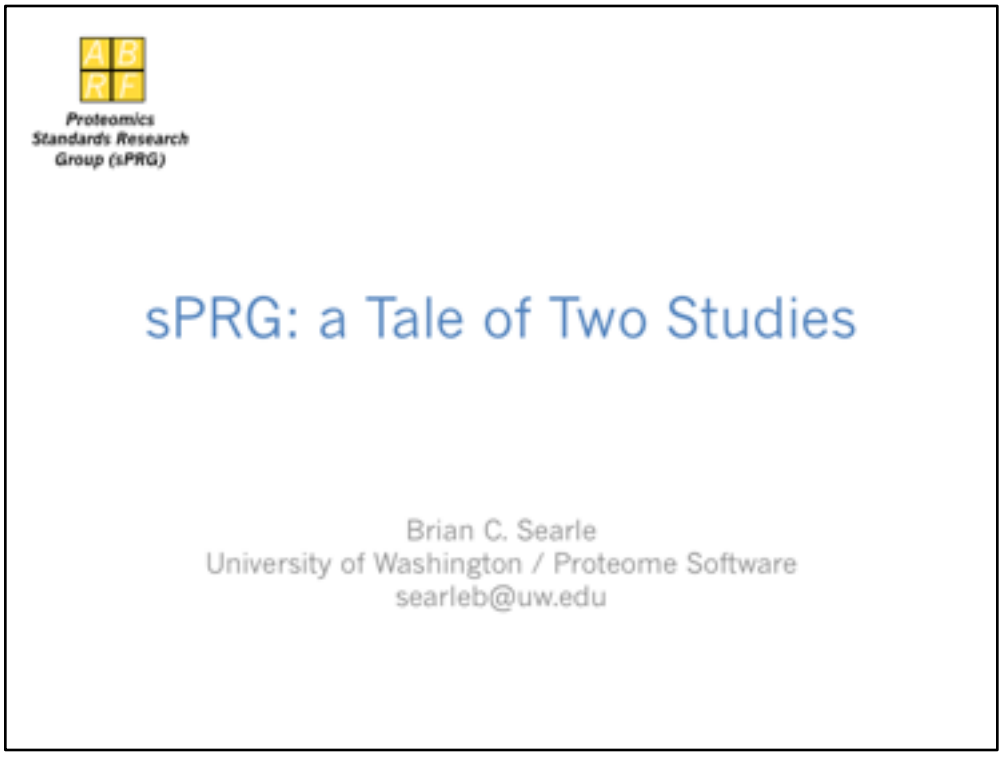

Hi I'm Brian Searle and I'd like to tell you about what the sPRG has been up to.

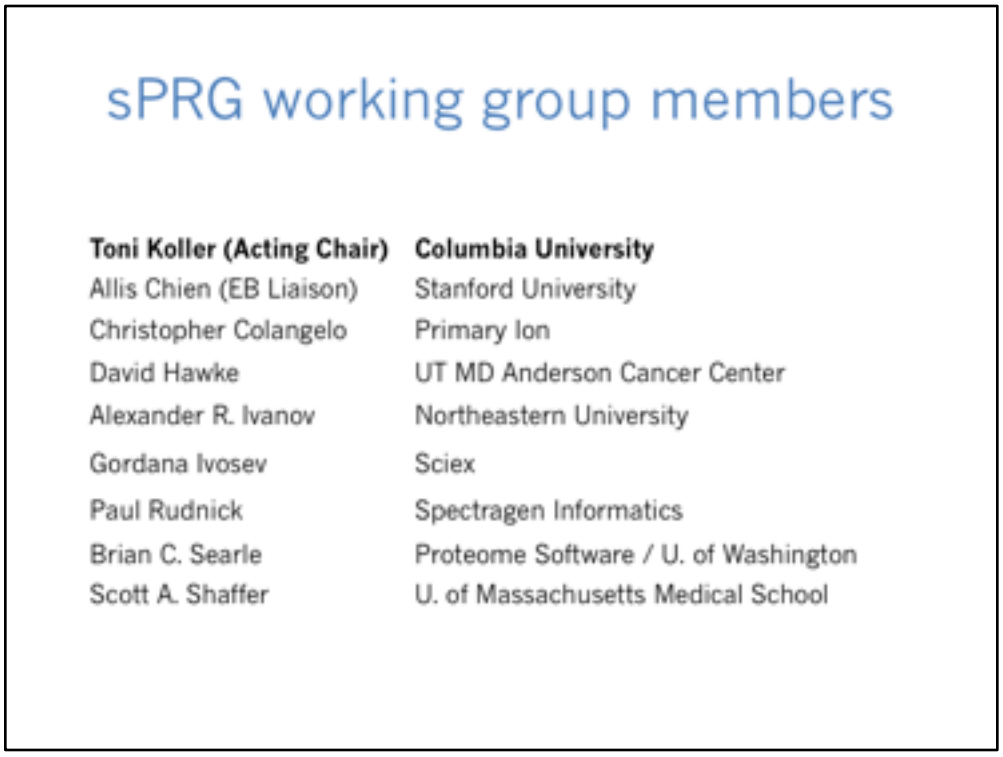

This year the sPRG working group is made up of a diverse group of ABRF members.

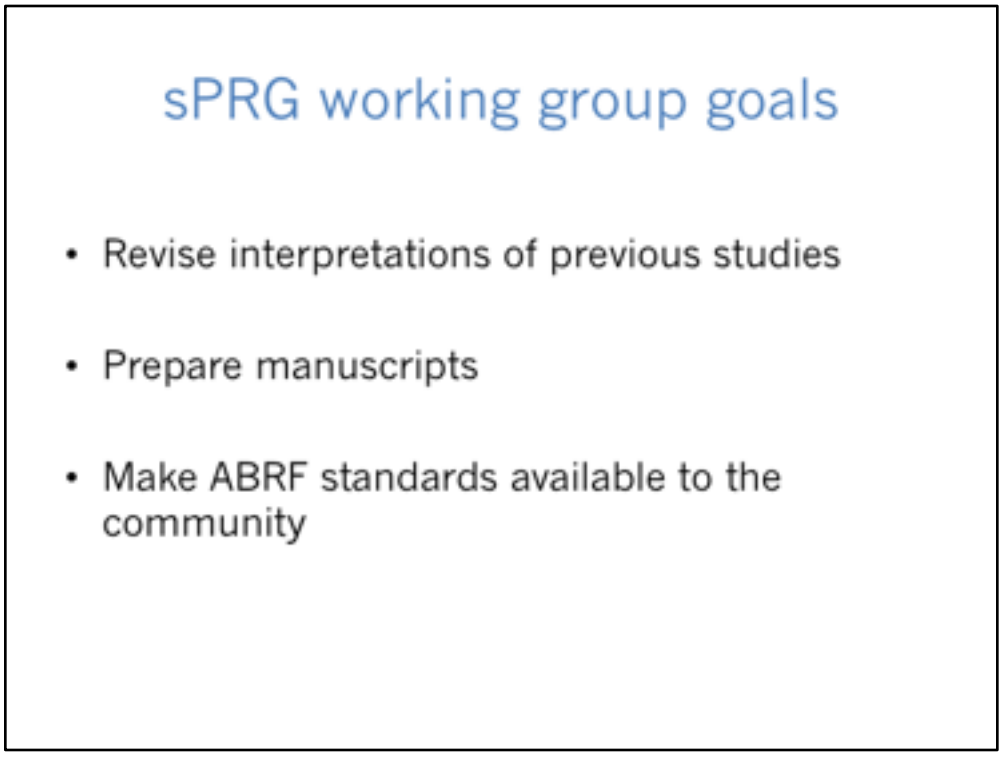

Our goals for this last year were not to launch a new study, but instead to organize and revise our analysis of previous studies, prepare manuscripts, and find ways to make our standards available to the community.

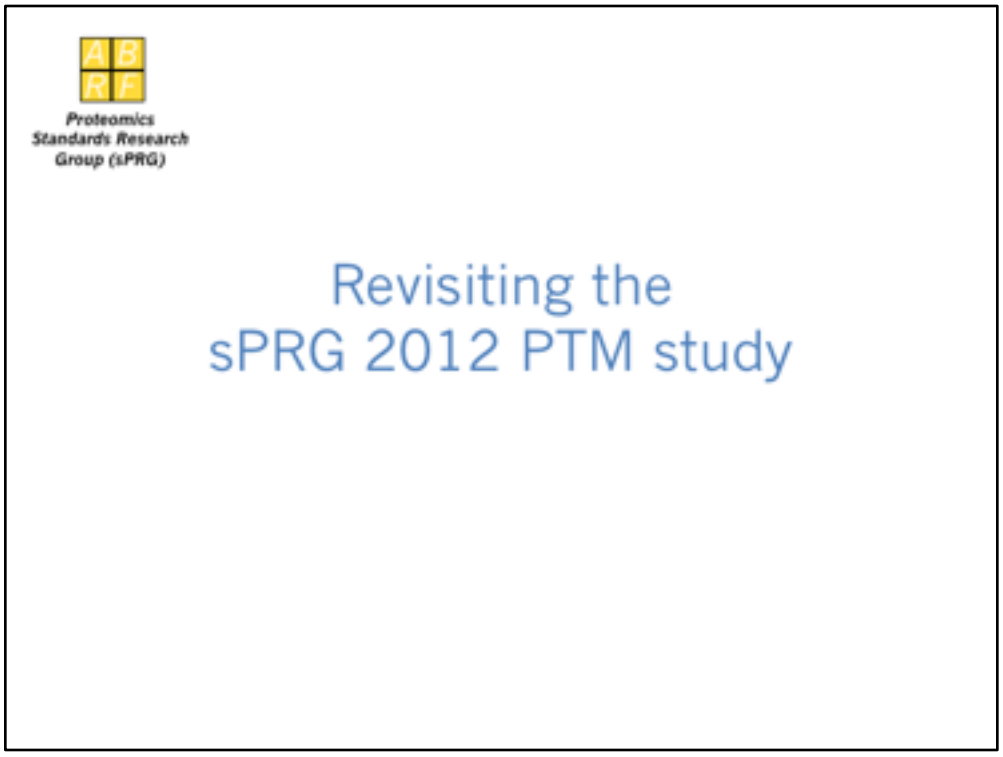

We've been working on finishing up two studies, the first is the 2012 study on PTMs, particularly phosphorylation.

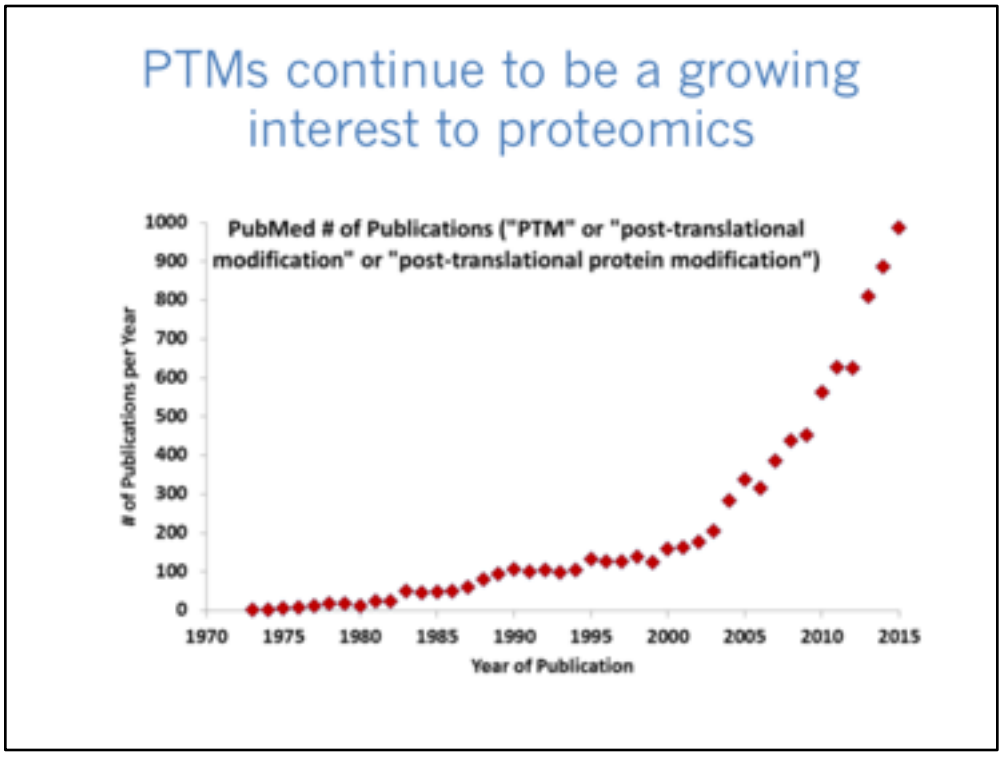

Interest in post-translational modifications has increased substantially over the past few decades, as indicated by the increase in publications per year on PTMs.

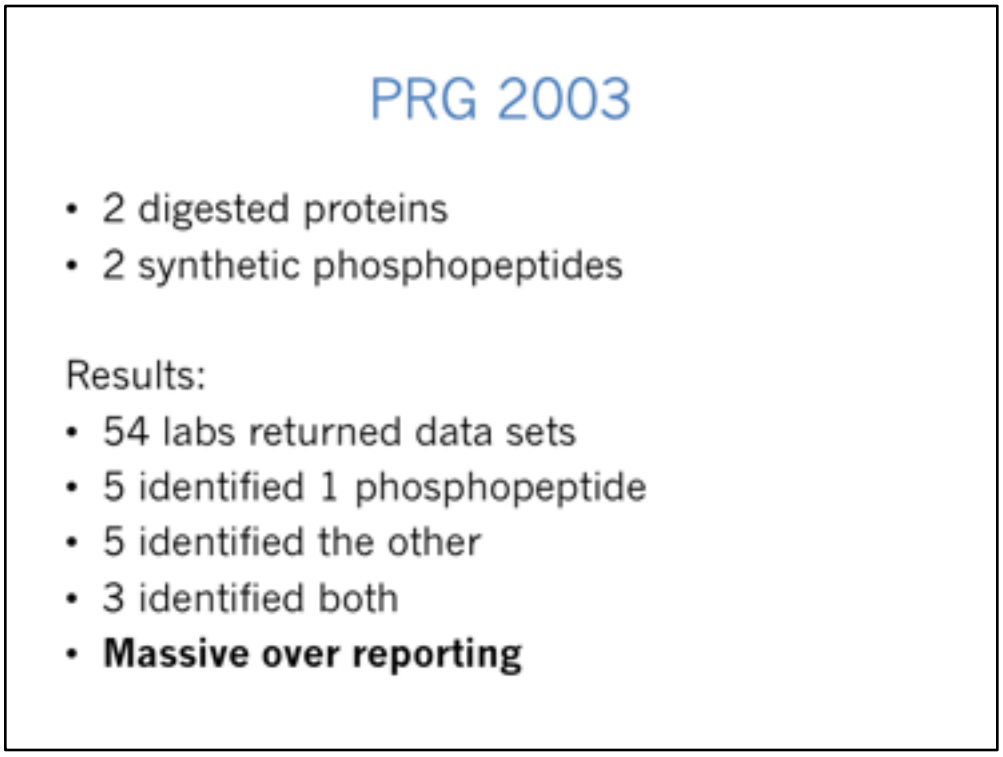

And the ABRF research group community has followed this trend with interest. The first PTM study we did was in 2003, where the PRG produced a sample containing 2 digested proteins and 2 synthetic phosphopeptides. Of the 54 labs that participated in this study, 5 labs could identify one peptide, another 5 could identify the other, and only 3 labs could identify both. This didn't stop people from reporting results, though, as was indicated by massive lists of incorrect peptides and proteins. The results were obviously pretty dismal and indicated a need for better methods and informatics.

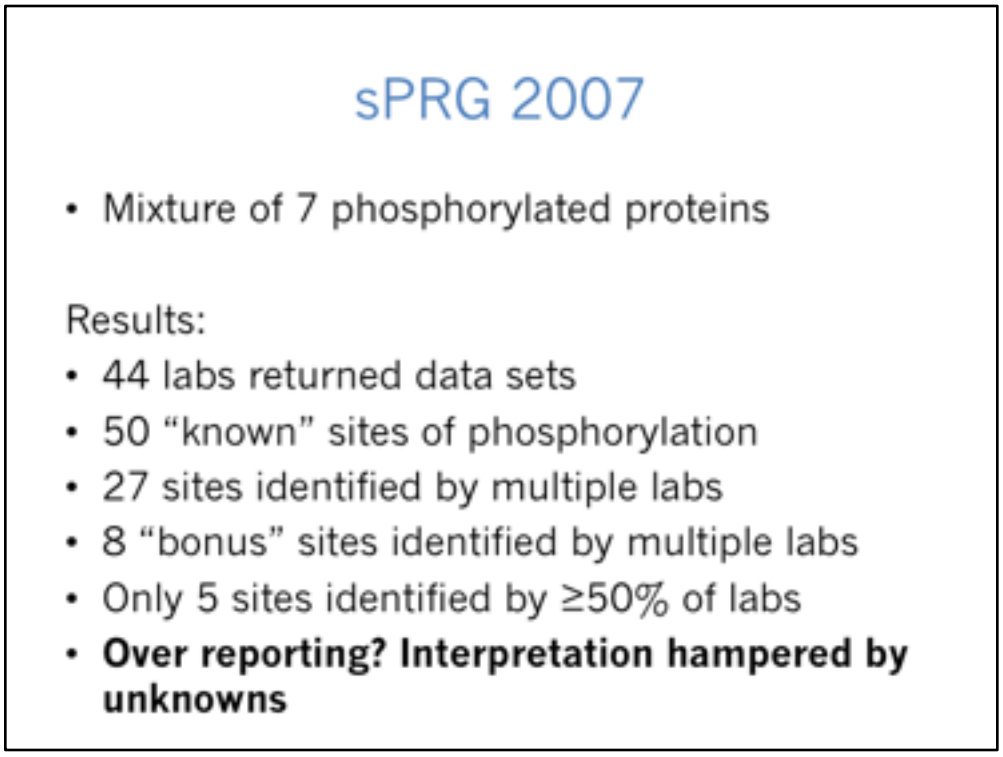

In 2007 the sPRG decided to try to follow this up by creating a standard mixture of 7 endogenously phosphorylated proteins. Of the 50 "known" sites of phosphorylation, only 27 were identified by multiple labs, where only 5 sites were identified by over 50% of the labs. Again, there was a ton of one-hit-wonder sites identified by individual labs, but since these were endogenous phosphorylations no one really knows what's in the standard.

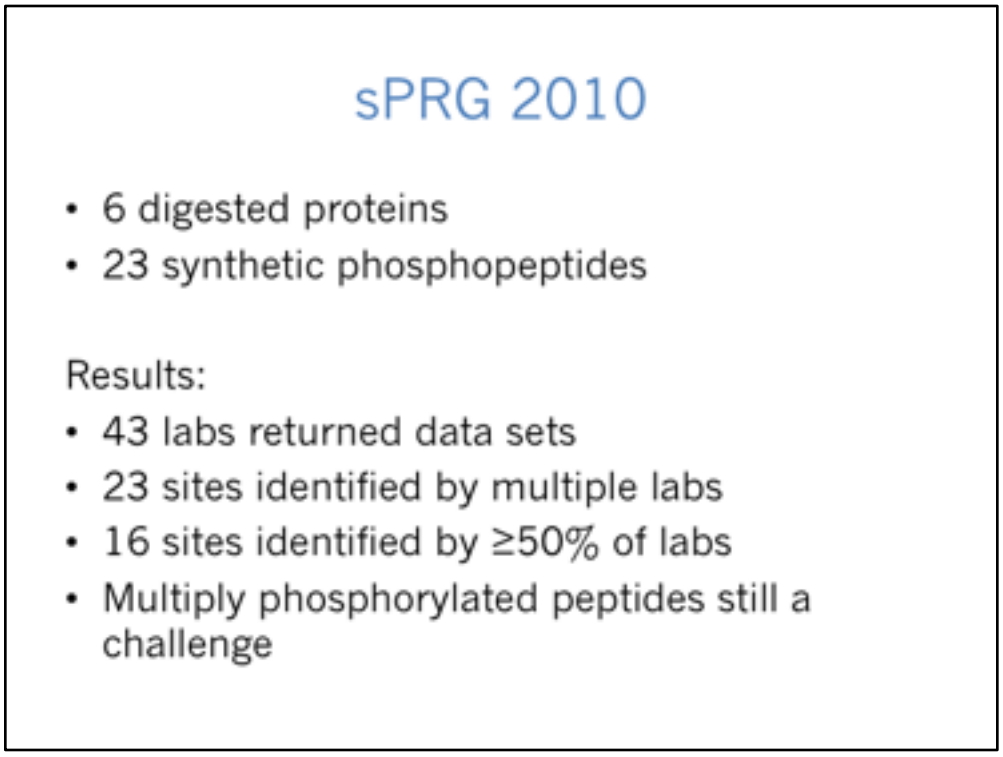

By 2010 things were getting better. Based on the issues with the previous study, the sPRG returned to 23 synthetic phosphopeptides and of those 16 could be identified by over 50% of labs. However, multiply phosphorylated peptides remained a major challenge.

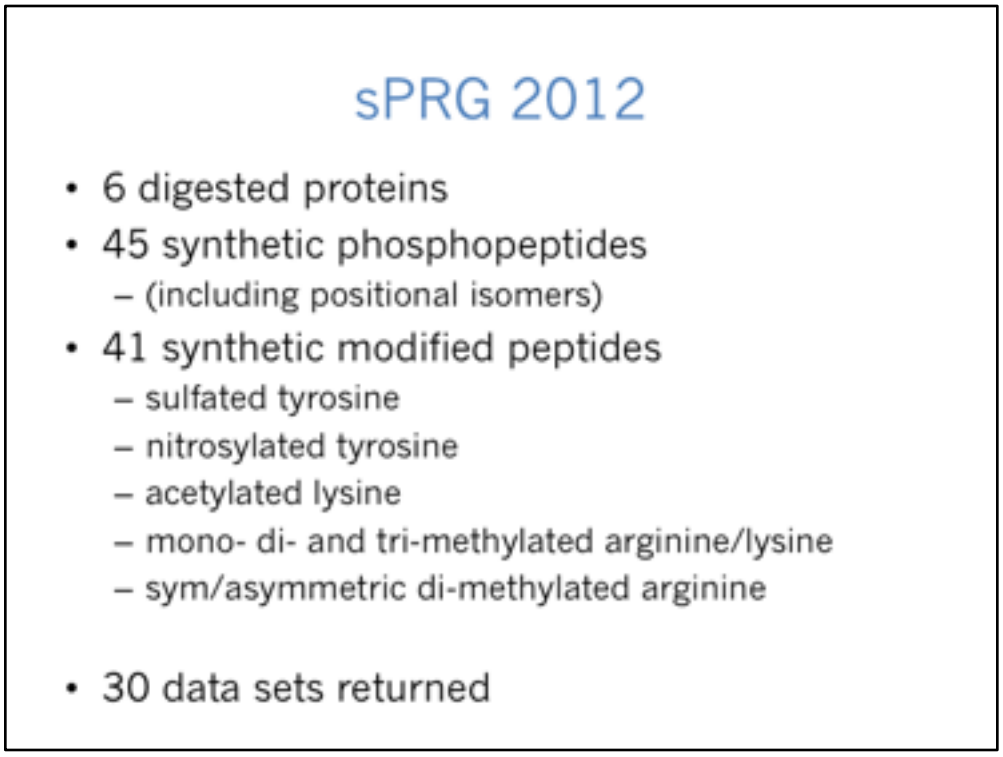

That brings us to the current study of 45 synthetic phosphopeptides, including positional isomers, and an additional 41 synthetic modifications of various types, including sulfation, acetylation, and methylation. 30 participants returned data sets.

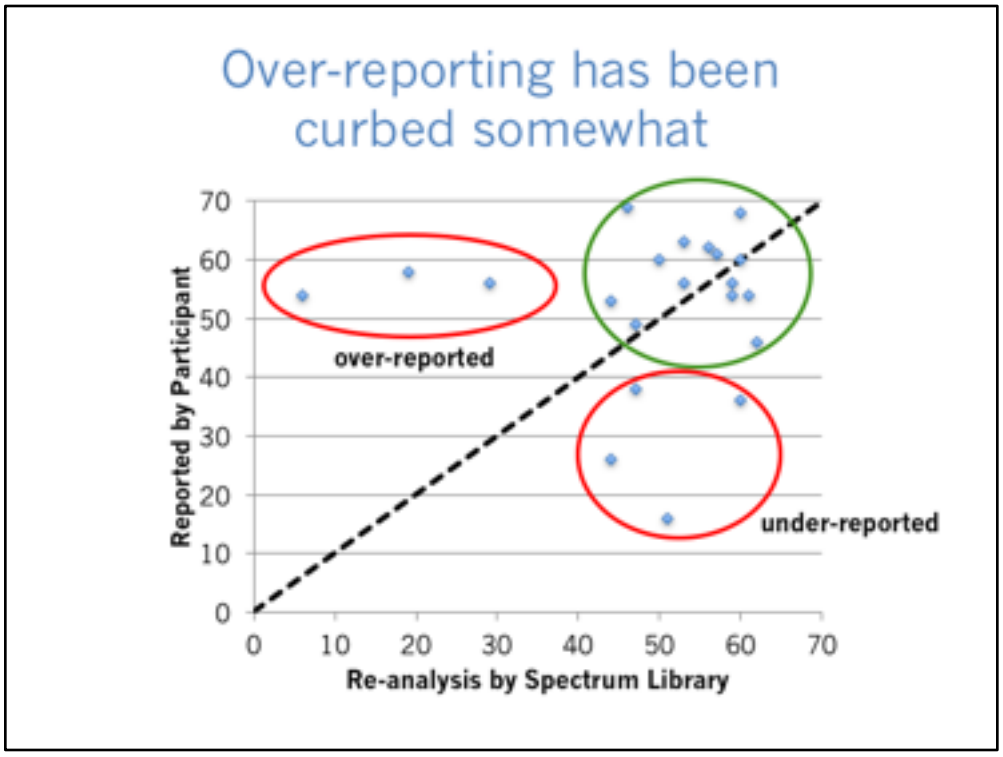

First off, one clear trend was that over-reporting was much less of a concern. Re-analysis of participants data in general produced the same number of peptide identifications as the participants reported, barring some under- and some over-reporting.

Cross study comparison shows general improvement

	PRG 2003	sPRG 2010	sPRG 2012
SVSpDYEGK	15%	40%	80%
THILLFLPKSpVSDYEGK	15%	62%	80%

Another clear trend is that as a community, we're getting better at identifying phosphopeptides. The two peptides from the 2003 study were carried forward into the 2010 and 2012 studies, and the percentage of participants that found those peptides increased dramatically over the decade between studies, culminating in 80% of participants in 2012 being able to identify those peptides.

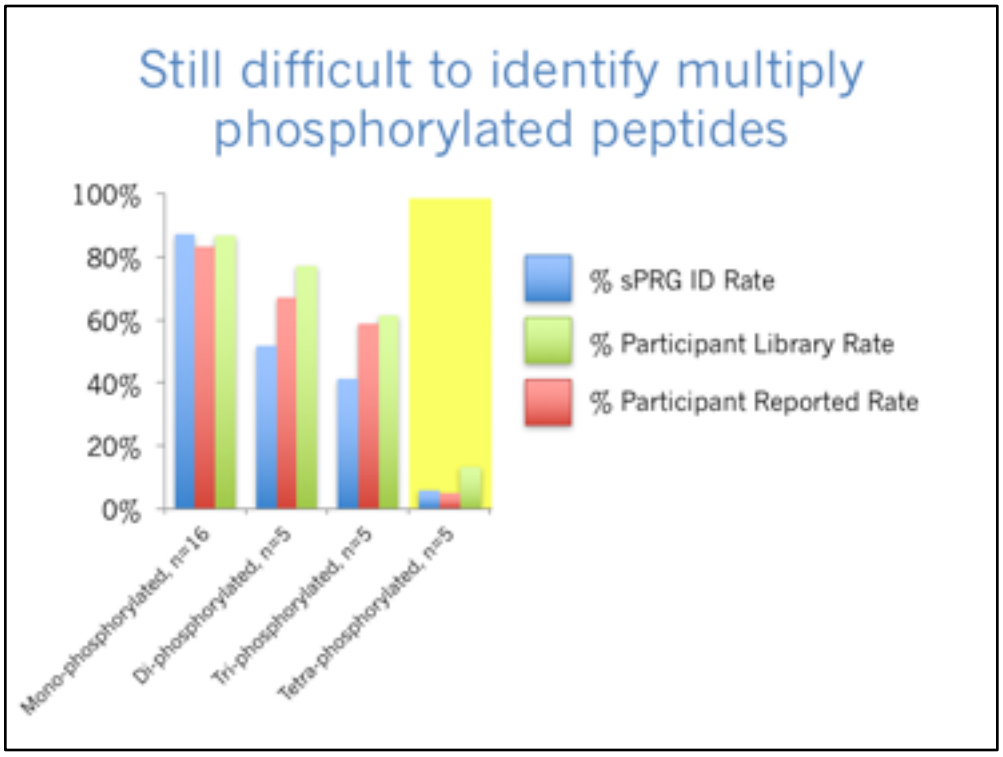

And in general that 80% identification rate held true for mono-phosphorylated peptides. While using a library search engine helped substantially, as you add more phosphates the likelihood of identification drops substantially.

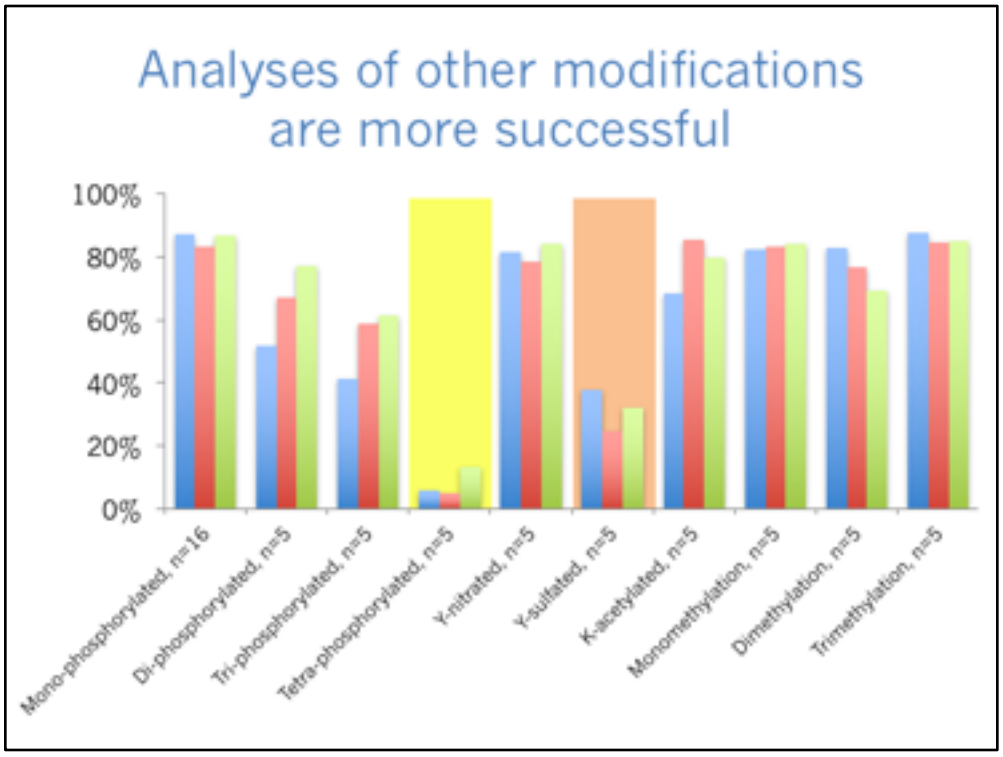

Looking at the other modification types, most of them are pretty consistently identified, except for sulfated peptides.

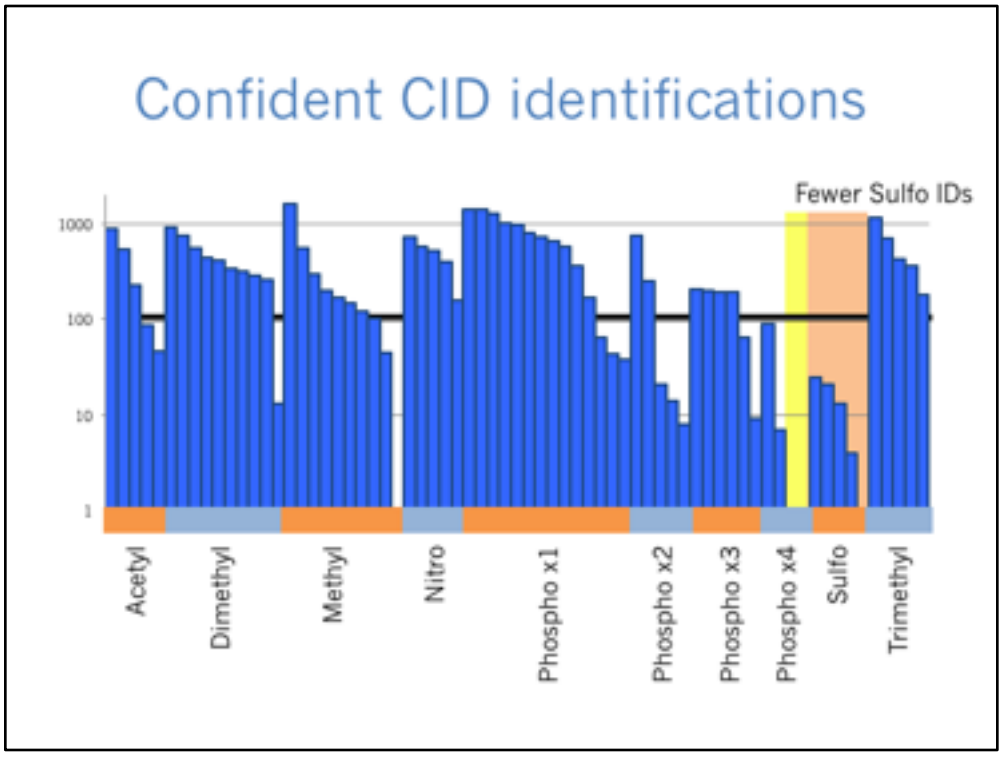

Looking at the individual peptides, here I'm plotting the number of confident identifications of CID fragmented peptides across all of the labs on a log scale. In general we're fairly confident of the fragmentation patterns for most modifications, except for the 5 sulfated peptides and two tetra-phosphorylated peptides that had no identifications.

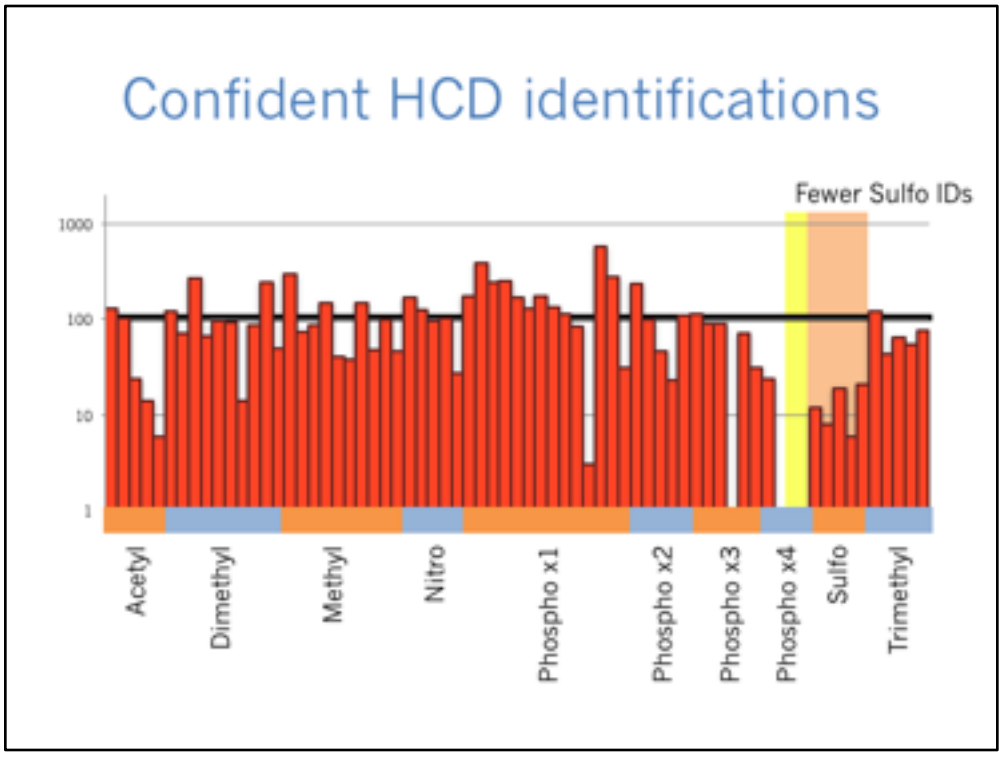

Looking at HCD fragmented spectra, the results seemed fairly consistent: an order of magnitude fewer sulfo peptides.

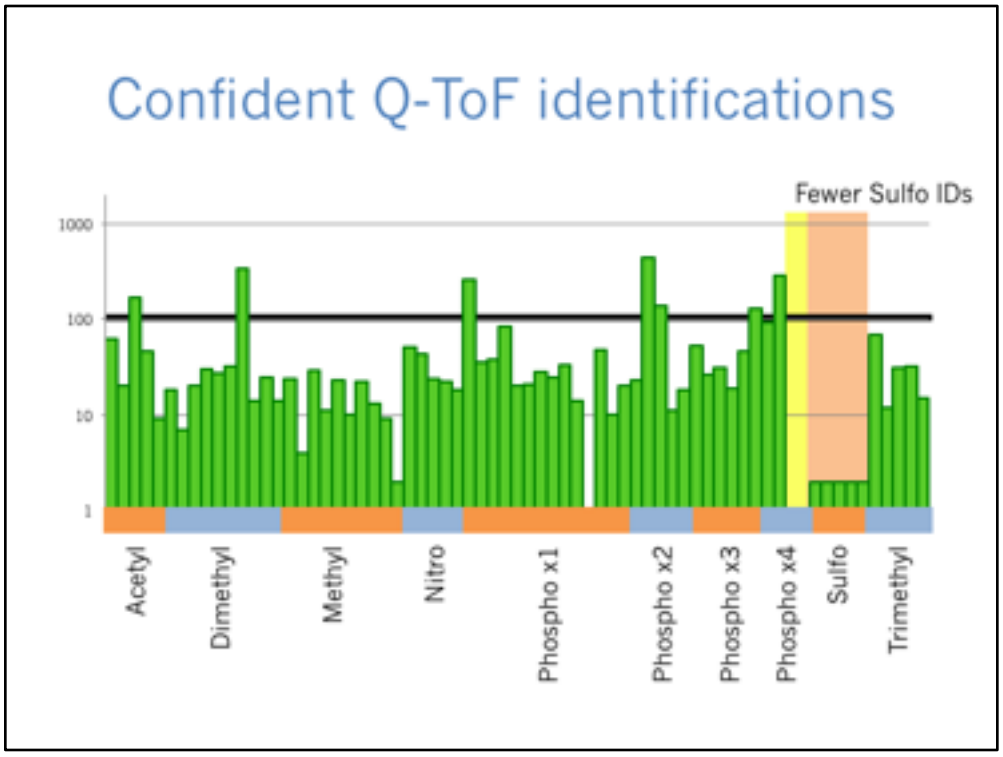

And again, while we had less Q-ToF data in general, the trend was similar.

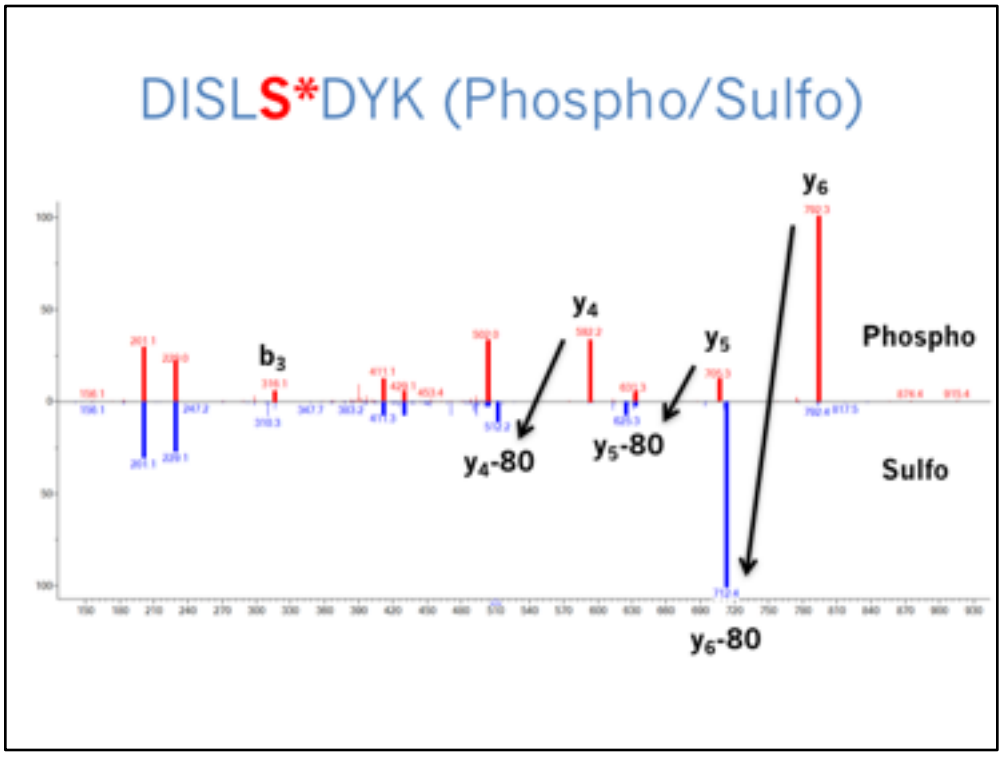

Given this high number of fragmentation spectra represented in the cross participant data set enables us to learn some cool trends about the peptide types. For example, this peptide was present in our sample in both phospho and sulfo forms, which are difficult to discriminate between based on precursor mass. However, considering the fragmentation patterns allows us to clearly differentiate between these two modifications, where the sulfo-forms often lost 80 AMU from the sulfate falling off.

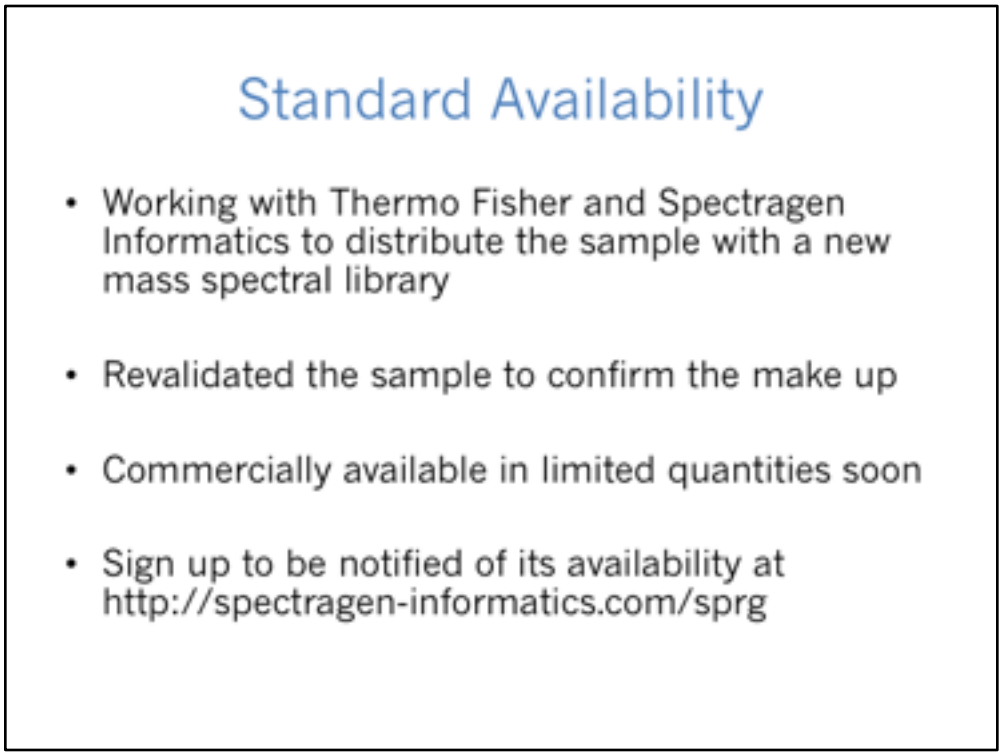

We've actively been working with both Thermo and Paul Rudnick's company, Spectragen Informatics, to distribute the sample. We've also revalidated the final sample to confirm the make up and try to confirm some of the hard-to-find peptides. We've made up several vials, which should be available in limited quantities very soon. We're not able to take orders yet, but if you're interested you can sign up for details at this website.

sPRG members involved in this study

Alexander R. Ivanov (Chair)	**Northeastern University**
Christopher Colangelo	Primary Ion
Craig Dufresne	Thermo Fisher Scientific
David Friedman	Vanderbilt University
Kathryn S. Lilley	University of Cambridge
Karl Mechtler	IMP Research Inst. of Molecular Pathology
Brett Phinney	University of California, Davis
Kristie Rose	Vanderbilt University
Paul Rudnick	Spectragen Informatics
Brian C. Searle	Proteome Software / U. of Washington
Scott A. Shaffer	U. of Massachusetts Medical School
Susan T. Weintraub	University of Texas HSC

This is obviously a large study and the product of effort from several people.

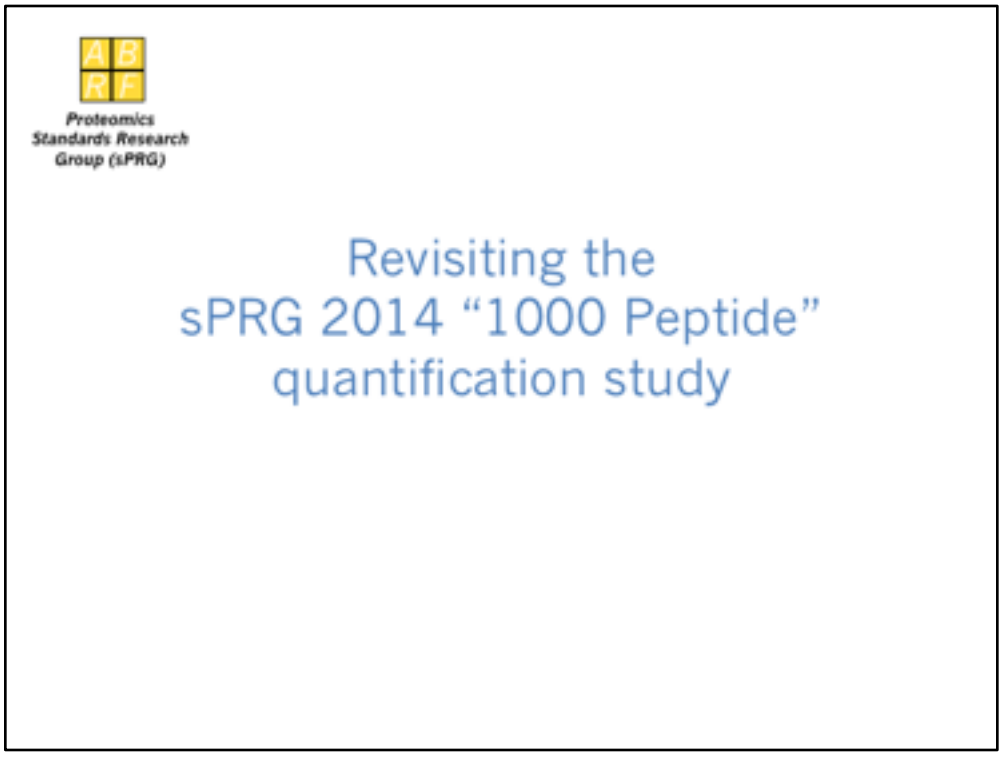

Having learned that it wasn’t too bad to do a large synthetic peptide study, we thought even bigger and designed a 1000 peptide standard for quantitative studies.

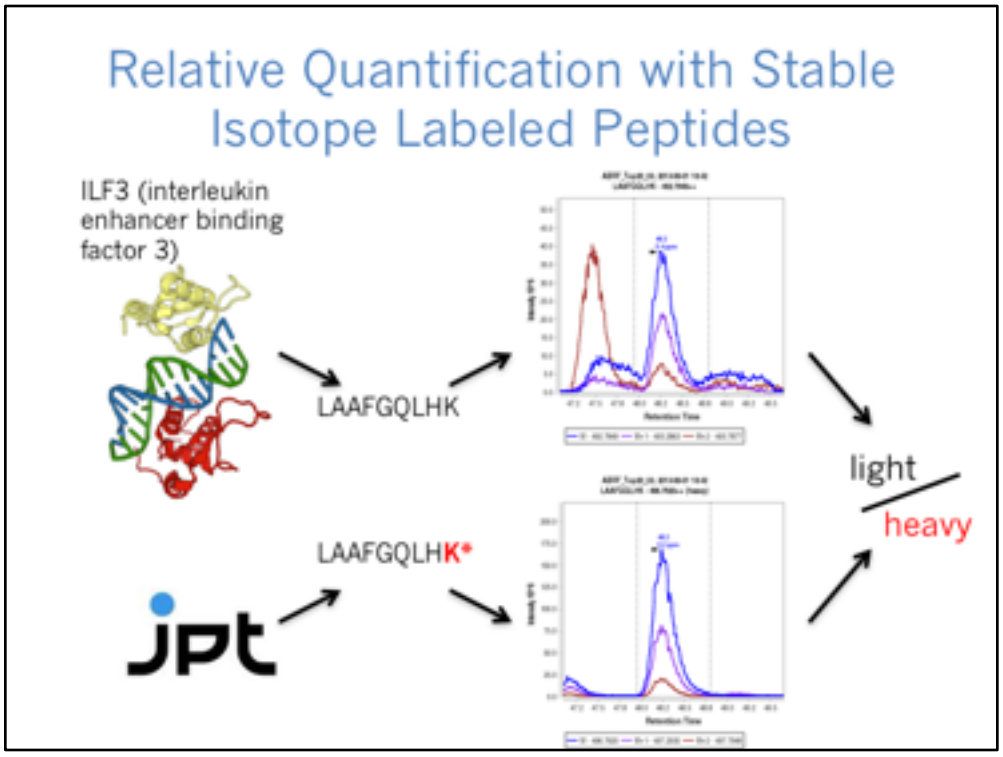

Say for example you wanted to quantify the transcription factor ILF3 based on a particular peptide, you might synthesize a matching stable isotope labeled peptide. Then when you quantify the signals from those peptides you can generate a ratio relative to the standard. Using the heavy standard as a normalizer, you can compare one biological sample to another.

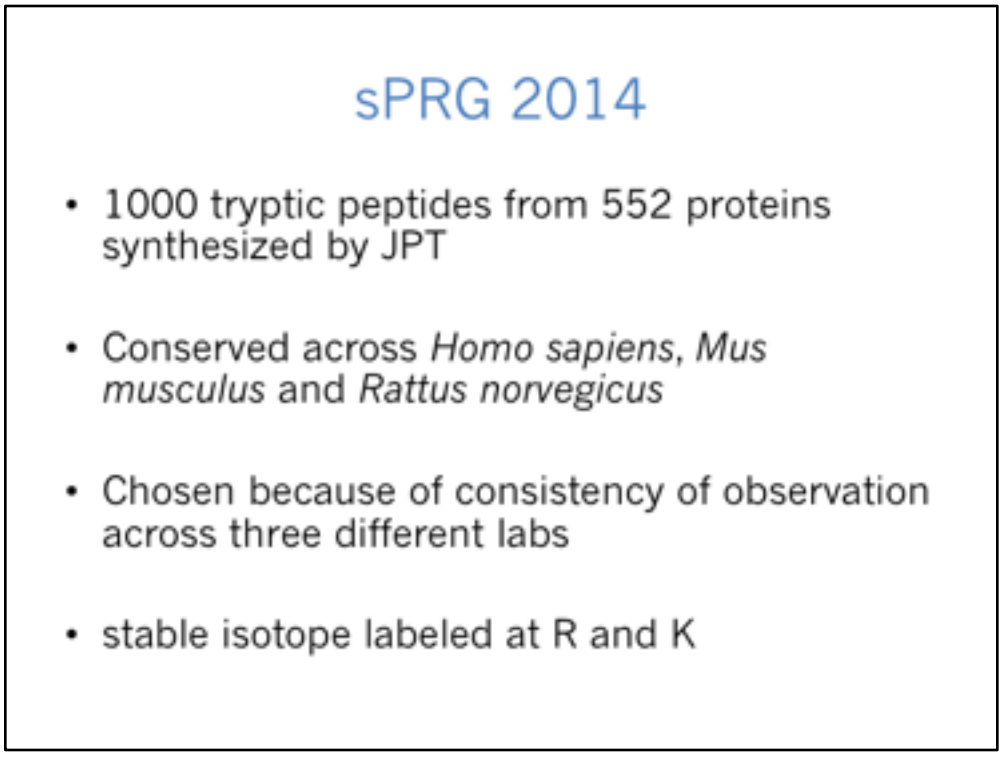

Based on that strategy, we designed a 1000 peptide standard from over 500 proteins. The peptides were conserved across human, rat and mouse, and chosen because they were found consistently across three of our labs in DDA experiments.

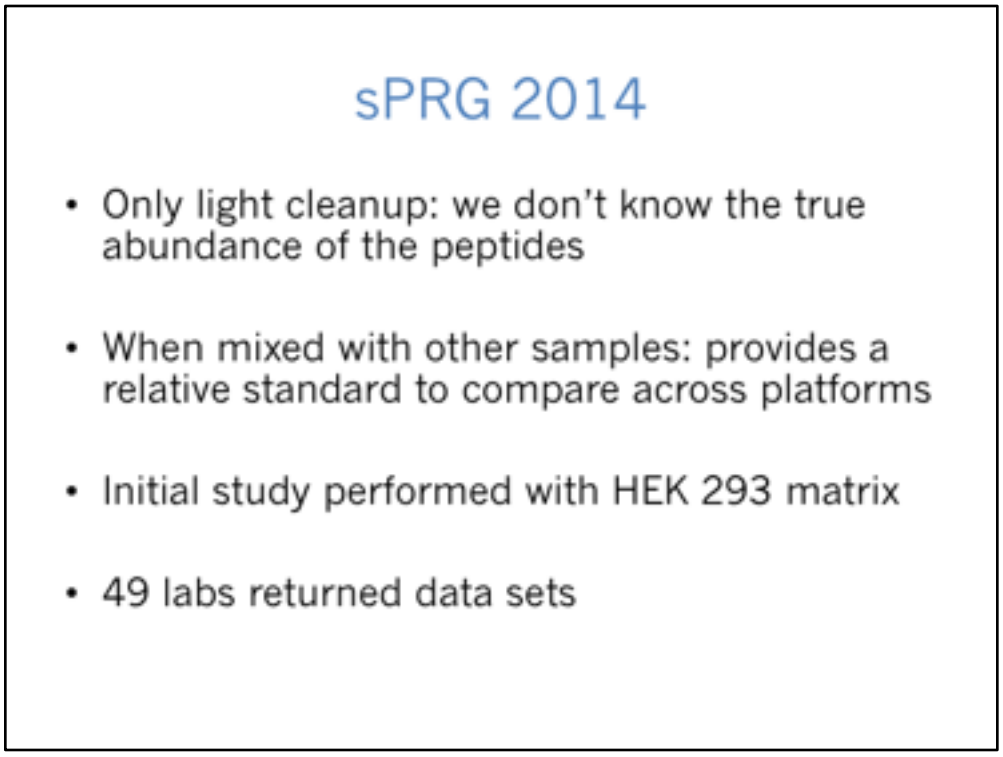

The peptides were crude synthesized and only lightly cleaned up, so we don't know the true abundance of these peptides. However, when mixed with other samples, it can be a useful relative standard. In our initial study we mixed this with H E K 2 9 3, distributed it, and received 49 data sets back.

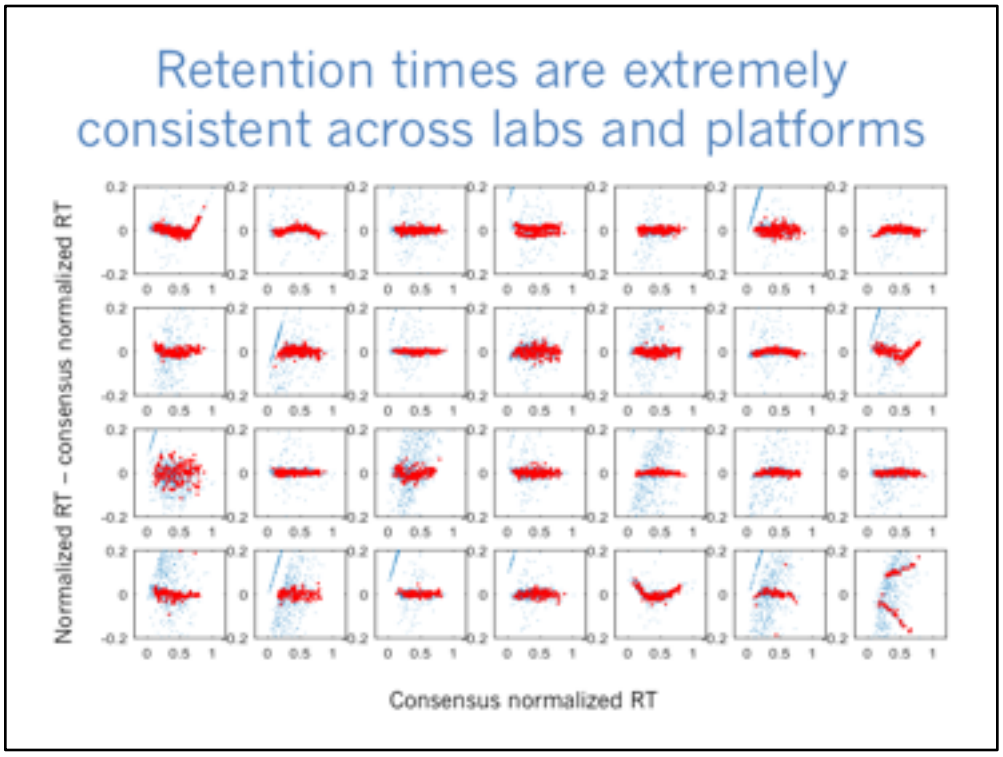

One of the things we learned is that retention times are extremely stable across platforms. Here I'm showing the delta retention time versus a consensus normalized retention time, and in most samples the delta change is flat at 0%.

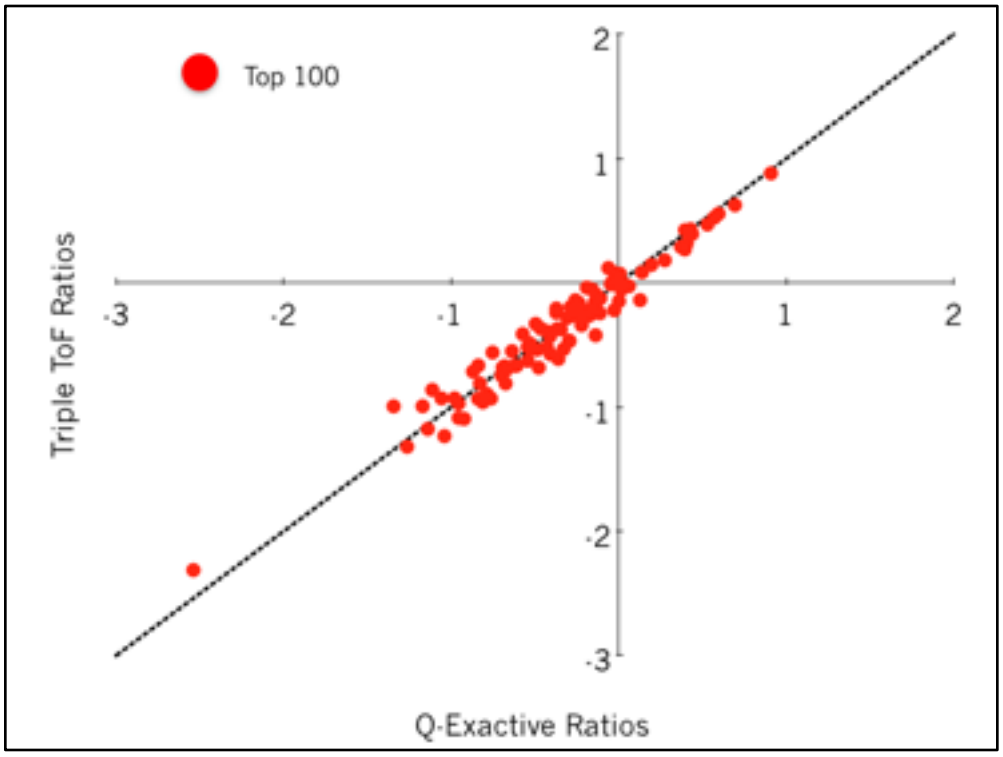

And looking at the most stable quantitative ratios, in general different instrument platforms provided the same answer. However, here I'm only showing you the top 100 of 1000 peptides.

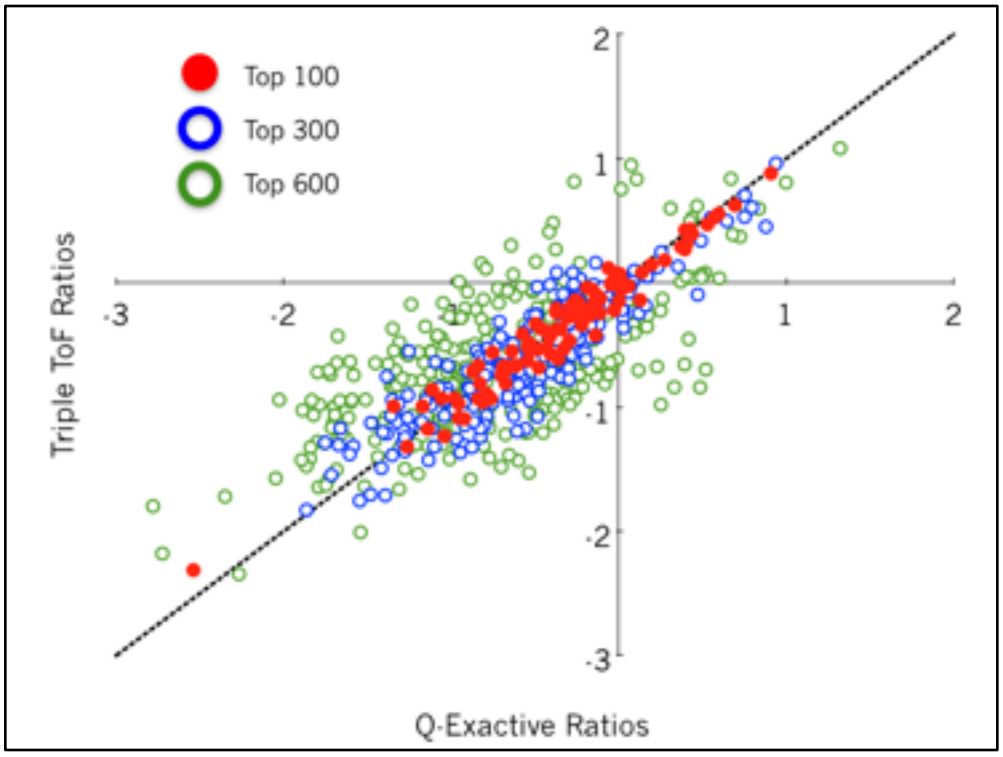

If I expand out to the top 300 or 600, the relative error rate increases substantially. I'm not going to show you the results for the bottom 400, it's basically a scatter shot.

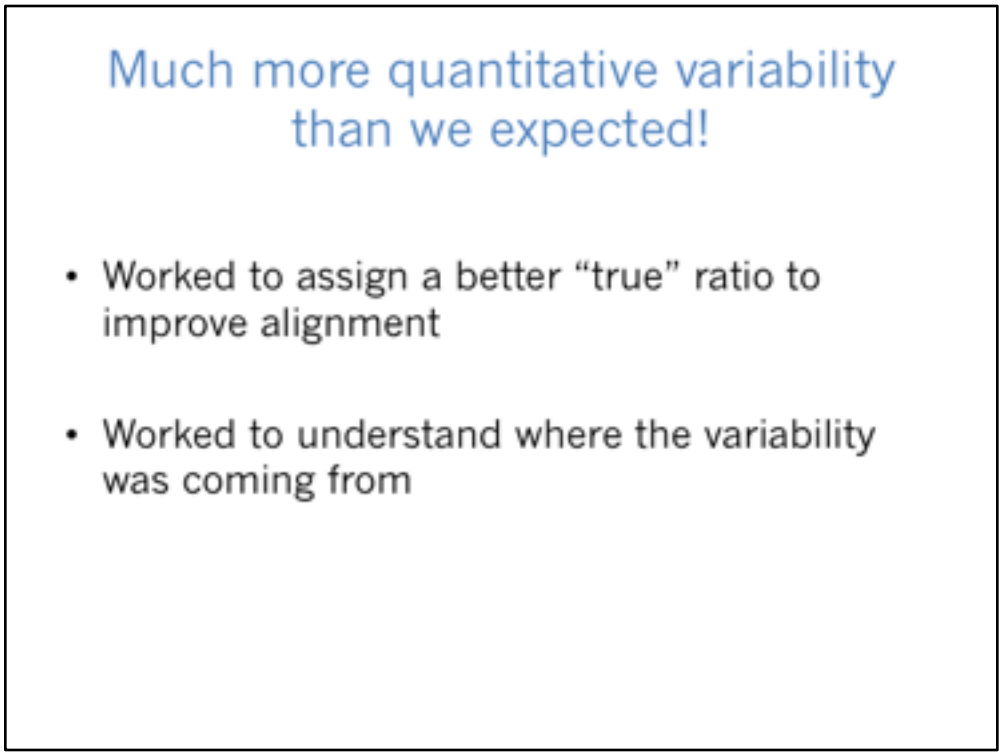

To be honest we were shocked by this level of quantitative variability. We decided to try to grapple with this from two ends: first we worked to assign a better "true" answer to potentially improve alignment, and second, we tried to understand the sources of variability to eliminate outliers.

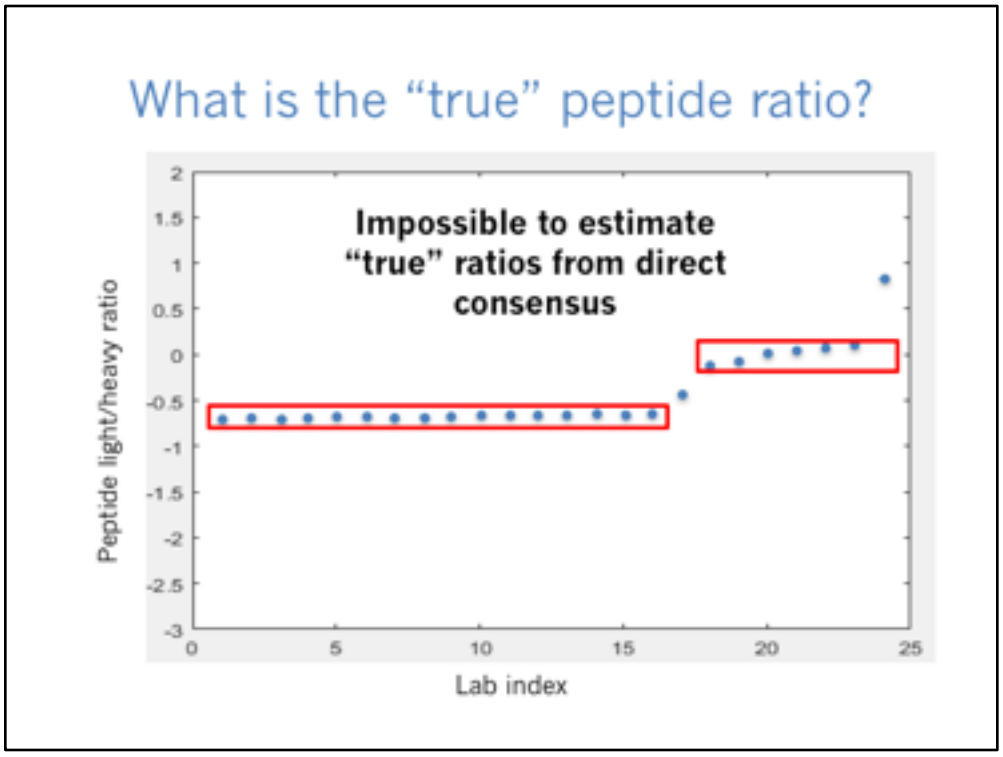

Originally we tried to estimate “true” ratios from direct consensus, however it quickly became clear that this didn’t work. This is a chart of the ratios from one peptide, across 25 labs. These ratios fall into two groups, so which is the consensus ratio? We needed a better strategy.

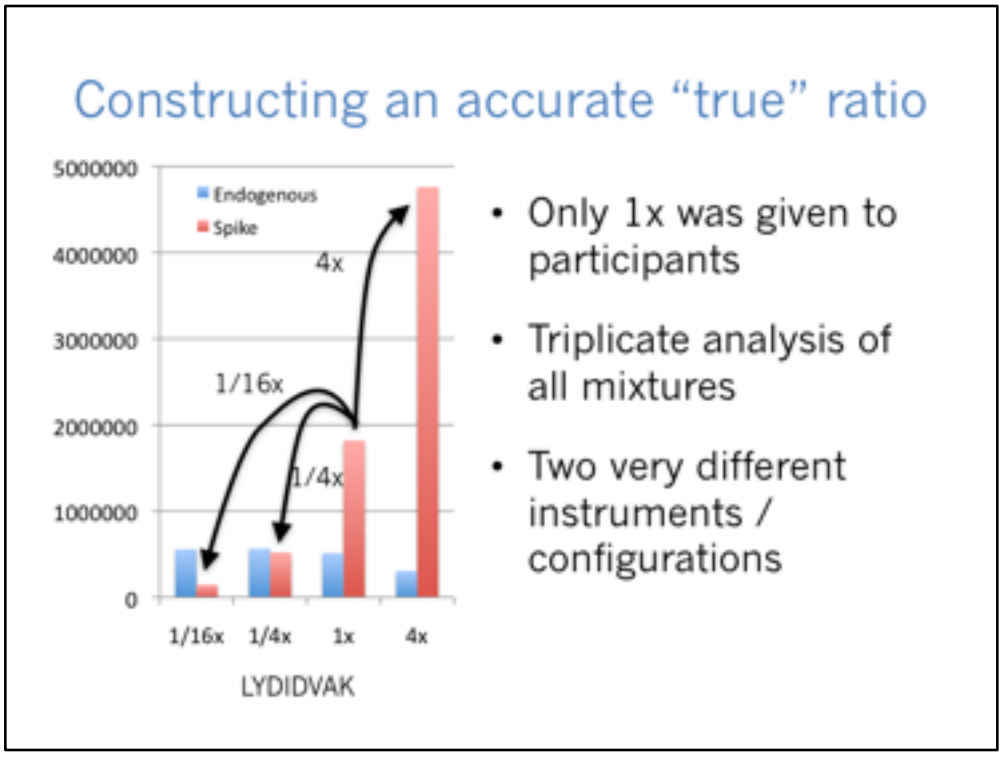

While we only provided participants with one mixture of standard to H E K, we actually generated four mixtures. This lets us do ratios of ratios, where the ratio of the 4x ratio should be 4 times larger than the 1x ratio. The 1x ratio should be 4 times larger than the ratio from the 1/4x mixture, etc, etc. We created a linear model to combine all these comparisons and generate a "true" ratio for the 1x sample.

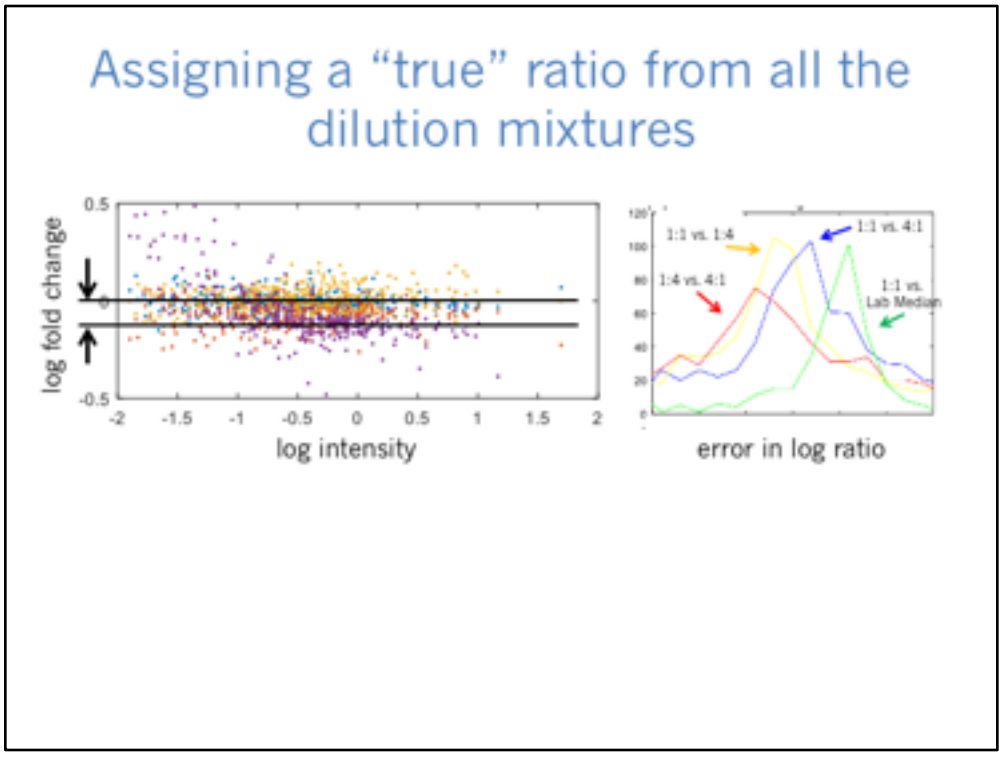

Without this level of normalization there was an up to 50% error between the dilution series samples, shown in the scatter plot, but also in the integrations on the right. After normalization we were able to estimate a more precise “true” ratio for each peptide.

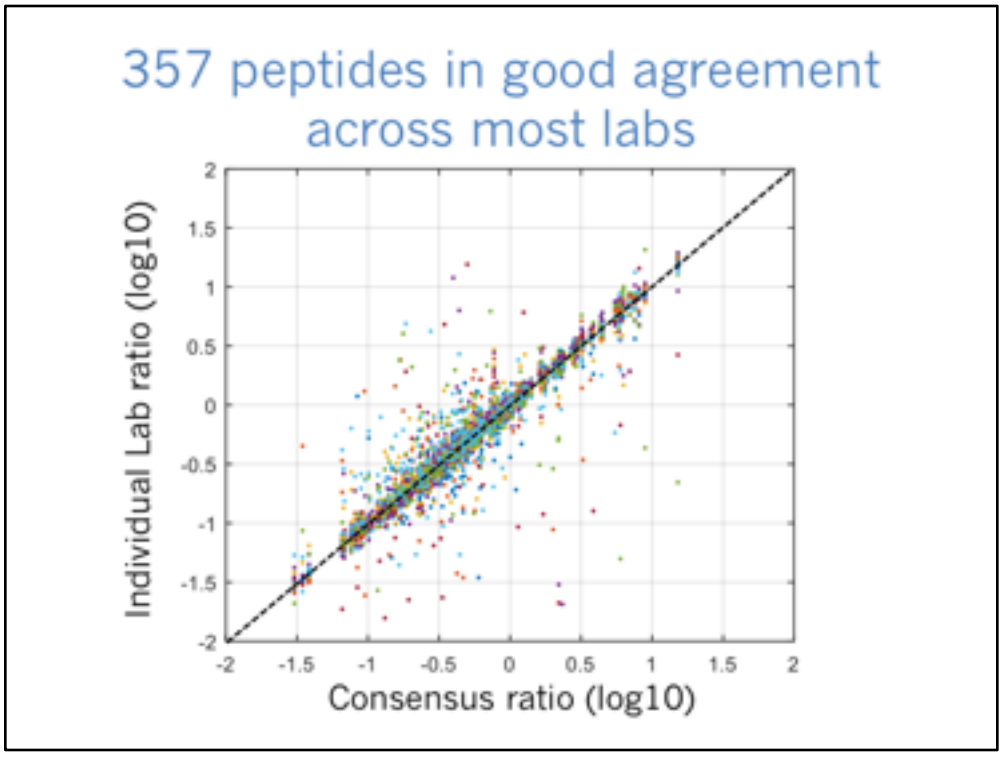

From here we found that there were about 300 peptides that were pretty consistent across most labs with a consensus. However we had to drop some samples.

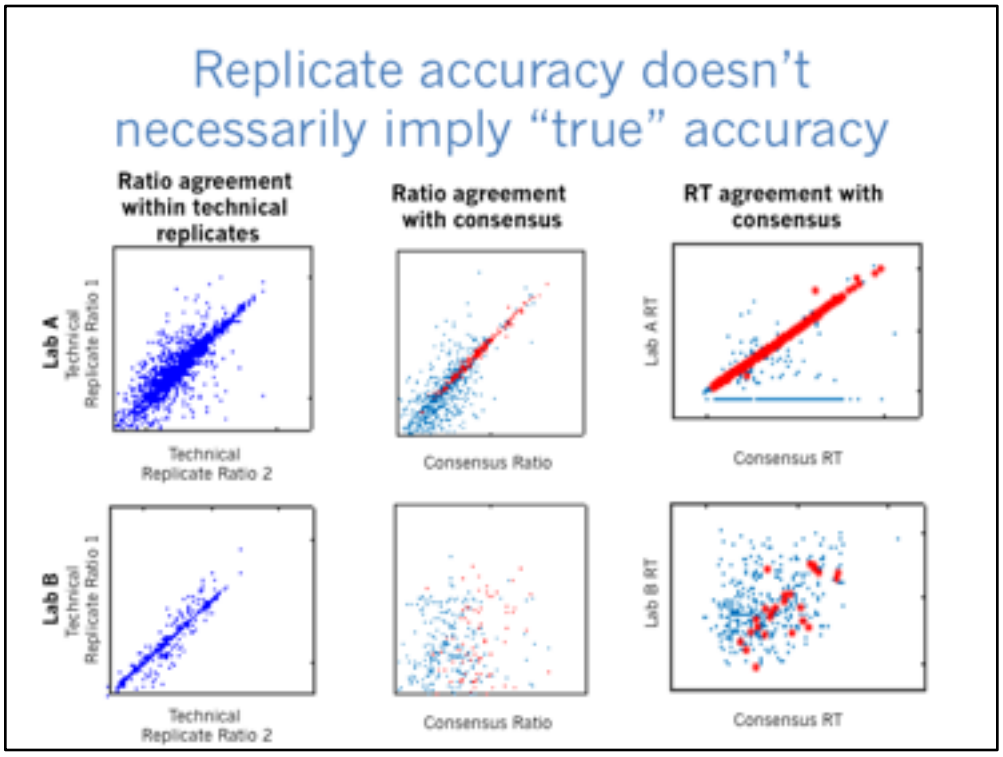

And the reason is actually pretty interesting. Here are replicate ratios for two labs, where the replicates largely agree with each other, implying that the quantitation is accurate. However, only one of the labs agrees with the consensus. This is further confirmed when we look at retention time accuracy with the consensus. This demonstrates that it is possible to reproducibly quantify the wrong thing!

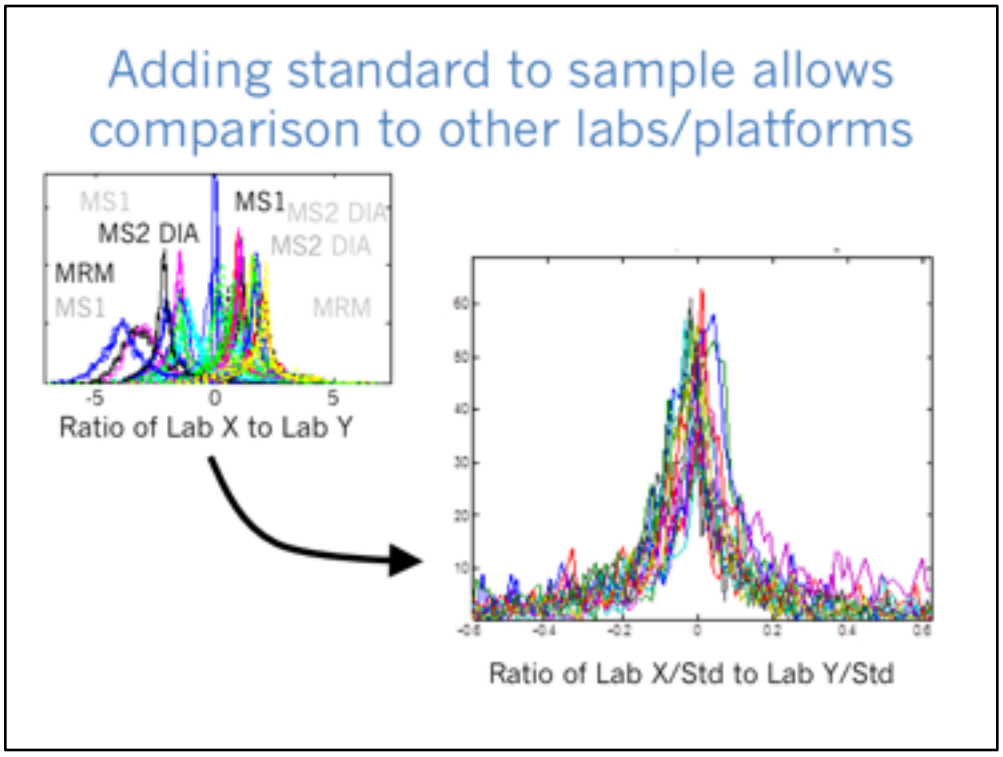

One cool take home of this study was that if you were to compare the quantitative values generated from one lab to another lab or one platform to another platform, they'd be all over the place. However, when you compare one labs results to a standard, and then another labs results to that same standard, the quantitative values were substantially more accurate.

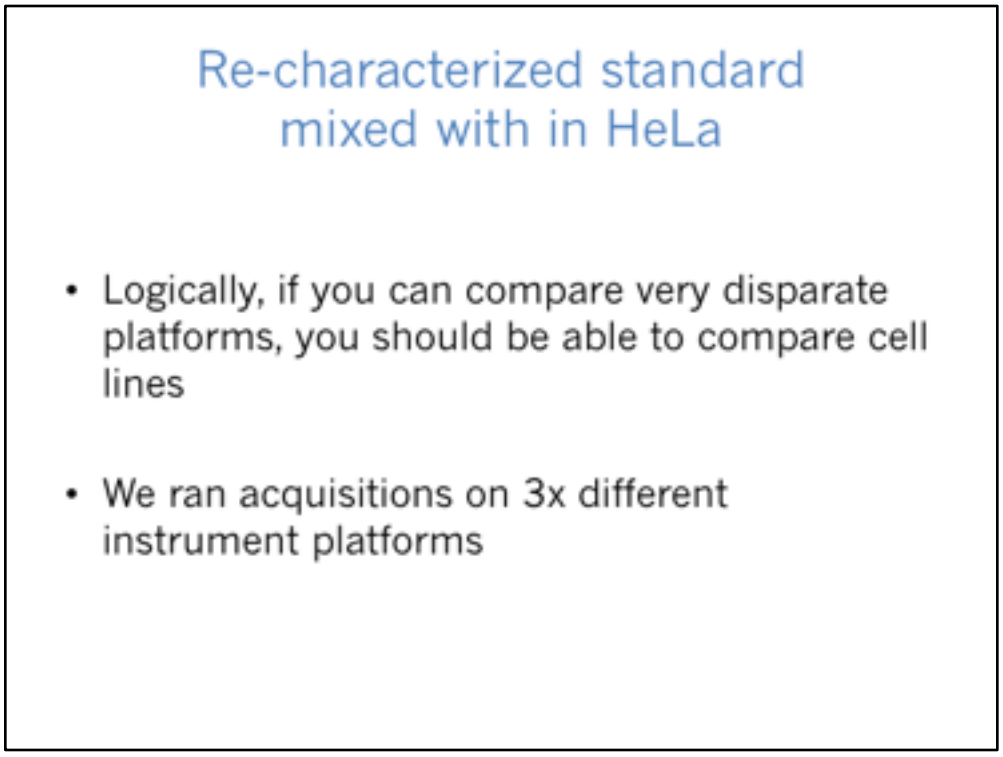

Logically it would make sense that if you could compare disparate platforms from triple-quad SRMs to ion-trap DDA, you should also be able to compare cell lines. We sought to re-characterize the standard when mixed with HeLa but found something interesting.

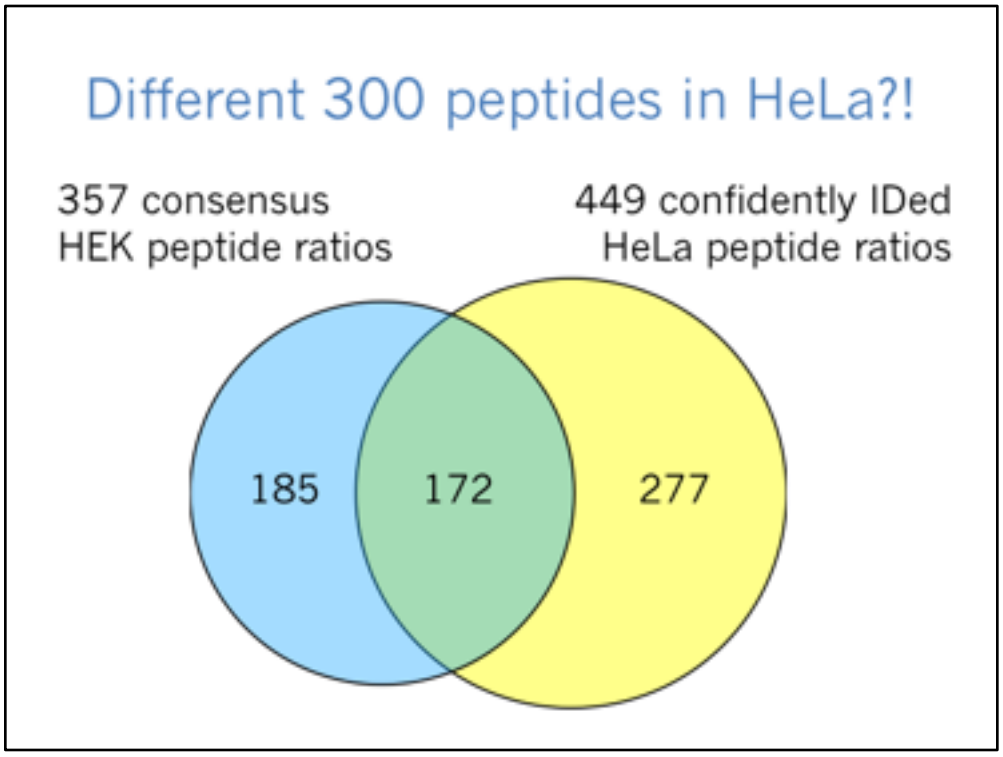

We ended up finding that, again, while we could identify all of the peptides, a totally different set of 300 quantitative peptide ratios were stable in HeLa. This was obviously surprising to us, but after thinking about it for a while it kind of makes sense.

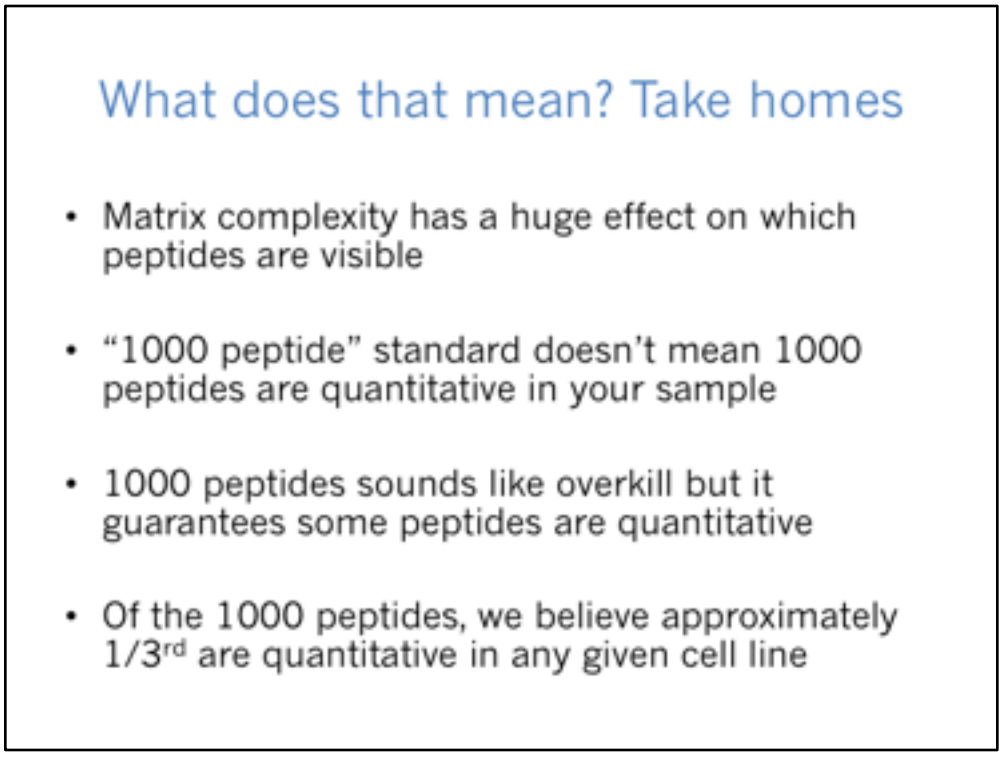

The complexity of your matrix has a huge effect on which peptides are visible, and while 1000 peptides originally sounded like massive overkill, it guarantees that at least 300 of them will be quantitative, enough to do statistics.

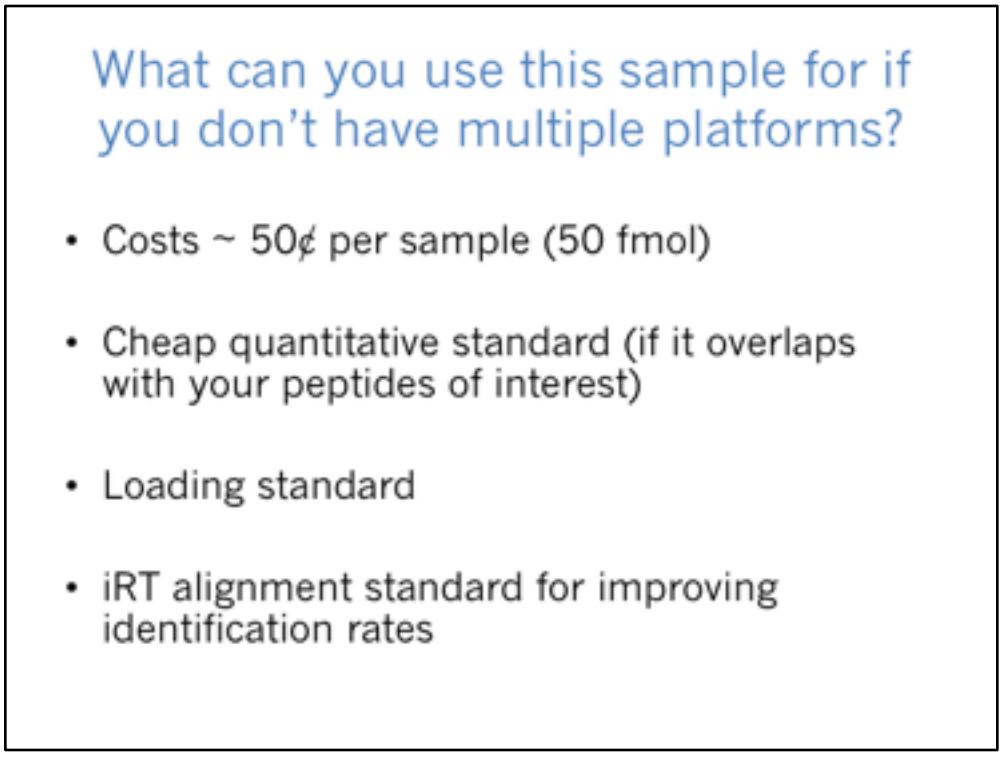

So why would you want to use this standard if you weren't doing a cross lab study? Well, for starters if you buy in bulk it's quite cheap. If you're interested in quantifying some of the proteins in our standard it would be a really inexpensive way to do that. But it can also be used as a loading standard to determine how much total protein was actually in your samples, or as a much tighter fitting iRT alignment standard than any other on the market because of the number of points in your curve.

As I said, the standard is available from JPT as the SpikeMix ABRF cross-species standard.

sPRG members involved in this study

Christopher Colangelo (Chair)	**Primary Ion**
Craig Dufresne	Thermo Fisher Scientific
David Hawke	UT MD Anderson Cancer Center
Gordana Ivosev	Sciex
Toni Koller	Columbia University
Brett Phinney (EB Liaison)	University of California, Davis
Kristie Rose	Vanderbilt University
Paul Rudnick	Spectragen Informatics
Brian C. Searle	Proteome Software / U. of Washington
Scott A. Shaffer	U. of Massachusetts Medical School
and:	
Brendan MacLean	U. of Washington
Vagisha Sharma	U. of Washington

Again, this was obviously a huge study that was put together by several people, including some volunteers from the Skyline team at University of Washington.

And that brings us to 2017. We've got a few ideas for a new study, but we need new members to help us put them into action!

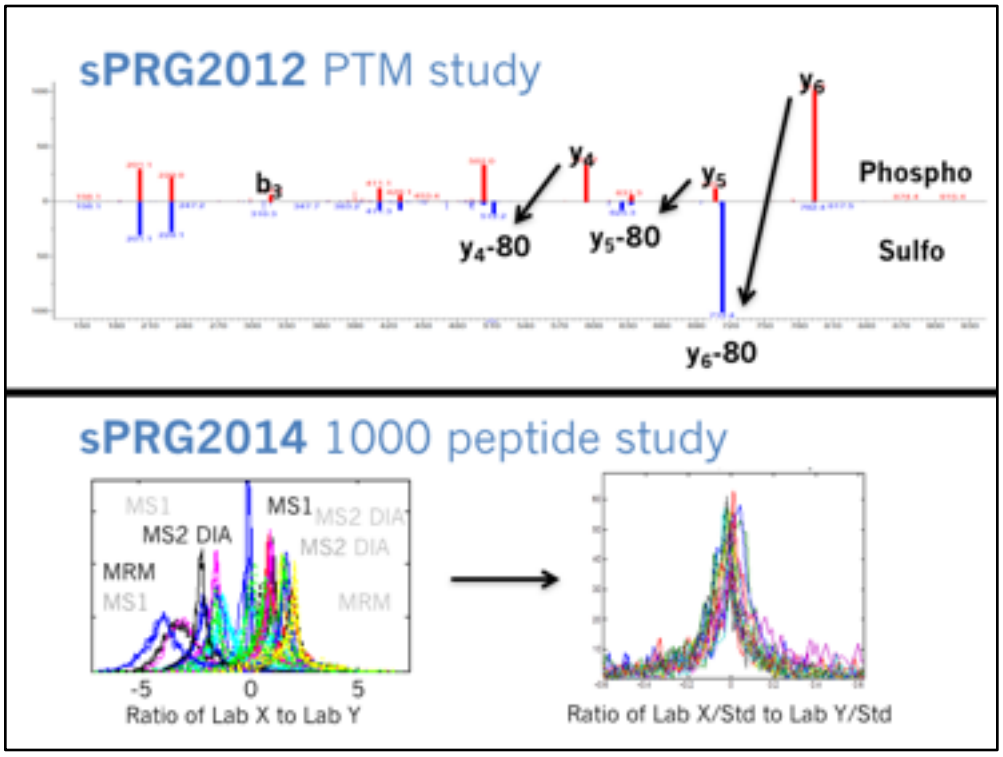

And with that I'd be happy to take questions on either of these studies.