Supplementary Material for:

The conceptual foundations of Network-Based Diffusion Analysis: choosing networks and interpreting results

Will Hoppitt

Further details on simulations

Testing the validity of an NBDA using a dynamic observation network

To test the performance of observation networks I simulated data from the model specified in Eqn. 3. I set $p_{learn:ij} = [0,0.05,0.1,0.15.0.2]$ and $B_j = 1$ in arbitrary time units, with performances of behaviour following a Poisson process (random in time). I set $p_{obs:ij} = 0.15$ to model a population with no social structure. I allowed a rate of asocial learning of 0.2, and a population size of 100 individuals. I recorded the order of learning and built a dynamic network, $a_{ij}(t)$, giving the number of times *i* has observed *j* prior to time *t*. I ran an NBDA (OADA and continuous TADA variants), recording whether social transmission was detected at the 5% significance level in a likelihood ratio test (LRT), and the proportion of events that were estimated to be due to social transmission, with 95% confidence intervals. I also recorded the 'real' proportion of events that occurred by social transmission. I repeated all simulations in this paper with population sizes of 20 and obtained similar results (but with lower statistical power) so in the main text I present only the results for a population size of 100. Results are described in the main text. Code for running these simulations is in the supplementary file "Simulations dynamic obs net error.R"

Testing the effect of error in a dynamic observation network

To simulate use of a proxy measure of observation, I re-ran each analysis varying the probability that each recorded observer was really an observer, $p_{obs|Robs}$= [1, 0.95, 0.9, 0.85, 0.8, 0.6, 0.4, 0.2, 0.15]. The probability each individual was recorded as an observer was set to 0.15. The probability that an individual not recorded as an observer was really an observer, $p_{obs|Nobs}$, was set to constrain the real probability of observation, $p_{obs:ij}$, to 0.15, i.e.

$$p_{obs|Nobs} = p_{obs:ij}\left(1 - p_{obs|Robs}\right)/\left(1 - p_{obs:ij}\right), \qquad \text{Eqn. S1}$$

with $p_{obs:ij} = 0.15$. This means that $p_{obs|Robs} = p_{obs|Nobs} = 0.15$ represents the extreme case where the proxy measure has no useful information about observers, whereas $p_{obs|Robs} = 1$ represents the case where all observers are reliably identified. $p_{learn:ij}$ was set to [0,0.2]. Results are described in the main text. Code for running these simulations is in the supplementary file "Simulations dynamic obs net error.R"

Testing the effect of noise in an association network

I simulated diffusion data from Eqn. 5, with λ_0 arbitrarily set to 0.01, and $s = B_j p_{learn:ij}/\lambda_0$ = [0, 0.1, 0.2, 0.3, 0.4, 0.5, 1, 1.5, 2, 3, 4]. I generated an association network with underlying population structure by grouping the population into 10 sub-groups, with strong associations (~Uniform(0.5,1)) within each sub-group and weaker associations (~Uniform(0,0.1)) among individuals from different sub-groups. I then simulated the diffusion as a Poisson process with rates determined by Eqn. 4. I then simulated an association network as 'recorded' by the researcher, by adding random noise to a_{ij}, taken from a normal distribution with mean 0, and SD= [0, 0.05, 0.1, 0.15, 0.2, 0.3]. Values <0 were capped at 0 and values >1 were capped at 1. I used these networks to run an NBDA (OADA and continuous TADA variants) for each level of random noise. I recorded whether social transmission was detected at the 5% significance level in a likelihood ratio test (LRT), the estimated value of *s*, with 95% confidence intervals, and whether the true value of s lay within the 95% confidence intervals. I repeated the simulation 1000 times for each value of *s*. Results are described in the main text. Code for running these simulations is in the supplementary file "Simulations association net noise.R"

Testing the effect of bias in an association network

I investigate the effect of such bias by repeating the simulations described above but with systematic bias in the 'recorded' network instead of random noise. I did this by transforming each a_{ij} as $\acute{a}_{ij} = a_{ij} + (a_{ij} - 0.5) \times bias$, with bias= [-0.25, -0.1, -0.05, 0, 0.05, 0.1, 0.25]. Here bias< 0 means small network connections are overestimated relative to large ones and bias> 0 means large network connections are overestimated relative to small ones. For each simulated diffusion, I re-ran the NBDA with each level of bias included in the network. Results are described in the main text. Code for running these simulations is in the supplementary file "Simulations association net bias.R"

Testing the use of transmission weights to account for variability in B_j

The questions remain as to a) whether inclusion of transmission weights increases statistical power by making the model more realistic; and b) whether an NBDA is still valid if B_j varies, but transmission weights are not available. I addressed these questions by repeating the simulations described above, but with noise added to B_j instead of a_{ij}. This was done by simulating a value for each individual with mean= 0 and SD= [0, 0.1, 0.25, 0.5, 1], and exponentially transforming the value to ensure it was positive. These values were divided by the mean across all individuals and multiplied by 2, to give a consistent average B_j of 2. The rate of transmission from j to i was taken to be sB_ja_{ij}. Since the weighted NBDA estimates p_{learn}/λ_0 and the un-weighted estimates B_jp_{learn}/λ_0, estimates of s for the latter should be 2x that for the former. I recorded Akaike's Information Criterion corrected for sample size (AICc) for both weighted and un-weighted models in addition to the statistics recorded above. Code for running these simulations is in the supplementary file "Simulations association net variation in Bj.R"

Type 1 error remained at 3-5% for both weighted and un-weighted NBDAs, showing that variation in the rate of performance of the behaviour does not result in an increased risk of detecting a spurious social transmission effect, even when the un-weighted model is used (see Fig S4). However, power in OADA was increased slightly by inclusion of transmission weights, especially when there was larger variation in B_j (see Fig. S4a). 95% C.I.s contained the true value of s in approximately 95% of the simulations, for both weighted and un-weighted models (see Fig. S5). As noise in B_j increased, and as p_{learn} increased, the difference in AICc between the weighted and un-weighted models increased in favour of the weighted models (see Fig. S6). This suggests that if transmission weights are available, they should be included in the analysis if they decrease AICc, as this indicates the model is more realistic and may result in better power to detect social transmission. If transmission weights are not available, but a researcher suspects variation in B_j, they can still use an un-weighted NBDA as a valid means to detect and quantify social transmission.

Supplementary Figures

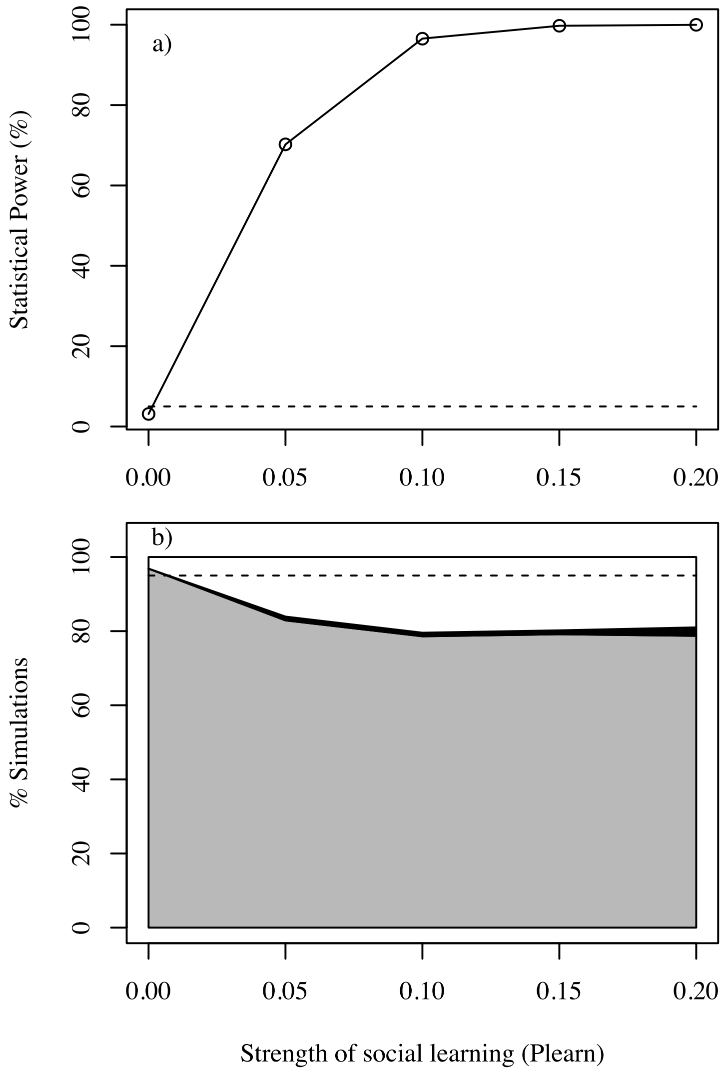

Figure S1. Performance of NBDA (OADA variant) using a dynamic observation network, when applied to data simulated from a realistic model of social transmission. a) Type 1 error rate and statistical power at 5% significance level as a function of p_{learn} (strength of social transmission). The dashed line shows 5%: the expected type 1 error rate (when p_{learn}=0). (b) Proportion of simulations for which the real proportion of social transmission events was within the 95% C.I. (light grey), underestimated (dark grey) or overestimated (white) as a function of p_{learn}. The horizontal dashed line shows 95%: ideally the true value should be within the 95% C.I. 95% of the time. Results were similar for the TADA variant of NBDA.

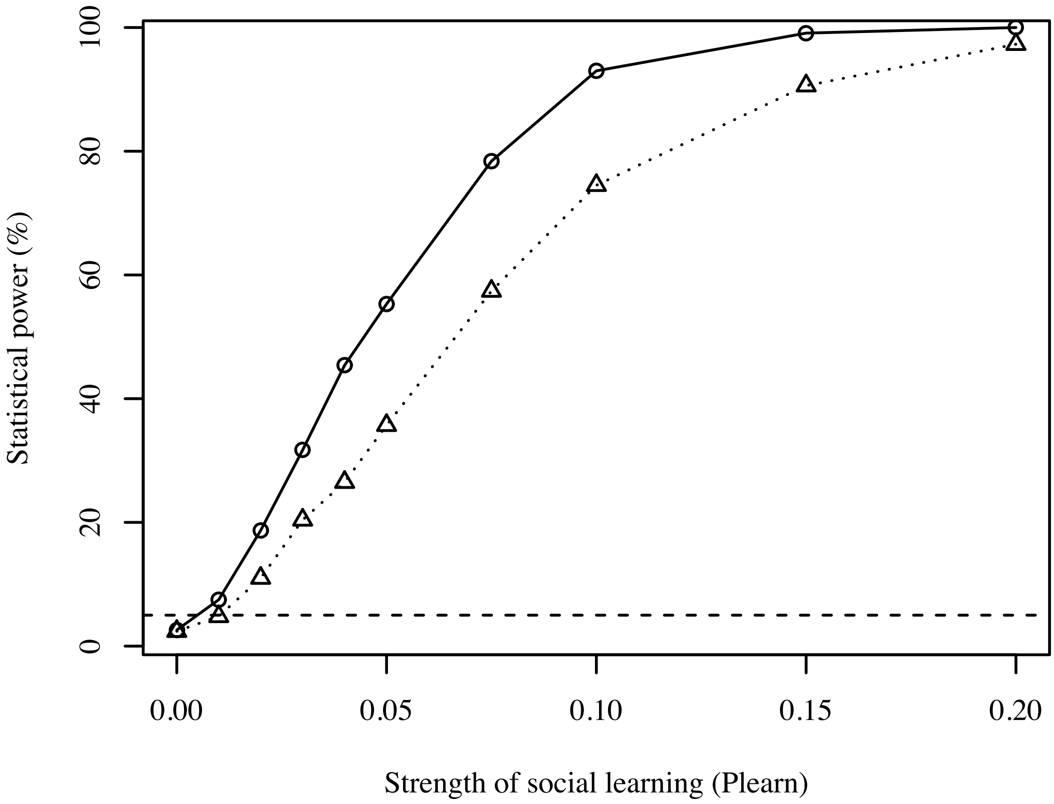

Figure S2. Comparison of the performance of NBDA (OADA variant) when using a dynamic observation network verus an association network when there is underlying social structure (see main text, section 4.2 for details of how the social network was generated). The plot shows type 1 error rate and statistical power at 5% significance level as a function of p_{learn} (strength of social transmission). The dashed line shows 5%: the expected type 1 error rate (when p_{learn}=0). Results were similar for the TADA variant of NBDA.

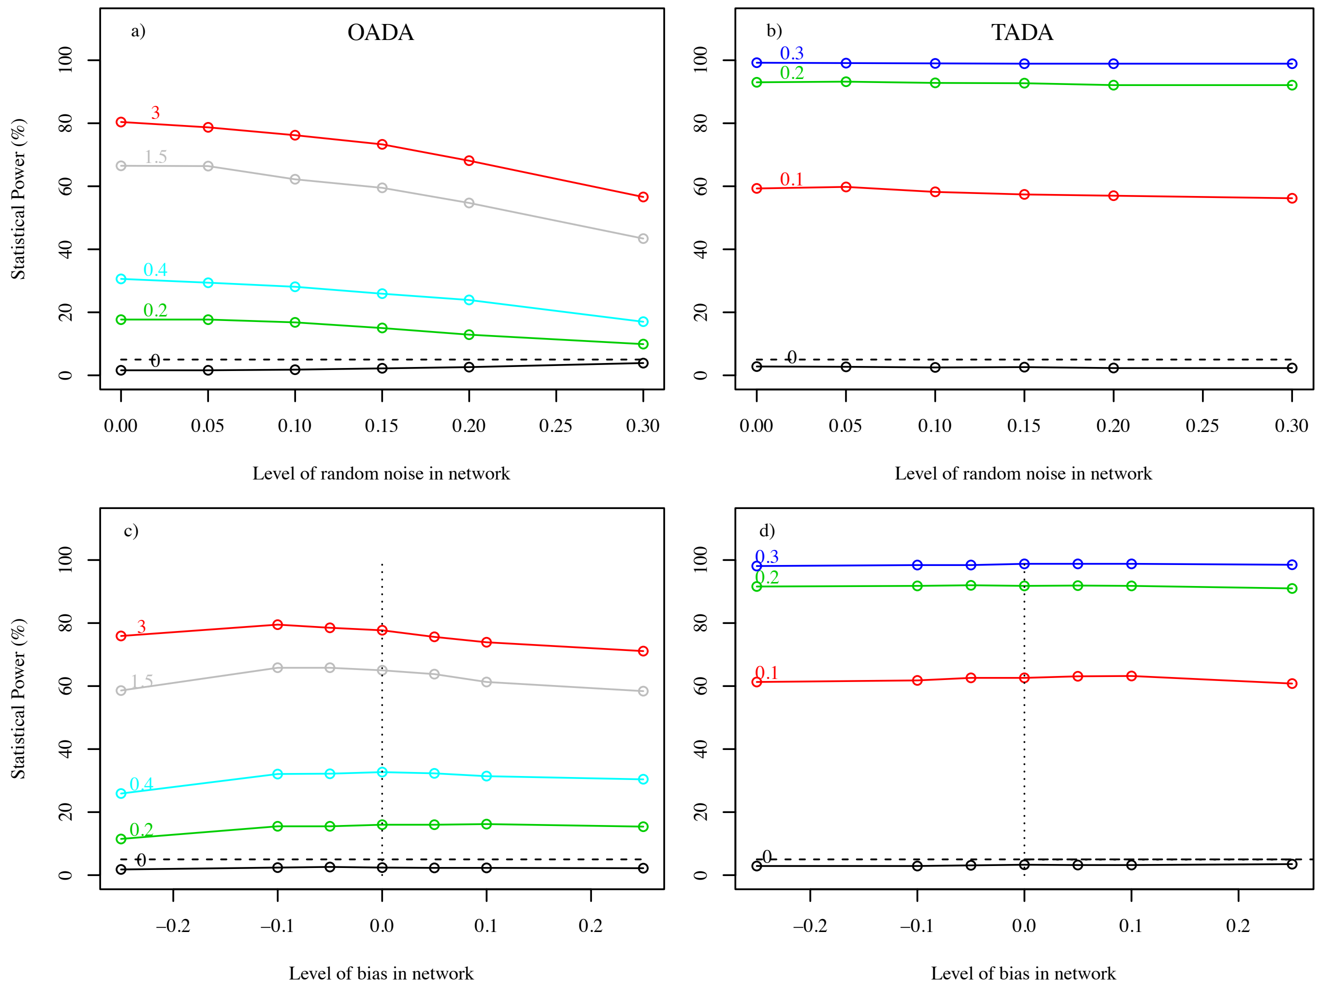

Figure S3. Statistical power of NBDA using an association network when there was noise (top panels) or bias (bottom panels) in the network. The left panels show the performance of the OADA variant of NBDA, the right panels show the performance of the TADA variant. Each plot shows type 1 error rate and statistical power at 5% significance level as a function of noise/bias, for different values of *s* (given by the numbers annotating each line). The dashed line shows 5%: the expected type 1 error rate (when s= 0). Positive bias means that large network connections are overestimated relative to smaller one, whereas negative bias means small network connections are overestimated.

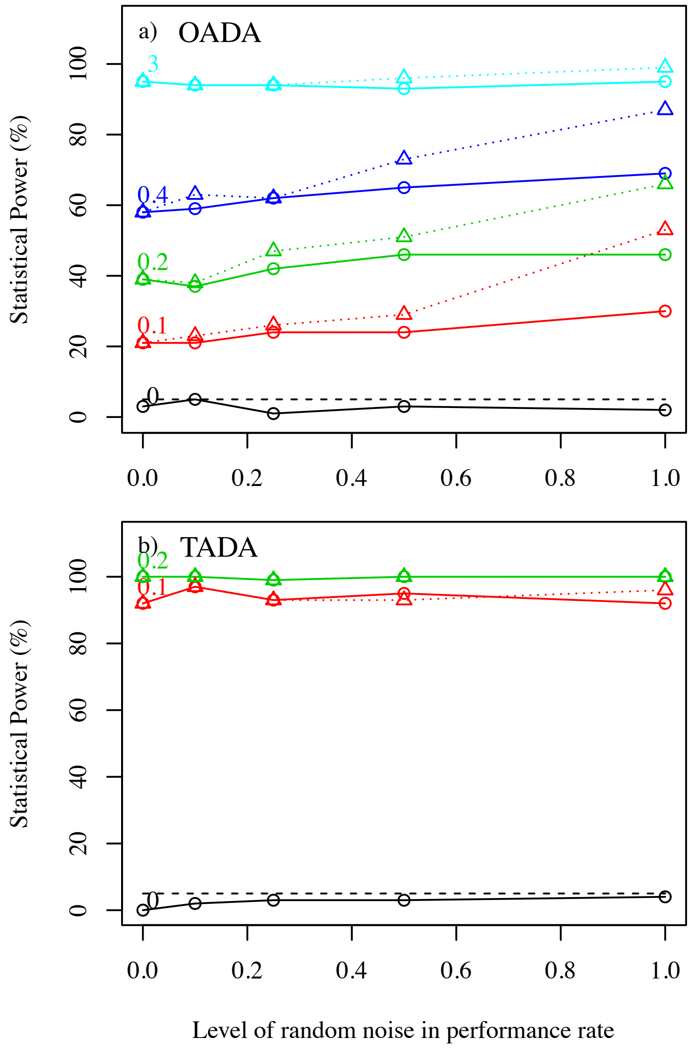

Figure S4. Statistical power of NBDA using an association network when there was variation in performance rate of the target behaviour by those that had learned it. a) Results for OADA variant of NBDA; b) results for TADA variant. Each plot shows type 1 error rate and statistical power at 5% significance level as a function of variation in performance rate (standard deviation on the log scale) for different values of s (given by the numbers annotating each line). The dotted lines give the performance of models including transmission weights allowing for variation in the rate of performance. The dashed line shows 5%: the expected type 1 error rate (when s= 0).

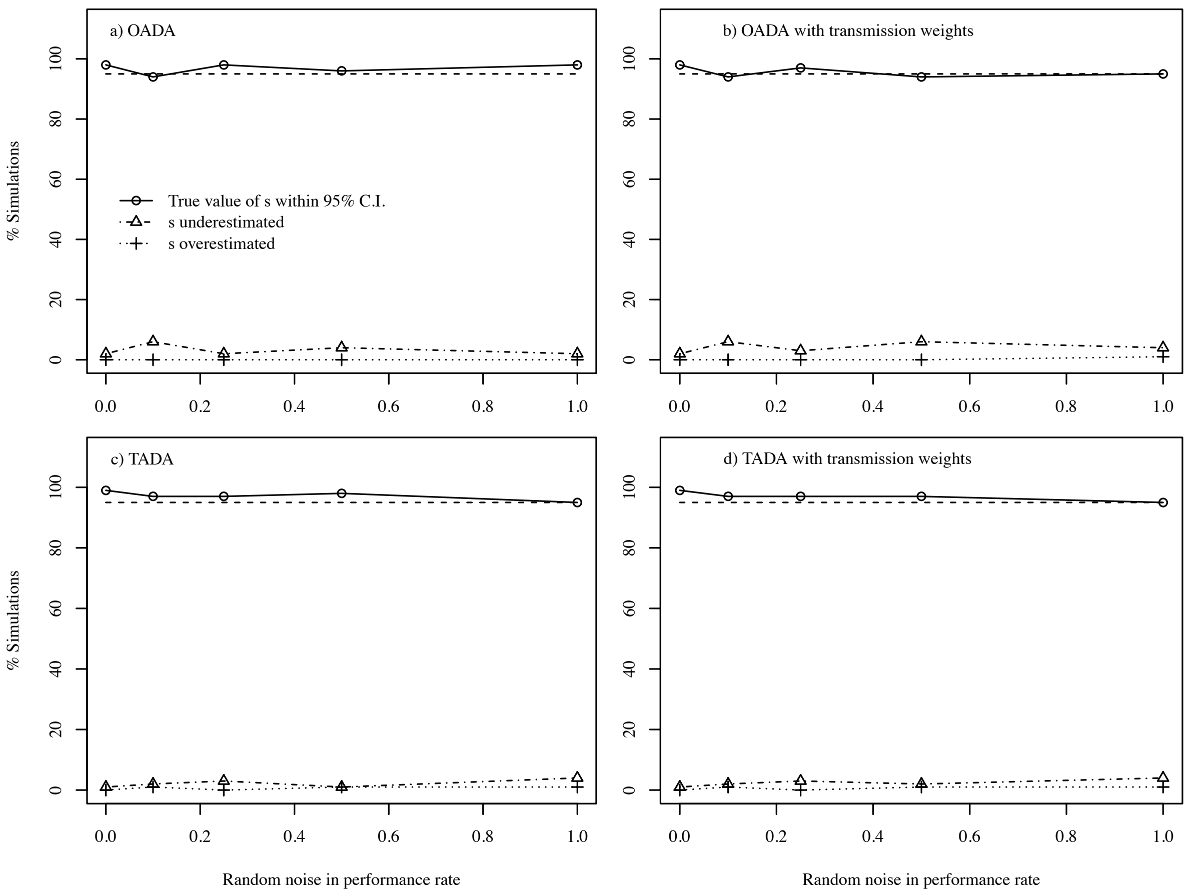

Figure S5. Performance of NBDA using an association network when there was variation in performance rate of the target behaviour by those that had learned it. The plots show the proportion of simulations in which the true value of the *s* parameter (giving the strength of social transmission) was within the 95% C.I. (circles), overestimated (crosses) or underestimated (triangles). The horizontal dashed line shows 95%: ideally the true value should be within the 95% C.I. 95% of the time. The top panels show the performance of the OADA variant of NBDA, the bottom panels show the performance of the TADA variant. The right panels show the performance of models including transmission weights allowing for varaition in rate of performance, the left panels show the performance of unweighted models.

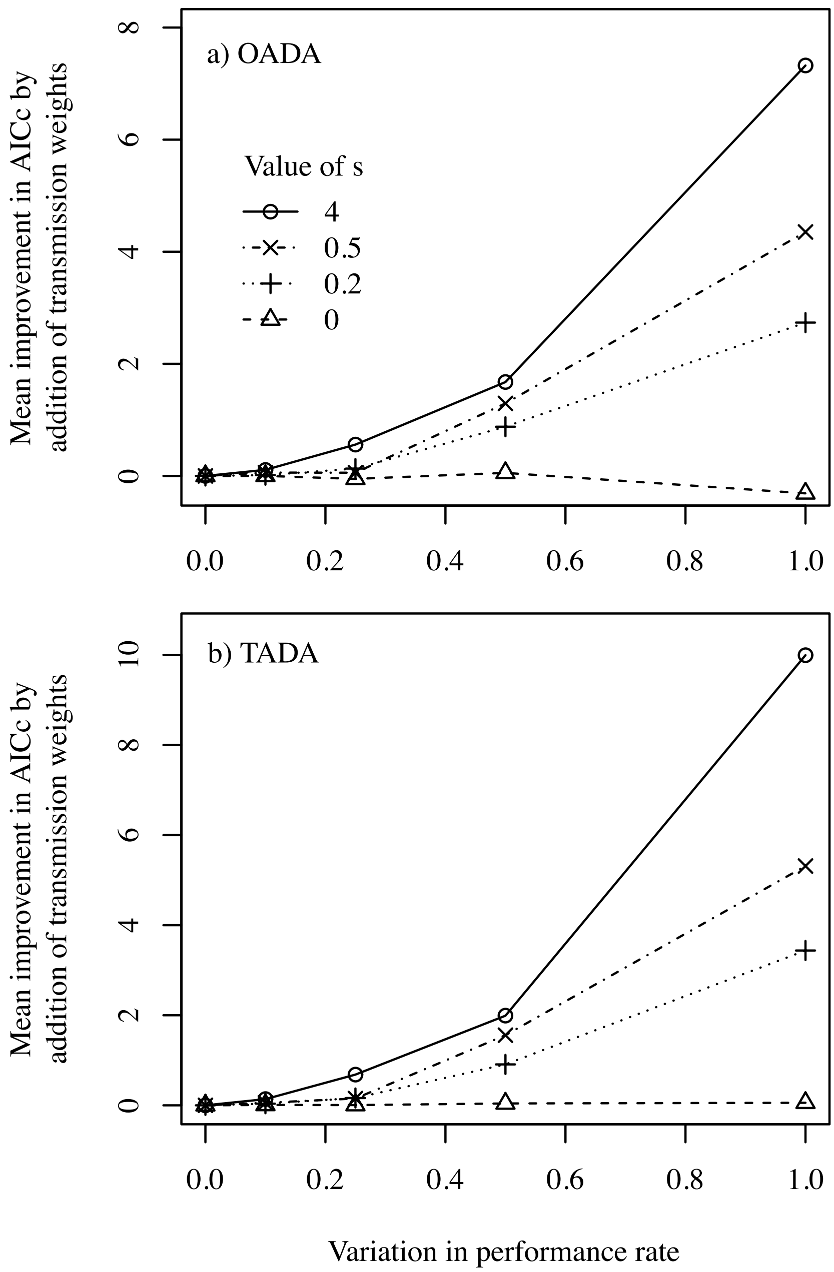

Figure S6. Comparison of predictive power of NBDA models with and without transmission weights as a function of variation in the rate at which individuals perform the target behaviour once they have learned it (standard deviation on the log scale). The plots show the average difference in Akaike's Information Criterion (AICc) in favour of weighted models, for four different values of *s* (strength of social transmission. a) shows the results for the OADA variant of NBDA, b) shows the results for the TADA variant.

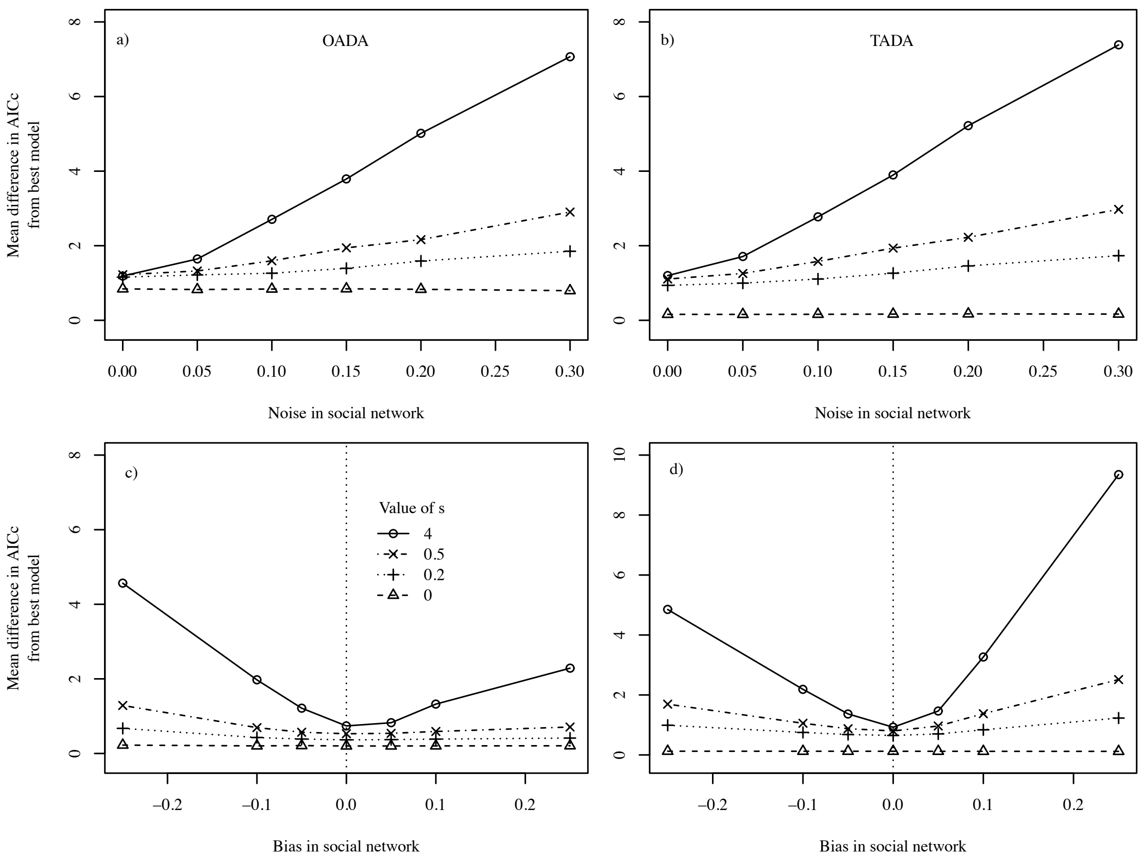

Figure S7. The predictive power of NBDA models using networks that differ from the true underlying network ($\boldsymbol{T_{ij}}$) as a result of noise (a-b) or bias (c-d). The left panels show the performance of the OADA variant of NBDA, the right panels show the performance of the TADA variant. The plots show the average difference in Akaike's Information Criterion (AICc) from the best model, for four different values of *s* (strength of social transmission). The results show that networks that more closely approximate $\boldsymbol{T_{ij}}$ are expected to have a lower AIC.