Lancaster University Management School
Working Paper
2006/030

On the bias of Croston's forecasting method

Ruud Teunter and Babangida Sani

The Department of Management Science
Lancaster University Management School
Lancaster LA1 4YX
UK

On the bias of Croston's forecasting method

Ruud Teunter [a], Babangida Sani [b]

[a] *Department of Management Science, Lancaster University, LA1 4YX, UK*
[b] *Department of Mathematical Sciences, Kano University of Science and Technology, Wudil, Nigeria*

Abstract

Croston's forecasting method has been shown to be appropriate in dealing with intermittent demand items. The method, however, suffers from a positive bias as shown by Syntetos and Boylan (2001, 2005) who proposed a modification. Unfortunately, the modification ignores the damping effect on the bias of the probability that a demand occurs. This leads to overcompensation and a negative bias, which can in fact be larger than the positive bias of the original method. Levén and Segerstedt (2004) also proposed a modified Croston method, but that suffers from an even more severe bias. Building on the results of Syntetos and Boylan (2001, 2005), we propose a new modification that takes the damping effect into account. A numerical study confirms that it considerably outperforms the existing methods. Moreover, the performance is better over the entire range of relevant parameters, which avoids the need to use different methods depending on the demand categorisation as suggested by Syntetos et al. (2005).

Keywords: Forecasting; Intermittent demand; Croston's method

1. Introduction

It is not easy to forecast intermittent demand due to its erratic and sometimes lumpy nature. Many organisations in the manufacturing and especially service industries simply use single exponential smoothing. However, as was first shown by Croston (1972), this generally leads to inappropriate stock levels. Croston proposed an alternative method that takes account of both demand size and inter-arrival time between demands. The method is now widely used in industry and it is incorporated in various best selling forecasting software packages (see Syntetos et al, 2005).

Croston's method has been assessed by several authors since 1972. The literature is reviewed in detail in Section 2. Most authors come to the conclusion that Croston's method is more suitable for intermittent demand than traditional methods such as moving average and single exponential smoothing. In fact, as shown by Teunter and Duncan (2006), contradicting results can be explained by the use of inappropriate performance measures.

A disadvantage of the original Croston method is that it is positively biased. Syntetos and Boylan (2001) noted this and proposed a modification. However, as we will show in Section 3, that modification over-compensates, leading to a negative bias instead. Levén and Segerstedt (2004) also proposed a modification, but their method is even more biased as we also show in Section 3. Building on the results of Syntetos and Boylan (2001, 2005), we propose a new modification of the Croston method in Section 4. A numerical study in Section 5 confirms that this new method considerably outperforms the existing methods. This paper ends with conclusions in Section 6.

2. Literature on Croston's method

We will review the contributions chronologically. Rao (1973) made corrections to several expressions in Croston's paper without affecting the final conclusions or the forecasting procedure. Schultz (1987) presented a forecasting procedure which is basically the Croston's method and suggested a base-stock inventory policy with replenishment delays. Willemain et al. (1994) compared Croston's method with exponential smoothing and concluded that Croston's method is robustly superior to exponential smoothing, although results with real data in some cases show a more modest benefit. Johnston and Boylan (1996b) obtained similar results, but further showed that Croston's method is always better than exponential smoothing when the average inter-arrival time between demands is greater than 1.25 review intervals. Sani and Kingsman (1997) compared various forecasting and inventory control methods on some long series of low demand real data from a typical spare parts depot in the UK. They concluded based on cost and service level, that the best forecasting method is moving average followed by Croston's method.

An important contribution is that by Syntetos and Boylan (2001). They show that Croston's method lead to a biased estimate of demand per unit time. They also propose a modified method and demonstrate the improvement in a simulation experiment.

Snyder (2002) critically assessed Croston's method with a view to overcome certain implementation difficulties on the data sets used. Snyder made corrections to the underlying theory and proposed modifications. Ghobbar and Friend (2003) compared various forecasting methods using real data of aircraft maintenance repair parts from an airlines operator. The data is sporadic in nature and they showed that moving average, Holt's and Croston's forecasting methods are superior to other methods such as the exponential smoothing. Willemain, Smart and Schwarz (2004) compared various forecasting methods using large industrial data sets. They showed that the bootstrapping method produces more accurate forecasts than both exponential smoothing and Croston's method.

In an attempt to develop a forecasting procedure that can handle both fast moving and slow moving items, Levén and Segerstedt (2004) proposed a modification of Croston's method which was thought to avoid the bias indicated by Syntetos and Boylan, 2001. The modification was shown to outperform exponential smoothing based on a simulation experiment. Eaves and Kingsman (2004) compared various forecasting methods using real data from the UK's Royal Air Force. They showed that the modified Croston's method by Syntetos and Boylan (2001) is the best forecasting method for spare parts inventory control.

In an attempt to further confirm the good performance of their modified Croston's method, Syntetos and Boylan (2005) carried out a comparison of forecasting methods including theirs and the original Croston's method. A simulation exercise was carried out on 3,000 products from the automotive industry with "fast intermittent" demand. It was shown that the modification is the most accurate estimator. In another study, Syntetos, Boylan and Croston (2005) analyzed a wider range of intermittent demand patterns and made a categorisation to guide the selection of forecasting methods. They indicated that there are demand categories that are better used with the original Croston's method and there are others that go well with the Syntetos/Boylan modification.

A recent comparison by Syntetos and Boylan (2006) shows overall superior performance of the Syntetos/Boylan modification, followed by simple moving average and the original Croston's method. Another comparative study was conducted by Teunter and Duncan (2006), using a large data set from the UK's Royal Air Force. Using a new performance measure that compares target to achieved service level, they showed that the original Croston's method as well as the Syntetos & Boylan and the Levén & Segerstedt variants outperform moving average and exponential smoothing.

3. Theoretical Background

Croston's original method forecasts separately the time between consecutive transactions p_t and the magnitude of the individual transactions z_t. At the review period t, if no demand occurs in a review period then the estimates of the demand size and inter-arrival time at the end of time t, $\hat{z}_t$ and $\hat{p}_t$ respectively, remain unchanged. If a demand occurs so that z_t >0, then the estimates are updated by

$$\begin{aligned}\hat{z}_t &= \alpha z_t + (1-\alpha)\hat{z}_{t-1} \\ \hat{p}_t &= \alpha p_t + (1-\alpha)\hat{p}_{t-1}\end{aligned},$$

where α is a smoothing constant between zero and one. Hence, the forecast of demand per period at time t is given as

$$C_t = \frac{\hat{z}_t}{\hat{p}_t}.$$

3.1. *Syntetos & Boylan modification*

Syntetos and Boylan (2001) pointed out that Croston's original method is biased. They showed that

$$E(C_t) = E\left[\frac{\hat{z}_t}{\hat{p}_t}\right] \approx \frac{\mu}{p}\left(1 + \frac{\alpha}{2-\alpha}\frac{p-1}{p}\right) \tag{1}$$

and, in particular, for α=1 that

$$E(C_t) = E\left[\frac{\hat{z}_t}{\hat{p}_t}\right] = E\left[\frac{z_t}{p_t}\right] = \mu\left[-\frac{1}{p-1}\ln\left(\frac{1}{p}\right)\right]. \tag{2}$$

Based on (1) and ignoring the term $\frac{p-1}{p}$, Syntetos and Boylan proposed a new estimator given as

$$SB_t = \left(1 - \frac{\alpha}{2}\right)\frac{\hat{z}_t}{\hat{p}_t}.$$

One can expect this new estimator to perform better as $\frac{p-1}{p}$ gets closer to one, i.e., as the probability $1/p$ of positive demand in a period gets smaller. This effect is illustrated in Figure 1, where the bias of the original Croston method and Syntetos & Boylan modification are compared. Note that the non-monotone behaviour is caused by the randomness of demand, as can also be seen from the differences for the two demand series.

INSERT FIGURE 1 ABOUT HERE

It is clear from Figure 1, and the above analysis explains why this also holds in general, that Croston's original method has a smaller (positive) bias if $1/p$ is large (few demands are zero), and the Syntetos/Boylan modification has a smaller bias if $1/p$ is small (many demands are zero). This also explains and confirms the findings by Syntetos, Boylan and Croston (2005) that were discussed in Section 1.

3.2. *Levén & Segerstedt modification*

Levén and Segerstedt (2004) modified Croston's method in an attempt to obtain a method that works for both slow and fast moving items. Their estimator is updated as follows

$$LS_t = \alpha \frac{z_t}{p_t} + (1-\alpha) LS_{t-1}.$$

They referred to the above summarized results in Syntetos and Boylan (2001) on the bias for Croston's original method, but remarked that their estimator does not suffer from such a bias. However, it does! Indeed, using (2) it follows that

$$E(LS_t) = \mu \left[-\frac{1}{p-1} \ln\left(\frac{1}{p} \right) \right].$$

It can easily be shown that there is always a positive bias and that this bias gets worse as the probability $1/p$ of a demand decreases. The bias is indeed very large; it is more than 50% when a demand occurs in 1 out of 3 periods or less. Figure 2 shows that if the probability of a demand is 0.2 and the average demand size is μ = 5, then the Levén & Segerstedt estimate fluctuates around 2 ($E(LS_t) = 2.02$ to be exact) whereas the expected demand per period is μ/p = 1.

INSERT FIGURE 2 ABOUT HERE

4. New method

As the Syntetos & Boylan method, the new method is based on (1), but it does not ignore the term $\frac{p-1}{p}$. Based on rewriting (1) as

$$E\left[\frac{\hat{z}_t}{\hat{p}_t} \right] \approx \frac{\mu}{p} \left(\frac{2p-\alpha}{2p-\alpha p} \right),$$

we suggest the following estimator:

$$\frac{\hat{z}_t}{\hat{p}_t}\frac{\left(2\hat{p}-\alpha\hat{p}\right)}{\left(2\hat{p}-\alpha\right)}=\frac{\hat{z}_t\left(2-\alpha\right)}{2\hat{p}-\alpha}=\left(1-\frac{\alpha}{2}\right)\frac{\hat{z}_t}{\hat{p}_t-\frac{\alpha}{2}}.$$

In the next section, we will compare the bias of this new method to that of the other methods in a numerical study.

We note that the Levén & Segerstedt method cannot be modified in a similar way, at least not without abandoning its core idea of only updating the demand per period.

5. Numerical comparison of new and existing methods

We compare the new modification to the original Croston's method, the Syntetos & Boylan modification and the Levén & Segerstedt modification in a numerical study. We employ a full factorial design, where we vary the smoothing parameter α (0.1, 0.2, 0.3), the probability of a demand 1/p (0.1, 0.3, 0.5, 0.7), and the type and variance of the demand distribution (normal with mean 1 and variance 0.1, normal with mean 1 and variance 0.3, discrete uniform between 1 and 2, discrete uniform between 1 and 10). We remark that, based on the results in Section 3, the main determinants of the bias are expected to be α and 1/p.

For each of the 3 · 4 · 4 = 48 experiments, a demand series of 10,000 periods is generated randomly. All methods are initialized with the correct values. The reported biases are averaged over all 10,000 periods.

Table 1 gives the complete results.

INSERT TABLE 1 ABOUT HERE

The results confirm our expectation that the smoothing parameter, α, and the probability of a demand, 1/p, are the main determinants of the biases of the different methods, and that type and variance of the demand distribution have little effect.

As expected based on the results in Section 3, the Levén & Segerstedt modification performs very poorly with an average bias of 71%. The original Croston method and the Syntetos & Boylan modification perform considerably better. However, the original Croston method still performs poorly when 1/p is small (with biases up to 18%), whereas the Syntetos & Boylan modification performs poorly when 1/p is large (with biases up to 12%). The performance of both methods deteriorates when α increases.

Although there are a couple of experiments where either the original Croston or the Syntetos & Boylan modification has the smallest (absolute) bias, the new method generally outperforms the existing methods for the entire range of considered parameter values. The average absolute bias is 1% for the new method as compared to 6% for the original Croston method and 5% for the Syntetos & Boylan modification.

6. Conclusions

Building on results in the literature, a new modification of Croston's method for forecasting intermittent demand was proposed. In a comparative numerical study, this new method was shown to significantly outperform existing methods. The average absolute bias for the new method was 1% as compared to 5%, 6% and 71% for the original Croston method, the Syntetos & Boylan modification and the Levén & Segerstedt modification, respectively. Furthermore, contrary to existing methods, performance is well for small as well as large demand intervals and does not deteriorate as the smoothing constant increases. The robustness implies that the new method can be used in all cases and avoids the need for demand categorization.

References

Croston, J.D., 1972. Forecasting and stock control for intermittent demands. Operational Research Quarterly 23, 289-303.

Eaves, A.H.C., Kingsman, B.G., 2004. Forecasting for the ordering and stock-holding of spare parts. Journal of the Operational Research Society 55, 431-437.

Ghobbar, A.A., Friend, C.H., 2003. Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. Computers and Operations Research 30, 2097-2114.

Johnston, F.R., Boylan, J.E., 1996a. Forecasting for items with intermittent demand. Journal of the Operational Research Society 47, 113-121.

Johnston, F.R., Boylan, J.E., 1996b. Forecasting intermittent demand: a comparative evaluation of Croston's method. International Journal of forecasting 12, 297-298.

Levén, E., Segerstedt, A., 2004. Inventory control with a modified Croston procedure and Erlang distribution. International Journal of Production Economics 90, 361-367.

Rao, A.V., 1973. A comment on: Forecasting and stock control for intermittent demands. Operational Research Quarterly 24, 639-640.

Sani, B., Kingsman, B.G., 1997. Selecting the best periodic inventory control and demand forecasting methods for low demand items. Journal of the Operational Research Society 48, 700-713.

Schultz, C.R., 1987. Forecasting and inventory control for sporadic demand under periodic review. Journal of the Operational Research Society 38, 453-458.

Snyder, R., 2002. Forecasting sales of slow and fast moving inventories. European Journal of Operational Research 140, 684-699.

Syntetos, A.A., Boylan, J.E., 2001. On the bias of intermittent demand estimates. International Journal of Production Economics 71, 457-466.

Syntetos, A.A., Boylan, J.E., 2005. The accuracy of intermittent demand estimates. International Journal of forecasting 21, 303-314.

Syntetos, A.A., Boylan, J.E., 2006. On the stock control performance of intermittent demand estimators. To appear in International Journal of Production Economics.

Syntetos, A.A., Boylan, J.E., Croston, J.D., 2005. On the categorisation of demand patterns. Journal of the Operational Research Society 56, 495-503.

Teunter, R.H., Duncan, L, 2006. Forecasting demand for slow moving inventory. Research Report, Lancaster University Management School.

Willemain, T.R., Smart, C.N., Schwarz, H.F., 2004. A new approach to forecasting intermittent demand for service parts inventories. International Journal of forecasting 20, 375-387.

Willemain, T.R., Smart, C.N., Shockor, J.H., DeSautels, P.A., 1994. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. International Journal of forecasting 10, 529-538.

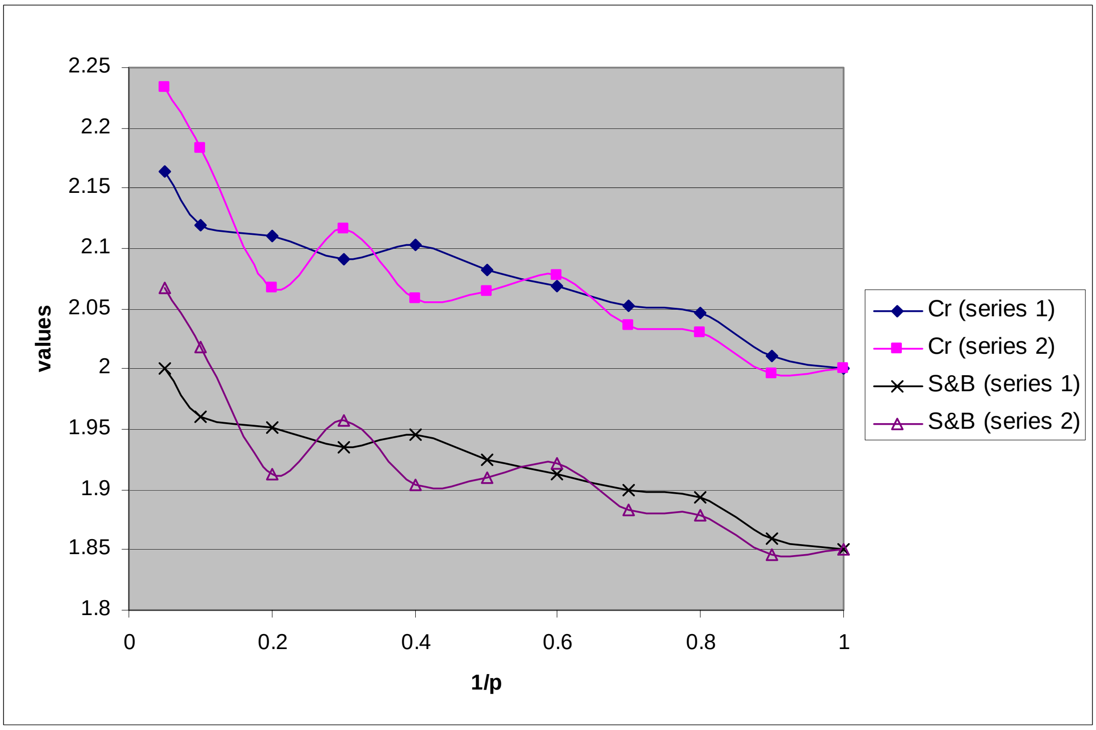

Figure 1. Comparison of the average bias over 10,000 periods for the Croston (Cr) and Syntetos & Boylan (S&B) methods for two randomly generated demand series (α = 0.15, μ/p = 2, σ = 0; both methods initialized using the correct values).

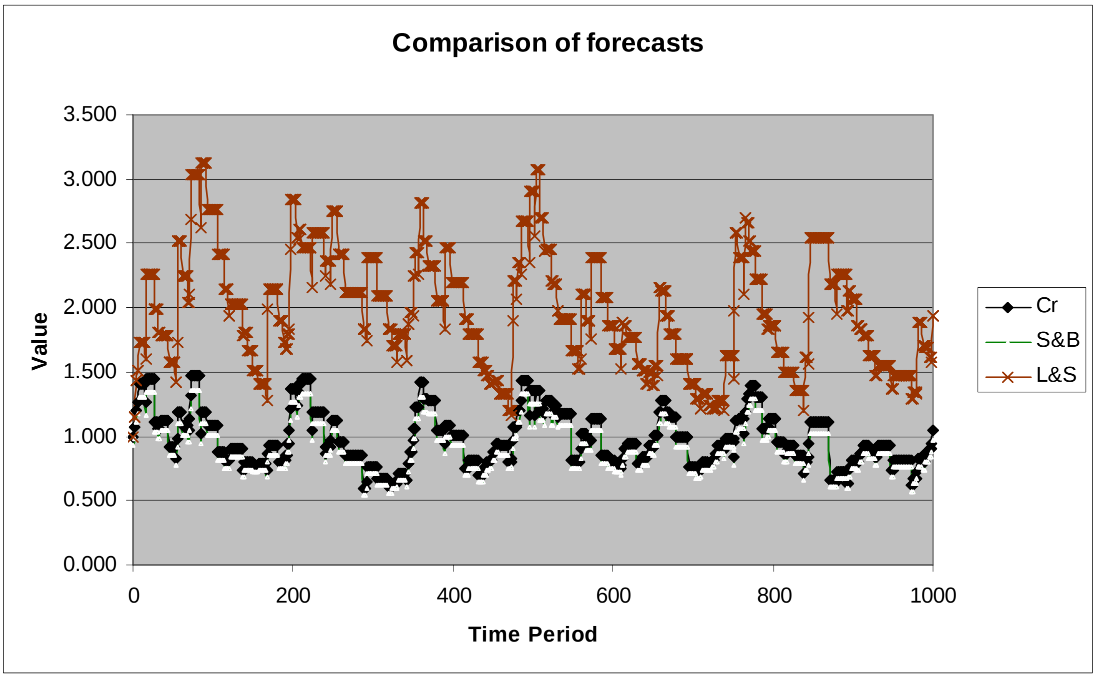

Figure 2. Comparison of the average bias over 10,000 periods for the Croston (Cr), Syntetos & Boylan (S&B) and Levén & Segerstedt (L&S) methods for a randomly generated demand series (α = 0.15, 1/p = 0.2, μ = 5, σ = 1; all methods initialized using the correct values).

`			**Demand**	**Value**	**Percentage Bias**			
S/N	**α**	**1/p**	**Distribution**	**Expected**	**Croston**	**Syntetos & Boylan**	**Levén & Segerstedt**	**New**
1	0.1	0.1	N(1,0.1)	0.100	1%	-4%	151%	**-3%**
2	0.1	0.1	N(1,0.3)	0.100	3%	-2%	153%	**-1%**
3	0.1	0.1	U(1,2)	0.150	**2%**	-3%	149%	-3%
4	0.1	0.1	U(1,10)	0.550	8%	**2%**	139%	3%
5	0.1	0.3	N(1,0.1)	0.300	3%	-3%	71%	**-1%**
6	0.1	0.3	N(1,0.3)	0.300	2%	-3%	73%	**-1%**
7	0.1	0.3	U(1,2)	0.450	4%	-1%	72%	**1%**
8	0.1	0.3	U(1,10)	1.650	3%	-2%	68%	**-1%**
9	0.1	0.5	N(1,0.1)	0.500	2%	-3%	37%	**0%**
10	0.1	0.5	N(1,0.3)	0.500	2%	-3%	41%	**-1%**
11	0.1	0.5	U(1,2)	0.750	3%	-2%	39%	**1%**
12	0.1	0.5	U(1,10)	2.750	4%	**-1%**	40%	2%
13	0.1	0.7	N(1,0.1)	0.700	2%	-3%	19%	**0%**
14	0.1	0.7	N(1,0.3)	0.700	2%	-3%	22%	**0%**
15	0.1	0.7	U(1,2)	1.050	2%	-3%	19%	**1%**
16	0.1	0.7	U(1,10)	3.850	**0%**	-5%	18%	-1%
17	0.2	0.1	N(1,0.1)	0.100	7%	-3%	146%	**-2%**
18	0.2	0.1	N(1,0.3)	0.100	**4%**	-6%	143%	-5%
19	0.2	0.1	U(1,2)	0.150	9%	-1%	142%	**0%**
20	0.2	0.1	U(1,10)	0.550	8%	-3%	166%	**-1%**
21	0.2	0.3	N(1,0.1)	0.300	7%	-4%	74%	**0%**
22	0.2	0.3	N(1,0.3)	0.300	7%	-4%	75%	**-1%**
23	0.2	0.3	U(1,2)	0.450	6%	-4%	71%	**-1%**
24	0.2	0.3	U(1,10)	1.650	6%	-5%	72%	**-1%**
25	0.2	0.5	N(1,0.1)	0.500	6%	-5%	39%	**1%**
26	0.2	0.5	N(1,0.3)	0.500	6%	-5%	41%	**1%**
27	0.2	0.5	U(1,2)	0.750	5%	-6%	38%	**0%**
28	0.2	0.5	U(1,10)	2.750	7%	-4%	40%	**2%**
29	0.2	0.7	N(1,0.1)	0.700	3%	-7%	19%	**0%**
30	0.2	0.7	N(1,0.3)	0.700	3%	-7%	21%	**0%**
31	0.2	0.7	U(1,2)	1.050	3%	-8%	18%	**0%**
32	0.2	0.7	U(1,10)	3.850	4%	-7%	19%	**1%**
33	0.3	0.1	N(1,0.1)	0.100	15%	-2%	141%	**0%**
34	0.3	0.1	N(1,0.3)	0.100	15%	-3%	165%	**-1%**
35	0.3	0.1	U(1,2)	0.150	18%	**0%**	165%	2%
36	0.3	0.1	U(1,10)	0.550	12%	-5%	159%	**-3%**
37	0.3	0.3	N(1,0.1)	0.300	12%	-5%	73%	**1%**
38	0.3	0.3	N(1,0.3)	0.300	14%	-3%	78%	**2%**
39	0.3	0.3	U(1,2)	0.450	12%	-5%	74%	**1%**
40	0.3	0.3	U(1,10)	1.650	12%	-5%	72%	**1%**
41	0.3	0.5	N(1,0.1)	0.500	6%	-10%	37%	**-2%**
42	0.3	0.5	N(1,0.3)	0.500	9%	-8%	41%	**1%**
43	0.3	0.5	U(1,2)	0.750	9%	-7%	39%	**1%**
44	0.3	0.5	U(1,10)	2.750	10%	-6%	40%	**3%**
45	0.3	0.7	N(1,0.1)	0.700	4%	-12%	18%	**0%**
46	0.3	0.7	N(1,0.3)	0.700	5%	-10%	22%	**1%**
47	0.3	0.7	U(1,2)	1.050	5%	-11%	19%	**1%**
48	0.3	0.7	U(1,10)	3.850	5%	-11%	19%	**0%**

Table 1. Comparison of biases for new and existing methods. The smallest bias for each of the 48 examples is indicated in bold.