Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Lieven Baele, Joost Driessen, Juan Londono and Oliver Spalt

Cumulative Prospect Theory and the Variance Premium

DP 12/2014-067

Cumulative Prospect Theory and the Variance Premium*

Lieven Baele† Joost Driessen‡ Juan M. Londono§ Oliver Spalt¶

December 2014

Abstract

Cumulative Prospect Theory (CPT) can explain the variance premium puzzle. We solve a simple equilibrium model with CPT investors and find that probability weighting plays a key role in generating a substantial variance premium, while loss aversion captures the equity premium. Using GMM on a sample of U.S. equity and index-option returns between 1996 and 2010, our estimate of the probability distortion parameter implies that real-world investors in option markets distort probabilities significantly, but less so than subjects in lab experiments. We also show that the CPT model prices the cross-section of out-of-the-money index options well. In a dynamic setting, probability weighting and time-varying equity return volatility combine to match the observed time-series pattern of the variance premium.

JEL Classification: C15, G11, G13

Keywords: Cumulative prospect theory, distorted probabilities, loss aversion, variance risk premium

*The authors greatly benefited from discussions with Nicholas Barberis, Sebastian Ebert, Lei Sun, and seminar participants at Tilburg University and the 2012 EFA meeting in Copenhagen.

†Finance Department, CentER, Tilburg University. E-mail: Lieven.Baele@uvt.nl.

‡Finance Department, CentER, NETSPAR, Tilburg University. E-mail: J.J.A.G.Driessen@uvt.nl.

§Division of International Finance, Federal Reserve Board of Governors. E-mail: juan-miguel.londono-yarce@frb.gov (corresponding author).

¶Finance Department, CentER, Tilburg University. E-mail: O.G.Spalt@uvt.nl.

1 Introduction

A central empirical success of prospect theory is its ability to explain major puzzles in financial economics. For example, Benartzi and Thaler (1995) show that prospect theory can help in understanding the equity premium puzzle, and Barberis, Huang and Santos (BHS, 2001) demonstrate how a dynamic prospect theory framework can simultaneously generate a high equity premium, predictability of equity returns, and excess volatility. In these and other settings, prospect theory offers a rigorous and testable alternative framework to more traditional asset pricing theories.

In this paper, we show that prospect theory can also help us understand one of the most important recent asset pricing puzzles: the variance premium. The variance premium, defined as the difference between the option-implied and the expected realized variance of stock returns, is strongly positive, on average, and time varying. Theoretically, the variance premium is a major puzzle because the standard consumption-based model with constant relative risk aversion (CRRA) preferences cannot generate a nonzero variance premium, irrespective of the risk aversion level and even when consumption variance varies over time (see, for example, Drechsler and Yaron (2011)).

Empirically, the variance premium is a first-order phenomenon. For example, Coval and Shumway (2001), Driessen and Maenhout (2007), and Eraker (2013) find that the Sharpe ratio for volatility-selling strategies, such as shorting straddles, which effectively bet on the variance premium, is at least twice the Sharpe ratio of the underlying equity index.[1] If the equity premium is a puzzle, then, by this metric the variance premium is an even bigger puzzle. A separate but closely related phenomenon, which we also address, is that prices of out-of-the-money put options seem to be much higher than what is predicted by standard option pricing models (see, for example, Bondarenko (2003a,b), Jones (2006), and Santa-Clara and Saretto (2009)).

Our new approach to explaining the variance premium puzzle and the pricing of out-of-the-money options builds on the work of BHS, who consider a representative agent model in which the agent's preferences are the sum of a CRRA utility function and a prospect theory value function. The value function is piecewise linear with a kink at the reference level, thus capturing loss aversion. Loss aversion is the main driver of the results in BHS. The main innovation in our paper is adding probability

[1]See also Bakshi and Kapadia (2003), Jiang and Tian (2005), Bakshi and Madan (2006), Carr and Wu (2009), Bollerslev et al. (2011), and Bollerslev et al. (2009), among others, for further evidence on the variance premium.

weighting, the second integral part of the cumulative prospect theory (CPT) of Tversky and Kahneman (1992), to the framework. Because our focus is on pricing options, we do not explicitly model consumption and dividends but directly depart from a lognormal distribution of stock market returns. The resulting model allows us to numerically solve for the equilibrium expected return on the stock market index and on a range of options on the index. From these equilibrium returns, we can then derive the model-implied variance premium.

We generate three new findings, which we detail below. First, comparative statics show that our simple model with CPT investors can generate a substantial variance premium, similar to empirically observed levels, while at the same time fitting the equity premium. Second, we use GMM and actual equity and option returns data to determine the optimal CPT preference parameters. Our estimates imply a substantial degree of probability distortion, and generate a very good fit of the observed variance premium and returns on out-of-the-money options. Third, if we allow for time-varying volatility, the CPT model generates substantial time variation in the variance premium, even with constant preference parameters, mainly because the return volatility and probability distortion interact in a nontrivial way.

In the first part of our analysis, to explore the mechanics of the CPT model, we examine the comparative statics of the equilibrium equity premium and variance premium with respect to the main input parameters. We find that the probability distortion has a strong effect on the variance premium, which makes intuitive sense: as the investor overweights extreme outcomes, the risk-neutral variance implied by option prices increases, which directly translates into a higher variance premium. The other CPT parameters (loss aversion and reference level) have a small effect on the variance premium, as they have only a negligible effect on how the representative investor values higher moments of the return distribution. In contrast, the equity premium is very sensitive to loss aversion, in line with Barberis et al. (2001). These different exposures of the equity and variance premiums to the various CPT parameters create the potential to fit both premiums with CPT preferences, which is an attractive feature of this framework.

In the second part of the paper, we turn to a formal estimation of the preference parameters from actual data. Using GMM, we estimate the parameters by fitting the following moment conditions: the equity premium, the variance premium, and expected call and put option returns across various strike prices. We use U.S. data on S&P 500 equity returns and S&P option prices from 1996 to 2010. We construct monthly call and put option returns for 13 different strike levels. The result from this GMM estimation is that, even though we fit 4 parameters to 28 moment conditions, we obtain a very good fit of the equity premium, variance premium, and cross-section of expected option returns. In addition, we find that, for our benchmark setup, the estimate of the Tversky-Kahneman probability distortion parameter is 0.79 with a standard error of 0.04, which clearly rejects the null hypothesis of no probability distortion (probability distortion parameter equal to 1). Our estimate of the probability distortion parameter implies that real-world investors in option markets distort probabilities significantly, but less so than subjects in lab experiments.

In the final part of the paper, we extend the static setup and explore the extent to which the model can generate time variation in the variance premium, focusing on three main sources: time variation in equity market volatility, probability distortion, and loss aversion. To the best of our knowledge, our paper is one of only a few attempts to examine how probability weighting among investors in asset markets changes over time. We find that, even when probability distortion is kept fixed at its benchmark value of 0.79, we still capture the dynamics of the variance premium rather well once we allow for time-varying volatility. The intuition for this result is straightforward: the power of probability weighting comes from overweighting the tails of the wealth distribution, and higher return volatility translates into larger tails, which, in turn, are then overweighted more. As a result, the variance premium is positively correlated with the volatility level, in line with empirical observations.

When we extend the model to also include time variation in probability distortion, distortion increases after losses, which in turn implies a higher variance premium after losses (which typically occur when volatility is high). When we include time variation in loss aversion, loss aversion also increases after losses. This dynamic behavior of loss aversion is one of the key assumptions in the model of Barberis et al. (2001), and our results thus provide further empirical support for their work.

Our work adds to the literature that studies the usefulness of prospect theory in finance and to a growing number of recent papers exploring the usefulness of probability weighting in asset pricing contexts (see Barberis (2013) for a recent survey). Related papers that use a prospect theory framework with probability weighting include Driessen and Maenhout (2007), who empirically investigate portfolio choice of CPT-style investors when options are part of the asset menu; Barberis and Huang (2008), who theoretically investigate the pricing of assets with positively skewed returns; and De Giorgi and Legg (2012), who examine dynamic portfolio choice with narrow framing and probability weighting. Several empirical asset pricing studies provide evidence that is consistent with a role for CPT's probability weighting feature (for example, Green and Hwang (2012) for IPOs, Boyer and Vorkink (2014) for stock options, and Ilmanen (2012) for a recent survey). Polkovnichenko and Zhao (2013) and Kliger and Levy (2009) are related papers that also use index options prices to identify CPT parameters. However, none of the above papers focus on the variance premium. To the best of our knowledge, our paper is the first to apply prospect theory to the variance premium.

We also contribute to a recent literature that studies potential explanations for the variance premium puzzle. Bollerslev et al. (2009), Drechsler and Yaron (2011), and Eraker (2009) follow the long-run risk approach of Bansal and Yaron (2004), which features recursive preferences and persistent time variation in consumption and dividend volatility. Drechsler (2013) extends this framework by incorporating Knightian uncertainty. Also following a long-run-risk model strategy, Londono (2014) provides a model to understand the sources of the volatility premium in an international setting. Bakshi and Madan (2006) show how the desire of agents to buy protection against extreme events is related to the variance premium. Bekaert and Engstrom (2010) show how a positive variance premium can emerge through a combination of preferences with habit formation and a non-Gaussian distribution for dividend and consumption growth. Gabaix (2012) proposes a model with time variation in the probability of disasters to explain various asset pricing puzzles, including the variance premium. Miao et al. (2012) analyze a setting with ambiguity-averse agents and regime switches in consumption and dividend growth.

A key difference between those approaches and this paper is that CPT preferences generate a substantial variance premium even when asset returns are i.i.d. lognormal. Relative to the CRRA-lognormal benchmark model, most of the previous papers modify preferences and return generating processes (by using Epstein-Zin preferences, long-run risks, and time-varying macroeconomic volatility, for example), whereas our approach can generate a substantial variance premium by modifying preferences only. Conceptually, our approach can be viewed as an attempt to understand the variance premium in a theoretical framework with a minimum number of additions to the CRRA-lognormal benchmark. Another important and potentially testable difference is that the above models are necessarily dynamic, whereas the CPT model can generate a variance premium even in a one-shot game with known probabilities. It is therefore possible that some of the variance premium is driven by a fundamentally different mechanism than suggested by the existing literature.[2]

The remainder of this paper is organized as follows. Section 2 introduces a static version of our model. In section 3, we perform a comparative static exercise that conveys the intuition for most of the results that follow. We then use a more formal GMM approach to match our CPT model to actual data in an unconditional setting. In section 4, we explore the ability of CPT to characterize the time-varying nature of the variance premium. Section 5 concludes.

2 The Baseline Model

We start our analysis by considering a static representative agent model—the most parsimonious setting that allows us to derive our key insights. We extend the framework to a time-varying setting in section 4.

2.1 Preferences

Let the representative agent's total utility function, $\Psi(W_T, X_T)$, have two additively separable components, which are defined below: (1) a standard concave utility function over final wealth W_T and (2) a standard prospect theory value function over gains and losses X_T:

$$\Psi(W_T, X_T) = U(W_T) + bV\left(X_T\right), \tag{2.1}$$

where b is a scaling term that governs the relative importance of the prospect theory part. The specification in equation (2.1) nests the pure CRRA utility case for $b = 0$. The utility over wealth, $U(W_T)$, is defined as a traditional CRRA function:

$$U(W_T) = \begin{cases} \frac{W_T^{1-\gamma}}{1-\gamma} & \text{for} \quad \gamma \neq 1 \\ \ln W_T & \phantom{\text{for}} \quad \gamma = 1 \end{cases}, \tag{2.2}$$

where γ is the risk aversion coefficient. The value function, $V(X_T)$, is piece-wise linear and captures the agent's concern about the return of her portfolio in terms of gains and

[2]Rare disaster models, such as the one in Gabaix (2012), present the polar opposite to our approach. These models maintain the CRRA nature of preferences, but change the return generating process. Like our model, those models could also, in principle, generate a variance premium in a static setting (albeit for very different reasons).

losses:

$$V(X_T) = \begin{cases} X_T & \text{for} \quad X_T \geq 0 \\ \lambda X_T & \phantom{\text{for}} \quad X_T < 0 \end{cases}, \tag{2.3}$$

where the parameter λ controls the degree of loss aversion (first-order risk aversion) and X_T represents gains or losses defined relative to a reference level W_{Ref}. In the static case, we define $X_T = W_T - W_{Ref}$, so that gains and losses are themselves a function of wealth.

There are two distinct advantages to using the preference structure in equations (2.1) through (2.3). The first is technical: because $\Psi(W_T, X_T)$ is a concave function in W_T, we can use standard maximization techniques to analyze its properties. Our setting is, therefore, not subject to the well-known problem whereby a pure prospect theory model with an s-shaped value function can lead to non-finite optimal weights in a portfolio choice context.[3] The second is conceptual: $\Psi(W_T, X_T)$ is the static analogue to the utility function of the representative investor in BHS, so we can both build on their insights and gauge what we gain by our extensions of their model.

While $\Psi(W_T, X_T)$ is motivated directly by BHS, we depart from their approach by incorporating probability weighting. This additional feature will be the key driver behind our novel results. In contrast to the investor in BHS, our investor does not take expectations using actual probabilities when evaluating an investment. Rather, probabilities are transformed into decision weights in a specific manner developed in Tversky and Kahneman (1992). Specifically, we assume that there are N discrete states of the world at time T, each occurring with objective probability p_i, and that each state is associated with a specific final wealth level $W_{T,i}$. Decision weights are then obtained as follows. First, states are ordered from worst to best relative to the investor's reference wealth level, W_{Ref}, according to: $W_1 \leq ... \leq W_{k-1} \leq W_{Ref} \leq W_{k+1} \leq ... \leq W_N$ (this notation implies that $W_k \equiv W_{Ref}$). Then, the subjectively distorted probability of

[3] Alternative ways to model aversion to extreme losses is to concavify the value function below a certain point, as in Gomes (2005), or to explicitly impose portfolio constraints to rule out extreme positions due to the convexity of the value function, as in Ait-Sahalia and Brandt (2001). Arguably, adding the concave CRRA part to investor preferences also reflects better the natural tendency of investors to become averse to portfolio positions that threaten their economic existence—an extreme case outside the scope of decisions for which prospect theory was designed (Kahneman and Tversky (1979), p. 278).

outcome i, π_i, is given by:

$$\begin{aligned}\pi_i &= w^-(p_1+...+p_i) - w^-(p_1+...+p_{i-1}) \text{ for } 2 \le i \le k \\ \pi_i &= w^+(p_i+...+p_N) - w^+(p_{i+1}+...+p_N) \text{ for } k+1 \le i \le N-1,\end{aligned}$$

where

$$w^-(p) = \frac{p^{c_1}}{\left[p^{c_1}+(1-p)^{c_1}\right]^{1/c_1}}, \tag{2.4}$$

$$w^+(p) = \frac{p^{c_2}}{\left[p^{c_2}+(1-p)^{c_2}\right]^{1/c_2}},$$

where $\pi_1 = w^-(p_1)$ and $\pi_N = w^+(p_N)$. Parameters c_1, $c_2 \in [0.28; 1]$ in equation (2.4) control the curvature of the weighting function for losses and gains, respectively.[4] This functional form for the decision weights implies that the tails of a distribution are subjectively overweighted more if c becomes smaller. If $c_1 = c_2 = 1$, there is no probability distortion.

In general, the decision weights, π_i, do not add up to 1 and cannot, therefore, be interpreted as distorted probabilities. To remedy this, we define the distorted probabilities as $\pi_i^* = \frac{\pi_i}{\sum_i \pi_i}$. Such a scaling of the decision weights does not matter for the investor's portfolio choice, as long as the decision weights do not depend on the portfolio chosen by the investor. As we show in section 2.2, in our setup, the investor holds the market portfolio in equilibrium, which directly fixes the ordering of wealth outcomes and the associated decision weights.

In this CPT setup, the investor then chooses a portfolio that maximizes the function

$$E^*\left[\Psi(W_T, X_T)\right] \equiv \sum_i \pi_i^* \Psi(W_{T,i}, X_{T,i}). \tag{2.5}$$

The model nests the expected utility benchmark as a special case ($c_1 = c_2 = 1$ and $b = 0$).

2.2 The Investor's Problem and the Expected Return on Equity

We consider a simple market with the following assets: (i) a risk-free asset with a constant and exogenously given gross return R_f; (ii) equity that pays x_i^E in state i

[4] The lower bound is a technical condition that ensures that decision weights are positive (Rieger and Wang (2006) and Ingersoll (2008)).

of the world at time T, which occurs with probability p_i; and (iii) M derivatives on equity, where derivative $m \in [1, M]$ pays $x_i^{D,m}$ in state i of the world at time T. The derivatives are in zero net supply, and we normalize the supply of equity to be one. The expected gross returns on equity, $E(R^E)$, and on each derivative m, $E(R^{D,m})$, are to be determined in equilibrium. The representative investor's problem is to choose a one-period optimal portfolio with positions in the risky asset, α_E, each derivative m, $\alpha_{D,m}$, and the risk-free asset such that she maximizes:

$$\max_{\alpha_E, \alpha_{D,m}} E^* \left[\Psi(W_T, X_T) \right]. \tag{2.6}$$

In any time-T state of the world i, wealth is given by

$$W_{T,i} = \left[(1 - \alpha_E - \sum_m \alpha_{D,m}) R_f + \alpha_E \frac{x_i^E}{s_0^E} + \sum_m \alpha_{D,m} \frac{x_i^{D,m}}{s_0^{D,m}} \right] W_0,$$

where W_0 is the investor's initial wealth and s_0^E and s_0^D denote time-zero prices of equity and the derivatives, respectively.

The key to deriving the equity risk premium in this model is to note that, as the risk-free asset and the derivatives are in zero net supply, market clearing requires $\alpha_E = 1$ and $\alpha_{D,m} = 0$ for all m. Substituting these equilibrium weights into the wealth constraint and the wealth constraint into the objective function yields the following first-order condition:

$$\sum_i \pi_i^* \left[\frac{x_i^E}{s_0^E} - R_f \right] \left[(W_0 \frac{x_i^E}{s_0^E})^{-\gamma} + b \left(1 + (\lambda - 1) \, I_{X_i < 0} \right) \right] = 0, \tag{2.7}$$

where π_i^* are the distorted probabilities and $I_{X_i<0}$ is an indicator function that equals 1 for losses. The first-order condition in equation (2.7) can be solved for the equilibrium price of equity, s_0^E, and the model-implied expected return on equity is then given by $E\left(R^E\right) = E\left(x^E\right) / s_0^E$. With an exogenous reference point, equation (2.7) provides a pricing kernel that we can use to price any asset in this economy that is in zero net supply—in equilibrium, the representative investor will not hold such assets in her portfolio. For example, we can price an option with payoffs $x_i^{D,m}$ by solving the following equation for the price $s_0^{D,m}$:

$$\sum_i \pi_i^* \left[\frac{x_i^{D,m}}{s_0^{D,m}} - R_f \right] \left[(W_0 \frac{x_i^E}{s_0^E})^{-\gamma} + b \left(1 + (\lambda - 1) \, I_{X_i < 0} \right) \right] = 0 \tag{2.8}$$

and the associated option return is then $E\left(R^{D,m}\right) = E\left(x^{D,m}\right)/s_0^{D,m}$.

2.3 The CPT-Implied Variance Premium

We can calculate the CPT-implied variance premium as follows. Define the actual probability measure P as the measure under which the actual equity payoffs x_i^E are generated, and Q as the risk-neutral measure under which equity and options are priced. The difference between P and Q is determined by the preferences of our representative agent. In our one-period discrete-state setting, the variance premium is defined as the difference between the risk-neutral variance and the actual variance of log equity returns:

$$\begin{aligned} VP &\equiv V^Q - V^P \\ &= \sum_i q_i \left[ln\left(\frac{x_i^E}{s_0^E}\right) - E^Q\left[ln\left(\frac{x_i^E}{s_0^E}\right)\right]\right]^2 \\ &\quad -\sum_i p_i \left[ln\left(\frac{x_i^E}{s_0^E}\right) - E^P\left[ln\left(\frac{x_i^E}{s_0^E}\right)\right]\right]^2, \end{aligned} \tag{2.9}$$

where q_i denotes the risk-neutral probability of state i, which is formally defined as

$$q_i = \frac{\pi_i^* \left[(W_0 \frac{x_i^E}{s_0^E})^{-\gamma} + b\left(1 + (\lambda - 1)\, I_{X_i<0}\right)\right]}{\sum_i \pi_i^* \left[(W_0 \frac{x_i^E}{s_0^E})^{-\gamma} + b\left(1 + (\lambda - 1)\, I_{X_i<0}\right)\right]}. \tag{2.10}$$

Equation (2.10) summarizes the preferences of the investor and illustrates how our setting differs from more traditional models. There are three defining components of q_i. First, the difference between the risk-neutral probability q_i and the actual probability p_i depends on the derivative of the CRRA part of the investor's utility function in equation (2.1). This is the standard component in traditional CRRA-based asset pricing models. Second, q_i depends on loss aversion and the reference point via the derivative of the value function (the second term inside the square brackets). Third, probability weighting transforms actual probabilities p_i into distorted probabilities π_i^*. The model collapses to the standard CRRA model for $c_1 = c_2 = 1$ and $b = 0$, in which case q_i is simply a function of marginal utility.

The variance premium can be calculated directly from equation (2.9). However, it is instructive to rewrite this expression. Exploiting a standard result from the option pricing literature, which we derive formally in appendix A, the risk-neutral variance,

V^Q, can be rewritten as a weighted average of puts and calls with prices $s_0^{D,m}$ across all possible strike prices. Hence, there is a direct link between understanding the variance premium and understanding the pricing of out-of-the-money options. Intuitively, one can "earn" the variance premium using a static option strategy. Specifically, by selling out-of-the-money put and call options with various strikes, one obtains a strategy that pays out $\frac{1}{R_f}V^Q$ today and has an expected payoff of $-V^P$ at time T. Essentially, this strategy is a portfolio of short positions in so-called "strangles," where each strangle is a combination of an out-of-the-money put and a call option. The variance premium, $V^Q - V^P$, represents the expected profit on this option portfolio. Hence, the variance premium reflects information about compensation for extreme outcomes on the underlying asset, as these outcomes determine the value of out-of-the-money puts and calls.

Rewriting equation (2.9) in terms of option prices shows that the variance premium aggregates expected returns on put and call options across all strikes into one number. As a result, a sufficient condition for our model to match the variance premium is to match the returns of out-of-the-money options. This has three main implications for our study. First, we can use this insight in our GMM procedure by adding option returns to the moment conditions. Second, because its construction is based on the rewritten expression for V^Q that we derive in appendix A, we can use the squared value of the VIX as a measure of V^Q in the data. Third, because fitting the variance premium by fitting options is a stricter test than fitting only the variance premium, we can implement a more informative test of the CPT model by incorporating option returns.

Finally, as an important point of reference for our subsequent analysis, note that the variance premium is zero for the CRRA case. If x^E and hence the terminal wealth of the investor W_T have a lognormal distribution and the representative agent has pure CRRA preferences ($b = 0$, $c = c_2 = 1$), Samuelson and Merton (1969) and Rubinstein (1976) show that call and put options are priced by the Black-Scholes formula, even though the investor is not allowed to trade during time 0 and T. This directly implies that the risk-neutral variance, as implied by option prices, equals the actual equity return variance and, hence, the variance premium is equal to zero.[5]

[5]Brennan (1979) in fact shows that CRRA preferences are the only preferences that lead to the Black-Scholes formula in a setting with discrete trading and lognormal returns. Thus, only CRRA preferences generate a zero variance risk premium. The intuition for this result is that CRRA investors, when faced with constant investment opportunities and lognormal equity returns, optimally choose a constant equity exposure even over longer horizons. Hence, these investors are not interested

3 CPT and the Variance Premium: Baseline Results

This section presents our results for the baseline static CPT setting. We start by performing a comparative static exercise that conveys the intuition for most of the results that follow. We then use a more formal GMM approach to match our CPT model to actual data.

3.1 Comparative Statics

We solve the model in section 2 for the equity and variance premiums under different input parameters. Our benchmark specification is as follows. We assume that the equity return, R_i^E, is lognormally distributed with parameters $\mu_E = E\left(R^E\right)$, which is to be determined endogenously, and volatility σ_E.[6] We assume a one-period constant risk-free rate of $R_f = 0.25\%$ and an unconditional volatility of equity returns of $\sigma_E = 5.18\%$. These are, respectively, the observed unconditional mean of the risk-free rate proxied by the U.S. 1-month T-bill rate and the unconditional monthly volatility of the S&P 500 index returns for the period running from January 1996 to October 2010. We parametrize preferences using standard values in the literature. Specifically, following Tversky and Kahneman (1992), we use a loss aversion parameter of $\lambda = 2.25$ and, to parameterize the probability weighting function in equation (2.4), we use $c_1 = c_2 = 0.65$. We use a benchmark coefficient of relative risk aversion of $\gamma = 1$ and a reference level R_{Ref} equal to the risk-free rate. Finally, we assume that $b = 0.65$, which, in combination with the other parameters in the benchmark specification, makes CPT's contribution to the value function equal to 50%.

3.1.1 The Variance Premium

Figure 1 shows the annualized CPT-implied variance premium as a function of the scale parameter, b, loss aversion, λ, and the probability distortion, c, in panels A to C, respectively.[7] As discussed in section 2, our model setup implies a variance premium of zero for the pure CRRA model ($b = 0$ and $c = 1$) irrespective of the assumed level

in dynamic trading over time. Given that options are equivalent to dynamic trading strategies, these investors thus price options in the same way irrespective of whether they are allowed to trade continuously (the Black-Scholes world) or not. Drechsler and Yaron (2011) show that CRRA preferences generate a zero variance premium also in a multi-period setting with long-run risk.

[6] Specifically, we use a discrete approximation of the lognormal distribution, taking a grid of 500 potential outcomes for the one-month-ahead equity return ranging from -50% to 150%.

[7] We do not vary the reference level, as it turns out to have little effect on the variance premium.

of risk aversion. By contrast, the main conclusion from figure 1 is that the CPT model can generate a substantial variance premium. Importantly, this ability to generate the variance premium derives almost exclusively from the probability distortion. Panel C shows that the variance premium tends toward zero as $c \to 1$ (no distortion), but reaches high levels of about 1000 for high levels of distortion ($c = 0.35$).

To understand the magnitude of the variance premium, it is interesting to compare the realized with the implied volatility. The actual annual variance in this setting, $12\sigma_E^2$, is equal to $322.0\%^2$ given that the actual monthly volatility is $\sigma_E = 5.18\%$. A variance premium of $1000\%^2$ thus implies a monthly risk-neutral volatility of $\sqrt{(1000 + 322)/12} = 10.5\%$, which is indeed well above the actual volatility. The effect of c on the variance premium is only slightly lower as we consider loss-averse agents ($\lambda \geq 1$), especially for highly distorted-probability scenarios ($c \to 0.35$), and holds even when the scale is set to 0—the case where agents have strictly CRRA-type preferences. Hence, our results suggest that, even in the absence of a CPT-type value function, the probability distortion has an economically meaningful impact on the implied variance premium.

The mechanism for the effect of the probability distortion on the variance premium is intuitive. In the presence of distorted probabilities, the state prices of extreme outcomes increase as the investor attaches a higher distorted probability to these events. Investors with probability weighting find insurance and lottery tickets attractive, and are willing to pay higher prices for out-of-the-money puts and calls. As discussed in section 2.3, the variance premium directly represents the expected return of a strategy that sells out-of-the-money put and call options, which implies that the variance premium increases with the degree of probability weighting. While the purpose of this comparative static exercise is simply to show the ability of the probability distortion to generate a substantial variance premium, rather than to pin down one specific value, the model-implied variance premium with $c = 0.65$ is $400\%^2$. This value is higher than what is typically observed in the data (in our sample it is 152.5^2, see section 3.2), which suggests that even moderate values of c will be sufficient to fit the data. We will investigate this issue more formally in the GMM framework in section 3.2.

Contrasting the strong effect of probability weighting, panels A and B show that the variance premium hardly moves with the scale or loss aversion parameters. For loss aversion, in panel B, we find at best a weakly negative relationship between λ and the variance premium, even if we give extremely high relative weight to the CPT part in the utility function. This result contradicts the simplistic intuition that loss-averse agents are willing to pay a higher price to cover themselves against the risk of extreme outcomes. However, as equation (2.7) shows, more loss aversion increases the equilibrium expected return on the risky asset, which makes extremely positive returns more likely and negative returns less likely. This in turn makes calls more expensive and puts less expensive, with a theoretically ambiguous overall effect on the variance premium. Empirically, the figure shows that the lower put prices dominate the higher call prices, which implies a lower variance premium as loss aversion increases. Quantitatively, however, the effect is tiny. The results from varying the scale parameter b in panel A underline that the ability of the model to generate the variance premium is largely independent of the shape of the utility function. That is, as long as we have probability weighting, it does not matter whether we give a high or low weight to the CRRA utility part.

In sum, then, the model can generate a substantial variance premium, mostly because of the probability weighting feature. Without probability weighting, both CRRA and CPT models cannot generate a variance premium. In particular, loss aversion plays no role for the variance premium, which highlights directly the usefulness of a model that incorporates probability distortion—an integral, but often neglected, ingredient of CPT.

3.1.2 The Equity Premium

Figure 2 replicates figure 1 for the expected equity return. Contrary to what we found for the variance premium, we find that the equity premium is most sensitive to loss aversion (panel B). The ability of loss aversion to generate a substantial equity premium underlies the results in Benartzi and Thaler (1995) and BHS. Figure 2 documents this very robust feature for our model. Panel B also shows that the equity premium increases more in loss aversion if we put a large weight on the CPT part of the utility function (high b). Higher λ makes the utility function more concave, which increases the effective risk aversion of the investor and, in turn, leads the equity premium to rise as loss aversion increases.

Panel A shows model-implied equity premiums obtained from varying the scale parameter b. While it is tempting from panel B to conclude that increasing the weight of the CPT part would necessarily increase the equity premium, because loss aversion matters only for the CPT part, this reasoning is incorrect. Specifically, because the utility function of the representative investor is a blend of CRRA and CPT parts, the overall effective risk aversion is a blend of the risk aversion implied by the individual components of utility. If risk aversion implied by the CPT part is lower than that of the CRRA part, increasing b can actually lead to a decrease in the equity premium, as the investor becomes more risk tolerant. The dashed line in panel A documents this finding for $\lambda = 1$.

Panel C shows that probability distortion alone cannot generate a meaningful equity premium. Without a CPT part in the utility function, or without loss aversion, the equity premium is essentially flat in c, which reflects the property that probability weighting overweights both left and right tails largely symmetrically under our assumption that $c_1 = c_2$. Interestingly, panels C and B reveal a complementarity between loss aversion and probability weighting: if the investor cares more about losses (larger λ), and if the investor cares more about the tails of the distribution (smaller c), then losses dominate and the required return for holding equity increases.

The results from figures 1 and 2 highlight the incremental usefulness of loss aversion and probability weighting. While probability weighting captures the variance premium, loss aversion captures the equity premium. Combining these two components in one model can therefore help us learn about the variance premium without losing CPT's ability, documented in Benartzi and Thaler (1995) and BHS, to fit the equity premium.

3.1.3 The Impact of Volatility

The previous results used a benchmark volatility level equal to the average historical one-month volatility of the S&P 500. A feature of actual data is that the level of volatility can vary substantially over time around this mean. It is therefore interesting to investigate how the previous results depend on the level of volatility.

Figures 3 and 4 show that both the model-implied variance premium and the equity premium depend strongly on the level of equity volatility, and increase whenever volatility increases. The model-implied variance premium (figure 3) nearly triples when monthly equity volatility is set at 10% (about 35% annually) rather than at the benchmark level of 5.18% (about 18% annually), and drops substantially for very low levels of volatility. This pattern is intuitive, as higher volatility leads to more extreme outcomes, which, in turn, are overweighted by CPT investors, leading to an even higher variance premium. The figure also shows that high volatility alone is not sufficient to generate a variance premium. Only when coupled with probability weighting does higher volatility lead to a higher variance premium (panel C). This complementarity is one we will exploit further in our time-varying setting in section 4. The patterns for alternative levels of loss aversion and relative weight of the CRRA and CPT components of the utility function remain unchanged, as those parameters play no role for the variance premium. Figure 4 presents a very similar pattern for the equity premium.

3.2 GMM Estimation

In this section, we provide more formal evidence on the ability of our CPT model to match the equity and variance risk premiums. We use GMM to estimate a set of model-implied parameters $\theta = [\lambda, R_{Ref}, b, c]'$. Following the evidence in section 3.1 on the importance of probability distortion to generate a substantial variance premium, we center the attention on the estimated degree of probability weighting, c. Because we are calculating the variance premium from a series of puts and calls with different strike prices, the moments we try to match include the returns on those options as well. This section will, therefore, highlight an additional advantage of our model, namely its ability to fit returns for out-of-the-money options. It is well known that these option returns are hard to explain with more traditional models. However, fitting option returns increases the hurdle for the model: while it is possible to match the variance premium well without matching any individual option return well, we show that our model can match both.

3.2.1 Data and Method

The option data we use run from January 1996 to October 2010. Our data contain daily closing midquotes of S&P 500 index options for various strikes and maturities obtained from the OptionMetrics database. Using these data, OptionMetrics creates a "surface" of interpolated option prices for fixed levels of the Black-Scholes delta (ranging from 0.2 to 0.8 for calls, and from -0.8 to -0.2 for puts) and fixed maturities (30 calendar days, 60 days, etc.). We focus only on the 30-day options because these are the most liquid options and the ones used more often in the literature. On each day and for each option, we construct the return on buying an option and holding it to maturity (i.e., for 30 calendar days). This gives a panel of 26 option returns across strikes (or deltas). Of these 26 options, 13 are puts with delta ranging from -0.8 to -0.2 (from in-the-money to out-of-the-money). The other 13 are calls with deltas ranging from 0.2 to 0.8 (from out-of-the-money to in-the-money).

The summary statistics for option returns are reported in table 1. The observed option returns are in line with the documented stylized facts in the literature (Coval and Shumway (2001); Jones (2006); Driessen and Maenhout (2007)). First, average call option returns are negative, in line with existing work of, for example, Bakshi et al. (2010), and in contrast to the standard CAPM, which implies positive call option returns (call options have a positive exposure to the underlying equity index, and a positive equity premium would thus imply positive expected call option returns). The negative average call returns suggests that call option prices are higher than what standard models predict (Bakshi et al. (2010)). Second, put options have strongly negative average returns. For example, for deep out-of-the-money put options, the average returns are about -55% per month. Of course, the CAPM predicts negative expected put returns because put options have a negative exposure to the equity index but, as we will show in section 3.2.2, the observed average put returns are well below the predictions of a standard CAPM-style pricing model.

We also construct the corresponding S&P 500 index returns for the 30-day holding periods at daily frequency. This gives an average return of 0.60% per 30 calendar days and a standard deviation of 5.18% per month. The average 1-month T-bill rate over the sample period is 0.25%, which gives an in-sample equity premium of 0.35% per month.

The variance premium is calculated as the difference between the risk-neutral variance and the expected actual variance of equity index returns. For the risk-neutral variance, V^Q, we use the square of the VIX, which is an implementation of the theoretical expression in equation (A.3).[8] To calculate the expected realized variance, we perform a regression of the realized variance on the one-month-lagged realized variance and the square of the VIX, where, at each point in time, the realized variance is calculated using daily returns over the last 22 trading days. At each point in time, the expected realized variance is then calculated as the forecast implied by this regression model (see Londono (2014)). In our sample, the average of the squared VIX equals $563.8\%^2$, while the average expected realized variance equals $410.4\%^2$ (annualized). This gives an average variance premium of $152.5\%^2$. These numbers are similar to those in the related literature; for example, the numbers in table 1 of Bollerslev et al. (2009) imply an annualized variance premium of $219.6\%^2$ and the average variance premium in Londono (2014) is $144.4\%^2$.

[8]The expression in equation (A.3) differs slightly from the formula used to calculate the VIX (Carr and Wu (2009), eq. (5)). The VIX is a measure of the integrated variance over a given time period, while equation (A.3) gives the variance of the log return over this time period. If returns are uncorrelated over time, both measures are equal to each other.

We use a GMM procedure to estimate the preference parameters using the following set of moment conditions:

$$m(\theta) = \begin{bmatrix} \mu_E(\theta) - T^{-1}\sum_1^T R_t^E \\ VP(\theta) - T^{-1}\sum_1^T VP_t \\ \mu_{p,1}(\theta) - T^{-1}\sum_1^T R_t^{p,1} \\ .. \\ \mu_{p,13}(\theta) - T^{-1}\sum_1^T R_t^{p,13} \\ \mu_{c,1}(\theta) - T^{-1}\sum_1^T R_t^{c,1} \\ .... \\ \mu_{c,13}(\theta) - T^{-1}\sum_1^T R_t^{c,13} \end{bmatrix},$$

where $\mu_E(.)$, $\mu_p(.)$, and $\mu_c(.)$ are the CPT-implied equilibrium expected returns for equity, calls, and puts (indexed by their strike prices), respectively, and $VP(.)$ is the model-implied variance premium, all matched to their empirical counterparts in our sample. Then, we find the set $\widehat{\theta}$ that minimizes the function $VF = m(\theta)'W_1 m(\theta)$, where W_1 is the weighting matrix for the moment conditions. We force the weighting matrix to give an equal weight of $1/4$ to the equity, call, put, and variance premium moments. We also calculate the Newey-West corrected standard errors for the estimated parameters to account for the potential autocorrelation and heteroskedasticity in the residuals due to overlapping returns as we construct holding-to-maturity returns each day. Thus, we can do inference on the estimated parameters knowing that

$$\widehat{\theta} \sim N(\theta, V/T),$$

where

$$V = [M'WM']^{-1}M'WSWM[M'WM']^{-1},$$

where $M = \delta m/\delta\theta$ and S is the standard Newey-West-corrected covariance matrix. To keep the estimation feasible, we exogenously fix risk aversion, γ, and, in most specifications, loss aversion, λ. Those curvature parameters are highly correlated with the scale parameter, b, and are, therefore, hard to jointly identify with sufficient precision. Specifically, we use the standard value of $\lambda = 2.25$ for loss aversion and CRRA risk aversion levels of 0.5, 1, 5, and 10.

3.2.2 GMM Results

Table 2 presents the GMM results. For each specification, we report CPT's contribution to the value function, the model-implied variance and equity risk premiums, and the value function evaluated at the optimum. We also report the level of risk aversion that would be needed in the no-CPT case ($b = 0$ and $c = 1$) to obtain the same model-implied equity premium.[9] We find the value function VF to be lowest when risk aversion, γ, is set to 0.5, and call this our benchmark specification (specification #1 in table 2). This specification fits the variance and equity premiums well: the model-implied premiums are $153.5\%^2$ and 0.30%, compared with in-sample estimates of $157.4\%^2$ and 0.35%, respectively. For higher levels of risk aversion, we still match the variance premium but increasingly overestimate the equity premium. This is because in our sample the equity premium is rather modest, while equity volatility is substantial. Hence, only a relatively low level of risk aversion is needed to fit the observed equity premium.[10] In sum, the first key finding is that the model can fit the observed variance premium remarkably well.

The second key finding in the table is that, for all levels of risk aversion, the probability distortion parameter, c, is estimated with high precision. Depending on the specification, we find an estimated degree of probability weighting that ranges between 0.72 and 0.79. This is somewhat higher than the Tversky and Kahneman (1992) benchmark value of $c = 0.65$, but it is remarkably close to the value of 0.71 reported in another well-known experimental study on the probability weighting function by Wu and Gonzalez (1996). We view this similarity between our results and those in the literature as a particular success for our model.

There is no guarantee that we can get anything close to experimentally observed values, and the fact that we do based on financial market data, rather than from experiments, provides further support for the view that probability weighting is a pervasive phenomenon in decision making more generally. Moreover, it is often argued that parameters based on laboratory studies would not hold in field settings with professional investors because professional investors would be less prone to behavioral biases. The fact that our estimates imply slightly less probability weighting than in lab experiments

[9]The contribution of CPT to the total utility function is calculated as $\frac{bE^*[V(X_T)]}{E^*[\Psi(W_T,X_T)]}$, see equation (2.6).

[10]In our approach, we exogenously match the observed level of equity volatility. We do not focus on explaining the relation between consumption volatility, dividend volatility, and equity volatility, as in BHS.

may indeed partly reflect that the marginal investor in option markets overweights extreme probabilities less. However, this does not take away from the fact that we still estimate substantial curvature in the weighting function, which is highly statistically significant—the standard errors on the estimate of c imply that we can reject it being equal to 1 at any reasonable significance level.

The third key finding is that we can not only match the variance premium, which aggregates the various underlying option returns, but also the underlying average option returns individually. Figure 5 shows the actual and model-implied expected returns for each of our 13 puts and calls. Clearly, the model fits the overall pattern in the data very well, producing only small estimation errors. Notably, even for the out-of-the money options, the fit is good. This finding is not trivial, as several studies have documented that prices of out-of-the-money put options are higher than predicted by the standard Black-Scholes model and that average put option returns are very low.[11]

To emphasize the latter point, figure 5 shows that the pure CRRA model fits options quite poorly, especially those far out of the money. In fact, the CRRA model yields that all calls have, on average, positive returns, in contrast to the observed call return pattern.[12] Thus, allowing for probability distortion also resolves the puzzlingly low returns on out-of-the-money call options.

The estimates discussed so far are based on a lognormal distribution for the underlying equity index returns. To assess whether our results are sensitive to the assumed return distribution, in table 3, we consider various alternative distributions. First, we use a normal distribution of returns instead of a lognormal distribution, where the volatility is again matched to the equity return data. This alternative specification (specification #2) generates very similar GMM estimates for the preference parameters. Next, we use the class of skewed student-t distributions, which are more flexible in matching the skewness and kurtosis that is observed in the data, through a skewness parameter ξ and a kurtosis parameter υ. The normal distribution is a special case of the skewed-t distribution with $\xi = 1$ and $\upsilon = \infty$. In appendix B, we describe in detail how we estimate the parameters of this distribution using Maximum Likelihood on the equity return data. If we allow both skewness and kurtosis to be different from the lognormal case (specification #3 in table 3), we find that the estimated return distri-

[11] See, among others, Bondarenko (2003b), Jones (2006), Bollerslev and Todorov (2011), and Constantinides et al. (2013).

[12] To understand why the pure CRRA model predicts positive call returns, one can apply a second-order Taylor approximation to the CRRA utility function, which gives mean-variance utility and CAPM pricing expressions. As call options have positive betas, the CAPM predicts positive expected returns.

bution exhibits negative skewness ($\xi < 1$) and substantial kurtosis ($\upsilon = 6.33$). Most importantly, for this alternative specification, the GMM estimates for the preference parameters are again similar to the benchmark case. In particular, the distortion parameter is estimated at 0.80 which is close to the benchmark estimate (0.79). The model-implied variance premium is very similar as well. The only difference is for the equity premium. Both negative skewness and higher kurtosis amplify the equity premium that is implied by the model.[13]

In sum, in this section, we show that a CPT model with probability weighting can generate a substantial variance premium, fit individual option returns, and produce precise and reasonable estimates of the probability weighting parameter in the Tversky and Kahneman (1992) probability weighting function.

4 Exploring the Time-Varying Nature of the Variance Premium

In the previous section, we showed that our model can match the unconditional level of the variance premium for a reasonable level of probability distortion. We now investigate whether it can also capture its substantial movement over time, focusing on time variation in three key parameters: equity market volatility, probability distortion, and loss aversion.

4.1 Time-Varying CPT Setting

As discussed in the previous section, figure 1 shows that the variance premium is, above all, sensitive to the level of probability distortion, while the equity premium is more sensitive to the degree of loss aversion. Figure 3, on the other hand, indicates that, for a fixed degree of distortion and loss aversion, the model-implied variance premium depends positively on the level of equity market volatility. Because the equity and variance premiums are most sensitive to them, we focus on time variation in these three parameters, keeping other parameters at their benchmark values. To the best of our knowledge, our paper is among the first to attempt to examine how probability weighting and loss aversion among investors in asset markets change over time.

[13]Obviously, for this alternative return distribution we could improve the fit to the observed equity premium by lowering the fixed values for the relative risk aversion and/or loss aversion parameters. The analysis in section 3.1 shows that this would not affect the fit of the variance premium.

We calculate the conditional monthly equity market variance as the integrated variance based on a simple GARCH(1,1) model estimated on daily returns over the 1996-2011 period.[14] The estimated GARCH(1,1) specification is

$$h_{t+1} = 1.38e - 06 + 0.91 h_t + 0.08 \left(r_t - \mu \right)^2 .$$

All estimated parameters are statistically significant at the 1% level. Consistent with previous evidence, we find the daily conditional variance process to be highly persistent but stationary. To calculate the conditional variance for the next month, we simply take the sum of the model-implied predictions of the variance for all days over the next month ("integrated variance").

We introduce time variation in probability distortion and loss aversion by letting these parameters vary with the recent performance of the representative agent's portfolio. Specifically, to stay as close as possible to the framework of BHS, we use the state variable z_t proposed by BHS as the sole driver of the time-series variation in the probability distortion and loss aversion:

$$c_t = \hat{c} + \kappa_c (z_t - \bar{z}_t),$$

$$\lambda_t = \hat{\lambda} + \kappa_\lambda (z_t - \bar{z}_t), \tag{4.1}$$

where $\hat{c}$ and $\hat{\lambda}$ are the unconditional estimates of the probability distortion and loss aversion parameters, respectively. $\bar{z}_t$ is the average of z_t and κ_c and κ_λ are, respectively, the sensitivity of probability distortion and loss aversion to recently realized gains and losses as measured by z_t. Following BHS, we assume that the representative agent compares the current price of the risky asset with a benchmark level, where the price benchmark level responds sluggishly to changes in the value of equity. The sluggishness is defined in BHS as follows: "when the stock price moves up by a lot, the benchmark level also moves up, but by less. Conversely, if the stock price falls sharply, the benchmark level does not adjust downwards by as much." Formally:

[14] Although we believe that the integrated variance from a GARCH process is a good approximation of the actual variance, in unreported results (available on request), we show that the results in this section are robust to alternative conditional variance methods, such as an exponentially weighted moving average or, even more simply, last month's realized variance.

$$z_{t+1} = \eta(z_t \frac{\overline{R^E}}{R^E_{t+1}}) + (1 - \eta), \tag{4.2}$$

where $\overline{R_E}$ is a fixed parameter calibrated to guarantee that half of the time the agent has prior gains ($R^E_t > \overline{R^E}$) and the rest of the time she has prior losses ($R^E_t < \overline{R^E}$). The parameter η measures the degree of sluggishness which can be interpreted as the agent's memory horizon. We assume $\eta = 0.5$, which implies that the half-life of the representative agent's memory is 1 month.[15]

4.2 Estimation Method and Results

In the estimation procedure, we center our attention on the time variation of the probability distortion and loss aversion. We extend the moments in the GMM procedure in section 3.2 to a time-varying setting as follows.

At each point in time t, we compare the model-implied equity and variance premiums to empirical estimates of the time-t conditional equity and variance premiums. The conditional variance premium at any time t is estimated as the difference between the square of the VIX at time t and the time-t expectation of the realized variance, as explained in section 3.2. To obtain an estimate for the conditional equity premium, we use two methods. In the first, the PD-based method, we perform a predictive regression of monthly S&P 500 returns on the price-dividend ratio and use the fitted values of this regression to obtain the empirically observed conditional equity premium at each point in time. In the second method, the PD and VP-based method, we add the variance risk premium to the predictive regression of S&P 500 returns.[16] The PD and VP-based procedure is also followed for the conditional expectations of the option returns. We thus obtain a panel of errors—at each point in time, we have model-implied and "empirically observed" values for the equity premium, variance premium, and option returns:

[15]In unreported results, we investigate the sensitivity of κ_c and κ_λ, the only parameters to be estimated, and the model fit to alternative values of η.

[16]The model in Bollerslev et al. (2009) implies that the variance premium has predictive power for equity returns, an implication for which they find empirical evidence.

$$
m_t(\kappa_c) = \begin{bmatrix} \mu_{E,t}(\kappa_c) - E_t(R^E_{t+1}) \\ VP_t(\kappa_c) - E_t(VP_{t+1}) \\ \mu_{p,1,t}(\kappa_c) - E_t(R^{p,1}_{t+1}) \\ .. \\ \mu_{p,13,t}(\kappa_c) - E_t(R^{p,13}_{t+1}) \\ \mu_{c,1,t}(\kappa_c) - E_t(R^{c,1}_{t+1}) \\ .. \\ \mu_{c,13,t}(\kappa_c,) - E_t(R^{c,13}_{t+1}) \end{bmatrix}.
$$

We then fix all parameters at their unconditional GMM estimates (section 3.2) except κ_c and κ_λ, which we estimate by minimizing the sum of the squared errors described above, thus summing across assets and over time using the same weighting of the equity premium, variance premium, and put and call returns as for the moment conditions defined in section 3.2. To avoid numerical problems, we use a grid of values for κ_c and κ_λ to identify the values that minimize the weighted sum of the squared errors for each specification.

Table 4 presents the results for alternative specifications of the time-varying CPT setting. For each specification, we also report the value function divided by the number of months and the model's fit for the equity premium, the variance premium, and an aggregate of put and call returns.

Before turning to the estimation results, we first discuss a simple case where only the equity return volatility is time varying, while the CPT parameters are fixed at their unconditional estimates and hence the time-varying parameters are not estimated. In such a model (specifications #1 and #2 in table 4), the equity and variance premiums vary over time only due to their dependence on the equity volatility level (see figures 3 and 4 and section 3.1.3). To get an idea of the model's fit, in figure 6, we show the model-implied and "empirically observed" variance and equity premiums for specifications #1 and #2 in table 4. The results suggest that, overall, our model fits the level and dynamics of the variance premium remarkably well, although it seems to overestimate the variance premium in late 2008, around the collapse of Lehman Brothers, because the equity return volatility is unusually elevated in this period. Our model also seems to fit quite well the unconditional level of equity premium, although it seems to miss some of its large movements, especially for the PD and VP-based equity premium, which displays much more time variation than the PD-based equity premium.

The model also fits the time variation in expected option returns well: the mean absolute error for option returns is 11.1% and 8.5% for puts and calls, respectively. In unreported results, we show that the residuals for option returns are of a similar magnitude irrespective of the degree of moneyness. However, our model seems to systemically underestimate the expected return of puts (and overestimate that of calls) in late 2008, precisely when we also overestimate the variance premium. In sum, when we keep all preference parameters constant over time, a model with time-varying equity volatility already generates a very good fit of the time-series variation in the variance premium.

Turning to the setting with time-varying probability distortion (specification #3 in table 4), we find that our estimate of κ_c is negative, which implies that the probability distortion increases (lower c) following losses (high z). However, the economic magnitude of this effect is small. Given that z_t ranges between 0.95 and 1.20 in our sample, with a sample average very close to 1, $\kappa_c = -0.05$ implies that the distortion parameter ranges between 0.78 and 0.79 over the sample period. Not surprisingly, the fit of the model with time variation in the distortion is very close to the fit of the model with constant parameters.

Finally, a setting with time-varying loss aversion yields a positive sensitivity of loss aversion to recent losses. Our estimates thus support one of the key assumptions in BHS, namely that loss aversion increases after losses. The estimate for κ_λ generates a substantial range in loss aversion, from about 1.1 to 6.5. However, this wide variation has a minor effect on the fit to the equity and variance premiums and to the option returns moments. The small effects on the variance premium and expected option returns are not surprising, as these are not very sensitive to the degree of loss aversion (see section 3.1). For the equity premium, the model with time-varying loss aversion still misses some of the large movements and the fit is slightly worse.

In sum, the main implication of the time-varying setting is that even a model with constant CPT parameters can generate substantial swings in the variance premium once we take into account the interplay between volatility and probability weighting. Hence, we can make progress explaining the time-series pattern of the variance premium without needing to use any additional degrees of freedom for making CPT parameters time dependent. All we need is a positive degree of probability weighting and time-variation in volatility observed in the actual data.

5 Conclusions

This paper investigates the potential of prospect theory to capture both the level and dynamics of the variance premium. We extend the representative investor model of Barberis et al. (2001), where the investor's preferences are the sum of a CRRA utility function and a prospect theory value function, by incorporating probability weighting. Probability weighting is an integral part of cumulative prospect theory (Tversky and Kahneman (1992)), and we document that the extended model is able to generate a positive variance premium. We show numerically that probability distortion is the key parameter driving the variance premium, while loss aversion is the main driver of the equity premium. The central new insight from our model is that loss aversion and probability weighting combine to jointly explain the equity premium and the variance premium.

When we estimate the preference parameters from our theoretical model using GMM on S&P 500 equity and options data, we find an estimate for the probability distortion parameter of 0.79. Hence, real-world investors in option markets distort probabilities significantly, but less so than subjects in lab experiments.

Finally, we extend the static baseline setting and explore to what extent the model can generate time variation in the variance premium. We find that the dynamic version of our model performs remarkably well once we allow for time-varying volatility, even when probability weighting is fixed. A second main insight from our paper is therefore that combining probability weighting—overweighting of extreme returns—and time-variation in volatility—the presence of extreme returns—yields a parsimonious model that matches the time-series behavior of the variance premium closely.

Our findings have a number of potentially important implications for future research. In particular, relative to the standard CRRA-lognormal model, which cannot generate a variance premium, we show that one can explain the variance premium by keeping returns i.i.d, and changing only the representative agent's preference structure. This is in contrast to most of the existing literature, which changes either preferences, the return generating process, or both. The cumulative prospect theory model can generate a variance premium even in a one shot game with known probabilities. It is therefore possible that some of the variance premium is driven by a fundamentally different mechanism than suggested by the existing literature.

References

Ait-Sahalia, Y., and M. Brandt, 2001, Variable selection for portfolio choice, *Journal of Finance* 56, 1297–1351.

Bakshi, G., and N. Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.

Bakshi, G., and D. Madan, 2006, A theory of volatility spreads, *Management Science* 52, 1945–1956.

Bakshi, G., D. Madan, and G. Panayotov, 2010, Returns of claims on the upside and the viability of u-shaped pricing kernels, *Journal of Financial Economics* 97, 130–154.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bansal, R., and A. Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

Barberis, 2013, Thirty years of prospect theory in economics: A review and assessment, *Journal of Economic Perspectives* 27, 173–196.

Barberis, N., and M. Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, *American Economic Review* 98, 2066–2100.

Barberis, N., M. Huang, and T. Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.

Bauwens, L., and S. Laurent, 2002, A new class of multivariate skew densities, with applications to garch models, *Working Paper,* CORE.

Bekaert, G., and E. Engstrom, 2010, Asset return dynamics under bad environment-good environment fundamentals, *Working Paper,* NBER.

Benartzi, S., and R. Thaler, 1995, Myopic loss aversion and the equity premium puzzle 110,, *Quarterly Journal of Economics* 110, 75–92.

Bollerslev, T., M. Gibson, and H. Zhou, 2011, Dynamic estimation of volatility risk premiums and investor risk aversion from option-implied and realized volatilities, *Journal of Econometrics* 160, 235–245.

Bollerslev, T., G. Tauchen, and H. Zhou, 2009, Expected stock returns and variance risk premiums, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *The Journal of Finance* 66, 2165–2211.

Bondarenko, O., 2003a, Statistical arbitrage and securities prices, *Review of Financial Studies* 16, 875–919.

Bondarenko, O., 2003b, Why are puts so expensive?, *working paper,* University of Illinois, Chicago.

Boyer, Brian H., and Keith Vorkink, 2014, Stock options as lotteries, *Journal of Finance,* Forthcoming.

Brennan, M. J., 1979, The pricing of contingent claims in discrete time models, *Journal of Finance* 34, 53–68.

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.

Carr, P., and L. Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311–1341.

Carr, Peter, and Dilip Madan, 2001, Towards a theory of volatility trading, *Option Pricing, Interest Rates and Risk Management, Handbook in Mathematical Finance* 458–476.

Constantinides, George M., Jens Carsten Jackwerth, and Alexi Savov, 2013, The puzzle of index option returns, *Review of Asset Pricing Studies* 3, 229–257.

Coval, J., and T. Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1009.

De Giorgi, E., and S. Legg, 2012, Dynamic portfolio choice and asset pricing with narrow framing and probability weighting, *Journal of Economic Dynamics and Control* 36, 951–972.

Drechsler, I., and A. Yaron, 2011, What's vol got to do with it?, *Review of Financial Studies* 24, 1–45.

Drechsler, Itamar, 2013, Time-varying fear, and asset prices, *Journal of Finance* 68, 1843–1889.

Driessen, J., and P. Maenhout, 2007, An empirical portfolio perspective on option pricing anomalies, *Review of Finance* 11, 561–603.

Eraker, B., 2009, The volatility premium, *Working paper,* University of Wisconsin.

Eraker, B., 2013, The performance of model based option trading strategies., *Review of Derivatives Research* 16, 1–23.

Gabaix, X., 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *The Quarterly Journal of Economics* 127, 645–700.

Gomes, F., 2005, Portfolio choice and trading volume with loss-averse investors, *Journal of Business* 78, 675–706.

Green, Clifton, and Byoung-Hyoun Hwang, 2012, initial public offerings as lotteries: Skewness preference and first-day returns, *Management Science* 58(2), 432–444.

Ilmanen, Antti, 2012, Do financial markets reward buying or selling insurance and lottery tickets?, *Financial Analysts Journal* 68, 26–36.

Ingersoll, J., 2008, Non-monotonicity of the tversky-kahneman probability-weighting function: a cautionary note, *European Financial Management* 14, 385–390.

Jiang, G., and Y. Tian, 2005, Model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

Jones, C., 2006, A nonlinear factor analysis of s&p 500 index option returns, *Journal of Finance* 61, 2325–2363.

Kahneman, D., and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Kliger, Doron, and Ori Levy, 2009, Theories of choice under risk: Insights from financial markets, *Journal of Economic Behavior and Organization* 71, 330–346.

Lambert, P., and S. Laurent, 2001, Modelling financial time series using garch-type models and a skewed student density, *Working Paper,* Universite de Liege.

Londono, J.M., 2014, The variance risk premium around the world, *Working Paper,* Federal Reserve Board.

Miao, Jianjun, Bin Wei, and Hao Zhou, 2012, Ambiguity aversion and variance premium, *Working Paper,* Federal Reserve Board.

Polkovnichenko, Valery, and Feng Zhao, 2013, Probability weighting functions implied in option prices, *Journal of Financial Economics* 107, 580–609.

Rieger, Marc Oliver, and Mei Wang, 2006, Cumulative prospect theory and the st.petersburg paradox, *Economic Theory* 28, 665–679.

Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *The Bell Journal of Economics* 7, 407–425.

Samuelson, Paul A., and Robert C. Merton, 1969, A complete model of warrant pricing that maximizes utility, *Industrial Management Review* 10, 17–46.

Santa-Clara, Pedro, and Alessio Saretto, 2009, Option strategies: Good deals and margin calls, *Journal of Financial Markets* 12, 391–417.

Tversky, A., and D. Kahneman, 1992, Advances in prospect theory, cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

Wu, George, and Richard Gonzalez, 1996, Curvature of the probability weighting function, *Management science* 42, 1676–1690.

Appendix A: Rewriting V^Q using Options

In this section, we present a more formal derivation of the statement in section 2 according to which V^Q can be restated in terms of options. For ease of notation, we switch to a setting with an infinite number of outcomes for the equity payoff x^E at time T. As shown by, for example, Bakshi et al. (2003) and Carr and Madan (2001), any continuous and twice-differentiable payoff function $G(x^E)$ can be written as

$$G(x^E) = G(\bar{x})+G'(\bar{x})(x^E-\bar{x})+\int_{\bar{x}}^{\infty} G''(K)max(x^E-K,0)dK+\int_{0}^{\bar{x}} G''(K)max(K-x^E,0)dK \tag{A.1}$$

for any value of $\bar{x}$. This equation shows that any payoff function can be replicated by an equity position and a position in an infinite number of call and put options with different strike prices. The risk-neutral expected value of the payoff function G can then be written as a function of call and put prices

$$E^Q\left[G(x^E)\right] = G(\bar{x})+G'(\bar{x})(R_f s_0^E-\bar{x})+\int_{\bar{x}}^{\infty} G''(K)R_f C(K)dK+\int_{0}^{\bar{x}} G''(K)R_f P(K)dK, \tag{A.2}$$

where we use the risk-neutral pricing equations for equity, $s_0^E = \frac{1}{R_f}E^Q\left[x^E\right]$, call options with strike K, $C(K) = \frac{1}{R_f}E^Q\left[max(x^E-K,0)\right]$, and put options with strike K, $P(K) = \frac{1}{R_f}E^Q\left[max(K-x^E,0)\right]$. Equation (A.2) shows that the risk-neutral expectation of any payoff function G can be written in terms of the prices of equity, call, and put options. To see the link with the risk-neutral variance, we choose the functions $G(x^E) = ln\left(\frac{x_i^E}{s_0^E}\right)$ and $G(x^E) = \left(ln\left(\frac{x_i^E}{s_0^E}\right)\right)^2$, and directly use (A.2) to write the risk-neutral variance as a function of option prices. Carr and Madan (2001) show that this gives

$$V^Q = E^Q\left[\left(ln\left(\frac{x_i^E}{s_0^E}\right)-F\right)^2\right] =$$

$$\int_{s_0^E e^F}^{\infty} \frac{2}{K^2}\left(1-ln(K/s_0^E)+F\right)R_f C(K)dK+\int_{0}^{s_0^E e^F} \frac{2}{K^2}\left(1-ln(K/s_0^E)+F\right)P(K)dK, \tag{A.3}$$

where the expression for $F = E^Q\left[ln\left(\frac{x_i^E}{s_0^E}\right)\right]$ is given in Carr and Madan (2001). This shows that the risk-neutral variance can be calculated directly from option prices. To follow the standard in the literature, we use this calculation approach for V^Q in our numerical and empirical work in sections 3 and 4, following Britten-Jones and Neuberger (2000), Bollerslev et al. (2009), and Carr and Wu (2009).[17]

[17]We have verified that the difference between the continuous measure and the discrete measure in equation (2.9) is very small.

Appendix B: Alternative Stock Returns Distribution: Skewed-t Distribution

In section 3.2.2, we investigate to what extent our baseline results are robust to using a skewed fat-tailed distribution instead of a standard lognormal distribution for stock returns. In this appendix, we first introduce the standardized skewed-t distribution (see, for example, Lambert and Laurent (2001) and Bauwens and Laurent (2002)). Then, we report the Maximum Likelihood estimates for the key parameters of these distribution using our sample of S&P 500 returns.

The return on the S&P 500, r_t, follows the process

$$r_t = \mu + \varepsilon_t \tag{B.1}$$

$$\varepsilon_t = \sigma \varsigma_t, \tag{B.2}$$

where the random variable ς_t is $SKST\left(0, 1, \xi, v\right)$ distributed; that is, it follows a standardized skewed-t distribution with parameters $v > 2$ (the number of degrees of freedom) and $\xi > 0$ (a parameter related to skewness). The density of this function is given by

$$f\left(\varsigma_t | \xi, v\right) = \begin{cases} \frac{2}{\left(\xi + \frac{1}{\xi}\right)} sg\left[\xi\left(s\varsigma_t + m\right)|v\right] & if \quad \varsigma_t < -m/s \\ \frac{2}{\left(\xi + \frac{1}{\xi}\right)} sg\left[\xi\left(s\varsigma_t + m\right)/\xi|v\right] & if \quad \varsigma_t \geqslant -m/s \end{cases}, \tag{B.3}$$

where $g\left(.|v\right)$ is a symmetric (zero mean and unit variance) Student-t density with v degrees of freedom, denoted $x \sim ST\left(0, 1, v\right)$, defined by

$$g\left(x|v\right) = \frac{\Gamma\left(\frac{v-1}{2}\right)}{\sqrt{\pi\left(v-2\right)}\Gamma\left(\frac{v}{2}\right)} \left[1 + \frac{x^2}{v-2}\right]^{-(v+1)/2} \tag{B.4}$$

where $\Gamma\left(.\right)$ is an Euler's gamma function.

The constants $m = m\left(\xi, v\right)$ and $s = \sqrt{s^2\left(\xi, v\right)}$ are, respectively, the mean and standard deviation of the non-standardized skewed-t density, $SKST\left(m, s^2, \xi, v\right)$, and are defined as follows:

$$m\left(\xi, v\right) = \frac{\Gamma\left(\frac{v-1}{2}\right)\sqrt{v-2}}{\sqrt{\pi}\Gamma\left(\frac{v}{2}\right)} \left(\xi - \frac{1}{\xi}\right) \tag{B.5}$$

$$s^2(\xi, v) = \left(\xi^2 + \frac{1}{\xi^2} - 1\right) - m^2. \tag{B.6}$$

The parameters ξ and v are related to the distribution's skewness and kurtosis. Because ξ^2 can be shown to be equal to the ratio of the probability masses above and below the mode, the distribution has zero skewness when $\xi = 1$, negative skewness when $\xi < 1$, and positive skewness when $\xi > 1$. The fatness of tails (kurtosis) decreases with the degrees of freedom parameter, v, but converges to a skewed normal as $v \to \infty$. When $v \to \infty$ and $\xi = 1$, this distribution collapses to a standard normal distribution.

Table B.1 reports estimates and standard errors for $\theta = (\mu, \sigma, \xi, v)$, using monthly log returns on the S&P 500 over the period running from 1960 to 2011. All parameters are statistically significant at the 1% level. ξ is significantly below 1, implying negative skewness. The ratio of the probability masses above and below the mode is equal to $\xi^2 = 0.71$, implying that the degree of negative skewness is also important in economic terms. The estimated degrees of freedom parameter, v, equals 6.33, indicating that the return distribution is not only left skewed but also fat tailed.

Table B.1: **Parameters of the Skewed-t Distribution**

This table reports the estimated parameters for the Skewed-t distribution described in equations (B.1) to (B.6). The standard errors are reported in parenthesis.

	μ	σ	ξ	v
Estimate	0.0077	0.0437	0.8448	6.3301
(s.e.)	(0.0017)	(0.0019)	(0.0487)	(1.5862)

Table 1: **Summary Statistics for Put and Call Option Returns**

This table shows the summary statistics for returns on holding S&P 500 index options with 30 calendar days to maturity, for puts and calls with various strike levels as measured by their Black-Scholes delta. Data are collected from OptionMetrics, which provides daily interpolated option prices for various deltas and maturities. The sample period runs from January 1996 to October 2010, with a daily frequency and overlapping 30-day returns. The t-statistic of the average return is calculated using Newey-West with 30 lags.

Put option returns													
Option delta	-0.800	-0.750	-0.700	-0.650	-0.600	-0.550	-0.500	-0.450	-0.400	-0.350	-0.300	-0.250	-0.200
Average return	-0.098	-0.116	-0.133	-0.151	-0.172	-0.194	-0.218	-0.244	-0.277	-0.320	-0.372	-0.442	-0.532
t-statistic average	-1.750	-1.884	-2.009	-2.144	-2.292	-2.456	-2.620	-2.802	-3.040	-3.370	-3.775	-4.337	-5.102
Return standard deviation	0.851	0.938	1.016	1.090	1.163	1.236	1.310	1.386	1.462	1.542	1.624	1.707	1.784
Call option returns													
Option delta	0.200	0.250	0.300	0.350	0.400	0.450	0.500	0.550	0.600	0.650	0.700	0.750	0.800
Average return	-0.288	-0.233	-0.190	-0.152	-0.122	-0.099	-0.078	-0.060	-0.046	-0.036	-0.027	-0.020	-0.011
t-statistic average	-2.640	-2.306	-2.017	-1.728	-1.485	-1.283	-1.082	-0.897	-0.749	-0.634	-0.518	-0.410	-0.257
Return standard deviation	2.036	1.819	1.644	1.499	1.375	1.266	1.165	1.073	0.987	0.905	0.825	0.746	0.662

Table 2: **Unconditional Setting, GMM Estimated Parameters**

This table reports GMM estimates for the degree of loss aversion, λ, reference level, R_{ref}, scale, b, and probability distortion, c, for levels of relative risk aversion of the CRRA component ranging from 0.5 to 10. The p-values to test the hypotheses $\lambda = 1$, $R_{ref} = 1$, $b = 0$, and $c = 1$ are reported in brackets. The p-values are calculated using Newey-West standard errors. For each specification, we also report CPT's contribution to the utility function, the model-implied variance and equity risk premiums, the value function evaluated at the optimum, and the level of risk aversion that would be needed in the no-CPT case to obtain the model-implied equity premium. These parameters are estimated applying GMM to returns on the S&P 500 and on S&P 500 call and put options with different strikes and 30-days to maturity and to the variance premium. We force the weighting matrix to give an equal weight of 1/4 to the equity, call, put, and variance premium moments. The sample period runs from January 1996 to October 2010.

	#1	#2	#3	#4	#5
γ (risk aversion)	0.5	1	5	10	0.5
λ (loss aversion)	2.25	2.25	2.25	2.25	9.12
					[0.49]
R_{ref} (reference)	1.01	1.01	1.00	1.00	1.01
	[0.06]	[0.27]	[0.41]	[0.36]	[0.00]
b (scale)	0.06	0.77	83,654.12	22,597.55	0.01
	[0.33]	[0.12]	[0.50]	[0.50]	[0.50]
c (distortion)	0.79	0.75	0.72	0.72	0.79
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
CPT contribution (%)	9.54	55.55	100	100	4.40
Implied Variance Premium (%^{2})	153.53	152.91	151.86	151.87	153.53
Implied Equity Premium (%)	0.30	1.27	2.12	2.12	0.29
VF	2.52	152.65	480.55	480.59	2.51
RA without CPT	1.11	4.83	8.17	8.17	1.11

Table 3: **Unconditional Setting, GMM Estimated Parameters, Alternative Equity Return Distributions**

This table reports GMM estimates for the reference level, R_{ref} , scale, b, and probability distortion, c, for equity return distributions alternative to the benchmark lognormal distribution. In particular, we compare the benchmark lognormal distribution with a normal distribution (column #2) and a skewed-t distribution (column #3). For the skewed-t distribution, the parameters driving the skewness and kurtosis—ξ and υ, respectively—are calibrated to match the observed monthly equity returns (see appendix B). In all specifications, the level of risk aversion, γ, is fixed at 0.5 and loss aversion, λ, is fixed at 2.25. For each specification, we also report CPT's contribution to the utility function, the model-implied variance and equity risk premiums, and the value function evaluated at the optimum. These parameters are fitted using GMM to returns on the S&P 500 and on S&P 500 call and put options with different strikes and 30-days to maturity and to the variance premium. We force the weighting matrix to give an equal weight of 1/4 to the equity, call, put, and variance premium moments. The sample period runs from January 1996 to October 2010. The p-values to test the hypotheses $\lambda = 1$, $R_{ref} = 1$, $b = 0$, and $c = 1$ are reported in brackets. These p-values are calculated using Newey-West standard errors.

	#1	#2	#3
	Benchmark	Normal	Skewed-t
γ (risk aversion)	0.5	0.5	0.5
ξ			0.84
υ			6.33
λ (loss aversion)	2.25	2.25	2.25
R_{ref} (reference)	1.01	1.01	1.00
	[0.06]	[0.06]	[0.31]
b (scale)	0.06	0.04	0.03
	[0.33]	[0.39]	[0.41]
c (distortion)	0.79	0.80	0.80
	[0.00]	[0.00]	[0.00]
CPT contribution (%)	9.54	6.06	5.04
Implied Variance Premium ($\%^2$)	153.53	153.05	150.48
Implied Equity Premium (%)	0.30	0.30	2.03
VF	2.52	2.34	21.00

Table 4: **Time-Varying Setting, GMM Estimated Parameters**

This table shows the estimated parameters and the model's fit for alternative specifications of the time-varying CPT setting. The specifications in columns #1 and #2 do not allow for time variation in the CPT parameters. The specification in column #1 uses the PD-based method to calculate the observed equity premium, wherein the expected equity returns are calculated using an in-sample forecast from the P/D ratio. In all other specifications, the observed equity premium is calculated as an in-sample forecast from the P/D ratio and the variance premium (PD and VP-based method). For every specification, we report the value function divided by the number of months (Avg. VF). The value function is calculated as a weighted average of the squared residuals for the variance and equity premiums as well as for the 10 options considered in the time-varying setting. We also report a decomposition of the model's fit into its variance premium, equity premium, and option return components. For each component, the fit is calculated as the mean absolute error. In all specifications, we assume $\gamma = 0.5$ (risk aversion) and $\eta = 0.5$ (agent's memory). The time-varying volatility is calculated as the integrated volatility based on a GARCH(1,1) method (see section 4.1).

	#1	#2	#3	#4
EP Method	PD	PD and VP		
κ_c (distortion)			-0.05	
κ_λ (loss aversion)				22.07
Avg. VF	87.45	87.47	87.44	85.20
EP fit (%)	0.39	0.45	0.46	0.48
VP fit (%^{2})	101.72	101.72	103.12	96.71
Put fit (%)	11.07	11.07	8.53	8.44
Call fit (%)	8.55	8.55	11.08	11.01

Figure 1: **CPT-Implied Variance Premium for Alternative Preference Parameter Values**

This figure shows the annualized CPT-implied variance premium for alternative values of the scale, b, loss aversion, λ, and probability distortion, c, parameters, in panels A to C, respectively. In the CPT setting, the representative agent's preference combines CRRA and CPT functions (equation (2.1)) and the probabilities are distorted, as explained in section 2. The variance premium is defined as the difference between the risk-neutral and the expected realized variance of equity returns, as explained in section 2.3. For the benchmark setup (the bold line) we set $c = 0.65$, $b = 0.65$, and $\lambda = 2.25$, respectively, unless otherwise indicated in the panel. We additionally assume $\gamma = 1$.

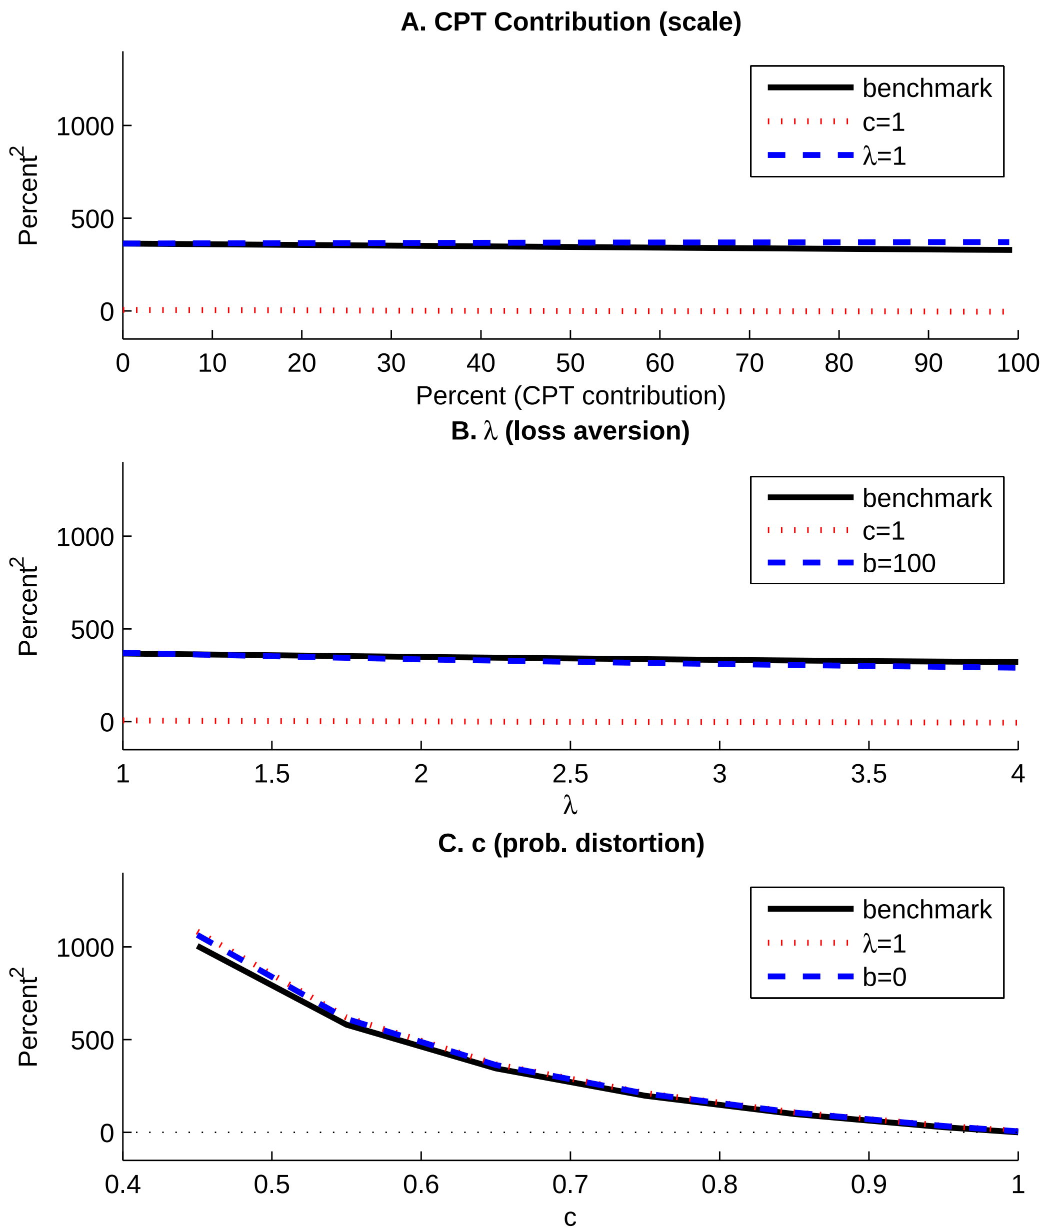

Figure 2: **CPT-Implied Equity Premium for Alternative Preference Parameter Values**

This figure shows the one-month CPT-implied expected equity return for alternative values of the scale, b, loss aversion, λ, and probability distortion, c, parameters, in panels A to C, respectively. The expected equity return is calculated as explained in section 2.2. For the benchmark setup (the bold line) we set $c = 0.65$, $b = 0.65$, and $\lambda = 2.25$, respectively, unless otherwise indicated in the panel. We additionally assume $\gamma = 1$.

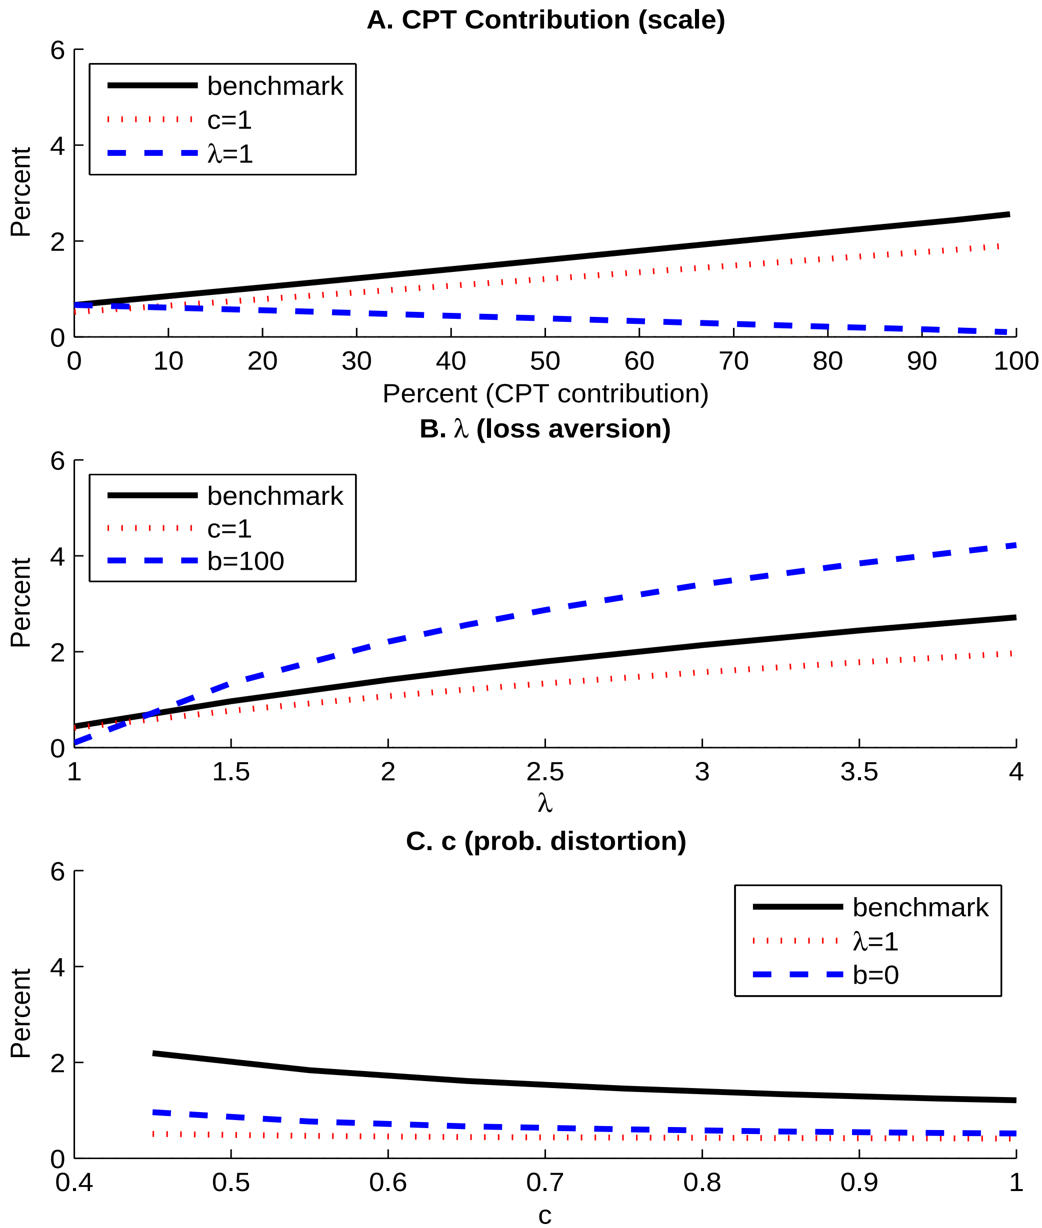

Figure 3: **CPT-Implied Variance Premium for Alternative Volatility Levels**

This figure shows the annualized CPT-implied variance premium (see figure 1) for three different values of monthly equity volatility: 2% (low), 5.18% (benchmark), and 10% (high). For the benchmark setup (the bold line) we set $c = 0.65$, $b = 0.65$, and $\lambda = 2.25$, respectively, unless otherwise indicated in the panel. We additionally assume $\gamma = 1$.

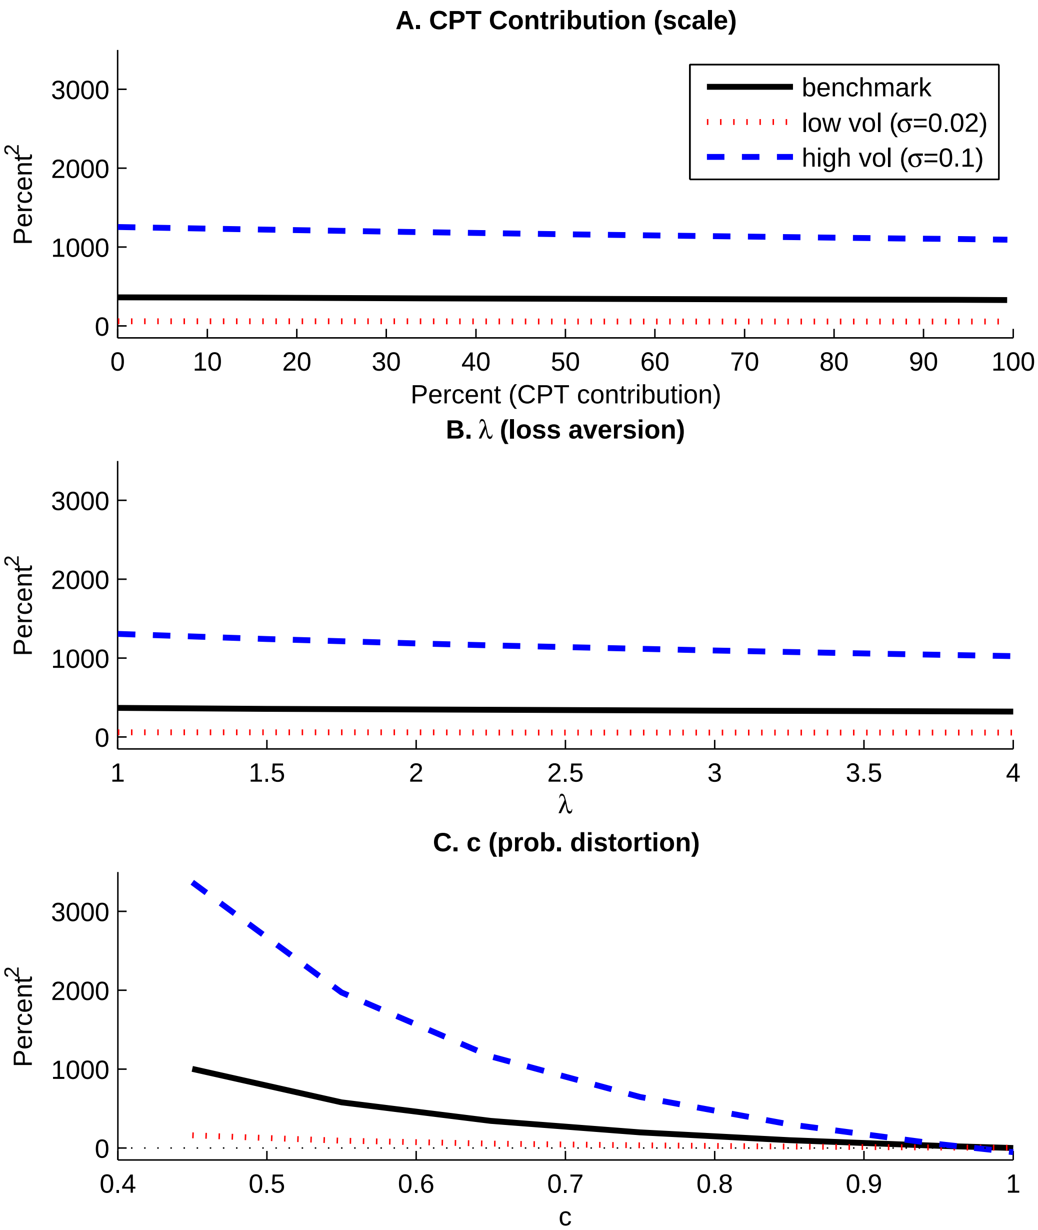

Figure 4: **CPT-Implied Equity Premium for Alternative Volatility Levels**

This figure shows the one-month implied expected equity return (see figure 2) for three different values of monthly equity volatility: 2% (low), 5.18% (benchmark), and 10% (high). For the benchmark setup (the bold line) we set $c = 0.65$, $b = 0.65$, and $\lambda = 2.25$, respectively, unless otherwise indicated in the panel. We additionally assume $\gamma = 1$.

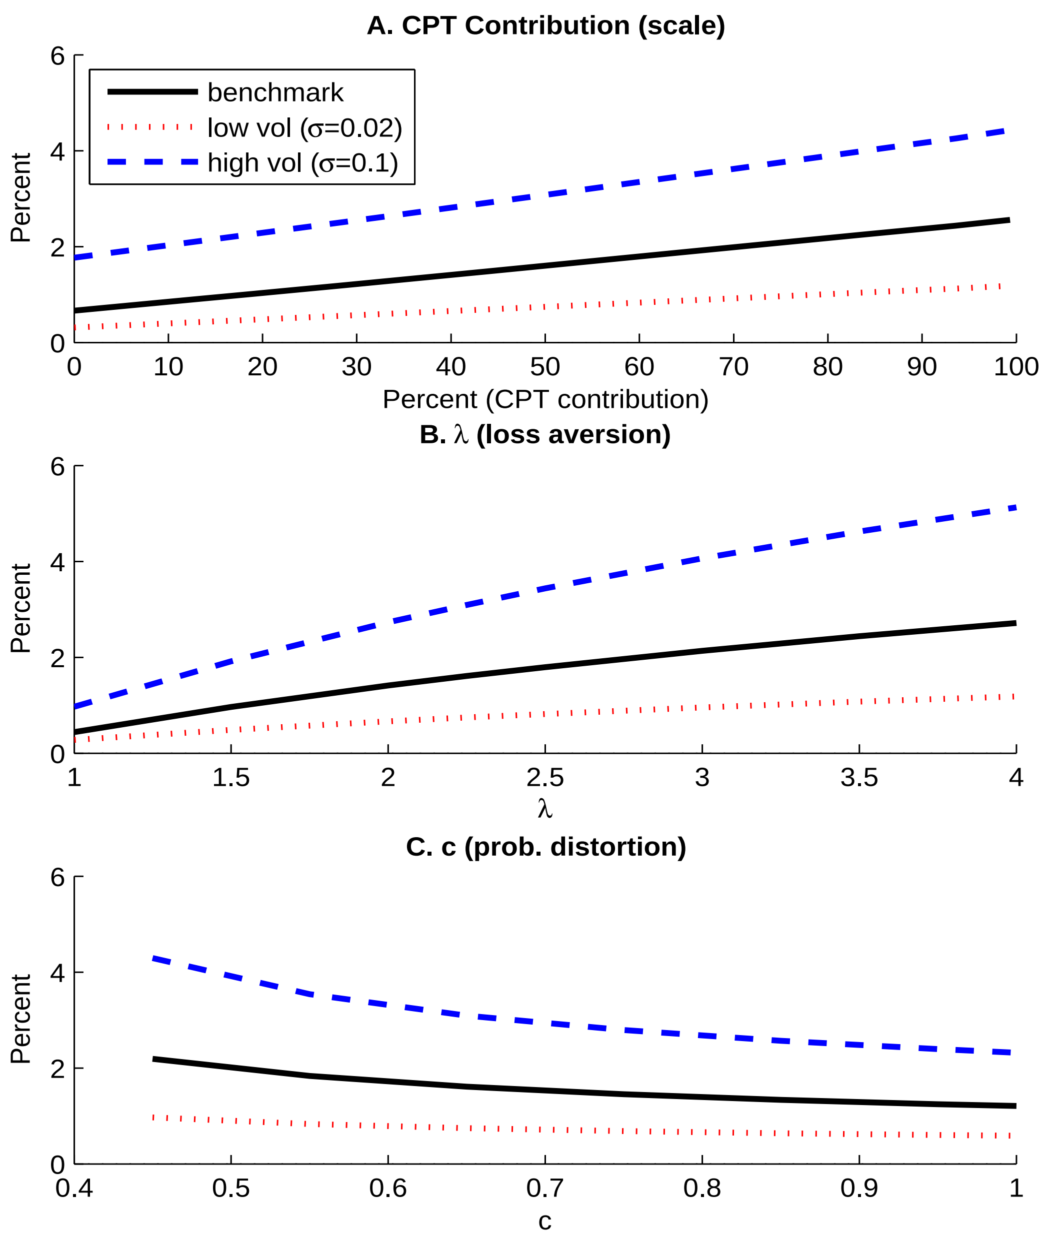

Figure 5: **Average Option Returns and Model Fit**

This figure compares the average observed 30-day return for all options in our sample with the ones fitted by our benchmark CPT representative agent model and by a CRRA model—a restricted CPT model with $c = 1$ and $b = 0$. Our sample includes 26 option returns grouped across strikes (or deltas). Of these 26 options, 13 are puts (panel A), with delta from -0.8 to -0.2 (from in-the-money to out-of-the-money). The other 13 are calls (panel B), with delta from 0.2 to 0.8 (from out-of-the-money to in-the-money).

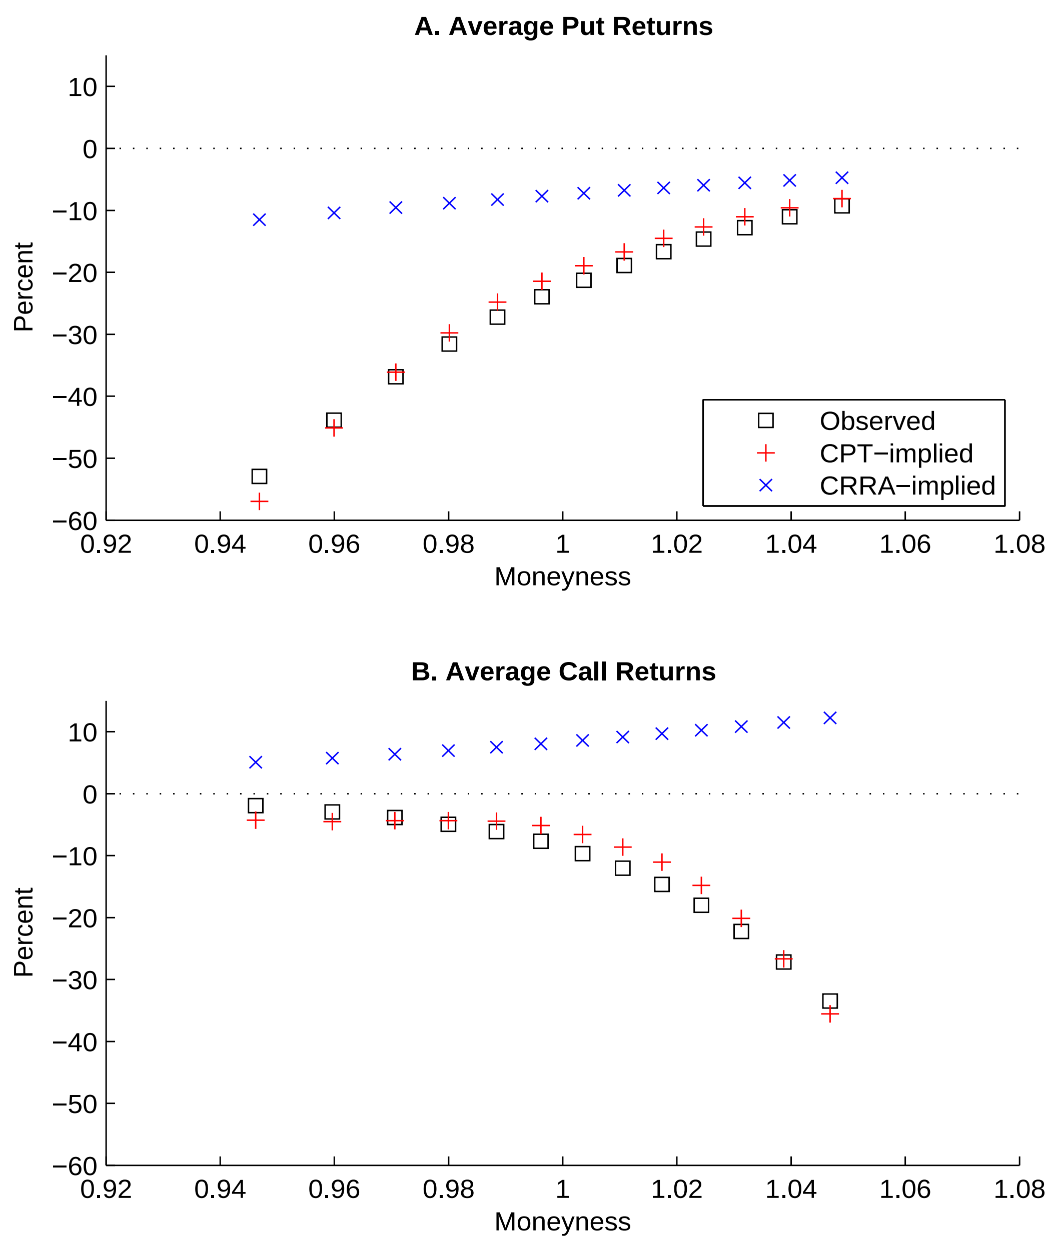

Figure 6: **Observed vs. CPT-Implied Equity and Variance Premiums**

This figure compares the observed equity and variance risk premiums with those implied by a CPT setting where the model's CPT parameters are constant and the stock return volatility is time varying (specifications #1 and #2 in table 4). The variance premium, in panel A, is calculated as the difference between the square of the VIX and the expected realized variance (see section 3.2). The expected realized variance is calculated as an in-sample forecast of the stock return variance using the lagged stock return variance and the square of the VIX. The CPT-implied equity and variance risk premiums are calculated as explained in section 2.3. In panel B, the observed equity premium is calculated as an in-sample forecast of equity returns from the price-dividend ratio, while, in panel C, we add the variance risk premium as a predictor of equity returns.