OBSERVATIONAL SELECTION EFFECTS WITH GROUND-BASED GRAVITATIONAL WAVE DETECTORS

HSIN-YU CHEN[1], REED ESSICK[2], SALVATORE VITALE[2], DANIEL E. HOLZ[1], AND ERIK KATSAVOUNIDIS[2]
[1] University of Chicago, Chicago, Illinois 60637, USA
[2] LIGO, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

ABSTRACT

Ground-based interferometers are not perfect all-sky instruments, and it is important to account for their behavior when considering the distribution of detected events. In particular, the LIGO detectors are most sensitive to sources above North America and the Indian Ocean, and as the Earth rotates, the sensitive regions are swept across the sky. However, because the detectors do not acquire data uniformly over time, there is a net bias on detectable sources' right ascensions. Both LIGO detectors preferentially collect data during their local night; it is more than twice as likely to be local midnight than noon when both detectors are operating. We discuss these selection effects and how they impact LIGO's observations and electromagnetic (EM) follow-up. Beyond galactic foregrounds associated with seasonal variations, we find that equatorial observatories can access over 80% of the localization probability, while mid-latitudes will access closer to 70%. Facilities located near the two LIGO sites can observe sources closer to their zenith than their analogs in the south, but the average observation will still be no closer than 44° from zenith. We also find that observatories in Africa or the South Atlantic will wait systematically longer before they can begin observing compared to the rest of the world; though, there is a preference for longitudes near the LIGOs. These effects, along with knowledge of the LIGO antenna pattern, can inform EM follow-up activities and optimization, including the possibility of directing observations even before gravitational-wave events occur.

Key words: gravitational waves – methods: observational

1. INTRODUCTION

The detection of binary black holes with the Laser Interferometer Gravitational-wave Observatory (LIGO; Collaboration et al. 2015; Abbott et al. 2016a, 2016e) has ushered in the age of gravitational wave astronomy. A particularly promising avenue for exploring new physics is multi-messenger astronomy, involving the joint detection of GW sources, electromagnetic (EM) signals, or astrophysical particles (Kulkarni et al. 1998; Gehrels et al. 2005). During the last years of initial LIGO and Virgo, as well as in advanced LIGO's first observing run (O1; The LIGO Scientific Collaboration et al. 2016a), a large consortium of EM observers followed up GW candidates (Abadie et al. 2012a; Aasi et al. 2014; Abbott et al. 2016f), and there has been substantial effort to plan and optimize EM follow-up. Previous work compared the advantages and disadvantages of different telescopes (Kasliwal & Nissanke 2014). Regardless of a facility's hardware, observatories at different locations will have systematically different opportunities to follow-up GW events due to properties of the GW detector network.

Even if astrophysical sources are distributed isotropically on the sky, GW detections with the two LIGO detectors will not be. Because of the detectors' locations (Hanford, WA and Livingston, LA) and duty cycles, detectable GW sources preferentially come from certain locations on the celestial sphere, and the preferred regions vary with the seasons. Follow-up of EM counterparts with emission timescales of less than a few weeks will be especially affected by these biases. In particular, the Earth's rotation limits ground-based EM follow-up facilities. Different sites can access different parts of the GW localization maps at different times, and therefore, have different expectations for the fraction of counterparts they can detect. These GW selection effects have other important ramifications for follow-up efforts, such as the amount of time before an average source will be accessible and the average air mass expected.

GW detectors do not operate continuously (Aasi et al. 2015; Abbott et al. 2016b). Previous studies considered their duty cycle when estimating detection rates (Singer et al. 2014; Abbott et al. 2016c; The LIGO Scientific Collaboration et al. 2016b). However, their operation is not uniformly distributed in time, and instead shows a strong preference for acquiring data during their local night. We describe this behavior quantitatively and assess its influence on GW detections as well as EM follow-up.

We first explain the sources of biases and their impact on GW detections in Section 2. We show the effects of the bias on ground-based EM follow-up facilities in Section 3, and conclude in Section 4.

2. OBSERVATIONAL BIAS INTRODUCED BY GROUND-BASED GRAVITATIONAL WAVE DETECTORS

Ground-based GW interferometers do not have isotropic sensitivity. Over time, the antenna pattern produces systematic preferences for the locations of detectable sources. These can be split into a dependence on the source's declination, which will not improve until additional detectors are added to the network, and a dependence on the source's right ascension, which can be mitigated with existing facilities.

2.1. Dependence on Declination

Because of projection effects, GW detectors are most sensitive to signals coming from above or below the plane defined by their arms (Thorne 1987). The two LIGO detectors, for example, are most sensitive, and equally sensitive, to sources directly above North America and above the Indian Ocean (Figure 1(a)). As the Earth rotates, the antenna pattern is swept across the celestial sphere, creating preferred bands in
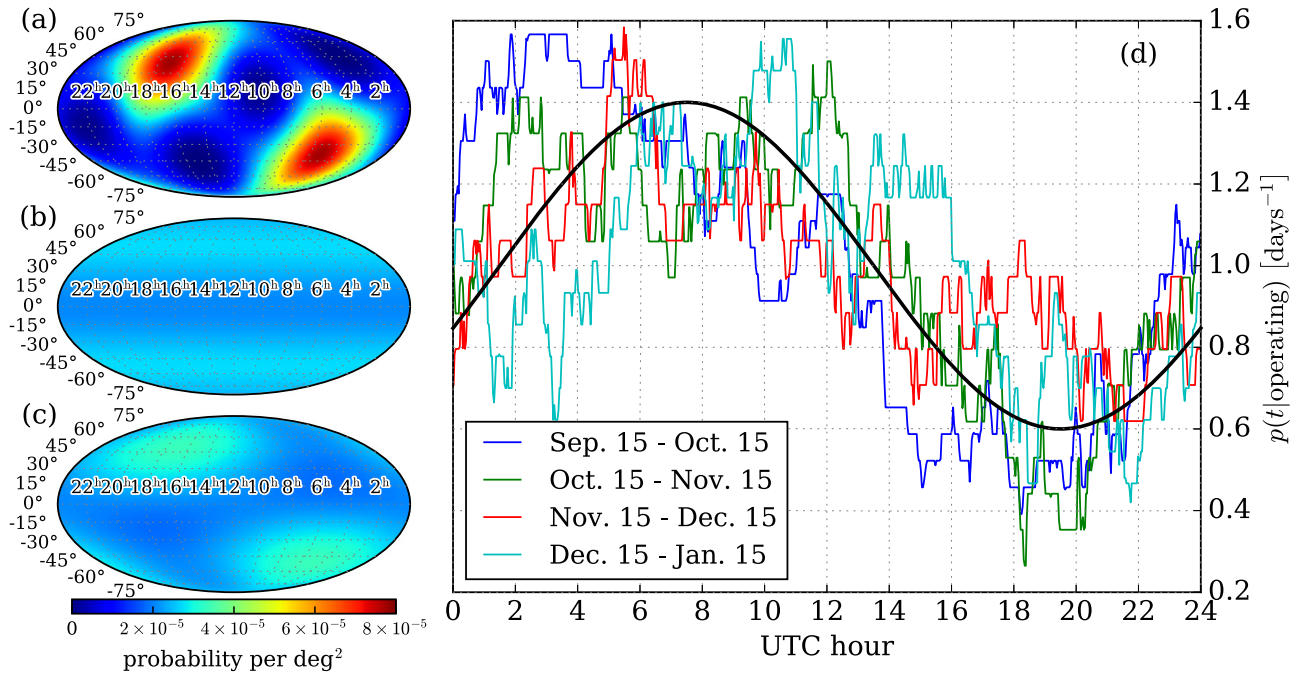
Figure 1. (a) The LIGO Hanford and Livingston network antenna pattern in equatorial coordinates at 00:29:18 UTC on 2015 September 14. The maxima lie above North America and the Southern Indian Ocean. (b) The antenna pattern swept over the celestial sphere assuming uniform operation throughout time, resulting in two maxima bands in mid-declinations. (c) The antenna pattern swept over the celestial sphere assuming a typical diurnal cycle. This produces a dependence on right ascension that persists over timescales of a few days to weeks. (d) The observed operation of the LIGO detectors during O1 and a sinusoidal model as a function of UTC time. We require both LIGO detectors to be operating at the same time; though, each detector shows similar behavior individually.

the mid-latitudes of both hemispheres (Figure 1(b)). This is determined primarily by the detectors' relative sensitivity and the geometry of the network. For the two LIGO detectors, we expect more detections at mid-declinations compared to low- or high-declinations.

2.2. Dependence on Right Ascension

The antenna pattern also introduces a dependence on longitude, which translates to a dependence on the right ascension. If the detectors operated uniformly in time, then this right ascension dependence would average away as the Earth rotates. However, because the duty cycle is not uniform in time, this engenders a net bias in favor of certain right ascensions.

Ground-based GW detectors' data acquisition shows a clear diurnal cycle. In previous runs (S6—Abadie et al. 2012b; Aasi et al. 2013—and earlier) this was due primarily to anthropogenic noise; the detectors observed lower ambient noise at night because humans were less active. During O1, we observed a similar diurnal cycle. Although anthropogenic noise was mitigated by improved seismic isolation, commissioning activities still preferentially occurred during the day at the detector sites. Furthermore, the detectors achieved an overall coincident duty cycle of 50% (Abbott et al. 2016). This generated a non-trivial preference for acquiring data during the site's night (Figure 1d). We model this preference with a probability distribution for the time of day during which both interferometers are likely to be operating:

$$p(t|\text{operating}) = \frac{1}{\text{day}}\left(1 + A\sin\left(\frac{2\pi t}{\text{day}} - \frac{\pi}{8}\right)\right). \qquad (1)$$

The amplitude (A) reflects the extent of the day/night bias and is typically near 0.4. This means that the detectors are 2.3 times

more likely to record data at their local midnight than at noon. We note that this bias may be decreased by reducing the amount of commissioning activity during the day or by increasing the overall duty cycle of the instruments (e.g., reducing the duration of downtime due to events uncorrelated with the diurnal cycle). As the duty cycle increases, the relative importance of the diurnal behavior will decrease (see the Appendix for more details). However, we always expect some small diurnal cycle to be present.

The diurnal cycle preserves the dependence on right ascension as the Earth rotates (Figure 1(c)), which persists over timescales of days to weeks. Figure 2 shows the dependence averaged over a month for several months throughout the year. However, over the course of a year, the Earth's orbit will average away this dependence. Nonetheless, as currently scheduled, GW detectors do not operate year-round, instead only recording data over a few consecutive months during observing runs. If observing runs are scheduled during the same season repeatedly, the right ascension dependence introduced by the diurnal cycle can persist for years.

Furthermore, the galactic plane intersects the preferred directions from May through September (Figure 2). Depending on the desired target, this could be an advantage or a hindrance. Detected compact binaries are expected to be extragalactic (Abbott et al. 2016a, 2016e) and the galactic plane will serve as a foreground for any EM follow-up, significantly complicating the removal of transient contaminants. However, detectable core-collapse supernova are expected to be primarily galactic (Gossan et al. 2016), and therefore GW detectors may have their best chance of observing such events during the north's summer. Regardless of the source, a full understanding of the distribution of detectable signals across both declination and
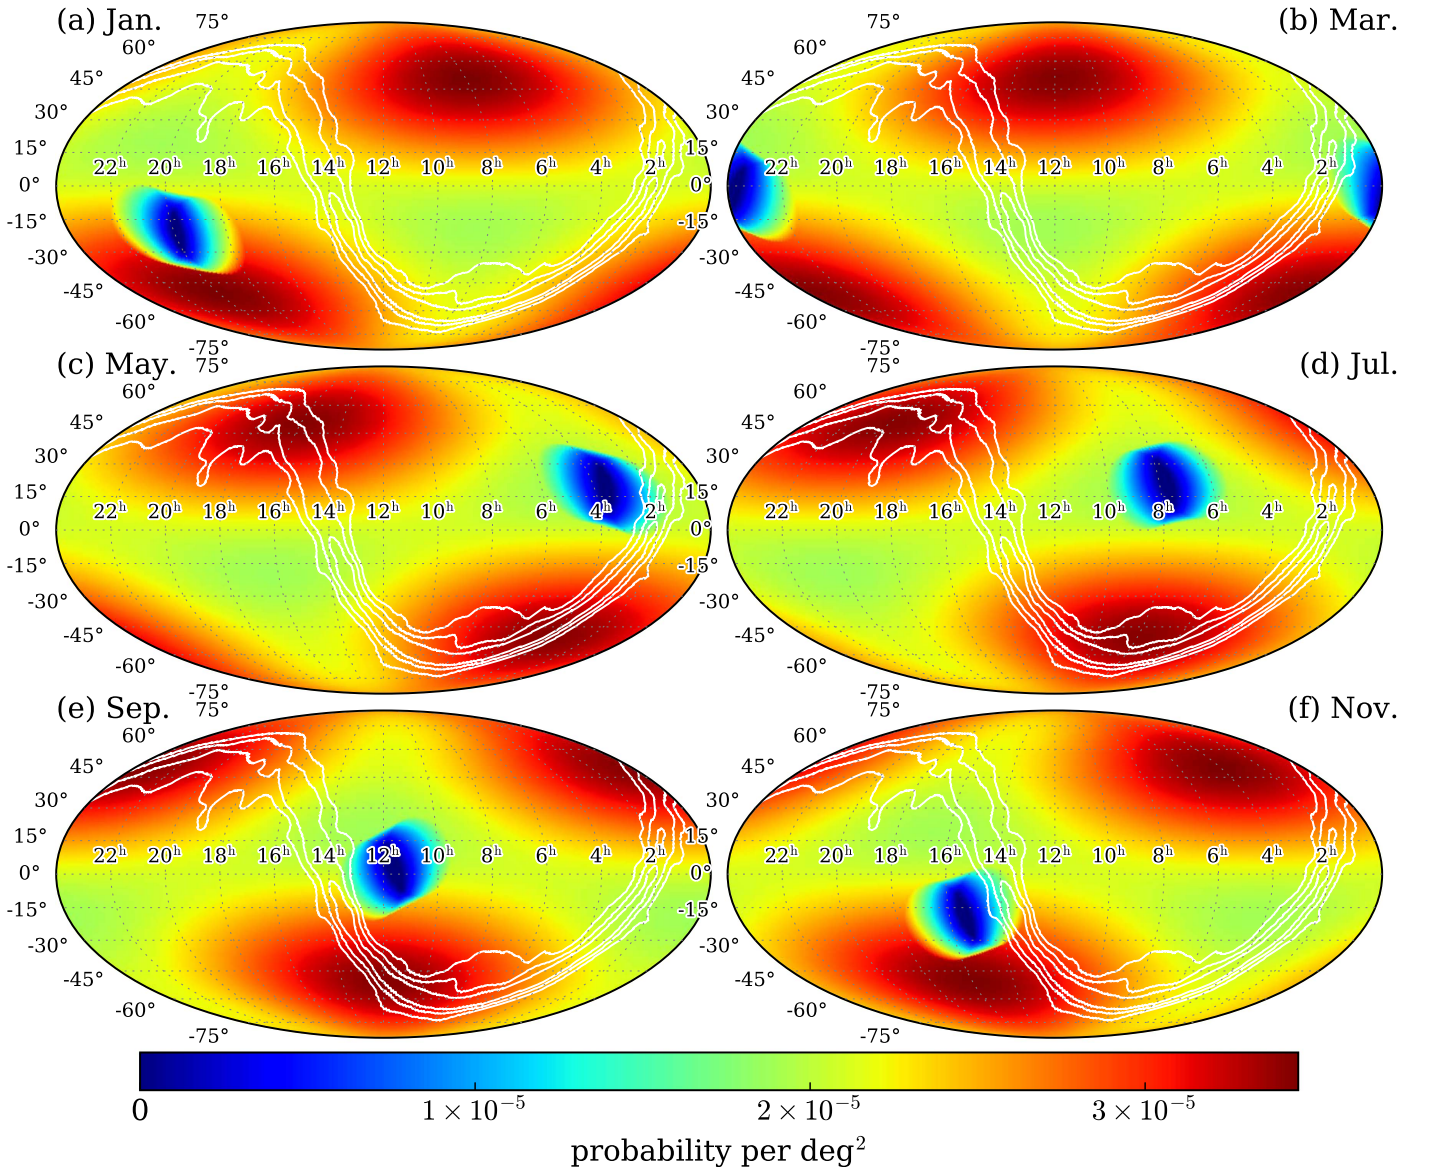
Figure 2. Monthly preferred regions for GW detections on the celestial sphere. The Sun occludes a region 18° in radius (assuming observatories can see down to their horizons), and this region is smeared out by averaging over a month. The white contours are the Planck occlusion masks (NASA/IPAC Infrared Science Archive 2016) and represent the Galactic plane.

right ascension will be crucial when considering any isotropy or homogeneity measurements using GW observations alone.

3. IMPACT ON GROUND-BASED EM FOLLOW-UP

Selection effects associated with ground-based GW detectors impact EM follow-up facilities. In particular, we focus on three possible effects: the localization probability that a telescope can survey (observable probability), the source's closest angle of approach to an observatory's zenith while the observatory can observe (mimimum zenith distance), and the time until a GW source becomes observable (delay time). We focus on EM follow-up timescales of up to a few days or weeks, and therefore neglect seasonal modulations of the sky over the duration of the EM observations for each individual GW event. This is an appropriate timescale for short gamma-ray burst afterglows (Berger 2014) and kilonovae (Barnes &

Kasen 2013), two promising EM counterparts of compact binary coalescences involving at least one neutron star. Radio transients may persist over longer timescales and therefore these effects may be less relevant. For these EM follow-up timescales, we consider low-latency GW alerts. In O1, it took a couple of days to issue alerts, but we expect this to be reduced to a few minutes in O2 and beyond (Abbott et al. 2016f). Throughout this paper, we assume the diurnal cycle modeled in Equation (1) as well as 18° of astronomical twilight (Patat et al. 2006). We also assume that the observatories can observe within 90° of their zenith.

In the limit of a large number of detections, the combined posterior distributions trace out the network antenna pattern. In what follows, we use the antenna pattern to approximate limiting distribution of many events (see the Appendix for more details). To quantify departures from this limit when only
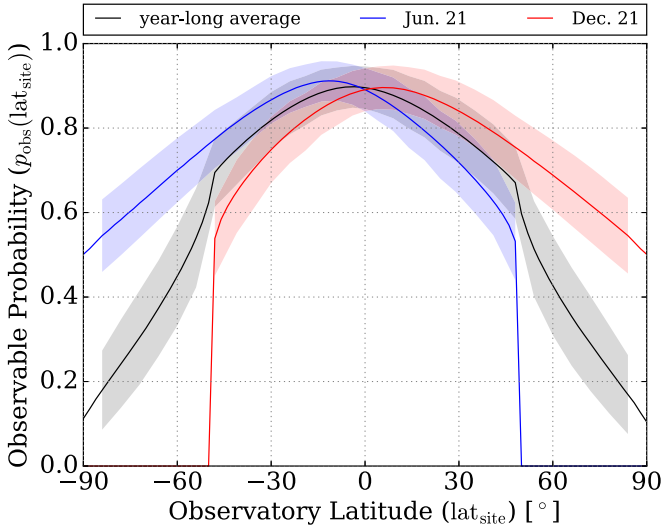
Figure 3. $p_{obs}(lat_{site})$ averaged over a year (black), near the northern summer solstice (blue), and near the northern winter solstice (red). Shaded regions correspond to the fluctuations from localization maps simulations ($\lim_{N_d\to\infty}(N_d/10)^{1/2}\sigma_{\hat{p}_{obs}}$).

Figure 4. D_{zen} averaged over a year (black), near the northern summer solstice (blue), and near the northern winter soltice (red). Shaded regions correspond to the fluctuations from localization maps simulations ($\lim_{N_d\to\infty}(N_d/10)^{1/2}\sigma_{\hat{D}_{zen}}$), and we note that $\hat{D}_{zen}$'s distribution is quite broad.

a few events are available, we simulate collections of events based on binary black hole localizations with two detectors (Essick et al. 2015; Essick 2015). We find good agreement.

3.1. Observable Probability

GW localizations are driven by triangulation and, for networks of two detectors, the localization is characterized by large rings that can span hundreds of square degrees regardless of source morphology. Furthermore, these error regions typically have support at antipodal points on the sky, making it difficult for a single EM observatory to access the entire skymap. This is compounded by solar occlusion, which renders certain parts of the sky inaccessible. We ignore lunar occlusion because, unlike solar occlusion, it is not thought to systematically correlate with when the detectors operate. Because we focus on timescales of a few days, we assume that the Earth will revolve at least once. Therefore, the observable region depends only on the observatory's latitude.

Figure 3 shows the observable probability as a function of latitude for a year-long average and for the solstices. We define

$$p_{obs}(lat_{site}) = \int dt \left(p(t|operating) \right.$$
$$\left. \times \int d\Omega\, p_{GW}(\Omega, t)\Theta_{obs}(\Omega, t, lat_{site}) \right), \quad (2)$$

where $p_{GW}(\Omega, t)$ is the probability density of the GW source coming from Ω (at time t), and Θ_{obs} is the region accessible from a particular latitude at t, respectively. We see sharp declines near latitudes of $\pm 50°$ corresponding to the Arctic/Antarctic circles and astronomical twilight, but otherwise p_{obs} follows a smooth distribution favoring equatorial observatories. This is because equatorial observatories can systematically access sources in both hemispheres, whereas other observatories may be confined to only the probability within their own. While equatorial facilities are favored overall, mid-latitudes
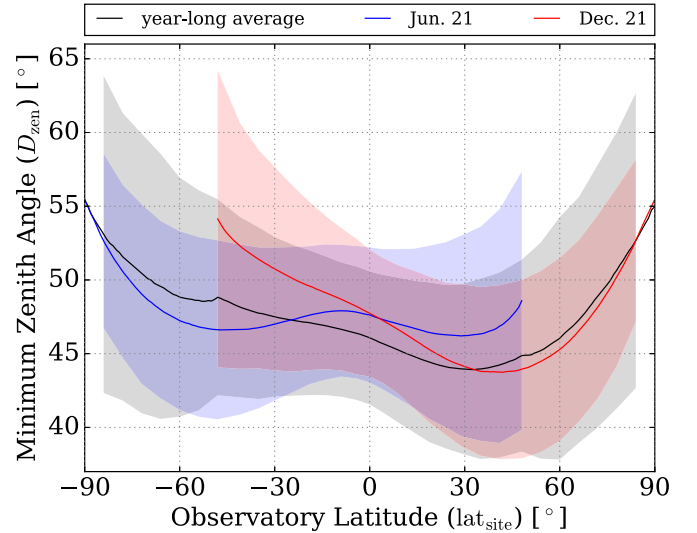
have larger p_{obs} than they would for isotropically distributed detections. This will persist regardless of the day/night cycle and is driven solely by the GW detectors' latitudes. We note that the optimal observing months of the northern and southern hemispheres are out of phase (Figure 3). For example, if LIGO operates from September to February, the northern hemisphere will have a better chance of observing counterparts than the southern hemisphere.

While p_{obs} does not depend on the longitude, the diurnal cycle can still introduce a systematic bias. This is because the Sun will be systematically out of phase with the northern maximum of the antenna pattern (Figure 2). We find that this typically produces an increase in p_{obs} of a few percent for mid-latitudes in the north, even though their analogs in the south have larger p_{obs} because of the shorter solar exposure over an entire year.

We note that p_{obs} reflects the amount of localization probability that is observable in the limit of many detections. Outside of this limit, we consider an analogous quantity ($\hat{p}_{obs}$) defined for a finite number of detections (N_d) with the correspondence $p_{obs} = \lim_{N_d\to\infty}\hat{p}_{obs}$ (see the Appendix for more details). Statistical fluctuations in $\hat{p}_{obs}$, caused by variations in which events occur, can be large, particularly with $\lesssim 10$ events. We expect the uncertainty in this estimate to scale inversely with $\sqrt{N_d}$, and Figure 3 reports $\lim_{N_d\to\infty}(N_d/10)^{1/2}\sigma_{\hat{p}_{obs}}$ as error bars along with the mean. Typical values are between 6% and 10%. We also note that the intrinsic distribution for single events ($N_d = 1$) may not be Gaussian but, in the limit of many detections, the distribution of the mean will be.

3.2. Minimum Zenith Distance

While equatorial observatories may be able to access the largest integrated probability, and thereby have the largest probability of being able to image a source, this does not necessarily imply that they will have the best conditions for observing. An important consideration is the closest approach

of each field to an observatory's zenith (minimum zenith distance). Sources at large angles from an observatory's zenith can be difficult to observe because of high air mass and mechanical limitations. In Figure 4, we present the mean minimum zenith distance, as a function of latitude, weighted by the probability that the source actually comes from each location:

$$D_{\rm zen}({\rm lat}_{\rm site}) = \frac{1}{p_{\rm obs}} \int dt \, (p(t|{\rm operating})$$
$$\times \int d\Omega \, p_{\rm GW}(\Omega, t) \Theta_{\rm obs}(\Omega, t, {\rm lat}_{\rm site})$$
$$\times D_{\rm zen}(\Omega, t, {\rm lat}_{\rm site})) \qquad (3)$$

$\Theta_{\rm obs}$ accounts for solar occlusion and $D_{\rm zen}(\Omega, t, {\rm lat}_{\rm site})$ incorporates when the EM facility will actually be able to observe. We find that observatories at extreme latitudes ($\pm 90°$) will almost always have large $D_{\rm zen}$, with a gradual transition to lower values at lower latitudes. There is a $\sim 10\%$ difference in $D_{\rm zen}$ between observatories at mid-latitudes, with northern sites preferred. This is because the diurnal cycle makes the Sun preferentially overlap the southern antenna pattern and forces southern facilities to observe closer to sunrise and sunset than their northern counterparts. This behavior is particularly evident at the northern winter solstice.

Figure 4 also shows the fluctuations in the analgous stastic defined for a finite number of detections ($\hat{D}_{\rm zen}$). We typically find $\lim_{N_d \to \infty} (N_d/10)^{1/2} \sigma_{\hat{D}_{\rm zen}} \sim 4$–$10°$ (see the Appendix for more details). Furthermore, $D_{\rm zen}$ corresponds to the mean of many events. For a single event, the mode of $\hat{D}_{\rm zen}(N_d = 1)$'s distribution falls near $20°$ at mid-latitudes in the north and $60°$ in the south, whereas the mode is near $50°$ at both poles.

3.3. Delay Time

For counterparts with timescales of days, EM observatories' longitudes can play an important role. This is because observatories will have to wait to begin observing until the source rises at their location. This can be exacerbated by the position of the Sun, which will systematically correlate with the diurnal cycle.

We expect that the amount of time an observatory must wait before commencing observations (delay time) will depend on both the observatory's longitude and latitude and define

$$D_{\rm del}(\Omega_{\rm site}) = \frac{1}{p_{\rm obs}} \int dt \, (p(t|{\rm operating})$$
$$\times \int d\Omega \, p_{\rm GW}(\Omega, t) \Theta_{\rm obs}(\Omega, t, {\rm lat}_{\rm site})$$
$$\times D_{\rm del}(\Omega, t, \Omega_{\rm site})) \qquad (4)$$

where $D_{\rm del}(\Omega_{\rm site}, t, \Omega)$ accounts for sunrise and sunset along with the source's relative position to the observatory. We restrict ourselves to only the parts of the skymap that are actually accessible from each site. Figure 5 shows the dependence on the observatory's longitude and latitude. Although the shape changes, we see a reasonably uniform distribution of $D_{\rm del}$ throughout the globe (there are slightly longer delay times in the north compared to their analogs in the south away from the blob over Africa) with the notable exception of the Southern Atlantic, Africa, and the Indian Ocean, particularly in the year-long average (Figure 5(e)).
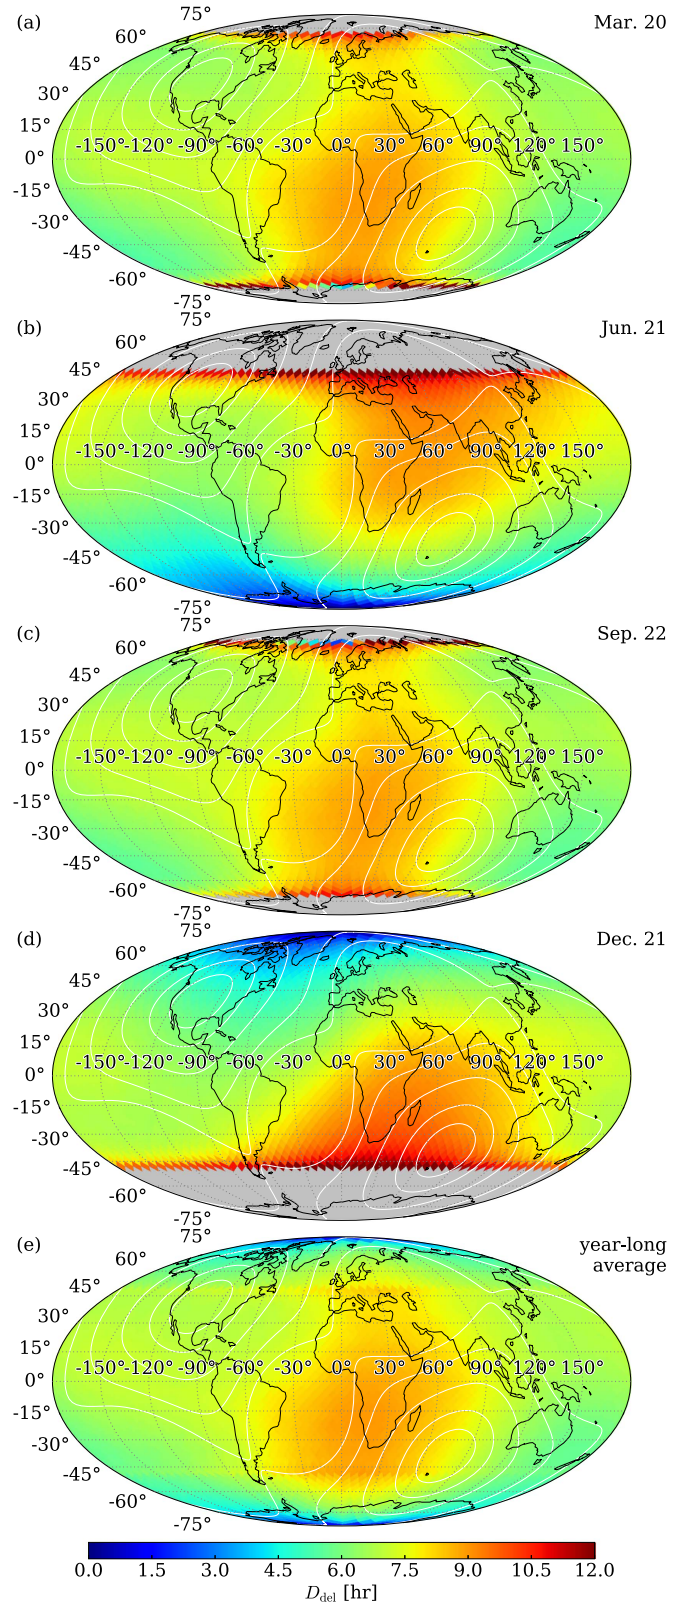

Figure 5. $D_{\rm del}(\Omega_{\rm site})$ for hypothetical observatories placed throughout the globe for times surrounding the northern (a) spring equinox, (b) summer solstice, (c) fall equinox, and (d) winter solstice. (e) A year-long average. Gray regions correspond to observatories that are in perpetual daylight or twilight at the corresponding time.

Because the majority of detections will occur during North America's night, these locations are likely to already be in daylight and will therefore have to wait for sunset.

Furthermore, observatories in the south will have to wait until their zenith is very close to the source's right ascension before they observe, which explains why the blob is wider in the south. From simulations outside the limit of many detections, we also note that the distribution of the analogous statistic $(\hat{D}_{\text{del}}(N_d = 1))$ is very non-Gaussian and very skewed to the right. However, the mean is still Gaussian as $N_d \to \infty$ with an associated variance $\lim_{N_d \to \infty}(N_d/10)^{1/2}\sigma_{\hat{D}_{\text{del}}} \sim 80\text{--}240$ minutes, with most values near 100 minutes, depending on the observatory's location (see the Appendix for more details).

We note that if the diurnal cycle were to be strongly reduced or eliminated, the blob above Africa would become less prominent and an analogous blob would appear in the Northern Pacific. The latter corresponds to observatories that are in daylight when a source is detected above the southern antenna pattern. The overall effect of removing the diurnal cycle, however, is toward a more uniform distribution of delay times across the entire globe.

4. DISCUSSION

In this paper, we enumerate three effects that imprint selection effects on the distribution of GW detections, which will be relevant for all types of GW sources: (1) the detector's locations introduce a preference for mid-declinations, (2) a diurnal cycle modulates when the detectors operate and produces a sky sensitivity with a dependence on the right ascension over short timescales, and (3) if detections are made only during relatively short observing runs, the right ascension bias can be imprinted over longer timescales. While the effects are important when modeling the expected distribution of GW detections on the sky, they can also have significant implications for EM follow-up.

In general, ground-based EM observatories located at latitudes comparable to those of the GW detectors are preferred. While equatorial sites can access the largest fraction of the sky, mid-latitude sites in the north will be able to observe more events closer to their zenith without sacrificing much coverage. Furthermore, there is a preference for EM observatories located near the same longitude as the GW detectors, or slightly west thereof, assuming the EM facilities can begin observing quickly and that counterparts decay with timescales of hours. This leads us to the conclusion that EM facilities located near the GW detectors are favored by how the GW detectors' actually operate.

In addition to informing follow-up efforts for GW candidates, our results may help guide EM activities in advance of GW observations. For example, because we know which parts of the celestial sphere are most likely to host detectable GW events, we can build up relevant templates for image subtraction before GW observations even begin, focusing particularly on the highest probability regions. Galaxy catalog constructions could be focused similarly. In addition, surveys may focus their observations near the peak of the instantaneous antenna pattern, thereby anticipating where detectable GW sources are most likely to occur and facilitating target of opportunity follow-ups. This will also increase the probability of serendipitous detection of prompt EM counterparts. We note that EM observatories located near GW detectors will naturally survey a maximum of the antenna pattern because their zenith lies near that maximum automatically.

The behavior of GW observatories can become more complicated as the global network of detectors expands. A larger network increases the uniformity of sensitivity to GW signals across the sky, and therefore reduces the selection effects in both declination and right ascension. Furthermore, detectors located around the globe will likely experience diurnal cycles that are out of phase, further reducing any preference for certain right ascensions. However, because all planned detectors lie within a relatively confined band of latitudes (Aso et al. 2013; Fairhurst 2014; Acernese et al. 2015), the bias on declination may persist at some level. In addition, the LIGO detectors will likely have the best sensitivity in the network for at least the next few years (Abbott et al. 2016c). This means that they will generally provide the dominant contribution to detections, and the observed distribution of sources will follow their antenna pattern. Thus, our analysis serves as a reasonable prediction for both the distributions of GW detections as well as their impact on EM follow-up efforts for the next few years.

APPENDIX A
CALCULATING MEANS IN THE LIMIT OF MANY DETECTIONS

Throughout the main paper, we use the antenna pattern to approximate limiting distributions of many events. We also study the distributions by drawing simulated localization maps. We repeatedly draw N_d maps, distributed through time according to $p(t|\text{operating})$, to obtain the distributions of our statistics for N_d detections. In this approach, the observable probability, the minimum zenith distance, and the delay time are calculated as

$$\hat{p}_{\text{obs}}(\text{lat}_{\text{site}}) = \frac{1}{N_d}\sum_{i=1}^{N_d}\int d\Omega \, p_{\text{sky},i}\Theta_{\text{obs}}(\Omega, t_i, \text{lat}_{\text{site}}), \qquad (5)$$

$$\hat{D}_{\text{zen}}(\text{lat}_{\text{site}}) = \frac{1}{N_d \, \hat{p}_{\text{obs}}}\sum_{i=1}^{N_d}\int d\Omega \, p_{\text{sky},i}\Theta_{\text{obs},i}D_{\text{zen}}(\Omega, t_i, \text{lat}_{\text{site}}),$$

$$\text{(6)}$$

$$\hat{D}_{\text{del}}(\Omega_{\text{site}}) = \frac{1}{N_d \, \hat{p}_{\text{obs}}} \sum_{i=1}^{N_d} \int d\Omega \, p_{\text{sky},i} \, \Theta_{\text{obs},i} D_{\text{del}}(\Omega, t_i, \Omega_{\text{site}}),$$

(7)

where $p_{\text{sky},i}$ is the localization map probability for the i_{th} detection and $\hat{p}_{\text{obs}}$ in Equations (6) and (7) is computed using the same set of skymaps as the explicit sum. We note that using the antenna pattern yields the mean of these statistics when simulating many detections. Namely, if we calculate χ using the antenna pattern and $\hat{\chi}$ using sets of simulated maps, we expect

$$\chi = \lim_{N_d \to \infty} \hat{\chi}$$
$$= \lim_{N_d \to \infty} \frac{1}{N_d} \sum_{i=1}^{N_d} \int d\Omega \, p_{\text{sky},i} \chi(t_i).$$

(8)

We note that $\chi(t_i)$ only depends on the time the event occurs through the position of the Sun; the dependence will be the same for all detections that occur at the same time. If we break the sum into small segments of time, we can write

$$\chi = \lim_{N_d \to \infty} \frac{1}{N_d} \sum_{j=1}^{N_t} \sum_{\substack{i=1 \\ t_i \in [t_j, t_j + \Delta t)}}^{N_{d,j}} \int d\Omega \, p_{\text{sky},i} \chi(t_i),$$

(9)

where $\sum_j N_{d,j} = N_d$. Now, when $N_d \to \infty$, we can make the segments as small as we like while maintaining a large number of detections in each bin. We then obtain

$$\chi = \lim_{\substack{N_t \to \infty \\ N_t \Delta t = T}} \sum_{j=1}^{N_t} \lim_{N_{d,j} \to \infty} \frac{1}{N_d} \sum_{\substack{i=1 \\ t_i \in [t_j, t_j + \Delta t)}}^{N_{d,j}} \int d\Omega \, p_{\text{sky},i} \chi(t_i)$$
$$= \lim_{\substack{N_t \to \infty \\ N_t \Delta t = T}} \sum_{j=1}^{N_t} \lim_{N_{d,j} \to \infty} \left(\frac{N_{d,j}}{N_d} \right) \int d\Omega \left(\frac{1}{N_{d,j}} \sum_{\substack{i=1 \\ t_i \in [t_j, t_j + \Delta t)}}^{N_{d,j}} p_{\text{sky},i} \right) \chi(t_i)$$
$$= \lim_{\substack{N_t \to \infty \\ N_t \Delta t = T}} \sum_{j=1}^{N_t} (p(t|\text{operating})\Delta t) \int d\Omega (p_{\text{GW}}(\Omega, t_j)) \chi(t_j)$$
$$= \int dt \, p(t|\text{operating}) \int d\Omega \, p_{\text{GW}}(\Omega, t) \chi(t)$$

(10)

where T is the length of the observing season. We assumed $t_i \to t_j$ for all events within each bin and used the fact that many localization posteriors stacked on top of one another will average to the network antenna pattern in equatorial coordinates ($p_{\text{GW}}(t_j)$), which depends on the bin's time. We have also used the fact that the fraction of events occuring in each bin is equal to the probability that the detectors are operating throughout that bin ($p(t|\text{operating})\Delta t$). By approximating these integrals, we obtain the limits of the means much more efficiently than through direct simulation. This proceedure, or an equivalent, is used in Figures 3–5.

We also note that we normalize by the total observable probability in Equations (3), (4), (6), and (7). This is because we restrict ourselves to only the fraction of the probability that is actually observable. The particular form of our normalization (dividing by $\hat{p}_{\text{obs}}$ for a set of events rather than each event separately) guarantees that we sample the antenna pattern in the limit.

APPENDIX B
DISTRIBUTIONS OF STATISTICS FOR INDIVIDUAL EVENTS

Section 3 presents the limit of our statistics when many detections are present. In this limit, the distributions of these statistics will be Gaussian, but it is also informative to examine the distributions when there are only a few events. This is particularly important when considering how an observatory could be impacted for any individual event, rather than for a collection of events. We also consider these distributions in the limit of extremely well localized sources: $p_{\text{sky},i} \to \delta(\Omega - \Omega_i)$. This point-source limit describes the distributions obtained when only observing the location of the true source.

These distributions were computed using bootstrapped simulations of BBH detections based off of Essick et al. (2015). The library developed to perform these simulations is publicly available (https://github.com/reedessick/selectionEffects) and readers are encouraged to use it to determine distributions for their favorite observatory.

B.1. $\hat{p}_{\text{obs}}$

Of all the statistics we consider, $\hat{p}_{\text{obs}}$ is the most Gaussian for small N_d for year-long averages. However, because it is bounded from above and below, the distribution does deviate at times. Figure 6 shows the distributions for a few latitudes and a few values of N_d.

We note that, in the limit $N_d \to \infty$, p_{obs} is the fraction of true counterparts that an observatory can observe. When considering extremely well localized sources, the observatory will either be able to observe the true source or not, and each trial will have nearly the same probability of success, modulo variations caused by the Sun's declination. Therefore, the fraction of true sources each observatory can observe will be nearly binomially distributed. Figure 7 plots the point-source limit of these distributions ($\hat{p}_{\text{obs}}^{(\text{pt src})}$), again with several values of N_d.

B.2. $\hat{D}_{\text{zen}}$

$\hat{D}_{\text{zen}}(N_d = 1)$ is also fairly Gaussian for some latitudes, but there can often be non-trivial deviations therefrom. In particular, mid-latitudes may show interesting skew right distributional shapes. Figure 8 shows these distributions for a few latitudes as a function of N_d. We note that these distributions are much narrower than the point-source limit in which all events are well localized ($\hat{D}_{\text{zen}}^{(\text{pt src})}$), shown in Figure 9. There is more "shot noise" for $\hat{D}_{\text{zen}}^{(\text{pt src})}$, which broadens the distributions, but the general distributional shapes are similar for both $\hat{D}_{\text{zen}}$ and $\hat{D}_{\text{zen}}^{(\text{pt src})}$.

B.3. $\hat{D}_{\text{del}}$

We note in Section 3.3 that the distribution of $\hat{D}_{\text{del}}$ is very skewed right when N_d is small. Figure 10 demonstrates this. Typically, there is an extremely large lobe near zero, corresponding to events that are immediately observable, and a long tail comprised of events that require waiting. When we consider the point-source limit of well localized events, this behavior is enhanced. Figure 11 demonstrates that. When we integrate over typical skymaps, instead of point sources, we find that the peak is smeared out to longer delay times. This is
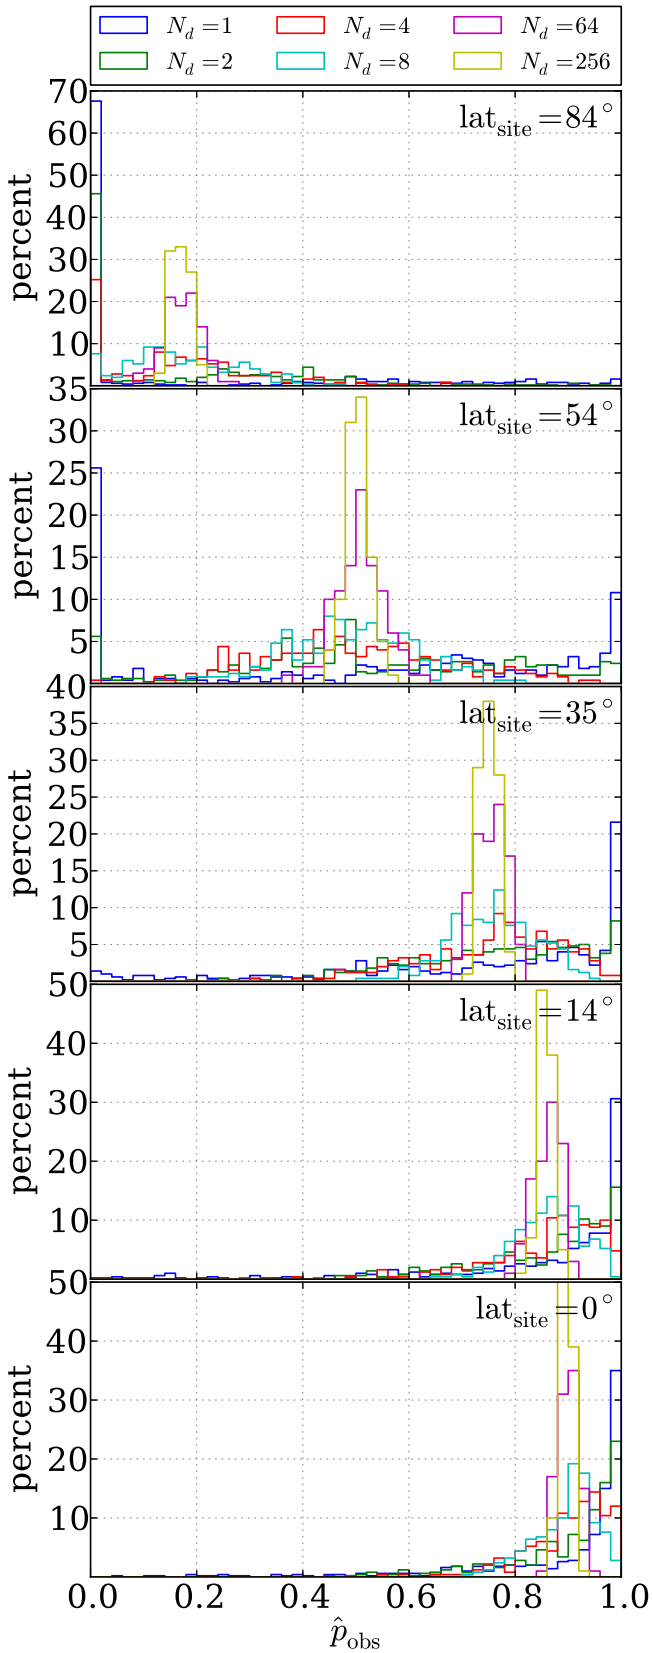
Figure 6. Distributions of $\hat{p}_{\mathrm{obs}}$ for a few latitudes and N_d for year-long averages. We note that the distributions with $N_d = 1$ may not be very Gaussian and display long tails. Nonetheless, as $N_d \to \infty$, the means of the distributions tend toward the values reported in Figure 3.

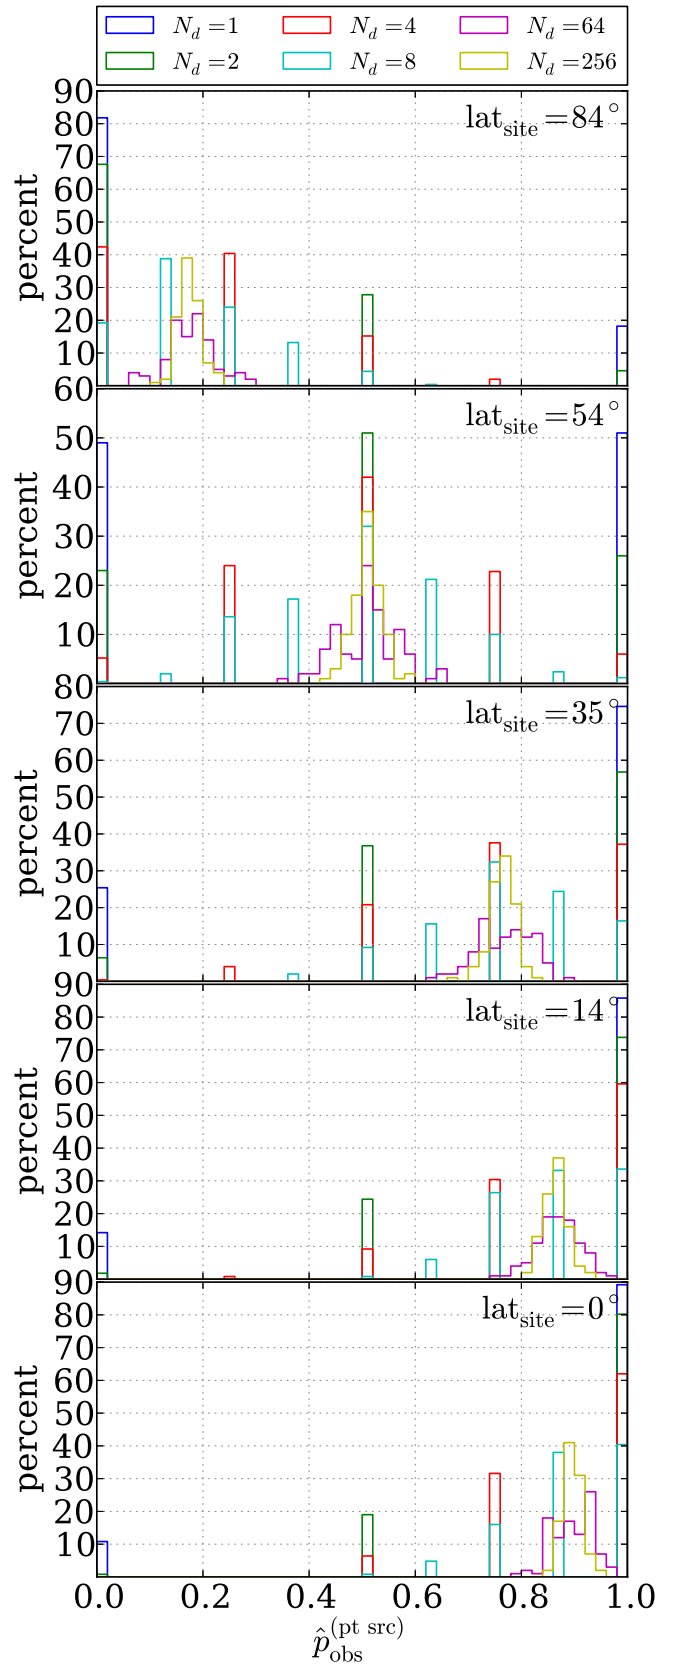**Figure 7.** Distributions of the fraction of observable point sources ($\hat{p}_{\mathrm{obs}}^{(\mathrm{pt\ src})}$) for a few latitudes and N_d. We note that the distributions tend toward means equivalent to those in Figure 6 as $N_d \to \infty$; though, their discreteness for small N_d somewhat obscures this.
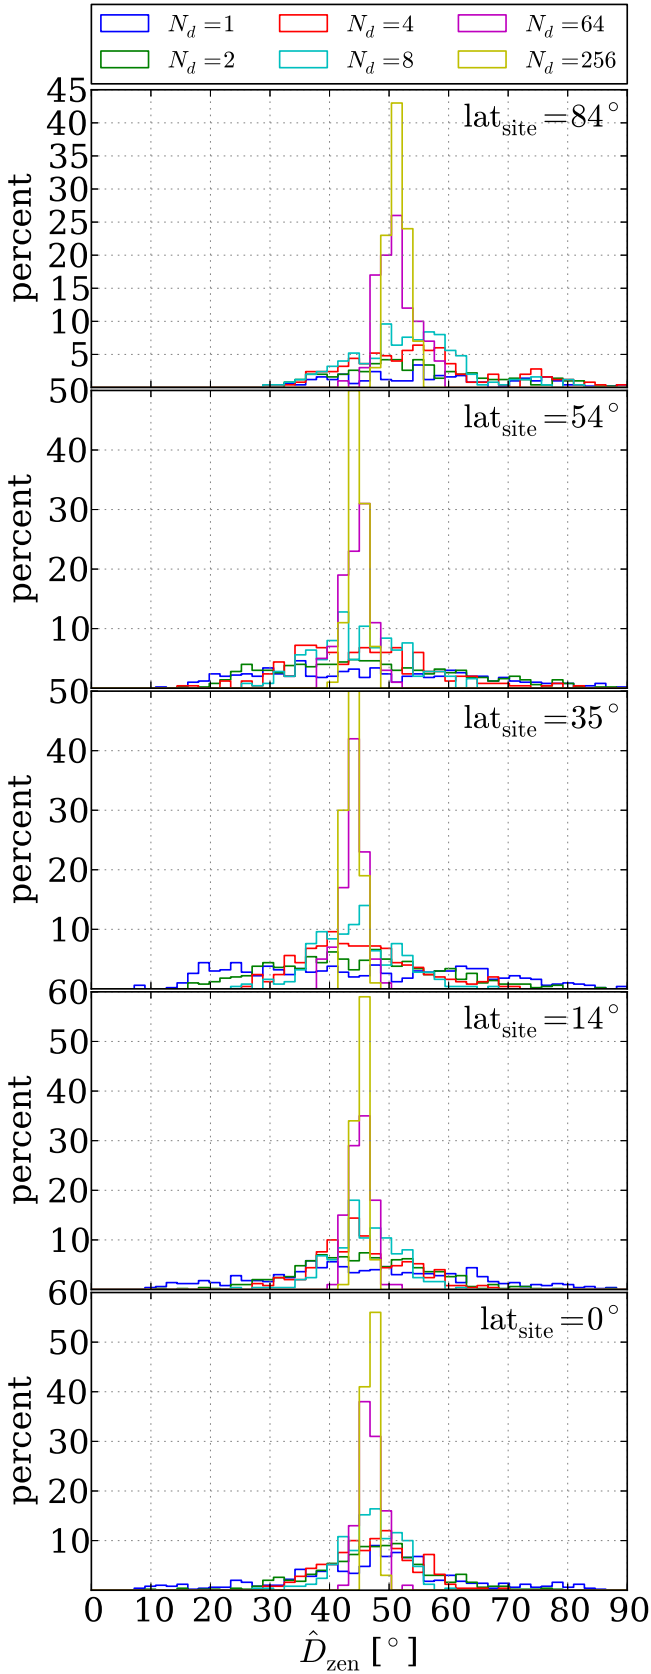
Figure 8. Distributions of $\hat{D}_{\mathrm{zen}}$ for a few latitudes and N_d for year-long averages. We note that the distributions are rather broad and all centered near similar values. This is reflected in the wide error bars and similar means in Figure 4. Nonetheless, the mean of the distribution collapses to the same values as in Figure 9 as $N_d \rightarrow \infty$.
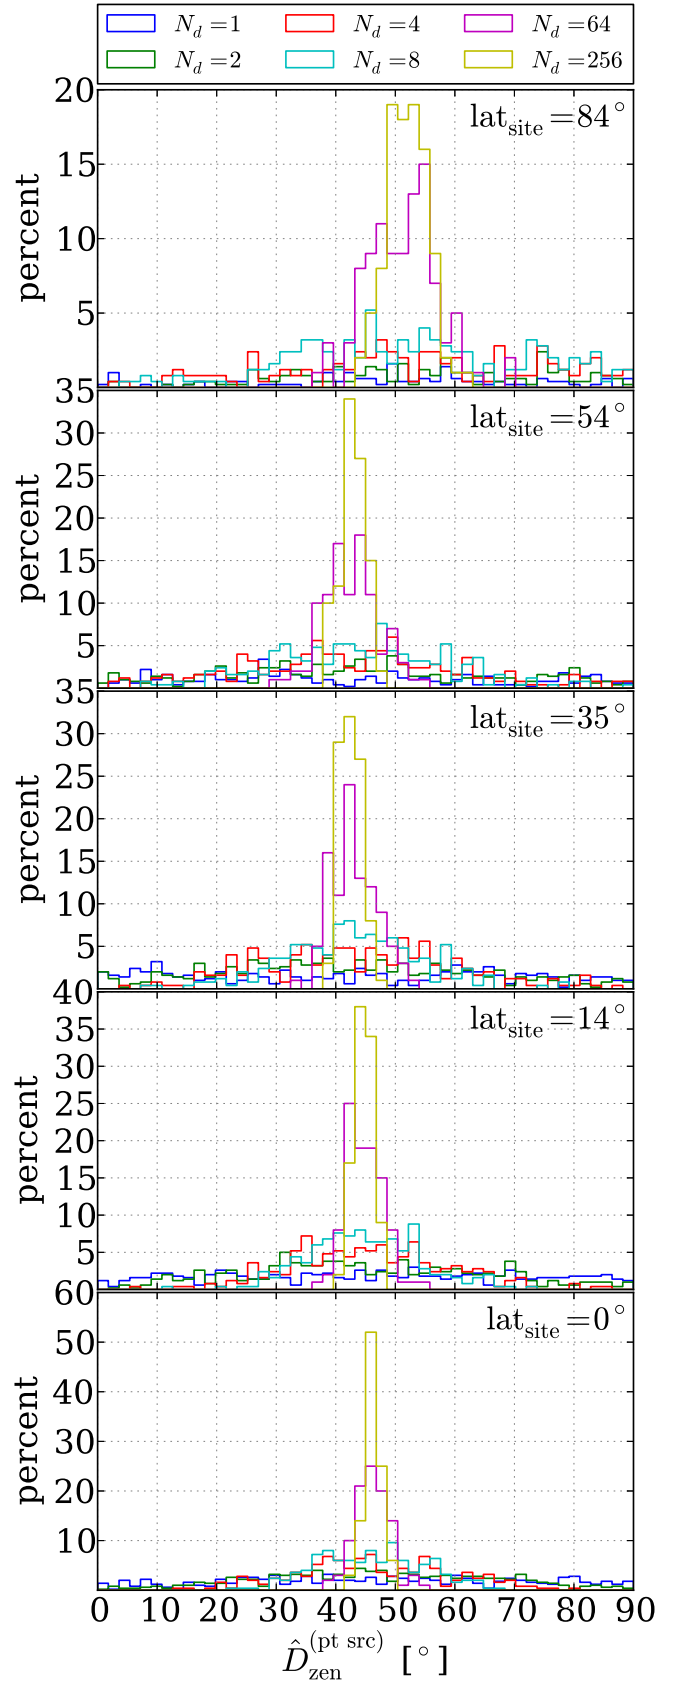
Figure 9. Distributions of $\hat{D}_{\mathrm{zen}}^{(\mathrm{pt\ src})}$ for a few latitudes and N_d for year-long averages. We note that these distributions have more shot noise than those in Figure 8, evident in the modes of the mid-latitude distributions. $\hat{D}_{\mathrm{zen}}(N_d = 1)$ is often larger than $\hat{D}_{\mathrm{zen}}^{(\mathrm{pt\ src})}(N_d = 1)$ because the triangulation rings can reach accross most of the antenna pattern with only two detectors.
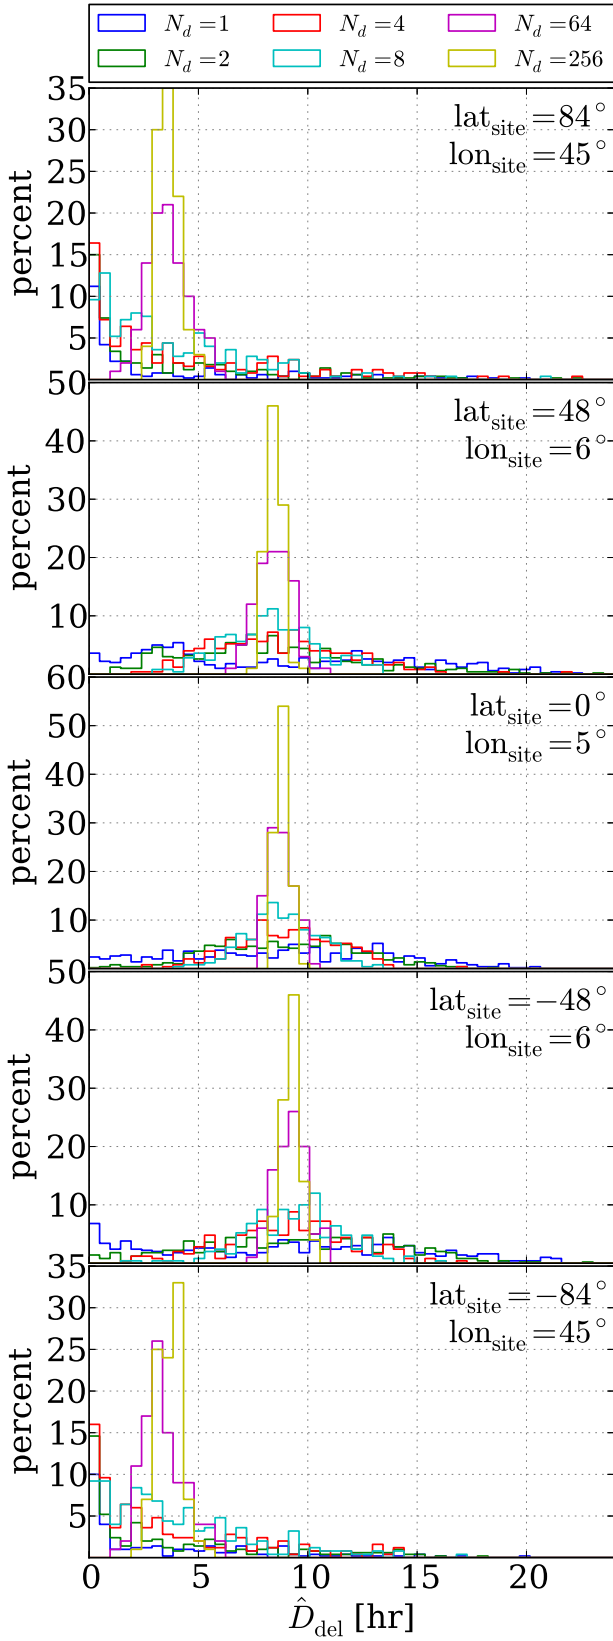
Figure 10. Distributions of $\hat{D}_{\mathrm{del}}$ for a few sites and N_d for year-long averages. We note that the $N_d = 1$ distributions have large modes near $\hat{D}_{\mathrm{del}} = 0$, corresponding to events that are immediately observable, as well as very broad support extending to long $\hat{D}_{\mathrm{del}}$. This is a particular evident near the poles.
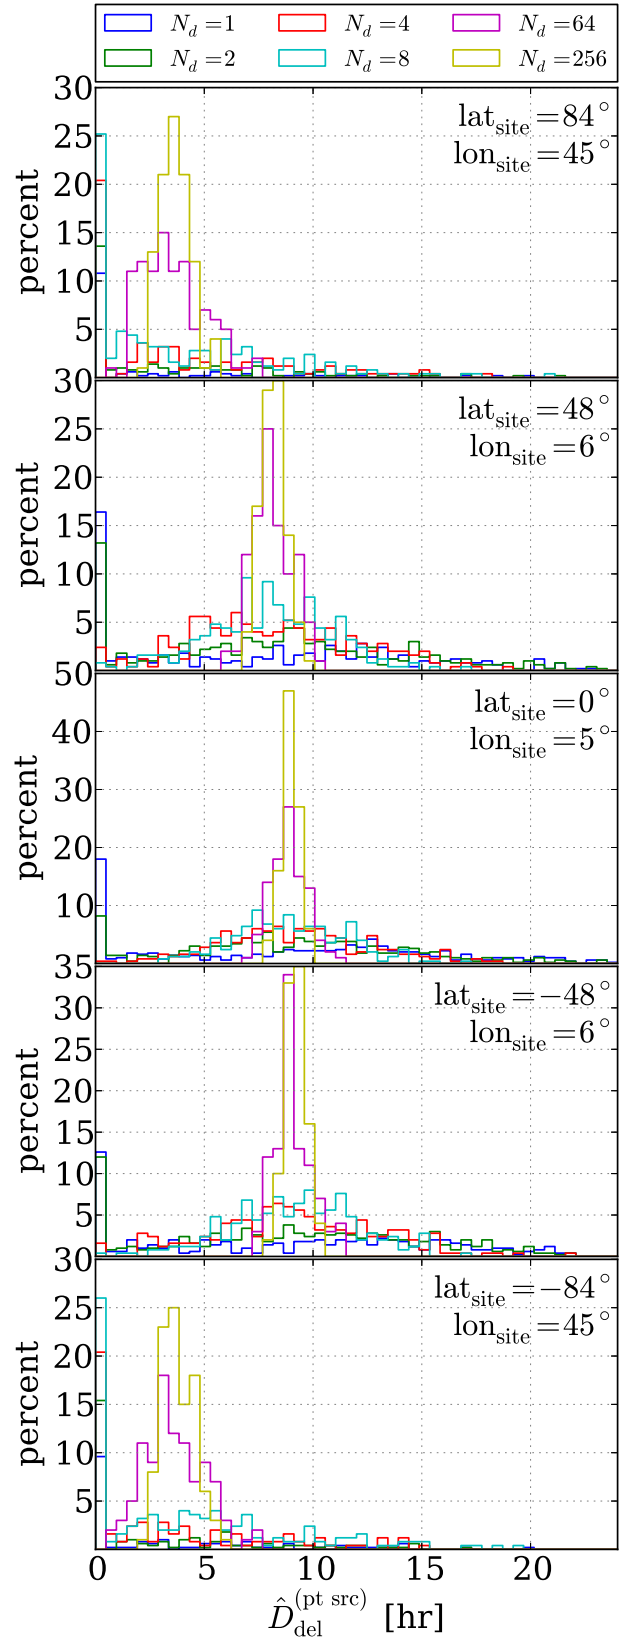
Figure 11. Distributions of $\hat{D}_{\mathrm{del}}^{(\mathrm{pt\ src})}$ for a few sites and N_d for year-long averages. The modes near $\hat{D}_{\mathrm{del}}^{(\mathrm{pt\ src})}(N_d = 1)$ are more pronounced than in Figure 10 because there is more shot noise in the point-source measurement. In fact, averaging over the skymap tends to broaden the distributions for $N_d \lesssim 5$.

because different parts of the skymap may become observable at different times, and that fuzz tends to smooth the distribution.

APPENDIX C
DERIVATION OF MODEL FOR $p(t|\text{operating})$

While Equation (1) is fairly self evident from the inspection of Figure 1, we can derive its form from more basic assumptions about how human activity may cause downtime. We posit two states of a network of detectors: *up* (u) and *down* (d), with science-quality data available only in the *up* state. Furthermore, we posit two possible causes for detectors being in the *down* state: *random* causes (r) that are uncorrelated with time and *cyclic* causes (c) which are correlated with time, usually through a diurnal cycle. We note that the *random* and *cyclic* models are not necessarily mutually exclusive and the detector could be down for mutliple reasons at the same time.

Therefore, we have

$$p(d|t) = p(d|c, t)p(c|t)$$
$$+ p(d|r, t)p(r|t) - p(d|c \cap r, t)p(c \cap r|t),$$

which implies

$$p(d) = \int dt p(t)p(d|t)$$
$$= \int dt\, p(t)[p(d|c, t)p(c|t) + p(d|r, t)p(r|t)$$
$$- p(d|c \cap r, t)p(c \cap r|t)]$$

Furthermore, because the *up* and *down* states are mutually exclusive and span the space of possible detector states at any single time, we have

$$p(u|t) = 1 - p(d|t) \Rightarrow p(u) = 1 - p(d)$$

and Bayes theorem yields

$$p(t|u) = \frac{p(u|t)p(t)}{p(u)} = \frac{(1 - p(d|t))p(t)}{1 - p(d)}.$$

We measure $p(u)$ and $p(t|u)$ in a straightforward manner from the data used in Figure 1. Typically, we assume some periodicity in $p(t|u)$ to generate a histogram with enough samples to be statistically meaningful, but we expect to be able to identify the periodic elements of $p(t|u)$ through Fourier analysis as well.

However, we are also interested in slightly different probabilities. Namely, we would like to know

$$p(c|d, t) = \frac{p(d|c, t)p(c|t)}{p(d|t)},$$

$$p(r|d, t) = \frac{p(d|r, t)p(r|t)}{p(d|t)}$$

and

$$p(c \cap r|d, t) = \frac{p(d|c \cap r, t)p(c \cap r|t)}{p(d|t)}$$

which express the probabilities that the detector is down due to a particular cause at a specific time. If we only care about long-term averages (over timescales much longer than the *cyclic* model's periodicty), then we can marginalize away the time dependence:

$$p(c|d) = \int dt p(t) \frac{p(d|c, t)p(c|t)}{p(d|t)}$$

$$p(r|d) = \int dt p(t) \frac{p(d|r, t)p(r|t)}{p(d|t)}$$

and

$$p(c \cap r|d) = \int dt p(t) \frac{p(d|c \cap r, t)p(c \cap r|t)}{p(d|t)}.$$

These equations should hold regardless of the specific form of the *cyclic* and *random* models.

By definition, we assume that the *random* causes are uncorrelated with time so that $p(d|r, t) = p(d|r)$. Furthermore, we assume some periodic function for the *cyclic* model so that $p(d|c, t) = p(d|c, t + \tau)$ for some τ. Specifically, we expand the periodic function in terms of the oscillating (AC) and constant (DC) components

$$p(d|c, t) = p_{DC}(d|c) + p_{AC}(d|c, t)$$

such that

$$p_{DC} = \int_0^\tau dt\, p(d|c, t).$$

Clearly, we require $p_{DC}(d|c) \geqslant -p_{AC}(d|c, t)\, \forall t$. We also assume the priors for the causes do not depend on time ($p(r|t) = p(r)$ and $p(c|t) = p(c)$) and that the priors are equal for the two causes ($p(r) = p(c) = p$).

Note that by measuring $p(d)$, we only extract the combination of $p_{DC}(d|c) + p(d|r)$ and cannot separate these terms further. However, by measuring $p(t|u)$ as well, we are able to determine $p_{AC}(d|c, t)$ from which we can determine $p_{DC}(d|c)$ by requiring that $\min_t\{p(d|c, t)\} = 0$. Any other DC component to the *cyclic* model is indistinguishable from the *random* model and therefore we lump it together with the *random* model.

If we allow the cause models to overlap but require them to be independent, we obtain $p(d|c \cap r, t) = p(d|c, t)p(d|r, t)$. This implies

$$p(r) + p(c) - p(r)p(c) = 1 \Rightarrow p(r) = p(c) = p = 1$$

and the two observables are

$$p(d) = p(d|r)p + (1 - p(d|r)p)p_{DC}(d|c)p$$

$$p(t|u) = p(t)\frac{1 - p(d|r)p - (1 - p(d|r)p)p_{DC}(d|c)p - (1 - p(d|r)p)p_{AC}(d|c, t)p}{1 - p(d|r)p - (1 - p(d|r)p)p_{DC}(d|c)p} = p(t)\left(1 - \frac{(1 - p(d|r)p)p}{1 - p(d)}p_{AC}(d|c, t)\right).$$

This yields

$$p(c|d, t)$$

$$= \frac{p_{AC}(d|c, t) + p_{DC}(d|c)}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

$$p(r|d, t)$$

$$= \frac{p(d|r)}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

$$p(c \cap r|d, t)$$

$$= \frac{(p_{AC}(d|c, t) + p_{DC}(d|c))p(d|r)p}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

and their marginalized counterparts

$$p(c|d) = \int dt\, p(t|d)$$

$$\times \frac{p_{AC}(d|c, t) + p_{DC}(d|c)}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

$$= \frac{p_{DC}(d|c)p}{p(d)}$$

$$p(r|d) = \int dt\, p(t|d)$$

$$\times \frac{p(d|r)}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

$$= \frac{p(d|r)p}{p(d)}$$

$$p(c \cap r|d) = \int dt\, p(t|d)$$

$$\times \frac{(p_{AC}(d|c, t) + p_{DC}(d|c))p(d|r)p}{(p_{AC}(d|c, t) + p_{DC}(d|c))(1 - p(d|r)p) + p(d|r)}$$

$$= \frac{p_{DC}(d|c)p \cdot p(d|r)p}{p(d)}.$$

A reasonable ansatz is $p(d|c, t) = B\left(1 + \sin\left(\frac{2\pi t}{\tau} - \phi\right)\right)$, in which case we obtain

$$p(d) = p(d|r)p + Bp(1 - p(d|r)p)$$

$$p(t|u) = p(t)\left(1 + \frac{(1 - p(d|r)p)pB}{1 - p(d)}\sin\psi\right)$$

$$= p(t)(1 + A\sin\psi),$$

which is exactly the form of Equation (1). We also see why a low duty cycle (small $p(u) = 1 - p(d)$) can amplify the amplitude of they day/night bias.

If we are interested in just the probability associated with the *cyclic* model regardless of the *random* model, we obtain

$$p(c|d) = \frac{1 - p(d)}{p(d)}\left(\frac{A}{1 - p(d) + (1 - p(d))A}\right) \approx 0.49,$$

which implies that nearly half the time *cyclic* causes were at least partly responsible for bringing down the detector network during O1. If we restrict ourselves to times when the *cyclic* causes were the sole cause of the downtime, we obtain

$$p(c|d) - p(c \cap r|d) = \frac{(1 - p(d))A}{p(d)} \approx 0.29,$$

which suggests that we could reduce the downtime by 30% if we completely removed *cyclic* causes of downtime.

REFERENCES

Aasi, J., Abadie, J., Abbott, B. P., et al. 2013, PhRvD, 87, 022002
Aasi, J., Abadie, J., Abbott, B. P., et al. 2014, ApJS, 211, 7
Aasi, J., Abadie, J., Abbott, B. P., et al. 2015, CQGra, 32, 115012
Abadie, J., Abbott, B. P., Abbott, R., et al. 2012a, A&A, 541, A155
Abadie, J., Abbott, B. P., Abbott, R., et al. 2012b, PhRvD, 85, 082002
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, PhRvL, 116, 061102
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, CQGra, 33, 134001
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, LRR, 19, 1A
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016d, PhRvL, 116, 131103
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016e, PhRvL, 116, 241103
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016f, arXiv:1602.08492
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001
Aso, Y., Michimura, Y., Somiya, K., et al. 2013, PhRvD, 88, 043007
Barnes, J., & Kasen, D. 2013, ApJ, 775, 18
Berger, E. 2014, ARA&A, 52, 43
Collaboration, T. L. S., Aasi, J., Abbott, B. P., et al. 2015, CQGra, 32, 074001
Essick, R. 2015, Localization of Short Duration Gravitational-wave Transients
    with the Early Advanced LIGO and Virgo Detectors Open Data Release,
    (Online; accessed 2016 March 24) (http://www.ligo.org/scientists/burst-
    first2years/)
Essick, R., Vitale, S., Katsavounidis, E., Vedovato, G., & Klimenko, S. 2015,
    ApJ, 800, 81
Fairhurst, S. 2014, JPhCS, 484, 012007
Gehrels, N., Sarazin, C. L., O'Brien, P. T., et al. 2005, Natur, 437, 851
Gossan, S. E., Sutton, P., Stuver, A., et al. 2016, PhRvD, 93, 042002
Kasliwal, M. M., & Nissanke, S. 2014, ApJL, 789, L5
Kulkarni, S. R., Frail, D. A., Wieringa, M. H., et al. 1998, Natur, 395, 663
NASA/IPAC Infrared Science Archive 2016, Planck Public Data Release 2
    Mission Ancillary Data
Patat, F., Ugolnikov, O. S., & Postylyakov, O. V. 2006, A&A, 455, 385
Singer, L. P., Price, L. R., Farr, B., et al. 2014, ApJ, 795, 105
Thorne, K. S. 1987, in Gravitational Radiation, ed. S. W. Hawking & W. Israel
The LIGO Scientific Collaboration, the Virgo Collaboration, Abbott, B. P.,
    et al. 2016a, arXiv:1606.04856
The LIGO Scientific Collaboration, the Virgo Collaboration, Abbott, B. P.,
    et al. 2016b, arXiv:1602.03842