Chapter 10

Results of the Search for Low Mass Compact Binary Coalescences in the First Year of Data from LIGO's Fifth Science Run

In this chapter we describe the results from the Search for Low Mass CBCs in the First Year of LIGO's Fifth Science Run (S5) Data. In section 10.1 we discuss the follow-up procedures we use to build confidence in a possible GW event. We discuss the loudest events from the search in section 10.2. Finally, we present upper limits on the rate of coalescences in section 10.3.

During the analysis, and prior to unblinding the nonplayground data, we discovered an error in the computation of the template metric. This metric is used in the placement of the bank and the coincidence test. The error caused the metric distance between templates to be overestimated for the higher mass signals. This has the effect of causing the template placement algorithm to over-cover the higher mass region (i.e., to produce a bank with less than the requested 3% loss in SNR). This increased the computational cost of the search, but did not significantly reduce the sensitivity. However, this error also affected the coincidence algorithm by overestimating the distance between triggers for high mass signals. Since the coincidence window was empirically tuned on software injections and time-shifted coincidences the impact on the sensitivity of the search was not significant. Consequently, the decision was taken to unblind the data using the original, suboptimal analysis in order to begin studying any possible detection candidates and to use this

result to compute the upper limit (in the absence of a detection). The decision was also taken to perform a complete reanalysis of the data with the corrected metric to verify the (non)detection statement from the original search by checking for any change in the list of the loudest in-time coincident triggers (the detection candidates).

10.1 Follow-up Procedure for Coincident Triggers

The distribution of IFARc in-time coincident triggers is examined to identify the events with the largest IFARc (the "loudest" events). Events with a FARc much smaller than 1 per observation time are deemed significant. However, this takes into account only the GW channel data filtered by the search pipeline. Information from other channels are only used coarsely in the definition of the category 2 and 3 DQ flags (section 6.8.1). Much more information is available in both the GW channel and the auxiliary channels, but it is not yet automatically folded into the detection statistic (this is a work in progress, which will be discussed more in section 11.1). We therefore have developed a rather detailed follow-up procedure to examine this additional information in detail, albeit, in a necessarily subjective manner. Automating this procedure for time-shifts and simulated GW signals (i.e., injections) to provide an improved, unbiased detection statistic is in progress.

In the meantime, we check our loudest coincident triggers with a checklist of tests designed to see if a statistically significant trigger is believable as a detection candidate. The methods employed in this checklist are tested against injection and time-shift coincidence triggers. Reference [139] describes the tests that we perform on the coincident triggers and the data surrounding them. At present, our standard tests include the following: We check the integrity of the data for corruption. We also check the status of the detectors and the presence of any data quality flags in the surrounding data. We assess whether there could have been environmental or instrumental causes found in auxiliary channels at the time of the trigger. We check the appearance of the data at the time of the trigger in the form of the strain, SNR, and χ^2 time-series, and time-frequency spectrograms (an example is shown in figure 10.1).

In addition, for any statistically significant candidate that survives the tests listed above, we plan

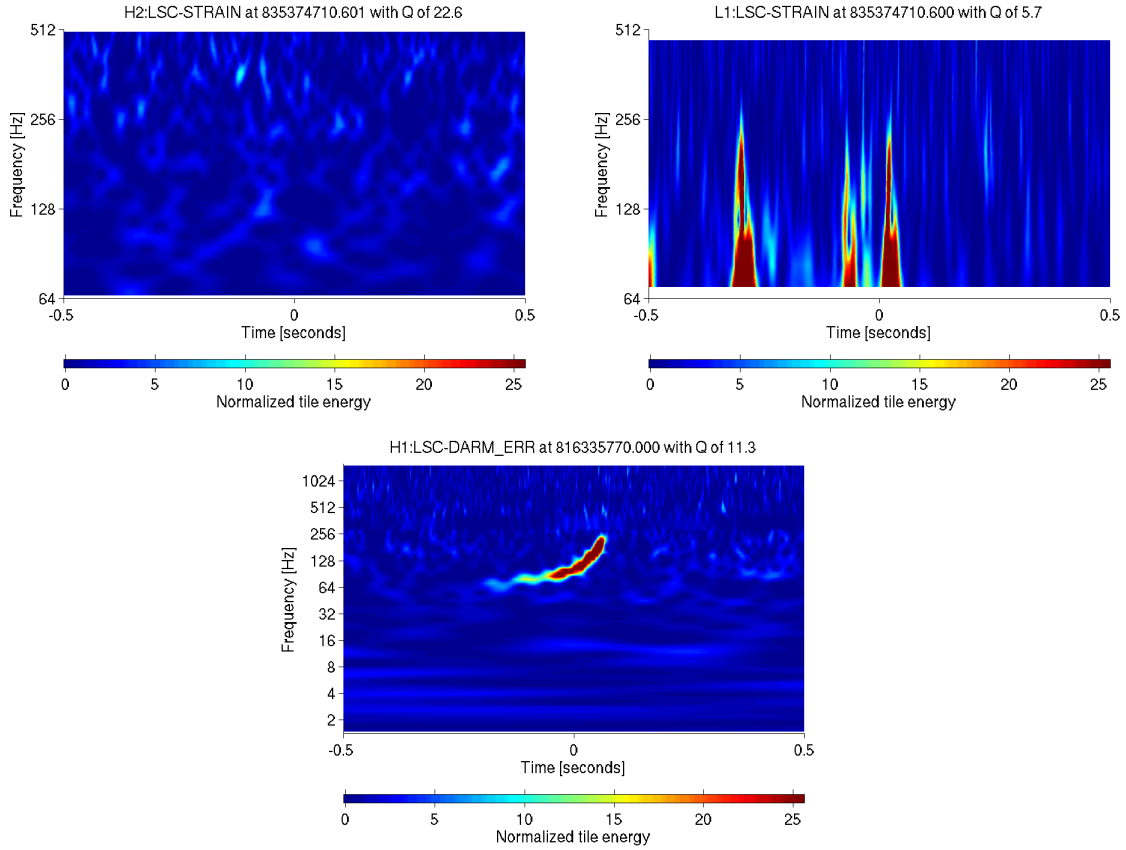

Figure 10.1: Comparison of Triggers from Loudest Event to an Injected Signal
A "Qscan" of the triggers in the loudest event (top) in this search after category 2 and 3 DQ flag vetoes have been applied and of an signal injected into the data (bottom). The Qscans show time-frequency content in the form of power in a Sine-Gaussian basis.

to do the following: Assess the coherence between the signals recorded by each individual detector operating at the time of the event. Verify the robustness of the trigger against small changes in the pipeline (i.e., changes in the adjacent Fourier transform boundaries or changes in the calibration of the data). Check the robustness across pipelines by employing other search techniques to analyze the same data (i.e., CBC pipelines using different templates or pipelines designed to search for unmodeled bursts). Finally, we will check for coincidence with external searches for gamma-ray bursts, optical transients, or neutrino events. (This last test is for information only, as a genuine GW event might or might not be accompanied by other signals.)

We examine the distribution of in-time triggers after each category of DQ flag veto is applied. In case there is a statistically significant outlier before category 2 or categories 2 and 3 DQ flag vetoes

are applied, we carry out a follow-up exercise to see if the veto that eventually rejected the event was rightfully applied. There are two reasons that this could be important. First, a very strong GW from within the Milky Way could cause an instrumental saturation of the sort that we use as a veto; this kind of problem would be easy to diagnose if it were to occur, since the signal would be strong enough for us to see in the moments leading up to the signal-induced saturation. Second, we want to guard against false dismissal of a candidate by other kinds of vetoes, which can have nonnegligible dead-time associated with them. Some of our vetoes are associated with recognizable forms of false signals; we check to be sure that a vetoed loud event looks like that kind of false signal, and not like a genuine coalescence signal. In this search, there was a single statistically significant outlier in the distribution of events before the application of categories 2 and 3 DQ flag vetoes. The follow-up exercise confirmed that the category 3 DQ flag that vetoed that event was correctly applied, as it occurred during a time that was especially glitchy and was recommended to be removed from analysis as the data was taken.

10.2 Loudest Events

At the end of our pipeline we are left with a set of coincident triggers that are potential detection candidates. The cumulative distribution of events above an IFARc versus IFARc for the different observation times is shown in figure 10.2; the H1H2L1 distribution is also shown in figure 7.10. This figure shows that the loudest candidates in all three observation times were consistent with the estimated background and thus were likely accidental coincidences. Thus, the search yielded no detection candidates. The results of the reanalysis were consistent with the original analysis and did not produce any plausible GW signals. We report an upper limit on the rate of coalescences in section 10.3.

As an exercise to prepare for future detections, we carried the loudest several events (such as the three loudest events that appear in each of the histograms in figure 10.2) through the follow-up procedure.

Even though we know our background is underestimated for H1H2 coincident triggers, we re-

viewed the two loudest H1H2 candidates using the detection checklist. In both of those cases, the waveforms from the two interferometers failed to match each other in detail, thus ruling them out as GW events.

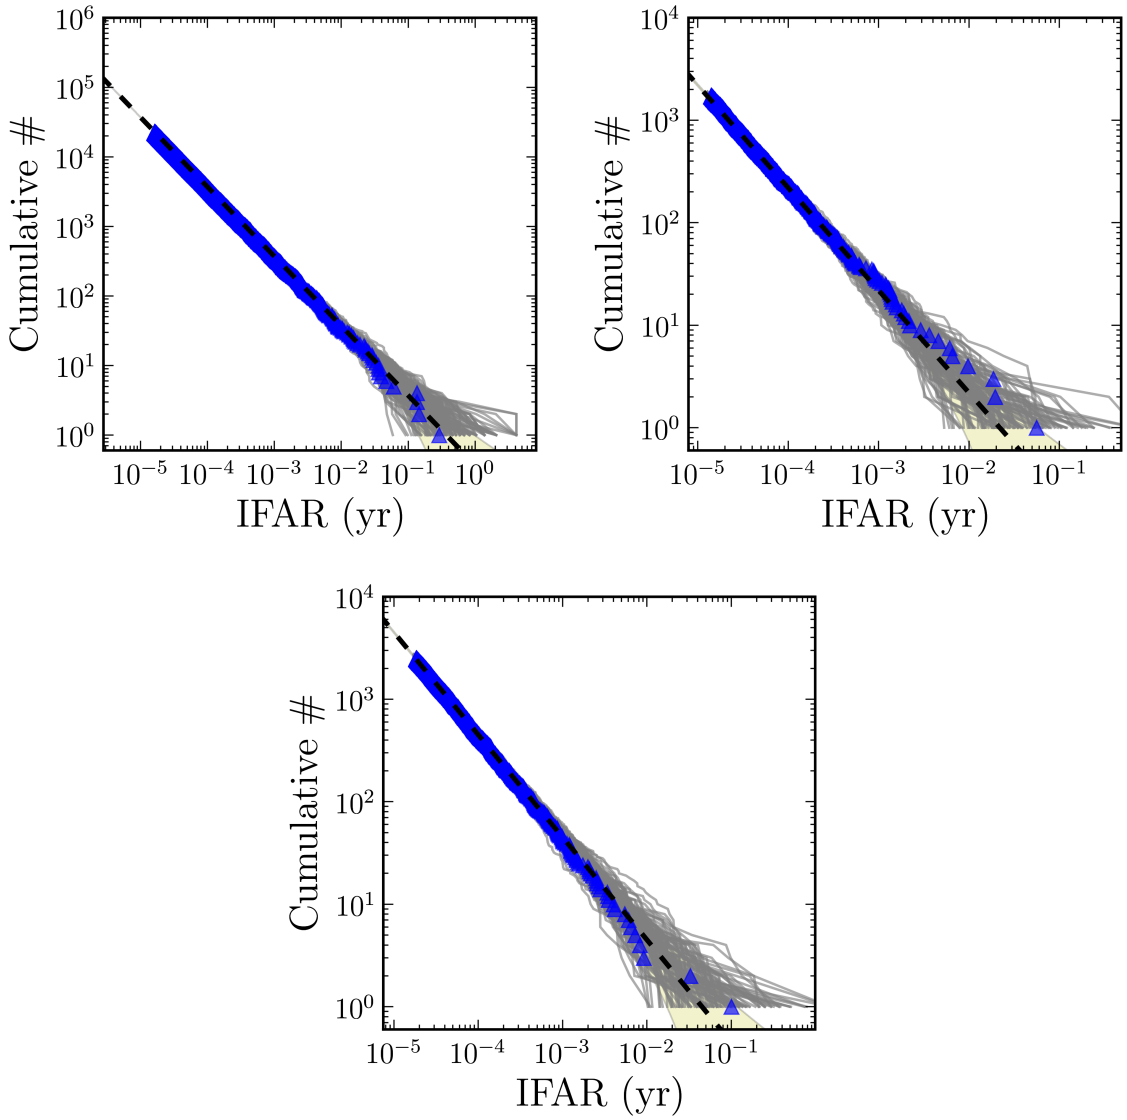

Figure 10.2: IFAR Loudest Events
The cumulative distribution of events above a threshold IFAR, for in-time coincident events, shown as blue triangles, from all coincidence categories for the observation times H1H2L1 (top left), H1L1 (top right), and H2L1 (bottom right). The expected background (by definition) is shown as a dashed black line. The 100 experimental trials that make up our background are also plotted individually as the solid grey lines. The yellow band shows the $N^{1/2}$ one-standard deviation errors on the average background estimate.

Table 10.1: Detailed Results of the BNS Upper Limit Calculation

BNS	$D_{\rm horizon} = 30\,{\rm Mpc}$		
Coincidence Time	H1H2L1	H1L1	H2L1
Cumulative Luminosity (L_{10})	250	230	120
Calibration Error	21%	3.9%	16%
Monte Carlo Error	5.4%	16%	13%
Waveform Error	26%	11%	20%
Galaxy Distance Error	14%	13%	6.1%
Galaxy Magnitude Error	17%	17%	16%
Λ [Eq. 9.2]	0.30	0.41	0.72
Marginalized Upper Limit $\left(\mathrm{yr}^{-1}L_{10}^{-1}\right)$	3.9×10^{-2}		

Summary of the search for BNS systems. The horizon distance, as defined in Eq. 5.38, for the 4 km detectors and averaged over the search. The cumulative luminosity is rounded to two significant figures. The errors in this table are listed as logarithmic errors in the luminosity multiplier based on the cited sources of error.

Table 10.2: Detailed Results of the BBH Upper Limit Calculation

BBH	$D_{\rm horizon} = 80\,{\rm Mpc}$		
Coincidence Time	H1H2L1	H1L1	H2L1
Cumulative Luminosity (L_{10})	4900	4900	1800
Calibration Error	23%	36%	25%
Monte Carlo Error	3.2%	11%	9.9%
Waveform Error	27%	34%	27%
Galaxy Distance Error	18%	18%	20%
Galaxy Magnitude Error	16%	16%	17%
Λ [See (9.2)]	0.59	1.2	1.4
Marginalized Upper Limit $\left(\mathrm{yr}^{-1}L_{10}^{-1}\right)$	2.5×10^{-3}		

Summary of the search for BBH systems. The rows of this table are defined as in Table 10.1.

10.3 Coalescence Rate Upper Limits

In the absence of detection, we set upper limits on the rate of CBCs per unit L_{10}, for several canonical binary masses, as a function of the total mass of the compact binary system, and as a function of the black hole mass for BHNS systems.

For each mass range of interest, we calculate the rate upper limit at 90% confidence level (CL) using the loudest event formalism [133, 130], described in chapter 9. We derive a Bayesian posterior distribution for the rate as described in chapter 9 and reference [130]. The cumulative luminosity for this search, calculated as described in chapter 9, can be found in Table 10.1, 10.2, and 10.3 for the BNS, BBH, and BHNS canonical mass CBC systems, respectively.

Table 10.3: Detailed Results of the BHNS Upper Limit Calculation

BHNS	$D_{\mathrm{horizon}} = 50\,\mathrm{Mpc}$		
Coincidence Time	H1H2L1	H1L1	H2L1
Cumulative Luminosity (L_{10})	990	980	390
Calibration Error	22%	47%	25%
Monte Carlo Error	3.6%	9.7%	11%
Waveform Error	26%	40%	27%
Galaxy Distance Error	17%	16%	17%
Galaxy Magnitude Error	17%	18%	18%
Λ [See (9.2)]	0.45	1.0	1.1
Marginalized Upper Limit $\left(\mathrm{yr}^{-1}L_{10}^{-1}\right)$	1.1×10^{-2}		

Summary of the search for BHNS systems. The rows of this table are defined as in Table 10.1.

10.3.1 Upper Limits Neglecting Spin

We apply the upper limit calculation to three canonical binary masses as well as calculating the upper limit as two functions of mass. Our three canonical binary masses are BNS ($m_1 = m_2 = (1.35 \pm 0.04)$ $M_\odot$), BBH ($m_1 = m_2 = (5 \pm 1)$ $M_\odot$), and BHNS ($m_1 = (5 \pm 1)$ $M_\odot$, $m_2 = (1.35 \pm 0.04)$ $M_\odot$). We use Gaussian distributions in component mass centered on these masses, with standard deviations following the $\pm$ symbols. The neutron star mass distribution was chosen according to an analysis of 50 known pulsars in binary systems at the time (reference [140]). However, the mass distribution for black holes is still uncertain, therefore we have chosen a canonical value to represent stellar mass black holes, which allows us to compare results across different science runs.

We combine the results of this search from the three different observation times in a Bayesian manner, described in section 9.1, and the results from previous science runs [21, 22] are incorporated in a similar way.

Assuming that spin is not important in these systems, we calculate upper limits on the rate of binary coalescences using our injection families that neglect spin (section 6.2). There are a number of uncertainties that affect the upper limit, including systematic errors associated with detector calibration, simulation waveforms, Monte Carlo statistics, and galaxy catalog distances and magnitudes [131]. We marginalize over these as described in section 9.4 and obtain upper limits on the rate of binary coalescences of

$$\mathcal{R}_{90\%,\mathrm{BNS}} = 3.9 \times 10^{-2}\,\mathrm{yr}^{-1}\mathrm{L_{10}}^{-1}\,, \tag{10.1a}$$

$$\mathcal{R}_{90\%,\mathrm{BBH}} = 2.5 \times 10^{-3}\,\mathrm{yr}^{-1}\mathrm{L_{10}}^{-1}\,, \tag{10.1b}$$

$$\mathcal{R}_{90\%,\mathrm{BHNS}} = 1.1 \times 10^{-2}\,\mathrm{yr}^{-1}\mathrm{L_{10}}^{-1}\,. \tag{10.1c}$$

We also calculate upper limits for two additional cases: as a function of the total mass of the binary, with a uniform distribution in the mass ratio $q = m1/m2$, and as a function of the mass of the black hole in a BHNS system, holding fixed the mass of the neutron star at $m_{\mathrm{NS}} = 1.35$ $M_\odot$(figure 10.3).

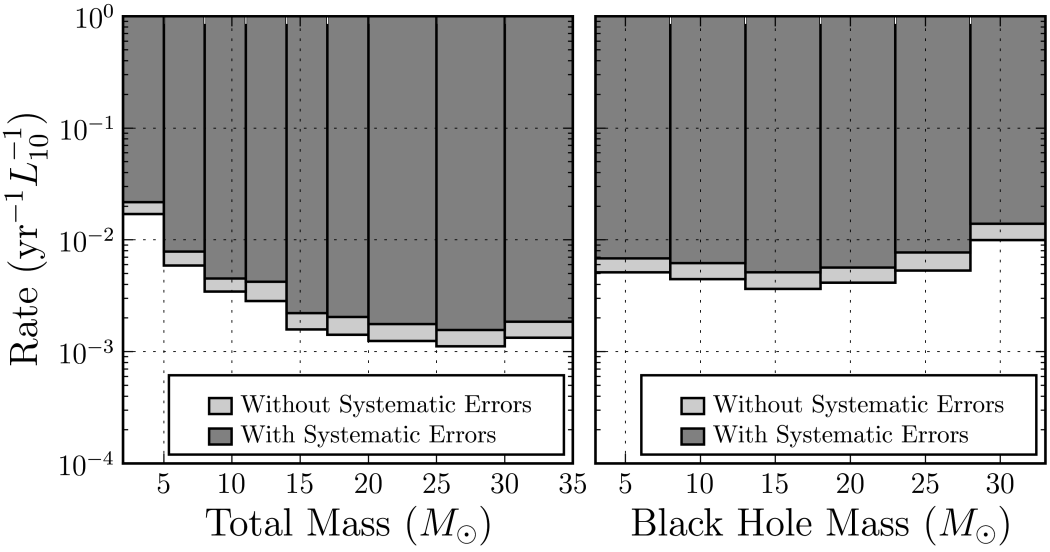

Figure 10.3: Upper Limits versus Mass
Upper limits on the binary coalescence rate per year and per L_{10} as a function of total mass of the binary system with a uniform distribution in the mass ratio (left) and as a function of the mass of a black hole in a BHNS system with a neutron star mass of $m_{\mathrm{NS}} = 1.35\ M_\odot$(right). The darker area shows the excluded region after marginalization over the estimated systematic errors. The lighter area shows the additional region that would have been excluded if the systematic errors of our most sensitive observation time had been ignored.

10.3.2 Upper Limits Including Spin

In the previous section, we reported upper limits on the rate of mergers for different classes of objects using injection waveforms generated assuming nonspinning objects. We can also evaluate the upper limits using injection waveforms that take into account the effects of spinning bodies.

Since the maximum possible rotational angular momentum S for a black hole of mass m is Gm^2/c, it is useful to describe the spin of a compact object in terms of the dimensionless spin parameter $\hat{a} = (cS) / \left(Gm^2\right)$. The distribution of black hole spin magnitudes within the range $0 \leq \hat{a} \leq 1$, as well as their orientations relative to binary orbits, is not well constrained by observations. To illustrate the possible effects of BH spins on our sensitivity to BBH and BHNS signals, we provide an example calculation using a set of injections using the SpinTaylor approximant, which calculates waveforms for systems whose component objects can include significant spin angular momentum. The spin magnitude $\hat{a}$ for this set of injections is uniformly distributed between 0 and 1 (section 6.2). However, assuming a canonical mass and uniform density for a component object, astrophysical observations of neutron stars show typical angular momenta corresponding to $\hat{a} \ll 1$ [141]. In addition, the spin effects are found to be weak for the frequency range of interest for LIGO [16] so the BNS upper limits in section 10.3.1 are valid even though we have ignored the effects of spin.

Using the above injections, we obtain marginalized upper limits on the rate of binary coalescences of

$$\mathcal{R}_{90\%,\mathrm{BBH}} = 3.2 \times 10^{-3}\,\mathrm{yr}^{-1}{\mathrm{L}_{10}}^{-1}\,, \tag{10.2}$$

$$\mathcal{R}_{90\%,\mathrm{BHNS}} = 1.4 \times 10^{-2}\,\mathrm{yr}^{-1}{\mathrm{L}_{10}}^{-1}\,. \tag{10.3}$$

We can see that the effect of spin on the relatively low mass systems targeted in this search is rather small, modifying the average distance we can see by a factor of order ∼5%–10%, and thus our upper limit by a factor that is three times larger (since the volume we are sensitive to is proportional to the cube of the distance) of order ∼15%–30%.

10.4 Summary of the Search

We can compare the above results to a prediction made for the upper limit on the rate of BNS coalescences from reference [131]. The example upper limit in that reference is computed assuming a horizon distance of 35 Mpc for the detector, an observation time of 1 yr, and a loudest event from the analysis with a combined SNR of 10. The result they predict, after marginalization over systemic errors is $7.3 \times 10^{-3}\,\mathrm{yr}^{-1}\,{\mathrm{L}_{10}}^{-1}$. If we adjust the parameters above to match the sensitivity (average horizon distance of 30 Mpc) and duration (0.37 yr of analyzed data) of this search, the predicted result becomes $3.1 \times 10^{-2}\,\mathrm{yr}^{-1}\,{\mathrm{L}_{10}}^{-1}$, which is very close to the upper limit we obtain in equation (10.1a).

Aside from possibly detecting GWs, this is what Initial LIGO was proposed to achieve and what it was built for. S5 was the first science run since the LIGO detectors reached their design sensitivity and also the first long science run LIGO has undertaken. The search for BNS GW signals, the flagship search, concluded with a result as sensitive as expected, a great achievement for LIGO and the LIGO Scientific Collaboration (LSC).

We can also extend this prediction to BBH systems by adjusting the predicted result by the factor $(M_{\mathrm{BNS}}/M_{\mathrm{BBH}})^{5/2}$, which comes from the mass dependence of the amplitude of a CBC GW in equation (2.43). This results in a predicted result for BBH systems of $1.2 \times 10^{-3}\,\mathrm{yr}^{-1}\,{\mathrm{L}_{10}}^{-1}$. This is further from the obtained result in equation (10.1b) than the BNS comparison, however this is expected due to the higher FAR for higher mass templates. This shows that the IFARc detection statistic prevents the contamination of the lower FAR for lower mass templates from the higher FAR from the higher mass templates.