


sRGB vs. CMYK in KDP Print: What the Numbers Actually Show



Walton Mendelson
walton@12on14.us

sRGB vs. CMYK in KDP Print: What the Numbers Actually Show

A measured look at color mode, shadow detail, highlight detail, and color fidelity in KDP Premium Color and black-and-white interiors.

Walton Mendelson
walton@12on14.us

Note

This study has taken a year of steady work. But it is the culmination of work that goes back fifteen years, to test books I designed and printed in 2010, with CreateSpace, the precursor to Kindle Direct Publishing.

I have relied on Claude to help. Here is one of several sets of equations for Delta E from Wikipedia (https://en.wikipedia.org/wiki/Color_difference):

The distance of a color (L_2^*, C_2^*, h_2) to a reference (L_1^*, C_1^*, h_1) is:^[15]

$$\Delta E_{CMC}^* = \sqrt{\left(\frac{L_2^* - L_1^*}{l \times S_L}\right)^2 + \left(\frac{C_2^* - C_1^*}{c \times S_C}\right)^2 + \left(\frac{\Delta H_{ab}^*}{S_H}\right)^2}$$
$$S_L = \begin{cases} 0.511 & L_1^* < 16 \\ \frac{0.040975 L_1^*}{1 + 0.01765 L_1^*} & L_1^* \geq 16 \end{cases} \quad S_C = \frac{0.0638 C_1^*}{1 + 0.0131 C_1^*} + 0.638 \quad S_H = S_C (FT + 1 - F)$$
$$F = \sqrt{\frac{C_1^{*4}}{C_1^{*4} + 1900}} \quad T = \begin{cases} 0.56 + |0.2 \cos(h_1 + 168^\circ)| & 164^\circ \leq h_1 \leq 345^\circ \\ 0.36 + |0.4 \cos(h_1 + 35^\circ)| & \text{otherwise} \end{cases}$$

CMC l:c is designed to be used with D65 and the CIE Supplementary Observer.^[16]

Some sets of calculations took hours to run, even at AI speeds, on just a quarter of the images; some runs hit Claude limits; and for most tasks, the calculations, the tests, or the art, had to be reworked to compensate for this or that. Run and re-run several times. Images were rescanned at least twice, sometimes three times, to get better quality against which to measure and analyze.

From 2010 until early last year, the thought of running calculations like this on over 300 individual images (swatches, pictures, and scans from sRGB and CMYK art), in whatever pairings—recognizing when the calculations produce errors suggesting the need to reassess the methodology, not once but several times to fine tune the process—was beyond me, and I rationalized it to be unnecessary.

Walton Mendelson

July 14, 2026

© 2026 Walton Mendelson
All rights reserved.

Photoshop, InDesign, and Illustrator are registered trademarks of Adobe Inc. CreateSpace (CS) and Kindle Direct Publishing (KDP) are registered trademarks of Amazon.com Inc. Affinity, Affinity Publisher, Affinity Photo, and Affinity Designer are registered trademarks of Serif Europe Ltd, owned by Canva. Ingram Spark (IS) is a registered trademark of Ingram. Word, and Microsoft fonts are copyrighted by or trademarks of Microsoft Corporation. Pantone is a trademark of Pantone LLC. All trademarks, these and others mentioned in the paper, are the property of their respective owners.

This book, One-Off Press, and the author are independent of any vendor, product, company, or person mentioned in this book. No product, company, service, or person mentioned or quoted in this book has in any way, explicitly or implicitly, endorsed, authorized, approved, or sponsored this paper.

Table of Contents

Why This Exists	7
<i>A synopsis of the study: If your choice is between using sRGB or CMYK images for printing with KDP, and you are not going to do test books and image adjustments, keeping your work in sRGB—and not using PDF/X-1a if that is an option—will yield slightly better color fidelity and much better shadow detail. This can be confirmed visually by looking at test images and empirically through quantitative testing and analysis.</i>	
Part One: How These Numbers Were Produced	9
<i>To understand how the numbers were produced. The major tests are Delta E and clip-rate. The selection of the images to test and compare, how they are prepared for analysis is laid out. In addition to continuous tone images, every page in the grayscale tests has a set of stepwedges, which provide clean data without the noise found in images.</i>	
What Six Proof Copies Can and Can't Tell You	11
<i>Six proof copies is a metaphor for some of the problems of assessing print quality from any print-on-demand printer.</i>	
Part Two: Tone and Detail Are Different Things	13
<i>Highlight and shadow detail, the difference between tone and detail, and the special problems in both printing and analyzing these represent.</i>	
Part Three: Shadow Detail—Mature, Replicated, Confirmed	15
<i>The loss of shadow detail confirmed in how KDP prints, and specifically, more extreme from CMYK art than sRGB.</i>	
Part Four: Highlight Detail—A Newer, Less Mature Finding, With an Important Correction	18
<i>The highlights benefit from using sRGB over CMYK.</i>	
Part Five: Maximum Black in Line Art	21
<i>Throughout the grayscale question runs the issue of color density; and the quest of this study was in large part to determine how best to insure what we upload to KDP is as black as we can make it because KDP does not print 100% K-only black.</i>	
Part Six: Color Fidelity—A Tally, Not an Average	25
<i>To answer the question about color fidelity and if it benefits from sRGB, we made color samples, vector shapes filled with named hex colors that tested both in- and out-of-gamut colors. Without the noise found in continuous tone images, testing showed the advantage of sRGB, although it does not hold consistently through all color families.</i>	
Part Seven: Color Photographs—What Actually Ends Up on a Page	27
<i>A more challenging analysis than the swatches, this section analyzes three full color images.</i>	
Part Eight: What This Means in Practice	30
<i>Based on the observations in this study, this is a quick overview of what can be taken away from it.</i>	
Known Problems, and What You Can (and Can't) Do About Them	32
<i>There were problems that became obvious only because of extensive testing and retesting. Often things that no one would catch even with professional expertise.</i>	
Appendix A	35
<i>All of the images used in the study</i>	
Appendix B	58
<i>The tables of Delta E numbers for all of the color swatches</i>	

Why This Exists


I've argued for years—in proof reviews, in posts, in conversations with other authors and designers—that sRGB outperforms CMYK for KDP print interiors.¹ I don't know of anyone else making that case independently on the CreateSpace or KDP forums; if it's out there, I haven't seen it. So this is my claim to defend, not received wisdom to confirm. Until now it was observational—looking at pages and proof copies and judging. This document replaces that judgment with measurement.

The short version: my advice—“use sRGB”—is a tad better, not a rule. Across every dataset in this project measured patch by patch or image by image—the Macbeth reference chart, eleven grayscale images, seven whole-photograph color comparisons, and 59 individually-scored swatches split across ten hue families and both in- and out-of-gamut sets—sRGB comes out ahead more often than CMYK does, and by enough to call it settled. On the swatch set specifically, the largest and most controlled of the three: sRGB wins 19 of 59 (32%), CMYK wins 16 (27%), the rest (41%) tie—and that gap narrows further once a color falls outside what the press can actually reproduce, where ties on their own more than double. Add the black-and-white shadow-detail work, where sRGB wins seven of ten tested images outright under every metric, splits two more (a desert fence, a foggy mountain ridge) that reverse under a whole-image correlation check, and meets one unexplained exception — and the advice holds. “Use sRGB” is the right default far more often than not; although it isn't the right call in every case, for most of us, our books and our decisions are either it's all sRGB or it's all CMYK—therefore between the two sRGB is the right choice. The rest of this document is the evidence.

The numbers undersell what this looks like in an actual proof copy, especially on the black-and-white side. Set the measurements aside for a moment: looking at sRGB and CMYK printings of the same black-and-white photograph side by side, the gap isn't subtle on anything but a high-key image—it reads as obviously, often dramatically, better in sRGB. That's a qualitative impression, offered honestly as one, not a measurement dressed up as something more rigorous. (The high-key exception lines up with Part Four: the highlight-side advantage is real, but smaller and less consistently dramatic than the shadow-side one.)

A note on which RGB

sRGB isn't the only option, and the natural next question is why not a bigger RGB—e.g., Adobe RGB, ProPhoto RGB—to capture more of what the source art has. Because the comparison that matters isn't sRGB against a bigger RGB, it's sRGB against the actual output gamut. sRGB happens to be roughly matched in overall size to a typical offset



CIE1931 Diagram: Cmglee, CC BY-SA 3.0, via Wikimedia Commons

¹ Tested through KDP only. IngramSpark requires CMYK and doesn't offer an inexpensive way to test-book before committing—the pattern documented here probably holds there, but it hasn't been confirmed.

CMYK print gamut: larger across most of the spectrum, but smaller than CMYK specifically in saturated cyans and yellow-oranges. Whether that theoretical gap shows up as a CMYK advantage in this project's own results is worth checking against Appendix B rather than assumed—the in-gamut cyan swatches, for instance, favored sRGB clearly in this dataset, which doesn't obviously match the theoretical picture. Adobe RGB and ProPhoto RGB extend mostly into greens and cyans the press already can't reach from sRGB either. A wider working space doesn't add printable color, because the press's gamut hasn't changed—it only adds more colors that exist in the file but reach neither the press nor, usually, a screen can display. That's the worst category this project found. Working wider doesn't shrink that risk. It grows it.²

What This Means in Practice—A Short Answer

If you're not going to do test books, analyze shadow detail, or spend time in Photoshop tweaking your images, here's the short version—based on the data in this document, plus on pushing several thousand images printed through these systems over the years:

Black and white: *sRGB has the advantage in roughly nine of ten low- and mid-key images tested—and the other tenth isn't usually CMYK winning outright, it's usually a wash or an unexplained exception. High-key images lean the same direction, on thinner evidence. There are exceptions, as there always are, and they're not predictable enough to plan around.*

Color: *sRGB wins more swatches than CMYK does within the press's reproducible range—14 of 30 in-gamut, against 9 for CMYK, with 7 tied. It isn't a landslide, and it isn't even across every hue. Once a color falls outside what the press can reproduce at all, the advantage mostly disappears, because both modes degrade toward the same poor result more often than either one wins.*

Either way: if you want one answer without doing any of the work above, the answer is sRGB. Submit a PDF/X-3 (or whatever your generic sRGB-to-print PDF workflow produces) rather than converting to CMYK (or PDF/X-1a, which automatically converts to CMYK) yourself.

A further, albeit small, confirmation, from a KDP Customer Support:³

For color “use RGB color space (NOT CMYK).”

For B&W interiors, “Images should be saved in grayscale or RGB . . . not CMYK.”

For PDFs use standard PDF format “not PDF/X-1a”

See Part Eight for more information.

² A rough check of relative gamut area on the CIE 1931 chromaticity diagram puts SWOP CMYK and sRGB within a few percent of each other—close enough, given that the chart is a flat 2D slice of a three-dimensional color volume, to call the two roughly comparable in size rather than one dwarfing the other. This cuts against the commonly repeated “16,000 colors for CMYK vs. 16,777,216 for RGB” comparison, neither side of which holds up: 16,777,216 is simply 256^3 , the count of addressable values in 8-bit RGB—a bit-depth fact, not a measurement of what the eye can see—and the 16,000 figure traces to unsourced print-industry marketing copy. The best available estimate of how many colors a person with normal color vision can actually discriminate, based on analysis of real hyperspectral image data against CIEDE2000 perceptual-difference thresholds, is roughly 2.3 to 2.5 million (Linhares, Pinto, & Nascimento, 2008; Flinkman et al., 2012).

³ April 21, 2026, [Poor Production Print Quality](#)

Part One: How These Numbers Were Produced

Every test in this document follows the same basic shape, whether it's a photograph or a flat color swatch:

1. Get the “before.” The true digital source file—what actually went into the printer, not a re-creation of it.
2. Get the “after.” A scan of the printed page, cropped to just the content being measured.
3. Line them up. The same point in the “before” image has to correspond to the same point in the “after” image before any comparison means anything. For photographs, this used feature-matching software that finds distinctive details (a corner of an eye, a vein in a leaf) in both images and calculates the geometric transform that aligns them. For flat color swatches, simpler box-detection was sufficient, cross-checked against the swatch's own known color value.
4. Check the alignment before trusting it. Two checks, every time: a visual crop comparison (do these two patches show the same thing?), and a difference image (subtract one from the other—a clean alignment shows only thin edge outlines, a bad one shows ghosting or doubled shapes). Several registrations in this project failed this check on the first attempt and were caught and fixed before being used.⁴
5. Measure what changed, in the right way for the question being asked. This document uses three different measurements, because three different things can go wrong:

Color distance (Delta-E).⁵ Convert both images to a color space called CIELAB, where equal numerical distances correspond reasonably well to equal perceived differences, and measure the distance between them. Bigger number, more visible shift. As a rough guide: under 1 is invisible to the human eye, 1–2 is visible only on close inspection, 2–10 is visible at a glance, and above 11 the shift reads as a genuinely different color, not just a shifted one.

Clip-rate (tone survival). Delta-E on an averaged patch can hide the most important failure in dark or light areas: if part of a patch crushes to solid black (or blows to solid white) while the rest survives, the average might barely move even though real information was destroyed. Clip-rate catches this directly: find every pixel darker (or lighter) than a threshold in the true source, then check what fraction of those same pixels in the printed scan got crushed to the floor (or ceiling).

Visual confirmation. Every number that looked unusual in this project—every outlier—was checked by eye against the actual image crops before being reported. Several were caught and dropped or corrected this way.

A methodological note worth stating outright: shadow detail and highlight detail aren't quite the same kind of thing as people usually mean when they talk about “detail.” There turns out to be two different questions hiding under that one word—and getting them confused is exactly what produced one of the more confusing dead-ends in this project before it got sorted out.

⁴ Sometimes images get slightly distorted in size, aspect ratio and skew on the page by KDP in the printing process.

⁵ <https://www.viewsonic.com/library/creative-work/what-is-delta-e-and-why-is-it-important-for-color-accuracy/> ; <https://everycalculators.com/delta-e-calculator.html> ; https://en.wikipedia.org/wiki/Color_difference


A note on the reference stepwedges

Every test page in this project carries three small calibration strips beneath the photograph: a 20-step reference wedge, built once as sRGB, once as K-only, once as full CMYK, printed and scanned along with everything else. They're not decoration—several claims in this document come directly from reading them. One example among several: none of the dedicated grayscale conversions tested (Photoshop, Affinity, GIMP, or XnView) ever reach a true K100 black. They cap around K87–91, every time.

The strips also gave a cleaner version of this document's central finding than any photograph did. The CMYK strip consistently loses all shadow separation from step 13 of 20 onward—about two-thirds down the ramp—while the sRGB and K-only strips hold full separation to the last step. Confirmed on four independently printed pages spanning the project's whole tonal range. Eight of the CMYK wedge's darkest steps become one indistinguishable block of ink. The other two builds never do that, not once.

One finding here earns its own line, because it's the kind you won't catch by looking. The identical "Convert to Grayscale" command in Photoshop, run on the same K-only source file, produced two different results—sometimes true DeviceGray, sometimes full CMYK values, visible both in Photoshop's own Info panel and carried through into the exported PDF. No trigger was ever identified, despite direct investigation. I only caught it at all because I was checking these values closely for this project—not something anyone checks running a book through KDP the normal way.

Which is the real point, and the point of this whole document: you have no control over what KDP's press does to your file once it's submitted—the batch drift and facility variance in Known Problems are proof of that. You do have control over what you send them. Checking that a "grayscale conversion" produced grayscale costs five minutes. No proof copy will announce that problem on its own.



Scan from sRGB



Original Test Page



Scan from CMYK

What Six Proof Copies Can and Can't Tell You

Six is a convenient number, and not quite an honest one. It's the literal count behind the grayscale work—three proofs, two batches—and a different six behind the swatch gamut test, three PDF/X-3 and three PDF/X-1a. Everything else in this document—the hijab portrait, the low-key woman in the hat, the paintings—came from a much larger and looser pool: several dozen proofs, pulled from at least three separate test books, accumulated over years rather than ordered for this project specifically. Treat the title as shorthand for “a deliberately small number of carefully chosen copies,” not as literal accounting. The real question it stands in for is bigger than sample size: how much can any number of printed copies, however carefully chosen, actually tell you?

Start with the part that's just arithmetic. KDP fulfills from multiple facilities, and nothing in this project ever tested whether facility itself is a variable—every comparison here should be read as “within whatever facility filled this particular order,” not as a claim that facility doesn't matter. The batch-to-batch density drift documented elsewhere in this PDF—one batch reading uniformly 30–40% darker than another, across unrelated content—is the same problem showing up in a number you can already see. Traditional offset printing manages that kind of drift with two separate mechanisms: continuous in-process adjustment while the run is happening, and culling bad copies before a customer ever sees them. A true print-run-of-one—which is what most POD orders actually are, even an order of fifty—forecloses both. There's no run to adjust mid-stream and no second copy to cull against. This isn't a claim that POD presses drift more than offset ones did; by most accounts the drift rate is comparable. It's a claim that POD has no mechanism to catch it once it happens. A batch larger than one reopens the possibility, but doesn't guarantee it: a 250-copy order, printed by the sheet rather than copy-by-copy, carried an identical scratch defect across every single copy, uncaught, until a human flagged it after delivery.

One safeguard worth naming on this project's own behalf: I have never ordered only a single proof copy of anything in this document, or in fact, for any test book. Every test got at least three, often more, often across multiple orders. That buys real protection against a fluke—a genuinely anomalous copy gets caught by eye before it gets mistaken for a finding. It does not buy protection against quieter, within-batch variation: only one copy per batch ever went through a scanner and a pixel-by-pixel analysis. Looking at three copies and measuring one isn't the same safeguard twice; it's two different, partial safeguards, and the gap between them is real.

That's the arithmetic. The harder limit isn't about counting at all.

For decades, my standard for whether something printed well was simple: how does it look? Not because I distrust rigor—because I'd already watched a confident, fully-calibrated system turn out to rest on components that weren't stable at all. When I finally bought a shutter-speed tester and checked every lens I owned, actual shutter speeds varied from the marked value by anywhere from roughly half to double or triple—not even consistently in one direction per lens, just unpredictably. The entire elaborate film-developer-paper testing regime built on top of an assumption—that careful calibration produces a reliable “if I do this, I'll get that”—was standing on sand—quick sand. Dan Margulis⁶ made a related point about offset presses in 1996, when

⁶Margulis, Dan. *Makeready: A Prepress Resource*. New York, MIS Press. 1996

desktop publishing and print-on-demand printing were in their infancy: a press, he wrote, is “a powerful, fast, and dangerous combination of machines”—not a precision instrument, and not predictable from one printing to the next. If that was true of fully-calibrated offset, there’s little reason to expect more predictability from a press running one book at a time, unsupervised.

So this document doesn’t trust a number over a look, or a look over a number. It uses both, and says so when they disagree—because they sometimes do, in both directions. This image is a literal demonstration, not a metaphor: the green segments in it are pixel-identical—measured directly, not assumed—whether they sit on blue or on yellow—R00 G166 B82. And yet no one looking at it will believe that without being told, because the segments on yellow read as a darker, richer green, and the segments on blue read as lighter, almost yellow-green. The measurement is correct. The perception is also, in its own terms, correct. Elsewhere in this document, the same kind of gap shows up the other way: a highlight region can carry a clip rate in the 90s—numerically close to total failure—while the printed page it came from doesn’t look like a failure at all, because there was barely anything in that tonal range to lose in the first place.



Neither instrument gets to referee the other. Where this document found them pointing in different directions, it says so, rather than quietly reporting whichever one sounded more authoritative.

Part Two: Tone and Detail Are Different Things

When people say a printed photo “lost highlight detail” or “lost shadow detail,” they’re usually describing one of two genuinely different failures, and the right way to test for one is the wrong way to test for the other.

Tone is a value. A patch of near-white (or near-black) either survives at roughly its original brightness, or it gets crushed toward the paper’s white (or the ink’s black)—and this can happen to a flat area that has no internal structure at all. A bridal portrait’s pale, featureless backdrop losing its gentle gradation is a tone problem. There’s no edge being destroyed, because there was no edge there to begin with.

Detail is an edge—a boundary between two regions that are different enough to be told apart. A cloud’s outline against the sky, or fine texture in fabric, is detail. What gets lost is the boundary itself: it blurs, flattens, or disappears.

In the high-key portrait, the backdrop is the tone example: a flat, near-white field with no edge in it to begin with, the same failure this section’s Tone paragraph describes. The bodice lacing and the seam at the shoulder strap are the detail example—real boundaries, between fabric and skin, that can blur or hold. The low-key portrait repeats the same pairing at the opposite end of the scale. The background is tone, for the identical reason the wedding backdrop is: nothing there to lose. The hat brim’s edge against that background, and the texture across the felt crown, are detail. Both photographs carry their tone and detail examples in a single frame on purpose—the point of this section is that the two are independent enough to succeed or fail separately within one image, not just across different images. (This same low-key photograph returns in Part Three, measured rather than described.)



High Key—Tone & Detail



Low Key—Tone & Detail

These need different measurement tools. Clip-rate (above) is the right tool for tone loss—it doesn’t care whether there’s any texture in the patch, only whether the value survived. Catching detail loss needs something that looks for contrast or edges surviving at a fine scale, which is a different—and, as of this writing, less fully built out—measurement.


These need different measurement tools. Clip-rate (above) is the right tool for tone loss—it doesn’t care whether there’s any texture in the patch, only whether the value survived. Catching detail loss needs something that looks for contrast or edges surviving at a fine scale, which is a different—and, as of this writing, less fully built out—measurement.

This distinction also resolves something that looked like an inconsistency early in this investigation: clip-rate showed massive tonal loss in some high-key images that, on inspection, had nothing one would normally call “detail” in the affected area at all (a plain bright backdrop, a bank of fog). The measurement wasn’t wrong—it was correctly answering the tone question. The mistake was reading a tone-loss number as if it were a detail-loss number, because both get called “loss of highlight detail” in ordinary speech.

The same logic cuts the other way when the affected region does contain real texture rather than a flat field: if a canyon’s rock face or a garment’s fold gets crushed to the floor, whatever edges lived in that tonal range are necessarily destroyed along with it. Clip-rate still can’t tell texture

from its absence directly—but where the source content is independently known to carry detail, tone loss becomes a meaningful, if indirect, proxy for detail loss. That’s the basis for treating clip-rate as evidence in Part Three’s textured photographs, distinct from the flat-backdrop case above.

A further wrinkle, worth stating honestly rather than glossing over: a measured definition of “detail”—does discriminable contrast survive?—is not quite the same as a viewer’s lived experience of a photograph. Two things sit outside what any pixel-level measurement here can speak to. First, simultaneous contrast: how visible a given patch of near-white or near-black tone appears depends on what’s next to it on the page, not on that patch alone—so the same printed dot pattern can read as present or absent depending on surrounding context, something a crop-level measurement can’t capture. Second, the visual system can register contrast below the threshold of conscious discrimination—a region could measure as having no surviving discriminable detail and still be contributing to how rich or deep an image feels to look at, without a viewer ever consciously naming what they’re seeing. Neither of these is something this project’s methodology was built to test, and they’re flagged here as a stated limit on what any of these numbers can claim, not as a problem with the numbers themselves.



Two motif types from the same 2011–12 test book, shown at matched tint stops: highlight side 0–7% K, shadow side 92–85% K. The bold medallion/quatrefoil squares resolve by roughly 7% (highlight) and 85% (shadow); the finer floral scrollwork is still effectively gone at those same percentages. Same tint, two different outcomes—the threshold depends on what’s being asked to survive, not where it sits on the brightness scale.


These extremes are where printing is most mechanically unforgiving. In shadow regions, dot gain—halftone dots spreading under press pressure until adjacent dots merge and detail disappears into solid ink—is the culprit. In highlight regions, the failure runs the other direction: dot loss (also called highlight dropout), where dots too small to hold ink reliably simply fail to transfer, and the area drops to bare paper. Both failure modes are well documented in prepress literature. Standard practice in professional prepress is to measure tone reproduction curves at every 1% from 1–15% in highlights and 85–100% in shadows—exactly the range shown here—precisely because these regions are where the press has the least tolerance for variation. That the bold squares survive and the scrollwork does not at the same tint percentages makes the point directly: detail survival is content-dependent.

Part Three: Shadow Detail— Mature, Replicated, Confirmed

This is the most heavily tested finding in the whole project, and the one with the strongest evidentiary standing: it has been replicated across two independent print batches, across ten separate black-and-white images. Seven win outright under every metric tested. Two—a desert fence and a foggy mountain ridge—win by this section's shadow-mask clip-rate but reverse under the whole-image correlation check reported in the coda after Part Four; both numbers are reported there in full, not smoothed over here. One, a hijab portrait, is an outright exception with no established mechanism.

sRGB consistently preserves shadow detail better than CMYK in continuous-tone black-and-white photographs printed through KDP's black-and-white interior option (uncoated stock).


The clearest single example is a canyon landscape photograph:



Scan from sRGB




Original Test Image



Scan from CMYK

This proof copy shows real shadow compression in the CMYK panel—the finding this image is here to illustrate. It also carries two unrelated printing defects: a faint vertical streak, and a position-locked density split roughly 1.75" from the left edge, where the left portion prints darker than the right. Neither is a shadow-detail effect. Both are confirmed present on every copy tested, in both batches—not a fluke of this particular scan—and are documented separately in Known Problems. The compression difference between panels is the finding; the streak and the density split are not.

	Batch 1	Batch 2
sRGB clip-rate	10.3%	23.2%
CMYK clip-rate	47.2%	68.3%
Gap	+36.9 points	+45.1 points



Scan from sRGB



Original Test Image



Scan from CMYK

The gap didn't just persist between print runs—it widened. That's a stronger signal than simple consistency would be: if the difference were print noise, a second run would be as likely to narrow the gap as widen it. A widening gap across batches is closer to what a real, reproducible content-dependent effect should look like.

The exception:



Scan from sRGB



Original Test Image



Scan from CMYK

Here the pattern reverses—the black headscarf and veil crush harder in the sRGB scan than in the CMYK scan: clip-rate 86.9% (sRGB) versus 60.7% (CMYK) against the same threshold used throughout this section, mean shadow value 8.3 versus 13.3. (Checked against the source image in this composite directly, not run through the full feature-matching alignment behind the canyon table—indicative, not held to that

same standard.) No mechanism for the reversal is established. It is the exception out of scores of test books and thousands of printed images. It's shown here rather than held back because "zero exceptions" only means something if it's zero exceptions found and reported, not zero because the one that broke it got filed somewhere quieter. Or simply put, stuff happens.

A few supporting findings from the same investigation, lower-stakes but worth stating for completeness: the four common ways of specifying “black” for editable text print indistinguishably from each other; a 100%-black solid swatch prints substantially darker than any text-black at the same nominal value (a feature-size effect—thin strokes and solid fills behave differently under any print process, not a sign that the press handles “black” inconsistently); and all color modes converge to statistically the same result at full black saturation. That's a claim about printed density, not file specification—it doesn't conflict with Part One's finding that grayscale conversions cap their ink-build spec around K87–91 rather than true K100. An 87% K-only build and a 300%-ink rich black are different specs on paper that can still print as the same solid black; this is that case. Color-mode divergence in black-and-white printing is specifically a midtone/shadow phenomenon—it does not show up at the black end of the scale at all.

Standing caveat: absolute print density is not stable between batches (a second print run measured uniformly 30–40% darker across unrelated content). This does not disturb the shadow-detail conclusion, because the shift was uniform and didn't change the relative comparison—but it's a reminder that absolute numbers from any single batch should never be treated as a fixed target.

Part Four: Highlight Detail—A Newer, Less Mature Finding, With an Important Correction


Highlight clip-rate testing came later in this project than shadow testing, using the same mirrored methodology on three black-and-white high-key images from the same source book: a bridal portrait, a desert landscape with a fence line, and a foggy mountain ridge.




Scan from sRGB



Original Test Image




Scan from CMYK




Scan from sRGB




Original Test Image




Scan from CMYK



Scan from sRGB



Original Test Image



Scan from CMYK

Image	Highlight mask = % of frame	sRGB clip-rate	CMYK clip-rate
Bridal portrait	92%	93.4%	97.9%
Desert landscape	23%	63.8%	70.8%
Fog/mountain	52%	63.9%	79.4%

CMYK clipped harder than sRGB in all three, mirroring the shadow-side direction.

The correction: the bridal portrait’s mask covers 92% of its frame—meaning that image was mostly already at or near white in the true source to begin with. Per the tone/detail distinction above, that’s a tone-loss measurement on an image that had very little distinct highlight detail to lose in the first place; it’s testing something closer to “did this already-bright image get pushed even brighter,” not “did fine highlight texture survive.” The fence and fog images are the cleaner illustration of genuine highlight loss, because they have real tonal range and, in the fence image’s case, actual edges (cloud shapes) within the highlight zone.

Standing caveat: this is one batch, three images—the same evidentiary stage the shadow-detail finding was at before its second-batch replication earned it confirmed status. The direction is consistent and worth reporting, but should not be presented with the same confidence as Part Three.

A Parallel Body of Evidence: Six More Subjects, a Different Metric

Six more black-and-white photo subjects were scanned in the same session as everything above—same scanner, same settings, changed only for the color/grayscale toggle. They weren’t measured with the shadow-mask clip-rate used everywhere else in this section; the metric here is whole-image correlation between each colorspace scan and the digital original—overall fidelity to the source, not specifically what survives in the shadows—using a different tool. Treat all of it as a

second, parallel line of evidence, not a replication of the numbers above. See Appendix A for all the images analyzed in this study.

Subject	sRGB corr.	CMYK corr.	Winner
Woman with bracelet	0.962	0.835	sRGB
Woman with cut eyebrow	0.977	0.945	sRGB
Canyon landscape	0.859	0.846	sRGB (barely)
Desert with fence	0.860	0.871	CMYK
Foggy mountain ridge	0.939	0.968	CMYK
Canyon desert	0.904	0.943	CMYK

Two of these—the bracelet and cut-eyebrow portraits—favor sRGB clearly and aren't discussed elsewhere in this document. The other four overlap with images tested earlier by the shadow- and highlight-mask clip-rate, and measured this way, some of them lean the other direction.


The canyon landscape still favors sRGB, but where the clip-rate gap was dramatic (37–45 points across two batches), the correlation gap nearly vanishes: 0.859 to 0.846. Same photograph, same direction, very different size of effect.

The fence and fog photographs reverse outright. By clip-rate, sRGB won both; by whole-image correlation, CMYK edges both—0.871 to 0.860 on the fence, 0.968 to 0.939 on the fog—and both at high confidence. The canyon-desert result also favors CMYK, but its sRGB registration is low-confidence and shouldn't be trusted without a manual check.

These aren't contradictions to explain away. Clip-rate and whole-image correlation ask different questions—one measures what survives in the shadows, the other measures overall fidelity to the source—and there's no reason a photograph has to win both. All four of the images that lean CMYK here are high-key or low-contrast, the same content category Part Four already flagged as less settled than Part Three's shadow work. That's a pattern worth noting, not a finding: "these images behave differently under a whole-image metric" is a smaller, safer claim than "CMYK wins on high-key content," which this data doesn't establish.

Part Five: Maximum Black in Line Art—A Single Finding, Not Yet Replicated

The obvious way to get the darkest possible line in 1-bit line art sounds like it shouldn't need testing: select the black strokes, fill with 100% K, confirm with the color picker. It was tried, verified two different ways—the histogram and the color picker both confirmed a genuine 0% gray at the checked point—and printed anyway.



It printed lighter.

Source	min	0.1 pctl	0.5 pctl	1 pctl
Corrected file (digital)	0	0.0	0.0	0.0
Corrected file, scanned print	43	63.0	72.0	87.0
Plain sRGB print	29	52.0	59.0	63.0
Plain CMYK print	32	52.0	59.0	64.0

A file that measured as genuinely, unambiguously black at the source came out 10–24 gray levels lighter than two ordinary, unenhanced prints of the same art, printed in the same batch—ruling out press or paper drift as the explanation.

The cause traces to how the correction was built, not to the correction itself failing on its own terms. Measuring ink-crossing width and count on matched crops explains it:

Source	Ink crossings	Mean width	Median width
Corrected print	19,428	11.80px	6.0px
Plain sRGB print	53,280	18.15px	4.0px
Plain CMYK print	55,518	17.64px	4.0px

The corrected version has roughly a third the ink-crossing count of the plain prints on matched crop size, and its crossings are narrower on average as well (11.80px mean against ~18px)—less ink by both measures, which is the direct explanation for why it prints lighter. The one place the corrected print runs higher is median width (6px against 4px), and that split is itself telling: the plain prints pair a low median with a much higher mean, meaning most of their crossings are thin but a long tail of wide ones pulls the average up. That wide tail is the soft, partial-coverage anti-aliased edge—present and printing in both plain versions, because that's simply how anti-aliasing renders. The corrected print's tighter, more uniform distribution is what's left after the color-range selection caught only the already-near-black core of each stroke and never touched that edge.


Worth stating plainly for anyone reaching for this as a fix: “select the black pixels, fill with 100% K, verify with the color picker” is not sufficient on its own for 1-bit line art with anti-aliased strokes. Done this way, it can measurably reduce total ink coverage relative to doing nothing—even while genuinely, verifiably fixing the specific point you checked. It's the kind of finding that looks like it must be a mistake until the ink-crossing numbers make the mechanism visible.

This is one test, one batch, not yet replicated—and unlike every other test page in this document, it doesn't carry the sRGB/K-only/CMYK reference stepwedge described in Part One. Late edits shifted the page layout and pushed the stepwedge off before print, and it wasn't caught in time. The test itself still has its own controls: same batch, same day, matched crop sizes, a digitally-verified 0% gray at the source. That same-batch pairing also rules out the batch-to-batch density drift documented elsewhere in this project—swings of 30–40% between print runs. What's left is ordinary page-to-page variation within a run, which two sample stepwedge pages elsewhere in this book measured within about 2% of each other—not a full survey of this book's roughly sixty pages, but not nothing. What hasn't been tried is the proposed alternative: converting to grayscale first, then applying the same Levels-verification discipline to the whole stroke's tonal range, rather than the color-range selection that caught only the anti-aliased core here. That's still the highest-priority follow-up open in this document.

A Note on Resolution and Screening in Line Art

Two smaller findings from the same ornament test, at different evidentiary tiers—one solid, one a single unreplicated data point worth reporting honestly as such.


Source resolution stops mattering earlier than intuition suggests. The same 1-bit ornament shape was rendered as a vector, and rasterized at 1200, 800, and 300 dpi, then printed and scanned. Edge sharpness—measured as the pixel width of the transition from white to black across many sampled edges per crop—came out like this:



Original Grayscale Art

Scans

Enlarged Detail of Scan



Original Grayscale Art

Scans

Enlarged Detail of Scan

Source	Mean edge width (px)	Median
Vector	20.42	21.0
1200 dpi	20.41	20.0
800 dpi	21.35	20.0
300 dpi	24.12	29.0

Vector and 1200dpi are indistinguishable once printed and scanned—0.01px apart, well within noise. 800dpi shows only marginal softening, about 5%. 300dpi is the first resolution with a real, measurable loss—18% wider transitions than vector, and the jump in the median suggests uneven stair-stepping rather than uniform blur. For this kind of 1-bit line art, resolution above roughly 800–1200dpi buys nothing once it goes through print and scan—the press and scan process are the bottleneck, not the file.

A 2014-vs-2026 comparison, on the other hand, is one sample and should be read as suggestive, not settled. The same edge-width measurement applied to a 2014-printed version of this ornament—same shape, same three shared resolutions—came out sharper on average, not softer: 17–24% narrower transitions across the board. Taken at face value, that runs against the general assumption that screening technology has improved since 2014.




But the two eras don't just differ in average—they differ in shape. The 2014 measurements are heavily skewed: the median (11.0px, identical across all three resolutions) sits far below the mean, with a standard deviation almost as large as the mean itself. The 2026 measurements track much tighter—mean and median close together, much smaller spread. That's consistent with two different underlying mechanisms: a periodic, grid-based (AM) halftone dot pattern, where edge sharpness depends on how a curve's angle happens to align with the dot grid—sharp in some places, soft in others, averaging out to a skewed distribution—versus a stochastic (FM) dot pattern, which doesn't care about angle and produces more consistent edges throughout.


Which framing matters more—2014 was sharper on average, or 2026 is more predictable—depends on what claim is being made, and this document isn't taking a side on that yet. One sample from each era isn't enough to confirm this is a real AM-to-FM signature rather than a one-off. It's reported here as a real, observed pattern worth watching, not as a conclusion.

Part Six: Color Fidelity—A Tally, Not an Average


Fifty-nine swatches, ten hue families, each scored in-gamut or out-of-gamut against its own labeled hex value—the controlled counterpart to Part Seven’s real photographs. Method: CIEDE2000, verified against the Sharma et al. (2005) reference cases before use. No averaging, at any level—every swatch is reported and counted individually; a “wins” tally is a count, not a mean. Ties are real and reported as their own category, not folded into either side: two ΔE_{2000} values within 1.0 of each other are treated as indistinguishable rather than arbitrarily awarded to whichever is numerically smaller—1.0 being the threshold most consistently cited across color-difference literature as the point below which a shift stops being reliably visible..




Scan from PDF/X-3—sRGB




Original Test Image in sRGB




Scan from PDF/X-1a—CMYK



Scan from PDF/X-3—sRGB



Original Test Image in sRGB



Scan from PDF/X-1a—CMYK

	sRGB wins	Tie	CMYK wins
In-gamut (30)	14	7	9
Out-of-gamut (29)	5	17	7
All (59)	19	24	16

Moving from in-gamut to out-of-gamut, sRGB's win count drops sharply—from 14 to 5, well over half—while CMYK's holds up better, easing only from 9 to 7. The real story is still the tie column: it more than doubles, from 7 to 17, becoming the single largest outcome out-of-gamut. Once a color falls outside what the press can physically reproduce, the two pipelines increasingly land in the same place: neither wins, because neither can—and sRGB has more to lose from that collapse than CMYK does, since more of its edge was concentrated in colors the press can't actually hold onto.

That collapse isn't uniform across hues. Green and cyan both favor sRGB cleanly in-gamut—but out-of-gamut they diverge: cyan collapses completely to ties (all 3), while green splits evenly across all three outcomes. Orange and red favor CMYK in-gamut and keep that lean out-of-gamut—the most stable pattern in the dataset. Blue also favors sRGB cleanly in-gamut and, unlike green and cyan, mostly holds that lean out-of-gamut too (2 of 3 still sRGB). Purple and pink don't show a clean in-gamut lean at this threshold to begin with; purple fades to a complete out-of-gamut tie, most visibly in Dark Pastel Purple (13.95 vs. 13.73) and Brilliant Lavender (14.94 vs. 14.88)—both the batch's worst performers for either pipeline and functionally identical to each other. Grays and browns sit close to a wash at every gamut position, the expected null result. Yellow shows a smaller, steadier sRGB-or-tie lean at both gamut positions—CMYK never wins a yellow swatch outright, in-gamut or out.

The practical reading: which pipeline wins depends on where in the color wheel a given color sits, and that dependency itself fades out once the color is one the press can't make cleanly regardless of format. This is diagnostic information for a specific troublesome image, not a book-level formatting rule—see Appendix B for the full per-swatch table.



Part Seven: Color Photographs—What Actually Ends Up on a Page

Everything in Part Six was isolated color patches—useful for isolating the mechanism, but nothing on a real cover or in a real photo book looks like a named swatch. This section is closer to that: seven color photographs and a 24-patch reference chart, scanned in the same session as the swatch work, same scanner, same settings except the color/grayscale toggle. Given that most book covers are color, and most image-heavy interiors are color, this is arguably where the whole project's findings matter most in practice—and it's also the newest, least-replicated data in the document. One batch. Worth saying plainly rather than letting the stakes imply more certainty than the evidence has.

A methodology note before the numbers: this section measures color distance with CIE76 Delta-E, not the CIEDE2000 used for the Part Six swatches. The two formulas don't produce comparable magnitudes—a CIE76 value here and a CIEDE2000 value in Part Six aren't on the same scale, even when they're measuring the same kind of shift.

Seven photographs, whole-image Delta-E against the true digital original:

Image	Conversion type	sRGB ΔE	CMYK ΔE	Winner
Montage	Full CMYK, perceptual intent	12.19	16.83	sRGB
Woman with leaf	Full CMYK	10.65	9.02	CMYK
Rossetti painting	Full CMYK	17.03	25.03	sRGB
Hijab portrait	Unconfirmed type	7.69	9.17	sRGB (see Part Three)
Hat	K-only	18.24	25.71	sRGB (low confidence)
Pier	K-only	8.21	10.26	sRGB
High-key woman	K-only	7.36	9.68	sRGB




Scan from sRGB

Original Test Image

Scan from CMYK

Hijab portrait: also carries the shadow-detail exception discussed in Part Three. Hat: low registration confidence—see appendix.

sRGB wins six of seven. The exception is the leaf photo, and it comes with its own history: this specific image required manual correction after an on-press geometric distortion was found between its sRGB and CMYK pages in an earlier print run. This new scan registered cleanly at high confidence, which argues against a severe repeat of that problem—but a clean registration isn’t the same as a manual visual check, and this is the one CMYK win in the batch. Worth a direct look before it’s cited as a clean counterexample rather than a residual print artifact.



Three of the seven—hat, pier, high-key woman—used K-only CMYK conversion rather than full four-color: a more constrained transform, closer in spirit to the K100-ceiling work in the black-and-white sections than to a full CMYK build. Individually, their CMYK penalties run +41%, +25%, and +32% over sRGB. The two full-CMYK photos not affected by the leaf’s distortion history run +38% and +47%. Averaged, the full-CMYK group’s penalty is actually the larger of the two—the opposite of what a same-day earlier read of this data concluded. On five images split three and two, that’s not a real group-level finding either way; it’s flagged here to correct the record, not to replace one unsupported claim with another.

The Macbeth chart—24 patches, same session:


	sRGB avg ΔE	CMYK avg ΔE	sRGB wins
All 24 patches	20.31	22.74	13/24
Color patches (18)	24.26	27.44	—
Neutral patches (6)	8.47	8.62	—

Direction matches an older Macbeth dataset already in the project record—sRGB ahead by a modest margin, neutrals far closer than colors in both. Absolute ΔE runs notably higher across the board in this new scan. That’s a real difference, not noise, and it’s not resolved here: could be scan-and-print generation variance, could be a genuinely different chart printing. The qualitative

pattern replicates; the specific numbers from the two scans shouldn't be treated as the same measurement.



Scan from sRGB



Original Test Image



Scan from CMYK

Read together with Part Six: the swatch work showed sRGB ahead in-gamut, converging toward parity once a color leaves what the press can reproduce. This photograph-level data doesn't contradict that—it extends it into the kind of content a cover or a photo book actually contains, where sRGB's advantage held on six of seven whole images. That's the closest thing in this document to a direct answer for the reader who just wants to know what to do with an actual book cover.

Part Eight: What This Means in Practice

If you're choosing one default and not planning to test further: submit sRGB—a PDF/X-3, or whatever your workflow's generic sRGB-to-print export produces—rather than converting to CMYK yourself. Not a universal rule, and this document doesn't claim it as one. Across every measurement here, it's the safer default more often than not. **Note:** I am not suggesting that if you do adjust your image, you use CMYK. I would stay in sRGB regardless. The exception would be if you work in CMYK, then I would stay in CMYK. If you submit work to Ingram Spark, it should be CMYK and PDF/X-1a. If you are working with any other printer ask for their specs.


1. **Black-and-white interiors.** Shadow detail favors sRGB—the strongest evidentiary standing in this document, replicated across two batches, ten images: seven win outright under every metric tested, two (a desert fence and a foggy mountain ridge) win by this project's primary shadow-detail metric but reverse under a whole-image correlation check, and one is an outright exception with no established mechanism. Highlight detail leans the same direction, but more softly: it's a newer finding, one batch—and that same correlation check, on those same two fence and fog photos, is what reverses it there too. Treat the highlight lean as real but not settled.

2. **Color.** In-gamut, sRGB comes out ahead more often than not. Out-of-gamut, the two pipelines converge toward parity —outright ties more than double once a color leaves what the press can reproduce, and neither mode reliably wins once gamut itself is the limit. There's a real hue-level pattern behind those numbers—some families lean sRGB, some lean CMYK, some collapse to a wash—but it isn't book-level advice, because a book isn't one hue. It's diagnostic: if a specific image's color is giving you trouble, that pattern (Part Six's appendix) is where to check whether the hue itself is the reason, before assuming the export format is.

3. **If one image is the actual problem,** a targeted proof of that image beats switching the whole manuscript's format. This is true in both directions—for an out-of-gamut color, no format will fix it; for an in-gamut color behaving worse than expected, a second export or a proof will tell you more than guessing.

4. **Curves can partially recover compressed tone, within limits**—data that still exists but sits crushed near the floor or ceiling. It cannot restore detail that was never resolved at the file or print-dot stage, and it cannot un-clip a color that fell outside the press's gamut to begin with. Full treatment, including a case where the obvious fix didn't do what it looked like it should, is in Known Problems.

This is not a how-to book, but it makes sense to show what curves can do. The only exception in these examples, where, working in Affinity, I made a selection and it did not respond to a Curves correction, but did to Levels The point is that we have control over highlight and shadow detail if there is some there in the first place.



Known Problems, and What You Can (and Can't) Do About Them

This project surfaced several distinct failure modes, not just the headline color-mode question. Some of them an author can compensate for after the fact. Others are properties of the print execution itself, sitting entirely outside the file you submit—no setting, no edit, no amount of care on your end changes them. Worth being honest about which is which, rather than leaving a reader to assume every problem in this document has a fix.

Compensable, within limits.

Shadow and highlight tone compression—documented in Parts Two through Four—can be partially clawed back with a curves adjustment, if the underlying data still exists and is merely compressed near the floor or ceiling rather than genuinely destroyed. What curves cannot do: restore detail that was never resolved at the file or print-dot stage to begin with, or un-clip a color that fell outside what the press can physically reproduce (see Part Six, where out-of-gamut ties more than double—neither mode wins reliably once gamut itself is the limit). Tone is recoverable. Resolution and out-of-gamut color are not.⁷ These are different failure modes, and only one of them responds to editing.

A caution belongs here too, not as a second technique but as a case where the obvious one didn't do what it looked like it should. Line art's strongest-possible-black test (Part Five) read as a straightforward compensation move—select the black strokes, fill with 100% K, verify with the color picker—and it genuinely, verifiably reached 0% gray at the checked point. It printed lighter than doing nothing at all. The correction wasn't the failure; the selection was too narrow, catching only the already-dark core of each anti-aliased stroke and leaving the softer edge untouched. Worth stating directly: a fix that verifies correctly on screen doesn't automatically behave the way it looks like it should once it reaches the press. Confirm on paper, not just in the file—the same standing caveat behind every finding in this document.

Not compensable—these are properties of the print, not the file.

Sample size and facility variance. Every measurement in this document comes from six proof copies across two batches. KDP fulfills orders from multiple facilities, and this project never tested whether facility is itself a variable—it can't claim facility-uniformity, and shouldn't be read as having shown it.

The POD quality-control gap. Traditional offset printing manages drift through two distinct mechanisms: continuous in-process checking with running ink-flow adjustment during a run, and post-run culling of bad copies before a customer sees them. A print-run-of-one—which is what most POD orders actually are, even a 50-copy order, produced one complete book at a time—forecloses both: there's no run to adjust mid-stream and no second copy to cull against.

⁷ Many times you can find an in-gamut color that looks like, even “matches,” the out-of-gamut one. This is usually trial and error, and can take awhile to find.


Batch size greater than one reopens the possibility of catching a defect mid-run, but doesn't guarantee it: a documented 250-copy order, printed by the sheet rather than copy-by-copy, carried an identical scratch defect across every single copy, uncaught, until a human flagged it after delivery. The batch size was there. The correction didn't happen anyway.

Two-page spread and gutter misalignment. Full-bleed images spanning the gutter printed cleanly for two decades under BookSurge and CreateSpace. Under KDP, they're unreliable—perhaps one or two in five comes out right, by rough estimate, never formally counted. KDP's own customer support has given mutually contradictory answers about why: that two-page images can't be printed at all, that KDP adds the inside white strip so don't add your own, that you should add it yourself, that KDP's handling of it is inconsistent, and that only one facility—reportedly handling only books grandfathered in from CreateSpace—actually knows how to do it correctly. None of that is independently verifiable; it's what support has said, not documented policy. What is independently corroborated: a public KDP Community thread from several years ago, with another author's tests matching this project's own, describing the same gutter-strip inconsistency.


Streaking and banding. Visible vertical banding—a thin line with a faint density shift on either side, or in worse cases a wide band suggesting several adjacent print-head nozzles affected—has shown up in roughly two or three of every five test books, also never formally counted. In twenty-five years printing color across BookSurge, CreateSpace, and KDP, this has not been observed on the AM-screened color books (BookSurge Color and CreateSpace Color, now KDP Premium Color)—though that's a personal track record, not a claim that it's structurally impossible there. It has been observed on the FM-screened side. Worth naming precisely: which press technology produces which screening pattern turns out not to be reliably knowable—even outside investigators looking closely at ostensibly identical settings have found inconsistent results—so this document anchors the claim to the screening pattern itself, which has actually been verified, rather than to a press-technology that hasn't.

Appendix A


Grayscale Printing




Scan from sRGB




Original Test Image




Scan from CMYK



Original Grayscale Art



Scans




Scan from sRGB




Original Test Image




Scan from CMYK




Original Grayscale Art




Scans




Scan from sRGB




Original Test Image



Scan from CMYK



Original Grayscale Art




Scans




Scan from sRGB




Original Test Image



Scan from CMYK



Original Grayscale Art




Scans



Scan from sRGB




Original Test Image



Scan from CMYK



Original Grayscale Art



Scans



Scan from sRGB




Original Test Image



Scan from CMYK



Original Grayscale Art



Scans



Scan from sRGB



Original Test Image




Scan from CMYK



Original Grayscale Art




Scans




Scan from sRGB




Original Test Image




Scan from CMYK



Original Grayscale Art




Scans




Scan from sRGB




Original Test Image




Scan from CMYK



Original Grayscale Art




Scans




Scan from sRGB




Original Test Image




Scan from CMYK



Original Grayscale Art




Scans




Scan from sRGB




Original Test Image




Scan from CMYK




Original Grayscale Art




Scans




Original Grayscale Art




Scans




Enlarged Detail of Scan




Original Grayscale Art




Scans




Enlarged Detail of Scan




Scan from sRGB





Original Test Image



Scan from CMYK





In Affinity

Upper Left: Pulled in guidelines to help locate a consistent reading.

Upper Right: Center the Color Picker Tool over the guidelines. The reading is Gray: 0%, maximum black.


Lower Left: In the Levels Panel push the black point/slider to at least the initial appearance of black in the Histogram window. You could also select the numerical value and move it up and down. It is at this point where the maximum black is 0%.

Converting to Grayscale first simplifies this.

PremiumColor




Scan from sRGB



Original Test Image




Scan from CMYK



Scan from sRGB



Original Test Image




Scan from CMYK



RGB

42

Scan from sRGB




Original Test Image



CMYK

43

Scan from CMYK




RGB

46

Scan from sRGB




Original Test Image



CMYK

47


Scan from CMYK



Scan from sRGB



Original Test Image



Scan from CMYK



Scan from sRGB



Original Test Image




Scan from CMYK



Scan from sRGB



Original Test Image



Scan from CMYK



Scan from sRGB




Original Test Image




Scan from CMYK


Named Hex Colors—PremiumColor




Scan from PDF/X-3—sRGB




Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB




Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB




Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB



Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB



Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB



Original Test Image in sRGB



Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB




Original Test Image in sRGB




Scan from PDF/X-1a—CMYK




Scan from PDF/X-3—sRGB




Original Test Image in sRGB



Scan from PDF/X-1a—CMYK



Scan from PDF/X-3—sRGB




Original Test Image in sRGB



Scan from PDF/X-1a—CMYK



Scan from PDF/X-3—sRGB



Original Test Image in sRGB



Scan from PDF/X-1a—CMYK

Appendix B

Hex Color Swatches—Analysis

Every swatch behind Part Six's tally, in the order the swatch images appear earlier in this appendix. ΔE_{2000} measured against each swatch's labeled hex value; winner assigned using this document's 1.0 tie threshold (see Part Six).

Cyan	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Tiffany Blue	in	5.11	8.20	sRGB
	Powder Blue	in	2.44	6.61	sRGB
	Pale Aqua	in	1.89	6.27	sRGB
	Turquoise	out	9.66	9.42	Tie
	Electric Blue	out	11.68	12.03	Tie
	Light Cyan	out	6.71	6.56	Tie
Green	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	MSU Green	in	4.02	9.92	sRGB
	Mint	in	5.30	6.99	sRGB
	Eton Blue	in	1.50	5.52	sRGB
	UP Forest Green	out	5.48	9.99	sRGB
	Caribbean Green	out	11.74	9.28	CMYK
	Moss Green	out	5.40	6.36	Tie
Blue	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Dark Slate Blue	in	3.76	6.51	sRGB
	Steel Blue	in	3.18	5.80	sRGB
	Baby Blue	in	2.52	6.34	sRGB
	UA Blue	out	7.65	7.09	Tie
	Azure	out	7.74	12.04	sRGB
	Pale Blue	out	7.35	8.70	sRGB
Brown	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Dark Taupe	in	4.85	7.50	sRGB
	Dark Tan	in	2.63	3.36	Tie
	Raw Sienna	in	5.64	3.81	CMYK
	Seal Brown	out	7.90	12.42	sRGB
	Sandy Taupe	out	5.38	3.59	CMYK
	Sand Dune	out	5.32	4.43	Tie
Gray¹	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Charcoal	in	4.36	7.20	sRGB
	Dark Gray	in	2.38	2.51	Tie
	Timberwolf	in	2.33	1.10	CMYK
	Smokey Black	out	2.72	2.59	Tie
	Gray	out	1.92	2.33	Tie

¹5 swatches—no third out-of-gamut color was selected for this family. The group is small, and there were only two out-of-gamut colors.

Orange	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Dark Coral	in	9.62	6.60	CMYK
	Bronze	in	7.39	2.74	CMYK
	Indian Yellow	in	3.91	2.68	CMYK
	Mahogany	out	10.00	4.28	CMYK
	Fluvous	out	5.03	3.84	CMYK
	Orange Peel	out	2.22	2.93	Tie
Pink	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Cerise	in	9.88	8.18	CMYK
	Coral Pink	in	4.80	5.70	Tie
	Pastel Pink	in	5.04	6.56	sRGB
	Bright Pink	out	9.73	10.25	Tie
	Bittersweet	out	7.37	5.95	CMYK
	Classic Rose	out	6.55	6.49	Tie
Purple	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Dark Lavender	in	4.12	6.44	sRGB
	Lavender Purple	in	7.30	6.72	Tie
	Thistle	in	5.24	4.87	Tie
	Royal Purple	out	7.65	7.96	Tie
	Dark Pastel Purple	out	13.95	13.73	Tie
	Brilliant Lavender	out	14.94	14.88	Tie
Red	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Persian Red	in	11.01	6.14	CMYK
	Brick Red	in	10.86	6.39	CMYK
	Razzmatazz	in	10.14	8.48	CMYK
	Rufous	out	7.84	3.68	CMYK
	KU Crimson	out	7.71	4.85	CMYK
	Rose	out	9.58	9.93	Tie
Yellow	Swatch	Gamut	sRGB ΔE	CMYK ΔE	Winner
	Satin Sheen Gold	in	3.14	2.42	Tie
	Dandelion	in	1.62	2.90	sRGB
	Light Khaki	in	1.80	2.24	Tie
	Chrome Yellow	out	1.86	2.67	Tie
	Canary Yellow	out	3.13	4.63	sRGB
	Pastel Yellow	out	5.36	5.77	Tie

