Many (all) the questions have been answered during the Q&A period. Nevertheless, we ask that you provide written answers below so students can come back to read them again. Thanks!

1. (Page 20) Is there an intuitive way to understand why the m(W) corrections depend quadratically on m(t) but logarithmically on m(H)? Does it follow simply from the coupling structure?

   This question is best answered by a theorist, but I will offer my understanding. First, about the strong dependence on the top quark mass. Recall how on page 6 of my presentation, we understood the transformations induced by W and Z bosons on the fermions "sitting" at the poles of the sphere. The charged W bosons swapped the fermions sitting on poles while the W3 boson (before its mixing with the hypercharge boson that produces the physical Z boson and photon as the superposed states) rotates the sphere around the equator. Thus, the symmetry between the charged W bosons and the W3 boson is exact if the fermions at the poles are exactly symmetric, ie. they have the same mass. Clearly, this symmetry is badly violated by the large difference between the top mass and the bottom quark mass. This asymmetry is expected to differentiate between the charged W bosons (which transform top and bottom quarks to each other) and the W3 boson (which does not change the type of quark).

   Regarding the quadratic dependence on the top quark mass - it is actually a dependence on the mass-squared difference between top and bottom quarks. One way to visualize this dependence is to consider a W boson emitting a left handed massless top quark. With a Higgs VeV insertion it transforms into a right-handed top quark. In order to be reabsorbed by the W boson (to complete the quantum loop fluctuation), there needs to be another Higgs VeV insertion so that the right-handed top quark transforms back into a left-handed top quark. Because of the two Higgs VeV insertions there is a dependence on the square of the top-Higgs Yukawa coupling. Similarly there will be two Higgs VeV insertions for the bottom quark as well. The proportionality to the Higgs VeV squared and the Yukawa-squared from the two Higgs insertions can be re-written as the mass-squared.

   Recall that such quantum loops in the W3 propagator are already accounted for in the measured Z boson mass (after accounting for the Z-photon mixing).

Therefore the quantum correction for the charged W boson mass, after taking the measured Z boson mass into account, only adjusts for the symmetry-breaking effect of the top mass and bottom mass difference.

Now about the logarithmic dependence on the Higgs boson mass. The Higgs field also has the pairing structure ("SU(2) doublet") but the difference from the top-bottom situation is that there is no explicit violation of this pairing symmetry. So the logic that explains the large M_top-squared - M_bottom-squared dependence does not apply here. The reason that there is a weak dependence on Higgs mass is because one component of the Higgs doublet develops a VeV, while the other does not (we know the vacuum is not electrically charged). So the Higgs VeV creates a symmetry-breaking effect within the Higgs doublet and this ends up creating the logarithmic dependence on the Higgs mass (due to the Higgs loop in the W propagator w.r.t the Z propagator).

Please check this intuition with a theorist.

2. Can you comment more generally about the missing higher order corrections in the version of RESBOS you employed as well the potential effect of newer PDFs on the extracted MW?

   Excellent question. There has been some mention of this in the media. There are two aspects of the generator that have some relevance to this measurement.

   The first is the transverse momentum (pT) of the bosons. Here, we are not relying on the a-priori generator parameters to be super-accurate. We adjusted the generator parameters so that the pT distribution of the Z boson agreed with the data. This is shown on slide 54. This makes the analysis insensitive to the higher-order improvements in the generator, because tuning on the data already incorporates the higher-order corrections, within the quoted uncertainties of the data.

   In the past, we did rely on the generator to transfer this tuning from Z bosons to W bosons. One might argue that the pT-scale of Z and W boson production are different, so this transference might not be sufficiently accurate. Therefore, in the latest version of the analysis, an additional constraint from the data was included on the generator; this time from the pT distribution of the W boson data directly. This second constraint is shown on page 82. Now we are insensitive to higher-order improvements on the scale variation as well. As an aside, no other hadron collider analysis of MW has used the latter constraint.

A newer version of RESBOS became available at some point, so the RESBOS authors performed a comparison between RESBOS1 and RESBOS2 and documented it in an archive paper. This is mentioned on page 53 of my slides. If you see Table 2 of their paper, you will see that they found a 1.5 MeV effect on MW, which is quite small and would have no material effect on our result and conclusion.

The second aspect of the generator is the polarization of the W boson as a function of pT(boson). At pT=0 the polarization is fixed by the V-A nature of the weak interaction. But as pT increases there are QCD effects that alter the W polarization in the laboratory frame of reference. The polarization of the W boson affects the polar angle distribution of the decay leptons, which affects the pT distributions of the latter.

The polarization of the bosons is expected to depend very weakly on the higher-order QCD processes - see for example papers based on the NNLOjet+RADISH generator. This expectation is further borne out by the RESBOS1 vs RESBOS2 comparison mentioned above.

About the PDFs: take a look at page 85 of my slides. The consistency between existing PDFs is about 2 MeV and well within the quoted PDF uncertainty of 4 MeV. So this is the level of potential effect of newer PDFs. Note that the PDFs are already pretty well-constrained in the x-range of 0.01-0.1 that is relevant for the Tevatron. Furthermore, at the Tevatron the W boson production is dominated by valence quarks (85% of the W boson events have at least one valence quark) which are pretty well constrained. The dependence on sea quarks and gluons and heavy flavor is small, unlike at the LHC. At the Tevatron, there is no b-quark induced production and the rate for charm-induced production is only about 2%. The strange-induced production is about 10%. These heavy-flavor fractions are much smaller than at the LHC.

3. Is there any (measured) mass difference between W+ and W- ?

   This is published in Table S10 of the SCIENCE paper (in the online supplemental). The measured difference is consistent with zero within the statistical error.

4. (Page 47) For COT alignment, are the wires assumed to be straight? Or do you allow for them to bow, whether that is due to gravity or electrostatic forces? How does this bow compare with hit resolution?

   A great question. The wires are not straight; there are indeed large deflections from a straight line due to both the gravitational sag and the electrostatic deflection. In fact, these two modes of deflection are coupled. Both deflection effects were calculated *a-priori* using a detailed finite-element analysis (FEA) of the COT. The shapes were further refined as part of the cosmic-ray alignment. You can find these details in the alignment paper mentioned on page 48-50 of my slides. There are additional details in the 2008 and 2014 PRD papers from CDF on the W mass measurement, as well as the dedicated papers on the COT that are referenced in the SCIENCE paper.

   Both the gravitational sag and the electrostatic deflection are larger than the hit resolution, and therefore they are calculated and measured carefully and incorporated in the alignment and track reconstruction procedures.

5. (Page 51) Does CDF COT only do alignment calibration using tracks passing through the center? Are radial misalignments taken into account? If so, it would seem a global fit is needed to simultaneously calibrate in (r, phi)?

   Another great question. The alignment is ultimately used for collider tracks that pass through the center. For such tracks there is a redundancy between radial alignment and transverse alignment. Therefore there is no direct benefit from performing a simultaneous (r,phi) alignment. In fact it could be dangerous since any mistake in radial alignment can induce an overall scale factor in the COT which can alter the momentum scale. The last thing we want to do in the alignment procedure is to unknowingly introduce a momentum scale shift.

   For this reason, cosmic rays passing close to the center are used for the alignment. However, your point is well taken. The cosmic rays passing far from the center are used to study the radial alignment after the transverse alignment is completed. These "radial alignment" corrections are computed and examined for any large deviations or systematic patterns. No large deviations or systematic patterns were found and it was deemed that there was insufficient motivation to make wire-by-wire radial corrections.

Finally, note that the precision survey measurements of the wire positions at the COT end-plates had been performed a-priori and these survey-based alignments are used as the starting point of the cosmic-ray based alignment. The accuracy of the survey is roughly the same in the transverse direction as in the radial direction. However the track parameters are far more sensitive to the transverse alignment, so it makes sense to fine-tune the transverse alignment. But given the observed survey precision of 50 microns, it is not worth it to try to improve the radial alignment more than this.

6. What do you think about the prospects of precision M_W measurements from D0, CMS and ATLAS?

   The D0 management can confirm this, but my understanding is that the tracking detector(s) of D0 have suffered radiation damage and the track's polar angle is not measurable at the level of accuracy needed for the M_W measurement. Note that the polar angle is needed for the M_Z measurement, it is not needed for the M_W measurement which uses only the transverse momentum components. However, when M_Z is used for momentum scale calibration, the polar angle becomes important for the M_W measurement ultimately. For this reason D0 has deemed the second half of the Run2 data to be unanalyzable for the M_W measurement.

   Now about the LHC: ATLAS and CMS certainly have a lot of data. At the LHC the W boson production mechanisms are more complicated - more QCD radiation, more sea-quark and gluon induced production at lower x-values than the Tevatron. Also, much more heavy-flavor induced production at the LHC than the Tevatron. The rapidity distribution is also broader at the LHC - though this is compensated by a larger rapidity coverage of the ATLAS and CMS trackers. The pileup is much larger at the LHC, but the trackers are designed for this. Nevertheless, the calorimeters will be sensitive to the pileup, and the jet backgrounds are larger at the LHC. Having said this, it ought to be possible to measure/constrain all these effects with the high-statistics collider data. It does take time to develop and understand all the analysis procedures carefully in order to perform all the calibrations with high precision. I am optimistic that this can be done.

One last comment. The four LEP experiments had published M_W measurements in 2004 whose average was consistent with the CDF measurement. Between 2004 and 2008, each LEP experiment revised the M_W value by about 80 MeV, resulting in an updated average that changed by about 40 MeV. It could be interesting to revisit the LEP measurements.

7. (Page 63) So what, then, of the 4 points in the high-pt end of the j/psi calibration? Is it possible that there is a nonlinear trend in the extrapolation?

   First, the invariant mass fits from which these points were obtained were examined carefully. There were no problems with the fits, so the points are kept no matter how they "look".

   Second, it is not quite correct to think of the momentum scale as an "extrapolation" of these points. The momentum scale is by definition a scale factor that is independent of curvature. Take a look at page 51 of my slides. You will see a polynomial response function of the tracker. The momentum scale is the c1 coefficient. You can convert this curvature response function to an invariant mass function that can be plotted on page 63. You will find that the c2 coefficient on page 51 becomes a slope on page 63. This c2 coefficient is caused by the particle's energy loss upstream of the COT - you can prove this analytically. So the tuning of the energy loss, as described in the paper, makes c2=0. Next, it is shown on page 51 that c0=0 with the final alignment. Therefore, c1 is the only coefficient left to fit for, which gives a flat line (a momentum scale) on page 63. Thus, there is no "extrapolation" involved. It is an extraction of the c1 coefficient after eliminating the c0 and c2 coefficients.

   We can try adding more parameters to the response function of page 51, as long as we maintain the function to be analytic. Within the statistical errors, there is no sensitivity to additional coefficients in the response function on page 51. Therefore, we leave these 4 points as they are and consider them a statistical fluctuation, regardless of appearances.

   You will note that there is a much larger statistical fluctuation in the third point from the right, but it does not get noticed as much.

   The last point I would make is a biased one - one cannot act on it, but I mention it because some people interpret these points as an indicator that the actual

M_W measurement should be lower than the one we quote. This guess actually goes in the wrong direction. If one wanted to make some adjustment so that these 4 points were to “fall in line” with the rest of the points, so that the plot would “look nicer”, then one would adjust the c0 and c2 parameters to flatten the points and then refit for the c1 coefficient. The end result of this fiddling would be that the M_W value would be even higher than the one we publish. This is why we are very careful in making analysis decisions only based on independent logics and not some perceived pattern in the a-posteriori data.

8. (Page 97) The new CDF measurement of M_W vs M_top seems to favor light SUSY over heavy SUSY. Can you quantify this preference? What range of SUSY mass scale is, say, 3-sigma from CDF measurement of M_W?

   Again, please confirm my answer with a theorist. My understanding is that SUSY particles in loops induce an M_W to M_Z difference for similar reasons as the top-quark and bottom-quark mass difference. In other words, the effect on M_W depends on the mass difference between the SUSY particles in an SU(2) doublet. But the effect depends on the fractional mass difference. So for light SUSY particles the mass difference can be proportionately small - however the SUSY particles can be heavy and if the mass difference is proportionately large, there is still the induced M_W effect. Thus, to my knowledge it is not possible to definitively isolate a SUSY particle mass scale.

   I understand there are papers from SUSY theorists that address your question. For example, it was mentioned in the Q&A session that Sven Heinemeyer, one of the co-authors of the SUSY scan shown in this plot (Figure 1 of the CDF SCIENCE paper) has written a paper that probably answers your question.