

Seasonally Varying Preferences: Theoretical Foundations for an Empirical Regularity

Mark J. Kamstra

Schulich School of Business, York University

Lisa A. Kramer

University of Toronto

Maurice D. Levi

Sauder School of Business, University of British Columbia

Tan Wang

Sauder School of Business, University of British Columbia;
Shanghai Advanced Institute of Finance, Shanghai Jiao Tong University

We investigate an asset pricing model with preferences cycling between high risk aversion and low EIS in fall/winter and the reverse in spring/summer. Calibrating to consumption data and allowing plausible preference parameter values, we produce returns that match observed equity and Treasury returns across the seasons: risky returns are higher and risk-free returns are lower or stable in fall/winter, and they reverse in spring/summer. Further, risky returns vary more than risk-free returns. A novel finding is that both EIS and risk aversion must vary seasonally to match observed returns. Further, the degree of necessary seasonal change in EIS is small. (*JEL* E44, G11, G12)

We are grateful for the helpful suggestions of Wayne Ferson and an anonymous referee, as well as Murray Carlson, James Choi, Vito Gala, Vincent Glode, Don Wook Lee, Erica Li, Elise Payzan-LeNestour, Monika Piazzesi, Piet Sercu, Jay Shanken, Tim Simin, Kevin Wang, conference participants at the 2013 China International Conference in Finance, the 2013 Frontiers in Finance conference, the 2013 Northern Finance Association conference, the 2012 American Finance Association meetings, the 2012 Leuven Finance Workshop at the National Bank of Belgium, the 2011 Financial Intermediation Research Society conference, the 2011 Queen's University Behavioral Finance conference, the Federal Reserve Bank of Chicago Summer 2009 Workshop on Money, Banking, and Payments, and Finance, the Xiamen University International Symposium on Risk Management and Derivatives, the 2009 Financial Management Association meetings, and seminar participants at Manchester Business School, National University of Singapore, Santa Clara University, the University of California at Davis, the University of California at Santa Barbara, the University of Colorado at Boulder, the University of North Carolina at Chapel Hill, the University of Oregon, the University of Utah, the University of Washington, and York University. We thank the Social Sciences and Humanities Research Council of Canada for financial support. Any remaining errors are our own. A previous version of this paper was titled "Seasonally Varying Preferences: Support from an Asset Pricing Model." Send correspondence to Lisa Kramer, University of Toronto, 105 St. George St., Toronto Ontario Canada, M5S 3E6; telephone 416-978-2496. E-mail: lkramer@rotman.utoronto.ca.

© The Author 2014. Published by Oxford University Press on behalf of The Society for Financial Studies.

All rights reserved. For Permissions, please email: journals.permissions@oup.com

doi:10.1093/rapstu/rau002

Advance Access publication March 12, 2014

A small, but growing, body of literature has emerged suggesting that the influence of seasonal variation in investor mood may lead to systematic seasonal effects in financial markets. For instance, Kamstra, Kramer, and Levi (2003) find that holding-period stock returns are higher during the six-month period when many individuals experience seasonal depression (the fall and winter),¹ varying conditionally across the seasons as much as 12% (annualized), even after controlling for well-known seasonalities, such as the January effect. Kamstra, Kramer, and Levi (forthcoming) show that average returns to U.S. Treasury notes and bonds vary countercyclically relative to the seasonal variation in equity returns, though with a smaller magnitude of variation.² Kamstra, Kramer, and Levi (2003) conjecture that empirical regularities such as these may arise due to seasonality in risk aversion produced by seasonality in mood, with seasonally depressed investors favoring safe securities over risky securities during the periods when they are depressed. The observed seasonal variation in the equity premium is comparable to the magnitude of the mean equity premium itself, making it larger than other established market regularities, including the January effect.

One of our primary findings is that allowing for seasonally varying investor risk aversion, alone, is not capable of generating returns with properties that match those of observed returns. Overall, we show in an equilibrium asset pricing model framework that we are able to match the qualitative and quantitative features of observed risky and risk-free returns only when we allow for both seasonally varying elasticity of intertemporal substitution (EIS) and seasonally varying risk aversion.³ In our model, a representative agent's preferences vary across two seasons: in one season the agent is depressed and has high risk aversion and low EIS and in the other he/she has low risk aversion and high EIS.

¹ Kamstra, Kramer, and Levi (2003) find that daily realized stock returns are significantly lower than average during the fall and significantly higher than average during the winter, leading to higher-than-average holding-period returns for investors who hold stock through the fall and winter.

² Additional papers have explored the influence of seasonal depression on other facets of financial markets, finding largely supportive results. Kamstra et al. (2013) investigate the flow of funds between safe and risky categories of mutual funds and find, controlling for other factors, that there are net flows out of risky funds and into safe funds in fall, and that the patterns reverse in winter, consistent with the seasonal depression hypothesis. Garrett, Kamstra, and Kramer (2005) explore seasonally varying risk aversion in an equilibrium asset pricing model that allows the price of risk to vary through the seasons, and they find evidence consistent with the seasonal depression hypothesis. DeGennaro, Kamstra, and Kramer (2008) study bid-ask spreads; Lo and Wu (2008) and Dolvin, Pyles, and Wu (2009) study analysts' stock earnings forecasts; Dolvin and Pyles (2007) study the underpricing of initial public stock offerings; and Pyles (2009) studies returns to real estate investment trusts. All find evidence consistent with the influence of seasonal depression on markets. Dowling and Lucey (2008) enlarge Kamstra, Kramer, and Levi's (2003) original study to thirty-seven countries and find similar results.

³ Intuitively, allowing a time-varying subjective discount rate may generate similar effects on asset returns as time-varying EIS. Indeed, Maurer (2012) finds several asset pricing puzzles are resolved in a model with shocks to the representative agent's subjective discount rate, which suggests a potentially fruitful avenue for research that seeks to explain asset return seasonality. In our model we employ Epstein and Zin (1989) utility, and so EIS is separated from risk aversion and the discount rate. We show that time-varying risk aversion and EIS go a long way in helping us understand seasonality in the returns to equity and the risk-free asset, leaving to future research the potential benefit of allowing for a seasonally varying discount rate.

The rationale for allowing seasonal variation in investor preferences arises from a combined literature in psychology, medicine, and economics (described in detail in Section 1), which links seasonality in mood with seasonality in willingness to take risk and seasonality in consumption smoothing (i.e., seasonality in EIS). First, many individuals experience severe depression during the fall and winter seasons, when daylight is relatively diminished, and the broader population experiences a milder degree of dampened mood during these seasons. Depressed individuals are more risk averse, implications of which have been well explored in the finance literature.

A second implication of depression which to date has not been explored in the finance literature is greater consumption impulsivity. That is, depression is associated with a tendency to prefer consuming in the present, which naturally results in less smooth consumption over time. In financial terms, this amounts to reduced elasticity of intertemporal substitution. Thus, during seasons when investors generally experience heightened depression, we expect to observe both increased investor risk aversion and reduced investor elasticity of intertemporal substitution. To the best of our knowledge, this paper is the first to consider the influence of seasonal variation in EIS on return seasonality.

The higher EIS and lower risk aversion associated with fall/winter depression may impact returns through the following channels. Reduced risk aversion in a given season translates into reduced demand for risky assets, which in turn results in lower prices for risky assets and higher prices for risk-free assets, and hence higher expected returns for risky assets and lower expected risk-free returns in that season. Furthermore, if an agent's EIS is lower in a particular season, then the agent is less willing to substitute today's consumption for future consumption. To induce the agent to invest more, the prices of all assets must drop, which means that expected asset returns must rise. To be clear, the seasonal variation in EIS induces a change in the agent's willingness to substitute today's consumption for tomorrow's consumption, whereas the effect of a change in risk aversion is to induce a change in the preference for substituting between risky and risk-free assets. Thus, a change in EIS induces the prices of all assets to move in the same direction, whereas a change in risk aversion causes the prices of risky assets and risk-free assets to move in opposite directions. The challenge for us is to see whether allowing for seasonal variation in agents' risk aversion and EIS gives rise to the specific seasonal asset return patterns documented by Kamstra, Kramer, and Levi (2003, forthcoming) and others. We find we can meet that challenge.

Specifically, the model we develop is able to match the following empirical regularities: (1) high risky asset returns during the season when the risk-free returns are low and low risky asset returns during the season when the risk-free returns are high (and hence a large change in the equity premium across seasons) and (2) much greater seasonal variation in risky asset returns than in risk-free asset returns. Whereas previous research on the links among

daylight, depression, and markets focused on the role of risk aversion in helping to explain asset returns, our work highlights the equal importance of EIS in matching moments of returns data. Further, to match those moments, the necessary degree of seasonal variation in both EIS and risk aversion is within ranges that are broadly accepted as reasonable. And notably, the amount of seasonal variation in EIS is very small relative to the degree of seasonal change in risk aversion. A change in EIS by as little as 0.005 (seasonally varying from 1.5 to 1.495, for instance) can influence whether or not statistical tests reject the model's ability to fit observed features of returns. We highlight the importance of this small degree of variation in EIS required to match observed returns: if instead we hold EIS constant across the seasons, while allowing seasonally varying risk aversion, the model is unable to match the signs and magnitudes of patterns in safe versus risky returns across the seasons; the source of the seasonal variation in the risk premium is split evenly between the variation in the equity and risk-free returns in that case, in sharp contrast to the data. Similarly, the model cannot match the observed characteristics of the data when EIS varies seasonally and risk aversion is constant.

In principle, variation in individuals' preferences over time and heterogeneity in preferences across individuals are not controversial notions (see, for instance, Blundell, Browning, and Meghir 1994; Attanasio and Browning 1995; Atkeson and Ogaki 1996; Barsky et al. 1997; Guiso, Sapienza, and Zingales 2013). Further, the consideration of seasonality in preference parameters to capture seasonality in the macroeconomy, as in this paper, is not unprecedented. For example, Chatterjee and Ravikumar (1992) explore the impact of seasonal fluctuation in demand (with a Christmas peak) and in production (with a winter slowdown) on the fit of macromodels, and Braun and Evans (1995) consider the degree of changes in consumption preferences, technology, and government purchases that are required to explain observed seasonal patterns in output and investment.

This paper contributes most closely to the literature that explores time variation in expected returns, arising because of time-varying preferences. Miron (1986) explores the impact of seasonal preference shocks on consumption, income, and asset returns in the context of a rational expectations version of the life-cycle-permanent income model, finding that seasonality in preference shocks can dramatically impact inference. Ferson and Harvey (1992) endogenize a seasonally adjusting marginal rate of consumption with a fixed habit parameter set in reference to past consumption from the same season, facilitating a better match between the strong seasonals in consumption with the relatively weak seasonals in returns. Ferson and Harvey (1993) also consider seasonal taste shocks in linear consumption models and, like Miron (1986) and Ferson and Harvey (1992), exploit consumption data that are not seasonally adjusted. This use of seasonally unadjusted

consumption data is in contrast to much of the consumption-based asset pricing literature, for which the use of seasonally adjusted data is standard.

Several other models of time-varying risk preferences have also focused on habit formation. Brandt and Wang (2003), for instance, consider a habit persistence model and a law of motion for log relative risk aversion, whereby risk aversion is shocked with news about both consumption growth and inflation. Bekaert, Grenadier, and Engstrom (2010) consider stochastic risk aversion as a preference shock, motivated by a habit persistence framework. There is also a significant body of literature that studies the relationship between time-varying risk and risk aversion. A strand of that literature focuses on the empirical observation that equity premia seem to be higher in recessions than in booms (see, for instance, Fama and French 1989). Campbell and Cochrane (1999) build a representative agent model and show that when the representative agent exhibits habit formation, his/her risk aversion is higher at business-cycle troughs than it is at peaks. As a result, equity premia are higher at business-cycle troughs than they are at peaks. Like us, Gordon and St-Amour (2000) adopt a representative agent model that allows for time-variation in preferences. Specifically, they capture “bear” and “bull” market cyclicity by allowing stochastic switches between high and low risk aversion states of the world. Piazzesi (2001) documents and discusses seasonal cross-correlations between consumption growth and asset returns in the context of Gabaix and Laibson’s (2002) paper on resolving the equity premium puzzle.

Although the focus of our primary analysis is on a calibration exercise based on seasonally adjusted consumption data, we also perform the exercise based on consumption data that have not been seasonally adjusted. This allows us to consider whether consumption seasonality or preference seasonality is better capable of explaining the data. Our results on this point are unambiguous: consumption seasonality is not able to explain the seasonal patterns in equity and risk-free asset returns, a result echoing the Ferson and Harvey (1993) finding that consumption seasonality cannot explain the equity premium puzzle. Specifically, we find that the seasonality evident in observed consumption pushes returns in the model to follow a counterfactual cycle, with both risky and risk-free returns becoming higher in the seasons when they are observed to be lower (and vice versa) and with risky returns exhibiting far less seasonal variation than they do in practice. Only when we allow risk preferences to vary seasonally, with lower EIS and higher risk aversion during periods when seasonal depression is most evident, is our model based on seasonally unadjusted consumption data able to match the observed characteristics of risky and risk-free returns.

1. Seasonal Depression

A sizable fraction of the population suffers from a seasonal form of clinical depression known as seasonal affective disorder (SAD); most or, according to

some recent research, perhaps all of the rest of the population experiences a milder condition known as winter blues.⁴ Kasper et al. (1989) explain that the distinction between SAD and the milder condition is not clear cut; they recommend that both conditions be viewed along a continuum of seasonal depression. Harmatz et al. (2000) demonstrate that depression varies across the seasons significantly, even among people who do not suffer from SAD, peaking in the fall/winter seasons, most markedly for women. Kramer and Weber (2012) find a similar result, documenting that depression peaks in the fall/winter seasons even among healthy individuals (regardless of gender).

Medical research has established that among the various possible environmental factors that might cause seasonal depression, length of daylight appears to be the primary cause.⁵ Individuals who suffer from seasonal depression typically begin experiencing depression in the early fall and recover by late spring. The peak point of onset across individuals is around October and the peak point of recovery is around April. We discuss this timing more fully in footnote 8.

As depression is a psychological/medical condition of the sufferers, it can manifest itself in the behavioral characteristics of the sufferers. Previous research in psychology has established a link between depression and both increased risk aversion and reduced consumption smoothing. We discuss each set of research in turn.

Depression is associated with increased risk aversion, including risk of a financial nature.⁶ With a substantial fraction of the population experiencing seasonal depression in the fall and winter months, Kamstra, Kramer, and Levi (2003) conjecture that the proportion of risk-averse investors is higher in those seasons. Risk-averse investors, they argue, begin to shun risky stocks in the fall as the length of day shortens, which has an immediate negative influence on stock prices, contributing to lower contemporaneous returns and higher expected future returns. As the amount of daylight rebounds through the winter months, investors begin to recover from their depression and become more willing to hold risky assets, which, Kamstra, Kramer, and Levi (2003) posit, has a positive influence on stock prices, contributing to higher contemporaneous returns and lower expected future returns.

⁴ As Mersch (2001) and Thompson, Thompson, and Smith (2004) note, specific estimates of the prevalence of seasonal depression vary considerably, depending on the diagnostic criteria and sample selection methods employed by the researchers. The nature, incidence, and cause of SAD are discussed in a wide range of articles in the medical and psychology literatures surveyed by Lee et al. (1998).

⁵ See, for instance, Molin et al. (1996) and Young et al. (1997).

⁶ Depressed individuals have been shown to be significantly more risk averse than nondepressed individuals. See Zuckerman (1984), Carton et al. (1995), and Kramer and Weber (2012), among others. Further, certain standard psychological measures of risk aversion have been shown to capture financial risk aversion. See, for instance, Harlow and Brown (1990) and Horvath and Zuckerman (1993). See Kamstra, Kramer, and Levi (2003) and Kamstra, Kramer, and Levi (forthcoming) for more details on these links. Additionally, Kramer and Weber (2012) study hundreds of individuals across the seasons, including individuals who suffer from seasonal depression and individuals who do not. They find that the depressed group is more averse to financial risk than the nondepressed group in all seasons, most markedly so in fall/winter.

Research in psychology has also established that depressed individuals tend to exhibit greater preference to consume in the present, which manifests in a condition termed “compulsive buying disorder.” Lejoyeux et al. (1997) and Lejoyeux, Haberman, and Adès (1999) find that the incidence of compulsive buying among depressed individuals is about 40%, compared to an incidence in the overall U.S. population around 6% (see Koran et al. 2006). The most common theory of compulsive buying is that low serotonin levels found in depression, including SAD, are associated with increased rates of impulsivity. Lejoyeux et al. (1997) find that treatment with antidepressants, including serotonin reuptake inhibitors (commonly known as SSRIs), can help to decrease the frequency and severity of uncontrolled buying. Thus, it may be that some individuals who suffer from compulsive buying disorder use the consumption experience as a way to self-medicate their depression. Compulsive buying disorder is characterized by a lack of impulse control over making purchases and hence smoothing consumption over time, or stated in terms more familiar to financial economists, low elasticity of intertemporal substitution. Berns, Laibson, and Loewenstein (2007) find that EIS is impacted by emotion and affect. Recall that SAD is itself a condition associated with seasonal changes in affect, specifically depression. Thus, just as Kamstra, Kramer, and Levi (2003) argue that seasonality in depression implies seasonality in risk aversion, so we hypothesize that seasonality in depression may also manifest itself in seasonally varying EIS, with low EIS in the season when risk aversion is high and vice versa.

2. Asset Prices and Returns

In this section we develop the theoretical framework for our calibration exercise. We assume a representative agent framework. Given the psychological and behavioral evidence reviewed earlier, it may seem most natural to consider a model with heterogeneous agents, some of whom suffer from seasonal depression and others who do not. In such a model, returns would vary across seasons such that all assets are held at all times, in the spirit of the CAPM. However, when the market is complete, the aggregation theory of Constantinides (1982) shows that a representative agent model applies without loss of generality as far as equilibrium asset prices are concerned. Thus, building on Shefrin’s (2008, chapter 14) point that the risk preferences of heterogeneous investors can be captured through a representative agent model, we assume that the representative agent has seasonally varying preferences. In addition, we assume that the representative agent has recursive utility as in Epstein and Zin (1989).⁷

⁷ Note that in a heterogeneous agent economy in which agents have recursive utility, the aggregate representative agent may not have recursive utility. Our assumption that the representative agent has recursive seasonally varying utility is for its analytical convenience.

2.1 Recursive utility with seasonally varying preferences

We assume that the representative agent has recursive preferences as developed by Epstein and Zin (1989), modified to allow for seasonally varying risk aversion and elasticity of intertemporal substitution,

$$U_t = \left[c_t^{(1-\gamma_t)/\theta_t} + \beta(E_t U_{t+1}^{1-\gamma_t})^{1/\theta_t} \right]^{\theta_t/(1-\gamma_t)} \tag{1}$$

where $\theta_t = (1 - \gamma_t)/(1 - 1/\psi_t)$. Here, γ_t is risk aversion; ψ_t is EIS; and c_t is consumption. Following Bansal and Yaron (2004) and others, we assume that $\psi_t > 1$, implying that $\theta_t < 0$. As discussed in Section 1, previous research in psychology has suggested connections between depression and greater risk aversion and between depression and lower EIS. Thus, during the seasons in which an individual suffers from seasonal depression, we expect that his/her EIS, ψ_t , will be lower, consistent with his/her decreased willingness to substitute his/her consumption today for consumption tomorrow. Assuming $\gamma_t > 1$, the lower ψ_t is associated with a lower θ_t . At the same time, the seasonal depression is associated with higher risk aversion, γ_t . Overall, the effect of seasonal depression on this individual is higher γ_t and lower θ_t . Conversely, in seasons in which the individual is not depressed, his/her behavior can be characterized by lower γ_t and higher θ_t . When $\theta_t = 1$ and γ_t is a constant, the result is the standard intertemporally additive expected utility.

With this model, the utility maximization problem of the representative agent is

$$J_t(W_t, x_t; \gamma_t, \theta_t) = \max_{c_t, \pi_t} \left[c_t^{(1-\gamma_t)/\theta_t} + \beta(E_t J_{t+1}^{1-\gamma_t})^{1/\theta_t} \right]^{\theta_t/(1-\gamma_t)},$$

subject to the constraint

$$W_{t+1} = (W_t - c_t)\pi_t(1+R_{t+1}),$$

where W_t is the wealth of the agent; x_t is the vector of state variables; π_t is the vector of portfolio weights; and R_t is the vector of asset returns (consisting of the risky market return, $R_{M,t}$, and the risk-free return, $R_{f,t}$).

We assume that the agent’s preferences oscillate across periods. In one period the agent has low risk aversion and high EIS, in the next he/she has high risk aversion and low EIS, and so on. We should emphasize that the representative agent knows his/her preferences will change over time and all of his/her decisions take that fact into account. That is, the representative agent is not myopic.

By an argument similar to that of Epstein and Zin (1989), it can be shown (see Equation (18) in Appendix A) that the price $P_{j,t}$ of asset j at time t is

$$P_{j,t} = \beta^{\theta_t} w_t(x_t)^{1-\theta_t} E_t \left[g_{t+1}^{-\gamma_t} (w_{t+1}(x_{t+1})+1)^{\frac{\theta_{t+1}(1-\gamma_t)}{(1-\gamma_{t+1})}-1} (P_{j,t+1}+d_{j,t+1}) \right], \tag{2}$$

where $g_{t+1} = \frac{c_{t+1}}{c_t}$ (the growth in consumption); $d_{j,t}$ is the dividend on asset j at time t ; and $w_t(x_t)$ is the price-dividend ratio of the market portfolio at time t . We denote the price of the risky market portfolio as $P_{M,t}$, and the price of the risk-free asset as B_t .

Equation (2) implies (see Equation (15) in Appendix A) that the price-dividend ratio of the market portfolio satisfies

$$w_t(x_t) = \beta \left(E_t \left[\left(g_{t+1} (1 + w_{t+1}(x_{t+1}))^{\frac{\theta_{t+1}}{(1-\gamma_{t+1})}} \right)^{1-\gamma_t} \right] \right)^{1/\theta_t}. \quad (3)$$

We are interested in the way the prices and returns of the risky and risk-free assets behave across seasons. Consider first the risky return. The return on the market portfolio is related to the price-dividend ratio through

$$R_{M,t+1} = \frac{P_{M,t+1} + d_{M,t+1}}{P_{M,t}} = \frac{(w(x_{t+1}) + 1)g_{t+1}}{w(x_t)}. \quad (4)$$

Taking the expectation yields

$$E_t[R_{M,t+1}] = E_t \left[\frac{P_{M,t+1} + d_{M,t+1}}{P_{M,t}} \right] = \frac{E_t[(w(x_{t+1}) + 1)g_{t+1}]}{w(x_t)}. \quad (5)$$

Recall that the price-dividend ratio depends on θ and γ , which vary over time; thus, the price-dividend ratio itself varies over time. Therefore, it is evident from Equation (5) that as the representative agent's preferences move back and forth between the seasons, the expected market return moves accordingly.

Next, we consider the price of the one-period risk-free asset. Equation (2) implies that the price of the one-period risk-free asset satisfies

$$B_t(x_t) = \beta^{\theta_t} w_t(x_t)^{1-\theta_t} E_t \left[g_{t+1}^{-\gamma_t} (w_{t+1}(x_{t+1}) + 1)^{\frac{\theta_{t+1}(1-\gamma_t)}{(1-\gamma_{t+1})} - 1} \right]. \quad (6)$$

Thus, as the representative agent's preferences move back and forth between the seasons, the bond price also moves accordingly.

The intuition behind the relationship between asset prices and EIS is straightforward. *Ceteris paribus*, when the agent is less willing to substitute between today's and tomorrow's consumption, he/she will consume more today and hence invest less, which implies the demand for all assets is lower and hence the returns on all assets higher.

The intuition of how the asset returns and the price-dividend ratio move as risk aversion varies is also straightforward. It works through two channels. The first one is the asset allocation channel, which is temporal in nature. As risk aversion increases, the agent shifts her investment away from the risky market portfolio toward the riskless asset. Thus, during periods in which risk aversion is high, the price of the risky asset $P_{M,t}$ is lower and the risky return is

higher than in low risk aversion periods. Because $P_{M,t}$ is equal to the price-dividend ratio multiplied by dividend, the price-dividend ratio is smaller in high risk aversion periods than in low risk aversion periods. Consequently, the magnitude of the price-dividend ratio varies countercyclically across periods relative to the cycles in risk aversion, and the risky asset return varies procyclically across periods relative to the cycles in risk aversion. The effect of time-varying risk aversion on the risk-free asset is exactly the opposite. The risk-free asset return varies countercyclically relative to risk aversion.

The second channel is the intertemporal substitution channel. As risk aversion increases, the certainty equivalent of future consumptions is lower. To establish this lower certainty equivalent level as the equilibrium level, it must be that the price of future consumption is higher relative to the price of today's consumption so that the representative agent does not want to substitute today's consumption for more future consumption. Because the risk-free rate is part of the price of future consumption (i.e., it is equal to risk neutral probability divided by $1+R_{f,t}$), the change in risk aversion will also have an effect similar to that of EIS on the risk-free rate through its intertemporal substitution effect.

It is worth emphasizing that, although a change in risk aversion has its effect through two different channels, its effect through the asset allocation channel is in contrast to the effect of EIS. An increase in risk aversion, through the asset allocation channel, leads to a lower price for the risky asset but a higher price for the risk-free asset, whereas a decrease in intertemporal substitution, either through a decrease in EIS or through an increase in risk aversion via the second channel outlined above, leads to lower prices of all assets.

2.2 Incorporating a Markov process for consumption growth

We assume that the agent cycles between two seasons with a different set of preferences in each season. We further assume that each of the periods (seasons) is six months in length.⁸ During one six-month period (season) every year, the set of preference parameters is $\{\gamma_{FW}, \theta_{FW}\}$ (for fall/winter), and in the other six-month season of every year, the set is $\{\gamma_{SS}, \theta_{SS}\}$

⁸ The periodicity we adopt in modeling the influence of seasonal depression on risk aversion and EIS is based on the clinical studies by Young et al. (1997) and Lam (1998), who document the timing of the clinical onset of and recovery from seasonal depression symptoms among North Americans known to be affected. They find that most people who suffer from seasonal depression experience their symptoms for about six months during the fall and winter seasons, with Lam (1998) finding the peak in onset during October and the peak in recovery during April (he reports monthly statistics) and Young et al. (1997) finding the peak in onset in the second week of October. (Young et al. 1997 report weekly data but do not report recovery statistics.) Thus, for the sake of parsimony, we opted to develop a two-season model based on the representative agent suffering from seasonal depression during the fall and winter (beginning in October, consistent with the timing observed by these clinical researchers) and not suffering from seasonal depression for the other six months of the year (beginning in April). (In a previous version of this paper, we performed robustness checks for which the six-month seasons begin in September and March, or begin in November and May, and in both cases we found results similar to those we report in this paper.) In principle, one might consider developing instead a more complex model with four distinct seasons: fall, winter, spring, and summer.

(for spring/summer). The fall/winter parameter set $\{\gamma_{FW}, \theta_{FW}\}$ corresponds to the time period in which the agent suffers from seasonal depression and consequently exhibits high risk aversion and low EIS (and thus low θ). In the other six-month period of every year, spring/summer, the agent exhibits low risk aversion and high EIS. To be clear, risk aversion and EIS vary deterministically as the seasons change. The representative agent is not myopic and treats the seasonal variation in preferences as known.

In contrast to the deterministic nature of the variation of risk aversion and EIS over time, we assume consumption growth varies according to a two-state Markov process akin to the Mehra and Prescott (1985) setting. Now the price dividend ratio, w_i^k , is indexed by $i = 1, 2$ for (consumption growth) state and by $k = FW, SS$ for the high risk aversion and inelastic consumption (fall/winter) and low risk aversion and elastic consumption (spring/summer) seasons in the model. Equation (3) then yields (see Equations (16) and (17) in Appendix A) the following system of equations for price-dividend ratios:

$$(w_i^{FW})^{\theta_{FW}} = \beta^{\theta_{FW}} \sum_{j=1}^2 \phi(i, j) g_j^{1-\gamma_{FW}} \left(1 + w_j^{SS}\right)^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}}, \quad (7)$$

$$(w_i^{SS})^{\theta_{SS}} = \beta^{\theta_{SS}} \sum_{j=1}^2 \phi(i, j) g_j^{1-\gamma_{SS}} \left(1 + w_j^{FW}\right)^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}}. \quad (8)$$

Here, w_i^k is the price-dividend ratio when the preference parameters are γ_k and θ_k ; the current state of consumption growth is i ; and ϕ is the Markov state-transition probability matrix.

In the case of a two-state Markov world and two sets of preference parameters, it follows from Equation (6) (see Equations (20) and (21) in Appendix A) that the one-period bond price satisfies the following system of equations:

$$B_i^{FW} = \beta^{\theta_{FW}} (w_i^{FW})^{1-\theta_{FW}} \sum_{j=1}^2 \phi(i, j) g_j^{-\gamma_{FW}} \left(1 + w_j^{SS}\right)^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}-1}, \quad (9)$$

$$B_i^{SS} = \beta^{\theta_{SS}} (w_i^{SS})^{1-\theta_{SS}} \sum_{j=1}^2 \phi(i, j) g_j^{-\gamma_{SS}} \left(1 + w_j^{FW}\right)^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}-1}. \quad (10)$$

We solve this system of equations for equilibrium price-dividend ratios, w_i^k , and bond prices, B_i^k . Then we integrate over growth states to determine the expected risky and risk-free rates of returns in the two seasons. Note that we occasionally suppress the qualifier “expected” when referring to the expected returns on the risky and risk-free assets.

2.3 Sensitivity of returns to preference parameters

One of our novel findings (to be shown in detail in later sections) is that to match observed seasonal characteristics of risky and risk-free returns, the necessary degree of seasonal variation in EIS is very small relative to that required in risk aversion. To demonstrate the relatively higher sensitivity of returns to changes in EIS, a natural starting point is an analytic approach. In Appendix B we present an analytic solution to evaluate the sensitivity of returns and the risk premium to changes in EIS, under the strong simplifying assumptions of iid consumption growth and nonseasonally varying preferences. However, for the more interesting general case with non-iid consumption growth and seasonally varying preferences, the solution is analytically intractable. Thus, we solve numerically the sensitivity of returns to seasonal changes in EIS. We present the results in Figure 1. As we now explain, the derivative of returns with respect to a change in EIS is represented by the slope of the lines shown in each of the Figure 1 panels.

Consider panel A of Figure 1. The vertical axis depicts seasonal variation in returns as the difference between fall/winter returns and spring/summer returns. The horizontal axis depicts the magnitude of seasonal variation in EIS ($\Delta_\psi = \psi_{FW} - \psi_{SS}$). The thin solid line represents the seasonal change in the risky return for different values of Δ_ψ , and the thick solid line represents the seasonal change in the risk-free return for different values of Δ_ψ . The dashed line represents the difference between the thin solid and thick solid lines; that is, it represents the seasonal change in the equity premium as we vary Δ_ψ . For all panels, risk aversion in the spring/summer season (γ_{SS}) equals 1.75 and risk aversion in the fall/winter season (γ_{FW}) equals 15.75. In panel A only, EIS in the spring/summer season (ψ_{SS}) equals 2. By moving along the horizontal axis, we can consider the impact on returns arising from different seasonal changes in EIS. Starting at the far right of panel A, Δ_ψ equals zero, and hence EIS in the fall/winter is identically equal to EIS in the spring/summer. In that case, we see from the plot that the seasonal change in the equity return (the thin solid line) is about 2, and the seasonal change in the risk-free rate (the thick solid line) is about -3 , both of which are at odds with the observed seasonal variation in risky and risk-free returns. (We provide more details on the observed seasonal variation in returns in the next section.) The seasonal variation in the equity premium at that point in panel A is about 5% (the dotted line). As we move to the left in this plot, we can evaluate the impact of allowing EIS to be lower in the fall/winter than in the spring/summer. Note that for values of Δ_ψ around -0.01 , the seasonal variation in the risky return, the risk-free return, and the equity premium all roughly match observed values. Specifically, the seasonal change in the risky return is around 5%, the seasonal change in the risk-free return is close to or just below zero, and the seasonal change in the risk premium is around 5%. Stated differently, the parameters allow for the best match with observed returns near the point at which the thin solid line crosses the dashed line. This also applies to panels B,




Figure 1

Sensitivity of seasonal changes in returns to EIS

In each panel, the vertical axis depicts seasonal changes in returns as the difference between fall/winter returns and spring/summer returns. The horizontal axis depicts the magnitude of seasonal variation in EIS ($\Delta\psi = \psi_{FW} - \psi_{SS}$). The dashed line represents the seasonal change in the equity premium; the thin solid line represents the seasonal change in the risky return; and the thick solid line represents the seasonal change in the risk-free return. In all cases, risk aversion in the spring/summer season (γ_{SS}) equals 1.75 and risk aversion in the fall/winter season (γ_{FW}) equals 15.75. EIS in the spring/summer season (ψ_{SS}) equals 2, 3, 4, and 6 in panels A through D, respectively.

C, and D, where we consider values of EIS in the spring/summer season (ψ_{SS}) of 3, 4, and 6, respectively. In panel B, we find the best match to observed data when the degree of seasonal change in EIS is around -0.02 ; in panel C this occurs around -0.04 ; and in panel D this occurs around -0.08 . Notice that in all cases, the required degree of seasonal change in EIS is very small, ranging from about -0.01 to -0.08 (depending on the assumed spring/summer value of EIS) compared with the magnitude of the required degree of seasonal change in risk aversion (15 in this particular example). In untabulated robustness checks, we explored a broad range of values of risk aversion and EIS preferences across the seasons, and we uniformly found that to match the seasonal characteristics of observed risky and risk-free returns, the necessary amount of seasonal variation in EIS is very small compared with that in risk aversion, uniformly near the range shown in this example.

3. Return Data

We now provide stylized facts on return seasonalities. We perform all of our analysis on a semiannual periodicity so that a semiannual risk-free rate, set at the beginning of the semiannual period, is the appropriate quantity to calibrate to the model's risk-free rate. Although a shorter-term instrument, such as the one-month Treasury-bill could be rolled over to produce a six-month return (and indeed we consider such a return for robustness and find qualitatively identical results), this would not be a true six-month risk-free rate because the one-month bill varies unpredictably month-to-month over the semiannual period. Unfortunately, six-month Treasury-bill rates are not available until 1959; thus, our calibration exercise relies primarily on data from 1959 and onward. For comparison, later in the paper we also discuss the features of data available beginning as early as 1926.

Our primary analysis is based on the 1959–2012 period. To get an overview of equity and Treasury market seasonalities, we present Table 1, which contains average inflation-deflated (real) six-month returns for equities (based on the CRSP value-weighted U.S. total-market index including dividends), average real six-month returns for risk-free securities (U.S. Treasury-bills), and equity premium values over the 1959–2012 period.

We calculate seasonal rates of return using the CRSP value-weighted returns, including dividends for the risky asset and the CRSP Fama Treasury-Bill Term Structure Files for the risk-free asset. There are two aspects related to our choice of series that warrant discussion. First, because these are nominal rates of return and our model describes real rates of return, we must deflate these return data. The use of real returns is consistent with related studies, including Mehra and Prescott (1985). We deflate the nominal return series following Mehra and Prescott (1985) using a deflator series produced by dividing real consumption of nondurables and services by the nominal consumption of nondurables and services.⁹ Second, the convention in the asset-pricing literature when performing calibration exercises is to consider value-weighted returns, as they better capture total returns to the typical investor. Thus, we employ value-weighted returns.

⁹ The Treasury-bill series are deflated using the predicted inflation rate, where we use an ARMA(1,1) time-series model to form our predicted inflation series. We use predicted inflation because the Treasury-bill rate is set using information available only at the beginning of the period; realized inflation is known only after the Treasury-bill matures. This regression, estimated on semiannual data, has an R^2 of 71% and removes evidence of autocorrelation to five lags (2.5 years). Coefficient estimates (with standard errors in parentheses) are as follows. Intercept: 0.0029 (0.0013), AR(1): 0.86 (0.058), MA(1): -0.097 (0.120). The realized inflation rate is used to deflate the equity return, because the realized inflation rate is known when the equity return is realized. The consumption data we use are sourced from the U.S. Department of Commerce, Bureau of Economic Analysis, via the Federal Reserve Bank of St. Louis FRED database. These data include personal consumption expenditures on services and nondurable goods, real and nominal, series IDs PCES, PCENDC96, PCEND, and PCESC96.

Table 1
Average realized U.S. real rates of return, 1959–2012

Period	Six-month VW equity return (%)	Six-month risk-free return based on holding six-month maturity security (%)	Six-month risk-free return based on rolling over thirty-day maturity security (%)
April–September	1.149	0.793	0.562
October–March	6.284	0.772	0.485
Seasonal change	5.134	−0.021	−0.077
Standard error of seasonal change	2.583	0.206	0.191

We report the average real six-month equity and risk-free returns for the 1959–2012 sample period. The six-month periods are spring/summer (April–September) and fall/winter (October–March). We consider two distinct risk-free returns series: one based on the return to holding a six-month Treasury-bill, and another based holding a thirty-day Treasury-bill rolled over each month to form a six-month semiannual return. The equity return data are CRSP U.S. total market (NYSE, NASDAQ, and AMEX) value-weighted returns, including dividends, deflated using realized inflation as described in footnote 9. The Treasury data are from the CRSP Fama Treasury-Bill Term Structure Files, deflated using predicted inflation as described in footnote 9. For each series, we consider average returns and the difference in average returns across the two seasons. Heteroscedasticity-consistent standard errors are calculated based on MacKinnon and White (1985).

The statistics we present are based on splitting the year into October–March and April–September six-month seasons. This split coincides with the clinical observation that October is the approximate point at which there is a peak in SAD diagnoses and April is the approximate peak point of recovery from SAD, as reported in footnote 8 above. In a previous version of this paper we performed a robustness check based on splitting the year into September–February and March–August halves and into November–April and May–October halves and found similar results.

The first noteworthy feature of Table 1 is that the average equity return is much higher in the fall/winter season (October–March) than in the spring/summer season (April–September), with the approximate 5% seasonal difference being statistically significant at the 5% level, two-sided, based on MacKinnon and White (1985) heteroscedasticity-consistent standard errors.

The second noteworthy feature of Table 1 pertains to the average Treasury-bill returns. We consider very short maturity (six-month) Treasury securities, and we find no strong seasonal return pattern at this maturity. We should emphasize that Kamstra, Kramer, and Levi (forthcoming) find an economically large and statistically significant seasonal pattern in longer-maturity Treasury securities (five-, seven-, ten-, and twenty-year Treasury notes/bonds). Specifically, they find longer-term Treasury returns are significantly lower than average during the fall/winter season when risky equity returns are higher than average, and vice versa for the spring/summer. They find the effect trails off monotonically as maturity shortens, presumably because monetary policy aims explicitly to remove seasonality in the money supply and is believed to have a large moderating influence on the shorter end

of the term structure.¹⁰ The moments in Table 1 show little or no seasonal variation in very short maturity Treasury security returns (consistent with the Federal Reserve's monetary policy objectives). The reader should keep in mind, however, the more general finding of significant seasonal variation in returns to Treasury notes and bonds, with significant countercyclical seasonal variation in those securities' returns relative to that found in equity returns: in general there is roughly a 1% seasonal difference in longer-maturity Treasury returns versus roughly a 5% seasonal difference in equity returns (see Kamstra, Kramer, and Levi forthcoming).¹¹

4. Using the Model to Match Returns

Our next task is to match model-predicted expected returns with observed return patterns. To do this, we must parameterize the consumption growth process of our model and determine the seasonal patterns of expected returns. In this section we explore two different calibrations to model consumption growth, both of which are based on seasonally adjusted consumption data commonly employed in the asset pricing literature. With both calibrations we can match the moments of the observed data when we allow for seasonal variation in EIS and risk aversion. Later, in Section 5, we base our calibration exercise on seasonally unadjusted consumption data, similarly finding that the model generates returns that match observed risky and risk-free returns only if we allow both EIS and risk aversion to vary seasonally.

4.1 Consumption data

Our first consumption growth calibration exercise is based on seasonally adjusted consumption data from January 1959 to December 2012.¹² We restrict our attention to data starting in 1959, due to unavailability of the risk-free six-month Treasury-bill return series and monthly consumption data prior to that year. The consumption data we employ for this exercise are real nondurables and services consumption (from the Bureau of Economic Analysis). We find positively autocorrelated consumption growth, and assuming a two-state

¹⁰ Gibson (1970, 442), referring to Treasury-bill rates, notes that an "aim of the Federal Reserve System is to accommodate seasonal swings in the financial needs of trade, and the System tries to do this by removing seasonal fluctuations from interest rates." Fama (2013, 198) is skeptical about the influence of such efforts, noting that "Fed actions with respect to its target rate have little effect on long-term interest rates, and there is substantial uncertainty about the extent of Fed control of short-term rates."

¹¹ One interpretation of the joint findings (based on our Table 1 and on the results of Kamstra, Kramer, and Levi [forthcoming]) is that equity returns exhibit a strong seasonal cycle and Treasury returns exhibit a smaller seasonal pattern that is countercyclical relative to the seasonal pattern in equity returns. Another interpretation is that equity returns exhibit a strong seasonal cycle, whereas the very short end of the Treasury security market does not exhibit a strong cycle (unlike the rest of the Treasury maturity spectrum), perhaps due to the influence of the Federal Reserve. Note that either interpretation is at odds with standard asset pricing models, which imply no seasonal variation in either set of returns. In our calibration exercises reported below, we attempt to match the return characteristics for the full sample, using the average return values shown in Table 1.

¹² Ending the sample in 2007 to avoid the influence of the financial crisis leads to qualitatively identical results.

Markov process for consumption growth, we derive $\phi(i, i) = 0.638$. The mean (annualized) real growth rate in consumption we find is 3.11%, together with a standard deviation of 1.71%. The second calibration of seasonally adjusted consumption growth is the classic parameterization of Mehra and Prescott (1985). This calibration exhibits negative autocorrelation of consumption growth ($\phi(i, i) = 0.43$), mean consumption growth equal to 1.80%, and standard deviation of growth equal to 3.60%. We also incorporate the use of levered equity (to help us match the unconditional level of returns) by following Campbell (1986), Abel (1999), and Lettau, Ludvigson, and Wachter (2008) and employing a scaling parameter (denoted λ by Abel 1999 and others). This parameter scales up the volatility of growth to reflect the effect of leverage. Unleveraged assets, such as those found in the Lucas (1978) fruit-tree model, have $\lambda = 1$, with values greater than one reflecting leveraged assets. The payoff in period t on the canonical asset is proportional to an observable variable, typically consumption, raised to the power λ . Lettau, Ludvigson, and Wachter (2008) set $\lambda = 4.5$, which we also adopt.

4.2 Grid for preference parameters

We determine the expected returns to holding risk-free and risky assets implied by the characteristics of consumption by searching over a grid of values of the risk aversion and EIS preference parameters. The grid we search over is informed by previous research. The parameter of relative risk aversion, γ , has been argued by Lettau, Ludvigson, and Wachter (2008) to be as high as possibly 30 (to match observed mean equity premia, dividend yields, and risk-free rates on postwar data). Small, negative values of the EIS parameter (denoted ψ) have been reported by authors, including Hall (1988), but for the most part the consensus in the literature is that the parameter is positive and, according to researchers, including Lettau, Ludvigson, and Wachter (2008), Bansal and Yaron (2004), and Vissing-Jorgensen and Attanasio (2003), greater than 1. Vissing-Jorgensen and Attanasio (2003) suggest values close to 1.5 but report estimates as high as 17.6 (see their Table I).

Because we explore seasonally varying risk aversion and EIS in this calibration exercise, each spring/summer (γ_{SS}, ψ_{SS}) pair has a corresponding fall/winter (γ_{FW}, ψ_{FW}) pair with $\gamma_{FW} \neq \gamma_{SS}$ and $\psi_{FW} \neq \psi_{SS}$. (We also consider varying only one of ψ or γ at a time.) In our full grid search, we explore values of γ_{SS} as low as 1.25 and as high as 30. We set $\gamma_{FW} = \gamma_{SS} + \Delta\gamma$, with $\Delta\gamma > 0$ and equal to as little as 0.25 and as much as 30. We consider values of ψ between 1 and 6.¹³ We set $\psi_{FW} = \psi_{SS} + \Delta\psi$, with $\Delta\psi < 0$ and equal to as

¹³ In a previous version of this paper we considered values of $\psi < 1$ as well. Doing so allows slightly better fit to equity and bond returns for the case of constant ψ but has little or no impact on fit when ψ is allowed to vary seasonally. Here, we restrict our attention to values of $\psi > 1$ because of the growing consensus in the literature.

little as -0.0001 and as much as -0.1 . In all cases, we set β no greater than 0.9849 for the semiannual frequency.^{14,15}

4.3 Calibration results

We demonstrate now that in a model with seasonally varying risk aversion and EIS, we can closely match the observed magnitudes and signs of changes in returns in our semiannual stock and bond data. We also show, in contrast, that if we allow only risk aversion to vary seasonally (with EIS held constant), we are able to replicate the signs of seasonal changes in risky versus risk-free asset returns, but we are not able to closely match the magnitudes. For the interested reader, we present results in Appendix C showing that if we allow only EIS to vary seasonally (with risk aversion held constant), we can match neither the signs nor the magnitudes of seasonal changes in risky and risk-free asset returns. Note that returns emerging from the calibration exercises that we report in the tables below are all expected returns, produced by solving the system of Equations (7)–(10) for the equilibrium values of w_i^k and B_i^k and integrating over growth states.

4.3.1 Calibrating to the Mehra and Prescott consumption growth data.

Consider Table 2, in which we present results based on the Mehra and Prescott consumption growth parameters. Panel A contains the stylized features of risky and risk-free returns that we attempt to match: much higher risky returns in fall/winter than in spring/summer, with a seasonal difference of about 5%, and slightly lower (or constant) risk-free returns in fall/winter versus spring/summer, with a seasonal difference of a few basis points.

Panel B contains calibration results for cases in which we allow both EIS and risk aversion to vary seasonally. Values for γ and ψ are in the first two columns, with results clustered in groups of two, the first corresponding to the April–September period (associated with low risk aversion and high EIS) and the second the October–March period (associated with high risk aversion and low EIS). We report the equity and risk-free rates of return produced by the a representative set of best-performing γ and ψ combinations, together with a test that the data sample moments reported in panel A are consistent with the model.¹⁶ In the top set of cells in panel B, γ varies from 1.5 in the

¹⁴ We set the semiannual rate of time preference of 0.9849 to match an annual rate of 0.97, based on the quarterly value of 0.9925 from Lettau, Ludvigson, and Wachter (2008).

¹⁵ There are (γ, ψ) pairs for which a simultaneous setting of $\beta = 0.9849$ can produce negative returns; when this occurs we lower β in steps by roughly 0.01 and restart. This is continued until returns are positive. Virtually all of the parameterizations have a β no less than 0.9746.

¹⁶ The model specification test has us comparing the model-implied return patterns with the return patterns observed in the actual data, following Gregory and Smith (1991) and Cogley and Nason (1995), among others. The details of this method are as follows. Using the γ and ψ values from the table, the consumption growth parameters reported in the paper, and Equations (7)–(10), we simulate 10,000 independent outcomes of 108 periods each (to match the 54 year semiannual actual data period to which we calibrate our model). The

Table 2

Calibration results: Using Mehra and Prescott consumption growth parameters (six-month rates of return)

Panel A: Stylized facts we attempt to match					
Period	Equity return (%)	Risk-free return (%)			
April–September	1.149	0.793			
October–March	6.284	0.772			
Seasonal change	5.134	−0.021			

Panel B: Calibration results for seasonally varying risk aversion and seasonally varying EIS					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.5	1.5	April–September	2.3	0.58	
7.5	1.495	October–March	7.8	0.38	0.73
2.25	1.5	April–September	3.1	0.53	
8.25	1.495	October–March	8.4	0.42	0.28
2	1.5	April–September	2.4	0.14	
8	1.495	October–March	7.9	0.09	0.38
1.5	4	April–September	2.6	0.94	
8.5	3.95	October–March	7.9	0.21	0.39
1.5	5.5	April–September	2.6	1.00	
8.5	5.4	October–March	8.0	0.45	0.94

Panel C: Calibration results for seasonally varying risk aversion and constant EIS					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.5	1.25	April–September	7.5	5.63	
7.5	1.25	October–March	8.1	0.49	<0.001
2.25	4	April–September	6.0	3.46	
9.25	4	October–March	8.4	0.29	<0.001
2.5	1.5	April–September	7.7	4.69	
8.5	1.5	October–March	8.6	0.45	<0.001
3	5.5	April–September	6.4	3.03	
10	5.5	October–March	8.9	0.44	<0.001

The data in panel A are from Table 1. In panel B we summarize results for which we allow both seasonally varying risk aversion and seasonally varying EIS, and in panel C we summarize results for which only risk aversion varies seasonally. We calibrate to consumption growth parameters from the Mehra and Prescott (1985) sample period: negative autocorrelation of consumption growth ($\phi = 0.43$), mean consumption growth is equal to 1.80%, and standard deviation of growth is equal to 3.60%. The consumption data are seasonally adjusted. Footnote 16 describes the specification test.

April–September season to 7.5 in the October–March season and ψ varies from 1.5 in the April–September season to 1.495 in the October–March season. The magnitudes of equity returns in each of the spring/summer and fall/winter seasons, 2.3% and 7.8%, respectively, and the difference of about

Markov transition probabilities introduce randomness to these 10,000 simulated outcomes. We then count the fraction of these simulated economies that yield results like those found in U.S. data. That is, we count how many simulated economies had low (high) risk aversion season equity returns that were no higher (lower) than the sample April–September (October–March) equity return of 1.149% (6.284%) and low (high) risk aversion season risk-free returns that were no lower (higher) than the sample April–September (October–March) risk-free return of 0.793% (0.772%).

5% across the seasons, come close to matching the observed values shown in panel A. The magnitude of risk-free rates in each season, 0.58% and 0.38%, and the difference of -20 basis points across the seasons, are also a fairly good match. We perform a formal test of the fit between the model and the data based on the specification test p -value in the last column, which shows that we fail to reject the model at conventional levels of significance. In the remaining cases shown in panel B we explore different values of risk aversion and EIS across the seasons, and the results show that we fail to reject the model, and we manage to capture the primary qualitative features of the data. We see a large seasonal swing in equity returns (as large as a 5.5% difference across the seasons), and we observe a small or nonexistent swing in risk-free returns (as small as a few basis points across the seasons), with the seasonal patterns in risk-free returns, where evident, typically being offset in timing relative to equity returns.

Panel C contains a representative set of the best-performing results for the Mehra and Prescott calibration, where we allow risk aversion to vary seasonally, but we restrict EIS to be constant across the seasons. That is, each (γ_{SS}, ψ_{SS}) spring/summer pair has a corresponding fall/winter (γ_{FW}, ψ_{FW}) pair with $\gamma_{FW} = \gamma_{SS} + \Delta_\gamma$ and $\psi_{FW} = \psi_{SS}$. The model generates moderate seasonal changes in risky returns of 1% to 3% instead of the 5% we seek to match. The seasonal change in the risk-free return is generally too large, around 3% to 5% instead of close to zero. And the specification test rejects the model's fit with the data in all cases.

Overall, the full set of Table 2 results highlights the importance of allowing seasonal variation in both EIS and risk aversion. When both are permitted to vary, the model is capable of generating risky and risk-free returns that match the characteristics of observed data. However, lack of seasonal variation in EIS causes gross failure of the model to match the data. In Appendix C we show that lack of seasonal variation in risk aversion similarly leads to model failure.

4.3.2 Calibrating to the 1959–2012 consumption growth data. Analogous results using the 1959–2012 consumption growth parameters appear in Table 3. Once again, the moments we seek to match appear in panel A, both EIS and risk aversion vary seasonally in panel B, and only risk aversion varies seasonally in panel C. We see in panel B that with seasonally varying EIS and risk aversion, we obtain the stylized seasonal patterns in returns, with large seasonal differences in risky returns (around 5% difference), much smaller seasonal differences in risk-free returns (5 to 16 basis points), and typically countercyclical seasonal variation in risk-free returns relative to risky returns. Notice that the seasonal changes in risk aversion required in Table 3 are generally larger than those that we saw in Table 2. The need for relatively larger changes in the coefficient of relative risk aversion for the

Table 3

Calibration results: Using the 1959–2012 consumption growth parameters (six-month rates of return)

Panel A: Stylized facts we attempt to match					
Period	Equity return (%)	Risk-free return (%)			
April–September	1.149	0.793			
October–March	6.284	0.772			
Seasonal change	5.134	–0.021			

Panel B: Calibration results for seasonally varying risk aversion and seasonally varying EIS					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.25	3	April–September	0.8	0.36	
16.25	2.98	October–March	5.9	0.28	0.22
1.75	4.5	April–September	0.9	0.16	
16.75	4.45	October–March	6.1	0.22	0.13
1.75	5.5	April–September	1.0	0.25	
16.75	5.43	October–March	6.0	0.09	0.26
2	5.5	April–September	0.9	0.11	
17	5.425	October–March	6.1	0.16	0.19
2	4.5	April–September	1.0	0.23	
16	4.455	October–March	5.8	0.13	0.01

Panel C: Calibration results for seasonally varying risk aversion and constant EIS					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.75	2.25	April–September	3.5	2.99	
21.75	2.25	October–March	5.5	0.05	<0.001
2	3	April–September	3.4	2.76	
22	3	October–March	5.8	0.07	<0.001
2	5	April–September	3.2	2.56	
22	5	October–March	6.1	0.10	<0.001
2.5	2.5	April–September	3.0	2.16	
11.5	2.5	October–March	4.2	0.28	<0.001

The data in panel A are from Table 1. In panel B we summarize results for which we allow both seasonally varying risk aversion and seasonally varying EIS, and in panel C we summarize results for which only risk aversion varies seasonally. We calibrate to consumption growth parameters from the 1959–2012 sample period: positive autocorrelation of consumption growth ($\phi = 0.638$), mean consumption growth is equal to 3.11%, and standard deviation of growth is equal to 1.71%. The consumption data are seasonally adjusted. Footnote 16 describes the specification test.

1959–2012 calibration arises because of the smaller volatility of consumption over this period relative to the Mehra and Prescott sample period. This feature of the 1959–2012 data makes it relatively more difficult to match the features of returns using plausible values of γ and ψ across the seasons. In the last row of panel B, we present an instance in which the model fails the specification test. Nevertheless, even in that particular case, the model comes fairly close to matching the level of returns in both seasons.

Turning to panel C, we see that with seasonal variation in risk aversion only, the returns are not a good match for the observed data and the

specification tests indicate rejection. The failure of the model comes from an equity rate of return that varies too little and a risk-free rate that varies too much relative to observed returns. Variation in EIS succeeds where fixed EIS fails by reducing all asset returns in the fall/winter and increasing all asset returns in the spring/summer, thus dampening the risk-free seasonal variation and amplifying the risky return variation. Again, this highlights the importance of allowing for seasonal variation in both EIS and risk aversion.

Overall, the results in Tables 2 and 3 show there exist combinations of parameters that generate equity and Treasury returns that match the properties of observed returns. We now turn to reporting on a full grid search across a range of values of γ and ψ to make general statements about the existence of ranges of parameter values that are capable of generating equity and Treasury returns with properties matching those observed in the data.

4.4 Full grid search over values of the preference parameters

In this section we consider plots of goodness-of-fit measures (i.e., p -values for the specification tests shown in Tables 2 and 3). These plots provide deeper intuition for the values of γ and ψ that yield the best match to observed returns over the 1959–2012 period and provide insight into the sensitivity of the match to the values of γ and ψ . We consider a broad set of values for seasonally varying γ and seasonally varying ψ . Results for the Mehra and Prescott calibration period appear in Figure 2, and those for the 1959–2012 calibration period appear in Figure 3. The higher the p -value on the y -axis (labeled “Fit”), the better the exercise succeeds in matching actual returns. A “Fit” value of zero, the lowest possible value on the scale of the y -axis, would correspond to a very poor fit, with the data rejecting the model very strongly. Terminating the top of the vertical scale at 0.10 is consistent with the fact that we do not typically reject a model based on standard levels of significance when we observe p -values of 10% or more. (Results are very similar based on either a 5% or 10% level of significance.) The values plotted on the horizontal and depth axes are the values of ψ and γ in the low risk aversion / high EIS (spring/summer) season, that is, ψ_{SS} and γ_{SS} . For each (ψ_{SS}, γ_{SS}) pair, the data plotted in the figure indicate the best fit over the range of $(\Delta\psi, \Delta\gamma)$ values we explored. The best model fit in Figure 2 is achieved for ψ_{SS} values between about one and six, and for γ_{SS} values up to about seven. The best model fit in Figure 3 is roughly similar, though the suitable range for γ_{SS} is more constrained depending on the value of ψ_{SS} .

4.5 Plausibility of preference parameters

One might reasonably wonder whether the parameter values emerging from this calibration exercise are plausible. That is, are the values of γ and ψ shown in panel B of Tables 2 and 3 reasonable relative to values documented in the literature? The EIS values are within standard ranges considered by




Figure 2
Best fit based on parameters from the Mehra and Prescott calibration period
 We plot goodness-of-fit measures for the case of seasonally varying EIS and seasonally varying risk aversion, based on consumption growth parameters from the Mehra and Prescott calibration period, with seasonally adjusted consumption data. The values on the horizontal and depth axes are γ and ψ in the low risk aversion / high EIS (spring/summer) season, that is, γ_{SS} and ψ_{SS} . The values on the vertical axis are the p -values described in footnote 16.




Figure 3
Best fit based on parameters from the 1959–2012 calibration period
 We plot goodness-of-fit measures for the case of seasonally varying EIS and seasonally varying risk aversion, based on consumption growth parameters from the 1959–2012 calibration period, with seasonally adjusted consumption data. The values on the horizontal and depth axes are γ and ψ in the low risk aversion / high EIS (spring/summer) season, that is, γ_{SS} and ψ_{SS} . The values on the vertical axis are the p -values described in footnote 16.

researchers, including Vissing-Jorgensen and Attanasio (2003) and Lettau, Ludvigson, and Wachter (2008). The coefficient of relative risk aversion values are well within plausible ranges for the low risk-aversion season, between 1.25 and 3. In the high risk aversion period, the coefficient of relative risk aversion rises to about 7 to 8.5 in Table 2 and as high as about 16 to 17 in Table 3. The highest values we consider are within the range considered by researchers including Lettau, Ludvigson, and Wachter (2008).

An additional consideration is whether we need to scale the parameter values to adjust for the possibility that only a fraction of the population experiences seasonally varying preferences. In response to this potential concern, we highlight the evidence mentioned in Section 1 that most everyone experiences some degree of seasonal depression and seasonally varying risk

preferences. That is, while people who suffer from severe seasonal depression experience the most dramatic seasonal changes in mood and risk preferences, even healthy individuals experience milder albeit marked changes. Nonetheless, even considering for a moment the counterfactual case in which only a small fraction of the population experienced seasonal changes in risk preferences, it would be the marginal trader influencing prices when risk preferences would be changing seasonally, and so it would not be appropriate to rescale parameters even under those extreme circumstances.

5. Consumption Seasonality

Now we consider whether observed seasonal patterns in equity and Treasury returns arise simply as a consequence of deterministically and seasonally varying consumption growth (see Miron 1986; Ferson and Harvey 1992; and Piazzesi 2001). Once again, following Mehra and Prescott, we assume a two-state Markov world, but we now allow the consumption growth rate to have a different mean and variance across seasons, as is evident in the seasonally unadjusted consumption data. The Markov transition matrix remains the same, $\phi(i, i) = 0.638$ (based on the 1959–2012 calibration period), but now the growth rate depends on both the season and the state. In the seasonally unadjusted consumption data, consumption has a lower growth rate and lower variance in the fall/winter season than in the spring/summer season, and we replicate this feature of the data in our calibration exercise. Later, we report specific values of the growth rate by season and by state.

Let us define the consumption growth mean and standard deviation values in the high risk aversion / low EIS season as $\{\mu_{FW}, \sigma_{FW}\}$ and the consumption growth mean and standard deviation values in the low risk aversion / high EIS season as $\{\mu_{SS}, \sigma_{SS}\}$. Further, the consumption growth rate realized in the high risk aversion / low EIS season is g_{FW} and the consumption growth rate realized in the low risk aversion / high EIS season is g_{SS} . Modifying Equation (3) to incorporate the state-dependent consumption growth rates then yields the following system of equations for price-dividend ratios:

$$(w_i^{FW})^{\theta_{FW}} = \beta^{\theta_{FW}} \sum_{j=1}^2 \phi(i, j) g_{j,FW}^{1-\gamma_{FW}} \left(1 + w_j^{SS}\right)^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}}, \tag{11}$$

$$(w_i^{SS})^{\theta_{SS}} = \beta^{\theta_{SS}} \sum_{j=1}^2 \phi(i, j) g_{j,SS}^{1-\gamma_{SS}} \left(1 + w_j^{FW}\right)^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}}. \tag{12}$$

Here, w_i^k has the interpretation of price-dividend ratios when the preference parameters are γ_k and θ_k and the current state of consumption growth is $i = 1, 2$, $k = FW, SS$ indexing the different sets of preference parameters over the fall/winter and spring/summer seasons, and ϕ is the Markov state-transition probability matrix.

In the case of a two-state Markov world and two sets of preference parameters, it follows from Equation (6) that the one-period bond price satisfies the following system of equations:

$$B_i^{FW} = \beta^{\theta_{FW}} (w_i^{FW})^{1-\theta_{FW}} \sum_{j=1}^2 \phi(i, j) g_{j, FW}^{-\gamma_{FW}} \left(1 + w_j^{SS}\right)^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}-1}, \quad (13)$$

$$B_i^{SS} = \beta^{\theta_{SS}} (w_i^{SS})^{1-\theta_{SS}} \sum_{j=1}^2 \phi(i, j) g_{j, SS}^{-\gamma_{SS}} \left(1 + w_j^{FW}\right)^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}-1}. \quad (14)$$

In Table 4, we provide annualized average growth rates and volatility for the seasonally unadjusted consumption data, estimated over 1959–2004 and over subperiods, covering the first and last halves of the sample.^{17,18} We find that the estimates (especially the means) are very stable across the full sample and the subperiods, and the seasonality across periods is very robust. Specifically, the consumption growth rate and its volatility are always lower in the October–March season.

We proceed now to the calibration exercise based on the adapted Epstein and Zin model. Let us consider first the case with time-invariant preferences, $\{\gamma_{FW}, \theta_{FW}\} = \{\gamma_{SS}, \theta_{SS}\}$. Searching over a wide range of parameter values (identical to those described in Section 3), we find without exception that the model delivers lower equity returns in the October–March season than in the April–September season, which is at odds with the observed return pattern over these seasons. The model also delivers, virtually without exception, lower bond returns in the April–September season. Panel B of Table 5 contains a representative set of the best-performing results for the time-invariant preference calibration exercise. In sum, the exercise yields results that are strongly inconsistent with observed data, yielding strong rejections of the model and leading us to conclude that seasonality in consumption growth does not explain the observed seasonal patterns in returns.

It is interesting to consider which degree of seasonal variation in consumption growth could possibly match the seasonal return patterns in a model with constant preference parameters. Unfortunately, even allowing large seasonal swings in consumption growth and volatility we cannot easily match the return data; the specification test *p*-values indicate that the data reject all models with the exception of models embedding counterfactually large seasonal changes in both the mean and volatility of consumption. For instance, we must push volatility of consumption to 8% (annualized) in the fall/winter

¹⁷ Our seasonally unadjusted consumption data is the sum of nondurables and services: Variables B004RU1 and B005RU1 from Table 8.1, Bureau of Economic Analysis, 1947:1-2004:4. (BEA has not updated this series recently, due to budget cutbacks.)

¹⁸ We deflate consumption with the consumer price index for all urban consumers, all items, series ID CPIAUCNS, U.S. Department of Labor: Bureau of Labor Statistics, using the CPI at the end of the quarter used as the deflator.

Table 4
Annualized average consumption growth rates and volatility for seasonally unadjusted consumption data

Period	Mean consumption growth (μ)	Volatility of consumption growth (σ)
Panel A: 1959–2004		
April–September	4.0%	3.1%
October–March	2.4%	1.3%
Panel B: 1959–1983		
April–September	4.3%	4.0%
October–March	2.4%	1.6%
Panel C: 1984–2004		
April–September	3.8%	1.7%
October–March	2.4%	0.8%

We report estimates of the mean consumption growth rate and the volatility of consumption growth for the April–September and October–March seasons over the 1959–2004 sample period and subperiods consisting of the first and last halves of the sample. The consumption data are seasonally unadjusted. The source of the consumption data is described in footnote 17, and the deflation method is described in footnote 18.

season (in contrast to the actual value of 1.3% in the data) and consumption growth to 19% (in contrast to the actual value of 2.4%) to achieve a match to the return data that is not rejected. These parameters generate an equity return of 3.4% in the spring/summer and 9.2% in the fall/winter and a risk-free return of 1.5% in the spring/summer and 1.3% in the fall/winter.¹⁹ Furthermore, to produce these results, we also need counterfactual seasonal variation in consumption growth; specifically, mean growth and volatility of growth would need to rise in the fall/winter, whereas we observe the opposite seasonal pattern in the data.

We turn next to allowing for both seasonality in consumption growth and seasonality in preferences. That is, we allow γ and θ to change with the semiannual seasons, and we also allow for consumption growth rates to have a lower mean and a lower variance in the fall/winter season, matched with observed data. In this calibration exercise, we find we are able to match the direction and, to a remarkable extent, the magnitude of changes in risky and risk-free returns. Representative best-performing results are presented in panel C of Table 5.

In Figure 4 we plot the goodness-of-fit measure for the case of seasonally varying EIS, seasonally varying risk aversion, and seasonally varying consumption growth and risk. The best fit is achieved for γ values up to about two and for ψ values greater than about three.

In sum, we find that when allowing for seasonality in consumption growth alone we are unable to explain the observed seasonal patterns in risky and risk-free returns. To capture the seasonal variability of returns, we must also allow for seasonally varying risk aversion and EIS.

¹⁹ This particular result was generated with $\gamma = 1.25$ and $\psi = 1.5$. We explored values of γ and ψ in the range of 1.1 to 10 and from 1.25 to 6, respectively, with similar results.

Table 5

Calibration results incorporating seasonal consumption growth: Using seasonally unadjusted consumption growth parameters (six-month rates of return)

Panel A: Stylized facts we attempt to match					
Period		Equity return (%)		Risk-free return (%)	
April–September		1.149		0.793	
October–March		6.284		0.772	
Seasonal change		5.134		−0.021	

Panel B: Constant γ and ψ					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.25	5.5	April–September	2.5	1.03	
1.25	5.5	October–March	2.0	1.46	<0.001
1.75	4	April–September	2.9	0.78	
1.75	4	October–March	2.2	1.43	<0.001
1.75	6	April–September	2.8	0.66	
1.75	6	October–March	2.1	1.32	<0.001
2.5	5.5	April–September	3.2	0.18	
2.5	5.5	October–March	2.3	1.16	<0.001

Panel C: Seasonally varying γ and ψ					
γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.01	5	April–September	1.2	0.11	
23.01	4.94	October–March	6.5	0.23	0.06
1.01	5	April–September	1.4	0.34	
26.01	4.95	October–March	6.5	0.07	0.15
1.05	5	April–September	1.4	0.32	
28.05	4.95	October–March	6.6	0.11	0.31
1.25	4	April–September	1.6	0.30	
31.25	3.97	October–March	6.6	0.15	0.13

The data in panel A are from Table 1. Results in panels B and C are from a model based on consumption growth parameters from the 1959–2004 sample period: seasonal consumption growth parameters of $\{\mu_{FW}, \sigma_{FW}\} = \{2.4\%, 1.3\%\}$ and $\{\mu_{SS}, \sigma_{SS}\} = \{4.0\%, 3.1\%\}$, calculated based on seasonally unadjusted consumption data. We also make use of the 1959–2012 seasonally adjusted consumption data to calculate the Markov growth transition matrix: positive autocorrelation of annual consumption growth ($\phi = 0.67$). In panel B we summarize results for which both EIS and risk aversion are constant, and in panel C we summarize results for which we allow both seasonally varying risk aversion and seasonally varying EIS.

6. Including Pre-1959 Return Data

As a supplement to our primary analysis of return data and stylized facts on seasonality, we now discuss extension of the sample period to include returns data prior to 1959, as far back as 1926, which is the earliest year for which CRSP returns data are available. Note that analysis using samples that include pre-1959 data rely on the one-month Treasury-bill rolled over to form a six-month return (analogous to that shown in the last column of Table 1 for the 1959–2012 sample) as a stand-in for the six-month Treasury-bill for observations prior to 1959.




Figure 4
Best fit for seasonally varying consumption growth and risk case, based on parameters from the 1959–2012 calibration period

We plot goodness-of-fit measures for the case with seasonally varying consumption growth and risk, and seasonally varying EIS and risk aversion, based on consumption growth parameters from the 1959–2012 calibration period, with ϕ based on seasonally adjusted consumption data and all other consumption growth characteristics based on seasonally unadjusted consumption data. The values on the horizontal and depth axes are γ and ψ in the low risk aversion / high EIS (spring/summer) season, that is, γ_{SS} and ψ_{SS} . The values on the vertical axis are the p -values described in footnote 16.

Mehra and Prescott (1985) report substantial subsample variability in the equity premium over periods prior to 1959, with an equity premium near 15% during 1919–1928, 0.18% during 1929–1938, 9% during 1939–1948, and above 18% during 1949–1958 (see their Table 1). The 1920s and 1930s were, of course, a remarkable period of time, encompassing the bubble of the late 1920s and the Great Depression of the 1930s. Consistent with this volatility, we find the seasonal change in the equity premium across the fall/winter versus the spring/summer is also very unstable over the 1920s and 30s. To see this, consider panel A of Figure 5. The vertical axis is the difference between the fall/winter equity premium and the spring/summer equity premium (where the equity premium is the difference between the value-weighted risky returns and risk-free returns). Each asterisk in the plot represents the seasonal change in the equity premium estimated over a five-year window. The farthest-left asterisk represents the seasonal change in the equity premium for 1926–1930, the next asterisk moves the window by a year, and so on to the farthest-right asterisk, which represents the seasonal change in the equity premium for 1945–1949. The influence of the Great Depression is evident over most of the sample period shown in panel A.

In panel B, we consider longer windows, with each rolling window no less than 20 years in length and with the windows expanding to a maximum of 54 years in length (matching the length of the 1959–2012 sample period we consider in the primary analysis). We omit data prior to 1941, acknowledging that the Great Depression was an unusual period of history during which there is little evidence of the type of seasonal variation in the equity premium that we find otherwise. In addition to plotting the seasonal change in the equity premium (again shown using asterisks), we now provide the lower bounds of

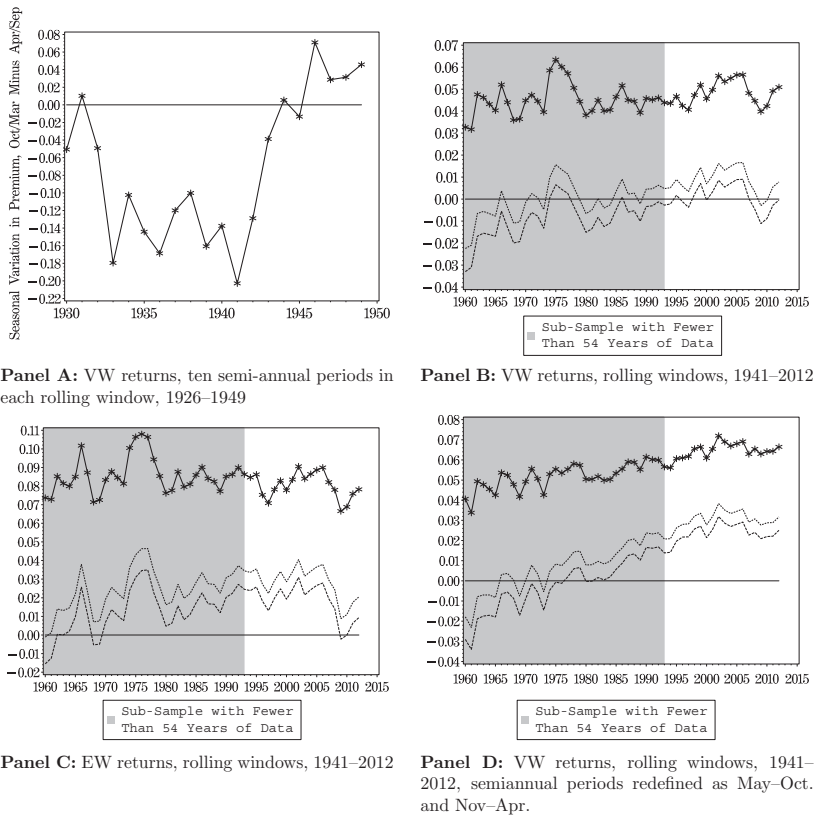


Figure 5

Rolling window estimation

In all panels, asterisks depict the seasonal difference in average risky returns minus average risk-free returns, that is, the equity premium in the fall/winter minus the equity premium in the spring/summer. In panels A, B, and D, the risky return is the value-weighted equity return, and in panel C the risky return is the equal-weighted equity return. In panel A, each rolling window used to calculate the average returns is ten semiannual periods (five years), with the first rolling window being 1926–1930 and the last rolling window being 1946–1949. In panels B, C, and D, each rolling window is no less than 20 years and no more than 54 years. 1960 has the minimum number of years of data, 20, and the number of years of data used increases each year beyond 1960 until 1993 when we achieve the first window with the maximum number of years, 54 years, which continues to be the window length until the end of the sample. That is, the first rolling window encompasses 1941–1960, the second 1941–1961, and so on until 1994 when the window encompasses 1942–1994. The last data point plotted uses the window 1959–2012. The shaded areas in panels B, C, and D highlight data points estimated with fewer than 54 years of data. In panels A, B, and C, the semiannual periods are April–September and October–March. In panel D, the semiannual periods are May–October and November–April. Confidence intervals are based on MacKinnon and White’s (1985) bootstrap heteroscedasticity-consistent standard errors.

90% and 95% confidence intervals, allowing us to test whether the seasonal changes arising from any rolling window is significantly greater than zero. The shaded area highlights data points estimated with fewer than 54 years of data and hence estimated with less information than we use in our primary analysis. Based on the 90% confidence bound, we find in the unshaded sample that the fall/winter equity premium is significantly greater than the

spring/summer equity premium with the exception of a brief period during the recent financial crisis. There are some additional exceptions based on the 95% confidence interval, but again the seasonal difference is broadly significant. The shaded area gives us some sense of the stability of the seasonal variation in the early part of the sample using shorter subsamples to estimate the seasonal variation in the equity premium. Although the confidence intervals are wider, as we might expect using shorter sample periods for estimation, the mean effect is fairly stable, mostly taking on values in the range of 3% to 6%.

Panel C is analogous to panel B, except it uses equal-weighted data for the risky returns, placing relatively more weight on small firms. With this modification, we find the seasonal change in the equity premium is close to everywhere significant based on the 90% confidence interval, and it is typically significant based on the 95% bound. Panel D is also analogous to panel B, and uses value-weighted risky returns, but the semiannual periods are redefined as May to October for spring/summer and November to April for fall/winter. With this shifted timing for the seasons, we find statistically significant seasonal changes in the equity premium for all rolling windows for most of the estimation windows we investigate, apart from the earliest windows, which use fewer observations to estimate the effect.

We provide the plot shown in panel D to highlight the fact that, of course, our findings are somewhat sensitive to the start and end dates of the six-month seasons. Given that research in psychology suggests the peak onset of seasonal depression is roughly in October, the ideal starting point for the fall/winter period may be close to October 15, but in absence of daily or weekly data, we must choose either October 1 or November 1 for the start of a period. For our primary analysis, we start the fall/winter period in October 1, due to its convenience in aligning with the reporting data for many government statistics, and consistent with the timing convention used in other studies of return seasonality related to depression such as Kamstra, Kramer, and Levi (forthcoming).

Overall, these data suggest fairly stable seasonality in the equity premium over the post-Depression era. Similar results obtain for seasonality in equity and Treasury returns, although consistent with Table 1, virtually all the seasonality comes from the equity return.

7. Conclusions

Past research posits that seasonal variation in investor mood impacts asset returns through the influence of depression on risk preferences. Any model of asset returns used to bridge the empirical-conceptual gap must be ultimately based on realistic preferences for consumption over time or preferences for risk. We explore the ability of a representative agent asset pricing model to match a remarkable empirical regularity, namely, that realized equity returns

vary conditionally across the seasons as much as 12% (annualized) and that average U.S. Treasury returns may vary countercyclically relative to equity returns, albeit on a much smaller scale.

On average, observed equity returns are higher and bond returns are lower in the fall/winter season (the time of year when many investors experience seasonal depression, lower EIS, and higher risk aversion). Further, the variation in bond returns across the seasons is much smaller than that in equity returns. In our model calibrated to consumption growth data, both EIS and risk aversion must vary seasonally in order to generate returns that match these distinct seasonal characteristics of risky and risk-free returns. We show that allowing only one or the other of EIS and risk aversion to vary seasonally is insufficient. Both must vary or the model generates counterfactual returns patterns and is strongly rejected by the data.

Our main analysis is based on calibrating to seasonally adjusted consumption data, but we also consider a model calibrated to unadjusted consumption data. Without seasonally varying EIS and risk aversion, the returns that emerge from this model exhibit characteristics grossly at odds with observed data, with too little seasonal variation in risky returns and wrongly signed movements in risky and risk-free returns across the seasons. Incorporating seasonally varying preferences in that model yields returns that closely match observed data.

Relative to prior work on the influence of seasonally varying preferences in asset pricing, our findings with respect to EIS are novel. Specifically, we find that seasonal changes in agents' willingness to postpone consumption play an important and previously overlooked role in explaining seasonal return dynamics. Further, the changes in the risk preferences required to match the observed seasonal changes in returns are such that the required seasonal change in EIS is very small relative to the required seasonal change in risk aversion. Changes in EIS as little as 0.08 generally suffice. In comparison, the change in risk aversion across the seasons must be around six or seven in one of our main calibration exercises in order to match the characteristics of observed returns.

Appendix

A. Recursive Utility with Seasonal Depression

Let us consider recursive utility with time-varying risk aversion γ_t and constant ϕ :

$$U_t = \left[c_t^{(1-\gamma_t)/\theta_t} + \beta(E_t U_{t+1}^{1-\gamma_t})^{1/\theta_t} \right]^{\theta_t/(1-\gamma_t)},$$

where ψ , defined through $\theta_t = (1 - \gamma_t)/(1 - 1/\psi)$, is the elasticity of intertemporal substitution. When $\theta_t = 1$ we get the standard intertemporally additive expected utility.

The utility maximization problem is

$$J_t(W_t, x_t) = \max_{c_t, \pi_t} \left[c_t^{(1-\gamma_t)/\theta_t} + \beta(E_t J_{t+1}^{1-\gamma_t})^{1/\theta_t} \right]^{\theta_t/(1-\gamma_t)},$$

subject to the constraint

$$W_{t+1} = (W_t - c_t)\pi_t(1+R_{t+1}),$$

where x_t is the vector of state variables and π_t is the vector of portfolio weights.

It is readily seen that the recursive utility is homothetic so that $J_t(W_t, x_t) = f_t(x_t)W_t$. The original utility maximization problem can be written as

$$\begin{aligned} f_t(x_t)W_t &= \max_{c_t, \pi_t} \left[c_t^{(1-\gamma)/\theta_t} + \beta(E_t[(f_{t+1}(x_{t+1})(W_t - c_t)\pi_t(1+R_{t+1}))^{1-\gamma}]]^{1/\theta_t} \right]^{\theta_t/(1-\gamma)} \\ &= \max_{c_t, \pi_t} \left[c_t^{(1-\gamma)/\theta_t} + \beta(W_t - c_t)^{(1-\gamma)/\theta_t} (E_t[(f_{t+1}(x_{t+1})\pi_t(1+R_{t+1}))^{1-\gamma}]]^{1/\theta_t} \right]^{\theta_t/(1-\gamma)}. \end{aligned}$$

The first-order condition for c_t is then

$$c_t^{(1-\gamma)/\theta_t-1} = \beta(W_t - c_t)^{(1-\gamma)/\theta_t-1} (E_t[(f_{t+1}(x_{t+1})\pi_t(1+R_{t+1}))^{1-\gamma}])^{1/\theta_t}.$$

Homogeneity suggests that $c_t = C_t(x_t)W_t$. Substitution into the utility maximization problem yields

$$f_t(x_t)^{(1-\gamma)/\theta_t} = C_t(x_t)^{(1-\gamma)/\theta_t} + \beta(1 - C_t(x_t))^{(1-\gamma)/\theta_t} (E_t[(f_{t+1}(x_{t+1})\pi_t(1+R_{t+1}))^{1-\gamma}])^{1/\theta_t}.$$

The first-order condition for c_t can be written as

$$C_t(x_t)^{(1-\gamma)/\theta_t-1} = \beta(1 - C_t(x_t))^{(1-\gamma)/\theta_t-1} (E_t[(f_{t+1}(x_{t+1})\pi_t(1+R_{t+1}))^{1-\gamma}])^{1/\theta_t}.$$

Combining the two equations yields

$$\begin{aligned} f_t(x_t) &= C_t(x_t)^{(1-\gamma-\theta_t)/(1-\gamma)}, \\ f_{t+1}(x_{t+1}) &= C_{t+1}(x_{t+1})^{(1-\gamma_{t+1}-\theta_{t+1})/(1-\gamma_{t+1})}, \end{aligned}$$

and

$$C_t(x_t)^{1-\gamma-\theta_t} = \beta^{\theta_t} (1 - C_t(x_t))^{1-\gamma-\theta_t} E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{1-\gamma_t}].$$

For the optimal portfolio choice, we examine

$$\max_{\pi_t} E_t[(f_{t+1}(x_{t+1})\pi_t(1+R_{t+1}))^{1-\gamma_t}].$$

The first-order condition is simply

$$0 = E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (R_{i, t+1} - r_t)],$$

which implies

$$E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (1+R_{i, t+1})] = E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (1+r_t)]$$

and

$$E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{1-\gamma_t}] = E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (1+r_t)].$$

Then

$$C_t(x_t)^{1-\gamma-\theta_t} = \beta^{\theta_t} (1 - C_t(x_t))^{1-\gamma-\theta_t} E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (1+r_t)]$$

and

$$C_t(x_t)^{1-\gamma-\theta_t} = \beta^{\theta_t} (1 - C_t(x_t))^{1-\gamma-\theta_t} E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (\pi_t(1+R_{t+1}))^{-\gamma_t} (1+R_{i, t+1})].$$

If there is only one risky stock, then

$$C_t(x_t)^{1-\gamma-\theta_t} = \beta^{\theta_t} (1 - C_t(x_t))^{1-\gamma-\theta_t} E_t[f_{t+1}(x_{t+1})^{1-\gamma_t} (1+R_{t+1})^{1-\gamma_t}].$$

A.1 The Market Portfolio

In equilibrium, the representative agent holds the market portfolio. The equation derived above can be written as

$$1 = \beta^{\theta_t} \left(\frac{1}{C_t(x_t)} - 1 \right)^{1-\gamma_t-\theta_t} E_t \left[f_{t+1}(x_{t+1})^{1-\gamma_t} \left(\frac{P_{M,t+1} + d_{t+1}}{P_{M,t}} \right)^{1-\gamma_t} \right],$$

where $P_{M,t}$ is the price of the market portfolio. Homogeneity suggests that $P_{M,t} = w(x_t)d_t$. Now note that in equilibrium $c_t = d_t$, and consumer wealth is equal to the stock price because there is only one share of the stock and the consumer does not have labor income. Thus, $W_t = P_{M,t} + d_t$ and $C_t(x_t) = c_t/W_t = 1/(w_t(x_t)+1)$. Substituting into the above equation yields

$$1 = \beta^{\theta_t} \left(\frac{P_{M,t} + d_t}{d_t} - 1 \right)^{1-\gamma_t-\theta_t} E_t \left[\left(\left(\frac{d_{t+1}}{P_{M,t+1} + d_{t+1}} \right)^{\frac{(1-\gamma_{t+1}-\theta_{t+1})}{(1-\gamma_{t+1})}} \right)^{(1-\gamma_t)} \left(\frac{P_{M,t+1} + d_{t+1}}{P_{M,t}} \right)^{1-\gamma_t} \right],$$

where we have used the expression for f_{t+1} derived earlier. Then

$$1 = \beta^{\theta_t} w_t(x_t)^{1-\gamma_t-\theta_t} E_t \left[(w_{t+1}(x_{t+1})+1)^{\frac{(\gamma_{t+1}+\theta_{t+1}-1)(1-\gamma_t)}{(1-\gamma_{t+1})}} \left(\frac{(w_{t+1}(x_{t+1})+1)g_{t+1}}{w_t(x_t)} \right)^{1-\gamma_t} \right].$$

Simplifying yields

$$w_t(x_t)^{\theta_t} = \beta^{\theta_t} E_t \left[g_{t+1}^{1-\gamma_t} (1+w_{t+1}(x_{t+1}))^{\frac{\theta_{t+1}(1-\gamma_t)}{(1-\gamma_{t+1})}} \right]. \quad (15)$$

Let us assume a finite-state Markov world. Now the price dividend ratio, w_j^k , is indexed by $i = 1, 2$ for state and by $k = FW, SS$ for the fall/winter and spring/summer seasons. Assuming there are only two states, Equation (3) yields the following system of equations:

$$(w_i^{FW})^{\theta_{FW}} = \beta^{\theta_{FW}} \sum_{j=1}^n \phi(i, j) g_j^{1-\gamma_{FW}} (1+w_j^{SS})^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}}, \quad (16)$$

$$(w_i^{SS})^{\theta_{SS}} = \beta^{\theta_{SS}} \sum_{j=1}^n \phi(i, j) g_j^{1-\gamma_{SS}} (1+w_j^{FW})^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}}. \quad (17)$$

A.2 The Bond Price

More generally, the price of any stock j should satisfy

$$1 = \beta^{\theta_t} w_t(x_t)^{1-\theta_t} E_t \left[(w_{t+1}(x_{t+1})+1)^{\frac{(\gamma_{t+1}+\theta_{t+1}-1)(1-\gamma_t)}{(1-\gamma_{t+1})}} ((w_{t+1}(x_{t+1})+1)g_{t+1})^{-\gamma_t} \left(\frac{P_{j,t+1} + d_{j,t+1}}{P_{j,t}} \right) \right].$$

Thus,

$$P_{j,t} = \beta^{\theta_t} w_t(x_t)^{1-\theta_t} E_t \left[g_{t+1}^{-\gamma_t} (w_{t+1}(x_{t+1})+1)^{\frac{\theta_{t+1}(1-\gamma_t)}{(1-\gamma_{t+1})}-1} (P_{j,t+1} + d_{j,t+1}) \right]. \quad (18)$$

It follows from Equation (18) that the one period bond price satisfies

$$B_t = \beta^{\theta_t} w_t(x_t)^{1-\theta_t} E_t \left[g_{t+1}^{-\gamma_t} (w_{t+1}(x_{t+1})+1)^{\frac{\theta_{t+1}(1-\gamma_t)}{(1-\gamma_{t+1})}-1} \right]. \quad (19)$$

Let us assume a two-state Markov world and that there are only two levels of risk aversion. Then Equation (19) yields the following system of equations for states $i = 1, 2$:

$$B_i^{FW} = \beta^{\theta_{FW}} (w_i^{FW})^{1-\theta_{FW}} \sum_{j=1}^n \phi(i, j) g_j^{-\gamma_H} (1+w_j^{SS})^{\frac{\theta_{SS}(1-\gamma_{FW})}{(1-\gamma_{SS})}-1}, \text{ for odd periods, } t = 1, 3, \dots \tag{20}$$

$$B_i^{SS} = \beta^{\theta_{SS}} (w_i^{SS})^{1-\theta_{SS}} \sum_{j=1}^n \phi(i, j) g_j^{-\gamma_{SS}} (1+w_j^{FW})^{\frac{\theta_{FW}(1-\gamma_{SS})}{(1-\gamma_{FW})}-1}, \text{ for even periods, } t = 2, 4, \dots \tag{21}$$

Having solved for w_i^k from Equations (16) and (17), we can now solve for the bond price and hence the risk-free rate. It is seen from these expressions that risk-free rates are stationary when the representative agent has Epstein and Zin (1989) utility.

B. Sensitivity of Returns to EIS

The bond price is given by

$$B_t(x_t) = E_t \left[\beta^{\theta_t} g_{t+1}^{-\gamma_t} w_t(x_t)^{1-\theta_t} (1+w_{t+1}(x_{t+1}))^{\frac{\theta_{t+1}(1-\gamma_t)}{1-\gamma_{t+1}}-1} \right].$$

Thus,

$$\frac{dB_t(x_t)}{d\theta_t} = E_t \left[\beta^{\theta_t} g_{t+1}^{-\gamma_t} w_t(x_t)^{1-\theta_t} (1+w_{t+1}(x_{t+1}))^{\frac{\theta_{t+1}(1-\gamma_t)}{1-\gamma_{t+1}}-1} \ln(1+w_{t+1}(x_{t+1})) \right] \frac{1-\gamma_t}{1-\gamma_{t+1}}.$$

Then

$$\frac{dR_{f,t}}{d\psi_{t+1}} = \frac{1}{(1+r_t)^2} E_t \left[\beta^{\theta_t} g_{t+1}^{-\gamma_t} w_t(x_t)^{1-\theta_t} (1+w_{t+1}(x_{t+1}))^{\frac{\theta_{t+1}(1-\gamma_t)}{1-\gamma_{t+1}}-1} \ln(1+w_{t+1}(x_{t+1})) \right] \frac{1-\gamma_t}{(\psi_{t+1}-1)^2}.$$

In the special case of iid growth and no seasonal variation in preferences, the risk-free rate is

$$1+R_{f,t} = \frac{1}{\beta^\theta E[g^{-\gamma}] w^{1-\theta} (1+w)^\theta} = \frac{1}{\beta E[g^{-\gamma}]} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-\gamma}.$$

Then

$$\frac{dR_{f,t}}{d\psi} = \frac{-1}{\psi^2} \ln \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right] \times \frac{1}{\beta E[g^{-\gamma}]} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-\gamma}.$$

Clearly the derivative also depends on the sign of $\ln \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]$.

When there is no seasonal variation in preferences and consumption growth is iid, the equity premium can be written as

$$\begin{aligned} E[R_{M,t}] - R_{f,t} &= E[g] \beta^{-1} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-1} - \frac{1}{\beta E[g^{-\gamma}]} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-\gamma} \\ &= \beta^{-1} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-1} \left\{ E[g] - \frac{E[g^{1-\gamma}]}{E[g^{-\gamma}]} \right\} \\ &= \beta^{-1} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-1} \frac{\text{cov}(g, -g^{-\gamma})}{E[g^{-\gamma}]} \end{aligned}$$

Then

$$\frac{dE[R_{M,t}] - R_{f,t}}{d\psi} = \frac{-1}{\psi^2} \ln \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right] \times \beta^{-1} \left[(E[g^{1-\gamma}])^{\frac{1}{1-\gamma}} \right]^{1/\psi-1} \frac{\text{cov}(g, -g^{-\gamma})}{E[g^{-\gamma}]}.$$

Table A1

Calibration results for seasonally varying EIS and constant risk aversion: Using Mehra and Prescott consumption growth parameters (six-month rates of return)

Panel A: Stylized facts we attempt to match

Period	Equity return (%)	Risk-free return (%)
April–September	1.149	0.793
October–March	6.284	0.772
Seasonal change	5.134	−0.021

Panel B: Calibration results for seasonally varying EIS and constant risk aversion

γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.5	4	April–September	2.3	0.70	
1.5	3.999	October–March	2.4	0.78	<0.001
2	3	April–September	2.5	0.37	
2	2.9999	October–March	2.5	0.38	<0.001
1.5	1.25	April–September	2.4	0.59	
1.5	1.248	October–March	6.9	5.00	<0.001
2	1.25	April–September	2.7	0.36	
2	1.247	October–March	9.0	6.49	<0.001
2.5	1.5	April–September	3.0	0.23	
2.5	1.4999	October–March	3.1	0.33	<0.001

The data in panel A are from Table 1. In panel B we summarize results for which we do not allow seasonally varying risk aversion but do allow seasonally varying EIS and for which we calibrate to consumption growth parameters from the Mehra and Prescott (1985) sample period: negative autocorrelation of consumption growth ($\phi = 0.43$), mean consumption growth is equal to 1.80%, and standard deviation of growth is equal to 3.60%.

Table A2

Calibration results for seasonally varying EIS and constant risk aversion: Using the 1959–2012 consumption growth parameters (six-month rates of return)

Panel A: Stylized facts we attempt to match

Period	Equity return (%)	Risk-free return (%)
April–September	1.149	0.793
October–March	6.284	0.772
Seasonal change	5.134	−0.021

Panel B: Calibration results for seasonally varying EIS and constant risk aversion

γ	ψ	Period	Equity return (%)	Risk-free return (%)	Specification test p -value
1.5	3.5	April–September	2.1	1.54	
1.5	3.498	October–March	2.4	1.78	<0.001
4	3.5	April–September	2.6	1.00	
4	3.499	October–March	2.7	1.11	<0.001
5	2.5	April–September	2.7	0.87	
5	2.499	October–March	3.0	1.11	<0.001
5	3.5	April–September	2.5	0.52	
5	3.494	October–March	3.1	1.15	<0.001
7	1.25	April–September	3.0	0.90	
7	1.2499	October–March	3.3	1.18	<0.001

The data in panel A are from Table 1. In panel B we summarize results for which we do not allow seasonally varying risk aversion but do allow seasonally varying EIS. We calibrate to consumption growth parameters from the 1959–2012 sample period: positive autocorrelation of consumption growth ($\phi = 0.638$), mean consumption growth is equal to 3.11%, and standard deviation of growth is equal to 1.71%.

Clearly the derivative depends on the sign of $\ln\left[\left(E[g^{1-\gamma}]\right)^{\frac{1}{1-\gamma}}\right]$.

In the more general case with seasonally varying preferences and non-iid consumption growth, the derivative is analytically intractable. We can however solve for the derivative numerically using Equations (5) and (6). See Figure 1.

C. Calibrations with Seasonally Varying EIS and Constant Risk Aversion

To supplement our primary analysis, we now consider the case in which EIS varies seasonally and risk aversion remains constant across the seasons. As shown in our analysis with constant EIS and seasonally varying risk aversion, we find models with constant risk aversion and seasonally varying EIS are not able to match the features of observed risky and safe asset returns.

Each (γ_{SS}, ψ_{SS}) spring/summer pair has a corresponding fall/winter (γ_{FW}, ψ_{FW}) pair with $\gamma_{FW} = \gamma_{SS}$ and $\psi_{FW} = \psi_{SS} + \Delta\psi$, where $\Delta\psi < 0$. As before, we search over a very large grid. We provide a representative set of the best-performing results based on the Mehra and Prescott consumption parameter values in Table A1 and based on the 1959–2012 values in Table A2. Panel A in each table contains the stylized returns, and panel B contains the model results. In both tables, we see that when ψ drops in the fall/winter season, the expected returns generally rise for both the risky asset and the risk-free asset. That is, the prices of both risky and risk-free assets drop lower than they would otherwise be for all assets to be held in equilibrium, and with the relatively lower prices come relatively higher expected returns for both asset classes. Overall, when we allow EIS to vary seasonally and hold risk aversion constant, equity and risk-free returns fail to demonstrate the empirically observed magnitudes and signs of returns across the seasons, with risk-free returns varying across the seasons as much as risky returns in many cases. The specification test p -values indicate rejection for all of the shown cases, and this is representative of the findings when we search over the full grid of parameters.

References

- Abel, A. B. 1999. Risk premia and term premia in general equilibrium. *Journal of Monetary Economics* 43:3–33.
- Attanasio, O., and M. Browning. 1995. Consumption over the life cycle and over the business cycle. *American Economic Review* 85:1118–37.
- Atkeson, A., and M. Ogaki. 1996. Wealth varying intertemporal elasticities of substitution: Evidence from panel and aggregate data. *Journal of Monetary Economics* 38:507–34.
- Bansal, R., and A. Yaron. 2004. Risks for the long-run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–509.
- Barsky, R., F. T. Juster, M. Kimball, and M. Shapiro. 1997. Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Survey. *Quarterly Journal of Economics* 107:537–79.
- Berns, G. S., D. Laibson, and G. Loewenstein. 2007. Intertemporal choice - Toward an integrative framework. *Trends in Cognitive Sciences* 11:482–88.
- Bekaert, G., S. R. Grenadier, and E. Engstrom. 2010. Stock and bond returns with moody investors. *Journal of Empirical Finance* 17:867–94.
- Blundell, R., M. Browning, and C. Meghir. 1994. Consumer demand and the life-cycle allocation of household expenditures. *Review of Economic Studies* 61:57–80.
- Brandt, M. W., and K. Q. Wang. 2003. Time-varying risk aversion and unexpected inflation. *Journal of Monetary Economics* 50:1457–98.
- Braun, R. A., and C. L. Evans. 1995. Seasonality and equilibrium business cycle theories. *Journal of Economic Dynamics and Control* 19:503–31.

- Campbell, J. Y. 1986. Bond and stock returns in a simple exchange model. *Quarterly Journal of Economics* 101: 785–804.
- Campbell, J. Y., and J. Cochrane. 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107:205–41.
- Carton, S., P. Morand, C. Bungenera, and R. Jouvent. 1995. Sensation-seeking and emotional disturbances in depression, relationships and evolution. *Journal of Affective Disorders* 34:219–25.
- Chatterjee, S., and B. Ravikumar. 1992. A neoclassical model of seasonal fluctuations. *Journal of Monetary Economics* 29:59–86.
- Cogley, T., and J. M. Nason. 1995. Output dynamics in real-business-cycle models. *American Economic Review* 85:492–511.
- Constantinides, G. M. 1982. Intertemporal asset pricing with heterogeneous consumers and without demand aggregation. *Journal of Business* 55:253–67.
- DeGennaro, R., M. J. Kamstra, and L. A. Kramer. 2008. Seasonal variation in bid-ask spreads. Working Paper, University of Toronto.
- Dolvin, S. D., and M. K. Pyles. 2007. Seasonal affective disorder and the pricing of IPOs. *Review of Accounting and Finance* 6:214–28.
- Dolvin, S. D., M. K. Pyles, and Q. Wu. 2009. Analysts get SAD too: The effect of seasonal affective disorder on stock analysts' earnings estimates. *Journal of Behavioral Finance* 10:214–25.
- Dowling, M., and B. M. Lucey. 2008. Robust global mood influences in equity pricing. *Journal of Multinational Financial Management* 18:145–64.
- Epstein, L. G., and S. E. Zin. 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* 57:937–69.
- Fama, E. F. 2013. Does the Fed control interest rates? *Review of Asset Pricing Studies* 3:180–99.
- Fama, E. F., and K. R. French. 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25:23–49.
- Ferson, W. E., and S. R. Foerster. 1994. Finite sample properties of generalized methods of moments in tests of conditional asset pricing models. *Journal of Financial Economics* 36:29–55.
- Ferson, W. E., and C. R. Harvey. 1992. Seasonality and consumption-based asset pricing. *Journal of Finance* 47: 511–52.
- . 1993. Seasonality and heteroscedasticity in consumption-based asset pricing: An analysis of linear models. *Research in Finance* 11:1–35.
- Gabaix, X., and D. Laibson. 2002. The 6D bias and the equity-premium puzzle. *NBER Macroeconomics Annual* 16:257–330.
- Garrett, I., M. J. Kamstra, and L. A. Kramer. 2005. Winter blues and time variation in the price of risk. *Journal of Empirical Finance* 12:291–316.
- Gibson, W. E. 1970. Interest rates and monetary policy. *Journal of Political Economy* 78:431–55.
- Gordon, S., and P. St-Amour. 2000. A preference regime model of bull and bear markets. *American Economic Review* 90:1019–33.
- Gregory, A. W., and G. W. Smith. 1991. Calibration as testing: Inference in simulated macroeconomic models. *Journal of Business and Economic Statistics* 9:297–303.
- Guiso, L., P. Sapienza, and L. Zingales. 2013. Time varying risk aversion. Working Paper, University of Chicago.
- Hall, R. E. 1988. Intertemporal substitution in consumption. *Journal of Political Economy* 96:339–57.

- Harlow, W. V., and K. C. Brown. 1990. Understanding and assessing financial risk tolerance: A biological perspective. *Financial Analysts Journal* 6:50–80.
- Harmatz, M. G., A. D. Well, C. E. Overtree, K. Y. Kawamura, M. Rosal, and I. S. Ockene. 2000. Seasonal variation of depression and other moods: A longitudinal approach. *Journal of Biological Rhythms* 15:344–50.
- Horvath, P., and M. Zuckerman. 1993. Sensation seeking, risk appraisal, and risky behavior. *Personality and Individual Differences* 14:41–52.
- Kamstra, M. J., L. A. Kramer, and M. D. Levi. 2003. Winter blues: A SAD stock market cycle. *American Economic Review* 93:324–43.
- Kamstra, M. J., L. A. Kramer, M. D. Levi, and R. Wermers. 2013. Seasonal asset allocation: Evidence from mutual fund flows. Working Paper, University of Toronto.
- Kamstra, M. J., L. A. Kramer, and M. D. Levi. Forthcoming. Seasonal variation in Treasury returns. *Critical Finance Review*.
- Kasper, S., S. L. Rogers, A. Yancey, P. M. Schulz, R. G. Skwerer, and N. E. Rosenthal. 1989. Phototherapy in individuals with and without subsyndromal seasonal affective disorder. *Archives of General Psychiatry* 46: 837–44.
- Koran, L. M., R. J. Faber, E. Aboujaoude, M. D. Large, and R. T. Serpe. 2006. Estimated prevalence of compulsive buying behavior in the United States. *American Journal of Psychiatry* 163:1806–12.
- Kramer, L. A., and J. M. Weber. 2012. Seasonal affective disorder and risk aversion in financial decision making. *Social Psychological and Personality Science* 3:193–9.
- Lam, R. W. 1998. Seasonal affective disorder: Diagnosis and management. *Primary Care Psychiatry* 4:63–74.
- Lee, T. M. C., E. Y. H. Chen, C. C. H. Chan, J. G. Paterson, H. L. Janzen, and C. A. Blashko. 1998. Seasonal affective disorder. *Clinical Psychology: Science and Practice* 5:275–90.
- Lejoyeux, M., V. Tassain, J. Solomon, and J. Adès. 1997. Study of compulsive buying in depressed patients. *Journal of Clinical Psychiatry* 58:169–73.
- Lejoyeux, M., V. N. Haberman, and J. Adès. 1999. Comparison of buying behavior in depressed patients presenting with or without compulsive buying. *Comprehensive Psychiatry* 40:51–56.
- Lettau, M., S. C. Ludvigson, and J. A. Wachter. 2008. The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies* 21:1653–87.
- Lo, K., and S. S. Wu. 2008. The impact of seasonal affective disorder on financial analysts and equity market returns. Working Paper, University of British Columbia.
- Lucas, R. 1978. Asset prices in an exchange economy. *Econometrica* 46:1429–45.
- MacKinnon, J. G., and H. White. 1985. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics* 29:305–25.
- Maurer, T. A. 2012. Is consumption growth merely a sideshow in asset pricing? Working Paper, Washington University in St. Louis.
- Mehra, R., and E. C. Prescott. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15: 145–61.
- Mersch, P. 2001. Prevalence from population surveys. In T. Partonen and A. Magnusson (eds.), *Seasonal affective disorder: Practice and research*. Oxford: Oxford University Press.
- Miron, J. A. 1986. Seasonal fluctuations and the life cycle-permanent income model of consumption. *Journal of Political Economy* 94:1258–79.
- Molin, J., E. Møllerup, T. Bolwig, T. Scheike, and H. Dam. 1996. The influence of climate on development of winter depression. *Journal of Affective Disorders* 37:151–5.

Piazzesi, M. 2001. Comment on Gabaix and Laibson. The 6D bias and the equity-premium puzzle. In B. Bernanke and K. Rogoff (eds.), *The NBER macroeconomics annual*. Cambridge: MIT Press.

Pyles, M. K. 2009. The influence of seasonal depression on equity returns: Further evidence from real estate investment trusts. *Quarterly Journal of Finance and Accounting* 48:63–83.

Shefrin, H. 2008. *A behavioral approach to asset pricing*. 2nd edition. NY: Elsevier.

Thompson, C., S. Thompson, and R. Smith. 2004. Prevalence of seasonal affective disorder in primary care; a comparison of the seasonal health questionnaire and the seasonal pattern assessment questionnaire. *Journal of Affective Disorders* 78:219–26.

Vissing-Jorgensen, A., and O. P. Attanasio. 2003. Stock market participation, intertemporal substitution and risk-aversion. *American Economic Review* 93:383–91.

Young, M. A., P. M. Meaden, L. F. Fogg, E. A. Cherin, and C. I. Eastman. 1997. Which environmental variables are related to the onset of seasonal affective disorder? *Journal of Abnormal Psychology* 106:554–62.

Zuckerman, M. 1984. Sensation seeking: A comparative approach to a human trait. *Behavioral and Brain Science* 7:413–71.