INTERNATIONAL ECONOMIC REVIEW
Vol. 59, No. 2, May 2018 DOI: 10.1111/iere.12279

MULTIATTRIBUTE UTILITY THEORY, INTERTEMPORAL UTILITY, AND CORRELATION AVERSION*

BY STEFFEN ANDERSEN, GLENN W. HARRISON, MORTEN I. LAU, AND E. ELISABET RUTSTRÖM [1]

Copenhagen Business School, Denmark; Georgia State University, U.S.A., and University of Cape Town, South Africa; Copenhagen Business School, Denmark, and Durham University, U.K.; Georgia State University, U.S.A., Örebro University, Sweden, Stockholm School of Economics, Sweden, and University of Cape Town, South Africa

Convenient assumptions about qualitative properties of the intertemporal utility function have generated counterintuitive implications for the relationship between atemporal risk aversion and the intertemporal elasticity of substitution. If the intertemporal utility function is additively separable, then the latter two concepts are the inverse of each other. We review a theoretical specification with a long lineage in the literature on multi-attribute utility and use this theoretical structure to guide the design of a series of experiments that allow us to identify and estimate intertemporal correlation aversion. Our results show that subjects are correlation averse over lotteries with intertemporal income profiles.

1. INTRODUCTION

Convenient assumptions about qualitative properties of the intertemporal utility function have generated counterintuitive implications for the relationship between atemporal risk aversion and the intertemporal elasticity of substitution. If the intertemporal utility function is additively separable, then these two concepts are the inverse of each other. This is no technical side issue: Untangling it is central to the general understanding of savings behavior (e.g., Hall, 1988), the analysis of insurance decisions by poor households in developing countries (e.g., Townsend, 1994), and the behavior of asset prices over time (e.g., Hansen and Singleton, 1983).

Our contribution is to demonstrate that weakly separable, nonadditive representations of preferences, built from a well-known theoretical foundation of intertemporal utility, can be estimated from controlled choices in a field experiment with adult Danes. Moreover, we show how one can incorporate different models of risk preferences that have empirical support from the modern experimental literature. Our approach is flexible in this regard, allowing one to examine popular modern alternatives to Expected Utility Theory (EUT), such as Rank Dependent Utility (RDU), for example. We can also easily examine popular behavioral alternatives to exponential discounting.

The core concept we investigate is known as *correlation aversion*. It arises from theoretical deviations from additively separable intertemporal utility functions, such as the deviations we employ here. Define the lottery α as a 50:50 mixture of $\{x, Y\}$ and $\{X, y\}$ and the lottery θ as a 50:50 mixture of $\{x, y\}$ and $\{X, Y\}$, where $X > x$ and $Y > y$. So α is a 50:50 mixture of both bad and good outcomes in time t and $t + \tau$, and θ is a 50:50 mixture of only bad outcomes or only

*Manuscript received April 2014; revised September 2016.

[1] We thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746, the Danish Social Science Research Council for research support under project 275-08-0289, and the Carlsberg Foundation under grant 2008-01-0410. We also thank two referees, the editor, Antoine Bommier, Anke Leroux, Harris Schlesinger, and seminar participants for many useful comments. Additional appendices are available in *CEAR Working Paper 2011-03* available at http://cear.gsu.edu/papers. Please address correspondence to: Glenn W. Harrison, Department of Risk Management and Insurance, Robinson College of Business, P.O. Box 40036, Atlanta, GA 30302-4036, U.S.A. E-mail: *gharrison@gsu.edu*.

good outcomes in the two time periods. These lotteries α and θ are defined over all possible "good" and "bad" outcomes. If the individual is indifferent between α and θ, we say that he is neutral to intertemporally correlated payoffs in the two time periods. If the individual prefers α to θ, we say that he is averse to intertemporally correlated payoffs: It is better to have a given chance of being lucky in one of the two periods than to have the same chance of being very unlucky or very lucky in both periods. The correlation averse individual prefers to have nonextreme payoffs *across* periods, just as the risk averse individual prefers to have nonextreme payoffs *within* periods. One can also view the correlation averse individual as preferring to avoid correlation-increasing transformations of payoffs in different periods.

Keeney (1971, 1972, 1973) develop the essential technical concepts of conditional utility functions and conditional risk aversion for two or more attributes. Richard (1975) formally states the main results in terms of multivariate risk aversion, to be contrasted with univariate risk aversion. Epstein and Tanny (1980) sharply define the specific concept of correlation aversion.[2] Since we interpret the two attributes as referring to different time periods, we refer to intertemporal risk aversion or intertemporal correlation aversion.[3] There are direct parallels in the older literature on multiattribute utility (Fishburn, 1965; Pollack, 1967; Keeney, 1971, 1972; Harvey, 1993). There are also parallels in the older literature on multivariate risk aversion (Kihlstrom and Mirman, 1974; Rothblum, 1975; Duncan, 1977; Karni, 1979), as demonstrated by Eeckhoudt et al. (2007) and Dorfleitner and Krapp (2007). And the concept of correlation aversion provides insight into important properties of univariate risk aversion, such as properness, prudence, and temperance (Denuit et al., 2010).

In Section 2, we review a simple theoretical specification, which has a long lineage in the literature on multiattribute utility, and demonstrate the critical role of intertemporal risk aversion or intertemporal correlation aversion. In Section 3, we use this theoretical structure to guide the design of a series of experiments that allows us to identify the core parameters of the latent structural models. We also discuss our specific experiments, conducted throughout Denmark using a representative sample of the adult Danish population. In Section 4, we review econometric models used to estimate the core parameters of the models.

Section 5 contains results, and *provides the first estimates of the extent of intertemporal correlation aversion.* Our results show that subjects are indeed correlation averse over lotteries with intertemporal income profiles and that the convenient additive specification of the intertemporal utility function is not an appropriate representation of preferences over time. We also show that *intertemporal correlation aversion contributes significantly to the overall risk premium that characterizes risky intertemporal profiles.* Finally, we demonstrate the flexibility of our approach by considering the "modular" effect on estimates of correlation aversion of allowing for popular specifications of nonexponential discounting and non-EUT risk preferences. Section 6 concludes.

Additional details are relegated to online appendices. Online Appendix A documents the experimental instructions; online Appendix B documents the experimental design and procedures; and online Appendix C provides detailed estimates allowing for observable heterogeneity in demographics and experimental procedures.

2. THEORY

We consider first an intertemporal decision model with weakly separable preferences under EUT. The intertemporal utility function at time $t = 0$ is written as

$$U(X_1, X_2, \dots, X_n) = E\left[\varphi\left(\sum_{t=1}^{n}(1/(1+\delta)^t)u(x_t)\right)\right], \tag{1}$$

[2] Keeney and Raiffa (1976, chapter 5) remains a superb exposition of the basic ideas and theorems.

[3] Everything we say about intertemporal risk aversion or intertemporal correlation aversion applies symmetrically to behavior that exhibits intertemporal risk loving or intertemporal correlation loving.

where $u_t \in X_t,\ u(x_t)$ is the atemporal utility of money at time t, and δ is the exponential discount rate. For now, let φ be an identity function; we return to it momentarily. The decision tasks in our experiments provide lotteries with payments at two different points in time, where the time horizon between sooner and later payments varied between 2 weeks and 12 months. We can simplify the model in (1) and consider decisions that involve payments at two different points in time. Assuming $u(0) = 0$, the intertemporal utility function is specified as

$$U(X_t, X_{t+\tau}), = E\left[\varphi\left(D_t u(x_t) + D_{t+\tau} u(x_{t+\tau})\right)\right], \tag{2}$$

where $D_t = 1/(1+\delta)^t$ is the discounting function with a constant discount rate δ.

Let the atemporal utility function be the constant relative risk aversion (CRRA) specification:

$$u(x_t) = x_t^{1-r}/(1-r), \tag{3}$$

for $r \neq 1$, where r is the CRRA coefficient, assumed for simplicity to be the same for periods t and $t+\tau$. With this functional form, $r = 0$ denotes risk neutral behavior, $r > 0$ denotes risk aversion, and $r < 0$ denotes risk seeking behavior, all defined over atemporal trade-offs in t or $t+\tau$.

Given the popularity of the CRRA function in the microeconomic and macroeconomic literature, it is natural to consider this alternative structural specification of the intertemporal utility function φ:

$$U(X_t, X_{t+\tau}) = E[(D_t u(x_t) + D_{t+\tau} u(x_{t+\tau}))^{(1-\eta)}/(1-\eta)] = E[\xi(x_t, x_{t+\tau})^{(1-\eta)}/(1-\eta)], \tag{4}$$

where η is the intertemporal relative risk aversion parameter ($\eta \neq 1$), and the expression for the weighted sum of atemporal utility flows, $\xi(x_t,\ x_{t+\tau})$, is useful below.[4] The intertemporal utility function is separable but not additive when $\eta \neq 0$ and collapses to $E[\xi(x_t,\ x_{t+\tau})]$ when there is intertemporal risk neutrality at $\eta = 0$.[5] We estimate the structural parameters in this model and discuss results in Section 5.

If the intertemporal utility function in (4) is additively separable, then the inverse of the intertemporal elasticity of substitution is equal to the coefficient of atemporal risk aversion. This assumption is one of convenience and is popular in models of intertemporal choice. The linear specification of intertemporal utility is then equal to a weighted sum of atemporal utility flows, where the weights are determined by the discount rate.

[4] See Epstein (1983), Epstein and Zin (1989), Farmer (1990), Weil (1990), Bommier (2006, 2007, 2013), Bommier and Rochet (2006), Epstein et al. (2014), and Bommier et al. (2017) for discussion of alternative specifications of the intertemporal utility function and implications for intertemporal risk aversion. We are not concerned here with preferences for the timing of resolution of risk, and all risk is resolved immediately in our experiments. We do not see this as a serious restriction for this initial approach, in light of the comment by Epstein et al. (2014, p. 2687ff.) that assumptions about preferences for early resolution may not be appealing. The specification in (4) can generate negative discount rates unless we assume that $r < 1$ and $\eta < 1$ to ensure that $\delta > 0$. This assumption is empirically innocuous, as we demonstrate later with estimated parameter values. Andersen et al. (2008a) add background consumption ω to a CRRA specification of atemporal utility and apply the function $u(x_t) = (\omega + x_t)^{1-r}/(1-r)$, where $u(\omega) > 0$. They show that this function is well behaved when the intertemporal utility function is additively separable, in the sense that the addition of background consumption is a sufficient condition to avoid negative discount rates. However, adding background consumption to experimental income is problematic in the estimation of the nonadditive function in (4) because one has to specify the atemporal utility of background consumption at every point in time and not only when the sooner and later payments are made.

[5] See Deaton and Muellbauer (1980, p. 137) for a discussion of strong separability and additive preferences. Tsetlin and Winkler (2009) provide a rich, general characterization of the class of utility functions that allow correlation aversion in the form of mixtures of exponential functions.

Following the exposition of Bommier (2007), we can define a number of important concepts using this structure. The marginal rate of substitution between money in periods t and $t+\tau$ can be defined as

$$MRS_{t,t+\tau} = (\partial U/\partial x_t)/(\partial U/\partial x_{t+\tau}). \tag{5}$$

The coefficient of relative risk aversion in period t can be defined by

$$RRA_t = -x_t[\partial^2 U/(\partial x_t)^2]/(\partial U/\partial x_t). \tag{6}$$

The (direct) elasticity of substitution between money in periods t and $t+\tau$ is

$$\sigma_{t,t+\tau} = \left\{1/x_t(\partial U/\partial x_t) + 1/x_{t+\tau}(\partial U/\partial x_{t+\tau})\right\} / \{a+b+c\}, \tag{7}$$

where $a = -[\partial^2 U/(\partial x_t)^2]/(\partial U/\partial x_t)^2$, $b = 2[\partial^2 U/(\partial x_t \partial x_{t+\tau})]/[(\partial U/\partial x_t)(\partial U/\partial x_{t+\tau})]$, and $c = -[\partial^2 U/(\partial x_{t+\tau})^2]/(\partial U/\partial x_{t+\tau})^2$. Finally, a coefficient of correlation aversion with respect to money flows in periods t and $t+\tau$ can be defined as

$$\rho_{t,t+\tau} = -2[\partial^2 U/(\partial x_t \partial x_{t+\tau})]/[\partial U/\partial x_t + \partial U/\partial x_{t+\tau}]. \tag{8}$$

Clearly, $\rho_{t,t+\tau} \geq 0$ if $\partial^2 U/\partial x_t \partial x_{t+\tau} \leq 0$, since $\partial U/\partial x_t \geq 0$ and $\partial U/\partial x_{t+\tau} \geq 0$, directly connecting this coefficient to the definition of correlation aversion under EUT proposed by Richard (1975).[6] The coefficient of correlation aversion is just one tractable way to measure the underlying concept of intertemporal risk aversion or correlation aversion, just as coefficients of absolute or relative risk aversion are just tractable ways to measure the concept of atemporal risk aversion.

With these concepts defined, there is a remarkable relationship between them noted by Bommier (2007, Proposition 1):

$$1/\sigma_{t,t+\tau}(1+MRS_{t,t+\tau}x_t/x_{t+\tau}) = (RRA_t + x_t/x_{t+\tau}MRS_{t,t+\tau}RRA_{t+\tau}) - \rho_{t,t+\tau}x_t(1+MRS_{t,t+\tau}). \tag{9}$$

From (9) we see formally that correlation aversion breaks the nexus between the intertemporal elasticity of substitution and atemporal relative risk aversion.

Given the parametric structure we have assumed, we can further state the marginal rate of substitution between money in periods t and $t+\tau$ using (5) as

$$MRS_{t,t+\tau} = D_t x_t^{-r} / \left(D_{t+\tau} x_{t+\tau}^{-r}\right), \tag{5$'$}$$

the relative risk aversion in period t using (6) as

$$RRA_t = (\eta/[\xi(x_t, x_{t+\tau})])D_t x_t^{1-r} + r, \tag{6$'$}$$

the (direct) elasticity of substitution between money in periods t and $t+\tau$ using (7) as

$$\sigma_{t,t+\tau} = 1/r, \tag{7$'$}$$

[6] The relations between correlation aversion, atemporal risk aversion, and the intertemporal elasticity of substitution are not limited to EUT. Bommier (2005, §5) provides a general graphical interpretation that does not rely on EUT, and we consider RDU risk preferences later.

and finally the coefficient of correlation aversion with respect to money flows in periods t and $t+\tau$ using (8) as

$$\rho_{t,t+\tau} = \left\{(2\eta/[\xi(x_t, x_{t+\tau})])D_{t+\tau}x_{t+\tau}^{-r}D_t x_t^{-r}\right\} / \left\{D_{t+\tau}x_{t+\tau}^{-r} + D_t x_t^{-r}\right\}. \tag{8'}$$

Hence $\rho_{t,t+\tau}$ is positive (negative) when η is positive (negative). The specific functional forms for these concepts will vary with the parametric assumption assumed, of course.

To elicit intertemporal risk aversion one would have to present subjects with choices over lotteries that have different income profiles over time.[7] Proper identification of intertemporal risk aversion (η) thus requires that one control for atemporal risk aversion (r) and the individual discount rate (δ). All three parameters are intrinsically, conceptually connected as a matter of theory, unless one makes strong assumptions otherwise. Our experimental design and econometric logic follow from this theoretical point.

3. EXPERIMENTS

The experimental procedures we adopt are a simple extension of those employed by Andersen et al. (2008a, 2013, 2014). Online Appendix B reviews the basic procedures, and we focus here on the extensions.

One task elicited atemporal risk attitudes for lotteries payable today as a vehicle for inferring the concavity of the atemporal utility function. Another task elicited time preferences over nonstochastic amounts of money payable at different times: in general, a smaller, sooner amount and a larger, later amount. In some cases the sooner amount was paid out today, and in some cases it was paid out in the future.[8]

A third task, new to this design, elicited intertemporal risk attitudes by asking subjects to make a series of choices over risky profiles of outcomes that are paid out at different points in time.[9] For example, lottery A might give the individual a 10% chance of receiving a larger amount L_t at time t and a smaller amount $S_{t+\tau}$ at time $t+\tau$, (L_t, $S_{t+\tau}$) and a 90% chance of

[7] Keeney (1977) illustrates this in a one-on-one, conversational elicitation with a decision maker. These lottery comparisons are used to justify the use of a correlation-neutral multi-attribute utility function and not to elicit the degree of correlation aversion. Delquie and Luo (1997) show how one can *test for* the tractable correlation-neutral class of functions using only two indifference judgments, rather than the set of lottery comparisons that are needed to *estimate the extent of* multi-attribute risk aversion.

[8] One maintained assumption is that subjects do not have access to perfect capital markets with which to arbitrage the investment opportunities offered by the experimenter. In that case one should consider the effect of censoring of responses in line with borrowing and savings interest rates that the subject has available outside the experiment, as in Coller and Williams (1999) and Harrison et al. (2002). Even when markets are perfect, understanding the quantitative importance of correlation aversion can be important for understanding life-cycle behavior (Bommier, 2013).

[9] Other experiments with real incentives have examined choices over time-delayed risk, but where the time delay is the same for all outcomes within a given lottery (Noussair and Wu, 2006; Baucells and Heukamp, 2010; Coble and Lusk, 2010; Abdellaoui et al., 2011). This is an important domain of choice to examine, but is not intended to address the issue of correlation aversion. The only previous experiments with real incentives that we know of that explicitly test for *atemporal* multiattribute risk aversion are due to von Winterfeldt (1980). He considers lottery choices of 18 subjects defined over 36 consumption bundles of gallons of gasoline and pounds of ground beef. For example, the lottery α might be a 50% chance of (16 gallons of gas and 10 pounds of ground beef) and a 50% chance of (no gas and no beef), and the lottery β defined as a 50% chance of (10 pounds of ground beef) and a 50% chance of (16 gallons of gas). He used three different methods of eliciting preference: direct statement of preference, including the option of indifference; a rating normalized between 0 and 100; and cash equivalents elicited using an incentive-compatible Becker et al. (1964) procedure with buying prices elicited from between USD0 and USD20. It appears that only the choices in the last elicitation procedure were played out for real, although the exposition is not completely clear (contrast the top of page 73 and the top of page 70). There were some response mode effects of the kind now known as "preference reversals" and some violations of nonsatiation. But the general conclusion is that "[m]ultiattribute risk aversion showed very clearly for all except two subjects" (p. 80). He also concluded that "[m]ultiattribute risk aversion appeared unrelated to ... single attribute risk aversion" (p. 81). Pliskin et al. (1980) and Payne et al. (1984) report experiments with *hypothetical* incentives that test for multi-attribute risk neutrality and respectively report results that support and reject that characterization.

receiving the smaller amount S_t at time t and the larger amount $L_{t+\tau}$ at time $t+\tau$ $(S_t, L_{t+\tau})$. Lottery B might give the individual a 10% chance of receiving L_t and $L_{t+\tau}$ and a 90% chance of receiving S_t and $S_{t+\tau}$. The subject picks A or B. We gave subjects 40 choices of this type, in four sets of 10 choices with prizes $S_t = S_{t+\tau}$ and $L_t = L_{t+\tau}$ in each set. Each set of 10 choices offered prizes with probability $p(L_t, S_{t+\tau}) = p(L_t, L_{t+\tau})$ starting at 0.1 and increasing by 0.1 until the last choice is between two degenerate lotteries. In this example, the last choice is a choice between receiving lottery A that pays $(L_t, S_{t+\tau})$ and lottery B that pays $(L_t, L_{t+\tau})$: Lottery B dominates lottery A and is a test of monotonic preferences.

We present each choice to the subject as a "pie chart" showing prizes and probabilities. We use the following four sets of larger (L) amounts and smaller (S) amounts: (A1: 3850, 100), (A2: 2000, 250), (A3: 2000, 75), and (A4: 4500, 50). Subjects were randomly assigned to one of these four prize sets. Each subject is presented with choices over four time horizons between 2 weeks and 1 year in the discount rate tasks, and the same four time horizons were applied in the intertemporal risk aversion tasks. We also varied the time delay to the early payments on a between-subjects basis, and this treatment followed from the discount rate tasks. If there was no time delay to the early payment in the discount rate tasks, then the early payments in the intertemporal risk aversion tasks were paid out immediately and similarly if the delay to the early payment option in the discount rate tasks was 1 month.

Between September 28 and October 22, 2009, we conducted experiments with 413 Danes. The sample was drawn to be representative of the adult population as of January 1, 2009, using sampling procedures that are virtually identical to those documented at length in Andersen et al. (2008a, 2013, 2014).

4. ECONOMETRICS

We consider nonadditive separable specifications of the intertemporal utility function and estimate the coefficient of intertemporal risk aversion. Equation (4) implies that the expected utility of option A in the intertemporal risk aversion task is given by

$$PEU_A = p(L_t, S_{t+\tau}) \times [\xi(L_t, S_{t+\tau})^{(1-\eta)}/(1-\eta)] + p(S_t, L_{t+\tau}) \times [\xi(S_t, L_{t+\tau})^{(1-\eta)}/(1-\eta)] \tag{10}$$

and the expected utility of option B is given by

$$PEU_B = p(L_t, L_{t+\tau}) \times [\xi(L_t, L_{t+\tau})^{(1-\eta)}/(1-\eta)] + p(S_t, S_{t+\tau}) \times [\xi(S_t, S_{t+\tau})^{(1-\eta)}/(1-\eta)], \tag{11}$$

where $p(L_t, L_{t+\tau})$ is the probability of receiving L_t in period t and $L_{t+\tau}$ in period $t+\tau$.[10] We can write out the likelihood function for the choices that the subjects made and jointly estimate the risk parameter r, the discount rate parameter δ, and the intertemporal risk parameter η. We employ the contextual error specification proposed by Wilcox (2011), and the latent index is specified by

$$\nabla PEU = ((PEU_B - PEU_A)/\lambda)/\mu^{SDR}, \tag{12}$$

where μ^{SDR} is a noise parameter for the ("stochastic discounting") intertemporal risk aversion choices. The normalizing term λ is defined as the maximum intertemporal utility over all prize profiles in this lottery pair $(L_t, L_{t+\tau})$ minus the minimum utility over all prize profiles in this lottery pair $(S_t, S_{t+\tau})$. The maximum intertemporal utility over all prize profiles in the lottery pair is $[D_t u(L_t) + D_{t+\tau} u(L_{t+\tau})]^{(1-\eta)}/(1-\eta)$, and the minimum intertemporal utility is $[D_t u(S_t) + D_{t+\tau} u(S_{t+\tau})]^{(1-\eta)}/(1-\eta)$.

[10] We assume here that the atemporal utility function is stable over time and is perceived ex ante to be stable over time. Direct evidence for the former proposition is provided by Andersen et al. (2008b), who examine the temporal stability of risk attitudes in the Danish population. The second proposition is a more delicate matter: Even if utility functions are stable over time, they may not be subjectively perceived to be, and that is what matters for us to assume that the same r that appears in (3) appears in (10) and (11).

The likelihood of the intertemporal risk aversion responses, conditional on the specification of intertemporal utility being true, depends on the estimates of r, δ, η, μ^{SDR}, μ^{RA}, and μ^{DR}, given the observed choices, where μ^{RA} is a noise parameter for the atemporal risk aversion choices and μ^{DR} is a noise parameter for the discount rate choices. The conditional log-likelihood is

$$\ln L(r, \delta, \eta, \mu^{RA}, \mu^{DR}, \mu^{SDR}; c, C) = \sum_i [(\ln \Lambda(\nabla PEU) \times I(c_i = 1)) + (\ln(1 - \Lambda(\nabla PEU)) \times I(c_i = -1))], \tag{13}$$

where $c_i = 1$ (–1) denotes the choice of option B (A) in intertemporal risk aversion task i, $\mathbf{C}$ is a vector of individual characteristics, and Λ is the cumulative logistic function.

The joint likelihood of the atemporal risk aversion, discount rate, and intertemporal risk aversion responses can then be written as

$$\ln L(r, \delta, \eta, \mu^{RA}, \mu^{DR}, \mu^{SDR}; c, C) = \ln L^{RA} + \ln L^{DR} + \ln L^{SDR}, \tag{14}$$

where L^{RA} is the conditional log-likelihood of the atemporal risk aversion responses, L^{DR} is the conditional log-likelihood of the discount rate responses, and L^{SDR} is defined by (13).

The nature of this joint likelihood function is matched by our experimental design. Ignoring the objective parameters of the tasks, the lottery choices over stochastic lotteries paid out today (RA) depend on r, η, and μ^{RA}; the discounting tasks over nonstochastic outcomes paid out today or sometime in the future (DR) depend on r, μ^{RA}, δ, and μ^{DR}; and the discounting tasks over stochastic outcomes paid out today or some time in the future (*SDR*) depend on r, μ^{RA}, δ, μ^{DR}, η, and μ^{SDR}. Putting the behavioral error terms aside, if we were to try to estimate r and δ using either the RA or the DR choices, we would be unable to identify both parameters. Similarly, if we were to try to estimate r, δ, and η using only two of three tasks, we would face an identification problem. These identification problems are inherent to the theoretical definitions of the discount rate and correlation aversion and demand an experimental design that combines multiple types of choices and an econometric approach that recognizes the complete structural model. The general principle is joint estimation of all structural parameters so that uncertainty about the parameters defining the utility function propagates in a "full information" sense into the uncertainty about the parameters defining the discount function and the intertemporal utility function. Intuitively, if the experimenter only has a vague notion of what $u(\cdot)$ is, because of poor estimates of r, then she simply cannot make precise inferences about δ or η. Similarly, poor estimates of δ, even if r is estimated relatively precisely, imply that one cannot make precise inferences about η.

Our inferential procedure about correlation aversion does not rely on the use of EUT or the CRRA functional form. Nor does it rely on the use of the exponential discounting function. The method generalizes immediately to alternative multiattribute models of decision making under risk, such as those presented in Miyamota and Wakker (1996). It also extends to specifications that use alternative discounting functions, as illustrated in Section 5.4.

5. RESULTS

5.1. *Basic Results.* Table 1 reports maximum likelihood estimates of the specification with the CRRA atemporal utility function.[11] There is evidence of a concave atemporal utility function ($r > 0$), with r estimated to be 0.35. The discount rate is estimated to be 11.4% on an annualized basis.[12] The main novelty here is evidence of intertemporal risk aversion, with η estimated to

[11] Virtually identical results are obtained if one uses the Expo-Power atemporal utility function.

[12] The estimates of r and δ are different from those reported for the specification in Andersen et al. (2008a) that sampled the same adult Danish population in 2003, but assumed intertemporal risk neutrality. In that case the point estimate of r = RRA was 0.74, and the discount rate δ was estimated to be 10.1% (Table III, p. 601). If we impose the constraint

TABLE 1
MAXIMUM LIKELIHOOD ESTIMATES ASSUMING CRRA ATEMPORAL UTILITY FUNCTION
(N = 49,560 OBSERVATIONS, BASED ON 413 SUBJECTS)

Parameter	Point Estimate	Standard Error	p-Value	95% Confidence Interval	
A. Atemporal Utility Function					
r	0.35	0.037	<0.001	0.28	0.42
μ^{RA}	0.18	0.011	<0.001	0.16	0.20
B. Discounting Function					
δ	0.114	0.012	<0.001	0.092	0.137
μ^{DR}	0.04	0.003	<0.001	0.03	0.04
C. Intertemporal Utility Function					
η	0.32	0.044	<0.001	0.23	0.40
μ^{SDR}	0.18	0.010	<0.001	0.16	0.20
D. Implied Estimates					
RRA	0.45	0.030	<0.001	0.39	0.51
1/RRA	2.21	0.147	<0.001	1.92	2.50
IES	2.85	0.305	<0.001	2.62	3.45
IES − 1/RRA	0.65	0.186	<0.001	0.28	1.01

be 0.32 and statistically significantly different from 0. The implication is that there should be a difference between the inverse of RRA and the intertemporal elasticity of substitution IES (which is equal to $1/r$), and this is confirmed by the implied estimates in Panel D of Table 1. The IES is estimated to be 2.85, with a standard error of 0.31, and exceeds the inverse of RRA by 0.65.[13] This difference between the IES and the inverse of RRA is statistically significant, with a p-value less than 0.001.

We can derive the certainty equivalents for each lottery in option A and option B of the intertemporal risk aversion tasks using (4) and then evaluate the risk premia associated with different prize sets.[14] Option A pays $(L_t, S_{t+\tau})$ with probability $p(L_t, S_{t+\tau})$ and $(S_t, L_{t+\tau})$ with probability $(1–p(L_t, S_{t+\tau}))$. The decision tasks are designed such that $S_t = S_{t+\tau}$ and $L_t = L_{t+\tau}$. If we assume, for simplicity, that the discount rate is equal to zero, then the present value of option A is $L_t+S_{t+\tau}$ kroner. If we define the certainty equivalent as either $(L_t, S_{t+\tau})$ or $(S_t, L_{t+\tau})$, then the risk premium is equal to zero for the lotteries in option A. However, if we define the certainty equivalent as the same certain amount to be paid out in both time periods, then the certainty equivalent is

$$CE_A = \left[\left(L_t^{(1-r)} + S_{t+\tau}^{(1-r)}\right)/2\right]^{1/(1-r)}, \tag{15}$$

$\eta = 0$ in our analysis, the log-likelihood drops significantly from –26,361.2 to –26,568.5, with r estimated to be 0.55 and δ is equal to 7.7%.

13 We evaluate RRA in the special case where $\mathrm{MRS}_{t,\,t+\tau}\ (x_t/x_{t+\tau}) = 1$. In this case $\mathrm{RRA}_t = \mathrm{RRA}_{t+\tau} = \eta(1-r)/2 + r$.

14 The *matrix* of risk premia in the multi-attribute case is characterized by Duncan (1977) and Karni (1979). Kihlstrom and Mirman (1974; §2.2) derived a "directional risk premium" that takes on as many values as there are possible "directions," and so is also multi-valued. But they pointed out that their measure allowed unique comparisons of utility functions representing the same ordinal preferences. The general point is that one cannot define risk premia as simply as when only considering the single-attribute case of univariate risk aversion.

where CE_A is paid out in period t and in period $t + \tau$. We can then define the risk premium of the lotteries in option A as

$$RP_A = (L_t + S_{t+\tau}) - 2CE_A. \tag{16}$$

The subject prefers smoothing consumption over time if the atemporal utility function is concave, which is just to say that the risk premium is positive when $r > 0$. Using the estimates from Table 1, the estimated risk premium is 911 kroner for prize set A1, 288 kroner for prize set A2, 440 kroner for prize set A3, and 1193 kroner for prize set A4.

If we allow the discount rate to be positive, then the certainty equivalent for option A is

$$CE_A = [((1 - \eta)PEU_A)^{1/(1-\eta)}(1 - r)/(D_t + D_{t+\tau})]^{1/(1-r)}, \tag{15$'$}$$

where PEU_A is the expected utility of option A given by (10). The risk premium is then derived as

$$RP_A = p(L_t, S_{t+\tau}) \times \xi(L_t, S_{t+\tau}) + (1 - p(L_t, S_{t+\tau})) \times \xi(S_t, L_{t+\tau}) - (D_t + D_{t+\tau}) \times CE_A. \tag{16$'$}$$

The estimated risk premium for prize set A1 varies between 855 kroner when $p(L_t, S_{t+\tau}) = 1$ and 868 kroner when $p(L_t, S_{t+\tau}) = 0.1$. The estimated risk premium for A2 is 271 and 274 kroner for $p(L_t, S_{t+\tau})$ equal to 1 and 0.1, respectively. Similarly, the risk premium interval is [413; 419] for A3 and [1119; 1139] for A4. Hence, the estimated risk premium for option A falls slightly when we consider positive discount rates.

The lotteries in option B pay $(L_t, L_{t+\tau})$ with probability $p(L_t, L_{t+\tau})$ and $(S_t, S_{t+\tau})$ with probability $(1 - p(L_t, L_{t+\tau}))$. The certainty equivalent of option B is

$$CE_B = [p(L_t, L_{t+\tau}) \times L_t^{(1-r)(1-\eta)} + (1 - p(L_t, L_{t+\tau})) \times S_t^{(1-r)(1-\eta)}]^{1/(1-r)(1-\eta)}, \tag{15$''$}$$

where CE_B is again a certain amount that is paid out in both period t and period $t + \tau$. This definition of certainty equivalence implies that CE_B is independent of the discount rate and is equal to L_t if $p(L_t, L_{t+\tau}) = 1$ and equal to S_t if $p(L_t, L_{t+\tau}) = 0$. The risk premium is then

$$RP_B = p(L_t, L_{t+\tau}) \times L_t + (1 - p(L_t, L_{t+\tau})) \times S_t - CE_B, \tag{16$''$}$$

which is equal to 0 if $p(L_t, L_{t+\tau})$ is equal to 0 or 1.

Figure 1 displays the estimated risk premium as a function of $p(L_t, L_{t+\tau})$ for each of the four prize sets in option B of the intertemporal risk aversion task. The solid line is based on the estimated parameter values for r and η in Table 1, and the dashed line is based on a constrained model in which we assume that η is equal to 0 and r is kept constant and equal to 0.35. Hence the risk premium when $\eta = 0$ and the decision maker is assumed to be correlation neutral (CN) derives entirely from the atemporal risk aversion r of the decision maker. When η and r are positive and the decision maker is correlation averse (CA) as well as being atemporally risk averse, the risk premium derives from both types of risk aversion. The results show that *intertemporal risk aversion accounts for a substantial amount of the estimated risk premium*. For example, the upper left panel shows that the risk premium for prize set B1 is equal to 1,522 kroner in the unconstrained model when $p(L_t, L_{t+\tau}) = 0.5$ and is equal to 910 kroner when the intertemporal risk aversion parameter η is constrained to be 0, so the difference of 612 kroner is due to correlation aversion.

5.2. *Observable Characteristics.* It is a simple matter to extend the econometric model to allow structural parameters to depend on observed demographics and experimental treatments. Online Appendix C does this and provides detailed estimates. The results imply considerable heterogeneity. Young subjects have lower estimated atemporal risk attitudes (r) than older

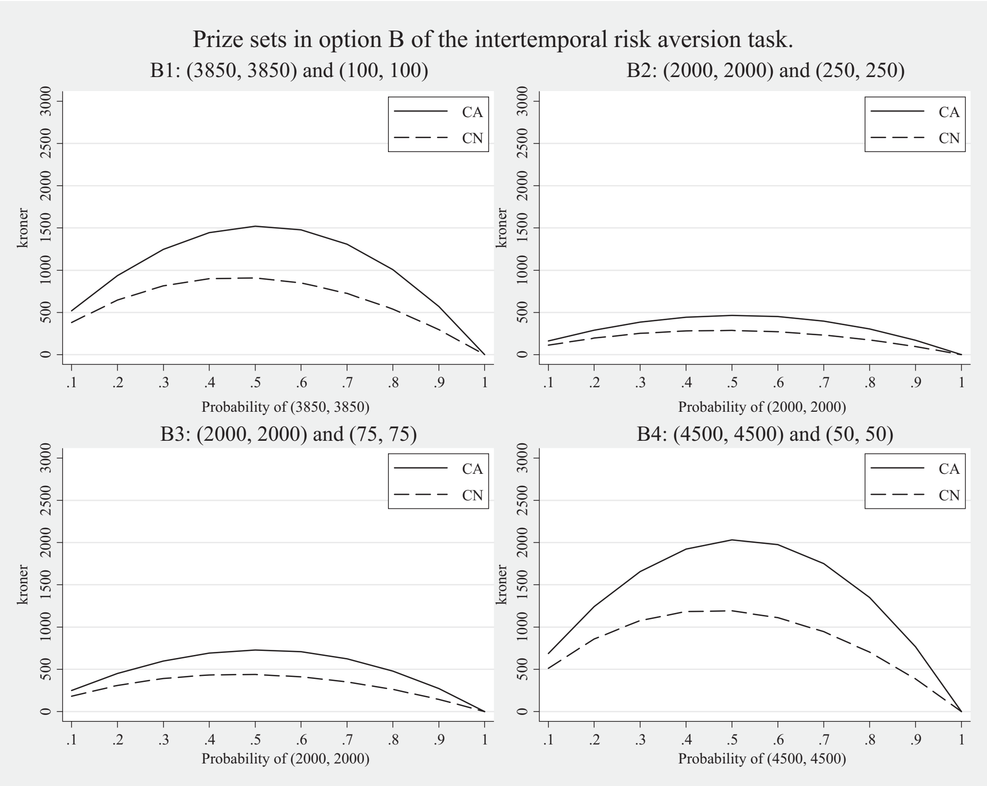

FIGURE 1

ESTIMATED RISK PREMIA

subjects, with an estimated marginal effect of –0.58 that is statistically significant with a p-value of 0.024. We find similar marginal effects for students and subjects with kids, and cannot reject the hypothesis that subjects with these characteristics have linear atemporal utility functions. The results also point to a negative correlation between relative risk aversion and income: Subjects with low income are significantly more risk averse than those with middle or high income, and subjects with high income are significantly less risk averse than those with middle income. It is noteworthy that there is *no significant gender effect on atemporal risk aversion*: Women exhibit greater risk aversion (+0.14) than men, but the 95% confidence interval on this estimate spans zero, and the p-value is only 0.24.

Despite the considerable variation in atemporal risk attitudes, we find that only one of the demographic characteristics is significantly correlated with individual discount rates (δ). Young subjects have higher discount rates than older subjects, and the marginal effect of 8.8% is significant with a p-value of 0.024.[15]

There is a significant age effect on intertemporal risk attitudes as measured by η. Young subjects are significantly more intertemporally risk averse than those above 40 years of age; the marginal effect is 0.25 with a standard error of 0.09 and a p-value of 0.006. We also find a gender effect, with women being more risk averse over lotteries with intertemporal payment

[15] One task characteristic is significantly correlated with discount rates: Information on interest rates in the discounting choices has a negative effect on inferred discount rates. The estimated effect is –4.3%, with a standard error of 2.3% and a p-value of 0.064.

profiles than men. The estimated η is 0.25 larger for women, and this estimate has a p-value of only 0.006.[16]

Table C2 in online Appendix C shows the implied estimates of relative risk aversion and the intertemporal elasticity of substitution. Younger subjects, students, and those with kids have linear atemporal utility functions, and the IES coefficient is effectively infinite for these subjects. We also infer significantly negative marginal effects on the RRA coefficient for these three types of subjects. The marginal effect on RRA for young subjects is –0.46, with a p-value of 0.079; it is –0.52 for subjects with kids (p-value of 0.066) and is –0.58 for students (p-value of 0.063). Finally, the results point to a significant positive marginal effect on RRA for subjects with low income and a significant negative marginal effect for subjects with high income. Hence, there is evidence that RRA declines with income, but we find no significant effect of income on the IES coefficient.

We also evaluate the *total effects* of several of the demographic characteristics on the estimated RRA and IES by estimating marginal effects without controls for other characteristics. We calculate total effects since many demographic characteristics co-vary in the population and therefore also in our sample. We find that women are more risk averse than men, with an estimated RRA of 0.45 for women and 0.33 for men. This difference in RRA between men and women is statistically significant, with a p-value of 0.026. However, we cannot reject the hypothesis that men and women have identical IES coefficients; the difference of 0.37 has a p-value of 0.740. There is an age effect on the two RRA and IES coefficients. The older age group has a higher RRA and lower IES than younger subjects, where the difference in RRA is significant with a p-value of 0.064, and the difference in IES has a p-value of 0.102. We cannot reject the hypothesis that the coefficients of RRA and IES are similar across income groups and education levels.

5.3. *The Coefficient of Correlation Aversion.* Figure 2 displays the predicted distribution of the coefficient of correlation aversion (CCA), evaluated using (8′) and the four prize sets in the intertemporal risk aversion tasks. The upper left panel shows predicted CCA values for the income profile with high sooner payments and low later payments. The predicted mean is 0.10 (÷ 1000), with a standard deviation of 0.07 (÷ 1000). There is clear evidence of correlation aversion in general, although roughly 5% of the sample exhibits correlation neutrality or correlation loving preferences. The lower left panel displays the predicted distribution for the income profile with low sooner payments and high later payments. If the individual discount rate is equal to 0, then the predicted values of CCA in the upper and lower left panels would be the same. The predicted mean of the individual discount rate for the sample is 0.10, with a standard deviation of 0.08, and we therefore see small differences between the estimated means and standard deviations across the two income profiles in the upper and lower left panels. The upper and lower right panels show the predicted distributions for the income profiles with two high payments and two low payments, respectively. The general pattern is the same as before, although the estimated means and standard deviations are lower (higher) for the income profile with two high (low) payments. Thus, we observe correlation aversion in general across the four income profiles in the intertemporal risk aversion tasks.[17]

[16] We find a significant effect of sex on relative risk aversion and discount rates when the intertemporal utility function is mis-specified to be additively separable: women appear to be significantly more risk averse and patient in monetary terms than men.

[17] We can also consider a more flexible specification of the intertemporal utility function and replace the η-parameter in (4) with a two-parameter version of the cumulative gamma distribution. There are very few a priori restrictions on the shape of the gamma distribution other than those of a cumulative density function. The estimated coefficients of r and δ are then different from before: The estimated IES is equal to 1.08 and the discount rate is estimated to be 1.3% on an annualized basis. However, we find that the estimated CCA is higher when the CRRA specification of (4) is replaced by the gamma distribution. We evaluate the CCA using the four prize sets in the intertemporal risk aversion tasks, as before, and find that CCA is 37.5 (÷ 1000) for the income profile with low sooner payments and low later payments, and 2.1 (÷ 1000) for the income profile with high sooner payments and high later payments. The estimated

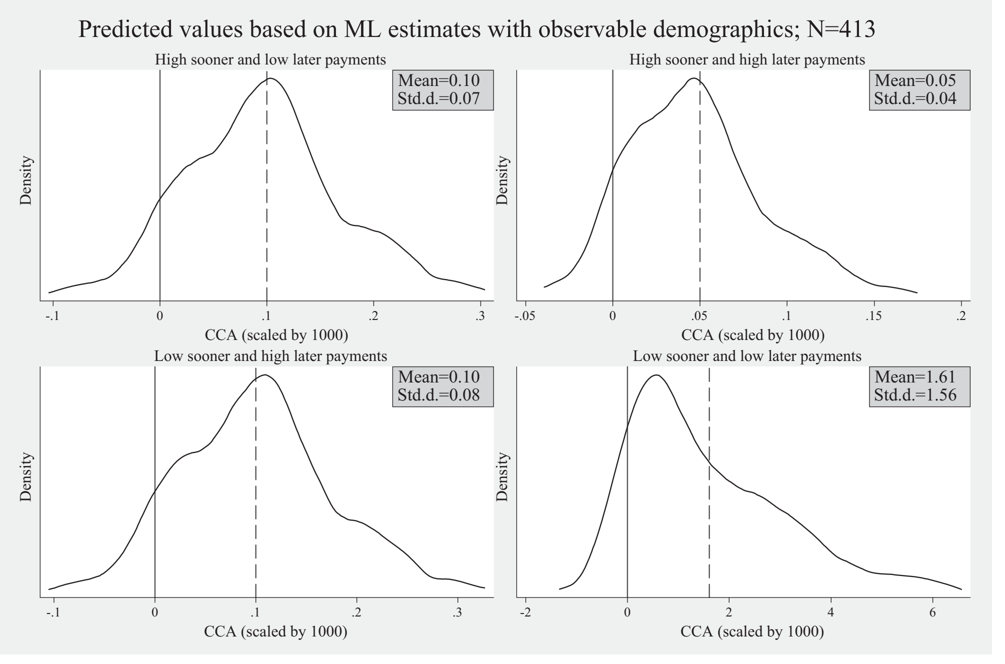

FIGURE 2

DISTRIBUTION OF THE COEFFICIENT OF CORRELATION AVERSION

5.4. *Nonexponential Discounting.* One of the advantages of our approach is that one can consider "modular" changes in the modeling of time preferences and risk preferences. This is desirable given the popularity of alternatives in the modern behavioral and experimental literature, surveyed by Starmer (2000) and Frederick et al. (2002).

Our objective here is to demonstrate that the concept of intertemporal correlation aversion does not depend on the use of the exponential discounting model. To illustrate the generality of the results, we consider the effect of using two popular alternative discounting models.[18] The exponential discounting model may be viewed as assuming a *constant variable utility cost* per time period of delay. The two alternatives we consider are the Quasi-Hyperbolic specification that allows for a *fixed utility cost* as well as a *constant* variable utility cost, and the Weibull specification that allows for a *nonconstant variable* utility cost.

The discount factor for the quasi-hyperbolic (QH) specification is defined as

$$D^{QH}(t) = 1 \quad \text{if } t = 0, \tag{17a}$$

$$D^{QH}(t) = \beta/(1+\delta)^t \quad \text{if } t > 0, \tag{17b}$$

where $\beta < 1$ implies quasi-hyperbolic discounting and $\beta = 1$ is exponential discounting. The defining characteristic of the QH specification is that the discount factor has a jump discontinuity at $t = 0$ and is thereafter the same as the exponential specification. The discount rate for the QH specification is the value of d^{QH} that solves $D^{\text{QH}} = 1/(1 + d^{\text{QH}})$, so it is

$$d^{QH}(t) = 1/[\beta/(1+\delta)^t]^{(1/t)} - 1, \tag{18}$$

CCA are significantly different from 0, and we continue to observe correlation aversion in general when the gamma function is used.

[18] There are many variants from the exponential model, and most are evaluated by Andersen et al. (2014) using a separable and additive intertemporal utility function.

for $t > 0$. Thus, for $\beta < 1$ we observe sharply declining discount rates in the very short run, and then discount rates asymptoting toward δ as the effect of the initial drop in the discount factor diminishes. The drop in the discount factor caused by β can be viewed as fixed utility cost of discounting anything relative to the present, since it does not vary with the horizon t once $t > 0$.

The QH model performs poorly in our model with intertemporal risk aversion, in the sense that the coefficient β is not significantly different from 1, which of course is the exponential case.[19] We estimate the value to be 1.003, with a 95% confidence interval between 0.989 and 1.018. The estimated values for r, δ, and η are virtually identical to those shown in Table 1 for the exponential specification, which is of course not surprising if $\beta \approx 1$.

The discount factor for the Weibull distribution from statistics[20] is defined as

$$D^W(t) = \exp(-\acute{r}t^{1/\acute{s}}) \tag{19}$$

for $\acute{r} > 0$ and $\acute{s} > 0$. For $\acute{s} = 1$ this collapses to the exponential specification, and hence the parameter $\acute{s}$ can be viewed as reflecting the "slowing down" or "speeding up" of time as subjectively perceived by the individual. This specification is due to Read (2001, p. 25, equation (16)). The discount rate at time t in the Weibull specification is then

$$d^W(t) = \exp(\acute{r}t^{(1-\acute{s})/\acute{s}}) - 1. \tag{20}$$

Thus one can again see the exponential emerge as a special case when $\acute{s} = 1$.

The Weibull model also performs poorly in our data, in the sense that the key parameter $\acute{s}$ is estimated to be 1.048, with a standard error of 0.141 and a 95% confidence interval between 0.77 and 1.32. This uncertainty in the estimate of $\acute{s}$ does translate into some uncertainty about discount rates in the short run, but not in an economically significant way. For horizons of 1 week the implied discount rate is 0.137, for 3 months it is 0.121, and for 1 year it is 0.113; the confidence intervals for these estimates are –0.013 $\leftrightarrow$ 0.288, 0.069 $\leftrightarrow$ 0.173, and 0.092 $\leftrightarrow$ 0.135, respectively. Again, the estimates of r, δ, and η are virtually identical to those shown in Table 1 for the exponential specification.

One implication of these nonexponential specifications, as noted by Backus et al. (2004, p. 328), is nonstationarity and time inconsistency. We do not want to rule that behavioral possibility out a priori, but others might want to.[21] If so, there exist general restrictions on the intertemporal utility function to ensure stationary preferences and time consistency, even when allowing for correlation aversion (e.g., see Bommier, 2013; Epstein et al., 2014; Bommier et al., 2017).

5.5. *Nonexpected Utility Theory.* Similarly, we can show that the concept of intertemporal correlation aversion does not depend on the use of the EUT model.[22] To concretely illustrate this, we consider the effect of using a popular alternative model due to Quiggin (1982), which

[19] The QH model performs poorly with these data even when one assumes intertemporal risk neutrality; see Andersen et al. (2014).

[20] Any probability density function $f(t)$ defined on $[0, \infty)$ can form the basis of a discounting function by taking the integral of $f(t)$ between t and ∞. Indeed, discounting functions are formally identical to the "survivor functions" that labor and health economists routinely estimate in duration models and are also known as "reliability functions" in the applied statistics literature on failure. Hence familiar and flexible families of probability density functions, such as the Gamma or Weibull, can be used to directly define discounting functions. This has the attraction of allowing the analyst to rely on a large literature in statistics on the properties of these functions for different inferential purposes.

[21] Indeed, we find it natural to think of there being a (statistical) mixture of data-generating processes at work, particularly in aggregate, representative agent models. One process might be time-consistent and another process might not be, as in Coller et al. (2012). To study such mixtures, one must be able to specify different discounting functions that allow such behavior.

[22] There have been extensions of multi-attribute utility theory to a wide range of non-EUT models, such as Fishburn (1984) and Miyamoto and Wakker (1996).

relaxes the Independence Axiom (IA). The RDU model posits probability weights based on some continuous function of the objective probabilities and then infers decisions weights from these probability weights. The probability and decision weights depend on the rank of the outcome, in a familiar manner, replacing the usual IA with a Comonotonic IA. If the atemporal expected utility function is

$$EU = [p(z) \times U(z)] + [p(Z) \times U(Z)], \tag{21}$$

then, if $Z > z$, we can rewrite the atemporal expected utility as

$$RDU = [(1 - \omega(p(Z)) \times U(z))] + [\omega(p(Z)) \times U(Z)] \tag{22}$$

for some probability weighting function $\omega(p)$. We use a general functional form proposed by Prelec (1998) that exhibits considerable flexibility:

$$\omega(p) = \exp\{-\zeta(-\ln p)^{\upsilon}\}, \tag{23}$$

defined for $0 < p \leq 1$, $\zeta > 0$, and $\upsilon > 0$.[23] Of course, EUT assumes the identity function $\omega(p) = p$, which is the case when $\zeta = \upsilon = 1$.

When the outcome is simply an amount of money, as in our atemporal lottery tasks, there is no complication calculating the rank to apply the RDU model. When the outcome consists of two time-dated amounts of money, as in our temporal lottery tasks, one has to be more careful. The natural quantity to base the rank on is then the present value of the atemporal utilities afforded by the two time-dated amounts of money. To see this explicitly, recall the expression for option A, referred to generically as lottery α in the definition of correlation aversion:

$$PEU_A = p(L_t, S_{t+\tau}) \times [\xi(L_t, S_{t+\tau})^{(1-\eta)}/(1-\eta)] + p(S_t, L_{t+\tau}) \times [\xi(S_t, L_{t+\tau})^{(1-\eta)}/(1-\eta)]. \tag{10}$$

Since the decision tasks are designed such that $S_t = S_{t+\tau}$ and $L_t = L_{t+\tau}$, we get that $\xi(L_{\mathrm{t}},S_{\mathrm{t}+\tau}) > \xi(S_{\mathrm{t}},L_{\mathrm{t}+\tau})$ for $\delta > 0$, and the rank-dependent utility, with just two outcomes, is

$$\begin{aligned} PRDU_A = {} & \omega(p(X, y) \times [\xi(L_t, S_{t+\tau})^{(1-\eta)}/(1-\eta)] + (1 - \omega(p(X, y))) \\ & \times [\xi(S_t, L_{t+\tau})^{(1-\eta)}/(1-\eta)]. \end{aligned} \tag{10$'$}$$

A similar construction applies for option B, and one can trivially identify the ranks on an a priori basis.

The estimates show evidence of greater correlation aversion when one allows for RDU preferences rather than EUT risk preferences. Figure 3 shows the estimated "S-shaped" probability weighting function and the decision weights implied for equi-probable reference lotteries (e.g., if there are four prizes each has probability ¼).[24] The estimated values of ζ and υ are 1.37 and 2.32, respectively, and one can easily reject the EUT hypothesis that $\zeta = \upsilon = 1$. There are some slight changes in the other core structural parameters from Table 1: r is now estimated to be

[23] Many apply the Prelec (1998, Proposition 1, part B) function with constraint $0 < \upsilon < 1$, which requires that the probability weighting function exhibit subproportionality (so-called "inverse-S" weighting). Contrary to received wisdom, many individuals exhibit estimated probability weighting functions that violate subproportionality, so we use the more general specification from Prelec (1998, Proposition 1, part C), only requiring $\upsilon > 0$, and let the evidence determine if the estimated υ lies in the unit interval. This seemingly minor point makes a major difference in our case: Constraining υ to the unit interval incorrectly leads to evidence of *no* probability weighting for the average adult Dane.

[24] In general, when there are more than two outcomes, there is a distinction between the probability weighting function $\omega(p)$ and the decisions weights $w(p)$ derived from them. When there are just two prizes, as here, the decision weights for the highest and lowest outcome are $\omega(p)$ and $1 - \omega(p)$.

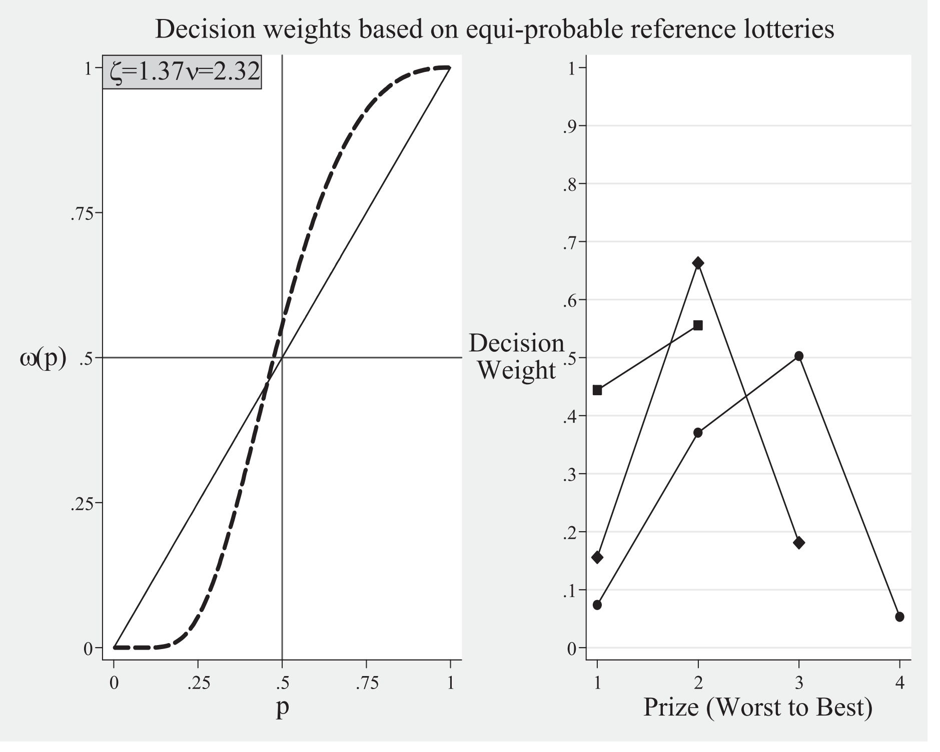

FIGURE 3

PROBABILITY WEIGHTING AND IMPLIED DECISION WEIGHTS

0.46, δ is estimated to be 0.093, and η is estimated to be 0.47. Yet, again we have evidence of statistically significant correlation aversion for the average adult Dane.

6. CONCLUSIONS

We elicit intertemporal risk attitudes from a representative sample of the adult Danish population using real economic commitments and a theoretical framework derived from well-known multiattribute theory. The results suggest that intertemporal risk aversion is a better characterization of the average Dane than intertemporal risk neutrality. This result implies that the convenient additive specification of the intertemporal utility function is not an appropriate representation of intertemporal preferences for the general Danish population. We also show that the characterization of intertemporal risk aversion leads to significantly different estimates of risk premia compared to the estimates obtained when one assumes it away. Our findings have important implications for the characterization of intertemporal preferences in life-cycle modeling, labor supply over time, retirement planning, and policy applications with varying time profiles of costs and benefits.

At a methodological level our approach to estimating correlation aversion complements the earliest work of Keeney (1971, 1972, 1973, 1977) and Keeney and Raiffa (1976) on the direct assessment of multi-attribute utility functions. Much of the literature on empirical demand systems in the 1950s and 1960s used utility structures exhibiting separability and additivity to facilitate estimation (Deaton and Muellbauer, 1980, chapters 3, 5), and early work on utility independence was explicitly motivated by the desire to *directly* elicit multi-attribute utility specifications in an efficient manner. Our approach is to *indirectly* elicit those specifications by estimating a latent structure consistent with observed binary choices over carefully

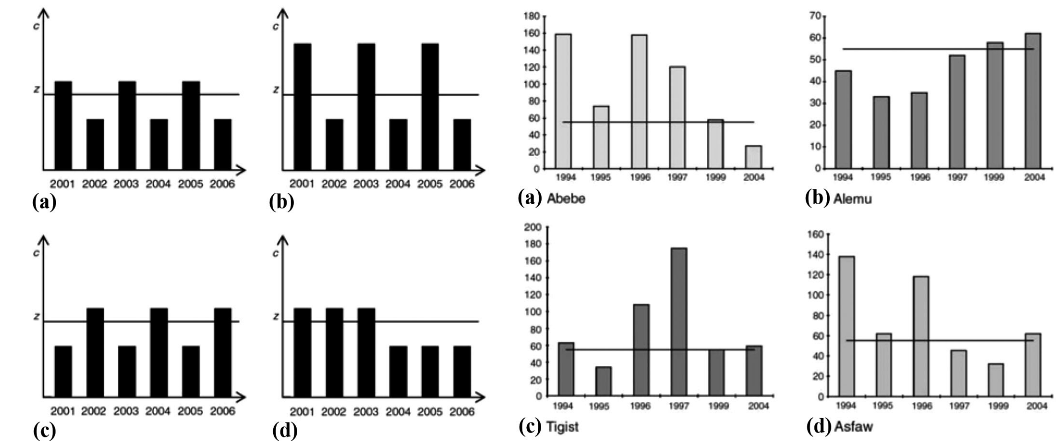

SOURCE: Calco and Dercon (2009, figures 2.1 and 2.2).

FIGURE 4

ILLUSTRATIVE AND EMPIRICAL EXAMPLES OF POVERTY SPELLS

selected time-dated lotteries. We still need to attend to the same general issues of identification but in a way that is arguably less demanding of the subject. Moreover, we are free to explore alternative models of time preferences and risk preferences, as demonstrated, whereas direct elicitation requires that one take a stance on those prior to observation of subject behavior.

Finally, the multi-attribute utility structure we employ is general and not a theoretical framework constructed just to address specific anomalies in intertemporal choice behavior. The concept of correlation aversion has much wider applicability than just to evaluating risk attitudes toward income streams over time.

For example, Bommier (2006, 2010, 2013) applies it to the idea that agents know that they have finite lives, but do not know *when* they will die. This is a naturally occurring and pervasive time-dated risk, and concepts of correlation aversion then have sharp implications for the characterization of discounting, portfolio choice, and life-cycle behavior. One could easily imagine extensions of this application to consider risk management choices in the form of investments in health or location to improve one's chances of living longer, or extensions to consider another attribute, the quality of life, as well as longevity per se.

In terms of atemporal correlation aversion, Keeney (1977) and Keeney and Raiffa (1976) document a wide range of policy applications in which these concepts naturally apply. For example, Gangadharan et al. (2015) consider donor preferences over aid to developing countries that is reasonably expected to have risky outcomes over multiple attributes, in their case access to water to meet basic needs on the one hand and the improvement of sanitation outcomes on the other hand. They find that private donors exhibit correlation *aversion* over these attributes when making donation decisions, in contrast to aid agencies, which tend to promote the complementarity of different attributes of aid of this kind that would be better suited to donors that were correlation *loving*. Thus the concept of correlation aversion would help in the design of better matches between philanthropic donors and aid providers.

As a final, normative example, consider the issue of evaluating the welfare cost of "poverty spells," defined as transient dips below the (absolute or relative) poverty line. Figure 4, extracted from Calvo and Dercon (2009), illustrates with different temporal patterns of consumption, with z denoting a poverty line. Consider the left four hypothetical scenarios and take scenario 4(a) as the benchmark. By comparison, do we say that period-long poverty is less in scenario 4(b) just because the periods of "compensation" provide more consumption than scenario 4(a)? If so, how are we to evaluate how much compensation of this kind would offset the episodic poverty? Is scenario 4(c) worse than 4(a) because the poverty spells arise sooner in time? Similarly is scenario 4(d) better than 4(a) since the poverty spells occur later in time, or is that offset by them being contiguous? And what if these four scenarios are equally likely at the outset and the specific levels of consumption within each year also random around these means? The right four scenarios are taken from field data for specific rural households in different locations in Ethiopia, showing that there can be considerable variety in poverty spells. Whether we are using the risk and time preferences of the households to evaluate these alternatives or some social welfare function reflecting the preferences of donors or governments, rigorous evaluation of these poverty spells requires a general structure that allows correlation aversion and that is flexible enough to accommodate a wide range of models of risk and time preferences.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix

REFERENCES

Abdellaoui, M., E. Decidue, and A. Öncüler, "Risk Preferences at Different Time Periods: An Experimental Investigation," *Management Science* 57 (2011), 975–87.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström, "Eliciting Risk and Time Preferences," *Econometrica* 76 (2008a), 583–619.

———, ———, ———, and ———, "Lost in State Space: Are Preferences Stable?" *International Economic Review* 49 (2008b), 1091–1112.

———, ———, ———, and ———, "Discounting Behavior and the Magnitude Effect," *Economica* 80 (2013), 670–97.

———, ———, ———, and ———, "Discounting Behavior: A Reconsideration," *European Economic Review* 71 (2014), 15–33.

Backus, D. K., B. R. Routledge, and S. E. Zin, "Exotic Preferences for Macroeconomists," *NBER Macroeconomics Annual* 19 (2004), 319–90.

Baucells, M., and F. H. Heukamp, "Common Ratio Using Delay," *Theory & Decision* 68 (2010), 149–58.

Becker, G. M., M. H. DeGroot, and J. Marschak, "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science* 9 (1964), 226–32.

Bommier, A., "Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion," Working Paper, CNRS-GREMAQ, University of Toulouse, 2005. Available at http://doi.org/10.3929/ethz-a-010574890.

———, "Uncertain Lifetime and Intertemporal Choice: Risk Aversion as a Rationale for Time Discounting," *International Economic Review* 47 (2006), 1223–46.

———, "Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion," *Economics Bulletin* 4 (2007), 1–8.

———, "Portfolio Choice under Uncertain Lifetime," *Journal of Public Economic Theory* 12 (2010), 57–73.

———, "Life-Cycle Preferences Revisited," *Journal of the European Economic Association* 11 (2013), 1290–319.

Bommier, A., and J. Rochet, "Risk Aversion and Planning Horizons," *Journal of the European Economic Association* 4 (2006), 708–34.

Bommier, A., A. Kochov, and F. Le Grand, "On Monotone Recursive Preferences," *Econometrica* 85 (2017), 1433–66.

Calvo, C., and S. Dercon, "Chronic Poverty and All That," in T. Addison, D. Hulm and R. Kanbur, eds., *Poverty Dynamics: Interdisciplinary Perspectives* (Oxford: Oxford University Press, 2009).

Coble, K. H., and J. L. Lusk, "At the Nexus of Risk and Time Preferences: An Experimental Investigation," *Journal of Risk and Uncertainty* 41 (2010), 67–79.

Coller, M., G. W. Harrison, and E. E. Rutström, "Latent Process Heterogeneity in Discounting Behavior," *Oxford Economic Papers* 64 (2012), 375–91.

———, and M. B. Williams, "Eliciting Individual Discount Rates," *Experimental Economics* 2 (1999), 107–27.

DEATON, A., AND J. MUELLBAUER, *Economics and Consumer Behavior* (Cambridge: Cambridge University Press, 1980).

DELQUIE, P., AND M. LUO, "A Simple Trade-off Condition for Additive Multiattribute Utility," *Journal of Multi-Criteria Decision Analysis* 6 (1997), 248–52.

DENUIT, M., L. EECKHOUDT, AND B. REY, "Some Consequences of Correlation Aversion in Decision Science," *Annals of Operations Research* 176 (2010), 259–69.

DORFLEITNER, G., AND M. KRAPP, "On Multiatttributive Risk Aversion: Some Clarifying Results," *Review of Managerial Science* 1 (2007), 47–63.

DUNCAN, G. T., "A Matrix Measure of Multivariate Local Risk Aversion," *Econometrica* 45 (1977), 895–903.

EECKHOUDT, L., B. REY, AND H. SCHLESINGER, "A Good Sign for Multivariate Risk Taking," *Management Science* 53 (2007), 117–24.

EPSTEIN, L. G., "Stationary Cardinal Utility and Optimal Growth Under Uncertainty," *Journal of Economic Theory* 31 (1983), 133–52.

———, AND S. M. TANNY, "Increasing Generalized Correlation: A Definition and Some Economic Consequences," *Canadian Journal of Economics* 13 (1980), 16–34.

———, AND S. E. ZIN, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57 (1989), 937–69.

———, E. FARHI, AND T. STRZALECKI, "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review* 104 (2014), 2680–97.

FARMER, R., "RINCE Preferences," *Quarterly Journal of Economics* 105 (1990), 43–60.

FISHBURN, P. C., "Independence in Utility Theory with Whole Product Sets," *Operations Research* 13 (1965), 28-45.

———, "Multiattribute Nonlinear Utility Theory," *Management Science* 30 (1984), 1301–10.

FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature* 40 (2002), 351–401.

GANGADHARAN, L., G. W. HARRISON, AND A. LEROUX, "Preferences over Risky Aid," Working Paper 2015-12, Center for Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2015.

HALL, R. E., "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96 (1988), 339–57.

HANSEN, L. P., AND K. J. SINGLETON, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy* 91 (1983), 249–65.

HARRISON, G. W., M. I. LAU, AND M. B. WILLIAMS, "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review* 92 (2002), 1606–17.

HARVEY, C. M., "Multiattribute Risk Linearity," *Management Science* 39 (1993), 389–94.

KARNI, E., "On Multivariate Risk Aversion," *Econometrica* 47 (1979), 1391–1401.

KEENEY, R. L., "Utility Independence and Preferences for Multiattribute Consequences," *Operations Research* 19 (1971), 875–93.

———, "Utility Functions for Multiattributed Consequences," *Management Science* 18 (1972), 276–87.

———, "Risk Independence and Multiattributed Utility Functions," *Econometrica* 41 (1973), 27–34.

———, "The Art of Assessing Multiattribute Utility Functions," *Organizational Behavior and Human Performance* 19 (1977), 267–310.

———, AND H. RAIFFA, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (New York: Wiley, 1976).

KIHLSTROM, R. E., AND L. J. MIRMAN, "Risk Aversion with Many Commodities," *Journal of Economic Theory* 8 (1974), 361–88.

MIYAMOTO, J. M., AND P. WAKKER, "Multiattribute Utility Theory Without Expected Utility Foundations," *Operations Research* 44 (1996), 313–26.

NOUSSAIR, C., AND P. WU, "Risk Tolerance in the Present and the Future: An Experimental Study," *Managerial and Decision Economics* 27 (2006), 401–12.

PAYNE, J. W., D. J. LAUGHHUNN, AND R. CRUM, "Multiattribute Risky Choice Behavior: the Editing of Complex Prospects," *Management Science* 30 (1984), 1350–61.

PLISKIN, J. S., D. S. SHEPARD, AND M. C. WEINSTEIN, "Utility Functions for Life Years and Health Status," *Operations Research* 28 (1980), 206–24.

POLLACK, R. A., "Additive Von Neumann-Morgenstern Utility Functions," *Econometrica* 35 (1967), 485–94.

PRELEC, D., "The Probability Weighting Function," *Econometrica* 66 (1998), 497–527.

QUIGGIN, J., "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization* 3 (1982), 323–43.

READ, D., "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty* 23 (2001), 5–32.

RICHARD, S. F., "Multivariate Risk Aversion, Utility Independence and Separable Utility Functions," *Management Science* 22 (1975), 12–21.

ROTHBLUM, U. G., "Multivariate Constant Risk Posture," *Journal of Economic Theory* 10 (1975), 309–32.

STARMER, C., "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk," *Journal of Economic Literature* 38 (2000), 332–82.

TOWNSEND, R. M., "Risk and Insurance in Village India," *Econometrica* 62 (1994), 539–91.

TSETLIN, I., AND R. L. WINKLER, "Multiattribute Utility Satisfying a Preference for Combining Good with Bad," *Management Science* 55 (2009), 1942–52.

VON WINTERFELDT, D., "Additivity and Expected Utility in Risky Multiattribute Preferences," *Journal of Mathematical Psychology* 21 (1980), 66–82.

WEIL, P., "Nonexpected Utility in Macroeconomics," *Quarterly Journal of Economics* 105 (1990), 29–42

WILCOX, N. T., "'Stochastically More Risk Averse:' A Contextual Theory of Stochastic Discrete Choice under Risk," *Journal of Econometrics* 162 (2011), 89–104.