15

Resolving the Resolution Issue

Many very different types of resolution are expressed by the same ambiguous acronym. At times, a high resolution is necessary. At others, it's a waste of disk space and computing power, if not an outright quality-killer.

Graphic arts is notorious for its ambiguous terminology. *Trap* can mean a prepress technique or a pressroom anomaly. *PMT* describes a variety of scanner or a type of positive proof. *Shadow* means one thing to a photographer and another to a retoucher. A typeface called *Gothic* is a bold sans-serif—unless it is a medieval blackletter, such as Old English. *Web,* to a printer, is quite a different animal from the worldwide variety. Even *red* means something different to a pressman than it does to the rest of the world.

But of all the semantic snares set for the unwary, the most insidious has to be the innocent-sounding *dpi*, which plagues us throughout our process and sometimes is used as a synonym for a second bugaboo, resolution.

Consider all the different things that DPI connotes, and understand why novices—and some experts—get confused. 300 DPI may be a type of scan or a laser printer. 2,400 DPI may denote a different type of scan or an imagesetter. 72 DPI may measure the number of phosphors in a monitor or a newspaper's screen ruling.

In choosing the various flavors of resolution for a certain job, it's silly to assume that bigger is better. Unnecessarily high resolutions, at best,

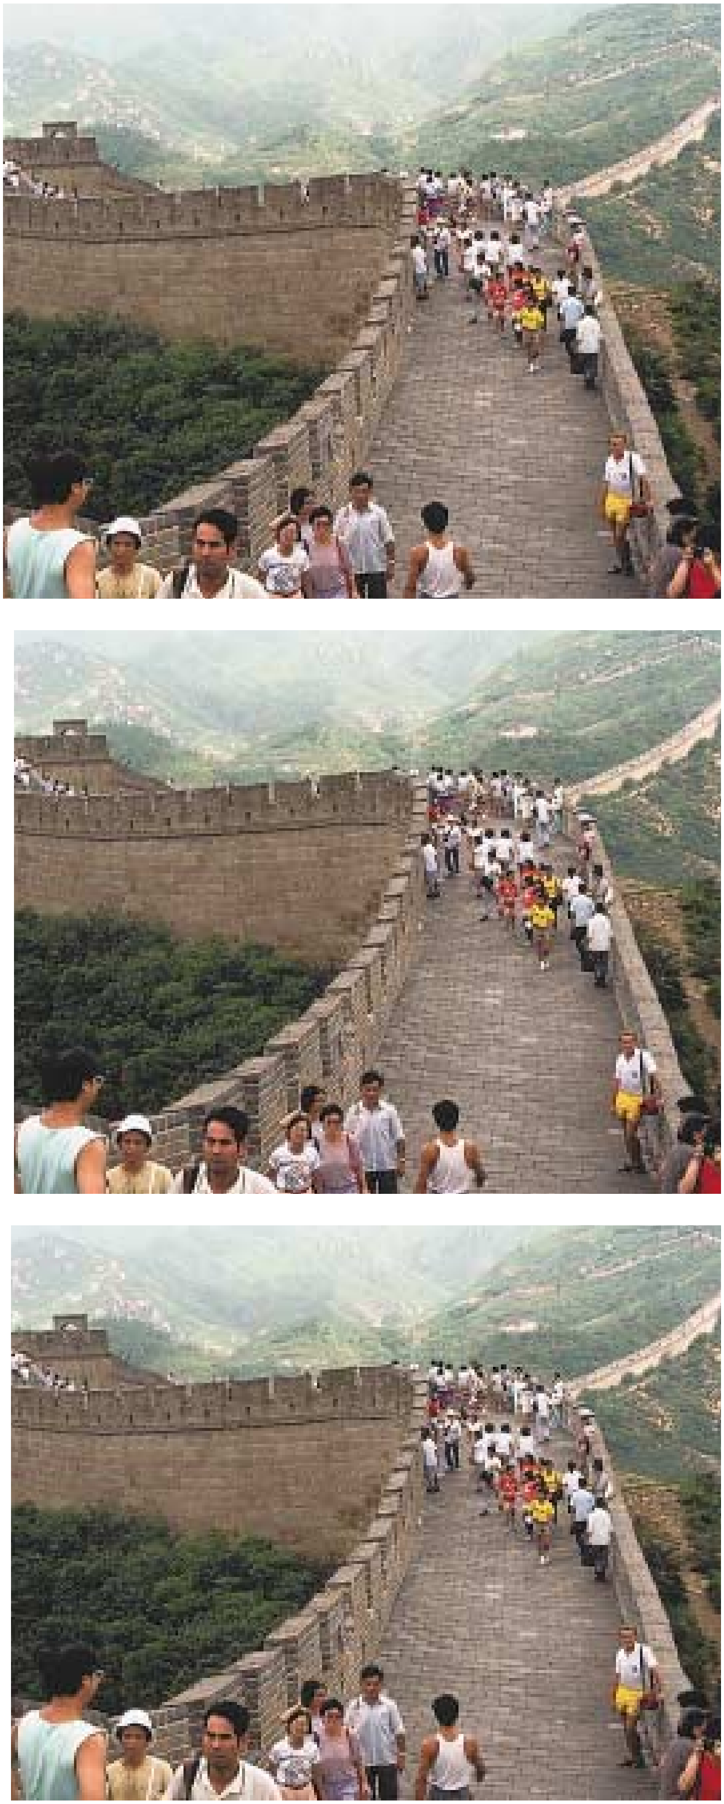

eat storage space and bog down networks. If we are not so lucky, they may bring a RIP to its knees, or worse, produce poorer quality than if we had used the proper resolutions to begin with.

And what are these correct resolutions? They depend on the job, but they also depend upon each other.

First of all, even defining *resolution* is not that easy. It means, more or less, how far apart the smallest distinct parts of the subject of discussion are. Frequently, these small parts are all of the same size, as in the individual pixels of a Photoshop file. But in nondigital parts of the process, they aren't, as in the case of film grain, which is, one might say, the resolution of film, or the halftone dot, which might be termed the resolution of a printing press.

Figure 15.1 suggests what happens when we print without enough press "resolution" —in other words, with a halftone screen that is too coarse. The grainy-looking center image is more appropriate for newspapers than for a book. The smaller the dots, the less obtrusive they are, and the more the final product looks like the original photograph it is supposed to recall.

Anyone who thinks that if a fine screen is good, then a finer one must be better, is a moron. The finer the screen ruling, the smaller the dots, but the smaller the dots, the harder they are to print properly. If they are on the cusp of what the press can tolerate, the following irritating things happen.

- Darker areas start to plug up, resulting in a perceived lower maximum shadow.
- The minimum acceptable highlight dot goes up; at some point, a dot simply gets

Figure 15.1 *Top, an image printed at press resolution (oops, screen ruling) of 150 dots per inch; center, at 65 DPI; and bottom, at 300 DPI.*

too tiny for the plate and the blanket to hold. Overall, detail in the highlight will become inconsistent.

- The image will begin to appear soft as transition areas become less distinct.
- Dot gain will appear to increase.

Now, at what point does all this unpleasantness start to kick in? In newspaper printing, as a rule, it happens at about 100, so most newspapers print with an 85-line screen, and 65-line is not uncommon. Some, however, do use a 100-line screen, and I know of at least one that has successfully used 120 lines.

As paper starts to get a little better and as we migrate to commercial presses, the tolerance goes up. Reasonable uncoated papers can easily hold a 120-line screen, and there can be some success with 133, especially if the printing is direct-to-plate. Magazines that use coated paper generally use 133, and some try 150.

High-quality commercial printing, such as annual reports, uses more expensive coated papers and usually is done at 175 lines, sometimes even 200. And waterless offset, a relatively new approach, appears to make it much easier to hold small dots. There have been reports, which I can't vouch for, of success with screens of over 1,000 lines.

The sad truth is that printers often overstate their own capabilities. In the last chapter, Figure 14.2 offered a closeup of what happens when the screen is too fine for the press. Many printing firms are in fact able to handle 175-line screens. Many others claim to be able to do so and produce muddier results than they would with 150.

Many people get fooled by their contract proofs, which are far easier to control than a press. The bottom version of Figure 15.1 looks just fine on mine. On press, assuming I can slip it by the printer's preflight department, I predict it will look like a cartoon: all brilliant colors. Where an ink is heavy, it will in fact print as solid, since the press won't be able to maintain dot integrity. The weakest colors, which would temper the brightness of these colors, will be missing altogether, as the tiny dots that theoretically compose them blow away.

This will be the first of several examples of how too much resolution can be harmful. Most people assume that the reason an excessive resolution would be counterproductive pertains exclusively to the press.

It doesn't. All resolutions depend upon one another. Too fine a screen may cause more quality problems with the imagesetter than the press.

Out of Spots, Dots

The term DPI stands for dots per inch. In printing, that's just what we are describing, dots. Infuriatingly, this is the one instance where the term DPI is *not* frequently used; most reserve it to describe situations in which it would be just as accurate to say bananas per inch as dots. But in talking about presswork, people don't say 65 DPI, but a 65-line screen, and they abbreviate it as 65 LPI. Yet we have dots, not lines.

To find out how those dots get there, we need to discuss another kind of resolution.

Whether from a $500,000 platesetter or a $300 inkjet printer, the halftone dots we've been talking about are made up of tinier dots. How tiny those tiny dots are governs how effectively the bigger dots can be drawn. The size of those tiny dots, which I will henceforth refer to as *spots,* represents the resolution of the device.

We've barely begun, and already the

terminology is tripping us up. This will never do. In the first edition of this book, after griping about it over several paragraphs, I caved in to the conventional and used DPI indiscriminately. No more. Since 1998, I've taken a hard line. Where appropriate acronyms didn't exist, I invented them. And, as you can see, I'm using SPECIAL TYPE to set them off.

I, therefore, use DPI to describe a dots-per-inch screen ruling, but never to describe measurements that don't involve dots. The rest of the world can use sloppy terminology and be damned. You don't like it, buy a different book.

End of rant. Imagesetters and platesetters made in the United States usually have a resolution of at least 2,400 DPI—oops, 2,400 *spots* per inch. If made elsewhere, the usual resolution is at least 2,540, which happens to be 100 spots per millimeter. Note that spots, unlike dots, never vary in size. They are either off or on; they're either there, or they're not.

The spots are too small for most of us to see, which may make them just about the right size to construct halftone dots with, as Figure 15.2 shows. But it may not: the size relationship between spot and dot is critical, and if it's out of whack, quality will suffer. To know what "out of whack" means, we have to consider yet another species of resolution, the ability of the human eye to resolve differences in color and tonality.

Nobody really knows how great that ability is. Some reputable sources have suggested that typical humans can only perceive around 2,500 different shades of color. On the opposite extreme are folks like myself who say that some individuals are capable of differentiating a million or more.

For realism, whatever printing method we choose should be able to

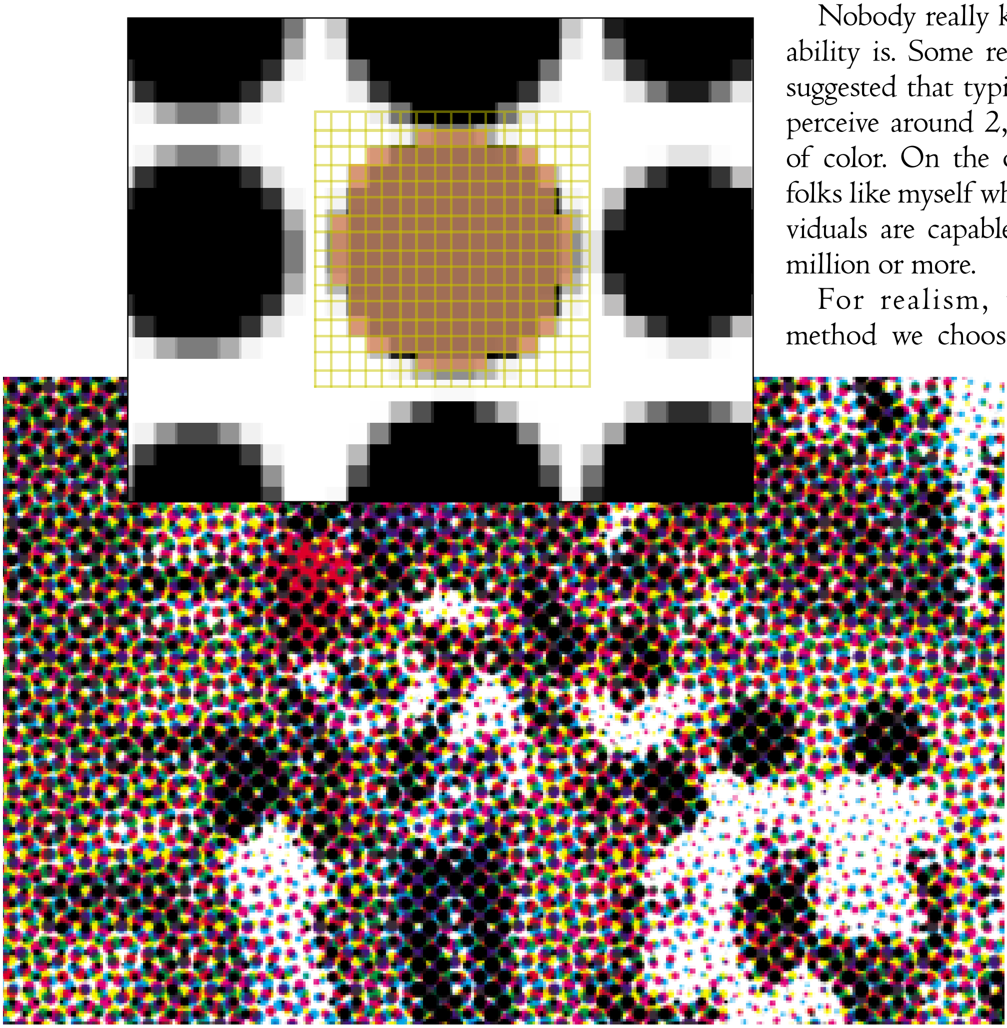

***Figure 15.2** Dots and spots: Below, a blowup of Figure 14.1 shows its dot structure clearly. The imagesetter constructs each dot, inset, by turning spots off and on in a grid. Here, the grid has 256 such spots available for each dot, which is just right—in theory.*

portray at least as many colors as typical humans can perceive, and preferably quite a few more, in case we start moving things around with curves. We need to be able to make very fine adjustments, therefore, in the size of the halftone dot. The smaller the imagesetter spot is, the more flexibility we'll have. On the other hand, constructing dots out of spots is not a trivial calculation. Having spots that are too small will snarl the most powerful RIP.

It certainly makes sense to have at least 200 different sizes of dot available, probably more. If the halftone dots are at 150 DPI and the imagesetter resolution is 2,400 SPI, this will be possible.

150 is 1⁄16 of 2,400. The spots that the imagesetter can potentially paint will be rectangles where each side is 1⁄2400 of an inch, or .00042". Those rectangles will be exactly 1⁄16 of the maximum width of the halftone dot. There will be 16 rows of rectangles across and 16 columns down, for a total of 256 rectangles. Depending upon how many of these actually get painted, there are 256 possible darknesses of the halftone dot, or 257, if you count zero.

256, is, coincidentally, a key number in yet another kind of resolution.

The Blind Man's Eyeglasses

An original photograph is said, rather inaccurately, to be *continuous-tone.* Digital files aren't. They can portray only a limited number of varieties of tone, frequently 256. Printable Photoshop files have 256 VOT *per channel,* which is why so many different colors are possible. An RGB or LAB file can accommodate 16,777,216 different colors, this number being 256 to the third power.

To add to the exponential chaos, a second term, *bit depth,* is often used in preference to VOT. A 256 VOT file is also known as an 8-bit file. In keeping with the spirit of the rest of this chapter, it gets a more precise name, 8 BPC, for bits per channel. This refers to the computer storage space required per pixel, in this case eight binary bits: eight zeros or ones. With eight zeros or ones, the total number of possible variations is two to the eighth power, which is, conveniently, 256.

Most scanners and some other devices operate, at least internally, at higher bit depths. A 12 BPC scan has 4,096 VOT. Some manufacturers would have us believe that a greater bit depth implies better scan

Dots & Spots: a Glossary

Much confusion is caused by the use of a single term, *dpi,* to refer to wildly different genera of resolution. In an act of rebellion against this practice, I'll use different abbreviations in this chapter. Unfortunately, in some cases, I have had to invent them. Here is an alphabetical list of the acronyms you'll find here. These acronyms don't agree with industry practice, and I don't seriously suggest that you use them. But at least they are more accurate than calling everything dpi.

BPC	Bits per channel, in a digital image file.
BPI	Black or white bits per inch, in a bitmapped graphic file.
DPI	Dots per inch, in a halftone screen.
PPI	Pixels per inch, in a digital file.
SPI	Spots per inch, the smallest area that can be marked by an output device such as an imagesetter or film recorder.
SSPI	Scanning samples per inch.
VOT	Varieties of tone, sometimes called levels of gray, the maximum number of shades of gray in a single channel of a digital file.

quality. Don't believe it. If a scanner can't see detail in shadow areas, more BPC won't help. We may have 4,096 VOT, but in this case, VOT stands for varieties of trash. To see why, let's compare the work of three very expensive pieces of hardware.

For more than 30 years, drum scanners have been the standard for those desiring the highest quality. They're still the best today, but not by much. Their photomultiplier tubes have certain advantages over the charge-coupled device technology used in flatbed scanners and most digital cameras. CCD devices are particularly vulnerable to a loss of detail in the darkest areas.

In 1996, I arranged a shootout between a drum scanner from the early 1980s and two professional-level (i.e., they cost about $50,000 apiece) CCD units. I refrain from naming the products, because all three vendors make better scanners today.

Anyway, the idea was to have expert operators of each scanner try to milk the most from a dozen moderate-to-difficult chromes. They were given identical printing and sizing specifications. If they did not like the scans for any reason, they could do them over. When they were satisfied, all 36 versions would be assigned random letters and proofed next to one another. The proofs would go to a panel of ten experts who, working in separate light booths, would evaluate which of the three versions of each original was the best, without knowing which ones came from where.

I expected this exercise to prove that the CCD scanners had basically caught up. It didn't. Of 120 first-place votes, Brand X (you and I know it's the drum scanner, but the jurors didn't) got 98, Brand Y 12, and Brand Z 10. In seven of the 12 contests, Brand X swept all ten votes.

Figure 15.3 is one of those in which the vote was unanimous. I'm showing a piece of it, then enlarging it, and then applying a drastic contrast-enhancing curve to it.

It's easiest to judge how well these beasts are holding shadow detail in the three right-hand versions. The problems are the center of the tree and the car beneath it.

Only Brand X is having any luck with the car. If you look hard at Figure 15.3E, you can even see that the taillights are red. In the tree, which is darker still, it's more of a struggle. Brand X is obviously at its wit's end, but it is still a drum scanner, and a cut above the other two in such shadow areas. Brand Y posterizes the inside of the tree, whereas Brand Z freaks out. The whole center of its tree in Figure 15.3I is featureless.

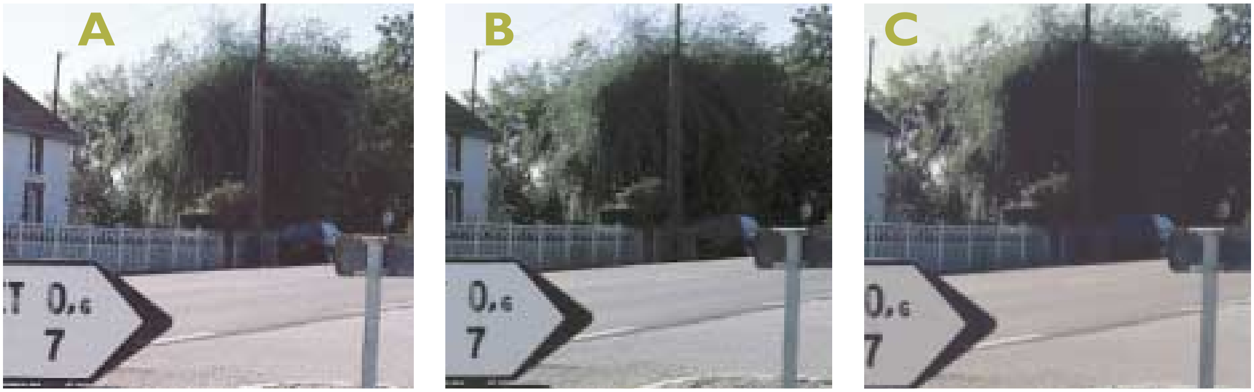

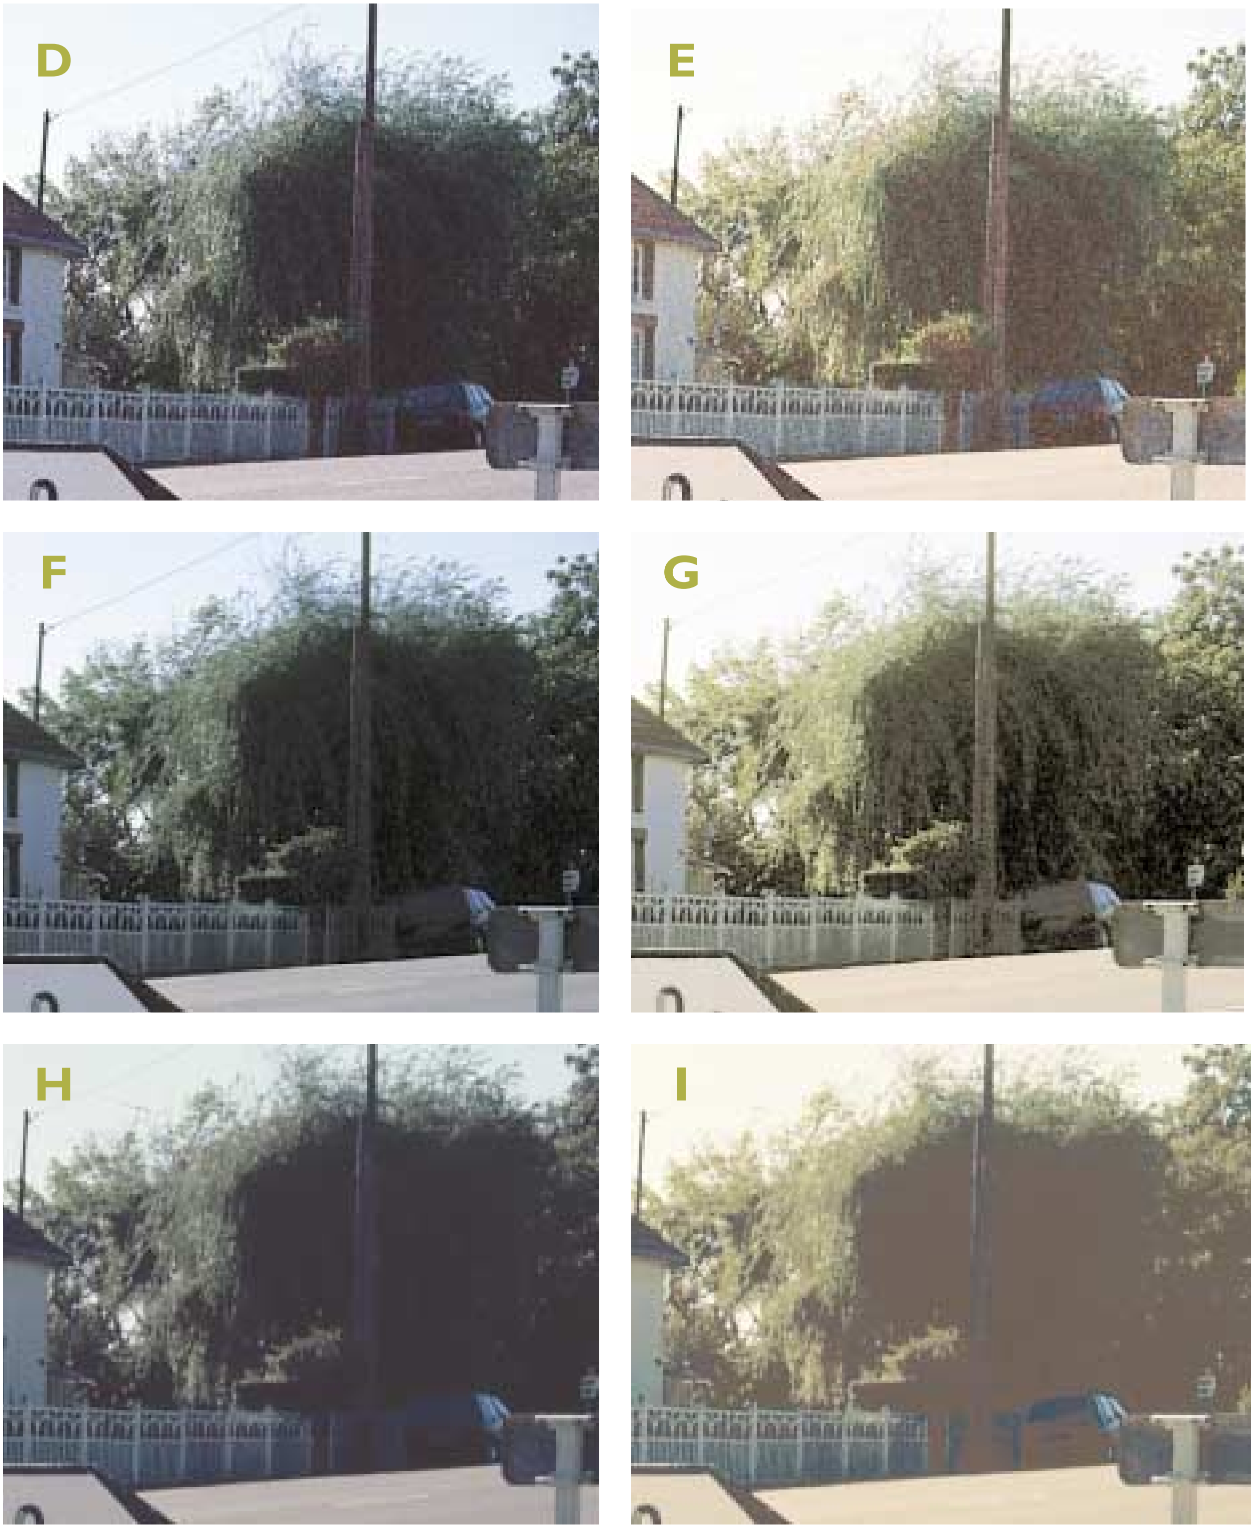

Figure 15.3 *The ability to retain shadow detail is a major test of a scanner. Opposite, normal-size reproductions of scans of the same original by, left to right, Brands X, Y, and Z. Above, the shadow areas magnified (left) and with contrast in the shadows greatly enhanced (right). Top to bottom, the Brand X, Y, and Z scans.*

The final output of these scanners is an 8 BPC file, but all three interpolate that down from an original with more data. Brands Y and Z are 12 BPC scanners. Brand X is analog; its original scan is a series of voltage readings that in principle carry an infinite variety of tonality. More modern versions of this scanner are fully digital and also give 12 BPC files.

So, Brand Z can portray 4,096 levels of tonality in each channel. A fat lot of good it did. In this image, endowing the scanner with extra bits is like handing a blind man a pair of eyeglasses. If Brand X were an 8-bit scanner, yea, verily, if it were a *six*-bit scanner capable of only 64 VOT, it would still have the best version here. If you're a buyer, forget this bit depth balderdash and look at a few tough samples.

Comes the Quintillion

The above discussion is no endorsement of drum scanners. Experts looking at these images closely found enough difference to say that Brand X was better. I concur, and so do you, presumably. The question is, though, how *much* better? Going back to the smaller ones that are the proper size for evaluation, Figures 15.3A, B, and C, I'd have to say, a little better, but very little.

So, while the drum scanners have a theoretical advantage, in real life it doesn't amount to much. In a race like this, the difference is the jockey, not the horse.

We've touched on three related factors that affect reproduction of these photos: the press screen (DPI); the capabilities of the imagesetter that produced the place (PPI); and the VOT found in the file itself. We now need to add a fourth, which is so important that it is commonly called the *resolution of the file,* as if the other three didn't exist.

The conventional wisdom, which is often wrong in resolution matters, says that the number of pixels per inch in the digital file should be roughly twice the line screen (er, DPI) of the output device. Actually, the conventional wisdom understands that one can get by with less, but desires to cut the designer a little slack.

The original files for Figure 15.3 were at 300 scanning samples per inch (SSPI), admirably double the 150 DPI screen of this book. If we were to use the scans here at same size, PPI and SSPI would be for the moment equivalent. In fact, though, even the smaller versions are placed at 125% of original scan size in my page layout program. Thus, there are no longer 300 pixels per inch. Instead, it's 300 pixels per inch and a quarter. The effective resolution is therefore 240 PPI, a ratio to the screen ruling of 1.6 to 1. The teeth of the conventional wisdom begin to chatter only at around 1.4 to 1, so we're still okay.

The six larger images, however, have effective resolutions of only 150 PPI—a fearsomely low ratio of 1 to 1. Do you see a disaster? I see some crummy-looking pictures, but not because of lack of resolution, *except* in Figure 15.3E. I find that one to be too noisy, and excessive noise is often a symptom of inadequate resolution. The conventional wisdom can go climb that tree; the needed resolution depends very much on the quality of the original.

***Figure 15.4** (opposite) How important is bit depth? One image is from a standard 8 bit per channel file, yielding about 256 varieties of tone per channel. There's also one where the bit depth has been cut to 7 BPC, or 128 VOT. Another cuts it to 64, and a fourth to 32, or 5 BPC. The four are in a random order. Can you tell which is which? The answers are in the box on page 318.*

A

B

C

D

Spots, Dots, & Tonality:
DPI vs. SPI vs. VOT

IMAGESETTER RESOLUTION (spots per inch)

SCREEN (dots per inch)	300	600	1200	2400	3600
65	21	85	256	256	256
85	12	50	199	256	256
100	9	36	144	256	256
120	6	25	100	256	256
133	5	20	81	256	256
150	4	16	64	256	256
175	3	12	47	188	256
200	2	9	36	144	256

***Figure 15.5** Realistic photographic images are impossible if the output device can't generate enough varieties of tone. For professional work, one needs at least 100 VOT, and many would say at least 200. Here's how many levels of tone various resolutions can theoretically produce at some common screen rulings.*

Figure 15.4 is a similar affront to the conventional wisdom. A full range of 256 VOT or thereabouts is often said to be critical to good reproduction. One of the versions has it. The other three had half, three-quarters, and seven-eights of their data discarded while still in RGB, via the Image: Adjustments>Posterize command. That is, one image had a maximum of 32 VOT per channel, one 64, one 128, and one 256. Can you tell which is which?

I've seen a Matchprint of the page, and was surprised at how close the results were—I expected the 32 VOT, and perhaps the 64 VOT one, to be garbage. I did not expect to be able to tell the other two apart. In fact, all four are perfectly acceptable. Because I know exactly what happens when images lose bits, I was able to pick the four out of a random proofing, but that's not to say any one is better than the others.

If there's not enough bit depth, the image will look somewhat harsh and jagged, just as a hypothetical fifth version, even at 256 VOT, would look if I cut its resolution to 100 PPI. Everything is related.

And the Band Played On

Banding in gradients has been a headache since the first days of PostScript. In the middle of what seems to be a smooth gradation, there is a systemic burp, a sudden, annoying jump from one tonality to another, ruining a job.

"How can this possibly have happened?" shrieks the hysterical artist, "when I specified 256 steps for the gradient?"

Two ways. The likeliest is that the gradient was originally made in RGB and involved colors close to the edge or outside of the CMYK gamut. The other is, of course, a resolution problem.

If a gradient goes from say, 10% to 20% in a certain channel, the output device will only have about 25 VOT available. It won't matter whether the input has 256 or 256 million VOT. Furthermore, some of the 25 are probably unavailable, due to rounding error, so banding is likely, especially if the gradient covers a wide area on the page.

In principle, an output device that can create 256 VOT itself has precisely enough resolution. In practice, that isn't quite true. The dots are angled, which can reduce the number of spots available. More important, although both file and imagesetter are at 256 VOT, the two 256s won't line up exactly. Certain values that are different in the original file will result in identical dots, and certain dots that are theoretically possible will actually be inaccessible.

So, how many VOT are actually available at 150 DPI on a 2,400 SPI imagesetter? Chances are, 220 or so. If photographs are the only thing we print, neither we nor anybody else will be able to tell the difference between 220 and 256 VOT. If, however, we start introducing fine gradations, it may make quite a difference indeed.

Banding can be defeated by adding a small amount of randomization (noise) to the digital file. Photos therefore rarely band; there is almost always enough natural noise to obliterate the problem.

Photoshop's separation algorithm has an anti-banding measure—a dither, a very fine amount of noise applied in areas where banding is likely to occur. Other methods of overcoming banding are to use a higher-resolution imagesetter or to cut the screen ruling. 150 DPI gradients on a 2,400 SPI imagesetter tempt fate. 133 DPI gradients are more reliable.

Figure 15.5 shows how many VOT are actually available at normal DPIs for some common SPIs of laser printers and imagesetters. I'd avoid any number under 150.

Inkjet printers usually are 1,200 SPI and use a dither rather than a conventional halftone screen, thus avoiding many problems. Try printing a halftone screen of even 85 DPI, however, on a 600 SPI laser printer, and you're likely to be dissatisfied.

And yet there's a resolution paradox. This book is imaged on a 2,400 SPI platesetter. If we were to substitute output from a 600 SPI laser printer, the images would be a joke—but few people would notice that anything was wrong with the *type*. How can the same resolution be so terrible in one case and so nearly undetectable in the other?

The Resolution That Isn't There

Unlike the images we've been considering so far, type contains no grays, only black letters on white paper. The type still has to be painted using the same imagesetter spots, but now it's much easier. The glory of PostScript is that it allows certain kinds of graphics to have an entirely different variety of resolution, to wit, none at all.

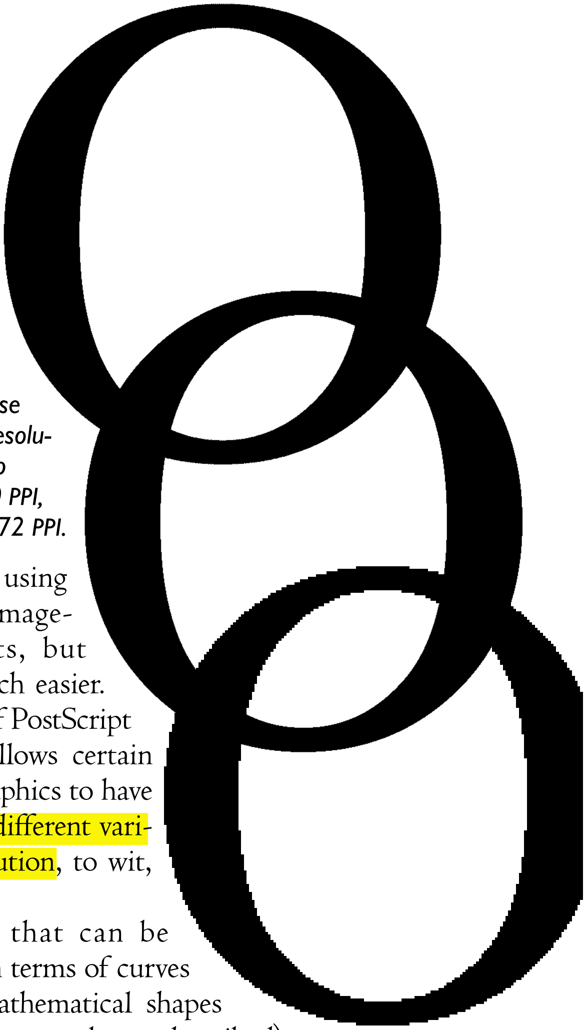

***Figure 15.6** How much resolution is needed for smooth-looking type and similar graphics? These letters have resolutions of, top to bottom, 1800 PPI, 300 PPI, and 72 PPI.*

Objects that can be described in terms of curves or other mathematical shapes (and typefaces can be so described) eventually need a resolution. An imagesetter doesn't print mathematical concepts, only spots. But a RIP's whole function is to map out those spots. When a file comes in saying, "I am a bunch of curves, map me however you think best for your imaging engine," the RIP does so more smoothly and undetectably than if the file already carried its own bitmap.

Since type and similar line graphics contain no grays, Figure 15.5 is irrelevant. The only question is, how many spots per inch does the printer need to construct the curves and fine lines of text type accurately?

Older laser printers generally have a resolution of 300 SPI and produce type that's obviously inferior to what you're

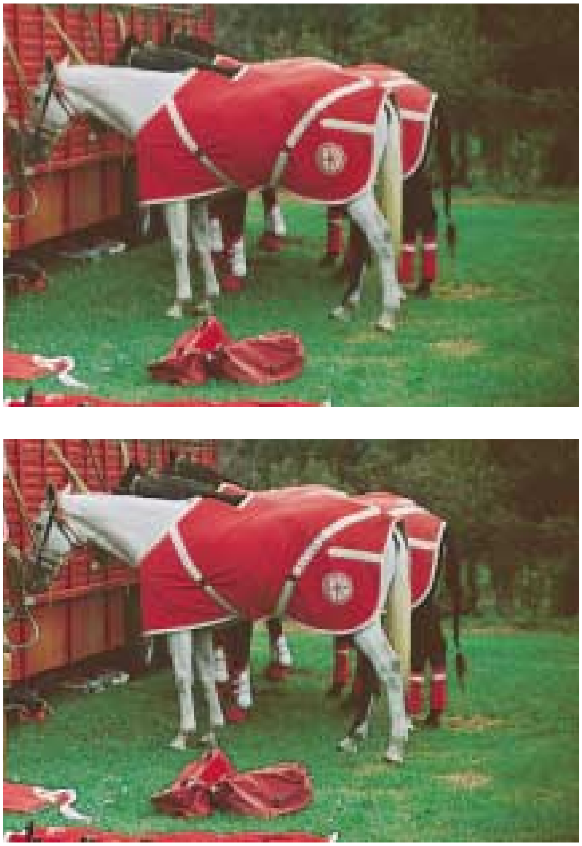

Figure 15.7 *Does resolution equal detail? The top version seems soft, even though it was scanned at three times the resolution of the bottom version, and takes up nine times as much disk space. In areas of one color, like the grass, the higher the resolution, the more even the color will become.*

reading here. At 600 SPI, the type is pretty good—one has to look closely to see the difference between it and type output from a 1,200 SPI imagesetter, which, in turn, is indistinguishable from 2,400 SPI without a loupe.

When it's necessary to scan type or other graphics because digital versions don't exist, yet another kind of resolution comes into play. We now have a file that can be expressed in bits per inch, each bit being either white or black. If a 600 BPI file is sent to a 600 SPI printer, the printer's RIP has to remap it. The results will not be quite as good as if a resolution-independent file were sent to the same printer.

So, at what resolution should one scan type and other line graphics? Half again the resolution of the output device is my rule, to a maximum of 1,800 SSPI.

If you've ever wondered why type always looks fuzzy in a photograph, it's that old devil resolution again. In Figure 15.6, you will observe that 300 BPI is inadequate for type—and most color images are scanned at less than 300 SSPI. The type in such images isn't quite as jagged as the example, because screening tends to soften images, but it still will be pretty bad.

One way to make it better, naturally, is to scan at a higher resolution. The higher the scanning resolution, the softer and smoother the curves will be. The problem is, so will everything else.

When scanning type only, excessive resolution eats up disk space, overburdens the imagesetter, clogs up communication, and is generally a complete waste of time. Other than that, it doesn't hurt. But with a photograph, too much resolution, in addition to the shortcomings enumerated above, actually *does* hurt. Which of the two images of grass in Figure 15.7 do you like best?

If we want something that looks like blades of grass and not AstroTurf, the bottom version seems clearly better. But it's the lower-resolution scan! Doesn't high resolution equate to more detail?

Of course it does. But here, we don't want detail, we want the *illusion* of detail. That's what the bottom one provides. Let me try to explain how and why.

The bottom image's resolution is approximately four scanning samples per halftone dot. That's in line with the conventional wisdom, which is that it should be between 1.5 and 2 times the screen ruling. The dots are roughly $\frac{1}{133}$" apart, the scanning samples roughly $\frac{1}{266}$". That squares up to four samples per dot: two across, two down.

The top image has three times this resolution. The samples are roughly $\frac{1}{700}$" apart. The file is nine times as large. There are now 36 scanning samples per halftone dot, rather than four.

The grass, obviously, is predominantly green. Parts are gray, black, yellow, or brown. But at either of these scan resolutions, probably three out of four samples will find green.

In the lower resolution image, therefore, which has four samples per halftone dot, the chances are that three of them will be green, but sometimes all four will be, and sometimes zero or one. In such cases, the resulting dots won't produce green.

In the higher resolution version, with 36 samples per dot, this effect is far less likely. It's conceivable that three out of four samples may not be green. It isn't conceivable that 27 out of 36 samples—which is the same ratio—will be something other than green. A rule of mathematics: the more samples, the less variance from what the law of averages predicts. If we flip a coin four times, it may well come up heads three out of the four, although two heads is more

Figure 15.8 *What passes for detail is often nothing more than variation. Pixels in the top two images (blowups of the two images of Figure 15.7) seem to have the same amount of variation, but this is an illusion. If the high-resolution version, top, is downsampled to match the resolution of the middle version, the result is the image at bottom, which is much softer.*

likely. If we flip four *hundred* times, 300 heads couldn't possibly happen.

The higher the resolution, the more uniform the color will be: the closer it will approach whatever the average color of the grass is. There is a lot more variation in the lower resolution version. That variation, or action, suggests the blades of grass that our imagination is telling us are actually there.

In scanning, moving to a lower resolution is a move toward action and variability. This is a fine concept, but if the resolution gets too low, the image will become harsh and jagged.

A higher resolution is a move toward smoothness and consistency, which are also laudable goals, in moderation. Too much resolution will make the image look soft and defocused.

It follows that there is no one "correct" scanning resolution. A woman's face generally should be scanned at a higher resolution than a man's, because we accept more roughness in a man's face. An image of furniture requires more resolution than does grass, because furniture has diagonal lines that shouldn't look too harsh. A damaged, noisy, or prescreened original also is helped by higher resolution. And, certainly, if you think there is a good chance you'll be upsizing the image, give yourself some extra resolution in the original scan.

Many people, refusing to believe that too much resolution can hurt, scan everything at 300 SSPI. This obstinacy explains why so many newspaper photographs look so soft. It also makes the vendors of disk drives very happy. File size increases with the square of the resolution. If, for magazine work, you go with 250 SSPI rather than 300, your files will only be two-thirds as large—and quality will probably be *better.*

Resampling and the Rogue Pixel

The foregoing discussion concerns scanning resolution, which is expressed in SSPI. The resolution of the Photoshop file is not necessarily the same kind of resolution. We express this in PPI—pixels per inch. A pixel is the smallest building block of a file. You can see them clearly in Figure 15.8.

At the moment a raw scan is opened directly into Photoshop, SSPI equals PPI. That equality does not necessarily continue, because at some point the scan may get *resampled.* Plus, with many desktop scanners, what seems to be the raw scan isn't any such thing: it may already have been interpolated.

Photoshop itself allows us to change the number of pixels in a file, using the dialog box shown in Figure 15.9, accessed by Image: Image Size.

When the Resample Image box is unchecked, changing the numbers changes only the nominal size of the image, not any data. A 4"×6" file at 150 PPI is exactly the same as a 2"×3" file at 300 PPI. One changes size without resampling for the sake of convenience. For example, most digicam captures have a nominal resolution of 72 PPI. That's more than enough pixels to use for this book—provided I place them in the page layout file at a quarter of their nominal size, which would make their effective resolution 288 PPI. To avoid this hassle, I change resolution to 250 PPI, without resampling. That, I know, will make the image close to the size I need.

Resampling *down*—that is, throwing some of the data away—is appropriate when there's more than enough for whatever use you intend. It's ridiculous to post a 5 MB image file for Web viewing, just as it's ridiculous to use a 15 MB file for the top half of Figure 15.7.

To resample down, check the resample box and enter a lower size, resolution, or both. But keep two things in mind. First of all, unless you're positive you'll never need to print the image at a larger size, save a copy of the original. Downsampling is a one-way street.

Second, realize that a downsampled image isn't equivalent to an original scan of the same PPI. The downsampled version will be softer, like the bottom third of Figure 15.8.

The lower the resolution, the more chance that a rogue pixel will appear, an area where the scanner picked up really atypical information. Such a thing usually translates into a halftone dot that looks out of place, almost like a speck of dust. The greater the density of scanning samples—resampled or not—the less likely this is to occur. And the higher the screen ruling, the less noticeable the effect of such a rogue pixel will be.

The conventional wisdom says that resolution should be 1.5 to 2 times screen ruling times magnification. This recommendation has some validity, but there are a whole lot of exceptions. If you are interested in playing fast and loose with low resolutions, you rate to get away with it under the following circumstances.

- **Digital cameras** produce files that are considerably purer, less noisy, than ones made by scanning film. High resolution is the enemy of noise. If the noise isn't there, there's no point in wheeling out its enemy.
- **Files that have been downsampled** have gone through an averaging process already. That reduces noise.
- **Inkjet printers** use a very fine, dithered dot that gives a softer appearance and suppresses noise naturally. If one dot is calculated incorrectly, it's not nearly as big a deal as with a conventional screen.
- **Finer screen rulings** don't need high resolution, the point of which would be to avoid halftone dots that are obviously out of place. The smaller the dot, the less offensive such an event will be. Nobody's ever studied the matter as far as I know, but I suspect that for typical scanned images a 1.2 ratio is adequate for 175 DPI screens and up, and 1.0 is probably acceptable if the file comes direct from a digicam.
- **The image itself** can let us cut corners. Figure 15.7 is a very soft subject. I've printed it successfully with the resolution as low as .8 of the screen ruling. For a picture as busy and full of fine detail as Figure 15.1, that would be disastrously low.

The Emperor's New Clothes

The principle that more data means smoothness, less means action carries over to a different type of resolution: the hotly debated topic of bit depth.

As noted earlier, many scanners and some digicams can record more than 256 VOT. Photoshop can work with them to a

Figure 15.9 *The Photoshop resampling dialog box, accessed under Image: Image Size.*

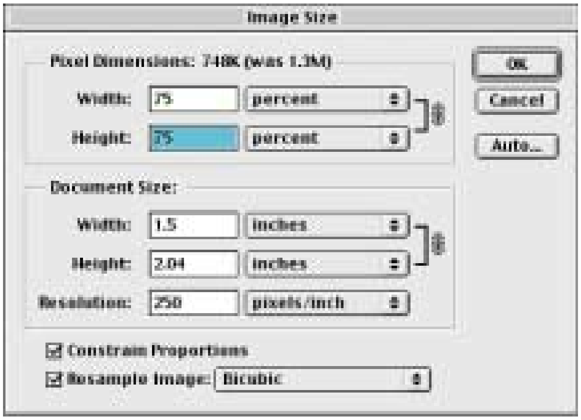

Pixel Dimensions: 748K (was 1.3M)

Width: 75 percent

Height: 75 percent

Document Size:

Width: 1.5 inches

Height: 2.04 inches

Resolution: 250 pixels/inch

☑ Constrain Proportions

☑ Resample Image: Bicubic

limited extent. We can apply curves, but only a few filters and retouching tools, before converting to the 8 BPC that all output devices require.

When we apply a curve, we reduce the image's VOT, for the following reason. Suppose that we have a 256 VOT B/W file. It happens to be a picture of a white cat, so we jack up the center point by, say, 10%. Originally, 128 tones fell below the midpoint, and 128 above it. But now, we have stretched the light tones and compressed the dark ones. Only about 115 real tones now fill the first 128 available spaces. On the other hand, 141 real tones are competing to fill the second 128 spaces. The surplus has to be discarded. Hence, only 243 real tonal values remain of the original 256.

If we work, instead, with the extra bits, this criticism will not apply. Photoshop's two options (under Image: Mode) are 8 or 16 BPC. While no scanner extant gives 16 meaningful bits, if we start to fool around in LAB or do various other things, it's possible that we may fill those extra bits up with something other than garbage, in the course of making our file size twice as large.

In addition to the bloated size, we also have, er, a bit more information. To be precise, each channel now has a resolution of 65,536 VOT. How many discrete CMYK possibilities can this produce? The answer is so impressive that I can't bear to use numerals, I have to just say it. Eighteen quintillion, 446 quadrillion, 744 trillion, 73 billion, 709 million, 600 thousand, that's how many.

If we apply a curve to such data, we still throw away a few of these possibilities. We will miss these about as much as Bill Gates misses the quarter he spends on his morning newspaper.

For those calibrationists to whom a good-looking histogram is more important than a good-looking image, this cinches it. We must work in 16 BPC whenever possible, to avoid that fearful bogeyman, data loss.

This belief is parroted in most textbooks and has spawned no end of seminars at industry conferences explaining how anyone who would even think of correcting in any other way is a rube, etc. In support of this dubious concept, viewers are treated to a display of worthless histograms, and also demonstrations that computer-generated gradients band if worked on in 8 BPC and not in 16.

As for the histogram, that and a billion dollars will put your net worth into ten figures. We sell images, not histograms. As for the gradient, just because an orange tastes good doesn't mean that a keyboard does. A gradient is original, first-generation art, from which any variation is an error. A digital photograph is anything but; it's already been mangled by whatever device captured the data, not to mention whatever screwup the user has added.

No, what we want to see is some 16 BPC photos where some series of real-world moves produce an image that's in any significant way better than it would have been if the file had just been converted to 8 BPC in the first place followed by the same moves. Remembering Figure 15.4, which displayed little loss of integrity even when three-quarters of the bits are gone, such a thing is rather difficult to find. When I've run tests in the past, corrections done in 8 BPC look sharper, those done in 16 BPC smoother, just as one might expect. So, one may theoretically work very slightly better on a given image than another, but the differences are very hard to detect.

In response to any number of queries from confused readers, I said publicly that I had seen nothing to suggest that this was either a good or bad thing to do, but it wastes a lot of disk space. I noted that none of the so-called authorities on these matters have ever, as far as I know, shown real images with real manipulations that showed any improvement. I requested that anybody having such images send them to me, and promised I would include them in magazine articles if they backed up the claims.

A number of skeptical users posed the same challenge to the "experts" in public forums. One actually offered a $100 reward for any natural photograph that could be shown to benefit from 16 BPC correction.

The reply was an indignant refusal to show images, almost as if the proponents had been so seduced by histograms that they had never tested the concept. Rather than admit this, they tried to bluff their way through, claiming that the difference was so dramatic that it wasn't worth the time to demonstrate. Here's a sampling. Four different voices are actually speaking here.

"16 bit capability is critical during all aspects of tone compression.…The difference CAN be seen in the final output very easily. Most definitely on the printed page, especially when using high-quality halftoning and even more so to a film recorder.…It's very easy to see that substantial color & tone editing will eventually result in data loss and banding.…If it means the difference between taking a 16-bit image capture and editing that to the final image and taking that same image in only 8-bit and editing that to the final image *then* there is a difference like between the day and the night.…Yes, if a histogram full of holes has no impact on final output, then throw away the graphs and just get on with the print run. However, all of us have Real World Output showing the superiority of superior data acquisition.…My advice? Take the information you've read here to the bank. Stop doubting and start applying what you've learned here.…If you really start out with a RAW image in high-bit form and a raw image downsampled to 8 bits, the difference really is night and day.…It's totally obvious to anyone who looks that it's very advantageous to do the big moves on high-bit data."

* * *

People sent me CDs full of 16 BPC images. During a cold and boring week in April 2002, I applied curve after curve, command after command, at different gammas, in different colorspaces. I concluded that there was *zero* meaningful difference, let alone the kind of enormous difference being claimed above, under any remotely conceivable circumstances.

I learned one very useful tidbit, which I will share in a moment, in the course of producing enough variant versions to fill this entire volume.

In the interest of sanity, I'll limit the exercise to three images and warn in advance that I have no intention of lavishing this much space on the topic in subsequent editions. Except as noted, each was corrected once in 16 BPC and then reduced to 8 BPC for printing, and once by converting the raw 16 BPC to 8 BPC immediately and then reapplying the same curves.

Since these are night and day differences, totally obvious, it's plainly a waste of space to identify which is which. Each set is in a random order. (If you have any difficulty, the answers are in the box on Page 318.)

Please ignore the question of whether the

Figure 15.10 *Drastic curves like the ones shown here allegedly favor the use of 16 BPC mode, particularly in the case of a professional chrome like this. Above, a reduced version shows the magnitude of the move. The two versions at right are at the full resolution of the midrange desktop scanner, 250 PPI at this size. Can you tell which one was corrected with curves applied in 16 BPC followed by conversion to 8 BPC as opposed to being converted first? Opposite page: At 1,000% magnification, we can see minor differences in the final green channels and to a lesser extent in the composite color image. But can you say which is better, or even which is which?*

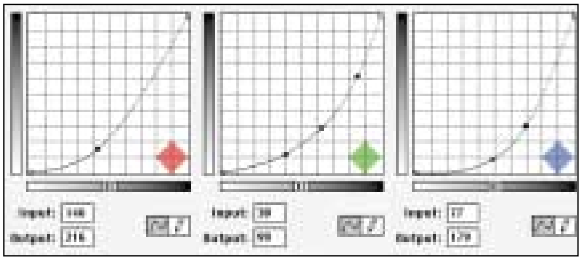

Input: 148
Output: 216

Input: 38
Output: 98

Input: 77
Output: 179

color correction is to your liking. We're evaluating texture here. Also, if you see any color variation between versions, it's caused not by curves but by press vagaries of the type described in Chapter 14.

In Search of Drastic Enough

If 16 BPC has an upside, it will show up in extreme moves, or a series of them. That's why the raw images are intentionally flat—if anything were going to show an advantage for 16 BPC, this type of scan would do it. All images are printed at 250 SSPI, without resampling. The 16 BPC originals and the curves I used to attack them are on the book's CD.

Figure 15.10 shows the effect of a single rather radical set of RGB curves. I have trouble seeing a difference between the two

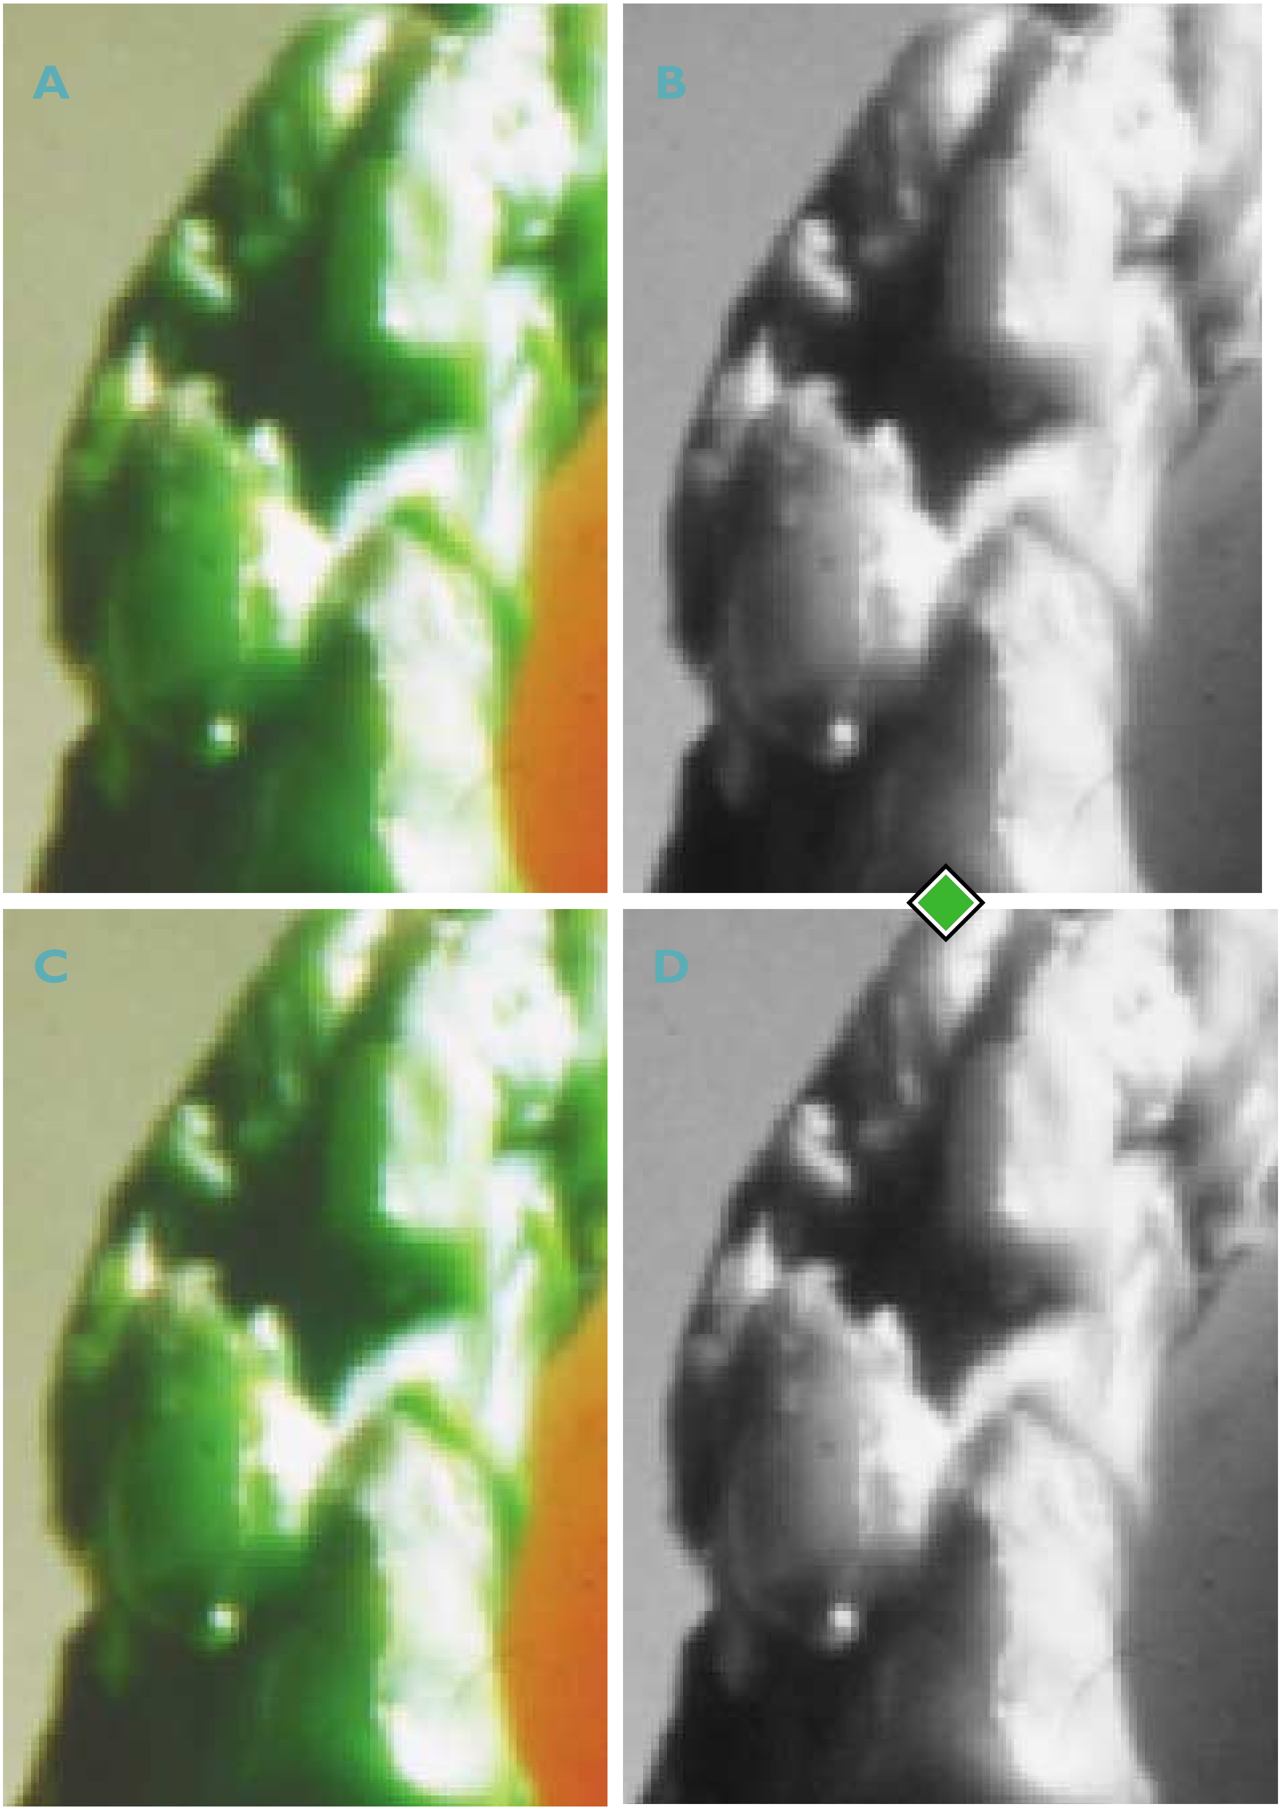
A
B
C
D

smaller corrected versions. Advocates of this workflow would retort that book printing isn't the most demanding process in the world. If this were going to a film recorder, they'd say, one might see the damage.

True enough. Going to a film recorder at same size, however, is not as demanding as printing at 1,000% magnification. At that size, can you see the difference yet? Do you think it affects quality?

Advocates say they need 16 BPC because the moves they make are very, very demanding. Personally, I feel that anyone in the habit of making corrections as substantial as that in Figure 15.10 has more workflow issues to worry about than the number of bits. But, some might conceivably say that they have to make even bigger changes.

To that dubious proposition, I say, get a load of Figure 15.11A.

Figure 15.11 *If a succession of steep curves can damage an 8 BPC image, this figure ought to prove it. The original, A, is so nearly nonexistent that no single curve could correct it. B and C therefore had very steep RGB curves applied, followed by equally steep ones in CMYK after conversion. One version was done in 16 BPC all the way. The other used the same curves but applied to an 8 BPC original. Which is which?*

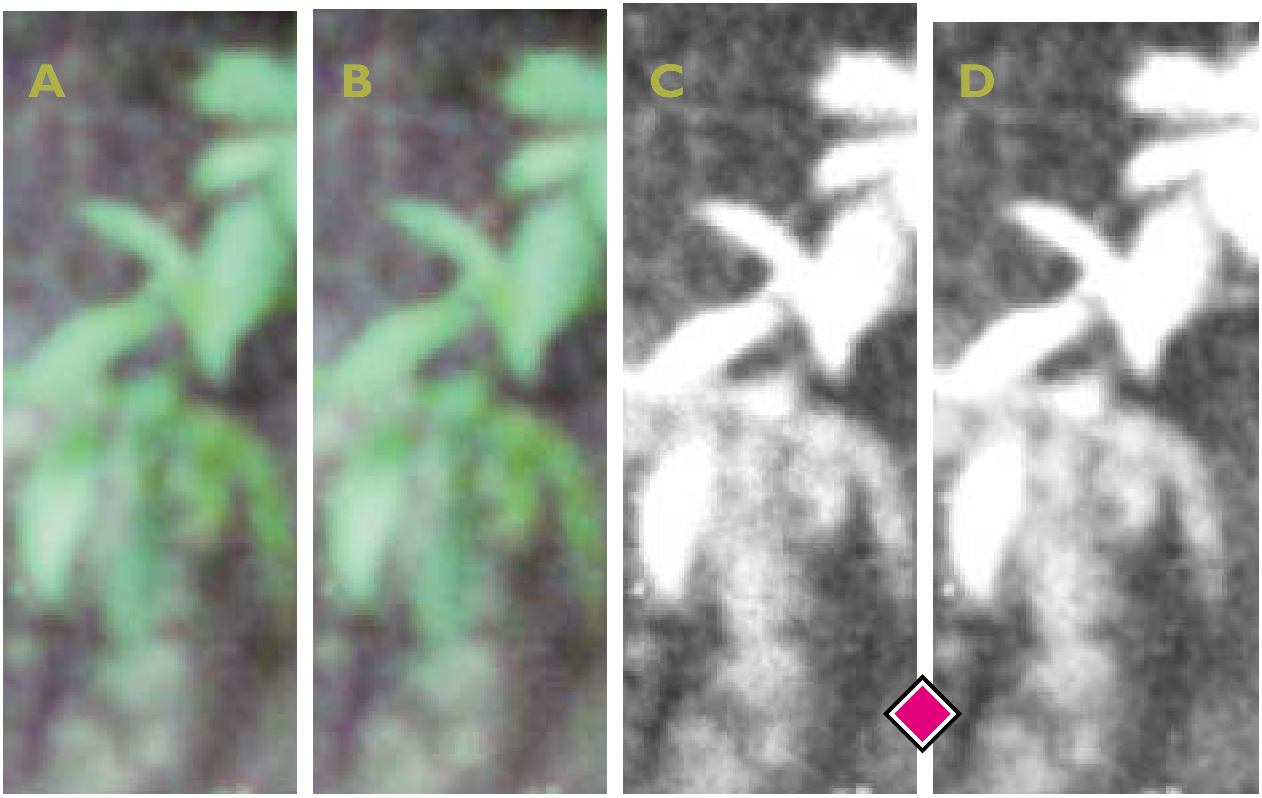

Figure 15.12 *At 400% magnification, the version done in 8 BPC looks a bit more noisy and variable. That's probably to our advantage in the leafy area above, which will look sharper, and a minus in the face below. The composite versions, and their magenta channels, are in a random order. Can you pick out the 8 BPC ones?*

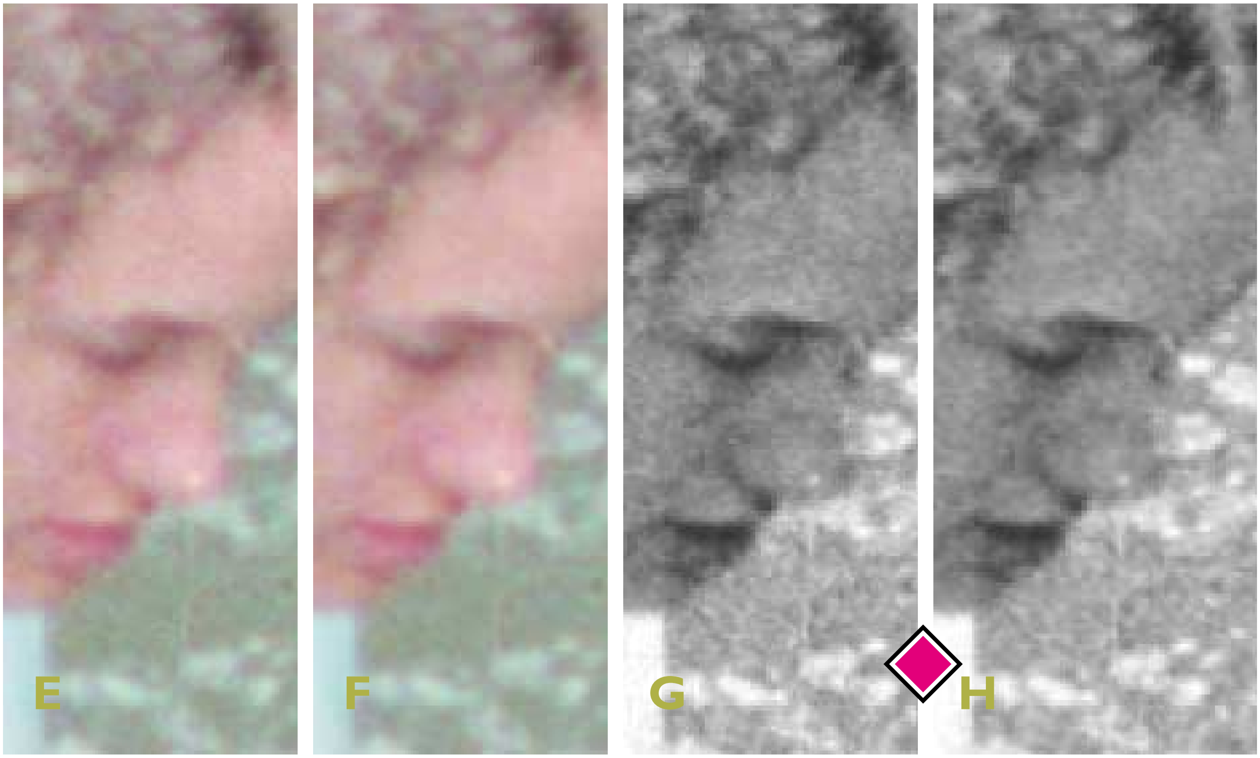


A


B


C


D

Nobody can get that mess into corrected form without at least two sets of immensely steep curves. I did one in RGB, followed by a second in CMYK. Can you tell the difference between the two versions? If you can, you probably like the one done in 8 BPC better.

Because of the enormity of the change, it's a lot easier to see the difference at 400% magnification in Figure 15.12 than it was at 1,000% in Figure 15.10. The 8 BPC correction appears grainier both in the face and in the leaves, which were cropped out of Figure 15.11. But that graininess actually helps at the final size of the image—if apparent sharpness is what we want.

Figure 15.13 *Dark blue gradients are notorious for banding under certain conditions. The original, top right, was deliberately scanned too dark to require extreme corrections. Facing page: one correction was done in 16 BPC, one was done in an 8 BPC file supplied by the scanner, not Photoshop, and one was done in Photoshop by changing the mode of the 16 BPC file to 8 BPC before correction. The ugly-looking green channel is from one of the three. This page, below: a torture test pitting 16 BPC against 8 BPC correction: four separate sets of drastic curves, applied in two different RGB definitions. In addition to the 16- and 8 BPC versions, there's a ringer below: it's the 8 BPC file saved as a Medium-quality JPEG and then resaved for printing here.*

Which Was Which?

Many of the figures in this chapter don't have their variants labeled, to see whether you can pick them out. In some cases, I couldn't, or at least not from proofs that the printer of this book provided.

In Figure 15.4, the test of various bit depths in an image, the one with 8 full bits, 256 VOT, was version C. The 7, 6, and 5 BPC versions were respectively B, D, and A. Working off a Matchprint, I identified all four correctly but consider that they are close enough for any purpose.

I didn't see proofs of the left side of Figure 15.10, the lamb plate, in which the top version was done in 16 BPC. I doubt that anyone could tell the images apart, since the right side, at 1,000% magnification, stumped me. A and B are the 16 BPC versions. Looking at a proof from an expensive photographic printer, I got this right after much labor. Some weeks later, from a Matchprint, I misidentified the B/W versions.

The extreme moves in Figure 15.11, the young woman in the tropical pool, to my mind favored the version done in 8 BPC, which is C. I had little difficulty identifying that and none whatever picking out the grainier 8 BPC versions in Figure 15.12. For the record, they are B, C, E, and G.

I was completely defeated by the left of Figure 15.13. I identified version A as being the 16 BPC because I thought it was the smoothest-looking on the page. I had a very tough time differentiating B from D.

Surprise! Version A is the one using the scanner 8 BPC, the one with the badly banded green channel shown in version C. The Photoshop 8 BPC is B, and C is the 16 BPC version. I was so floored by this that I checked the images several times and pulled another proof, but the result was the same.

In the more extreme lipstick move, F is the 16 BPC, G the JPEGged 8 BPC, and H the uncompressed 8 BPC. I identified the 16 BPC, but mixed up the other two, expecting a slightly softer look from the JPEG. I prefer the background of the 16 BPC version in this image, although I think it has the worst metal.

The flip side, of course, is that when smoothness is at a premium, the 16 BPC correction may look marginally better. But the smoothing effect is much less pronounced than that offered by a higher scanning resolution.

Scanner and Photoshop 8 BPC

The lipstick of Figure 15.13 features a semigradient in the background, which might be expected to favor 16 BPC. A reader sent the digital files, along with an Epson proof that apparently showed just the kind of banding in the blue that might justify statements about night and day differences.

He provided both 8 and 16 BPC original files. My lightening curves produced the same result: ugly banding in the 8 BPC, smoothness in the 16 BPC.

Noticing that the two files weren't quite the same color, my paranoid side took over. I made a copy of the 16 BPC file, converted it to 8 BPC in Photoshop, and applied the same curves. Bingo! Banding gone.

It developed that the reader's 8 BPC file was interpolated from the 16 BPC not by Photoshop but by his scanner. During conversions, Photoshop applies very fine noise as an antibanding measure. I wouldn't have thought it would be quite so effective. But since then, I've heard from two other readers who confirm the same results with their capture devices. On the basis of this, I recommend that, if your scanner or digicam *can* give you a 16 BPC file, take it. Thereafter, I see no point in maintaining it in 16 BPC, but it doesn't hurt either.

My second shock came when looking at proofs of the four left-hand

Quick & Dirty

SOME RESOLUTIONS ABOUT RESOLUTION

✓ Many different kinds of resolution are often described by one ambiguous term, *dpi.* Photoshop users have to know what each kind of dpi means.

✓ Graphics and type prepared in vector programs such as Adobe Illustrator, Macromedia FreeHand, and CorelDraw are resolution-independent: the output device will draw the graphics in the optimal fashion.

✓ The choice of screen ruling shouldn't outstrip the capabilities of either the press or the imagesetter. If the imagesetter isn't capable of at least 150 varieties of tone at a given screen, go to a coarser screen.

✓ The industry consensus is that scan resolution for an image that is to be printed at same size should be 1.5 to 2 times screen ruling. As the screen ruling goes up, the effect of inadequate resolution goes down. Also, digital camera captures, being very consistent, don't require as much resolution as scanned artwork.

✓ Scanning at a relatively high resolution guarantees smoothness and consistency. If overdone, however, images become soft. Lower resolution is a move toward action and variability. If overdone, though, images become harsh and jagged.

✓ Excessive resolution is often as harmful as not having enough. This caution applies to the screen ruling of a printed image, and to scanning, where too much resolution yields an overly soft result.

✓ Type and other line graphics should generally be scanned at 1.5 times the printer or imagesetter resolution, or greater.

✓ Many manufacturers of scanners and digital cameras trumpet how many bits per channel they capture. This is very interesting but has little bearing on the critical question, which is how accurately the device sees into shadow areas.

✓ It's fashionable in some circles to forcefully advocate doing corrections to 16-bit files rather than the conventional 8-bit. This doubles file size, and is inconvenient. There is no reason to believe that there's any gain in doing so. The number of levels needed for quality output is much lower than the 256 allowed by 8 bits.

✓ Inkjet printers, which tend to give softer-looking reproductions than presses do, don't need as much resolution in the input file.

versions of Figure 15.13. They came from a very expensive photographic device, which required that my CMYK book files be converted to RGB, among other atrocities that may have included an anti-banding measure of one kind or another.

As Figure 15.13C shows, there really is serious banding going on in the reader's original 8 BPC file. I expected that file to print much worse than either the 16 or the Photoshop 8 BPC versions, which I thought would be very close.

Instead, two of the versions were close, and one was *better.* I assumed that this one must be the 16 BPC version. It wasn't. It was the reader's own file that had banded so badly on an Epson printer. I was so flabbergasted that I put the prints under a microscope to verify that the smooth-looking result really came from the banded file.

Finally, in the right half of Figure 15.13, I tortured the image with four sets of curves and four conversions between differing flavors of RGB. This time, I favored the 16 BPC version as having a less grainy background. But for the actual pressrun, there's a ringer: a version of the 8 BPC version that was compressed into a medium-quality JPEG before being resurrected for printing.

The 16 BPC CMYK file, much of which has been cropped out, weighs in at almost 20 MB. The JPEG version is around 800K. There's a gradient here, and the magnitude of the curves is vastly, inconceivably, and incontestably greater than one would ever encounter in the real world. Which version is which?

Some Final Resolutions

The genesis of this chapter was a magazine column of many years ago, which I reluctantly wrote in response to the requests of novices who didn't get the difference between the "DPI" of scans and laser printers. On reading the first draft, however, it struck me that this is the most baffling and confusing topic I've ever written about—even though it doesn't even touch monitor resolution, the resolution of Web images, or large-format printers. That convinced me to highlight each occurrence of phrases like yet another species of resolution.

Like many professionals, I had made the mistake of assuming that this was all intuitive. After observing just how many highlighted phrases there were (and after reading an extraordinary volume of correspondence after the column appeared), I had a much better understanding of why some artists find themselves buried under a blizzard of soft images, jagged edges, strangled networks, and unhappy RIPs.

Meanwhile, you've just seen how a resolution question has completely buffaloed some of the more distinguished names in the Photoshop world, who have been championing an inconvenient method that seems to do no good at all. Before relying on a histogram or some other magic charm, think back on how enormous the differences they thought would show up in the pictures we've just looked at, and recall that quantity of data is one thing and quality another.

The way to resolve one's resolution difficulties can be simply stated: don't ask for too much, don't provide too little.

You may find it easier to do that if you refrain from using that deceptive DPI term.

Declaring that we will never let those three deadly letters pass our lips is probably impractical. I don't really advocate that. But even if you *say* DPI, don't *think* DPI. Keeping the true meaning in mind is one of the best resolutions you can make.