European Finance Review **6:** 277–290, 2002.

Risk Preferences Heterogeneity: Evidence from Asset Markets *

DORON KLIGER and ORI LEVY
The University of Haifa, Israel. E-mails: kliger@econ.haifa.ac.il; orilevy@econ.haifa.ac.il

Abstract. Using asset market data, as well as theoretical relations between investors' preferences, option-implied, risk-neutral, probability distribution functions (PDFs,) and index-implied, actual, PDFs, this paper extracts a time-series of investors' relative risk aversion (RRA) functions. Based on results recently derived by Benninga and Mayshar (2000), these functions are used to recover the evolution of risk preferences heterogeneity. Applying non-parametric estimation on European call options written on the S&P500 index, we find that: (i) the RRA functions are decreasing; and (ii) the constructed risk preferences heterogeneity series is positively correlated in a static, as well as a dynamic, setup with a prevalent proxy for investors heterogeneity, namely, the spread between auction- and market-yields of Treasury bills.

Keywords: preferences heterogeneity, relative risk aversion, S&P500 index options.

JEL classification: D81, G12, G13.

1. Introduction

We use asset market data, applying theoretical relations between investors' preferences, option-implied, risk-neutral, Probability Distribution Functions (PDFs,) and index-implied, actual, PDFs, to extract a time-series of investors' relative risk aversion (RRA) *functions*.[1] We then use these RRA functions to recover the time evolution of investors' risk preferences heterogeneity.

Our estimated RRA functions are used to test Benninga and Mayshar's (2000) theory, asserting that in an economy populated by agents with constant, yet different coefficients of RRA: (i) preferences of the "pricing-representative" consumer exhibit decreasing RRA (their Proposition 2); and (ii) in "high" ("low") aggregate

* We wish to thank Jens Carsten Jackwerth and Mark Rubinstein for providing the option data, seminar participants at Tel-Aviv University, and especially Simon Benninga and an anonymous referee, for helpful discussions and comments. All errors are ours.

[1] Bick (1990), He and Leland (1993), Wang (1993), Dybvig and Rogers (1997), and Cuoco and Zapatero (2000) discuss the recoverability of preferences from observed asset prices and an agent's consumption choice. Ait-Sahalia and Lo (2000) and Jackwerth (2000) empirically extract risk aversion functions using theoretical relations between state price densities, actual probability distribution functions and risk preferences.

consumption states, the RRA coefficient of the "pricing-representative" consumer resembles the RRA coefficient of the consumer exhibiting the weakest (strongest) aversion toward risk (their Proposition 3.)[2]

The empirical analysis we conduct is twofold. First, we investigate whether aggregate RRA functions are indeed decreasing. Second, we use the slopes of the RRA functions to characterize the dynamics of risk preferences heterogeneity, and examine the correlation of the RRA slope with a commonly used proxy for investors' heterogeneity, namely, the spread between auction- and market-yields of Treasury bills. We do so in a static setup, where contemporaneous RRA slopes and yield spreads are compared, as well as in a dynamic setup, where RRA slopes are compared with a history (time series) of yield spreads. The idea of measuring investors' heterogeneity by the steepness of demand and supply curves was advanced by, e.g., Bagwell (1992) who uses elasticities of upward-sloping supply curves firms face when repurchasing their shares in Dutch auctions, and by Kandel et al. (1999) who use demand elasticities in the market for initial public offerings. Also, Wang (1996) shows that risk-preferences heterogeneity is positively correlated with the dispersion of short-term bond yields.[3]

Our empirical study of implied RRA functions and risk-preferences heterogeneity exploits the relation between investors' preferences, risk-neutral PDFs, and actual PDFs. We estimate expiration-date option-implied, risk-neutral, PDFs by employing Breeden and Litzenberger's (1978) non-parametric approach on European call options written on the S&P500 index.[4] To estimate the corresponding expiration-date index-implied, actual, PDFs, we use a non-parametric approach consisting of sampling with replacement ("bootstrapping") of actual index-return realizations throughout the options' remaining time to expiration.

The rest of the paper is organized as follows. Section 2 starts with the introduction of theoretical relations between implied risk aversion, risk-neutral PDFs, and actual PDFs, turning then to the specific approaches used for deriving option- and index-implied PDFs. Section 3 describes the data, and discusses and motivates the applied statistical model. Our empirical analysis is presented in Section 4. Section 5 concludes.

---

[2] Benninga and Mayshar (2000) define a "pricing-representative" agent to be (p. 10,) "one whose tastes are such that if all agents in the economy had tastes identical to his, then the equilibrium state prices in the economy would remain unchanged".

[3] See Heaton and Lucas (1996), Brav et al. (2000), Storesletten et al. (2000), and references therein for the fast-growing empirical work (using microeconomic data) in a heterogeneous-agent environment.

[4] Since in practice the spacing between strike prices is discrete and the total number of traded options, especially those that are far in- and out-of-the-money, is relatively small, there could be no unique empirical SPD consistent with observed option prices. Bates (1991), Soderlind and Svensson (1997), Posner and Milevsky (1998), and Rubinstein (1998) propose parametric approaches for estimating entire SPDs, whereas Shimko (1993), Rubinstein (1994), Jackwerth and Rubinstein (1996), and Ait-Sahalia and Lo (1998) propose non-parametric approaches. Bahra (1997) and Jackwerth (1999) provide comprehensive surveys.

2. Implied Risk Preferences

Theoretically, investors' preferences, risk-neutral PDFs, and actual PDFs are tightly related. The purpose of this Section is to make precise the way actual and risk neutral probabilities are being used to extract the usually unobserved RRA functions. First, theoretical relations between risk preferences and shapes of risk-neutral and actual PDFs are set forth. Next, we introduce the non-parametric approach used for estimating option-implied PDFs. Finally, we describe the statistical approach for estimating the relevant index-implied PDFs.

2.1. RISK PREFERENCES, RISK-NEUTRAL PDFS, AND ACTUAL PDFS

Our study exploits a well-known relation between investors' preferences, risk-neutral state probabilities, and actual state probabilities. In particular,

$$U_t'(S, \tau) \cdot Q_t(S, \tau) \propto P_t(S, \tau), \tag{2.1}$$

where $P_t(S, \tau)$ and $Q_t(S, \tau)$ are the time-t risk-neutral and actual probabilities of state S at time $t + \tau$, respectively, and $U_t'(S, \tau)$ is the time $t + \tau$ marginal utility. Differentiating 2.1 with respect to S, and dividing by the same equation yields:

$$\frac{Q_t'(S, \tau)}{Q_t(S, \tau)} + \frac{U_t''(S, \tau)}{U_t'(S, \tau)} = \frac{P_t'(S, \tau)}{P_t(S, \tau)}. \tag{2.2}$$

Multiplying by S and rearranging yields the coefficient of RRA, $\theta_t(S, \tau)$:

$$\theta_t(S, \tau) \equiv -S \cdot \frac{U_t''(S, \tau)}{U_t'(S, \tau)} = S \cdot \frac{Q_t'(S, \tau)}{Q_t(S, \tau)} - S \cdot \frac{P_t'(S, \tau)}{P_t(S, \tau)}. \tag{2.3}$$

Note that $\theta_t(S, \tau)$ is locally identified from the shapes of the risk-neutral and actual PDFs around S.[5]

2.2. OPTION-IMPLIED RISK-NEUTRAL PDFS

Breeden and Litzenberger (1978) propose an approach to extract unobserved expiration-date risk-neutral PDFs of the underlying asset from prices of European

[5] Our empirical analysis is based on post-1987 crash data. It is worthwhile noting that although there are no crashes or large runups in the estimation period, the RRA estimates are not prone to be affected by a "Peso problem". To clarify this claim, assume that the estimated index-implied PDFs are conditional on the monthly returns being moderate, and that the probability of an extreme monthly return is ε. If this is the case, then our estimate of the (unconditional) actual probability, $Q_t(S, \tau)$, is upward biased by a factor of $(1 - \varepsilon)$. But the slope of that probability, $Q_t'(S, \tau)$, is upward biased by the same factor, leaving the ratio $Q_t'(S, \tau)/Q_t(S, \tau)$, and thus also the estimated coefficient of RRA unaffected.

options. We henceforth briefly demonstrate their result, and accommodate it to the discrete finite set of strike prices available in the market.[6]

Consider the option pricing interpretation of Cox and Ross (1976), summarized as follows,

$$\begin{aligned} C &= \exp(-R_F\tau)E^*(\max(0, S-K)) \\ &= \exp(-R_F\tau)\int_{-\infty}^{\infty} P(S)\max(0, S-K)\,dS, \end{aligned} \tag{2.4}$$

where C is the current price of a European call option with strike price K, R_F is the risk-free rate till the option's expiration, τ is the time to expiration, E^* denotes the conditional expectation operator according to the *risk-neutral* PDF, $\max(0, S-K)$ is the option's expiration-date payoff, and $P(S)$ is the expiration-date risk-neutral probability. Differentiating (2.4) with respect to K yields (see, e.g., Soderlind and Svensson (1997)),

$$\partial C/\partial K = -\exp(-R_F\tau)\int_K^{\infty} P(S)\,dS. \tag{2.5}$$

Differentiating with respect to K once more yields the risk-neutral PDF:

$$P(S)|_{S=K} = \exp(R_F\tau)\cdot\partial^2 C/\partial K^2. \tag{2.6}$$

To accommodate $\partial^2 C/\partial K^2$ in (2.6) to a discrete set of strike prices, suppose that an investor simultaneously sells two European call options with a strike price K_j at a price of $C(K_j)$ each, and buys two European call options with strike prices $K_j+\Delta$ and $K_j-\Delta$ at the prices $C(K_j+\Delta)$ and $C(K_j-\Delta)$, respectively.[7] Clearly, this strategy (known as a "butterfly spread", centered around K_j) approximates an expiration-date payoff of Δ conditional on S being around K_j. Dividing by Δ to normalize the payoff to unity, the cost of the strategy is:[8]

$$\frac{C(K_j+\Delta)-C(Kj)}{\Delta} - \frac{C(K_j)-C(K_j-\Delta)}{\Delta}. \tag{2.7}$$

Dividing by Δ again and letting $\Delta \to 0$ approximates $\partial^2 C/\partial K^2$.[9]

If the call option pricing function decreases monotonically in the strike price and is convex (which in essence rules out arbitrage opportunities,) it is guaranteed that

[6] For ease of presentation, we suppress the time index.

[7] Prominent examples of non-parametrically extracting risk-neutral PDFs are Jackwerth and Rubinstein (1996), and Ait-Sahalia and Lo (1998).

[8] Since three successive strikes are required for each butterfly spread, only J risk-neutral probabilities may be recovered in a cross-section with $J+2$ quotes. The slope of the risk-neutral probability is separately estimated between each two adjacent probability levels, leaving us with $J-1$ point estimates.

[9] Note that the method of Breeden and Litzenberger (1978) for estimating PDFs does not depend on equal spacing of the observed strike prices. In particular, for any three successive strike prices K_{j-1}, K_j, and K_{j+1}, the strategy cost becomes: $\frac{C(K_{j+1})-C(K_j)}{K_{j+1}-K_j} - \frac{C(K_j)-C(K_{j-1})}{K_j-K_{j-1}}$.

the risk-neutral probabilities are positive, since all the above-mentioned butterfly spreads have positive prices. Admittedly, an inherent problematic feature of the proposed approach is that nothing forces the probabilities to sum to unity. Indeed, for the commonly observed range of strike prices, the sum is smaller. Consequently, the approach is unable to allocate perceived probabilities in the tails of the PDF. Nonetheless, as we have seen, the proposed approach for estimating implied risk preferences is not hampered by this shortcoming since investors' RRA at K_j can be estimated using the risk-neutral and actual PDFs *around* K_j *only*.

2.3. INDEX-IMPLIED, TIME-VARYING, ACTUAL PDFS

To estimate the expiration-date index-implied actual PDFs, we bootstrap actual S&P500 index returns (Efron (1979).)[10] Specifically, actual return realizations throughout the remaining time to expiration of each cross-section of options are randomly sampled with replacement, to generate a time-t dependent actual PDF.[11]

Let τ be the number of days until the options' expiration (in our sample, τ ranges from 23 to 43 days) and let $N < \tau$ be the number of actual trading days in that period. In addition, let $X_t, \ldots, X_{t+N}$ be price realizations of the S&P500 index and $\log\left(\frac{X_{t+n}}{X_{t+n-1}}\right)$, $n = 1, \ldots, N$, the daily (logarithmic) rates of return, assumed to be independent and identically distributed. Sampling the trading days with replacement and summing the daily returns of each N draws yields an empirical realization R_t^*,

$$R_t^* \equiv \sum_{n=1}^{N} \log\left(\frac{X_{t+k_n}}{X_{t+k_n-1}}\right), \tag{2.8}$$

where k_n denotes the nth randomly-sampled trading day. We then re-express R_t^* in terms of simple (i.e., non-logarithmic) return R_t,

$$R_t = \exp(R_t^*) - 1. \tag{2.9}$$

We repeat that procedure to produce 100, 000 empirical return realizations, which we henceforth denote by $R_{t,i}$, $i = 1, \ldots, 100{,}000$.

The actual (empirical) probability may now be approached. Define state j by means of a set of return values around K_j, commensurate with the butterfly spread in Subsection 2.2. For simplicity, we denote the probability of these returns by $Q_t(K_j)$, which is estimated as follows,

$$Q_t(K_j) = \frac{1}{100,000} \sum_{i=1}^{100,000} I(R_{t,i} \in \text{state } j), \tag{2.10}$$

[10] S&P500 index returns are taken from the University of Chicago's Center for Research on Security Prices (CRSP) files.

[11] For a comprehensive survey of the theoretical and methodological aspects of the bootstrap, see, e.g., Hall (1994).

where $I(\cdot)$ is an indicator function that receives the value 1 whenever $R_{t,i}$ falls within state j and zero otherwise.

Jackwerth (2000) is the first to provide empirical estimates of agents' risk preferences using information regarding shapes of risk-neutral and actual PDFs. Interestingly, the recovered risk aversion functions are inconsistent with standard assumptions usually made in utility theory. In particular, in the post-October 1987 crash, implied risk aversion functions occasionally exhibit increasing absolute risk aversion and coefficients of ARA implying local convexity of the utility function in various returns. Jackwerth (2000) assumes that actual PDFs can be approximated by time-invariant index-implied PDFs based on monthly non-overlapping returns over a four-year period. The approach adopted in this paper extends his contribution, by allowing for an estimation of *time-varying* actual PDFs, required for a dynamic analysis of preferences. It is shown below that the particular choice of actual PDFs' estimation may be consequential for the implied RRA functions.

3. The Data and Statistical Model

In this Section we provide a concise description of the data we use. We start by describing the data for estimating the option-implied PDFs. Then, we describe the proxy of investors' heterogeneity-the spread between the yield of Treasury bills at the secondary market and auctions' weighted average yield from multiple-price auctions. We conclude the Section by motivating our choice of a statistical model, in which temporal RRA functions are cross-sectionally fitted to the option and index data.

3.1. OPTION SAMPLE STATISTICS

For estimating one-month-ahead option-implied PDFs, we use prices (across strike prices) sampled on a monthly basis of European call options written on the S&P500 index, traded on the Chicago Board Options Exchange. The data we use are from Jackwerth and Rubinstein (1996) and Jackwerth (2000) and so will only be briefly discussed.[12] Option prices are averages of bid and ask quotes adjusted by a penalty parameter (see Jackwerth and Rubinstein (1996) for a thorough discussion on the penalty method.) The dataset extends from December 1987 (i.e., immediately after the October 1987 crash) through December 1995 (the end of the Jackwerth–Rubinstein database) and includes 74 independent cross-sections of option premia for which there were no arbitrage violations. The selection of our sample's starting date is motivated by empirical evidence of a "regime shift" (non-stationarity) in option prices following the October 1987 crash (see, e.g., Rubinstein (1994), Jackwerth and Rubinstein (1996), and Christensen and Prabhala (1998).) Note also that until March 1986 options written on the S&P500 were American, rather than European.

[12] We are grateful to Jens Carsten Jackwerth and Mark Rubinstein for providing us with their data.

Table I. Data description

	No. of options	Time to expiration	Volatility	Yield spread level	Yield spread ratio
Mean	15.41	31.61	15.73	0.17	2.84
Median	15.00	30.00	14.26	0.17	2.88
Max	26.00	43.00	35.72	0.36	4.30
Min	6.00	23.00	9.17	0.05	1.32
Std. dev.	4.58	4.06	5.42	0.08	0.67
No. obs.			74		

The table describes the 74 cross sections of options, as well as yield spreads in the data set. 'Volatility' is the annual, at the money volatility implied by the Black and Scholes (1973) option pricing equation. 'Yield Spread Level' is the difference between the yield of Treasury bills at the secondary market and auctions' weighted average yields from multiple-price auctions. 'Yield Spread Ratio' is the yield spread level divided by the secondary market yield. The data on auction- and secondary market-yields are obtained from the Chicago Federal Reserve. The sample extends from December 1987 through December 1995.

Table I describes the selected sample. Our dataset consists of 1140 options, with 6 to 26 options in a cross section and an average of 15.4. The cross-sectional time-to-expiration ranges from 23 to 43 days, with an average of 31.6. At-the-money volatility, implied by the Black and Scholes (1973) pricing equation ranges from 9.2% to 35.7% (in December 1987,) averaging 15.7%. Risk-free nominal interest rates are the implied interest rates imbedded in the put-call parity, a relation that must hold in the absence of arbitrage opportunities.[13]

Figure 1 plots the panel of recovered RRA coefficients on rates of return, as well as a non-parametric curve fitted to these values.

Apparently, the fitted RRA curve is decreasing over most of the return domain. The curve becomes negative at a (monthly) rate of return of about 3%, suggesting plausibility of models in which preferences exhibit risk seeking behavior at high stakes (see, e.g., Camerer (1995) for a comprehensive survey.)

[13] The put-call parity states that,

$$C_t + K \exp(-R_t^F \tau) = P_t + S_t,$$

where C_t is the price of a European call option, P_t is the price of a European put option with the same strike price K and time to expiration τ, R_t^F is the risk-free rate of return, and S_t is the current underlying asset's price.

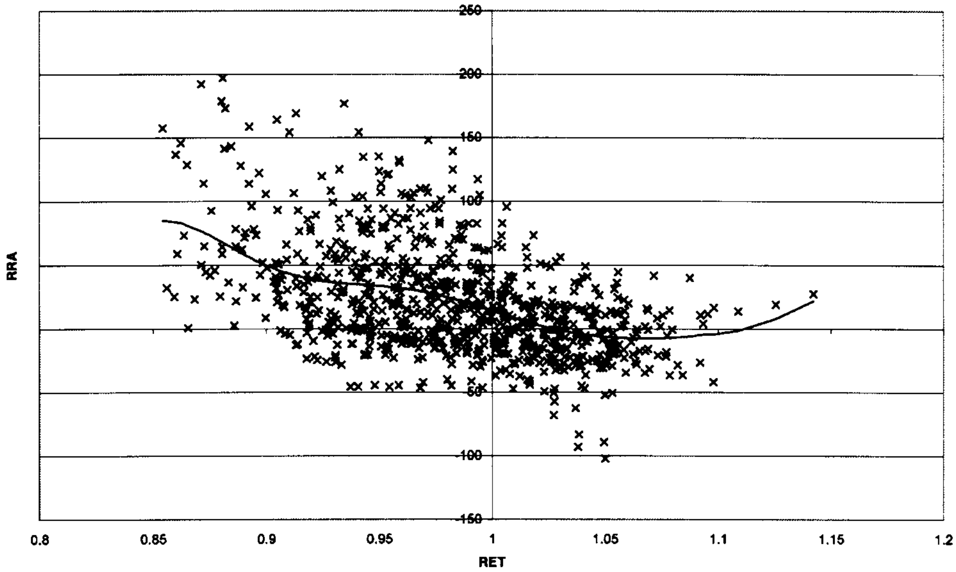

Figure 1. The figure plots coefficient of RRA estimates on rates of return ('RET'), as well as a non-parametric curve fitted to these values. Time- and state-dependent coefficients of RRA are estimated using (2.3). Risk-neutral probabilities are option-implied and estimated using the non-parametric approach outlined in Subsection 2.2. Actual probabilities are index-implied and estimated using the non-parametric approach outlined in Subsection 2.3. The panel of coefficients of RRA includes estimates for the period December 1987 (post October 1987 market crash) through December 1995.

3.2. SPREAD CALCULATION

The yield spread we use, termed *Spread* (see Table I) is defined using (i) the yield of Treasury bills at the secondary market, and (ii) auctions' weighted average yields from multiple-price auctions. We use two alternative spread specifications. The first, *Yield Spread Level*, is the difference between (i) and (ii), and the second, *Yield Spread Ratio*, equals the difference, divided by the secondary market yield. The data on auction- and secondary-market yields are obtained from the Chicago Federal Reserve. The secondary market yields are provided on a constant maturity basis, as constructed by the Treasury Department, based on the most actively traded marketable Treasury securities. Yields on these issues are based on composite quotes reported by US government security dealers to the Federal Reserve Bank of New York. To obtain the constant maturity yields, a yield curve is constructed at each business day and the yield values are read from the curve at fixed maturities. Auctions' weighted average yields are calculated using the accepted bids in multiple-price Treasury auctions of 13-week bills. Currently, the auctions are held each Monday for bills that will be issued on the ensuing Thursday, in the absence of holidays.

To arrive at the time-series of yield spreads, we match each cross-section of options with the closest preceding auction date. This matching procedure results in time gaps usually not exceeding four days (but for one case of a 7-day gap.)

3.3. THE STATISTICAL MODEL

As we have emphasized, we estimate option-implied PDFs for each cross-section of option premia and thus have estimates of 74 independent implied ex-ante risk aversion functions. Nonetheless, every cross-section t, $t = 1, \ldots, 74$ is possibly contaminated by idiosyncratic "noise" that can be the result of, e.g.: the bootstrap procedure, affecting estimates of actual probabilities; mis-measurement of option prices, affecting risk-neutral probabilities (clearly, the two need not be mutually exclusive.) This motivates our choice of a statistical model which assumes that a vector of parameters is randomly drawn in every cross-section, a framework known in the literature as the "random coefficients model".[14] In particular, we let the empirical linear equation be of the following form,

$$\theta_t(x_j) \equiv f(x_j; \Psi_t) = \alpha_t + \beta_t \cdot x_j + \varepsilon_{t,j}, \tag{3.1}$$

where x_j is the index return commensurate with the jth state, $\Psi_t' \equiv (\alpha_t, \beta_t)$, and $E(\varepsilon_t) = 0$. The estimation results in a time-series of estimated risk aversion functions $\theta_t(\cdot)$, $t = 1, \ldots, 74$.[15]

In Figure 2 we plot the 74 temporally fitted RRA functions, $\hat{\theta}_t \equiv E[f(\hat{\Psi}_t)]$, for the period December 1987 through December 1995. Conspicuously, the random coefficients model fits the data rather well. Indeed, the statistical model is capable of explaining 84% of the overall variation in the coefficient of RRA estimates.

Furthermore, these results are clearly consistent with the DRRA hypothesis, since virtually all of the estimated aggregate recovered RRA functions are decreasing (i.e., $\beta_t < 0$, for all t). The latter finding is *in line* with Benninga and Mayshar (2000) (their Proposition 2), claiming that if individuals are endowed with preferences of the CRRA type, aggregate (observed) risk preferences should exhibit DRRA.

[14] For a comprehensive survey of the theoretical and methodological aspects of random coefficient models, see, e.g., Chow (1984).

[15] Ait-Sahalia and Lo (2000) *aggregate* S&P500 index option premia over time (specifically, over the year 1993) to provide a single time-averaged risk-neutral PDF of the underlying asset, and thus a single risk aversion function. As we have noted, our approach is fundamentally different in that we estimate a separate option-implied PDF for each cross-section of option premia (a total of 74). Consequently, we have 74 separate estimates of implied ex-ante RRA functions.

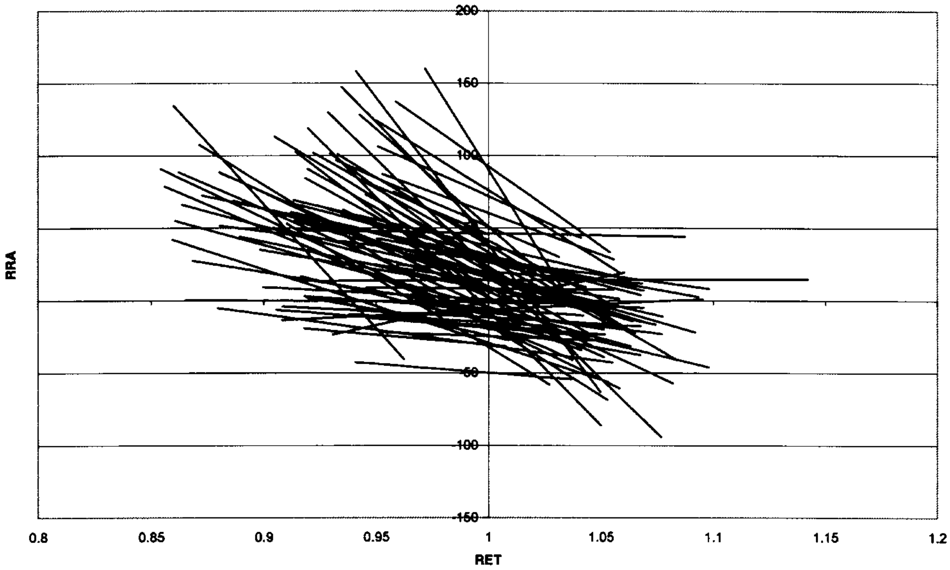

Figure 2. The figure plots forecasts of RRA functions using the log-linear specification in (3.1), on rates of return ('RET'). The sample extends from December 1987 (post October 1987 market crash) through December 1995 and includes 74 independent implied ex-ante risk aversion functions.

4. Empirical Analysis

Building on Benninga and Mayshar (2000) (Proposition 3),[16] we aim at testing whether the slope of the RRA function, β_t, may serve as a measure of investors' time-t risk preferences heterogeneity. Recall that the yield spread ($Spread_t$) serves as a commonly used proxy for investors' heterogeneity. If, indeed, β_t is a valid heterogeneity proxy, then β_t and $Spread_t$ should exhibit mutual variation over time. Specifically, we conjecture that the steeper the RRA function (the more *negative* is β_t), the wider the spread, i.e., we expect β to be *negatively* correlated with *Spread*. We thus specify the following equation:

$$\beta_t - (A + B \cdot Spread_t) = \varepsilon_t, \tag{4.1}$$

where ε_t is a mean-zero error term, and conjecture that $B < 0$.

In addition to heterogeneity in risk attitudes, the yield spread may capture other differences among the market participants (such as information asymmetry and taxation.) The existence of various heterogeneity sources might hamper our ability to find support for the above conjecture.

[16] Asserting that if $C \to \infty$ ($C \to 0$), then $\gamma^* \to \mathrm{Min}_i\{\gamma_i\}$ ($\mathrm{Max}_i\{\gamma_i\}$), where C denotes aggregate consumption, γ^* is the pricing-representative consumer's degree of relative risk aversion, and γ_i is agent's i constant coefficient of RRA.

Table II. Regression analysis

Model Specification	Single period Level	Single period Ratio	Multiple period Level	Multiple period Ratio
A	0.054	0.052	0.150	0.196
	(0.420)	(0.426)	(0.207)	(0.228)
B	−0.179	−0.150	−0.276	−0.309
	(0.006)	(0.056)	(0.002)	(0.007)
C			0.398	0.546
			(0.001)	(0.000)
χ^2	6.414	6.543	5.662	5.549
Prob.	0.093	0.088	0.059	0.062
No. obs. (adjusting for endpoints)	73			

The table presents GMM estimation results, as discussed in Section 4, where definitions of the two cases ('Single Period' and 'Multiple Period') as well as economic interpretations of the estimated coefficients are given. 'Level' and 'Ratio' are the two alterative yield spread definitions. Two-tail p-values appear in parentheses.

Equation (4.1) is estimated using the Generalized Method of Moments (GMM) of Hansen (1982).[17] Throughout, we use the two alternative definitions ('level' and 'ratio') of $Spread_t$. Since both β and *Spread* are non-stationary in our sample (i.e., one cannot reject the unit root hypothesis using either the Dickey–Fuller or the Phillips–Perron unit root tests), estimation is performed using the detrended series obtained after applying the Hodrick and Prescott (1997) filter and standardizing the series to have a standard deviation of one.[18] The estimation results are collected in Table II.

As can be seen in the table, B is negative as predicted. For the *Yield Spread Level* specification, the estimated B is -0.18, and significantly different from zero (two-tail p value of 0.006,) suggesting that a one standard deviation increase in *Spread* is associated with 0.18 standard deviations decrease in the (negative) slope of the RRA function. The picture emerging from the *Yield Spread Ratio* specification is similar (B estimate of -0.15, and p value of 0.056.) Furthermore, testing the over-identifying restrictions using the J statistic indicates that both specifications cannot be rejected (probability of 0.093 and 0.088, for the 'level' and 'ratio' specifications respectively.)

[17] The pool of instrumental variables used in the estimation consists of a constant, and current and lagged variables. As long as the variables appearing in the estimated equation were included, the estimation results were hardly sensitive to the addition of higher-order lags. We thus report only the results obtained with current and first-lag values as instruments.

[18] We use a smoothing parameter equal to 100.

Having established the validity of the RRA slope (β) as a measure of investors' heterogeneity, a more detailed look in a dynamic (multiple-period) context is in place. Note that the commonly used heterogeneity proxy, *Spread*, is mainly determined by auction (i.e., primary market) bids. Therefore, it may be viewed as a 'flow', rather than a 'stock' variable. On the other hand, β is determined by (secondary) *market* prices, and thus may better represent the aggregate heterogeneity.

Let g be a measure of heterogeneity depending on the history (time-series) of spreads:

$$\begin{aligned} &g(Spread_t, \ldots, Spread_{t-L}, \ldots, \Gamma) = A + B \cdot Spread_t + \\ &C \cdot B \cdot Spread_{t-1} + \cdots + C^L \cdot B \cdot Spread_{t-L} + \cdots \end{aligned} \tag{4.2}$$

where $\Gamma' \equiv (A, B, C)$. The coefficient C provides the relative weights of the temporal spreads, where $0 \leq C < 1$ is consistent with a decaying effect of past spreads. The case $C = 0$ reduces to the static (single-period) context of Equation (4.1).

To test whether the RRA slope captures the information regarding heterogeneity contained in the history of spreads, i.e., whether β and g exhibit mutual variation over time, we specify the following empirical equation:

$$\beta_t - g(Spread_t, \ldots, Spread_{t-L}, \ldots; \Gamma) = \epsilon_t, \tag{4.3}$$

where ϵ_t is a mean-zero error term. For empirical estimation, we use the first lag of Equation (4.3) to represent lagged *Spread* values, and obtain:[19]

$$\beta_t - (A \cdot (1 - C) + B \cdot Spread_t + C \cdot \beta_{t-1}) = \xi_t, \tag{4.4}$$

where $\xi_t = \epsilon_t - C \cdot \epsilon_{t-1}$. We conjecture that $0 \leq C < 1$ and, as before, that $B < 0$.

The GMM estimation results are collected in Table II.[20] Our parameter estimates turn out to be consistent with the maintained hypothesis: in both specifications ('level' and 'ratio'), the estimates of B are negative (-0.276 and -0.309, respectively) and significantly different from zero at the 1% significance level. Furthermore, the estimates of C are positive and smaller than one (respective values of 0.398 and 0.546,) and significantly different from zero at any significance level, suggesting that the RRA slope dynamically captures market heterogeneity. Lastly, testing the over-identifying restrictions using the J statistic indicates that both specifications cannot be rejected at the 5% significance level.

5. Conclusions

Using *asset market data*, we attempted to shed light on investors' risk preferences heterogeneity. Theoretical results recently derived by Benninga and Mayshar

[19] Taking $t - 1$ to be the same for β and *Spread* allows this iterative substitution.

[20] Recall that we have detrended β and *Spread* by applying the Hodrick and Prescott (1997) filter, and standardized the series to have a standard deviation of one.

(2000) show that preferences heterogeneity determine the slope of the representative consumer's RRA function. To test this theory, we estimated RRA functions, exploiting theoretical relations between investors' preferences, option-implied PDFs, and index-implied actual PDFs.

The RRA functions were estimated using European call option prices written on the S&P500 index, and sampled on a monthly basis during the period December 1987 to December 1995. The slopes of these RRA functions were then compared to a prevalent proxy for investors' heterogeneity, namely, the spread between auction and market Treasury bill yields.

Our statistical analysis yielded the following results: (i) The aggregate RRA functions are decreasing, in line with Benninga and Mayshar's (2000) prediction; and (ii) The constructed risk preferences heterogeneity series is positively correlated in a static, as well as a dynamic setup, with the spread between auction and market Treasury bill yields.

To sum, our results provide corroborative empirical support for Benninga and Mayshar's (2000) suggested way of inferring the usually unobserved risk-preferences heterogeneity from *aggregate* market data.

References

Ait-Sahalia, Y. and Lo, A. W. (1998) Nonparametric estimation of state-price densities implicit in financial asset prices, *J. Financ.* **53**(2), 499–547.

Ait-Sahalia, Y. and Lo, A. W. (2000) Nonparametric risk management and implied risk aversion, *J. Economet.* **94**, 9–51.

Bagwell, L. S. (1992) Dutch auction repurchases: An analysis of shareholder heterogeneity, *J. Financ.* **47**(1), 71–105.

Bahra, B. (1997) Implied risk-neutral probability density functions from option prices: Theory and application, Working paper, Bank of England, London.

Bates, D. S. (1991) The crash of 87: Was it expected? The evidence from options markets, *J. Financ.* **46**(3), 1009–1044.

Benninga, S. and Mayshar, J. (2000) Heterogeneity and option pricing, *Rev. Derivat. Res.* **4**, 7–27.

Bick, A. (1990) On viable diffusion price processes of the market portfolio, *J. Financ.* **45**(2), 673–689.

Black, F. and Scholes, M. (1973) The valuation of options and corporate liabilities, *J. Pol. Econom.* **81**, 637–654.

Brav, A., Constantinides, G. M., and Geczy, C. C. (2000) Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, Working paper.

Breeden, D. and Litzenberger, R. (1978) Prices of state-contingent claims implicit in options prices, *J. Business* **51**, 621–651.

Camerer, C. (1995) Individual decision making, in John Kagel and Alvin Roth (eds.), *The Handbook of Experimental Economics*, Princeton University Press, Princeton, New Jersey.

Chow, G. (1984) Random and changing coefficient models, in Zvi Griliches and Michael D. Intriligtor (eds.), *Handbook of Econometrics* Vol. II, Elsevier Science Publishers B.V.

Christensen, B. J. and Prabhala, N. R. (1998) The relation between implied and realized volatility, *J. Financ. Econom.* **50**, 125–150.

Cox, J. C. and Ross, S. A. (1976) The valuation of options for alternative stochastic processes, *J. Financ. Econom.* **3**, 145–166.

Cuoco, D. and Zapatero, F. (2000) On the recoverability of preferences and beliefs, *Rev. Financ. Stud.* **13**(2), 417–431.

Dybvig, P. H. and Rogers, L. C. G. (1997) Recovery of preferences from observed wealth in a single realization, *Rev. Financ. Stud.* **10**(1), 151–174.

Efron, B. (1979) Bootstrap methods: Another look at the jackknife, *Ann. Statist.* **7**, 1–26.

Hall, P. (1994) Methodology and theory for the bootstrap, in Robert F. Engle and Daniel L. McFadden (eds.), *Handbook of Econometrics*, Vol. IV, Elsevier Science Publishers B.V.

Hansen, L. P. (1982) Large sample properties of generalized method of moments estimators, *Econometrica* **50**, 1029–1054.

He, H. and Leland, H. (1993) On equilibrium asset price processes, *Rev. Financ. Stud.* **6**(3), 593–617.

Heaton, J. and Lucas, D. J. (1996) Evaluating the effects of incomplete markets on risk sharing and asset pricing, *J. Pol. Econom.* **104**, 443–487.

Hodrick, R. J. and Prescott, E. C. (1997) Postwar US business cycles: An empirical investigation, *J. Money, Credit and Banking* **29**, 1–16.

Jackwerth, J. C. (1999) Option-implied risk-neutral distributions and implied binomial trees: A literature review, *J. Derivat.* 66–82.

Jackwerth, J. C. (2000) Recovering risk aversion from option prices and realized returns, *Rev. Financ. Stud.* **13**(2), 433–451.

Jackwerth, J. C. and Rubinstein, M. (1996) Recovering probability distributions from option prices, *J. Financ.* **51**, 1611–1631.

Kandel, S., Sarig, O., and Wohl, A. (1999) The demand for stocks: An analysis of IPO auctions, *The Rev. Financ. Stud.* **12**(2), 227–247.

Posner, S. and Milevsky, M. (1998) Valuing exotic options by approximating the SPD with higher moments, *J. Financ. Engineer.* **7**(2), 109–125.

Rubinstein, M. (1994) Implied binomial trees, *J. Financ.* **49**(3), 771–818.

Rubinstein, M. (1998) Edgeworth binomial trees, *J. Derivat.* **5**(3), 20–27.

Shimko, D. (1993) Bounds on probability, *Risk* **6**, 33–37.

Soderlind, P. and Svensson, L. (1997) New techniques to extract market expectations from financial instruments, *J. Monet. Econom.* **40**, 383–429.

Storesletten, K., Telmer, C. I., and Yaron, A. (2000) The welfare cost of business cycle revisited: Finite lives and cyclical variation in idiosyncratic risk, NBER Working paper No. 8040, National Bureau of Economic Research, Cambridge, MA.

Wang, J. (1996) The term structure of interest rates in a pure exchange economy with heterogeneous investors, *J. Financ. Econom.* **41**, 75–110.

Wang, S. (1993) The local recoverability of risk aversion and intertempoal substitution, *J. Econom. Theory* **59**, 333–363.