

Analytical Results
Supplementary Material for Modeling the Spread of Innovation in Wild Birds, Journal of the
Royal Society Interface, 2017

Thomas R. Shultz
Department of Psychology and School of Computer Science
McGill University
thomas.shultz@mcgill.ca

Marcel Montrey
Department of Psychology
McGill University

Lucy M. Aplin
Edward Grey Institute, Department of Zoology
University of Oxford

To complement our simulations and gain additional insight into our results, we construct an analytic model. We represent each of the learning algorithms as a system of differential equations, where the probability of a given conversion from one agent state to another is represented as a rate of change. This analysis is primarily intended to guide intuitions about each algorithm's average behavior in large populations, where multiple interactions occur concurrently.

1. Differential Equations Model

Let x represent the proportion of agents in a population currently using the x solution, y the proportion currently using the y solution, and u the proportion who are currently undecided but open to conversion (i.e., non-Byzantine). To derive a system of differential equations, we take the probability of a given conversion and represent it as a rate of change over time.

In CIU, an agent can only be converted to x if it is undecided. The probability of this occurring is the joint probability of an x agent being selected as the initiator (x) and a u agent being selected as the recipient (u), which is xu . Represented as a derivative with respect to time (i.e., rate of change), this yields

$$\dot{x}_{CIU} = xu \tag{S1}$$

Similarly, the probability of an agent being converted to y is the joint probability of a y agent being selected as the initiator (y) and a u agent being selected as the recipient (u), which is yu . This gives the derivative with respect to time as

$$\dot{y}_{CIU} = yu \tag{S2}$$

Because there are no conversions back to u from a decided state in CIU, the only changes in the proportion of u agents are conversions to x or y . The probability of such a conversion is the joint probability of x or y being selected as the initiator ($x + y$) and a u agent being selected as the recipient (u), which is $xu + yu$. When represented as a derivative with respect to time, we take the negative of this value, because it describes the rate at which u agents are lost:

$$\dot{u}_{CIU} = -xu - yu \tag{S3}$$

AC differs from CIU in that x agents can convert y agents, and vice versa. The probability of an agent being converted to x is therefore the joint probability of an x agent being selected as the

initiator (x) and a y or u agent being selected as the recipient ($u + y$), which is $xu + xy$. However, x agents can now be lost by being converted to y . This occurs whenever a y agent is selected as the initiator (with probability y) and an x agent is selected as the recipient (with probability x), yielding a joint probability yx . When represented as a rate of change, the probability of a conversion to y is subtracted from the probability of a conversion to x :

$$\dot{x}_{AC} = xu + xy - yx \quad (\text{S4})$$

Similar reasoning applies to y agents, who convert u agents to y (with probability yu) and x agents to y (with probability yx), but are themselves converted to x (with probability xy). Represented as a rate of change, this gives

$$\dot{y}_{AC} = yu + yx - xy \quad (\text{S5})$$

Notice however that the xy and yx terms cancel out via addition. On average then, conversions between x and y in AC negate one another. As a result, AC exhibits the same dynamics as CIU:

$$\dot{x}_{AC} = xu \quad (\text{S6})$$

$$\dot{y}_{AC} = yu \quad (\text{S7})$$

$$\dot{u}_{AC} = -xu - yu \quad (\text{S8})$$

In AM, both x and y agents convert u agents to their respective solutions with the same probability as in AC and CIU (xu and yu , respectively). However, instead of losing x and y agents to one another, both solutions now lose agents to an undecided state. The probability of an x agent becoming undecided is given by the joint probability of a y agent being selected as the initiator (y) and an x agent being selected as the recipient (x), which is yx . When represented as a rate of change, losses to u are subtracted from conversions of u agents to x :

$$\dot{x}_{AM} = xu - yx \quad (\text{S9})$$

Similarly, y loses an agent whenever an x agent is selected as the initiator (with probability x) and a y agent the recipient (with probability y), yielding a joint probability xy . These losses are subtracted from the probability of a u agent being converted to y (yu), giving the following rate of change:

$$\dot{y}_{AM} = yu - xy \quad (\text{S10})$$

De-conversions from x and y now provide u agents with a source of growth. The probability of an agent being de-converted is the probability of an x agent or y agent being selected as the initiator while an agent of the opposing solution is the recipient, which is $xy + yx$. Represented as a rate of change, and subtracting out the probability of a u agent being converted to x or y , this gives

$$\dot{u}_{AM} = 2xy - xu - yu \quad (\text{S11})$$

With this analytical model in hand, we revisit each of the phenomena covered in our agent-based simulation results. To allow visual comparisons, we also integrate the system to produce analytical prediction plots for each set of phenomena (using equations S1-S3 for CIU, S6-S8 for AC, and S9-S11 for AM). Seeding values for initial agent states are identical to those in our agent-based simulations, but expressed in terms of proportions. In this math model, each time step represents a whole population and corresponds to numerous simulation cycles, as each time cycle in simulations represents only a single interaction between a pair of agents. In analytical plots, we roughly match the starting and ending conditions to the simulations. In each case, the analytical predictions closely match results in our computer simulations, which in turn closely match results in the bird experiments (1).

2. Diffusion of Knowledge

Our initial knowledge diffusion simulations use only the majority x solution, with $y = 0$. The predicted number of solvers over time is plotted in Figure S1, for treatment and control conditions. As in our simulations, the treatment condition is seeded with 2 knowledgeable solvers while the control condition is seeded at a later start with a single agent who discovers the solution by that time. Notice that without any y states, all three algorithms reduce to the same dynamics: $\dot{x} = xu$ and $\dot{u} = -xu$, just as the algorithms themselves become equivalent in simulation. Hence, one plot suffices for all three algorithms.

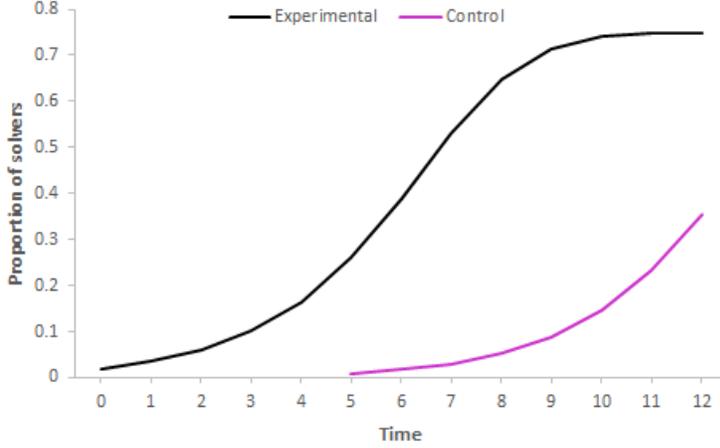


Figure S1. Analytically predicted average number of solvers over time for the AC, CIU, and AM algorithms in the treatment (2 knowledgeable agents) and control (1 individual learner) conditions of the diffusion experiment. Compare to simulation results in Figure 2.

Our simulations indicate that all three algorithms reach consensus at about the same time (Figure 1). The rate of change in u should therefore be equal across algorithms, so long as only one solution is present in the population. Setting $y = 0$ shows that this is indeed the case in the mathematical model:

$$\dot{u}_{AC} = \dot{u}_{CIU} = -xu - yu = -xu \quad (\text{S12})$$

$$\dot{u}_{AM} = 2xy - xu - yu = -xu \quad (\text{S13})$$

On average, consensus emerges at the same time in all three algorithms because the only possible changes are conversions from u to x . Consistent with these analytical predictions, diffusion of knowledge occurs at about the same rate in all three simulations (Figure 2).

Furthermore, statistical testing suggests that simulated diffusion takes a sigmoidal shape. We check analytically if diffusion is indeed sigmoidal. For simplicity, we restrict our analysis to non-Byzantine agents. By assuming $x + y + u = 1$, we derive \dot{x} in terms of x and y alone for AM

$$\dot{x} = x - x^2 - 2xy \quad (\text{S14})$$

and for AC and CIU

$$\dot{x} = x - x^2 - xy \quad (\text{S15})$$

Because only one solution is present, we set $y = 0$, and so for all three algorithms:

$$\dot{x} = x - x^2 \quad (\text{S16})$$

We rewrite this as a recurrence relation

$$x[t + 1] = 2x[t] - x[t]^2 \quad (\text{S17})$$

for which we find the closed-form solution

$$x[t] = 1 - e^{-c2^t} \quad (\text{S18})$$

This is indeed a sigmoid function, where c is a constant that can be calculated from starting conditions

$$c = \ln(1 - x[0]) \tag{S19}$$

3. Naïve Agents Joining an Experienced Group

As in the bird data (1), simulations of the AC, CIU, and AM algorithms show a boosted rate of acquisition when the initial population is seeded with many solvers (Figure 4). This effect is also illustrated using differential equations in Figure S2. Again, because there is only one solution, all three algorithms are expected to produce the same average behavior.

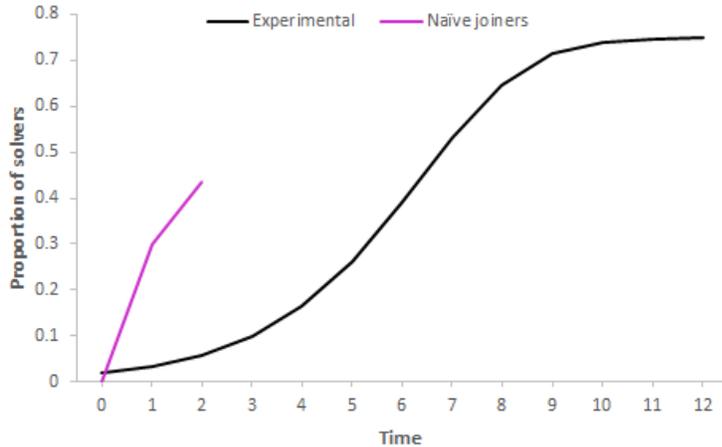


Figure S2. Analytically predicted average number of solvers over time for all three algorithms, for naïve agents joining an experienced group, contrasted with the treatment condition from the diffusion experiment. Compare to simulation results in Figure 4.

4. Groups with Both Majority and Minority Solutions

When simulations are seeded with both solutions, all three algorithms behave similarly with regards to the majority solution (Figure 6). Where they differ is in the minority solution. We illustrate this finding analytically in Figure S4. For the AC and CIU algorithms, the minority (y) solution grows until it stabilizes at its initial proportion. For the AM algorithm, the minority solution grows a bit but then is driven to extinction.

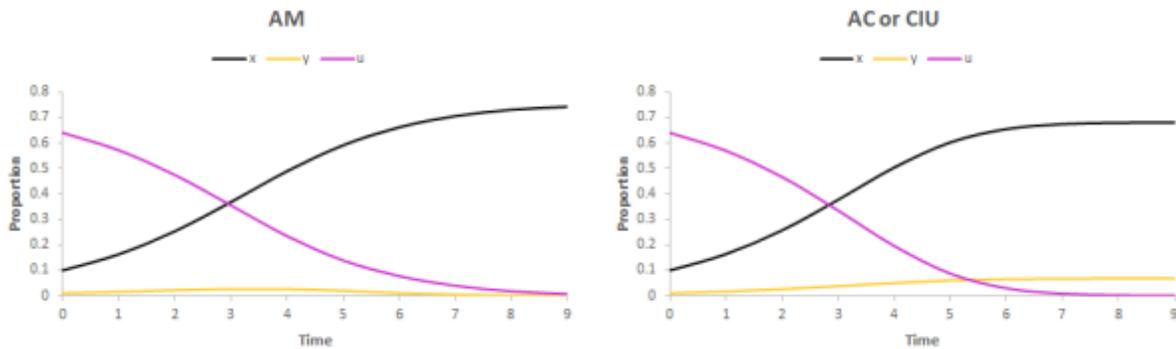


Figure S4. Analytically predicted average proportions of AM and AC or CIU agents over time. The minority solution (y) grows until it stabilizes. Compare to simulation results in Figure 6.

To understand why the minority solution does not grow at the same rate in AM as in AC and CIU, we can compare the rate of change of a solution (x or y) across algorithms. In AM, there is a negative term (xy) that is absent in AC and CIU (compare Equations S14 and S15). This represents more conversions back to u in AM, which slow the growth rate of both majority and minority solutions.

However, the minority solution not only grows at a slower rate in AM, but is also driven to extinction. We illustrate these dynamics with simplex plots (2) in Figures A5 and A6. These plots show the trajectories of change in three proportions (x , y , and u) that always sum to 1. In AC and CIU, the system's dynamics are relatively straightforward: x agents and y agents increase continually as a function of their own frequency and the frequency of u agents, until u agents go extinct (Figure S5). This results in an equilibrium, where the number of x agents being converted to y agents equals the reverse conversions in AC, and where conversions can no longer occur in CIU. This explains the more erratic simulation behaviour of x in AC than in CIU (Figure 7). For both algorithms, the equilibria are: $x = y = 0$ and $u = 0$. Because we assume some solution is present in the population ($x > 0$ or $y > 0$), the only meaningful equilibrium is $u = 0$. The simplex plot for AC and CIU illustrates that both the majority and minority solution can be present in the population once u agents go extinct (i.e., the entire x - y line is an equilibrium).

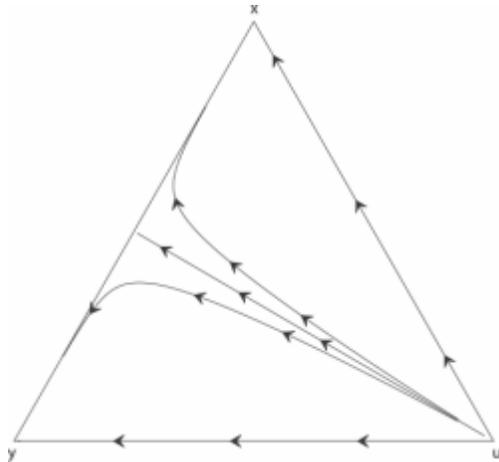


Figure S5. Simplex plot illustrating the dynamics of AC and CIU. Arrows indicate the trajectories along which the system evolves. As x and y recruit u agents, u is driven to extinction. The entire x - y line is an equilibrium, representing whatever mix of minority and majority solutions is reached when u goes extinct.

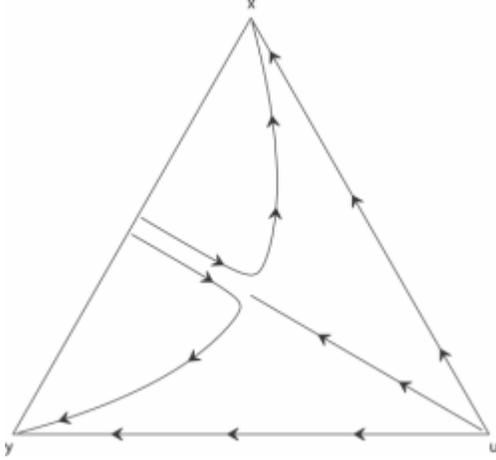


Figure S6. Simplex plot illustrating the dynamics of AM. Arrows indicate the trajectories along which the system evolves. Because u is replenished whenever both minority and majority solutions are in the population, the x - y line is not an equilibrium. Instead, stable equilibria are only reached when the entire group is composed of x or y agents. Note that the equilibrium at $x = y = u$ is unstable.

In AM, the system's dynamics are complicated by the fact the u population is continually replenished by interactions between x agents and y agents. The minority solution increases when there are more u agents than majority agents, but decreases when the reverse is true (Figure S6). Taken together, this means that the system evolves until u and the minority solution both go extinct. This can also be seen by solving for the system's equilibria. These are at: $x = y = 0, u = x = y, u = x = 0$, and $u = y = 0$. We once again ignore the equilibrium at $x = y = 0$, as this is of no interest. We also ignore the equilibrium at $u = x = y$, because it is unstable. This leaves two stable equilibria, both of which involve u and one of the solutions (x or y) being driven to extinction. Thus, unlike AC or CIU, AM requires one solution (usually the minority) to go extinct for a stable equilibrium to be reached.

5. Proportion of Seeded Option

When two solutions are present in the population, the proportion of birds using the seeded majority solution increases over time. Simulations indicate that AM, but not AC or CIU, produces this result (Figure 7). Our analytical model makes similar predictions, as shown in Figure S7. Analogous to the simulations and bird experiments, the proportion of the seeded option is here computed as $x / (x + y)$, where x and y are themselves proportions. In AC and CIU, both solutions convert the same population of u agents, which is never replenished, and so the relative proportions remain constant. This is not true of AM, where u agents can be created via interactions between x and y agents. One of the solutions is driven to extinction in AM, and so the proportion of agents using the majority solution increases.

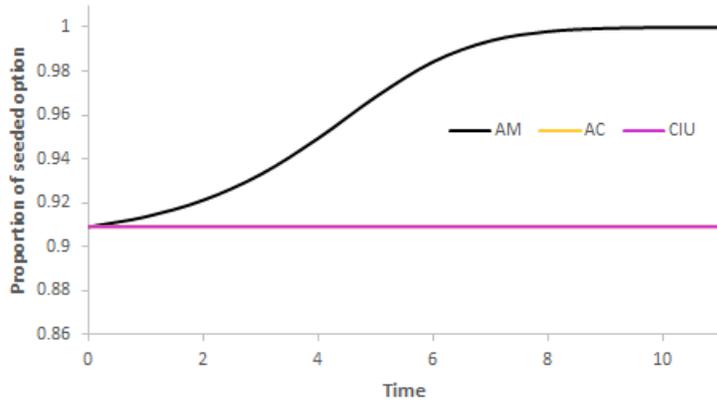


Figure S7. Analytically predicted average proportion of the seeded majority solution over time. AC and CIU overlap. The proportion of majority solvers increases only in AM. Compare with simulation results in Figure 7.

6. Percent Changes in Majority and Minority States

A similar, but more detailed, measure of conformity bias is the percent change in majority and minority solutions. While all three algorithms show a large percent increase in the majority solution, simulations indicate that only AM does not yield a lasting increase in the minority solution (Figure 8). We show this result analytically in Figure S8. As in the simulations, percent increase is computed with Equation 5. The minority solution (here, y) increases only so long as $u > x$, and then decreases when $u < x$. Eventually, the percent change in y reaches minus 100, as the minority solution is driven to extinction.

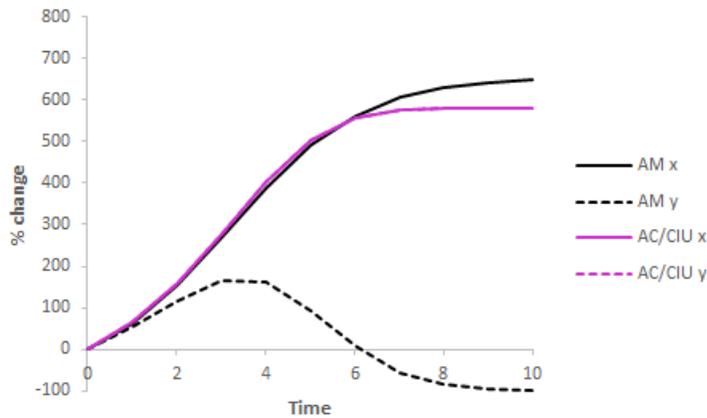


Figure S8. Analytically predicted average percent change in solvers over time, where x is the majority solution and y is the minority solution. Unlike in AC and CIU, the minority solution increases in AM only so long as $u > x$, then decreases until it is driven to extinction. Compare with simulation results in Figure 8 and with novel bird data in Figure 10. In AC and CIU, the x and y plots overlap.

7. Compound States

In the binary agreement model (BAM), an agent can be converted to x from either the undecided state u or the compound state xy . However, an xy agent can also be selected as an initiator. Consequently, an xy agent can convert a u agent to x , given that it exhibits the x solution, which occurs with probability $.5$. An xy agent can also convert another xy agent to x , given the same probability. Meanwhile, x agents continue to be lost to interactions with y agents, though they now enter the compound state xy rather than becoming undecided. Finally, x agents may now be lost to agents in the compound xy state, if the latter exhibits the y solution, which occurs with probability $.5$. Putting this all together as a rate of change in x , where $[xy]$ denotes the proportion of agents in the compound state:

$$\dot{x} = xu + x[xy] + .5[xy]u + .5[xy][xy] - yx - .5[xy]x \quad (\text{S20})$$

Similar reasoning applies to y agents, whose rate of change is given by:

$$\dot{y} = yu + y[xy] + .5[xy]u + .5[xy][xy] - xy - .5[xy]y \quad (\text{S21})$$

Unlike AM, undecided agents no longer have a source of growth. They lose agents both to interactions with x agents and y agents. However, unlike in AC and CIU, undecided agents are also lost to interactions with xy agents. These losses, when put together, yield a negative rate of growth:

$$\dot{u} = -xu - yu - [xy]u \quad (\text{S22})$$

This brings us to the compound state xy . Agents convert to xy whenever an x and y agent interact. Conversions also occur when an xy interacts with an x or y and happens to exhibit the opposite solution. Finally, agents in the xy state may be lost either to interactions with x or y agents, or to interactions with one another. Represented as a rate of change, this yields:

$$[\dot{xy}] = 2xy + .5[xy]x + .5[xy]y - x[xy] - y[xy] - [xy][xy] \quad (\text{S23})$$

Simplifying, the dynamics of BAM can be described by the following set of differential equations:

$$\dot{x} = xu + .5x[xy] + .5[xy]u + .5[xy]^2 - yx \quad (\text{S24})$$

$$\dot{y} = yu + .5y[xy] + .5[xy]u + .5[xy]^2 - xy \quad (\text{S25})$$

$$[\dot{xy}] = 2xy - .5x[xy] - .5y[xy] - [xy]^2 \quad (\text{S26})$$

$$\dot{u} = -xu - yu - [xy]u \quad (\text{S27})$$

With this system of equations in hand, we can now examine whether BAM exhibits the same end states as AM. Solving for the system's equilibria, we find three of possible interest. The first of these is at $u = 0$, $x = y = [xy]$, where $[xy] > 0$. As in AM, this equilibrium is unstable and of no further interest. The other two equilibria are stable and occur when $u = y = [xy] = 0$, where $x > 0$; and when $u = x = [xy] = 0$, where $y > 0$.

In short, stable equilibria in this system only occur when one of the two solutions, the undecided state, and the intermediate state have all gone extinct. As in AM, the system cannot come to rest until every state other than either x or y has been driven to extinction. This suggests that building consensus is not dependent on the intermediate state being an undecided one. Visual results comparing BAM to AM are provided in Figures S9-S12. Note that, when u is extinct, BAM can be represented by the same simplex plot as AM (Figure S6), except with compound state xy in the place of u .

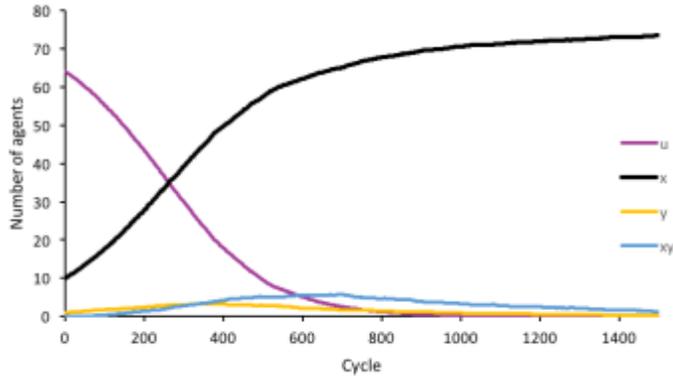


Figure S9. Analytically predicted proportions of BAM agents over time. Both the minority solution (y) and the compound solution (xy) increase a bit and then extinguish. Compare to simulation results in Figure 13 and to AM analytic results in Figure S4. Extinction takes longer with BAM than with AM.

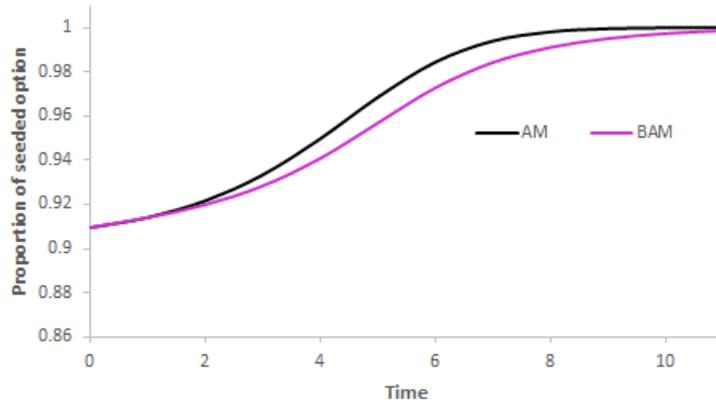


Figure S10. Analytically predicted proportions of agents adopting the majority solution in BAM, compared to AM. Both converge to 1, but BAM takes longer.

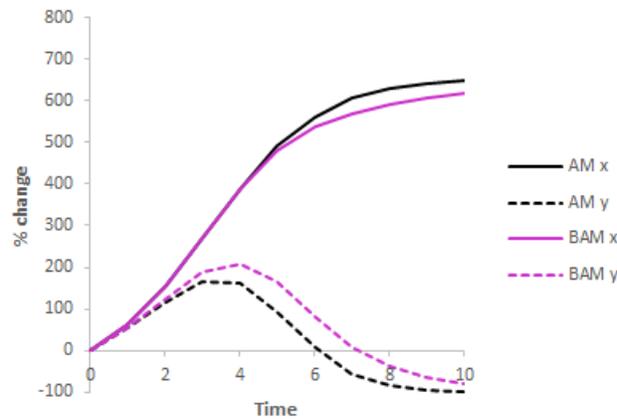


Figure S11. Analytically predicted percent change for majority (x) and minority (y) BAM agents, compared to AM. Results are qualitatively similar, but a bit less extreme for BAM.

8. Conclusions

Our analytical, differential-equations models produce essentially the same results as our agent-based computer simulations based on the various algorithms that we tested: AM, AC, CIU, and BAM. The AM agent-based simulations provide an abstract version of the bird experiments, and our analytical models provide further abstraction of the simulations by ignoring the stochastic process of selecting initiators and recipients for individual interaction, thus implementing a mean field theory approximation. The analytical models also provide further insights into similarities and differences between these learning algorithms and their relative ability to explain innovation diffusion in wild birds.

References

1. Aplin LM, Farine DR, Morand-Ferron J, Cockburn A, Thornton A, Sheldon BC. Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature*. 2015;518:538–541.
2. McElreath R. Baryplot R package. 2008. p. 1–2.