

# ASSURING FINITE MOMENTS FOR WILLINGNESS TO PAY IN RANDOM COEFFICIENT MODELS

Andrew Daly, RAND Europe and ITS Leeds  
Stephane Hess, ITS Leeds  
Kenneth Train, University of California, Berkeley

## Abstract

Random coefficient models such as mixed logit are increasingly being used to allow for random heterogeneity in willingness to pay (WTP) measures. In the most commonly used specifications, the distribution of WTP for an attribute is derived from the distribution of the ratio of individual coefficients. Since the cost coefficient enters the denominator, values of the cost coefficient that are close to zero induce large values of WTP, with unboundedly large values of WTP resulting from cost coefficients arbitrarily close to zero. In this paper, we identify a criterion to determine whether the distribution of WTP has finite moments. Using this criterion, we show that some popular distributions used for the cost coefficient in random coefficient models, including normal, truncated normal, uniform and triangular, imply infinite moments for the distribution of WTP, even if truncated or bounded at zero. We also point out that relying on simulation approaches to obtain moments of WTP from the estimated distribution of the cost and attribute coefficients can mask the problem by giving finite moments when the true ones are infinite. We identify several approaches that analysts can utilise to assure that the distribution of WTP has finite moments.

## 1. Introduction

Discrete choice models with random coefficients are often used to estimate the distribution of willingness to pay (WTP) for attributes of goods or services. For example, mixed logit models (see e.g. McFadden & Train, 2000; Hensher & Greene, 2003; Sillano & Ortúzar, 2005), which are the most widely used form of random coefficient choice models, have been used extensively in transportation research to estimate various types of WTP, including travellers' WTP for changes in travel time. Under the standard approach, the analyst specifies the distribution of the cost and attribute coefficients and estimates the parameters of this distribution. The distribution of WTP is then derived from the estimated distribution of the coefficients. This practice, while accurate in theory, can be problematic in practice because the cost coefficient enters the denominator of WTP. A value of the cost coefficient that is arbitrarily close to zero results in an arbitrarily large WTP. As a result, the moments of the WTP distribution might not exist for a given distribution of the cost coefficient. However, it is these moments of the WTP distribution, and especially the mean, that are of crucial interest, for example, in policy appraisal. The analyst must be especially careful, therefore, when specifying the distribution for the cost coefficient to assure that the resulting distribution of WTP is useful and meaningful.

In this paper, we provide a theorem that identifies, under certain conditions, whether or not the moments of WTP exist for a given distribution of the cost coefficient. Using the theorem, we show that many commonly used distributions for the cost coefficient, including normal, uniform, and

triangular (even if truncated or bounded at zero) imply that the distribution of WTP has undefined (i.e. infinite) moments. We point out, and illustrate with examples, that simulation of the WTP distribution from draws of the cost and attribute coefficients can mask the problem, providing (incorrect) finite moments even though the true moments are infinite. The problem is only masked further when relying on censoring of draws during simulation.

We conclude by discussing several ways that the analyst can assure finite moments for WTP distributions in random coefficient models. These methods include (1) using a cost coefficient distribution that meets the requirements of the theorem, (2) re-parameterising the model in WTP-space, so that the distribution of WTP is specified directly rather than derived, and (3) using finite distributions rather than continuous ones, with each mass point for the cost coefficient bounded away from zero.

The remainder of this paper is organised as follows. Section 2 gives the theorem and applies it to several well-known distributions. Section 3 discusses simulation of WTP moments and shows how simulation obtains finite moments for distributions whose moments are known not to exist. Section 4 provides another caution, namely, that attempts to avoid problems by using the distribution of individual-level conditional means as an indication of the population distribution can mask undefined moments because it misapplies the relation between conditional and unconditional distributions. Section 5 provides a number of ways for the analyst to assure finite moments of WTP. Section 6 concludes.

## 2. Existence of moments of WTP distribution

In a typical model specification, the utility function for an alternative is a linear function of the various attributes of an alternative, such as travel time and travel cost, multiplied by their coefficients. For example, the utility obtained by a given person from alternative  $j$  may be represented as:

$$U_j = \theta c_j + \beta a_j + \text{other terms} + \varepsilon_j \quad (1)$$

where  $c_j$  is the cost of the alternative and  $a_j$  is a non-cost attribute. In the discussions in this paper, we focus on a specification that interacts the coefficients and the attributes in a purely linear fashion, but the issues highlighted here apply similarly to specifications incorporating non-linear interactions.

A random coefficients specification is obtained by considering  $\theta$  and  $\beta$  to be random with a specified distribution in the population, whose parameters are estimated. The final error term  $\varepsilon$  is also considered random, most commonly iid extreme value, so that the model is a mixed logit.

The WTP for an attribute of alternative  $j$  is, by definition, the ratio of the marginal utility of the attribute to the marginal utility of its cost, which in the case of linear-in-parameters utility is the ratio of the attribute coefficient to their cost coefficient:

$$\text{WTP} = \frac{\partial U_j / \partial a_j}{\partial U_j / \partial c_j} = \frac{\beta}{\theta}$$

The distribution of  $\beta$  and  $\theta$  in the population induces a distribution of WTP. The question that we are examining is whether the implied distribution of WTP has finite moments, which is certainly a desirable property and might even be considered a necessary property for a meaningfully specified model. For use in appraisal, the mean is clearly essential.

Our theorem is applicable to situations in which  $\beta$  and  $\theta$  are independent. If  $\beta$  and  $\theta$  are not independent, then a more complex analysis is needed than our theorem provides (which constitutes an important topic for future research).<sup>1</sup> If  $\beta$  and  $\theta$  are independent, then the moments of WTP are the product of the moments of  $\beta$  and  $1/\theta$ . In particular, the  $k$ -th moment of WTP is

$$E[(\beta/\theta)^k] = E(\beta^k) * E(1/\theta^k)$$

The analyst in specifying the distribution of  $\beta$  can directly assure that  $E(\beta^k)$  exists. The relevant question, then, is whether the inverse moments  $E(1/\theta^k)$  exist for a specified distribution of  $\theta$ . The following theorem provides the necessary guidance.

If a random variable  $\theta$  has an absolutely continuous probability density  $f(\theta)$ , then for any positive integer  $k$  the inverse moment  $E(1/\theta^k)$  exists if and only if  $\lim_{\theta \rightarrow 0} \frac{f(\theta)}{\theta^h}$  exists for some  $h > k - 1$ .

The proof is given in the Appendix. The theorem has the following Corollary, also proved in the Appendix..

If a random variable  $\theta$  has an absolutely continuous probability density  $f(\theta)$  defined on the positive half-line and  $\lim_{\theta \rightarrow 0} f(\theta) > 0$ , then none of the inverse moments  $E(1/\theta^k)$  exists.

The Corollary has the following applications to some commonly used distributions.

- For a uniform distribution bounded at 0, all inverse moments are non-existent.
- For a normal distribution truncated at 0, as with the uniform distribution, all the inverse moments are non-existent.

The main theorem can be used to investigate other commonly used distributions.

- For a triangular distribution bounded at 0,  $f(\theta) = \lambda\theta$  for  $\theta$  less than the mean. Then  $f(\theta)/\theta^1 = \lambda$  so that the limit is finite for  $h = 1$  and the mean of  $1/\theta$  exists. However, for

---

<sup>1</sup> In the special case of nonindependent  $\beta$  and  $\theta$  being jointly normal, our results still apply, since joint normals can be re-expressed in terms of independent normals. In particular,  $\beta$  can be re-expressed as  $\beta = \alpha_1\beta^* + \alpha_2\theta$  where  $\beta^*$  and  $\theta$  are independent normals. Then wtp is  $\beta/\theta = \alpha_1(\beta^*/\theta) + \alpha_2$ , and the distribution of wtp depends on the distribution of two independent terms. In general, however, no such transformation is available.

$h > 1$ ,  $f(\theta)/\theta^h = \lambda/\theta^{h-1}$  and the limit does not exist. The variance of  $1/\theta$  and its higher moments do not exist.

- Both lognormal and Johnson's Sb distribution have all inverse moments in their basic specifications (i.e. with the domains between 0 and infinity for the lognormal, and 0 and 1, or any other positive number, for the Johnson Sb). Of course, this result is already known for the lognormal, but the theorem can also be used to show it: since the lognormal density approaches zero faster than any power function, the limit exists for all  $h$  and so all inverse moments exist. The Johnson's Sb distribution tends to a constant times the lognormal as it approaches zero, which implies that its inverse moments also exist.
- For the gamma distribution with shape parameter  $\nu$ , inverse moments exist for  $k < \nu$  but no higher. This is because the gamma distribution tends to a constant times  $\theta^\nu$  as its argument  $\theta$  approaches zero. The negative exponential distribution is a gamma distribution with  $\nu = 0$ , and so it has no inverse moments.

Further, no inverse moments exist for distributions with strictly positive density at zero. This result applies to the commonly used normal distribution, as well as to bounded distributions (e.g., uniform, triangular, Johnson's Sb, and lognormal) if they are offset such that they straddle 0.

The theorem also implies that if  $f(\theta) = 0$  in an interval around zero, then all inverse moments exist. This result means that all inverse moments can be assured to exist by setting bounds on the distribution such that the density is zero within an interval around zero. These bounds can be set a priori (if some means of determining them is available) or estimated. For example,, the uniform and triangular distributions can be used and still assure finite inverse moments by shifting the distributions away from zero, through an extra parameter that is estimated or specified *a priori*.

### 3. Simulation of WTP ratios

The most common practice when dealing with induced distributions is to estimate their moments through simulation (see e.g. Hensher & Greene, 2003, for an in-depth discussion). In the case of WTP, numerous draws are taken from the distribution of  $\beta$  and  $\theta$ , and the ratio of  $\beta$  to  $\theta$  is calculated for each draw. The ratios are draws from the distribution of WTP, and the mean and variance of the draws of the ratios are used as estimates of the mean and variance of WTP in the population. For distributions with finite moments, simulation of the moments is an appropriate and useful procedure. However, when the moments do not exist, simulation can serve to mask their non-existence, providing finite simulated moments when the true moments are infinite. The reason is straightforward: the true moments are infinite because values of the denominator can be arbitrarily close to zero, resulting in unboundedly large values for the ratio. However, then simulating the ratio, it is highly unlikely that a draw of the denominator will be obtained that provides a value of the ratio that is larger than the computer can recognise. In fact, a particular set of draws of the denominator might not result in any unreasonably large ratios, such that the mean ratio seems not only finite but reasonable. Of course, if the simulation is repeated, quite different results will probably be obtained in different simulations, and this variance over simulations might alert the analyst to the problem. However, simulation is usually performed once rather than repeatedly, and an analyst obtaining a finite, and perhaps even reasonable, simulated mean WTP would not have any reason to suspect that the true moments for the specified model are infinite.

Even though the point (that simulation gives finite moments even when the true moments are infinite) is fairly obvious, it is perhaps useful to provide some illustration of the phenomenon, in

order to emphasise its importance. For these illustrations, we simulate the value of time (VOT, aka, the WTP for time reductions) with time measured in minutes and cost in £. To focus the situation, we specify the travel time coefficient to be fixed, with a value of -0.05. The cost coefficient is given the following distributions:

- Uniform, distributed between 0 and -1, i.e. with mean of -0.5
- Triangular, distributed between 0 and -1, i.e. with mean and mode of -0.5
- Normal, distributed with a mean of -0.5 and a standard deviation of 0.3, such that 95% of the mass is on the negative side of zero.

In each simulation, we use  $10^7$  draws (ten million) of the cost coefficient for each distribution. We calculate the mean and variance of the simulated draws of VOT. Note that  $10^7$  is a far larger number of draws than is used in most simulations; we chose such a large number of draws in order to show that finite moments are obtained even with an extremely large number of draws. Experiments making use of smaller numbers of draws confirmed our results while also highlighting even greater risk of problems with infinite moments being masked; these results can easily be reproduced by the reader.

We repeated the simulations 10 times, with a different set of  $10^7$  random draws each time. Table 1 gives the mean and standard deviation over the ten simulations of the simulated mean and variance.

**Table 1: Simulation results**

		Mean	Variance
<b>Uniform</b>	<b>Mean</b>	60.23	$2.05 \cdot 10^9$
	<b>Std. Dev.</b>	13.92	$3.03 \cdot 10^9$
<b>Triangular</b>	<b>Mean</b>	8.32	252.02
	<b>Std. Dev.</b>	0.01	38.53
<b>Normal</b>	<b>Mean</b>	1.88	$5.73 \cdot 10^8$
	<b>Std. Dev.</b>	6.57	$8.51 \cdot 10^8$

The uniform distribution produces very large values of simulated mean VOT, while the normal distribution produced very low values, even though in both cases the mean is actually undefined. Of course, in both cases, the standard deviation over simulations is large, especially given the huge number of draws in each simulation. For the triangular distribution, the true mean VOT is actually finite, such that it is amenable to simulation. As expected, the standard deviation of the mean over simulations is exceedingly small (0.01) for the triangular distribution, since the true mean exists and each simulation uses a very large number of draws.

The simulated variances are very large for the uniform and normal distributions, which should provide a clue to the analyst that something is amiss. However, for the triangular distribution, the simulated variance of VOT, while large, is not nearly as great as for the other two distributions, even though the true variance is infinite under all three of them. As with the simulated means, the simulated variances do not provide a reliable guide as to whether the true variances are defined, and, indeed, one would not expect simulation to be useful for this purpose.

Analysts have occasionally advocated the use of censoring in order to account for the fact that the simulation results are unduly influenced by a very small number of extreme draws. Using the example of the normal distribution, we simulated the mean and variance of VOT using different degrees of censoring of the distribution, ranging between one percent and ten percent with symmetrical censoring. The results are shown in Table 2, where once again, we present summary

results across the ten different sets of simulation draws. The results for the full sample (i.e. 0% censoring) are given in the first column.

**Table 2: Impacts of censoring**

		0%	1%	2%	5%	10%
<b>Mean</b>	<b>Mean</b>	1.88	7.24	7.24	7.23	7.22
	<b>Std. Err</b>	6.57	0.02	0.01	0.01	0.01
<b>Variance</b>	<b>Mean</b>	$5.73 \cdot 10^8$	1,955.10	961.61	368.39	171.11
	<b>Std. Err</b>	$8.51 \cdot 10^8$	14.18	4.47	0.86	0.30

The results suggest that a small amount of censoring leads to a high degree of stability across runs, so that the variation across runs is indeed just a result of a few extreme values. From this perspective, censoring may appear to be a very desirable solution. However, this conclusion neglects the fact that the true mean is in fact infinite: the censoring, by creating stability over simulations, serves to further mask the reality of the situation. This fact is shown more vividly with the simulated variance. With each additional degree of censoring, there is a reduction in the simulated variance of VOT. The analyst essentially selects the variance of VOT by selecting a degree of censoring, rather than estimating the variance of VOT from the data. The basic problem is that the model, as specified, implies infinite mean and variance of VOT, and so the solution is to re-specify the model rather than censor under the existing specification.

Another observation can be made at this point. While the moments of a distribution may not be defined, the percentiles always exist. As an illustration, Table 3 shows the simulated lower and upper 1%, 5% and 10% points for the ten sets of draws for the three choices of distributions above. As can be seen from these results, the percentile estimates are very stable across the ten runs. This finding may be useful in scenarios where the median WTP can be used (£6/hr for uniform and triangular in this case, and £5.60/hr for normal), or where we are interested in knowing what share of the sample population have a WTP above some threshold value of interest. However, while the temptation may exist to consider the mean WTP to be approximated by the midway point of the 5-percentile and 95 percentile values (or some similarly symmetric percentiles), this concept is not correct. In fact, we observed scenarios with the normal distribution where the *simulated* mean exceeded the 99% percentile. And of course, the actual mean is infinite.

**Table 3: Simulating percentiles of WTP distribution**

		1%	5%	10%	50%	90%	95%	99%
<b>Uniform</b>	<b>Mean</b>	3.03	3.16	3.33	6.00	29.99	59.99	299.67
	<b>Std. Err</b>	0.00	0.00	0.00	0.00	0.02	0.07	0.90
<b>Triangular</b>	<b>Mean</b>	3.23	3.56	3.86	6.00	13.42	18.97	42.42
	<b>Std. Err</b>	0.00	0.00	0.00	0.00	0.01	0.01	0.07
<b>Normal</b>	<b>Mean</b>	-90.63	2.21	3.04	5.60	16.11	26.86	107.44
	<b>Std. Err</b>	0.27	0.00	0.00	0.00	0.01	0.02	0.31

#### 4. On the use of conditional distributions

Train (2003) provides a procedure for estimating the distribution of coefficients for each individual in a dataset conditional on the choices that the person made. It may be tempting to use these conditional distributions to provide alternative distributions of WTP when the estimated models

imply unreasonable distributions. In particular, it is possible to calculate the conditional mean coefficients for each sampled individual, take the ratio of these conditional means, and interpret the mean and variance of these ratios over individuals as the mean and variance of WTP in the population. Generally, the variance of the ratio of conditional means is smaller than the variance of WTP calculated from the unconditional distribution of coefficients, and so the former may be considered (erroneously) an improvement in estimation when the later are unreasonably large.

There are two problems in this approach. First, the ratio of conditional means is not the mean of the conditional distribution of the ratio, since  $E(\beta/\theta) \neq E(\beta)/E(\theta)$ . Second, the procedure mistakenly ignores the fact that conditional distributions are derived from, and must aggregate to, the unconditional distribution. In particular, the (unconditional) population variance is equal to the variance of the conditional means PLUS the variance of the conditional distribution around the conditional means. The variance of conditional means is less than the unconditional variance not because it is a more reasonable estimate, but rather because it incorrectly excludes the variance around the conditional means.

## 5. Proposed solutions

There are several paths that analysts can take in assuring that their models have distributions of WTP with finite moments. A few of the most prominent are:

- In selecting a distribution for the cost coefficient, the analyst can use the theorem above to determine whether the implied distribution of WTP has finite moments. As highlighted above, the existence of inverse moments may, for some distributions, depend on the actual estimated parameters of the distribution.
- The analyst can require that the distribution of the cost coefficient be bounded away from zero. As shown in appendix B, inverse moments exist for any distribution that does not have support arbitrarily close to zero. The bound can be set by the analyst (which entails a degree of arbitrariness that may be deemed unacceptable, cf. Hess et al., 2005) or it can be treated as an additional parameter to estimate.
- Another possible solution is the use of a finite mixture models, including latent class models (see e.g. Hess et al., 2007) and the non-parametric estimation procedures suggested by Train (2008). In these specifications, the continuous distribution for coefficients discussed above is replaced with a finite distribution, i.e., a distribution that has mass at a finite number of coefficient values. If all the points for the cost coefficient are either estimated or constrained to be away from zero, the implied WTP distribution is also finite with defined values.
- Finally, the model can be re-parameterised in WTP space, as suggested by Train and Weeks (2005). That is, utility in equation 1 can be re-written as

$$U_j = \theta c_j + \theta \lambda a_j + \text{other terms} + \varepsilon_j \quad (2)$$

where  $\lambda = \beta/\theta$  is the WTP for the attribute. Instead of specifying distributions for  $\theta$  and  $\beta$ , the analyst specifies distributions for  $\theta$  and  $\lambda$ . In estimation, the coefficient of the attribute is calculated as the product of  $\theta$  and  $\lambda$  rather than as one coefficient in itself. Any model specified as in equation (1) can be re-expressed in the form of equation (2), and vice versa, and so the re-expression is simply a re-parameterisation rather than a new model. However, in equation (2), the analyst specifies the distribution of WTP directly, rather than specifying the distribution of  $\beta$  and  $\theta$  and deriving the distribution of WTP. The analyst can therefore be

sure to specify a reasonable distribution for WTP that has finite moments, without needing to refer to our theorem.

## 6. Conclusions

The majority of discrete choice models estimated by academics and a growing share of models estimated by non-academic practitioners now utilise random coefficients for cost and non-cost attributes. The distribution of WTP for an attribute (which is the coefficient of that attribute divided by the cost coefficient) is derived from the distribution of the coefficients. Since the cost coefficient enters the denominator, values of the cost coefficient that are close to zero induce large values of WTP, with unbounded large values of WTP resulting from cost coefficients arbitrarily close to zero. In this paper, we identify a criterion to determine whether the distribution of WTP has finite moments. Using this criterion, we show that some popular distributions used for the cost coefficient in random coefficient models, including normal, truncated normal, uniform and triangular, imply infinite moments for the distribution of WTP. We also point out that relying on simulation approaches to obtain moments of WTP from the estimated distributions of the cost and attribute coefficients can mask the problem by giving finite moments when the true ones are infinite. Similarly, using conditional distributions is inappropriate, and percentiles cannot be used to infer any information on the moments (since the moments don't exist). We identify several approaches that analysts can utilise to assure that the distribution of WTP has finite moments, including the use of continuous distributions that satisfy the conditions in our theorem, with appropriate constraints if required; using finite mixture such as latent class models and a form of non-parametrics; and by re-parameterising the model in WTP space such that the distribution of WTP is estimated directly rather than derived.

## Acknowledgements

The second author acknowledges the financial support of the Leverhulme Trust in the form of a "Leverhulme Early Career Fellowship".

## References

- Hensher, D. & Greene, W. (2003), 'The Mixed Logit Model: The State of Practice,' *Transportation* 30(2), 133–176.
- Hess, S., Bierlaire, M. & Polak, J.W. (2005), 'Estimation of Value of Travel-Time Savings using Mixed Logit Models,' *Transportation Research Part A*, 39(2-3), pp. 221-236.
- Hess, S., Bierlaire, M. & Polak, J.W. (2007), 'A Systematic Comparison of Continuous and Discrete Mixture Models,' *European Transport* 37, pp. 35-61.
- McFadden, D. (1998) 'Measuring Willingness-to-Pay for Transportation Improvements,' in T Gärling, T Laitila, K Westin (eds.), *Theoretical Foundation for Travel Choice Modelling*, Pergamon.
- McFadden, D. & Train, K. (2000), 'Mixed MNL Models for Discrete Response,' *Journal of Applied Econometrics* 15, 447–470.
- Sillano, M. and Ortúzar, J. de D. (2005), 'WTP estimation with mixed logit models: some new evidence', *Environment and Planning A* 37 (2005)

Train, K.E. (2008), 'EM Algorithms for Nonparametric Estimation of Mixing Distributions,' *Journal of Choice Modelling*, 1(1), pp. 40-69.

Train, K. and Weeks, M. (2005). 'Discrete Choice Models in Preference Space and Willingness-to-Pay Space,' in R. Scarpa and A. Alberini (eds), Application of Simulation Methods in Environmental and Resource Economics, Springer Publisher, Dordrecht, The Netherlands, chapter 1, pp. 1-16.

**APPENDIX : Proof of theorem in Section 2**

If a random variable  $\theta$  has an absolutely continuous probability density  $f(\theta)$ , then for any positive integer  $k$  the inverse moment  $E(1/\theta^k)$  exists if and only if  $\lim_{\theta \rightarrow 0} \frac{f(\theta)}{\theta^h}$  exists for some  $h > k - 1$ .

We first prove the following Lemma.

Let  $f(\theta)$  be the probability density function of a random variable  $\theta$ , where  $f(\theta)$ :

1. is absolutely continuous
2. has support only on the positive half-line
3. is monotonic (either non-decreasing or non-increasing) in an interval  $(0, r)$  for some  $r$

Then for a positive integer  $k$ ,  $E(1/\theta^k)$  exists if and only if  $\lim_{\theta \rightarrow 0} \frac{f(\theta)}{\theta^h}$  exists for some  $h > k - 1$ .

Suppose  $\lim_{\theta \rightarrow 0} \frac{f(\theta)}{\theta^h} = H$  for some non-negative value of  $h$ . Certainly the limit exists for  $h = 0$  because of the continuity of  $f$ .

Define  $S_n^k(r) = \int_{r/n+1}^{r/n} \theta^{-k} f(\theta) d\theta$ . Then, because of the monotonicity property, we know that  $S_n^k(r)$  is contained in a closed interval

$$S_n^k(r) \in \left[ \left( \frac{r}{n} - \frac{r}{n+1} \right) \left( \frac{r}{n+1} \right)^{-k} f\left( \frac{r}{n+1} \right), \left( \frac{r}{n} - \frac{r}{n+1} \right) \left( \frac{r}{n} \right)^{-k} f\left( \frac{r}{n} \right) \right]$$

i.e.

$$S_n^k(r) \in \left[ \left( \frac{1}{n} \right) \left( \frac{n+1}{r} \right)^{k-1} f\left( \frac{r}{n+1} \right), \left( \frac{1}{n+1} \right) \left( \frac{n}{r} \right)^{k-1} f\left( \frac{r}{n} \right) \right] = [L_n^k(r), U_n^k(r)], \text{ say}$$

and

$$\lim_{n \rightarrow \infty} L_n^k(r) = \left( \frac{1}{n} \right) \left( \frac{n+1}{r} \right)^{k-1} H \left( \frac{r}{n+1} \right)^h = H \left( \frac{1}{n} \right) \left( \frac{n+1}{r} \right)^{k-h-1}$$

$$\lim_{n \rightarrow \infty} U_n^k(r) = \left( \frac{1}{n+1} \right) \left( \frac{n}{r} \right)^{k-1} H \left( \frac{r}{n} \right)^h = H \left( \frac{1}{n+1} \right) \left( \frac{n}{r} \right)^{k-h-1}$$

Then, providing  $k < h + 2$ , this interval shrinks to zero and, because of the continuity and limit properties of  $f$ ,

$$\lim_{n \rightarrow \infty} S_n^k(r) = \left( \frac{H}{r^{k-h-1}} \right) n^{k-h-2}$$

Then define  $T_N^k = \sum_{n=1}^N S_n^k$ . If the limit exists,  $\lim_{N \rightarrow \infty} T_N^k = \int_0^r \theta^{-k} f(\theta) d\theta$ . Note that if  $k \geq h + 2$ ,  $S_n^k(r)$  does not tend to zero, so that the series  $T$  clearly diverges.

It is a classical result that  $\sum n^\alpha$  converges if and only if  $\alpha < -1$ , so that the series  $T$  converges if and only if  $(k - h - 2) < -1$ , i.e.  $k - 1 < h$ , proving the Lemma.

The Lemma has the following Corollary

If a random variable  $\theta$  has an absolutely continuous probability density  $f(\theta)$  defined on the positive half-line and  $\lim_{\theta \rightarrow 0} f(\theta) > 0$ , then none of the inverse moments  $E(1/\theta^k)$  exists.

The corollary follows immediately by noting that the limit in the Lemma fails to exist for any  $h > 0$  and so the moments do not exist for any  $k \geq 1$ .

We can now conclude that the condition of monotonicity in the Lemma is not required. If  $\lim_{\theta \rightarrow 0} f(\theta) > 0$ , then the inverse moments and the limit fail to exist, as in the Corollary. However, if  $\lim_{\theta \rightarrow 0} f(\theta) = 0$ , then because of continuity the function must be non-decreasing in a neighbourhood of 0 (it must remain non-negative).

Suppose the function  $f$  is defined over the negative half line. For negative  $\theta$ , the limit is defined for integer  $h$  only but with that reservation the same existence result then applies to the negative half-line as for the positive half-line.

Finally, if  $f$  is defined over the whole line then the  $k^{\text{th}}$  moment exists for the whole line only if the moments over both positive and negative half-lines exist, completing the proof of the theorem.

Note that in the case when  $\lim_{\theta \rightarrow 0^-} \frac{f(\theta)}{\theta^h} \neq \lim_{\theta \rightarrow 0^+} \frac{f(\theta)}{\theta^h}$ , i.e. there is some sort of 'kink', then if both limits exist the  $k^{\text{th}}$  moment exists for  $k < h - 1$ ; but if one does not exist then the  $k^{\text{th}}$  moment does not exist.