$$\boxed{\textbf{4}}$$

RESTORATION 2

Using levels to assess images
- White and black points

More on Layers
- Layer modes
- Curves tool
- Sharing layers between images

Channels
- The different types of information
- Viewing Channels
- Using Channels

Masks

Unsharp masking

Using Levels to assess an image

It's not uncommon to come across images that have been lit unevenly and this can make what should otherwise be a relatively levels adjustment somewhat tricky. Remember we're often dealing information which is just discernible from the background. It can also be difficult sometimes to spot exactly where the brightest and darkest parts of such an image are. Create a Levels adjustment layer just above the background, or a duplicate thereof, (yes, you don't have to use an adjustment layer but just in case you click OK...) and, in the Levels dialog box, hold down the `Alt` key while dragging the white level slider to the left, this will show the lightest pixels first. Note that the pixels will show in their colour i.e. if the red channel is the lightest then red pixels will show before green or blue. To find the darkest pixels in an image drag the black slider while holding down the Alt key, this time the darkest colours pixels will show first. This technique shows up uneven illumination fairly well. A similar effect can be had with a Threshold adjustment layer but it ignores the colour information. If need be we can then go on and create a mask/transparency layer based on the luminance of the image (see Channels later).

The image on the next page shows the use of the `Alt` key with the levels sliders to show the lightest and darkest areas of the image.

On the workstation see images <Levels_light_dark.png>, <RGB_difference_1.png> and < RGB_difference_1a.png>.
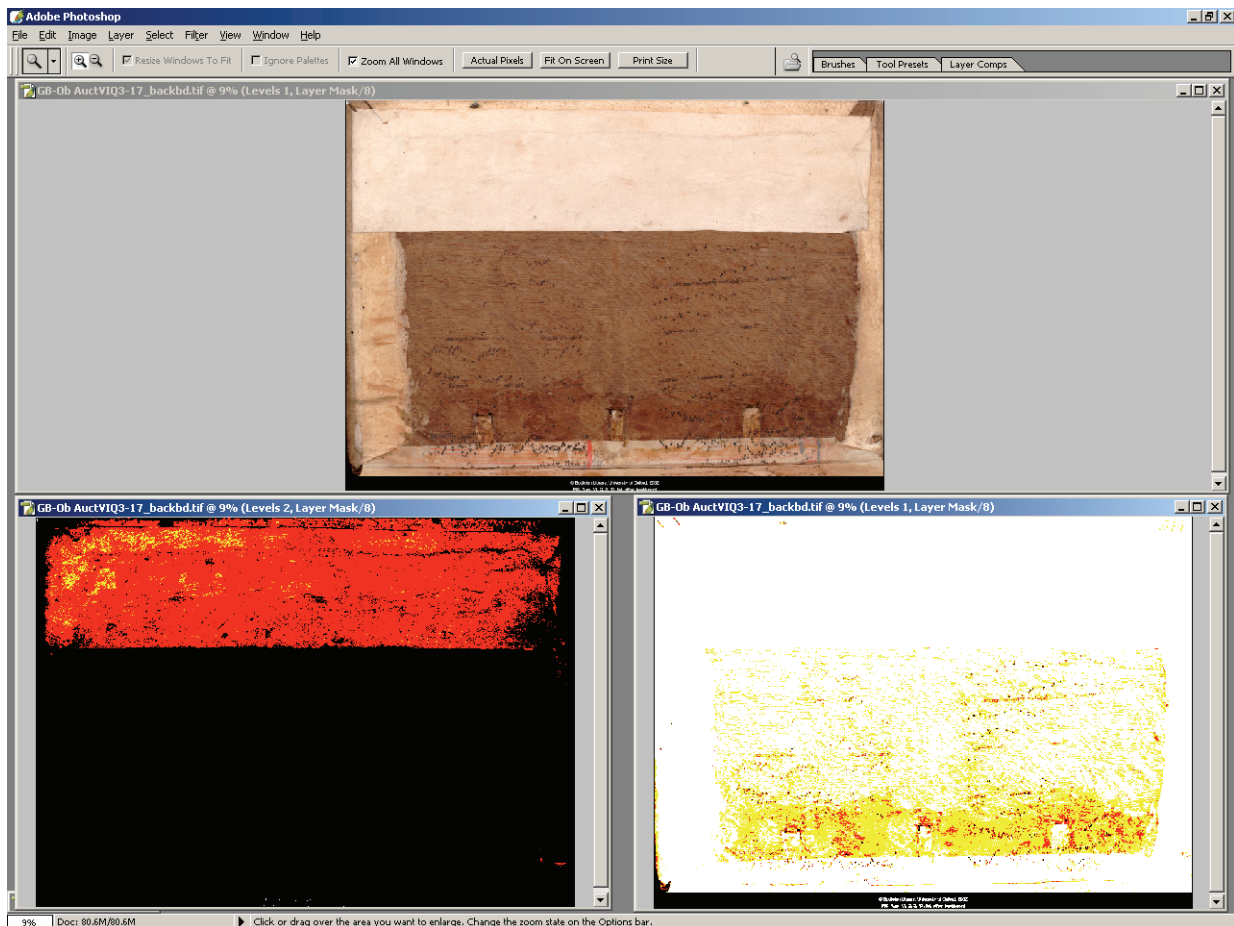
> 👆 Remember that the middle eye dropper in the Levels dialog will adjust whatever colour is clicked on to a neutral grey, so if you have an image that you suspect has a not-too-obvious colour problem you can run a quick check by clicking on the grey on the colour card you shot with the subject (you did, didn't you?) and observing what happens to the overall colour balance. `Cancel` out of the dialog so as not to save any changes, unless you want to of course.

More on Layers

Layer modes

One of the advantages of using layers is that the way in which the layer interacts with the image (and other layers) below it can be altered. The two main ways we are most likely to use in the restoration process are the blending mode and opacity.

The blending modes fall into a number of grouped functions, again we are most likely to use just two of these (though others may well have their part to play, particularly the Contrast group). These two groups are the darkening and lightening modes.

The Darkening modes consist of:

- Darken: lighter pixels are replaced with darker ones when the two images (i.e. current layer and those below) are compared.

- Multiply: darkens the entire image, useful to reinforce faint highlight detail in overexposed images.

- Color Burn: increases contrast and darkens image.

- Linear Burn: combination of Multiply and Color Burn. Dark pixels can be forced to pure black.

The Lightening group consists of:

- Lighten: darker pixels are replaced with lighter pixels in the comparison of the two images.

- Screen: lightens entire image, used to bring out shadow detail in underexposed images.

- Color Dodge: black values are preserved while there is a decrease in contrast in pixels lighter than 50% grey.

- Linear Dodge: combination of Screen and Color Dodge, light areas can be forced to pure white.

As mentioned in Restoration 1 the strength of the effect of the active layer can be adjusted by the Opacity slider in the Layers palette.

☝ Note that the effects of layers is cumulative, if the effect of one adjustment layer isn't strong enough then you can duplicate that layer to increase the effect, repeat as neccesary. If the last duplication is just a touch too strong then use the opacity slider to fine tune the effect. In this way we can get 1.5x, 1.7x, 2.3x or whatever factor we require by a combination of compounding layers and adjusting their opacity.

Open the image <LayerModeMultiply.psd>. As you can see parts of this image are a little overexposed: while we can't restore information that isn't there we can make the best of what we have using another layer with blending mode set to Multiply. There are other 'issues' with this image that we may be able to deal with if time permits.

Open the file <LayerModeScreen.psd>. You will notice that there are two colour samplers placed on the image (yes, Photoshop very usefully saves the samplers in its PSD files). The first sampler is in the upper left corner of the page, the second is on the upper half of the white paper. Make sure both adjustment layers are turned off (they should be already, this is something else that Photoshop saves the state of). Now make the 'Levels adjusted' layer visible, try and assess the image and make a note of the two sampler values. Turn this layer off again. Repeat the exercise with the 'Screen mode' layer. Do you notice any difference?

💣 A 'gotcha': when using some tools (e.g. Magic wand, Paint bucket, Blur and Sharpen) on images with layers a little check box appears on the options bar for 'Use all layers'. What this does is pretty much what is says but it can all too easily lead to a tool not working as you expect, common problems are:

- you have just created a new layer and you click with the Magic wand to make a selection only to have the whole image selected. The tool is working correctly it just that the 'Use all layers' box is unchecked – it only takes the sample from the from the current layer. To correct this just check the box.

- you want to blur background detail, you have already isolated the foreground objects on a separate layer, you have even remembered to select the background layer (copy, as it's a destructive process) but for some odd reason the tool is blurring your foreground objects too – in this case the 'Use all layers' box is checked when you don't want it to be.

Once you get used to the power of layers it is all to easy to build up quite a collection in an image. This is fine, if you have the RAM!, but can lead to a confusion of layers ("Now just what does this one do?"). There are a couple of things we can do to help reduce this confusion. One is something you will have seen this in most of the example images and that is layer naming – the other is layer sets.

Layers are given a default name by Photoshop e.g. Layer 1; Levels 4; Curves 3 etc. to rename the layer to something more descriptive you can: right-click on the layer (PC) or double click on the layer name (Mac) or go to the palette options button ⬛ (top right of the palette well) and choose Layer Properties, this will bring up a simple dialog box where you can edit the name.

When you have more than a few layers it can be useful to organise them into groups related by function: this is where 'layer sets' comes in. To use this feature create a new set by clicking on the 'Create Layer set' ⬛ button at the bottom of the layers palette, the new set will have a default name (starting at 'Set 1') which you may wish to change, the process mirrors that of renaming layers. You can now create new layers within that set by selecting the set before creating the layer. Existing layers may be added to the set simply by dragging and dropping (note that this appears not to work with an empty set for some, no doubt logical, reason). Quite often this leads to the dropped layer not ending up in the right position in the set but, just like ordinary layers, they may be dragged to a new position. Sets can also be created in, or dragged to, another set. A very useful feature of sets is that the whole group of layers within it can be disabled by clicking on the 'eye' 👁 icon to the left of the set. Oh, and to tidy up the palette sets can be 'collapsed' so that only the set name is showing by clicking on the triangle icon to the left of the set name, clicking it again opens the set up.


Curves tool
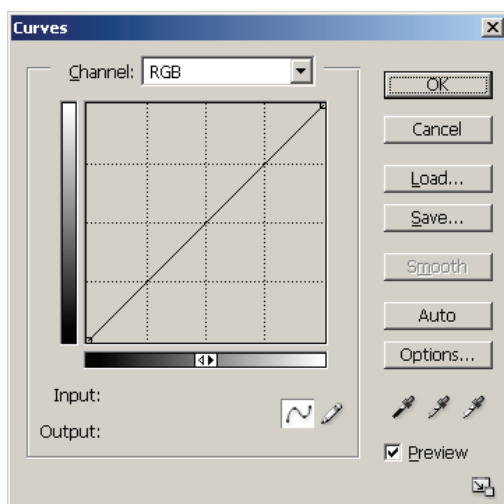
This tool, like Levels, is used to re-map pixel values to new ones. The main difference is that where the Levels dialog has three sliders (plus input and output levels) the Curves tool can have up to nine sliding points. This does make it a more complex tool to get to grips with, the bonus for doing this is the flexibility it gives you. Unfortunately the dialog box is really too small for such a tool but this gives us an opportunity to look at how keyboard control can make things much easier. But first off we'll take a quick look at the dialog box and *some* of its settings.

The image here is the Curves dialog as it first opens. Note that the 'curve' is a straight line, at the moment there will be no effect on the image (you can easily see effects if you ensure that the 'Preview' box is checked). Most of the controls are similar to the ones you are familiar with in the levels control. The Options button leads on to a sub-menu where tool defaults can be set (and reset) and other options such as snapping neutral midtones, clipping levels and so on. The Auto button uses soem of these settings to make a 'best guess' adjustment to the image. it may give you what your after but it is much more likely that you'll will need to take manual control.

In this image there are no numbers by the Input and Output labels as there is no control point for them to echo the values of (excepting the points at the end of the curve).
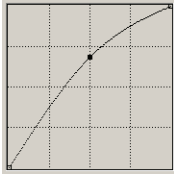
Clicking anywhere in the central grid places a control point on it that pulls the line to it. In the example on the left the midtone values are made lighter, the effect here is like moving the midtone slider in the levels control to the left – so far no major advantage.

So let's take a look at how this tool might offer benefits over levels...
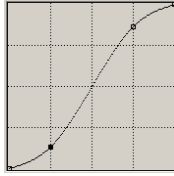
Here a second control point has been added and the first one moved up and to the right. Now we have a curve! This curve will have the effect of: lightening and compressing highlights, darkening and compressing shadows while enhancing midtone contrast. Since we can have up to nine control points the affected the potential for both finely-tuned correction and strange effects is high (try moving a control point significantly out of position). These points can be on a per-channel basis not just RGB.

Note that the active control is denoted by being solid black instead of an outline.

This tool allows you to bring out detail in narrow sections of the tonal range in a way that levels can't manage, though it does make it trickier to use (think multiple gamma controls). Part of its trickiness comes from the control points not being very easy to control(!) with a mouse so we'll have a look at how the keyboard can make this easier...

To start with there is accurate placement of the intial point, if you're not quite sure where in the tonal range to place it you can click on the eyedropper tool (it's in the options bar for this tool) and click somewhere in the image that contains a tone you want, you'll notice that a control point appears on the curve at the tonal point you have just sampled, release the click and it disappears – not neccessarily what you want. To fix the point on the curve press `Ctrl/Cmd` while you click.

Once you have the point(s) you need then moving it with the mouse gives only fairly coarse control, to make things easier click on the point you wish to move, now use the `arrow` keys on your keyboard to move the point in the direction you want to move it. Note that the numbers by Input and Output change by one for each press, using the Shift in conjunction accelerates the change to 15.

☞   These number-change and 'accellerator' keystrokes work in the same way in any number box, though the accelerator is usually 10. It is often a much easier and more reproducible way to change values than using the mouse, another advantage is that it's easier to keep your eye on the image and not get led astray by histograms etc.

The three eyedroppers on the bottom right of the dialog box set black, grey and white points respectively. The grey point eyedropper can be very useful in neutralising colour casts: click on the Set Grey Point eyedropper to activate and then click on a point in the image that contains a known neutral grey (preferably middle tonal range), depending on how well you chose the reference grey the colour cast should be removed.

There is a lot more to the Curves tool but hopefully, especially if you've not felt comfortable with it before, this short introduction will encourage you to experiment with it.

Sharing Layers between images

It quite often happens that when there are problems with an image, other images shot at the same time have the same (or at least similar) problems. If this is the case then try and pick the most representative of the group and work on getting that right. Once you have the adjustment layer(s) providing the corrections you're after these can be easily shared with the other images – saving a lot of duplication of work.

To share the layer(s) have both the image containing the layers to share (source) and the destination image open, make sure that both are visible and, with the source active, then quite simply drag the required layer(s) from the Layer palette onto the destination *image*. That's it. If the effect isn't quite right it may be that adjusting the layer opacity (and/or duplicating the layer) will give sufficient fine tuning.

The same process can be applied to layer sets too.

☝ Note that adjustment layers are scaleable, they don't care about differences in the aspect ratio or size of the two images. This lends itself to speeding up work on big images: downsample (a copy of!) the image and work on this, note that this will be for overall adjustments not fine detail. Because the memory requirements can be drastically reduced the processing can run somewhat faster, especially if there are lots of layers. When you've finished, save the image for safekeeping, and just copy the layer(s) across to the full size version as above.

Channels

To quote Adobe: "Channels are greyscale images that store different types of information."

What are these different types of information and how can we make use of them?

Starting with almost the simplest image: the single channel just contains the greyscale data for that image. If we move on to basic RGB colour images then the 4 channels 'created' contain the values of each of the colour channels, 3 of them show the greyscale version of either Red or Green or Blue data, the fourth channel shows the combined (or composite) data and is shown in colour.

Note that for CMYK images an extra channel is needed making a minimum of five.

Another type of channel is the alpha channel, these allow you to create, manipulate, save and load masks. These are used for protecting areas from, or enabling areas for, edits. They also are used to create transparency effects by allowing pixels to show through or not according to the pixel values in the alpha channel. Note that alpha channels are greyscales, this means that the 'strength' of a particular effect is applied according to these greyscale values (black means no effect, white means full effect).

Viewing channels:

☝ You will need to have an image open in order to do this.

Open the channels palette by either clicking on the Channels tab in the layer palette if open or `Window > Channels`, if you prefer to have both the Layers and Channels visible at the same time, then you can click-drag the Channels palette out of the Layers palette. Assuming a simple RGB image you can click on the eye icon to the left of the channel (this works in the same way as the eye icon in the layers palette) to view that channel's information – note they are greyscales. You can also click on the R, G and B channels in turn (after first clicking on RGB if its not already active), to view another channel in combination with the first, say R and B, then hold down `shift` while clicking on the second channel. To view all channels in combination just click on the RGB channel. As usual there are useful keyboard shortcuts to change channels: `Ctrl/Cmd-~` shows all channels, `Ctrl/Cmd-1` shows Red, `Ctrl/Cmd-2` shows Green and `Ctrl/Cmd-3` shows the Blue

channel. Using the keyboard in this way makes it much easier to observe any differences on the image between channels, as your eyes can more easily stay on the subject.

☞ The `New Window` for <image name> option on the `Window` menu allows you have more than one view of an image, these may be: two (or more, this applies to all the following) different zooms; two areas of the image; two individual or combined channels and so on.

These views may be synchronised for zoom (not relevant in first example!) and/or location. Go to `Window > Arrange` and choose which option you want.

How can we use Channels?

Sometimes image information can be seen more readily in single channel view than in the RGB in combination. For example if a dark stain is predominantly a mixture of red and green then the blue channel may show objects (or show them more clearly) that are obscured by that stain in RGB (see the example on the next page: compare the RGB and B channels in the stained area). It may then be easier to use this channel for making selections than the combined view. Also more advanced manipulations (easier in later versions of Photoshop) can be used to extract the data from one channel and use it to 'rebuild' data in other channel(s) – a sort of cross-channel cloning.
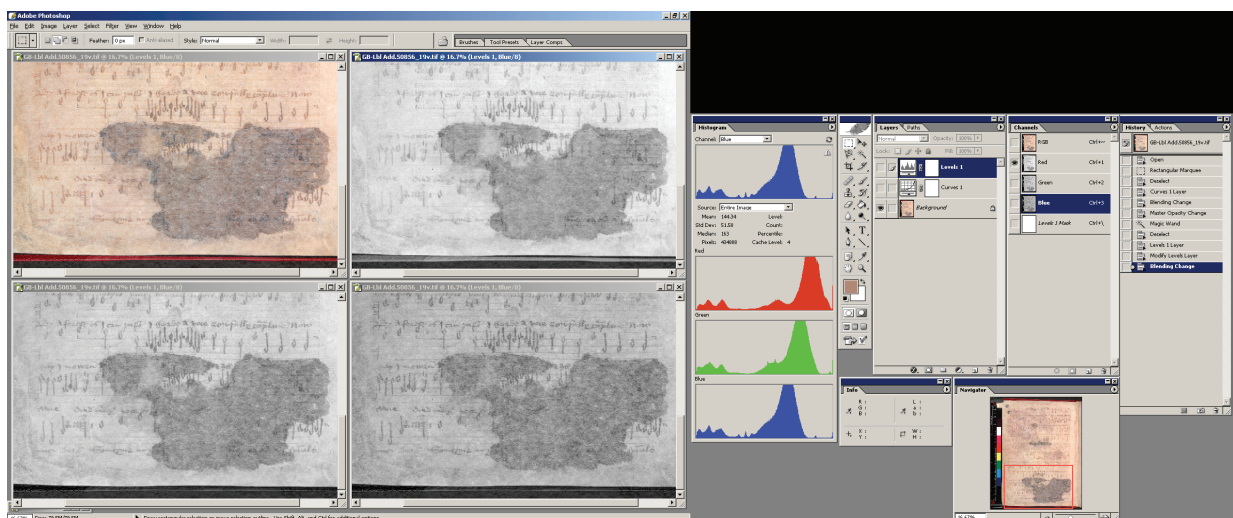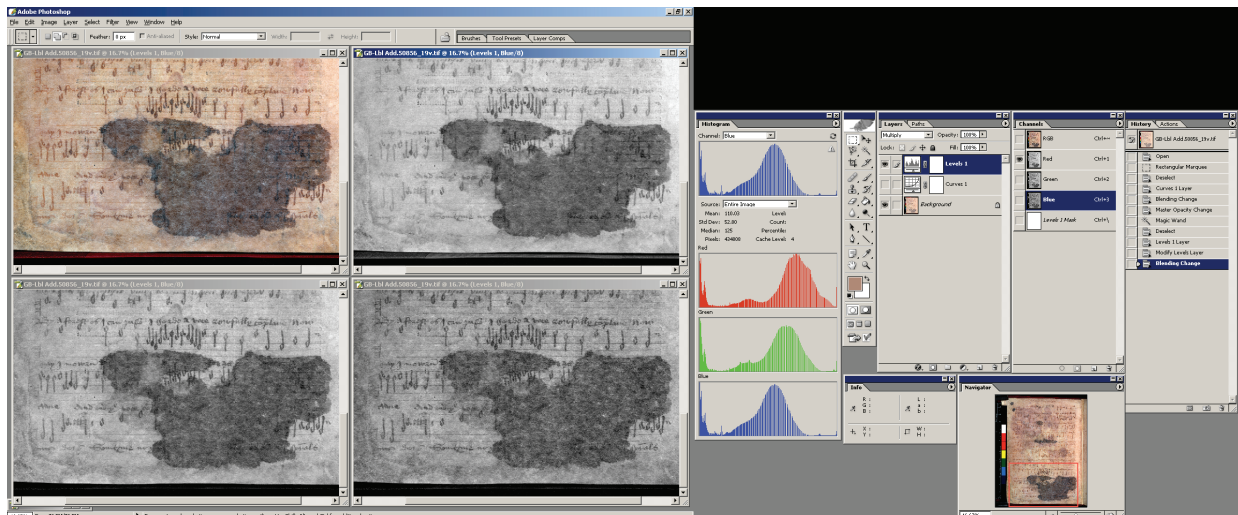

RGB_Difference_1

This is the image as it was shot – the adjustment layers are turned off.

Note the overall shape and positon of the histograms (this is Photoshop CS so the R, G, B are viewable separately.

The dual screen setup allows the tools and palettes to be placed on the second screen allowing more room for the image itself. The dark area at top right is due the second monitor running at a lower resolution.

> **RGB_Difference_1a**
>
> Now a levels adjustment layer has been turned on: the black point has been moved up in order to darken the image.
>
> Look at the histograms again, note the gaps that have appeared due to tonal stretching.
>
> The four views of the image were created by menus:
>
> `Window > Arrange > New Window for <image>`, then
>
> `Window > Arrange > Tile` , clicking on the appropriate channel and finally
>
> `Window > Arrange > Match Zoom and Location`.

A number of filters are also quite happy to operate on single or double channels. For instance, consumer digital cameras are particularly prone to noise (graininess) in the blue channel and one of the ways of ameliorating this is to blur the blue data – this is easy to do, select the blue channel (and thereby disable the red and green) and `Filter > Blur`, or for more control `Filter > Gaussian Blur`.

> ☞ The blue channel rarely contains *detail* that we can use as the eye does not respond well to these wavelengths (and not at all in the very centre of the eye). This is made use of in lossy image compression algorithms such as JPEG, PhotoCD etc. where the blue information can be discarded to a much greater extent than either the red or green data

Masks

Masks, as their name implies, protect (or enable) areas of an image from (or for) editing and viewing. As they are greyscale images they can be used to alter the strength of the effect of other tools: black means no effect, white means full and mid-grey somewhere in between. They can be created in a number of ways:

- Using `Quickmask` 🔲 (found on the main Tools palette) after making a selection with any selection tool

- Creating a mask from a Threshold adjustment layer

- Painting directly onto an alpha channel (click the `New Channel` 🔲 button in the Channels palette)

Selections were covered earlier in this workbook; couple them with channels and the power and versatility increase enormously. Selections can be made into masks, masks into selections. To save a selection: click on the `Save Selection` 🔲 button in the Channels palette. Conversely to load a mask (alpha channel) as selection: click on the channel you wish to load then click on the `Load Selection` ⬚ button.

As has already been mentioned these are greyscale images, that means they can be operated on by any tool that operates on greyscale images: sharpen; blur; add or remove noise; curves; levels and so on.

Painting on masks is an easy way to refine selections, again there are a number of approaches to this. The easiest is to use the Quickmask tool:

- Create a selection using an appropriate selection tool

- Click on Quickmask 🔲 (press Q), a rubylith overlay will appear (its colour and transparency can be changed).

- Paint with any of the painting tools to either add to the mask (black), subtract from the mask (white) or a mid-grey to alter transparency.

- To return to Standard mode editing click on the `Standard mode` 🔲 button (or press Q). Note how the mask has become a selection again with the edits incorporated.

Having refined your mask you may want to save it as a Channel Mask so that it is not lost and is available for further editing.

☞   It is a matter of personal preference but for a number of people graphics tablets make certain aspects of image work very much easier. They are much more 'natural' to use for painting, drawing, outlining (as with the Magnetic Lasso) and other such tasks than a mouse. In addition, most are pressure sensitive and can therefore alter some aspect of the tool's behaviour as the user changes pressure – with the brush tool, for example, the most common is to increase the width, increase the saturation of a colour or a combination of the two.

Masks can interact with another mask/selection. For example we may have may a selection that has very distinct borders (e.g. the edge of a page) that we want to mask from the same set of operations as another area that needs the the boundary between it and the rest of the page softened. In this case we would create both selections separately with the appropriate tools, save them as masks and then combine them.

Adding these two masks together load one as a selection (click `Load Selection` button in a channels palette), hold down `Ctrl+Shift` (PC) or `Cmd+Shift` (Mac) and click the channel you wish to add. Of course this new selection can be saved as a mask too (there is a limit of 56 channels). Mask combination is not restricted to addition, they can be subtracted and intersected as with the selection tools themselves.

Layer Masks are masks associated with a layer and are used to restrict the effects of an adjustment layer to areas where the mask is transparent.

☞   Note that a Layer Mask may not show in the Channels palette unless either the mask already existed as a Channel or the layer with which the mask is associated is active.

There is a lot more to channels, masks and layer masks than this workbook attmpts to cover. For some idea of their power and flexibility open an image containing an adjustment layer, make a selection (anything will do as you won't be saving the image) then `right-click` (PC) or `alt-Click` (Mac) on, firstly, the Layer Mask Thumbnail and look at the

range of options; secondly `right-click` (PC)/`alt-Click` (Mac) on the Layer Thumbnail and note the range of options here. Close the image or save under another name and start exploring the range of options available.

☞   A Layer Mask can be temporarily disabled by `shift-clicking` (PC/Mac) on its icon in the Layer palette. `Shift-clicking` again will re-enable it.

As mentioned earlier in the section on layers we can usefully create a mask based on the luminance of an image, this can be used as a way of dealing with uneven illumination problems. To create such a mask click on the RGB Channel in the Channels palette or press `Ctrl/Cmd+~` then hold down the `Ctrl/Cmd` key and click the RGB channel in the Channels palette. You now have a luminance-based selection which can be saved as a luminance mask, used with an adjustment layer (or both.) Note that this mask can be manipulated like any mask to enhance, diminish or otherwise vary its effects.

Unsharp Masking

This subject has intentionally been left until last – it is the final operation you should to an image (other than save the file). There are two main reasons for this: firstly, you won't know how much sharpening you will need until you know the output size and destination (printer or screen); secondly, and more importantly, the process is destructive.

Photoshop has a few sharpening tools, they're found under the Filters menu (`Filters > Sharpen`), they consist of

- Sharpen,
- Sharpen Edges,
- Sharpen More,
- Unsharp Mask.

We'll ignore the first three as they are blunt instruments with no user settable parameters. The Unsharp Mask however does have settings that the user can vary (as indicated by the ellipses after the menu entry). The settings are:

- Amount: the extent to which the image will be sharpened, too high a value leads to halo effects around edges.
- Radius: the 'area' that will be sampled to calculate the sharpening effect. Setting this too high will also make the sharpening obvious.
- Threshold: the difference in value that adjacent pixels need to be before they are seen as an edge. Use higher values on noisy images so that the noise does not get emphasised by the sharpening process.

The Unsharp Mask only works on a single layer at a time so you may need to merge layers (or flatten the image) to achieve the effect you're after – of course it may be that you only want to sharpen selective bits of the image. You can also do this by creating an appropriate selection and using this to define the area to which the effect of the filter should be limited. It is also possible to apply the Unsharp Mask to selected channels, so we could sharpen just the red and green channels while leaving the blue channel as it is (which, as you may remember, tends to be noisier).

A couple of 'general rules':

- images for print will need more sharpening than those destined for screen use. The only real test here is to make proof prints on the printer the image is to be printed on – remember to turn off print driver enhancements as these can mask the changes you are trying to assess.

- the more an image has been reduced in size the more sharpening it will need to give the *appearance* of detail.

you need to view the image at 100% to see the effects, any lower zoom factor will almost certainly lead to the image being oversharpened.

The easiest way to learn about the interaction of the settings is to adopt the process of changing just one parameter at a time, pushing it to extremes (use the Preview option if your computer can update fast enough), then change parameter and repeat.