NATURALIZING NATURAL SALIENCE

JACOB VANDRUNEN AND DANIEL A. HERRMANN

Abstract. In the paradigm of Lewis-Skyrms signaling games, the emergence of linguistic conventions is a matter of equilibrium selection. What happens when an equilibrium has "natural salience" – stands out as uniquely attractive to the players? We present two models. We find that the dynamics of natural salience can encourage the learning of more successful signaling conventions in some contexts, and discourage it in others. This reveals ways in which the supposed worst-case scenario – a lack of natural salience – might be better than some cases in which natural salience is present.

> Darwin sees some kind of *natural salience* operating at the origin of language. At that point signals are not conventional, but rather the signal is somehow naturally suited to convey its content. Signaling is then gradually modified by evolution. (Skyrms, 2010, 20)

1. Natural Salience

The Epicureans, Rousseau, and Quine all argued that linguistic conventions could not arise without recursive reliance on prior linguistic conventions. David Lewis (1969) argued the contrary position. He proposed the *signaling game*, in which the fundamental linguistic problem of determining which signal means what can be represented as a game-theoretic coordination problem between a sender and a receiver, with no direct appeal to prior convention. But this reduces the previous problem into one of equilibrium selection: how do agents come to agree on a solution to the coordination problem in the first place?

One of Lewis's suggestions was the mutually-recognized *salience* of some particular signaling system among players. Salience can take on different forms. A paradigmatic example which was influential to Lewis's account is Schelling's idea of a *focal point* (1960). But Schelling focal points rely on mutual expectation of particular strategies among agents. Throughout the present work, we will be focusing on salience in settings in which agents have significantly lower-rationality capacities than is assumed for economic humans. The example from Darwin's *Descent of Man* – the context for the epigraph – was of an "unusually wise ape-like animal" imitating the "growl of a beast of prey" to send a warning signal to conspecifics (quoted in Skyrms 2010, 20). In this case, there is a *natural salience* in making a growl-like noise, mutually acknowledged by the hypothetical simians at the level of instinct. In the literature up to this point on how signaling games may be used to model the evolution and learning of language, little has been said about natural salience. To understand why – and in turn why natural salience might yet

University of California, Irvine; Irvine, CA 92697, USA
E-mail address: jvandrun@uci.edu and daherrma@uci.edu.
Date: Draft of May 27, 2022.
This paper has been accepted for presentation at PSA 2022.

be important – we must first understand how coordination problems are typically solved in the Skyrmsian paradigm.

Skyrms (1996, 2003, 2010) appeals to the symmetry-breaking properties of evolution and learning, arguing that coordination problems might be solved without any appeal to salience whatsoever. Instead, in a way suggested schematically by Democritus and Hume, players might happen upon a successful strategy profile, which would then ultimately be carried to fixation by an adaptive process. The idea is to give a worst-case analysis: even when there are no saliences, conventions can still emerge, as it were, *ex nihilo*.

In the case of signaling, however, the story is not all gumdrops and roses. Under both evolutionary and learning dynamics, populations at equilibrium can end up in signaling conventions with information bottlenecks: relevant states of the world fail to be picked out by signals, while some signals might be synonyms for states already picked out by other signals. These are the *partial pooling equilibria*. Symmetry can be broken without salience, but it is not guaranteed to break perfectly. To that end, LaCroix (2020) reintroduces salience into a Skyrms signaling game via a *salience parameter*. This, he claims, allows for a smooth transition between Lewis and Skyrms. He shows by simulation that his salience games avoid partial pooling equilibria significantly more often than vanilla Skyrms games do, and reach a signaling system significantly more quickly.

LaCroix salience begins to fill in an important part of the story surrounding the evolution of language. But more work is needed to pick out the phenomenon of *natural* salience that Darwin had in mind. Learning agents in LaCroix games use a kind of greedy choice rule to amplify the likelihood of choosing acts which have been most successful in the past. This tracks a kind of salience (in Lewis's categorization, a salience built on *precedent* (1969, 39)) which develops out of the initial learning process, and speeds up convergence. This "subjective" salience, however, does not cleanly map onto Darwin's case, in which a particular signal is appropriated from some existing cue which is objectively correlated with the state of nature.[1] When these naturally correlated cues exist in the environment, under what conditions might agents come to recognize them as salient and use them as signals?

In broad outline, the answer we propose here goes like this. There exists some feature of the world which serves as a possible channel over which agents might communicate. This channel is *conditionally correlated* with some state of the world relevant to successful action – that is, if no agents are manipulating it, it is at-least-partially informative about which state of the world obtains, such that an agent might condition her action on it and do better on average than if she had guessed what the relevant state was. An example of this might be the aforementioned predator-prey situation. An ape which takes a growling noise to be indicative of the presence of a lion will do better than one who ignores the (presence or absence of a) growl and simply guesses as to whether or not a lion is nearby. A proto-receiver

[1]Cubitt and Sugden (2003) diverge from Skyrms in their reconstruction of Lewis. For Lewis, they claim, the appeal to salience is not an explanation of how linguistic conventions arise *ex nihilo*. Rather, it is the salient analogy between the states and actions of consecutive coordination games which grounds the maintenance and replication of conventions based on precedent. Thus, while LaCroix games do represent a transition from Lewis to Skyrms on Cubitt and Sugden's account, our present concern is rather to investigate particular considerations within the Skyrmsian paradigm orthogonal to other issues in the debate surrounding salience.

happens upon this feature and learns to condition her actions on it. At some point, a proto-sender who has the ability to manipulate that same feature, happens upon it, and in turn learns to send the signals which will cause the (proto-)receiver to act appropriately.[2] This is a story about the self-assembly of a signaling game, in which one agent learns to take on the role of a receiver in order to do better than chance, and then another learns the role of a sender in order to benefit from the dispositions of the first. But it is also the story of how cues acquire a natural salience by virtue of being initially correlated with some relevant state of the world, and how that salience might make the process of equilibrium selection in the signaling game more reliable.

In the present paper, we study the emergence of natural salience more concretely with two models. In the first (§2), we show how natural salience can arise in the context of the dynamics of the early stages of an evolutionary game. In the second (§3), we give an account of how those early dynamics themselves might emerge out of a game with further degrees of freedom. The conclusion is mixed. Conditional correlation can encourage the learning of more successful signaling conventions in some contexts, and discourage it in others. This reveals ways in which Skyrms's worst-case scenario – a lack of natural salience – might be *better* than some cases in which natural salience is present.

2. Conditionally Correlated Channels

A traditional $N \times N \times N$ Lewis-Skyrms signaling game has N possible states of nature, N possible signals, and N possible acts, each one appropriate to a different state of nature. On a given round of play, nature chooses a state at random. The *sender* observes the state of nature, and sends a signal based on a set of dispositions. The receiver observes the sender's signal and performs an act based on another set of dispositions. If the act corresponds to the state of nature, the players receive a payoff. Otherwise, they do not. Players' dispositions are modified over time under an adaptive dynamics. Of interest is whether or not the players will coordinate on a bijective map of states of nature to acts. A perfectly successful map of states to signals to acts is called a *signaling system*.

One such adaptive process is characterized by the *simple reinforcement learning* dynamics.[3] The sender has N urns, each corresponding to a state of nature. The sender draws from the corresponding urn after observing nature. The urn has balls of N different colors, each corresponding to a possible signal. The sender sends the signal corresponding to the ball drawn. The receiver in turn has N urns corresponding to the possible signals, each with balls of N colors corresponding to the possible acts. When the sender and receiver are jointly successful, they return the balls drawn to their respective urns, and add an additional ball of the same color. When they are not successful, they return the balls to the urns without adding any new balls.

[2]Extending the predator-prey example to this case suffers from infelicities. While the ape might learn to imitate the lion's growl when her friend is unaware of the presence of the lion, she cannot manipulate the roar/no-roar channel in the other direction: canceling out the lion's roar. A more salient example of this negative action might be found in the nervous system, in which neurons can inhibit the action of other neurons.

[3]This learning model is motivated both psychologically and by its formal relation to the evolutionary replicator dynamics. See (Skyrms, 2010, ch. 7-8) for background on the model and its application to evolutionary signaling games.

Assuming there is symmetry in the players' initial dispositions, each of the $N!$ possible signaling systems are equally attractive. For the case of simple reinforcement learning, the initial symmetry comes from each urn beginning with one ball of each color. Under simple reinforcement learning, Hu et al. (2011) prove that every signaling system has a positive probability of being attained in the limit. But, as far as we know, it is not guaranteed that the players will arrive at a signaling system in the long run. Barrett (2006) shows by simulation that in the medium run players in signaling games with $N > 2$ sometimes end up in partial pooling equilibria, and that this effect increases with larger N.[4] Natural salience removes this initial symmetry, making some strategies more likely than others. The hypothesis is that natural salience will dry up the basins of attraction for partial pooling equilibria, leading to the more reliable attainment of signaling systems.[5] One way this might be accomplished is by adding extra initial balls of certain colors to certain urns in order to jump-start the learning process favorably. Contrastingly, the model we present here shows how the initial symmetry might be broken via the dynamics of the early game, eliminating the ad hoc stipulation of favorable initial conditions. The model we present in §3 adds further degrees of freedom in an attempt to provide the most general account of the process.

2.1. **Model A.** We consider a modified $N \times N \times N$ Lewis-Skyrms signaling game with $N = 16$ to ensure a large number of viable partial pooling equilibria. Sender and receiver both begin with one ball of each type in their urns, but the sender does not play in the early rounds of the game. For an initial sequence of K rounds, the signaling channel is instead directly correlated with nature in the following way. With probability p, the signal sent maps perfectly onto the state of nature according to one of the $N!$ possible signaling systems (pre-selected at the outset of the game). With probability $p - 1$, the signal sent is chosen uniformly at random. This can be thought of as the receiver imperfectly observing the actionable state of nature via some correlated cue. It falls into the category which Barrett and Skyrms (2017) call a *cue-reading game*, albeit with only partially-informative cues.[6] During the initial rounds, the receiver learns by simple reinforcement. For all rounds after the first K, the signaling channel is manipulated by the sender's play instead. At that point, the sender begins to learn by simple reinforcement as well. In our simulations, we test $p \in [0, 1]$ with increments of 0.1 and $K \in \{10^2, 5 \times 10^2, 10^3, 5 \times 10^3, 10^4\}$, with 10^7 regular rounds following the initial rounds. We run 10^3 simulations of each experimental condition.

2.2. **Results.** Figure 1 shows the results for Model A. The most relevant quantity to consider is the "alignment with nature." This tracks how closely the sender's signaling dispositions align with the signaling system that the channel was correlated with during the initial rounds:

$$(1)\quad \text{Alignment} = \frac{\#\text{ of sender urns with modes corresponding to initial bijection}}{N}.$$

[4]For the $N = 2$ case, Argiento et al. (2009) prove that there is long-run convergence to a signaling system with probability one.

[5]Skyrms writes: "In some cases there may well be natural salience, in which case the amplification of pre-existing inclinations into a full fledged signaling system is that much easier" (2010, 21).

[6]More specifically, the mutual information between the signal s with the state of nature σ in the initial rounds is $I(s; \sigma) = [(N - 1)p + 1] \log[(N - 1)p + 1]$.

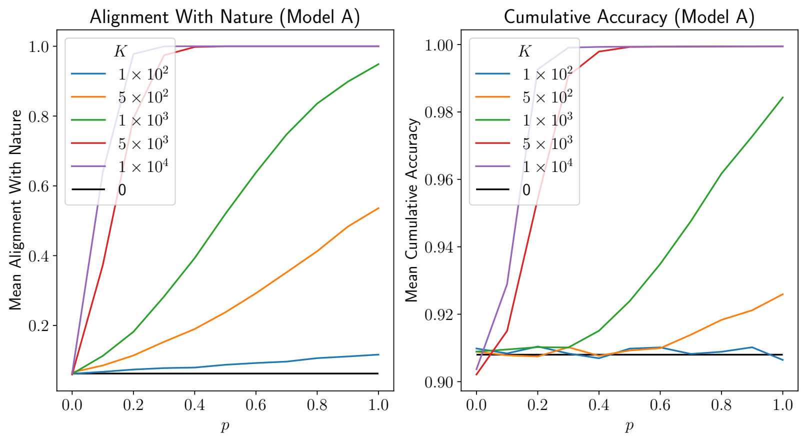

FIGURE 1. Results for Model A showing the alignment of the final signaling convention with the mapping given by the initial correlation of the signaling channel (left), and the cumulative accuracy of the players after 10^7 plays (right). Means are taken out of a sample size of 10^3 simulations for each condition.

The baseline for this quantity when $K = 0$ is 0.062, which is approximately 1/16.[7] The quantity is sensitive both to increasing p and increasing K, but is more sensitive to the latter. With $p = 0.2, K = 10^4$, the mean alignment is 0.978, which is significantly higher than $15/16 \approx 0.938$. With $p = 1.0, K = 10^2$, the mean alignment is only 0.116, which is still less than $2/16 = 0.125$. With $p = 1.0, K = 5 \times 10^2$, however, this quantity jumps to 0.536, which is around $8.5/16 \approx 0.531$. We conclude from this that senders are indeed more likely to send signals corresponding to the initial correlation of the signaling channel with nature when the receiver has had the opportunity to learn using that initial correlation.

The other relevant quantity is the cumulative accuracy which players attain. The baseline at $K = 0$ is 0.908. Once again, this quantity is more sensitive to an increase in K than to an increase in p. With $K = 10^2$, the players on average never do better than the baseline, even when $p = 1.0$. With $p = 1.0, K = 5 \times 10^2$, this jumps slightly to 0.926, and it again jumps to 0.984 when K is increased to 10^3. By comparison, when $K = 5 \times 10^3$, the mean cumulative accuracy reaches 0.991 when $p = 0.3$, and it reaches 0.993 when $p = 0.2, K = 10^4$. But with large K ($K > 10^3$) we observe another effect when there is no correlation of the signaling channel in the initial rounds ($p = 0.0$). With $K = 5 \times 10^3$ the mean cumulative accuracy is below the baseline at 0.902. With $K = 10^4$, it is also below the baseline at 0.904. So, with a large number of initial plays, a signaling channel which is perfectly uncorrelated with nature can be more of a help than a hindrance, but with only a small correlation, it becomes a significant help to successful coordination on a signaling system. On the other hand, with a small number of initial plays, the initial correlation (even when it is high) has much less of an effect overall.

[7]Because in this case there are no initial rounds, the value of p is irrelevant.

2.3. **Discussion.** Model A provides an account of Darwin's natural salience within the framework of Skyrms signaling games. The objective initial correlation of the signaling channel produces a subjective salience of particular acts for particular observed signals on the receiver's part. The sender, in turn, learns to mimic the signals sent during the initial rounds, playing what Barrett and Skyrms (2017) call a *sensory-manipulation game* given the receiver's initially-learned dispositions. This salience exemplifies two properties. First, it is dynamically acquired from initially-symmetric conditions: all that was required was for there to exist a conditionally correlated signaling channel which the receiver happens to be attuned to. Second, as a basis for Lewisian "common knowledge," it is weak – perhaps about the weakest one could go. Lewis writes that salience in general is a weak basis for common knowledge because "the salience of an equilibrium is not a very strong indication that agents will tend to choose it" (Lewis, 1969, 57). In this case, in fact, the salient equilibrium is not even salient to both players. The Skyrmsian adaptive dynamics is still doing the heavy lifting for establishing the signaling convention, but natural salience is getting it off the ground.

The largest remaining gap in the story is how the agents coordinate on the conditionally correlated channel in the first place. How does the receiver become attuned to the relevant feature of the world, and how does the sender come to manipulate that feature? In the next section, we show how this can come about by means of an adaptive process as well. This will also shed light on how the weak salience we are studying can produce failures of coordination as well as successes.

3. Coordinating on the Correlated Channel

In order to study the modular composition of signaling games, Herrmann and VanDrunen (2022) introduce the *attention game*. In an attention game, the receiver has multiple possible signaling channels she can choose to condition her action on. She must learn not only which signal means what, but which partition of the world to interpret as a signal in the first place. In the standard attention game, features of the world which serve as potential signaling channels are correlated with the sender's action. In our modified attention game, we stipulate that signaling channels are conditionally correlated with nature. We also add a process by which a sender can choose which channel to manipulate. This allows us to examine how a sender and receiver might learn to coordinate on a particular signaling channel based on its conditional correlation with nature.

3.1. **Model B.** We consider a modified Herrmann-VanDrunen attention game with N states of nature, N corresponding acts, and 2 possible signaling channels (features) f_0 and f_1, both of which can take on N different states. In addition to two sets of N urns mapping states of nature to signals (one set for each feature), the sender also has a *manipulation urn* with balls of 2 colors corresponding to the different signaling channels he could choose to manipulate. On a round of the game, the sender draws a ball from his manipulation urn, then draws a ball from his signaling urn corresponding to the state of nature and the chosen signaling channel, and sends the corresponding signal.

The signaling channel which the sender does not manipulate takes on a value based on the following algorithm. If it is f_0, it takes on any value uniformly at random – that is, f_0 is perfectly uncorrelated with nature when the sender is not manipulating it. If it is f_1, however, then it takes on a value using the same

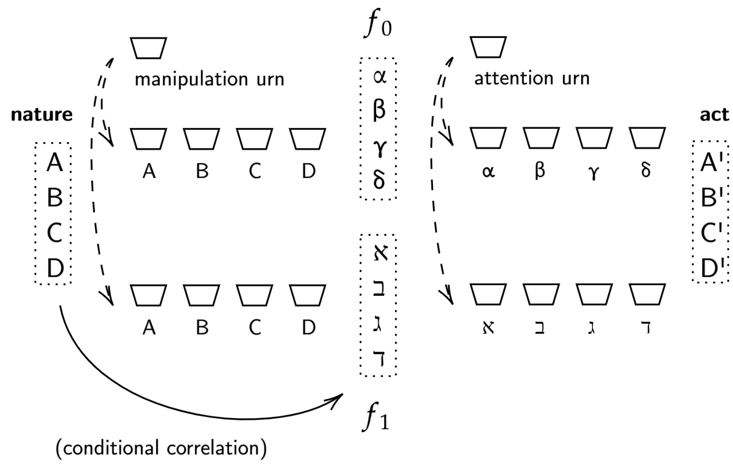

FIGURE 2. Model B.

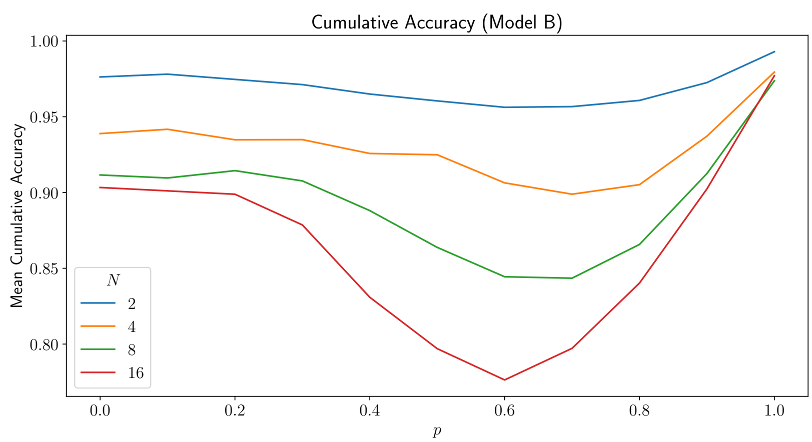

FIGURE 3. Results for Model B showing the cumulative accuracy of the players.

correlation procedure described for the initial rounds of Model A: with probability p it takes on a value corresponding to a fixed bijective map from states of nature onto signals, and with probability $1 - p$ it takes on a value uniformly at random.

The receiver likewise has two sets of N action urns corresponding to the two possible signaling channels, and an *attention urn* with which she selects which channel to condition her act on. Sender and receiver reinforce at the end of the round according to the simple reinforcement dynamics described in §2. The complete model is illustrated in Figure 2. Once again, we run 10^3 simulations for 10^7 rounds of play, but this time there are no initial rounds. We test $p \in [0, 1]$ with increments of 0.1 and $N \in \{2, 4, 8, 16\}$.

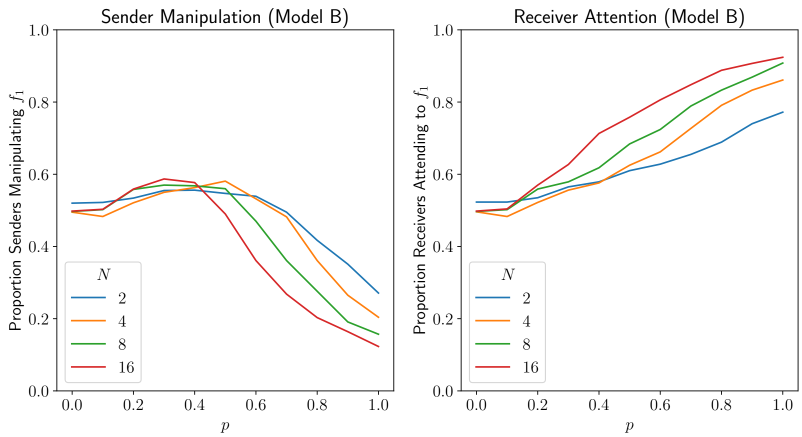

FIGURE 4. Results for Model B showing the proportion of senders (left) and receivers (right) attuned to the conditionally correlated channel f_1.

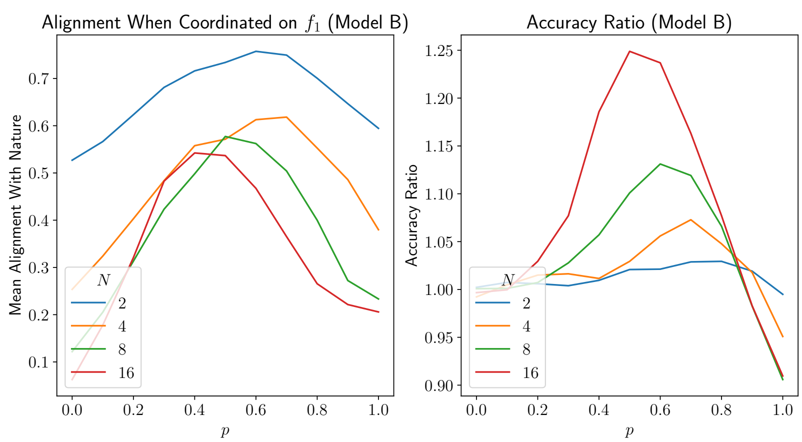

FIGURE 5. Results for Model B showing the alignment of the final signaling convention when both sender and receiver are coordinated on f_1 (left), and the ratio of the accuracy of players when coordinated on f_1 to the accuracy of players under all other outcomes.

3.2. **Results.** Understanding what is going on in Model B is a more subtle undertaking. Consider first the results shown in Figure 3 which parallel the results from Model A. Here, because senders have two sets of signaling dispositions – one for each channel – and only one channel is correlated with a pre-determined bijection, the alignment with nature cannot be calculated absolutely as in Model A. So, we begin by examining the cumulative accuracy. We see that accuracy decreases for middling values of p to a nadir at around $p = 0.6$ (for $N \in \{2, 16\}$) or $p = 0.7$ (for $N \in \{4, 8\}$), before rising to levels significantly above the baseline when $p = 1.0$. Unlike in Model A, accuracy is not an approximately monotonic function of p, but follows a much different pattern. What is the explanation of this?

We begin to get an idea of what's going on by examining the behavior of senders and receivers individually. Figure 4 shows the proportion of senders and receivers who are more likely to manipulate/attend to the conditionally correlated feature f_1. As expected, the receivers display the clear trend of being more likely to attend to f_1 as p increases (the increase is monotonic for all N except $N = 4$, for which it is only monotonic on $p \in [0.1, 1.0]$). Perhaps surprisingly, however, the senders exhibit a more complex trend. After a slight increase in likelihood of manipulating f_1 as p increases, the likelihood declines to a quantity far below the baseline. Both the initial increase and the final decrease is once again larger for greater N. We see, then, that despite an initial attraction of senders to f_1 with low-but-nonzero values of p, *miscoordination* becomes more common as p approaches 1.

The presence of possible miscoordination suggests further analysis to get a clear picture of how the conditional correlation of a channel influences the success of those senders and receivers who actually do coordinate on it. Figure 5 plots two natural statistics for the groups (in each condition) of agents who successfully coordinate on f_1. The first (left) is the alignment with nature. The effect we observe is that, as p increases, alignment goes up, but then it falls back (although still above the baseline at $p = 0.0$) when p reaches 1. This effect is greater with larger N, but peaks earlier: at $N = 2$ the peak value is at $p = 0.6$, at $N = 4$ it is at 0.7, at $N = 8$ it is at $p = 0.5$, and at $N = 16$ it is at 0.4. The second statistic (right) is a ratio of accuracies calculated in the following way. Let G be the group that coordinates on f_1, and G^C its complement. Then,

$$\text{Accuracy Ratio} = \frac{\text{Mean Cumulative Accuracy } (G)}{\text{Mean Cumulative Accuracy } (G^C)}. \tag{2}$$

The accuracy ratio reveals that, although the accuracy of players in aggregate is declining for middling values of p, it is actually increasing for those players who successfully coordinate on f_1. And, although the accuracy of players in aggregate increases toward 1 for high values of p, the accuracy ratio decreases to below 1. Further analysis reveals that this is not because the accuracy of the players who successfully coordinate on f_1 decreases, but because of the increase in accuracy of those who do not. When we examine the accuracy ratio when G is the group who coordinates on f_0 (the conditionally *uncorrelated* channel), a nearly-identical picture emerges, except that neither the proportional increase nor decrease is as extreme.[8] So, it is largely the coordination *simpliciter*, and not the coordination on the conditionally correlated channel in particular, that is having the largest effect on accuracy.

3.3. **Discussion.** The basic conclusion is that the presence of a signaling channel which has a small conditional correlation with nature can improve both the chances of coordination and the effectiveness of signaling when players coordinate on it. This supports our findings from Model A in §2. But this basic conclusion is tempered by the further result that as conditional correlation increases, the chances

[8]More specifically, for increasing N, the accuracy ratio peaks at $1.03, 1.05, 1.10, 1.19$ respectively, and takes on minimum value (at $p = 1.0$) of $1.00, 0.98, 0.95, 0.92$ respectively.

of miscoordination also increase. For middling values of p, this is to the detriment of the success of the players as a whole – the receiver is attracted to the conditionally correlated channel, but without sender intervention is not capable of doing better than the conditional correlation allows. But when p approaches 1, the receiver is capable of near-perfectly-successful action even without sender intervention.

That the receiver will do better without sender intervention is obviously the case at $p = 1.0$, but is less obviously the case even at $p = 0.9$. When $p = 0.9, N = 16$ the greatest possible accuracy (without sender intervention) is $9/10 + (1/10)(1/16) \approx 0.906$ – slightly less than the average accuracy of 0.908 attained in the regular $16 \times 16 \times 16$ Lewis-Skyrms signaling game (Model A, $K = 0$). But initially – before the sender has learned any signaling conventions – the receiver will do better without sender intervention as long as $p > 0$. It is this factor which allows miscoordination to take hold in the learning process.

There are two kinds of symmetry-breaking that must happen in attention games: *within-channel* coordination of which signals mean what, and *between-channel* coordination of which partition of the world to treat as a signal in the first place. Conditional correlation can lead to a natural salience that affects both processes. It can facilitate within-channel coordination, as seen both in Model A and in Model B for middling values of p. It can also facilitate between-channel coordination, as seen for small values of p in Model B. But it can also impede the process of between-channel coordination, as illustrated in Model B with middling-to-high values of p. The sender is rewarded based on whether the receiver is successful, and not whether the receiver is successful *by attending to the channel the sender is manipulating*. So, there is an initial pragmatic cost to the sender's choice of breaking the conditional correlation if the receiver could be doing better than chance without sender intervention.

Skyrms assumes that the case without something like natural salience is the worst case for the emergence of coordinated signaling conventions. We have seen that this is true for within-channel coordination, but Model B shows that there is an even worse case for between-channel coordination. This is when there are possible signals which are *too* "naturally suited" to convey their content. In such cases, it is not worth it for the sender to signal at all.

4. Conclusion

With this initial work, a complex picture of natural salience emerges. The conditional correlation of signaling channels provides a basis for natural salience, which in turn can facilitate the learning of efficient signaling conventions. But when the conditional correlation is too great, it can distract to the detriment of communication. In this case, the same dynamics which drive natural salience can also drive miscoordination. Three avenues for future work naturally present themselves.

First, more work is needed to characterize the full range of possible outcomes in Model B. For example, what happens when both channels are correlated to some degree? We are presently in the process of answering this question.

Second, one might extend the receiver's attentional capacities to allow her to condition her action on signals from multiple channels. In this case, a conditionally correlated channel might not be a distraction but an opportunity for redundancy. This would parallel models in (Barrett, 2021) and (Barrett and VanDrunen, 2022), in which senders and receivers learned to take advantage of redundancy in their signaling conventions.

Finally, in pursuit of more realistic models of natural salience, one might consider cases in which an agent has more influence over certain states of the signaling channel than others. An ape might be able to imitate a predator's growl, but it cannot make the predator un-growl. It may be that choosing to signal is always worthwhile in certain cases of such asymmetrical influence, and that it is paradoxically the sender's fine-grained control over a signaling channel which discourages effective communication.

References

Argiento, R., Pemantle, R., Skyrms, B., and Volkov, S. (2009). Learning to signal: Analysis of a micro-level reinforcement model. *Stochastic processes and their applications*, 119(2):373–390. Publisher: Elsevier.

Barrett, J. A. (2006). Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar.

Barrett, J. A. (2021). Self-Assembling Games and the Evolution of Salience. *The British Journal for the Philosophy of Science*. Publisher: The University of Chicago Press.

Barrett, J. A. and Skyrms, B. (2017). Self-assembling Games. *The British Journal for the Philosophy of Science*, 68(2):329–353. Publisher: The University of Chicago Press.

Barrett, J. A. and VanDrunen, J. (2022). Language Games and the Emergence of Discourse. *Synthese*.

Cubitt, R. P. and Sugden, R. (2003). Common knowledge, salience and convention: A reconstruction of David Lewis' game theory. *Economics & Philosophy*, 19(2):175–210. Publisher: Cambridge University Press.

Darwin, C. (1871). *The Descent of Man, and Selection in Relation to Sex*. John Murray.

Herrmann, D. A. and VanDrunen, J. (2022). Sifting the signal from the noise. *British Journal for the Philosophy of Science*.

Hu, Y., Skyrms, B., and Tarrès, P. (2011). Reinforcement learning in signaling game. *arXiv:1103.5818 [math]*. arXiv: 1103.5818.

LaCroix, T. (2020). On salience and signaling in sender–receiver games: partial pooling, learning, and focal points. *Synthese*, 197(4):1725–1747.

Lewis, D. K. (1969). *Convention*. Harvard University Press, Cambridge.

Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge.

Skyrms, B. (1996). *Evolution of the social contract*. Cambridge University Press.

Skyrms, B. (2003). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge.

Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford University Press.