

What Makes the XmR Chart Work?

How does it separate the signals from the noise?

Donald J. Wheeler

There are two basic ideas or principles that need to be respected when creating a chart for individual values and a moving range (an XmR Chart). This article will explain and illustrate these two principles for effective XmR Charts.

The first principle for an effective XmR Chart is that successive values need to be logically comparable. The second is that the moving ranges need to isolate and capture the local, short-term, routine variation that is inherent in the data. When these principles are ignored the XmR Chart can miss signals that it would otherwise detect.

While the use of the time-order sequence of the data will usually be sufficient to satisfy these two principles for an effective XmR Chart, there are times when a careful consideration of the structure in the data will require a different organization. As a case in point consider the Camshaft Bearing Diameters shown in Figure 1. The strict time-order for these data is:

Camshaft No. 1; Bearing 1, Bearing 2, Bearing 3, then
Camshaft No. 2; Bearing 1, Bearing 2, Bearing 3, etc.

<i>Camshaft</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>10</i>	<i>11</i>	<i>12</i>	<i>13</i>
Bearing 1	49	51	51	52	56	52	51	52	50	49.5	51	52	53
Bearing 2	50	50	52	49	52	49	49	50	48	48	49	48	50
Bearing 3	50	46	52	51	42	50	50.5	44	48	49	43	49.5	50
<i>Camshaft</i>	<i>14</i>	<i>15</i>	<i>16</i>	<i>17</i>	<i>18</i>	<i>19</i>	<i>20</i>	<i>21</i>	<i>22</i>	<i>23</i>	<i>24</i>	<i>25</i>	<i>26</i>
Bearing 1	52	51	51	55	51	50.5	49	51	51	52.5	50	50	50
Bearing 2	49	49	51	51	52	50	50	48	49	50	50	48	50
Bearing 3	46	49	43	42	45	42	43	43	45	42	46	47	45
<i>Camshaft</i>	<i>27</i>	<i>28</i>	<i>29</i>	<i>30</i>	<i>31</i>	<i>32</i>	<i>33</i>	<i>34</i>	<i>35</i>	<i>36</i>	<i>37</i>	<i>38</i>	<i>39</i>
Bearing 1	53	52	50	53	52	52	51.5	51	49.5	52	51	51.5	51
Bearing 2	47	48	49	48	52	52	53	53	51	51	51.5	49	54.5
Bearing 3	42	49	49	52	46	50	51	50	51	50	52	52	54
<i>Camshaft</i>	<i>40</i>	<i>41</i>	<i>42</i>	<i>43</i>	<i>44</i>	<i>45</i>	<i>46</i>	<i>47</i>	<i>48</i>	<i>49</i>	<i>50</i>		
Bearing 1	51	50	50.5	51	51	51	56	50	50	52.5	57		
Bearing 2	50	48	50	47	49	49	48	50	52	48	48		
Bearing 3	52.5	54	51	51	51	49.5	52	49	49	50	50		

Figure 1: The Camshaft Bearing Diameters

When the data are arranged in time order the Average Moving Range is 3.10. With an Average of 49.81 we obtain the limits are shown in Figure 2.

$$UNPL = 49.81 + 2.66 (3.10) = 58.1$$

$$LNPL = 49.81 - 2.66 (3.10) = 41.6$$

$$URL = 3.268 (3.10) = 10.1$$

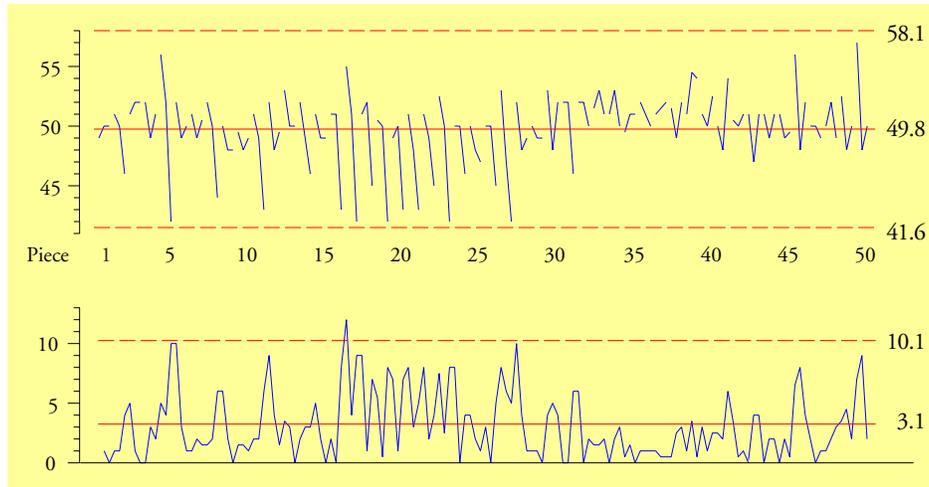


Figure 2: An *XmR* Chart for the Camshaft Bearing Diameters in Time Order

Except for one moving range, all of the values in Figure 2 fall within the computed limits. However, with this arrangement of these data the moving ranges represent the differences between Bearing One, Bearing Two, and Bearing Three. Since the three bearings were produced by three separate processes operating in parallel there could be systematic differences between the three bearings that have nothing to do with routine variation. Even though the three bearings are supposed to be the same, the fact that they are produced by parallel operations makes it a bit naive to assume that they are indeed the same.

The idea of checking your data for homogeneity that is the essence of the use of a process behavior chart is based on a skeptical view of homogeneity, not a gullible one. Therefore, rather than using the strict time order shown in Figure 2, a more rational approach in this case is to organize these data according to bearing number, and then to use the time order within each bearing to create moving ranges.

Bearing No. 1; Camshaft 1, Camshaft 2, ..., Camshaft 50, then
 Bearing No. 2; Camshaft 1, Camshaft 2, ..., Camshaft 50, then
 Bearing No. 3; Camshaft 1, Camshaft 2, ..., Camshaft 50

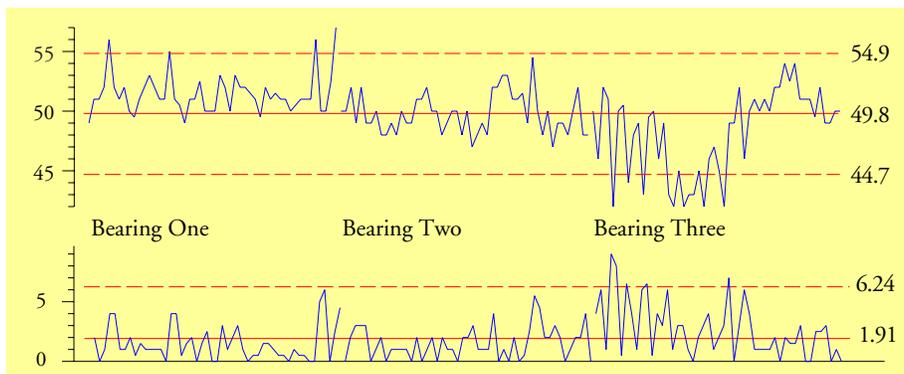


Figure 3: An *XmR* Chart for the Camshaft Bearing Diameters Organized by Bearing Number

When these data are organized in this manner the moving ranges will represent the natural, short-term, routine variation within each production process rather than the differences, if any, between the three processes. Now the Average Moving Range is 1.91, giving the limits of 44.7 to 54.9 shown in Figure 3. This organization and these limits allow us to see the differences between the operations producing these three bearings with greater clarity.

Thus, when there is a structure within your data, it is imperative that you consider that structure when organizing the data for an *XmR* Chart. If there are logical partitions or subsets in your data, isolate those subsets from each other so that successive values will be logically comparable and the moving ranges can characterize the routine variation rather than the differences between the subsets.

In fact, for the Camshaft Bearing Diameters we could take the next step and compute a separate set of limits for each bearing. These *XmR* Charts are in Figure 4. Bearing One has an average of 51.46 and an average moving range of 1.47. Bearing Two has an average of 49.78 and an average moving range of 1.53. Bearing Three has an average of 48.20 and an average moving range of 2.69.

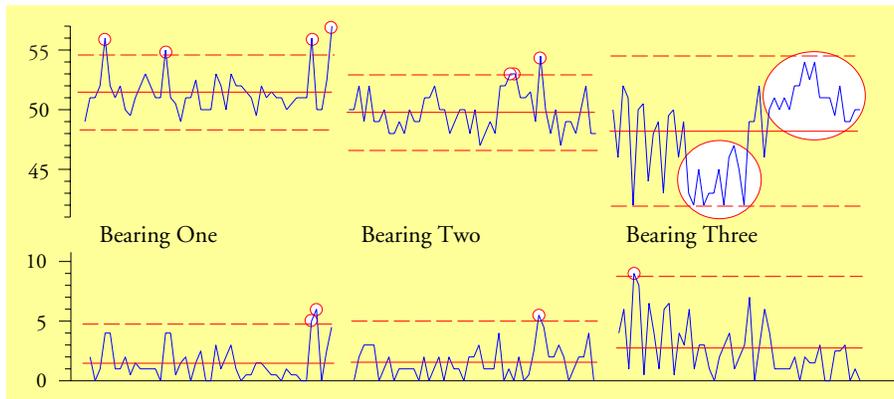


Figure 4: Three Separate *XmR* Charts for the Camshaft Bearing Diameter Data

Not only are these three processes operating at different averages, but each process shows evidence of unpredictable operation. Moreover, the operation producing Bearing Three is seen to be qualitatively different from the other two bearing operations. This ability to show qualitative as well as quantitative differences is one of the real advantages of a process behavior chart.

Both Figures 3 and 4 are superior to Figure 2 simply because they respect the two principles for effective *XmR* charts: They organize the data so that successive values are logically comparable and the moving ranges capture the routine variation in each of the three production operations. Once you have organized the data in a rational manner, there may be several ways to use the limits to tell the story that is contained within the data. The objective is understanding and insight rather than computing a particular value. There is an element of judgment involved in using a process behavior chart, and this element cannot be removed. It is essential to an effective analysis.

WHAT DO WE GAIN FROM THE MOVING RANGE CHART?

It has been suggested that the Moving Range Chart adds so little to the Chart for Individual

Values that you should not bother to show it—“Simply show the *X* Chart and forget the *mR* Chart.” The basis for this recommendation seems to be the documented fact that the combined *XmR* Chart does not have an appreciably greater ability to detect signals than does the *X* Chart alone.

However, this mathematical analysis overlooks the interpretative benefits to be gained by including the *mR* Chart. In Figure 3 the moving ranges confirm the impression that the process for Bearing Three has more variation than the other two processes. In Figure 4 consider how the point above the upper range limit for Bearing Three identifies a change in the process between Camshaft 4 and Camshaft 5. This change is too large to have occurred by chance alone. Something was changed, and since the target value is 50, this change was detrimental. Thus, the Moving Range Chart will, on occasion, provide new information in addition to reinforcing the message of the *X* Chart.

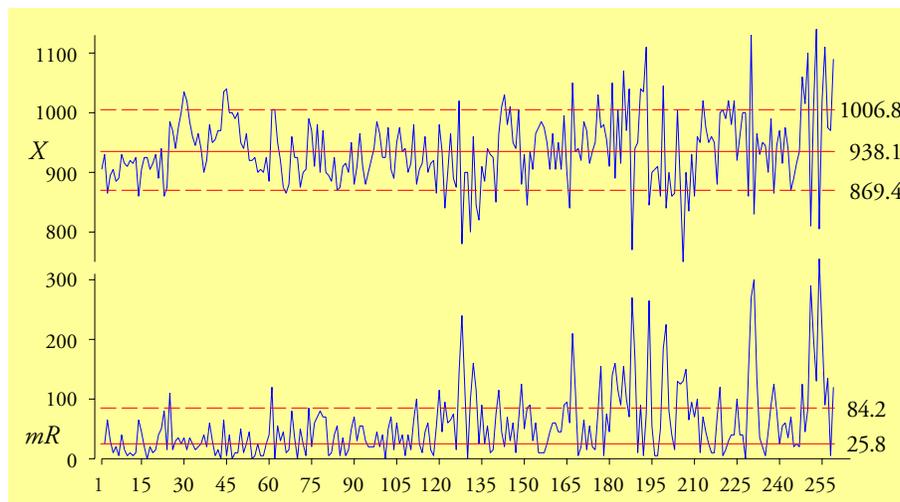


Figure 5: The *XmR* Chart for the Batch Weight Data

When computing the limits for the Batch Weight Data in Figure 5, I only used the first 58 moving ranges to compute the Average Moving Range. Why did I do this? Inspection of the *mR* Chart shows an increase in the process variation after value 60. Since the objective is to compute limits that characterize what the process is capable of doing, the first 58 moving ranges do this better than the rest.

How did I detect this shift in variation in Figure 5? At the scale used here it is hard to see, but only 22 of the 61 values between value 60 and value 120 are below the central line. While we do not use the traditional run tests with a moving range chart, shifts such as this one which are shown by a substantial number of values may still be interpreted as changes in the process variation. For example, beginning around value 120 there is a second upward shift in the process variation in Figure 5. And around value 180 there appears to be a third upward shift. Since these shifts all make sense in the context of these data, it is reasonable to interpret them as being real.

Why do we not use the traditional run test rules with the moving range chart? While the usual run tests such as eight in a row on one side of the central line, two out of three beyond two-sigma, and four out of five beyond one sigma may be used with the *X* chart whenever the order

of the points makes sense, you should not use these tests with the moving range values. This prohibition is due to the nature of the computation of the moving ranges. Since each individual value is used to create two moving ranges, the computations can create correlations within the moving range values. To see how this happens look at the first three X values that are above the upper natural process limit on the chart for Bearing One in Figure 6.

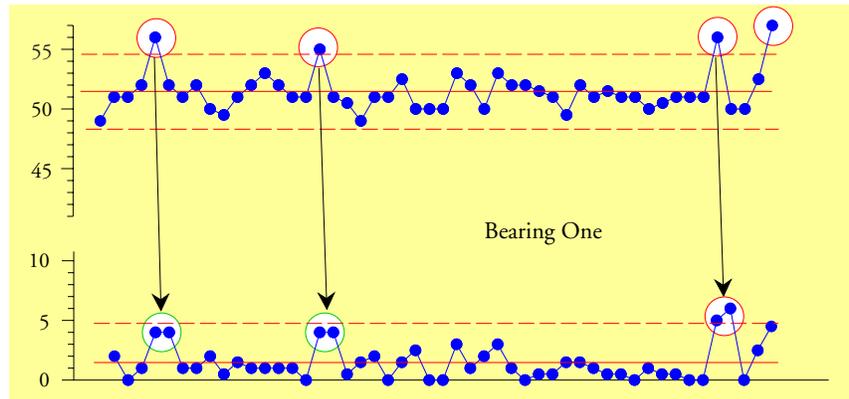


Figure 6: The XmR Chart for Camshaft Bearing One

Each one of these values generated two large moving range values. Since this correlation structure is an artifact of the computations, it is something that we do not want to interpret as a signal. To avoid false alarms from run tests due to this artifact of the computations it is best to avoid using any of the traditional run tests with the moving range values. Moving Ranges above the upper range limit will denote breaks in the original data, and as such are valid signals. To interpret shifts and long runs on the moving range chart you should use substantially more data than the traditional run tests use.

ADDITIONAL REASONS TO USE THE MOVING RANGE CHART

The second reason that I cannot agree to the suppression of the mR Chart is that there are many people, and many software packages, that actually compute *three standard deviation limits* rather than *three-sigma limits*. If you are shown a naked X Chart you will have no way of knowing if the limits have been computed correctly. However, if you are shown an XmR Chart, you will immediately have a higher level of confidence that the limits have been computed correctly. Moreover, by using the central line of the mR Chart, you can quickly check to see if the limits are indeed correctly computed. Thus, the mR Chart is the secret handshake of those who know the correct way of computing limits for an X Chart. Omit it and your readers cannot be sure that you are a member of the club.

Finally, the mR Chart allows you and your audience to check for the problem of chunky data. This is a problem that occurs when the data have been rounded to the point that the variation is lost in the round-off. When this happens the many moving ranges that get rounded to zero will deflate the average moving range, which will tighten the limits. At the same time that the limits are being tightened the round-off will restrict the number of values for both the original data and the moving ranges. This restriction will keep the running record from shrinking with the limits and as a result, the chart will start to show an increased number of false alarms. For a more

complete treatment of the problem of chunky data see my column for December of 2011. Since the only way to check for chunky data with an X chart is to use the moving range chart, good form requires that the mR chart be shown along with the X chart.

SUMMARY

Thus, while on the basis of mathematical theory there may be little justification for showing the mR Chart, there are three practical reasons to do so, any one of which is sufficient to justify the inclusion of the mR Chart with your X Chart. Displaying the mR chart along with your X chart is full disclosure. Anything less is suspect.