Camera Dynamic Range (DR)
and Signal to Noise Ratio (SNR)

There are several ways to view dynamic range (DR). At base ISO the sensor can record the highest useful value or signal. If we increase the exposure we will pass the "full well" capacity of the sensor.

The useful range means that we can easily get a linear response from the sensor as we reduce the exposure. Once the signal is no longer proportional to the amount of light being recorded the camera is at its maximum analog limit. The bottom of the analog range might be a way to define the DR.

With the lens cap in place, there will be no useful signal, only dark current or residual noise. By the time we reach that level we will probably have passed the point where a linear response is possible. We could also define DR based on that lower limit but by then the amount of noise will be quite visible.

The level of visible noise we can tolerate is subjective. So we can also define DR as the range from the brightest to the darkest image that is, not only linear, but also above the noise level that is acceptable.

Each of these definitions of DR is different but they all start at base ISO. Increasing the ISO does not change the DR of the camera.

When we raise the ISO setting above base we no longer record the largest possible signal. To keep the digital value from exceeding the binary upper limit (16383 for a 14-bit raw file) the exposure needs to be limited. This reduces the signal and amplifies the noise. The signal to noise ratio drops.

Many of us use measurements compiled by DXOMARK (DXO) to see how a camera reacts to changes in ISO. Here is an example for the Nikon Z7:

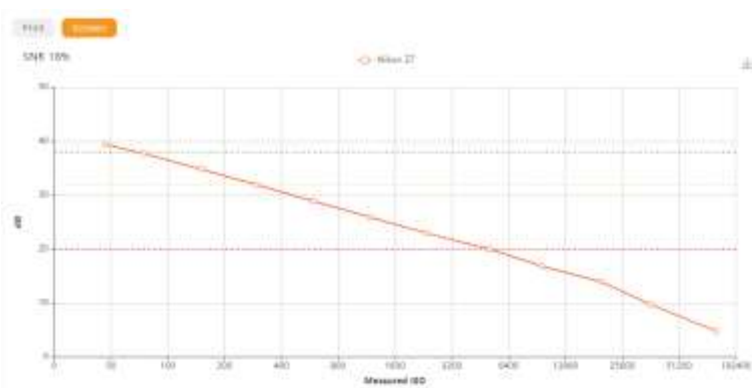
[1]

The vertical scale is in decibels (dB). Note that the change in SNR is almost linear.

But SNR is different from DR. SNR is the ratio of the signal, in this case 18% gray (EC+0), to a particular level of noise. If we know the signal that produces 18% gray (close to 14-bit raw 1024 at base ISO or 2^10) then we can calculate the noise DXO assumed at that level.

A DR plot shows similar results but the vertical scale is in base 2 logarithms, Ev.

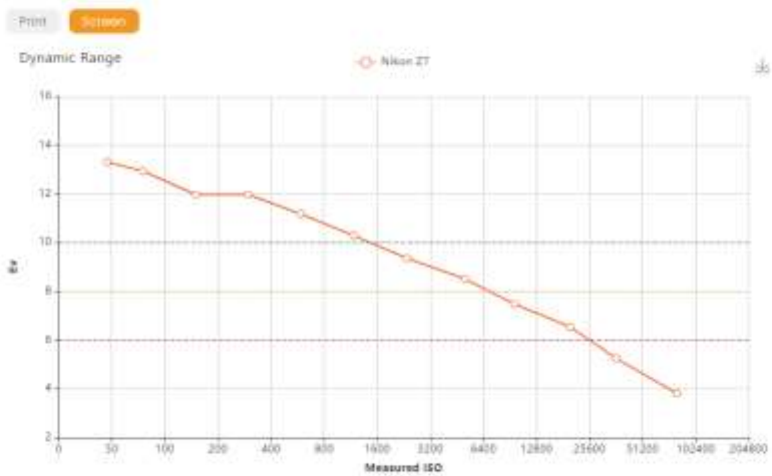

2

The dB and Ev scales are directly proportional, 1 Ev equals 6.0206 dB.

This appears to show that something is happening in the camera due to a change in how the gain is handled for low and high ISO ranges. But this gain change is not reflected in the SNR plots. It does not affect the signal or the noise used for SNR.

The horizontal scale shows a progression of ISO steps. The DXO label is "Measured ISO" for the plotted points. The manufacturer's ISO is higher. The relevance of DXO 's measured ISO is not clear.

There is also something different about the two DXO DR plot for the Z7. It is not strictly linear. It has a break in continuity between ISO 200 and 400.

A more precise presentation of DR can be found at Photons to Photos (P2P) which displays values in 1/3 stop increments. The dynamic range is calculated a little differently. The points are plotted at the manufacturer's ISO setting. The solid points are for the normal analog range. Open points above or below the normal range are manipulated to artificially extend the results.

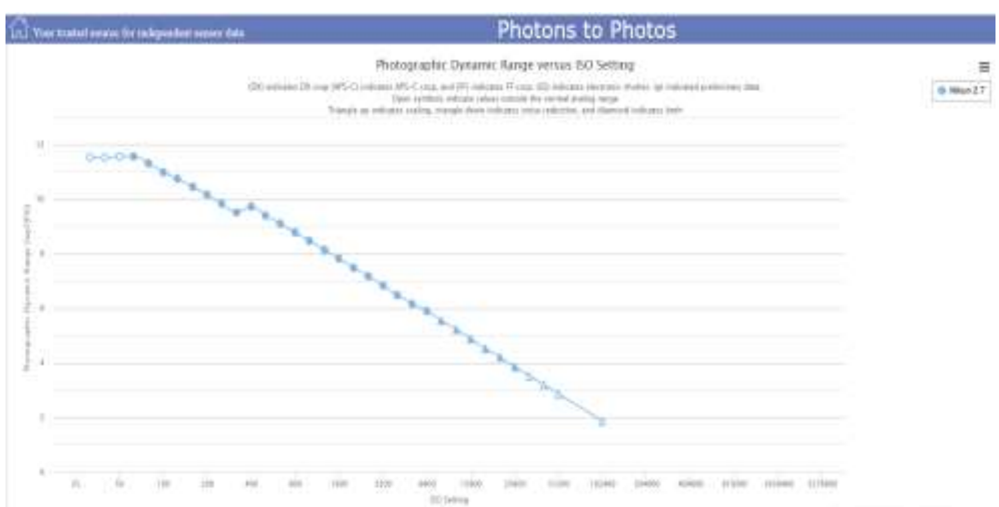

3

The P2P plot narrows down the break in linearity to between ISO 320 and 400.

The information displayed by P2P is based on data compiled by DXO.

P2P also provides a wealth of information about what is happening under the covers as the information recorded by the sensor is saved as a digital record.

Here is another example for a Sony A7 III:

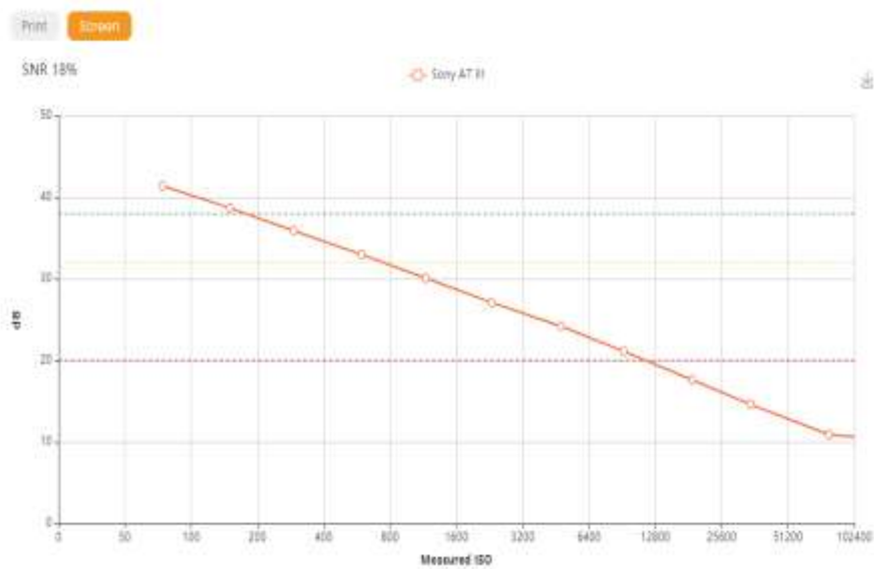

4

There are no discontinuities in the SNR plot. Since both the signal and SNR lines are straight, the noise line must also be straight, almost constant. Nevertheless:

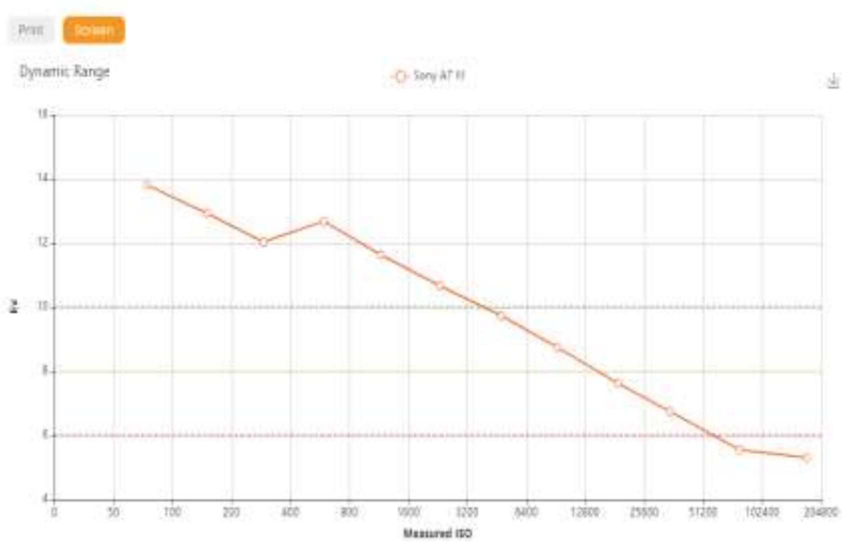

5

The DR plot might show a discontinuity for cameras with dual gain processors.

The primary purpose of these plots is to provide a simple way to compare the noise characteristics of two or more cameras. Here are two comparisons from for the Nikon Z7 and the Sony A7 III. By selecting the “Print” tab the cameras with different resolutions can be fairly compared.

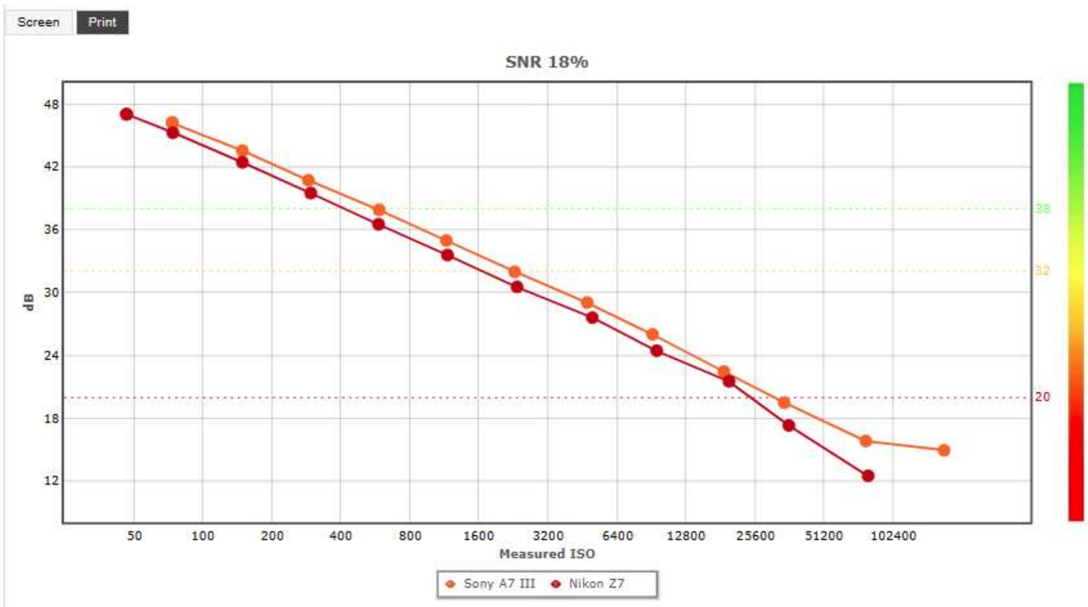

6

The 24MP Sony shows a slightly better SNR performance than the 45MP Nikon.

But the DR comparison shows something different:

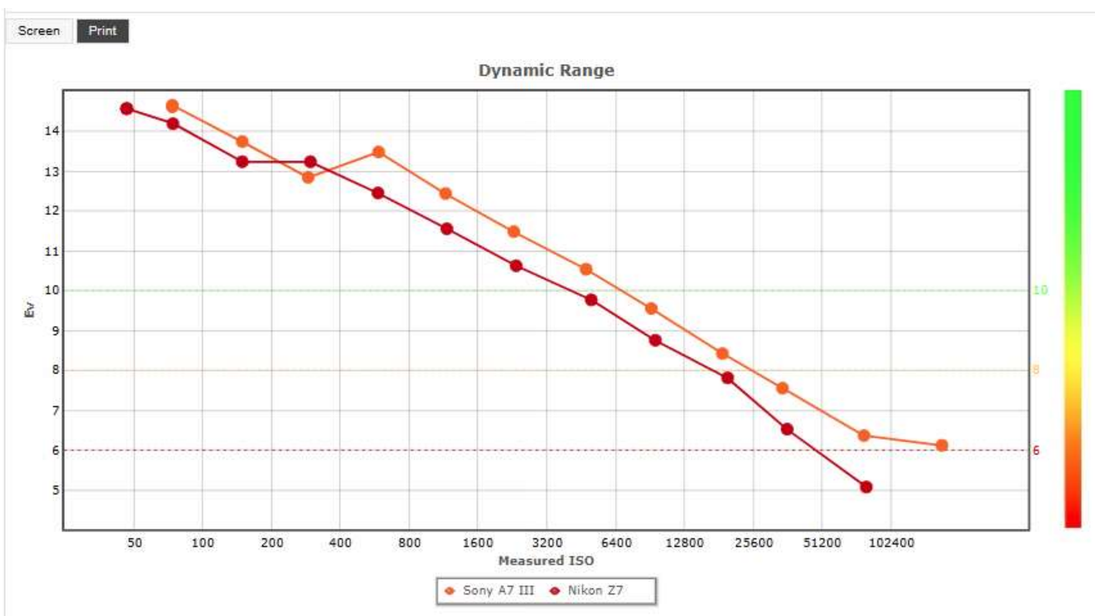

7

Here is the DR information from P2P for the A7 III:

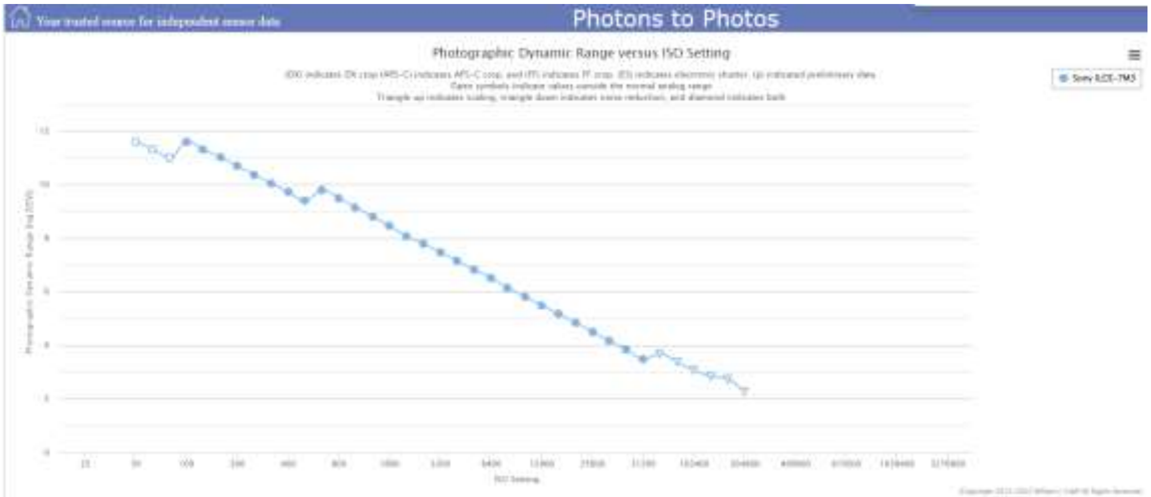

8

In P2P the discontinuity happens between ISO 500 and 640.

P2P only shows DR, no SNR information. The P2P definition of DR more conservative than DXO's.

A sudden change in read noise may not translate into a discontinuity in the DR. As an example, see [Sony A7S DR-Pix Read Noise (photonstophotos.net)]

The article illustrates how read noise changes between ISO 1600 and 2000.

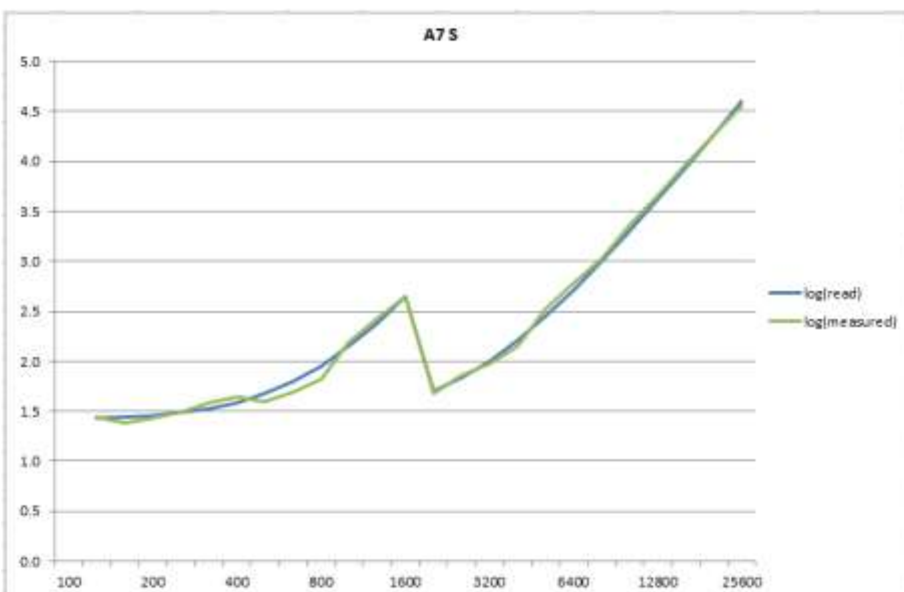

9

Yet DXO does not show a significant discontinuity.

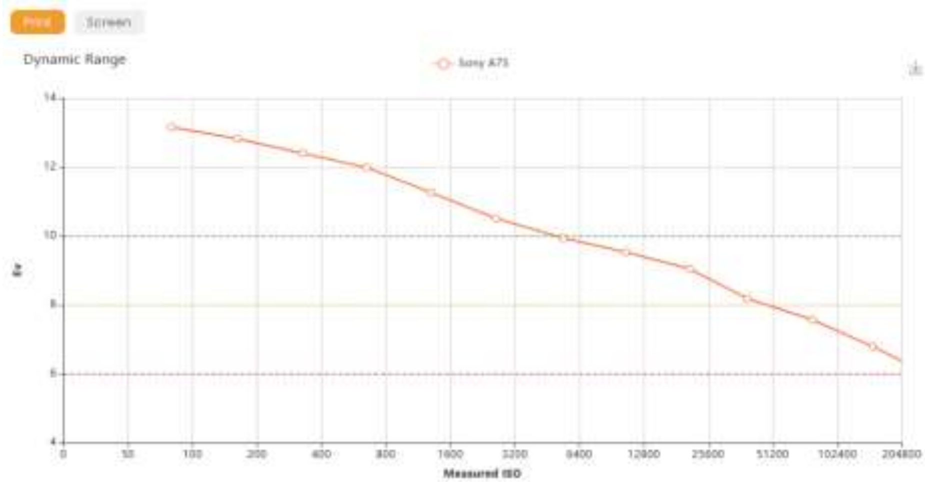

10

The DXO data is used by P2P.

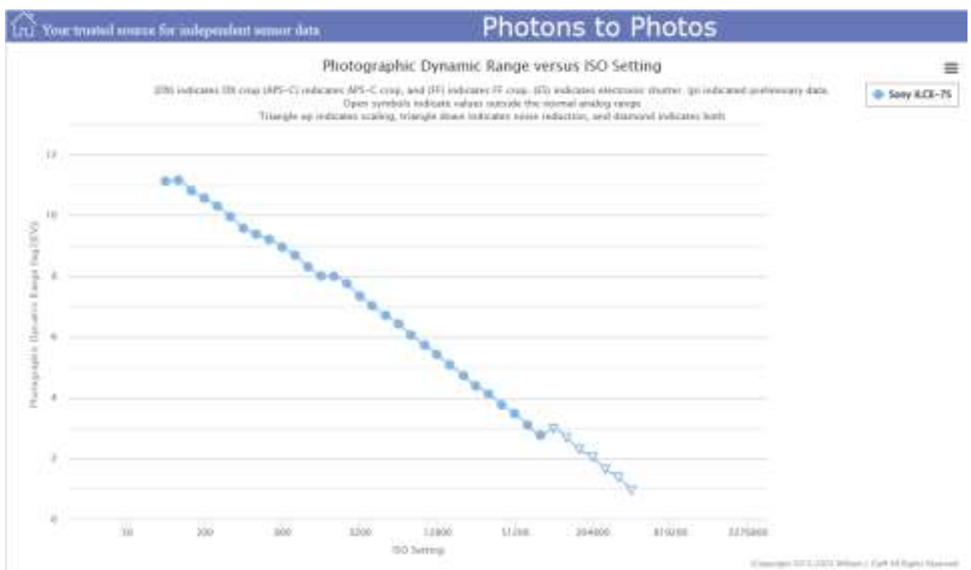

11

The large change in read noise does not result in much change in DR. But the A7S III shows it more clearly at both DXO and P2P.

Reason for Discontinuities in the DR Plots

The explanation for the discontinuities shown in DR for the Z7 and the A7 III is that both cameras use dual gain sensors with a different gain for low ISO settings than for high affecting the read noise. The lower ISO range does not need much more DR but the higher range can benefit from the lower read noise.

There are no discontinuities in DXO's SNR plots because the assumed noise is not base on read noise. They only appear in the original DR plots.

The read noise shows a distinct change where the gain is switched in the camera. The read noise is measured with the lens cap in place and the camera shielded from light.

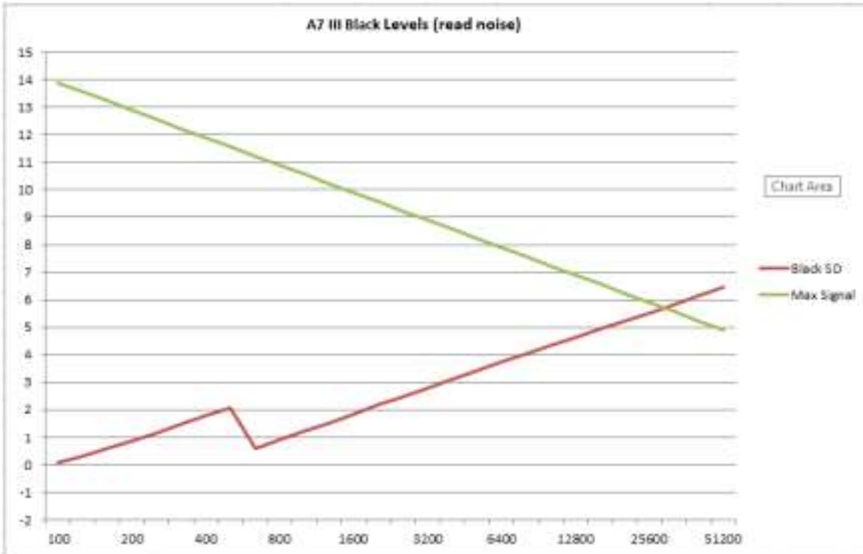

12

The change is also visible for the Z7.

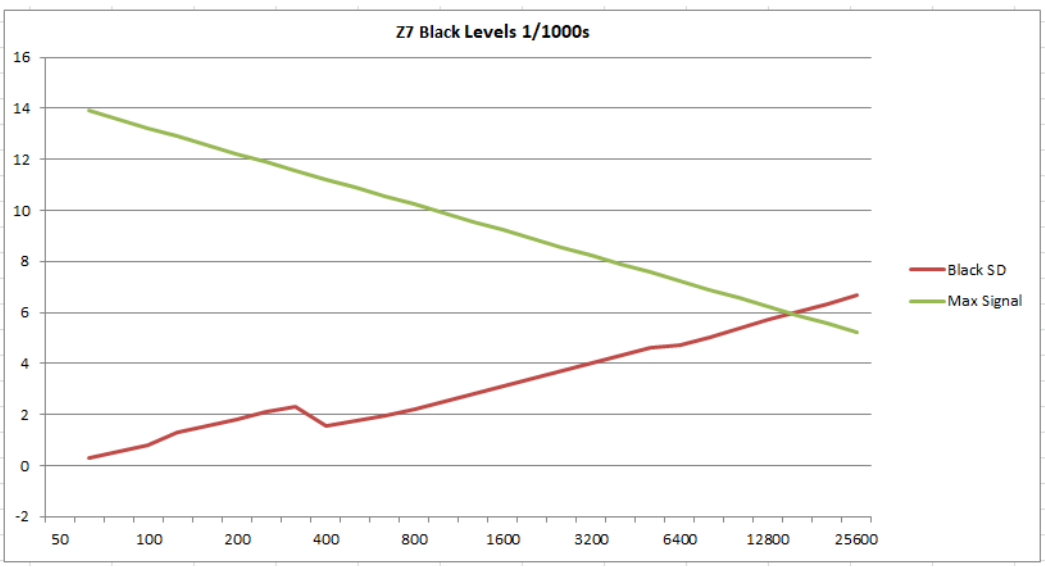

13

We can see why DXO and P2P DR plots display discontinuities.

Here is a plot for the A7 II which does not use dual gain.

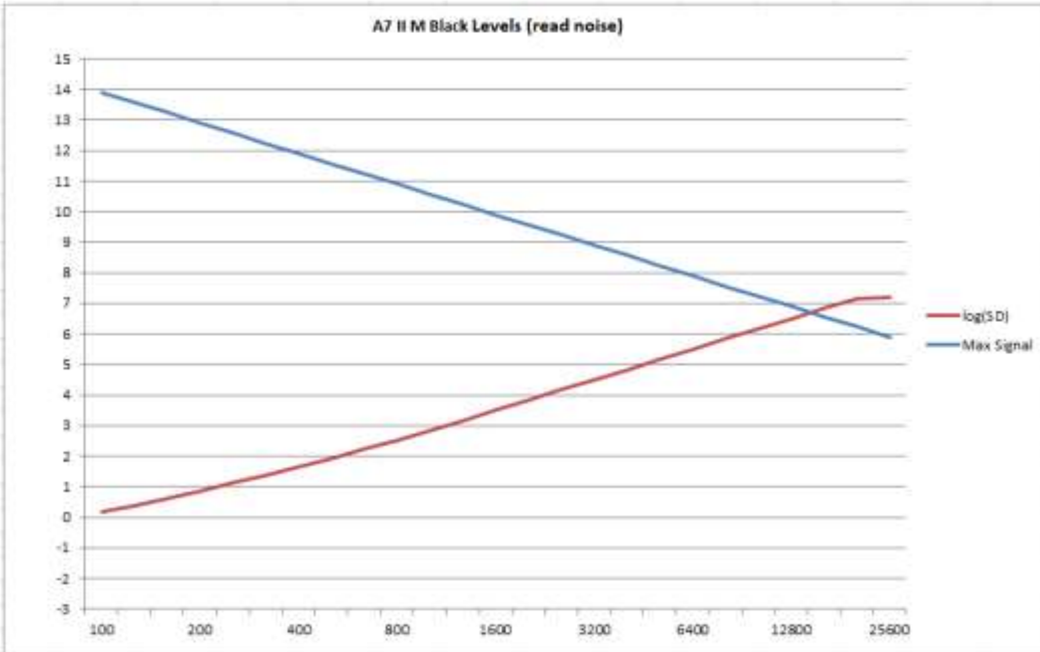

14

There is no discontinuity but the final SD value looks wrong so the analog range might end at ISO 20000.

By over and underexposing the target we can obtain these results to plot of the SD for the A7 III:

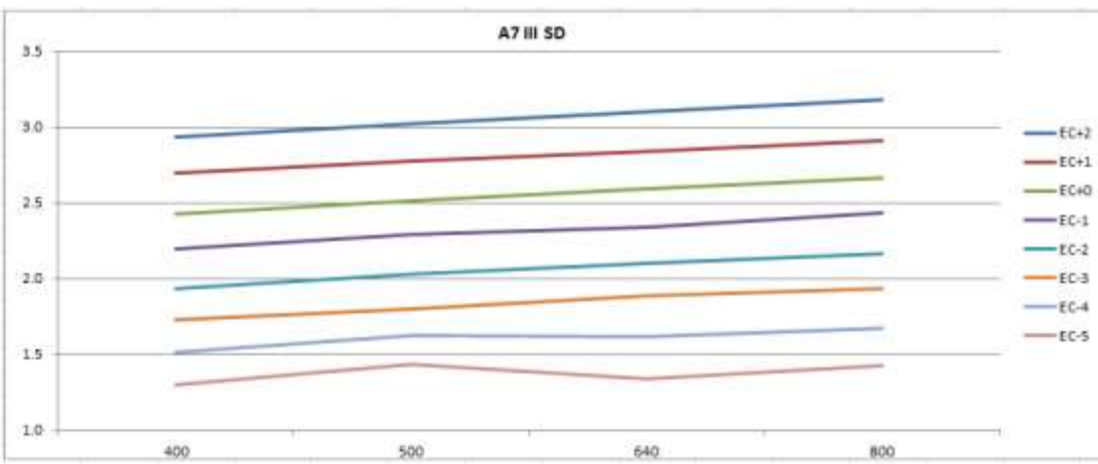

15

It's harder to show for the Z7:

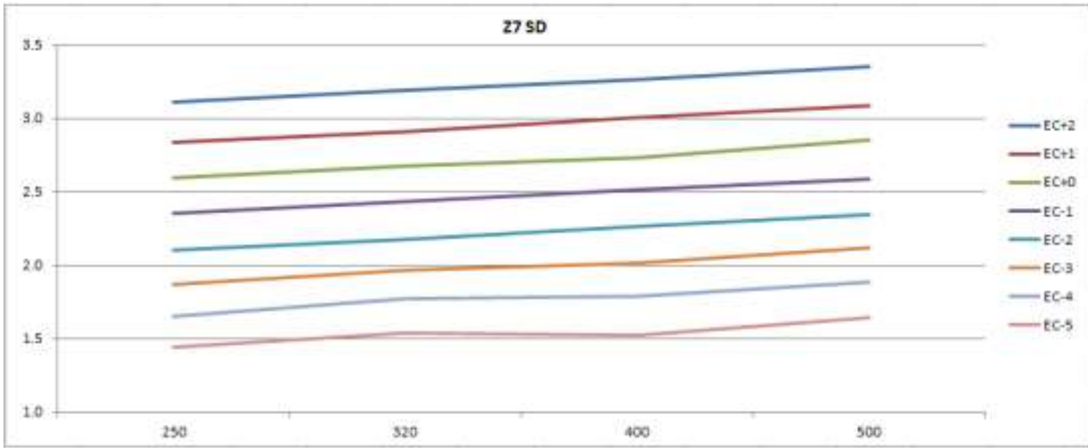

16

These are plots of the standard deviations. Discontinuities are not as pronounced as they are in the read noise plots because the total noise is higher than the read noise itself.

In both cameras the DR does not show a significant discontinuity until EC-5. That's five stops lower than middle gray. There are still about 3 stops above middle gray before the raw highlights blow out. Since each stop represents a range of +/- ½ of a stop this covers an overall range of nine stops.

We don't reach the black level itself until EC-9. At that level the image contains nothing of value.

Even at EC-5 the unadjusted image is almost pitch black. It is only during shadow recovery that the really low raw values at base ISO can be lightened enough for the noise to be seen. Using higher ISO settings makes that happen sooner. The problem with shadow recovery is that it lightens everything else below middle gray proportionately.

Remember the distinction that DXO makes between Print and Screen DR. If you are looking at an image on your computer screen at 100%, you are likely to see noise sooner for a high MP camera than for a lower MP camera. But when you print the same image from both cameras they will look more reasonable. That's because every time you increase the MP of a camera the noise shows up more easily. In other words, the DR is dependent on the degree of magnification. A 24MP camera will appear to have about a half stop advantage in DR over one with 45-50MP when you look at them at 100%. But that half stop goes away when both cameras capture the same image and it is printed.

Another consideration affecting DR is the upper limit. In order to not blow the highlights in an image and leave room for the black level adjustment. For the Z7 it's 1008 and at base ISO we need to use an upper limit of 13.9 rather than 14 for the log value. But at ISO 25600 we need to lower the upper limit to about 13 because the DR of at that level is much higher. The way to find this upper limit is to determine the highest exposure where RawDigger does not report any blown highlights.

We can plot the raw values at base ISO while varying the exposure, converting them to base 2 log values.

Here is a plot of raw values for the A7 III.

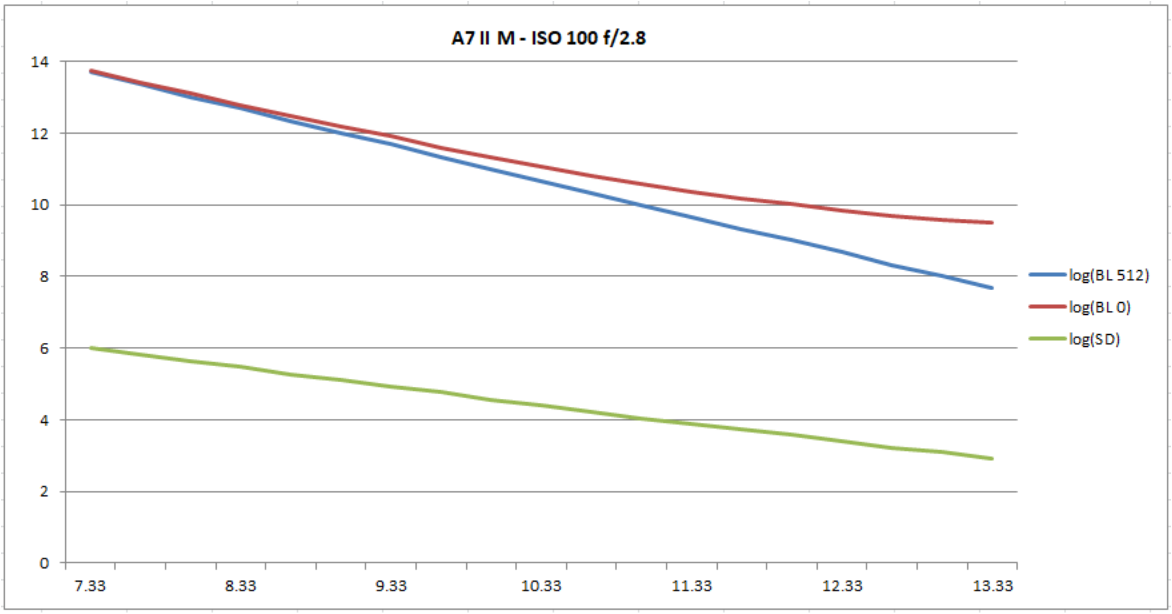

17

The log(SD) shows the standard deviation for the exposure values.

The log(BL 0) represents the sensor recorded levels including the black level over a six stop range.

BL is an arbitrary integer used to facilitate the recording of dark values with more precision than would be possible if they were stored directly. *It is not the same as the dark (read) noise which is measured with the lens covered.* A 3 or 4 digit integer is more precise than one using only 1 or two digits. It also allows recording values close to zero (like standard deviation that represents noise) without discarding values that would otherwise be too close to 0. For example, the two A7 bodies use a BL of 512 (2^9) and the Z7 BL is 1008 (almost twice as high). Neither BL value has anything to do with DR. Some older cameras (D610, Df) use a BL of 0. The dark parts of their analog ranges are less precisely noted in the raw file. This might look slightly noisier. The use of the BL value can make it easier to recover shadow information.

When you subtract the black level (512) within the camera's analog ISO range you end up with a straight line, proportional to the light value (LV). A wider range of exposures does not change the fact that the blue and green lines essentially remain straight although very low exposures they might deviate a little.

Sony (and Nikon, etc.) deliberately keep the signal proportional to the exposure to conform to the exposure triangle. They also concentrate on keeping the noise down so that it does not become visible until the exposure is drastically reduced.

As stated in the opening paragraph, DR for the sensor is defined at base ISO. All other values within the normal analog range of the sensor are something other than dynamic range.

How to Measure SNR

The DXO data may have been compiled by putting the camera on aperture priority.

A series of images can be captured with EC set to 0. The target should have a smooth tone with a controlled brightness. A white display on a calibrated monitor works well. To avoid any unwanted patterns, the camera should be manually focused at infinity and the camera held close to the screen. Images should be recorded using spot metering. Measurements should be taken for a small selection at the center of the image using the green channel since they represent half of the Bayer color array (a little more for X-trans).

As the ISO is increased the shutter speed gets shorter to keep middle gray at the same raw and JPEG level. It is primarily *the reduction in exposure* that drops the SNR making the noise more visible.

We are familiar with the progression of ISO settings in whole stop increments – 100, 200, 400, 800 … But what about 1/3 stop increments? They are ISO 100, 126, 159, 200, 252, 317, 400 … The intermediate steps are not whole numbers since they are based on 100x(2^(n-1)/3) where **n** can be any integer, even 0 or negative.

We may also be familiar with the progression of shutter speeds in seconds – 1, ½, ¼, 1/8, 1/15, 1/30 and 1/60 … The last three are actually 1/16, 1/32, 1/64 … The intermediate steps are also based on powers of 2 in 1/3 step increments.

Here is a typical sequence for ISO and shutter speed:

#	ISO	1/ss
1	100	12.70
2	126	16.00
3	159	20.16
4	200	25.40
5	252	32.00
6	317	40.32
7	400	50.80
8	504	64.00
9	635	80.63
10	800	101.59
11	1008	128.00
12	1270	etc.

These values were used to test the A7 III and Z7 using a standard target spot metered to render a middle gray value at the center of the image.

Measurements were made using a 150x100 pixel selection from the center of the target image viewed in RawDigger:

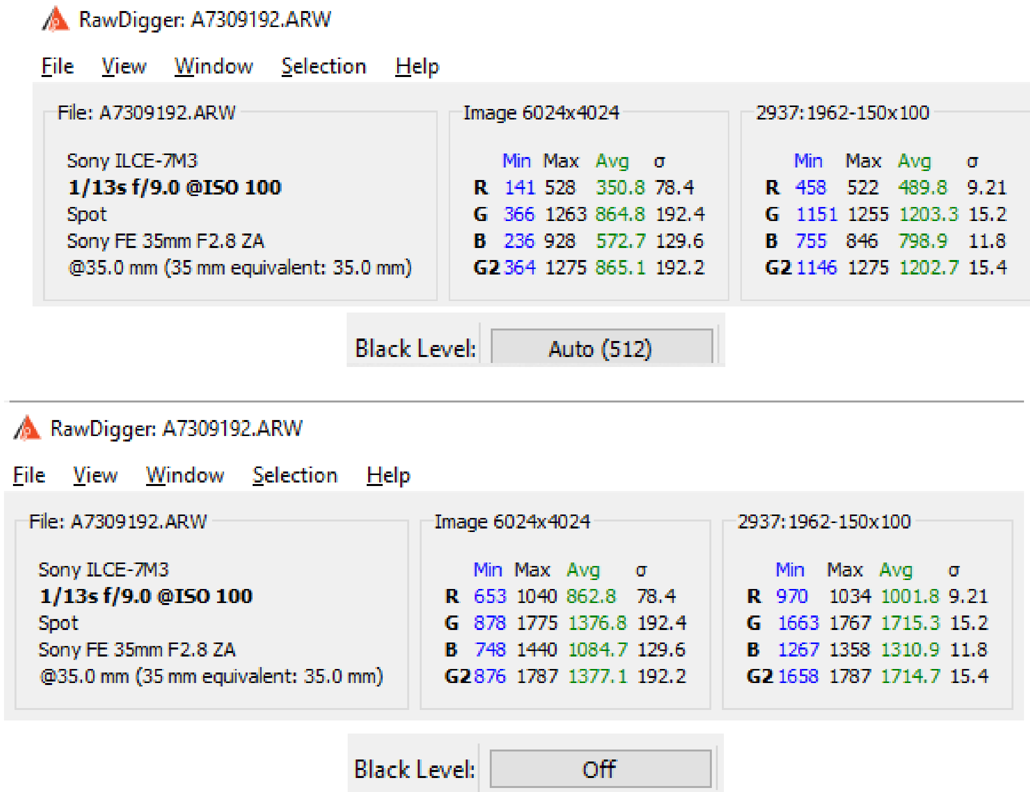

18

The average raw values and SD (σ) were recorded for the two green channels. The test was run at base ISO by varying the shutter speed.

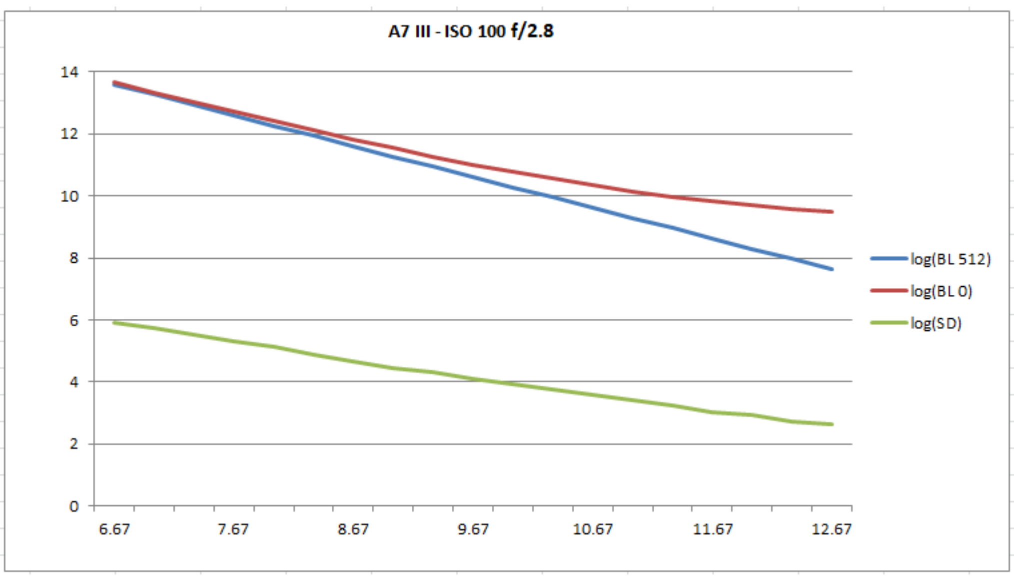

19

The blue line is the plot of base 2 logarithm for middle gray after subtracting the black level of 512.

DXO shows the SNR in decibels (dB). Converted to base 2 logarithms:

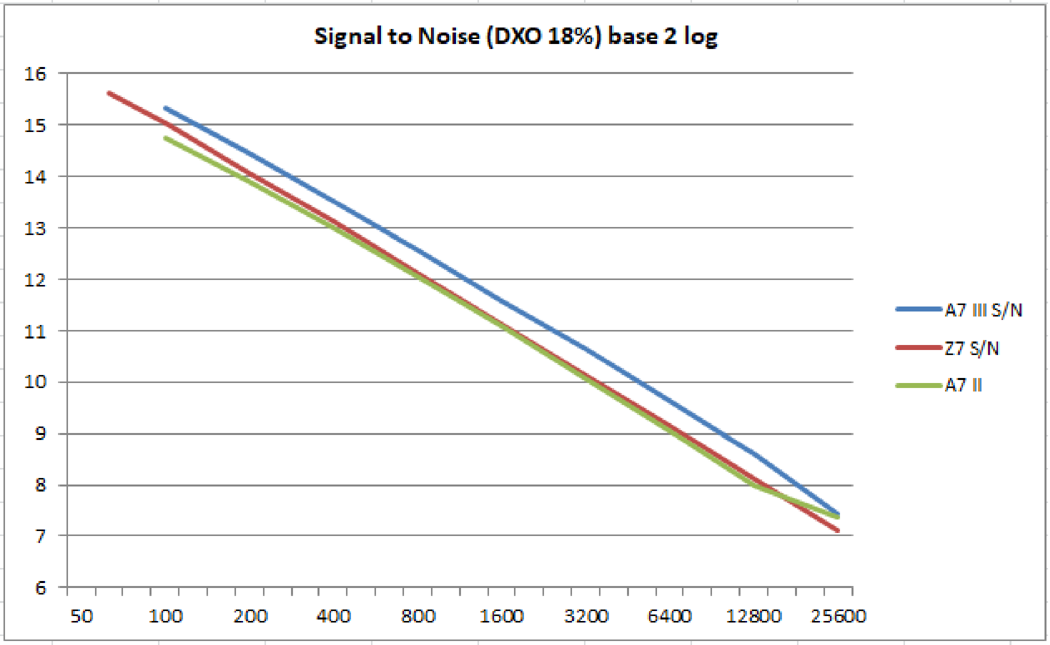

20

The SNR that DXO shows is based on the ratio of the signal to the dynamic range (σ) from a smooth portion of an image. It includes both the black noise and all of the other noise.

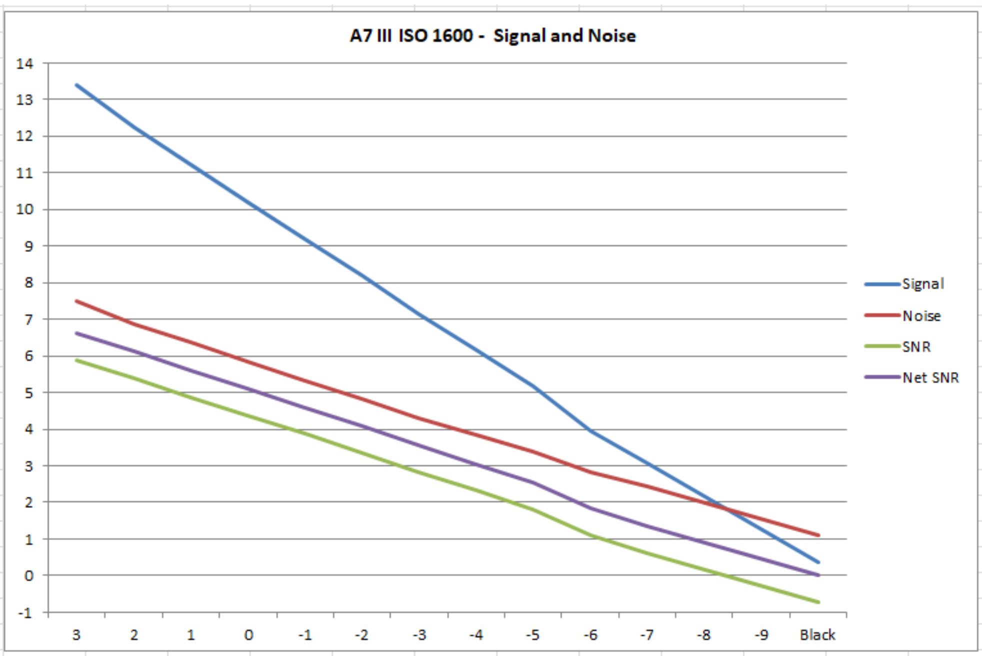

21

The vertical scale here is the base 2 log of the values reported by RawDigger for a small selection at the center of a test image. The horizontal scale represent exposure compensation (EC) where 3 is three stops above middle gray, 0 is middle gray, etc. Net SNR represents the noise level for the image if we could remove the black noise.

This is a much simpler way to show how exposure affects SNR. In a single image, only one ISO setting is used. This plot shows that noise is related to exposure, not ISO.

DXO publishes the SNR values using a decibel (dB) vertical scale (6.0206 times the base 2 log scale).

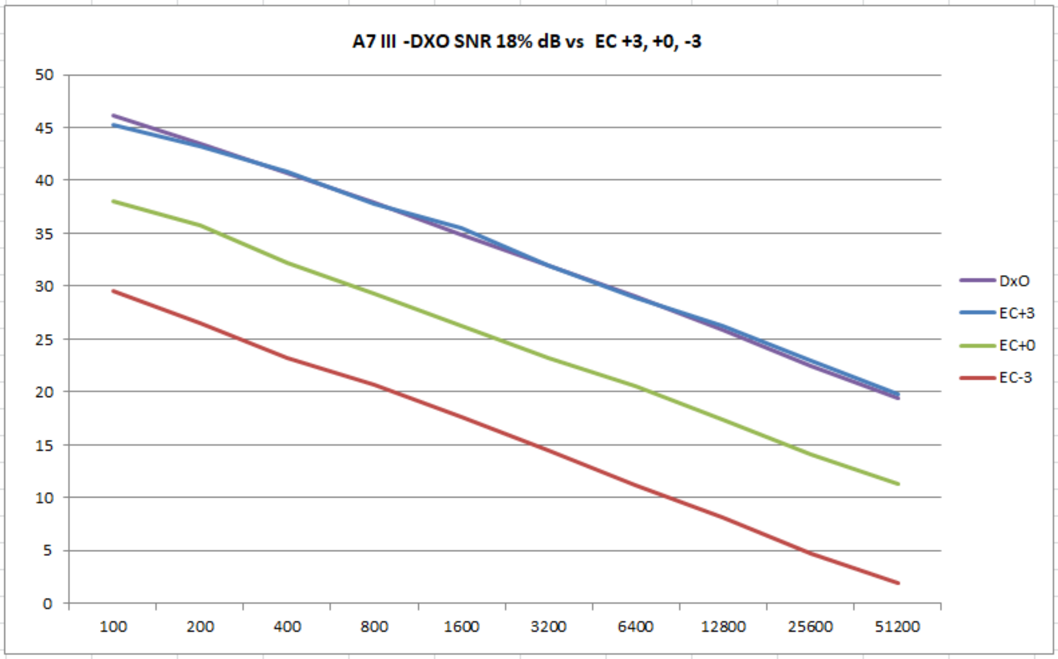

22

This plot is overlays the DXO Print version of SNR over the plots for EC+3, EC+0 and EC-3. Note that the DXO line actually aligns with the EC+3 plot which is actually 3 stops brighter than middle gray.

Conclusion

Both DXO and P2P do a good job comparing two or more cameras.

DXO uses a simpler definition of noise to compute SNR than is does for DR.

Both DXO and P2P base DR on the difference between the maximum recorded brightness and the black noise. P2P is more conservative than DXO in that it does not use the full difference by reducing it by about 2.7 units but this does not work as well for cameras with dual gain sensors.

Neither of them addresses the aesthetic value of noise levels in determining DR.