Replication of the Equity Premium & Risk-Free Rate Puzzles under a CRRA-Utility Framework

Austin Bender
abender@uoregon.edu

Presented to the Department of Economics for fulfillment of the honors requirement.

Professor David Evans

University of Oregon
Eugene, Oregon
March 2018

Abstract

In this paper I seek to replicate aspects of the paper, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles" by Bansal and Yaron (2004). To do this, I pull the core growth processes from Bansal and Yaron (2004) and modify them to fit a power, constant relative risk aversion (CRRA), utility model that is then utilized to showcase the equity risk premium and risk-free rate puzzles apparent in using such a method.

Table of Contents

Introduction

Throughout history, the U.S. capital markets have produced on average an annual equity risk premium of roughly 6% over treasury bills. This risk premium is considered high in regard to the level of real interest rates and degree of risk aversion present in representative agents. In fact, the risk premium is considered to be so high that it has been a widely researched phenomenon in academic finance since the 1980s when Mehra and Prescott (1985) published a paper called "The Equity Premium: A Puzzle." Since then the topic has been researched to try and solve the apparent puzzle, and Bansal and Yaron (2004) seek that in their paper, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles."

The paper, Bansal and Yaron (2004), utilizes an approach to the equity premium puzzle based on Epstein and Zin recursive preferences (1989), persistent expected growth rates, and economic uncertainty (time-varying consumption volatility). It was originally intended for myself to build up to this model of preferences after working with a model containing Constant Relative Risk Aversion (CRRA), but due to the time constraints, it was soon realized that doing recursive preferences under Epstein and Zin (1989) would require another term of research replication. As a result, I focus my methodology on modeling the equity risk premium in a power utility framework while utilizing the processes in Bansal and Yaron (2004) with modified parameters to generate data to fit this framework and effectively show the puzzle. Additionally, Bansal and Yaron (2004) split their paper up into two cases. Case 1 consists of a persistent expected growth rate component, while case 2 incorporates all of case 1 plus stochastic volatility. For my replication I will focus on case 1.

Data Generation

The goal of the data generation phase was to generate data across 840-time periods, which are interpreted as months, that will be annualized to 70 years of consumption and dividend growth rate data. The reasoning behind monthly data is due to the fact that the decision interval of the agent (assuming representative agents) is monthly according to Campbell and Cochrane (1999), and that aggregating the months up into years leads to desired features found in the real world economic data. Consumption data is assumed to be real per-capita nondurables and services data.

To generate the data for the consumption and dividend growth rates contained in the paper, I utilized three different approaches to the AR(1) process contained within both rates:

1. Creation of the AR(1) process via a proprietary code of the equation for an AR(1) process, given below in equation (1). Then create the two versions of the dividend and consumption growth rates:
    a. Compounded
    b. Additive
2. Simulate the rouwenhorst procedure in Julia to generate equation (1). Then use the results in a compounding framework for the growth rate processes.

All of these approaches may seem a bit redundant on the surface as they should do the same thing, but that is essentially the purpose—confirm my program for an AR(1) process is accurate in my growth rate data equations. Analysis of the three approaches will examine the difference, if any, of the rates for consumption and dividends generated via the different persistent component approaches and how each compare with one another. The equations for consumption and dividend growth rate processes that incorporate the persistent component x_t are below as (2) and (3), respectively.

$$x_{t+1} = \rho x_t + \varphi_e \sigma e_{t+1} \tag{1}$$

$$g_{t+1} = \mu + x_t + \sigma n_{t+1} \tag{2}$$

$$g_{d,t+1} = \mu_d + \phi x_t + \varphi_d \sigma u_{t+1} \tag{3}$$

$$e_{t+1}, n_{t+1}, u_{t+1} \sim N.i.i.d.(0,1)$$

Persistence Without Recursion, Issues

The AR(1) process of x_{t+1} contains a very high persistence parameter at ρ = 0.979, which means the previous value of x_t is nearly passed through to the next period completely. A paper by Barsky and DeLong (1993) have shown that a value of 1 is reasonable for ρ, but Bansal and Yaron (2004) go with a parameter value less than one to maintain stationarity. Still, the persistence parameter is relatively large, and that is necessary to capture the volatility of price-dividend ratios and hence the risk premium on equities. This would be reasonable based on the conclusions the authors find in the paper, but only if I was using Epstein and Zin preferences (1989). Given the fact that I used a CRRA setup, the model does not fit the data very well for backing out equity risk premiums intended for Epstein and Zin preferences (1989). In fact, it produces very small if not negative results. This is largely due to the fact that agents do not care about the persistence aspect about rates, which will impact their long-term wealth and, hence, their decisions today. Instead, the agents care about the near-term shocks to their consumption and it how it impacts their consumption behavior. To address this issue, I shift g_{t+1} from dependence on x_t to x_{t+1} to capture an expectation of the growth rate of x under a probability distribution, and also, end up changing ρ = 0.979 to ρ = 0. This eliminates the persistence aspect of x, and makes the autocorrelation coefficients generated in Bansal and Yaron (2004) very different from mine, as I essentially eliminated autocorrelation moments from my processes. To maintain the ergodic standard deviation of the distribution previously dependent on a high level of persistence, I set the distribution's standard deviation to,

$$\varphi_e = \sqrt{\frac{(\varphi_e^2)}{(1-\rho^2)}} = 0.216 \tag{4}$$

$$N(\rho x_t, \sigma^2) \rightarrow N(0, \frac{\sigma^2}{(1-\rho^2)})$$

The problem still exists however, that the random shock component on dividends and consumption is independent of one another and reduces the correlation between consumption growth rates and dividend growth rates. As I will explain in the methodology below, it is critical that the two processes move closely together to generate equity risk premiums. This leads to changing the parameters again to fit the power utility model, as opposed to the Epstein and Zin (1989) preference model. To maximize the equity risk premium, I shutdown both the independent random shock components attached to dividends and consumption, and instead feed shocks through the system via x_{t+1} to consumption and dividends. Next, I select values for the leverage parameter (ϕ), which is the degree of leverage on expected consumption growth according to Abel (1999), in the dividend process so that it matches the standard deviation of dividend growth. Selecting a value of 4.12 is required for this to match, and although it is above that given in the paper, I don't think it is unreasonable as there have been ϕ values as high as 6.2 in Kiku (2006), but that was on value stocks only and whether or not 4.12 is too high for an aggregate market dividend ϕ parameter is a topic for another analysis. Finally, the standard deviation component in x, must be matched to that of the standard deviation of consumption growth. Doing these parameter changes help the processes move together, and therefore drive an equity risk premium that stems from the covariance. The computation of the new parameters is as follows,

$$\varphi_e = \frac{\frac{\sigma_g}{12}}{\sigma} = 0.353$$

Note that the original value of σ_g (not listed in the table) was used to compute this new value for φ_e and that σ is a given parameter in the paper that is multiplied with φ_e. A similar approach for ϕ was used to arrive at ϕ = 4.12.

Approach 1 – *Compounded*

Generating the AR(1) process via my own proprietary code consisted of generating a process that contained 1840 periods and 1000 simulations of that process. Then I selected the last 840 periods of each simulation so that the process would not be starting off from zero, but rather starting at a point where the benefits of large samples are realized. Using this AR(1) process, I then generated the monthly data for consumption, via a matrix that consisted of 840-time periods (every component replicated in this paper is ran through a Monte Carlo simulation with 1000 trials, hence the matrix would have 1000 rows). The next step is to compound the rates, which due to everything being in log form, everything is expressed as decimals, so a 1 needs to be added to every rate to make it compoundable. Otherwise the system will shrink to zero, and the economy and dividends would become non-existent. Finally, I take all the compounded rates from $t_1 \rightarrow t_{840}$ and generate a function to pull every 12$^{\text{th}}$ rate and divide it by the previous 12-month rate to create a vector with 70 entries of consumption growth rates, g_{t+1}, which are now considered annual. This same approach was used for generating the dividend growth process, $g_{d,t+1}$, just with the equation for the dividend data instead.

The results found that the average annual growth rate of consumption, about 1.8%, was able to be matched by my simulations, with the exception of some fluctuation in the hundredths of a percent from time-to-time when I would run the model. To eliminate this, I attempted to change the number of simulations from 1000 to 2000, 5000, and 10,000. However, each increase in the number of simulations yielded no noticeable change, as I could easily hit the same value with 1000 simulations at the same amount of consistency as 10,000—as a result I did not feel the need to run statistical significance tests given the outcomes.

Approach 1 – Additive

The second version of the first approach based on data across time, without given states of the world, was just a simpler version of the compounded method described above. To do this, I ran through the simulations of consumption growth and dividend growth in a monthly format, but then instead of compounding the rates, I pulled the sum of every set of 12 periods and placed them into a vector until I had 70 entries. This result was similar to the previous compounded example and was really another confirmation of the reasonability behind output generated in the compounded process, with the exception of some slightly different rates and moments given the fact there was no compounding.

Figure 1 – *Values based on the following parameters,*

$$\rho = 0, \mu = \mu_d = 0.0015, \sigma = 0.0078,\ \varphi_e = 0.353,\ \varphi_d = 0, \phi = 4.12$$

Variable	Historical Average	Mean (model)	95%	5%	p-Value *(% of data that is above the estimate)*
$\sigma_{(g)}$	2.93	2.90	3.35	2.52	0.466
$\mu_{(g)}$	1.8	1.83	2.40	1.29	0.525
$\sigma_{(gd)}$	11.49	3.96	4.50	3.41	0.0
$\mu_{(gd)}$	n.a.	1.81	2.58	1.04	n.a.
Corr(g,g_d)	0.55	.3288	0.505	0.140	0.012

Note: All rates in this paper are considered real, not nominal.

Figure 1 shows the historical values of the moments of interest, and my estimates for those moments in my generated data. There is one thing to note, the standard deviation of dividend growth ($\sigma_{(gd)}$), is much lower in my data than that of the historical average. This is due to the shutdown of the random shock component attached to the dividend process, which instead receives stochastic shocks via x_{t+1}. One other moment of interest is the correlation, and although mine is a little lower than the average, as I miss the average barely with my interval, my model still produced results consistent with research findings on the equity risk premium under power utility.

Approach 2 – States-Based Markov Process

Approach 2 focuses on generating the growth rates based on states of the world that will used in simulations for the equity premium puzzle. The reason for generating states is because it is assumed that returns and consumption depend on the state of the world, such as an expansion or recession. To generate the data based on states of the world, I utilized the rouwenhorst procedure in Julia. The rouwenhorst procedure creates an AR(1) process for a variable, and a Markov chain for that variable. I utilized 10 states for the world in this procedure. These 10 states range anywhere on a spectrum from a boom to a bust in the economic cycle. Changing the number of states to 2 and then to 100 made little to no difference, similar to that of changing the number of simulations, and so I thought it was reasonable to assume 10 states given the complexity of economic cycles.

In utilizing rouwenhorst as an approach to model the growth rates based on states of the world, I would simulate a path of states of the world then pull the corresponding growth rates for each state, and eventually aggregate them up like in the time-only dimension constructed in approach 1. Upon aggregation to annual data, the results to this process were nearly identical to those found in approach 1, making **figure 1** is still relevant for the growth process data generated via the rouwenhorst method.

Methodology

Asset Pricing Intuition

Modeling the equity risk premium begins with looking at what agents want to do in a CRRA model, and in their life in general—maximize expected utility. To do so, I first provide a utility function of the following form (Cochrane, 2001),

$$U(c_t, c_{t+1}) = U(c_t) + \beta \mathbb{E}_{t+1}[U(c_{t+1})] \tag{5}$$

c_t denotes consumption at time t, and c_{t+1} for consumption at time t+1.

The CRRA utility model is (Cochrane, 2001),

$$U(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma} \tag{6}$$

(Note as $\gamma \to 1, U(c_t) = \ln(c_t)$, however this will serve no purpose as I never use 1 for γ.)

In this basic model, it is unknown what consumption will be tomorrow, and nor does the agent know their wealth tomorrow. This model does not capture any wealth preferences for the long term, but rather just focuses on how consumption makes the agent feel more or less wealthy in the short term. Using power utility of the form in equation (6) captures the fact that agents experience more utility as consumption increases, but for additional units of consumption, the benefit of each unit is less than the previous unit. It also models the behavior that agents want a smooth consumption stream over time, but are risk averse and impatient. Beta (β) contains the degree of impatience and gamma (γ) is the measure of risk aversion and the inverse of γ measures intertemporal elasticity of substitution (IES). Risk aversion is defined as the desire to maintain consumption stability across different states of the world, while IES is the degree to which an agent wants to maintain smooth consumption across time. Decreasing β will increase impatience, and increasing γ will lead to an increased desire to smooth consumption across time and states of the world (Cochrane, 2001). Mathematically, increasing γ, will lead to increased concavity in the utility function, which is explained the narrative in the previous sentence.

Introducing assets into the mix, agents (investors) want assets because they help to provide consumption smoothing over time. Assume a payoff of x_{t+1} (note this x_{t+1} is not the same as the AR(1) x_{t+1} talked about previously) at time *t+1*, and a purchase/sell amount of j. Then essentially the goal the is to accurately price this payoff based on how it affects consumption—at what price is the equilibrium where traders stop exchanging claims on the asset. To see this, define the values of consumption in the following ways (Cochrane, 2001),

$$c_t = e_t - p_t j \tag{7}$$

$$c_{t+1} = e_{t+1} + x_{t+1} j \tag{8}$$

Knowing that agents want to maximize their utility, the equation for utility needs to be maximized under a first-order condition (FOC) (Cochrane, 2001),

$$\max_j U(c_t) + \beta \mathbb{E}_{t+1}[U(c_{t+1})] \tag{9}$$

Maximizing this equation under j after substituting in for consumption leads to (Cochrane, 2001),

$$P_t U'(c_t) = \mathbb{E}_{t+1}[\beta U'(c_{t+1}) x_{t+1}] \tag{10}$$

Then solve for p_t, and the price of the asset with payoff x_{t+1} becomes (Cochrane, 2001),

$$p_t = \mathbb{E}_{t+1}\left[\beta \frac{U'(c_{t+1})}{U'(c_t)} x_{t+1}\right] = \mathbb{E}_{t+1}\left[\beta \frac{{c_{t+1}}^{-\gamma}}{{c_t}^{-\gamma}} x_{t+1}\right] \tag{11}$$

This equation now relates prices of assets with payoff x_{t+1} to the change in marginal utility, which is considered the growth rate of consumption, and to the value of β. To make this equation more general, and back out the traditional core asset pricing equation (Cochrane, 2001),

$$m_{t+1} = \beta \frac{U'(c_{t+1})}{U'(c_t)} \tag{12}$$

Note that m is sometimes referred to as the stochastic discount factor (SDF). Also, $x, m,$ and $\mathbb{E}$ are assumed to be in t+1, so the equation can be written without time subscripts (Cochrane, 2001),

$$p = \mathbb{E}(mx) \tag{13}$$

The general intuition behind this simple equation could be thought of as discounting expected cash flows—just as one would use a rate in the appropriate time period to discount them, I do the same thing here based on consumption growth (Cochrane, 2001). Given this general form of pricing assets, $p = \mathbb{E}(mx)$, it is clear that the price of an asset can be broken down into 2 components by covariance rules (Cochrane, 2001),

$$cov(m, x) = \mathbb{E}(mx) - \mathbb{E}(m)\mathbb{E}(x) \tag{14}$$

$$p = \mathbb{E}(mx) = \mathbb{E}(m)\mathbb{E}(x) + cov(m, x) \tag{15}$$

As a result of this covariance framework, the general intuition that the riskiness of assets depends on a covariance relationship with the SDF is evident. This says that if the covariance between the SDF and payoff, x, is positive, then the price will be higher than when $cov(m, x) = 0$. The reason is to follow, but first, the $cov(m, x) = 0$ when the payoff is certain (the covariance between a constant and a variable is 0). This means that the payoff must be that of a risk-free bond. Setting the risk-free payoff of x = 1**,** yields $p = \mathbb{E}(m)$, so now the price is just that of the agent's SDF, and the return would be the payoff, 1, divided by the price (Cochrane, 2001),

$$R_f = \frac{1}{\mathbb{E}(m)} \tag{16}$$

The intuition behind what drives risk-free rates in this model can be shown by expanding (16) into the following (Cochrane, 2001),

$$R_f = \frac{1}{\mathbb{E}\left[\beta \frac{{c_{t+1}}^{-\gamma}}{{c_t}^{-\gamma}}\right]} = \frac{1}{\beta}\frac{{c_{t+1}}^{\gamma}}{{c_t}^{\gamma}} \tag{17}$$

This relationship shows that the risk-free rate rises when β is small, and people are impatient. The logic is that to prevent agents from consuming a lot due to their impatience, interest rates must be high enough to postpone consumption. In addition, when consumption growth is high, the risk-free rate must increase to contain consumption growth. For the degree of risk aversion, there is a positive relationship with the risk-free rate because as risk aversion (γ) increases, the risk-free rate increases. Lastly, the IES becomes smaller as γ increases since it is the inverse of risk aversion, however, when IES becomes smaller, rates must rise to still postpone the increased desire to consume today.

Assets that don't have $cov(m, x) = 0$, will either have a price higher or lower than the expected discounted payoff ($\mathbb{E}(m)\mathbb{E}(x)$). For assets with positive covariance, the price is higher and therefore the return is lower. The logic comes from the fact that if an asset pays well when marginal utility is high, then the risk is lower. This stems from marginal utility being high in bad states of the world, like recessions, because each additional unit of consumption has such a high value, so consumption must be relatively low. This means that the asset will pay well when the agent is in a bad state of the world (think insurance) and is not able to consume as much as usual, and the insurance counteracts that decline in consumption by paying out in bad states of the world—smoothing consumption (Cochrane, 2001). This convention of insurance means the agent is willing to pay a higher price to have that sense of stability in their consumption and receive a lower return on the asset. A negative covariance implies a lower price, and this would come from assets whose payoffs vary positively with consumption, and negatively with marginal utility. The investor is not willing to pay as much for these assets because they provide more consumption when marginal utility is already low, and hence consumption is high, so each additional unit has less value to the agent. In bad states of the world, when the agent is feeling worse off than normal due to low consumption, these assets only amplify the effect instead of reducing it since they have poor payoffs in bad states of the world, and good payoffs in good states of the world—consumption is not being smoothed under these assets.

To show the relationship when $cov(m,x) \neq 0$ and the price is assumed to be 1 in a concrete form including the risk-free rate (R_f) and a risky asset (R_e), see below (Cochrane, 2001),

$$1 = \mathbb{E}(mR_e) = \mathbb{E}(m)\mathbb{E}(R_e) + cov(m, R_e) \tag{18}$$

$$1 = \frac{\mathbb{E}(R_e)}{R_f} + cov(m, R_e) \tag{19}$$

$$1 - cov(m, R_e) = \frac{\mathbb{E}(R_e)}{R_f} \tag{20}$$

$$R_f - \frac{cov(m, R_e)}{\mathbb{E}(m)} = \mathbb{E}(R_e) \tag{21}$$

This shows that the expected return on the risky asset is equivalent to the risk-free rate less any asset that positively covaries with the stochastic discount factor. This is intuitive because when a person buys an asset that has negative covariance, that means that marginal utility is high when returns are low. As a result, the counter is true: if marginal utility is high, then consumption is low—investors need additional returns to compensate them for holding such an asset that pays poorly in bad states of the world since they want to smooth consumption.

Generating Simulations for Returns

To generate returns and essentially the equity risk premium, I wanted to model everything based on states of the world and based on different levels of risk aversion/IES.

For constructing the actual model in Julia, I used the rouwenhorst procedure again, and wanted to model the basic equation $p = \mathbb{E}(mx)$ but in a stationary way since prices are expected to rise over time, making the path nonstationary. To compensate for this, I utilize the price-to-dividend ($\frac{p}{d}$) ratio to maintain stationarity within the data. As prices rise, their increases are fueled by growth in expected dividends as prices are just the discounted expected payoffs (dividends in this case) of the asset. Thus, the two move together and I am able to back out a stationary process by dividing prices by dividends (Cochrane, 2001). In doing this, I was able to back out the desired equations that were simulated above. The breakdown is shown below in equations (22-24) (Sargent & Stachurski).

$$\frac{p_t}{d_t} = \mathbb{E}_{t+1}\left[m_{t+1}\left(\frac{d_{t+1}}{d_t} + \frac{p_{t+1}}{d_t}\right)\right], \qquad where\ x_{t+1} = d_{t+1} + p_{t+1} \tag{22}$$

$$v_t = \frac{p_t}{d_t} = \mathbb{E}_{t+1}\left[m_{t+1}\frac{d_{t+1}}{d_t}\left(1 + \frac{p_{t+1}}{d_{t+1}}\right)\right] \tag{23}$$

$$v_t = \mathbb{E}_{t+1}\left[{g_{t+1}}^{-\gamma}(g_{d,t+1})(1 + v_{t+1})\right] \tag{24}$$

Since the equations g_{t+1} and $g_{d,t+1}$ are in log form, I take the exponential of them to get the actual rates,

$$V_t = \mathbb{E}_{t+1}\left[\exp({g_{t+1}}^{-\gamma})\right] * \mathbb{E}_{t+1}[\exp(g_{d,t+1})] * \mathbb{E}_{t+1}[1 + V_{t+1}] \tag{25}$$

The equations (22 – 25) above were the basis for moving into the states-oriented version just as I did in the data generation phase. An analysis of the $\frac{p}{d}$ ratio at time *t* shows that it depends on g_{t+1} and $g_{d,t+1}$, which both are dependent on x, the AR(1) process, at time *t+1* after changing x in the data generation phase to be in the same period as the growth rates. This results in the general states version of the equations having a format of that listed below in (Bansal & Yaron, 2004),

$$V(s) = \mathbb{E}[\exp(g(s')^{-\gamma})] * \mathbb{E}[\exp(g_d(s'))] * \mathbb{E}[1 + V(s')] \tag{26}$$

$$V(s) = \mathbb{E}[\exp(\mu + x(s) + \sigma n(s'))^{-\gamma}] * \mathbb{E}[\exp(\mu_d + \ \phi x(s) + \varphi_d \sigma u(s'))] * \mathbb{E}[1 + V(s')] \tag{27}$$

where s is the state at time t+1, and s' is the state in the next period

To simplify the model for solving, rearrange the equation as,

$$V(s) = \ \mathbb{E}[a(s)b(s)(1 + V(s')], \tag{28}$$

$$where\ a(s)\ is\ \exp(g^{-\gamma})\ and\ b(s)\ is\ \exp(g_d)$$

Then taking the expectation over $g^{-\gamma}$ and over g_d for each state and value of γ, I used Gauss-Hermite quadrature due to the two equations containing random normally distributed components (the shocks). This was still a necessary process after shutting down their individual shocks in the data generation phase. Then, taking expectations over $V(s')$, I used the transition matrix given the rouwenhorst process. Doing so leads to (Sargent & Stachurski),

$$V(s) = \sum_{s'} a(s)b(s)P(s,s')\big(1 + V(s')\big) \tag{29}$$

$$h(s,s') = a(s)b(s)P(s,s') \tag{30}$$

$$V(s) = \sum_{s'} h(s,s')(1 + V(s')) \tag{31}$$

From this point, one can assume a steady state for the $\frac{p}{d}$ ratio and therefore $\frac{p}{d}$, at time *t* is equivalent to the discounted value of $\frac{p}{d}$, at time *t+1* so that now it is possible to solve for $\frac{p}{d}$ (Sargent & Stachurski).

$$V = H(1 + V) \tag{32}$$

$$V = H1 + HV \tag{33}$$

$$V - HV = H1 \tag{34}$$

$$(I - H)V = H1 \tag{35}$$

$$V = (I - H)^{-1}H1 \tag{36}$$

After backing out a solution for the $\frac{p}{d}$ ratio based on each state and value of risk aversion, I generated a matrix that contained the value of the risk-free rate in each state of the world and for each level of risk aversion so that I would have the risk-free rate to use in calculating the risk premium. This came in the form of the intuition introduced earlier in equation (16).

Finally, with the equations defined for each state of the world and value of risk aversion, I was able to run simulations of the $\frac{p}{d}$ ratio, risk-free rate, dividend growth rate, and therefore returns and the risk premiums connected to those returns. Below in the results section is a chart showing the equity risk premium, risk-free rate, as well as the volatility of the equity return (risk-free has no volatility due to their being no stochastic components in this model for it).

Results

Figure 2

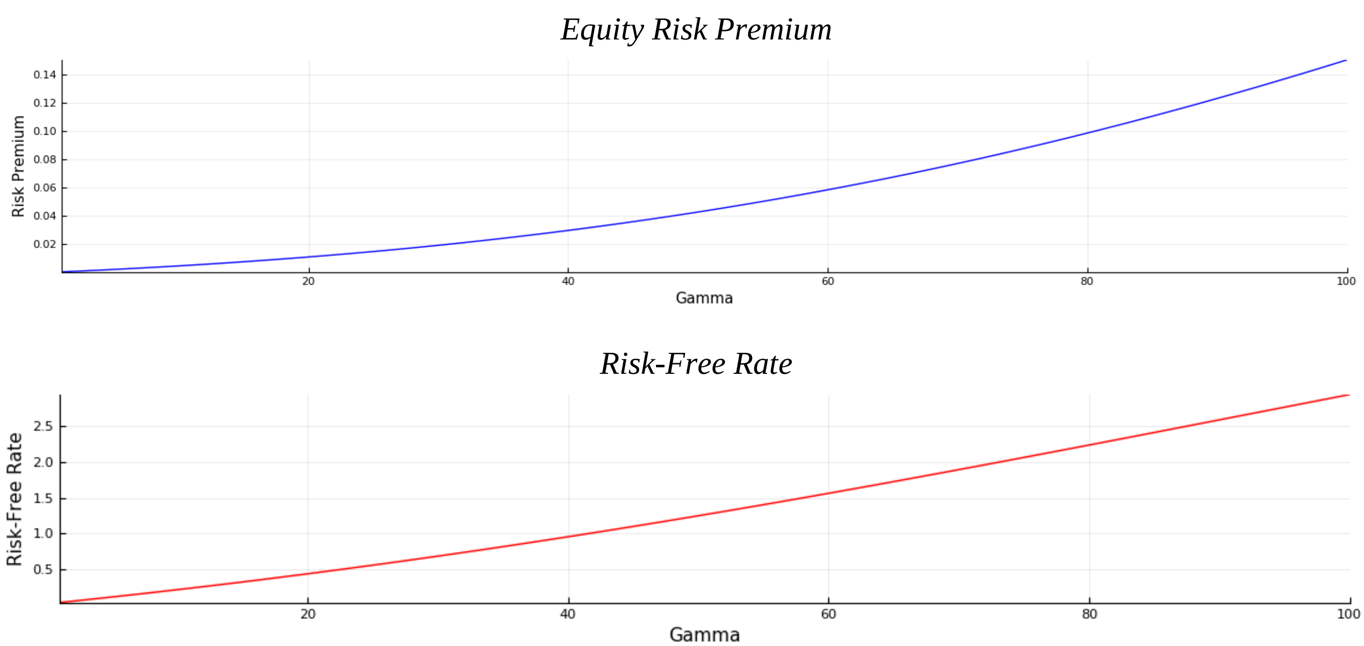

Equity Return Volatility

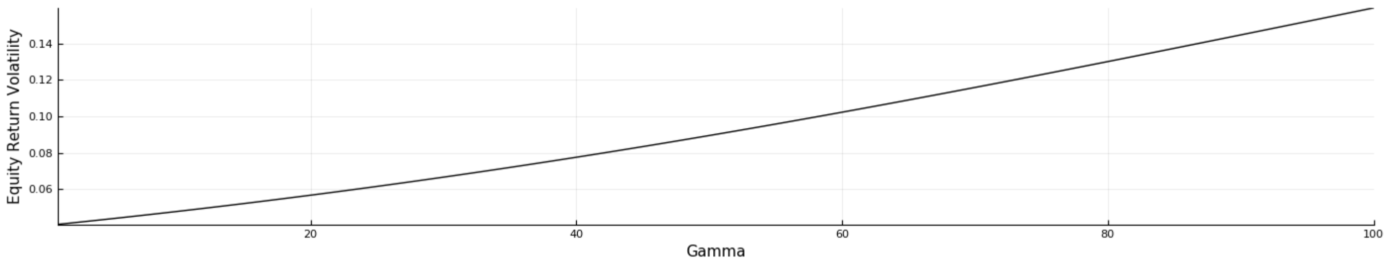

The interpretation of the output values in each of these charts is in log form, so *1.5* on the risk-free rate axis implies a rate of 150%. Additionally, the equity return volatility chart is easily confused with the risk premium chart, however, they are different. Looking at the charts it might not be quite as clear as the table 1 below, which shows the real dilemma: to reach a premium within the range of the historical interval, risk aversion must exceed 20, and even then, looking at a risk aversion of 50, the premium hasn't even touched the upper bound of the interval for historical risk premiums. This value of risk aversion is problematic because Mehra and Prescott (1985) suggest that a reasonable maximum for risk aversion is 10.

To reach a premium within the interval, risk-free rates are going to exceed 44.2%. But that is far to unrealistic for anything seen in the U.S. treasuries history because if that was the case the U.S. would have debt cost issues and inflation would be very large. Even if one found the highest historical nominal rate in the United States it would never come near 44.2%. An even more interesting question to pose is that given today's interest rate environment and the Federal Reserve's limited flexibility in traditional monetary policy through the federal funds rate, how are equity returns still this high and sometimes even higher?

The disparity that has become clear in the output of the model as it showcases two puzzles: one is termed the *equity premium puzzle* and the second is the *risk-free rate puzzle*. Essentially the puzzles are two sides of the same coin. The equity premium puzzle views the dilemma from a point that to generate such a high return implies unfathomable values of risk aversion that are not realistic, and hence cannot exist (Equity). The risk-free rate puzzle on the other hand deals with the question of why the real rate of interest is so low, as investors could buy equities to bring down the equity return and boost the real interest rate (Risk-Free). But even with these two puzzles at play there is a third component that the model cannot equate to historical averages, and that is volatility of equity returns.

Volatility in this model is much lower than historical averages. For a risk premium of about 6%, agents would have to have a risk aversion value of 60, but even at that, volatility although increasing over γ, is only 10% compared to a historical estimate of 17% (Gilchrist). For the volatility the model does produce that is close to this historical estimate, risk aversion would have to be roughly 100 or more. The reason behind why the volatility chart shows volatility of the equity return increasing in risk aversion is due to the following equation setup for returns in my simulation,

$$\left(\frac{V_{t+1}+1}{V_t}\right)\exp(g_{d,t+1}) \tag{37}$$

This equation has V increasing in risk aversion but remaining constant under i.i.d. shocks. The dividend growth rate on the other hand, is constant under risk aversion, so the equity return volatility will just scale up with risk aversion due to V being an increasing function of it.

Table 1

	Historical Average	γ = 2	γ = 5	γ = 10	γ = 20	γ = 50
Risk Premium	3.5% - 5.5% (Maverick)	0.065%	0.195%	0.442%	1.06%	4.22%
Risk-Free Rate (real)	1% (Cochrane, 2001)	6.17%	12%	22.1%	44.2%	125%

Conclusion

This paper took the processes generated in Bansal and Yaron (2004) and transformed them from being applicable to Epstein and Zin (1989) preferences to being applicable to a power utility model under constant relative risk aversion. In doing so, the separation of the intertemporal elasticity of substitution and risk aversion is blurred and becomes a single parameter in the model and labeled, γ. This leads to the study of a representative agent's behavior under changing values of γ, as changing the other major parameter of interest, β, is left fixed at what Bansal and Yaron (2004) select for the degree of impatience, believed to be reasonable.

The output of the model showcases an empirical issue in the historical value of the equity risk premium and leads one to what is conceived to be a puzzle of two-dimensions: the ridiculously high equity premium given that risk aversion cannot take on value above 10 and be reasonable, and a huge risk-free rate that has never before been witnessed in U.S. treasury security history to justify the risk premium. This begs the question that if risk aversion cannot be that big, and risk-free rates never reach that high of a level, why hasn't there been a balance between the risk-free rate and the premium so that investors exert downward pressure on equity returns and upward pressure on the risk-free rate so that the premium declines within a reasonable level of risk aversion? Unfortunately, beyond the conclusion of a puzzle being apparent in the data based on this model, no other clear conclusions can be drawn as they are beyond the scope of the power utility model. Epstein and Zin (1989) preferences utilized in Bansal and Yaron (2004) are one part of an approach to explain the puzzle found in this model, but that is for separate analysis in another paper.

Works Cited

Abel, A. B. (1990). Asset Prices Under Habit Formation and Catching up with the Joneses. *American Economic Review 80*, 38-42.

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

Barksy, R & Delong, B. J. (1993). Why does the Stock Market Fluctuate? *Quarterly Journal of Economics 108*, 291-312.

Campbell, J. Y. & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggreagate stock market behavior. *Journal of Political Economy 107,* 205-251.

Cochrane, J. H. (2001). *Asset Pricing.* Princeton, New Jersey: Princeton University Press.

Epstein, L & Zin, S. (1989). Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica 57*, 937-969.

Equity Premium Puzzle – EPP. (2018). Retrieved from: https://www.investopedia.com/terms/e/epp.asp

Gilchrist, S. Lecture 1: Asset Pricing and the Equity Premium Puzzle [PowerPoint slides]. Retrieved from: http://people.bu.edu/sgilchri/teaching/EC%20745%20Fall%202013/Lecture%20Slides/lecture1_equity_premium.pdf

Kiku, D. (2006). Is the Value Premium a Puzzle? *Job Market Paper* [PDF]. Retrieved from: http://econ.ucsb.edu/~mkapicka/RG/kiku.pdf

Maverick, J.B. What is the Historical Market Risk Premium? [Website]. Retrieved from: https://www.investopedia.com/ask/answers/040715/what-historical-market-risk-premium.asp

Mehra, R & Prescott, E. C. (1985). The Equity Premium: A Puzzle. *Journal of Monetary Economics 15, 145-161.*

Sargent, T. J. & Stachurski, J. *Asset Pricing I: Finite State* Models [Website]. Retrieved from Lectures in Quantitative Economics Website: https://lectures.quantecon.org/jl/markov_asset.html

Risk-Free Rate Puzzle – RFRP. (2018). Retrieved from: https://www.investopedia.com/terms/r/rfrp.asp