Why does my cyan have the blues?[1]

John Seymour

When I started in the print industry as an apprentice to Gutenberg, I noticed that the folks in the press room called the inks red, yellow, and blue. This confused me. Everything I had read in color theory books said that cyan, magenta, and yellow were the subtractive primaries. These were the primaries that you use to make a wide range of colors with pigments and filters. Pigments and filters work by subtracting certain wavelengths of light. On the other hand, red, green, and blue were the additive primaries, and these were used to make all the colors when you are mixing light, as in a TV or computer monitor.

Polaroid snapshot of me working at my first job

Why were those silly printers using some of the additive and some of the subtractive primaries? Didn't they realize that this reduced their gamut? That was the theory, anyway[1].

Just a naming issue?

Anyone who knows me, or who loves me[2] can attest to the fact that I am a firm believer that ignorance is the main explanation for every cultural and scientific phenomenon. In this case,my previous blog about counting colors provides a clue as to the sort of ignorance that might explain why magenta is so curiously called red.

The eleven people who read my previous blog learned that there are only eleven basic one-word color names in our active vocabulary. Neither cyan nor magenta made that list[3]. Clearly the folks on press were calling the inks "red" and "blue" because they have no other words to describe the colors.

Cyan ink is blue, and magenta is red

In my normal incisive way, it took me a few years to realize that the pressmen were not quite as ignorant as I thought they were. I guess I spent too much time running for buckets of halftone dots to actually put my head in a bucket of ink. When I finally did put my head in a bucket of ink (as part of a hazing[4] experiment) I could see that cyan ink is blue, and that magenta ink is red when you look at them in a bucket.

---

[1] *This article appeared in FlexoGlobal eZine, Feb 12, 2013*

Cyan and magenta inks are blue and red in the can

Cyan and magenta inks are cyan and magenta on paper

So, this confusion is obviously beyond my original explanation. Just like when a fellow accidently calls his wife by the name of a former girlfriend, you can bet there is something deeper going on.

Beer's law revisited

In yet another very popular[5] blog of mine, I provided a charming explanation of Beer's law. This blog post is a prerequisite for the following exciting discussion.

Let's just say that we have a perfect magenta ink. A perfect magenta ink will reflect all the red light and all the blue light that hits it. As for the green light, a light shade of magenta might reflect about 10% of the green. A rich shade of magenta will reflect about 1%.

Now we bring in Beer's law. Let's say we start with that light magenta and add another layer of the same ink. Beer's law would predict that the reflectance would multiply. Since perfect magenta reflects 100% of red and blue light, Beer's law predicts that the double layer of magenta will reflect 100% of the red and blue light. Beer's law would further predict that the green light would reflect at only 10% X 10%, which is 1%. A double layer of light magenta becomes a rich magenta.

Key point here: for this perfect magenta ink, the hue is still that of magenta. It still reflects most of the red and blue light, and absorbs most of the green light.

Let's just say that we now switch over to a magenta that is less pure. Let's just say that for some inexplicable reason, the publishers of *Schlock* magazine are unwilling to spend $100,000 per gallon for their ink. The bargain ink they decide to use does not reflect quite as much blue light as we would hope; maybe it only reflects 40% of the blue light when we put a thin film down, and maybe 10% of the green light. Let's say that the red light is still reflected at 100%.[6]

What happens when we double the amount of ink on the paper? Beer's law takes over, and we see that blue light is reflected at 40% X 40% = 16%. Green light? The reflectance goes from 10% down to 1%. Red light stays at 100%. The table below summarizes the Beer's law estimation.

	Blue	Green	Red
Thin layer	40%	10%	100%
Thick layer	16%	1%	100%

From this table, it would *seem* that the thick layer of magenta is a lot closer to red. The plot below shows the actual spectra of two magenta patches, one at a larger ink film thickness than the other. The plot leads one to the same impression – that a thick layer of magenta is closer to red in hue than a thin layer.

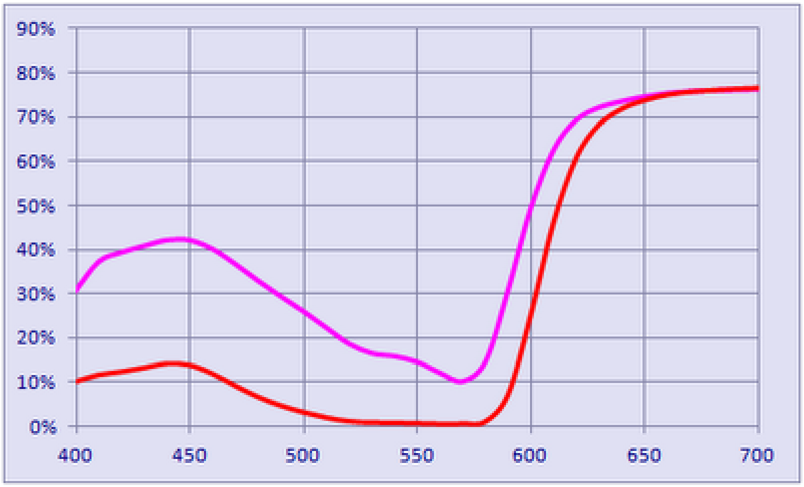

Spectrum of a magenta ink, normal thickness and thick

The tentative conclusion is that magenta turns red when it is thick because it is impure, or more accurately, because there are several different reflectance levels in the spectrum. When Beer's law kicks in, the areas of the spectrum where the reflectance is "mid-level" (i.e. 40% reflectance) are grossly effected by the ink film thickness.

The plots below are the spectra of cyan and yellow inks. If the previous rule applies, then we would expect that cyan ink will have an appreciable change in hue as it gets thicker. From the plot of cyan ink, we see that the reflectance values between 500 nm and 600 nm are "intermediate", somewhere between the highest value and the darkest value. This is the green range. As cyan ink gets thicker, we would expect the amount of green light reflected to drop.

Thus, based on Seymour's rule of ink hue shift, a quick look at the plot below would suggest that thick cyan ink will be blue, just like thick magenta ink will be red. Yellow ink has very little in the way of intermediate values. It basically has either 75% reflectance or 3%. From that, you would guess that yellow ink will not change in hue. Note that a bucket of yellow ink does indeed look yellow.

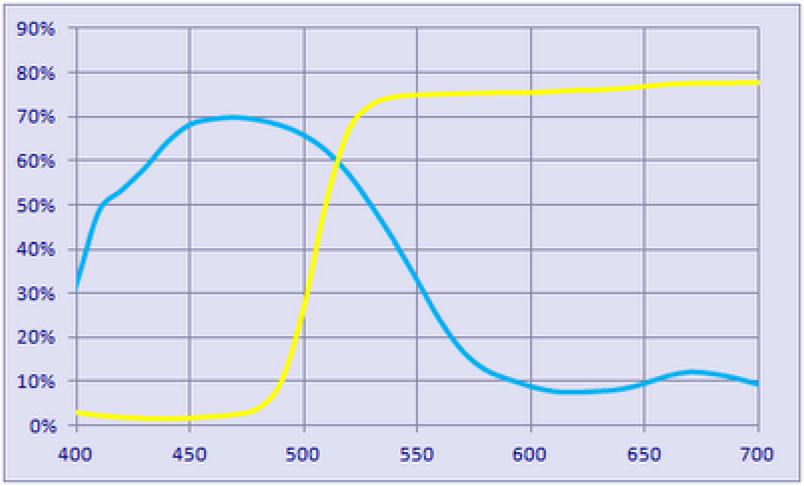

Plots of cyan and yellow ink

But spectra can be a bit misleading when trying to discern color. I don't know many people who can look at a spectrum and tell what the color is. So, I offer a little computational experiment to further validate Seymour's rule of ink hue shift.

First, I will show the results. Then I will explain how I got them. The chart below shows the a*b* values of a set of ten magenta patches with increasing ink film thickness. These values are the ten blue diamonds in the plot. There is clearly a strong hook. The first five are pretty much along a line without much hue shift. The sixth one goes around the bend, and the last four are changing a lot more in hue than they are in chroma.

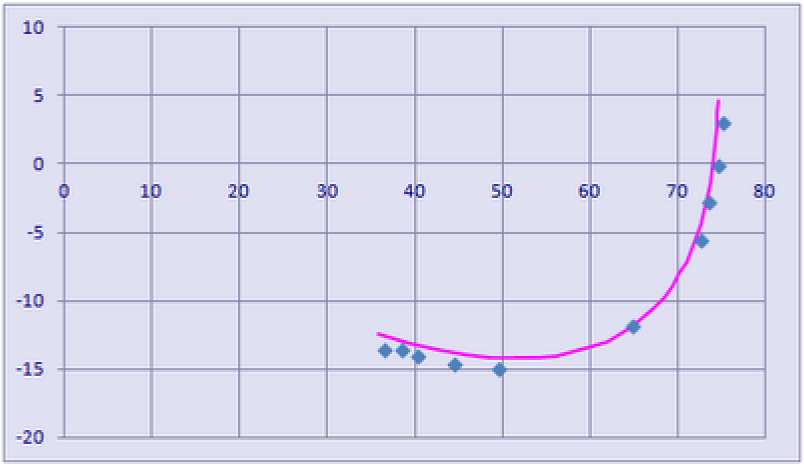

The magenta hook, real and estimated

For the other views of this data, I have published an addendum to this blog post.

The magenta colored line in the plot is a prediction of what I call the "ink trajectory". This is the set of all L*a*b* values that an ink will go though as you change the ink film thickness. To compute this estimated trajectory, I started with the spectrum of the sixth patch and that of the paper. (You will note that the magenta line goes right through that point.) I loaded these spectra into a spreadsheet, and used Beer's law to estimate the spectrum over a range of ink film thickness. You will note that the estimated trajectory comes reasonable close to predicting actual measured values, and definitely predicts the hook.

For those who want more detail, I have a little more description below. This is excerpted from a paper I presented at TAGA in 2008.

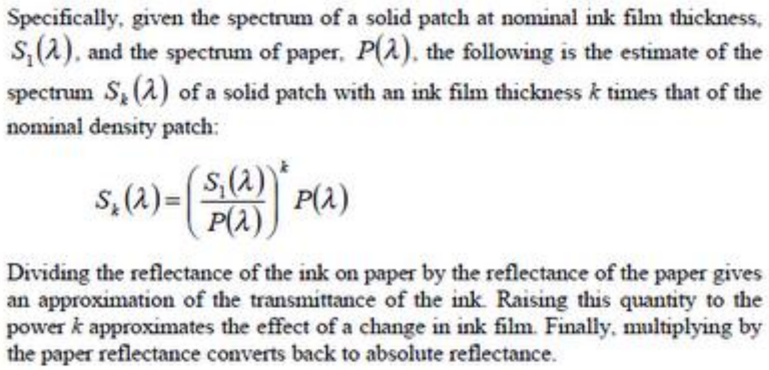

Specifically, given the spectrum of a solid patch at nominal ink film thickness, $S_1(\lambda)$, and the spectrum of paper, $P(\lambda)$, the following is the estimate of the spectrum $S_k(\lambda)$ of a solid patch with an ink film thickness k times that of the nominal density patch:

$$S_k(\lambda) = \left(\frac{S_1(\lambda)}{P(\lambda)} \right)^k P(\lambda)$$

Dividing the reflectance of the ink on paper by the reflectance of the paper gives an approximation of the transmittance of the ink. Raising this quantity to the power k approximates the effect of a change in ink film. Finally, multiplying by the paper reflectance converts back to absolute reflectance.

Excerpt from *Building a Bridge from Dense City to Colorimetropolis*

This pretty well settles it in my mind. Magenta ink on paper is magenta. Magenta ink in a bucket is red. I have explained this with some simple ciphering with Beer's law. This led me to define Seymour's rule of ink hue shift, which allows you to tell (just by looking at a spectrum), whether an ink will have an appreciable hook.

I then showed some really, really impressive results that show that, armed with just the spectrum of your paper and that of your ink on that paper, you can determine the magenta hook. This is clearly a triumph of modern science.

I have come a long way since I was ransacking the printing plant to find those elusive halftone dots!

Caveats

This is where I admit to some of the lies in the previous section.

First off, Beer's law is only an approximation. It makes the simplistic assumption that a photon will either pass right through the ink, or get absorbed. It does not make allowances for photons that reflect directly from the surface, or for photons that bounce around a bit in the ink and maybe come out of the ink without ever having visited the paper.

Despite those simplifications is does fairly well. For the standard process inks. I do not have data to see whether it works for Pantone inks. *If anyone has a cup of data to spare…*

One limitation that I glossed over is that it does not do well at predicting the reflectance of a double layer of ink. Us folks in the know like to say that ink is “sub-additive”, which means that Beer’s law does not do well at predicting the reflectance of a double layer of ink. It will, however, give you a spectrum that is attainable, however. Just not at that particular ink film thickness.

Well, that was kind of a lie as well. There are limitations, especially when you get up to the very high densities. You will note that my hook graph fits the data pretty decently, but it would not be nearly so good if I tried to predict the lightest density from the darkest, or the other way around.

There is one more lie, or one more pair of lies actually, but they are subtle. I demonstrated two ways of deciding whether the spectra of magenta showed a hue change. The first way was kind of hand-wavy. “Look at the spectra and see that it looks a lot like red. Ignore the little bump behind the curtain at 450 nm.”

Well, this argument may fly for someone who has not spent thousands of hours looking at spectra. But, if you have devoted a lifetime to deciphering spectra, you would know that sometimes the stuff happening down at the dark end is important. That little bump at 450 nm might just have a big effect on the color.

In this case, it didn’t. Converting to CIELAB demonstrated that the magenta is definitely turning red.

Or did it turn red? This is where the lie gets very subtle. We are trained from childhood to believe that colors with the same CIELAB hue angle are actually the same hue. But I have stubbornly disagreed with this all along. My first grade teacher almost flunked me over this point. I was glad to come upon a paper by Nathan Moroney where he made an off-hand comment that agreed with me.

The issue has to do with the fact that the CIELAB formula performs a nonlinear function on the XYZ values, which are a linear combination of the actual sensors in the eye, but which probably don’t actually exist in the eye or the brain. But that is grist for another blog.