Regarding the benefit of zero-dimensional noise

Daniel H. Baker   Department of Psychology, University of York, York, UK

Tim S. Meese   School of Life & Health Sciences, Aston University, Birmingham, UK

Baker and Meese (2012) (B&M) provided an empirically driven criticism of the use of two-dimensional (2D) pixel noise in equivalent noise (EN) experiments. Their main objection was that in addition to injecting variability into the contrast detecting mechanisms, 2D noise also invokes gain control processes from a widely tuned contrast gain pool (e.g., Foley, 1994). B&M also developed a zero-dimensional (0D) noise paradigm in which all of the variance is concentrated in the mechanisms involved in the detection process. They showed that this form of noise conformed much more closely to expectations than did a 2D variant.

Allard and Faubert (2013) (A&F) criticized this work on several grounds. There are several aims to our reply:

1. To reinforce our point that masking from 2D noise does not derive purely from the influence of variability: Suppression is also involved. This means that 2D noise masking experiments might not be revealing the processes that their designers intended. This is a valuable conclusion from our original work using 0D noise and receives further experimental support here (Appendix A).
2. To point out the shortcomings in several of A&F's arguments and address several errors and misunderstandings that emerge.
3. To present our own critique of the 0D noise masking paradigm.
4. To further illustrate the value of the 0D noise-masking paradigm in an experimental context (Appendix A).

We begin by correcting a mistake in the second sentence of A&F's critique, where they mischaracterize the concept of 0D noise. It does not involve jittering just the target contrast, as they state. Instead, it involves adding a pedestal to both intervals of two-alternative forced-choice (2AFC) trial with zero mean contrast (negative contrasts are carried by a 180° switch in phase). This means there is an independent source of noise in each of the 2AFC intervals. It seems that this mischaracterization of our stimulus does not lead them to further misunderstandings, but we feel bound to alert readers to this error.

In 2AFC, the responses from two intervals can be compared within the framework of signal detection theory (SDT). One contains noise alone; the other contains signal plus noise. This is true for 0D external noise, regardless of the noise level. Thus when the external noise level is zero, the limiting noise is entirely internal. When the external 0D noise is sufficiently high, it dominates the internal noise, which becomes irrelevant. Determining the external noise level at which this transition occurs is the basis of the EN paradigm (Pelli, 1990). But regardless of either noise level, the observer compares a signed response from each interval, each of which contains noise and only one of which contains signal. Operationally, the task is exactly the same in each case (compare two responses and choose the interval that contains the target), and the mathematics and application of SDT are identical. Thus, 0D noise has exactly the characteristics we want from it and the general processing strategy is the same, regardless of the level of the noise.

We wondered whether some of A&Fs misgivings on this point had been driven by the intuitive concern that the 0D noise task "feels" more like a contrast discrimination task than a contrast detection task. We accept that this might be the case, so long as we keep in mind that the discrimination is not based on *absolute* contrast, but signed (i.e., phase dependent) contrast (where negative phase counts as negative contrast). Nevertheless, their point is of little relevance. Whilst it is obvious that for 0D noise there is typically nonzero stimulus contrast in each of the two intervals, exactly the same is true of 2D noise so long as the noise level is above its own detection threshold. Furthermore, even for 2D noise it must also be the case that there is typically nonzero activity within the detecting mechanism(s) in each interval, otherwise masking could not be attributed to the direct effects of noise in that

paradigm at all! Essentially, detection in any type of noise (and hence all detection) *is* discrimination, even when the noise is purely internal. This is because there are nonzero responses in each interval, not deterministic "detect" versus "not-detect" states as implied by long-discredited high-threshold assumptions (e.g., Nachmias, 1981). Note that the instruction to our experienced observers was: "Indicate whether the target was presented in the first or second interval," with the caveat that targets were always in positive cosine phase relative to fixation. As pointed out by a reviewer, this is equivalent to saying: "Select the interval with the brightest horizontal bar at fixation." These instructions were clearly sufficient for observers to perform the 0D noise task in the fovea. Whether the same is true for peripheral presentation, where phase information is often lost (e.g., Solomon, 2002) remains to be determined.

On a related tack, A&F complained that 0D noise masking does not lend itself to being used in a yes/no (y/n) detection task and used this to try and claim that unlike the 2D case, the task is no longer one of detection in the 0D case. Their position on this is wrong. In a suitably designed y/n experiment, including catch trials to assess false alarms, the tenets of SDT can be brought to bear in the analysis of the 0D noise mask in exactly the same way as for the 2D noise mask. We do not see that the 0D noise experiment would be impractical, as A&F claim. A&F's objection seems to be that in a y/n task, "... the observer almost always perceives a target." This statement is wrong. It should say, "... there is almost always external drive to the detecting mechanism." Whether the observer judges that response to derive from a *target* or not will depend on how they set their internal criterion for detection in the task. This is no different from when 2D noise is used, or when there is no noise. In all cases, when the response is much higher than the criterion for detection, the observer will typically say "yes," and there is a high probability that there was a target, in which case the observer scores a hit. When the response is only just above the criterion then the observer still says "yes," but if the target contrast was zero, then that is a false positive. It should be clear that observers would be wise to set their criterion high if the noise level is high (regardless of whether that noise is internal or external, or 0D or 2D).

Of course, one problem with high levels of 0D noise (and 2D noise) is that the combination of signal and noise in the detecting mechanism can result in the response of the detecting mechanism being driven into a different part of its operating characteristic. If sufficiently high, this might involve response compression, equivalent to self-suppression from contrast gain control (Foley, 1994). We refer readers back to our original paper for further discussion of this point.

The treatment above considers only a single detecting mechanism. Although A&F don't spell this out, one argument might be that contrast detection involves multiple mechanisms with receptive fields in different positions across the image, for example. Then we must ask how these mechanisms might be combined. One widely used approach is to suppose that the observer constructs a template of the target, and this is used to weight the filtered signal and noise at each location in the image. The contributions from each mechanism "beneath" the template are then summed to construct the decision variable (see Meese & Summers, 2012 for a recent example of this). In this case, when the external noise is zero, performance is limited by multiple sources of independent internal noise, the summation of which results in a single noise source. When the external noise is high enough to be the limiting noise, then for the 2D case, we have the same situation as before. The situation in the 0D case is slightly different, in that the multiple noise sources are 100% correlated. This means there is no longer any benefit in combining information from multiple mechanisms. Nevertheless, in all cases, the observer's decision variable is constructed from a comparison of signal and noise with noise alone.

A&F commented that, "... if different processing strategies underlie contrast detection in no and 0D noise, then this compromises the application of the EN paradigm." However, our contention is that this is exactly the problem with 2D noise—it invokes suppression from the broadly tuned gain control mechanism. 0D noise, on the other hand does not, other than the caveat about self-suppression mentioned above. Thus, 0D noise addresses the noise in only a single detector. If 2D noise is used in an attempt to try and estimate the EN for multiple mechanisms then it will be confounded by the suppressive effects from the contrast gain pool. This illustrates clear practical limitations to the EN paradigm.

There are other criticisms that might be levelled at the use of 0D noise. For example, it is possible that 0D noise reduces uncertainty, whereas this is less likely for 2D noise. However, A&Fs criticisms do not appear to have this type of detail in mind.

A&F complain that 0D noise is not suitable for measuring the tuning properties of the detecting mechanism. The fact that it cannot was so obvious we saw no need to mention it. However, we might add that we are skeptical about the general approach of using external noise to assess the excitatory tuning properties of visual mechanisms. The problem is not specific to noise, but to masking in general. The logic behind most masking studies is that performance is disturbed when the mask stimulus excites the detecting mechanism. In the noise-masking paradigm, the disturbance is an increase in response variance, in the pattern-masking paradigm, the disturbance is a result of response
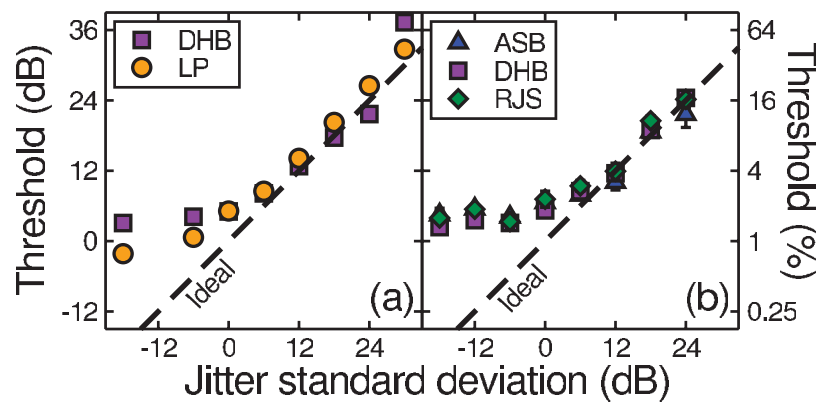
Figure 1. 0D noise masking functions from (a) Baker & Meese (2012) and (b) the experiment described in Appendix A. The oblique dashed line represents the prediction of the ideal observer.

compression (Wilson, McFarlane, & Phillips, 1983). However, the central difficulty is that there are (at least) two critical degrees of freedom in a well-constructed model of the process: the tuning of the detecting mechanism and the tuning of the suppressive gain pool. It is difficult to design experiments that provide suitable constraints on these two parameters (though see Foley, 1994) and many studies don't even attempt to, ignoring the suppressive effects of the contrast gain pool altogether. We consider this risky at best (see Meese & Holmes, 2010). As the same criticisms are likely to apply to 2D noise masking when used as A&F propose, we advise that it be done with caution, at the very least.

A further criticism of 0D noise advanced by A&F is that it cannot reveal any properties of the detection process because thresholds in high-variance 0D noise correspond to those expected for the ideal observer (they show this for their one subject, RA). Figure 1 shows data from two experiments measuring detection in 0D noise (those in Panel A are replotted from B&M as absolute detection thresholds, having been refitted using a cumulative Gaussian instead of a Weibull function). The prediction of the ideal observer is given by the oblique dashed line $x = y$ for thresholds at $d' = 1$. Our data confirm that human observers approach the ideal prediction in this task, within the range of experimental error.

The observation above tends to reinforce A&F's claim that 0D noise is of no value because it conforms to the behaviour of the ideal observer. For example, this means that the EN can be estimated from the unmasked detection threshold and the ideal prediction without having to gather any noise-masking results at all. However, this misses an important point about the role of 0D noise experiments. For example, the main aim of our earlier paper (Baker & Meese, 2012) was to demonstrate the shortcomings of 2D pixel noise: that it is not the research tool it is often assumed to be. A&F do not dispute our position on 2D noise, and we consider our claims to be safe. That study involved

comparing results from 2D noise experiments with those from 0D noise. We found that 2D noise failed to meet predictions for noise masking in four independent experimental tests. In stark contrast, 0D noise did meet those predictions, illustrating that the limitations were not simply attributable to experimental protocol or other difficulties in linking a paradigm that originates from the study of electrical hardware with neural wetware (Pelli, 1981).

Regardless of the value of 0D noise within the equivalent noise paradigm, it has wider applicability as a method for introducing variance into an observer's responses. If participants behave like noisy ideal observers (as A&F propose), the double-pass procedure (Burgess & Colborne, 1988) offers an alternative method for estimating the ratio of internal to external noise. Yet previous studies had always found that double-pass data produced different estimates of internal noise to the equivalent noise method (e.g., Lu & Dosher, 2008), necessitating additional model parameters or mechanisms (such as induced internal noise) to explain this. The discrepancy is easily understood when the suppressive effects of noise masks are taken into account. (This also requires additional degrees of freedom in the model of course, but ones well supported by independent studies such as Foley, 1994, and Meese & Holmes, 2007.) The 0D noise stimulus produces much stronger double pass consistency than 2D noise, even when the two noise types are equated for the level of threshold elevation that they produce (experiment 2 of B&M). Thus, it is a useful tool for assessing internal noise levels in a range of situations which would otherwise be difficult to compare, e.g., across sensory modalities (Neri, 2010).

In other work, described in Appendix A, we have found 0D noise masking to be a valuable tool for investigating combined influences of pattern masking (suppression from gain control) and noise masking. We did this by performing masking experiments that contained a pure noise mask component (0D noise) in

addition to a pattern mask. Our aim was to determine whether pattern masking includes a component of noise masking (from induced noise; e.g., Burgess & Colborne, 1988) or whether it is purely suppressive. To do this, we needed a noise mask for which we were confident there was little or no effect from suppression. 0D noise is the only noise mask to meet this criterion. By measuring and modelling 0D noise masking in the presence of a pattern mask we were able to determine that (a) pattern masking does not involve induced noise and (b) the two forms of masking are not additive. Neither of these conclusions could have been reached using traditional noise masks which themselves contribute suppression. Thus, far from being the irrelevant stimulus for characterizing the processing properties of detection mechanisms that A&F claim, we have found it to be a useful tool in furthering our understanding of exactly that.

In sum, we agree with A&F that 0D noise is not suitable for measuring mechanism selectivity, but this observation is so plainly true we consider it trivial. Furthermore, this does not mean that 0D noise is of no value, as pointed out above and in Appendix A. Nevertheless, we are grateful to A&F for replicating our finding about the reduction of the slope of the psychometric function in 0D noise, and for providing us the opportunity to discuss the issues they raise, since we think it is possible that other readers might share their misgivings and misunderstandings.

Acknowledgments

Commercial relationships: none.
Corresponding author: Daniel Hart Baker.
Email: daniel.baker@york.ac.uk.
Address: Department of Psychology, University of York, UK.

References

Allard, R., & Faubert, J. (2013). Zero-dimensional noise is not suitable for characterizing processing properties of detection mechanisms. *Journal of Vision, 13*(10):26, 1–3, http://www.journalofvision.org/content/13/10/26, doi:10.1167/13.10.26.

Baker, D. H., & Meese, T. S. (2012). Zero-dimensional noise: The best mask you never saw. *Journal of Vision, 12*(10):20, 1–12, http://www.journalofvision.org/content/12/10/20, doi:10.1167/12.10.20. [PubMed] [Article]

Burgess, A. E., & Colborne, B. (1988). Visual signal detection. IV. Observer inconsistency. *Journal of the Optical Society of America A, 5,* 617–627.

Foley, J. M. (1994). Human luminance pattern-vision mechanisms: Masking experiments require a new model. *Journal of the Optical Society of America A, 11,* 1710–1719.

Holmes, D. J., & Meese, T. S. (2004). Grating and plaid masks indicate linear summation in a contrast gain pool. *Journal of Vision, 4*(12):7, 1080–1089, http://www.journalofvision.org/content/4/12/7, doi:10.1167/4.12.7. [PubMed] [Article]

Klein, S. A., & Levi, D. M. (2009). Stochastic model for detection of signals in noise. *Journal of the Optical Society of America A, 26,* B110–B126.

Lu, Z.-L., & Dosher, B. A. (2008). Characterizing observers using external noise and observer models: Assessing internal representations with external noise. *Psychological Review, 115,* 44–82.

Meese, T. S., & Holmes, D. J. (2007). Spatial and temporal dependencies of cross-orientation suppression in human vision. *Proceedings of the Royal Society B, 274,* 127–136.

Meese, T. S., & Holmes, D. J. (2010). Orientation masking and cross-orientation suppression (XOS): Implications for estimates of filter bandwidth. *Journal of Vision, 10*(12):9, 1–20, http://www.journalofvision.org/content/10/12/9, doi:10.1167/10.12.9. [PubMed] [Article]

Meese, T. S., & Summers, R. J. (2012). Theory and data for area summation of contrast with and without uncertainty: Evidence for a noisy energy model. *Journal of Vision, 12*(11):9, 1–28, http://www.journalofvision.org/content/12/11/9, doi:10.1167/12.11.9. [PubMed] [Article]

Nachmias, J. (1981). On the psychometric function for contrast detection. *Vision Research, 21,* 215–223.

Nachmias, J. (1993). Masked detection of gratings: The standard model revisited. *Vision Research, 33,* 1359–1365.

Neri, P. (2010). How inherently noisy is human sensory processing? *Psychonomic Bulletin & Review, 17,* 802–808.

Pelli, D. G. (1981). *Effects of visual noise* (Unpublished doctoral dissertation), University of Cambridge, Cambridge, UK. Retrieved from http://psych.nyu.edu/pelli/pubs/pelli1981thesis.pdf.

Pelli, D. G. (1990). The quantum efficiency of vision. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 1–24). Cambridge, UK: Cambridge University Press.

Solomon, J. A. (2002). Noise reveals visual mechanisms of detection and discrimination. *Journal of Vision*, *2*(1):7, 105–120, http://www.journalofvision.org/content/2/1/7, doi:10.1167/2.1.7. [PubMed] [Article]

Wilson, H. R., McFarlane, D. K., & Phillips, G. C. (1983). Spatial tuning of orientation selective units estimated by oblique masking. *Vision Research, 23*, 873–882.

Appendix A

Nonadditivity of stochastic and deterministic masks

To illustrate the benefit of 0D noise, we considered how suppression from cross-channel grating masks combines with threshold elevation due to external noise. We measured noise masking functions for 0D noise masks in three observers, using the procedures described in B&M. Stimuli (horizontal 1 c/° log-Gabor patches) were displayed using a Bits++ on a gamma corrected NEC MultiSync P1150 monitor. The results are shown for each observer separately in Figure 1b, and the average is given by the blue function in Figure A1a. The 0D masking function closely resembles those described previously.

We then repeated the experiment with an additional superimposed grating mask at three times the target frequency (3F), with an orientation of −45° from vertical and a contrast of 32% (30 dB). The 3F grating mask elevated detection thresholds in weak 0D noise by around a factor of four (12 dB), consistent with previous work (Holmes & Meese, 2004). But for strong 0D noise (18 & 24 dB), thresholds converged with those in the absence of the 3F mask. This demonstrates the nonadditivity of these two forms of masking—a very different phenomenon from the linear combination of two grating masks (Holmes & Meese, 2004).

Cross-channel masking is typically attributed to suppression from a gain pool (cf. Nachmias, 1993). In Figure A1b, we show that a standard divisive suppression model (Foley, 1994) predicts the arrangement of both masking functions, including their convergence. But an alternative explanation is that the
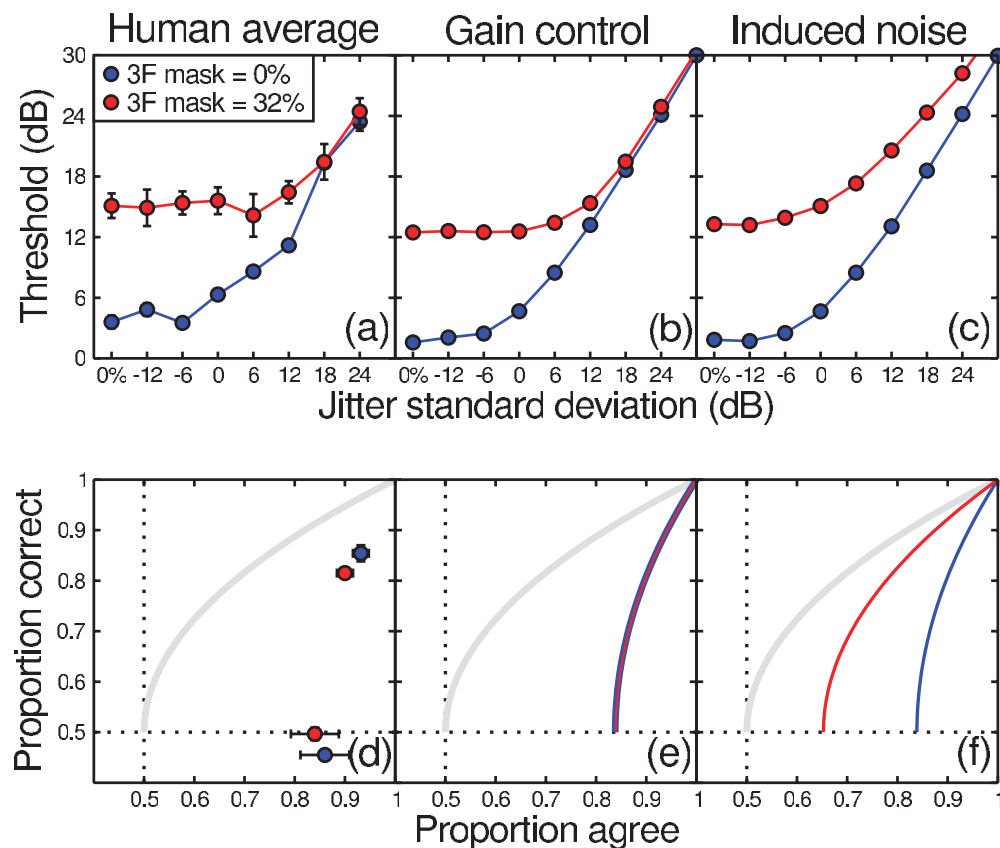

Figure A1. Data and model predictions for 0D noise masking with and without a cross-channel mask. (a) Masking functions averaged across three observers. (d) PCPA data from a double pass experiment using the same stimuli, comprising 400 MCS trials per observer. The gray curve shows the expected agreement in the absence of external noise (Klein & Levi, 2009). (b), (c), (e), (f) Predictions of two models for the same conditions, as described in the text. Error bars give ±1 *SEM* across observers.

3F mask increases the level of internal noise in the system. This "induced noise" (e.g., Lu & Dosher, 2008) can produce masking resembling that from divisive suppression (Figure A1c; details of model simulations are given in Appendix B).

Although both models make similar predictions regarding the masking functions, they differ in one critical way. Because internal noise remains constant in the gain control model, it predicts that the consistency of observer responses in a double pass experiment (see Burgess & Colborne, 1988) will be unchanged when a 3F mask is added (red and blue functions superimpose in Figure A1e). In contrast, if the 3F mask increases the level of internal noise, then it should reduce observer response consistency dramatically (red and blue functions diverge in Figure A1f).

We collected double pass data for two target contrast levels (0% and 22%) at a jitter standard deviation of 16% (24dB), both with and without the 32% contrast 3F mask. The results (Figure A1d) show unambiguously that response consistency is not affected by the 3F mask. This allows us to reject induced noise as a source of cross-channel masking. Note that the use of 0D noise was critical for reaching this conclusion, as any variety of noise that itself introduced suppression would have combined differently with the 3F mask and affected the results.

Appendix B

Model details

The standard gain control model is given by the equation:

$$resp = \frac{C^{2.4}}{Z + C^2 + wX^2} + G_\sigma, \tag{1}$$

where C is the contrast in the detecting mechanism tuned to the target, X is the summed contrast of a gain pool with weight $w = 0.05$, Z is a saturation constant with a value of two, and G represents internal noise with standard deviation $\sigma = 0.5$ (parameter values were chosen to produce illustrative behavior approximating that of the human data and are consistent with values obtained elsewhere by fitting). The induced noise model is similar, except that the cross-channel mask produces additional internal noise proportional to its contrast:

$$resp = \frac{C^{2.4}}{Z + C^2} + G_{wX} + G_\sigma, \tag{2}$$

with all terms retaining their previous meanings and values. The induced internal noise term (G_{wX}) is equivalent to that proposed in previous studies (e.g., Burgess & Colborne, 1988; Lu & Dosher, 2008) in that it is summed with the additive internal noise following the transducer. Although previous implementations have made the induced noise proportional to the noise contrast, the purpose of induced noise is to account for the mask effects outside of the detecting channel (Lu & Dosher, 2008, point out the similarity to contrast gain control). So, in Equation 2 the induced variance (G_{wX}) is proportional to the activity in nontarget mechanisms (e.g., the mask contrast) but not to that within the detecting channel.

We simulated 5,000 trials per target contrast level to produce a full psychometric function at each mask level. An independent sample of external noise was added to the target contrast on each interval of every trial (in the null interval the target contrast was zero). Thresholds were estimated from these simulated data using Probit analysis. To produce the double pass predictions, the first half of the simulated trials used the same samples of external noise as the second half, but different samples of internal noise.