CHAPTER 12

DISCUSSION OF RESULTS

This chapter will explore various aspects of the results presented in Chapter 11. First, Section 12.1 will compare the observed 2σ limits to the expected 2σ limits and discuss relevant differences. In Section 12.2, correlations among the WCs will be studied, and two-dimensional scans for selected pairs of WCs will be presented. Section 12.3 will discuss the relative importance of the various bins and categories to the sensitivity to each WC. Next, Section 12.4 will explore the impacts of the uncertainties on the limits. Section 12.5 will compare the observed limits to the results of other analyses. Lastly, the limits will be interpreted in terms of the energy scale Λ in Section 12.6.

12.1 Comparison of observed limits to predicted limits

Figure 12.1 shows a comparison between the observed 2σ limits and the expected SM 2σ limits (obtained by fitting to Asimov data). We expect statistical fluctuations to cause differences between these limits. The differences between the observed and the expected 2σ ranges vary by WC from about 2% to about 70%. The anomalously large 70% difference corresponds to the $c_{t\varphi}$ fit in which the other WCs are profiled. As can be seen in the one-dimensional scans presented in Chapter 11, the second minimum for $c_{t\varphi}$ crosses the 2σ threshold in the Asimov fit, but not in the fit to the observed data; the widths of the 2σ intervals for this comparison are calculated from the extrema of the 2σ limits, thus explaining the large difference for this WC. Another noticeable difference is the $c_{t\varphi}$ limit in the case where all other WCs are

fixed to 0. While the widths of the intervals are similar, the range is significantly shifted (since the minimum in the Asimov fit is at $c_{t\varphi} = 0$, while the minimum in the data fit is near $c_{t\varphi} = -5$). This feature is similarly visible in the one-dimensional scans presented in Chapter 11.

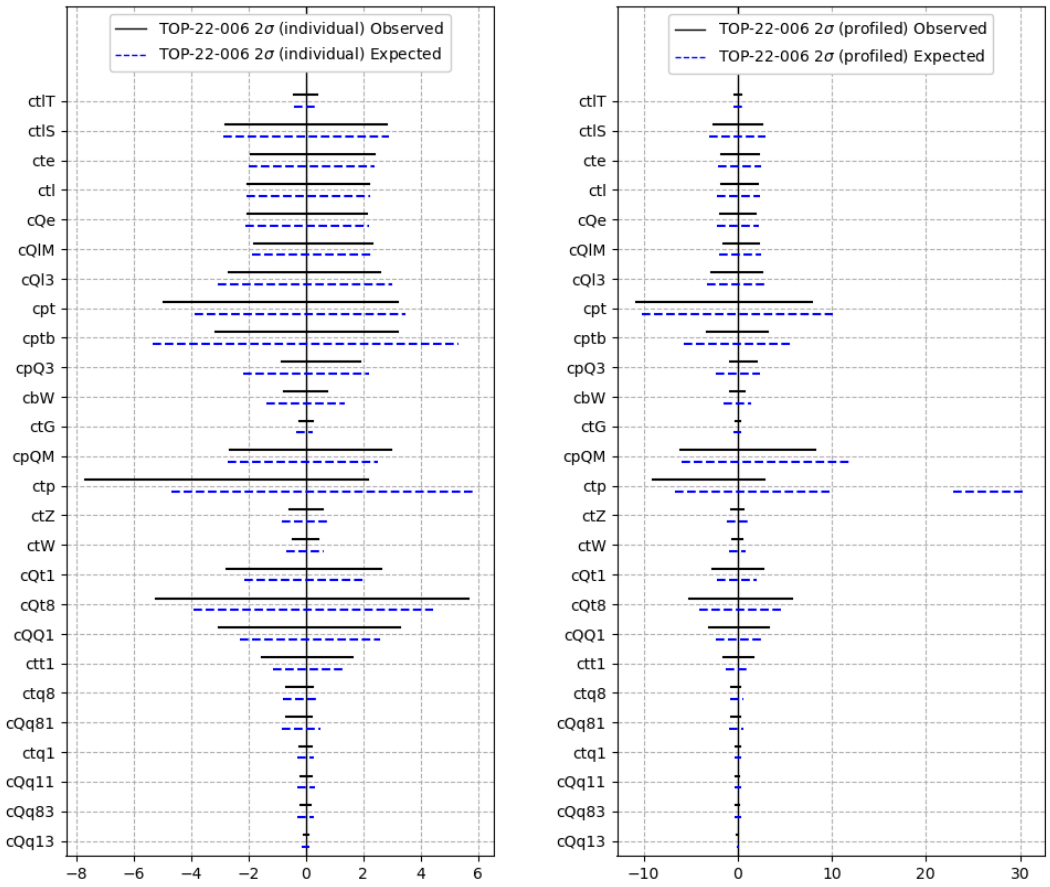

(a) (b)

Figure 12.1. The 2σ observed limits (solid black line) compared against the 2σ Asimov limits (dashed blue lines). Figure (a) shows the results of the likelihood fits in which a single WC is fit with all other WCs fixed to their SM values of zero. Figure (b) shows the results of the likelihood fits in which the other WCs are profiled. In these plots, $\Lambda = 1$ TeV.

12.2 Exploration of correlations among WCs

In Section 11.2, we explored the likelihood along one-dimensional directions of the full 26-dimensional surface. In principle, it would be interesting to study the shape of the full 26-dimensional surface; however, in practice it is very difficult to explore high-dimensional spaces. In lieu of the full 26-dimensional visualization, we can at least explore two-dimensional slices of the space. Similar to the one-dimensional scans described in Chapter 11, we can scan over two WCs and profile the remaining 24 WCs. However, there would be a total of 325 unique pairs of WCs to consider, and two-dimensional scans are significantly more computationally expensive than one-dimensional scans. For example, each the one-dimensional scans shown in Figure 11.3 was performed with 100 scan points, so a two-dimensional scan with the same granularity would require $100 \times 100 = 10,000$ scan points.

For these reasons, it would be beneficial to determine ahead of time a subset of the pairs of WCs that would be interesting to explore. Pairs of WCs are considered to be interesting if they have some correlation (i.e. if the value of one WC has some non-trivial relationship to the value of another WC). In Section 12.2.1, we will step through the method by which we identify potentially interesting pairs of correlated WCs. Then, in Section 12.2.2, the pairs of correlated WCs will be discussed and the two-dimensional scans for the pairs will be presented. Finally, Section 12.2.3 will discuss various factors that can affect the correlations.

12.2.1 Methodology of identification of correlations

In order to identify pairs of potentially interesting WCs, we might first look to the correlation matrix from the likelihood fit. While this provides information about correlations near the best fit point, the correlations may in principle vary across the 26-dimensional WC space. Thus, in order to explore the relationships among WCs throughout the full range of values explored within the 2σ intervals, we will examine

the values of the profiled WCs for each scan.

For example, let us consider the one-dimensional scan over the $c_t^{S(\ell)}$ WC, with the other 25 WCs profiled. While there are 25 profiled WCs to explore, let us first consider c_{tG}. For each of the scan points along the $c_t^{S(\ell)}$ direction, we can plot the value taken on by the profiled parameter c_{tG}, as shown in Figure 12.2 (a). In this plot, the profiled values of c_{tG} are plotted on the y axis for each of the points along the $c_t^{S(\ell)}$ likelihood scan, which are shown on the x axis. For a given value of the scanned parameter $c_t^{S(\ell)}$, the profiled value of c_{tG} corresponds to the value of c_{tG} that maximizes the likelihood at the given value of $c_t^{S(\ell)}$. Since the profiled values of c_{tG} are zero across the full range of the $c_t^{S(\ell)}$ scan, we may conclude that c_{tG} and $c_t^{S(\ell)}$ are not correlated.

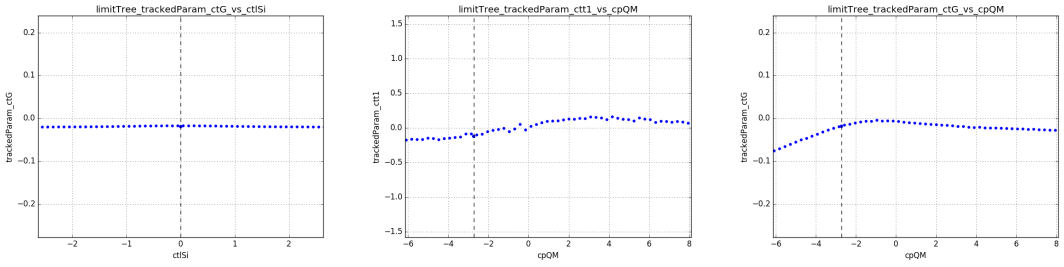

(a) (b) (c)

Figure 12.2. Plots showing profiled WC values (on the y axis) vs the value of the scanned parameter (on the x axis). In (a), the scanned parameter is $c_t^{S(\ell)}$ and the profiled WC shown is c_{tG}. In (b), the scanned WC is $c_{\varphi Q}^-$ and the profiled WC shown is c_{tt}^1. In (c), the scanned parameter is again $c_{\varphi Q}^-$, but this time the profiled parameter shown is c_{tG}. Figure (a) shows essentially no signs of interplay between the WCs, (b) shows very minimal interplay, and figure (c) shows potentially moderate interplay.

We may then repeat this procedure for each of the other 24 profiled WCs in the $c_t^{S(\ell)}$ scan, and similarly repeat the procedure for the scans over the other 25 WCs as well. In total, this results in $26 \times 26 - 26 = 650$ total plots to consider (not including the 26 plots that would show a WC vs itself). The majority of the 650 plots are either entirely flat (e.g. Figure 12.2 (a)), nearly flat (e.g. Figure 12.2 (b)), or show only moderate signs of correlation (e.g. Figure 12.2 (c)).

However, some pairs of WCs show clearly significant indications of interplay between the WCs. For example, in Figure 12.3 (a), we see that the value of the profiled parameter has a feature in its shape; a feature like this may arise if the profile fit is avoiding a "hill" in the likelihood surface, as discussed in Chapter 10. Figure 12.3 (b) shows an example of two WCs with a linear correlation. The pairs of WCs with features like the examples shown in Figure 12.3 are cases that may be interesting to explore in more detail (i.e. with two-dimensional scans).

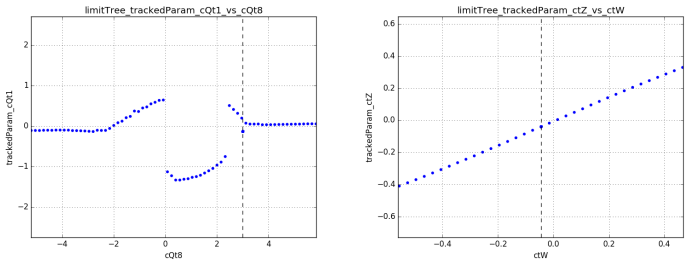

(a) (b)

Figure 12.3. Plots showing profiled WC values (on the y axis) vs the value of the scanned parameter (on the x axis). In (a), the scanned parameter is c_{Qt}^{8} and the profiled WC shown is c_{Qt}^{1}. In (b), the scanned WC is c_{tW} and the profiled WC shown is c_{tZ}. Figure (a) shows that there is some non-trivial correlation between the two WCs, and (b) also shows a correlation between the WCs (in this case the correlation is linear).

All 650 of the profiled WC vs scanned WC plots are shown together in Figure 12.4. In this "matrix" of plots, the first six rows/columns correspond to the two-light-two-heavy WCs, the next four correspond to the four-heavy WCs, the next nine correspond to the two-heavy-plus-boson WCs, and the final seven correspond to the two-heavy-two-lepton WCs. The diagonal of the matrix contains the plots that show a WC vs itself; by definition a WC must be correlated with itself, so the plots along the diagonal do not provide us with any new information. However, examining the off-diagonal plots, we can identify pairs of WCs that have potentially interesting correlations by looking for plots that show non-trivial relationships between the WCs.

In Figure 12.4, the background color of the individual plots indicates the approximate degree to which the WCs are correlated. Plots with a white background color show essentially no signs of correlation, as exemplified in Figure Figures 12.2 (a). The light grey and light blue plots show minimal or moderate signs of interplay, as exemplified by Figures 12.2 (b) and (c), respectively. The orange plots show pairs of WCs where there is clearly interplay between the WCs, e.g. as shown in Figure 12.3 (a), while the red plots show pairs of WCs with a linear correlation, e.g. as shown in Figure 12.3 (b). In the following section, we will explore the pairs of WCs that show significant signs of correlation.

12.2.2 Correlated pairs of WCs

The pairs of WCs that have clear signs of interplay are indicated with red or orange backgrounds in Figure 12.4). In this section, we will discuss these WCs and examine two-dimensional scans for these pairs. The full set of these correlated pairs of WCs is listed in Table 12.1.

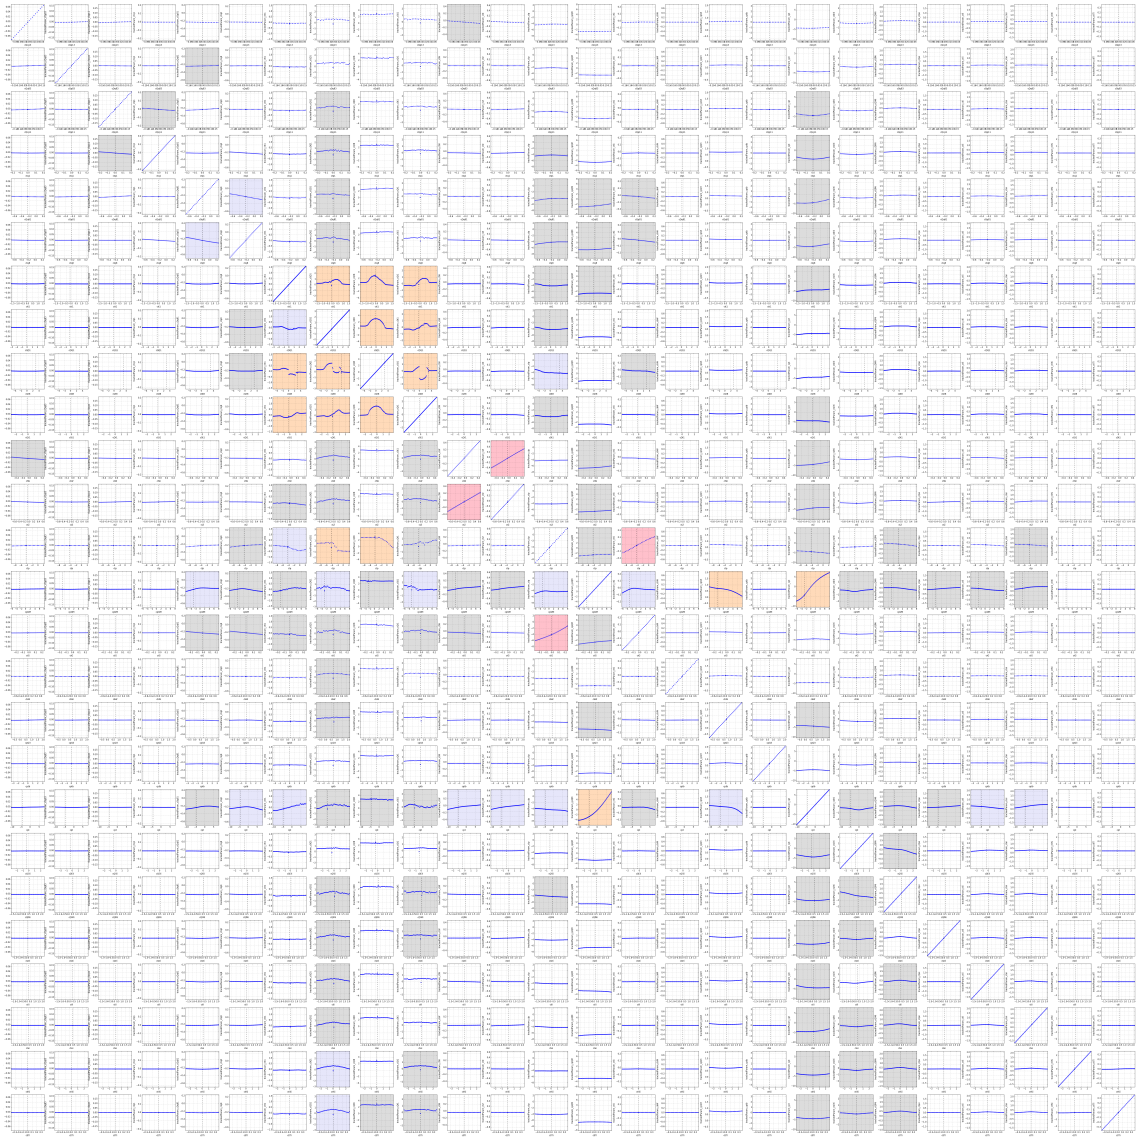

Figure 12.4. All 676 profiled WC vs scanned WC plots. The plots without a background color show no signs of correlation between the WCs (i.e. the plot is essentially flat), grey indicates minimal correlation, blue indicates potentially moderate correlation, orange indicates significant correlation between the WCs, and red indicates a linear correlation. The order of the WCs in the rows and columns is as follows: two-light-two-heavy WCs (c_{Qq}^{31}, c_{Qq}^{38}, c_{Qq}^{11}, c_{tq}^{1}, c_{Qq}^{18}, c_{tq}^{8}), four-heavy WCs (c_{tt}^{1}, c_{QQ}^{1}, c_{Qt}^{8}, c_{Qt}^{1}), two-heavy-plus-boson WCs (c_{tW}, c_{tZ}, $c_{t\varphi}$, $c_{\varphi Q}^{-}$, c_{tG}, c_{bW}, $c_{\varphi Q}^{3}$, $c_{\varphi tb}$, $c_{\varphi t}$), and two-heavy-two-lepton WCs ($c_{Q\ell}^{3(\ell)}$, $c_{Q\ell}^{-(\ell)}$, $c_{Qe}^{(\ell)}$, $c_{te}^{(\ell)}$, $c_{t}^{S(\ell)}$, $c_{t}^{T(\ell)}$).

TABLE 12.1

PAIRS OF CORRELATED WCS BASED ON FIGURE 12.4.

Four heavy - four heavy:
c^1_{QQ}- c^1_{Qt} , c^1_{QQ}- c^8_{Qt} , c^1_{QQ}- c^1_{tt} , c^1_{Qt}- c^8_{Qt} , c^1_{Qt}- c^1_{tt} , c^8_{Qt}- c^1_{tt}
Four heavy - two heavy with bosons:
c^1_{QQ}- $c_{t\varphi}$, c^8_{Qt}- $c_{t\varphi}$
Two heavy with bosons - two heavy with bosons:
$c^-_{\varphi Q}$- $c_{\varphi t}$, $c^-_{\varphi Q}$- $c^3_{\varphi Q}$, c_{tZ}- c_{tW} , c_{tG}- $c_{t\varphi}$

All six pairs of four-heavy WCs show signs of correlations. Since these WCs all impact $t\bar{t}t\bar{t}$, it is not unexpected to observed correlations. The two-dimensional scans over each of these pairs of WCs is shown in Figure 12.5 (a) through (f). The color scale on the plot indicates the NLL (as explained in Chapter 10), and the overlaid black markers show the paths of the one-dimensional scans (within the 2σ range of the scan) for each of the WCs shown in the plot. Two of the four-heavy WCs also have some interplay with $c_{t\varphi}$. The $c_{t\varphi}$ WC affects $t\bar{t}t\bar{t}$ (as can be seen in Figure D.17), so correlation with the four-heavy WCs is not unexpected. The two-dimensional scans for each of these pairs of WCs is shown in (g) and (h) of Figure 12.5.

There are also many correlations among the two-heavy-with-boson WCs. The correlations seem to arise primarily among WCs that have a strong effect on $t\bar{t}Z$ (e.g. $c^-_{\varphi Q}$, $c_{\varphi t}$, and c_{tZ}), but c_{tG} and $c_{t\varphi}$ (which primarily affect $t\bar{t}H$) also show correlations. The two-dimensional scans over each of these pairs of WCs are shown in Figure 12.6.

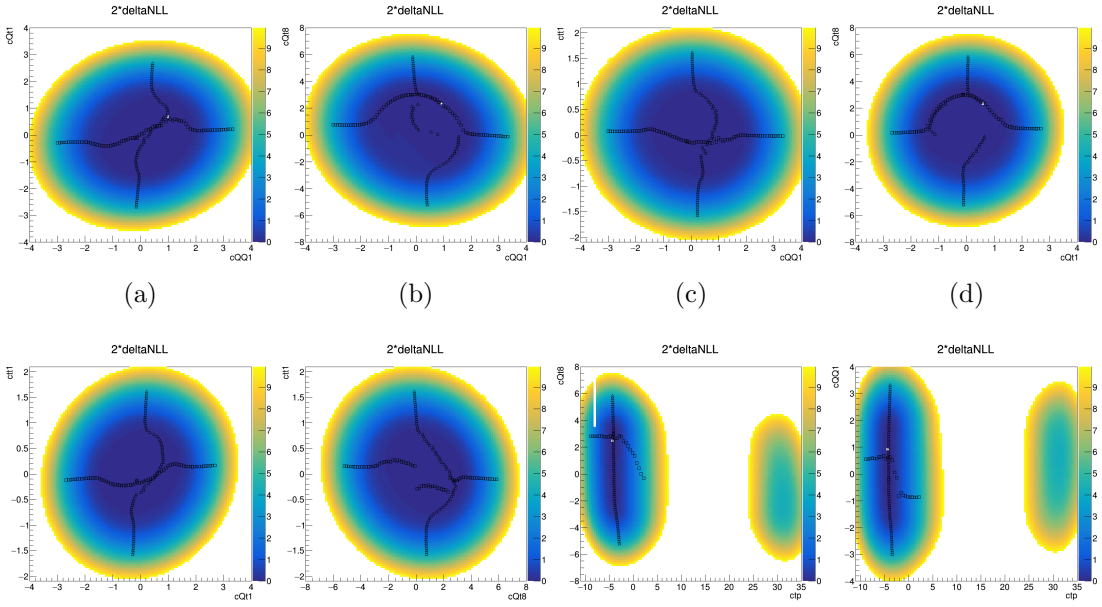

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 12.5. Two-dimensional scans over the pairs of WCs from Table 12.1 involving four-heavy WCs. The other 24 WCs are profiled in the likelihood fit. The color indicates the NLL, and the overlaid black markers show the paths of the one-dimensional scan (within the 2σ range of the scan) for each of the WCs shown in the plot. The path of the one-dimensional scan along the x-axis WC is shown with square markers, and the path of the one-dimensional scan along the y-axis WC is shown with triangular markers. Figure (a) shows c^1_{QQ}-c^1_{Qt}, (b) shows c^1_{QQ}-c^8_{Qt}, (c) shows c^1_{QQ}-c^1_{tt}, (d) shows c^1_{Qt}-c^8_{Qt}, (e) shows c^1_{Qt}-c^1_{tt}, (f) shows c^8_{Qt}-c^1_{tt}, (g) shows $c_{t\varphi}$-c^8_{Qt}, and (h) shows $c_{t\varphi}$-c^1_{QQ}.

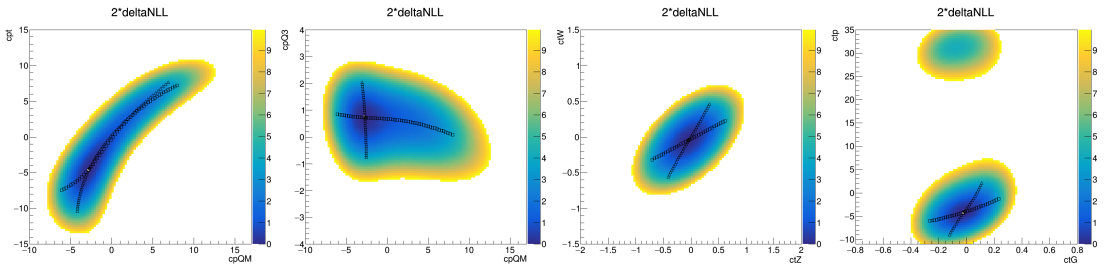

(a) (b) (c) (d)

Figure 12.6. Two-dimensional scans over the pairs of WCs from Table 12.1 involving only two-heavy-with-boson WCs. Figure (a) shows $c_{\varphi Q}^{-}$-$c_{\varphi t}$, (b) shows $c_{\varphi Q}^{-}$-$c_{\varphi Q}^{3}$, (c) shows c_{tZ}-c_{tW}, and (d) shows c_{tG}-$c_{t\varphi}$. All other relevant details of the plots are described in the caption to Figure 12.5.

Although this section has focused primarily on correlated WCs, we should keep in mind that the vast majority of pairs show little to no signs of correlation. Figure 12.7 (a) and (b) shows two examples of WCs with only a small amount of interplay (i.e. an example of pairs that are indicated with a blue background color in Figure 12.4). The pairs shown in this example are $c_{t\varphi}$-$c_{\varphi t}$ and c_{tG}-$c_{\varphi Q}^{-}$. Examples of WCs with essentially no correlation (i.e. a white background color in Figure 12.4) are shown in Figure 12.7 (c) and (d). The pairs in this example are $c_{Q\ell}^{-(\ell)}$-$c_{Qe}^{(\ell)}$ and $c_{\varphi tb}$-$c_{Q\ell}^{3(\ell)}$.

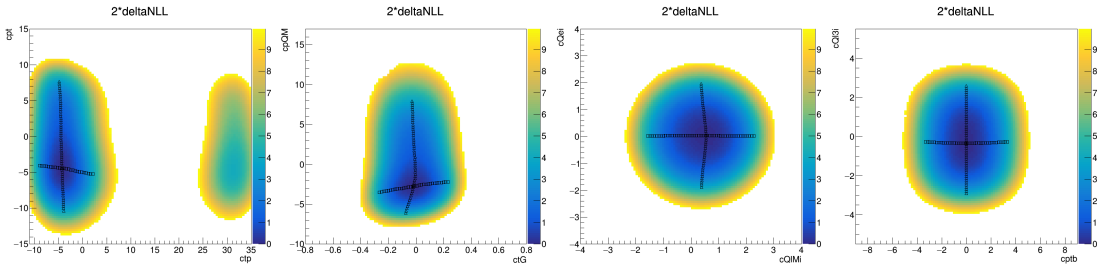

(a) (b) (c) (d)

Figure 12.7. Two-dimensional scans over the pairs of WCs that do not show signs of strong correlations. Figure (a) shows $c_{t\varphi}$-$c_{\varphi t}$, (b) shows c_{tG}-$c_{\varphi Q}^{-}$, (c) shows $c_{Q\ell}^{-(\ell)}$-$c_{Qe}^{(\ell)}$, and (d) shows $c_{\varphi tb}$-$c_{Q\ell}^{3(\ell)}$. All other relevant details of the plots are described in the caption to Figure 12.5.

12.2.3 Discussion of factors affecting correlations

In the previous section, we identified and plotted pairs of correlated WCs. In this section, we will explore what these correlations may imply, and discuss various factors that can lead to such correlations.

First, to specify the terminology, a correlation between two WCs simply means that there is some relationship between their values in the likelihood fit. This relationship could be caused by a number of factors. For example, if two WCs have primarily linear effects on the same set of bins, they may exhibit a linear correlation or anti-correlation (depending on the signs of the interference terms) since as the fit increases the value of one WC, the fit is free to simultaneously increase or decrease the value of the other WC to compensate for the effects of the first WC. Significant interference between two WCs may also lead to a correlation between the WCs when the interference term (which may be positive or negative) is able to compensate for the quadratic pieces (which may only be positive). The c_{tW} and c_{tZ} WCs represent an example of this type of correlation, which will be discussed further in Section 12.3.

It is also interesting to consider cases in which the WCs in question are dominated by purely quadratic components (i.e. cases in which all interference terms are small compared to the purely quadratic contributions). In this scenario, the WCs will not be able to compensate for one another, as turning any WC to a non-zero value can only increase the yield. Thus, in the SM scenario (where the prediction is exactly equal to the observation), the likelihood can never be improved by turning a second WC to a non-zero value, so the WCs will appear to be completely uncorrelated in the likelihood fit. However, it is interesting to note that this absence of a correlation does not necessarily imply that the fit is able to distinguish between the effects of the WCs in question; when the purely quadratic terms dominant, the WCs will be uncorrelated even if their effects are completely degenerate. Thus, we should be careful to avoid the assumption that uncorrelated WCs are necessarily non degenerate.

As a likelihood fit scans across a WC, the relative contributions of the linear and quadratic pieces will change. For some scans, we may begin in a region where linear and interference terms are relevant, but eventually transition into a region where the purely quadratic contributions dominate. The scans over the four-heavy WCs provide examples of this effect. Near zero, interference terms are relevant and the WCs may play off of one another (as illustrated in the two-dimensional scans shown in Figure 12.5). However, as the value of the scanned parameter grows larger, the purely quadratic terms become dominant, causing the yield to increase monotonically, which ultimately drives the likelihood past the 2σ threshold. This explains why the profiled vs tracked parameter plots for the four-heavy WCs (e.g. Figure 12.3 (a)) show signs of correlations near zero, but do now show signs of correlations at values far from zero.

It is also important to consider cases where the observation deviates from the prediction, as these effects can further complicate the interpretation of the correlations. For example, if the observed yield is larger than the predicted yield, effects from multiple WCs may combine to match the observation. Thus, even if two WCs have no interference (i.e. their effects may only increase the predicted yield), they may become correlated with one another in cases where the observation is larger than the prediction. If the observation is less than the prediction, interference effects (both with the SM and among WCs) will be important as well, since these terms can be negative and are thus able to compensate for positive contributions arising from purely quadratic effects.

In conclusion, there are multifarious factors that may influence correlations among WCs. When exploring and interpreting the sensitivity to the WCs (as we will do in the following section, 12.3), it is important to keep in mind interferences, degeneracies, and how the observation compares to the prediction for the relevant bins in the likelihood fit.

12.3 Interpretation of sensitivity to WCs

In this section, we will discuss the sensitivity to the WCs (defined by the expected 2σ limits from the profiled likelihood fits) with the goal of understanding which bins or categories of bins provide the most important contributions to the sensitivity for each WC. While all 178 analysis bins contribute to the sensitivity to the 26 WCs, the relative contribution of each bin varies by WC. Organizing the WCs based on the relative contributions of the categories of analysis bins, the WCs may be classified into seven main groups, summarized in Table 12.2.

In order to arrive at the categorization presented in Table 12.2, we have employed two complimentary approaches. Starting from the groups of operators, we can first consider their vertices and the processes they affect in order to gain insight into the analysis bins that may be particularly impacted by the given WCs. In parallel to this "top-down" approach, we also study individual bins, comparing the prediction to the observation at the 2σ limits and taking into account the uncertainty on the prediction in order to identify the bins that contribute most significantly. While this section will not focus on the methodology through which the categorization was determined, the technical details of the "bottom-up" approach can be found in Appendix H. This section will instead aim to discuss the conclusions of the study.

It should be emphasized that the conclusions summarized in Table 12.2 represent a simplified picture of the interpretation of the sensitivity; while there are indeed some cases where the majority of the sensitivity to a WC is derived from relatively clear subset of the analysis bins, the sensitivity to many of the WCs is provided by a diverse combination of bins across all selection categories. Furthermore, when characterizing relevant bins, it is also important to keep in mind interference and correlations among WCs. The following sections will step through each of the groups of WCs outlined in Table 12.2, discussing the subsets of bins that provide the leading contributions to the sensitivity.

TABLE 12.2

SUMMARY OF CATEGORIES THAT PROVIDE LEADING CONTRIBUTIONS TO THE SENSITIVITY FOR SUBSETS OF THE WCS.

Grouping of WCs	WCs	Lead categories
Two heavy two leptons	$c_{Q\ell}^{3(\ell)}$, $c_{Q\ell}^{-(\ell)}$, $c_{Qe}^{(\ell)}$, $c_{t\ell}^{(\ell)}$, $c_{te}^{(\ell)}$, $c_{t}^{S(\ell)}$, $c_{t}^{T(\ell)}$	3ℓ off-Z
Four heavy	c_{QQ}^{1}, c_{Qt}^{1}, c_{Qt}^{8}, c_{tt}^{1}	2ℓss
Two heavy two light "$t\bar{t}l\nu$-like"	c_{Qq}^{11}, c_{Qq}^{18}, c_{tq}^{1}, c_{tq}^{8}	2ℓss
Two heavy two light "$t\bar{l}\bar{l}q$-like"	c_{Qq}^{31}, c_{Qq}^{38}	3ℓ on-Z
Two heavy with bosons "$t\bar{t}l\bar{l}$-like"	c_{tZ}, $c_{\varphi t}$, $c_{\varphi Q}^{-}$	3ℓ on-Z and 2ℓss
Two heavy with bosons "tXq-like"	$c_{\varphi Q}^{3}$, $c_{\varphi tb}$, c_{bW}	3ℓ on-Z
Two heavy with bosons with significant impacts on many processes	c_{tG}, $c_{t\varphi}$, c_{tW}	3ℓ and 2ℓss

12.3.1 WCs from the two-heavy-two-lepton category of operators

Beginning with the WCs in the two-heavy-two-lepton group, the 3ℓ off-Z channels provide the majority of the sensitivity for these WCs. To quantify the contributions of the off-Z channels, a fit is performed with only this subset of bins included; the resulting 2σ profiled limits show that the expected sensitivity is only degraded by

about 6% compared to the results when all bins are included. It is perhaps unsurprising that the off-Z bins provide the dominant contributions to our sensitivity to this group of WCs, as these WCs are associated with four-fermion vertices that produce pairs of leptons without an intermediate Z.

12.3.2 WCs from the four-heavy category of operators

The next group of WCs are those associated with the two-light-two-heavy operators. The sensitivity to these WCs is provided primarily by the 2ℓss bins, with leading contributions from the bins requiring at least three b tags. Since the $t\bar{t}t\bar{t}$ process contributes significantly to these bins and the four-heavy WCs strongly impact $t\bar{t}t\bar{t}$, it is expected that these bins would contribute significantly to the sensitivity. To obtain a quantitative characterization of the sensitivity provided by the 2ℓss bins, we performed a fit with only these bins included; the resulting 2σ limits are only degraded by about 5% (with respect to a fit with all bins included), showing that the 2ℓss bins indeed represent the dominant source of sensitivity to the four-heavy WCs.

12.3.3 WCs from the two-heavy-two-light category of operators

The next set of WCs are those associated with the two-heavy-two-light category of operators. Four of these WCs (c_{Qq}^{11}, c_{tq}^{1}, c_{Qq}^{18}, and c_{tq}^{8}) primarily affect the $t\bar{t}l\nu$ process, so bins populated significantly by $t\bar{t}l\nu$ are expected to provide important contributions to the sensitivity to these WCs. Performing a fit with only the 2ℓss bins included, the expected 2σ limits are degraded by only about 5-15%. The 2ℓss bins thus provide the primary source of sensitivity for these WCs, though other bins (e.g. from the off-Z channels) also contribute to the sensitivity.

The remaining two WCs from the two-heavy-two light group (c_{Qq}^{31} and c_{Qq}^{38}) are distinct from the other two-heavy-two-light WCs in that they feature t$-$b$-$q$-$q$'$ vertices. These vertices allow c_{Qq}^{31} and c_{Qq}^{38} to significantly impact the $t\bar{l}q$ process in

3ℓ on-Z bins with two b tags and low jet multiplicity (as discussed in Section **??**). The on-Z bins thus contribute significant sensitivity to these WCs. While the 2ℓss and off-Z categories also contribute to the sensitivity to these WCs, the 3ℓ on-Z bins provide the leading contribution; the expected 2σ intervals for these WCs loosen by more than 30% when the 3ℓ on-Z bins are excluded from the fit.

12.3.4 WCs from the two-heavy-with-bosons category of operators

The final set of WCs are those associated with the two-heavy-with-bosons category of operators. These nine WCs impact a broad range of processes, leading to diverse effects across the full set of 178 analysis bins and making it challenging to definitively characterize subsets of bins that provide dominant contributions to the sensitivity. However, the WCs can be classified into three main groups (as listed in Table 12.2) based on the processed they impact most significantly.

The c_{tZ}, $c_{\varphi Q}^{-}$, and $c_{\varphi t}$ WCs feature t−t−Z EFT vertices and primarily affect the $t\bar{t}l\bar{l}$ process; the on-Z bins are thus important for these WCs. However, these WCs also impact other processes (e.g. $t\bar{t}t\bar{t}$), meaning other categories of bins can also provide important sensitivity. Furthermore, $t\bar{t}l\bar{l}$ also significantly populates the 2ℓss bins (making up about 20% of the total expected yield), so the $t\bar{t}l\bar{l}$ effects can also be relevant in the 2ℓss bins. Thus, the 3ℓ on-Z bins and 2ℓss bins are important for these WCs. The 3ℓ off-Z bins provide a smaller (though non-zero) contribution to the sensitivity; performing a fit with these bins excluded results in an approximately 6% degradation of the expected 2σ confidence intervals.

Next, let us consider $c_{\varphi Q}^{3}$, $c_{\varphi tb}$, and c_{bW}. These WCs primarily impact $tl\bar{l}q$ and tHq, and their sensitivity arises from multiple categories of analysis bins. The 3ℓ on-Z bins represent the leading (though not overwhelmingly dominant) contribution. Performing a fit with only the 3ℓ on-Z bins included, the expected 2σ limits for these WCs loosen by an average of only about 10% compared to a fit with all bins included.

The final three WCs from the two-heavy-with-bosons group are c_{tG}, $c_{t\varphi}$, and c_{tW}. Impacting multiple processes, these WCs gain sensitivity from the full spectrum of analysis bins. For example, c_{tG} impacts $t\bar{t}H$ (so the 2ℓss and 3ℓ off-Z bins are important as $t\bar{t}H$ significantly populates these bins) but also strongly impacts $t\bar{t}l\bar{l}$ (so the on-Z and 2ℓss bins also play an important role). The $c_{t\varphi}$ WC significantly impacts $t\bar{t}H$, tHq, and $t\bar{t}t\bar{t}$; most of the analysis bins provide sensitivity to this WC, though the on-Z bins provide only minor contributions (dropping the on-Z bins only results in about a 5% effect on the expected 2σ profiled limits for $c_{t\varphi}$). Finally, the c_{tW} WC impacts all signal processes and derives important sensitivity from a variety of analysis bins. Further complicating the picture, c_{tW} has significant interference with c_{tZ}, and the two WCs have a strong linear correlation in the profiled fit (as shown in Figure 12.6). Thus, when we consider the 2σ profiled limits for c_{tW}, it is important to recall that the c_{tZ} operator is also set to a non-zero value, so bins that c_{tZ} affects can also be important when considering the sensitivity to c_{tW}.

12.4 Impacts of uncertainties

The precision of this analysis is limited by both statistical and systematic uncertainties, with the relative importance of each source of uncertainty varying by WC. In Section 12.4.1, we will discuss how the total effect of the systematic uncertainties compares to the statistical uncertainty. In Section 12.4.2, the dominant sources of systematic uncertainty will be explored.

12.4.1 Relative importance of statistical and systematic contributions

To understand whether the analysis is dominated by statistics or systematics, we will begin by exploring the likelihood scans with Asimov data, in the case where all of the other 25 WCs are fixed to their SM values of zero. This simple case will allow us to probe the effects of the uncertainties without complications due to correlations among WCs or statistical fluctuations in the data. The resulting 1σ confidence intervals extracted for each WC essentially correspond to the 1σ error bars on the best fit value of the WC (which must be zero in this case, since we are fitting to Asimov data). Let us refer to this total uncertainty as σ_{Tot}.

The total uncertainty σ_{Tot} includes both statistical and systematic components; the goal of this study is to determine the relative contribution of statistical and systematic sources to σ_{Tot}. Assuming the statistical and systematic components are independent, σ_{Tot} should correspond to the quadrature sum of the components:

$$\sigma_{\text{Tot}} = \sqrt{\sigma^2_{\text{Stat}} + \sigma^2_{\text{Syst}}}, \tag{12.1}$$

where the statistical uncertainty is σ_{Stat} and systematic uncertainty is σ_{Syst}. Although σ_{Tot} is known, σ_{Stat} and σ_{Syst} are not yet known. Fortunately, it is relatively straightforward to obtain σ_{Stat}, and once this is known we can simply solve Eq. 12.1 to obtain σ_{Syst} as well. To find σ_{Stat}, we can perform the likelihood fit without the

systematic uncertainties (i.e. nuisance parameters) included in the fit. Since the statistical uncertainties are the only sources of uncertainty in this fit, the resulting 1σ confidence intervals for each WC from will correspond to the σ_{Stat} for each WC. We can then solve Eq. 12.1 to find σ_{Syst} for each WC as well.

Following this procedure, we obtain σ_{Stat} and σ_{Syst} for all 26 WCs. We can thus compare the size of σ_{Stat} and σ_{Syst} for each WC in order to understand if the analysis is dominated by statistical or systematic uncertainties. The results of this comparison show that for the majority of the WCs, σ_{Stat} is at least as large or larger that σ_{Syst} (by up to a factor of $\sim$4, depending on the WC). However, there are also cases where σ_{Syst} is larger than σ_{Stat} (by up to a factor of about $\sim$2 to about $\sim$3, depending on the WC). The WCs for which the systematic uncertainty is significantly larger than the statistical uncertainty are c_{Qq}^{18} and c_{tq}^{8} (from the two-heavy-two-light category of WCs) and $c_{t\varphi}$, $c_{\varphi Q}^{-}$, c_{tG}, and $c_{\varphi t}$ (from the two-heavy-with-bosons category of WCs). Section 12.4.2 will explore the dominant sources of systematic uncertainty for each of these WCs.

12.4.2 Impacts of the systematic uncertainties

In this section, we will step through the leading systematic uncertainties for each WC (as determined by the impact of the systematic on the best fit value of the WC, described in Appendix I).

For all four-heavy WCs, the leading systematic is the uncertainty on the NLO $t\bar{t}t\bar{t}$ cross section to which the LO $t\bar{t}t\bar{t}$ samples are normalized. Other systematics that have large impacts for the four-heavy WCs are the ISR systematic and the renormalization systematic. However, as discussed in Section 12.4.1, the four-heavy WCs are not dominated by systematics. Rather, for these WCs, the statistical uncertainty is larger than the total systematic uncertainty by a factor of about two.

For the two-heavy-two-lepton WCs, the leading systematic uncertainties vary by

WC, with the uncertainty on the fake estimation, the electron scale factor systematic, and the diboson N_{jet} systematic frequently making the largest impacts. However, like the four-heavy WCs, this group of WCs is not dominated by systematic uncertainties. The statistical uncertainty is generally larger than the systematic contribution by a factor of approximately three.

Next, let us consider the two-heavy-two-lepton WCs. One of the leading systematic uncertainties for these WCs is the uncertainty on the NLO $t\bar{t}l\nu$ cross section (to which we normalize the LO $t\bar{t}l\nu$ sample). Other theoretical uncertainties, such as renormalization and factorization, rank fairly high as well. Some experimental sources of uncertainty (primarily the uncertainty on the fake estimation and the diboson N_{jet} uncertainty) also have relatively large impacts. While the statistical uncertainty represents the dominant contribution for most of the two-heavy-two-light WCs, the systematic contribution is dominant for c_{Qq}^{18} and c_{tq}^{8} (by approximately a factor of two). For c_{Qq}^{18} and c_{tq}^{8}, the leading systematic is the uncertainty on the NLO $t\bar{t}l\nu$ cross section.

The final group of WCs are those associated with the two-heavy-with-bosons category of operators. The dominant systematic uncertainty varies by WC, but the uncertainty on the NLO $t\bar{t}l\nu$ or $t\bar{t}l\bar{l}$ cross section often represents the leading impact. The diboson N_{jet} uncertainty, the uncertainty on the electron SF, and the uncertainty on the fake estimation also have relatively large impacts. This group of WCs is evenly split between cases where the statistical contribution is larger, and cases where the systematic contribution is larger. As reported in Section 12.4.1, the systematic uncertainty dominates for $c_{\varphi t}$, $c_{\varphi Q}^{-}$, c_{tG}, and $c_{t\varphi}$. For the $c_{\varphi t}$, $c_{\varphi Q}^{-}$, and c_{tG} WCs, σ_{Syst} is larger than σ_{Stat} by a factor of approximately three, and the leading systematic is the NLO $t\bar{t}l\bar{l}$ cross-section uncertainty. For the $c_{t\varphi}$ WC, σ_{Syst} is larger than σ_{Stat} by a factor of almost two, and the leading systematic is the $t\bar{t}l\nu$ cross-section uncertainty.

12.5 Comparison of results to other analyses

In this section, the WC limits obtained by the analysis presented in this thesis will be compared against the results of other analyses. Section 12.5.1 will discuss CMS results, while Section 12.5.2 will explore the results presented by a global theory combination. For ease of notation throughout this section, the analysis presented in this thesis will be referred to as TOP-22-006.

12.5.1 Comparisons against other CMS analyses

A summary of EFT limits obtained by CMS analyses is collected in Ref. [82]. For each WC, this section will compare the tightest limit from Ref. [82] against the limit obtained by TOP-22-006. All comparisons will use the 2σ limits from the likelihood fits in which all other WCs have been fixed at their SM value of zero. The goal of the comparison is to contextualize the results of TOP-22-006 in order to help us to understand what value this analysis brings to the field and to identify opportunities for potentially beneficial combinations with other analyses.

Before examining the results of the comparison, there are two caveats that should be mentioned. First, this comparison will include the results of the recent Ref. [13] analysis, which are not yet included in [82]. Secondly, results from Ref. [19] (the direct predecessor of TOP-22-006) will be excluded from the comparison presented in this section; TOP-22-006 supersedes the results of [19], and there is no possibility of performing a combination with [19], so including [19] would not align with the goals of the comparison presented in this section; rather, Appendix J.1 provides a direct comparison of the TOP-22-006 results against the Ref. [19] results and discusses the factors contributing to the improvements in the limits.

Figure 12.8 summarizes the comparison of the TOP-22-006 results to the best results of other CMS analyses. TOP-22-006 is the only analysis to probe the two-heavy-two-lepton WCs (apart from Ref. [19], which is excluded from this comparison)

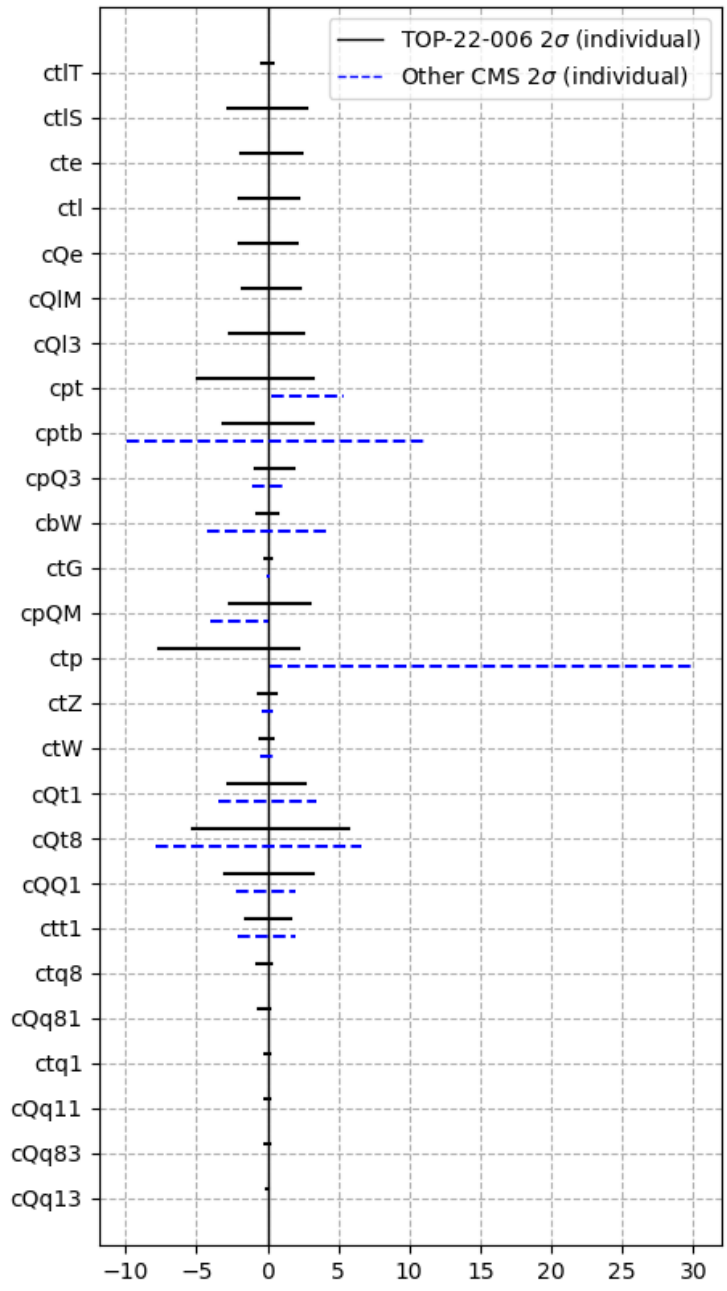

Figure 12.8. The 2σ individual limits of this analysis (indicated with the solid black line labeled TOP-22-006) compared against the individual limits obtained by other CMS analyses. The other CMS results for each WC are as follows: $c_{\varphi t}$ and $c^{-}_{\varphi Q}$ [9], c_{tW} and $c^{3}_{\varphi Q}$ [10], c_{tZ} [11], c_{tG} [12], c_{bW} $c_{\varphi tb}$ and $c_{t\varphi}$ [13], and all four-heavy WCs [14]. The details of each referenced analysis are discussed in the text. In this plot, $\Lambda = 1$ TeV.

and is the only analysis to probe the two-heavy-two-light WCs, so Figure 12.9 does not include comparisons for these WCs. For the other WCs, the TOP-22-006 2σ limits are generally comparable to or better than the best limits from other CMS analyses (by up to a factor of ∼6). However, for five of WCs (c^1_{QQ}, c_{tZ}, $c^-_{\varphi Q}$, $c_{\varphi t}$, and c_{tG}), the TOP-22-006 limits are approximately 40-60% looser than the best limits from other CMS analyses.

There are many differences between the analyses that may contributed to the variations in the observed results (e.g. the amount of data included, final states considered, EFT approach used, etc.). In cases where the analyses study a non-overlapping set of data events, a combination between the analyses may be interesting to pursue, as the combined result would be stronger than the results obtained by either analysis independently. A brief description of the specifics of each of the analyses featured in Figure 12.8 is provided below (where the WCs for which the analysis provides limits in Figure 12.8 are indicated parenthetically):

Ref. [14] (c^1_{Qt}, c^1_{QQ}, c^1_{tt}, c^1_{tt}): This analysis presents a search for $t\bar{t}t\bar{t}$ production in multilepton final states using data collected during 2016 (35.8 fb^{-1}). The limits from this analysis are looser than the TOP-22-006 limits (by about 30%) for all but the c^1_{QQ} WC (for which TOP-22-006 is looser by about 50%). This analysis makes use of single-lepton events and events with pairs of opposite-sign leptons, so its signal selection is orthogonal to TOP-22-006 and a combination would be potentially useful. However, the EFT approach used in this analysis (a reinterpretation of the cross section measurement) is very different than the detector-level approach used in TOP-22-006, so a combination may be technically difficult to implement.

Ref. [12] (c_{tG}): This analysis presents a measurement of the top quark polarization and $t\bar{t}$ spin correlations using data collected in 2016 (35.9 fb^{-1}). The analysis finds a better constraint on c_{tG} than TOP-19-001 (by about 50%). Opposite-sign pairs of leptons are selected, so the data does not overlap with TOP-22-006 and a combination would be potentially interesting, though the EFT approach of this analysis (a reinterpretation of a differential cross section measurement) is significantly different from the EFT approach used in TOP-22-006, so a combination may be technically challenging.

However, in future iterations of the TOP-22-006 analysis, it may be interesting to explore ways of exploiting the effect of c_{tG} on $t\bar{t}$ in order to improve the sensitivity to this WC (as well as to other WCs that strongly impact $t\bar{t}$, e.g. the two-light-two-heavy WCs).

Ref. [11] (c_{tZ}): This analysis presents a measurement of $t\bar{t}\gamma$ inclusive and differential cross sections using data collected in 2016-2018 (137 fb^{-1}). The limit on c_{tZ} from this analysis is better than the limit obtained by TOP-22-006 (by about 40%). The analysis selects events with a single lepton, so the data does not overlap with the data used in TOP-22-006, and a combination could be potentially interesting. However, the EFT approach of this analysis (a reinterpretation of the cross section measurement) is significantly different from the TOP-22-006 approach, so a combination may be technically challenging. Instead, it may be interesting to target $t\bar{t}\gamma$ events in a future iteration of the TOP-22-006 analysis in order to potentially gain additional sensitivity to the c_{tZ} WC (as well as other WCs that impact $t\bar{t}\gamma$).

Ref. [9] ($c^{-}_{\varphi Q}$, $c_{\varphi t}$): This analysis presents a measurement of top pair production in association with a Z in multilepton final states. The analysis uses data collected during 2016 and 2017 (77.5 fb^{-1}). The limits on $c^{-}_{\varphi Q}$ and $c_{\varphi t}$ obtained by this analysis are better than the limits obtained in TOP-22-006 by about 40%. The EFT approach of this analysis differs from TOP-22-006; in this analysis, the EFT effects are simulated at generator level and applied to SM samples at the detector level. The event selection overlaps with TOP-22-006, so a combination may not be useful. Instead, we may consider techniques presented in this analysis that may be benificial to implement in subsequent iterations of the TOP-22-006 analysis. For example, this analysis makes use of an angular differential variable (corresponding to the cosine of the angle between the direction of the Z in the detector frame and the direction of the negatively charged lepton in the Z frame), which may help to provide additional sensitivity to these WCs.

Ref. [10] ($c^{3}_{\varphi Q}$, c_{tW}): This analysis presents a search for new physics impacting top quarks produced in association with a Z. The analysis uses data collected durring 2016-2018 (138 fb^{-1}). The limits for $c^{3}_{\varphi Q}$ and c_{tW} are similar to the limits obtained by TOP-22-006 (within $\sim$15%). Like TOP-22-006, this analysis parametrizes the predicted yields in terms of the WCs in order to study the EFT effects directly at detector level. The event selection overlaps with TOP-22-006, so a combination may not be useful. However, since the two analyses study essentially the same data, use similar techniques, and derive similar results, the analyses represent a useful cross-check of one another.

Ref. [13] (c_{bW}, $c_{\varphi tb}$, $c_{t\varphi}$): This analysis presents a search for new physics impacting events in which top quarks are produced in association with a boosted Z or Higgs. The analysis uses data collected from 2016-2018 (138 fb^{-1}).

For $c_{t\varphi}$ and $c_{\varphi tb}$, the limits obtained by TOP-22-006 are tighter than the limits obtained by this analysis by a factor of $\sim$3; for c_{bW}, the TOP-22-006 limit is tighter by a factor of $\sim$6. This analysis targets events with a single lepton, so the selection is orthogonal to the TOP-22-006 selection. Furthermore, this analysis makes use of the same EFT approach as TOP-22-006 (parametrizing the event weights in terms of the WCs in order to target EFT effects directly at detector level). A combination between this analysis and TOP-22-006 would thus be useful and potentially fairly straightforward to implement; in fact, such a combination is being pursued currently. For a direct comparison of the TOP-22-006 results to the results of Ref. [13], please see Appendix J.2.

12.5.2 Comparisons against global theory combination

In this section, we will compare the TOP-22-006 results against the limits presented in the 2021 SMEFiT global theory combination [15]. Combining more than 300 measurements from approximately 50 papers, the SMEFiT analysis integrates a large number of processes and final states. Although global combinations can provide powerful constraints on a wide variety of EFT effects, there are many challenging aspects of such analyses. For example, it is important to ensure that a given dataset is used to extract only one measurement, or to ensure that cases of double counting are handled with a proper statistical treatment. Furthermore, it is important to keep in mind how the WCs impact processes that are considered to be backgrounds of the cross section measurements included in the combination; many WCs impact multiple processes, so challenges may be presented by cross section measurements that assume all background processes (including processes affected by the WCs of interest) are equal to the SM prediction.

Figure 12.9 shows the 2σ limits from TOP-22-006 compared against the 2σ limits obtained by the SMEFiT combination. In the case where the other WCs are fixed to zero (Figure 12.9 (a)), we see that the TOP-22-006 limits are comparable (approximately within a factor of two) to or tighter than the SMEFiT analysis for the majority of the WCs that are common to both analyses. However, for several of the WCs from the two-heavy-with-bosons category (c_{tZ}, c_{tW}, $c_{t\varphi}$, c_{tG}, and $c^3_{\varphi Q}$),

the SMEFiT limits are significantly tighter than the TOP-22-006 limits (by factors ranging from approximately three to approximately eight).

The largest difference between the TOP-22-006 and SMEFiT individual limits is associated with the c_{tZ} WC. For this WC, the SMEFiT limit is tighter than the TOP-22-006 limit by a factor of eight. The cause this large difference in sensitivity is not clear. However, it is interesting to note that the sensitivity of the SMEFiT combination to this WC improved significantly from the 2019 version of the analysis [83] to the 2021 iteration of the analysis [15]; the limit obtained by the 2021 combination is more than a factor of 50 tighter than the limit obtained by the 2019 combination. For other similar WCs (e.g. $c_{\varphi t}$ and $c^{-}_{\varphi Q}$, which also primarily impact the $t\bar{t}Z$ process), the improvement from the 2019 to 2021 SMEFiT results is not as large (a factor of $\sim$3). It would be interesting to explore the factors that contribute to the 2021 SMEFiT combination's sensitivity to c_{tZ}.

Moving from the individual likelihood fits to the profiled likelihood fits, Figure 12.9 (b) shows that the largest difference between the profiled limits is a factor of $\sim$3 (for the c_{tW} WC). It is interesting that the TOP-22-006 and SMEFiT limits are more similar for the profiled fits than for the individual fits. Since the profiled fits are in principle less comparable than the individual fits, it is possible that multiple competing effects work together to wash out the differences.

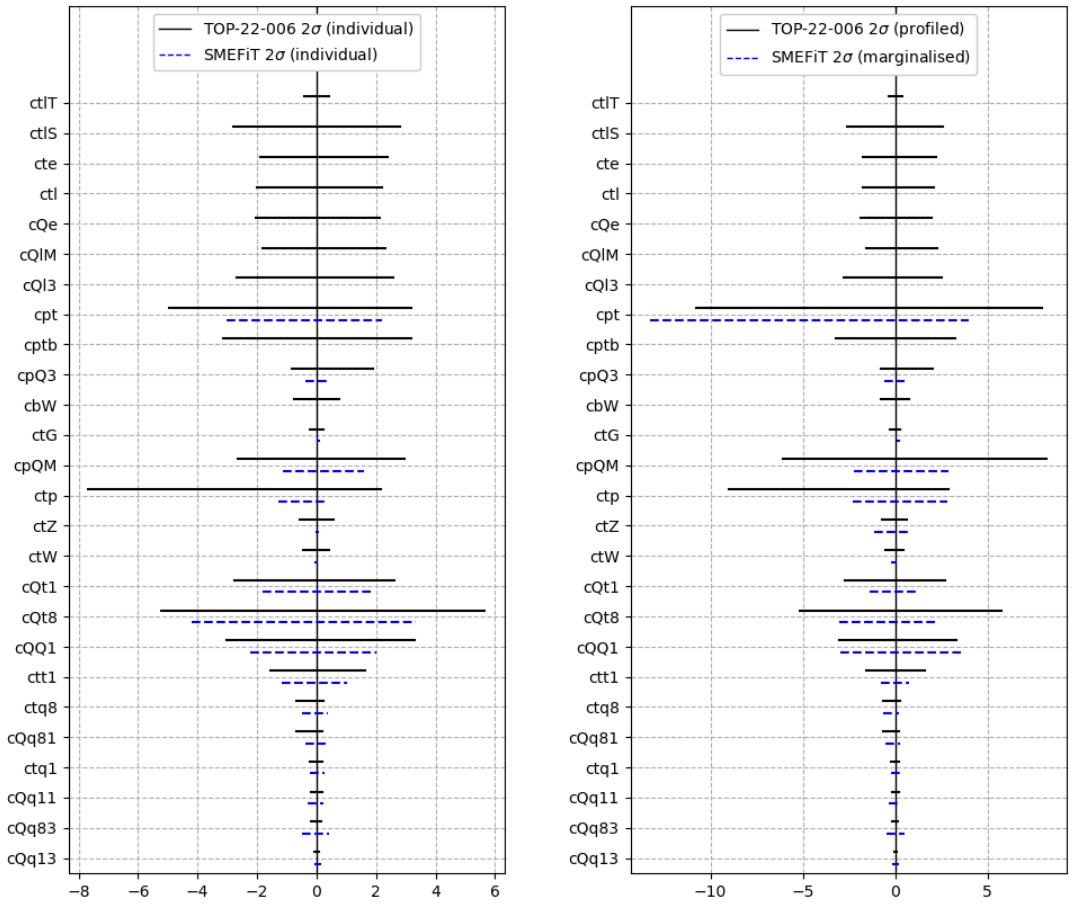

(a) (b)

Figure 12.9. The 2σ profiled limits of this analysis (indicated with the solid black line labeled TOP-22-001) compared against the limits obtained in the 2021 SMEFiT global theory combination [15] (indicated with the dashed blue lines). Figure (a) shows the results of the likelihood fits in which a single WC is fit with all other WCs fixed to their SM values of zero. Figure (b) shows the results of the likelihood fits in which the other WCs are profiled. In these plots, $\Lambda = 1$ TeV.

12.6 Interpretation of results in terms of energy scale

While EFT limits are often discussed in terms of the value of the WC c, we should recall that the variable on which we set a limit is actually c/Λ^2, where Λ is the scale of the new physics, as defined in Eq. 1.1. The scale Λ is often conventionally set to 1 TeV (e.g. in the dim6top model used in this analysis discussed in Chapter 3); this choice of Λ makes c/Λ^2=c (at least in terms of the numerical value, the units of c and c/Λ^2 of course still differ), explaining why the Λ^2 is sometimes dropped.

In fixing Λ to some assumed value, we are able to look at the implications for c; however, we can also approach the situation from the opposite perspective, assuming a value for c and exploring the implications for the scale Λ. We can express Λ in terms of c and the observed limit on c/Λ^2 as follows:

$$\text{Observed limit} = \frac{c}{\Lambda^2} \quad \Longrightarrow \quad \Lambda = \sqrt{\frac{c}{\text{Observed limit}}} \ . \tag{12.2}$$

Taking the 2σ limits from Tables 11.1 and 11.2, we can thus solve for Λ under various assumptions for c. Figure 12.10 shows the resulting values of Λ for three different assumptions for the value of c. The darkest colored bars correspond to a "small" value of c, taken to be 0.01. The medium bars correspond to $c = 1$. Finally, the lightest colored bars correspond to $c = (4\pi)^2$.

The bars in Figure 12.10 essentially show the energy scale to which this analysis has probed. Since the observed limits are inversely proportional to Λ, stronger limits correspond to longer bars on the plot (i.e. a larger excluded region). For example, for a c of 1, Figure 12.10 shows that energies up to about 1 TeV have been probed by this analysis, implying that if new physics (impacting top quarks produced in association with additional charged leptons) were to exist below a range of ∼1 TeV, this analysis should have been sensitive to its effects. In other words, the unexplored region beyond the ∼1 TeV frontier may yet hold new physics to be discovered.

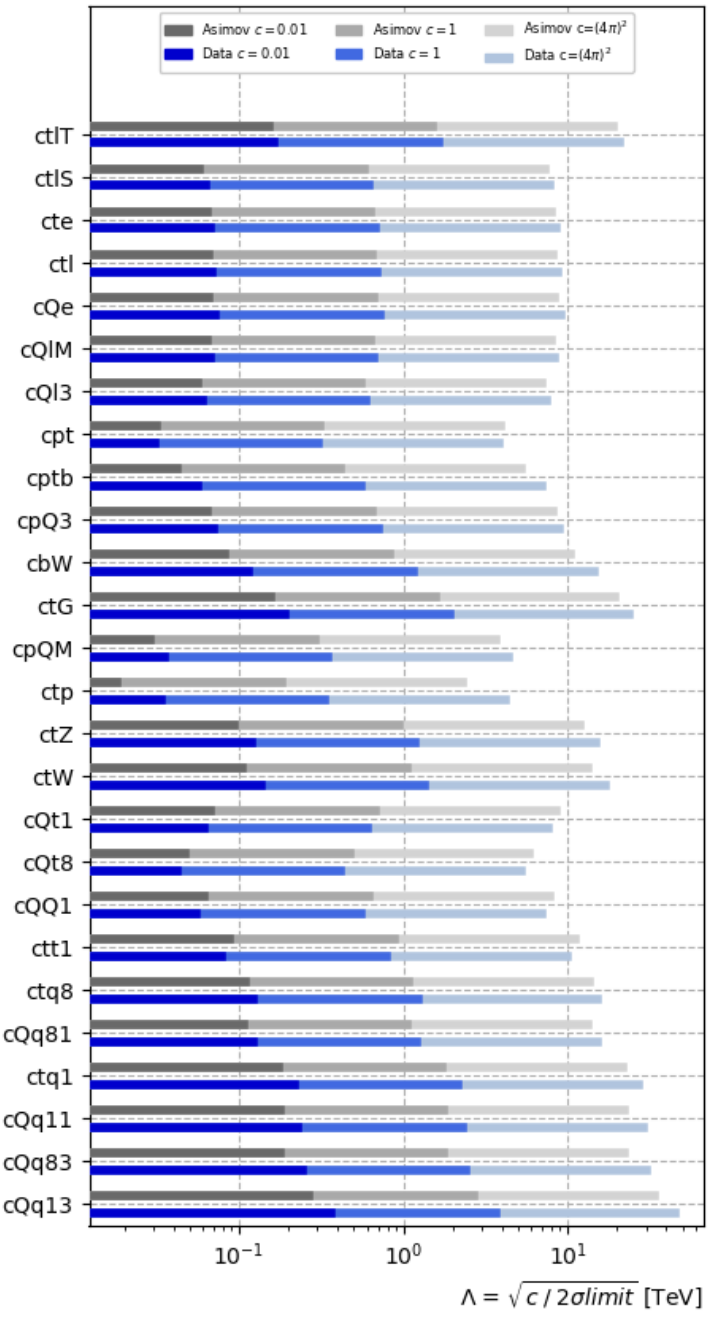

Figure 12.10. The 2σ profiled limits obtained by this analysis interpreted in terms of the energy scale Λ for fits to Asimov data (in grey) and the real data (in blue) for three different assumptions for the value of the WC as indicated in the legend and described in the text. The Asimov and observed limits are taken from the 2σ limits presented in Tables 11.1 and 11.2; for asymmetric $+2\sigma$ and -2σ limits, we have taken the absolute value of the looser limit.