Option-implied risk attitude under rank-dependent utility*

Maik Dierkes[§]

First version: November 2008
This draft: December 2009

Abstract

Some psychologists and neuroscientists suggest that humans are badly calibrated when dealing with small probabilities and/or tail events. Rank-dependent utility, one of the most important extensions of expected utility, addresses this issue with an inverse-S shaped probability weighting function. In this paper, we employ rank-dependent utility to infer a representative investor's risk preferences from option prices and the time series of the underlying. This preference calculus can explain a pricing kernel's non-monotonicity, violations of second-order stochastic dominance in option markets, and differences in risk-neutral distributions depending on whether they are estimated from option cross sections or from the underlying's time series via Girsanov's theorem. In line with the findings of psychologists and behavioral economists, we fail to reject an inverse-S shaped probability weighting function.

Keywords: Rank-dependent utility, probability weighting, pricing kernel puzzle, risk aversion smile, option pricing.

JEL: G12, G13

EFM Classification Codes: 310, 320, 410, 450, 720

*I am indebted to Nicole Branger, Jens Jackwerth, Thomas Langer, seminar participants at the Finance Center Münster, participants at the German Finance Association meeting 2009, and the First International Summer School in Behavioral Economics and Retirement Savings 2008 held in Münster for valuable comments and insights. All remaining errors are my own. I acknowledge financial support from FAG Münster.

[§]Finance Center Münster, University of Münster, Universitätsstraße 14-16, 48143 Münster, Germany. Email: Maik.Dierkes@uni-muenster.de

There is a growing strand of literature that employs option prices to infer the economy's aggregated risk attitude. Among others, Rosenberg and Engle (2002) and Bliss and Panigirtzoglou (2004) emphazise the advantages of option prices over consumption data. Bliss and Panigirtzoglou (2004) infer the representative investor's risk attitude under the assumption of constant relative and constant absolute risk aversion and find reasonable parameters. However, once this parametric restriction is weakened, some puzzling facts about option implied preferences are reported. For example, Jackwerth (2000) nonparametrically estimates the risk-neutral distribution from option prices and the data-generating distribution from the historical time series of the underlying. This enables him to calculate the Arrow-Pratt measure of absolute risk aversion as a function of the economy's wealth. He finds the puzzling result of risk aversion smiles over the wealth domain: risk aversion is first decreasing and then increasing in the economy's wealth. Even more strikingly, the representative agent is risk prone for moderate wealth levels. Aït-Sahalia and Lo (2000) confirm the U-shaped risk aversion function over the economy's wealth. However, they do not find a risk-loving agent.[1] Rosenberg and Engle (2002) fit an asymmetric GARCH process to the underlying's time series and Chebyshev polynomials to the pricing kernel. Their results are in line with Jackwerth (2000). Ziegler (2007) tries to explain these U-shaped risk aversion functions with several approaches including heterogeneous risk aversion among agents, heterogeneous beliefs, and misestimated beliefs caused by stochastic volatility and jumps. However, none of these approaches is able to explain risk aversion smiles. Bakshi and Madan (2008) and Bakshi, Madan, and Panayotov (2009) also consider heterogenous beliefs to explain U-shaped pricing kernels. Furthermore, Bakshi, Madan, and Panayotov (2009) provide ample evidence of U-shaped pricing kernels.

Researchers then began to question the assumptions imposed so far, like the S&P 500 being a good proxy for the economy's wealth or the representative agent deriving state-

[1] Jackwerth (2000, p. 443) argues that this may be due to the use of overlapping returns in Aït-Sahalia and Lo (2000).

independent utility. For example, Brown and Jackwerth (2004) suggest that utility depends on the volatility level, but conclude that this can not be the only omitted variable. Chabi-Yo, Garcia, and Renault (2008), among others, assume a state-dependent pricing kernel. They explain the observed non-monotonicity of the pricing kernel with non-standard preferences like external habit formation with state dependent beliefs.

In this paper, we assume a representative agent economy in which the agent derives state independent utility over terminal wealth. Furthermore, we assume that the S&P 500 proxies the economy's wealth. However, we relax the assumption of expected utility theory (EUT) and thereby a rational agent.[2] From a behavioral point of view, option markets are particularly susceptible to irrationalities. To price out-of-the-money options, traders must deal with rather small probabilities and extreme outcomes, a situation in which humans are particularly miscalibrated, as psychologists and behavioral economists have found. More specifically, extreme outcomes and small probabilities are typically overweighted (see, for example, Kahneman and Tversky (1979) and Tversky and Kahneman (1992)). Descriptive theories of decision making accommodate such behavior by employing rank dependence via probability weighting. Interestingly, recent research from neuroeconomists shows that probability weighting has its origin in the human brain (Hsu, Krajbich, Zhao, and Camerer (2009)). fMRI scans suggest that expected reward from gambles is non-linear in probabilities. Thus, it is no surprise that even professional option trader exhibit this bias (see Fox, Rogers, and Tversky (1996) and the description of their study below). We wonder, however, whether probability weighting survives a competitive market environment like the S&P 500 options market.

This paper is the first attempt to estimate a represenitive investor's option-implied risk attitudes under rank-dependent utility (RDU), one of the most prominent generalizations of EUT. We assume a one-parameter extension of EUT in most parts of this paper to provide a

[2]We take the stand that a violation of EUT's independence axiom is irrational.

parsimonious model, but also rely on nonparametric estimates for the probability weighting function. Our results indicate that (1) agents significantly deviate from EUT and that (2) this deviation is in the same direction as suggested by the psychological literature, i.e. option traders exhibit probability weighting such that extreme outcomes are overweighted. This implies an inverse-S shaped probability weighting function. Further, this inverse-S shape can explain the non-monotonicity of the option-implied pricing kernel.

Chew, Karni, and Safra (1987) and Ryan (2006) provide a thorough analysis of risk attitudes under RDU. Under mild restrictions on the probability weighting function and the assumption of a concave and increasing utility function, second-order stochastic dominance is violated if and only if the probability weighting function has both convex and concave parts. Thus, the typically observed inverse-S shape of the probability weighting function implies a violation of second-order stochastic dominance. Accounting for transaction costs, Constantinides, Jackwerth, and Perrakis (2009) observe exactly this violation of second-order stochastic dominance in option prices on the S&P 500, one of the world's most liquid option markets. However, evidence of direct arbitrage opportunities is negligible as Bakshi and Panayotov (2008) find. Whereas previous research was primarily concerned about the lower tail of risk-neutral distributions, Constantinides, Jackwerth, and Perrakis (2009) moreover document problems with out-of-the-money calls indicating problems in the upper tail. An inverse-S shaped weighting function addresses these issues by overweighting both tails of the distribution.

We emphasize that the concept of probability weighting is different from having beliefs distinct from the objective distribution. Subjects who conduct probability weighting report the true probabilities when asked, but internally they process the distorted distribution. Hsu, Krajbich, Zhao, and Camerer (2009) support this view as it is the human brain that processes probabilities non-linearly in the pattern predicted by an inverse-S shaped probability weighting function. Note that this behavior may explain the differences in the risk-neutral distribution based on options cross sections and based on a time series estimate (by employing Girsanov's theorem) as found by Aït-Sahalia, Wang, and Yared (2001)[3]: whereas investors may agree on the data-generating distribution, which can be inferred from past returns, they form a distorted distribution when they process the historical distribution and price options. These different distributions deliver different risk-neutral measures when discounted with the stochastic discount factor (or pricing kernel).

Two consequences are noteworthy here. First, the observationally equivalent hypothesis of correctly processing the true data-generating distribution with an abnormal pricing kernel is rejected by Aït-Sahalia, Wang, and Yared (2001). Second, our hypothesis of an inverse-S shaped probability weighting function seems to be empirically indistinguishable from a corresponding belief model. A belief-based model, however, has to argue why rational investors have persistently biased beliefs whereas probability weighting (by definition) does not have to do so. For example, Ziegler's (2007) most promising approach considers three investor types, one type with optimistic and two types with either pessimistic or extremely pessimistic beliefs. In the long run, however, it is hard to justify the fact that all types of investors do not learn and instead stick to their biased beliefs. Note that an inverse-S shaped probability weighting function mimics optimistic and pessimistic beliefs simultaneously by overweighting both tails of the underlying's distribution.

The point we want to make is not the representative investor's imprecise knowledge about the underlying's distribution although ambiguity may amplify the magnitude of probability weighting. Liu, Pan, and Wang (2005) consider uncertainty aversion toward rare events and thereby account for this imprecise knowledge. Gollier (2006) shows that smooth ambiguity aversion (see Klibanoff, Marinacci, and Mukerji (2005)) can explain the pricing kernel's non-monotonicity. Again, we assume for simplicity that the representative investor knows the underlying's distribution perfectly. However, unconsciously he does not

[3] Although accounting for downward jumps with an unreasonably high probability helps to approximate both risk neutral densities, differences remain.

process this distribution, but a distorted one. And this distortion is predictable and thus testable. Following a growing strand of research, we stress the limits of arbitrage (Shleifer and Vishny (1998), Santa-Clara and Saretto (forthcoming)) and propose a behavioral explanation for observed option prices (see, for example, Stein (1989), Poteshman (2001), Poteshman and Serbin (2003), Coval and Shumway (2005), Haigh and List (2005), and Han (2008)).

In our analysis, we will focus exclusively on probability weighting and will not consider further behavioral concepts like reference dependence, loss aversion, and a reflection effect as, for example, incorporated in Cumulative Prospect Theory (CPT, Tversky and Kahneman (1992)). We concentrate on RDU theory and ignore additional behavioral features because, in liquid markets like the S&P 500 option market, the rational paradigm is the null hypothesis. Therefore, we would like to incorporate as little deviation from EUT as possible.[4] In a CPT framework, Barberis and Huang (2008) focus on probability weighting and explicitly state (but do not test) the hypothesis that right-skewed payoff profiles like out-of-the-money options are too expensive due to probability weighting. In a portfolio choice context, Driessen and Maenhout (2007) find that only RDU with an inverse-S shaped probability weighting funtion generates positive portfolio weights on an at-the-money straddle and on an out-of-the-money put whereas CPT and EUT fail to do so.

Given the liquidity and high stakes in the competitive S&P 500 option market, it might be questionable that "small stake theories" (those motivated by lab experiments) like RDU or CPT predict market prices. However, Coval and Shumway (2005) and Haigh and List (2005) successfully apply CPT concepts like loss aversion and the reflection effect to the S&P 500 option market. More related to our hypothesis, Fox, Rogers, and Tversky

[4]The results by Gurevich, Kliger, and Levy (2009), who estimate CPT parameters from individual equity options in a discrete state space framework, support our view. For example, their loss aversion coefficient only slightly deviates from one.

(1996) ask professional option traders on the floors of the Pacific Stock Exchange and the Chicago Board Options Exchange to participate in an experiment on option pricing with an adequate incentive scheme. These traders exhibit non-linear probability weights when asked for derivatives on Microsoft, IBM, or General Electric shares. Furthermore, Etchart-Vincent (2004) examines the influence of low and high stakes on probability weighting. She finds similar inverse-S shaped probability weighting functions in both cases. Snowberg and Wolfers (2007) analyze a proprietary data set on sports betting markets which enables them to discriminate EUT-based risk proclivity from probability weighting as an explanation for the favourite longshot bias. In this presumably less competitive market environment, their results favor the concept of probability weighting, although Sonnemann, Camerer, Fox, and Langer (2009) show that other behavioral biases like partition dependence influence horse race betting markets owing to the specific payoff structure in these markets. The payoff structure of plain vanilla options makes it unlikely that our results are confounded by the partition dependence bias.

The remainder of the paper is organized as follows. Section 1 gives a more detailed introduction to RDU. Section 2 derives the implication of RDU for option prices and risk aversion functions. Section 3 describes the used data. The results of three different estimation approaches, each with different assumptions, are presented in Section 4. The first approach employs Pan's (2002) model accounting for state-dependent distributions via stochastic volatility and jumps whose intensity is correlated with volatility. The second one follows Bliss and Panigirtzoglou (2004) and assumes stationary preferences instead of stationary distributions. Risk preferences are found by maximizing the forecast ability of the risk-adjusted state price densities amended with risk preferences. The third estimation approach uses the estimation techniques of Aït-Sahalia, Wang, and Yared (2001) and recovers a nonparametric probability weighting function without parametric assumptions on the investor's utility function. Section 5 concludes.

1 RDU

RDU with its axiomatic foundation was introduced by Quiggin (1981) for decision making under risk (the probabilities are known) and by Schmeidler (1989) for decision making under uncertainty (the probabilities are unknown). However, to make our analysis feasible we (have to) focus on decision making under risk and assume that the representative investor has perfect knowledge of the relevant odds.

The intuition behind RDU is that risk aversion is not only displayed in diminishing marginal utility, but also in a probabilistic risk attitude (see Wakker (1994)). For example, most people prefer gambles like winning \$1000 with a probability of 0.001 or winning nothing to a fixed payment of \$1. The famous Allais paradox exploits this type of behavior. By transforming the cumulative distribution function of the risk at hand, this behavior can be modeled. Instead of the true cumulative distribution function F, the representative investor uses the following distorted one:

$$\tilde{F}(x) = w(F(x)) = w \circ F(x), \tag{1}$$

where $w : [0, 1] \to [0, 1]$ is a strictly increasing mapping with the only additional restrictions that $w(0) = 0$ and $w(1) = 1$. Utility for the risk given by the distribution F is then calculated with the usual additive functional $V_{u,w}(F) = \int u(x)d(w \circ F)(x)$.

A lot of effort has been spent on eliciting individual probability weighting functions w in lab experiments. The vast majority of the literature in this field favors an inverse-S shape for w which implies overweighting of extreme outcomes. This pattern was found by parametric estimation (see e.g. Tversky and Kahneman (1992), Camerer and Ho (1994), Tversky and Fox (1995), Bleichrodt, van Rijn, and Johannesson (1999)) as well as nonparametric estimation (see e.g. Wu and Gonzalez (1996), Gonzalez and Wu (1999), Abdellaoui (2000), Bleichrodt and Pinto (2000)). Several parametric forms are suggested in the literature. In a CPT framework, Stott (2006) conducts an extensive study to analyze the predictive power of these functional forms. He favors the one introduced by Prelec (1998). We follow his advice (although we do not employ CPT) and focus on $w(p) = e^{-(-ln(p))^{\delta}}$. Another argument for this parametric form is its strict monotonicity in contrast to the probability weighting function proposed by Tversky and Kahneman (1992) (see e.g. Rieger and Wang (2006)). Thus, we use a one-parameter extension of EUT. The typical inverse-S shaped form is revealed for $\delta < 1$. In addition to this functional form, we also estimate the probability weighting function in a nonparametric manner as described below. Figure 1 depicts the density functions of a standard normal distribution and the corresponding distorted distribution based on $w(p) = e^{-(-ln(p))^{\delta}}$ with $\delta = 0.6$ and $\delta = 0.8$. Evidently, this probability weighting function overemphasizes both tails, but especially the lower tail of the distribution. This leads to higher (perceived/evaluated) kurtosis and left-skewness. This is exactly in line with Aït-Sahalia, Wang, and Yared's (2001) results. They find the option-implied risk-neutral distribution to be more left-skewed and to have higher kurtosis compared with the time-series-implied risk-neutral distribution. Furthermore, this particular probability weighting scheme with a utility function exhibiting constant relative risk aversion (CRRA) can explain the extreme negative skewness of the risk-neutral distribution as it introduces additional negative skewness and excess kurtosis in the (perceived/evaluted) data-generating distribution (see Bakshi, Kapadia, and Madan (2003)). Claims that we should have seen more large stock movements in the past to rationalize option prices (see Bates (2000)) may thus find an answer in RDU preferences.

2 Implications of RDU for option pricing

We now turn to the implications of probability weighting for option pricing. We assume a representative agent with monotonically increasing utility function u and monotonically

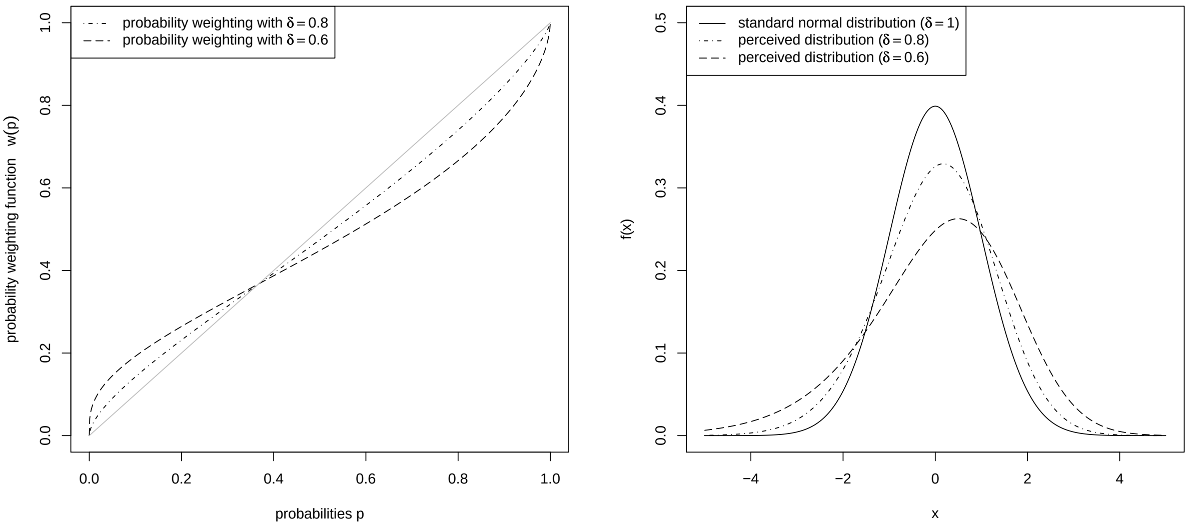

Figure 1: Density function of the standard normal distribution and its perception with probability weighting.
The left panel shows the probability weighting function $w(p) = e^{-(-ln(p))^{\delta}}$ with two different parameters $\delta = 0.6$ and $\delta = 0.8$ (gray line corresponds to the identity function). The right panel shows in addition to the density function of the standard normal distribution two density functions of the distorted normal distribution. Distortion is conducted with probability weighting according to the weighting function in the left panel.

increasing probability weighting function w. Chapman and Polkovnichenko (2009) discuss the conditions for the existence of a representative agent under RDU preferences. Both u and w are assumed to be smooth (twice continuously differentiable). All wealth in the economy is proxied by an index with S_t and S_T denoting wealth today and in the future, respectively. Let f_P and f_Q denote the density functions of the data-generating process and the risk-neutral measure with corresponding cumulative distribution functions F_P and F_Q, respectively. Under EUT (i.e. linear w), economic theory says that

$$f_Q(S_T) = f_P(S_T) \cdot \beta \frac{u'(S_T)}{u'(S_t)} \tag{2}$$

where β is a normalizing constant. If we estimate f_Q from a cross section of option prices and f_P from the historical index prices, then we can infer the representative investor's risk preferences measured by the Arrow-Pratt coefficient of absolute risk aversion as

$$-\frac{u''(S_T)}{u'(S_T)} = \frac{f'_P(S_T)}{f_P(S_T)} - \frac{f'_Q(S_T)}{f_Q(S_T)}. \tag{3}$$

Reasonable estimates for the quantity in equation (3) should be non-increasing in S_T and non-negative, thereby indicating a risk-averse investor whose risk aversion does not increase with wealth. However, several studies document smiling risk aversion coefficients with significant parts being negative (Aït-Sahalia and Lo (2000), Jackwerth (2000), Rosenberg and Engle (2002)). Clearly, under non-expected utility preferences, the Arrow-Pratt measure is meaningless and hence may smile. Under RDU (non-linear w), the investor prices options as if the data-generating process's distribution function is $F_{\tilde{P}}(x) = w(F_P(x))$. Thus, the corresponding density is given by

$$f_{\tilde{P}}(S_T) = w'(F_P(S_T)) \cdot f_P(S_T)), \tag{4}$$

and our hypothesis is that the option-implied risk-neutral measure is generated under RDU preferences, i.e. according to

$$f_Q(S_T) = w'\big(F_P(S_T)\big) \cdot f_P(S_T) \cdot \beta \frac{u'(S_T)}{u'(S_t)}. \tag{5}$$

The same steps that led from equation (2) to equation (3) then yield

$$\frac{f_P'(S_T)}{f_P(S_T)} - \frac{f_Q'(S_T)}{f_Q(S_T)} = \Big(-\frac{w''(F_P(S_T))}{w'(F_P(S_T))} f_P(S_T)\Big) + ARA_u(S_T), \tag{6}$$

where $ARA_u(S_T)$ denotes the absolute risk aversion function across index levels S_T associated only with the agent's utility function u. The term $-\frac{w''(F_P(S_T))}{w'(F_P(S_T))} f_P(S_T)$ on the right-hand side of (6) displays the probabilistic risk attitude. The denominator is always positive due to the strictly increasing weighting function. Hence, observed risk aversion as a function of the economy's wealth is decomposed under RDU into a probabilistic counterpart to the Arrow-Pratt measure of absolute risk aversion and the Arrow-Pratt measure itself. Under the proclaimed inverse-S shaped form of the probability weighting function, this term is positive for low wealth levels and steadily decreases until it is negative for high wealth levels. The intuition here is that concave parts of the probability weighting function increase the observed risk aversion and convex parts decrease it.

3 Data

In sections 4.3 and 4.4, we use bid and ask prices and the last trade price of European put and call options on the S&P 500 traded on the Chicago Board Options Exchange (CBOE). The data include option expiries from January 1997 to December 2007. To estimate risk-neutral distributions, we calculate closing prices as follows. If the last trade price is between the last bid and the last ask price, then the last trade price is taken as the closing price. If the last trade price is below the last bid price, then the last bid price is taken. Similarly, the last ask price is taken as the closing price if the last trade price exeeds the last ask price. If no trade occurred on a specific day, we discard the option from consideration.

We compute the implied volatility based on the Black and Scholes option pricing fomula. For an option to be included in our estimation procedure, its implied volatility has to be below 100% and its price has to exceed \$0.20. Thus, too deep out-of-the-money options whose prices are unreliable are excluded. Furthermore, we exclude options that violate simple arbitrage bounds like $max(0, S_t - Ke^{-r_{t,\tau}\tau}) \leq C_t \leq S_t$, where K, C_t, S_t and $r_{t,\tau}$ denote strike price, call price, underlying's price and riskless interest rate at time t with time to maturity τ, respectively. We discard option cross sections with less than five options.

To circumvent the unobservability of expected dividend yields $D_{t,\tau}$, we employ the usually most liquidly traded at-the-money put and call options with equal strike price K to infer the futures price $F_{t,\tau} = S_t e^{(r_{t,\tau} - D_{t,\tau})\tau}$ via the put-call parity

$$C_t(S_t, K, \tau, r_{t,\tau}, D_{t,\tau}) + Ke^{-r_{t,\tau}\tau} = P_t(S_t, K, \tau, r_{t,\tau}, D_{t,\tau}) + F_{t,\tau}e^{-r_{t,\tau}\tau}. \qquad (7)$$

With the same arguments given by Bliss and Panigirtzoglou (2004), we use the annualized three-months' treasury bills (secondary market) taken from the Federal Reserve Board's website to proxy the riskless interest rate.

Section 4.4 moreover uses historical prices of the S&P 500 that is sampled every 5 minutes. We extrat this data from tick data for the period January 1997 to December 2007.

4 Results

4.1 Overview

To test our theory, we base our results on three different approaches. First, we employ Pan's (2002) model. She fits a rich option-pricing model inspired by Bates (2000) that accounts for stochastic volatility and jumps whose jump intensity depends on the current volatility level. Her results are particularly useful since she simultaneously uses option cross sections and time series returns for estimation. Second, we employ the data described in section 3 and follow the ideas in Bliss and Panigirtzoglou (2004). They nonparametrically infer state price densities from cross sections of option prices. Amending the state price density with a stochastic discount factor and probability weighting enables us to predict future index levels. Third, we use the methodology in Aït-Sahalia, Wang, and Yared (2001) and estimate option-implied and time-series-implied state price densities. These estimates enable us to infer a nonparametric probability weighting function without parameric assumptions on the representative investor's preferences.

4.2 Results for Pan's (2002) model

For option prices and index prices, Pan (2002) proposes a model with stochastic volatility and jumps. The jump intensity is correlated with the current volatility level. She estimates risk premia simultaneously with cross sections of option prices and index returns. She proposes the following process for the underlying S_t:

$$dS_t = [r_t - q_t + \eta^S V_t + \lambda V_t(\mu - \mu^*)]S_t dt + \sqrt{V_t} S_t dB_t^1 + dZ_t - \mu S_t \lambda V_t dt, \tag{8}$$

$$dV_t = \kappa_v(\bar{v} - V_t)dt + \sigma_v \sqrt{V_t}\big(\rho dB_t^1 + \sqrt{1-\rho^2} dB_t^2\big), \tag{9}$$

where V_t denotes the instantaneous variance that follows a square-root process with long-run mean $\bar{v}$, mean reversion rate κ_v, and volatility σ_v. Random innovations are introduced by two independent Brownian motions B_t^1, B_t^2, and a poisson process Z_t, whose jump intensity is λV_t and thus perfectly correlated with volatility V_t. The logarithm of the relative jump size conditional on a jump occurring is normally distributed with mean $\mu_J = \ln(1+\mu) - \sigma_J^2/2$ and variance σ_J^2. The riskless rate r_t and the dividend yield q_t both follow a square-root process with long-run means $\bar{r}$ and $\bar{q}$, mean reversion rates κ_r and κ_q, and volatility coefficients σ_r and σ_q, respectively.

The risk-neutral dynamics evolve according to

$$dS_t = (r_t - q_t)S_t dt + \sqrt{V_t} S_t dB_t^1(Q) + dZ_t^Q - \mu^* S_t \lambda V_t dt, \tag{10}$$

$$dV_t = [\kappa_v(\bar{v} - V_t) + \eta^v V_t]dt + \sigma_v \sqrt{V_t}\big(\rho dB_t^1(Q) + \sqrt{1-\rho^2} dB_t^2(Q)\big), \tag{11}$$

where V_t plays the same role as above, but with mean reversion rate $\kappa_v^* = \kappa_v - \eta_v$, long-run mean $\bar{v}^* = \kappa_v \bar{v}/\kappa_v^*$ and volatility coefficient σ_v. $B_t^1(Q)$, $B_t^2(Q)$ and Z_t^Q are two independent Brownian motions and the poisson process under the risk-neutral measure Q, respectively. Again, the jump intensity is λV_t. The logarithm of the jump size conditional on a jump occurring is normally distributed with mean $\mu_J^* = \ln(1+\mu^*) - \sigma_J^2/2$ and variance σ_J^2. Pan (2002) estimates all parameters in this model including the three risk premia $\eta^S V_t$, $\eta^v V_t$, and $\lambda V_t(\mu - \mu^*)$ for the diffusion, volatility, and jump risk, respectively (see Pan (2002), Table 3 and Table 6). We extract the state price densities and data-generating densities via transform inversion for different times to maturity (details can be found in Ziegler (2007)).

Given the data-generating density f_P, we can try to match the state price density f_Q by imposing different pricing kernels and probability weighting schemes. Matching is done by minimizing the variation distance Δ between the two densities[5]

$$\Delta(f_{\hat{Q}}, f_Q) = \frac{1}{2}\int_0^\infty |f_{\hat{Q}}(S_T) - f_Q(S_T)| dS_T, \tag{12}$$

where $f_{\hat{Q}}(S_T) = f_P(S_T) \cdot w'\big(F_P(S_T)\big) \cdot \beta \frac{u'(S_T)}{u'(S_t)}$ is the adjusted data-generating density. The considered class of utility functions u is restricted to power functions $u(S_T) = \frac{S_T^{1-\gamma}}{1-\gamma}$. However, other usually employed utility function like, e.g. the exponential function, do not alter our results significantly. As noted by Ziegler (2007), the implicit risk attitudes in the estimated pricing kernel are not in line with any standard assumption on a rational, risk-averse representative investor, since the implicit absolute risk aversion inferred from equation (3) is positive for low wealth levels and becomes more and more negative for higher wealth levels. Therefore, we prefer a standard utility function with a probability weighting function as an additional degree of freedom. The employed utility function exhibits CRRA. The probability weighting function is restricted to having the one-parameter Prelec (1998) form.

Note that our model provides only two degrees of freedom, one in the utility function and one in the probability weighting function. Pan's (2002) model, however, allows for three different risks to be priced. Thus, our preference structure is unlikely to be consistent with Pan's (2002) model in general. Given her estimates for the risk premia, however, our RDU model does a remarkably good job in explaining the time series and risk neutral distributions simultaneously (see below). This result militates in favor of the reasonable preference structure we impose.

Figure 2 depicts the results for time to maturity of $\tau = 1$ year. The left panel indicates a remarkably close fit of the densities when amending the data-generating density with a

[5] We use a discretized version for the numerical evaluation of the integral and integral limits were 0.1 and 2, since density functions calculated via transform analysis are vitually 0 for values S_T outside the interval $[0.1, 2]$.

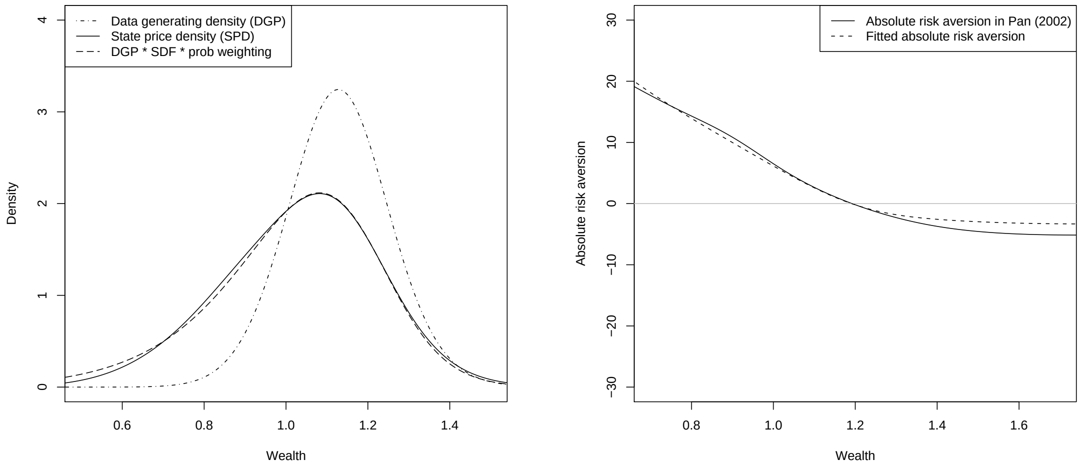

Figure 2: Model implied densities, fitted density, and corresponding absolute risk aversion smiles in Pan's (2002) model.
The left panel shows the data-generating density, state price density, and fitted density for time to maturity $\tau = 1$. Data-generating and state price density are calculated with Fourier inversion with model parameters estimated by Pan (2002) (see Table 3 and Table 6 in Pan (2002)). By amending the data-generating density with a CRRA pricing kernel (i.e. $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$) and a probability weigting of the Prelec (1998) form (i.e. $w(p) = e^{-(-ln(p))^{\delta}}$) such that the variation distance to the state price density is minimized we yield the fitted density. The fittes parameter values are $\gamma = 3.3691$ and $\delta = 0.7390$. Under RDU, $\delta < 1$ and $\delta > 1$ indicate an inverse-S shaped and S shaped probability weighting function, respectively. The right panel shows the corresponding absulute risk aversion functions over wealth.

CRRA pricing kernel and probability weighting of the Prelec (1998) form. Fitted parameter values are $\gamma = 3.3691$ and $\delta = 0.7390$. Thus, we have a considerable overweighting of tail events as hypothesized by an inverse-S shaped probability weighting function. Indeed, $\delta = 0.7390$ is only slightly higher than observed in lab experiments where median individual probability weighting parameters are typically aroung 0.69.

Ziegler (2007) calculates risk aversion functions implied by Pan's (2002) model estimates for times to maturity from one year to one month. They exhibit a simlar shape as in Figure 2, i.e. the implicit absolute risk aversion is positive for low wealth levels and becomes more and more negative for higher wealth levels (see footnote 11 on page 876 in Ziegler (2007)). From equation (6) it is clear that this pattern is exactly in line with an inverse-S shaped probability weighting function.

4.3 Predicting index levels with risk-adjusted state price densities

Assuming CRRA or CARA utility functions, Bliss and Panigirtzoglou (2004) provide an excellent framework to infer the option-implied risk preferences under EUT. For a given forecast horizon, they nonparametrically estimate the risk-neutral density from a cross section of options on the S&P 500 future with time to maturity equal to the forecast horizon. According to their presumed utility function, they risk adjust the risk-neutral density f_Q by simple manipulation and renormalization of equation (2). Equipped with this theoretical data-generating density f_P, they forecast the future price at the date of expiry. Moreover, by optimizing over the risk preference parameter, they maximize the forecast ability, thereby estimating the representative investor's risk aversion in terms of constant relative or constant absolute risk aversion. We adopt their ideas and additionally allow for probability weighting. Before turning to the results, we describe the procedure of estimating the risk-neutral density and optimizing the forecast ability of the risk-adjusted risk-neutral density in more detail.

For each forecast horizon τ ($\tau = 2$ weeks to 6 weeks), we follow Aït-Sahalia and Lo (1998) and Aït-Sahalia, Wang, and Yared (2001) and estimate the state price density implicit in option prices as follows. We assume a semiparametric model, namely the Black-Scholes model with an implied volatility function $\sigma(K/F_{t,\tau}, \tau)$ which depends on the ratio of the strike price K and the futures price $F_{t,\tau}$ and on the time to maturity τ. The goal is then to provide a nonparametric estimate $\hat{\sigma}(K/F_{t,\tau}, \tau)$ for this volatility function.

To obtain the risk-neutral density for a fixed time to maturity τ, we determine the futures price $F_{t,\tau}$ as described in Section 3 and predict implied volatilies $\hat{\sigma}(K/F_{t,\tau}, \tau)$ on a fine grid of equidistant strike prices $K = K_1, \ldots, K_n$. With the help of the Breeden and Litzenberger (1978) formula

$$f_Q(S_T) = e^{-r_{t,\tau}\tau} \frac{\partial^2 C_t(S_t, K, \tau, r_{t,\tau}, D_{t,\tau})}{\partial K^2}\Big|_{K=S_T} \tag{13}$$

we numerically derive the risk-neutral density from call prices $C_t(S_t, K, \tau, r_{t,\tau}, D_{t,\tau})$.

We obtain the nonparametric estimate $\hat{\sigma}(K/F_{t,\tau}, \tau)$ as follows. For a given time to maturity τ and expiration date $t + \tau$, we collect all options on the S&P 500 that have the same expiration date $t + \tau$ and whose time to maturity deviates from τ by at most five days. We discard in-the-money options and convert put options into call options via the put-call parity. For all remaining n options we calculate for each option i the Black-Scholes implied volatity σ_i. Now, we apply the Nadaraya-Watson kernel estimator and obtain

$$\hat{\sigma}(K/F_{t,\tau}, \tau) = \frac{\sum_{i=1}^{n} k_{K/F_{t,\tau}}\left(\frac{K/F_{t,\tau} - K_i/F_{t_i,\tau_i}}{h_{K/F_{t,\tau}}}\right) k_\tau\left(\frac{\tau - \tau_i}{h_\tau}\right)\sigma_i}{\sum_{i=1}^{n} k_{K/F_{t,\tau}}\left(\frac{K/F_{t,\tau} - K_i/F_{t_i,\tau_i}}{h_{K/F_{t,\tau}}}\right) k_\tau\left(\frac{\tau - \tau_i}{h_\tau}\right)} \tag{14}$$

We use the kernel functions

$$k_{K/F_{t,\tau}}(x) = \frac{1}{2\pi}\exp(-\frac{x^2}{2}), \tag{15}$$
$$k_{\tau}(x) = \frac{1}{2\pi}\exp(-\frac{x^2}{2}). \tag{16}$$

We choose bandwidths $h_{K/F_{t,\tau}}$ and h_τ such that we obtain a good fit of actual data and reasonably smooth state price density functions. Visual inspection of fitted implied volatilities revealed that we rather underestimate implied volatilities in the tails and thereby produce conservative results in terms of our hypothesis of an inverse-S shaped probability weighting function.

Once we have the risk-neutral distribution f_Q, we can risk adjust it and test the forecast ability of this risk adjusted density. To do so, we pick settlement values from the CBOE's website. For a given forecast horizon τ, we now have a time series of pairs consisting of an estimated forecast distribution $\hat{f}_t$ and a realization, i.e. the settlement value $X_{t+\tau}$. Under the null hypothesis, saying that the settlement values $X_{t+\tau}$ are independent and our estimated density $\hat{f}_t$ reasonably approximates the true distribution f_t, the inverse probability transformation of realizations $X_{t+\tau}$

$$y_t = \int_{-\infty}^{X_{t+\tau}} \hat{f}_t(S_T)dS_T \tag{17}$$

is independently and uniformly distributed. Berkowitz (2001) introduced a method to test uniformity and independence jointly. He backtransforms the values y_t and considers the time series $z_t = \Phi^{-1}(y_t)$, where Φ is the standard normal cumulative distribution function. The reasoning behind this is that the equivalent null hypotheses, namely that z_t is independently and standard normally distributed, is easier to test than the independence and uniformity of y_t. Specifically, Berkowitz (2001) proposes to fit an AR(1) process

$$z_t - \mu = \rho(z_{t-1} - \mu) + \epsilon_t \tag{18}$$

with $\epsilon_t \sim$ i.i.d. $\mathcal{N}(\mu, \sigma^2)$ by maximum likelihood estimation. Let $L(\mu, \sigma^2, \rho)$ denote the log-likelihood function, then the likelihood ratio test statistic $LR_3 = -2(L(0,1,0) - L(\hat{\mu}, \hat{\sigma}, \hat{\rho}))$ has a $\chi^2(3)$ large sample distribution under the null hypothesis of $(\mu, \sigma^2, \rho) = (0,1,0)$. Bliss and Panigirtzoglou (2004) conducted a simulation study to compare this likelihood ratio test with the chi-squared, Kupier, and Kolmogorov-Smirnov tests. They favor the Berkowitz (2001) procedure described above.

To determine risk preferences, we minimize the test statistic LR_3 over risk preferences. Under RDU, we need to account for diminishing marginal utility and probability weighting in a parametric form. We make the parametric assumptions that the utility and probability weighting functions are given by $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ and $w(p) = e^{-(-ln(p))^\delta}$, respectively.[6] Index prices and hence settlement values $X_{t+\tau}$ are distributed according to the data-generating density f_P and the connection to the risk-neutral density f_Q is given by equation (5). Dividing both sides by the pricing kernel and integrating up to $X_{t+\tau}$ yields

$$y_t = \int_0^{X_{t+\tau}} \frac{f_Q(S_T)}{\beta \frac{u'(S_T)}{u'(S_t)}} dS_T = w\big(F_P(X_{t+\tau})\big). \tag{19}$$

Clearly, to account for probability weighting in maximizing the forecast ability, we have to apply the normal distribution transformation not to y_t, but to $w^{-1}(y_t)$. Hence, we fit the

[6] In an unreported robustness check, we also consider power probability weighting functions $w(p) = p^\delta$ with power utility. This results in negative estimates for the relative risk aversion coefficient γ. Therefore, we focus on Prelec's (1998) probability weighting function.

AR(1) process in equation (18) to the time series

$$z_t = \Phi^{-1}\big(w^{-1}(y_t)\big). \tag{20}$$

We minimize the resulting likelihood ratio statistic LR_3 over different preference combinations (γ, δ). Thereby, we optimize the forecast ability, or in other words, we come closer and closer to the data-generating distribution F_P.[7] The best parameter combination (γ, δ) then reflects the RDU risk preferences of the representative investor under our parametric assumptions on the utility and probability weighting function. For comparison reasons, we also infer the constant relative risk aversion coefficient under EUT, i.e. we fix $\delta = 1$ in this case.

Table I shows the estimated risk preference under EUT and RDU. First, the estimated utility function parameters γ are all in a reasonable range of roughly 1 to 4.5. Further, the values for γ are pretty close under EUT and RDU. Looking at the probability weighting parameter δ estimated under RDU, we confirm the inverse-S shaped (i.e. $\delta < 1$) probability weighting function found in the previous section in Pan's (2002) model. The standard errors for δ based on 1,000 bootstrapped samples indicate a significant inverse-S shaped probability weighting function. Given the close-by estimates for γ under EUT and RDU, we interpret the results as evidence that accounting for an inverse-S shaped probability weighting function helps to improve the forecast ability significantly. Indeed, the reported p-values for the difference of both maximized LR_3 statistics support this view.

[7]Our focus here is on the existence of an inverse-S shaped probability weighting function and not on the forecast ability itself. Clearly, allowing for more freedom in risk adjusting risk-neutral distributions by a more complex preference structure has to improve the in-sample forecast ability.

Table I: Estimated risk preferences under EUT and RDU.

γ is the coefficient of constant relative risk aversion for the assumed power utility function. Under RDU, $\delta < 1$ and $\delta > 1$ indicate an inverse-S shaped and S shaped probability weighting function, respectively. For $\delta = 1$, both models coincide. Reported parameter estimates are based on maximizing the forecast ability, i.e. maximizing the LR_3 statistic. Standard errors based on 1,000 bootstrapped samples are reported in parentheses. We calculate the differences of the maximized LR_3 statistics and report p-values under a $\chi^2(1)$ distribution as the RDU model has one additional parameter of freedom. N denotes the number of used cross sections.

forecast horizon	EUT	RDU		Difference in LR_3 statistics
	γ	γ	δ	p-value
2 weeks ($N = 131$)	3.129	2.716	0.805	0.001
	(2.127)	(1.947)	(0.047)	
3 weeks ($N = 132$)	4.295	4.073	0.874	0.041
	(1.920)	(1.808)	(0.060)	
4 weeks ($N = 131$)	2.439	2.386	0.876	0.047
	(1.559)	(1.460)	(0.067)	
5 weeks ($N = 131$)	2.425	2.422	0.860	0.029
	(1.347)	(1.267)	(0.059)	
6 weeks ($N = 131$)	1.326	1.239	0.844	0.012
	(1.186)	(1.142)	(0.056)	

4.4 A nonparametric estimate of the probability weighting function without parametric assumptions on the utility function

Two questions motivate the estimation procedure of this section. First, are our estimates for the probability weighting function confounded by our specific assumptions on the representative investor's utility function? For example, Post and Levy (2005) apply several stochastic dominance criteria and analyze the cross section of stock returns. They cannot discriminate the two hypotheses of either an inverse-S shaped utility function, i.e. one with risk proclivity in high wealth regions, or a utility function exhibiting risk aversion on the entire wealth region paired with an inverse-S shaped probability weighting function. Second, we used specific parametrizations for the probability weighting function employed in our estimation processes so far. How does this confound our results?

In this section we set out to provide a nonparamtric estimate for the probability weighting function without making parametric assumptions on the representative investor's utility function. To do so, we employ the ideas of Aït-Sahalia, Wang, and Yared (2001). In addition to the state price density estimated from option prices in Section 4.3, we estimate the state price density from the time series of the underlying by using the very same semi-parametric model as in Section 4.3, i.e. we assume that the S&P 500 follows the Itô diffusion

$$\frac{dS_t^{(P)}}{S_t^{(P)}} = \mu(S_t^{(P)})dt + \sigma(S_t^{(P)})dW_t^{(P)} \tag{21}$$

with Brownian motion $W^{(P)}$. The state price density then results from an application of Girsanov's Theorem as

$$\frac{dS_t}{S_t} = (r_{t,\tau} - D_{t,\tau})dt + \sigma(S_t)dW_t. \tag{22}$$

The crucial point here is that the volatility function $\sigma(S_t)$ should be the same as the option-implied volatility function $\sigma(K/F_{t,\tau}, \tau)$ from Section 4.3. However, according to our hypothesis, we believe that only the option-implied state price density is distorted by probability weighting simply because humans are badly calibrated when dealing with small probabilities as is necessary when pricing out-of-the-money options. Aït-Sahalia, Wang, and Yared (2001) indeed prove a significant difference between both volatility functions and thus between the option-implied and time-series-implied state price density. They find that the option-implied state price density exhibits more kurtosis and left-skewness. We argue that an inverse-S shaped probability weighting function may be a key driver for their result.

By slightly changing our notation from Section 2, we rewrite Equation (5) as

$$f_{Q_{\text{option-implied}}}(S_T) = w'\big(F_P(S_T)\big) \cdot f_P(S_T) \cdot \beta \frac{u'(S_T)}{u'(S_t)} \tag{23}$$

and define the time-series-implied state price density function which is not affected by probability weighting as

$$f_{Q_{\text{time series}}}(S_T) = f_P(S_T) \cdot \beta \frac{u'(S_T)}{u'(S_t)}. \tag{24}$$

Once we estimated $f_{Q_{\text{option-implied}}}$ and $f_{Q_{\text{time series}}}$, we calculate the ratio

$$\frac{f_{Q_{\text{option-implied}}}(S_T)}{f_{Q_{\text{time series}}}(S_T)} = w'\big(F_P(S_T)\big) \tag{25}$$

and with the knowledge of $w(0) = 0$, $w(1) = 1$, and $F_P(S_T)$ by estimating the process (21) we obtain a nonparametric estimate for the probability weighting function w. At a first glance, the estimation of the mean return μ appears to be the Achilles heel of our results. A high (low) μ corrsponds to significant (moderate or even negative) risk aversion in the economy. However, our results are surprisingly robust with respect to a specific choice of μ. Reported results here use the historical average return in our data set, i.e $\mu = 0.073$.

However, we have redone the analyses with μ ranging from 0.12 to 0.00. The latter value $\mu = 0.00$ corresponds to a negative excess return and thus induces risk proclivity in the economy. Furthermore, we also tested a perfect foresight of the realized μ. All results are pretty similar and reveal the hypothesized inverse-S shaped probability weighting function.

Unfortunately, nonparametric estimation techniques demand high sample sizes for reliable results. Therefore we sample our data similar to Aït-Sahalia, Wang, and Yared (2001). We consider all expiration dates in March, June, September, and December in our option data (44 expiration dates). For a given expiration date, we then collect all option prices with time to maturity less than twelve and more than ten weeks. We extract the state price density from these prices as described in Section 4.3. Furthermore, we collect S&P 500 prices every five minutes beginning ten weeks prior to the expiration and ending on the expiration date. This high frequency data is used to derive a nonparametric estimate $\hat{\sigma}(S_t)$ for the time series' volatility function $\sigma(S_t)$. This function is then used to simulate 1,000 paths of the processes (21) and (22). A kernel estimate with a Gaussian kernel function then provides the time-series-implied state price density function. The bandwidth in the kernel estimation is set to 1.5 times the bandwidth suggested by Silverman's (1986) rule of thumb. Note that this oversmoothing results in conservative estimates for our hypothesis as it implies fatter tails in the estimate of the time series state price density.

Following Aït-Sahalia, Wang, and Yared (2001) we estimate $\hat{\sigma}^2(S_t)$ with the consistent kernel estimator

$$\hat{\sigma}^2(S) = \frac{\sum_{i=1}^{N-1} k_{\text{time series}}\left(\frac{S_{t+\tau\cdot i/N}-S}{h_{\text{time series}}}\right) \cdot N \cdot \left(S_{t+\tau(i+1)/N} - S_{t+\tau\cdot i/N}\right)^2}{\sum_{i=1}^{N} k_{\text{time series}}\left(\frac{S_{t+\tau\cdot i/N}-S}{h_{\text{time series}}}\right)}. \tag{26}$$

We use the kernel function $k_{\text{time series}} = \frac{1}{2\pi}\exp(-\frac{x^2}{2})$. The bandwidth $h_{\text{time series}}$ is chosen such that $\hat{\sigma}(S_t)$ at the actual index price S_t equals the at-the-money option-implied volatility. Note that this choice again mitigates the effect of probability weighting as the time series realized volatility is usually lower than the option-implied volatility (see, for example, Christensen and Prabhala (1998)). Note that Aït-Sahalia, Wang, and Yared (2001) match the first two moments of the two state price densities and focus only on the third and fourth moments. We refrain from matching the first two moments as that would distort our estimate for the aggregate probability weighting function.

Figure 3 shows the average probability weighting function (solid line) for our 44 estimates. Additionally, we graph ± 0.5 times the empirical pointwise standard deviation (dashed lines). In line with our hypothesis we see a clear and notable inverse-S shape. However, the average probability weighting function intersects the identity line (gray) at higher probabilities than Prelec's (1998) parametric probability weighting function which crosses the identity line at $\frac{1}{e} \approx 0.368$.

5 Conclusion

There is ample evidence such as, for example, the Allais paradox that people's risk attitudes cannot be fully described by marginal utility. We therefore inferred the option implied risk preferences under RDU, one of the most prominent extensions of expected utility, which introduces probabilistic risk attitudes via probability weighting. Specifically, in most cases, we used a one-parameter extension to EUT to provide a parsimonious preference structure. Our results document an inverse-S shaped probability weighting function which is in line with psychological literature.

Many papers successfully fit other non-standard preferences to option prices (see, for example, Garcia, Luger, and Renault (2003), Liu, Pan, and Wang (2005), Benzoni, Collin-Dufresne, and Goldstein (2007), and Chabi-Yo, Garcia, and Renault (2008)). We did not make an attempt to quantitatively distinguish our hypothesis of an inverse-S shaped probability weighting function from other preferences. In fact, it is likely that this is an

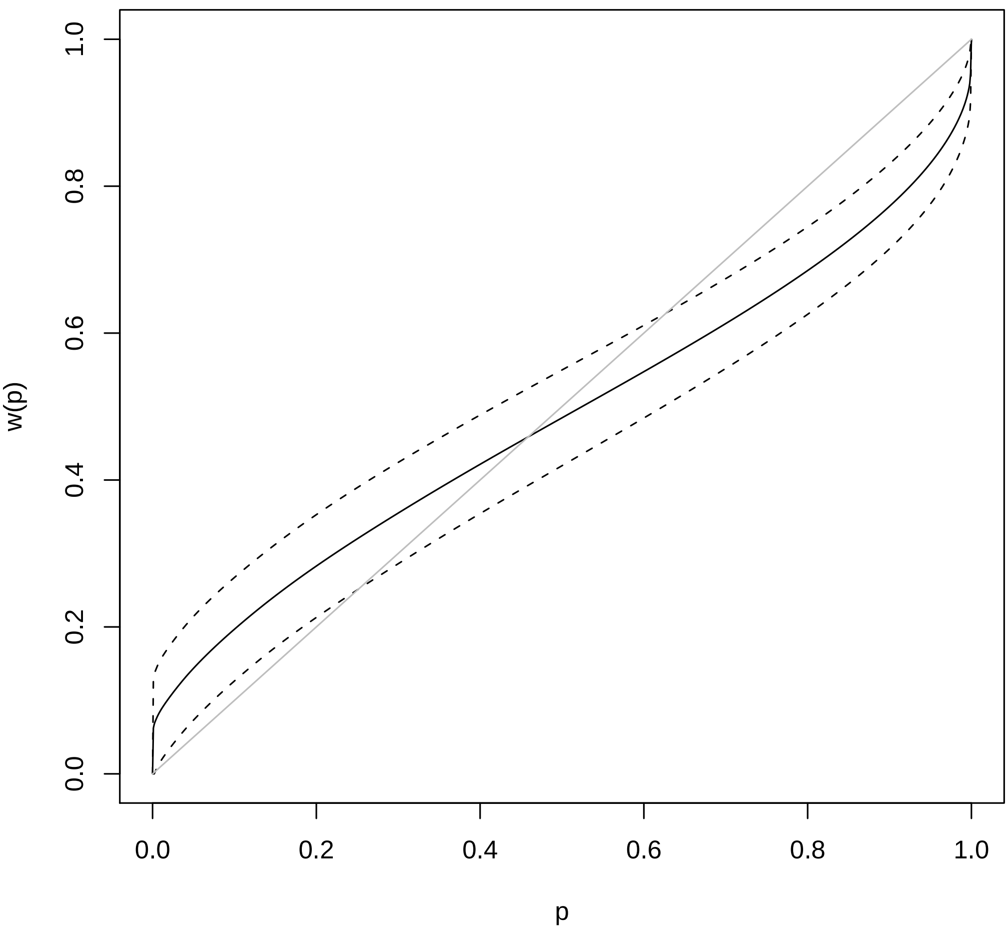

Figure 3: Average nonparametric estimate of the probability weighting function without parametric assumptions on the representative investor's utility function.
We estimated 44 state price densities from option prices written on the S&P 500 and the corresponding 44 state price densities from a high frequency time series of the S&P 500 with the methodologies in Aït-Sahalia, Wang, and Yared (2001). Under our hypothesis of RDU preferences, the ratio of the option-implied and time-series-implied state price density equals $w'\big(F_P(S_T)\big)$, i.e. the first derivative of the probability weighting function w at the cumulative distribution function $F_P(S_T)$ of the data-generating density. The average probability weighting function (solid line) is calculated with the knowledge of $w(0) = 0$ and $w(1) = 1$. Dashed lines represent the average probability weighting function ± 0.5 times the empirical pointwise standard deviation. The gray line corresponds to the identity function.

impossible mission as options data might well equally fit to different models, especially if these models make use of latent variables that are unobservable to the researcher. The virtue of our hypothesis is twofold. First, our model is extremely parsimonious. Second and probably more important, we motivate our research by findings on the micro-level, i.e. lab experiments and research from neuroscientists. Even more suggestive, even professional option traders exhibit the bias we exploit for our hypothesis (see Fox, Rogers, and Tversky (1996)). All it needs now to justify our model, is the fact that behavioral biases survive a competitive market environment like the S&P 500 index options market. And this assumption appears to be warranted by many studies in the finance literature (Stein (1989), Shleifer and Vishny (1998), Poteshman (2001), Poteshman and Serbin (2003), Coval and Shumway (2005), Haigh and List (2005), and Han (2008)). A Bayesian would formulate this argument as follows. Under a diffuse prior many preference models may explain observed option prices equally well. However, there are some arguments (see Kahneman and Tversky (1979), Tversky and Kahneman (1992), Fox, Rogers, and Tversky (1996), Hsu, Krajbich, Zhao, and Camerer (2009)) that the prior and thereby the posterior should favor our hypothesis of an inverse-S shaped probability weighting function. However, we take a frequentistic viewpoint here and conclude that we fail to reject RDU preferences with a well-behaved utility function and an inverse-S shaped probability weighting function. The subsequent work of Polkovnichenko and Zhao (2009) confirms our results.

We want to reiterate that our hypothesis relies on the specific nature of out-of-the-money options, which induce a specific framing of consumption risk. Here, traders have to deal with low probabilities and extreme outcomes. This situation is especially susceptible to an inverse-S shaped probability weighting function. Therefore, our hypothesis is silent about the equity premium puzzle (see Mehra and Prescott (1985)). Indeed, Epstein and Zin (1990) do not generate sufficient risk premia under RDU to explain the differences between consumption data and asset prices. A different type of framing and the behavioral concept of loss aversion that we did not incorporate might be the more relevant and driving factors here (see Barberis, Huang, and Santos (2001)). However, our results provide a unifying way to understand several puzzling facts about option prices.

References

Abdellaoui, Mohammed, 2000, Parameter-free elicitation of utilities and probability weighting functions, *Management Science* 46, 1497–1512.

Aït-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Aït-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Aït-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do option markets correctly price the probabilities of movement of the underlying asset?, *Journal of Econometrics* 102, 67–110.

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

Bakshi, Gurdip, and Dilip Madan, 2008, Investor heterogeneity and the non-monotonicity of the aggregate marginal rate of substitution in the market index, *Working Paper.*

Bakshi, Gurdip, Dilip Madan, and George Panayotov, 2009, Returns of Claims on the Upside and the Viability of U-Shaped Pricing Kernels, *Working Paper, University of Maryland.*

Bakshi, Gurdip, and George Panayotov, 2008, A framework for studying option mispricing: A new test and empirical evidence, *Working Paper.*

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

Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and assets prices, *The Quarterly Journal of Economics* 116, 1–53.

Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics* 94, 181–238.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2007, Explaining Pre- and Post-1987 Crash Prices of Equity and Options within a Unified General Equilibrium Framework, *Working Paper, University of Minnesota.*

Berkowitz, Jeremy, 2001, Testing density forecasts with applications to risk management, *Journal of Business and Economic Studies* 19, 465–474.

Bleichrodt, Han, and Jose Luis Pinto, 2000, A parameter-free elicitation of the probability weighting function in medical decision analysis, *Management Science* 46, 1485–1496.

Bleichrodt, Han, Jaco van Rijn, and Magnus Johannesson, 1999, Probability weighting and utility curvature in QALY-based decision making, *Journal of Mathematical Psychology* 43, 238–260.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal of Finance* 59, 407–446.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contigent claims implicit in options prices, *Journal of Business* 51, 621–651.

Brown, David, and Jens Carsten Jackwerth, 2004, The Pricing Kernel Puzzle: Reconciling Index Option Data and Economic Theory, *Working Paper, University of Wisconsin, Madison.*

Camerer, Colin F., and Teck-Hua Ho, 1994, Violations of the betweenness axiom and nonlinearity in probability, *Journal of Risk and Uncertainty* 8, 167–196.

Chabi-Yo, Fousseni, Renè Garcia, and Èric Renault, 2008, State dependence can explain the risk aversion puzzle, *Review of Finacial Studies* 21, 973–1011.

Chapman, David A., and Valery Polkovnichenko, 2009, First-order Risk Aversion, Heterogeneity, and Asset Market Outcomes, *Journal of Finance* 64, 1863–1888.

Chew, Soo Hong, Edi Karni, and Zvi Safra, 1987, Risk aversion in the theory of expected utility with rank dependent probabilities, *Journal of Economic Theory* 42, 370–381.

Christensen, B. J., and N. R. Prabhala, 1998, The relationship between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.

Constantinides, George M., Jens Carsten Jackwerth, and Stylianos Perrakis, 2009, Mispricing of S&P 500 Index Options, *Review of Financial Studies* 22, 1247–1277.

Coval, Joshua D., and Tyler Shumway, 2005, Do behavioral biases affect prices?, *Journal of Finance* 60, 1–34.

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

Epstein, Larry G., and Stanley E. Zin, 1990, First-order' risk aversion and the equity premium puzzle, *Journal of Monetary Economics* 26, 387–407.

Etchart-Vincent, Nathalie, 2004, Is probability weighting sensitive to the magnitude of consequences? An experimental investigation on losses, *Journal of Risk and Uncertainty* 28, 217–235.

Fox, Craig R., Brett A. Rogers, and Amos Tversky, 1996, Options traders exhibit subadditive decision weights, *Journal of Risk and Uncertainty* 13, 5–17.

Garcia, Renè, Richard Luger, and Èric Renault, 2003, Empirical Assessment of an Intertemporal Option Pricing Model with Latent Variables, *Journal of Econometrics* 116, 49–83.

Gollier, Christian, 2006, Does ambiguity aversion reinforce risk aversion? Applications to portfolio choices and asset prices, *Working Paper, University of Toulouse.*

Gonzalez, Richard, and George Wu, 1999, On the shape of the probability weighting function, *Cognitive Psychology* 38, 129–166.

Gurevich, Gregory, Doron Kliger, and Ori Levy, 2009, Decision-making under uncertainty - A field study of cumulative prospect theory, *Journal of Banking and Finance* 33, 1221–1229.

Haigh, Michael S., and John A. List, 2005, Do professional traders exhibit myopic loss aversion? An experimental analysis, *Journal of Finance* 60, 523–534.

Han, Bing, 2008, Investor sentiment and option prices, *Review of Financial Studies* 21, 387–414.

Hsu, Ming, Ian Krajbich, Chen Zhao, and Colin F. Camerer, 2009, Neural Response to Reward Anticipation under Risk Is Nonlinear in Probabilities, *Journal of Neuroscience* 29, 2231–2237.

Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect Theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.

Liu, Jun, Jun Pan, and Tan Wang, 2005, An Equilibrium Model of Rare Event Premia and Its Implication for Option Smirks, *Review of Financial Studies* 18, 131–164.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Polkovnichenko, Valery, and Feng Zhao, 2009, Probability Weighting Functions Implied in Option Prices, *Working Paper, University of Texas at Dallas.*

Post, Thierry, and Haim Levy, 2005, Does risk seeking drive stock prices? A stochastic dominance analysis of aggregate investor preferences and beliefs, *Review of Financial Studies* 18, 925–953.

Poteshman, Allen M., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market, *Journal of Finance* 56, 851–876.

Poteshman, Allen M., and Vitaly Serbin, 2003, Clearly irrational financial market behavior: Evidence from the early exercise of exchange traded stock options, *Journal of Finance* 58, 37–70.

Prelec, Drazen, 1998, The probability weighting function, *Econometrica* 66, 497–527.

Quiggin, John, 1981, Risk perception and risk aversion among Australian farmers, *Australian Journal of Agricultural Economics* 25, 160–169.

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

Rosenberg, Joshua V., and Robert F. Engle, 2002, Empirical pricing kernels, *Journal of Financial Economics* 64, 341–372.

Ryan, Matthew J., 2006, Risk aversion in RDEU, *Journal of Mathematical Economics* 42, 675–697.

Santa-Clara, Pedro, and Alessio Saretto, forthcoming, Option strategies: Good deals and margin calls, *Journal of Financial Markets*.

Schmeidler, David, 1989, Subjective probability and expected utility without additivity, *Econometrica* 57, 571–587.

Shleifer, Andrei, and Robert Vishny, 1998, The limits of arbitrage, *Journal of Finance* 52, 35–55.

Silverman, Bernard W., 1986, *Density estimation*. Chapman and Hall.

Snowberg, Erik, and Justin Wolfers, 2007, Explaining the favorite-longshot bias: Is it risk-love or misperceptions?, *Working Paper*.

Sonnemann, Ulrich, Colin F. Camerer, Craig R. Fox, and Thomas Langer, 2009, Partition dependence in financial binary option and prediction markets: Field and lab evidence, *Working Paper, University of Muenster*.

Stein, Jeremy, 1989, Overreactions in the options market, *Journal of Finance* 44, 1011–1022.

Stott, Henry P., 2006, Cumulative Prospect Theory's functional menagerie, *Journal of Risk and Uncertainty* 32, 101–130.

Tversky, Amos, and Craig R. Fox, 1995, Weighing risk and uncertainty, *Psychological Review* 102, 269–283.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in Prospect Theory: Cumulative Representation of Uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

Wakker, Peter, 1994, Separating marginal utility and probabilistic risk aversion, *Theory and Decision* 36, 1–44.

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

Ziegler, Alexandre, 2007, Why does implied risk aversion smile?, *Review of Financial Studies* 20, 859–907.