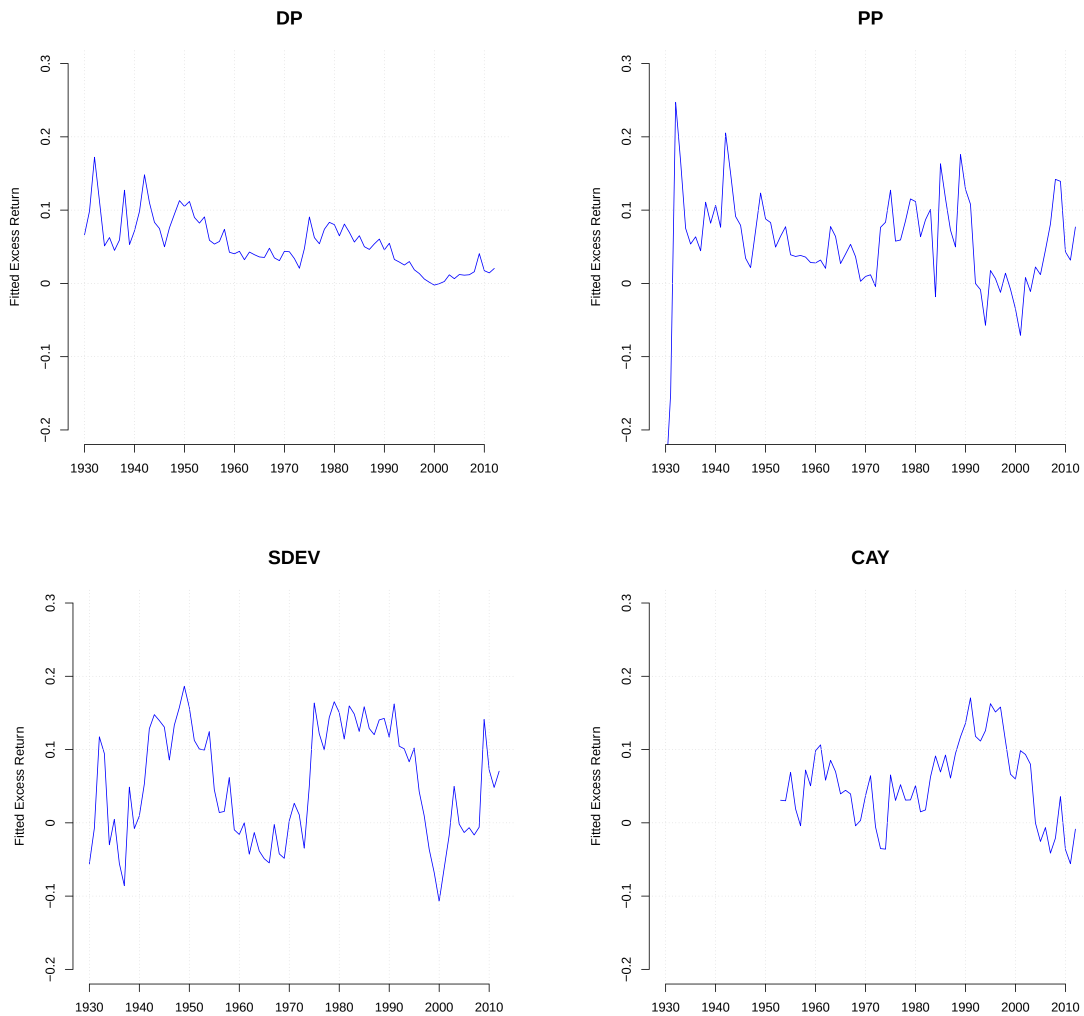

Notes - Figures display fitted values of US excess stock returns using the dividend-price ratio (DP), payout-price ratio (PP), SDEV, and CAY. Fitted values correspond to the one-year regression in Table 1.

Figure 5: Recursively Estimated Slope of Stock Wealth-Consumption Relation

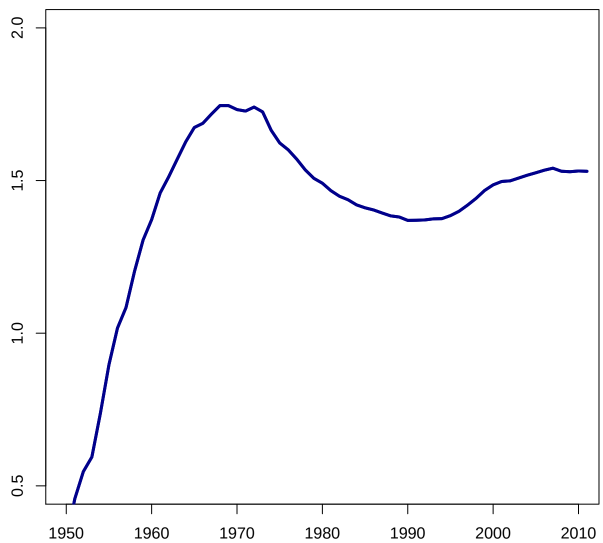

Notes - Figure displays recursive estimates of β in the regression: $s_t = \alpha + \beta c_t + SDEV_t$. Data is annual. Estimates from 1929:Q1-1939:Q4 are not displayed.

Figure 6: International Fitted Excess Returns

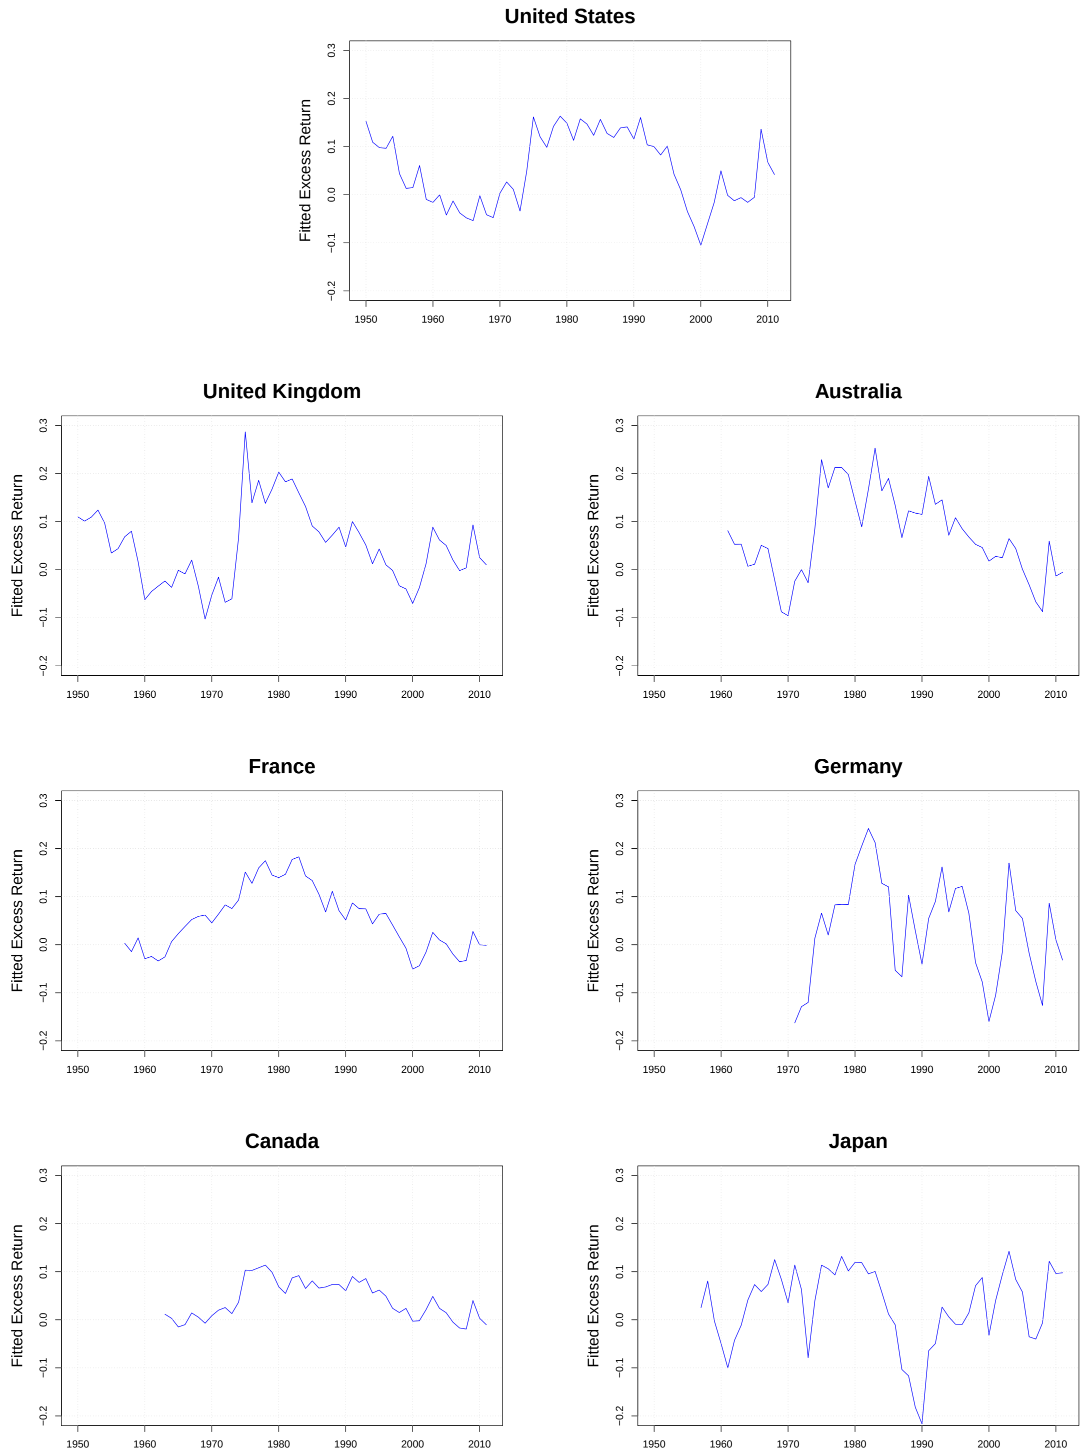

Notes - Figures display fitted values of international excess stock returns. U.S. fitted values correspond to the one-year SDEV regression in Table 1. All other fitted values correspond to the excess return (real returns less the real U.S. risk-free rate) regressions from Table 4.

Bonds, Aggregate Wealth, and Stock Market Risk

Mark Rachwalski

March 26, 2014

Abstract

Bond returns predict consumption growth after controlling for equity returns, which suggests that bonds capture important information about aggregate wealth. Consistent with this, bond risk is priced in the cross section of stocks. Bond risk partially explains momentum profits and the flat cross-sectional relation between stock index beta and returns. The results suggest stock indices are an insufficient proxy for aggregate wealth and that bond risk is an important component of consumption risk. Also, term structure-based return predictability often declines with risk loadings, suggesting that time-variation in risk premia is not the sole driver of such predictability.

1 Introduction

According to the ICAPM (Merton (1973)), expected returns are determined by the covariances of asset returns with aggregate wealth and other state variables, which define the investment opportunity set. While this theory provides a succinct and intuitively appealing characterization of risk and return, applying the ICAPM is difficult. State variables are not identified by the model. Therefore, in the absence of additional restrictions, the ICAPM can be used to justify any variable known to predict returns (see Cochrane (2001)). A priori, aggregate wealth is likely to play an important role. However, even models that focus on aggregate wealth (i.e. the Sharpe-Lintner CAPM) are difficult to implement, as aggregate wealth (or market portfolio) returns are difficult, if not impossible, to measure (Roll (1977)). Then, use of an observable market portfolio proxy likely yields an incomplete measure of risk (e.g. if the value-weighted stock index is an insufficient market portfolio proxy, then beta to the stock index is an insufficient measure of risk).

Researchers have adopted a number of models of expected returns, most prominently the three-factor model of Fama and French (1993). Although this model is successful in capturing many known return patterns, the economic source of the returns of the size and (especially) value factors remains controversial because the factors are not directly connected to investor consumption or preferences (see Lakonishok, Shleifer, and Vishny (1994) and Fama and French (2004)). Instead, the factors are motivated by known patterns in returns.

In this paper, I search for priced consumption-based risk. However, instead of directly measuring an asset's consumption risk as covariance with consumption growth, I take an indirect approach, related to the maximum correlation portfolio of Breeden, Gibbons, and Litzenberger (1989), where a portfolio of assets is used to proxy for consumption growth and asset risk is determined by covariance with the proxy. Breeden (1979) demonstrates that the ICAPM can be collapsed into a single-factor model, where risk is measured as the covariance of returns with consumption. This suggests that a good way to identify ICAPM state variables is to find variables that are correlated with consumption growth. Empirically, bond returns predict consumption growth, even after controlling for stock returns. This suggests that bond returns capture important information about state variable innovations. Also, bond returns predict labor income growth. This suggests that bonds proxy for the labor component of aggregate wealth.

Motivated by these results, I test for priced bond risk in the cross section of stock returns. Because bonds proxy for a state variable and covariance with state variables is priced, stock returns should depend on covariance (or beta) with bond returns. This approach, where Breeden's consumption CAPM is used to motivate relevant risks, can be viewed as a disciplined way to search for variables that should be priced in the cross section. I find that, consistent with the consumption results, bond beta is positively priced, even after controlling for factors that are often interpreted as corrections for risk (i.e. the market, size, and book-to-market factors of Fama and French (1993)). This relation does not appear to be limited to the U.S., as I also find evidence supporting such a relation when using U.K. data.

The positive bond beta-return relation is strongest among large stocks, which suggests that the relation is not driven by well-known microstructure-related return patterns known to be important for small stocks (e.g. short-term reversals), but is instead reflective of risk. Also, the relation is strongest when controlling for time variation in the correlation of stock and bond returns, which is important because such variation is substantial. Interestingly, the aggregate stock-bond correlation has declined over time, which suggests that stock returns are becoming less informative about market portfolio returns (bonds capture important information about aggregate wealth, and stock returns are becoming less informative about bond returns). Then, stock indices may be becoming an increasingly inadequate proxy for the market portfolio.

Because bond risk is empirically tied to consumption risk, it is easy to interpret the high returns of high bond-risk stocks as compensation for risk. In contrast, for many other characteristics known to be related to stock returns, such an interpretation is less straightforward (e.g. the value premium may be due to risk or mispricing). Then, the return patterns documented in this paper can be used to interpret other anomalies. For example, if high book-to-market stocks tend to have high bond betas, then the value premium likely reflects, at least in part, compensation for risk. I find no evidence that the size or value premiums can be explained by bond risk. However, somewhat surprisingly, a portion (22%) of momentum profits can be explained by bond risk.

Bond risk can also help to explain the surprisingly flat cross-sectional relation between stock index beta and returns, which is an important empirical failure of the CAPM (see Fama and French (2004), Lewellen and Nagel (2006), and Frazzini and Pedersen (2013)). Although the time-series average of stock index excess returns is significant and positive, bearing stock index-beta risk seems to offer little reward in the cross section (i.e. stocks with high stock index betas have low alphas and the reverse). This apparent contradiction can be partially resolved by bond risk; stocks with high stock index betas tend to have low bond risk while stocks with low stock index betas tend to have high bond risk. Therefore, generalizing the risk-adjustment model can explain a good amount of this puzzling lack of a relation (in particular, this reduces the difference in alpha between high and low stock-risk stocks by 30-35%).

I further examine connections between stock and bond markets by considering term structure-based return predictability of stock portfolios sorted by bond and stock risk.[1] Term structure variables are known to be excellent predictors of excess bond returns (see, for example, Cochrane and Piazzesi (2005)). Then, there may be important cross-sectional

[1]This follows Baker and Wurgler (2012), who find that the term structure factor of Cochrane and Piazzesi (2005) predicts the excess returns of bond-like stocks (Baker and Wurgler define bond-like stocks as stocks that have bond-like characteristics, such as large size and low volatility). In contrast, this paper examines the return predictability of stock portfolios formed by sorting on bond risk.

differences in stock return predictability, related to the level of bond risk. To examine this, I sort stocks by stock and bond risk and examine term structure-based return predictability of the sorted stock portfolios. This may help to identify the economic quantities that drive return predictability (e.g. if the returns on stocks with high bond risk are highly predictable, but returns on stocks with low bond risk are not, then predictability is more likely attributable to variation in the price of bond risk).

I find no evidence that the returns on stocks with high bond risk are more predictable than the returns on stocks with low bond risk. However, the returns on stocks with low stock risk (i.e. stocks with a low beta to the value-weighted stock index) are far more predictable, based on the term structure factor, than the returns on stocks with high stock risk. Therefore, the results suggest that the returns of low risk (both stock and bond risk) stocks are at least as predictable, and sometimes more predictable, than the returns of high risk stocks. Because consumption growth is positively related to both stock and bond returns (so that sorts on stock and bond risk can be interpreted as a partial sort on consumption risk), the results suggest that the returns of low consumption-risk stocks are often more predictable than the returns of high consumption-risk stocks.

This finding is difficult to reconcile with term structure-based return predictability primarily driven by time variation in stock, bond, or consumption risk premia (if variation in risk premia is driving predictability, and risk exposures are sufficiently stable, then high risk stocks should have the most predictable returns[2]). Although the results of this paper cannot definitively isolate the source of term structure-based return predictability, the results are consistent with some return predictability driven by sentiment or periodic flight-to-quality episodes (see Baker and Wurgler (2006, 2012)). Interestingly, under this interpretation, the results suggest that high-quality stocks (i.e. stocks that comove with bonds in response to shifts in sentiment) are better defined as having low stock risk than high bond risk. Then,

[2] I find that time-varying betas are not likely to reconcile this prediction with the empirical evidence.

flight-to-quality episodes seem to be characterized by a shift in investor preferences from high stock-risk stocks to low stock-risk stocks and bonds.

This paper can be viewed as an attempt to model risk under the Roll (1977) critique. Although aggregate wealth returns may remain unobservable, use of a broader market portfolio proxy should yield more accurate measures of risk. Note that a good aggregate wealth proxy need not include all assets. Instead, as emphasized by Shanken (1987), the quality of a proxy will depend on the correlation between the proxy and omitted assets. The results of this paper suggest that stock index returns are an insufficient proxy for aggregate wealth returns.[3]

This paper is related to a growing literature that examines the joint role of stocks and bonds in asset pricing. Historically, researchers have often examined stocks and bonds separately (e.g. the three factors of Fama and French (1996) are constructed entirely from excess stock returns). Recently, researchers have increasingly attempted to model aggregate stock and bond returns jointly[4]. The results of this paper complement such work by showing that priced bond risk is connected to consumption risk and accounts for some of the the cross-sectional variation in expected stock returns. Moreover, this component of stock risk is not explained by familiar risk-adjustment models.[5]

[3]This is consistent with prior research. Shanken (1987) generally rejects the joint hypothesis that the CAPM is valid and the correlation between a multivariate proxy (stocks and government bonds) and the market portfolio exceeds 0.7. Fama and French (2004) and Lewellen and Nagel (2006) discuss the poor empirical performance of the CAPM and conditional CAPM (respectively); one explanation for this poor performance is a poorly measured market portfolio. Stambaugh (1982) estimates an equity share of the market portfolio of 20-30%, but does not consider the value of human capital when computing shares. Accounting for human capital yields an estimated wealth share of 7-10% (see also Jagannathan and Wang (1996)).

[4]See, for example, Bekaert, Engstrom, and Xing (2009), Bekaert, Engstrom, and Grenadier (2010), and Campbell, Sunderam, and Viceira (2013).

[5]Other researchers have examined relations between bonds and the cross section of stocks, often as a way to examine the value premium. Baker and Wurgler (2012) find that bonds and stocks with bond-like characteristics (e.g. large size, low volatility) exhibit comovement and overlapping predictability. Lettau and Wachter (2011) consider a common shock to stock and bond risk premia that can explain equity, term, and value premia. Koijen, Lustig, and Van Nieuwerburgh (2012) propose a model that jointly prices stocks sorted by book-to-market and treasury bonds of various maturities, as well as time-series variation in bond returns.

This paper contains four main contributions. First, I establish that bond returns predict consumption growth. Predicting consumption growth is itself interesting. Hall (1978) hypothesized that consumption should follow a random walk (provided real rates are constant). Empirically, many studies have found that consumption follows (at least approximately) a random walk (in addition to Hall, see Cochrane (1994), Lettau and Ludvigson (2001a), and Lettau and Ludvigson (2005)). However, recent asset pricing models often admit the possibility of consumption growth predictability, which can be caused by time-inseparable preferences (see Weil (1990)) or by directly assuming a consumption process that implies consumption growth predictability (as in Bansal and Yaron (2004), see also Beeler and Campbell (2012)).[6] Second, I show that bond risk, which is empirically linked to consumption risk, is priced in the cross section of stock returns. This serves to link the cross section of expected stock returns, bond risk, and consumption risk. This link makes bond beta easily interpretable as risk. In contrast, other characteristics known to be related to returns in the cross section (e.g. size and book-to-market) are not as easily interpreted as risk. Third, this paper shows that the anomalous returns of momentum and stock index beta-sorted portfolios can be partially explained by bond risk. Therefore, forming a more general measure of risk partially addresses two important asset pricing anomalies. Finally, this paper shows that term structure-based return predictability is often stronger for low risk assets. This suggests that such predictability is not primarily driven by variation in risk premia.

The paper proceeds as follows. Section 2 examines predictive regressions of consumption and labor income growth. Section 3 establishes that bond risk is priced in the cross section of stocks. Section 4 examines several well-known anomalies before and after controlling for bond risk. Section 5 examines term structure-based predictability of risk-sorted stock

[6] Other explanations of predictable consumption growth include time-varying real discount rates, consumption adjustment constraints, and consumer's behavior departing from the usual solution to the utility maximization problem (see Campbell and Mankiw (1990)).

portfolios. Section 6 concludes.

2 State Variables, Consumption, Labor Income, and Unemployment

In this section, I consider predictive regressions of consumption growth, labor income growth, and the change in the unemployment rate. This section establishes that bond returns predict consumption growth. This empirical result is of independent interest. Also, in later sections I find that exposure to bond risk (beta) is priced in the cross section of stocks. Because bond returns predict consumption growth, bond returns likely capture information about aggregate wealth or other state variable innovations. Then, bond beta measures covariance with a state variable and should capture information about expected returns. So this section serves to empirically link bond beta and consumption risk. Finally, this section demonstrates that bond returns capture information about labor income growth and the unemployment rate, which suggests that bond returns capture information about returns to human capital. Because human capital is a large part of aggregate wealth, this is consistent with a positive relation between bond and aggregate wealth returns.

2.1 Identifying State Variable Proxies

I want to determine whether bond returns (and other asset returns) serve as a proxy for aggregate wealth. Conceptually, this could be accomplished by regressing aggregate wealth returns on contemporaneous asset returns. However, this regression is not feasible because aggregate wealth is not observed. Alternatively, one could regress the returns of assets with large aggregate wealth shares (e.g. human capital) or proxies for these assets (e.g. labor income growth) on contemporaneous asset returns. However, such regressions suffer from lack of data and mismeasurement concerns (despite these shortcomings, I consider predictive regressions of labor income growth later in this section).

In this paper, I identify state variable proxies using predictive regressions of consumption growth. The ICAPM of Merton (1973) says that expected stock returns are determined by the covariances of returns with state variables. Breeden (1979) demonstrates that the ICAPM can be collapsed into a single-factor, consumption-based model. In Breeden's model, expected returns are determined by an asset's covariance with consumption. Breeden's result suggests that one way to identify state variable proxies is to find variables that are correlated with consumption growth.[7] Then, covariance with state variable proxies should capture information about risk. This can be viewed as a disciplined way to identify risks that should affect the cross-section of expected stock returns.

2.2 Empirical Setup

I consider predictive regressions of temporally aggregated real consumption growth on lagged real consumption growth and asset returns. For example, the four-quarter consumption growth regression takes the form:

$$\Delta c_{t+1,t+4} = \beta_0 + \beta_c \Delta c_t + \beta_S r_{s,t} + \beta_L \Delta l_t + \beta_{BAA} r_{BAA,t} + \epsilon_{t+1,t+4}. \quad (1)$$[8]

In this specification, Δc is consumption growth, r_s is the value-weighted stock market return, Δl is labor income growth (a proxy for human capital returns), and $r_{BAA,t}$ is the BAA corporate bond index return. Other regressions follow this pattern. With this specification, one can establish that asset returns predict changes in relevant state variables.[9]

[7]The use of predictive regressions of consumption growth to identify proxy assets is closely related to the maximum correlation portfolio (MCP) of Breeden et al. (1989). The MCP is formed by projecting consumption growth onto asset returns. Proxy assets, as identified in this paper, play an important role in such a projection.

[8]In this specification, β_c may be expected to be positive due to time-aggregation in consumption data (see Breeden et al. (1989)). Including lagged consumption growth as an explanatory variable controls for the effects of time aggregation.

[9]To examine the contemporaneous relation between asset returns and state variable innovations, rather than the predictive relation, one could reduce the lag of the asset returns by one (in the above example,

Use of temporally aggregated consumption growth, rather than one-period consumption growth, allows for an examination of longer-run consumption growth predictability. Parker and Julliard (2005) find that the use of temporally aggregated consumption growth improves the ability of the consumption CAPM to explain the cross-section of stock returns. There are several reasons why a researcher may wish to consider temporally aggregated consumption growth. First, consumption may respond with a lag to innovations in wealth. Second, Weil (1990) notes that, if utility is not time separable, contemporaneous consumption may not be a sufficient statistic for marginal utility and Hall's random walk hypothesis need not hold. Therefore, models featuring time-inseparable preferences (Epstein and Zin (1989, 1991), Weil (1990), and Campbell and Cochrane (1999)) may lead to consumption growth predictability. Third, there may be a persistent component to consumption growth (as in long-run risk models, see Bansal and Yaron (2004)). Fourth, slowly adjusting consumption could be caused by a systematic failure to solve the utility maximization problem (see Campbell and Mankiw (1990)) or adjustment constraints.

Consumption predictability results are reported for consumption including durable goods and consumption excluding durable goods. Durable goods are consumed over many periods, which complicates measurement of the associated utility stream. This suggests the use of ex-durable consumption as a measure of utility. However, durable goods affect utility. For this reason ex-durable consumption may be an inadequate measure of utility (see Yogo (2006) for evidence that utility is not separable in durable and nondurable consumption). Because it is not clear which consumption measure is a better measure of utility, both are examined.

An additional benefit of temporal aggregation of consumption growth is to alleviate the durable goods problem. The temporal mismatch between durable goods purchases and the associated utility stream should be less severe as the measurement interval increases (i.e.

---

asset returns and labor growth would have subscript $t+1$).

a durable good may not be consumed immediately, but much of the utility stream may lie within the consumption growth measurement interval). As noted by Breeden, Gibbons, and Litzenberger (1989), nondurable goods become more "durable" when consumption is measured over a short interval. Conversely, durable goods become less "durable" when consumption is measured over long intervals.

In all consumption regressions, the dependent variable is temporally aggregated real per capita consumption growth. In the labor income growth regressions, the dependent variable is temporally aggregated real per capita labor income growth. Four aggregate wealth proxies are considered: stocks, government bonds, corporate bonds, and labor income.[10] Labor income growth is an appealing market portfolio return proxy because human capital is likely a large part of aggregate wealth. Stock and bond returns are appealing proxies because these assets are part of aggregate wealth, and the returns of these assets likely serve as timely proxies for the returns of omitted or mismeasured assets[11].

2.3 Consumption Growth Results

Table 1 reports predictive regressions of annual consumption growth on lagged consumption growth and a single asset return.[12] All asset returns are separately significant when predicting one-year consumption growth. For both consumption measures, corporate and government bond yield changes are highly significant predictors of consumption growth. All explanatory variables are standardized in these regressions, so the parameter estimates can

[10] Value-weighted stock index returns are used to measure stock returns. Following Jagannathan and Wang (1996) and Lettau and Ludvigson (2001b,a), per capita labor income growth is used to proxy for the return on human capital. The negative of the change in the ten-year treasury and BAA yield is used to proxy for bond returns. Results are similar when using government and corporate bond returns, which are available at this frequency (quarterly). The negative of the yield change is used because this proxy for bond returns is available at a daily frequency for a reasonably long time series, and, for this reason, is used in later sections. Bond results are also similar when using the AAA yield.

[11] Omitted assets include privately-held firms, durable goods, real estate, and commodities. Human capital returns are likely imperfectly measured by labor income growth.

[12] In unreported regressions, I find that consumption growth is sometimes positively related to lagged term spreads (see Harvey (1988)), although the results and conclusions of this section are not meaningfully altered by including term spreads as a regressor.

be interpreted as the change in consumption growth associated with a one standard deviation change in the explanatory variable. A one standard deviation increase in the value of BAA corporate bond assets is associated with an increase in the one-year consumption growth rate (including durable goods) of 0.55%. This is approximately 25% of the unconditional mean of consumption growth (2.21%). The sensitivity of consumption growth excluding durable goods is somewhat smaller, where a one standard deviation increase in the value of BAA assets is associated with an increase in the consumption growth rate of 0.28%. The bond return parameters are generally of the same magnitude as the corresponding stock return parameter and always larger than the labor parameter; this suggests that bonds are a relatively important explanatory variable.

The corporate bond parameter point estimates are larger than the government bond parameter estimates (although the difference is not statistically distinguishable). This is an appealing result. Although both corporate and government bonds likely capture information about aggregate wealth, market portfolio cash flows are almost certainly not as secure as government bond cash flows, which suggests that corporate bonds may be a superior market portfolio proxy. For this reason, and the slightly stronger empirical relationship between corporate bonds and consumption, in the remainder of the paper I often focus on corporate, rather than government, bonds. However, results are generally similar when using government bonds.

The choice of one year of temporal aggregation of consumption growth is arbitrary. Table 2 reports the ability of the asset returns to predict consumption growth at different horizons by varying the amount of temporal aggregation from one to twelve quarters. In these regressions, all asset returns are considered simultaneously in multiple regressions.

The corporate bond parameter is significant in every regression where consumption is measured with durable goods, including the two- and three-year regressions, where stock returns are not significant. The explanatory variables are standardized, so the parameter estimates are comparable. In every regression, the point estimate of the corporate bond return parameter is larger than the point estimates of the other asset return parameters (and in some cases, substantially larger). To get a sense of the economic significance of the results, at the two-year horizon, a one standard deviation increase in the value of BAA corporate bonds is associated with an increase of 0.55% in expected consumption growth; this is an increase of 12% from the average 2-year consumption growth rate of 4.45%. Results are similar, although slightly weaker, in the ex-durable consumption regressions.[13] Overall, Tables 1 and 2 demonstrate that bond returns are an important predictor of consumption growth, even after controlling for stock returns and labor income growth. Therefore, the results suggest that bond returns contain important information about state variable innovations that is not captured by the other asset returns.

2.4 Labor Income and Unemployment Results

In the consumption growth regressions, labor income growth is used as a proxy for human capital returns. However, human capital returns depend on the innovations in contemporaneous labor income growth and expected future labor income growth. Then, if labor income growth is predictable, contemporaneous labor income growth may be an insufficient proxy for human capital returns. In this section, I consider predictive regressions of one-, two-, and three-year labor income growth. Assets with significant explanatory power in these regressions likely capture information about returns to human capital. If labor income growth is predictable, these asset returns may serve as a timely proxy for human capital returns. Also, because the wealth share of human capital is large, an asset that predicts labor income growth likely captures important information about aggregate wealth.

The labor income (and unemployment) specifications mirror the consumption specifications. Labor income regression results are reported in Panel 1 of Table 3. Empirically,

[13] One explanation for the weaker predictability of ex-durable consumption is that durable goods purchases are more sensitive to changes in aggregate wealth than nondurable goods and service purchases.

stock and bond returns predict one-year labor income growth. However, at the two-year horizon, only bond returns significantly predict labor income growth. At this horizon, a one standard deviation increase in the value of BAA corporate bond assets is associated with an increase of 0.65% in labor income growth (an increase of 19% from average two-year labor income growth, 3.34%). At the same horizon, a one standard deviation increase in the value of stocks has small negative (although insignificant) association with labor income growth. Overall, the results suggest that bond returns capture timely information about human capital returns that is not captured by stock returns or lagged labor income growth.

Panel 2 of Table 3 reports predictive regressions of unemployment rate changes. Results are similar to the consumption and labor income results; bond returns capture information about future unemployment rate changes that is not captured by stock returns or contemporaneous labor income growth. A high unemployment rate suggests low per capita labor income. Then, the labor income and unemployment results both suggest that bond returns are positively associated with human capital returns. More generally, the unemployment rate may be an important state variable (or state variable proxy). If bond returns predict unemployment rate changes, then bond returns likely capture timely information about state variable innovations.

2.5 Discussion of Consumption and Labor Income Results

Empirically, corporate bond returns are at least as important as stock returns when predicting consumption growth. This suggests that corporate bonds capture important information about aggregate wealth or other state variables. The empirical importance of corporate bonds may be somewhat surprising; many researchers would probably contend, a priori, that stocks should be more important. One explanation for the empirical results is that corporate bonds capture information about the omitted (and large) labor component of the market portfolio. This explanation is supported by the labor income growth and unemployment rate results. Also, this explanation is sensible; when economic conditions deteriorate, the value of corporate bonds and human capital are both likely to decline. More generally, if some omitted components of the market portfolio are better represented by a senior claim on firms' assets (corporate bonds) than a ju nior claim (stocks), then corporate bond returns are likely to capture important information about market portfolio returns.

According to the ICAPM, expected stock returns are determined by the conditional covariances of stocks with state variables. The above results suggest that both stock and bond returns serve as important proxies for state variable innovations. In this case, covariance with stock and bond returns should be priced risks. The next section examines the relation between bond risk and expected returns in the cross section of stocks.

3 Bond Risk and the Cross-Section of Stock Returns

3.1 Bond Beta as a Measure of Risk

Consider a multivariate proxy (see Shanken (1987)) for the market portfolio consisting of stocks and bonds. Such a proxy can be represented by a linear regression:

$$r_m = \pi_s r_s + \pi_b r_b + \epsilon. \tag{2}$$

In this regression, π_s and π_b are the projection coefficients associated with stocks and bonds, respectively. The projection coefficients will reflect correlations with omitted assets. Given this proxy, the CAPM[14] risk-return relation for an arbitrary asset can be written as

$$E[r_i] = \gamma \pi_s Cov[r_i, r_s] + \gamma \pi_b Cov[r_i, r_b], \tag{3}$$

[14] I focus on a CAPM (rather than an ICAPM) setting for the sake of simplicity, and because market portfolio risk is likely a large part of consumption risk.

where γ is the price of market risk (defined as the expected return of the market divided by the variance of the market).

Given $\pi_b \neq 0$, $Cov[r_i, r_b]$ will capture information about expected returns (assuming $\gamma \neq 0$). From the consumption section, I infer $\pi_b > 0$. If the price of market risk is positive ($\gamma > 0$), then $\pi_b > 0$ implies

$$\frac{\partial E[r_i]}{\partial Cov[r_i, r_b]} > 0. \tag{4}$$

An asset with higher bond risk (i.e. a larger bond beta) will have a higher expected return. This hypothesis will be empirically tested using the cross section of stocks.

3.2 Data

Data sources are standard. The stock sample consists of all NYSE, Amex, and NASDAQ stocks. Momentum portfolio returns are calculated as the equal-weighted return of prior winners less prior losers (using a formation period of months $t-6$ through $t-2$ and extreme decile portfolios). Month $t-1$ is skipped to alleviate the effects of bid-ask bounce and short-term reversals (see Jegadeesh (1990) and Lehmann (1990)). To reduce the microstructure effects and other return patterns associated with small and low-priced stocks, stock-month observations with a lagged stock price less than two are deleted. Return data for the Fama and French (1996) factors is obtained from from Ken French's website. Bond yield data is obtained from the St. Louis Federal Reserve website. Daily corporate bond yield data is available starting in 1986, which limits the time series used to examine the price of corporate bond risk. Daily treasury bond data is available starting in 1962, which allows for a longer time series.

3.3 Estimating Bond Betas

Each month, stocks are sorted into four portfolios based on historical bond betas. For each stock, the following regression is estimated monthly using the prior two years of daily data[15]:

$$r_{i,t} = \alpha_i + \beta_i(-\Delta_{BAA,t}) + \epsilon_{i,t}. \tag{5}$$

$-\Delta_{BAA,t}$ is the negative of the change in the BAA corporate bond yield. β_i is the estimated corporate bond beta for stock i.

The use of daily data is critical when estimating corporate bond betas. This can be seen in Table 4, which reports correlations between historical corporate bond beta using daily data (β_D), historical corporate bond beta using monthly data (β_M), and (subsequent) realized corporate bond beta (β_R, calculated using daily data). A high correlation between historical and realized betas indicates that historical beta is a good predictor of subsequent realized beta. Pooling all observations, the Pearson correlation between β_D and β_R is 0.364, while the correlation between β_M and β_R is 0.012.[16] The monthly estimates are noisy (and not very useful). In contrast, the estimates based on daily data are informative about subsequent realized betas. However, the correlation is not particularly high even when using daily data. This likely occurs because firm-level bond beta estimates are noisy and may change over time. For this reason portfolios are used to examine the bond risk-return relation (portfolio beta estimates are more precise than firm-level beta estimates).

Table 4 also reports the correlation between historical and realized corporate bond beta by stock size. There is a substantial increase in the β_D-β_R correlation as size increases (the

[15] The historical window used to calculate bond betas is set to two years because the correlation between estimated and realized betas (see Table 4) is lower for windows less than two years, which suggests that short historical windows yield noisy beta estimates. Use of more than two years of data does not substantially improve the estimates, but reduces the length of the sample used to evaluate portfolio returns.

[16] The monthly estimate uses a longer historical window than the daily estimate (five years of prior data vs. two years). This is necessary to generate reasonably precise estimates. Varying the window did not substantially improve the estimate.

β_D-β_R Spearman correlation is 0.243 for the smallest size quartile and 0.557 for the largest size quartile). This is important when examining the relation between bond risk and returns. If the sort on expected bond betas is noisy (which will happen when the β_D-β_R correlation is low), then it may be difficult to form portfolios with cross-sectional dispersion in expected bond risk. One way to combat this noise is to use value-weighted portfolio returns, which assign larger weights to stocks with more precise beta estimates. The use of value-weighted returns carries additional benefits, as cross-sectional differences in returns are less likely to be driven by microstructure-related return patterns that are more prevalent in small, illiquid stocks (e.g. bid-ask bounce, short-term reversals). An alternative approach, also used, is to calculate equal-weighted returns after eliminating small stocks from the sample.

3.4 Descriptive Statistics

Descriptive statistics of corporate bond beta-sorted quartile portfolios are presented in Table 5. Value-weighted mean portfolio returns increase monotonically as corporate bond beta increases. This is consistent with positively priced corporate bond risk (and corporate bond returns positively correlated with aggregate wealth innovations). The high minus low hedge portfolio mean value-weighted return is 41.9 basis points per month. This hedge portfolio will be referred to as CORP in subsequent analysis.

Equal- and value-weighted realized corporate bond betas are reported for each portfolio (these are post-formation betas, not the betas used to form the portfolios). Betas are calculated from a regression of stock returns on bond yield changes, and can be interpreted as the average stock return (in percent) given a one percent change in bond yields (not bond returns). For example, the value-weighted corporate bond beta of the high bond beta portfolio is 2.41. This indicates that, given a one percent decline in bond yields, this stock portfolio will tend to return 2.41%. The same decline in bond yields is associated with a -2.05% return in the low bond beta portfolio and a 4.46% return in the high-low hedge portfolio.

Realized betas increase monotonically across the quartile portfolios. This suggests that the sorting procedure is successful in generating post-formation period dispersion in corporate bond betas. Therefore, differences in corporate bond risk are a plausible explanation for differences in mean portfolio returns. The equal-weighted beta dispersion is smaller than the value-weighted beta dispersion (3.28 vs. 4.46), which is consistent with value-weighting reducing some of the noise associated with estimating small stock bond betas.

Although value-weighted portfolio returns indicate that high corporate bond beta stocks earn high returns, equal-weighted returns do not. As discussed above, this may be due to imprecise corporate bond beta estimates for small stocks. However, later in the paper I find that both equal- and value-weighted hedge portfolios earn positive conditional alphas when controlling for commonly-used risk factors and time-varying risk loadings.

The strong presence of the corporate bond beta-return relation in large stocks is particularly interesting because this suggests that the relation cannot be explained by return patterns that are more important for small stocks (e.g. bid-ask bounce, short-term reversals, short-sale constraints). Instead, the dispersion in returns is more likely related to differences in risk. Also, CORP is an economically interesting portfolio because the dispersion in returns is present in stocks that are a large part of the value-weighted stock market. This means that corporate bond risk is likely relevant for the average investor.

3.5 Time-Series Regressions of CORP on Contemporaneous Factors

Time-series regressions of CORP on the value-weighted stock market (MKT), the Fama-French three-factor model (FF3), and the Fama-French three-factor plus momentum (WML) model are presented in Table 6. Results are reported for the full sample and large stock subsamples (formed by progressively eliminating stocks from size quartiles 1, 2, and 3, using NYSE breakpoints).

Value-weighted alphas are economically large (0.55-0.80% monthly) and highly significant when using the full sample or the large stock subsamples. However, equal-weighted alphas are not significant when using the full sample, although equal-weighted alphas increase in magnitude and significance as small stocks are progressively eliminated from the sample. When focusing on the largest size quartile (on average, 81.8% of market capitalization), equal-weighted alphas are significant in all regressions. Therefore, corporate bond risk appears to be especially important in understanding the returns of large stocks (and value-weighted returns). Later in the paper, I find that the conditional alphas of equal-weighted hedge portfolios, which allow for time-varying risk exposures, are often positive and similar to the conditional alphas of value-weighted hedge portfolios.

Two-sided tests are used to determine significance. However, theory and the predictive regressions of consumption suggest that CORP's raw portfolio returns should be positive (see Equation 4), which implies a one-sided test). Therefore, one can reasonably interpret the reported significance levels as conservative. While theory need not imply a one-sided test if the portfolios load on other risk factors, theory suggests that, after adjusting for other relevant risks, higher corporate bond beta should be associated with higher expected returns. Then, by controlling for other risks, factor models may yield a cleaner test of Equation 4. Finally, results are similar when examining the mean returns of the four bond beta-sorted stock portfolios via GMM rather than only hedge portfolio returns (results not tabulated).

3.6 Government Bonds and Time-Varying Risk

In this section, I examine the returns of a hedge portfolio formed by sorting on government bond beta (referred to as GOV). This is done partly as a robustness check. Because government bonds are, on average, positively related to subsequent consumption growth (see Table 1), GOV should, on average, earn positive returns (this does not imply that conditional expected returns must always be positive). Also, use of government bond yields allows for a longer time series. Daily BAA yield data is available starting in 1986, while daily ten-year treasury yield data is available starting in 1962.

The GOV hedge portfolio is formed identically to the CORP hedge portfolio, except that the ten-year treasury yield is used in place of the BAA corporate bond yield. Time-series regressions of GOV on the three Fama-French factors are presented in Table 7. From 1988-2011, use of government bonds yields similar results to the corporate bond results described above. Over this time period, the value-weighted monthly three-factor alpha of GOV is 70 basis points. However, use of the early (1964-1987) or full (1964-2011) samples yields alphas that are small in magnitude and insignificant.

There is a simple explanation for the low and insignificant alphas of the 1964-1987 sample. From 1964-1987, mean excess monthly returns of corporate, intermediate-term government, and long-term government bonds were -0.03%, 0.07%, -0.07%, respectively[17] (interest rates increased over this time period), while from 1988-2011, mean returns were 0.43%, 0.21%, and 0.45%, respectively. All of the 1988-2011 mean returns are significantly greater than zero, while all of the 1964-1987 mean returns are insignificant and small in magnitude. Therefore, bonds realized low returns over the early subsample and high returns over the late subsample. Then, the GOV portfolio likely realized an insignificant alpha from 1964-1987 because realized bond returns were low. This could occur even if the ex ante risk premium was positive.

The risk exposure of the GOV portfolio appears to have shifted over time. From 1964-1987, GOV loaded positively on the market factor, with a loading of 0.147. However, from 1988-2011, the factor loading is -0.559. This change in risk exposures (which is statistically significant), as well as the time-varying alphas of Table 7, suggests that a straightforward application of the three-factor model may struggle, at times, with correctly pricing bond

[17]These mean returns are calculated using data obtained from the Ibbotson Associates Yearbook.

risk. Although the three-factor model is generally highly successful in capturing certain return patterns in the cross-section of stocks, the model is not constructed to price bonds.

3.7 Conditional Factor Regressions

This section examines the effects of changing stock-bond covariances on the above factor regressions. This is important because time-variation in stock-bond covariance is substantial, and could easily affect the expected returns of the bond-beta hedge portfolios examined in this paper. I find that controlling for time-variation in risk exposures reveals a positive alpha for the GOV portfolio over the full sample.

Consider the expected return of the bond beta-sorted hedge portfolio under the stock-bond CAPM (as in Equation 3),

$$E[r_h] = \gamma\pi_s Cov[r_h, r_s] + \gamma\pi_b Cov[r_h, r_b]. \tag{6}$$

A ceteris paribus increase in $Cov[r_h, r_b]$ will lead to higher expected hedge portfolio returns (γ and π_b are assumed to be positive). While theory suggests that covariances, rather than correlations, are relevant for asset pricing, some of the following discussion will focus on correlations, which are easier to interpret. As a practical matter, choosing to examine correlations rather than covariances may not matter much, as these quantities are strongly related (the time-series correlation of stock-government bond correlation and covariance is 0.78).

Allowing for time variation in $Cov[r_h, r_b]$ may be important because stock-bond covariances appear to be quite volatile.[18] The mean stock-treasury bond correlation is 0.25 from 1964-1987 and 0.03 from 1988-2011 (the respective mean covariances are 0.00021 and

[18] This paper does not explore the causes of changes in the stock-bond correlation, which remains an interesting question (see Baele, Bekaert, and Inghelbrecht (2010)). Instead, this paper is focused on the asset pricing implications of changing correlations.

-0.00020, with the difference highly significant).[19] This change in correlation carries potentially important implications for researchers interested in modeling risk. A low stock-bond correlation, combined with an a priori belief that stock and bond risk premia are both positive, suggests that a model that explains the mean returns of stocks may struggle with explaining the mean returns of bonds (if stock and bond returns are uncorrelated, then covariance with a common factor cannot explain nonzero stock and bond returns). While the stock-bond correlation is never particularly high, the above results suggest that stocks are becoming an increasingly inadequate proxy for bonds. Of course, it remains possible that a stock-based model may successfully capture bond risk using the cross-section of stocks (e.g. the HML and SMB factors may reflect bond risk). However, Table 6 suggests that three-factor model, often used to control for risk, does not adequately capture bond risk.

In addition to the long-run decline in the stock-bond correlation[20], there is substantial month-to-month variation. Within the 1964-1987 and 1988-2011 subsamples, the time series standard deviation of the stock-bond correlation is 0.25 and 0.43, respectively. The substantial time variation in the stock-bond correlation (both low- and high-frequency), combined with the risk-return relation of Equation 6, suggests that modeling time variation in stock-bond correlations may be critical when examining the interaction of stock and bond risk and return.

To explore the effects of time-varying stock-bond correlations, I consider a factor regression that allows for time-varying risk exposures. Each month I calculate the historical

[19] The correlations and covariances discussed in this section are calculated monthly, using daily data from the previous three months. I use three months of data because correlations from shorter samples are less informative about subsequent correlations (this was determined by examining regressions of correlation on lagged correlations calculated from different sample lengths). See Campbell, Sunderam, and Viceira (2013) and d'Addona and Kind (2006) for additional evidence that stock-bond correlations have declined.

[20] Although long-run changes in the covariances of major assets are likely to be important in understanding historical returns, researchers have generally not focused on such low-frequency changes (however, see Campbell, Sunderam, and Viceira (2013) and Baele, Bekaert, and Inghelbrecht (2010)). Instead, researchers have mostly considered the higher-frequency dynamics of stock-bond correlations (see Connolly, Stivers, and Sun (2005)).

covariance between the hedge portfolio and bond yield changes.[21] I construct a dummy variable that takes the value one when the lagged hedge portfolio-bond covariance is less than the median, and zero otherwise.[22] I then run the following factor regression

$$\begin{aligned} r_t =&\alpha + \alpha_D * D + (\beta_{mkt} + \beta_{mkt,D} * D) * MKT + (\beta_{hml} + \beta_{hml,D} * D) * HML \\ &+ (\beta_{smb} + \beta_{smb,D} * D) * SMB + \epsilon, \end{aligned} \tag{7}$$

where D is the dummy variable. This can be interpreted as a specification that allows for betas to vary with observable state variables, as in Shanken (1990) (in this case the state variable is just the dummy variable, which is determined by lagged covariance). According to Equation 6, the expected return of the hedge portfolio should be high when the conditional covariance is high ($D = 0$), and low when the conditional covariance is low ($D = 1$).

Factor regression results are reported in Table 8. Results are reported for three samples using government bonds (1964-2011, 1964-1987, and 1988-2011), the corporate bond sample (1988-2011), and a full-sample implementable trading strategy that uses government bonds (which calculates the median, used to generate the dummy variable, recursively). Consistent with the predicted relations, hedge portfolio alphas tend to be high when the conditional covariance is high, and low when the conditional covariance is small. For example, in the full sample, the monthly hedge portfolio alpha is 82 basis points when hedge portfolio bond risk is high, and -46 basis points when bond risk is low. In every specification, the alpha is lower when bond risk is low ($D = 1$). Also, in every specification except the early sample (1964-1987), the hedge portfolio earns a significant positive alpha when bond risk is high. This is consistent with Equation 6, and suggests that the insignificant full sample alpha of

[21] This requires computing daily hedge portfolio returns.

[22] Using predictive regressions, I find that lagged hedge portfolio-bond covariance is informative about subsequent covariance, and is therefore a reasonable instrument for the expected covariance. I do not use contemporaneously measured covariance, which may be correlated with the error term of the regression. I use a dummy variable, rather than the level of historical covariance, to reduce the influence of outliers (stock-bond covariance is highly leptokurtic). In untabulated results, I find that the hedge portfolio-stock covariance is relatively stable, and that little is gained from allowing for time-variation in this covariance.

Table 7 is not due to an insignificant price of bond risk, but can instead be attributed to time-variation in the risk of the GOV hedge portfolio. Overall, the conditional regressions provide further evidence that the three factor model does not adequately capture bond risk.

More broadly, the recent decline in the stock-bond correlation suggests that accounting for bond risk may becoming more important. Tables 1 and 2 suggest that stocks and bonds each capture important information about aggregate wealth. If stock and bond returns are correlated, stock returns may capture much of the aggregate wealth-relevant information in bond returns, and beta to a stock index may serve as a good measure of the consumption risk of an asset. Then, the three-factor model may be sufficient to price bond risk. However, the recent decline in stock-bond correlations suggests that stock returns are becoming less informative about aggregate wealth returns, and that stock index beta is becoming an increasingly inadequate measure of risk. This is consistent with the time-varying alphas of Table 7.

3.8 International Evidence

I repeat the above analysis using a sample of U.K. stock and bond returns[23] and find similar results. I calculate a bond beta for each U.K. stock monthly, using U.K. government bond returns and daily data over the preceding year. I then rank the stocks into bond beta quintiles and examine the returns of these quintile portfolios. From December 1984-October 2013, I find that the mean monthly equal-weighted returns of the low to high quintiles are 0.93%, 0.97%, 1.19%, 1.12%, and 1.11%, respectively. The mean return of the high-less-low hedge portfolio is 0.18%, with a standard error of 0.08% (t-statistic of 2.25). In a regression of the high-less-low hedge portfolio on excess aggregate U.K stock market returns, the intercept is 0.18% with a standard error of 0.08% (t-statistic of 2.20). Therefore, the U.K results mirror the U.S. results. High bond beta stocks earn, on average, high returns.

[23] I obtain the broadest possible sample of U.K. stocks available from Datastream.

4 Anomalies and Bond Risk

4.1 Size, Value, and Momentum Profits

This section considers time-series regressions of hedge portfolios formed by sorting on size, value, and prior return. Such portfolios are known to exhibit nonzero returns after controlling for covariance with the value-weighted stock market (see Fama and French (1992, 1996), Jegadeesh and Titman (1993)). However, interpretation of these returns remains controversial, especially for the value premium (e.g., see Lakonishok, Shleifer, and Vishny (1994), Fama and French (1996, 2004), and Liew and Vassalou (2000)). Nonzero returns of these portfolios may be compensation for an omitted risk factor. Alternatively, these returns may be due to mispricing. In this section, I attempt to determine whether these nonzero returns can be attributed to bond risk, which is related to consumption risk. Then, this section may help to determine the economic cause of an anomaly's nonzero returns. To do this, I simply regress the returns of traded factors known to exhibit nonzero returns on contemporaneous factors, including CORP. If, for example, the value premium can be attributed to bond risk, then the HML portfolio should load positively on CORP, and have an insignificant alpha when CORP is included as a regressor.

Results are reported in Table 9. There is little evidence that CORP attenuates the returns of the SMB or HML portfolio. In Panel 1, a regression of SMB on MKT and CORP has a larger intercept than a regression of SMB on MKT. However, neither intercept is significant; this is likely due to the short length of the time series (1988-2011). Similarly, there is little evidence that CORP can explain the value premium. In Panel 2, the magnitude of the HML intercept is nearly unchanged after including CORP. These results are consistent with Table 6, where CORP generally loads negatively on both SMB and HML.

However, corporate bond risk can explain a portion of momentum profits. Under the three-factor model, the alpha of the momentum portfolio is 1.318% per month. Under the three-factor plus CORP model, the alpha is 1.024% per month. The intercept attenuation associated with including CORP as a regressor is highly significant (a test of equal intercepts in the model with and without CORP models yields a p-value 0.003[24]). Therefore, CORP appears to explain about 22% of momentum profits. This is an interesting finding because interpretation of the momentum anomaly has proven difficult. In contrast, because corporate bond beta is connected to consumption risk, the nonzero returns of the CORP portfolio can be easily understood as compensation for risk. This interpretation of CORP, and the empirical relationship between CORP and WML, suggests that the momentum anomaly is partly due to cross-sectional differences in risk.[25]

4.2 The Cross-Sectional Price of Stock Risk

Corporate bond risk can partially explain the surprisingly flat cross-sectional relation between stock index beta and returns. Although the time-series average of value-weighted excess stock index returns is significant and positive, bearing additional stock index risk in the cross section appears to offer little, if any, incremental return (see Fama and French (2004), Lewellen and Nagel (2006), and Frazzini and Pedersen (2013)). This is examined in Table 10, which reports time series regressions of stock index beta-sorted stock portfolios.

Panel 1 of Table 10 demonstrates that, under both the market and three-factor model, the intercept of the low stock risk (beta) portfolio is greater than the intercept of the high beta portfolio. I reject equal intercepts with a p-value of 0.000 under both models. Therefore, in this sample, the cross-sectional price of bearing stock risk is inconsistent with the mean return of the aggregate market.[26]

[24]This test was conducted by simultaneously estimating the two models via GMM while imposing the restriction that the intercepts are equal.

[25]Jegadeesh and Titman (2002) show that cross-sectional variation in unconditional expected returns (which are presumably reflective of risk) is unlikely to explain momentum profits. However, it remains possible that variation in conditional expected returns can explain a substantial portion of momentum profits.

[26]Time-varying betas could yield nonzero intercepts (if this variation is not modeled), although Lewellen and Nagel (2006) show that the variation would have to be implausibly large to explain asset pricing anoma-

In Panel 2, CORP is included as an additional explanatory variable. I still reject tests of equal extreme portfolio alphas, as well as all alphas equal to zero. However, the use of CORP substantially attenuates hedge portfolio returns (this attenuation is statistically significant, with a p-value less than 0.01). For example, adding CORP to the MKT model reduces the intercept difference from 0.742% to 0.484% (monthly). Similarly, adding CORP to the FF3 model reduces the difference from 0.643% to 0.438%. Therefore, adjusting for corporate bond risk attenuates, by 30%-35%, the surprisingly flat relation between stock index beta and returns. This occurs because stocks with high stock index betas tend to have low corporate bond betas, while stocks with low stock index betas tend to have high corporate bond betas. Interestingly, more of the attenuation appears to be attributable to an increase in the negative alpha of the high-beta portfolio than a decrease in the positive alpha of the low-beta portfolio.

Prior empirical studies often focus on the excess zero-beta rate, which theory suggests to be equal to zero but is generally estimated to be positive (see Fama and French (2004)). This is a consequence of the high average returns of stocks combined with the flat cross-sectional relation between stock index beta and returns. The results of Table 10 imply that controlling for bond risk will yield a lower zero-beta rate. Stocks with high stock index beta tend to have low bond risk, while stocks with low stock index beta tend to have high bond risk. Then, controlling for bond risk will tend to yield a steeper relation between stock index beta and returns and a lower zero-beta rate.

These results suggest that the value-weighted stock index is an insufficient proxy for the market portfolio. A more general risk-adjustment model substantially alleviates the flat cross-sectional relation between stock index beta and returns, and suggests that the CAPM may not be such a poor model of risk and return. However, I still reject the CAPM, as alphas are not universally equal to zero even under the generalized risk-adjustment model.

---

lies.

It is possible that an even better market proxy will further attenuate this alphas (e.g. including bonds in addition to a bond-risk sorted hedge portfolio of stocks). Additionally, there are a number of alternative explanations that, in conjunction with the generalized risk-adjustment model described here, may fully explain the flat cross-sectional beta-return relation (see Black (1972), Baker, Bradley, and Wurgler (2011), and Frazzini and Pedersen (2013)).

5 Bond Risk and Return Predictability

Bond returns are known to be highly predictable by variables related to the term structure of interest rates (see, for example, Cochrane and Piazzesi (2005)). This paper finds that there is important cross-sectional variation in stocks' exposure to bond risk. Together, these results suggest that there may be substantial cross-sectional heterogeneity in stock return predictability. For example, term structure variables may predict the returns of high bond-risk stocks but not low bond-risk stocks. This section examines this conjecture, and may help to isolate the source of return predictability (e.g. if only high bond-risk asset returns are predictable, then variation in bond risk premia is likely an important part of term structure-based predictability).

Baker and Wurgler (2012) examine the predictability of bond-related sorted stock portfolios, and find that term structure variables that predict bond returns also predict the returns of bond-like stocks (where bond-like stocks are defined as stocks with bond-like characteristics, such as low volatility). In contrast to the work of Baker and Wurgler (2012), this paper examines the predictability of stock portfolios sorted by bond risk (beta). This is done to determine whether time variation in bond risk premia is driving predictability, a question that is naturally addressed by examining the predictability of assets with variation in bond risk.[27]

5.1 Predictability Theory and Evidence

A natural conjecture is that term structure variables capture information about bond risk premia, and that high bond risk stocks will be more predictable than low bond risk stocks. For example, consider the expected return of an arbitrary asset under the stock-bond CAPM,

$$E[r_i] = \beta\pi_s Cov[r_i, r_s] + \beta\pi_b Cov[r_i, r_b]. \tag{8}$$

Consider two assets, the first with a high exposure to bond risk ($Cov[r_i, r_b]$), and the second with zero exposure to bond risk. Given a ceteris paribus increase in bond risk premia (equal to $\beta\pi_b Var[r_b]$), the expected return of the first asset will increase, while the expected return of the second will not. Then, expected returns of high bond-risk assets will be more sensitive to changes in bond risk premia than those of low bond-risk assets, and high bond-risk asset returns may be more predictable based on variables that capture information about bond risk premia.

However, it remains possible that expected returns are, in general, positively correlated. Variables that predict the returns of high bond-risk assets may also predict the returns of low bond-risk assets, which may have exposure to other risks (e.g. stock and bond expected returns may be positively correlated). Indeed, Cochrane and Piazzesi (2005) find that term structure variables predict stock returns as well as bond returns, while Fama and French (1989) find that stock and bond expected returns share common components. For this reason, it is not clear, a priori, whether high bond-risk assets should be more predictable

[27] There is also an important methodological difference between the approach of this paper and that of Baker and Wurgler (2012). Baker and Wurgler (2012) examine regressions of the form $r_{p,t} = \alpha + \beta r_{m,t} + \gamma CP_{t-1}$, where r_p is the portfolio return and r_m is the market return. This paper omits the $r_{m,t}$ term. This difference is meaningful because aggregate stock returns appear to be predictable by term structure variables (see Cochrane and Piazzesi (2005) and the evidence presented in this paper), so stock and bond predictability may overlap.

than low bond-risk assets when using term structure-based predictive variables.

Table 11 reports the results of predictive regressions of excess stock and bond returns on the term structure factor of Cochrane and Piazzesi (2005) (CP), which is known to be an excellent predictor of excess bond returns.[28] Results are reported using corporate bond beta-sorted stock portfolios from 1988-2011, government bond beta-sorted stock portfolios from 1967-2011, the value-weighted stock index, and several bond indices. Consistent with the results of Cochrane and Piazzesi (2005), the CP factor predicts bond returns.[29] Also, as demonstrated by Cochrane and Piazzesi (2005), the CP factor predicts stocks returns from 1967-2011. However, from 1988-2011, the CP factor is not a significant predictor of stock returns. This may be attributable to low power, as this is a short period to examine return predictability.

Corporate and long-term government bond returns, as well as equal-weighted returns of the bond risk-sorted stock portfolios, are often strongly predictable in the 1988-2011 subsample. In this sample, parameter estimates are similar across the bond risk portfolios (for example, the equal-weighted corporate bond portfolios parameter estimates are 4.46, 4.17, 4.65, and 5.00.) and always significant. Consistent with this, I find no evidence of high-low hedge portfolio (e.g. EW4-EW1) predictability (not tabulated). This suggests that, in the cross section of stocks, term structure-based predictability does not vary with bond risk exposure. In the 1967-2011 sample, every bond risk-sorted portfolio exhibits at least marginal evidence of predictability, and all bond returns are predictable. Again, in the cross section, predictability does not appear to vary with bond risk exposure. Overall, the CP factor is often informative about expected returns, but CP-based predictability is not confined to high bond risk assets.

[28] In this paper, the CP factor is constructed using data described in Gurkaynak, Sack, and Wright (2007).

[29] Somewhat surprisingly, the CP factor is not significant when predicting the average excess return of two-, three-, four-, and five-year discount bonds from 1988-2011. This likely reflects some instability in the parameters of the predictive relation, as the CP factor is generated using data from the full sample (1967-2011).

Next, I examine CP-based predictability using sequentially-sorted stock portfolios (sorted by bond risk then stock market risk, and the reverse) in Table 12. Table 12 does not reveal any obvious relation between bond risk and CP-based predictability. However, Table 12 suggests that CP-based predictability is related to stock risk. The returns of high stock-risk stocks (the SB4 portfolios) appear to be largely unpredictable. However, there is often strong evidence of predictability for low stock-risk stocks (the SB1 portfolios). The stronger statistical evidence of predictability in low stock-risk stocks appears to be due to both higher parameter point estimates (which tend to decrease as stock risk increases) and lower standard errors (which tend to increase with stock risk).

5.2 Interpreting the Predictability Results

Overall, the results suggest that term structure-based predictability is not limited to assets with high exposure to bond risk. Indeed, bond risk appears to be largely unrelated to predictability, while stock risk appears to be negatively related to predictability. Therefore, the returns of low consumption-risk assets appear to be at least as predictable, and sometimes more predictable, than the returns of high consumption-risk assets (importantly, the consumption growth evidence presented in Tables 1 and 2 suggests that both stock and bond risk are positively related to consumption risk). There are several potential explanations for this result, although it is difficult to reconcile the results with term structure-based predictability caused by variation in risk premia.

First, assets with low stock or bond risk may have time-varying betas, so that these assets have high risk exposures when risk premia are high. However, I find little evidence to suggest that time-varying betas play an important role. For example, the low stock-risk quartile of stocks has an average stock beta of 0.24. This beta declines, on average, by 0.03 when CP increases by one standard deviation. This relation, although not significant, points in the wrong direction for a time-varying beta explanation of the high predictability of low stock-risk stocks. I find that time-varying betas likely play, if anything, a small role in understanding the predictability of low bond-risk stock portfolios.[30]

Second, the broad evidence of term structure-based predictability may be explained by positively correlated cross-asset risk premia (i.e. expected stock and bond returns tend to be high at the same time). However, low-risk assets are often at least as predictable as high-risk assets, even when controlling for other risks (e.g. predictability tends to decrease with stock risk, even when controlling for bond risk). This suggests that much of return predictability is not driven by time variation in risk premia, at least as prescribed by the stock-bond CAPM (see Equation 8). While an alternative specification of risk and expected returns could reverse this conclusion, it seems unlikely that high risk assets under the stock-bond CAPM will turn out to be low risk assets under an alternative specification (and the reverse), which is required to explain the negative relation between stock risk and CP-based predictability evident in Table 12. Overall, it is difficult to explain the predictability results within a purely risk-based framework. Of course, this does not rule out time-varying risk premia playing some role in return predictability.[31]

Term structure-based return predictability may be related to changes in sentiment or periodic flight-to-quality episodes, where investors become more willing to hold government bonds and high-quality stocks (see Baker and Wurgler (2006, 2012)). As investors become more willing to hold government bonds and high-quality stocks, the expected returns of these assets will decline (and the reverse). Interestingly, the results suggest that, under this interpretation, high-quality stocks are better defined as having low exposure to stock

[30] Using both the corporate and government bond sample, I find that low bond-risk stock portfolios have a negative exposure to bond risk, which tends to remain negative when stock and bond risk premia are high. Then, a high CP factor (which positively predicts risk premia) should negatively predict the returns of low bond-risk stock portfolios. However, the relation is positive.

[31] This paper is focused on bond risk, and therefore examines a term structure-based predictive variable that is known to predict bond returns. In untabulated results, I find no easily interpretable pattern between dividend yield-based predictability and consumption risk. Predictability associated with the consumption-wealth ratio (Lettau and Ludvigson (2001b,a)) tends to be strongest for high stock- and/or bond-risk stocks, which suggests that variation in risk premia is driving such predictability.

market risk than high exposure to bond risk.

6 Conclusion

In this paper, I find that bond returns positively predict consumption growth, even after controlling for equity returns and labor income growth. This suggests that bonds serve as an important and timely proxy for a relevant ICAPM state variable (likely human capital returns). The ICAPM says that the covariance between an asset's return and a relevant state variable is a priced risk. Therefore, assets with high bond betas should earn high returns. I empirically verify this prediction. The relation is most apparent when focusing on large stocks, in recent years (after 1988), and when accounting for time variation in risk exposures.

Bond beta can be traced to consumption risk. For this reason, the nonzero returns of bond risk-sorted hedge portfolios can be easily understood as compensation for risk. This contrasts with commonly-used risk adjustment factors (i.e. the three factors of Fama and French (1996)), where interpretation of factor returns remains controversial. I attempt to explain well-known anomalies with bond risk. There is no evidence that the size and value premiums can be explained by bond risk, although bond risk can partially explain momentum profits and the surprisingly flat cross-sectional relation between stock index beta and returns.

I find that, when using term structure variables known to predict bond risk premia, the returns of high bond-risk stocks are no more predictable than those of low bond-risk stocks (although both are often predictable). Somewhat surprisingly, the returns of low stock-risk stocks are often strongly predictable, while high stock-risk stock returns are often mostly unpredictable. These results suggest that term structure-based return predictability is not solely driven by time variation in risk premia.

More generally, the results suggest that bonds are an important part of the market portfolio, bond risk is an important part of market portfolio and consumption risk, and commonly-used risk-adjustment models, such as the three-factor model of Fama and French (1996), do not fully capture consumption risk.

References

Baele, L., G. Bekaert, K. Inghelbrecht, 2010. The Determinants of Stock and Bond Return Comovements. Review of Financial Studies 23(6), 2374–2428.

Baker, M., B. Bradley, J. Wurgler, 2011. Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly. Financial Analysts Journal 67.

Baker, M., J. Wurgler, 2006. Investor Sentiment and the Cross-Section of Stock Returns. Journal of Finance 61(4), 1645–1680.

——— , 2012. Comovement and Predictability Relationships Between Bonds and the Cross-Section of Stocks. Review of Asset Pricing Studies 2(1), 1–31.

Bansal, R., A. Yaron, 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. Journal of Finance 59(4), 1481–1509.

Beeler, J., J. Y. Campbell, 2012. The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment. Critical Finance Review 1(1), 141–182.

Bekaert, G., E. Engstrom, S. R. Grenadier, 2010. Stock and Bond Returns with Moody Investors. Journal of Empirical Finance 17, 867–894.

Bekaert, G., E. Engstrom, Y. Xing, 2009. Risk, Uncertainty, and Asset Prices. Journal of Financial Economics 91, 59–82.

Black, F., 1972. Capital Market Equilibrium with Restricted Borrowing. The Journal of Business 45, 444–455.

Breeden, D. T., 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities. Journal of Financial Economics 7(3), 265–296.

Breeden, D. T., M. R. Gibbons, R. H. Litzenberger, 1989. Empirical Tests of the Consumption-Oriented CAPM. Journal of Finance 44(2), 231–262.

Campbell, J. Y., J. H. Cochrane, 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. Journal of Political Economy 107(2), 205–252.

Campbell, J. Y., G. N. Mankiw, 1990. Permanent Income, Current Income, and Consumption. Journal of Business & Economic Statistics 8(3), 265–279.

Campbell, J. Y., A. Sunderam, L. M. Viceira, 2013. Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds. Working Paper.

Cochrane, J. H., 1994. Permanent and Transitory Components of GNP and Stock Prices. Quarterly Journal of Economics 109(1), 241–265.

——— , 2001. Asset Pricing. Princeton University Press, Princeton.

Cochrane, J. H., M. Piazzesi, 2005. Bond Risk Premia. American Economic Review 95(1), 138–160.

Connolly, R., C. Stivers, L. Sun, 2005. Stock Market Uncertainty and the Stock-Bond Return Relation. Journal of Financial and Qunatitative Analysis 40(1), 161–194.

d'Addona, S., A. H. Kind, 2006. International Stock-Bond Correlations in a Simple Affine Pricing Model. Journal of Banking & Finance 30, 2747–2765.

Epstein, L. G., S. E. Zin, 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57(4), 937–970.

——— , 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. Journal of Political Economy 99(2), 263–286.

Fama, E. F., K. R. French, 1989. Business Conditions and Expected Returns on Stocks and Bonds. Journal of Financial Economics 25, 23–49.

——— , 1992. The Cross-Section of Expected Stock Returns. Journal of Finance 47(2), 427–465.

——— , 1993. Common Risk Factors in the Returns on Stocks and Bonds. Journal of Financial Economics 33, 3–56.

——— , 1996. Multifactor Explanations of Asset Pricing Anomalies. Journal of Finance 51(1), 55–84.

——— , 2004. The Capital Asset Pricing Model: Theory and Evidence. Journal of Economic Perspectives 18, 25–46.

Frazzini, A., L. H. Pedersen, 2013. Betting Against Beta. Working Paper.

Gurkaynak, R. S., B. Sack, J. H. Wright, 2007. The U.S. Treasury Yield Curve: 1961 to the Preset. Journal of Monetary Economics 54(7), 2291–2304.

Hall, R. E., 1978. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. Journal of Political Economy 86(6), 971–987.

Harvey, C. R., 1988. The Real Term Structure and Consumption Growth. Journal of Financial Economics 22, 305–333.

Jagannathan, R., Z. Wang, 1996. The Conditional CAPM and the Cross-Section of Expected Returns. Journal of Finance 51(1), 3–53.

Jegadeesh, N., 1990. Evidence of Predictable Behavior of Security Returns. Journal of Finance 45(3), 881–898.

Jegadeesh, N., S. Titman, 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. Journal of Finance 48(1), 65–91.

——— , 2002. Cross-Sectional and Time-Series Determinants of Momentum Returns. Review of Financial Studies 15(1), 143–157.

Koijen, R. S., H. Lustig, S. Van Nieuwerburgh, 2012. The Cross-Section and Time-Series of Stock and Bond Returns. Working Paper.

Lakonishok, J., A. Shleifer, R. W. Vishny, 1994. Contrarian Investment, Extrapolation, and Risk. Journal of Finance 49(5), 1541–1568.

Lehmann, B., 1990. Fads, Martingales, and Market Efficiency. Quarterly Journal of Economics 105(1), 1–28.

Lettau, M., S. C. Ludvigson, 2001a. Consumption, Aggregate Wealth, and Expected Stock Returns. Journal of Finance 56(3), 815–849.

——— , 2001b. Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying. Journal of Political Economy 109(6), 1238–1287.

——— , 2005. Expected Returns and Expected Dividend Growth. Journal of Financial Economics 76(3), 583–626.

Lettau, M., J. A. Wachter, 2011. The Term Structures of Equity and Interest Rates. Journal of Financial Economics 101, 90–113.

Lewellen, J., S. Nagel, 2006. The Conditional CAPM does not Explain Asset-Pricing Anomalies. Journal of Financial Economics 82(2), 289–314.

Liew, J., M. Vassalou, 2000. Can Book-to-Market, Size and Momentum be Risk Factors that Predict Economic Growth?. Journal of Financial Economics 57, 221–245.

Merton, R. C., 1973. An Intertemporal Capital Asset Pricing Model. Econometrica 41(5), 867–887.

Parker, J. A., C. Julliard, 2005. Consumption Risk and the Cross Section of Expected Returns. Journal of Political Economy 113(1), 185–222.

Roll, R., 1977. A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. Journal of Financial Economics 4(2), 129–176.

Shanken, J., 1987. Multivariate Proxies and Asset Pricing Relations: Living with the Roll Critique. Journal of Financial Economics 18(1), 91–110.

——— , 1990. Intertemporal Asset Pricing: An Empirical Investigation. Journal of Econometrics 45(1-2), 99–120.

Stambaugh, R. F., 1982. On the Exclusion of Assets from Tests of the Two-Parameter Model: A Sensitivity Analysis. Journal of Financial Economics 10, 237–268.

Weil, P., 1990. Nonexpected Utility in Macroeconomics. Quarterly Journal of Economics 105(1), 29–42.

Yogo, M., 2006. A Consumption-Based Explanation of Expected Stock Returns. Journal of Finance 61(2), 539–580.

Table 1: Predicting One-Year Consumption Growth

Panel 1: Consumption Growth

Dependent	R^2	Intercept	$r_{s,t}$	Δl_t	$-\Delta BAA_t$	$-\Delta UST_t$	Δc_t
$\Delta c_{t+1,t+4}$	0.179	2.219***	0.513***				0.735***
		0.199	0.121				0.175
$\Delta c_{t+1,t+4}$	0.124	2.209***		0.228*			0.713***
		0.208		0.137			0.225
$\Delta c_{t+1,t+4}$	0.193	2.213***			0.546***		0.818***
		0.197			0.091		0.181
$\Delta c_{t+1,t+4}$	0.162	2.243***				0.436***	0.933***
		0.202				0.094	0.196

Panel 2: Ex-Durable Consumption Growth

Dependent	R^2	Intercept	$r_{s,t}$	Δl_t	$-\Delta BAA_t$	$-\Delta UST_t$	Δcxd_t
$\Delta cxd_{t+1,t+4}$	0.224	2.076***	0.253***				0.666***
		0.144	0.087				0.140
$\Delta cxd_{t+1,t+4}$	0.197	2.071***		-0.017			0.735***
		0.148		0.119			0.170
$\Delta cxd_{t+1,t+4}$	0.234	2.073***			0.281***		0.716***
		0.144			0.067		0.141
$\Delta cxd_{t+1,t+4}$	0.220	2.090***				0.185***	0.763***
		0.146				0.072	0.145

Notes - Table reports regressions of annual percent change in per capita real consumption (Δc) and per capita real ex-durable consumption (Δcxd) on the lagged value-weighted stock return (r_s), lagged change in per capita labor income (Δl), the negative of the lagged BAA corporate bond and ten-year treasury yield change ($-\Delta BAA$ and $-\Delta UST$, respectively), and lagged consumption growth. All explanatory variables are standardized. Data is quarterly and spans 1953-2011. For each regression, estimated coefficients are reported above Newey-West standard errors. ***,**,* indicates a significant slope parameter at the 1%, 5%, and 10% level, respectively.

Table 2: Predicting Consumption Growth, Multiple Horizons

Panel 1: Consumption Growth

Dependent	R^2	Intercept	$r_{s,t}$	Δl_t	$-\Delta BAA_t$	Δc_t
Δc_{t+1}	0.152	0.545	0.134***	0.212***	0.163***	-0.006
		0.050	0.048	0.067	0.046	0.111
$\Delta c_{t+1,t+2}$	0.249	1.091	0.248***	0.301***	0.268***	0.230**
		0.083	0.074	0.099	0.067	0.108
$\Delta c_{t+1,t+3}$	0.165	1.643	0.267***	0.274***	0.337***	0.280
		0.141	0.098	0.088	0.078	0.200
$\Delta c_{t+1,t+4}$	0.141	2.199	0.352***	0.298**	0.405***	0.229
		0.202	0.126	0.133	0.097	0.230
$\Delta c_{t+1,t+8}$	0.059	4.445	0.136	0.368*	0.553***	0.074
		0.451	0.184	0.206	0.181	0.316
$\Delta c_{t+1,t+12}$	0.040	6.714	0.122	0.405*	0.523**	0.168
		0.692	0.206	0.230	0.216	0.304

Panel 2: Ex-Durable Consumption Growth

Dependent	R^2	Intercept	$r_{s,t}$	Δl_t	$-\Delta BAA_t$	Δcxd_t
Δcxd_{t+1}	0.226	0.514	0.074**	0.113***	0.086***	0.137***
		0.031	0.035	0.035	0.029	0.041
$\Delta cxd_{t+1,t+2}$	0.225	1.032	0.114**	0.132**	0.155***	0.267***
		0.060	0.055	0.067	0.051	0.073
$\Delta cxd_{t+1,t+3}$	0.204	1.558	0.136*	0.092	0.162***	0.418***
		0.097	0.074	0.083	0.060	0.109
$\Delta cxd_{t+1,t+4}$	0.180	2.089	0.206**	0.071	0.204***	0.469***
		0.142	0.091	0.100	0.073	0.145
$\Delta cxd_{t+1,t+8}$	0.075	4.226	0.148	0.070	0.219	0.528*
		0.351	0.148	0.174	0.141	0.288
$\Delta cxd_{t+1,t+12}$	0.052	6.384	0.212	0.062	0.150	0.583
		0.569	0.170	0.206	0.173	0.368

Notes - Table reports regressions of percent change in per capita real consumption (Δc) and per capita real ex-durable consumption (Δcxd) on the lagged value-weighted stock return (r_s), lagged change in real per capita labor income (Δl), negative of the lagged BAA corporate bond yield change (-ΔBAA), and lagged consumption growth. All explanatory variables are standardized. Data is quarterly and spans 1953-2011. For each regression, estimated coefficients are reported above Newey-West standard errors. ***,**,* indicates a significant slope parameter at the 1%, 5%, and 10% level, respectively.

Table 3: Predicting Labor Income and Unemployment Growth

Panel 1: Predicting Labor Income Growth

Dependent	R-squared	Intercept	$r_{s,t}$	$-\Delta BAA_t$	Δl_t
$\Delta l_{t+1,t+4}$	0.047	1.672	0.238*	0.371**	-0.062
		0.236	0.144	0.150	0.189
$\Delta l_{t+1,t+8}$	0.036	3.343	-0.082	0.645**	0.068
		0.481	0.234	0.278	0.303
$\Delta l_{t+1,t+12}$	0.017	5.049	0.088	0.517	0.046
		0.749	0.271	0.327	0.320

Panel 2: Predicting Unemployment Growth

Dependent	R-squared	Intercept	$r_{s,t}$	$-\Delta BAA_t$	Δl_t	UR_t
$\Delta UR_{t+1,t+4}$	0.252	0.054	-0.268***	-0.098**	-0.111***	-0.332***
		0.086	0.062	0.053	0.065	0.091
$\Delta UR_{t+1,t+8}$	0.286	0.063	-0.154	-0.277***	-0.117**	-0.749***
		0.186	0.094	0.092	0.107	0.137
$\Delta UR_{t+1,t+12}$	0.316	0.074	-0.002	-0.302***	-0.160	-1.080***
		0.254	0.098	0.108	0.144	0.164

Notes - Table reports regressions of percent real per capital labor income growth (excluding income derived from financial assets) and unemployment rate change on the lagged value-weighted stock return (r_s), lagged change in per capita labor income (r_l), and the negative of the lagged corporate bond yield change (ΔBAA). Unemployment rate regressions also include the lagged unemployment rate as an explanatory variable. All explanatory variables are standardized. Data is quarterly and spans 1947-2011 (labor) and 1948-2011 (unemployment). For each regression, estimated coefficients are reported above Newey-West standard errors. ***,**,* indicates significance slope parameters at the 1%, 5%, and 10% level, respectively.

Table 4: Predicting Corporate Bond Beta

		Pearson		Spearman	
Size Quartile	Variable	β_D	β_M	β_D	β_M
ALL	β_R	0.364	0.012	0.391	0.043
	β_M	0.118		0.150	
1 (Small)	β_R	0.235	0.006	0.243	0.021
	β_M	0.098		0.117	
2	β_R	0.444	-0.002	0.475	0.034
	β_M	0.112		0.145	
3	β_R	0.476	0.011	0.519	0.061
	β_M	0.147		0.187	
4 (Large)	β_R	0.502	0.085	0.557	0.140
	β_M	0.224		0.277	

Notes - Table reports correlations between historical and realized corporate bond beta. For each calendar year from 1988-2011, corporate bond betas are calculated using daily data from the current calendar year (β_R), daily data from the prior two years (β_D), and monthly data from the prior five years (β_M). Table reports pooled correlations. Results are reported for all stocks and size quartiles (using NYSE quartile breakpoints).

Table 5: Corporate Bond Beta-Sorted Portfolio Descriptive Statistics

Portfolio	EW Ret.	VW Ret.	Cap. Share	ME	MB	PRET6	EW β_{BAA}	VW β_{BAA}
ALL	0.726	0.569		12.434	0.526	7.411	0.15	0.75
1	0.738	0.307	0.158	12.136	0.497	8.072	-1.47	-2.05
2	0.744	0.498	0.238	12.593	0.511	6.358	-0.40	-0.18
3	0.765	0.663	0.297	12.696	0.522	6.743	0.69	1.09
4	0.658	0.726	0.307	12.311	0.558	8.472	1.81	2.41
CORP	-0.008	0.419					3.28	4.46

Notes - Table reports time-series averages of descriptive statistics of corporate bond beta-sorted portfolios. I report equal- and value-weighted returns, lagged log market capitalization (ME), lagged log market-to-book (MB), prior return over months $t-6$ to $t-2$ (PRET6), and equal- and value-weighted realized corporate bond beta. Realized corporate bond betas are calculated annually using portfolios formed at the end of the prior calendar year. All other statistics are calculated monthly (using portfolios reformed monthly). Results are reported for corporate bond beta quartile portfolios 1-4, all stocks, and the CORP hedge portfolio (4-1). Data spans 1988-2011.

Table 6: Corporate Bond Beta-Sorted Portfolio Time-Series Regressions

Panel 1: All Stocks

Equal-Weighted					Value-Weighted				
Int.	MKT	SMB	HML	WML	Int.	MKT	SMB	HML	WML
0.125	-0.371				0.711***	-0.528			
0.210	0.077				0.241	0.074			
0.284	-0.386	-0.297	-0.388		0.764***	-0.464	-0.416	-0.060	
0.201	0.068	0.064	0.088		0.235	0.074	0.088	0.108	
0.123	-0.349	-0.336	-0.348	0.122	0.550**	-0.415	-0.469	-0.007	0.163
0.228	0.065	0.067	0.093	0.047	0.254	0.069	0.095	0.106	0.055

Panel 2: NYSE Size Quartiles 2-4

Equal-Weighted					Value-Weighted				
Int.	MKT	SMB	HML	WML	Int.	MKT	SMB	HML	WML
0.313	-0.391				0.721***	-0.536			
0.227	0.078				0.246	0.075			
0.430*	-0.376	-0.338	-0.258		0.766***	-0.471	-0.404	-0.039	
0.224	0.073	0.075	0.100		0.241	0.075	0.089	0.109	
0.264	-0.338	-0.379	-0.217	0.126	0.554**	-0.422	-0.456	0.013	0.161
0.254	0.071	0.081	0.105	0.055	0.260	0.071	0.096	0.107	0.056

Panel 3: NYSE Size Quartiles 3 and 4

Equal-Weighted					Value-Weighted				
Int.	MKT	SMB	HML	WML	Int.	MKT	SMB	HML	WML
0.555**	-0.416				0.758***	-0.544			
0.235	0.079				0.258	0.077			
0.624***	-0.377	-0.338	-0.123		0.790***	-0.477	-0.382	-0.009	
0.234	0.077	0.081	0.105		0.254	0.077	0.092	0.112	
0.420	-0.330	-0.388	-0.072	0.154	0.576**	-0.428	-0.434	0.044	0.162
0.261	0.071	0.089	0.105	0.061	0.272	0.074	0.100	0.110	0.058

Panel 4: NYSE Size Quartile 4

Equal-Weighted					Value-Weighted				
Int.	MKT	SMB	HML	WML	Int.	MKT	SMB	HML	WML
0.734***	-0.498				0.785***	-0.571			
0.272	0.079				0.284	0.080			
0.766***	-0.449	-0.301	-0.026		0.802***	-0.502	-0.360	0.030	
0.273	0.079	0.087	0.107		0.282	0.082	0.097	0.116	
0.590**	-0.409	-0.344	0.018	0.133	0.596**	-0.455	-0.410	0.081	0.156
0.295	0.078	0.093	0.110	0.055	0.300	0.080	0.105	0.115	0.058

Notes - Table reports regressions of corporate bond-beta sorted hedge portfolio returns (CORP) on traded factors. MKT, SMB, and HML are the three factors from Fama and French (1996). WML is calculated as the equal-weighted portfolio return of past winners (top decile) minus losers (low decile), using the prior return from month $t-6$ to $t-2$. Portfolios are formed monthly. Parameter estimates are reported above heteroskedasticity-robust standard errors. ***, **, * indicates intercepts significant at the 1%, 5%, and 10% levels, respectively (using two-sided tests). Data spans 1988-2011.

Table 7: Government Bond Beta-Sorted Portfolio Time-Series Regressions

Int.	MKT	HML	SMB
	1964-2011		
0.178	-0.154**	0.020	-0.188**
0.189	0.060	0.081	0.079
	1964-1987		
-0.214	0.147**	-0.151	0.101
0.188	0.059	0.096	0.092
	1988-2011		
0.705***	-0.559***	0.085	-0.426***
0.253	0.077	0.094	0.093

Notes - Table reports regressions of government bond-beta sorted hedge portfolio value-weighted returns on contemporaneous traded factors. MKT, SMB, and HML are the three factors from Fama and French (1996). Portfolios are formed monthly. Parameter estimates are reported above heteroskedasticity-robust standard errors. ***, **, * indicates intercepts significant at the 1%, 5%, and 10% levels, respectively (using two-sided tests).

Table 8: Government and Corporate Bond Beta-Sorted Portfolio Time-Series Regressions, with Time-Varying Risk

Weighting	Int.	Int. D	MKT	MKT D	HML	HML D	SMB	SMB D
Government Bond-Beta Hedge Portfolio, 1964-2011								
VW	0.820***	-1.279***	-0.320***	0.445***	-0.087	0.244	-0.217*	0.044
	0.309	0.361	0.083	0.103	0.121	0.153	0.129	0.156
EW	0.370	-0.651**	-0.261***	0.334***	-0.431***	0.621***	-0.200*	0.128
	0.264	0.309	0.077	0.093	0.114	0.138	0.118	0.132
Government Bond-Beta Hedge Portfolio, 1964-1987								
VW	0.217	-0.850**	0.131*	0.003	-0.240*	0.185	0.187	-0.146
	0.299	0.378	0.079	0.118	0.136	0.190	0.137	0.187
EW	-0.162	-0.181	0.222***	-0.067	-0.327***	0.168	0.145	-0.042
	0.212	0.264	0.060	0.081	0.082	0.127	0.102	0.137
Government Bond-Beta Hedge Portfolio, 1988-2011								
VW	1.114***	-1.030**	-0.727***	0.541***	0.119	0.143	-0.457***	0.096
	0.359	0.495	0.090	0.173	0.141	0.198	0.141	0.190
EW	0.908***	-1.072***	-0.797***	0.752***	-0.099	0.358**	-0.383***	0.137
	0.292	0.381	0.072	0.126	0.119	0.164	0.130	0.160
Corporate Bond-Beta Hedge Portfolio, 1988-2011								
VW	1.203***	-1.097***	-0.625***	0.519***	-0.119	0.317	-0.460***	0.144
	0.341	0.443	0.092	0.126	0.169	0.205	0.126	0.164
EW	0.733**	-1.019***	-0.608***	0.664***	-0.453***	0.463	-0.335***	0.183
	0.302	0.351	0.086	0.109	0.166	0.182	0.116	0.131
Government Bond-Beta Hedge Portfolio, 1988-2011, Implementable Trading Strategy								
VW	1.083***	-1.339***	-0.640***	0.571***	0.097	0.264	-0.481***	0.227
	0.323	0.459	0.087	0.146	0.115	0.191	0.132	0.181
EW	0.643**	-0.748*	-0.610***	0.442***	-0.142	0.357	-0.369***	0.139
	0.285	0.393	0.081	0.136	0.106	0.161	0.107	0.142

Notes - Table reports regressions of corporate and government bond-beta sorted hedge portfolios (CORP and GOV) on contemporaneous traded factors and an intercept interacted with a dummy variable. MKT, SMB, and HML are the three factors from Fama and French (1996). The dummy variable takes the value one if the lagged difference in the historical covariance of the extreme portfolios is less than the full-sample median. Regressions take the form: $r_t = \alpha + \alpha_D * D + (\beta_{mkt} + \beta_{mkt,D} * D) * MKT + (\beta_{hml} + \beta_{hml,D} * D) * HML + (\beta_{smb} + \beta_{smb,D} * D) * SMB + \epsilon$. Results are also reported for an implementable trading strategy, where the dummy variable is calculated recursively. Portfolios are formed monthly. Parameter estimates are reported above heteroskedasticity-robust standard errors. ***, **, * indicates intercepts significant at the 1%, 5%, and 10% levels, respectively (using two-sided tests).

Table 9: Regressing SMB, HML, and Momentum Returns on Contemporaneous Factors

Panel 1: SMB Returns

α	MKT	HML	SMB	CORP	R^2
0.077	0.179				0.056
0.187	0.041				
0.255	0.047			-0.250	0.150
0.196	0.052			0.064	

Panel 2: HML Returns

α	MKT	HML	SMB	CORP	R^2
0.349*	-0.175				0.060
0.191	0.060				
0.324	-0.157			0.036	0.063
0.201	0.059			0.064	

Panel 3: WML Returns

α	MKT	HML	SMB	CORP	R^2
1.229***	-0.185				0.017
0.351	0.115				
1.318***	-0.300	-0.327	0.324		0.086
0.339	0.115	0.216	0.215		
1.024***	-0.122	-0.304	0.484	0.385	0.143
0.369	0.088	0.207	0.206	0.143	

Notes - Table reports regressions of SMB, HML, and momentum returns on contemporaneous traded factors. MKT, SMB, and HML are the three factors from Fama and French (1996). Construction of CORP is described in the text. Momentum returns are calculated as the equal-weighted portfolio return of past winners (top decile) minus past losers (low decile), using the prior return from month $t-6$ to $t-2$. Portfolios are formed monthly. Parameter estimates are reported above heteroskedasticity-robust standard errors. ***, **, * indicates intercepts significant at the 1%, 5%, and 10% levels, respectively (using two-sided tests). Data spans 1988-2011.

Table 10: Stock Beta Hedge Portfolio Returns Regressed on Contemporaneous Factors

Panel 1: CORP Excluded

	MKT Model		FF3 Model			
	α	MKT	α	MKT	HML	SMB
$\beta 1$ (Low)	0.437	0.486	0.291	0.469	0.328	0.417
	0.136	0.042	0.114	0.037	0.058	0.032
$\beta 2$	0.363	0.804	0.186	0.778	0.390	0.530
	0.139	0.041	0.090	0.030	0.048	0.039
$\beta 3$	0.193	1.089	0.017	1.037	0.358	0.649
	0.143	0.041	0.071	0.022	0.036	0.042
$\beta 4$ (High)	-0.305	1.635	-0.352	1.473	-0.054	0.857
	0.204	0.052	0.119	0.042	0.062	0.064
$\alpha_1 = \alpha_4$	0.000			0.000		
$\alpha = 0$	0.002			0.007		

Panel 2: CORP Included

	MKT, CORP Model			FF3, CORP Model				
	α	MKT	CORP	α	MKT	HML	SMB	CORP
$\beta 1$ (Low)	0.447	0.479	-0.013	0.223	0.510	0.333	0.454	0.088
	0.136	0.043	0.040	0.103	0.031	0.053	0.033	0.032
$\beta 2$	0.383	0.790	-0.027	0.108	0.825	0.397	0.572	0.101
	0.140	0.046	0.040	0.081	0.027	0.043	0.038	0.024
$\beta 3$	0.274	1.029	-0.115	-0.012	1.055	0.360	0.665	0.038
	0.147	0.043	0.045	0.071	0.022	0.035	0.043	0.020
$\beta 4$ (High)	-0.037	1.436	-0.377	-0.215	1.389	-0.065	0.782	-0.179
	0.201	0.052	0.062	0.122	0.035	0.057	0.062	0.045
$\alpha_1 = \alpha_4$		0.013				0.005		
$\alpha = 0$		0.012				0.065		

Notes - Table reports estimated alphas, factor loadings, and associated tests of regressions of stock index beta-sorted portfolio equal-weighted returns on factor models. Factors are MKT, HML, SMB, and CORP. Estimated α and factor loadings are reported above corresponding standard errors. P-values of tests of parameter restrictions are also reported. $\alpha_1 = \alpha_4$ tests for equality of α of portfolios $\beta 1$ and $\beta 4$. $\alpha = 0$ tests for α of all portfolios jointly equal to zero. Data is monthly and spans 1988-2011.

Table 11: Predictive Regressions of Stock, Bond, and Bond Risk-Sorted Stock Portfolio Returns

	CORP (1988-2011)			GOV (1967-2011)		
	Intercept	beta	$\bar{R}^2$	Intercept	beta	$\bar{R}^2$
EW1	0.045	4.459**	0.045	0.067	2.824*	0.023
	0.051	2.236		0.042	1.682	
EW2	0.048	4.170**	0.071	0.064	2.965**	0.043
	0.040	1.737		0.032	1.347	
EW3	0.045	4.647***	0.101	0.057	3.038**	0.048
	0.037	1.610		0.033	1.405	
EW4	0.024	5.001***	0.112	0.049	2.710*	0.030
	0.037	1.628		0.037	1.665	
VW1	0.022	1.964	0.008	0.021	2.789**	0.035
	0.059	2.514		0.034	1.347	
VW2	0.025	3.355*	0.056	0.028	3.267***	0.078
	0.041	1.794		0.028	1.182	
VW3	0.050	2.731*	0.056	0.025	2.868***	0.079
	0.032	1.508		0.022	1.006	
VW4	0.063	2.289	0.034	0.026	2.578*	0.045
	0.032	1.567		0.029	1.353	
MKT	0.038	2.141	0.023	0.022	3.041***	0.073
	0.040	1.711		0.026	1.125	
Corp.	0.013	2.950***	0.235	-0.012	3.745***	0.304
	0.013	0.678		0.011	0.576	
Int. Gov.	0.017	0.646	0.022	0.003	1.518***	0.129
	0.011	0.660		0.008	0.415	
Long-Term Gov.	0.021	2.448***	0.160	-0.009	3.363***	0.240
	0.014	0.766		0.012	0.702	
Discount Bond Avg.	0.013	0.419	0.033	0.000	0.991***	0.157
	0.007	0.375		0.005	0.261	

Notes - Table reports predictive regressions of asset returns on the term structure factor of Cochrane and Piazzesi (2005). Assets are equal- and value-weighted bond beta-sorted stock portfolios, the value-weighted stock market portfolio, corporate bonds, intermediate- and long-term government bonds, and an equal-weighted portfolio consisting of two-, three-, four-, and five-year zero coupon bonds (this portfolio is used to construct the CP factor). Returns are one-year excess returns. Newey-West standard errors (with twelve lags) and the regression adjusted r-squared are also reported.

Table 12: Predictive Regressions of Sequentially-Sorted Stock Portfolio Returns

Panel 1: Sort by Stock then Corporate Bond Beta (1988-2011)

	Equal-Weighted Returns				Value-Weighted Returns			
	SB1	SB2	SB3	SB4	SB1	SB2	SB3	SB4
BB1	6.554***	5.935***	4.463**	4.412	2.520	2.555*	2.285	1.884
	2.062	2.019	2.193	3.236	1.694	1.431	1.639	3.714
	0.149	0.114	0.057	0.020	0.035	0.041	0.021	0.001
BB2	6.089***	4.518***	4.676**	2.836	2.346*	3.316**	4.065**	1.345
	1.795	1.749	1.861	2.677	1.288	1.357	1.735	3.291
	0.162	0.093	0.081	0.012	0.048	0.082	0.076	0.001
BB3	5.378***	4.460**	4.553**	3.974	3.056**	4.649***	3.392**	0.803
	1.499	1.746	1.812	2.683	1.449	1.442	1.610	2.349
	0.186	0.093	0.081	0.027	0.068	0.149	0.064	-0.002
BB4	5.265***	5.280***	4.529**	5.980**	3.520***	2.849*	2.760*	3.575
	1.412	1.999	2.048	2.928	1.090	1.618	1.550	2.545
	0.202	0.096	0.059	0.052	0.144	0.048	0.039	0.027

Panel 2: Sort by Corporate Bond then Stock Beta (1988-2011)

	Equal-Weighted Returns				Value-Weighted Returns			
	SB1	SB2	SB3	SB4	SB1	SB2	SB3	SB4
BB1	5.898***	4.822**	5.014**	3.643	3.069	3.724**	2.093	0.716
	2.146	2.193	2.555	3.519	1.896	1.787	1.998	3.859
	0.105	0.067	0.046	0.011	0.039	0.055	0.013	-0.003
BB2	5.122***	4.674***	3.883**	4.274*	3.458***	3.435**	4.112**	2.923
	1.902	1.783	1.909	2.439	1.330	1.481	1.678	2.708
	0.105	0.093	0.056	0.035	0.085	0.074	0.076	0.020
BB3	5.786***	4.882***	5.069***	4.282*	2.755**	3.884**	3.955***	2.016
	1.803	1.695	1.723	2.183	1.383	1.566	1.532	1.945
	0.141	0.114	0.112	0.045	0.057	0.099	0.104	0.015
BB4	6.239***	4.636***	4.939**	5.867**	3.985***	2.881**	1.238	3.390
	1.386	1.649	2.033	2.834	1.199	1.364	1.510	2.294
	0.242	0.114	0.082	0.059	0.121	0.063	0.008	0.036

Panel 3: Sort by Stock then Government Bond Beta (1967-2011)

	Equal-Weighted Returns				Value-Weighted Returns			
	SB1	SB2	SB3	SB4	SB1	SB2	SB3	SB4
BB1	4.139**	2.879	2.162	1.760	3.426***	3.283***	2.741**	3.560*
	1.610	1.850	1.900	2.140	1.276	1.212	1.242	1.960
	0.071	0.025	0.011	0.004	0.071	0.062	0.040	0.029
BB2	4.850***	3.851***	2.928*	2.309	4.108***	3.732***	4.358***	2.397
	1.223	1.346	1.593	1.963	0.943	1.060	1.293	1.781
	0.133	0.075	0.034	0.011	0.171	0.119	0.117	0.018
BB3	4.699***	3.937***	3.412**	2.256	4.428***	3.374***	3.613***	2.392
	1.211	1.452	1.652	1.948	1.021	1.067	1.282	1.672
	0.150	0.077	0.046	0.013	0.189	0.104	0.088	0.020
BB4	3.784***	3.353**	3.262*	3.431	3.420***	3.704***	3.043**	3.302*
	1.219	1.463	1.890	2.224	0.900	1.117	1.313	1.868
	0.110	0.054	0.035	0.025	0.168	0.119	0.058	0.031

Panel 4: Sort by Government Bond then Stock Beta (1967-2011)

	Equal-Weighted Returns				Value-Weighted Returns			
	SB1	SB2	SB3	SB4	SB1	SB2	SB3	SB4
BB1	5.898***	4.822**	5.014**	3.643	3.069*	3.724**	2.093	0.716
	2.146	2.193	2.555	3.519	1.896	1.787	1.998	3.859
	0.105	0.067	0.046	0.011	0.039	0.055	0.013	-0.003
BB2	5.122***	4.674**	3.883**	4.274*	3.458***	3.435**	4.112**	2.923
	1.902	1.783	1.909	2.439	1.330	1.481	1.678	2.708
	0.105	0.093	0.056	0.035	0.085	0.074	0.076	0.020
BB3	5.786***	4.882***	5.069***	4.282**	2.755**	3.884**	3.955***	2.016
	1.803	1.695	1.723	2.183	1.383	1.566	1.532	1.945
	0.141	0.114	0.112	0.045	0.057	0.099	0.104	0.015
BB4	6.239***	4.636***	4.939**	5.867**	3.985***	2.881**	1.238	3.390
	1.386	1.649	2.033	2.834	1.199	1.364	1.510	2.294
	0.242	0.114	0.082	0.059	0.121	0.063	0.008	0.036

Notes - Table reports predictive regressions of sorted stock portfolios on the term structure factor of Cochrane and Piazzesi (2005). Stock portfolios are sequentially sorted by bond beta (BB) and stock beta (SB). Returns are one-year excess returns. Slope estimates from the predictive regression are reported above Newey-West standard errors (with twelve lags) and adjusted r-squared.