2016

Random Expected Utility and Certainty Equivalents: Mimicry of Probability Weighting Functions

Nathaniel Wilcox
Chapman University, nwilcox@chapman.edu

Random Expected Utility and Certainty Equivalents: Mimicry of Probability Weighting Functions

Comments
Working Paper 16-14

Random Expected Utility and Certainty Equivalents: Mimicry of Probability Weighting Functions

by

Nathaniel T. Wilcox
Economic Science Institute
Chapman University
Orange, CA 92866
nwilcox@chapman.edu

November 28, 2016
(version 1.2)

Abstract: For simple prospects routinely used for certainty equivalent elicitation, random expected utility preferences imply a conditional expectation function that can mimic deterministic rank dependent preferences. That is, an agent with random expected utility preferences can have expected certainty equivalents exactly like those predicted by rank dependent probability weighting functions of the inverse-s shape discussed by Quiggin (1982) and advocated by Tversky and Kahneman (1992), Prelec (1998) and other scholars. Certainty equivalents may not nonparametrically identify preferences: Their conditional expectation (and critically, their interpretation) depends on assumptions concerning the source of their variability.

I am grateful to the Economic Science Institute and Chapman University for their ongoing support. Jose Apesteguia, Matias Del Campo Axelrod, Miguel A. Ballester, Sudeep Bhatia, Michael Birnbaum, Daniel Cavagnaro, Mark DeSantis, Glenn W. Harrison, Anthony Marley, John P. Nolan and Mark Schneider provided help or commentary, though none are responsible for remaining errors.

Elicitation of certainty equivalents has become routine in laboratory measurement of preferences under risk and uncertainty (Tversky and Kahneman 1992; Tversky and Fox 1995; Wu and Gonzales 1999; Gonzales and Wu 1999; Abdellaoui 2000; Abdellaoui, Bleichrodt and Paraschiv 2007; Halevy 2007; Bruhin, Fehr-Duda and Epper 2010; Vieider et al. 2015). While elicitation methods vary across such studies, the formal empirical interpretation of elicited certainty equivalents is invariably the same. The subject is assumed to have a unique and fixed preference order, implying (under unchanged conditions of background wealth, risk and so forth) a unique and fixed certainty equivalent for each prospect. *Elicited* certainty equivalents are then interpreted as this unique and fixed certainty equivalent plus some error of banal origin with standard properties.

Such added error, or something like it, is necessary: In repeated elicitations using exactly the same prospect, elicited certainty equivalents vary within subjects (Tversky and Kahneman 1992, p. 307; Krahnen, Rieck and Theissen 1997, p. 477; von Winterfeldt et al. 1997, p. 422; Gonzalez and Wu 1999, pp. 144-146; Pennings and Smidts 2000, p. 1342) and other evidence also suggests inherent variability of elicited certainty equivalents (Butler and Loomes 2007; Loomes and Pogrebna 2014). Luce (1997, pp. 81-82) argued that theory and empirical interpretation need to take a position on such response variability. Adding mean zero error to an otherwise deterministic model of certainty equivalents is clearly one option here, and I call this the *standard model* of an elicited certainty equivalent.

Alternatively, one might assume that the individual subject's preference order is a random variable, and that each certainty equivalent elicited from that subject is fully determined by a single realization of that random variable: Call this a *random preference model* of an elicited certainty equivalent. Interest in random preference models is both long-standing and contemporary (Becker, DeGroot and Marschak 1963; Eliashberg and Hauser 1985; Hilton 1989; Loomes and Sugden 1995, 1998; Regenwetter and Marley 2001; Gul and Pesendorfer 2006; Regenwetter, Dana and Davis-Stober 2011; Ahn and Sarver 2013; Apesteguia and Ballester 2016; Karni and Safra 2016), particularly in the realm of discrete choice. Here, I examine implications of this model for elicited certainty equivalents and find a significant complication of their empirical interpretation.

Specifically, random model expected utility preferences (or more simply *random EU* as Gul and Pesendorfer 2006 call it) imply expected certainty equivalents that can *mimic* those implied by standard model rank-dependent preferences (or more simply *standard RDU*). That is, a random EU agent can have expected certainty equivalents that appear to reveal rank dependent probability weighting functions of the inverse-s shape discussed by Quiggin (1982) and advocated by Tversky and Kahneman (1992) and other scholars. Indeed, one may derive the celebrated Prelec (1998) weighting functions from certainty equivalents governed by random EU. Hilton (1989) first showed that certainty equivalents have some unexpected properties under random EU; additionally, recent work by Navarro-Martinez et al. (2015) contains a strong suggestion of my direction here. My results come first, with some intuition to follow; and to conclude I argue, contra widespread suggestions to the contrary, that elicited certainty equivalents may *not* nonparametrically identify preferences, since their conditional expectation (and critically, their interpretation) depends on the source of their variability.

1. Formal results

Consider *simple prospects* (W, p) with money outcomes $z = W > 0$ with probability p and $z = 0$ with probability $1 - p$. Simple prospects figure prominently in discussions of rank-dependent utility (RDU) and cumulative prospect theory (CPT) because their certainty equivalents are thought to reveal the *probability weighting function* of the rank-dependent preference family when the utility or value of outcomes is linear (Tversky and Kahneman 1992; Prelec 1998). To see this, let the utility or value of outcomes have the power form $v(z) = z^{1/x}$, where I write the power as $1/x$ for convenience and assume that $x \in (0, \infty)$. The rank dependent utility or RDU of (W, p) will then be $\pi(p|\omega) \cdot W^{1/x}$ (given specific x and ω) where $\pi(p|\omega)$ is a probability weighting function depending on preference parameters ω. The certainty equivalent of (W, p) is then $\pi(p|\omega)^x \cdot W$, but divide this by W to free it of dependence on W and let $C^{rd}(p|x, \omega) \equiv \pi(p|\omega)^x$ be the RDU *normalized certainty equivalent* of any simple prospect (given specific x and ω). Notice that when $x = 1$ (that is for a linear value of outcomes), one has $C^{rd}(p|1, \omega) \equiv \pi(p|\omega)$, so normalized certainty equivalents of simple prospects are thought to reveal RDU (or CPT) probability weighting functions when $x = 1$. Expected utility or EU is the special case where $\pi(p|\omega) \equiv p$, so also define $C^{eu}(p|x) \equiv p^x$ as the EU normalized certainty equivalent of any simple prospect (given specific x).

Let ce be the certainty equivalent for (W, p) elicited from a subject and let $c = ce/W \in [0,1]$ be the normalized version of it. Commonly, if only implicitly, the empirical specification for observed normalized certainty equivalents is $c = E(c|p) + \varepsilon$, where $E(c|p)$ is the *conditional expectation function* or c.e.f. of c and one assumes that errors ε have conventional properties (e.g. $E(\varepsilon) = E(\varepsilon|p) = 0$). In a *standard RDU model*, the c.e.f. is $E(c|p) = C^{rd}(p|x,\omega) \equiv \pi(p|\omega)^x$, yielding the specification $c = \pi(p|\omega)^x + \varepsilon$ in which the error term ε is thought to arise from banal sources such as "carelessness, hurrying, or inattentiveness" (Bruhin, Fehr-Duda and Epper 2010, p. 1383). Estimation of (x, ω) can then proceed using various estimators. Bruhin, Fehr-Duda and Epper use maximum likelihood, while Tversky and Kahneman (1992) use nonlinear least squares. Since c is a limited dependent variable (most elicitation methods enforce $c \in [0,1]$), the distribution of ε cannot be wholly independent of p. But one may accommodate this while keeping $E(\varepsilon|p) = 0$ in various ways (see e.g. Bruhin, Fehr-Duda and Epper; Gonzalez and Wu 1999). A brief Monte Carlo study in my appendix looks at similar (and other) estimators.

For many weighting functions $\pi(p|\omega)$, x and ω cannot be wholly identified solely from the normalized certainty equivalents of simple prospects. For instance suppose $\pi(p|\omega)$ is the celebrated 2-parameter Prelec (1998) weighting function $\exp(-\beta[-\ln(p)]^\alpha)$ where α and β are strictly positive parameters. Then the normalized certainty equivalent will be $C^{rd}(p|x,\alpha,\beta) = \exp(-x\beta[-\ln(p)]^\alpha)$, and only α and the product $x\beta$ will be estimable. Scholars know this quite well, so experimental designs meant to separately estimate all three parameters always contain at least some "non-simple" prospects. I focus on simple prospects because of their tractability and their simple interpretation under standard RDU: $C^{rd}(p|1,\omega) \equiv \pi(p|\omega)$, so normalized certainty equivalents of simple prospects reveal weighting functions at linear $v(z)$ (Tversky and Kahneman 1992; Prelec 1998). The Monte Carlo study in my appendix employs an experimental design (that of Gonzalez and Wu 1999) containing both simple and non-simple prospects, and results of that study echo my formal results for simple prospects.

In general, a *random preference model* might take both x and ω to be realizations of nondegenerate random variables X and Ω within an individual. However, as far as I am aware, existing random preference estimations treat any weighting function parameters ω as fixed within any individual (e.g. Loomes, Moffatt and Sugden 2002; Wilcox 2008, 2011), and contemporary random preference theory seems to be confined to treatment of X as random only (e.g. Gul and Pesendorfer 2006; Apesteguia and Ballester 2016). Therefore, all of my random preference analysis treats only x as the realization of a random variable X; and in any case I confine my analysis of the random model to the random EU special case.

In any distinct elicitation trial, an independent random EU model, or simply random EU for short, assumes that an independent and identically distributed realization x of X occurs and fully determines the normalized certainty equivalent $C^{eu}(p|x) = p^x$. Assume that a probability density function $f(x|\psi)$ of X with support $(0,\infty)$ lies within an individual, with parameters ψ governing moments or location, scale and so forth. Then define $\mathbb{C}^{eu}(p|\psi)$ as

$$(1) \qquad \mathbb{C}^{eu}(p|\psi) \equiv E_X[C^{eu}(p|x)] = \int_0^\infty p^x f(x|\psi)dx = \int_0^\infty \exp(-x\tau)f(x|\psi)dx,$$

where the final integral (which becomes useful shortly) lets $\tau = -\ln(p)$ and rewrites p^x as $\exp(-x\tau)$. The function $\mathbb{C}^{eu}(p|\psi)$ is the expected normalized certainty equivalent of a random EU agent for simple prospects (W,p), given her underlying p.d.f. $f(x|\psi)$ with parameters ψ. Under random EU, we again have an empirical model $c = E(c|p) + \xi$ of the same form as the standard RDU model. However, the random EU c.e.f. is $E(c|p) = \mathbb{C}^{eu}(p|\psi)$ as given by eq. 1, and the errors ξ are just $p^x - \mathbb{C}^{eu}(p|\psi)$. The eq. 1 definition implies that these new errors ξ also satisfy the usual properties ($E(\xi) = E(\xi|p) = 0$), so we have a close resemblance between the random EU model $c = \mathbb{C}^{eu}(p|\psi) + \xi$ and the standard RDU model $c = C^{rd}(p|x,\omega) + \varepsilon$ and can estimate both using the same variety of estimators.

Allow a brief digression on elicitation methods. There is another way of thinking about the p.d.f. $f(x|\psi)$ of X in the Random EU model. Suppose an experimenter uses some method M to elicit certainty equivalents ce from a subject, normalizing them as $c = ce/W$, and computes $x(c,p)$ to solve $p^x = c$; that is, let $x(c,p) \equiv \ln(c)/\ln(p)$. Suppose that in repeated elicitations using method M, across various values of p, the empirical c.d.f. of

$x(c,p)$ is observed to be $\hat{F}_M(x|p)$ which converges to $F_M(x|p)$ as the sample of observations grows. If $F_M(x|p)$ is in fact independent of p and so just $F_M(x)$, the variability of the normalized certainty equivalents observed by the experimenter could be interpreted as arising from a random EU model of the kind assumed here, where a p.d.f. $f_M(x|\psi)$ is derived from $F_M(x)$. This suggests ways in which one might test versions of the random EU model (or versions of a random RDU model, and later I will return to this), but additionally indicates that the results here only require that elicitation methods satisfy two key assumptions: (1) repeated trials using the method yield variability in elicited certainty equivalents; and (2) this variability is consistent with the assumptions of a random EU model—namely, that $F_M(x)$ is independent of p. Neither assumption rules out dependence of $f(x|\psi)$ on the elicitation method M; and debates about the relative merits of various elicitation methods need not impinge on these assumptions in any necessary way.

The close resemblance of the standard RDU and random EU models suggests two possible types of *mimicry*. First, since $C^{rd}(p|1,\omega) \equiv \pi(p|\omega)$ in standard RDU, it will be troubling if $\mathbb{C}^{eu}(p|\psi)$ can "look like" a stereotypical $\pi(p|\omega)$, that is, can have properties like those that scholars believe are empirically characteristic of RDU weighting functions. I will refer to this as *weak mimicry* (of standard RDU by random EU). Second, it may happen that for some well-known and specific $\pi(p|\omega)$, there exists a specific $f(x|\psi)$ such that $\mathbb{C}^{eu}(p|\psi)$ is a *re-parameterization* of $C^{rd}(p|x,\omega)$. Let D be the set of possible parameter vectors (x,ω), and let Ψ be the set of possible parameter vectors ψ; and suppose that, for some $f(x|\psi)$, there exists a function $H_f\colon D \to \Psi$ such that $\mathbb{C}^{eu}\left[p|H_f(x,\omega)\right] \equiv C^{rd}(p|x,\omega) \equiv \pi(p|\omega)^x$: Then one may say there is *strong mimicry* (of standard RDU by random EU) for $f(x|\psi)$. Notice that strong mimicry implies weak mimicry but not vice versa.

Since $-\ln(p) > 0 \;\forall\, p \in (0,1)$, so that $\tau > 0$ too, the final integral in eq. 1 is the one-sided Laplace transform $\mathcal{L}\{f\}(\tau)$ of the p.d.f. $f(x|\psi)$—provided it exists; and below I only use p.d.f.s for which the existence and form of $\mathcal{L}\{f\}(\tau)$ have been demonstrated and derived by others. In such instances, these known Laplace transforms $\mathcal{L}\{f\}(\tau)$ of a p.d.f. $f(x|\psi)$ make it simple to derive various examples of $\mathbb{C}^{eu}(p|\psi)$, using the relationship $\mathbb{C}^{eu}(p|\psi) = \mathcal{L}\{f\}[-\ln(p)]$. Some examples follow.

Example 1. Suppose *X* has the Gamma p.d.f.

$$f(x|k,\theta) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-(x/\theta)} \text{ for } x, k, \text{ and } \theta \in (0,\infty).$$

It's widely known that this has the Laplace transform $\mathcal{L}\{f\}(\tau) = (1+\theta\tau)^{-k}$, implying that

(2) $\quad \mathbb{C}^{eu}(p|k,\theta) = (1-\theta\ln(p))^{-k}$.

Figure 1-A shows this Gamma c.e.f. for $k = 0.75$ and θ =2.79. At these parameter choices, it has the "inverse-s" shape many believe is characteristic of weighting functions $\pi(p|\omega)$ and the fixed point $p \approx e^{-1}$ which is also characteristic of Prelec's (1998) 1-parameter weighting function; so this is an instance of weak mimicry. One needs to say that this Gamma c.e.f. *can* (not *must*) weakly mimic this characteristic shape. Figure 1-B shows the Gamma c.e.f. for $k = 0.75$ and θ =0.9: Here we see the "optimist" shape discussed by Quiggin (1982), and also the plurality shape of individually estimated weighting functions in Wilcox (2015). Such shape flexibility is also characteristic of 2-parameter weighting functions found in the literature on RDU and CPT, where this flexibility is usually regarded as a feature rather than a weakness.

Example 2. Suppose *X* has the Inverse Gaussian (Wald) p.d.f.

$$f(x|\mu,\lambda) = \sqrt{\frac{\lambda}{2\pi}} x^{-3/2} exp\left[\frac{-\lambda(x-\mu)^2}{2\mu^2 x}\right] \text{ for } x, \mu \text{ and } \lambda \in (0,\infty).$$

This has the Laplace transform $\mathcal{L}\{f\}(\tau) = exp\left[(\lambda/\mu)\left(1-\sqrt{1+2\mu^2\tau/\lambda}\right)\right]$ (Seshadri 1993 p. 41), implying that

(3) $\quad \mathbb{C}^{eu}(p|\mu,\lambda) = exp\left[(\lambda/\mu)\left(1-\sqrt{1-2\mu^2\ln(p)/\lambda}\right)\right]$.

Figure 2 shows this Inverse Gaussian c.e.f. for $\mu = 11$ and $\lambda = 0.55$. This also has the inverse-s shape and is also an instance of weak mimicry. Again, different parameter values will produce a wide variety of different shapes of this Inverse Gaussian c.e.f.

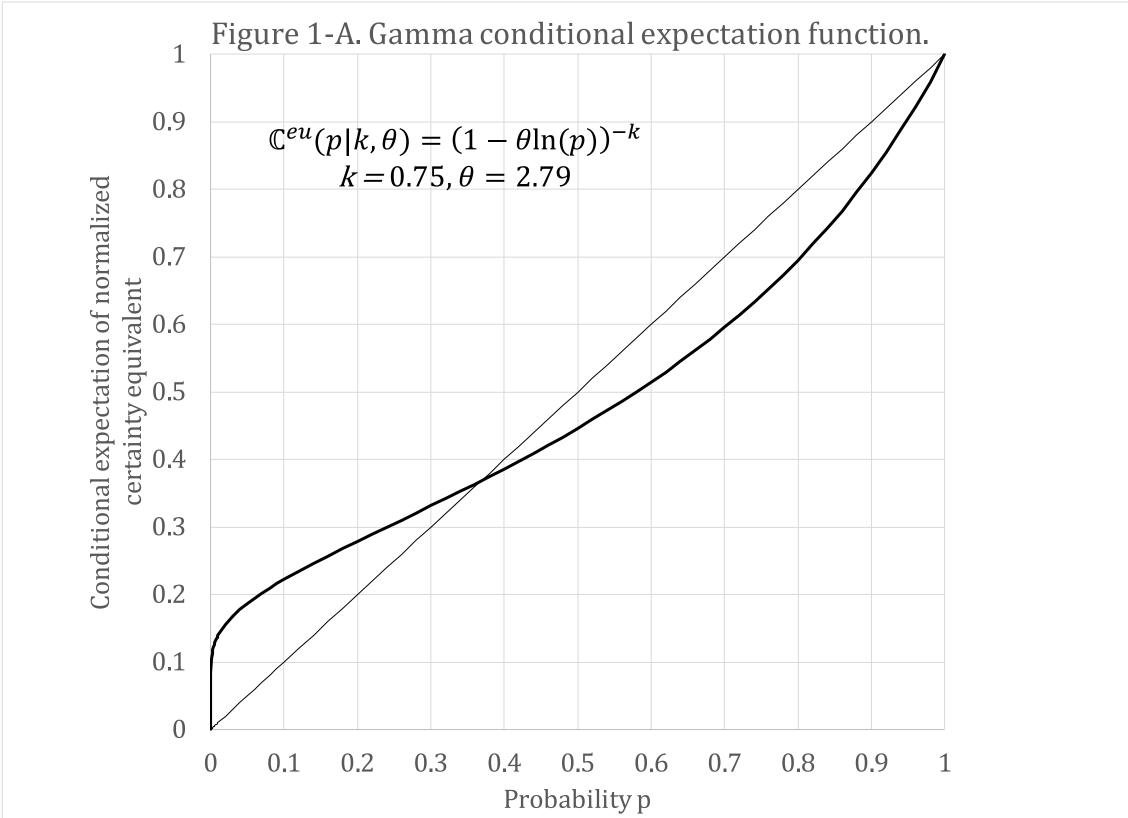
Figure 1-A. Gamma conditional expectation function.
$\mathbb{C}^{eu}(p|k,\theta) = (1-\theta\ln(p))^{-k}$
$k = 0.75, \theta = 2.79$
Conditional expectation of normalized certainty equivalent
Probability p
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

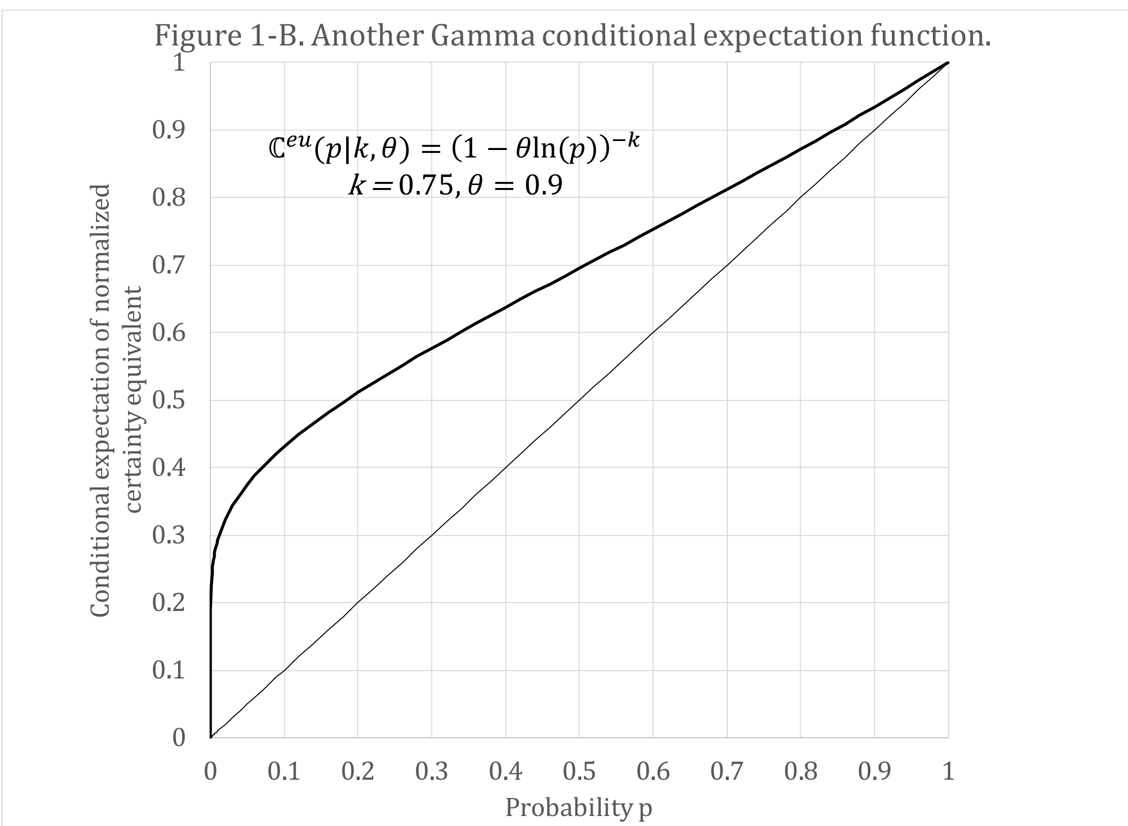
Figure 1-B. Another Gamma conditional expectation function.
$\mathbb{C}^{eu}(p|k,\theta) = (1-\theta\ln(p))^{-k}$
$k = 0.75, \theta = 0.9$
Conditional expectation of normalized certainty equivalent
Probability p
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Figure 2. Inverse Gaussian conditional expectation function.

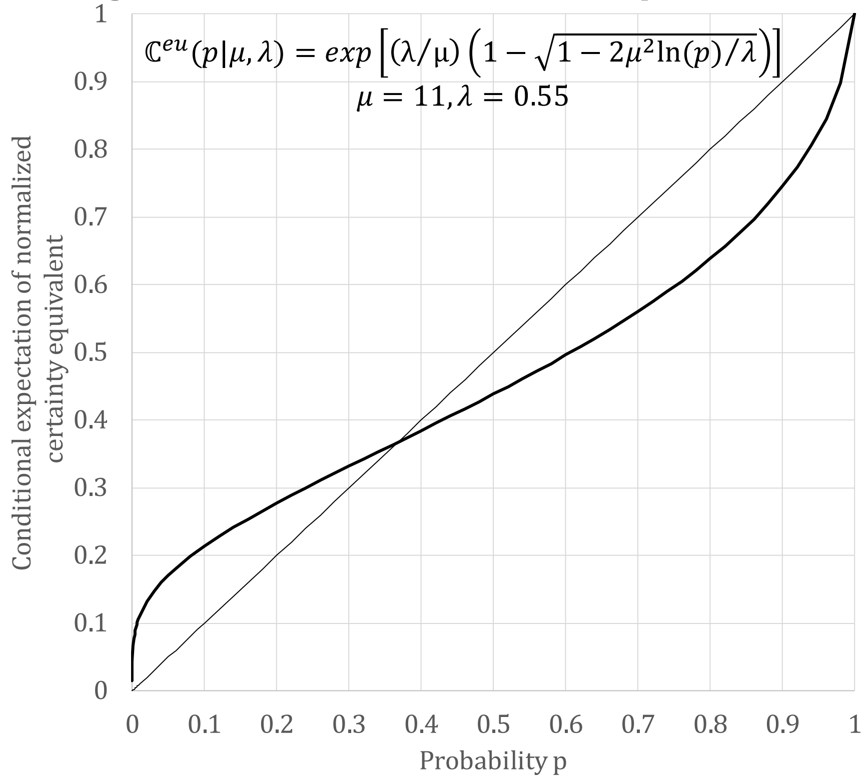

Example 3. Suppose X has the (unshifted) Lévy p.d.f.

$$f(x|\delta) = \frac{\delta}{2\sqrt{\pi}} x^{-3/2} exp\left(\frac{-\delta^2}{4x}\right) \text{ for } x \text{ and } \delta \in (0,\infty).$$

This has the Laplace transform $\mathcal{L}\{f\}(\tau) = exp\left(-\delta\tau^{1/2}\right)$ (González-Velasco 1995, p. 537), implying that

$$\text{(4)} \qquad \mathbb{C}^{eu}(p|\delta) = exp\left(-\delta[-\ln(p)]^{1/2}\right).$$

Earlier I noted that in the case of the Prelec (1998) 2-parameter weighting function, $C^{rd}(p|x,\alpha,\beta) = \exp(-x\beta[-\ln(p)]^{\alpha})$. Clearly, this is identical to eq. 4 if we set $\delta = x\beta$ and require that $\alpha = 1/2$. This is very close to being a case of strong mimicry, but not quite, since eq. 4 can only mimic Prelec weighting functions when α just happens to be $1/2$. Empirically, estimates of α have a wider range than a small neighborhood of $1/2$.

However, this result for the Lévy distribution provides a strong and fruitful hint. The Lévy distribution is a specific instance of the Lévy Alpha-Stable distributions, also known more simply as the Stable distributions. Except for special cases (Normal, Cauchy and Lévy), Stable random variables X have no p.d.f. expressible in terms of elementary functions, but their Laplace transforms exist as relatively simple expressions. For Stable random variables with support $(0,\infty)$, Nolan (2017, p. 109) shows that the Laplace transform exists and is $\mathcal{L}\{f\}(\tau) = \exp\left\{-\gamma^a \left(sec\frac{a\pi}{2}\right)\tau^a\right\}$, where $\gamma > 0$ is a scale parameter and $a \in (0,1)$ is called the *index of stability* or *characteristic exponent* (see Feller 1971 and Hougaard 1986 for similar forms parameterized differently). Therefore, for Stable distributions of X on $(0,\infty)$, we have

$$(5) \qquad \mathbb{C}^{eu}(p|a,\gamma) = \exp\left\{-\gamma^a \left(sec\frac{a\pi}{2}\right)[-\ln(p)]^a\right\}.$$

Eq. 5 is identical to $C^{rd}(p|x,\alpha,\beta) = \exp(-x\beta[-\ln(p)]^\alpha)$ when we set $a = \alpha$ and $\gamma = \left[\beta x\left(cos\frac{\alpha\pi}{2}\right)\right]^{1/\alpha}$, so we have strong mimicry of standard RDU with the 2-parameter Prelec (1998) function—provided that $\alpha < 1$. Since this is both characteristic of most empirical estimates of α and indeed yields the characteristic inverse-s shape, this is strong mimicry of a well-known and widely used probability weighting function in the relevant part of the parameter space.

2. Intuition

For intuition behind the formal results, it helps to use the more usual representation of the power utility function, specifically $u(z) = z^\sigma$, thinking now of σ as having a distribution with support $\Sigma \subseteq (0,\infty)$ under the random preference model. Then under EU the normalized certainty equivalent of a simple prospect (given any value of σ) will be $p^{1/\sigma}$, whose second derivative with respect to σ is

$$(6) \qquad \frac{\partial^2}{\partial\sigma^2}\, p^{1/\sigma} = \frac{-\ln(p)p^{1/\sigma}}{\sigma^4}[-\ln(p) - 2\sigma] > 0 \ \text{ for all } \ \sigma < -\frac{1}{2}\ln(p).$$

As p approaches zero, eq. 6 shows that the normalized certainty equivalent $p^{1/\sigma}$ approaches being convex in σ at all $\sigma \in (0, \infty)$. Moreover, if Σ is in fact bounded above, there will be sufficiently small p such that $p^{1/\sigma}$ is convex $\forall\ \sigma \in \Sigma$: In that event, Jensen's Inequality implies that $E\left(p^{1/\sigma}\right) > p^{1/E(\sigma)}$. Assuming that $E(\sigma) \geq 1$, then, we have $E\left(p^{1/\sigma}\right) > p$ for sufficiently small p. That is: If Σ is bounded above and $E(\sigma) = 1$ (the mean value function $v(z)$ is linear), mean normalized certainty equivalents will exceed p when p is small enough. We have apparent overweighting of small enough probabilities.

As p approaches 1, on the other hand, eq. 6 shows that $p^{1/\sigma}$ becomes concave at almost all σ, and the argument above flips around: If Σ is bounded below away from zero, there will be p sufficiently close to 1 such that $p^{1/\sigma}$ is concave $\forall\ \sigma \in \Sigma$, and Jensen's Inequality then implies that $E\left(p^{1/\sigma}\right) < p^{1/E(\sigma)}$. Assuming that $E(\sigma) \leq 1$, we have $E\left(p^{1/\sigma}\right) < p$ for p sufficiently close to one. That is: If Σ is bounded below away from zero and $E(\sigma) = 1$ (the mean value function $v(z)$ is linear), mean normalized certainty equivalents will fall short of p when p is high enough: We have apparent underweighting of high enough probabilities.

Figure 3 illustrates this intuition. Assume that the agent has a binomial distribution on σ such that $E(\sigma) = 1$: Specifically she has $\sigma = 0.5$ with probability $2/3$ and $\sigma = 2$ with probability $1/3$. Figure 3 shows the function $p^{1/\sigma}$ for $\sigma \in (0,3]$ for two values of p. The upper heavy curve is for $p = 0.95$ and, as can be seen, this curve is overwhelmingly and strongly concave: In this case, $E\left[0.95^{(1/\sigma)}\right] < 0.95$, so this agent appears to underweight high probabilities. The lower heavy curve is for $p = 0.05$ and, as can be seen, this curve is first convex and, for σ beyond about 1.5, very gently concave: Here, $E\left[0.05^{(1/\sigma)}\right] > 0.05$, so this agent appears to overweight low probabilities.

This story does not completely explain the formal results: All of the examples in section 1 involve p.d.f.s with support $(0, \infty)$, so this story (which is told by appealing to a support bounded above and bounded below away from zero) is only an aid to intuition, not any sort of demonstration. Therefore, the formal results are needed. The intuition does, however, explain why one may easily derive the characteristic inverse-s shape from many p.d.f.s $f(x|\psi)$ underlying a random EU model.

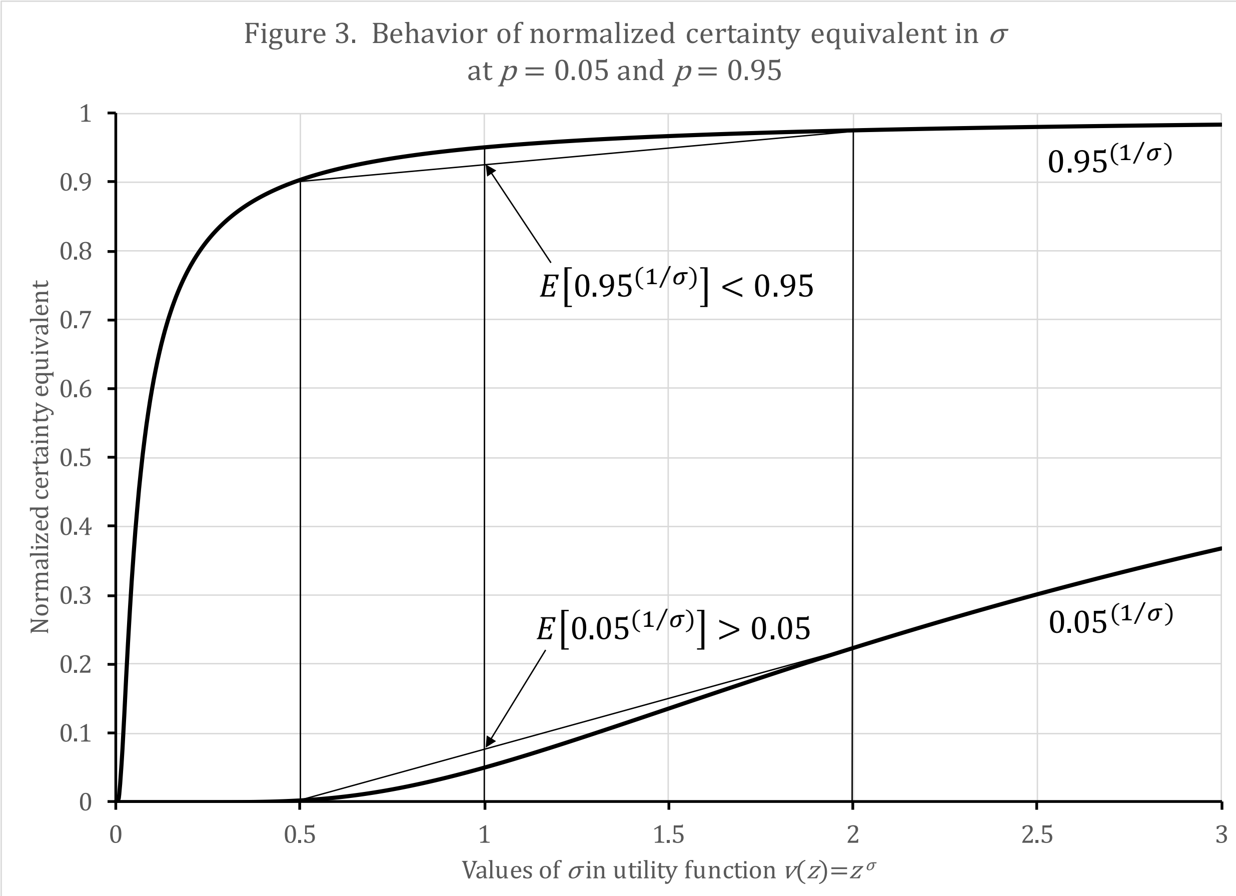

3. Discussion and conclusions

My results complicate interpretation of elicited certainty equivalents. However, I say 'complicate' rather than 'undermine' for several reasons. First, I have not shown that random preference EU and standard RDU certainty equivalents are indistinguishable. The formal results are entirely about conditional expectations and say nothing about conditional medians or conditional variances and other moments; and one might test both random EU and random RDU on the basis of these other characteristics.

For instance, recall that $C^{rd}(p|x,\omega) \equiv \pi(p|\omega)^x$ and suppose we now assume that x is a realization of a random variable X, giving a random RDU model. Let CV denote coefficient of variation; then under random RDU, assuming that any weighting function parameters are fixed (not themselves random variables), we have

(7) $$CV(-\ln(c)) \equiv \frac{\sqrt{V(-\ln(c))}}{E(-\ln(c))} = \frac{\sqrt{V(X)[-\ln(\pi(p|\omega))]^2}}{E(X)[-\ln(\pi(p|\omega))]} = \frac{\sqrt{V(X)}}{E(X)} = CV(X)$$

This says that for any given individual, the coefficient of variation of $-\ln(c)$ will be equivalent to the coefficient of variation of X and, moreover, independent of the particular W and p of any simple prospect (W, p), regardless of whether the weighting function is an identity function (EU) or not (RDU). This immediately suggests a test of both random preference EU and RDU based on multiple (more than two) certainty equivalent elicitation trials for several different simple prospects. To my knowledge, such data are scarce but more could be gathered with appropriate experimental designs. The key point, however, is that for certainty equivalents, the random preference hypothesis can make strong refutable predictions about higher moments that are independent of the form or even the presence of any rank-dependent weighting function.

Second, discrete choice experiments already suggest that random EU cannot be a complete model of discrete choices (e.g. Loomes and Sugden 1998). Under the random preference hypothesis, much of what EU predicts concerning pairs of related discrete choice problems remains unchanged relative to what EU predicts in its deterministic form (Loomes and Sugden 1995; Gul and Pesendorfer 2006; Wilcox 2008). This implies that many well-known discrete choice violations of EU also violate random EU. Here I showed once more (see Hilton 1989) that certainty equivalents are a different matter: Under random EU, the expected values of certainty equivalents can mimic predictions of standard RDU and CPT. The upshot of this fact is that when one estimates risk models from certainty equivalents, part of the estimates (perhaps substantial parts) may reflect random preference heterogeneity as well as any underlying mean preference.

Third, the formal results only complicate estimation based on conditional expectation functions. While this is the overwhelmingly common basis for *estimation*, some of the empirical literature on RDU and CPT uses pooled sample conditional medians of certainty equivalents for *description* (Tversky and Kahneman 1992, pp. 309-311; Gonzalez and Wu 1999, p. 144-145). It may be that conditional median estimation (that is, least absolute deviation or LAD estimators) can solve the problem uncovered here. Recall the key role played by Jensen's Inequality in Section 2 where I discussed intuition: There is no counterpart of Jensen's Inequality for medians. No individual-level conditional median estimations (based on elicited certainty equivalents, using LAD estimation) of either RDU or CPT models are available. (Tversky and Kahneman 1992 do estimate a pooled sample weighting function from pooled sample conditional medians, but using a nonlinear least squares estimator.) My appendix looks at a LAD estimator and finds encouraging results for random EU data, but *not* for standard model EU data. I know of no estimator that correctly identifies weighting functions regardless of the true probabilistic model generating the data; finding such an estimator would be a nice contribution to decision research.

However, meaningful preference measurement may not be possible without strong assumptions concerning the random part of decision behavior (Wilcox 2008; Blavatskyy and Pogrebna 2010; Wilcox 2011; Apesteguia and Ballester 2016). Some say that elicited certainty equivalents permit "nonparametric" identification and estimation of preferences (Gonzales and Wu 1999; Abdellaoui 2000; Bleichrodt and Pinto 2000; Abdellaoui, Bleichrodt and Paraschiv 2007) and many others repeat it (e.g. Prelec 1998; Luce 2000; Nielson 2003; Fox and Poldrack 2009; Wakker 2010). "Nonparametric" means different things, but many econometricians divide discussion of models in two parts: (1) a conditional expectation function, or perhaps a conditional median function, and (2) the error, the random part that remains once such a function has been removed in a way that makes the expectation (or median) of the error zero. In the preference measurement literature, scholars who say their estimation is "nonparametric" mean they are making few or no assumptions about the form of preference entities (utilities or values, and probability weights, and so forth) that appear in standard model conditional expectation functions. However, they routinely make the strong assumption of the standard model itself, and an old and well-developed alternative (the random preference model) has complicating consequences.

The essence of the standard model assumption is that the c.e.f. has an obvious interpretation—the intended interpretation being that of algebraic (deterministic) decision theory. Hendry and Morgan (2005, p. 23) argue that when we speak of model identification, we have things in mind beyond the original Cowles Foundation meaning—including "correspondence to the desired entity" and "satisfying the assumed interpretation (usually of a theory model)." Estimation of preferences from elicited certainty equivalents is complicated in just these senses. The standard model is but one probabilistic model assumption, and under a venerable and contemporary alternative—the random preference model—the c.e.f. in part reflects the underlying distribution of preferences within the individual, in ways that can mimic "the desired entity," the preference entity called the probability weighting function. I do not know whether certainty equivalents can nonparametrically identify such entities: This question needs a good answer. However, at this time there is certainly no rigorous reason to think elicited certainty equivalents free scholars of critical interpretive assumptions. As is true of discrete choices, we choose a probabilistic model the moment we interpret certainty equivalents.

References

Abdellaoui, M., 2000, Parameter-free elicitation of utility and probability weighting functions. Management Science 46, 1497-1512.

Abdellaoui, M., H. Bleichrodt and C. Paraschiv, 2007, Loss aversion under prospect theory: A parameter-free measurement. Management Science 53, 1659-1674.

Ahn, D. S. and T. Sarver, 2013, Preference for flexibility and random choice. Econometrica 81:341-361.

Apesteguia, and M. Ballester, 2016, Monotone stochastic choice models: The case of risk and time preferences. Journal of Political Economy (forthcoming).

Becker, G. M., M. H. DeGroot and J. Marschak, 1963, Stochastic models of choice behaviour. Behavioral Science 8, 41–55.

Blavatskyy, P. and G. Pogrebna, 2010, Models of stochastic choice and decision theories: Why both are important for analyzing decisions. Journal of Applied Econometrics 25, 963-986.

Bleichrodt, H. and J. L. Pinto, 2000, A parameter-free estimation of the probability weighting function in medical decision analysis. Management Science 46:1485-1496.

Bruhin, A., H. Fehr-Duda and T. Epper, 2010, Risk and rationality: Uncovering heterogeneity in probability distortion. Econometrica 78:1375-1412.

Butler, D. and G. Loomes, 2007, Imprecision as an account of the preference reversal phenomenon. American Economic Review 97, 277-297.

Eliashberg, J. and J.R. Hauser, 1985, A measurement error approach for modeling consumer risk preference. Management Science 31, 1-25.

Feller, W., 1971, An Introduction to Probability Theory and Its Applications Vol. 2, 2nd ed. New York: Wiley.

Fox, C. R. and R. A. Poldrack, 2009, Prospect theory and the brain. In P. Glimcher, C. Camerer, E. Fehr and R. Poldrack, eds., Neuroeconomics: Decision Making and the Brain. London, UK: Academic Press.

González-Velasco, E. A., 1995, Fourier analysis and boundary value problems. San Diego: Academic Press.

Gonzalez, R. and G. Wu, 1999, On the shape of the probability weighting function. Cognitive Psychology 38, 129-166.

Gul, F. and W. Pesendorfer, 2006, Random expected utility. Econometrica 74,121-146.

Halevy, Y., 2007, Ellsberg revisited: An experimental study. Econometrica 75, pp. 503-536.

Hendry, D. F. and M. S. Morgan, 2005, The foundations of econometric analysis. Cambridge: Cambridge University Press.

Hilton, R.W., 1989. Risk attitude under random utility. Journal of Mathematical Psychology 33, 206–222.

Hougaard, P., 1986, Survival models for heterogeneous populations derived from stable distributions. Biometrika 73:387-396.

Karni, E. and Z. Safra, 2016, A theory of stochastic choice under uncertainty. Journal of Mathematical Economics 63:164–173.

Krahnen, J. P., C. Rieck and E. Theissen, 1997, Inferring risk attitudes from certainty equivalents: Some lessons from an experimental study. Journal of Economic Psychology 18:469-486.

Loomes, G., P. Moffatt, R. Sugden, 2002, A microeconometric test of alternative stochastic theories of risky choice. Journal of Risk and Uncertainty 24, 103-130.

Loomes, G. and G. Pogrebna, 2014, Measuring individual risk attitudes when preferences are imprecise. Economic Journal 124, 569-593.

Loomes, G. and R. Sugden, 1995, Incorporating a stochastic element into decision theories. European Economic Review 39, 641-648.

Loomes, G. and R. Sugden, 1998, Testing different stochastic specifications of risky choice. Economica 65, 581-598.

Luce, R. D., 1997, Some unresolved conceptual problems in mathematical psychology. Journal of Mathematical Psychology 41:79-87.

________, 2000, Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches. London, UK: Erlbaum.

Navarro-Martinez, D., G. Loomes, A. Isoni and D. Butler, 2014, Boundedly rational expected utility theory. University of Warwick working paper.

Nelder, J. A. and R. Mead, 1965, “A simplex method for function minimization.” Computer Journal 7, 308–313.

Nielson, W. S., 2003, Probability transformations in the study of behavior toward risk. Synthese 135:171–192.

Nolan, J. P., 2017 (forthcoming), Stable Distributions—Models for Heavy Tailed Data. Boston: Birkhauser.

Pennings, J. M. E. and A. Smidts, 2000, Assessing the construct validity of risk attitude. Management Science 46:1337-1348.

Powell, M. J. D., 1992, A direct search optimization method that models the objective and constraint functions by linear interpolation. Technical Report DAMTP 1992/NA5, Department of Applied Mathematics and Theoretical Physics, University of Cambridge.

Prelec, D., 1998, The probability weighting function. Econometrica 66, 497-527.

Quiggin, J., 1982, A theory of anticipated utility. Journal of Economic Behavior and Organization 3, 323-343.

Regenwetter, M., J. Dana and C. P. Davis-Stober, 2011, Transitivity of preferences. Psychological Review 118:42–56.

Regenwetter, M. and A. A. J. Marley, 2001, Random relations, random utilities and random functions. Journal of Mathematical Psychology 45:864-912.

Seshadri, V., 1993, The inverse Gaussian distribution—A case study in exponential families. Oxford: Clarendon Press.

Tversky, A. and C. R. Fox, 1995, Weighing risk and uncertainty. Psychological Review 102, 269-283.

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

Vieider, F. M., M. Lefebvre, R. Bouchouicha, T. Chmura, R. Hakimov, M. Krawczyk and P. Martinsson, 2015, Common components of risk and uncertainty attitudes across contexts and domains: Evidence from 30 countries. Journal of the European Economic Association 13, 421-452.

von Winterfeldt, D., N.-K. Chung, R. D. Luce, and Y.-H. Cho, 1997, Tests of consequence monotonicity in decision making under uncertainty. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 406–426.

Wakker, P., 2010, Prospect Theory: For Risk and Ambiguity. Cambridge, UK: Cambridge University Press.

Wilcox, N., 2008, Stochastic models for binary discrete choice under risk: A critical primer and econometric comparison. In J. C. Cox and G. W. Harrison, eds., Research in Experimental Economics Vol. 12: Risk Aversion in Experiments pp. 197-292. Bingley, UK: Emerald.

Wilcox, N., 2011, 'Stochastically more risk averse:' A contextual theory of stochastic discrete choice under risk. Journal of Econometrics 162, 89-104.

Wilcox, N., 2015. Unusual estimates of probability weighting functions. Chapman University, Economic Science Institute Working Paper #15-10.

Wu, G. and R. Gonzalez, 1999, Nonlinear decision weights in choice under uncertainty. Management Science 45, 74-85.

Appendix: A brief Monte Carlo illustration of the problem

Simulated data sets for this brief Monte Carlo analysis of several estimation methods are based on the experimental design of Gonzalez and Wu (1999). Certainty equivalents were elicited from their subjects for $t = 1, 2, \ldots, 165$ distinct two-outcome prospects $(p_t, h_t; 1 - p_t, l_t)$. These were constructed by fully crossing fifteen distinct pairs of high and low outcomes (h_t, l_t) with eleven distinct probabilities p_t of receiving the high outcome h_t (and corresponding probabilities $1 - p_t$ of receiving the low outcome l_t). The eleven probabilities are $p_t \in \{.01, .05, .10, .25, .40, .50, .60, .75, .90, .95, .99\}$; and the fifteen high and low outcome pairs are $(h_t, l_t) \in \{(25{,}0), (50{,}0), (75{,}0), (100{,}0), (150{,}0), (200{,}0), (400{,}0), (800{,}0), (50{,}25), (75{,}50), (100{,}50), (150{,}50), (150{,}100), (200{,}100), (200{,}150)\}$. These same 165 prospects (88 simple prospects and 77 non-simple prospects) are the "input" to the simulated subjects I create in the Monte Carlo data sets. Let $Z = \{0, 25, 50, \ldots, 800\}$ denote the set of the nine distinct outcomes found in these 165 prospects.

Each simulated subject $s = 1, 2, \ldots 1000$ in the first data set is given a random EU certainty equivalent for each of the 165 prospects. Each subject s is endowed with parameters k^s and θ^s of the Gamma distribution p.d.f. as given in Example 1 of Section 1. The parameter k^s is drawn *once* for each subject from a Lognormal distribution with mean $E(k) = 0.75$ and variance $\mathrm{V}(k) \approx 0.16$. The parameter θ^s is then chosen (given the drawn k^s) so that $\mathbb{C}^{eu}(e^{-1}|k^s, \theta^s) = (1 + \theta^s)^{-k^s} = e^{-1}$. This endows each simulated subject s with a random EU c.e.f. having the fixed point e^{-1}, as is characteristic of the 1-parameter Prelec (1998) weighting function, but also creates heterogeneity in the degree of curvature of subjects' c.e.f.s. Then for each subject s, $t = 1{,}2, \ldots{,}165$ values x_t^s are independently drawn from the Gamma distribution with that subject's parameters k^s and θ^s. These create the 165 simulated elicited certainty equivalents $ce_t^s = \left[p_t {h_t}^{1/x_t^s} + (1 - p_t){l_t}^{1/x_t^s}\right]^{x_t^s}$ for each subject s. Repeating this 1000 times yields the "random EU" data set.

For comparison, I create a second data set of 1000 simulated subjects who are given standard EU certainty equivalents for each of the 165 prospects. Each simulated subject s is endowed with a *fixed* value x^s, drawn *once* for each subject from a Gamma distribution with the parameters $k = 0.75$ and $\theta = 2.79$. For each subject s, this x^s then creates 165 expected certainty equivalents $E(ce_t^s) = \left[p_t {h_t}^{1/x^s} + (1-p_t){l_t}^{1/x^s}\right]^{x^s}$: These are standard EU c.e.f.s, and one must somehow add standard model errors to them. To do this, notice that each expected certainty equivalent may be rewritten as a proportion of the interval $[l_t, h_t]$, that is as $\Delta_t^s = (E(ce_t^s) - l_t)/(h_t - l_t)$. One may then interpret this proportion as the mean of a Beta distribution on the interval (0,1) and define parameters of that Beta distribution as $\alpha_t^s = \varpi\Delta_t^s$ and $\beta_t^s = \varpi(1-\Delta_t^s)$. (Beta distributions may be parameterized in terms of their mean $\Delta \in (0,1)$ and an inverse dispersion parameter $\varpi > 0$, from which their more usual parameterization $\alpha \equiv \varpi\Delta$ and $\beta \equiv \varpi(1-\Delta)$ may be had. I chose $\varpi = 6$ to give the resulting simulated certainty equivalents ce_t^s in this simulated Standard EU data conditional variances resembling those found in the simulated Random EU data.) Then one may draw a beta variate y_t^s on (0,1) using these parameters, and the simulated certainty equivalents with their standard model error become $ce_t^s = l_t + (h_t - l_t)y_t^s$.

I consider four estimation methods. The first two methods use a standard RDU model of the c.e.f. of the ce_t^s, that is $E(ce_t^s|v^s, w^s) = (v^s)^{-1}\left[w^s(p_t)v^s(h_t) + \left(1 - w^s(p_t)\right)v^s(l_t)\right]$; the corresponding empirical model is then $ce_t^s = E(ce_t^s|v^s, w^s) + \varepsilon_t^s$. I make the standard assumptions about the error, those being $E(\varepsilon_t^s) = E(\varepsilon_t^s|p_t, h_t, l_t) = 0$, but also adopt the assumption of Bruhin, Fehr-Duda and Epper (2010) that $Var(\varepsilon_t^s)$ is proportional to $(h_t - l_t)^2$ for each subject. (This assumption happens to be true for the simulated Standard EU data.) This implies a "weighted error" $\epsilon_t^s = [ce_t^s - E(ce_t^s|v^s, w^s)]/(h_t - l_t)$, and the first two estimation methods optimize a function of these weighted errors.

The first estimation method combines a nonlinear least squares estimator with lean 1-parameter forms of the functions w^s and v^s, $w^s(p) = p^{\gamma^s}/\left[p^{\gamma^s} + (1-p)^{\gamma^s}\right]^{1/\gamma^s}$ and $v^s(z) = z^{\sigma^s}$. This is the estimation method of Tversky and Kahneman (1992): I'll call it NLS-M-L (for "<u>n</u>onlinear <u>l</u>east <u>s</u>quares, <u>m</u>oney errors, <u>l</u>ean parameterization"). The second estimation method combines a maximum likelihood estimator with the same $v^s(z) = z^{\sigma^s}$, but a more expansive 2-parameter weighting function $w^s(q) = \delta^s p^{\gamma^s}/\left[\delta^s p^{\gamma^s} + (1-p)^{\gamma^s}\right]$. The weighted error ϵ_t^s is assumed to have a Normal distribution with zero mean and constant variance. This estimation method is inspired by Bruhin, Fehr-Duda and Epper (2010), but I will always estimate at the individual subject level whereas they estimated finite mixture models of the subject population and included prospect-specific error variance terms (which cannot be done in the case of individual estimation). I'll call this method ML-M-C (for "maximum likelihood, money errors, common parameterization"). The power utility function, combined with some 2-parameter weighting function, is quite common in the literature on risk preference estimation.

The third method writes an estimating equation in utility rather than money terms, and the parameterizations of v^s and w^s are maximally expansive. There are nine distinct outcomes in Z, so there are nine distinct values of $v^s(z)$. Since the RDU value function is an interval scale, one can choose $v^s(0) = 0$ and $v^s(800) = 1$, leaving seven unique and distinct values of $v^s(z)$ as seven parameters to estimate. Similarly, the eleven distinct probabilities in the experiment become eleven distinct parameters $w^s(p_t)$ to estimate. Now linearly interpolate $v^s(ce_t^s)$ from the parameters $v^s(z)$ in the following manner. Let $lub(ce_t^s) = \min_{z \in Z} z \mid z \geq ce_t^s$ and $glb(ce_t^s) = \max_{z \in Z} z \mid z \leq ce_t^s$ be the least upper bound and greatest lower bound (among the nine outcomes in the experiment) on ce_t^s, with values given by the parameter values $v^s(lub(ce_t^s))$ and $v^s(glb(ce_t^s))$. Then define

$$\tilde{v}^s(ce_t^s) = \frac{[lub(ce_t^s) - ce_t^s]v^s\big(glb(ce_t^s)\big) + [ce_t^s - glb(ce_t^s)]v^s\big(lub(ce_t^s)\big)}{lub(ce_t^s) - glb(ce_t^s)},$$

a linear interpolation of $v^s(ce_t^s)$. This estimation method then assumes that the c.e.f. of $\tilde{v}^s(C_t^s)$ is the RDU of prospect t, that is $E(\tilde{v}^s(ce_t^s)|v^s, w^s) = w^s(p_t)v^s(h_t) + \big(1 - w^s(p_t)\big)v^s(l_t)$, and one may then think of $\tilde{v}^s(ce_t^s) - E(\tilde{v}^s(ce_t^s)|v^s, w^s)$ as a "utility error." Following Wilcox (2011), assume the variance of these utility errors is proportional to $[v^s(h_t) - v^s(l_t)\,]^2$. Then $\zeta_t^s = [\tilde{v}^s(C_t^s) - E(\tilde{v}^s(C_t^s)|v^s, w^s)]/[v^s(h_t) - v^s(l_t)\,]$ is a weighted utility error that becomes the object of nonlinear least squares estimation. I call this the NLS-U-E estimation (for "nonlinear least squares, utility errors, expansive parameterization"). It is inspired by Gonzalez and Wu's (1999) estimation method, though there are several differences between their method and this one (see Gonzalez and Wu 1999, pp.146-148, for details).

Finally, I consider an estimation method that may sidestep the issue identified in the text. Rather than taking $(v^s)^{-1}\big[w^s(p_t)v^s(h_t) + \big(1 - w^s(p_t)\big)v^s(l_t)\big]$ to be the conditional mean of ce_t^s, this last estimation method takes this to be the conditional median of ce_t^s:

That is, let $Med(ce_t^s|v^s, w^s) = (v^s)^{-1}\big[w^s(p_t)v^s(h_t) + \big(1 - w^s(p_t)\big)v^s(l_t)\big]$, and let weighted money errors be $\epsilon_t^s = [C_t^s - Med(C_t^s|v^s, w^s)]/(h_t - l_t)$. Although these errors have exactly the same form as the errors in the first two methods, the fact that we wish to estimate a conditional median function (rather than a c.e.f.) implies that least squares is not the appropriate estimator: Rather, we want a least absolute deviation or LAD estimator. Combined with the same lean parameterization used for the first method, I call this the LAD-M-L estimation (for "least absolute deviation, money errors, lean parameterization").

With the exception of the NLS-U-E estimation method, the well-known simplex algorithm of Nelder and Mead (1965) was used to optimize objective functions. For the NLS-U-E estimation method, I imposed monotonicity constraints on the estimated $v^s(z)$ and $w^s(p_t)$ (one difference versus Gonzalez and Wu 1999) and this requires a different optimization algorithm: Powell's (1992) COBYLA algorithm is used for this estimation instead. All estimations were performed using the SAS procedure "NLP" (nonlinear programming) in the SAS 9.4 version of the SAS/OR software.

Rather than providing tabular results of these four estimation methods as applied to the two data sets, I provide a sequence of eight figures. The features of each figure are identical. Estimated weighting functions for the first 250 subjects in each data set are plotted as quite thin, light greyscale lines on a black background: This has the effect of representing the behavior of each method as a light cloud of lines. A heavy light grey identity line shows the (linear, identity) weighting function of an EU agent; deviations from this line represent both sampling variability and possible bias in the estimations. Finally, a heavy dashed white line plots the mean estimated probability weight (across all 1000 subjects in each simulated data set) at each of the eleven values of p_t in the experimental design: Since all simulated subjects in both data sets are EU agents with identity weighting functions, deviations of this heavy dashed white line from the identity line illustrate the bias of each estimation method in each data set.

The figures come in pairs on each page that follows. Each page presents the results for one estimation method, with the top and bottom figures showing results for the Standard EU and Random EU data sets, respectively. The pair of Figures A1-a and A1-b show results for the NLS-M-L estimation method; Figures A2-a and A2-b show results for the ML-M-C method; Figures A3-a and A3-b show results for the NLS-U-E method; and Figures A4-a and A4-b show results for the LAD-M-L method.

None of these four estimation methods are bias-free for both the Standard EU and Random EU data sets, and this is the primary finding of this appendix. The method NLS-U-E is biased towards finding inverse-s probability weighting for *both* data sets: In the case of the Standard EU data I suspect this is because this method is just too parametrically expansive for the sample size. By contrast, the NLS-M-L and ML-M-C methods are virtually unbiased for Standard EU data, while they show the predicted bias when applied to the Random EU data. As speculated, the LAD-M-L method provides unbiased (and astonishingly tight) estimates for the Random EU data, but displays a pronounced bias in the Standard EU data in a direction opposite to inverse-s probability weighting. In sum, none of these four estimation methods are robust to the underlying source of randomness in the data generating process.

Figure A1-a: NLS-M-L Weighting Estimates, Standard EU Data

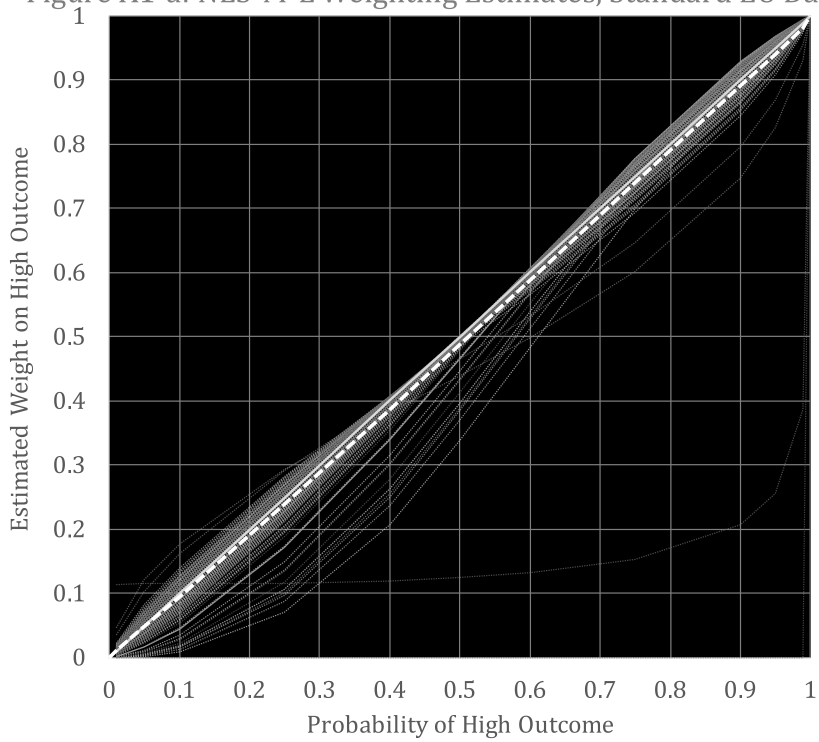

Figure A1-b: NLS-M-L Weighting Estimates, Random EU Data

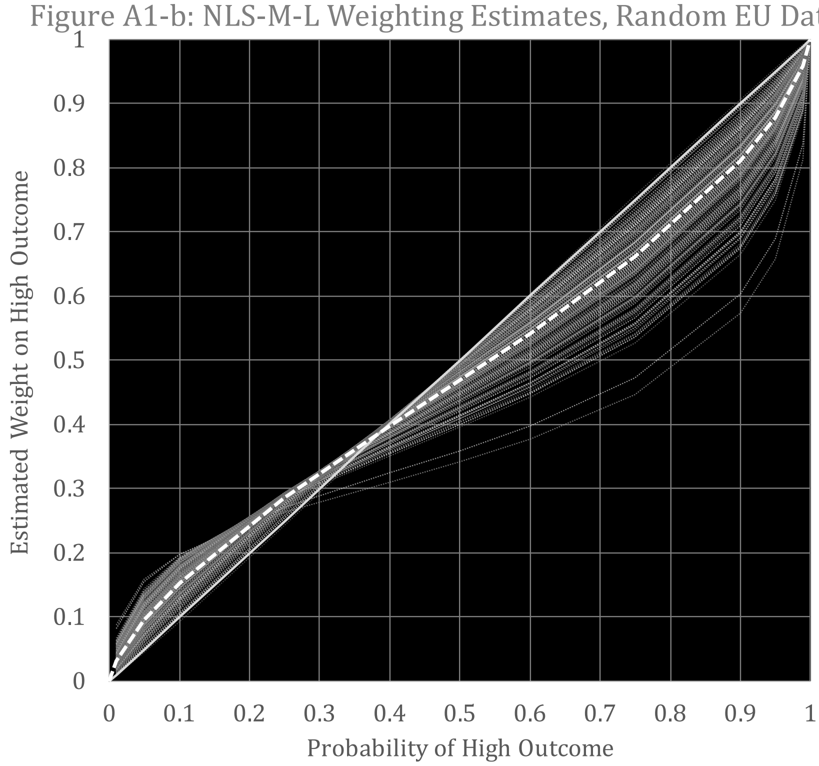

Figure A2-a: ML-M-C Weighting Estimates, Standard EU Data

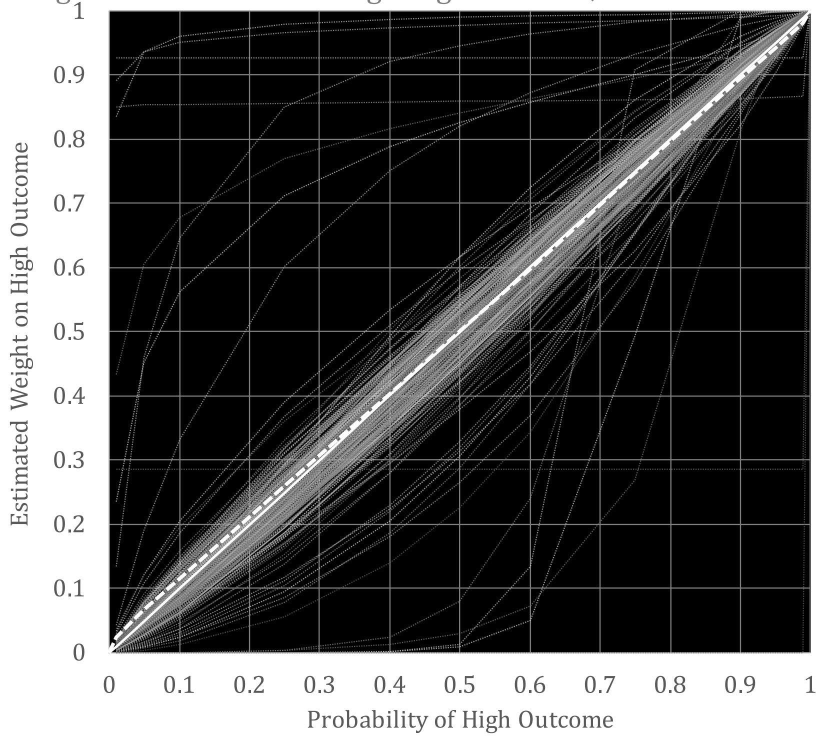

Figure A2-b: ML-M-C Weighting Estimates, Random EU Data

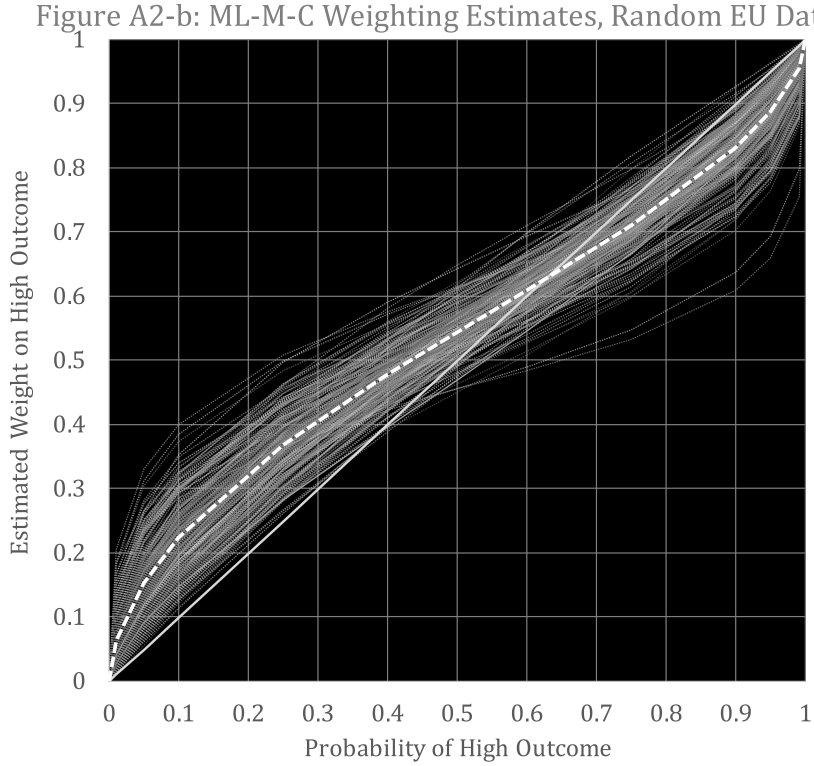

Figure A3-a. NLS-U-E Weighting Estimates, Standard EU Data

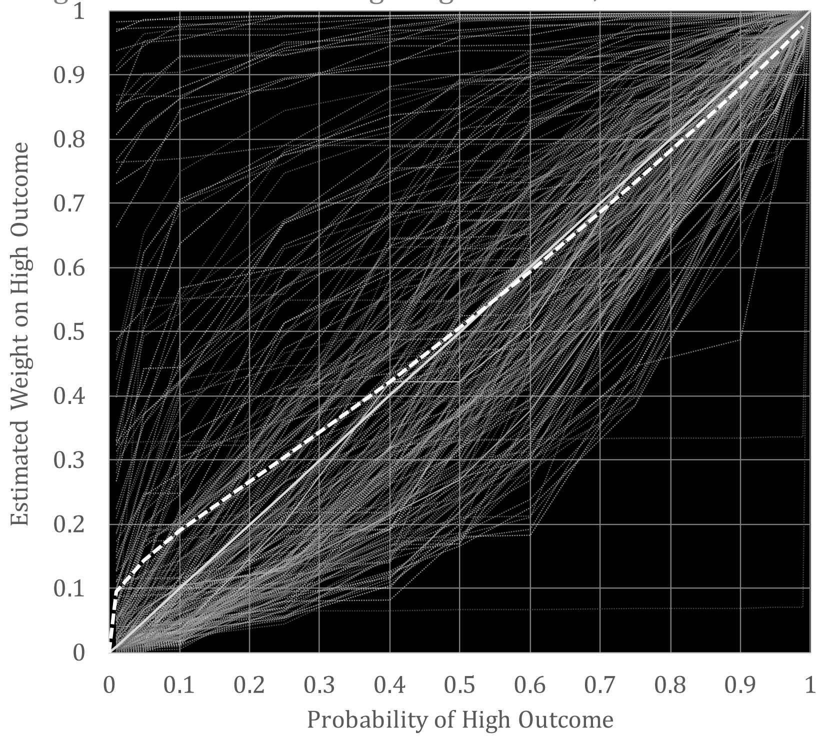

Figure A3-b. NLS-U-E Weighting Estimates, Random EU Data

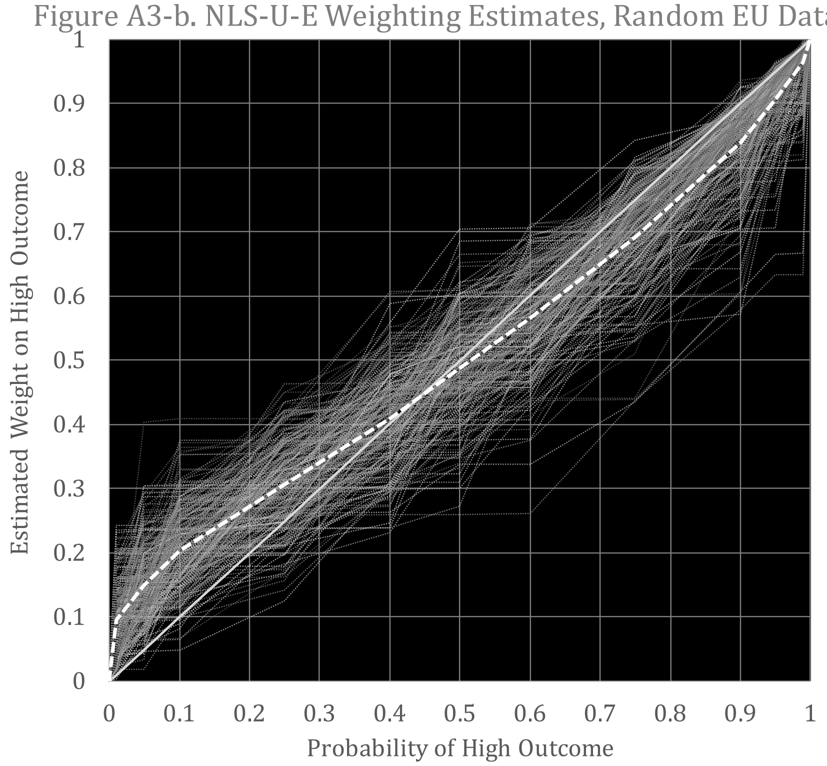

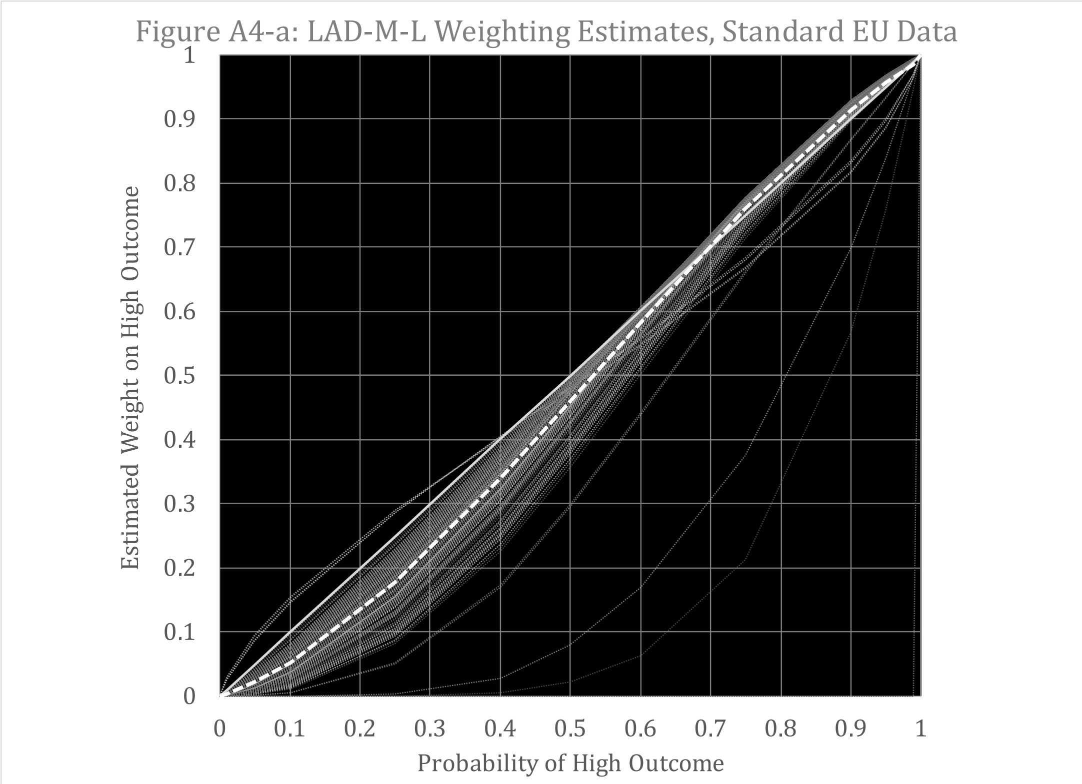
Figure A4-a: LAD-M-L Weighting Estimates, Standard EU Data
Estimated Weight on High Outcome
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Probability of High Outcome

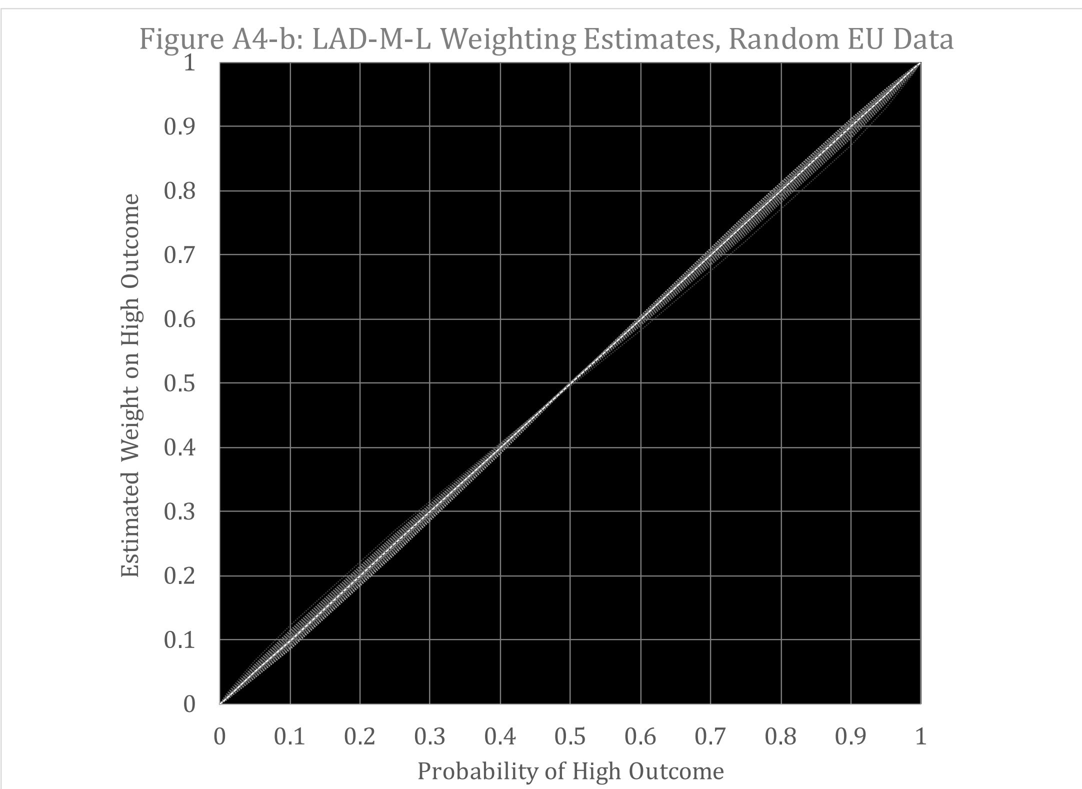
Figure A4-b: LAD-M-L Weighting Estimates, Random EU Data
Estimated Weight on High Outcome
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Probability of High Outcome