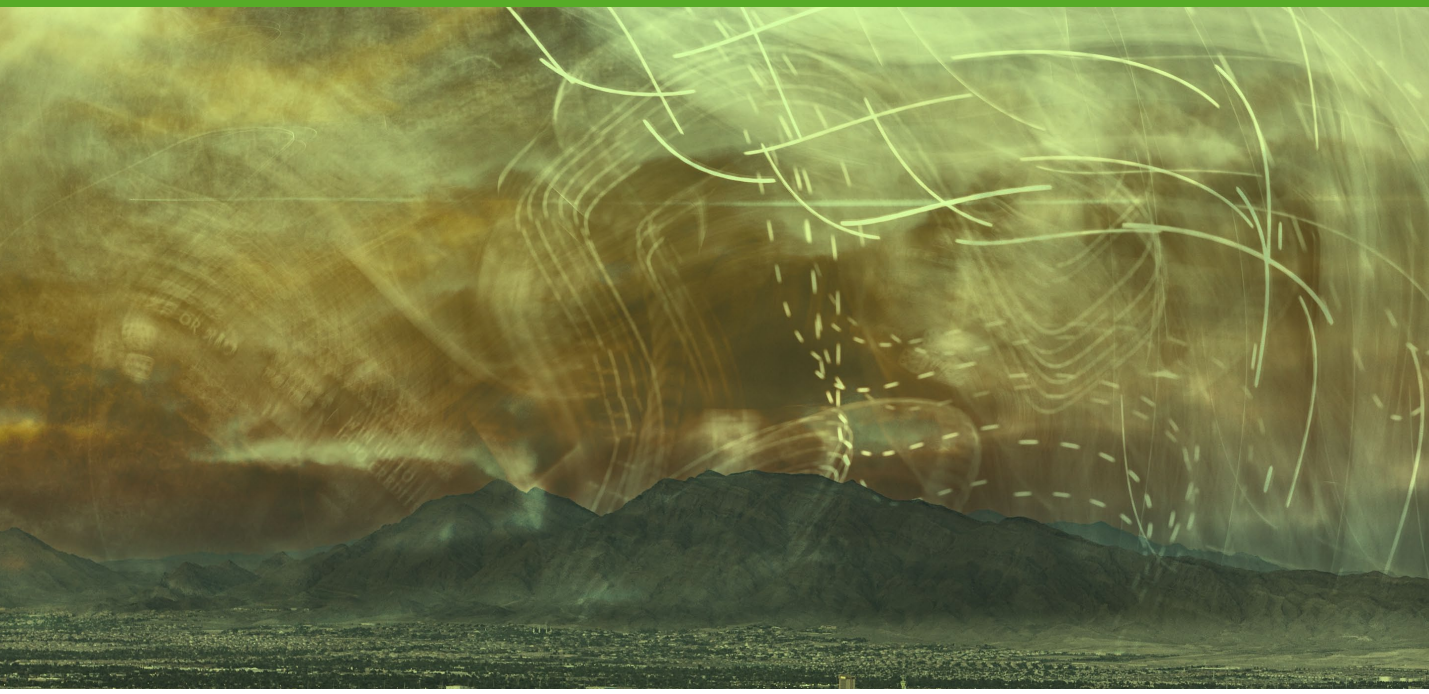

Chapter 7

Extending the dynamic range

For many years now, everyone has become preoccupied with counting the numbers of pixels in a digital capture as if this were the one benchmark of image quality that mattered above all else. Yet size isn't everything and it is really the quality of the pixel capture we should be concerned with most. Digital SLR cameras tend to have better quality sensors than compact digital cameras and the large high-end camera backs have features such as built-in cooling mechanisms that help them produce the very best in image capture quality. The one thing people haven't focussed on so much is the dynamic range of a camera sensor. Dynamic range refers to the ability of a sensor to capture the greatest range of tones from the minimum recordable shadow point to the brightest highlights and this is what we are going to focus on here in this chapter.

Camera Raw Smart Objects

A Smart Object stores the raw pixel data within a saved PSD or TIFF image (you'll learn more about Smart Objects in Chapter 9). This means you have the freedom to re-edit the raw data at any time. Although it is possible to apply the technique shown here to JPEG images, this won't bring you any real benefit compared to processing a raw image original. The important thing to stress here is that this technique really applies to editing raw files only.

Multiple raw conversions

Most of the time it will suffice for you to use the localized correction tools in Camera Raw to dodge or burn a raw photo and bring out more tonal detail where it is needed. However, if you want to extend the dynamic range of an image capture with greater precision than this you can make two or more conversions using Camera Raw, combine them together as layers in a single document and blend them to produce a composite image.
In the technique I describe here, I have shown how you can open a Camera Raw image as a Smart Object, process it in two different ways and blend the two versions together using a pixel image layer mask.

Alternatively, if you are able to shoot with the camera on a tripod, another option would be to combine two separate exposures. This would allow you to extend the dynamic range of your camera and blend the results using the manual method described here. This approach also brings you the benefit of being able to edit the individual Camera Raw Smart Object layers.

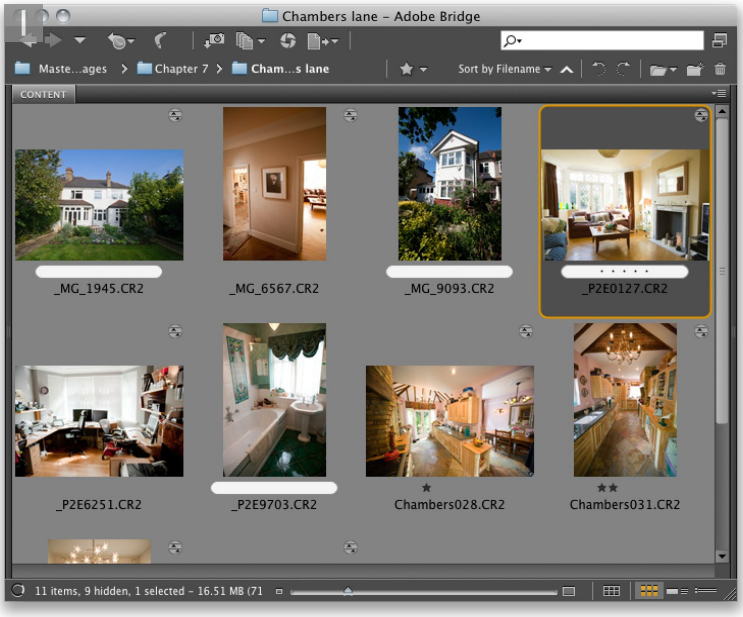

1 To begin with I went to Bridge and selected a raw image that I wanted to edit and opened this photo via Camera Raw, using File ⇨ Open in Camera Raw... (⌘R ctrl R).

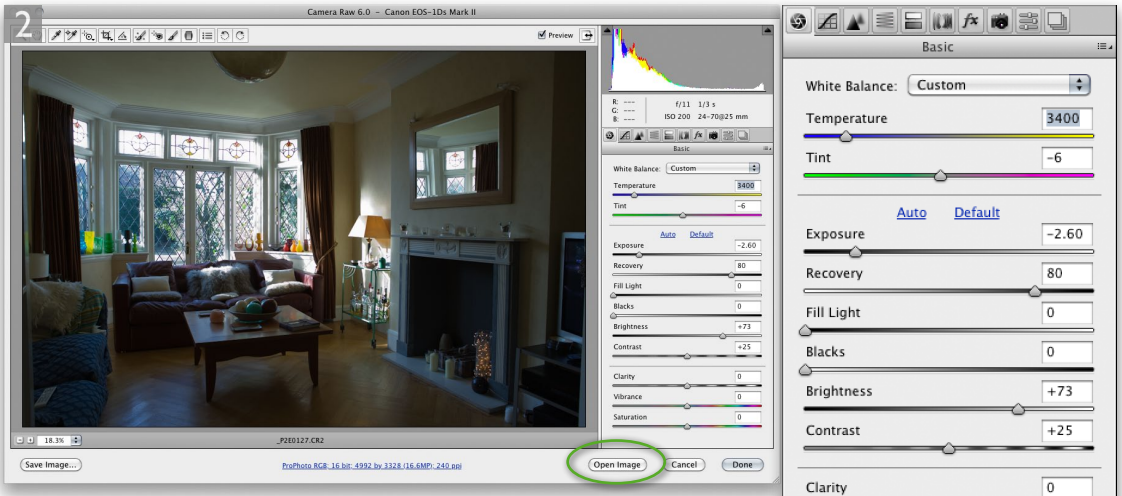

2 In the Camera Raw dialog I adjusted the Camera Raw settings to achieve the best White balance, Exposure and Recovery adjustments to reveal the detail outside the windows. Once I was happy with these settings, I held down the *Shift* key and clicked the 'Open Image' button (circled) to open this as a Smart Object in Photoshop.

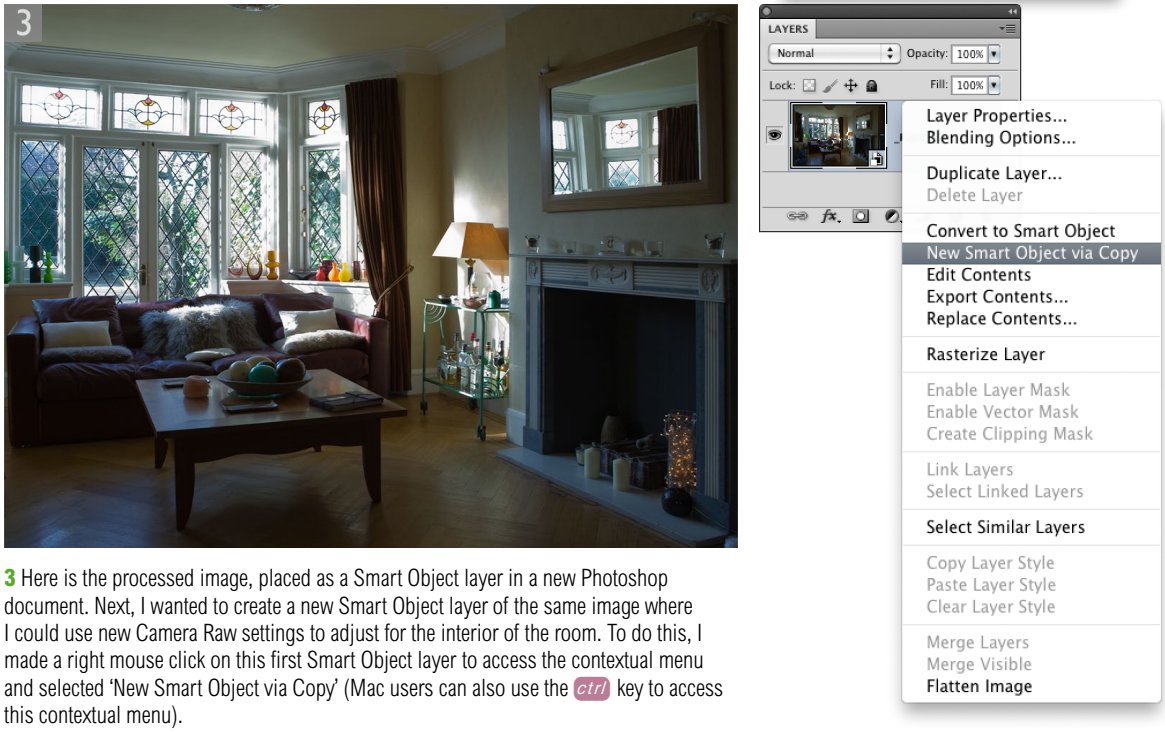

3 Here is the processed image, placed as a Smart Object layer in a new Photoshop document. Next, I wanted to create a new Smart Object layer of the same image where I could use new Camera Raw settings to adjust for the interior of the room. To do this, I made a right mouse click on this first Smart Object layer to access the contextual menu and selected 'New Smart Object via Copy' (Mac users can also use the *ctrl* key to access this contextual menu).

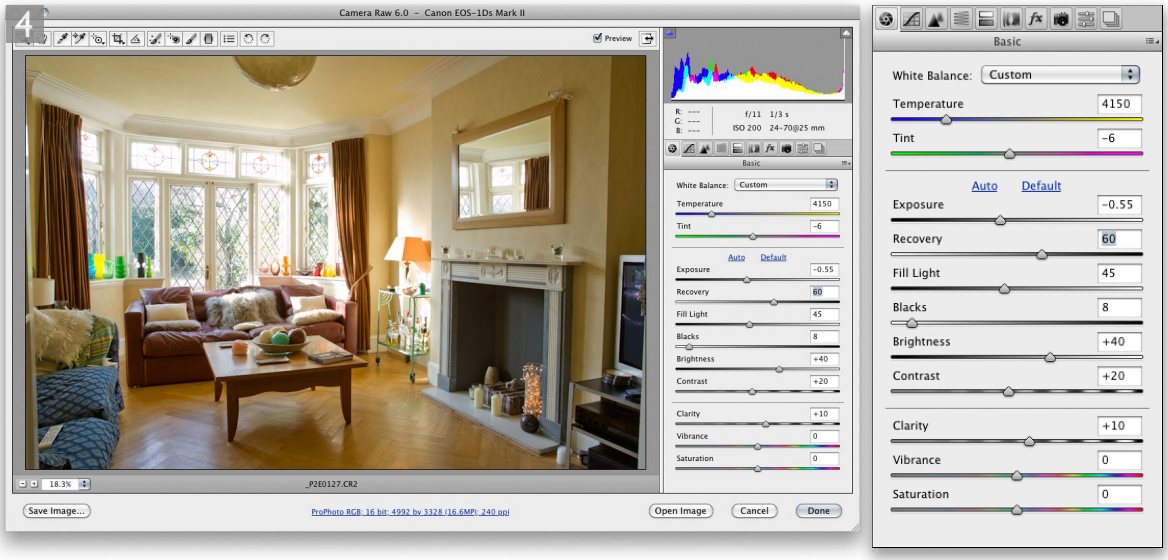

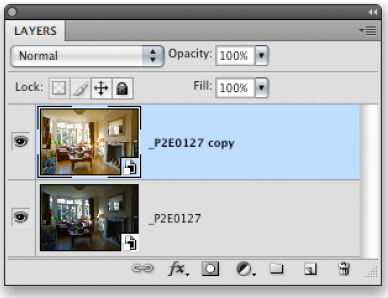

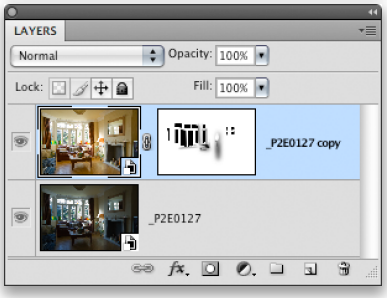

4 This duplicated the original layer and allowed me to edit the new, copied Smart Object layer. The easiest way to do this was to double-click the Smart Object layer thumbnail. This opened the Camera Raw dialog again, where I was able to use the Basic panel controls to apply a lighter adjustment to bring out more detail in the room interior. When I was finished I clicked Done to OK the adjustment, which updated the Smart Object copy layer in the master image document.

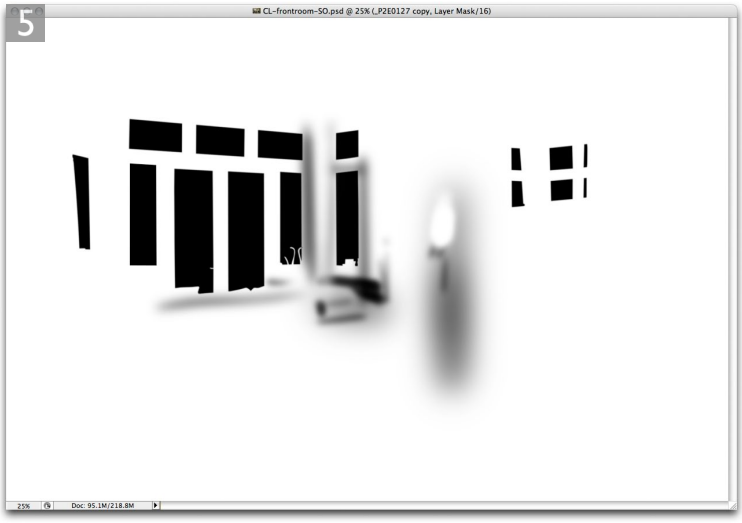

5 I then added a layer mask to the Smart Object copy layer and edited the mask to hide the windows and brush in more highlight detail in some of the brighter areas of the room interior. This screen shot shows a full image view of the layer mask that was applied to the Smart Object copy layer.

6 This final image shows the lighter processed Smart Object layer overlaying the darker processed Smart Object layer using a blend opacity of 90%. The carefully drawn mask around the window frames allowed the darker processed version to show through the windows. With this approach you can endlessly fine-tune the Camera Raw processing on each layer until you are happy with the balance achieved for the Camera Raw settings on both layers.

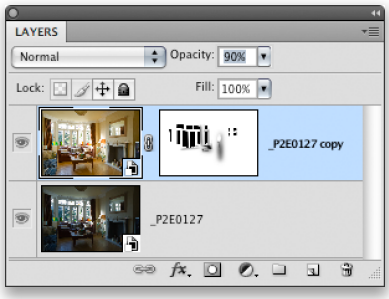

Smart Object layer blending

This technique can be adapted in various ways. You can use the layer blending options to change the layer blend mode, but you can also double-click the Smart Object layer to open the Layer Style dialog where you can adjust the 'Blend if: This Layer' options to adjust the transition between the two Smart Object layers (see page 385 for an example of where I adjusted the Layer Style 'Blend If' sliders.

Place-A-Matic script

The steps I just outlined can be carried out more easily by running Dr. Brown's Place-A-Matic script for Bridge. You can download this script for free via the Russell Brown Show website. Go to: www.russellbrown.com and then go to the Photoshop Tips & Techniques page and scroll down to the Scripts section where you can download 'Dr. Brown's Services'. This, among other things, includes the Place-A-Matic script. Follow the installation instructions and when you next launch Bridge you should see this appear as one of the new script options in the Tools menu. Basically, the Place-A-Matic script can automate almost all the steps I have just described here and offers a quicker way to create double raw settings conversions via Smart Objects (note that you will still have to add a layer mask and decide how best to blend the two Smart Object layers).

Other HDR applications

32-bit image editing is also used extensively to create the realistic CGI effects you see in many movies and television programs. These are all created using a 32-bit color space to render the computer-generated characters. It is necessary to do this in order to make them interact convincingly with the real world film footage. What usually happens is a light probe image is taken of the scene in which the main filming takes place. This is an omnidirectional HDR image which can consist of a sequence of six or seven overlapping exposures shot of a mirrored sphere. The resulting light probe image contains all the information needed to render the shading and textures on a computer generated object with realistic-looking lighting. Paul Debevec is a leading expert in HDR imaging and his website www.debevec.org contains a lot of interesting information on HDR imaging and its various applications.

High dynamic range imaging

It is interesting to see how camera sensor technology has evolved over the last few years and speculate what the future might have in store for us in the years to come. In time we may see camera sensors become available that are able to capture high dynamic range scenes in a single exposure. However, HDR cameras are not that common yet, so currently it is all about capturing bracketed sequences of images and blending these together to create single high dynamic range images that can contain the entire scenic tonal scale.

Right now there are certainly a lot of photographers who are interested in exploring what can be done using high dynamic range image editing. For example, by using the Merge to HDR Pro feature in Photoshop CS5 you can combine two or more images that have been captured with a normal digital camera, shot at different exposures and blend these together to produce a 32-bit floating point, high dynamic range image. You can then convert this 32-bit HDR file into a 16-bit per channel or 8-bit per channel low dynamic range version which can then be edited further in Photoshop.

Basically, high dynamic range image editing requires a whole new approach to the way image editing programs like Photoshop process the high dynamic range image data. Because of this the Photoshop team had to rewrite a lot of the Photoshop code so that some of the familiar Photoshop tools could be made to work in a 32-bit floating point image editing environment. Photoshop now offers a limited range of editing controls such as layers and painting and these features are available for all current versions of Photoshop CS5, rather than the extended version only (which was previously the case with CS3 and CS4).

HDR essentials

Traditionally, most camera sensors have been designed to record the light that hits the individual photosites. I won't complicate things with a discussion of the different sensor designs used, but essentially the goal of late has been to design sensors in which the photosites are made as small as possible and crammed ever-closer together so as to increase the number of megapixels. Camera sensors have also been made more efficient so that they can capture images over a wide range of ISO settings without

Figure 7.1 One approach to increasing the dynamic range of a camera sensor is to have two different sized photosites working together to capture the light hitting the sensor. One can liken this to having a wide rimmed glass and a narrow rimmed glass side by side. When water is poured, the wide rimmed glass is able to capture more water than the narrow rimmed glass, and can easily detect small amounts of water, but when more water is poured than the wide rimmed glass can handle, the narrow rimmed glass is able to pick up the slack and keep measuring the amount of water poured.

generating too much electronic noise in the shadow areas or at the higher ISO settings. The problem all sensors face though is that at the low exposure extreme there comes a point where the photosites are unable to record any usable levels information over and above the random noise that's generated in the background. At the other extreme, when too much light hits a photosite it becomes oversaturated and is unable to record the light levels beyond a certain amount. One way to look at this is to imagine photosites as being like glasses ready to be filled with water. When only a little water is added it can be hard to accurately measure how much water is in there, and if you fill the glass to the brim, the water will overflow and you'll be unable to measure any extra water that is poured (see Figure 7.1).

Fuji Super CCD

So far, Fuji are the only camera company to come up with a new approach to this problem. The Fuji-designed Super CCD comprises of two sets of photosites laid out in an octagonal pattern, where you have a standard size photosite that can record the same dynamic range as a normal photosite and a smaller photosite next to it that can record any detail that is brighter than what the standard sized photosite is able to record. The data captured using

this type of sensor can be extrapolated to create a raw capture image which represents a dynamic range that is wider than most other digital cameras. Fuji classify the two photosites as a single photosite sensor, which means that the latest Fuji S5 camera has in fact a total of 6 million 'combined' photosites, but because of the way the photosites are laid out on the sensor chip and how the data is interpreted, the effective number of megapixels is rated as 12.1 MP. Fuji have had a problem convincing everyone that this can be regarded as a genuine megapixel rating, but wedding photographers in particular have been very impressed with the dynamic range offered by this range of cameras, which is especially important when shooting white wedding dresses alongside dark suits under what can sometimes be less than ideal shooting conditions. Unfortunately some wedding photographers have been sold the idea that JPEG shooting is the best approach for weddings and in turn lost all the advantages that can be gained from high dynamic range raw capture!

Alternative approaches

Other high dynamic range sensor technologies are in the pipeline. One method relies on the ability of a sensor to quickly record a sequence of images in the time it takes to shoot a single exposure. By varying the exposure time value for each of these exposures the camera software can extract a single high dynamic range capture. The advantage of this approach is that it might be feasible to capture a high dynamic range image using a fast shutter speed, although maybe not with a high speed strobe flash unit.

Bracketed exposures

Until we have true HDR cameras, we will have to rely on using bracketed exposures instead (see Figure 7.2). The aim here is to capture a series of exposures that are far enough apart in exposure value so that we can extend the combined range of exposures to encompass the entire scenic tonal range as well as extend beyond the limits of the scenic tonal range. The advantage of doing this is that by overexposing for the shadows we can capture more levels information and this can result in cleaner, noise-free shadows. Exposing beyond the upper range of the highlights can also be useful when trying to recover information in certain tricky highlight areas. Shooting bracketed exposures is the only way most

of us can realistically go about capturing all of the light levels in any given scene and merge the resulting images into a single HDR file. When this is done right you have the means to create a low dynamic rendered version from the HDR master that can reproduce most if not all of the original scenic tonal range detail.

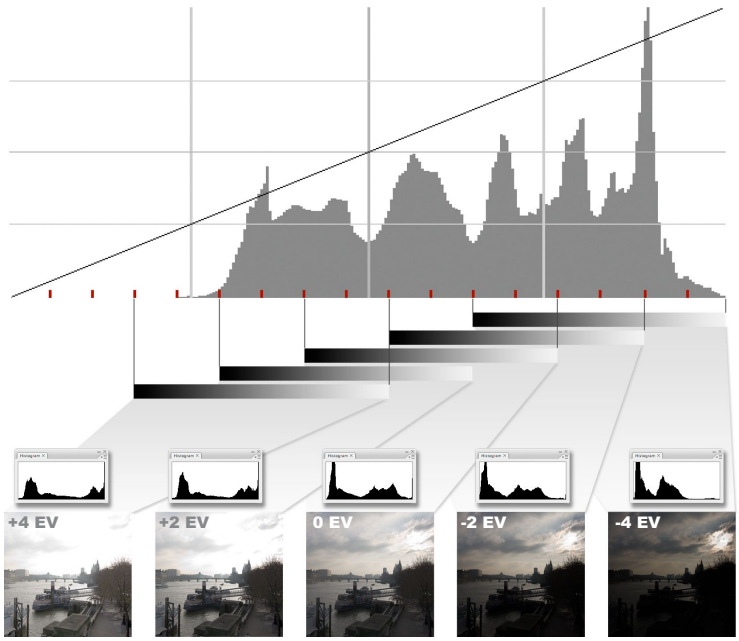

Figure 7.2 This diagram illustrates how individual bracketed exposures when merged to form a single High dynamic range image can extend the histogram scale to encompass the entire luminance of the subject scenic range.

Displaying deep-bit color

It is hard to appreciate the difference between 8-bit per channel and 16-bit per channel images, let alone 32-bit images when all you have to view your work with is an 8-bit per channel display. However, display technology is rapidly improving and in the near future we may see the introduction of displays that use a combination of LEDs and LCDs to display images at greater bit depths and over a higher dynamic range. Dolby already supply such specialist high dynamic range displays and the difference is remarkable. In the future, such displays may allow us to see the images we are editing in greater tonal detail and over a much wider dynamic range.

Capturing a complete scenic tonal range

The light contrast ratio from the darkest point in a scene to the brightest will vary from subject to subject, but in nearly every case it will certainly exceed the dynamic range of even the best digital cameras. Our human vision is able to differentiate between light and dark over a contrast ratio of 1:10,000, which in photography terms is equivalent to about 14 exposure values (EV). Meanwhile, most digital cameras can only capture a tonal range of around 6–8 EV. For the most part we have to choose our exposures carefully and decide in advance whether we wish to expose for the shadows, or for the highlights, or somewhere in between. We also know from experience that we don't always need to record every single tone in a scene in order to produce a good-looking photograph. It is after all OK to deliberately allow some highlights to burn out or let the shadows go black. However, if we wish to capture every level of tonal information in a scene, the only practical solution right now is to shoot a succession of bracketed exposures (Figure 7.2), From this we can create a single image that is capable of capturing the entire scenic tonal range.

Therefore, when capturing a high dynamic range image the objective is to make sure you capture the entire contrast range in a scene from dark to light. You can do this by taking spot meter readings and manually work out the best exposure bracketing sequence to use, plus how many brackets are required. An alternative (and simpler) approach is to use a standard method of shooting in which you first measure the best average exposure (as you would for a single exposure) and bracket either side of that using either 3, 5 or 7 bracketed exposures at 2 EV apart. This may not be so precise a method, but a 5 bracketed sequence should at the very least double the dynamic range of your camera.

There are several benefits to capturing a high dynamic range. First of all you can potentially capture all the light that was in the original scene and edit the recorded information any way you like. Secondly, providing you manage to capture all the individual brackets without any subject movement, a merged HDR image should contain smoother tonal information in the shadow regions. This is because more levels are captured at the bright end of the levels histogram (see 'Digital exposure' on page 186). Because of this the overexposed brackets will have more levels with which to record the shadow detail. When you successfully capture and create an HDR image, there should be little or no noise in the

shadows and you should have a lot more headroom to edit the shadow tones without the risk of banding or lack of fine detail that is often a problem with normally exposed digital photos.

HDR shooting tips

The first thing you want to do is to set up your camera so that it can shoot auto bracketed exposures. This can usually be done via the camera controls. Some cameras only allow you to shoot just three bracketed exposures, others more. With the Canon EOS range you should find that by tethering your camera to the computer you can use the Canon camera utilities software to set the default to five or more exposure brackets. The bracketing should be done based on varying the exposure time. This is because the aperture must always remain fixed so that you don't vary the focus between captures. Next, you want the camera to be kept still between exposures. It is possible to achieve this by shooting the pictures with a hand held camera and keep as still as possible, but for best results you should use a sturdy tripod with a cable release. Even then you may have the problem of mirror shake to deal with. This is where the flipping up of the mirror on an SLR camera can set off a tiny vibration which can cause a small amount of image movement during the exposure. However, this is mostly only noticeable if using a long focal length lens. When shooting on a tripod this can be a problem, but if you shoot hand held, the vibrations are usually dampened by your hands holding the camera. So apart from using a cable release, do enable the mirror up settings on your camera if you can. As a Canon user it has been frustrating going through the custom function menu options to set the camera to mirror up mode, but setting the mirror lock up has been made easier with the latest EOS 1Ds MkIII camera.

The ideal exposure bracket range will vary, but an exposure bracket of five exposures of 2 EV apart should be enough to successfully capture most scenes. You can use just three exposures that are 2 EV apart and get good results, but you won't be recording as wide a dynamic range. As you shoot a bracketed sequence watch out for any movement between exposures such as people moving through the frame, cars or where the wind may be causing movement. Sometimes it can be hard to prevent everything in the scene from moving and there are some software programs that are capable of removing some ghosting effects, but it's best to avoid this if you can.

Really still, still life

The Merge to HDR Pro dialog can automatically align the images for you, but it is essential that everything else remains static. You might just get away shooting the HDR merge images with a hand-held camera using an auto bracket setting, but if so much as just two of the pictures fail to register you won't be able to create a successful HDR merge. If you do resort to hand holding the camera, use a fast motor drive setting, raise the ISO setting at least two stops higher than you would use normally for hand held shooting and try to keep the camera as steady as possible (hand holding the camera should absorb some of the mirror shake movement).

If you shoot three or five exposures and separate these by 2 exposure values (EV), this should allow you to capture a wide scenic capture range efficiently and quickly. You can consider narrowing down the exposure gap to just 1 EV between each exposure and shoot more exposures. This can make a marginal improvement to edge detail in a merged HDR image, but can also increase the risk of error if there is movement in any of the individual exposures.

HDR File formats

True high dynamic range images can only originate from a high dynamic range capture device or be manufactured from a composite of camera exposures using a method such as the Merge to HDR pro option (which is described over the following pages). Photoshop's 32-bit mode also uses floating point math calculations (as opposed to regular whole integer numbers) to describe the brightness values, which can range from the deepest shadow to the brightness of the sun. It is therefore using a completely different type of image mode to describe the luminance values in an image.

If you want to save an HDR created image out of Photoshop you are offered a choice of formats. You can use the Photoshop, Large Document format (PSB) or TIFF format to save an HDR image file. These file formats can store Photoshop layers or adjustment layers, but the downside is the file sizes are at least four times that of an ordinary 8-bit per channel image. However, there are ways to make 32-bit HDR files more compact. You can use the Open EXR or Radiance formats to save your HDR files more efficiently and the Open EXR format will very often be only slightly bigger than an ordinary 8-bit version of an image. The downside is you can't save Photoshop layers using OpenEXR, but this could still be considered a good format choice for archiving flattened HDR images, despite the fact that it is utilizing less of the data than a full 32-bit per channel format such as PSB or TIFF.

How to fool Merge to HDR

Some people have asked if it is possible to take a standard single shot image, create versions of varying darkness and merge these together as an HDR image. The thing is, you can't fool Merge to HDR Pro since it responds to the camera time exposure EXIF metadata information in the file rather than the 'look' of the image. However, you can now use Image ⇨ Adjustments ⇨ HDR Toning to create a

fake HDR look from a normal dynamic range image. The way it does this is to convert an 8-bit, or ideally a 16-bit per channel image to 32-bits per channel mode and then pops the HDR Toning dialog shown in Figure 7.3 which allows you to apply HDR toning adjustments as if it were a true HDR original. Note, this only works if you are editing an image that is in RGB or Grayscale mode and has been flattened first. This isn't true HDR to LDR photography, but it does provide a means by which you can create an 'HDR look' from photographs that weren't captured using a bracketed exposure sequence.

Figure 7.3 This shows an example of HDR Toning being applied to a normal 16-bit per channel image (top) to produce the fake high dynamic range effect shown here (bottom).

Response curve

Each time you load a set of bracketed images, Merge to HDR Pro automatically stores a response curve in Photoshop's preferences for every camera it encounters. As you merge more images from the same camera, Merge to HDR Pro updates the response curve to improve its accuracy. If consistency is important when using Merge to HDR Pro to process files over a period of time, you might find it useful to save a response curve (see Step 3) and reuse the saved curve when merging images in the future.

Merge to HDR Pro

Now that you've learnt what a high dynamic range image is, it's time to put the experience into practice and go out and shoot some pictures. To avoid disappointment, I suggest you choose an easy subject to shoot with first and follow the advice on the previous pages about bracketing and using a tripod with a cable release. You can certainly get successful results from shooting JPEG images, so don't feel you have to use raw, but in my view raw gives you more options, like the ability to pre-sharpen correctly and ensure the white balance is synchronized. The Merge to HDR Pro command can be accessed via the File ⇨ Automate menu in Photoshop or via the Tools ⇨ Photoshop menu in Bridge. I usually find it best to open via Bridge, since the image alignment is applied there automatically.

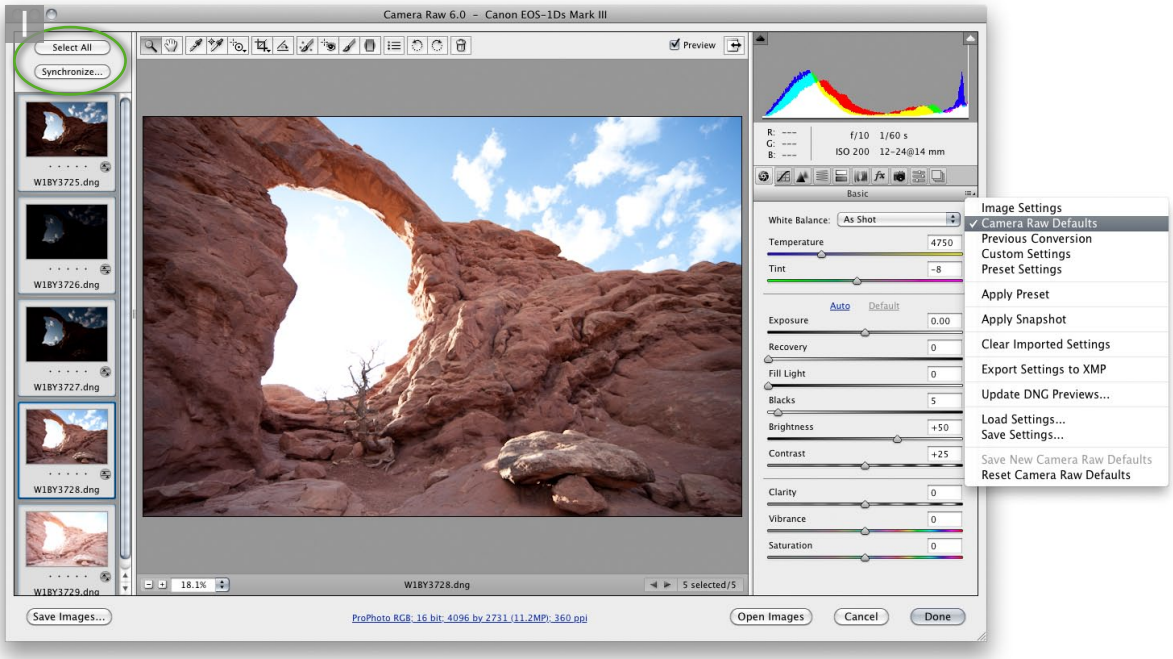

1 The original pictures were bracketed using different time exposures at two exposure values (EV) apart. I began by opening a selection of five raw digital capture images via Camera Raw. It was important that all auto adjustments were switched off. In this example, I made sure the Camera Raw Defaults were applied to the first image and synchronized this setting across all the other selected images (you'll definitely need to check the white balance is included in the synchronization if the camera was set to use an auto white balance setting).

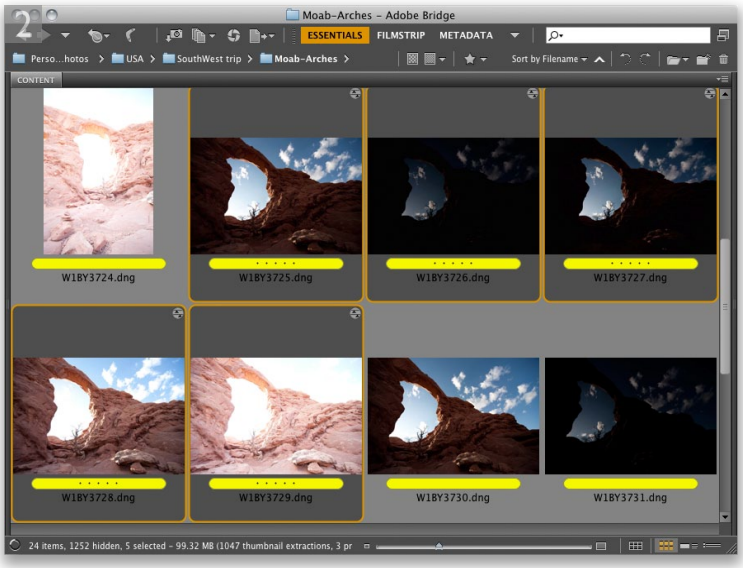

2 I kept the images selected in Bridge and went to the Tools menu and chose Photoshop ⇨ Merge to HDR Pro.

3 This shows the Merge to HDR Pro dialog in 16-bit mode. Providing the 8-bit or 16-bit mode are selected you will see the HDR toning options shown here. These allow you to apply an HDR to LDR conversion in one step (the HDR toning controls are described more fully on pages 410–413). If you prefer at this stage to simply save the image as a 32-bit master HDR file, you should select the 32-bit mode, where the only option available is to adjust the exposure value for the image preview. There is also a fly-out menu in the Merge to HDR Pro dialog (circled below) where you can deselect the Automatic Response curve mode and also choose to save or load a custom response curve (see sidebar opposite).

Merge to HDR Pro script

There is a 'Merge to HDR' script you can load from the Presets/Scripts folder that allows you to open files or folders of images to process via Merge to HDR Pro. It does not allow you to process layered files, although there are hooks present that could allow this to be scripted.

Exposure and Gamma

You can use the Exposure slider to compensate for the overall exposure brightness and the Gamma slider to (effectively) reduce or increase the contrast. These controls are rather basic, but they do allow you to create a usable conversion from the HDR image data.

Highlight Compression

The Highlight Compression simply compresses the highlights, preserving all the highlight detail. It can render good midtones and highlights at the expense of losing detail in the shadows.

Equalize Histogram

The Equalize Histogram option attempts to map the extreme highlight and shadow points to the normal contrast range of a low dynamic range Photoshop image, but this will usually make for a rather blunt conversion.

Tone mapping HDR images

After you have created a merged 32-bit per channel HDR image, you can save the HDR master using the PSB or TIFF format to preserve maximum image detail plus any layers. Or, you can use the EXR format, which as I explained earlier is a more efficient, space saving file format for storing 32-bit images (but lossy). You can if you like skip saving the merged HDR image and jump straight into the tone mapping stage by selecting the 16-bit per channel or 8-bit per channel option in the Merge to HDR dialog. I think you will find though that there are some definite advantages to preserving a master image as an HDR file. There is a real art to tone mapping an image from a high dynamic range to a normal, low dynamic range state and you won't always make the best judgement on your first try. It therefore makes sense to save the HDR file first as a 32-bit master image and then use the Image ⇨ Mode in Photoshop to convert from 32-bits to 16-bits or 8-bits per channel. This pops the HDR Toning dialog (Figure 7.5), which offers four methods of converting an HDR image to a low dynamic range version (see the sidebars on the left and the section below). With each of these the aim is the same: to squeeze all of the tonal information that is contained in the high dynamic range master down into a low dynamic range version of the image. Here I am mainly going to concentrate on the Local Adaptation method.

Local Adaptation

The Local Adaptation method is designed to simulate the way our human eyes compensate for varying levels of brightness when viewing a scene. For example, when we are outdoors our eyes naturally adjust and compensate for the difference between the brightness of the sky and the brightness of the ground. The difference in relative brightness between these two areas accounts for the 'global contrast' in the scene. As our eyes concentrate on one particular area, the contrast we observe in say, the clouds in the sky, or the grass on the ground is contrast that is perceived at a localized level. The optimum settings to use in an HDR conversion will therefore depend on the image content. In Figure 7.4 we have a photograph of a scene that has a high dynamic range. The global contrast would be the contrast between the palm tree seen in silhouette against the brightly lit buildings in the background, while the localized contrast would be the detail contrast within both the bright and dark regions of the picture (magnified here).

The Radius slider in the Local Adaptation HDR Toning dialog Edge Glow section (Figure 7.5) is said to control the size of the glow effect, but I prefer to think of this as a 'global contrast' control. Basically, the tone mapping process lightens the shadows relative to the highlights and the tone mapping is filtered via a soft edge mask. Increasing the Radius amount widens the halos. At a low setting you'll see an image in which there may be a full tonal range from the shadows to the highlights, but the image looks rather flat. As you increase the Radius this widens the halos, which softens the underlying mask and this is what creates the impression of a normal global contrast image. You can then use the Strength slider to determine how strong you want the effect to be. At a zero Strength setting the picture will again look rather flat. As you increase the Strength amount, you'll see more contrast in the halos that are generated around the high contrast edges in the image. The Glow Strength slider can therefore be used to soften or strengthen the Radius effect, but you do need to watch for ugly haloes around the high contrast edges.

When the Gamma slider is dragged all the way to the left, there is no tone compression between the shadows and highlights. As you drag the other way to the right, this compresses the shadows and highlights together. The Exposure slider can then be used to compensate for the overall exposure brightness. Note that this slider adjustment can have a strong impact as it is applied after the tone mapping stage rather than before.

The Detail slider works a bit like the Clarity slider that's found in Camera Raw and you can basically enhance the localized contrast by adding more Detail. The Shadow and Highlight sliders are fine-tuning adjustments. These sliders can be used to independently adjust the brightness in the shadows or highlight areas. For example, the Shadow slider can be used to lighten the shadow detail in the darkest areas only. HDR Local adaptation conversions typically mute the colors so you can use the Vibrance and Saturation sliders to control the color saturation.

Finally, we come to the Toning Curve and Histogram. You can use this to apply a tone enhancing contrast curve as a last step in the HDR conversion. The histogram displayed here represents that for the 32-bit image, but you'll find the Histogram panel in Photoshop more useful when gauging the outcome of a conversion. When you are done you can click on the OK button for Photoshop to render a low dynamic range version from the HDR master.

Figure 7.4 Here is an example of a subject that has a wide dynamic range and strong global contrast.

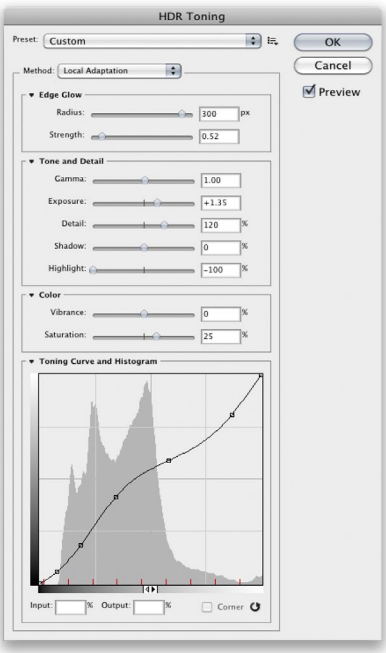

Figure 7.5 The Local Adaptation tone mapping method (also displaying the Tone Curve and Histogram options).

Photomatix Pro

Photoshop's tone mapping methods are designed to help you create natural looking conversions from an HDR to an LDR image. Photomatix Pro has proved extremely popular with photographers because it offers excellent photo merging (sometimes with more accurate image merging of hand held shots), ghosting control and above all, more extensive tone mapping options. Tone mapping with Photomatix Pro is much easier and also allows you to create those illustration-like effects that are often associated with a high dynamic range image look (which we should really call 'HDR to LDR converted images'). One explanation for the difference between the Photoshop HDR conversion method and Photomatix Pro may be because Photoshop uses a bilateral filter for the tone mapping, while Photomatix Pro uses a gaussian filter which produces more noticeable-looking halos. I like using Photomatix Pro and find it quick to work with, but I do also find the new Merge to HDR Pro in CS5 works extremely well and I now find myself split between working with both.

Removing ghosts

It is important to minimize any movement when shooting exposures, which is why it is best to shoot using a sturdy tripod and cable release. Even then there remains the problem of objects that may move between exposures such as tree branches blowing in the wind. To help address this the Merge to HDR Pro process in Photoshop CS5 utilizes a new ghost removal algorithm which automatically tries to pick the best base image to work with and discards the data from other images in those areas where there is movement. When the 'Remove Ghosts' option is checked in the Merge to HDR Pro dialog you'll see a green border around whichever thumbnail has been selected as the base image. You can override this by clicking to select an alternative thumbnail and make this the new base image. For example, if the moving objects are in a dark portion of the photograph then in these circumstances it will be best to select a lighter exposure as the base image. In practice I have found the ghost removal to be very effective on most types of subjects, although moving clouds can still present a problem. Skies are also tricky to render because the glow settings can produce a noticeable halo around the sky/horizon edge. This problem can usually be resolved by substituting the medium exposure sky image as a separate layer with a layer mask based on the outline of the sky.

How to avoid the 'HDR' look

It has to be said that HDR to LDR converted images can sometimes look quite freaky because there is a temptation to squeeze everything into a low dynamic range. Just because you can preserve a complete tonal range does not mean you should. It really is OK to sometimes let the highlights burn out or let the shadows remain black. The Photomatix Pro program has proved incredibly popular with HDR enthusiasts, and in the process, spawned the classic 'HDR' look, which I personally think has been done to death now. Besides, Photomatix Pro can actually be used to produce nice, subtle tone mapped results, but I suppose people don't notice these types of HDR photos quite so much. The Photoshop approach also lets you produce what can be regarded as natural-looking conversions and I think you'll agree that the Figure 7.6 example shows how the HDR to LDR image process can result in a photo that looks fairly similar to a normal processed image, but with much improved image detail in the shadow regions.

Figure 7.6 On this page you can see a comparison between a single edited image, using an optimum exposed photograph (top), processed via Camera Raw and output as a 16-bit file. Below you can see the HDR edited version that was converted to make a 16-bit low dynamic range image. At first glance the difference is quite subtle. I did after all try to get the two images to match as closely as possible, but you should notice better tone and detail contrast in the rocks in the HDR converted version. The difference was more noticeable though when I examined the shadow areas. In the enlarged close-up views you can see there is much more image detail and virtually no shadow noise in the bottom image.

Exposure slider

It is impossible to represent an HDR image on a standard computer display, which is why the Exposure slider is available as a slider at the bottom of the document window. It allows you to inspect an HDR image at different brightness levels. Since the display you are using is most likely limited to a bit depth of 8-bits, this is the only way one can actually 'see' what is contained in a high dynamic range 32-bit image.

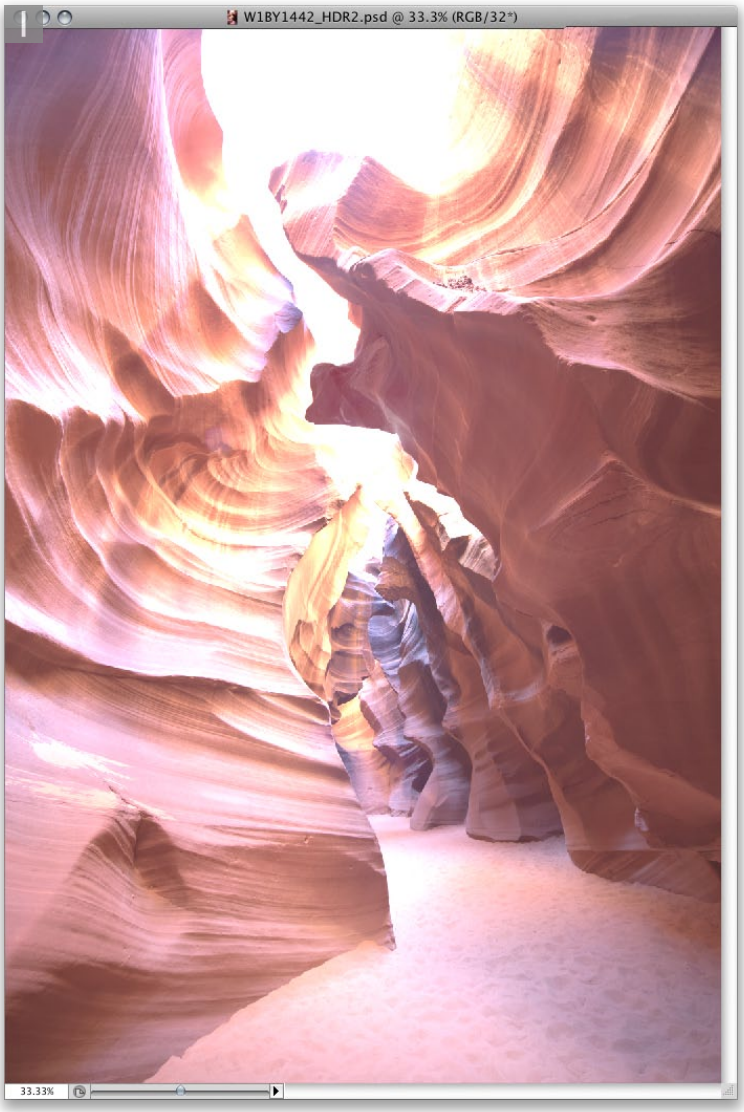

1 I began here with an HDR image that was produced by merging together a bracketed sequence of 3 photographs, shot at 2 EVs apart. To reduce any movement in the individual exposures, the camera was mounted on a tripod and the pictures were taken using a cable release. This first screen shot shows the merged HDR file as it appeared in Photoshop in 32-bit per channel mode with the Exposure slider adjusted so that I got to see a reasonably good view of the HDR image. In 32-bit per channel mode one can use the Exposure slider to preview the HDR image at different levels of brightness. Here, the Exposure slider allowed me to adjust the preview as I focussed on retouching different parts of the picture using the clone stamp tool prior to making the HDR tone conversion.

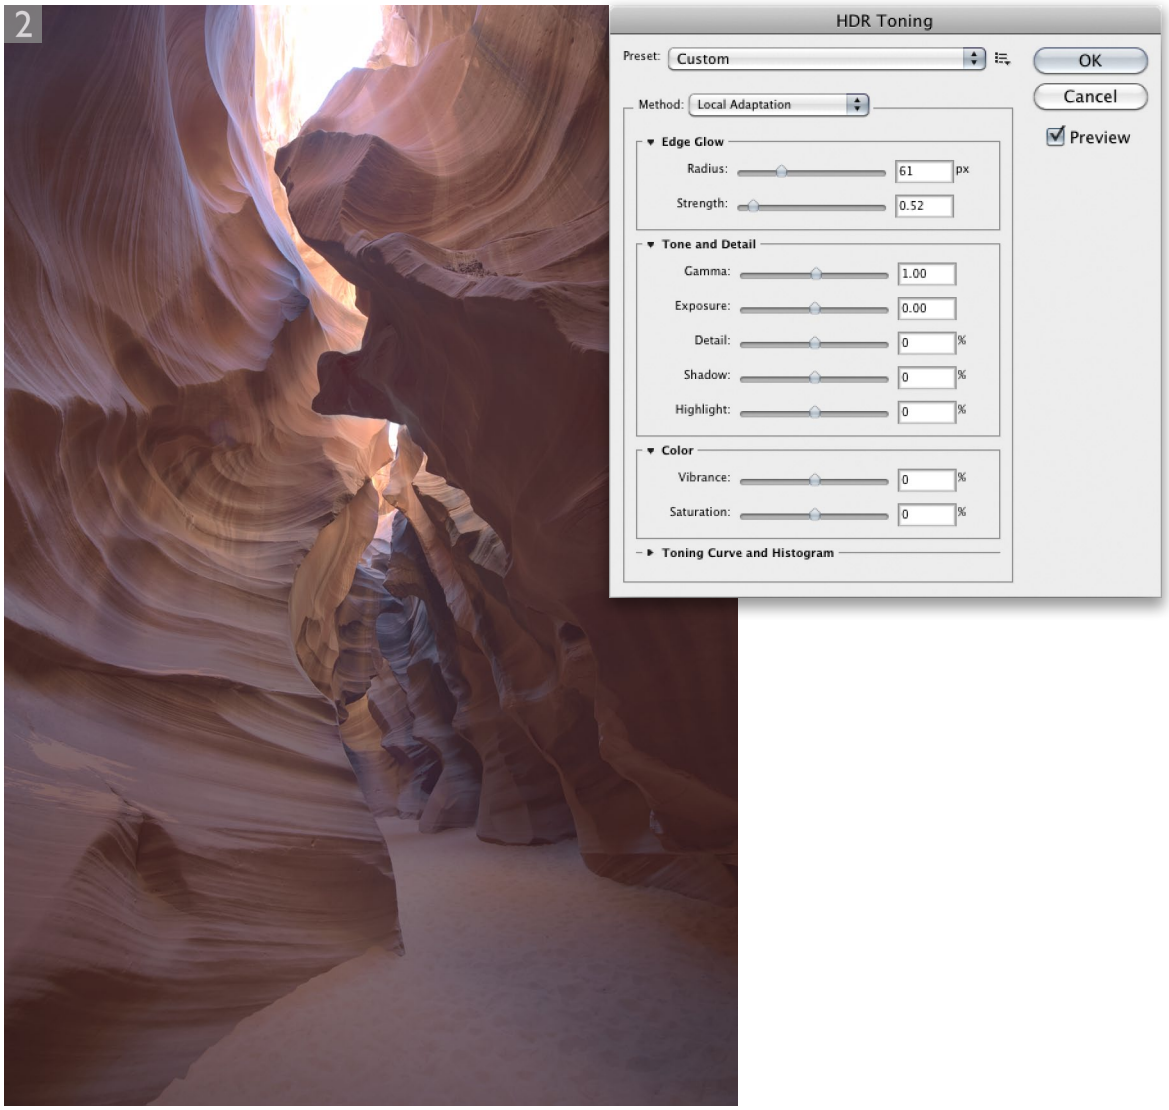

2 To convert this high dynamic range image into a low dynamic range version I had two options. I could go to the Image ⇨ Mode submenu and chose 16-bits / Channel, or choose Image ⇨ HDR Toning… Either of these methods would open the HDR Toning dialog shown here. Of the four tone mapping options that are available in this dialog, I find that the Local Adaptation method usually works the best and in this screen shot I left all the sliders at their default positions. Although the image doesn't look all that great just yet, this is certainly quite an improvement upon how the default HDR converted image preview looked in previous versions of Photoshop.

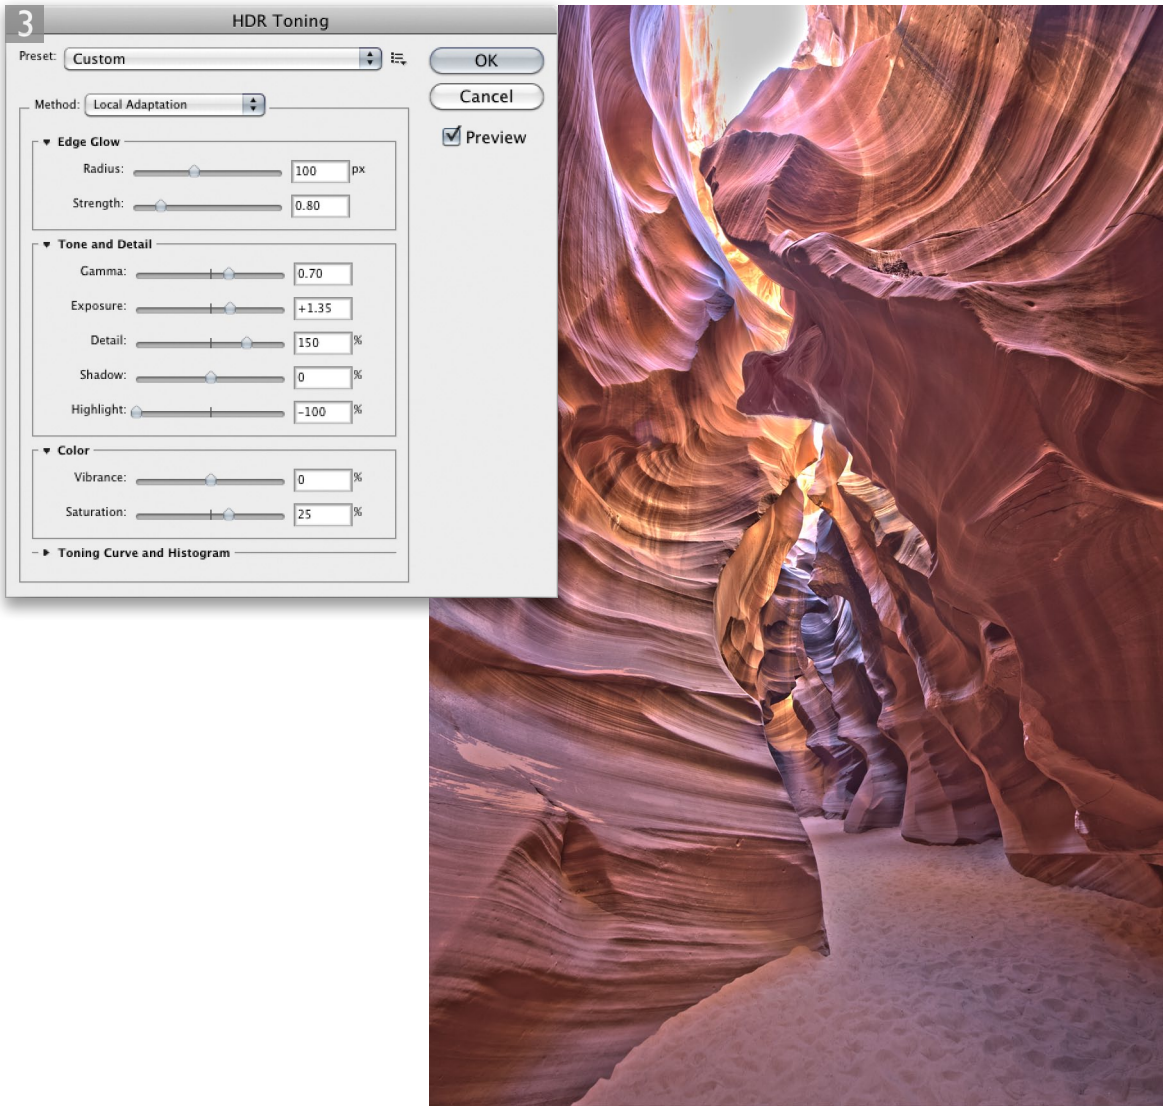

3 Here I adjusted the HDR Toning sliders to produce what might be called the 'illustration look', that is favoured by many HDR photography enthusiasts. If this is the type of effect you are after I don't think the Photoshop HDR Toning adjustment is really as capable as say, Photomatix Pro and nor is it as simple to configure, but Photoshop CS5 can now be made to produce the rather obvious 'HDR toned effect'. Looking at the settings shown here, I set the Radius slider to 100 pixels and raised the Strength to 0.80. I took the Gamma slider to 0.7, set the Exposure slider to +1.35 and the Detail slider to 150%. I then reduced the Highlight slider to -100% and increased the Saturation slightly, setting it to 25%.

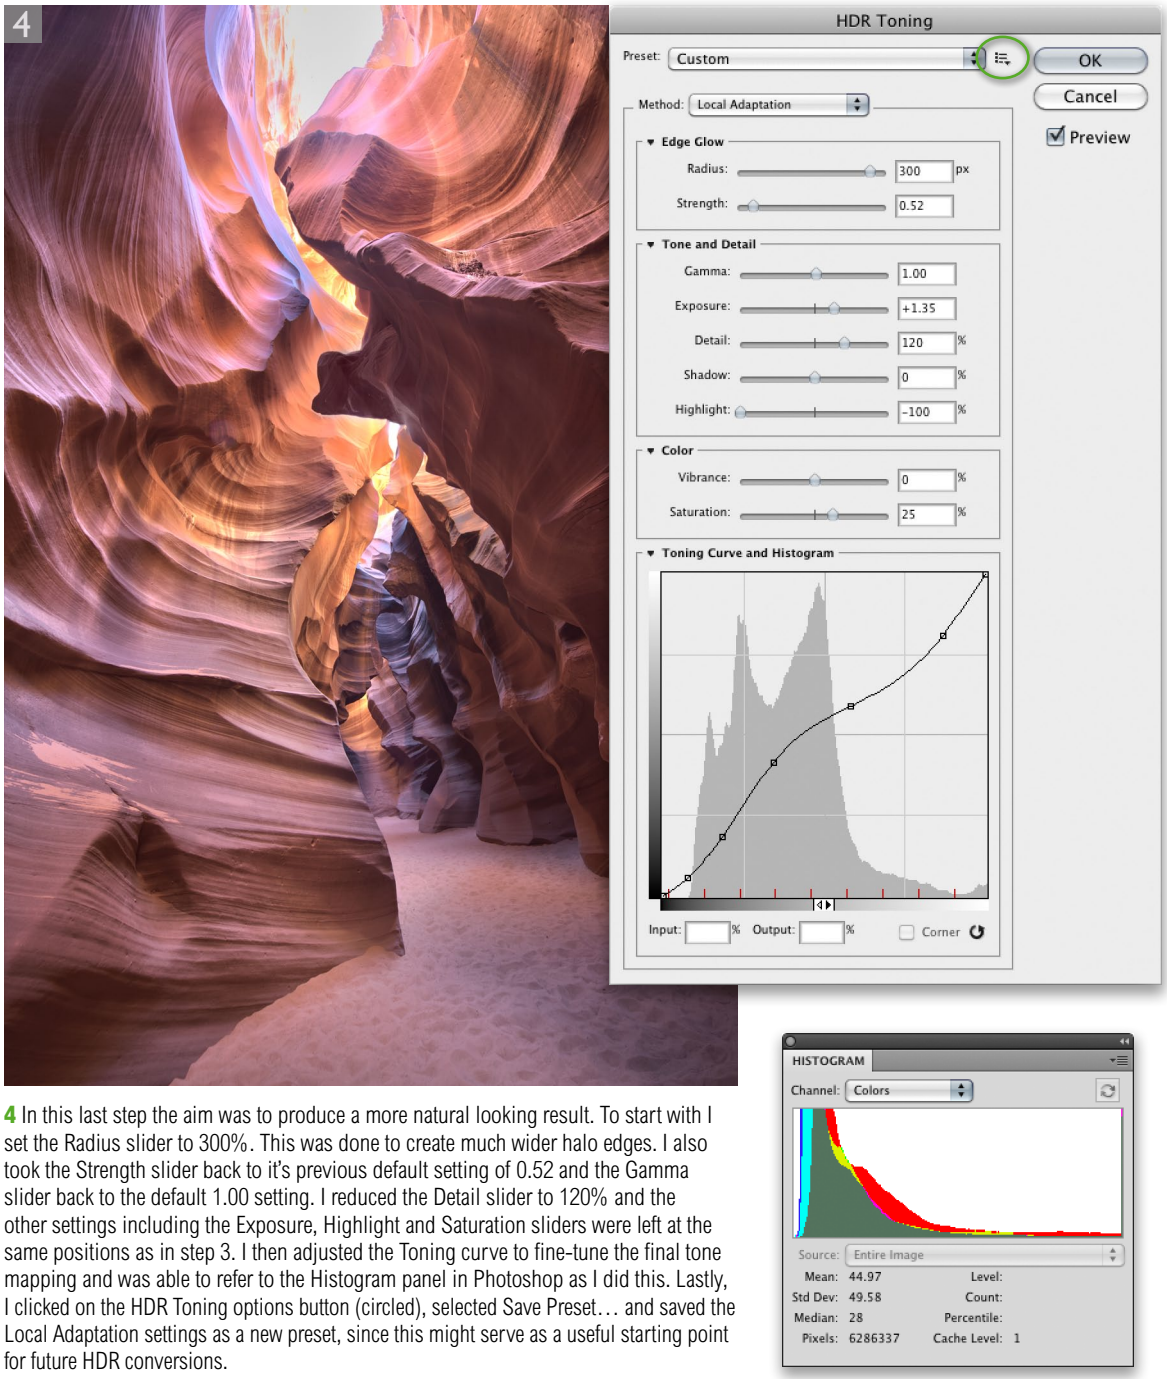

4 In this last step the aim was to produce a more natural looking result. To start with I set the Radius slider to 300%. This was done to create much wider halo edges. I also took the Strength slider back to it's previous default setting of 0.52 and the Gamma slider back to the default 1.00 setting. I reduced the Detail slider to 120% and the other settings including the Exposure, Highlight and Saturation sliders were left at the same positions as in step 3. I then adjusted the Toning curve to fine-tune the final tone mapping and was able to refer to the Histogram panel in Photoshop as I did this. Lastly, I clicked on the HDR Toning options button (circled), selected Save Preset… and saved the Local Adaptation settings as a new preset, since this might serve as a useful starting point for future HDR conversions.

Exposure adjustments

Although you can use the Exposure adjustment to edit 16-bit or 8-bit images, it is really designed as an image adjustment for working on 32-bit images.

Work from the top downwards

When you add a new adjustment layer, it is best to add the newest layer at the bottom of the layer stack just above the Background layer. This reduces the length of time you'll see the 'Building histograms' dialog appear on the screen.

Manual tone mapping

There is yet another way to tone map an HDR image. You can add adjustment layers to a 32-bit image and use these to edit the tone and color. Previously, the ability to paint and add layers was only available in the extended versions of the program, but this is something that you can now do in all versions of Photoshop CS5. In Figure 7.7, you can see a preview of an image that was in 32-bit per channel mode, where I used a combination of Exposure adjustment layers to adjust the contrast and brightness and Levels adjustment layers to adjust the color and masked each of these individual adjustment layers to selectively apply these adjustments to specific areas of the image. By building up a succession of adjustment layers, I was able to tone map the original HDR image just the way I wanted, but without having to render it as an LDR version just yet. If you wish to preserve all the layers when saving such a file, you will need to use the Photoshop, TIFF, Portable Bit Map or PSB format. Preserving the layers means that the master image file size will be increased, but direct HDR image editing does offer a lot of flexibility.

Figure 7.7 You can use adjustment layers to manually tone map the 32-bit data while keeping the HDR image in it's original 32-bit mode.