

The secret life of photographic exposure metering

Douglas A. Kerr

Issue 7
December 2, 2014

ABSTRACT

In photographic exposure metering we make some photometric measurement of the scene or its illumination environment, from which, combined with knowledge or presumption of the sensitivity of the film or digital sensor, is determined a photographic exposure (shutter speed and aperture) we hope will fulfill our "exposure strategy".

There are many subtleties to the concept in its various forms, and many tricky details in its execution. There are many misconceptions and misunderstandings afoot about the area.

In this article I describe the principles of the various types of exposure metering and their implications, and try to rectify some of the misunderstandings. I also discuss certain special arrangements used in exposure meters. The article, however, is not a treatise on exposure metering practice.

Extensive background is given in many matters that are predicates of the process. The presentation is somewhat technically detailed, and some basic algebra is involved here and there. The various concepts are presented in layers, a given layer possibly being visited more than once, so that at any stage the reader will hopefully have all the necessary background to follow the presentation.

1. INTRODUCTION

1.1 Photographic exposure metering

Photographic exposure metering is the process by which we make certain photometric measurements of the scene or its environment (with a free-standing instrument or an instrument integrated into our camera) and from the result, combined with knowledge or presumption of the sensitivity of the film or digital sensor, develop a recommended photographic exposure (shutter speed and aperture) we hope will fulfill our "exposure strategy".

Two principal approaches are widely used with free-standing exposure meters: *reflected light* and *incident light* exposure metering.

1.2 Reflected-light exposure metering

Here, in the basic and traditional form, essentially we determine the average luminance of the scene over a certain "field of view", usually corresponding approximately to the field of view of the camera.

From this value, together with knowledge or presumption of the sensitivity of the film or digital sensor, a photographic exposure (combination of the effect of exposure time and aperture) is recommended. I will discuss later the implications of doing this, after we have laid in some background on some of the photometric terms and concepts that are involved.

1.3 Incident-light exposure metering

Here our instrument measures the illumination that is incident on the subject (exactly what photometric property of it is determined I will discuss at length shortly). From that value, combined with knowledge or presumption of the sensitivity of the film or digital sensor, a photographic exposure (combination of the effect of shutter speed and aperture) is recommended for use in the "shot".

1.4 Inbuilt automatic exposure control systems

Many cameras have what is formally described as "automatic exposure control systems". These basically make a determination of the appropriate photographic exposure, based on a possibly-elaborated form of reflected light exposure metering, and then set a shutter speed and aperture that will produce that photographic exposure.

There are many special wrinkles that can be in play in such systems (especially in their modern forms), and I will discuss such systems separately after discussing incident light exposure metering as practiced by a freestanding exposure meter.

By the way, I will generally hereafter refer to these inbuilt automatic exposure control systems as "exposure metering systems". They are that, just with some automation accompanying them.

2. PHOTOMETRIC BACKGROUND

2.1 Introduction

The central issue here is "measurement of light", in particular, of its "potency" (in various senses), and we must keep in mind the precise significance of various quantities we encounter in that matter.

I will only speak here of the SI ("metric") units of these quantities. Some of the quantities we will not directly encounter in our work, but we must travel through them on our way to the quantities with which we will be concerned.

2.2 Luminous flux

Luminous flux is the "stuff of light". It is denominated in the unit *lumen* (lm). The symbol for luminous flux is ϕ (lower-case Greek *phi*).

2.3 Luminous intensity

Luminous intensity is the measure of the "potency" of the light emitted, in a certain direction, by a point source—that is, by a source whose dimensions are very small compared to the distance from which we observe it.

Luminous intensity has the dimensionality of *luminous flux per unit solid angle*. Why is solid angle involved? We might think that the luminous intensity in a certain direction from a "point source" would just be the amount of luminous flux emitted in that direction.

But in fact, no luminous flux is emitted in a "certain direction". The line that is the manifestation of a certain direction has zero cross-section, and thus can contain no luminous flux (just as a pipe of zero diameter can convey no flow of water).

Rather, we must think in terms of a very thin "cone" with its apex at the point source and ask, "how much luminous flux flows within that cone?" We take the ratio of the amount of luminous flux in the cone to the solid angle of the cone, in the limit as the size of the cone approaches zero. That is the *luminous intensity* along the direction of the axis of that (now infinitesimal) cone.

Conceptually, luminous intensity would be denominated in the unit *lumen/steradian* (lm/sr), where the steradian is the unit of solid angle. But there is a special name for that composite unit: the *candela* (cd). The symbol for luminous intensity is *I*.

2.4 Luminous flux density

Luminous flux density is the measure of the potency of a "beam" of light at a particular place in its travels. If we take some certain small cross-section of the beam and determine how much luminous flux there is within it, then the ratio of that amount of luminous flux to the area of that cross section is the luminous flux density of the beam.

The dimensionality of luminous flux density is *luminous flux per unit area*. It is denominated in the unit *lumen per meter squared* (lm/m²). There is no usual symbol for this quantity. (We most often encounter it only on the way to a determination of *luminance*—see below). I usually use the symbol *e*.

2.5 Illuminance

Illuminance is the measure of the impact of a light beam falling on a surface. Its dimensionality is *luminous flux per unit area*. It is conceptually denominated in the unit *lumen per meter squared* (lm/m^2), but there is a special name for that: the *lux*. The symbol for illuminance is E^1 .

If a beam of light with luminous flux density $e \text{ lm}/\text{m}^2$ arrives at a surface at angle Θ (measured with respect to a line perpendicular to the surface, called the "normal to the surface", or just the "normal"), it will cause an illuminance on the surface of $e \cos \Theta$ lux.

The factor $\cos \Theta$ does not come from any mysterious physical mechanism. It merely recognizes the fact that if we, for example, consider a part of the beam with a cross-sectional area of 1 mm^2 , which contains a certain quantity of luminous flux, and the beam strikes the surface at an angle Θ , that portion of the beam "lands" on an area of the surface whose area is $1/\cos \Theta \text{ mm}^2$ (a larger area). Thus the areal density of deposit of luminous flux on the surface (the illuminance) is $\cos \Theta$ times the luminous flux density in the beam.

2.6 Luminance

Luminance is a metric of the emission of luminous flux from a surface of finite area. We may consider such a surface to be populated by an infinity of point sources, each with a certain luminous intensity. If we consider the collective luminous intensity per unit area of the surface, that is the quantity *luminance*.

The dimensionality of luminance is *luminous intensity per unit area*, or *luminous flux per unit solid angle per unit area*. Its unit is *candela per meter squared* (cd/m^2). The symbol for luminance is L .

Luminance can be thought of colloquially as the metric of "brightness", although in fact, formally, *brightness* has a different, but closely related, meaning.

2.7 Lambertian reflection

Before we can proceed, we must look into the meaning of *Lambertian reflection*.

There are many kinds of reflection. At one extreme of the continuum is specular reflection, the kind of reflection exhibited by a mirror. There, if the light that strikes the reflecting surface is a narrow beam, what leaves the surface is also a narrow beam. The angle at which the

¹ Mnemonic: think, "Ee-luminance".

beam departs (the angle of reflection—measured with respect to the normal) is exactly the same as the angle at which the beam arrived (the angle of incidence, measured that same way). (Of course the departure direction is opposite the arrival direction.)

At the other end of the spectrum is *diffuse reflection*. Essentially, here if the light arrives as a narrow beam, it will nevertheless depart "spread over all directions" away from the surface (not generally uniformly). There can be many different variations of diffuse reflection.

We will in particular consider the classical example of diffuse reflection, a surface that exhibits *Lambertian reflection*. If we have such a surface illuminated with a certain illuminance, then:

- a. Its luminance is the same regardless of the direction from which it is observed (not from behind it, of course).
- b. Its luminance does not depend on the angle of incidence of the illuminating beam.

With regard to (b), we may be tempted to say, "but wait a minute—for illumination by a beam of a given flux density, the illuminance depends on the angle of incidence." Yes. But I said here not "illuminated with a beam of a **certain flux density**", but rather "illuminated with a **certain illuminance**." Thus we have already taken the effect of angle of incidence into account.

So exactly how the surface received that "certain illuminance"—whatever combination of *beam flux density* and *angle of incidence* produced it—does not affect the resulting luminance of the surface.

The uniformity of luminance with angle of observation turns out to mean that the amount of luminous flux per unit solid angle per actual unit of area of the surface must decline as the angle of departure increases. That's because, although I did not emphasize this earlier, luminance (a visual property, after all) is the amount of luminous flux per unit solid angle per unit area **as we see that area**—that is, the projection of that area from our vantage point. If we see the surface obliquely, at an angle Θ , what seems to us to be a certain area is actually $1/\cos \Theta$ times that area on the surface (a greater area).

Thus, if the emission of luminous flux from the surface declines with angle of departure as $\cos \Theta$, these two cosine effects cancel out, and the luminance we observe remains independent of angle of observation.

If the surface were a radio antenna, we would say that its *directivity* followed a *cosine pattern*.

Are all the surfaces of our photographic subjects always Lambertian reflectors? Certainly not. But we often proceed as if they were.

2.8 Reflection

A reflecting surface can be Lambertian but not reflect all the luminance flux incident on it (in fact, it would be rare if it did). We use the factor R (the *reflectance* of the surface) to describe the fraction of all the luminous flux incident on the surface that is reflected (in whatever direction).

If we have a Lambertian reflecting surface illuminated by an illuminance E , then the observed luminance, L , (from any direction from in front of the surface) will be (in SI units):

$$L = \frac{1}{\pi} RE \quad (1)$$

2.9 Photometric exposure

Photometric exposure is defined as the product of the luminance on the film or sensor and the time it exists (the exposure time).² Simplistically, it is the physical quality to which film or a digital sensor responds.³ Its dimensionality is *luminance times time*. Its unit is the *lux-second* (lx-s). Its symbol is H.

2.10 Photographic exposure

This is not actually a photometric quantity, but it is a critical parameter of the exposure process, and this seems a handy place to discuss it.

The "scene" being photographed, from an exposure standpoint (and let's assume a "monochrome" camera so that chromaticity issues are not involved) presents to the camera as a mosaic of varying *luminance*, with a certain overall range. The lens transforms this into a mosaic of *illuminance* upon the film or sensor.

² More generally, given that the illuminance may vary during the exposure interval, *photometric exposure* is the time integral of the illuminance.

³ However, film or a digital sensor may be subject to *reciprocity failure*, which means that for a long exposure time, the response may not exactly follow the photometric exposure.

For focus at infinity, the relationship of the illuminance on the sensor, E_s , for some on-axis spot in the scene to the luminance, L_s , of that spot on the scene follows this relationship:

$$E_s = \frac{\pi}{4} \frac{1}{N^2} T L_s \quad (2)$$

where N is the lens aperture, expressed as an f-number, and T is the transmission of the lens (the fraction of the light entering the lens that actually emerges to form the image). Thus the photometric exposure on the sensor for that spot is:

$$H_s = \frac{\pi}{4} \frac{t}{N^2} T L_s \quad (3)$$

where t is the exposure time (shutter speed).

The quantity $\frac{t}{N^2}$ is called the *photographic exposure*. There is no uniform symbol for it.

2.11 Exposure result

When we are done with the shot, we have perhaps a photographic negative or a digital image file. In the film negative, we often use the *density* of any spot on the image as the metric of exposure result there. In a JPEG (or other "developed", not "raw") digital image file, the digital color coordinates for the pixel of interest are the salient indicator of exposure result.

But we often use as the metric of exposure result what I will call the "relative luminance" implied by those coordinates. For example, if the color space of the image as we receive it is sRGB, then pixel color coordinates 180, 180, 180 (on a scale of 0-255) imply a luminance that is 45.6% of the luminance implied by the maximum possible color coordinates (255, 255, 255).

2.12 ISO speed of the sensor

The *ISO speed* is the traditional metric of the sensitivity of the film or digital sensor. We will concentrate here on digital cameras. For digital sensors, *ISO speed*, S , is defined thus^{4 5}:

⁴ In international standard ISO 12232

⁵ This is in particular what is called the "saturation-based ISO speed". There is another form of the ISO speed metric, based on signal-to-noise ratio. We will not work with it here.

$$S = \frac{78}{H_{sat}} \quad (4)$$

where H_{sat} is the *saturation photometric exposure* of the sensor system in its current sensitivity setting (the photometric exposure above which changes in photometric exposure do not result in very much change in the response).

3. EXPOSURE METERING—A CLOSER LOOK

3.1 Exposure strategy

The objective of exposure metering is to recommend or set a photographic exposure that will fulfill our *exposure strategy*. What do we mean by that?

The "scene" being photographed, from an exposure standpoint (and let's assume a "monochrome" camera) presents to the camera as a mosaic of varying *luminance*, with a certain overall range. The lens transforms this into a mosaic of *illuminance* upon the film or sensor.

We would like the range of illuminance in that mosaic to be "planted" in an appropriate place on the usable photometric exposure range of the film or sensor.

But there are several strategies we might choose for "appropriate". Two commonly-chosen ones are:

- A. "Expose to the right"⁶. Here we seek to have the "brightest" spots in the scene receive photometric exposure that is "close to saturation"—that is, close to the photometric exposure above which changes in photometric exposure do not result in very much change in the response.
- B. "Reflectance-based".⁷ Here we seek to map the portions of the scene having different reflectances approximately onto proportional values of photometric exposure (on a scale that runs to the saturation photometric exposure).

An advantage of (A) is that the range of the film or sensor is best exploited with regard to such performance properties as dynamic range and noise performance. A disadvantage of (A) is that if we

⁶ So called because "to the right" is the direction of increase in photometric exposure, exposure result, and such in various charts, histogram displays, and so forth.

⁷ This very much follows the underlying concept of the Zone System, a doctrine of exposure planning devised and promoted by Ansel Adams and others.

achieve it, the image of a "gray cat on an ash pile" (nothing else in the scene) will look like the image of a "white cat on a snowdrift".

An advantage of (B) is that, following the metaphor above, the images will reveal the various objects (cats, what the cats sit on) as we expect to see them. Stuff we know to be "gray" will look "gray", regardless of the overall scene content; Stuff we know to be "white" will look "white".

3.2 Reflected light exposure metering

3.2.1 *The inputs*

A basic external reflected-light exposure meter measures the average luminance of the scene (over a certain field of view, which may or may not closely conform to the field of view of the camera as it will be used to photograph the scene). We also feed into the meter an *exposure index*, which in basic practice would be the advertised *ISO speed* of the film or digital sensor system.

3.2.2 *The exposure metering equation*

The meter then offers us a repertoire of combinations of shutter speed and f-number, all of them amounting to the same photographic exposure. Under the international standard for free-standing exposure meters, ISO 2720, the recommendation should be developed according to this *reflected light exposure metering equation*:

$$\frac{t}{N^2} = \frac{K}{L_a S} \quad (5)$$

where L_a is the average luminance as measured, S is the exposure index, and K is a calibration constant which defines the exact relationship. ISO 2720 allows a substantial range for the calibration constant, K (why this is I will discuss later). However, it is quite common today to give K the value 12.5.

In traditional "analogue" exposure meters, the exposure metering equation is usually worked by a circular slide rule integrated into the meter.

Figure 1 shows a classic example (*ca.* 1946—a model my father had).



Figure 1. GE DW-58 exposure meter

Image © James Ollinger

3.2.3 The result on the film or sensor

Assume that:

- The field of view of the meter is the same as that of the camera (so the scene represented by the image on the sensor is the same as the scene whose average luminance is measured by the meter).
- S is set to the actual *ISO speed* of the film or digital sensor, as defined by the applicable international standard.
- We set the photographic exposure as recommended by the meter.

Recall that the average photometric exposure on the sensor, H_a , is given by:

$$H_s = \frac{\pi}{4} \frac{t}{N^2} TL_s \quad (3a)$$

where L_a is the average luminance of the scene imaged on the sensor.

Since we have set the photographic exposure to that recommended by the meter, we can substitute equation 5 into equation 3a, giving us:

$$H_a = \frac{\pi}{4} \frac{K}{S} T \quad (6)$$

Then, substituting equation 4 into equation 6, we get:

$$H_a = \frac{\pi}{312} KTH_{sat} \quad (7)$$

We further assume:

- K in the meter has the value 12.5 (very common in current practice).
- $T=1$ (of course it is never that high, but assuming this allows me to get a numerical result that is widely cited).

Then we get:

$$H_a = 0.126H_{sat} \quad (8)$$

Thus, for a shot following the meter's recommendation as to photographic exposure:

The average photometric exposure on the film or sensor will be about 12.6% of the saturation photometric exposure.⁸

If the camera does not do any intervention in the "tonal scale", we would expect that to lead to an exposure result with an average relative luminance of 12.6%.

Which of the two exposure strategies mentioned above does this fulfill? Neither. What is the point then of using this scheme? That we can do it with the simple and straightforward process I just described—basic reflected light exposure metering.

In fact, with this type of metering, unless we intervene, the image of a "white cat on a snowdrift" or a "black cat on a coal heap" will both look like the image of a "gray cat on an ash pile".

3.2.4 *Implication of this result*

What is the implication of this result? We can describe it in terms of an outlook upon which the trail of international standards is based:

If we have a uniformly-illuminated scene for which the average reflectance is 18% of the maximum reflectance, then on the film or sensor, the maximum photometric exposure will be very nearly

⁸ It is just an accident of the various constants involved that the "percentage of saturation" value of the average photometric exposure and the value of K are so nearly the same (it is because $\pi/312$ is almost equal to $1/100$).

"1/2 stop" below the saturation value of photometric exposure (that is, about 0.707 times the saturation value).⁹

So, in that situation, this comes close to fulfilling exposure strategy (A), "expose to the right". But it falls short by "1/2 stop". Why?

Well, we have no assurance that in any given scene the average reflectance is 18% of the maximum reflectance. If that ratio is lower, then the maximum photometric exposure on the film or sensor will be greater than "1/2 stop below saturation". For example, if the average reflectance were 9% of the maximum reflectance, then for a metered exposure, the maximum photometric exposure would be "1/2 stop" **above** the saturation value. We would thus have "clipping of the highlights" in the scene.

We are "safe" so long as the average reflectance of the scene is about 12.6% or more of the maximum reflectance (in which case the photometric exposure for the "brightest" spot in the scene would be just at saturation).

The value 1/2 stop is often called, by parallel with a similar consideration in audio recording, the "headroom" of the standard exposure metering plan.

3.2.5 *The liberal range for K in the standard*

Why the liberal range for K (10.6 to 13.4) in the standard? We note that there is no such thing as the "correct" exposure equation for reflected light exposure metering. For example, the process cannot inherently reliably fulfill either of the two common "exposure strategies."

Keeping that in mind, at the time that exposure meter behavior was first being codified in standards, different exposure meter manufacturers took different views as to the exposure equation that would bring their users the greatest overall satisfaction with their results in using the meter. And this was honored in the standard by its embrace of a range of values of *K*.

Incidentally, there is no "tolerance" specified upon the value of *K*. I suppose that once one is told that one can pick the value to taste, there is hardly any point in specifying a tolerance.

⁹ This expression of the relationship is based on the notion that for the "typical" scene, the average reflectance will be about 18% of the maximum reflectance. That is of course a meaningless statistic, but nevertheless it is a landmark of the traditional theology of exposure metering.

3.2.6 *Statements about "exposure meter calibration"*

We often hear it said that a "normal exposure meter is calibrated to 12.6%" (or some nearby number), or "12.6% reflectance", or "18%", or "18% reflectance".

None of those statements actually describe an exposure meter "calibration". They are all "shorthand" for the same rather complex situation—the one I described above (where K is 12.5).

You will of course see in the discussion above where the numbers "12.5%" (or thereabouts) and "18%" were found by those composing the "shorthand".

3.2.7 *"Spot" and "quasi-spot" reflected light metering*

Some photographers wish to, to some extent at least, follow "exposure strategy (B)". They may use for the purpose a free-standing reflected light exposure meter whose field of view is substantially less than the probable field of view of the camera (a "spot meter"). Thus such a meter can measure the average luminance of a certain limited region of the scene.

Common fields of view of such meters are 5° and 1° (and are almost always circularly conical—we can think of them as "circular"). A 5° field of view has a diameter about $1/5$ the diagonal field of view of a camera with a full-frame 35-mm frame size using a 100 mm lens. A 1° field of view has a diameter about $1/25$ the diagonal field of view of such a camera. These meters generally have a viewfinder through which the photographer can see exactly how the field of view of the meter lies on the scene.

Typically (although not universally), the calibration of these meters (that is, the scaling of their exposure equation) follows the same value of K we might expect in a "full scene" reflected light meter—often 12.5.

As to technique, if for example in the scene of interest there is a "gray door" which the photographer feels should be rendered in the image with a relative luminance (as an exposure result) about 12.6% of the maximum recordable luminance, then he may point the meter at that gray door, read the meter, and use its recommendations "as given" for setting the photographic exposure.

This is sometimes described as "choosing that gray door to be the 'example' for 'mid gray' in the tonal scale of exposure result", which makes sense if you think of 12.6% relative luminance as "mid gray".

Some people think that about 18% relative luminance should be considered "mid gray"¹⁰. In that case, if you want that gray door to be the "example" for mid gray, you need to point the meter at it, note its recommendations for exposure, "bump that up" by about 1/2 stop, and use that for the shot.

Further discussion of this technique is beyond the scope of this article.

3.3 Inbuilt exposure metering systems

3.3.1 Introduction

Inbuilt automatic exposure control systems (which I generally just call exposure metering systems) are usually reflected light metering systems with the following special features:

- Rather than delivering a photographic exposure recommendation, they actually put into place a certain photographic exposure, by setting the shutter speed and the lens aperture.

A given photographic exposure can of course be realized by many combinations of shutter speed and aperture. The particular pair of these parameters to be put in place to attain the photographic exposure ordained by the metering system is determined by an algorithm in the metering system, often called the "program line" (because of the way it may be graphically specified).

- Rather than the user feeding into the meter the exposure index (typically the ISO speed of the film, or sensor in use), here that value is automatically fed in when the user sets the sensitivity of the camera to one of the available choices (ISO 400, for example).
- There is usually provision for the user to ask the metering system to make its photographic exposure settings a bit higher or lower than usual, in order to get a certain photographic effect, or to "outwit" an expected bad decision by the metering system in a "difficult case".¹¹

¹⁰ One argument for that outlook is that in the "Gray gamma 2.2" monochrome color space, for a relative luminance of 18%, the single coordinate of the color space, "K", which we can think of as the metric of "grayness", has a value not far from 50% (the scale running from 0% for white to 100% for black). Another argument is that in the L*ab color space, a gray color with relative luminance 18% has an L* value of about 50 (on a scale running from 0 for black to 100 for white). So our "18%" luminance is about "mid-scale" in those two color spaces.

¹¹ For example, we may use this to make the image of a white cat on a snowdrift come out looking like that, not like a gray cat on an ash heap.

This is formally called *exposure bias*, but the common name (typically used on the camera control) is *exposure compensation*.

(In some fancy free-standing exposure meters, especially of the digital persuasion, there is a similar functionality, but otherwise the user can fake this by intentionally setting the exposure index to higher or lower than the actual ISO speed involved.

- Almost always today the metering is done through the camera lens. Thus differences in lens transmission, T , do not cause errors in the process, and we need not feel guilty about ignoring it in our analyses.
- The international standard (ISO 2721) does not (except for 8-mm and Super-8 motion-picture cameras) provide for "manufacturer's choice" of a value of K —a fixed value of 12.5 is specified, in a very indirect way (K itself is never mentioned). But then there is a tolerance of ± 1 stop, so K can legitimately be from 6.25 through 25.0! (Compliance within ± 0.5 stop is urged— K from 8.8 through 17.7.)
- The metering system may use a more intelligent premise than working from the measured average scene luminance (discussed in the next section).

3.3.2 "Intelligent" reflected light exposure metering

We have seen that the basic concept of reflected light exposure metering tries to do a lot with very little. From knowledge of only the average luminance of the scene, plus the sensitivity of the film or sensor, it attempts to follow some exposure strategy (and it has no idea which one of several recognized ones we would like it to use). But in fact it can't reliably "deliver" on any of the recognized exposure strategies.

In modern times, especially in connection with digital cameras, camera manufacturers have put a great deal of effort into what we can consider "intelligent" reflected light exposure metering.

Typically, in these systems, the luminance of the scene is determined at multiple points (in some cases, very many), perhaps guided to some degree by hints of where the important subjects may be, and then from this tries to predict the entire range of illuminance and then choose a photographic exposure that will hopefully fulfill some desirable exposure strategy.

This is a very complicated matter, and in fact the details are often closely-held trade secrets of the camera manufacturers.

In any case, we can generally think of the process as intelligently developing from the multiple measurements a *proxy average scene luminance* value that is fed into a classical exposure metering equation to develop the photographic exposure to be put into place.

Commonly, the scaling of this process is that if we have the camera regard a scene of uniform luminance (a test card perhaps), the *proxy average scene luminance* will be exactly the actual average scene illuminance. Thus the exposure result in such a case would be consistent with the result using a "basic" reflected light metering procedure.

3.4 The ISO SOS (standard output sensitivity)

As we saw earlier, the basic reflected light exposure metering equation, for what is considered a "typical" scene, will give the brightest object spot a photometric exposure of about 1/2 stop short of saturation. Thus "headroom" is to militate against "overexposure" in the event that the scene has a substantially lower ratio of average to peak luminance assumed by the metering equation model.

But with the onset of "intelligent" reflected light exposure metering systems, the risk of such overexposure is somewhat reduced. In light of that, the camera manufacturers began to regard the 1/2-stop "headroom" as "capability left on the table". If we could exploit that part of the sensor range ("eat the headroom"), we could improve the camera's dynamic range and noise performance.

Thus the manufacturers considered consistently "bumping" the photographic exposure result of the metering system up by about 1/2 stop. This could be have been done in (at least) these two ways:

- The value of K in the implicit exposure metering equation could be increased from 12.5 to about 17.7.

The disadvantage of this is that now the exposure put into place by the camera would disagree with that suggested by an external exposure meter, perhaps leading sophisticated photographers to complain of "inaccuracy" of the inbuilt system.

- The value of ISO speed assigned to the various sensitivity settings could be discounted to about 71% of its actual measured value.

Again there would be a similar disadvantage, in that in certain work the sophisticated photographer would recognize that this was not the real ISO speed of the sensor.

In order to resolve this conundrum, several camera manufacturers (I believe spearheaded by Canon) arranged to have added to the ISO

standard for the "sensitivity" of digital cameras (ISO 12232) a new metric of sensitivity, an alternative to the *ISO speed*, called the *ISO standard output sensitivity (ISO SOS)*.

This metric is defined not in terms of photometric exposure but rather in terms of digital output (as found in such color spaces as sRGB). Nevertheless, if we assume that the camera does not "tamper" with the tonal scale, we can give a photometric exposure based definition of ISO SOS that can be directly compared with the definition of S_s , the *ISO speed* metric:

$$S_{SOS} = \frac{55.2}{H_{sat}} \quad (9)$$

Thus, (compare with equation 4) we find that:

$$S_{SOS} = 0.708 S_s \quad (10)$$

Therefore, if we plug into an exposure metering equation, which is expecting S_s , S_{SOS} instead, we find the meter's photographic exposure recommendation will be almost exactly 1/2 stop "hotter" than before.

And by this process we "spend" the headroom without anybody being able to cry "wrong!".

Ah, the mind of Canon.

3.5 Incident light exposure metering

3.5.1 *The concept*

We see that basic reflected light exposure metering cannot reliably fulfill for us either exposure strategy (A) or (B). But an alternate metering technique, *incident light exposure metering*, can in many cases rather reliably fulfill for us exposure strategy (B).

Putting aside some important wrinkles, we can think of the first step in this technique as measuring the *illuminance* upon the objects in the scene from the incident light. We do this by placing the incident light exposure meter at the subject location, with its "receptor" parallel to the plane with respect to which we wish to determine the illuminance of the incident illumination (typically a plane "facing" the camera).

Figure 2 shows the noted Turkish cinematographer Erkan Umut making an incident light measurement of Sibel Can, the famous Turkish popular singer (1996). The meter is a Minolta Autometer IIIF.



Figure 2. Incident light exposure measurement

As before, we feed into the meter an *exposure index*, which again in basic practice would be the advertised ISO speed of the film or digital sensor system.

3.5.2 *The incident light metering exposure equation*

The meter then offers us a continuum of combinations of shutter speed and f-number, all of them amounting to the same *photographic exposure*. Per international standard ISO 2720, the recommendation is developed according to this *exposure metering equation*:

$$\frac{t}{N^2} = \frac{C}{E_s S} \quad (11)$$

where E_s is the illuminance upon the scene as measured, S is the exposure index, and C is a calibration constant which defines the exact relationship. The international standard for free-standing exposure meters allows a substantial range for the calibration constant, C . A common value today is 270.

3.5.3 *The result*

Let us assume for a bit that:

- the objects in the scene are all flat and face the camera
- the object surfaces are all Lambertian reflectors

Applying equation 1 to this situation:

$$L_p = \frac{1}{\pi} R_p E_s \quad (1a)$$

we find that the luminance of each point of the scene, L_p , is proportional to the reflectance of that point, R_p .

Since we have set the photographic exposure, t/N^2 , as recommended by the meter, we find that the photometric exposure on the sensor for our point will be as given by:

$$H_p = \frac{\pi}{4} \frac{t}{N^2} T L_p \quad (2a)$$

Combining equations 1a and 2a, we get:

$$H_p = \frac{1}{4} \frac{t}{N^2} T R_p E_s \quad (12)$$

And substituting equation 11 into that, we get:

$$H_p = \frac{1}{4} \frac{C}{S} T R_p \quad (13)$$

If we substitute equation 4 into that, we get:

$$H_p = \frac{1}{312} C T R_p H_{sat} \quad (14)$$

If we assume:

- C is 312 (nicely within the range allowed by the standard), and
- T is 1 (not of course realistic).

then this becomes

$$H_p = R_p H_{sat} \quad (15)$$

That is, any spot on the scene will receive a photometric exposure that is its reflectance times the saturation photometric exposure. We can consider this an ideal implementation of exposure strategy (B).

Note that this provides no "headroom". Ideally none is needed. This process does not depend on any (possibly untrue) assumptions about scene reflectance. And we might reasonably expect that the greatest reflectance we will encounter in a scene (with all Lambertian reflective objects) would be 1.0.

3.5.4 *The impact of the ISO SOS*

Incident light exposure meters usually expect the exposure index to be generally set to the *ISO speed* of the film or sensor involved.

In many modern digital cameras (most Canon models, for example), the "ISO number" to which we set the sensitivity is actually the *ISO SOS*, which is about 0.708 times the *ISO speed*.

If we set the exposure index of the meter to the *ISO SOS*, we can expect a photometric exposure for any given object spot that is about 1/2 stop "hotter" than that we would expect based on the analysis above. Is this good? Not if there are high-reflectance objects in the scene.

3.5.5 *The cosine directivity function*

We conducted our discussion of incident light exposure metering based on the assumption that all important elements in the scene are flat Lambertian reflectors and face the camera. I will continue those assumptions.

We then assumed that the incident light exposure meter actually determined the incident illuminance on the scene objects. This latter requires the following:

- The "response" of the light receptor in the meter to a beam containing a certain amount of luminous flux must vary as the cosine of the angle of incidence of the beam (over the range from -90° to $+90^\circ$). This pattern of variation of response with angle will here be called the *directivity pattern* of the receptor. This exactly parallels the definition of the illuminance created by a beam of a certain luminous flux density with certain angle of incidence.
- The meter must be oriented so that its receptor is parallel to the surface of the objects. This is consistent with the definition of the illuminance created on such object surfaces.

If all the above is true, then no matter how the illumination on the scene is composed (perhaps of illumination components from multiple directions), the entire photometric drama I described earlier will play out, and we will end up with a photometric exposure on the sensor (and thus an exposure result) for each object proportional to its reflectance.

In figure 3 are a polar plot and a rectangular plot of a cosine directivity pattern.




Figure 3. Cosine directivity pattern

3.5.6 *More generally*

On a more general front, the idealized concept of incident light exposure measurement as I described it only plays out fully if every surface element of the subject receives the same illuminance (and of course if all those elements are Lambertian reflectors).

3.5.7 *In real life*

But in reality there may be significant departures from that. For one thing, not all the object surfaces of interest face in the same direction. For example, if we consider the face of a human model, only a very small part of its surface faces the camera. If the subject is facing the camera, the forehead faces the camera, but the cheeks face at various angles away from the camera.

This diversity of orientation interacts with the fact that the luminous flux density of the incident light may not be the same for all directions of arrival, as a consequence of the arrangement of the light sources.

Thus the different surface areas may not receive the same illuminance, and accordingly, their luminances (as seen from the camera) will not be consistently proportional to their reflectances. Therefore we cannot hope to exactly achieve the presumed underlying objective of the incident light exposure metering technique (per exposure strategy B). And no metering scheme can overcome this.

3.6 An advanced incident light exposure meter

3.6.1 *Don Norwood's development*

During the late 1930s, photographer Donald W. (Don) Norwood was regularly involved with a situation in which the lighting was intentionally not uniform from all directions: the *key-fill* lighting scheme, often used for the photography of a human face. Here, one light source (the *key light*) is directed at the subject from some angle intended to give the desired "sculpting" of the face through shadowing of the features. Another light source, usually less in potency (the *fill light*), often located at or near the camera, gives

general lighting to the face to "dilute" the shadowing so it will have the desired degree for the artistic effect desired.

Norwood realized that in such cases, a basic incident light exposure meter (perhaps responsive to the true illuminance upon the plane of its receptor), whether oriented toward the camera or toward the key light source, could not be relied upon to consistently deliver a photographic exposure recommendation leading to a desirable exposure result.

In fact, common practice was to make two measurements of the incident light, with the meter oriented first toward the key light and then toward the fill light, and combine the readings mathematically before feeding that result into the exposure calculator. This was of course time-consuming for the photographer.

Norwood introduced an incident light exposure meter in which the receptor surface was hemispherical. He found that the use of such a meter, on the basis of a single measurement, gave exposure recommendations that, regardless of the location of the key light, more consistently led to good exposure results.

3.6.2 *A more economical implementation*

Norwood recognized that an actual hemispherical receptor is difficult to manufacture. He determined that a thin translucent hemispherical dome, placed over a flat receptor, could accurately emulate a hemispherical receptor. In fact today a "Norwood principle" incident light exposure meter generally uses this implementation. We sometimes call such an element a "collector".¹²

3.6.3 *The cardioid directivity pattern*

We earlier discussed that for a basic incident light meter to actually determine the *illuminance* on the plane of the meter's receptor, the directivity pattern of the receptor must be a cosine function.

For an ideal hemispherical receptor (or its emulation with an "ideal" translucent dome), the theoretical directivity pattern, plotted in polar coordinates, follows a mathematically-defined curve called a *cardioid*.¹³

This comes about from the fact that, if we have a hemispherical receptor, we find that the area it presents, seen from a point at some

¹² It is actually a specialized form of what is described in metering theory as an *incident light diffuser*.

¹³ The name, from the Latin, means "having the form of a heart". The allusion here is to the kind of "heart" we see on greeting cards.

angle with respect to the axis of the hemisphere, follows the mathematical definition of the cardioid function (here in generic form):

$$R = \frac{1 + \cos \Theta}{2} \quad [-180^\circ \leq \Theta \leq 180^\circ] \quad (16)$$

where (applied to our case) R (*the radius on the polar plot*) tells us the relative projected area as seen from a vantage point at an angle Θ from the axis of the hemisphere. This is demonstrated in Appendix A.

Accordingly, this is also the equation for the directivity pattern of our meter, where now R becomes the relative sensitivity for light arriving from a direction at angle Θ to the polar axis of the hemisphere.

In figure 4 we see a polar plot and a rectangular plot of such a cardioid directivity pattern.




Figure 4. Cardioid directivity pattern

We see that here the response declines more gradually with increasing angle than for the cosine pattern (seen in figure 3).

3.6.4 *Why is that a good behavior of the meter?*

It is widely considered by professional cinematographers and photographers that, over a range of photographic lighting situations, a hemispherical receptor meter usually gives (in one measurement) an "appropriate" exposure recommendation, meaning one that leads to an "appropriate" exposure result.

There have been offered many facile explanations of why the meter should give this desirable performance, but none are supportable by any credible physical-mathematical model.

And even if we had such a model, a problem is that there is no generalized objective criterion for what is an "appropriate" exposure result. Even if we thought to embrace "strategy B" as our ideal, we realize, from fundamental theoretical considerations, that rarely can we actually attain that.

Among these explanations, we begin with a key passage from Norwood's basic patent on the hemispherical receptor meter:

One of the particular objects of the invention is to provide an exposure meter which is substantially uniformly responsive to light incident upon the photographic subject from practically all directions which would result in the reflection of light to the camera or other photographic register.

That has a lovely plausible sound to it. But there are two problems:

- We cannot contrive a model that would show why that property would lead to a consistently "appropriate" exposure meter recommendation of photographic exposure (even if there was a clear criterion of what that might be).
- It's not true. The theoretical response of a hemispherical receptor is not "substantially uniform" with angle of incidence. Theoretically, there is a 2:1 difference in the response to light arriving "head on" vs. light arriving from 90° to the side (see figure 4).

Next we consider the commonly-given explanation that the hemispherical receptor is a proxy for the parts of the surface of a human head that are visible to the camera (which is of itself quite reasonable). If we look at the photometric implications of that concept, we find that the meter "reading" is essentially an indicator of the average (by area) illuminance on the camera-visible surface of the subject.

That sounds very nice. But again, we cannot construct a model that suggests why this measurement should lead to a photographic exposure recommendation that should be "appropriate" over a range of lighting situations (whatever that is).

3.6.5 *A helpful outlook*

In 1950, Don Norwood published a paper before the Society of Motion Picture and Television Engineers ("Light Measurement for Exposure Control", *J SMPTE 1950, 54:585-602*) that gave a helpful outlook into that mystery, not through an abstract mathematical model but rather through analysis of empirical observation in a test program. The presentation is riddled with (to me disappointing) lapses of rigor (perhaps even of candor), but fortunately these do not invalidate the practical conclusion.

I discuss (and critique) this paper in some detail in **Error! Reference source not found.**

Briefly, Norwood found that, in a key-fill lighting setup, for each of several angular positions of the key light, there was a certain

photographic exposure (greater than the exposure used for a comparison shot with the key light at the camera) which produced an image which observers adjudged to be “comparable in appearance” to the comparison shot (whatever that might be).

Norwood then went on through several stages to demonstrate that the response of a hemispherical-collector meter vs. the angle of the light hitting it¹⁴ would be such that the meter would give an exposure indication that would exactly be the exposure which the subjective tests had shown was needed to produce a consistent “visual appearance” of the image.

Sadly, the development of this conclusion is riddled with the kind of gaffes that would have caused the paper to be sent back by any credible peer review board.¹⁵

But the good news is, despite the lack of forensic credibility created by these gaffes, the ensuing numerical discrepancies are not large at all, and overall this paper still demonstrates that the readings of a Norwood system meter are a good guide to photographic exposure over a range of situations of key-fill lighting.



Figure 5. Norwood Director exposure meter, Model B

¹⁴ This is technically referred to as the *directivity pattern* of the meter.

¹⁵ This matter is discussed in detail by the author’s companion article, “Norwood’s dome: a revolution in incident light photographic exposure metering”, probably available at the same place you got this article.

3.6.6 *Norwood's legacy*

Since 1946, the firm founded (in part) by Norwood (and its successors in interest) have made a wide range of well-respected and widely-used incident light meters, almost all of them with a prominent dome.

In figure 5, we see a Norwood Director, Model B exposure meter (made by American Bolex, *ca.* 1948, visual design by Alpheus Maple).

The directivity of an almost-identical meter, the Norwood Director Model C, as measured here, is seen in figure 6. Only half of the pattern is shown; it is symmetrical.




Figure 6. Directivity pattern of Norwood Director Model C exposure meter

The theoretical cardioid pattern is shown for comparison. We see that the meter's directivity pattern is a reasonable match to the cardioid pattern.

3.7 The Weston Invercone

3.7.1 *Initial concept*

In about 1948, Weston introduced a new incident light metering accessory for their Weston Master II exposure meter (basically a reflected light meter). This incident light diffuser, which I suspect was developed by Denis Connelly of Sangamo Weston (UK), was called the "Invercone" from the inverted cone that was an important aspect of the diffuser design.



Figure 7. Weston Invercone incident light diffuser (original type)

In figure 7 we see an Invercone diffuser of the original type (mounted on a Weston Master II meter).

3.7.2 *The later design Invercone*

In 1965, Weston introduced a new design of the Invercone, which could be fitted to both the then-current Weston Master V meter and as well to the earlier Weston Master IV. In figure 8, we see it in place on a Weston Master IV light meter.



Figure 8. Weston Master IV meter with Invercone (newer type)

On figure 9, we see (in black) the directivity of this type of Invercone (measured here) on a Weston Master V meter). Note that this is a close approximation to the classical cardioid curve (shown for comparison).

We do not know what motivated this significant change in the directivity pattern of the Weston meters in incident light mode.

Perhaps Weston found that cinematographers, in particular, seemed to prefer the performance of the competing Norwood-concept meters with their cardioid directivity pattern and wanted to offer a product that would follow that behavior.

3.7.3 Directivity patterns

In figure 9, we see in red the polar plot of the directivity pattern (measured here) of an "original type" Invercone (I call it "Type 1B"), mounted on a Weston Master V meter. (This is a different item than we see in figure 7—"Type 1A"— because of differences in the mounting arrangements between the Weston Master II and III and the Weston Master IV and V.)




Figure 9. Directivity pattern of Invercone (two types) on Weston Master V exposure meter

For comparison, the theoretical "cosine" directivity pattern is also shown, in green. We see that the directivity pattern of the meter with the Type 1B Invercone in use is a very close match to the cosine pattern.

This figure also shows (in black) the directivity pattern of the newer (Type 2) Invercone and (in blue) the theoretical cardioid directivity pattern. We see that the directivity pattern of the meter with the Type 2 Invercone in use is a very close match to the cardioid pattern.

3.8 Dual-mode exposure meters

Many general-purpose photographic exposure meters offer both reflected light and incident light modes. Their basic receptors are intended for reflected light work (reporting the *luminance* of what they regard), and they are not directly suited for incident light measurement. To fit them for such, it is necessary to place an *incident*

light diffuser in front of them. These can often be slid into place for the incident light mode.

If the meter thus equipped has a cosine response (not necessarily so), we can describe the situation this way:

The diffuser accepts the total luminous flux landing on its surface and, on its rear, presents for observation by the receptor a luminous disk whose luminance is proportional to that total luminous flux.

This the meter would respond to true luminance upon the plane of the diffuser.

Often these diffusers are dome-like. Does that mean that the meter exhibits a cardioid (rather than cosine) directivity pattern in its incident light mode? In the case of the more sophisticated ones, probably. Otherwise, maybe not quite.

Figure 10 shows the Miranda Cadius dual-mode exposure meter (*ca.* 1963).

On the top edge we can see the small low dome diffuser (it is seen here slid to the left so it is not in play, the meter now being in reflected light mode). Its directivity pattern in the incident light mode is a fair approximation of the cosine pattern.



Figure 10. Miranda Cadius exposure meter

Image from KEH camera

3.8.1 *Incident light metering variations*

In fact, sophisticated meters dedicated to incident light measurement often have two selectable directivity modes. In one mode, the receptor

is essentially flat. In that mode, the directivity of the receptor is essentially a cosine function.¹⁶ In the other mode, a dome is in place. In that mode, the directivity of the receptor is essentially a cardioid function.

What is the purpose of these two directivity modes?

The "cardioid" directivity mode (with the dome) is intended for use for general-purpose incident light exposure measurement, following the concept discussed just above, as introduced by Norwood.

But sophisticated photographers may occasionally wish to make true illuminance measurements for such reasons as:

- Determining the actual illuminance on a wholly flat subject to be photographed, such as a painting or document, for exposure determination purposes.
- Determining the relative *luminous flux density* (at the subject location) of individual light sources in a studio setting so that a certain planned "balance" between their influence on the subject can be attained. (Note that to measure luminous flux density with an "illuminance" meter, we need only orient the receptor perpendicular to the line of arrival of the "beam"—"face it toward the source".) Of course, the term "luminous flux density" is never seen in discussions of this technique. But that is what is actually of interest here.
- In a non-photographic sense, such as determining the ambient illuminance upon an office desk surface for workspace illumination planning purposes.

The cosine directivity mode of these meters is intended for such tasks. In this mode, the meter actually determines the illuminance of the incident light with respect to the plane of the meter's receptor.¹⁷

Sometimes the cosine directivity mode is spoken of as the "2D incident light mode" (as it is, for one thing, intended for photography

¹⁶ In fact, an actual simple flat receptor will generally not accurately exhibit a cosine response for larger angles of incidence. Thus certain special optical features (perhaps an array of tiny lenses or prisms, or a clever scheme with a hemispherical dome below a circular aperture) are used so the receptor will actually "act flat".

¹⁷ In fact, most of the commercial meters made under Norwood's patent, and their descendants, had three "collectors" to be placed over the receptor proper: a hemispherical dome, for exposure measurement *à la* Norwood; a "cosine diffuser", to provide for true illuminance measurement; and a "grid" that equipped the receptor for reflected light measurement with a controlled field of view.

of flat objects) and the cardioid directivity mode as the "3D incident light mode" as it is intended for the photography of not-flat objects.

3.8.2 *An interesting implementation*

In certain incident light exposure meters made by Sekonic and others (and developed by Sekonic), an ingenious arrangement is used to allow either the cardioid or cosine directivity pattern to be put in place without the need to keep track of two loose "front ends" (perhaps a dome and a "cosine diffuser") to be mounted as needed.

Here, the dome (normally used to attain the cardioid directivity pattern) is always present, but for the cosine pattern mode it is retracted (typically by rotating the surrounding ring) until its peak is flush with, or just projects a little bit above, the plane of the "rim" of the "incident head" (this differs by model). The result is that the meter exhibits (at least approximately) a cosine directivity pattern. It is hard to imagine how this might happen. Sekonic politely declines to explain, citing trade secrets. Appendix B gives my current best guess as to how this works.



Figure 11. Sekonic L-358 exposure meter

In figure 11, on the left, we see the "incident light head" of a Sekonic model L-358 meter with the dome "up" (cardioid directivity). On the right, we see the same model with the dome "down" (cosine directivity).

3.8.3 *In the ISO standard*

International standard ISO 2720 gives the requirements for (free-standing) general purpose photographic exposure meters. It recognizes both the reflected light and incident light types.

For the incident light mode, it provides separate specifications for the cosine and cardioid pattern forms.

For each, it specifies the "directivity pattern" (in a curious way, like everything else in these standards).

As I mentioned before, this standard allows a "liberal" range of values for the calibration constant, C , which characterizes the incident light exposure metering equation.

Before we proceed, note that, given that the sensitivity of the meter varies with the angle of incidence of the light, the metering equation as stated applies to a "head on" incidence.

But the standard prescribes a different permissible range of values of C for the cosine and cardioid forms of the meter. The range limits for the cardioid form are about 4/3 those for the cosine form. What's with that?

We recall that a larger value of C leads, for a given illuminance measurement, to a **greater** recommended photographic exposure.

If the meter actually used values of C for the two modes that followed that ratio from the standard¹⁸, it means that in a situation where all the light on the subject came from "in front of the meter", the cardioid pattern mode has a lower response than the cosine pattern mode—it would give a greater photographic exposure recommendation..

What is the reason for this difference in "head-on" sensitivity? I suspect it comes from some empirical observations about what values of C for the two meter directivity seemed to produce "equivalent" results over some range of scene types.

3.9 Incident light metering without an incident light exposure meter

3.9.1 Introduction

We may wish to employ the concept of incident light metering but without an actual incident-light meter. There are two common ways to do this.

3.9.2 The "gray card" method.

In this method, we utilize the inbuilt exposure metering system of our camera—which essentially operates in the "reflected light mode"—to

¹⁸ And for certain meters, for which the actual values of K and C are stated in the published specifications, that is indeed so. For example, for the Sekonic L-408, we have $C_{\text{cardioid}} = 340$, $C_{\text{cosine}} = 250$ —almost exactly a 4:3 ratio. $C_{\text{cosine}} = 270$ is also common today.

regard a "gray card" placed in the scene before we take the actual shot. This card hopefully:

- Is a Lambertian diffuse reflector
- Has an accurately-known reflectance
- Has the same reflectance at all visible light wavelengths (is "chromatically neutral").

Then, the luminance exhibited by the card will be proportional to the illuminance upon the card (which of course is dependent on the orientation of the card.) Because of the latter consideration, we should orient the card parallel to the object surface for which we want the most accurate photographic exposure recommendation.¹⁹

Recall that the luminance of an illuminated Lambertian diffuse reflecting surface is given (in SI units) by:

$$L = \frac{1}{\pi} RE \quad [1]$$

We now recall the two exposure metering equations per ISO 2720:

$$\text{Reflected light: } \frac{t}{N^2} = \frac{K}{L_a S} \quad [5]$$

$$\text{Incident light: } \frac{t}{N^2} = \frac{C}{E_s S} \quad [11]$$

If we solve the system of these three equations for C, we find that:

$$C = \pi \frac{K}{R} \quad (17)$$

So if our inbuilt reflected light meter has a K of 12.5 (as it should), and our gray card has R of 0.18 (true for a widely-used one)²⁰, then in effect we have created an incident light meter with C of 218.

What would be the exposure implications of using this metering technique if we use, for the actual shot, the photographic exposure put into effect by the inbuilt meter upon observation of this gray card?

¹⁹ The usual recommendation is to face it toward the camera, which misses the point completely.

²⁰ Often called an "18% gray card", not a good idea, since in graphic arts practice "18% gray" means R=0.82 (its "grayness" is 18% of the way to full black, which is said to be "100% gray").

Well, the working of an inbuilt exposure meter should produce an average photometric exposure on the sensor of about 0.126 of H_{sat} . Thus when we have the camera observe the gray card, the photographic exposure it puts in place would (if used for a photograph of the card) result in that photometric exposure for the entire card.

If we lock in that photographic exposure and use it for the shot of the actual scene, we would expect a scene element with R of 0.18 to receive a photometric exposure of 0.126 H_{sat} . But our real objective (under exposure strategy b) is to have an object with a reflectance of 0.18 receive a photometric exposure of about 0.18 H_{sat} . Thus we must "bump" the exposure determined when regarding the gray card by about 1/2 stop²¹ when we actually shoot.

If we apply such a 1/2-stop "bump", we then can consider our simulated incident light metering rig to have a C of about 308²².

Now, if our camera operates on the basis of $K=12.5$ but uses the *ISO SOS* as S in the exposure metering equation (not the *ISO speed*), that "bump" will be inherently imposed, and we can use the photographic exposure determined when regarding the gray card as the exposure for the actual shot.

3.9.3 *The measurement diffuser method*

Here we place immediately in front of our camera lens (it usually mounts like a filter) a translucent *measurement diffuser*. Its job is to receive the light incident on it from various directions (only from in front, of course) and homogenize it into a luminous disk on its rear, which the camera's inbuilt metering system regards.

Ideally, the different rays of light incident on the face of the diffuser are "weighted" in their contribution to the luminance of the disk according to the cosine of their angle of incidence. Thus the observation by the camera's meter of the luminance of the exit disk will be consistently proportional to the illuminance on the front of the diffuser.

If, in preparation for the actual shot, we place the camera-*cum*-diffuser at the location of our subject, with the diffuser parallel to the subject surface of primary interest from an exposure metering viewpoint, then

²¹ In fact, the instructions that come with the famous Kodak gray card ($R=0.18$) at one time made just that recommendation.

²² Note that this is essentially the same as the value of C , 312, implied by the work in section 3.2.3 (see equation (7)).

we have a situation exactly comparable to having the camera regard a gray card correspondingly placed.

The ideal "transfer" of incident illuminance to exit disk luminance can be characterized thus (in SI units):

$$L = \frac{1}{\pi} JE \quad (18)$$

which we recognize as parallel to the equation for a Lambertian reflecting surface. The transfer parameter J (my notation) is exactly parallel to the reflectance, R , of such a surface. And in fact, many such diffusers are made so that J is 0.18, such that the exposure planning practice used with a gray card with R of 0.18 can be followed directly with them.

Again, if we have a diffuser with a J of 0.18, and the inbuilt meter uses the standard exposure metering equation with K of 12.5, and S is set to the *ISO SOS* for the sensor, and we use the metered exposure, as determined with the diffuser, for the shot, we can expect a result in accordance with exposure strategy (B).

It turns out that a simple translucent disk does not typically have the "cosine response" needed for this process. Thus actual measurement diffusers, just like typical "flat" incident light receptors, often have clever arrangements, perhaps involving tiny lenses or prisms, to being about a cosine response.

3.9.4 *An observation*

Note that neither of these two techniques follows the Norwood concept of dealing with a key-fill lighting situation, as do most serious incident light exposure meters.

4. ACKNOWLEDGMENT

Special thanks to Carla Kerr for her insightful copy editing of this difficult manuscript in its earlier forms. Responsibility for any errors in its current form is entirely mine.

#

APPENDIX A

Derivation of the cardioid response of the hemispherical dome

The cardioid curve

We first recall that the expression for a cardioid curve, in polar coordinate form, normalized to a maximum value of 1.0, is:

$$R = \frac{1 + \cos \Theta}{2} \quad (19)$$

The directivity response of the hemispherical receptor

We assume that the "directivity response" of a hemispherical incident light metering receptor (including as implemented with a flat receptor covered by a translucent hemispherical dome) is proportional to the projected area of the dome as seen from the angle of interest. (That area determines how much luminous flux the dome will capture from a beam of any given luminous flux density.)

The projected area of a hemisphere from various angles of observation

We will work from figure 12.

Panel a—"head on" observation

In panel a of the figure, we see the projected area of the dome as we would see it from a point on its axis. Here Θ , the angle of observation, is 0. The cosine of Θ is 1.0. We use A to represent the projected area as seen from $\Theta=0$ (that is, as seen in this panel). A will mean that very same area in future panels.

In order to set the stage for our future work, I divide the projected area into two equal portions by a vertical dotted line. The area of each portion is $A/2$.

Note that in this case, the boundary of the projected area is in fact identical to the "rim" of the hemisphere as seen from our vantage point. Accordingly, in this view, the area of each half of the projected area of the hemisphere is of half the area of the circle defined by the rim of the hemisphere. This is in turn determined by the radius of the hemisphere, R .

Panel b— observation from an angle of 60°

In panel b, we have moved our vantage point to the right by 60°, so that Θ , the angle of view of the hemisphere, is 60°. $\cos \Theta$ is 0.5.

The left boundary of the projected area is no longer the left half of the rim of the hemisphere, which has moved "around back"—just the

leftmost "limb" of the hemisphere. But the right boundary of the projected area is still the right half of the rim, which has now moved "around to the front". That half of the rim is a semi-circle, but, since we see it from an angle to its plane, we see it foreshortened as a semi-ellipse.




Figure 12. Projected area of the hemisphere

As a result of this foreshortening, the horizontal semidiameter of that projected ellipse is $R \cos \Theta$. Said another way, the width of that semicircular area is reduced by the factor $\cos \Theta$. And thus the area itself is reduced by the factor $\cos \Theta$.

Therefore the area embraced by that right-hand semi-ellipse is $(A/2) \cos \Theta$, or $A/4$. Thus the entire projected area of the hemisphere, the sum of the two sections, is $3A/4$. (That is shown in bold.)

Panel c— observation from an angle of 60°

In panel c our view is from 90° to the right. We note that for $\Theta = 90^\circ$, $\cos\Theta = 0$.

Now the "near half" of the rim of the hemisphere is seen "head on", and collapses to a vertical line; we do not see it.

Accordingly, the projected area of the hemisphere is just $A/2$.

Panel d— observation from an angle of 60°

In panel d, our view is from 120° to the right. We note that for $\Theta = 120^\circ$, $\cos\Theta = -0.5$.

As in panel b, the right boundary of the projected area is the projection to us of the "near" half of the rim of the hemisphere, now "flipped" left of the dotted centerline. Again, its horizontal semidiameter is $R \cos\Theta$ (but, to be rigorous, since $\cos\Theta$ is negative, we must state that (positive) distance as the absolute value of $R \cos\Theta$).

Thus, the total projected area, A' , is the "left portion" area, $A/2$, diminished by the area in the semiellipse, $(A/2) \cdot |\cos Q|$ (which comes to $A/4$), a net area of $A/4$.

Summary

We see that in every case, geometrically, the net projected area of the hemisphere is an area of $A/2$ to which we add an area of $(A/2) \cos\Theta$ (noting that for $Q > 90^\circ$, $\cos\Theta$ is negative, so that area then would actually be subtracted).

Algebraically, then, the projected area of the hemisphere from a point at angle Θ is consistently given by:

$$A' = \frac{A}{2} + \frac{\cos\Theta A}{2} \quad (20)$$

or

$$A' = \frac{1 + \cos\Theta}{2} A \quad (21)$$

Thus the relative sensitivity of the receptor, s , which we have assumed is proportional to the projected area of the hemisphere from the angle of interest, is:

$$s = \frac{1 + \cos\Theta}{2} \quad (22)$$

which is identical with the expression, in polar coordinates, for a cardioid curve:

$$R = \frac{1 + \cos \Theta}{2} \quad [19]$$

Quod erat demonstrandum.

A departure

Almost certainly, in the usual implementation, for angles of incidence beyond 90° there would be some obscuration of the dome by the meter housing. Thus we might expect for such greater angles the actual response would decline faster than as predicted by the cardioid curve.

#

APPENDIX B

The retracting dome scheme

Some serious incident light exposure meters (notably various models made by Sekonic), which in their normal mode use a dome to attain the "proxy for the human head" conceit (with the attendant cardioid directivity response pattern), in order to attain the cosine pattern needed for true luminance measurement, allow the user to retract the dome into a well (see figure 11).

The exact way in which this gives a cosine pattern is considered a trade secret by Sekonic, and is not, for example, disclosed in any patents.

In this appendix I present a partial concept I have of how that might work. Refer to figure 13.

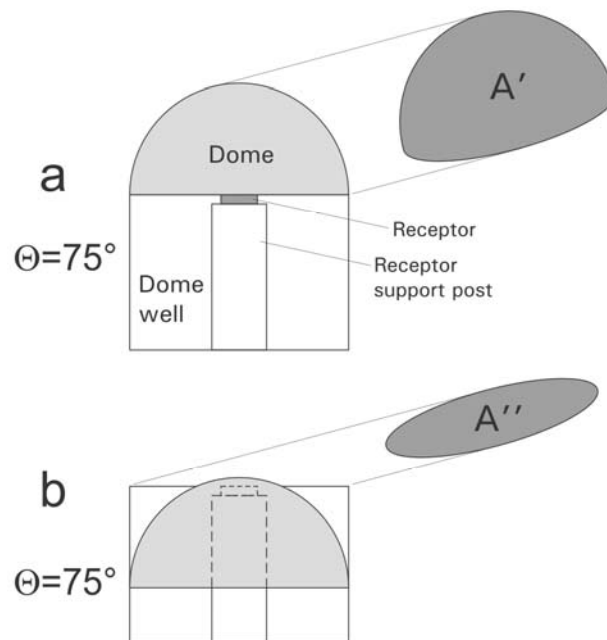


Figure 13. Retracting dome scheme

The cardioid pattern mode

For continuity with the earlier discussions, as well as to give some conceptual details of the physical layout of the Sekonic arrangement, in panel a, we review the operation of the dome in its extended position (giving a "cardioid" response pattern). The location of the receptor is such that the system operates much like the classical "integrating dome". Accordingly, we presume that the illuminance on

the receptor is proportional to the total luminous flux striking the surface of the dome.

We consider an incident light component arriving at an angle, Θ , from the axis of the hemisphere, of 75° . We see graphically (in dark gray) the projected area of the dome from that angle, which follows this "cardioid" relationship (see Appendix A):

$$A' = \frac{1 + \cos \Theta}{2} A \quad (23)$$

where A' is the projected area of the dome and A is the area of the dome as it would be seen "head on". In this case, with $\Theta = 75^\circ$, A' would be 0.629 times A .

The cosine pattern mode

In panels b, we see the dome retracted into its well for "cosine" operation.

The mouth of the well is the only means of entry of luminous flux to the dome and receptor. The effective size of this "port" with respect to a light component is the projected area of the mouth from the assumed direction of incidence of the component.

We here again specifically consider an incident light component arriving at an angle, Θ , of 75° . We see the projected area of the mouth of the well, A'' , from that angle, which follows this "cosine" relationship:

$$A'' = A \cos \Theta \quad (24)$$

In this case, with $\Theta = 75^\circ$, A'' would be 0.259 times A .

So if indeed a consistent fraction of all luminous flux that entered the mouth of the well ended up striking the receptor, we would have our cosine response. How does that happen?

Beats me.

We must conjecture that somehow the combination of the exact location of the dome with respect to the mouth of the well, the implications of the receptor now being not on the "equator" of the interior of the dome but nearer its "North pole", and the reflective properties of the wall of the well adjacent to the dome, somehow bring this about. There is no doubt more art than science at work here.