70th

birthday series

Exposure Meter Calibration

Douglas A. Kerr, P.E.

Issue 1
May 6, 2006

ABSTRACT

We often hear that “a standard photographic exposure meter [or automatic exposure system] is calibrated to a reflectance of 18% [or maybe 13%, or some other nearby number].” Sometimes the word “gray” appears in the description. What does this mean, and why the large variation in the numerical value? In this article, we look at several relevant ISO standards and see how a “standard exposure meter calibration” is implied by their interaction. In an appendix, we look into the calibration situation for Canon digital SLR cameras, as inferred from a test recommended by the manufacturer. We also discuss the related issue of incident light metering, including by way of the use of a “gray card”.

SUMMARY

We often hear that “a standard photographic exposure meter [or automatic exposure system] is calibrated to [a reflectance of] 18% [or maybe 13%, or some other nearby number].” What does this mean? This expression of course does not really describe any “calibration” situation. Rather, it is just a shorthand for an underlying concept.

If we have a digital camera whose automatic exposure system is calibrated per the applicable international standard (ISO 2721), and whose rating of the “ISO sensitivity” of its sensor system is done per the applicable international standard (ISO 12232), and we use it to take a picture of a uniform-luminance test target, the exposure on the sensor should be 12.8% of the “saturation exposure” (that is, the highest exposure for which there is a unique digital result).

This has the effect that, for a real scene whose actual average reflectance is not less than 12.8%, the exposure given an object in the scene with a reflectance of 100% (which we might think of as the “lightest possible natural object”) is not over the saturation exposure.

This also means that for a real scene whose actual average reflectance is 18% (often thought to be “typical”), the exposure given an object in the scene with a reflectance of 100% will be 1/2 stop below the saturation exposure—which we can think of as giving a “1/2 stop exposure cushion” for such a scene.

These two outlooks on the same calibration situation are the source of the two numbers we often see in “the statement”.

Many modern digital cameras offer an exposure metering system that is more sophisticated than the one contemplated by the above exposure control doctrine, which only measures the average luminance of the scene. These intelligent exposure metering systems can do a better job of predicting the actual maximum luminance in the scene. Thus, it can be argued that the exposure control doctrine described above is too conservative: it gives an exposure that is unnecessarily low in most cases. Therefore, perhaps, in such a camera, we would wish to use an exposure control doctrine that results in a greater “standard exposure” than 12.8% of the saturation exposure.

That could be implemented by doing one or both of two things. We could use a calibration for the exposure metering system that gives a higher exposure than the one called for by ISO 2721. Or we could rate the “ISO sensitivity” of the camera’s sensor at a lower value than called for under ISO 12232.

Doing the former would be problematical, since if a photographer used a properly-calibrated external exposure meter it would lead to different exposure settings than would the camera’s own exposure control system, a likely cause of consternation.

Thus, the manufacturer may well elect to rate the ISO sensitivity of the sensor at a lower value than would be called for under ISO 12232.

For Canon digital SLR cameras, we get some insight into what has been done from an exposure control test recommended by Canon. It implies that the “standard exposure” planned for these cameras is about 17.3% of saturation exposure.

We believe that for these cameras the exposure meter calibration itself is closely in line with ISO 2721. Thus, we must conclude that Canon has rated the “ISO sensitivity” of their sensor systems at about 74% of what would be determined in accordance with ISO 12232—that is, a sensitivity that Canon rates “ISO 100” would probably be rated as ISO 135 if actually determined in accordance with ISO 12232.

INTRODUCTION

We often hear that "a standard photographic exposure meter (or automatic exposure system) is calibrated to [a reflectance of] 18% (or maybe 13%, or some other nearby number)." Sometimes the word "gray" appears in the description. What does this mean, and why the large variation in the numerical value?

Of course that phrase doesn't actually state any calibration relationship. Rather, it is shorthand for a concept that is one part of the complicated story of exposure metering and the calibration of exposure meters. And the different numbers often appearing in the phrase may not relate to two different calibration situations, but perhaps to two ways of looking at the same one. In this article, we will explore that story from end to end.

Although the concepts apply equally to film and digital cameras, certain details differ between the two situations. In this article, we will specifically investigate the concepts in the context of digital cameras.

Note also that in the body of this article we will be speaking of basic "reflected light" metering systems: those that operate by measuring the luminance of the scene and, in particular, the average luminance. An important alternative is "incident light" metering, which operates by measuring the illuminance of the light falling on the scene. This will be discussed in detail in Appendix C.

THE TERM "EXPOSURE"

Before I can really get started, I need to clarify a matter of terminology.

The term "exposure" is legitimately used in photographic science to refer to two different quantities. To avoid any misunderstanding, in this article I will use two coined words, "exposure1" and "exposure2" for these two distinct quantities, as follows:

Exposure1: The joint effect of an effective relative aperture (effective f/number) and an exposure time (shutter speed). There is no standard scientific symbol for this quantity; I often use the symbol *X*. It is this quantity that in its APEX (logarithmic) form is designated "Ev"[1]

---

[1] Do not be confused here by the common, but incorrect, usage of "Ev" to denote scene luminance.

Exposure2: The phenomenon to which film or a digital sensor responds, the product of the illuminance on the film or sensor and the exposure time. The standard scientific symbol for this quantity is *H*. [2]

A third quantity I will frequently mention is what I call *exposure result*. This is what is caused by an exposure2—what is "left for us" in the film or digital image. In the case of black and white negative film, the specific physical property is that of *density*. In the case of a digital imaging chain, the property is the digital representation of the color of a particular pixel (*e.g.,* in an RGB color space, the triplet of R, G, and B values).

IMPLICATIONS OF THE ISO STANDARDS

Automatic Exposure Control Systems

International standard ISO 2721-1982 defines the "calibration" of camera automatic exposure control systems. We will start there, rather than with the ISO standard for "handheld" exposure meters, because the story is tidier. The equivalent story for handheld exposure meters is given in Appendix A.

The standard essentially prescribes as follows:

> If the camera, with automatic exposure control system in effect, regards a scene of uniform luminance, it should set an exposure1 such that the resulting exposure2 created on the sensor will be:
>
> $$H_u = \frac{10}{S_i} \quad \text{lux•seconds} \qquad (1)$$
>
> where H_u is the exposure2 on the sensor (*u* being evocative of "uniform scene") and S_i is the *exposure index*: what we have told the exposure control system is the ISO sensitivity of the sensor system.[3] H_u is sometimes called the "standard exposure".

---

[2] In some cases the symbol *E* is used, but this is not advisable since this is the standard scientific symbol for illuminance (think "e-luminance"). The present symbol "*H*" was probably inspired by the fact that the classical curve of film response, plotting *D* (density) against *E* (exposure2), is often called the "H&D curve" in honor of Hurter and Driffield, two early researchers in this field. Many people thought that since "D" was the symbol for density, "H" must be the symbol for exposure2. And so it was indeed eventually adopted for that!

[3] In digital cameras, if we set the camera to use a certain ISO sensitivity, the exposure control system is automatically advised of that value. But that is no guarantee that this is actually the "proper" rating of the sensitivity as defined by the ISO standard for sensitivity. Thus my cautious language!

ISO sensitivity

ISO 12232-1998 defines the determination of the "ISO speed" (that's actually the official term; I prefer to call it the "ISO sensitivity" since that term seems more apt) for a digital camera sensor system.

That definition, in effect, states:

> The ISO sensitivity of a digital camera sensor system is defined by:
>
> $$S = \frac{78}{H_{sat}} \qquad (2)$$
>
> where S is the ISO sensitivity and H_{sat} is the exposure2 (in lux•seconds) at which the camera delivers the greatest possible digital output (that is, the *saturation value* of exposure2). [4]

Note that the ISO sensitivity of different classes of film (prescribed by other standards) is determined based on a different criterion; one cannot infer any "saturation" value of H from those ratings.

Putting that all together

If we combine those two definitions, we get the following for a digital camera whose automatic exposure system is calibrated in accordance with ISO 2721, when the automatic exposure system has been correctly advised of the ISO sensitivity of the sensor system as defined according to the saturation basis of ISO 12232:

> If the camera, with automatic exposure control system in effect, regards a scene of uniform luminance, it should set an exposure1 such that the resulting exposure2 created on the sensor, H_u, will be:
>
> $$H_u = 0.128 H_{sat} \qquad (3)$$
>
> where again H_{sat} is the saturation value of exposure2 for the sensor system.

It is the achievement of this relationship that is really the underlying objective of the ISO 2721 calibration standard, although it is never

[4] The standard actually provides two different bases for rating the ISO sensitivity of a digital sensor system. One is based on the "saturation exposure2" and the other on noise performance. The latter is not applicable for cameras in which the image output is in JPEG form. We will be considering here only the "saturation" basis.

mentioned directly in the standard, but must be discerned by considering the interaction of that standard with ISO 12232.

Note that we can quantify this relationship by way of the ratio H_u/H_{sat}:

$$\frac{H_u}{H_{sat}} = 0.0128 \qquad (4)$$

and that ratio will be a convenient way to describe other results of the combination of exposure meter calibration and ISO sensitivity rating we may encounter with regard to digital cameras. We could imagine this particular relationship being called, in shorthand, "calibration to 12.8%" (although we don't encourage that usage).

The implications

The above calibration situation can be interpreted this way:

> If the camera regards a scene with an average reflectance of 12.8%, then the exposure2 created by an object in the scene with a reflectance of 100% (which we will assume is the "lightest" natural object we might encounter) will cause an exposure2 on the sensor that is just at the saturation level.

It is this outlook that is responsible for the statement, "a standard photographic exposure meter (or automatic exposure system) is calibrated to (a reflectance of) 12.8%". Often, the number cited is 12%, 12.5%, or 13%, probably just convenient approximations of the number we have derived here.[5]

Or the same calibration situation can be interpreted this way:

> If the camera regards a scene with an average reflectance of 18%, then the exposure2 created by an object in the scene with a reflectance of 100% (which we will assume is the "lightest" natural object we might encounter) will cause an exposure2 on the sensor that is 1/2 stop below the saturation exposure2.

This outlook reflects an underlying philosophical basis for the establishment of the "standard" calibration we are discussing here. That basis unfolds in these steps:

---

[5] In fact, when we look at the ISO standard for hand-held exposure meters, we find that there is a substantial range of "calibrations" allowed, which can lead to a considerable range of the ratio H_u/H_{sat}.

- We will assume that the average reflectance of the scene is 18% (widely thought to be "representative").
- We will assume that the "lightest" reflective object that might appear in a scene has a reflectance of 100%.
- We wish to be certain that the exposure2 for such an object does not exceed the saturation exposure2 (H_{sat}).
- Thus we could choose a calibration such that, when the camera regards a scene with an average reflectance of 18%, an exposure1 would be indicated such that an object in the scene with a reflectance of 100% would receive an exposure2 of H_{sat}.
- However, If the actual scene being photographed has an average reflectance less than the assumed 18%, then the exposure1 indicated by the meter will result in a "100% reflectance" object being given an exposure2 above saturation (with the consequence that highlight detail will be "clipped").
- To protect against that, we choose a calibration that will indicate an exposure1 of 1/2 stop less than with the calibration postulated just above, so as to give a 1/2 stop "cushion" against the possibility of a 100% reflectance object getting an exposure2 above saturation. (This calibration is sometimes said to give 1/2 stop of "headroom".)

It is this outlook that is responsible for the statement, "a standard photographic exposure meter (or automatic exposure system) is calibrated to [a reflectance of] 18%."

Note that the two outlooks are just that: two different ways of looking at exactly the same calibration situation.

We can summarize that entire rationale as follows: we are betting that in only the rarest of cases will the brightest object in the scene have a luminance more than 7.8 times the average luminance (1.0/0.128).

But the landmark study of the distribution of luminance in photographed exterior scenes (Jones and Condit, 1941), upon which much exposure metering planning has been based, showed that in most cases the maximum scene luminance was not over 3.6 times the average luminance.

If we were willing to "bet on" that situation, then we could have established an exposure metering calibration doctrine that, in connection with the definition of ISO sensitivity, would lead to an

H_u/H_{sat} of about 0.278. If we still built in a 1/2 stop cushion, that would suggest an H_u/H_{sat} of about 0.194—a “19.4%” calibration.

Can we test how our camera deals with this?

We cannot verify that an exposure meter, or a camera automatic exposure control system, conforms to the calibration prescribed by ISO 2721 (as shown in equation 1), or that the “rating” of the ISO sensitivity for the camera’s sensor system conforms to the definition given by ISO 12232 (as shown in equation 2), without having a test light source whose luminance is accurately known (or an accurate photometer).

However, without such a source or photometer, we can still readily ascertain fairly well if the joint result of the calibration and the ISO sensitivity rating conforms to the situation described by equation 3.

In appendix B, we will follow the principles of such a test as applied to Canon digital SLR cameras.

IN THE DELIVERED IMAGE

We have so far discussed life at the face of the sensor, where the phenomenon of interest is H, the value of exposure2.

Unless we are in a camera testing laboratory, we actually have no direct visibility of the exposure2. All we can see is its effect on the exposure result: the relative luminance value represented by the digital codes for various pixels in the digital image. Yet, as we will see in Appendix B, we often use that as an indicator of H_u/H_{sat}. How can we make that connection?

If we know the definition of the “color space” in which the digital image is represented, we can determine the relative luminance represented by any given digital code. For example, if the image is in the sRGB color space, then an RGB code value of 128,128,128 represents a luminance that is 0.2159 of the luminance represented by an RGB code value of 255,255,255 (the maximum code value for that color space). This ratio, the “relative luminance”, can be expressed as Y/Y_{max}. (where Y is the symbol for luminance, on an arbitrary scale, in this case).

In general, we can assume that the maximum possible code value is associated with an exposure2 of H_{sat}. Then does the ratio Y/Y_{max} directly equate to H/H_{sat}?

It does if the camera’s digital image strictly follows the color space it is said to follow. But in general that is really not exactly true. Usually there is "in camera image processing", which does such things as adjust the contrast of the image. In that case, then strictly, the output image does not really follow the definition of the "announced" color space (perhaps the "sRGB" color space).[6]

More to the point, it means that we cannot, by examination of the digital image from a camera, make a precise determination of the ratio H_u/H_{sat} produced by the automatic exposure system. Still, the assumption that they are approximately equivalent is reasonable in interpreting the behavior of a camera.

WHAT ABOUT "GRAY"

At the outset, we quoted the often-heard statement about exposure meters and mentioned that sometimes the statement includes the word "gray", perhaps as follows:

"A standard photographic exposure meter is calibrated to 12.8% gray."

Remember, this statement does not, in almost any of its forms, actually describe a calibration situation—it is just a widely-used "shorthand" for the situation we discuss at length in this article.

Note that one implication of the number in the statement (in the form I just quoted) is that it is meant to be (expressed as a percentage) the ratio H_u/H_{sat}, which of course we can usually only see in terms of the corresponding relative luminance of the exposure result, Y_u/Y_{max}. Because tests of exposure meter calibration and its implications are often conducted with a "chromaticity-neutral" ("gray") test scene, it is common to describe a Y_u/Y_{max} of, say, 0.0128 as a "12.8% gray".[7] Thus the appearance of that word in the famous statement.

"Mid-gray"

Often, it is said that a "standard" camera, regarding a scene of uniform luminance, will record that scene with a "mid gray" exposure result. What might that mean?

---

[6] Many people will say that such an image is in fact recorded in the sRGB color space but has just had its luminance values tampered with before the image was encoded.

[7] Of course, in the normal discussion of a "grayscale image", that relative luminance is described as "87.3% gray", so we have to be very cautious here.

Some relate it to the CIE L*a*b color space, in which an L* value of 50 (on a scale of 0-100) represents a relative luminance of 0.184 (18.4%).

ALTERNATIVE METERING SCHEMES

Alternative metering patterns

So far, we have discussed exposure metering (or automatic exposure control) that operates by measurement of the average luminance of the scene (recognizing that the field of view observed by the metering instrument may not exactly match the field of view used for the image).

Many cameras today offer some alternative metering "patterns". Here are some of them, with the one we have been assuming listed first for completeness of the comparison:

- **Scene-wide average metering**, in which the average scene luminance is observed over a field of view which, for "through the lens" metering, spans all or almost all of the field of view for the image itself.

- **Center-weighted average** metering, in which the scene luminance is observed over a field of view spanning almost the entire field of view used for the image, and the results averaged, but with the readings at the center given a larger "weight" in determining the overall average, the weight smoothly declining as we move from the center of the frame.

- **Spot** metering, in which the meter observes only a central circular part of the scene, the luminance being averaged over that area.[8]

In many cameras, the scaling of that output for the different patterns is such that, when observing a scene of uniform luminance, the exposure1 set by the automatic exposure system will be the same for any of them. This is not necessarily ideal from an exposure strategy standpoint, but is done to provide a comforting consistency. (Of course, for scenes not having a uniform luminance, the different patterns will generally lead to different values of exposure1; if they didn't, then there would clearly be no need to have different patterns available.)

---

[8] Most Canon digital SLR cameras have a spot metering mode called "Partial", in which the field of view of the meter is about 8% of the field of view for the image. Some Canon models also have a spot metering mode called "Spot", whose field of view is about 3% of the field of view for the image.

Why one or the other of these schemes might produce the most desirable overall exposure result for a particular scene is beyond the scope of this article.

Intelligent metering

Many cameras today offer an "intelligent" metering system.[9] Such a system typically separately measures the luminance of the scene in each of a number of regions across the field of view. This suite of readings is then evaluated by an intelligent algorithm, whose objective is to determine, or at least predict, the actual distribution of luminance across the whole scene.

By doing so, we move away from reliance on an assumed average scene reflectance and an assumed maximum object reflectance. Accordingly, an optimal exposure1 can be selected under a doctrine that does not require building in any "cushion" to deal with "surprises" in the matter of actual average and maximum scene reflectance. Thus, in many cameras, the exposure2 produced by an intelligent metering system on a uniform luminance scene will be higher than the standard exposure2 discussed earlier in this article.

Further, extending the concept of "comforting consistency" mentioned earlier, in cameras having an intelligent metering system option the calibrations of the basic metering patterns are often "pushed up" to produce a result, for a uniform luminance scene, consistent with that given by the intelligent metering system.

Increasing the standard exposure above that implied by the two ISO standards can be done through either of the following or both:

a. The calibration of the metering system itself can be increased above that implied by ISO 2721.

b. The sensitivity rating of the sensor system can be decreased from that which would be assigned under ISO 12232.

If the manufacture decides to use method (a), a consequence is that a photographer using an external exposure meter will be advised to use a different exposure1 for a certain scene than would be chosen by the camera's automatic exposure system, a likely cause of consternation. Thus, more typically, the camera manufacturer will apply method (b). We will see an example of this in Appendix B.

---

[9] Examples are the "Matrix" metering system offered by Nikon and the "Evaluative" metering system offered by Canon.

THE EXPOSURE INDEX

The concept

Earlier, I gave glancing mention to the term "exposure index". Here is some further information on that concept.

I will illustrate the concept with a parable from the world of film photography. Suppose I am using film rated at ISO 400. I will be shooting scenes in a low-light situation, and the required shutter speed is worrisomely-slow from a standpoint of concern with motion blur.

However, I set the ISO dial on my exposure meter (or the camera's automatic exposure control system) to "ISO 800". Thus the meter will indicate (or the automatic exposure system will enact) an exposure[1] that is one stop less than if I had set the actual ISO sensitivity of the film. This can of course involve a faster shutter speed.

Then I take the film to a custom laboratory for processing, asking that they develop it with a "one-stop push".

That means that they use a special developing technique that results in the film sensitivity being essentially that which would be exhibited by "ISO 800" film with normal development (at least insofar as the relationship between exposure[2] and density in the "lower" portions of the response curve).

The result (hopefully) will be a set of usable negatives from which we can prepare nice prints. There is a "price": the effect of grain in the image may be exacerbated, and the dynamic range may have been reduced. But we will assume that these penalties are more than overcome in the particular case by the ability to use a faster shutter speed.

How do we describe this scenario

Usually, we describe this approach by saying that "I shot these pictures at ISO 800". A more exact technical expression would be, "I shot these picture at an exposure index of 800."

Exposure index, in this context, is what we "tell" the exposure meter or automatic exposure system is the ISO sensitivity of the film, even if that is not so.

With regard to the processing, some would say that "the push processing gave the film an ISO sensitivity of ISO 800". But that's not correct. The ISO sensitivity of a film type is measured in a certain way, with the film processed in the standard way recommended by the manufacturer (or, for some basic film types, using a standard process defined in the ISO standard). That rating remains applicable regardless of how we might actually process the film (or even if we don't process it at all!).

Rather, we can more accurately say that the film has been processed "to an exposure index of 800."

Exposure index, in this context, describes the sensitivity of a film (in accordance with some criterion), when processed in some particular way, on the same scale as is used for true ISO sensitivity.

Historical background

Before there were industry standards for the "speed" of photographic films, exposure meters had a dial on which the photographer would set the "exposure index" for the film being used, on an arbitrary scale. The meter manufacturers would then provide, with the meter, a pamphlet giving what the manufacturer felt was an appropriate exposure index for the various kinds of commonly available film.

Each manufacturer had their own exposure index scale (and basis for rating film speeds under it).

When a standard for the speed of film was published by the American Standards Association (the "ASA" speed rating), this was in fact a standardized exposure index with a specific definition, to be measured under carefully controlled conditions. The scale was chosen so as to fall about halfway between the scales used for the exposure index scales used on two popular brands of exposure meter, the Weston and the General Electric. The ASA rating was about 125% of the Weston rating for the same film type.

Today's ISO scale is essentially comparable to the ASA scale, although the current ISO standards have different details as to how the rating is to be measured.

#

APPENDIX A

Free-standing exposure meters

The calibration of free-standing exposure meters is prescribed by ISO 2720-1974. Here, the situation is a bit more complicated than for integrated automatic exposure control systems.

For one thing, the indication of the meter must be passed in a standard form (f/number and shutter speed, collectively a value of exposure1) to the camera. For another, the exposure meter doesn't actually know how the effective f/number of the particular lens compares to the actual f/number. The effective f/number is greater than the actual f/number because of the effect of:

a. bellows factor (if the camera is not focused at infinity)[10]

b. transmission of the lens not being 100%

The meter cannot really take into account factor (a); for a subject at a fairly close distance, where this matter is of significance, the photographer may take that into account in choosing the aperture to be used based on the one indicated by the meter. But perhaps the calibration of the meter will make some arbitrary allowance for factor (b).

Then, on top of all that, the standard allows a fairly wide discretionary range for the manufacturer in the matter of calibration.

Because of the need for the meter's output to be expressed in terms of exposure1, the calibration is defined in terms of the reflected light exposure metering constant, thus:

$$\frac{t}{n^2} = \frac{K}{L_a S_i} \tag{5}$$

where t is the exposure time (shutter speed) in seconds, n is the relative aperture (as an f/number), L_a is the measured average luminance of the scene (in candelas/m^2), and S_i is the exposure index: what we tell the exposure meter is the ISO sensitivity of the film or digital sensor system. K is the *reflected light exposure metering*

---

[10] Although we often treat the f/number of the lens as being the indicator of its impact on exposure, that is only an approximation, strictly true only for focus at infinity. Application of the *bellows factor* gives us an *effective f/number* which can be used in the same way as f/number but will give the proper indication of exposure impact of the lens for the actual focus distance in effect.

constant, and it is through choice of this that the meter manufacturer sets a "calibration" for the meter (including the accommodation made for the assumed transmission of the camera lens).

Note that the quantity t/n^2 is the numerical value of exposure1. It can be expressed in APEX (logarithmic) terms as Ev, which is defined thus:

$$\text{Ev} = -\log_2 \frac{t}{n^2} \tag{6}$$

or, as more conventionally expressed:

$$\text{Ev} = \log_2 \frac{n^2}{t} \tag{7}$$

ISO 2720 prescribes that K may have a value between 10.6 and 13.4.[11]

By combining the calibration equation with the fundamental photometric relationship between scene luminance and image illuminance, we find that if the camera regards a scene of uniform luminance, which is also observed by the exposure meter, and the exposure1 is set as indicated by the meter, the resulting exposure2, H_u, created on the sensor will be given by:

$$H_u = \frac{\pi}{4} z \frac{K}{S_i} \quad \text{lux•seconds} \tag{8}$$

where again S_i is the exposure index (what we have told the meter is the ISO sensitivity of the sensor system), and z is the lens transmission.

If we compare this exposure2 with the exposure2 that would be produced for an automatic exposure system operating per ISO 2721 (see equation 1), we find that the exposure2 would be the same in the two cases if:

$$K = \frac{12.7}{z} \tag{9}$$

In fact, if we ignore the matter of lens transmission (let $z = 1$), then an exposure meter with K of 12.7 will indicate an exposure1 consistent with that used by an "ISO 2721" automatic exposure system. And a K

[11] The predecessor US standard, ANSI PH3.49-1971 essentially prescribed a K value of 12.5, and that value is widely observed today.

of about that value is indeed widely used. (One inference that can be drawn from that is that the meter manufacturers have declined to make any assumption about lens transmission, probably a prudent choice, since that varies so much between different lens designs.)

The reader may note the “coincidence” of the value of K for a free-standing exposure meter corresponding to the ISO 2721 calibration for an integrated automatic exposure system (12.7) being almost the same as the ratio H_u/H_{sat} expressed in percent (12.8). This comes about since the ratio 10/78 (the two constants appearing in the ISO 2721 and ISO 12232 definitions, respectively) is almost identical to the ratio $\pi/4$ (appearing in equation 8).

This is likely an accident. But some have suggested that the choice of the constant “78” in ISO 12232 may have been influenced by enthusiasm for this tidy coincidence.

#

APPENDIX B

The calibration of Canon digital SLR cameras

It is widely reported that Canon single lens reflex (SLR) cameras, both film and digital, utilize a value of *K*, the reflected light exposure metering constant, of about 12.5 for their integrated exposure control systems.

While, as we mentioned earlier, we cannot verify that without the use of a calibrated luminance source or an accurate photometer, the technical information office of Canon US has suggested a practical field test by which one can verify whether a Canon digital SLR camera indeed fulfills the overall exposure relationship for which the camera is designed (the relationship given, for "ISO standard" calibration and sensitivity rating, by equation 3).

Essentially, in the test, the camera, with its automatic exposure system in operation, takes an image a of scene of constant luminance. The image is inspected using Adobe Photoshop, in a mode in which the image is virtually converted to the Photoshop "gray gamma 2.2" grayscale color space. This color space has only a single color coordinate, which we may call "blackness", identified by the symbol "K". [12] [13]

If the K value for the center of the image is 55%, the overall exposure metering chain is deemed to be operating as intended by Canon.

What value of H_u/H_{sat} does a K value of 55% represent? In the "gamma 2.2 grayscale" color space, the relative luminance, Y, represented by a K value is given by:

$$Y = \left(\frac{100 - K}{100}\right)^{2.2} \qquad (10)$$

where Y is on a scale of 0-1. The formulation:

[12] Do not confuse the color coordinate "K" with the reflected light exposure metering constant, K.

[13] That symbol is presumably borrowed from the "CMYK" color space, which relates to "four-color" printing of an image. There, the coordinate "K" represents the level of black ink.

$$\left(\frac{100-K}{K}\right) \tag{11}$$

accommodates the fact that K is in percentage units and works in the "downward" direction.

Thus, a K value of 55% corresponds to a relative luminance of 0.173.[14] Does that imply an H_u/H_{sat} of 0.173? Yes, if the camera did not make significant adjustments in the "tonal scale" of the delivered image for such reasons as adjusting contrast. Tests here on a Canon EOS 20D suggest that there is not much of a problem from that cause in the part of the scale that is relevant to the test.

In any case, assuming that Canon's target value of H_u/H_{sat} is in fact 0.173, we note that this is 0.43 stop "hotter" than the value implied by the ISO standards, 0.128.

As discussed in the body of the article, the adoption of this "above standard" value of H_u/H_{sat} may well have been a suggested by the use, in these cameras, of an intelligent metering system, which avoids the need for the famous "1/2 stop cushion".

If that is in fact Canon's intent, let's look at how they did it.

Recall that the value of H_u/H_{sat} results from the interaction of two things:

a. The exposure metering calibration used by the camera's automatic exposure system (which may or may not conform with the calibration implied by ISO 2721).

b. The way in which the camera manufacturer decided to "rate" the ISO sensitivity of the sensor system for its various sensitivity settings (which may or may not conform with the rating basis given by ISO 12232).

As mentioned before, we cannot determine either of these things separately for a camera we have at hand without a test light source of known luminance or an accurate photometer.[15] However, there is considerable anecdotal evidence suggesting that, with regard to factor (a) (metering calibration), the Canon digital SLR camera metering systems indeed conform fairly closely to the calibration implied by

---

[14] In an sRGB color space basis, for a chromaticity neutral ("gray") scene), this would correspond to RGB values of about 113,113,113.

[15] And I don't have either!

ISO 2721. That's not surprising. Had they adopted a substantially non-standard metering calibration, then a photographer who preferred to use a free-standing exposure meter would get a significantly different exposure indication for a certain scene than the camera's own metering system would give—a likely cause of consternation.

So that leaves Canon with only the possibility of achieving their target value of H_u/H_{sat} by using a non-standard rating of ISO sensitivity—a rating that is about 0.74 that which would be determined under ISO 12232. In other words, the sensitivity that is designated "ISO 100" by Canon would probably be rated at about ISO 135 under ISO 12232.

#

APPENDIX C

INCIDENT LIGHT METERING

We look to an exposure meter or automatic exposure system to fulfill for us some objective with regard to the overall exposure result for an image. But what might that objective be? Here are two that are widely held:

1. We would like the brightest object in the scene to receive an exposure2 that is almost at saturation. This is often called the "expose right" philosophy, from the fact that the high end of the scale of exposure2 is ordinarily to the right on an *exposure histogram*, a graphical presentation of the distribution of exposure result in a digital image. This approach is often encouraged on the basis that it best exploits the camera's noise performance. (However, to cite a famous "tough case" for exposure metering, this approach may result in the image of a gray cat on a coal pile looking like a white cat on a gray background.)

2. We would like the relative luminance implied by the exposure result for each object to generally correspond to its reflectance. This is essentially the underlying premise of the Zone System, a famous discipline for exposure planning. This approach is often encourages on the basis that it produces the "most realistic" rendition of the scene "out of the camera". (In the famous example, this approach results in the image of a gray cat on a coal pile looking like a gray cat on a black background.)

Unhappily, the general type of exposure metering we have been discussing here—*reflected light* exposure metering—in its basic form is incapable of consistently attaining either of these objectives over a range of real scenes.

But we can consistently attain objective 2 with a different type of exposure metering, *incident light* exposure metering. There, instead of measuring the (average) **luminance** of the scene about to be photographed (the light that the scene reflects, hence the name), we measure the **illuminance** of the light falling on the scene ("incident on the scene"). If "we" (meaning the metering system) know that, we can determine the exposure1 needed to attain objective 2.

As for reflected light exposure meters, the relevant ISO standard (ISO 2720 again) defines the calibration of an incident light exposure meter in terms of an equation involving a constant (in that case, C). That equation is:

$$\frac{t}{n^2}=\frac{C}{ES_i} \tag{12}$$

where again t is the exposure time (shutter speed) in seconds, n is the (relative) aperture, as an f/number, E is the illuminance on the scene, in lux, S_i is the exposure index (what we tell the meter is the ISO sensitivity of the film or sensor), and C is the incident light exposure metering constant.

The standard allows the value of C to be from 240 to 400 (a generous range!). (In the earlier ASA standard, the value was allowed to be from 218 to 326.) The geometric mean of the two ISO limits is 310.

Perhaps we can see where that comes from.

The photometric equation that indicates the exposure2, H, on the focal plane for an object of luminance L is:

$$H=\frac{\pi}{4}L\frac{t}{n^2} \tag{13}$$

(Note that this has nothing to do with exposure meter calibration or ISO sensitivity—just the laws of physics.)

If we substitute for t/n^2 from equation 12, we get:

$$H=\frac{\pi}{4}L\frac{C}{ES_i} \tag{14}$$

The luminance, L, for a perfectly-diffuse reflecting surface with reflectance R, illuminated by illuminance E, is given by:

$$L=\frac{1}{\pi}RE \tag{15}$$

If we substitute that expression for L in equation 14, we get:

$$H=\frac{RC}{4S_i} \tag{16}$$

Now the definition of the true ISO sensitivity, S, in terms of the saturation exposure2, H_{sat}, is:

$$S=\frac{78}{H_{sat}} \tag{17}$$

If we assume we have entered the real ISO sensitivity into the meter ($S_i = S$), then combining equations 16 and 17 gives us:

$$H = \frac{RCH_{sat}}{312} \qquad (18)$$

Now, if we wish the metering operation to fulfill exposure objective 2, then for R=1 (a hypothetical "100% reflectance object) we want $H = H_{sat}$ (that is, the exposure2 would be just at saturation). Plugging those relationships into equation 18 and simplifying, we get:

$$C = 312 \qquad (19)$$

Thus we see that if our incident light exposure meter is to fulfill exposure objective 2 in connection with an "ISO standard" camera, it must be calibrated to $C = 312$. This is almost exactly the geometric mean of the ISO limits for *C*. How about that!

But in fact we hear that typical incident light exposure meters have a rather lower calibration factor, perhaps 250.

Gray card metering

If we do not have (or do not find it convenient to use) an incident light exposure meter, we can perform the same technique by using the camera's integrated exposure control system and having it regard a test surface of known reflectance, such as a "photographic gray card", that is exposed to the illumination on the scene proper.

Let us first imagine an "ISO standard" camera (with respect both to the calibration of the exposure meter and with respect to its assignment of ISO sensitivity ratings to its sensor system).

Imagine that the test target has a reflectance of 12.8%. Then, if we have the camera regard the test target, and then hold the indicated exposure1 when we actually photograph the scene, the exposure2 for any 100% reflectance object should fall just at H_{sat}, the saturation exposure2.

Often we have available a test target with a nominal reflectance of 18%. In that case, if the scene is photographed with an exposure1 that is 1/2 stop greater than indicated by the automatic exposure system (such as by setting an exposure bias of +1/2 stop), we will have the same result described just above.

What about "cushion" against the possibility of an over-saturation exposure2 for high-reflectance objects? We don't really need any cushion—the result we seek is not dependent on any assumption as to average scene reflectance, and will be reliably obtained for any scene.

What about our Canon camera?

As mentioned in Appendix B, we can expect a Canon digital SLR camera to operate on the basis of an automatic exposure system calibration consistent with ISO 2721 but with an ISO sensitivity rating that is about 0.43 stop low compared to a rating under ISO 12232.

If we follow the above guidelines for the use of "gray card" metering, it will probably result in an exposure2 for a 100% reflectance object that is almost 1/2 stop above saturation.

If we wish to avert this possibility for a scene containing high-reflectance objects, than we should probably meter from a test target with a reflectance of 18% and use the exposure indicated by the automatic exposure system "as is".

#