Review of Financial Economics 22 (2013) 169–179

Contents lists available at ScienceDirect

Review of Financial Economics

journal homepage: www.elsevier.com/locate/rfe

Asset pricing under quantile utility maximization☆

Bruno C. Giovannetti *

Department of Economics, University of Sao Paulo, Brazil

ARTICLE INFO

Article history:
Received 14 September 2012
Received in revised form 29 March 2013
Accepted 28 May 2013
Available online 11 June 2013

JEL classification:
G11
G12

Keywords:
Asset prices
Downside risk
Quantile utility

ABSTRACT

"Focus on the downside, and the upside will take care of itself" is a famous quote among professional investors. By considering an agent who follows this advice, we reproduce the first and second moments of stock returns, risk-free rate and consumption growth. The agent's behavior toward risk is analogous to a relative risk aversion of about 3 under expected utility, the elasticity of intertemporal substitution is about 0.5 and the time discount factor is below 1. In particular, the proposed model separates time and risk preferences in an innovative way.

1. Introduction

A famous quote among professional investors is "Focus on the downside, and the upside will take care of itself". In this paper, we consider a consumer–investor who follows this advice. Surprisingly, the consumption-based asset pricing model that emerges from this idea explains the main existing puzzles found within the asset pricing literature. These include the equity premium and the risk-free rate puzzles, the countercyclicality of the equity premium and the procyclicality of the risk-free rate.

In the proposed model, the consumer–investor is concerned with the so-called downside risk. This is done by replacing the standard setting of expected utility optimizing agents with the concept of quantile utility. Under this framework, the agent summarizes a risky situation using a worst-case scenario which is a function of his downside risk aversion. The more downside risk averse the agent, the worse the worst-case scenario he considers. The τ quantile of a continuous random variable can be interpreted as the worst possible outcome that can occur with probability $1-\tau$. Hence, instead of maximizing the expected value of his utility function, the agent maximizes a given τ quantile of it. As we will see, τ defines his downside risk aversion: the lower τ, the higher the downside risk aversion.[1]

The crucial difference between quantile and expected utility is straightforward. Under expected utility, an agent, when facing a situation where he has to choose among uncertain alternatives, picks the one that maximizes the expected value of his utility function. However, under quantile utility, the agent picks the one that maximizes some given quantile of the utility distribution, instead of its mean. For instance, the given quantile can be the median of the utility distribution, or the 0.25 quantile. In the case of the 0.25 quantile for example, when evaluating an uncertain situation, he looks at the worst outcome that can occur with 75% probability (i.e., the chance of the realized scenario being better than the scenario he considers is 75%).

The choice rule of quantile utility was axiomatized by Rostek (2010). It nests the famous maxmin and maxmax decision criteria. Indeed, decision makers who select an alternative that offers the highest minimal or maximal payoff can be viewed as maximizing the lowest or the highest quantile, respectively. Maxmax and maxmin have been applied in broad literature, as game theory, robust control, individual and social choice, bargaining, and voting. However, these criteria have been commonly criticized for basing choice on what may be extreme and unlikely outcomes (maxmin agents would not invest, would not drive, and so on). The quantile utility captures more moderate preferences while preserving the qualitative properties

☆ I would like to thank Dennis Kristensen for invaluable support on this work. I also thank Andrew Ang, Pierre-André Chiappori and Marcelo Moreira for constant and crucial advice, and Ricardo Brito, John Donaldson, Guilherme Martins, Marcos Nakaguma, Walker Hanlon, Ricardo Reis, Bernard Salanié, and the participants at the Columbia Econometrics Colloquium, Applied Micro Colloquium and Finance Colloquium, and the Economics Seminars at Fundação Getúlio Vargas (EESP), Insper and Universidade de São Paulo for important comments and discussions. All errors left are only mine.

* Tel.: +55 11965253300.
E-mail address: bcg@usp.br.

[1] One could say that the agent's objective function is given by the value at risk (VaR) of his utility. However, since τ here is a free parameter defining preference toward risk, it is not restricted to being close to zero (as in standard VaR applications).

1058-3300/$ – see front matter
http://dx.doi.org/10.1016/j.rfe.2013.05.008

of maxmin and maxmax that the expected utility does not exhibit, such as ordinality and robustness, as discussed in Rostek (2010).

We present a novel extension of the static decision-theoretical framework axiomatized by Rostek (2010) for a dynamic asset pricing setting. In a two-period standard economy with one risky and one risk-free asset, we can derive an arbitrage-free asset pricing model, where both main characteristics of the canonical expected utility consumption-based approach (Hansen and Singleton (1982), Mehra and Prescott (1985), hereinafter, the canonical model) are modified. The equity premium is no longer based on the covariance between the risky return and the consumption growth. Instead, it is a linear function of the risky return standard deviation. In addition, risk aversion and elasticity of intertemporal substitution (EIS), which are linked throughout a single parameter in the canonical model, are automatically disentangled in a simple way.

These two endogenous changes are the main drivers of the good empirical results. Since stock returns historically have a high standard deviation, the price of such a risk, i.e., the level of downside risk aversion, will not have to be high to match the empirical excess returns. Moreover, the attitude toward intertemporal substitution is not polluted by risk preferences.

To reproduce (i) the first and second moments of the risk-free return, the equity premium, and the consumption growth, (ii) the low covariance between risky return and consumption growth, (iii) the countercyclical risk premium, and (iv) the procyclical risk-free rate that we see in data, our model requires only three parameters related to preferences: a downside risk aversion (τ) of about 0.43, an EIS (ψ) of about 0.5 and a time discount factor (β) of less than 1. A downside risk aversion of such a magnitude is reasonable in that it produces reasonable certainty equivalents for bets on continuously distributed random variables (stock indexes, for example). By comparing certainty equivalents under quantile and expected utility maximization, an agent with this level of downside risk aversion is analogous to an expected utility agent with a relative risk aversion coefficient of 3. According to Mehra and Prescott (1985) reasonable values for such a parameter would be between 1 and 10. An EIS of about 0.5 is also an acceptable value. In a recent work using microdata, Engelhardt and Kumar (2009) estimate the EIS to be 0.74, with a 95% confidence interval that ranges from 0.37 to 1.21. Using macrodata and separating stockholders from nonstockholders, Vissing-Jorgensen (2002) estimates the EIS around 0.4 and 0.9 for these respective groups.

To illustrate the main differences between the predictions of our framework and the predictions of the canonical model, we first derive equations in closed-form for the risky return, the risk-free rate, and the equity premium. These equations come from combining the Euler equations of the quantile agent with the standard assumption of joint lognormality of returns and consumption growth. In order to replicate the well-evidenced existence of predictability in future excess returns, we then allow for time-varying economic uncertainty in the aggregate economy dynamics. From this, a countercyclical risk premium and a procyclical risk-free rate are produced. Taking the model to data, we first perform simulation exercises, matching the first and second moments of consumption growth, risk-free rate and excess returns. Then, we estimate the model free of distributional assumptions using a simple two-step procedure.

Since the quantile agent summarizes a risky situation using a worst-case scenario, our model considers the fact that people care asymmetrically about good and bad outcomes. Therefore, it belongs to the class of models related to asymmetric preferences, such as Epstein and Zin (1990, 2001), Bekaert, Hodrick and Marshall (1997), Barberis, Huang and Santos (2001), Routledge and Zin (2010), and Feunou, Jahan-Parvar and Tédongap (2013).

The good empirical results from, for instance, Barberis, Huang and Santos (2001) and Routledge and Zin (2010), indicate that consideration of asymmetric preferences over good and bad outcomes is a promising path for theories on choices and, in particular, for a well-accepted resolution of the asset pricing puzzles. Nevertheless, such models have a large number of preference-related parameters, which is crucial for their success, and this is a delicate issue.[2]

First, it is not easy to translate the models into a comprehensive view of the whole process. Second, it is hard to assign precisely the corresponding importance of each parameter to the obtained results. Finally, and perhaps most problematic, matching data by augmenting the parametric dimension is subject to the standard over-fitting critique. According to this critique, the larger number of parameters may simply describe better the noise in the data, rather than the underlying economic relationships. In other words, these models could be providing spurious data-fitting.[3]

The present paper helps to clarify such issues. The developed model is quite parsimonious, requiring only three preference-related parameters: the time discount factor; the EIS; and the downside risk aversion. At the same time, it solves the main asset pricing puzzles addressed by Barberis, Huang and Santos (2001) and Routledge and Zin (2010). Given that, this study makes two important contributions to the literature. Given its ability to explain the financial puzzles parsimoniously, it (i) offers a simpler view regarding the relationship between asymmetric preferences and financial data, and (ii) provides evidence that the good empirical results obtained by the studies employing asymmetric preferences are not due to over-fitting.

The rest of this work is organized as follows. Section 2 presents the quantile utility agent in its general form and derives some basic results of asset pricing under quantile maximization. Section 3 solves the model under lognormality and simulates from it. Section 4 discusses how to estimate the model free of distributional assumptions and presents the results. Section 5 concludes.

2. Quantile utility maximization and asset pricing

In this section, we first present the elements of the quantile utility model, following Manski (1988) and Rostek (2010). Then, we apply this theoretical-decision framework to asset pricing.

2.1. Quantile utility maximization elements

A general choice theory for quantile maximizing agents was developed recently. Rostek (2010) is the first study to axiomatize the quantile utility agent. Notwithstanding, the quantile maximization model for decision making under uncertainty was first proposed 23 years ago by Manski (1988).

The main idea is simple. An agent, when facing a situation where he has to choose among uncertain alternatives, picks the one that maximizes some given quantile of the utility distribution instead of its mean, as in the expected utility model. In this framework, the agent cares about the worst outcome that can happen with a given probability. For instance, the given quantile can be the median of the utility distribution, or the 0.25 quantile. In the case of the 0.25 quantile for example, when evaluating an uncertain situation, he looks at the worst outcome that can occur with 75% probability (i.e., the chance of the realized scenario being better than the scenario he considers is 75%).

The quantile of concern is an intuitive measure of pessimism. If agent *A* looks at the worst that may happen in 90% of the situations, i.e., quantile 0.10, and agent *B* looks at the worst that may happen

[2] Barberis, Huang and Santos's (2001) model has six parameters related to preference. Routledge and Zin's (2010) model has five.

[3] This tense relationship between the augmentation of the expected utility framework with additional parameters and the over-fitting critique is raised, for instance, by Zin (2002). Based on that article, Watcher (2002) claims that "behavioral models leave room for multiple degrees of freedom in the utility function. Taken to an extreme, this approach could reduce structural modeling to a tautological, data-fitting exercise" and "I believe that parsimony lies at the root of what Zin refers to as reasonableness. A parsimonious model is a model in which the number of phenomena to be explained is much greater than the number of free parameters."