Veiling Glare AKA lens flare

Veiling glare is stray light in lenses and optical systems caused by reflections between surfaces of lens elements and the inside barrel of the lens. It is a strong predictor of lens flare— image fogging (loss of shadow detail and color) as well as "ghost" images— that can degrade image quality in the presence of bright light sources in or near the field of view. It occurs in every optical system, including the human eye.

Veiling glare can only be measured reliably from RAW images, converted to
16 or 48-bit files with gamma = 1. Processed images (JPEGs from cameras, etc.) handle dark tones (critical to veiling glare measurements) in a variety of ways that affect results.

Lenses with low flare have been traditionally known for their excellent color performance. Color saturation is higher, especially in shadows. Low flare may be as responsible as sharpness for the exalted reputations of many classic Leica and Zeiss lenses. Even though lost color saturation and contrast can be recovered with digital processing, lenses with low flare will always have an edge in quality. There are several ways of dealing with lens flare.

- In lens design, veiling glare is controlled by using high quality lens coatings (multiple coatings are best), baffling in the lens barrel, and careful design.
- In the field it is controlled by lens hoods and anything that can shield the lens surface from the sun. Ansel Adams used his hat, as well as a bellows lens shade, made famous in a Datsun (Nissan) commercial.
- In the studio it is controlled by "barn doors" on light sources.

Obligatory bad photo showing severe lens flare caused by the sun (far left), and moderate flare with the sun blocked (near left).
Canon EOS-20D, 24-70mm f/2.8L, 24mm, f/8, no hood.

Taken during Colorado's notorious "Winter of 2007″
(five feet of snow in three weeks in December and January)

The veiling glare measurement in Imatest Stepchart is designed to give a good prediction of a lens's susceptibility to lens flare in challenging situations: the larger the number, the worse it's likely to be. But it's not perfect in every detail. For example, it can't predict ghosting (visible behind the Stop sign in the image on the left). And it is not reliable for processed images; only RAW images give reliable measurements. And RAW images are often not available for inexpensive cameras and camera phones.

The measurement is fairly simple to perform, but several pitfalls can cause inaccurate or misleading results unless measurements are made with care. Some degree of exposure control— either a manual setting or exposure compensation— is recommended for best accuracy. Exposure control may not be available with camera phones and other simple devices.

Veiling glare is measured by photographing a perfectly black object against a uniform white field that extends well beyond the image frame. It is defined as the ratio of the light reaching the sensor from the black object (which emits no light) to the light reaching the sensor from the surrounding white field.

> ***Veiling glare = V = L(black object) / L(white surface)****, where* ***L*** *is the illuminance at the sensor.*

Since digital image sensor output is linear, this measurement is straightforward

- ***if*** you have access to the linear (RAW) output of the sensor,
- ***if*** the white area is unsaturated,
- ***if*** the signal in the black area is large enough to be well above the noise, and
- ***if*** the black area is perfectly black.

We deal with the last of these issues first: how to create a perfectly black object— a ***black hole***. Don't worry: it's more environmentally friendly than the type that sucks planets into its core.

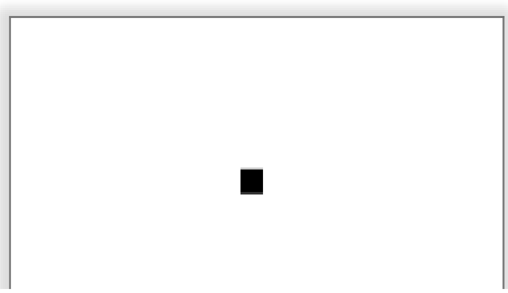

Preparing the target

Construction of the “black hole”

The black hole

The darkest surfaces— materials or pigments— reflect about 0.5 to 1% of the incident light, i.e., they are are far from perfectly black. To obtain a surface suitable for measuring veiling glare you will need to construct a surface that little light reaches— a ***black hole***— the darkest possible object that can be photographed in a bright environment.

The black hole should be constructed inside a box or tube that is approximately 3x4 inches on its top (or 4 inches in diameter) and 4-8 inches deep. The bottom (inside back surface) should be lined with black velvet — the darkest material you can buy. The sides can be lined with any matte black material. We used black art paper because it was easy to work with. The top is a piece of black foam board with a 1x2 inch opening cut in the center.

The structure should be kept as lightweight as possible. If necessary it can be constructed entirely out of black foam board, with the black velvet in the inside back surface.

The black hole is mounted next to a standard step chart such as the Kodak Q-13 or Q-14, which performs two functions. It allows you to measure the camera’s tonal response (pixel level as a function of scene luminance) and it provides a reference for determining the white level from a deep gray patch, so you don’t have to measure it directly. This is important because processed files from many cameras have nonlinear response in the bright regions, i.e., the response curve has a (film-like) “shoulder”— a good thing pictorially because without it digital cameras have a strong tendency to burn out highlights.

An alternative: Bart van der Wolf sent this link, which suggests that a cone painted glossy black can be used to make a superior black hole.

Gamma (image contrast) for measuring veiling glare is calculated by one of several methods, and can be rather different from the average gamma shown in the figures.

No Veiling glare

calculation (default)	
Veiling glare; gamma from darker zones (old)	Used prior to Imatest 3.9. Inaccurate with subtle glare (a problem with the semigloss surface of the Q=13/Q-14 charts) reduces shadow contrast.
Veiling glare; gamma from lighter zones (better)	Better if there are any problems with glare in the shadows, but not so good if there is a “shoulder” in the tonal response curve. Good with RAW files.
Veiling glare; gamma = 1 for RAW FILES!!! (best)	Best with Bayer RAW files that have been accurately converted with gamma = 1.

The white point is inferred from the region where the step chart pixel level is between 0.1 and 0.6 of the minimum-to-maximum pixel range, i.e., the brightest areas, which are frequently nonlinear (part of the tonal response curve “shoulder”), are excluded. The white reference is inferred from the patch in the middle of this region. For example, suppose it is patch 11. Since the Q-13/Q-14 has density steps of 0.1, patch 11 is 1.0 density units darker than (1/10 as bright as) patch 1, which is the reference white, i.e., we ***infer*** the white level from the measurement of patch 11: the inferred white pixel level would be (10 times the pixel level of patch 11)^{gamma}. Since nonlinearity or even clipping of highlights is quite common, this technique results in a more accurate veiling glare measurement than you’d get by measuring the white region directly. Tonal response for an in-camera JPEG and a linear (simple gamma curve) RAW conversion are shown here. Here is the modified veiling glare equation.

> ***Veiling glare = V = (Pixel level (black object) /Pixel level (white surface, inferred))***$^{1/gamma}$

Assembling the target

Typical framing for measurements

A Kodak Q-13 step chart and the black hole box are mounted on a piece of mat board, roughly twice the width of the Q-13 chart— its dimensions are not critical. A 3/4x1 inch hole is cut in the mat board, just to the right of where of where the Q-13 chart will be mounted. The black hole box is mounted behind it. The Q-13 is mounted to slightly overlap the opening. The mat board is attached to the 40x60 inch white foam board with Velcro so it can be easily removed for safe storage.

Black hole location detail

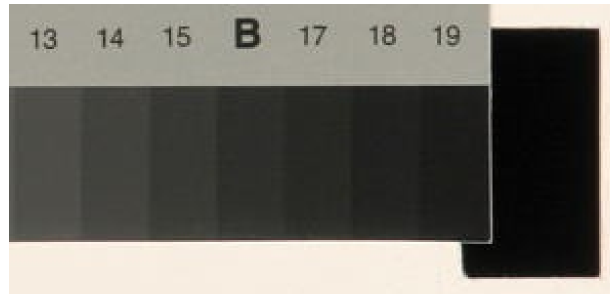

The entire target

Measuring veiling glare

- Illuminate the target evenly, with no more than about ±20% variation in illumination across the target, as described in [Imatest test lab](). Keep the ambient light near the camera as low as possible ti minimize the light entering the “black hole.”
- Note whether you are using a lens hood or a filter (like the ubiquitous UV filter most photographers leave on their lenses for protection). Both affect veiling glare.
- Set the camera to the lowest available ISO speed to minimize the noise in dark patches.

- Adjust the exposure (or exposure compensation) so the white level is close to saturation, i.e., bias it in the direction of overexposure. It doesn't hurt if some highlights are blown out: the Imatest veiling glare algorithm is insensitive to highlight saturation.
- RAW output is strongly preferred, though JPEG may be perfectly fine if the response is reasonably linear (i.e., can be described by a simple exponential; does not have a complex tonal response curve, especially in the "toe" region). Unfortunately we've seen camera phones where this assumption completely failed. Here are the recommended dcraw settings:

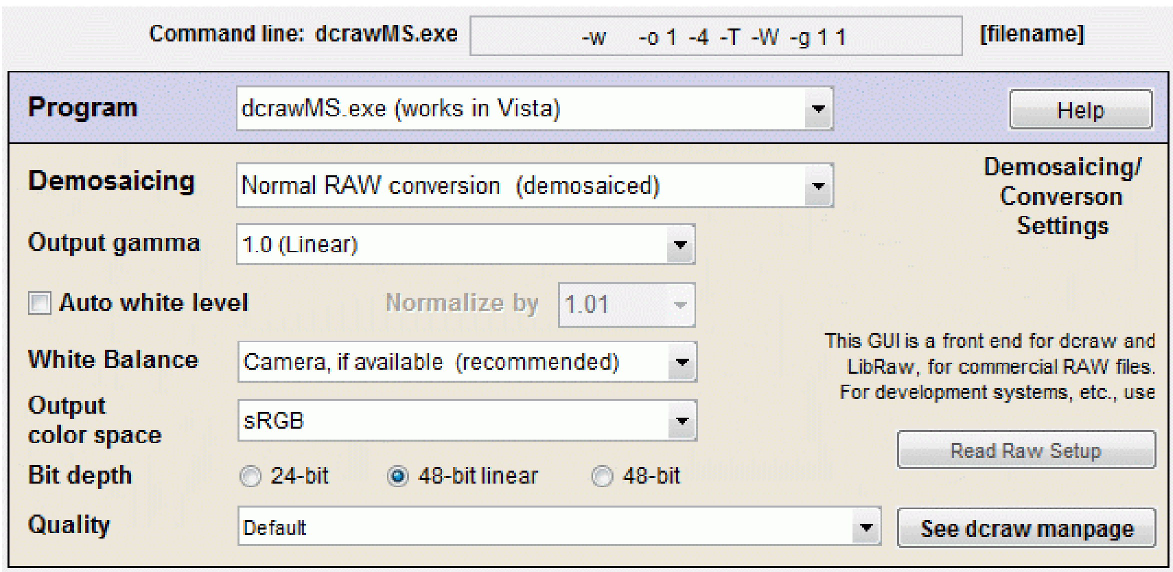

Recommended dcraw settings for veiling glare measurements

- Photograph the central portion of the target, leaving some white area around the Q-13 chart and "black hole." Typical framing is shown above, to the right of ***Assembling the target***. The length of the white target to the left and right of the frame should be about equal the frame itself. Similarly, the length of the white target above and below the frame should be about equal to the frame itself. (Light entering the lens from outside the image frame is an important contributor to lens flare.)
- Run Imatest and select . Open the image file. If Automatic Zone detection is unchecked and the image is the same size as the previous run you'll be asked if you want to repeat the same ROIs (Regions of Interest). If you enter , you'll be asked to make an ROI selection. It's OK to make a rough selection: another window will enable you to refine it.
- After you've made the rough selection, the Stepchart input dialog box, shown below, appears. To measure veiling glare, Automatic Zone detection must be unchecked, the number of patches must be set to one greater than the number of patches in the step chart: 21 in this case, and Veiling glare must be set appropriately. For RAW images decoded according to the above recommendations, we recommend **Veiling glare; gamma = 1**.

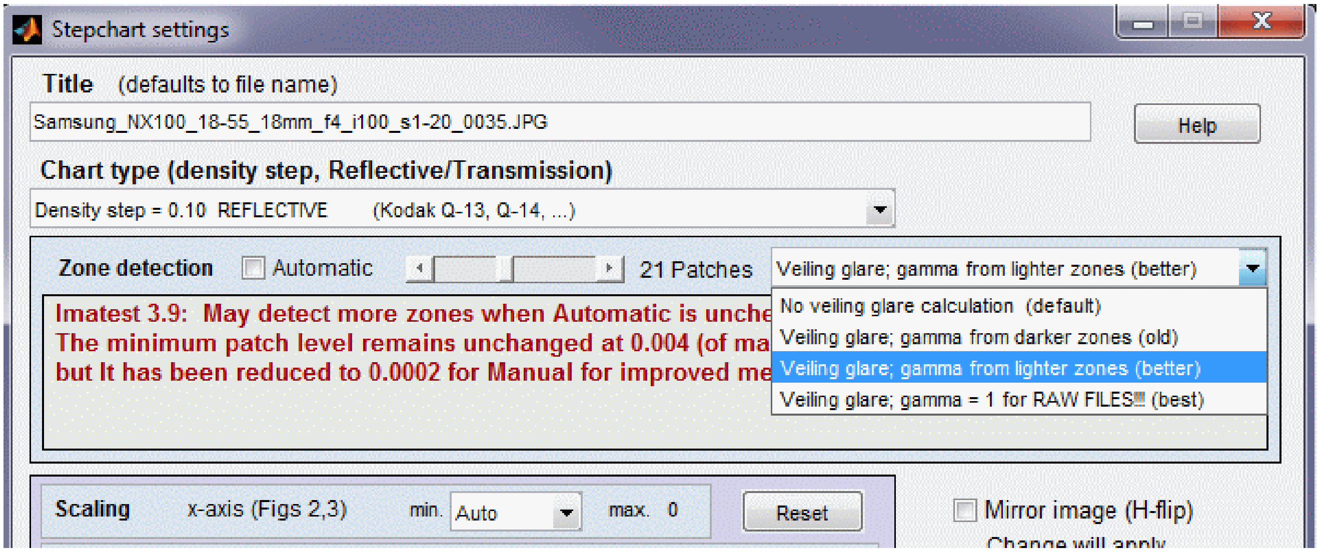

Portion of Stepchart settings box, showing veiling glare settings.

- Click . The fine adjustment box appears. Press , just to the right of to view the dark patches clearly. The button changes to when the image is in lightened mode. It also helps to maximize the window and click the radio button (middle, left).

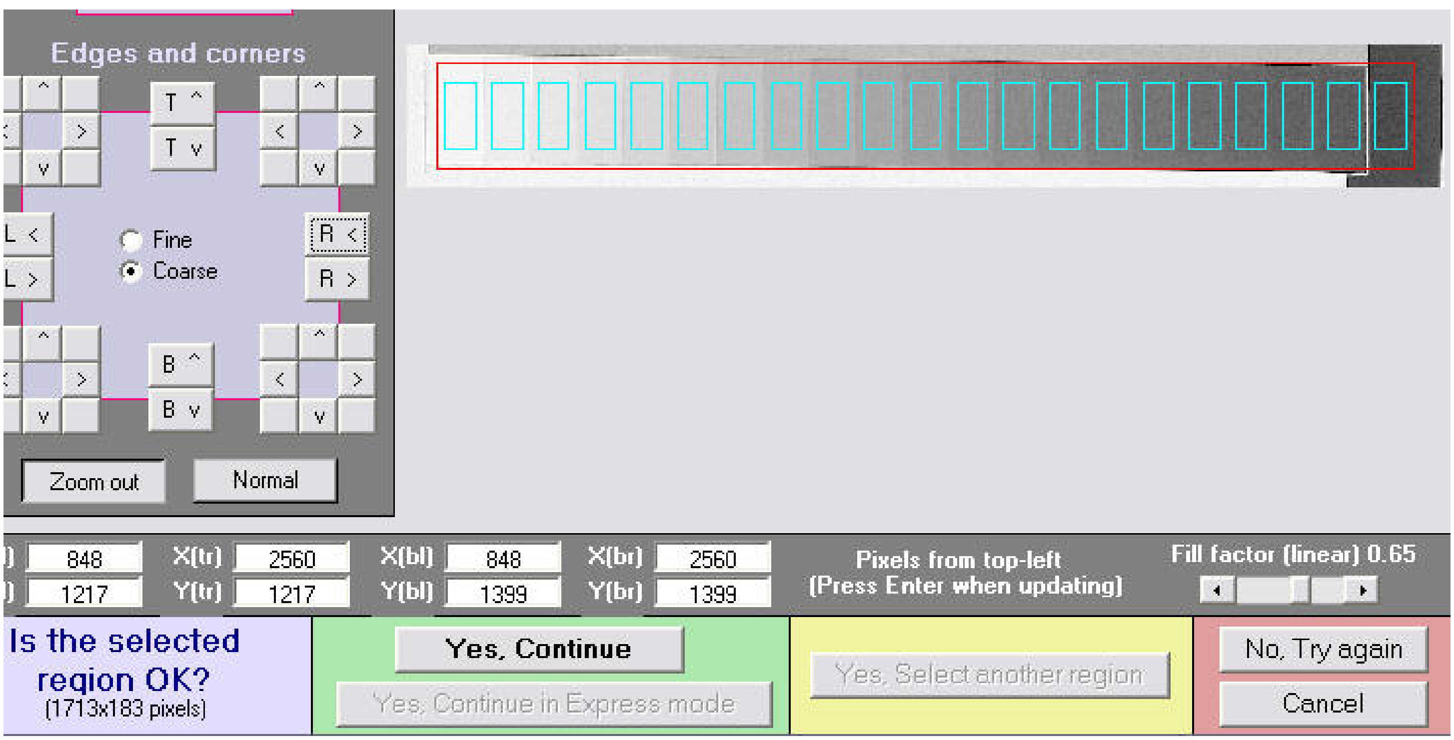

Portion of Fine ROI adjustment window; Lightened view (box is displayed)

- Click to run Stepchart and calculate veiling glare.

Results

The first illustration shows the results for the Canon EOS-20D with the 24-70mm f/2.8L lens set at 70mm, f/8. Veiling glare is shown on the middle-right of the first figure.

with UV filter (not multicoated)

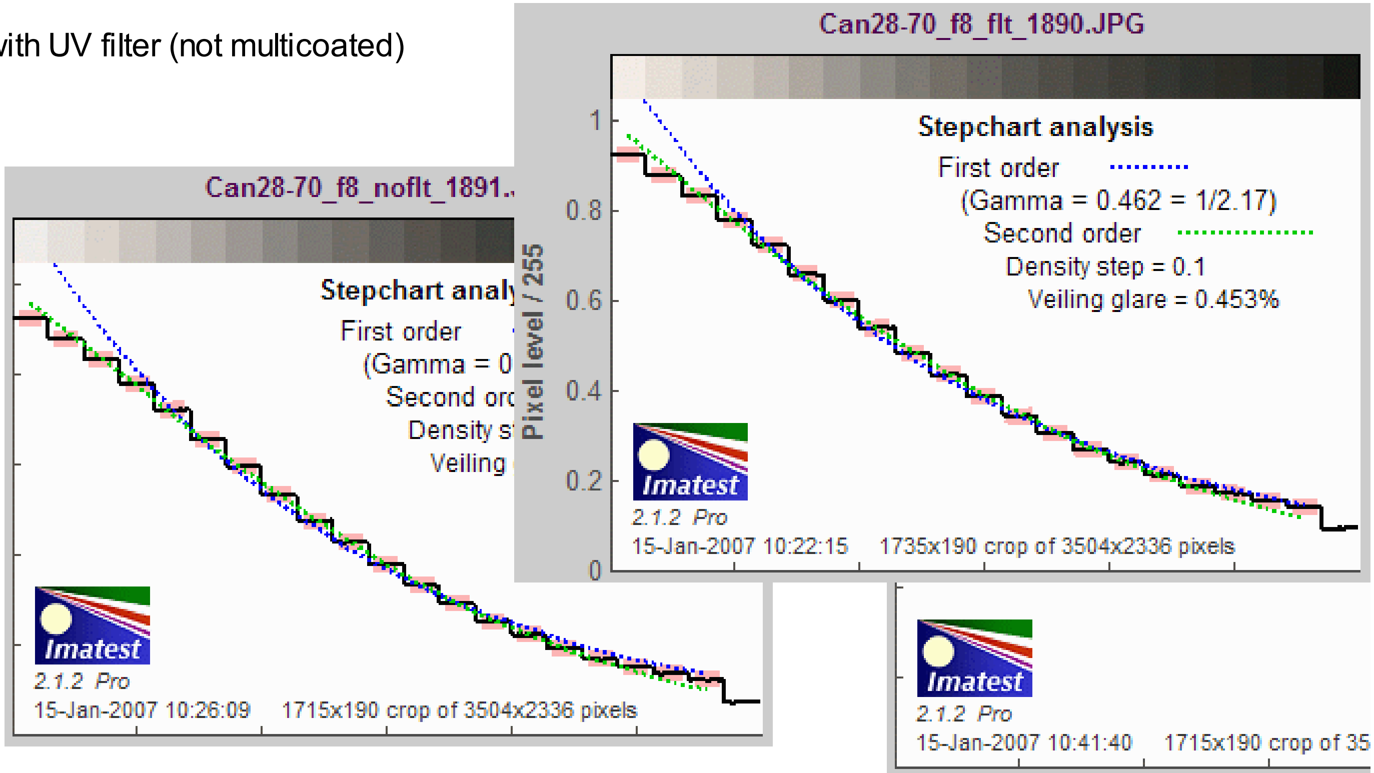

No filter: 23% lower veiling glare

Overexposed about 2 f-stops, v
Results are insensitive to ove

These plots show the difference in linearity between JPEG and RAW (to TIFF) density response. The EOS-20D only deviates slightly from linearity (a straight gamma curve).

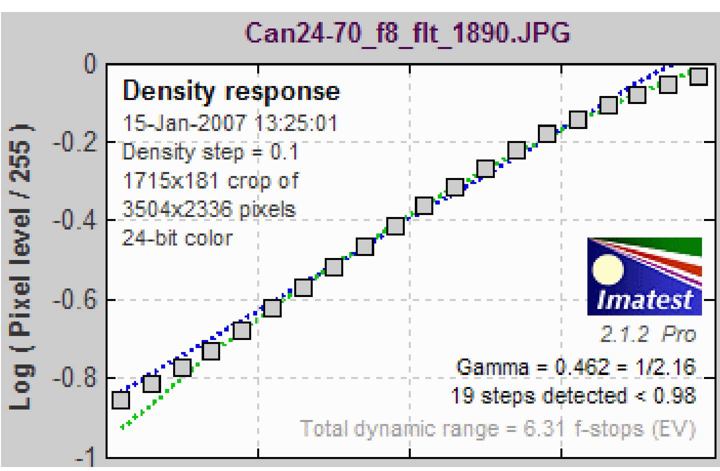

JPEG density response,
showing nonlinearity in highlights

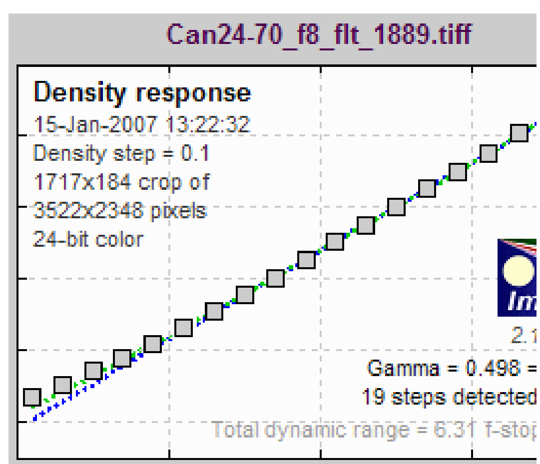

RAW/TIFF density response,
showing good linearity in highlights

Table of results: Canon EOS-20D JPEG data						
Lens	Elements	Groups	Focal length	f-stop	Veiling glare	Comments
Canon 24-70mm f/2.8L USM (no hood; single-coated UV filter)	16	13	70	8	0.453%	A premium zoom, known for excellent sharpness. Results are strongly affected by the single-coated UV filter, but they are insensitive to overexposure. When working from RAW (CR2) images, veiling glare was 0.564% with the filter and 0.391% with without it. Different results; similar trends.
” (no filter)					0.348%	
” (overexposed 2 stops, with filter)					0.464%	
Canon 90mm f/2.8 TS-E (no filter)	6	5	90	8	0.291%	A very high quality tilt/shift lens with relatively few elements and groups. Has beautiful tonality and color quality. Low veiling glare expected.
Sigma 18-125mm f/2.5-5.6 DC (hood, multi-coated UV filter)	15	14	77	8	0.293%	A remarkably fine lens for the price. Surprisingly low veiling glare. (If you’re thinking of buying one: it has poor autofocus performance; it’s much better on manual.)
Canon 28-80mm f/3.5-5/6 (no filter)	10	9	80	8	0.634%	(1991 version) A cheap “kit” lens, designed for low

						cost. Expectations were low. Strangely reddish white balance.
Canon 70-200mm f/4L USM (UV filter, hood)	16	13	70	8	0.396%	Excellent lightweight lens. The IS version has 20 elements in 15 groups, which will increase flare.
The table at the bottom shows the relationship between the number of groups (sets of attached elements) and the number of reflections that contribute to veiling glare.						

Most of the results are not surprising. The Canon 90mm has the lowest flare, while the cheap Canon 28-80 is the worst. The only surprise is the excellent performance of the Sigma 18-125: a lens know for its fine performance and value, though its autofocus performance is mediocre. I use mine on manual.

Limitations

The veiling glare may be underestimated for telephoto lenses— especially if they are measured without lens hoods— because telephoto lenses form an image of only a small fraction of the light reaching the front element (the portion is much larger with normal and wide angle lenses). The target, as shown above, may not be large enough to simulate all the light that reaches the lens. If a hood is used, this error is considerably reduced.

If there are nonlinearities in the camera response at low light levels, the results may be incorrect (although relative results, i.e., comparisons, will still be valid). It's always safest to work with RAW images and convert them with a "linear" (i.e., simple gamma curve) setting. You can read many RAW formats into Imatest, using the dcraw converter.

Lens elements, groups, reflections, and flare— some geeky math fun

Lens flare is closely related to the number of secondary reflections in a lens, i.e., light entering the lens that is reflected off one lens surface, then another, then back to the image plane. For uncoated air-to-glass surfaces, about 4% of the incident light is reflected. Simple coating reduces the reflection to around 2%; muti-coating reduces it to 1% or less. For a simple coating, the amount of the secondary reflection is 0.02 * 0.02 = 0.0004, which doesn't seem like much until you calculate the number of reflections.

A lens consists of N elements in M groups, where a group may

Groups M	**Reflec-tions R**
1	1
2	6
3	15
4	28
5	45
6	66
7	91
8	120
9	153
10	190
11	231

consist of several elements cemented together. It's the number of groups *M*— actually the number of air-to-glass surfaces *2M*— that really counts. The first surface has no secondary reflection. The second surface has 1: light that bounces off the second surface, then the first, then back to the image plane. Continuing with this reasoning, we see that the *m*(th) surface has *m*-1 secondary reflections, i.e., Total reflections = $R = (2M-1) + (2M-2) + (2M-3) + \ldots + 1 = M(2M-1) = 2M^2-M$ If you add a filter to a lens with *M* groups, you increase the number of reflections *R* by $4M+1$.	12 13 14	276 325 378

Zooms typically have more elements than primes. Examples (easy to locate in the Canon Museum): The Canon 24-105mm f/4L IS USM has 18 elements in 13 groups. The Canon 50mm f/1.4 USM has 7 elements in 6 groups. The Canon 90mm f/2.8 TS-E has 6 elements in 5 groups. Now you know one reason primes are still used (others are large apertures, light weight, and excellent sharpness (MTF)), as well as why some photographers avoid using filters (though a UV filter is useful protection for field work).

ISO 9358

Veiling glare, as measured by Imatest, is very similar to the ***veiling glare index (VGI)*** specified in the integral (black patch) test of the ISO 9358 standard, as described in Controlling Veiling Glare in an Optical Imaging System by Amber Czajkowski (University of Arizona). The key differences are

- The ISO standard calls for more rigorous testing conditions. In particular, an integrating sphere must be used. This means that there will be more light from outside the field of view, which will sometimes result in higher readings.
- The ISO standard does not include a step chart in the image. This means that only raw files (or files where the tonal response, i.e., OECF, is very well characterized) should be used.

Czajkowski' report contains a description of a really neat do-it-yourself integrating sphere project, based on stainless steel balls from http://www.gazingballoutlet.com, which sells balls up to 30 inches (0.75 meters) in diameter! A 16 inch (40 cm) ball is under $100 USD.