Effect of E. coli mitigation on the proportion of time primary contact minimum acceptable state concentrations are exceeded: Technical note.

Prepared for Ministry for the Environment

December 2016

Prepared by:
Sandy Elliott
Amy Whitehead

For any information regarding this report please contact:

Sandy Elliott
Principal Scientist
Catchment Processes
+64-7-859 1839
Sandy.Elliott@niwa.co.nz

National Institute of Water & Atmospheric Research Ltd
PO Box 11115
Hamilton 3251

Phone +64 7 856 7026

NIWA CLIENT REPORT No: 2016135HN
Report date: December 2016
NIWA Project: MFE16209

Quality Assurance Statement		
	Reviewed by:	G. McBride
A. Bartley	Formatting checked by:	A. Bartley
	Approved for release by:	Neale Hudson

Contents

Tables

Figures

Executive summary

This report addresses the question “by how much would mitigation measures decrease the proportion of time that river reaches exceed the primary contact (‘swimming’) minimum acceptable state *E. coli* concentration of 540 per 100 ml”. This information was requested by the Ministry for the Environment (MfE) to support policy analysis regarding suitability of water for primary contact, relating to the National Objective Framework (NOF) in the National Policy Statement for Freshwater Management (NPS-FM) (New Zealand Government 2014). The particular mitigation measures were taken from a previous national modelling study (Elliott et al. 2016), and relate to mitigations beyond the those currently in place. The results are summarised by the proportion of stream length nationally that exceeds a concentration of 540 per 100 ml for specified fractions of the time (5%, 10%, 20% and 30%). Stream length in this case relates to streams in the River Environment Classification (REC). While the NOF minimum acceptable state for primary contact relates to a concentration that is exceeded 5% of the time (that is, the 95th-percentile concentration), there was interest in how often the concentration of 540 is exceeded. For example, even if mitigation measures don’t result in a concentration of 540 being met for 95% of the time, could they go some way to reducing the proportion of time that the concentration of 540 is exceeded? This proportion could serve as a measure of progress towards achieving primary contact goals. This report provides technical documentation of the methods used in this analysis and key summary results.

The analysis combines models of: a) the current fraction of time that a concentration is greater than 540 per 100 ml (which we will call the exceedance of 540, or G_{540}); b) load and concentration reductions associated with new mitigation measures (beyond those already in place); and c) statistical models of the probability distribution of concentrations at a site to determine how G_{540} reduces in response to concentration reductions. The calculations were performed for each reach nationally, and results were then summarised as the fraction of total national stream length exceeding G_{540} cut-off levels of 5%, 10%, 20% and 30% (fraction of stream length exceeding a concentration of 540 per 100 ml for specified fractions of the time). The analysis is very approximate in nature, relying on multiple models and assumptions. Therefore the results must be seen as exploratory rather than definitive, giving an indication of the overall magnitude and patterns of change rather than accurate predictions.

Mitigation measures were predicted to have a small effect (<3.9% reduction) on the proportion stream length that exceeds a concentration of 540 per 100 ml for specified fractions of the time. This result arises from the modest changes in load (mean value of 4.0%) in conjunction with the wide spread of concentration values at each site. The response was greater at the 30% exceedance level (reduction from 28.5% of stream length to 24.8% of stream length), compared with the 5% exceedance level (reduction from 75.9 of stream length to 75.2% of stream length). Hence, progress towards target concentration levels would be more demonstrable at the 30% exceedance level. Greater load reductions (associated with more extreme mitigation measures or more optimistic evaluations of their effectiveness) would result in larger reductions in the proportions of stream lengths exceeding a given G_{540} level.

1 Methods

The analysis is based on multiple steps:

1. The current fraction of time that that the concentration is greater than 540 per 100 ml (which we will call the exceedance of 540, or G_{540}) is based on a statistical model fitted to the NEMAR3 (National Environmental Monitoring and Reporting version 3) dataset[1] and applied to all REC Version 1 reaches nationally. The statistical model was a random forest model that used the same predictors and methods as in previous NEMAR work (Unwin and Larned 2013). The resulting G_{540} values were supplied to NIWA by MfE (pers. comm., Ton Snelder, 14 Dec. 2016).

2. The fraction reduction in loading associated with new mitigation measures was obtained from a national model of *E. coli* loads, which takes the current extent of mitigation and the efficacy of various mitigation measures into account (Elliott et al. 2016).

3. It was assumed that a given fraction reduction in loading for a stream reach would result in concentrations being reduced by the same fraction, across the full range of concentrations in that reach. It was also assumed that reductions would be equally effective for concentrations at the low or high end of the concentration range. Storm-flow loading of *E. coli* can dominate over baseflow-loading, so that reductions in storm-flow loading might have little effect on baseflow concentrations. To avoid this situation, the assessment of the effectiveness of mitigation measure focussed on baseflow conditions relevant to swimming, rather than storm conditions where swimming is not likely.

4. The reduction in G_{540} for a reach was based on the reduction in concentrations and a log-normal parametric probability distribution for the concentrations. The log-normal distribution was fitted to the concentration data for each site, using the fitdist function in fitdistrplus package within the R statistical software[2].

5. The log-sd parameter of this distribution, (σ , standard deviation of log-transformed concentrations) was modelled nationally using a random forest model, with predictors as in previous NEMAR work. The other parameter of this distribution, the log-mean (μ), was derived from the given G_{540} and the estimated log-sd, and the formula for the log-normal distribution (in which $ln\ C$ follows a normal distribution with mean $ln\ \mu$ and standard deviation σ)[3]:

   $$F(lnC;\sigma,\mu) = 0.5 + 0.5erf\left(\frac{lnC-\mu}{\sqrt{2}\sigma}\right)$$

   where, C is the concentration, and erf is the error function. G is by definition 1- F, and C in our case is 540 per 100 ml. Hence, inverting the previous equation, the formula for μ is:

   $$\mu = \ln 540 - \sqrt{2}\sigma erf^{-1}(1-2G_{540})$$

   A concentration reduction factor D reduces lnC by an increment lnD. Hence, the new $G_{540}$for a concentration reduction factor D is:

---

[1] https://dc.niwa.co.nz/niwa_dc/srv/eng/metadata.show?id=285&currTab=simple. This dataset covers the period 1992 to 2012. For later analysis steps, the NRWQN dataset was extended to December 2016, to enable more robust analysis.

[2] https://www.r-project.org/

[3] https://en.wikipedia.org/wiki/Log-normal_distribution

$$G'_{540} = 0.5 - 0.5erf\left(\frac{\ln 540 - \mu - \ln D}{\sqrt{2}\sigma}\right)$$

The performance of this method was assessed using an empirical determination of G_{540} for measurement sites, where all concentration measurements were reduced by a factor D.

6. The reduction in G_{540} was then determined for every reach nationally, and the results were summarised as a proportion of reach length exceeding values of 5%, 10%, 20% and 30%, as requested by MfE.

2 Results and discussion

2.1 Fit of lognormal distribution to concentration data

Sites with more than 5% low-censored data (14 sites) were removed from the analysis, leaving 467 sites, with a median of 84 samples. The log-normal distribution was found to be an acceptable probability distribution in most cases (e.g., Figure 2-1). The distribution tended to be wider than the normal distribution at the upper end of the distribution (long-tailed compared with the normal distribution), for about half the sites (e.g., Figure 2-2). More complex distributions such as the GEV distribution fitted the data marginally better than the normal distribution at the tails of the distribution. But, considering that these more complex distributions introduce additional parameters which would be difficult to predict nationally, and we are not primarily interested in the extremes of the distribution, we retained the log-normal distribution for subsequent analysis.

For the example distribution shown in Figure 2-1, G_{540} for this site is 0.54 (54%). If concentrations are reduced by 10% (D= 0.9), then the concentrations are shifted uniformly to the left in the figure, and G_{540} reduces to 50% based on empirically shifting the data, with similar number for shifting the fitted distribution. This limited sensitivity to D results from the wide distribution of concentrations. The sensitivity is less at the tails of the distribution (due to the nature of the normal distribution).

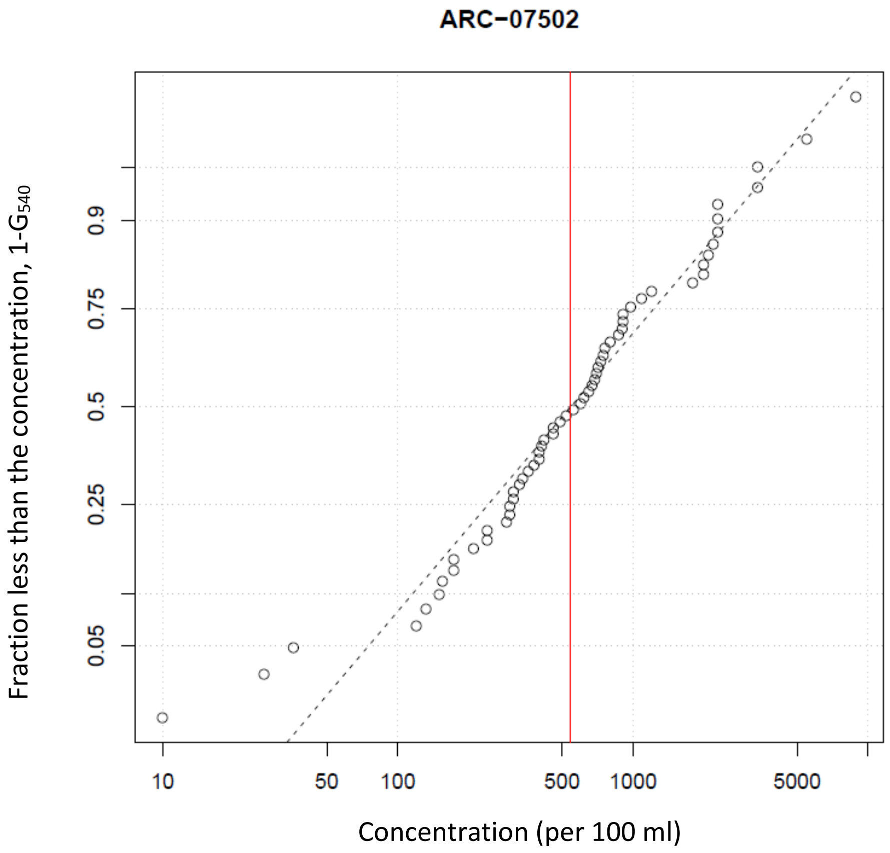

Figure 2-1: An example probability distribution, showing the log-normal fit. The vertical red line is at a concentration of 540 per 100 ml. The dashed line is a log-normal fit to the data. The horizontal scale is log-transformed, and the vertical scale is on a normal distribution scale.

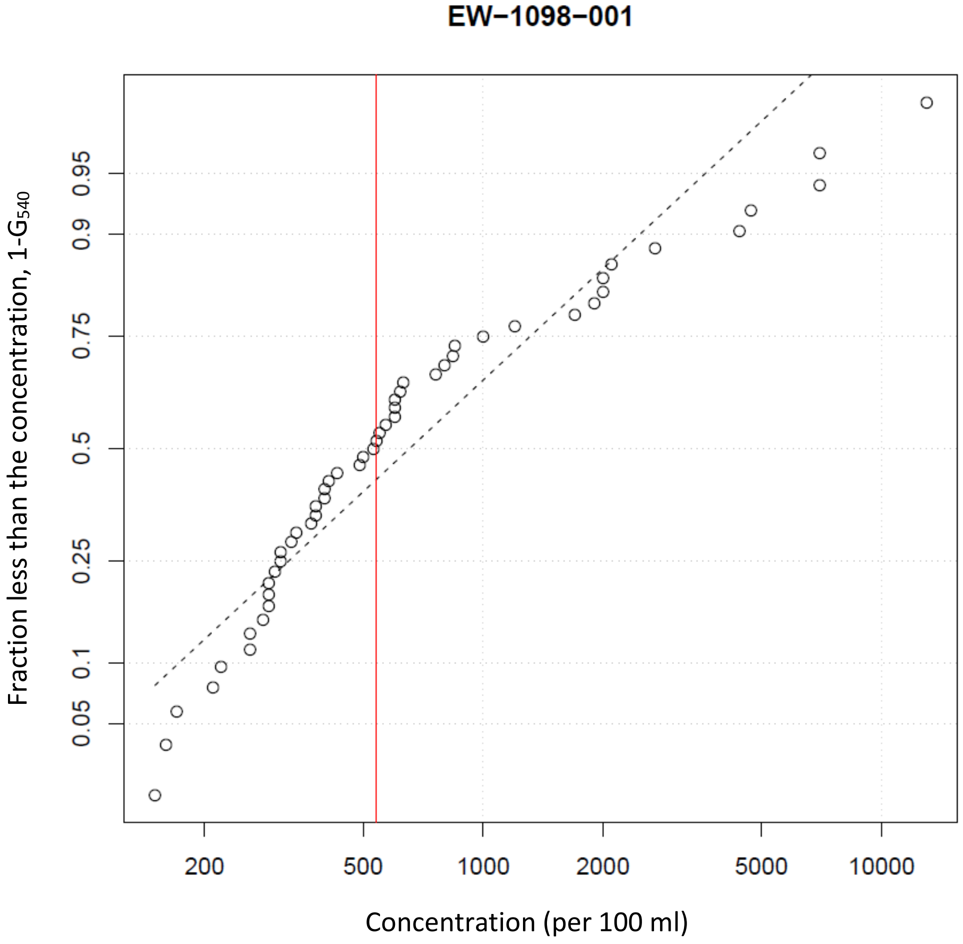

Figure 2-2: A further probability distribution example, showing the log-normal fit in a case where the fit is not as good as in the previous example (Figure 2-1). The vertical red line is at a concentration of 540 per 100 ml. The dashed line is a log-normal fit to the data. The horizontal scale is log-transformed, and the vertical scale is on a normal distribution scale.

2.2 Sensitivity of G_{540} to given percentage reduction in concentration at measurement sites.

To assist with interpretation of the results, we calculated the reduction in G_{540} for a 10% concentration reduction for all measurements (without fitting distributions), see Figure 2-3 and Figure 2-4. This concentration reduction is in the middle of the range expected for *E. coli* mitigation in the load modelling exercise (Table 2-1 and Figure 2-7). To avoid the influence of local variability in the data, the data were smoothed using lowess (locally-weighted linear regression), with a span of 0.1 of the data.

The reduction in G_{540} is fairly modest, with a maximum reduction of about 0.07. For sites with current $G_{540}$of about 0.3 (exceeding a concentration of 540 per 100ml for 30% of the time), the reduction in G_{540} is about 0.06, whereas for lower G_{540} values of 0.05, the reduction in G_{540} is smaller.

The proportion of sites with $G_{540} > 0.2$ decreases from 48% to 44% with a 10% reduction in concentration. This small response is due in part to the wide distribution of concentrations as mentioned above, but also arises because sites with high concentrations (G_{540} much larger than 0.2) will continue to have high concentrations after the 10% reductions, and sites with low concentrations (G_{540}<0.2) will remain with G_{540} <0.2 after the concentration reduction.

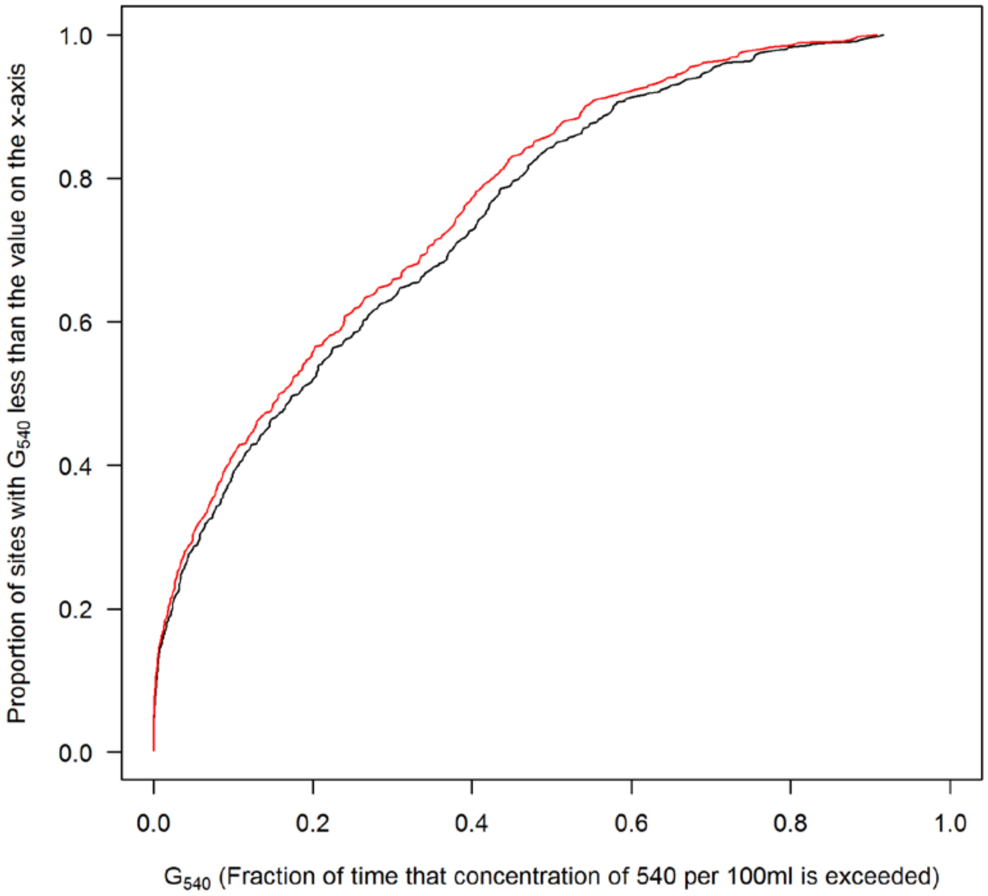

Figure 2-3: Distribution of G_{540} across measurement sites for a 10% reduction in concentration, based on measurement sites. The black line is the distribution before concentrations are reduced, while the red line is the new G_{540}.

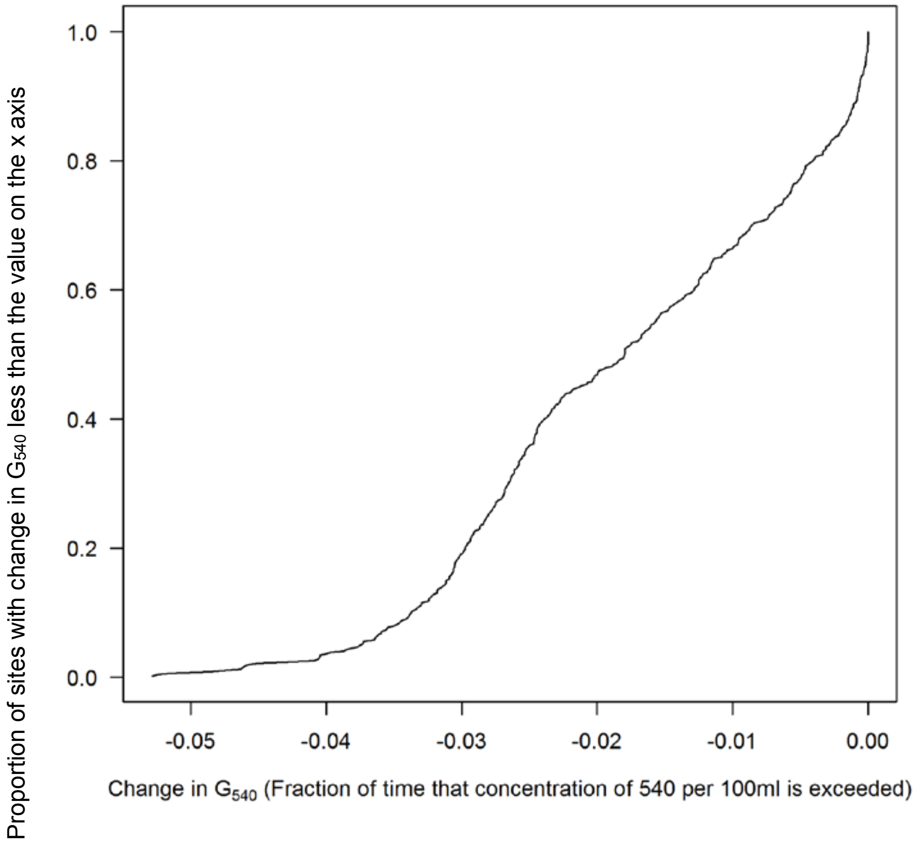

Figure 2-4: Distribution of reduction in G_{540} across measurement sites for a 10% reduction in concentration.

2.3 Evaluation of suitability of the log-normal distribution for determining changes in G_{540}.

The performance of the log-normal method for determining changes in G_{540} is assessed for a given 10% change in concentration (D = 0.9). This reduction is in the middle of the range expected for *E. coli* mitigation in the load modelling exercise, for locations where reductions occur (Table 2-1). For the data-based method, the change in G_{540} was assessed by shifting all the original data by the factor D and interpolating G_{540} from the new empirical distribution. For this interpolation, lowess curve smoothing with a span of 0.1 was used, to avoid effects of very local variations in the data. This data-based approach was compared with the log-normal parametric method, where the current G_{540} is taken empirically from the data and the new G_{540} is determined from the method described in Section 1, and σ is from the fitted distribution.

The results of this comparison are shown in Figure 2-5. The comparison shows that the log-normal distribution provides a reasonable estimate of the sensitivity of $G_{540}$to changes in concentration, provided that the current G_{540} and σ are known. There is considerable scatter for greater reductions (more negative values). Part of this is due to the local wiggles in the cumulative distribution of the data, which were only partially damped by the smoothing.

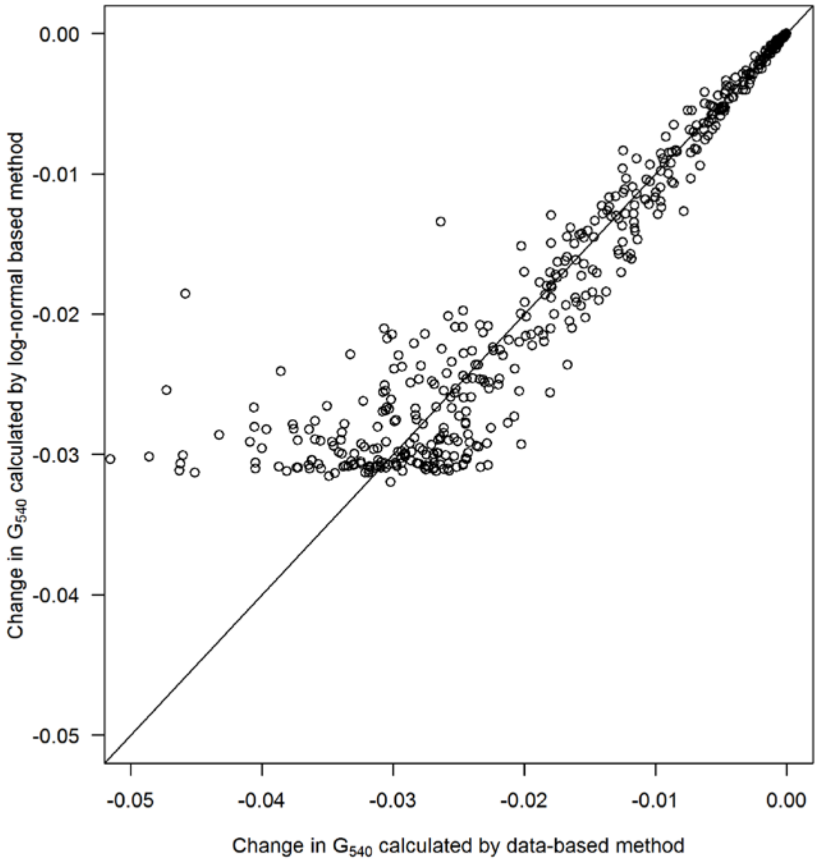

Figure 2-5: Comparison between data-based and log-normal based method for change in G_{540} for a 10% reduction in concentration.

2.4 Estimation of the log-sd parameter, σ.

The random forest model was not able to successfully predict variations in σ. The root-mean-square-error of the predictions was comparable to the standard deviation of the data.

Hence, we decided to base subsequent analysis on the empirical distribution of σ from the measurement sites (Figure 2-6). Note that σ did not vary systematically with the log-mean concentration, so we do not need to apply different σ for different concentrations or G values, and there is also fairly small range of σ (and hence good prospects for aggregate results that we are ultimately interested in having low variability).

To take account of variability in σ, we conducted a Monte Carlo assessment of the effect of variations of this parameter. We applied the Monte Carlo method for determining the proportion of stream length that exceeds specified values of G_{540} , the key output metric of interest. For each Monte Carlo iteration, we took values of σ randomly from a normal distribution (mean 1.34, standard deviation 0.25), with separate values for each stream reach, and determined the proportions of stream length that exceeded the G_{540} value of interest. This process was iterated and the resulting distribution of stream length proportions was analysed. We found that variability in σ resulted in very little bias or variability in the proportion of stream exceeding the given G_{540} value.

This occurred because smaller changes in G_{540} associated with high σ were counterbalanced by larger changes in G_{540} associated with low σ, and there were a large number of reaches involved giving stable proportions of stream length. Hence we report values only for the mean proportion of stream length.

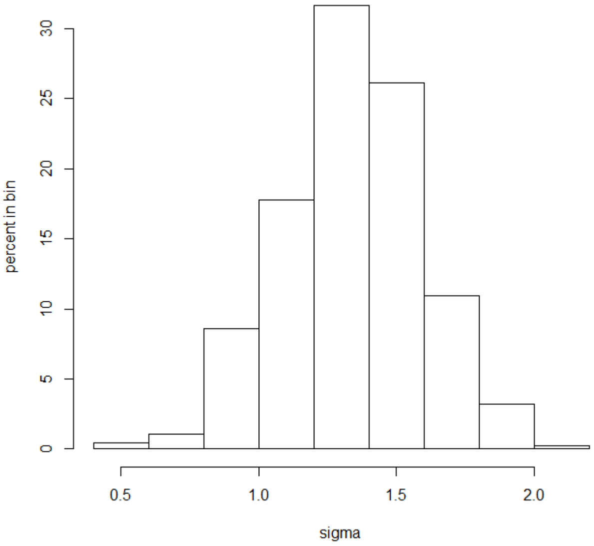

Figure 2-6: Distribution of log-sd (σ) based on the measurement sites.

2.5 Effect of mitigation on the proportion of stream length with G_{540} greater than specified cut-off levels.

2.5.1 Changes in loading from mitigation measures

The percentage reductions of loading associated with mitigation measures, provided as input to the load model (Elliott et al. 2016) are modest (Table 2-1). This is because many of the key mitigation measures such as stock exclusion from waterways and effluent management have already been implemented widely, and some mitigation measures are unsuitable for intensive sheep and beef areas.

The resulting reduction in load from mitigation measures in the catchment model, which takes account of the distribution of land uses and cumulative effects down the catchment, is consequently modest (Figure 2-7), with a mean reduction of 4%. The curve in Figure 2-7 flattens out at 14% reduction, corresponding to intensive sheep and beef, and is near vertical near that point because there are a number of catchments with largely sheep and beef land use. There is a very slight increase beyond that point associated with the small amount of deer land use.

Table 2-1: Load reduction for new mitigation measures. These values relate to mitigation beyond that already in place.

Land use	Percentage load reduction
Dairy	3
Dairy Canterbury	3
Deer	25
Hill sheep and beef	0
Intensive sheep and beef	14

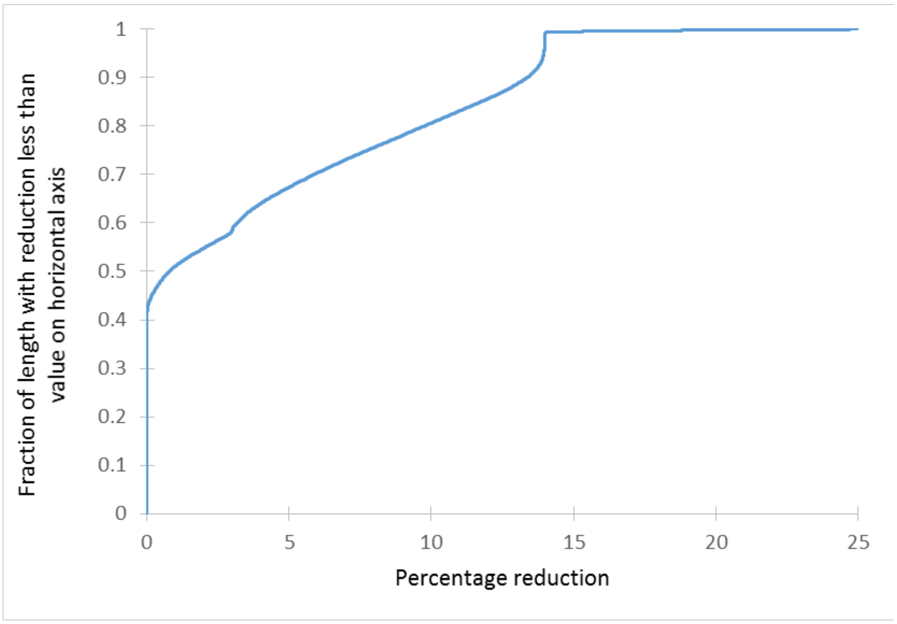

Figure 2-7: Cumulative distribution in percentage reduction in load associated with new mitigation measures, from the catchment model. The fraction of reaches will be a close approximation to the fraction of stream length.

2.5.2 Reduction in proportions of stream length with G_{540} greater than specified cut-off levels.

The proportion of stream length with G_{540} greater than cut-off values of 5%, 10%, 20%, and 30% is shown in Table 2-1 and Figure 2-8, with and without mitigation. Note that the cut-off value of 5% corresponds to the minimum acceptable state for primary contact under the NOF.

As the cut-off value of G_{540} increases, the proportion of stream length exceeding that cut-off value reduces. So, for current conditions, 75.9% of stream length exceeds a concentration of 540 per 100 ml for more than 5% of the time, but only 28.7% of stream length exceeds 540 per 100 ml for more than 30% of the time.

New mitigation is predicted to have only a small effect (<3.9% reduction) on the proportion of stream length that exceeds a given G_{540}. This is due to: a) the small degree of mitigation on average; b) the wide nature of the concentration distribution at a site, leading to low sensitivity of G_{540} to changes in concentration at a site; and c) the small proportion of stream reaches with G_{540} both a) currently greater than the cut-off value and b) close enough to the cut-off to move below the cut-off when G_{540} is reduced.

Mitigation has a larger effect for larger G_{540} cut-off values. At a cut-off of 5%, there is only a 1.0% reduction due to mitigation, while for a cut-off of 30%, there is a 3.9% response to mitigation. This relates to the flatter nature of the normal probability distribution towards the tails of the distribution. From the sigmoid nature of the normal distribution, values which are exceeded 5% of the time are at the tails, whereas values which are exceeded 30% of the time occur more towards the central steeper part of the distribution. Moreover, sites with only 5% exceedance of 540 per 100 ml may be more likely to occur in catchments with low degrees of catchment development, where the degree of mitigation is smaller.

The implication of this result is that the proportion of streams that exceed a concentration of 540 per 100 ml for 30% of the time or less provides a more sensitive measure of effects of mitigation, compared with using the minimum acceptable state of 5% of the time. As the minimum acceptable state is approached, the goal becomes more elusive. This diminishing-returns effect is related to the long tails of concentration distributions.

Introducing larger or more widespread mitigation will result in a larger effect on G_{540}. For example, doubling the amount of mitigation (doubling the percentage reductions from Table 2-1) resulted in approximate tripling of the changes in proportions of stream length with G_{540} greater than specified cut-off values (approximate tripling of the values in the right column of Table 2-2).

Table 2-2: Proportion of stream length with G_{540} greater than specified cut-off values. G_{540} is the fraction of time that a stream location exceeds a concentration of 540 per 100 ml.

Cut-off value of G_{540}	Proportion of stream length with G_{540} greater than cut-off value		
	Without new mitigation	With new mitigation	Reduction due to mitigation
5%	75.9%	75.2%	0.7%
10%	60.9%	59.9%	1.0%
20%	45.2%	43.6%	1.6%
30%	28.7%	24.8%	3.9%

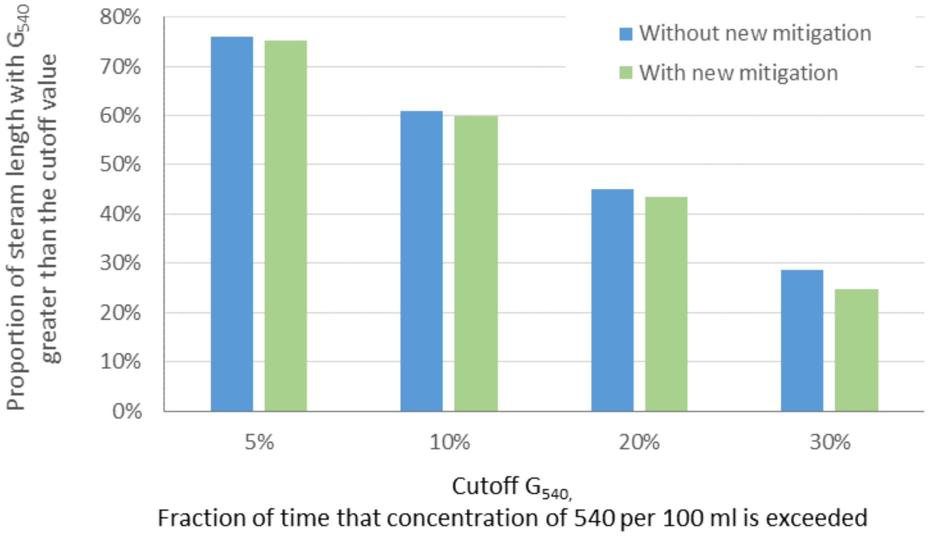

Figure 2-8: Proportion of stream length with G_{540} greater than cut-off levels, with and without new mitigation. G_{540} is the fraction of time that a stream location exceeds a concentration of 540 per 100 ml.

2.6 Caveats regarding model uncertainty

This analysis is very approximate in nature, relying on multiple models and assumptions, each with their own uncertainty, and we are not able to quantify many aspects of individual or combined uncertainty. Therefore the results must be seen as only indicative of the magnitude of change. Assumptions and inaccuracies include: the current exceedance of 540 per 100 ml is based on predictions of statistical models fitted to measurement sites; the current load is based on a mass accounting type model, with source parameters derived from statistical model fitting; reductions in loads are based on uncertain information regarding the efficacy and current extent of mitigation; mitigation effectiveness and types of mitigation are limited to those assessed in by Elliott al. (2016) (for example, no load reduction for sheep and beef areas); it is assumed that a given percentage reduction in loading will result in concentrations being reduced by that same percentage across the range of concentrations; and changes in concentration exceedance associated with a reduction in concentration is determined with a probability distribution that entails uncertainty both in the accuracy of fit and in the distribution parameters. Despite these uncertainties, the results will be of interest as an indication of the sensitivity of 540 per 100 ml to concentrations in general, and to specific scenarios of mitigation measures from the catchment model.

3 References

Elliott, S., Wadhwa, S., Whitehead, A., Snelder, T., Muirhead, R., Monaghan, R. (2016) Modelling national land-use capacity. Exploring bottom lines and headroom under the NPS-FM 2014. Update Report. Prepared for the Ministry for the Environment June 2016. *NIWA Client Report* 2016103HN, National Institute of Water and Atmospheric Research.

New Zealand Government (2014) National Policy Statement for Freshwater Management 2014. New Zealand Government.

Unwin, M.J., Larned, S.T. (2013) Statistical models, indicators and trend analyses for reporting national-scale river water quality (NEMAR Phase 3). *NIWA Client Report* CHC2013-033, National Institute of Water and Atmoshperic Research.