Utility in Reliability and Fusion and Statistics

Nozer D. Singpurwalla
The George Washington University
Washington, D.C. 20052, USA

We start with the following ***questions***:

- Why must decision makers in the DoD be <u>interested</u> in this topic?
- Similarly with control theorists, information fusers, censor fusers, survival analysts and statisticians?
- Should the utility of survivability be a concave, convex, S-shaped, or reverse S-shaped <u>increasing function</u> of survivability?

•To answer these questions, we need to set reliability and survivability in a broader context. We therefore start with the following as a preamble.

Managing Risk

- Reliability (and Survival Analyses), done by engineers and (biostatisticians), are yardsticks for <u>quantifying</u> the random nature of future life-times.

- We need to quantify this randomness for the purpose of <u>managing the risk</u> due to failure.

- By managing risk, we mean <u>making decisions</u> under uncertainty that minimize the loss due to failure.

- In the context of engineering, such decision making pertains to choosing between competing designs, or managing maintenance, or an acceptance/rejection of manufactured lots.
- In the context of biomedicine, such decision making pertains to treatment options (medication or surgery) and other choices that impact survivability and/or the quality of life.

Thus a purpose of reliability and survival analyses is to help make a certain class of decisions. The same is also true of censor fusion and control.

Coherent Decision Making.

- Decision making is a well studied topic, more so by those in business and economics, than those in statistics, though fundamental contributions to decision theory can be attributed to mathematicians and statisticians like Ramsey, von Neumann, Wald, and Savage Also some old timers like Blackwell, Rubin (Herman, that is) and Chernoff.
- Coherent decision theory rests on two pillars and one principle, namely, *probability* and *utility,* and the principle of *maximization of expected utility (MEU).*

The Structure of a Decision Problem

- In order to gain an appreciation of the utility of reliability, it is helpful if we review the structure of a decision problem, as seen by decision theorists.
- The structure consists of the following elements:
- $\boldsymbol{A} = (a_1, \ldots, a_n)$ is a set of n actions or choices;
- $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_k)$ is a set of k outcomes, $k > (=) < n$.
- $(U(\theta_1), \ldots, U(\theta_k))$ is a set of utilities, where $U(\theta_i)$ is the utility of the outcome θ_i; $i = 1, \ldots, k$.

- *D* denotes a *Decision Maker*, who is tasked with choosing one among the *n* possible actions.

- When *D* is contemplating which action to take, *D* is unaware as to which of the *k* outcomes (also known as the *states of nature*) will be true (or come about). That is *D* is uncertain about the happening of each θ_i.

- *D*'s uncertainty about the happening is encapsulated by $P(\theta_i)$, $i = 1, \ldots, k$, where $P(\theta_i)$ is *D*'s *personal probability* that when nature reveals itself, the outcome will be θ_i.

- The utility of θ_i, $U(\theta_i)$ encapsulates *D*'s personal preferences about the θ_is. Essentially, if $U(\theta_i) \geq U(\theta_j)$ then *D* prefers θ_i over θ_j, or is indifferent between θ_i and θ_j.

- In writing the above, we have tacitly assumed that $U(\theta_i)$ is a numerical quantity and all the θ_is can be *preference ranked* vis a vis. *D*'s preferences.

- Indeed, the axiomatics of the utility theory lead to the result that $U(\theta_i)$ – the utility of outcome θ_i – is a number between 0 and 1, both inclusive, and that –

Utility is a probability and it obeys the calculus of probability

- Since utilities are probabilities, the assessed utilities must *cohere.* The focus of this talk is how to coherently assess *D*’s utilities,.

The Utility of Reliability.

- With the reliability (and the survival function) interpreted as the unknown state of nature in a decision theoretic set-up, our next task is to assess the utility of reliability, namely the utility of, say $G(x|\boldsymbol{\theta})$.

- To do so, we anchor on two points, $G(x|\boldsymbol{\theta})=1$ and $G(x|\boldsymbol{\theta}) = 0$, and set the utilities of these anchor points, denoted $U[G(x|\boldsymbol{\theta})]$, as $U(1)=1$ and $U(0)=0$. This step is necessary part of utility theory as enunciated by von Neumann and Morgenstern.

- Having set up the above anchor points, we ask what is $U[G(x|\boldsymbol{\theta})]$ for any $G(x|\boldsymbol{\theta}) \in (0, 1)$, for D ?
- The essence of D's utility of $G(x|\boldsymbol{\theta})$ is based on the following gamble:

- Suppose D were to be offered the choice of obtaining $G(x|\boldsymbol{\theta})$ for sure, versus obtaining $G(x|\boldsymbol{\theta}) = 1$ with *chance* p, and $G(x|\boldsymbol{\theta}) = 0$ with *chance* $(1-p)$. Here chance ≡ *propensity*.
- Then D's $U[G(x|\boldsymbol{\theta})]$ is that value of p at which D is ***indifferent*** between the two choices: certainty versus uncertainty of obtaining 1 or 0.

- In other words, *D*'s utility is the propensity at which *D* is indifferent between the two offered choices. Thus the claim that utility is a probability, and utilities must cohere like how probabilities do.

- With $U[G(x|\boldsymbol{\theta})]$ anchored between 0 and 1, an archetypical utility function could take the forms shown below:

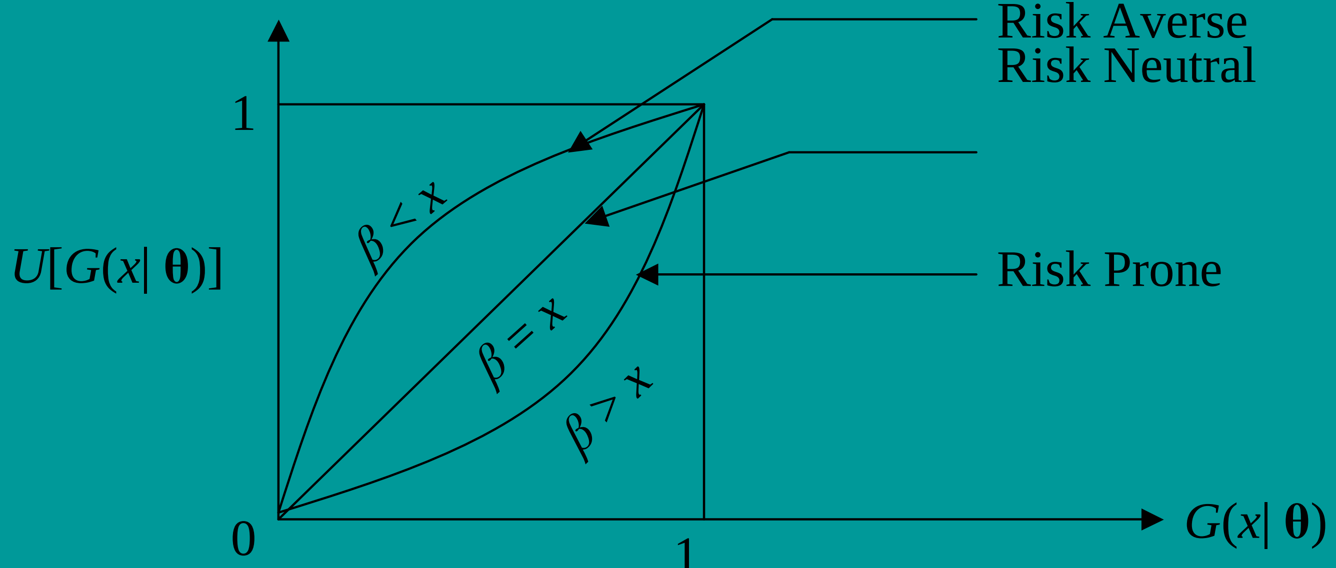

$U[G(x|\,\boldsymbol{\theta})] = [G(x|\,\boldsymbol{\theta})]^{\beta/x}$

Discussion on the Utility Function

- There are other possible shapes that the utility function can take. The concave and convex forms are just two extreme cases. For example, *D* can be risk prone for small values of the reliability and risk averse for large values. This would make $U(\cdot)$ and S-shaped function.
- In actuality, it may be difficult for *D* to arrive upon an indifference probability *p*. However, for any pre-chosen *p* and $G(x|\theta)$, it may be easier for *D* to choose certainty over uncertainty.
- Indeed, this is the basis of the utility elicitation approach we propose here.

- Even when *D* is able to arrive upon an indifference probability *p*, there is no assurance that *D* will be coherent in his (her) specifications. The process of elicitation must also ensure coherence. *Linear programming* and *regression* techniques have been proposed as a way of ensuring coherence.
- In our development thus far, the cost (or the ***disutility***) associated with achieving a certain reliability has not been considered.
- This needs to be incorporated, because there is a price to be paid for attaining a high reliability, or achieving a high survivability.

- The matter of disutility due to high survivability is particularly germane in *quality of life studies*.
- Certain medical procedures – such as surgery – may prolong life (i.e. increase survivability) but incur a disutility by way of costs and discomfort.
- An archetypical disutility function could be of the form

$$1-\exp\left[-\frac{\delta G(x\,|\,\boldsymbol{\theta})}{1-G(x\,|\,\boldsymbol{\theta})}\right],$$

so the *net utility of survivability* (or *reliability*) could be

$$[G(x\,|\,\boldsymbol{\theta})]^{\beta/x}-\left[1-\exp\left(-\frac{\delta[G(x\,|\,\boldsymbol{\theta})]}{1-[G(x\,|\,\boldsymbol{\theta})]}\right)\right].$$

There could be two approaches for the incorporation of disutility in our analysis.

1. Elicit a utility <u>without</u> regard to costs and inconveniences.

   Elicit a disutility <u>without</u> regard to benefits and pleasures.

   Combine the two elicitations above to obtain a *nett utility*

2. When eliciting utilities ask the subject to bear in mind his/her disutilities when declaring choices and preferences. Under 1 above, the elicited utilities will be monotonic; not necessarily so under 2.

Utility Assessment via Item Response Models.

- It was suggested to me by a colleague in France that the *Rasch Model* of Item Response Theory, popular in Quality of Life Studies, can be a potentially useful device for eliciting utilities.
- There is merit to this suggestion because the structure of the Rasch Model contains the germ of an idea which enables us to develop an Item Response Theory type model that may be more appropriate for assessing the utility of reliability.
- To see how, we start with an examination of the Rasch Model.

The Rasch Model: An Overview

- Let X_{ij} denote the binary response of an individual i, $i = 1,\ldots, N$, on item j, $j = 1,\ldots,n$.
- Set $X_{ij} = 1(0)$ if i's response to j is correct (incorrect).
- Let α_i – an unknown parameter – encapsulate i's innate *ability* or *capability* to respond correctly to items.
- Let β_j – an unknown parameter – encapsulate the inherent *difficulty* in responding correctly to item j.
- The parameters α_i and β_j are known as the ability and difficulty parameters, respectively.

- Then the Rasch Model

$$P(X_{ij} = x_{ij} \mid \alpha_i, \beta_j) = \frac{\exp(x_{ij}(\alpha_i - \beta_j))}{1 + \exp(\alpha_i - \beta_j)}$$

gives the probability that an individual with ability α_i yields x_{ij} as a response to an item (or question) whose difficulty is encapsulated by β_j; $x_{ij} = 1$ or 0.

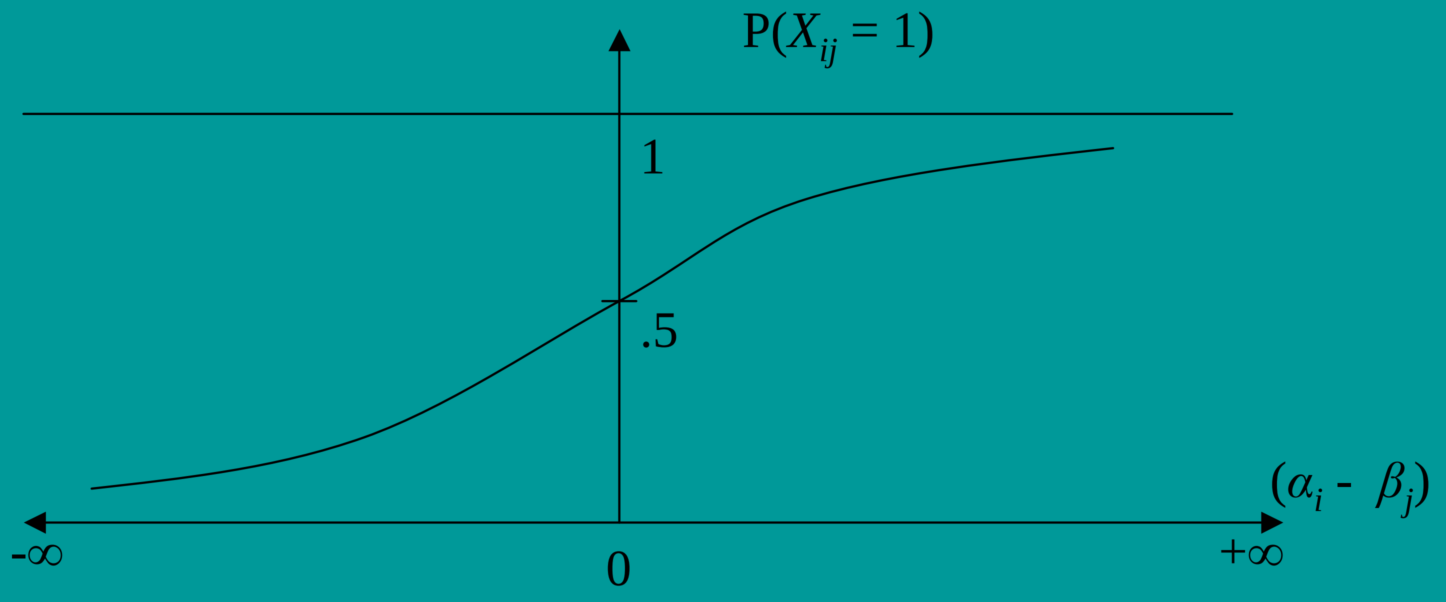

- The model captures the notion that greater the ability the higher the probability of a correct response.
- The model is commonly used in educational testing and sometimes in consumer preference evaluation.
- Its attractive features are that the model is indexed by the ***disparity*** between ability and difficulty, and that subjects are required to make only binary choices.

- The Item-Response model scenario of the Rasch Model and its essential features shares a commonality with the scenario of utiliy elicitation, wherein *D* is required to make a binary choice and can under certain circumstances, experience a difficulty in making his (her) choice.

Eliciting the Utility of Reliability

- We now come to the main purpose of this talk, namely, eliciting $U[G(x|\ \boldsymbol{\theta})]$, D's utility of reliability $G(x|\ \boldsymbol{\theta}) \in [0, 1]$, with anchor points $U(0) = 0$ and $U(1) = 1$. For convenience set $G(x|\ \boldsymbol{\theta}) = c$, $0 \leq c \leq 1$.
- In order to obtain $U(c)$, for $0 < c < 1$, D is asked to make a choice between:

  i) Receiving c for sure, or

  ii) Receiving 1 with propensity p, and 0 with propensity $(1-p)$.
- Recall that $U(c)$ is that value of p, say p^* , at which D is <u>indifferent</u> between the two choices. Our goal is to elicit p^* from D. To find p^* we may proceed in one of the two directions:

Strategy **1 –** *The Conventional Approach*

- Fix c, and interrogate D over a range of values of p, until D is able to converge upon a p^* at which D is indifferent between the choices.
- This strategy is cumbersome – if not difficult – to implement. Furthermore it could also result in D being incoherent over different values of c. Finally, in extracting a p^* from D, there is no measure of uncertainty about the p^*.
- What if D is indifferent over a range of values of p?

Strategy 2 – *Endowing a Probability of Choice*

- Fix $c = c_i$ and $p = p_{ij}$ and ask *D* to make a binary choice between the p_{ij}-gamble and the sure-c_i, for $j = 1,\ldots, n_i$, say.
- Set $X_{ij} = 1$, if *D* opts for the p_{ij}-gamble; $X_{ij} = 0$, otherwise. Repeat the above for i = 1,…, *k*.
- Endow the elicitation set-up with a probability of choice model, involving *c*, *p*, and other parameters – something akin to the Rasch Model – and using X_{ij}, $i = 1,\ldots, k$, $j = 1,\ldots, n_i$ as data, estimate the model parameters, using Bayesian or other methods.
- For any *c*, find that *p* at which the probability of choosing the sure-*c* (or the *p*-gamble) is 0.5.

Developing a Probability of Choice Model

- In order to develop a meaningful probability of choice model for *D*, we need to explore aspects of *D*'s choice processes, in particular how easy or hard it is for *D* to make a choice between the gamble and the sure thing. Our motivation for doing so is prompted by the ability-difficulty disparity used in the Rasch Model.
- The following are some *boundary conditions* for D's choices:

- When $p = 1$, D will choose $X = 1$, for all $c < 1$;
- When $p = 0$, D will choose $X = 0$, for all $c > 0$;
- When $p = c$, a *risk neutral* D will chose $X = 1$ or $X = 0$ with probability ½.
- By contrast, when $p = c$, a *risk prone* (*averse*) D will choose $X = 1$ with probability > (<) 0.5.
- Let $P(X = 1) = \pi$ denote the probability of D choosing the gamble over sure c.

Note: We have here to probabilities, p and π. The former is a propensity and the latter a personal probability of D.

- The illustration below highlights the boundary conditions of the choices of a risk neutral but coherent *D*.
- Here, *D*'s indifference between choices tantamounts to π being a half.

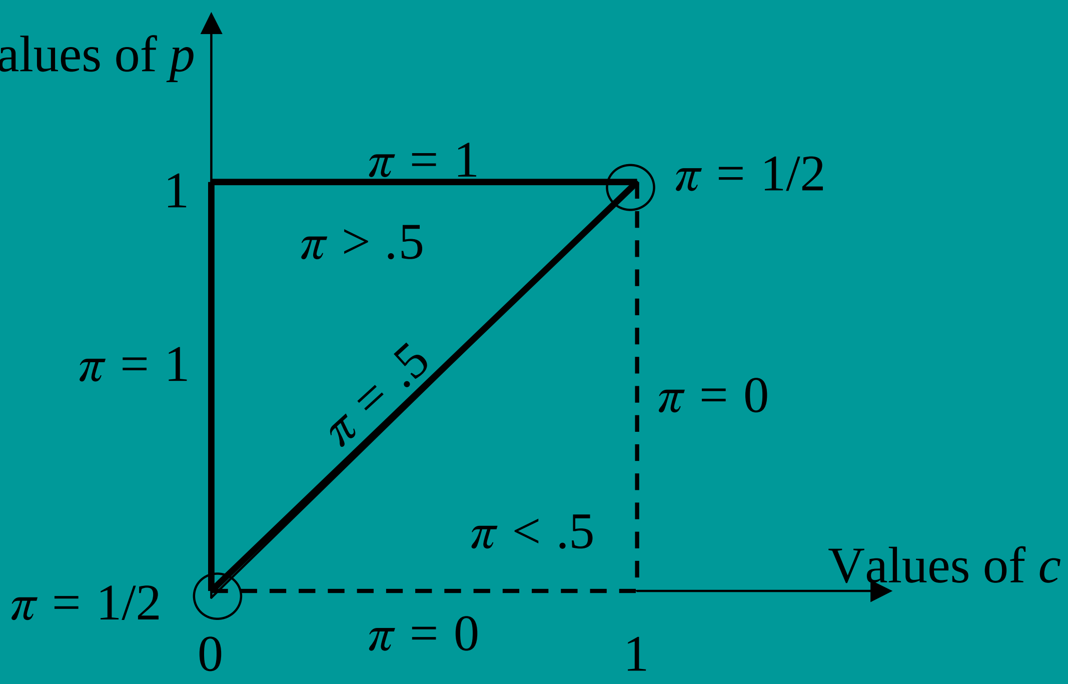

- Clearly, *D*'s choice are the <u>easiest</u> when p =1 or 0, irrespective of what c is.

- *D*'s choices <u>get difficult</u> as p and c get close to each other becoming most difficult when $p = c$.

- Making the decision to be indifferent is, in this context, a difficult one.

- The above features motivate us to suggest that a model for *D*'s choice probabilities possess the features illustrated below, for some $\beta \geq 0$.

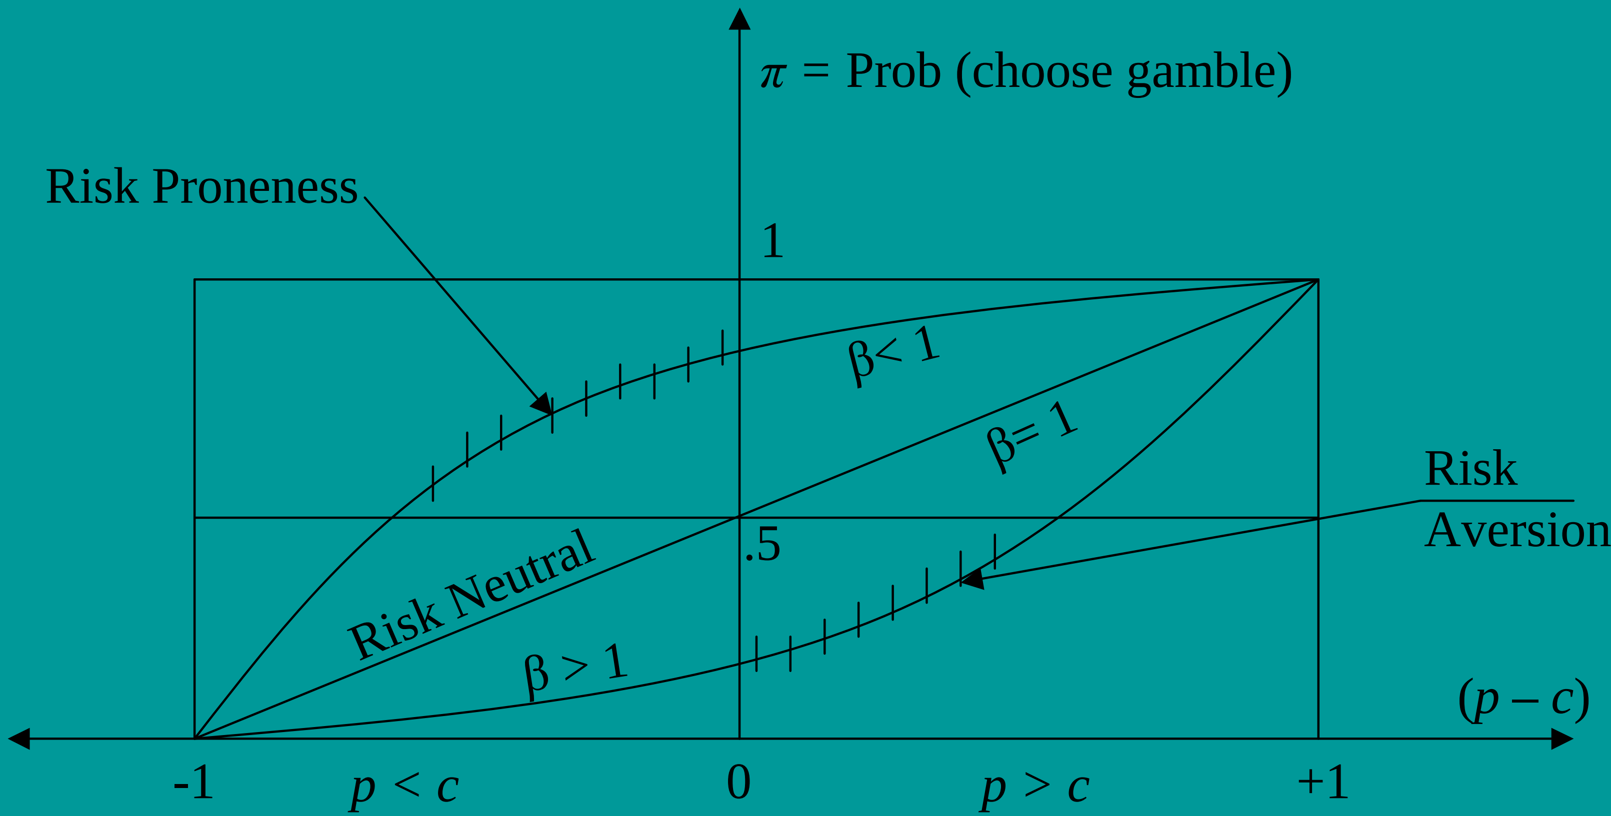

- The hatched segments of the above choice probability curves characterize the feature of *D*'s *Risk Aversion* or *Risk Proneness*. The diagonal line encapsulates *Risk Neutrality.*

- A relationship between π and $(p - c)$ which serves as a probabilistic model for D's choice of $(X=1)$ can be of the form

$$P(X=1 \mid \beta; p, c) = \left(\frac{(p-c)+1}{2} \right)^{\beta} \stackrel{def}{=} \pi$$

  Which for $\beta = 1$ is linear in $(p - c)$, and concave (convex) in $(p - c)$ for $\beta < (>)\ 1$.

- An enhancement of the above model to incorporate varying degrees of risk aversion and proneness – see the next three illustrations – is achieved by the inclusion of an additional parameter $\alpha > 0$, so that

$$P(X=1 \mid \alpha, \beta; p, c) = \left(\frac{(p-c)^{\alpha}+1}{2} \right)^{\beta} \stackrel{def}{=} \pi$$

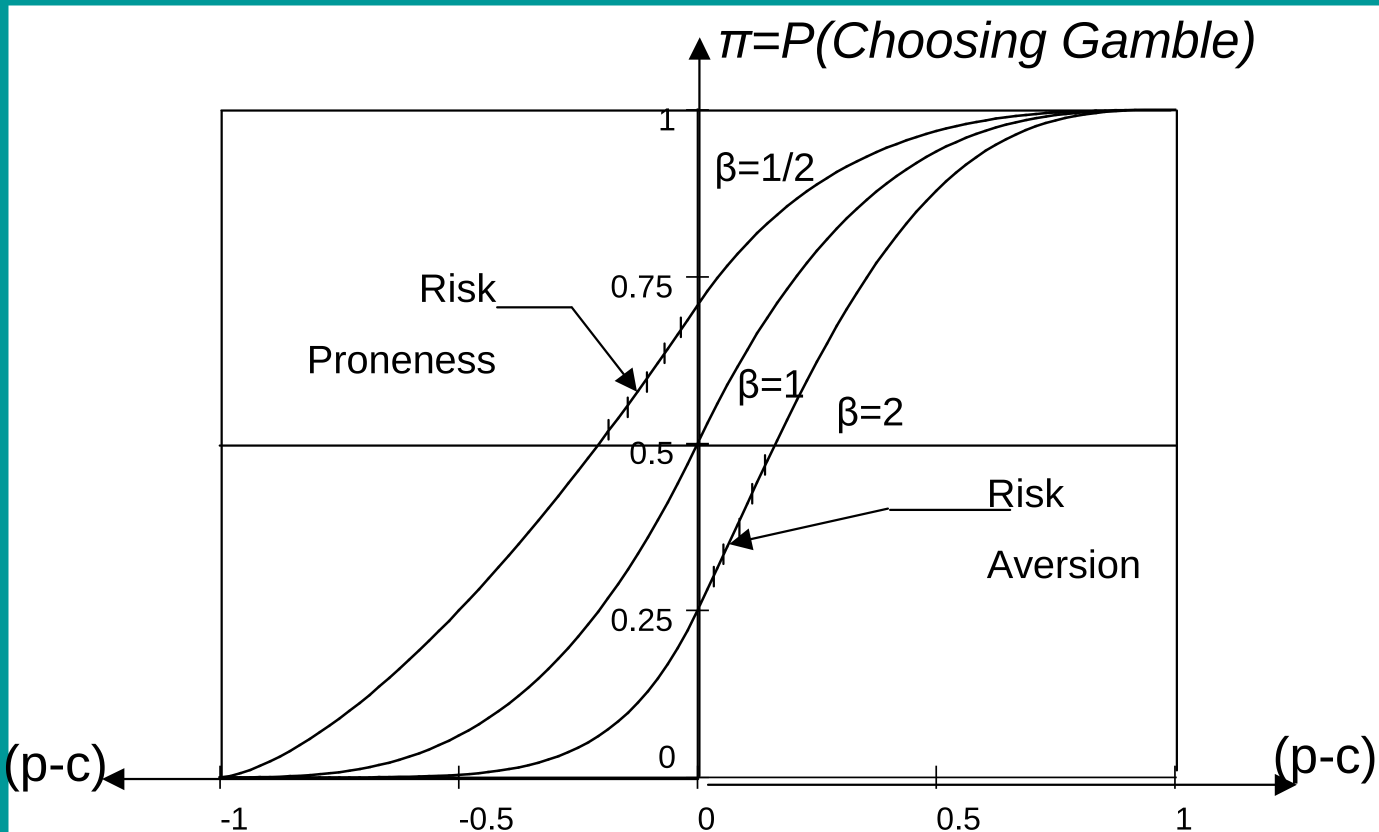

$$\pi = \left(\frac{(p-c)^{1/3} + 1}{2} \right)^{\beta}$$

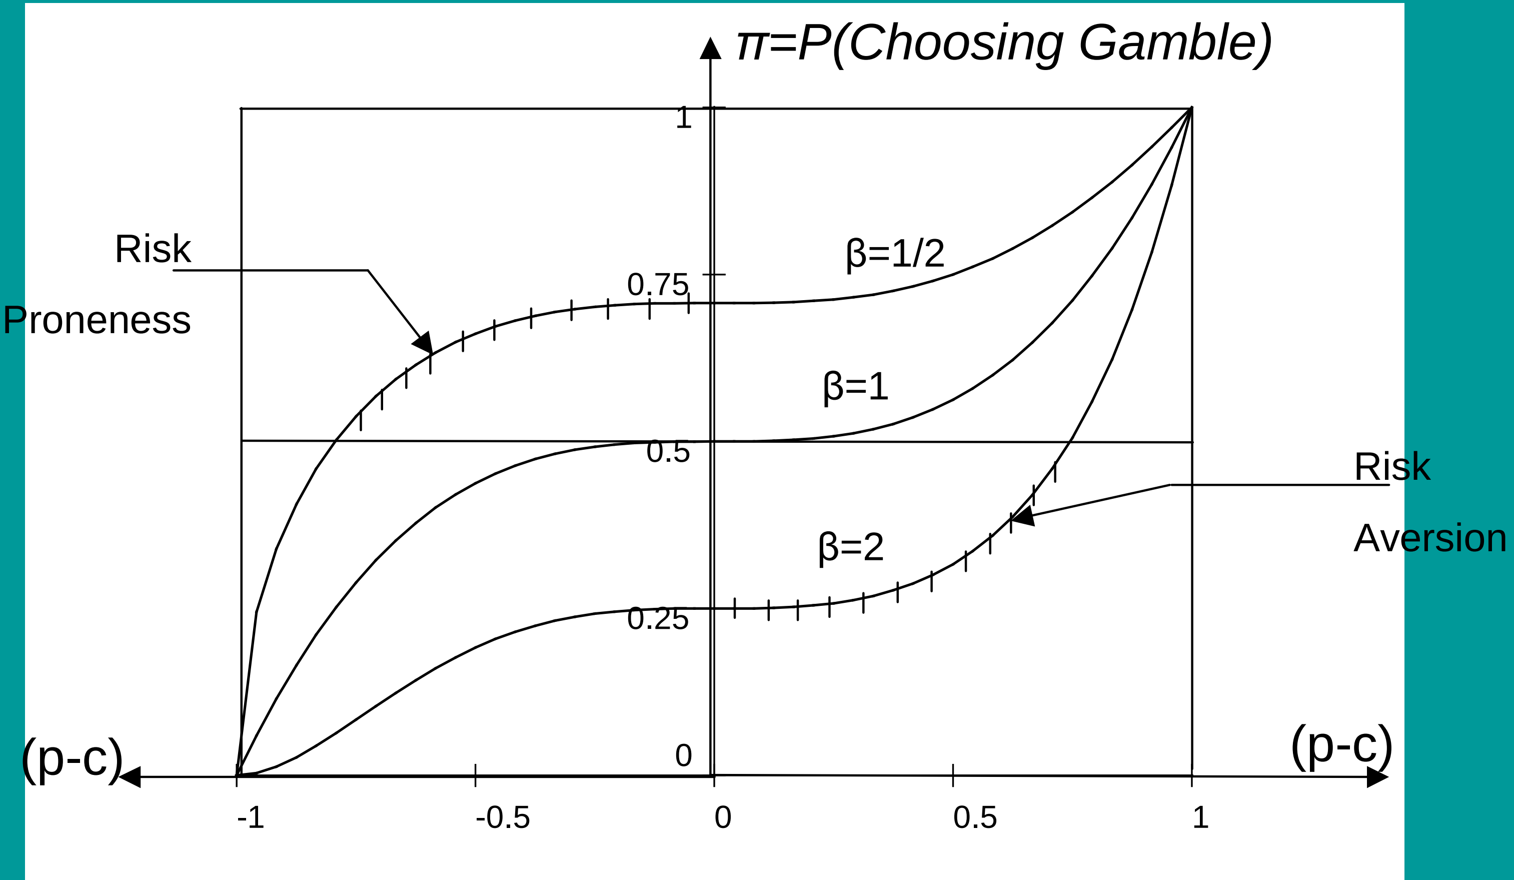

$$\pi = \left(\frac{(p-c)^3 + 1}{2} \right)^{\beta}$$

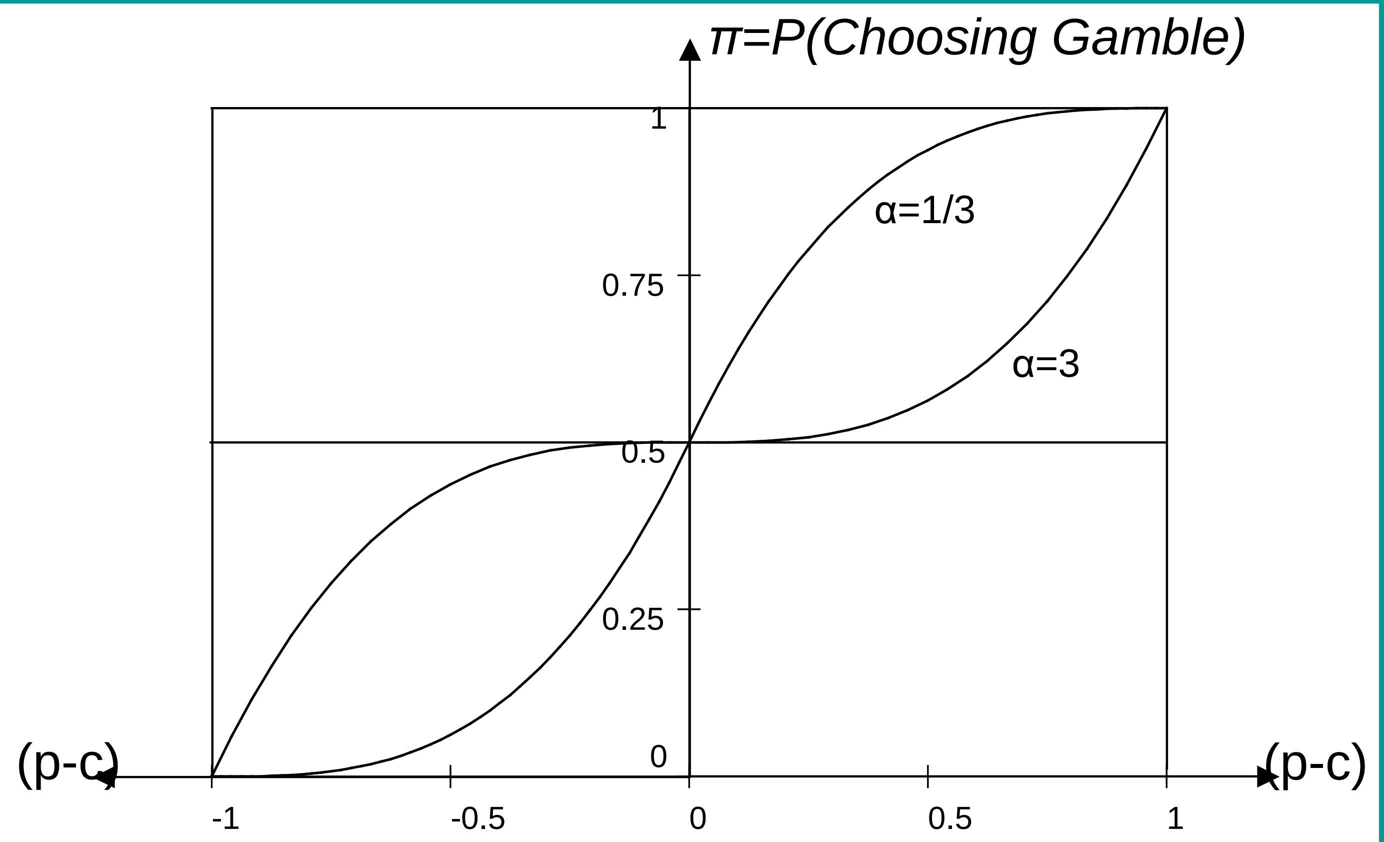

$$\pi = \frac{(p-c)^{\alpha} + 1}{2}$$

- To summarize, our proposed Item Response (or Choice Theory) based model for eliciting utilities, when the state of nature *c* is confined between 0 and 1 – such as reliability and survivability of an item – is of the form:

$$P(X=1 \mid \alpha, \beta; p, c) = \left(\frac{(p-c)^{\alpha}+1}{2} \right)^{\beta},$$

where α and β are unknown parameters, and *p* the propensity of a gamble that yields 1 with probability *p*, and 0 with probability (1 – *p*), versus the choice of a sure *c*.

Model Refinement

- Whereas the proposed model for *D*'s responses is intuitively appealing, it suffers from a technical deficiency. Specifically, in order to avoid complex roots, only certain values of α and β are admissible. Such restrictions pose technical difficulties with statistical inferences about α and β.

- In view of the above, we consider a refinement of the model, wherein

$$P(X=1\,|\,\alpha,\beta;p,c)=\left(\frac{1+sign(p-c)\,|\,p-c\,|^{\alpha}}{2}\right)^{\beta},$$

where $sign(x)$ = -1 (+1) [0] if x <(>)[=] 0.

Deploying the Elicitation Model.

- Our aim is to find that value of $(p - c)$ for which *D*'s, $P(X=1|\alpha, \beta; c, p) = ½$; i.e. *D* is indifferent between a certain *c* and a *p*-gamble.
- For this we need to know *D*'s α and β, and to assess these we subject *D* to a series of questions pertaining to *D*'s choice of certainty vs. a gamble.
- Specifically, we fix *c* at say c_i, and choose a range of values of *p*, say $p_{i1}, p_{i2}, \ldots, p_{ini}$ for $n_i \geq 1$.
- Then for each (c_i, p_{ij}) combination, $j = 1, \ldots, n_i$, we elicit from *D* a response $X_{ij} = 1$ or 0.

- Using (c_i, p_{ij}, X_{ij}) as inputs, $j = 1,\ldots, n_i$, we estimate D's model parameters as $\underline{\alpha}_i$ and $\underline{\beta}_i$. The method of maximum likelihood is a good illustrative vehicle.
- Plugging in $\underline{\alpha}_i$ and $\underline{\beta}_i$ in the equation below, we solve for $w_i = (p - c_i)$ in

$$\left(\frac{1 + sign(w_i)\,|\,w_i\,|^{\underline{\alpha}_i}}{2}\right)^{\underline{\beta}_i} = \frac{1}{2},$$

as

$$\underline{w}_i = sign(\underline{\beta}_i - 1)[sign(\underline{\beta}_i - 1)(2^{1-\frac{1}{\underline{\beta}_i}} - 1)]^{\frac{1}{\underline{\alpha}_i}};$$

$\underline{w_i} \in [-1, +1]$.

- The *utility* of c_i, $U(c_i)$, is of the form

$$U(c_i) = \begin{cases} \min(1, c_i + \underline{w}_i) & \text{if } \underline{w}_i > 0, \\ \max(1, c_i + \underline{w}_i) & \text{if } \underline{w}_i < 0, \\ c_i & \text{if } \underline{w}_i = 0 \end{cases}$$

- Thus *D* is risk prone (averse) [neutral] for c_i when $\underline{w}_i$ < (>) [=] 0.
- We repeat the above exercise for a range of values of *c*, namely, c_1, c_2, …, c_i,…, c_k, obtaining *D*'s utility for reliability (or survival) as $U(c_1)$, $U(c_2)$,…, $U(c_k)$.

Bayesian Analysis – An Outline.

- A simple – albeit naïve – strategy for performing a Bayesian analysis of the proposed elicitation model would be to assign independent gamma priors for α and β, say $\Pi(\alpha)$ and $\Pi(\beta)$.
- Then given the input (c_i, p_{ij}, X_{ij}), $j = 1,\ldots, n_i$, we may obtain the joint posterior of α, β, say $\Pi(\alpha_i, \beta_i; \cdot)$ via $L(\alpha_i, \beta_i; \cdot)$ the *likelihood*

$$\prod_{j=1}^{n_i}[.5+.5sign(p_{j`}-c_i)\,|p_j-c_i|^{\alpha}]^{\beta X_{ij}} \cdot$$

$$[1-(.5+.5sign(p_{j`}-c_i)\,|\,p_j-c_i\,|^{\alpha})^{\beta}]^{(1-X_{ij})}$$

- Once the above is at hand, we may assess $w_i = (p - c_i)$ via $\hat{w}_i$, where $\hat{w}_i$ is the solution to

$$\int_0^{\infty}\int_0^{\infty} \pi(\alpha_i, \beta_i; \cdot)\left[\left(\frac{1 + sign(w_i)\,|\,w_i\,|^{\alpha_i}}{}\right)^{\beta_i} - \frac{1}{2}\right] d\alpha_i\, d\beta_i = 0.$$

- This has been done numerically by Philip Wilson (who also proposed the model refinement considered herein).
- With $\hat{w}_i$ at hand, $i = 1,\ldots, k$, we may obtain as before the utilities $U(c_1),\ldots, U(c_k)$.

Discussion – Open Questions

- The general architecture proposed here has the advantage that it simplifies *D*'s task from pinning down an indifference propensity p to one of making binary decisions.
- However the analysis is guided by a specific *choice model* whose structure may be too rigid in actuality.
- More important, there is no assurance that the elicited $U(c_i)$, $i = 1,\ldots, k$ are monotone increasing in c_i, and also coherent in the sense that utilities should be, vis a vis the assumption of independence.

- Finally, the matter of questionnaire design in terms of what the n_i and the k should be remains to be addressed. Whereas increasing the n_i and the k enhance our appreciation of the utility function, large values of k will also increase the chances of encountering incoherence.