

Supplemental Information: Discussion

Distinction between response bias in 2-AFC tasks and choice bias in m-ADC tasks

Two alternative forced choice (2-AFC) or Yes/No tasks, for example, those that employ response probes (Hawkins et al. 1990; Luck et al. 1996; Pestilli et al. 2011; Wyart et al. 2012) permit constructing a 2x2 contingency table that can be analyzed with one-dimensional signal detection models. Such tasks permit quantifying sensitivity and *response bias* at cued and uncued locations, and do not require the m-ADC model framework for analysis. However, there is an important difference between response probe tasks and m-ADC tasks *in terms of the bias that they quantify*. This distinction is highly relevant to the present study and dictated our choice of task and model. We elaborate on this difference next.

In response probe tasks subjects are required to detect or discriminate events, specifically, at the response-probed location on each trial. The most relevant location for attention is the location that is likely to be probed for subsequent response; this is typically indicated to the subject *apriori* with a relevance cue (e.g. Pestilli et al. 2011, Wyart et al. 2012). Importantly, signal probability at each location is decoupled from other locations, by task design. On every trial, signal events (e.g. change events) occur independently across locations, such that, on any trial, signals can happen at multiple locations, simultaneously. In addition, the subject makes independent decisions (e.g. Yes/No) at one of the several locations on each trial. Consequently, in response probe tasks, probability of each type of signal is irrelevant as an attention cue.

In contrast, in m-ADC attention tasks (a variant of the standard Posner cueing task, Fig. 1, main text), subjects are required to detect and localize events (e.g. change) at one of multiple locations. The likely location of signal events is cued with probabilistic, spatial cues. Since only one signal event occurs on each trial, a higher signal probability at one location (cued location) implies a lower probability at other locations (uncued locations). On every trial, signal events are not independent across locations, such that, on any trial, the occurrence of a signal event at one location precludes the occurrence of a signal event at all other locations. Subjects make a single localization decision on each trial. The most relevant location for attention is, therefore, the location with the highest signal probability (location where the signal event is most likely to occur). Thus, in the m-ADC attention task *signal probability is the only relevant attention cue*.

These key differences between the two task designs have important consequences for the relationship between attention, sensitivity and bias in each task type. *Specifically, criteria in the two task types differ in the nature of the bias that they quantify.*

Criteria in response probe tasks, as measured with one-dimensional 2-AFC models, reflect a response bias for reporting change versus no change (e.g. Wyart et al. 2012) or a bias for reporting one alternative (interval) over another (e.g. Pestilli et al. 2011). Such biases typically reflect the relative probabilities of each signal type, independently, at each location. Relative values of 2-AFC criteria across locations do not typically represent the bias (or weightage) for sensory evidence at one location over another for making decisions. In contrast, relative values of decision criteria -- as measured in m-ADC attention tasks -- indeed represent this latter kind of bias. The m-ADC bias represents a *spatial bias for localizing events*, which, by definition, is not independent across locations.

We illustrate the essential distinction between the biases in the two task types with an example.

Consider a response probe task involving change detection (Fig. A-1A, below) in which subjects make decisions about change events at different locations (e.g. cued versus uncued) on different trials. On each trial, the subject must make a Yes/No decision at the response probe location independent of sensory evidence at the other location. On no trial does the subject need to compare sensory evidence across different locations for making a decision. Criteria measured with the 2-AFC model in this type of task represent the bias for detecting a signal versus no signal event at each location, independently of other locations. Therefore, the criterion at each location would align closely with signal prior probability (expectation) at that location, independent of other locations, and independent of cue relevance. Thus, 2-AFC criteria in response probe tasks do not measure biases linked to attention. *These 2-AFC response criteria likely reflect differential signal or differential reward expectation at each location, rather than attention mechanisms*, a conclusion also supported by many other studies (Wyart et al. 2012; Summerfield and Egner 2009).

We illustrate this further in Figure A-1B. For example, for the response probe design shown in the figure, signal probabilities can be equally high (e.g. 0.8/0.8), equally low (e.g. 0.2/0.2) or completely unrelated (e.g. 0.9/0.4) across both locations; the response probe task design does not preclude this. As a consequence, an equally high or equally low response bias may occur at all locations (cued and uncued) in response probe tasks (Fig. A-1B). Similarly, changes in response criteria may be induced by selectively rewarding one type of response over another, or by varying the probabilities of reward for each type of response at each location (e.g. Luo and Maunsell, 2015, 2018). For example, an equally low criterion (favoring Yes responses) can be induced at both stimulus locations by providing higher rewards for hit relative to correct rejections, at both locations. Consequently, 2-AFC measures of criteria and bias in response probe tasks do not reflect attention mechanisms.

In contrast, in the m-ADC attention task (Fig. A-2A) subjects must decide whether a change event occurred at the cued location or at one of the uncued locations on every trial. Decisions are not made independently across the different locations: *Rather, the subject must directly compare sensory evidence across both cued and uncued locations on every trial to make a single localization decision*. Therefore, the relative values of decision criteria, as measured by the m-ADC model represent a spatial choice bias for localizing change events at the cued versus uncued locations. Specifically, a larger difference between the criteria at the cued versus uncued locations (e.g. ADI-bias, Fig 3D (bottom), main text), reflects a greater weight (spatial bias) for sensory evidence at the cued location, when making the decision to localize the change (see Fig. A-2 caption). This spatial bias corresponds to biasing the competition in favor of sensory evidence at the cued location when making the localization decision. *This interpretation of the relative values of m-ADC criteria as a spatial choice bias is possible because of the multidimensional space (Fig. A-2 B-C) in which decisions are modeled in the m-ADC attention task*.

To summarize: The m-ADC bias reflects a decision policy that affords greater weight to signal information at the cued versus uncued locations, biasing the competition in favor of the sensory evidence at the cued location for making a localization decision on each trial. To further demonstrate that the m-ADC spatial bias is linked to attention, rather than signal expectation, we show a close relationship between m-ADC bias and a novel measure of spatial attention bias: differential risk curvature (lines 888 - 988 and Figure 5, main text).

