Comprehensive Radar Cross-Section "Target Typing" Investigation for Spacecraft

M.D. Hejduk[1]
SRA International, Colorado Springs, CO, 80916

D. DePalma[2]
ITT Corporation, Colorado Springs, CO, 80916

A study presented last year revealed that the standard target type models used for power allocation by spacetrack radars—namely, the Swerling models, which are two specific forms of the two-parameter gamma distribution—are actually not particularly representative of the radar cross-section (RCS) histories of actual space objects; they in fact could account for only 35% of the objects in a typical operational sample. The present study re-examines this result with a more robust dataset and attempts to identify additional model candidates to represent the unmodeled remainder, trying two- and three-parameter versions of the gamma, Weibull, and lognormal distributions and using rigorous parameter estimation and goodness-of-fit tests as validation.

This investigation yielded a number of findings. First, the Swerling I/II distribution (Rayleigh distribution) performs surprisingly poorly, suggesting that Rayleigh scattering is actually a smaller component of spacetrack radar backscatter than previously believed. Second, the introduction of targeted additional gamma distribution models, specifically one with a shape parameter of 1.5, improved the modeling notably. Third, the use of three- rather than two-parameter fits, for gamma, Weibull, and lognormal distributions, did not appreciably improve the situation beyond that for the two-parameter fits; it is thus reasonable to presume that spacecraft RCS history probability density functions begin at the origin. Fourth, only the two-parameter lognormal distribution emerged as a significant contender to serve as an additional model; and even then the addition of two specified lognormal distributions accounted for only 12% of the overall cases. Finally, a descriptive statistical analysis of the cases that resist representation by the standard distributions reveals bimodality, unusual skewed distributions (in dBsm space!), and a minority of reasonably well-behaved leptokurtic distributions. In short, there is no obvious "magic" solution or distribution to improve the target type modeling, but some progress could be made by introducing one additional gamma distribution and two specific lognormal distributions, attempting model matching with a standard leptokurtic distribution (perhaps a student's *t*-distribution), and developing/deploying an outlier test that does not make the assumption of an underlying Gaussian distribution.

I. Introduction

TO operate efficiently in tracking space objects, space radars need to have a reasonable characterization of the PDF of the objects' radar cross-section (RCS) history in order to determine the proper amount of power to apply to achieve a desired probability of detection. Based on Peter Swerling's early study of aircraft RCS data, two forms of the gamma statistical distribution have been presumed as appropriate for spacecraft and guided the software development for most space radars: an exponential distribution (the Swerling I/II type) and an Erlang (chi-square with X degrees of freedom) distribution (the Swerling III/IV type).[1] A study presented last year[2] focused on assessing the adequacy of these two distribution types to represent RCS histories for spacecraft and found that, taken together, they were able to account for only 35% of a broad sample of space objects. A proposed replacement of the entire Swerling paradigm with a use of the two-parameter lognormal distribution was also investigated and found to

[1] Adjunct Scientist, 1795 Jet Wing Drive, #160, Senior Member AAS
[2] Systems Engineer, 4410 Fountain Blvd, Member AAS

be unsuccessful. The study concluded with a suggestion to try more elaborate versions of the gamma distribution family and to investigate other canonical distribution types to see if a more comprehensive distribution set could be assembled for spacetrack radars.

The present effort attempts to provide this suggested expansion in the following manner. First, it improves on the previous dataset used for the investigation by enforcing a more rigorous radar calibration investigation, retaining only those data taken during the most restrictive calibration-compliant periods. Second, it examines a few natural extensions to the Swerling PDF family to determine any increase in comprehensiveness by a controlled proliferation of the present canonical distribution type. Third, it investigates whether using the three-parameter version of the Swerling PDF family (Swerling himself limited the distributions to two parameters, setting the location parameter to zero) brings any increase in comprehensiveness. Fourth, it attempts to align the remaining distributions to some other suggested types, such as the Weibull and an expanded lognormal, to see if this somewhat expanded repertoire can account for the remainder. Finally, it characterizes these further resistant types with descriptive statistics to help to understand their nature and assess their likelihood of being accounted for by any of the standard distributions.

II. Dataset

The dataset of individual RCS hit values used in the previous study was composed of six months of spacetrack data from the Eglin FPS-85 radar, a large, UHF tracking radar located in the Florida panhandle. This was a quite voluminous dataset, consisting of twenty million radar hits taken from September 2008 to March 2009, thus including objects generated from the two recent debris-producing events (Chinese ASAT and Iridium-COSMOS collision). Usual site-based calibration controls were in place during the data collection, but subsequent maintenance efforts at the Eglin radar have uncovered some problems with the Eglin RCS calibration procedure. To remedy these difficulties and ensure that a well-calibrated dataset was used for this present investigation, results of the tracking of RCS calibration satellites were examined (which are tracked routinely throughout each day of Eglin's operation), and the dataset was winnowed to retain only those radar hits taken during periods when all of the calibration satellite RCS data remained within 1.5 dBsm of each calibration satellite's theoretical value. This criterion did reduce the dataset significantly, but enough data remained to allow adequate sampling for the study. Table 1 gives the number of space objects, by object type, with sample densities as indicated.

Table 1. Dataset size by object type

Object Type	# of Radar Hits		
	> 0	> 200	> 1000
Payloads	2033	1876	1471
Rocket Bodies	1362	1143	691
Debris	8034	6252	3291
Total	11441	9282	5463

For all of the subsequent analyses, only objects with more than 1000 radar hits were used. This would typically correspond to at least twenty separate tracks, which should give a quite comprehensive survey of the different satellite-sensor aspect orientations for each object.

II. Existing and Expanded Swerling Types

Swerling chose to define his target types with reference to the chi-squared distribution, but its more general form is the three-parameter gamma distribution PDF, given below:

$$f(x;\gamma,\beta,m) = \frac{1}{\beta^m \Gamma(m)} (x-\gamma)^{m-1} \exp\left(\frac{-(x-\gamma)}{\beta}\right) \quad (1)$$

In this formulation m is the shape parameter, which governs the overall appearance of the PDF and thus the type of gamma distribution employed. If m is an integer, the doubling of m indicates the number of degrees of freedom of the distribution if one wishes to consider it to be a chi-squared distribution. β is the scale parameter, for which the maximum likelihood estimator (MLE) is the sample mean divided by m. γ is the location parameter for this

distribution and was set to zero by Swerling. In such a case, setting *m* to 1 produces an exponential (Rayleigh) distribution of the form

$$f(x;\beta) = \frac{1}{\beta}\exp\left(\frac{-x}{\beta}\right) \tag{2}$$

which is the PDF for the Swerling I and II types; and setting *m* to 2 produces the following Erlang distribution

$$f(x;\beta) = \frac{1}{\beta^2}x\exp\left(\frac{-x}{\beta}\right) \tag{3}$$

which is the PDF for the Swerling III and IV types. However, there is no particular reason to limit *m* to integer values; and one is led to ask whether there would be an advantage to some "intermediate" and expanded Swerling types, such as the distribution in (1) with *m* set to 0.5, 1.5, 2.5, and perhaps values even larger than 2.5. Expanding the Swerling target type set was explored at a high level in the previous study but not rigorously enough to allow for any emerging recommendations. So the first activity here is to evaluate the new dataset for conformity to the canonical Swerling types and a modest expansion of them through the addition of fractional *m*-values. In this paper, the nomenclature "S#," such as S1.5, S2, &c., will be used to refer to a gamma distribution with location parameter 0 and shape parameter (*m*) of the number given. Figure 1 gives the gamma distribution PDFs for a number of different shape parameters, as well as a typical lognormal distribution for comparison.

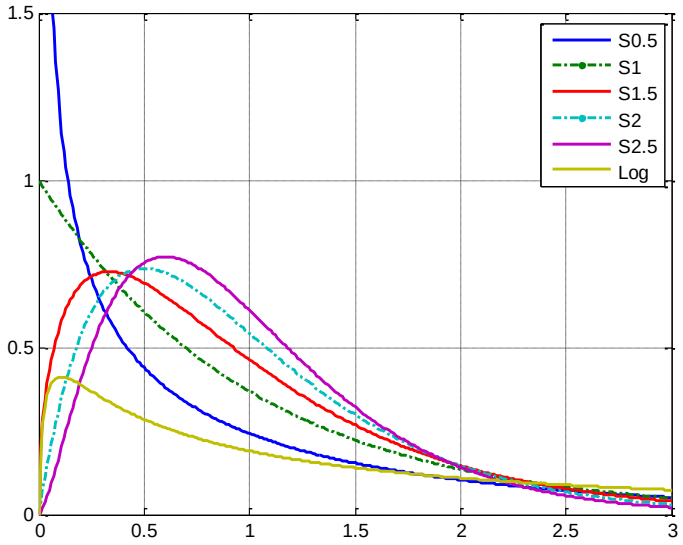

Figure 1. Gamma distributions of different shape parameters

The goodness-of-fit (GOF) evaluation approach is similar that used in the previous study but with some necessary expansions; so only an abbreviated description of the basic technique will be given here, while the expansions will be treated in somewhat more detail. Empirical distribution function (EDF) techniques examine the difference between the empirical CDF of the actual distribution and the CDF of the hypothesized distribution, sum the differences between them, manipulate this summed value into a more robust statistical factor, and compare the factor to Monte-Carlo-generated tables of percentile points in order to determine the *p*-value for which the actual distribution could be considered to conform to the hypothesized distribution. This technique is desirable in the present analysis because it can be used to test a number of different distributions under different conditions, and it works well when percentile point tables are readily available for the hypothesized distribution and for the particular parameters estimated from the sample. For example, for the Swerling cases given above, the gamma shape parameter is given *a priori* but the scale parameter (β) is estimated from the sample; a percentile point table for this specific situation is required. However, some of the distributions to be tested, such as the three-parameter gamma, Weibull, and lognormal distributions, do not have published tables, so a different approach must be chosen. Typically, the approach used is to divide the dataset into two groups, using one for the parameter estimation and another for the GOF test; the dataset used for evaluation is thus not the same as that used for estimation, and percentile point tables for a fully-specified distribution (no parameters estimated) can be employed. The power of the test is somewhat diminished under this circumstance, but it is still quite serviceable.[3,4]

Because GOF tests tend to miscarry when applied to very large sample sets, it is common to use a resampling approach, testing each sample and then noting the number of samples that "pass" the GOF test for a given p-value. To investigate the conformity to the canonical and expanded Swerling types, 1000 pairs of 100-point samples (without replacement, one for parameter estimation and one for GOF testing) were drawn from the calibration-winnowed RCS histories for each object, and the percentage of the 1000 samples that met or exceeded a p-value of 0.05 was noted; and 75% of the samples meeting the p-value was selected as the standard for which one may say that the hypothesized distribution reasonably represents the actual distribution. P-value tables exist for the so-called "case 1" of the gamma distribution, in which the location parameter is 0, the shape parameter is dictated *a priori*, and the scale parameter (beta) is estimated from the sample. Some commentators might consider this p-value to be overly restrictive, as p-values of 0.02 or 0.01 are often used; others might consider the 75% threshold to be too permissive, as one in four samples would fail to achieve the desired p-value. To the present authors, the combination of this slightly-demanding p-value and slightly-permissive percent-of-samples level seemed a reasonable amalgamation.

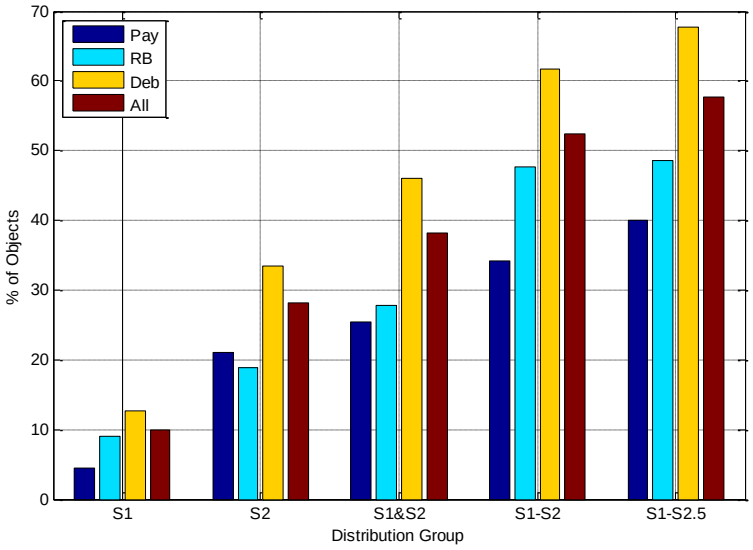

Figure 2. Swerling-type distribution GOF test results

A selected presentation of the results is given in Figure 2. The figure shows the number of satellites for which 75% of the samples achieved a p-value of 0.05 when tested against the shown distribution. The "higher" Swerling distributions of shape parameters three or greater made only a marginal contribution and are thus not shown here.

A number of interesting items emerge from these results. First, it is surprising that the S1 case, which represents the Swerling I and II distributions and is the darling of the traditional target-typing family, fares relatively poorly, accounting for only 10% of the cases—one-third of what the S2 distribution (Swerling III and IV) can accomplish. This implies that satellites, whether they be intact spacecraft with sharp edges and booms or irregularly shaped pieces of debris, do not exhibit the random scattering described by a Rayleigh distribution; the systematic effects of smoother, more continuous facets predominate. Second, the intermediate case of S1.5 accounts for about 12.5% of the tested objects and about 20% of the rocket bodies and debris—well more than the S1 distribution. Adding the S2.5 distribution, however, accommodates only about another 5% of the objects; and the expanded Swerling types with larger shape parameters accommodate even less. The S1.5 distribution thus seems to be a neglected contributor. For spacetrack radars that already use the Swerling models and thus perform gamma-distribution GOF tests, it is clear that the addition of the S1.5 distribution would be a straightforward and helpful addition. With such a distribution added, the Swerling family can be said to account for a little over half of the space objects sampled, although it is important to point out that they represent only about one-third of the intact payloads but almost 60% of the debris. Regardless, it is clear that the Swerling family, at least as sampled here, still leaves almost half of the space objects unmodeled.

III. Broader Range of Gamma Shape Parameter

One potential approach to modeling these objects that cannot be represented by the S1-S2 expanded Swerling set is simply to continue to expand the set through a full range of gamma distribution shape parameters. Approximately 2600 objects in the sample set did not meet the GOF threshold for the S1-S1.5-S2 distributions; this group of satellites was thus evaluated against full-range run two-parameter gamma distribution parameter estimation and then the appropriate GOF tests. Four hundred and seventy-seven objects met the GOF standard (75% of samples achieving a 0.05 p-value), about 18% of this group. Figure 3 gives a CDF plot of the shape parameters (*m*-values) for this set of objects. One observes that there are few successful cases below *m*-values of 2.25 or so, as

these tend to be represented adequately by the S1, S1.5, and S2 distributions; but after that the representation increases more rapidly until an m-value of about 4 or so, at which point the curve starts to flatten. From this graph, one can see that about 60% of the cases are contained between m-values of 2.5 and 4, and perhaps 75% of them could be represented by a collection of S2.5, S3, S3.5, and S4 distributions. Adding these four distribution types would properly account for 200 more satellites; and while every little bit could be considered helpful, this is a relatively small marginal gain for introducing four new distribution types: The S2 type alone accounted for 28% of the entire sample, whereas S2.5-S4 would account for only about 4%.

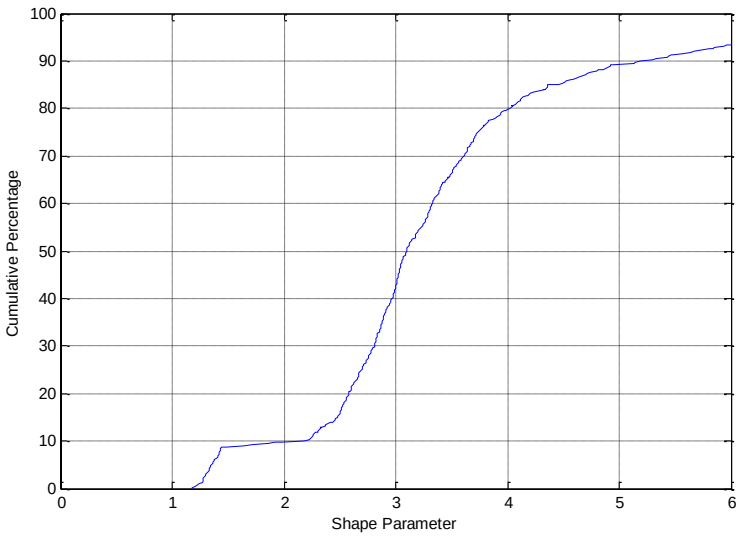

Figure 3. Estimated shape parameters for full-range two-parameter gamma distribution exploration.

IV. Three- versus Two-Parameter Gamma Distributions

Swerling postulated that any RCS history PDF would begin at the origin, presumably because he believed that a random-scattering component would always be a notable contributor to the overall scattering function. This would explain his embrace of the Rayleigh distribution, in which all of the radar return is small-component scattering, and the Erlang distribution, which could be seen as a "dilution" of the Rayleigh distribution by including the systematic effects of larger, more continuous facets. However, if the facets are large and continuous enough, it is quite possible that the PDF will not in fact tend towards the origin but instead seek out some non-zero positive value; this value would represent the minimum RCS that the more regular, smoother object would produce. An *a priori* commitment to Swerling's Rayleigh-scattering conviction, combined with the difficulty of performing parameter estimation for the three-parameter gamma distribution, have kept this form of the distribution from consideration.

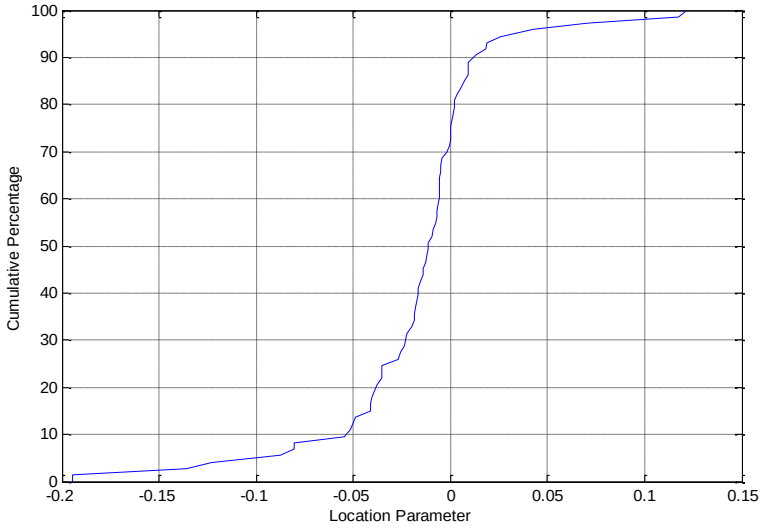

Figure 4. Estimated location parameters for three-parameter gamma distribution.

The set of cases not adequately modeled by the S1-S2 distributions were fitted for three-parameter Gamma distributions using the Cohen-Whitten modified moment estimator technique, an approach more resistant to third-moment sampling errors and more stable than the iterative solutions required to perform MLE estimates. With these estimated parameters, 73 objects met the GOF criteria for the three-parameter gamma distribution while not meeting the criteria for the two-parameter gamma distribution. Figure 4 gives a CDF of the mean location parameter for these objects.

About 70% of the location parameters are less than zero,

specifying a distribution set that is physically impossible for radar returns (as no object can have an RCS less than zero); in such cases, one would in practice set the location parameter to zero and re-estimate the two remaining parameters, which is in effect a return to the two-parameter situation. The remaining 30%, for which the location parameter is greater than zero, could be modeled by the three-parameter estimate; but it represents such a small number of objects (0.4% of the entire sample) that it is not worth pursuing or characterizing further.

V. Three-Parameter Weibull Distribution

Some studies have suggested that the Weibull distribution can serve as a useful PDF to represent the RCS histories of certain object types.[5] A test for the three-parameter Weibull distribution, whose PDF is of the form

$$f(x;\gamma,\beta,m) = \frac{m}{\beta^m}(x-\gamma)^{m-1}\exp\left[\left(\frac{-(x-\gamma)}{\beta}\right)\right]^m \tag{4}$$

was performed for all objects not modeled by the S1-S1.5-S2 set, using Cohen-Whitten parameter estimation and the same GOF approach described previously; and objects were culled who met the 75%/0.05 percentile/p-value standard when either the 2- or 3-parameter gamma did not. One hundred seven cases met this criterion, but only nine of them had positive location parameters, making the Weibull distribution an insignificant contributor. To satisfy research curiosity, however, the shape parameters for these nine cases were examined and are plotted in Figure 5.

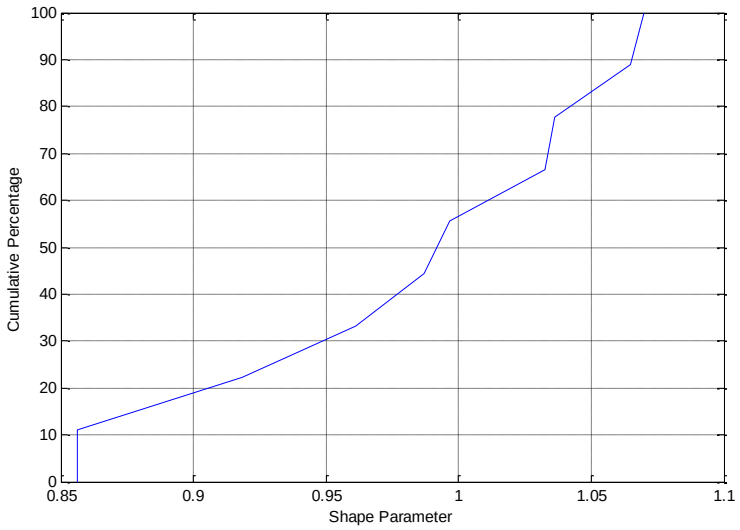

Figure 5. Location parameters for three-parameter Weibull distributions that represent realistic RCS situations

It is interesting that such cases manifest shape parameters very close to unity, making them quite similar to the exponential (S1) distribution. The exponent on the Weibull distribution's exponential term must alter the PDF enough to allow a somewhat better fit than the S1 distribution.

VI. Lognormal Distribution

The lognormal distribution is, of course, the principal alternative to the gamma-based Swerling set. It was examined as a contender in the previous study, with the findings that 1) it could not serve as a replacement for the gamma distribution and 2) it could account for only 10% of the overall examination sample that could not be modeled by the S1 and S2 distributions. However, the present analysis brings a more reliable dataset, and it tested explicitly for the three-parameter lognormal distribution rather than presuming that the lognormal location parameter would be always zero. A good place to begin is an examination of the performance of the three-parameter versus two-parameter case.

The parameter estimation approaches differ substantially between these two cases. The two-parameter approach is extremely simple and stable: one performs a log-transform of the data and then simply calculates the mean and standard deviation. The three-parameter estimation approach must choose among Pearson moment estimators, which are notoriously inaccurate; maximum likelihood estimators, which are notoriously unstable; or modified-moment estimators, which are more accurate than moment estimates and more stable than MLE but still subject to the problems of both. Modified moment estimators were selected for use here; but given that their approach differs from MLE, one expects to get different estimates for the case in which the location parameter is close to zero than if one simply took the log transform and estimated normal distribution parameters.

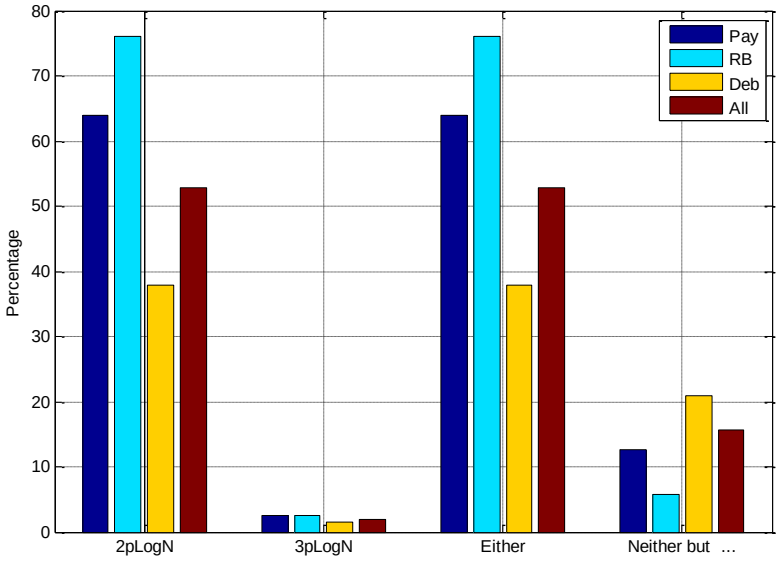

Figure 6. Two and three-parameter lognormal results

Figure 6 shows, for the satellites that cannot be adequately represented by the S1, S1.5, and S2 distributions, the performance of the two- and three-parameter lognormal distributions. It is clear right away that very few satellites produce realistic three-parameter lognormal results (that is, with a location parameter that is positive. Overall, about 50% of the non-Swerling cases appear to be representable by a lognormal distribution, which is about 25% of the total samples; this is a more encouraging result than that which was obtained in the previous study, which reported only around 13% of the cases being so situated. Of course, both analyses allowed both parameters (mean and variance) of the lognormal distribution to be estimated; in order to make the distribution suitable for radar use, certain fixed values of the scale parameter will need to be determined.

To help to determine what these should be, Figure 7 below gives a CDF plot of the standard deviations (square root of the variances) observed. The majority of the objects (60% of them) produce a standard deviation between 0.8 and 1.2. It is expected that using two lognormal distributions, one with a 0.9 standard deviation and one with a 1.1 standard deviation (corresponding essentially to variances of 0.8 and 1.2), would adequately model these 60%—representing about 12% of the total. Interestingly, even though admittedly the full range of possible lognormal distributions has been narrowed down to two, one has returned to the 10-13% figure from the previous study. One should further consider whether a distribution should be added that can account for only 6% or so of the total objects surveyed.

The "Neither but …" column from Figure 6 indicates situations in which a lognormal distribution (of any variance) failed to represent the object but some other non-Swerling distribution, such as a large-shape-parameter gamma or Weibull, did. The sum of this column and the "Either" column is about 70%, indicating that 30% of the non-Swerling objects, which is about 15% of the total sample, cannot be said to conform to any of the distributions tested. It is about this set that one should like to know more.

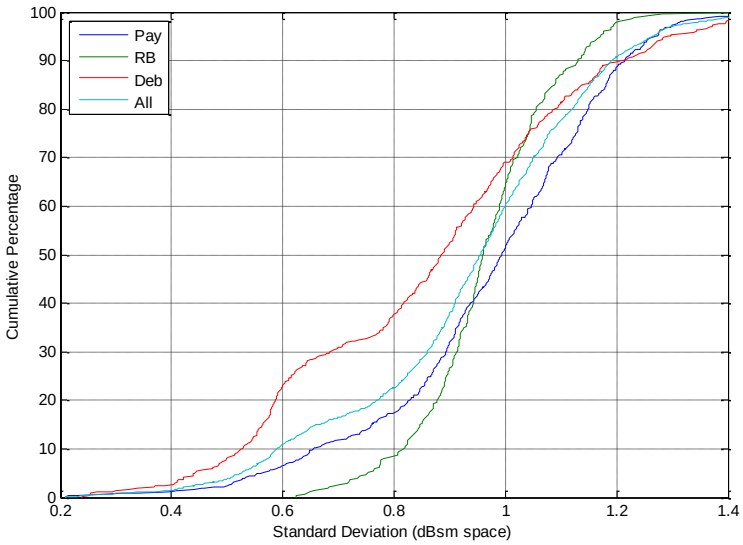

Figure 7. Standard deviations for successful lognormal fits (Gaussian fits in dBsm space)

VII. Non-Canonical Distribution Cases

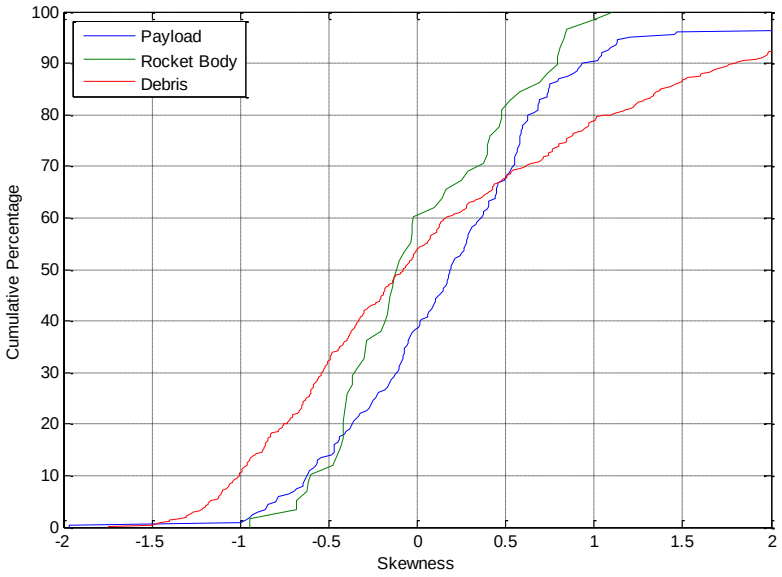

Figure 8. Skewness results for the non-canonical cases

To understand these objects' RCS histories, it makes sense to begin with standard descriptive statistics. The examination takes place in dBsm space, since the impossibility of negative RCS values in meters-squared space virtually guarantees skewed distributions in that reference frame. The descriptive statistics that will yield the most immediate information about the distributions are the standardized third and fourth Pearson moments, usually called the skewness and kurtosis. The skewness is a measure of the symmetry of the distribution. A perfectly symmetric distribution will have a skewness value of zero; a skewed distribution will have a large positive or negative value, depending on which half-plane contains the distribution's tail. Typically, skewness values of less than 0.5 (absolute value) are said to define symmetric distributions.

Figure 8, a CDF of the skewness results for the non-canonical-distribution cases, gives several interesting results. First, the most symmetrical distributions belong to rocket bodies; about 70% of the cases fall within the -0.5 to 0.5 range in skewness. Payloads are second with 60% of the cases in that range. Debris is a noticeable third with 40%. The usual thinking has been that debris objects, especially those resulting from on-orbit fragmentation, would be of irregular and multi-faceted shape and thus present a more symmetrical response about the mean value, but that is not what is observed here.

Kurtosis is a measure of the "peakedness" of a distribution, with a low value (less than 3) indicating a rounded appearance with short tails ("platykurtic") and a high value (greater than 3) indicating a peaked appearance with long tails ("leptokurtic"). These descriptions have less intuitive meaning when the distribution is not symmetric, since there is correlation between large skewness values and large kurtosis values. Figure 9 gives the kurtosis CDF for the symmetric (left side) and asymmetric (right side) cases.

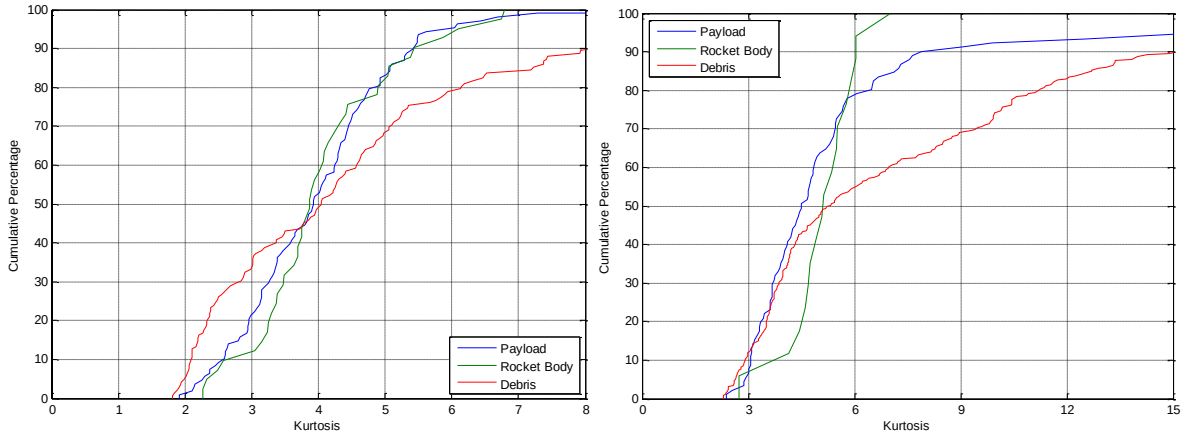

Figure 9. Kurtosis results for the symmetric (left graph) and asymmetric (right graph) cases.

The left graph shows that the great majority of the cases are leptokurtic (large kurtosis value), even when only the symmetric distributions are considered; the asymmetric cases (right graph) show the expected leptokurtoses. In order to understand the particular situations presented by the combinations of skewness and kurtosis values, individual cases were examined; and some heuristic categorizations and conclusions were reached based on this examination. The main division will be between symmetric and asymmetric cases, with subdivisions within each based on kurtosis value.

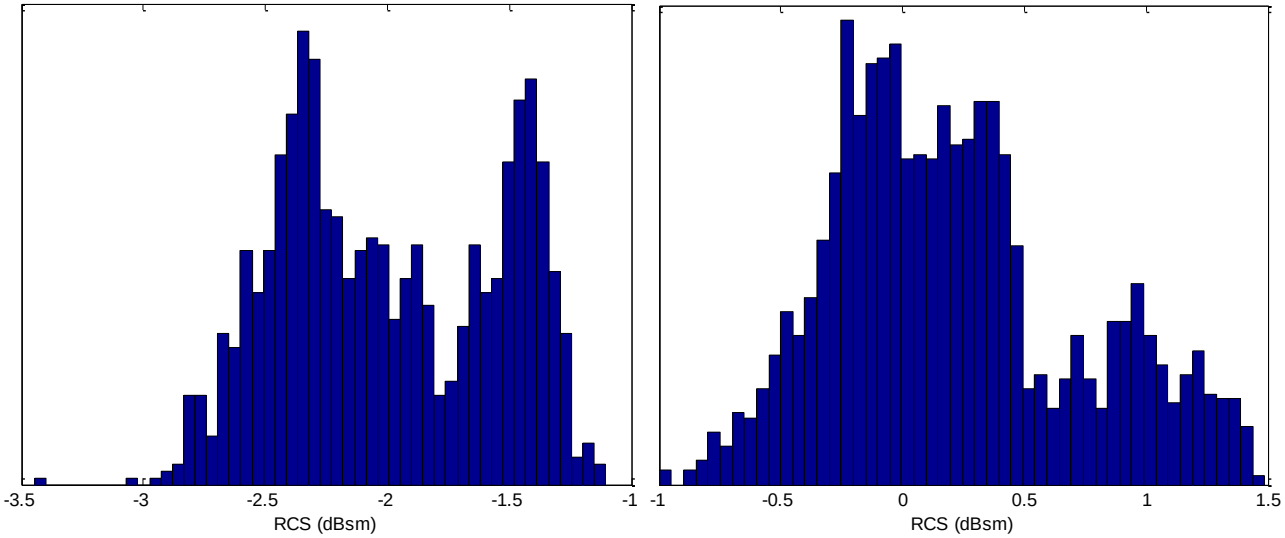

Figure 10. Symmetric platykurtic cases exhibiting two different strains of bimodality

Of the symmetric cases, about one-third (12% of the entire set) are also platykurtic; for RCS histories, nearly all of these are bimodal; and as the kurtosis increases to three, the bimodality becomes weaker. The two graphs in Figure 10 show a strongly bimodal case (k=1.8) and a weaker one (k=2.6). As one can see, these distributions are not really symmetric; rather, their bimodality is such that they happen to bring about a symmetric skewness value.

The remaining two-thirds (25% of whole sample) are leptokurtic; as the kurtosis increases, they proceed through an "inequality" phase, to a classic leptokurtic phase, to a "detritus" phase, to an extreme leptokurtic phase. Figure 11 shows these four states, clockwise from upper left. The first (k=3.96) is an extreme version of the dampened bimodality of the previous figure; the second (k=4.57) is a more traditional leptokurtic case; the third (k=5.32) is much like the second but with some outliers that are raising the kurtosis value; and the fourth (k=12.15) is an extreme leptokurtic distribution.

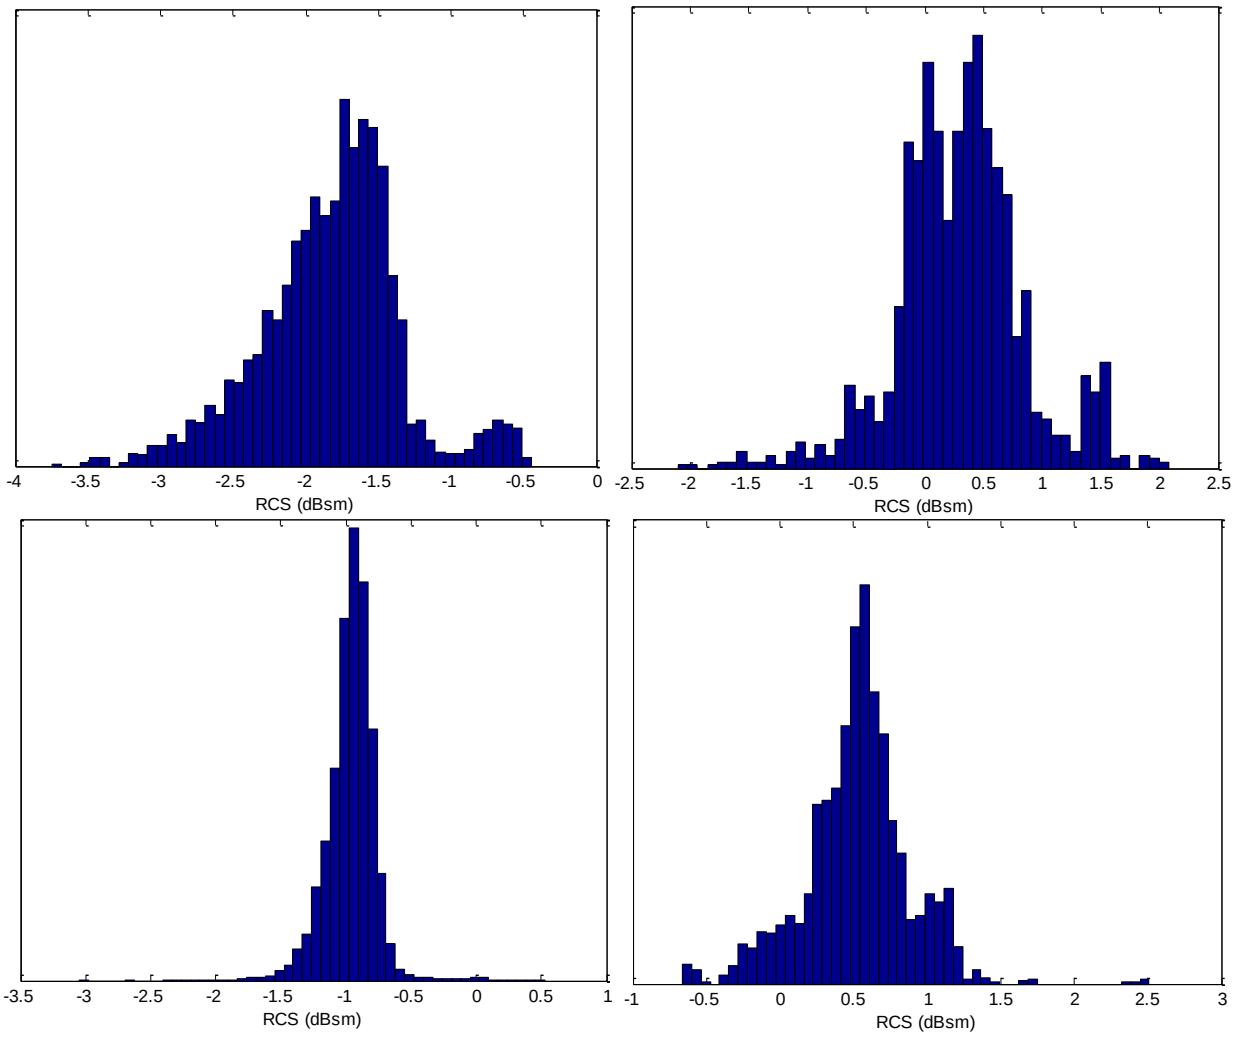

Figure 11. Symmetric leptokurtic results.

For the asymmetric cases, when the kurtosis is less than three (about one-sixth of the cases), bimodality is still observed; but it is now lopsided enough to produce a skewness value that is outside of the threshold. Figure 12 gives an example of this situation (sk=-0.69, kurt=2.75).

As the kurtosis grows, the observed results depend on the degree of skewness and kurtosis. The different possibilities are shown in the collection of graphs in Figure 13. Moving clockwise from upper left, one finds the beginnings of a right-skewed tail (sk=1.23, kurt=4.67); a symmetric, leptokurtic distribution with a small number of polluting outliers, producing extremely large skewness and kurtosis values (sk=8.28, kurt=117.42); a fairly symmetric distribution with a right-skewed tail (sk=0.83, kurt=5376),

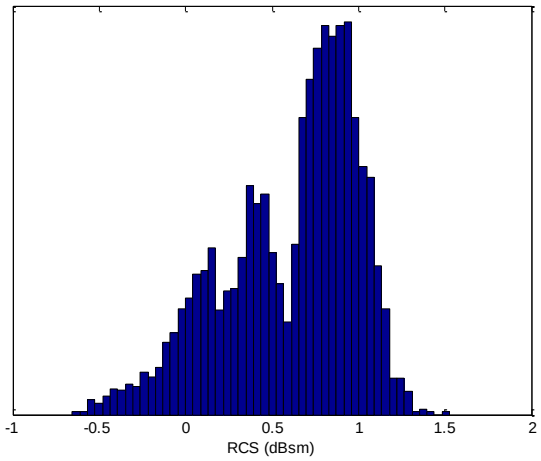

Figure 12. Asymmetric platykurtic case

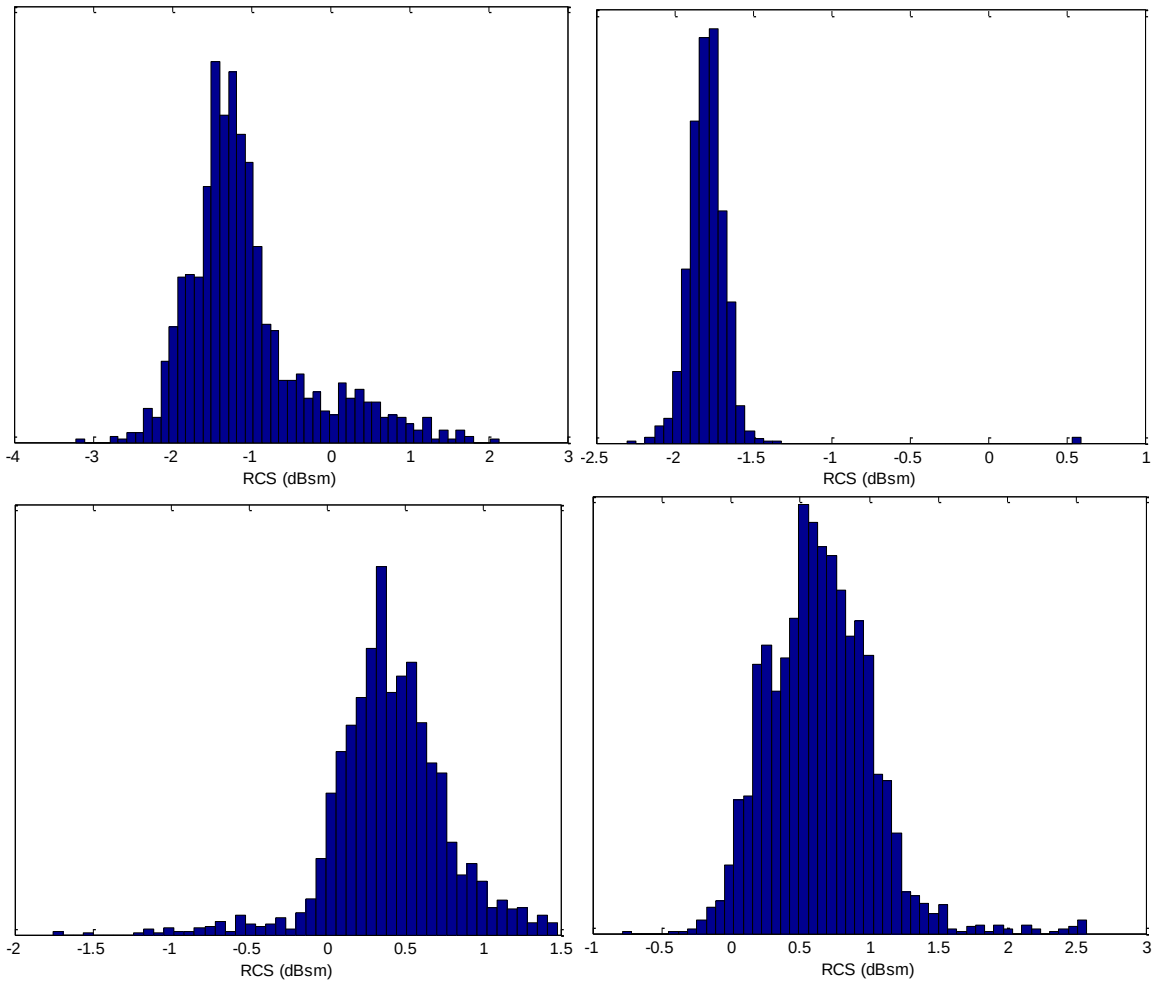

Figure 13. Asymmetric leptokurtic cases.

and a fairly symmetric distribution with a left-skewed tail (sk=-0.57, kurt=6.07). In the second graph of this figure, it is reasonable to conclude that an outlier is present; and perhaps with a properly-crafted specialty test this situation could be identified and the RCS history reprocessed (actually a non-trivial task, given that there is very little theory developed for outlier identification for non-Gaussian cases)[6]; the other three situations are likely to resist the imposition of any of the canonical distributions. It is also difficult to generate simple physical explanations for these response types.

VIII. Conclusions and Recommendations

The previous study offered two conclusions: the S1 and S2 distributions appropriately model only about 35% of the space objects, and the lognormal distribution (of any variance set) cannot replace these two distributions. The present aim was to investigate a much broader distribution set to see if a different distribution type could contribute significantly to the standing set of models, and to try to gain a better understanding of the situations in which none of the typical model types serve as viable contenders. As part of this investigation, the following specific findings and recommendations are offered:

- The S1 distribution, which corresponds to the Swerling I and II types, is actually not very broadly useful or representative, modeling only 10% of the objects in the sample set. Rayleigh scattering is obviously a much smaller contributor to the backscatter function than was previously believed. If only two standard models are to be used, there may be other superior choices to this model.
- One contender for a superior choice is the S1.5 distribution, which is a gamma distribution with location parameter 0 and shape parameter 1.5. This distribution models about 15% of the objects in the sample set. If one is looking to add a third to the two standard Swerling model that will improve the situation notably, this is the recommended addition: it is of the same family as the other Swerling models and, if included, will extend the overall object representation ability beyond 50%.

- Adding additional strains of the two-parameter gamma distribution is an exercise in diminishing return. The S2.5 distribution accommodates only about an additional 5% of the sample set objects, and this number gets even smaller with each subsequent addition.
- The three-parameter distributions investigated—gamma, Weibull, and lognormal—did not accommodate very many additional objects; many of the parameter estimates produced physically impossible situations, and those that did not were small in number. It thus seems that, when working in meters-squared space, it is reasonable to presume that RCS history PDFs begin at the origin.
- The two-parameter lognormal distribution does have a potential role in RCS history modeling. However, when the variances are fixed in order to define one-parameter lognormal distributions, the two most promising distributions would accommodate only an additional 12% of the overall objects, 6% each. In the previous analysis, a threshold of 10% was defined as the point at which adding an additional model brought enough added benefit to warrant the addition; that standard is not met here.
- For the objects that cannot be modeled by any of the standard statistical distributions, about half produce symmetric PDFs: rocket bodies the most symmetric and debris the least. Nearly all the platykurtic cases show some form of bimodality; and, as such, there is little hope of developing a model for these. The leptokurtic cases can be either fairly uncorrupted, symmetric leptokurtic PDFs (the minority) or, more commonly, leptokurtic situations corrupted with outliers or skewness. One might do something to identify the cases that conform to a model and try to use a leptokurtic distribution there (such as a student's t-distribution), but this effort would be unlikely to bear much fruit. Similarly, one could apply tests to try to remove outliers, allowing the remaining measurements to fit the hypothesized distributions more successfully. However, there is very little outlier exclusion theory developed for non-Gaussian distributions.

IX. Acknowledgements

This research was supported by Air Force Space Command, Studies and Analysis Division (A9A). The authors would like to thank Col. Cynthia Visel for her support and assistance. They would also like to thank Air Force Material Command, 850th ELSG, for making available the RCS dataset used for the study.

X. References

[1] Swerling, P., "Probability of Detection for Fluctuating Targets," *IRE Transactions on Information Theory Special Monograph Issue,* IT-6 (April 1960), pp. 269-308.
[2] Hejduk, M.D., "Improved Radar Cross-Section "Target Typing" for Spacecraft," AAS 09-301, *Advances in the Astronautical Sciences,* Vol. 135 (2009), American Astronautical Society.
[3] Stephens, M.A., "Use of the Kolmogorov-Smirnov, Cramér-von Mises and Related Statistics without Extensive Tables," *Journal of the Royal Statistical Society* B 32 (1970), pp. 115:122.
[4] Braun, H., "A Simple Method for Testing Goodness of Fit in the Presence of Nuisance Parameters," *Journal of the Royal Statistical Society* B 42 (1980), pp.53-63.
[5] Peebles, P.Z., *Radar Principles*, John Wiley & Sons, New York, 1998, Chapter 5.
[6] D'Agostino, R. and Stephens, M. *Goodness-of-Fit Techniques*, Statistics: Textbooks and Monographs, Volume 68. New York: Marcel Dekker, Inc., 1986., Chapter