INFORMATION AND MEANING IN THE EVOLUTION OF COMPOSITIONAL SIGNALS

TRAVIS LACROIX

Abstract. This paper provides a formal treatment of the argument that syntax alone cannot give rise to compositionality in a signalling game context. This conclusion follows from the standard information-theoretic machinery used in the signalling game literature to describe the informational *content* of signals.

Keywords — Signalling games; Information transfer; Communication systems; Semantic meaning; Compositionality; Reflexivity

1. Introduction

The signalling game (Lewis, 1969; Skyrms, 2010) is appealed to as a useful model for explaining the evolution of conventional meanings for arbitrary signals. However, there is a gap between the simple communication systems for which the signalling game model accounts and the linguistic communication systems of *Homo sapiens* (LaCroix, 2020a,b). Attempts to bridge this gap have focused on the evolution of *compositional* signals on the assumption that, because compositionality is an apparently unique feature of human-level linguistic communication systems, explaining of the evolution of compositional signalling would constitute significant progress toward explaining the evolution of *language*.

Several models have been proposed to explain the emergence of compositional signals using the signalling-game framework (Barrett, 2006, 2007, 2009; Franke, 2016; Scott-Phillips and Blythe, 2013; Steinert-Threlkeld, 2016, 2020; Barrett et al., 2020). However, these models often focus on the *syntactic* composition of individual signals. Some researchers have suggested that syntax alone cannot give rise to compositionality (Franke, 2016; Steinert-Threlkeld, 2016; LaCroix, 2020a).

Using an information-theoretic approach to understand the *meanings* of syntactically compositional signals, this paper provides a formal treatment of the argument that syntactic signalling cannot be compositional—at least in the robust

Department of Philosophy, Dalhousie University

E-mail address: `tlacroix@dal.ca`.

Date: Unpublished draft of June, 2022. *Please cite published version, if available.*

This paper has been accepted for presentation at PSA 2022.

sense required for an adequate explanation of the evolution of compositional signals. Section 2 provides some formal background for the main claim of this paper, introducing the signalling game framework (2.1), some concepts from information theory (2.2), and how the latter has been used to elucidate a notion of *semantic information*—or *informational content*—in the context of a signalling game (2.3). In Section 3, I provide a model that explains why syntactic signalling is not compositional. This model shows that the receiver still interprets syntactically compositional signals atomically. Therefore, the evolution of compositional syntax does not give rise to systematicity, which is a requirement for linguistic compositionality. Section 4 concludes.

2. Information and Meaning

This section provides some formal machinery that will be useful for the main argument in Section 3. I introduce the signalling game (2.1) framework and a simple dynamic that is often employed to analyse the evolution of signalling. I then introduce and discuss some formal concepts from information theory (2.2). Finally, I highlight how the latter formalism has been used to provide a concept of semantic meaning in the signalling game framework (2.3).

2.1. **Signalling Games.**

The simplest signalling game is one in which there are two players (called the Sender and Receiver), two states of the world (s_0 and s_1), two possible signals or messages (m_0 and m_1), and two possible actions (a_0 and a_1). This is referred to as a 2×2 signalling game (Skyrms, 2010). The sender observes the state and sends a signal to the receiver. The receiver observes the signal and chooses an action. Both players receive some payoff if they coordinate on states and actions. A formal definition is given in Definition 2.1.[1]

Definition 2.1: *Signalling Game*
Let $\Delta(X)$ be a set of probability distributions over a finite set X. A *Signalling Game* is a tuple,

$$\Sigma = \langle S, M, A, \sigma, \rho, u, P\rangle,$$

where $S = \{s_0, \ldots, s_k\}$ is a set of *states*, $M = \{m_0, \ldots, m_l\}$ is a set of *messages*, $A = \{a_0, \ldots, a_n\}$ is a set of *acts*, with S, M, and A nonempty. $\sigma : S \to \Delta(M)$, is a function from states to a probability distribution over the set of messages that defines a *sender*, $\rho : M \to \Delta(A)$ is a function from messages to a probability distribution over actions that defines a *receiver*, $u : S \times A \to \mathbb{R}$ defines a *utility function*, and $P \in \Delta(S)$ gives a probability distribution over states in S. Finally, σ and ρ have a common *payoff*, given

[1]For further details, see discussion in Huttegger (2007); Steinert-Threlkeld (2016); LaCroix (2020a).

by

$$\pi(\sigma, \rho) = \sum_{s \in S} P(s) \sum_{a \in A} u(s,a) \cdot \left(\sum_{m \in M} \sigma(s)(m) \cdot \rho(m)(a) \right).$$

◊

The payoff, $\pi(\sigma, \rho)$, for a particular combination of sender and receiver strategies gives an expectation of the utilities of state-act pairs (given by $u(s,a)$) weighted by the relative probability of a particular state, provided by $P(S)$. This is referred to as the *communicative success rate* of the strategies σ and ρ. The extensive form of the 2×2 signalling game is given in Figure 1.

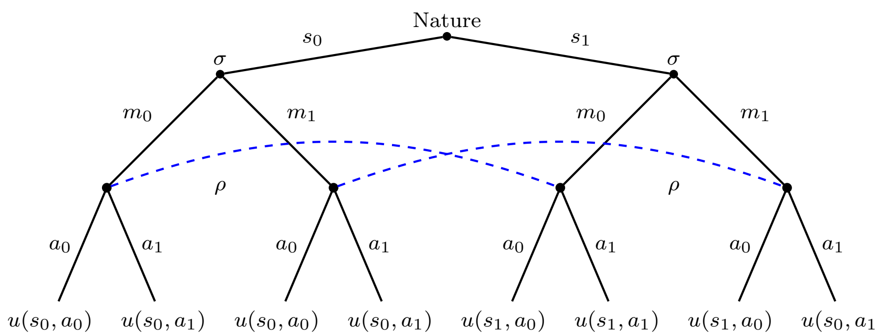

FIGURE 1. The extensive form of the simple 2×2 signalling game. Each node denotes a choice point for a given player, and each branch denotes the possibilities available to her at that point. The dotted lines indicate the receiver's information set.

Following the notation of Steinert-Threlkeld (2016), we can introduce the further definition of an *atomic* signalling game—where states, messages, and actions are equinumerous, the utility function is 1 when the act matches the state and 0 otherwise, and nature is unbiased. See Definition 2.2.

Definition 2.2: *Atomic n-Game*
The *Atomic n-Game* is a signalling game, Σ, with the following restrictions:

(1) $|S| = |M| = |A| = n$,
(2) $u(s_i, a_j) = \delta_{ij}$, where δ_{ij} is the Kronecker delta,

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases},$$

and
(3) $P(s) = \frac{1}{n}$ for all $s \in S$. ◊

A *signalling system* describes a situation in which the sender and receiver strategies lead to perfect coordination and maximal payoff. The atomic 2-game has exactly two signalling systems, shown in Figure 2. Following the formal specification

(A) Signalling system 1 (B) Signalling system 2

FIGURE 2. The two signalling systems of the 2×2 signalling game

in Definition 2.1, the signalling systems of a signalling game can be defined formally as in Definition 2.3.

Definition 2.3: *Signalling Systems*
A signalling system in a signalling game is a pair (σ, ρ) of a sender and receiver that maximises $\pi(\sigma, \rho)$. $\Diamond$

In an evolutionary model, a dynamic explains how sender-receiver strategies change over time. One common dynamic is simple reinforcement learning, described by the following urn-learning metaphor. We assume the sender has urns labelled s_0 and s_1. Similarly, the receiver has urns labelled m_0 and m_1. At the outset, each sender urn is equipped with one ball for each message—labelled m_0 and m_1. Similarly, each receiver urn contains a ball for each action—labelled a_0 and a_1. See Figure 3. In each play, the state is chosen at random. The sender selects a ball

(A) Sender Urns (B) Receiver Urns

FIGURE 3. Simple reinforcement learning model

at random from the urn corresponding to the state of the world and sends that message to the receiver. The receiver then chooses a ball at random from the urn corresponding to the message received. If the action matches the state of the world, then the sender and the receiver both reinforce their behaviour by returning the ball to the urn from which it was chosen and adding another ball of the same type to the urn from which the original ball was chosen. If the action does not match the state, each player returns the drawn ball to the urn from which it was drawn. The game is then repeated for a newly chosen state.

The dynamic shifts strategies to the extent that adding balls to an urn for a successful action shifts the relative probability of picking a ball of that type on a future play of the game. Adding balls to a particular urn changes the conditional probabilities of the sender's signals (conditional on the state) and the receiver's acts (conditional on the signal). Thus, the conditional probabilities of the sender's signals and the receiver's actions change over time, and the players become more likely to perform previously successful actions.

In the next section, I provide some formal background from information theory before describing how a mathematical notion of information has been used to elucidate the content of signals in the signalling game.

2.2. **Shannon Entropy and Relative Entropy.** Shannon entropy measures the degree of randomness in some data set. Higher entropy means a higher degree of randomness, and less entropy means *higher predictability*. Suppose X is a discrete random variable (RV) with alphabet $\mathcal{X}$ and probability mass function $p(x) = p_X(x) = \Pr\{X = x\}, x \in \mathcal{X}$.[2] The definition for Shannon entropy is given in 2.4.

Definition 2.4: *Shannon Entropy*:
The *entropy*, $H(X)$, of a discrete random variable, X, is defined by

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log_b p(x). \tag{1}$$

◊

The base of the logarithm, b, determines the unit of measure. For $b = 2, e, 10$, the unit of information is given by *Bit*, *Nat*, or *Hart*, respectively. We assume that $0 \log 0 = 0$. The entropy of a discrete RV does not depend on the alphabet since it is a function of the *distribution* of X; therefore, it depends solely upon the probabilities underlying this distribution. The entropy of an RV, in general, is characterised as a measure of how much information is required, on average, to describe the RV fully. For example, if we consider the set of states in the atomic 2-game as a discrete RV, $S = \{s_0, s_1\}$ with $p(s_0) = p(s_1) = 1/2$, $H(S)$ tells us that we need, on average, 1 bit of information to describe S.

Relative entropy—also known as Kullback-Leibler (KL) Divergence—is understood as a measure of the similarity of two probability distributions, p and q.

---

[2] In this case, $p(x)$ and $p(y)$ refer to two different RVs—indeed, two different probability mass functions, $p_X(x)$ and $p_Y(y)$. See discussion in Cover and Thomas (2006).

Definition 2.5: *Relative Entropy (Kullback-Leibler Divergence)*:
The *relative entropy*, or the *Kullback-Leibler distance*, between two probability mass functions $p(x)$ and $q(x)$ is defined as

$$
\begin{aligned}
D(p \parallel q) &= \sum_{x \in \mathcal{X}} p(x) \cdot (\log_b p(x) - \log_b q(x)) \\
&= \sum_{x \in \mathcal{X}} p(x) \log_b \frac{p(x)}{q(x)}
\end{aligned} \tag{2}
$$

$\Diamond$

With these definitions in place, the next section describes how KL-divergence has been used to describe the *semantic information* of a signal in a signalling game.

2.3. **Semantic Information and Signalling.** Entropy (Definition 2.4) is not equivalent to, or a measure of, information in the colloquial sense—e.g., the *content* of a signal or message. Since entropy (H) is an average, every message in a repertoire 'has' the same entropy value. However, each message in the repertoire may be *about* different things—i.e., messages may have different *meanings* or *contents*. Thus, the entropy of distinct signals may be identical though the 'information' those signals carry, in the colloquial sense, is different.

Entropy depends upon discrete RVs. However, we note that the elements of the signalling game, described in Definition 2.1, can be understood as a set of discrete RVs, $\{S, M, A\}$. S is a static RV with some probability distribution—uniform, in the atomic case. At a signalling system, the signals are entirely informative, and the receiver has complete information about the state. Therefore, she can act as though she had observed the state directly. The 'key quantity', described by Skyrms (2010), depends upon a comparison between the (conditional) probability that a particular state obtains *given that* a signal was sent, and the likelihood that we are in that state simpliciter:

$$\frac{p(s_i|m_j)}{p(s_i)}$$

We can define the *quantity* of information a signal, m_j, carries (i.e., about a particular state, s_i) as

$$H(m_j) = \log_2 \frac{p(s_i|m_j)}{p(s_i)}.$$

When signals are random, they carry no information. At a signalling system in the atomic 2-game, each signal carries exactly 1 bit of information, corresponding to a reduction of uncertainty from two possible states to one, conditional on the signal.

Skyrms (2010) highlights that signals may carry information about different states. Taking a weighted sum of the probabilities of being in any particular state conditional upon the specific signal, we obtain the following measure of the quantity of information carried by a particular signal, m_j, about the states:

$$I_{states}(m_j) = \sum_{i=1}^{|S|} p(s_i|m_j) \cdot \log_2 \left(\frac{p(s_i|m_j)}{p(s_i)} \right) \tag{3}$$

This is just the KL-Divergence (Definition 2.5) of the two probability distributions $P = p(s|m), Q = p(s)$. Signals can also carry information about the acts:

$$I_{acts}(m_j) = \sum_{i=1}^{|A|} p(a_i|m_j) \cdot \log_2 \left(\frac{p(a_i|m_j)}{p(a_i)} \right) \tag{4}$$

In this context, the relative entropy of a particular signal can be understood as a measure of *additional bits gained* by moving from a prior to a posterior distribution, in a Bayesian sense. Equations 3 and 4 give the *quantity* of information in a signal. On Skyrms' (2010) account, the quantity of information carried by a signal can be used to define the signal's informational *content*; this is a vector that specifies the information that the signal gives about each state. This vector is given by

$$\left\langle \log_2 \left(\frac{p(s_1|m_j)}{p(s_1)} \right), \log_2 \left(\frac{p(s_2|m_j)}{p(s_2)} \right), \cdots, \log_2 \left(\frac{p(s_n|m_j)}{p(s_n)} \right) \right\rangle \tag{5}$$

for the content about the states of a particular signal, m_j.

Suppose there are four initially equiprobable states. In this case, the informational content about the states of each signal is given by the following vectors.

$$\begin{aligned} I(m_1) &= \langle 0, 0, 0, 0 \rangle \\ I(m_2) &= \langle 0, 0, 0, 0 \rangle \\ I(m_3) &= \langle 0, 0, 0, 0 \rangle \\ I(m_4) &= \langle 0, 0, 0, 0 \rangle \end{aligned} \tag{6}$$

None of the signals carries any information about the states, so their content is empty everywhere. If we further suppose that the sender and receiver evolve to a signalling system where signal i is sent only in state i, then the informational content of each signal at that signalling system is given by the following vectors.

$$\begin{aligned} I(m_1) &= \langle 2, -\infty, -\infty, -\infty \rangle \\ I(m_2) &= \langle -\infty, 2, -\infty, -\infty \rangle \\ I(m_3) &= \langle -\infty, -\infty, 2, -\infty \rangle \\ I(m_4) &= \langle -\infty, -\infty, -\infty, 2 \rangle \end{aligned} \tag{7}$$

Now, each signal carries precisely 2 bits of information about the state of nature. The $-\infty$ components tell us which signals end up with probability 0, conditional on the states. Skyrms (2010) suggests that the traditional account in the philosophy of language—where the (declarative) content of a signal is a proposition, and a proposition is a set of possible worlds—is *contained* in this richer information-theoretic account of the content of a signal.[3]

With the formal machinery of Sections 2.1, 2.2, and 2.3 in place, we are now able to understand why syntactic signalling cannot be compositional.

3. Measuring Compositionality

Given the formal definition of semantic information discussed in Section 2.3, we can make exact the argument that syntax alone does not give rise to compositionality. This captures the complaints of Franke (2016); Steinert-Threlkeld (2016), that composite signals are interpreted atomically and so cannot be compositional in the sense that they do not capture intuitions about generalisability conditions for compositional signalling.

Suppose we have a 4×4 syntactic signalling game (Barrett, 2006, 2007, 2009), with two senders, σ_A and σ_B, and one receiver, ρ. Each sender can send one of two messages, and the receiver is sensitive to which sender sent which message. Suppose further that the senders and receiver have evolved a signalling system so that each sender's signal partitions nature into two sets—$\{s_0, s_1\}$ and $\{s_2, s_3\}$ for σ_A, and $\{s_0, s_2\}$ and $\{s_1, s_3\}$ for σ_B. The signal combinations determine the state via the intersection of these sets. See Figure 4.

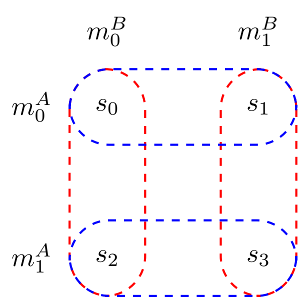

Figure 4. Fully partitioning states via set intersection

[3]Note that some authors have criticised and extended this account. For example, the above characterisation of informational content depends upon how probabilities are *moved* (Skyrms, 2010). Godfrey-Smith (2011) suggests that the content of the signal should say something about *the world* rather than how much the probability of a particular state was moved by the signal's being sent. Birch (2014) highlights that Skyrms' account of informational content falls prey to the *problem of error* (in the same way as the information-theoretic approach to content in Dretske (1981)), when we consider what it means for a signal to have *false* propositional content; see also Fodor (1984); Godfrey-Smith (1989); Crane (2003). Although these insights are theoretically valuable, we will ignore them for now. However, see further discussion in Skyrms and Barrett (2019); Shea et al. (2018); LaCroix (2020a).

We know the maximal entropy of the system from Definition 2.4, given by

$$\begin{aligned} H(S) &= -\sum_{s \in \mathcal{S}} p(s) \log_2 p(s) \\ &= -\log_2 \left(\frac{1}{4}\right) \\ &= 2 \text{ bits.} \end{aligned}$$

Thus, an entirely informative length-two signal carries 2 bits of information because it reduces the possible states from 4 to 1.

We defined the *informational content* of a particular signal with respect to the states as a vector. Therefore, we can give the entire informational content of all of the signals explicitly as a matrix. Each row is the informational *content*, as described in Section 2.3, of a particular message; see Table 1.

TABLE 1. Complete informational content about the states at a signalling system in a 4 × 4 syntactic signalling game.

		States			
		s_0	s_1	s_2	s_3
Informational Content	m_0^A	1	1	$-\infty$	$-\infty$
	m_1^A	$-\infty$	$-\infty$	1	1
	m_0^B	1	$-\infty$	1	$-\infty$
	m_1^B	$-\infty$	1	$-\infty$	1

Further, we can see that a particular state is *wholly determined* by all and only the messages that carry information about that state. Therefore, s_0 is entirely determined by the combination of m_0^A and m_0^B, rather than, e.g., the combination of m_0^A and m_1^B, because the latter carries no information about state 0 when in combination with m_0^A. The syntactic combination of a syntactic length-two signal carries complete information about a particular state; see Table 2.

TABLE 2. Complete informational content in simple signals about the acts at a signalling system in a 4 × 4 syntactic signalling game.

		States			
		s_0	s_1	s_2	s_3
Informational Content	$m_0^A \frown m_0^B$	2	$-\infty$	$-\infty$	$-\infty$
	$m_0^A \frown m_1^B$	$-\infty$	2	$-\infty$	$-\infty$
	$m_1^A \frown m_0^B$	$-\infty$	$-\infty$	2	$-\infty$
	$m_1^A \frown m_1^B$	$-\infty$	$-\infty$	$-\infty$	2

Now, suppose that σ_B spontaneously changes her signal m_0^B to a new signal, $m_?^B$. We might imagine two distinct explanations for this change:

(1) σ_B simply uses a novel signal in lieu of m_0^B.
(2) B forgets the meaning of signal m_0^B.

These two situations might be modelled in various ways, but these will be functionally equivalent under the assumption that the meaning of the *other* signal does not change for σ_B, as we shall see.

In case (1), the arbitrary signal, $m_?^B$, has simply replaced m_0^B; they mean the same thing. Under the urn-learning metaphor described in 2.1, this can be modelled by taking every ball labelled m_0^B in each of the state urns for σ_B and re-labelling them $m_?^B$. This re-labelling does not change that the senders already convened upon a signalling system that perfectly partitions the states of nature; however, the receiver must now learn the meaning of $m_?^B$. This situation is a cue-reading game (Barrett and Skyrms, 2017).

In case (2), σ_B needs to *re*-coordinate so the new signal successfully partitions nature when combined with σ_A's signal. This is similar to a normal signalling context since the σ_B must re-learn when to send this novel signal (given the meanings of all the other signals are fixed, the correct strategy is to send the new signal in the same context as that in which the prior signal was used), and ρ must additionally learn the meaning of the novel signal. 'Forgetting' the meaning of signal m_0^B, as in case (2), can be modelled by 'emptying' all of the balls from the s_0 and s_2 urns for σ_B and adding a novel ball labelled $m_?^B$ to those urns. However, since the meaning of m_1^B is fixed, it follows that even if we 'reset' the urns for s_0 and s_2 with one each of $m_?^B$ and m_1^B, the conditional probability that s_0 obtains given that m_1^B is sent is effectively 0. Therefore, the informational content vectors about the states remain unchanged under either interpretation. This is shown in Table 3.

TABLE 3. Complete informational content about the states at a signalling system in a 4×4 syntactic signalling game with a novel signal identical to the old signal.

		States				
		s_0	s_1	s_2	s_3	
Informational Content	m_0^A	1	1	$-\infty$	$-\infty$	
	m_1^A	$-\infty$	$-\infty$	1	1	
	$m_?^B$	1	$-\infty$	1	$-\infty$	$\leftarrow$ Novel Signal
	m_1^B	$-\infty$	1	$-\infty$	1	

However, the signals also carry information about the acts. Assuming that m_0^B is replaced with $m_?^B$, this can be modelled by effectively throwing out the receiver urns that have a token of m_0^B and *creating* new urns that are labelled identically to the old urns, except with each token of m_0^B replaced with $m_?^B$. Therefore, we can calculate the information that each of the concatenated signals contains about the acts, as in figure 4. That is to say, any composite signal containing a token of the

TABLE 4. Complete informational content in simple signals about the acts at a signalling system in a 4×4 syntactic signalling game.

		Acts			
		a_0	a_1	a_2	a_3
	$m_0^A \frown m_?^B$	0	0	0	0
Informational	$m_0^A \frown m_1^B$	$-\infty$	2	$-\infty$	$-\infty$
Content	$m_1^A \frown m_?^B$	0	0	0	0
	$m_1^A \frown m_1^B$	$-\infty$	$-\infty$	$-\infty$	2

novel signal *carries no information* about the acts.

If the concatenated signals were compositional, this should not happen. Consider that, regardless of the new signal's meaning, m_0^A is only sent for a_0 or a_1. Therefore, the conditional probability that a_2 or a_3 should obtain, given that the receiver has received a length-two string starting with m_0^A, is 0. Similarly, for a_3. The probability of a particular act being appropriate simpliciter is still the chance probability, 0.25. What does this mean for the informational content of the concatenated signal? It is given by

$$\left\langle \log_2 \left(\frac{p(a_0 | m_0^A \frown m_?^B)}{p(a_i)} \right) \right\rangle, \qquad i \in \{1, 2, 3, 4\}.$$

Substituting the values for the conditional and unconditional probabilities, we have

$$\left\langle \log_2 \left(\frac{1/2}{1/4} \right), \log_2 \left(\frac{1/2}{1/4} \right), \log_2 \left(\frac{0}{1/4} \right), \log_2 \left(\frac{0}{1/4} \right) \right\rangle,$$

which resolves to the informational content vector

$$\langle 1, 1, -\infty, -\infty \rangle .$$

However, this makes no sense: m_0^A alone gives us 1 bit of disjunctive information—namely, about $a_0 \vee a_1$. If the ρ interprets the concatenation of m_0^A and $m_?^B$ compositionally, indeed, the second part of the length-two signal would not give her any *new* information regarding the disjunction $a_0 \vee a_1$—namely, unlike before, where the novel signal provides disjunctive information so the union of the two signals uniquely determines a single state. There is no reason why changing the second signal should take information away from the entire composite signal. The receiver interprets the signal as an atomic whole, which provides no information about the act.

This shows that the signals are not *interpreted* compositionally. However, it also highlights that they are compositional for the senders (or, for the states, if you prefer). This is because there is a notion of *independence*—concerning the information that the signal carries about the states—that does not hold for the information that the signal carries about the acts.

We assumed that the states were fixed in the previous example, and only one of the signals changed its meaning. We saw that this has no effect on the informational content of the signal concerning the states, but the receiver counter-intuitively loses the information that should have been contained in the unchanged signal. The same argument holds if, instead of supposing the lexicon is altered, it is merely extended—i.e., if a novel state, a novel signal to represent that state, and a novel action to perform in that state are introduced into the signalling game.

To tell an intuitive story, we might suppose that σ_A sends a verb, and sender B sends a *noun*. Suppose there are two distinct action contexts and two distinct object contexts. Thus, we have the 4×4 syntactic signalling game, as before. Suppose now that a *novel object* context is added to the game. The noun-sender accommodates this by adding a new signal to her lexicon and sending that in the novel context. The receiver must learn what is appropriate given this new signal; however, given that the verb context has not changed, she should gain *some* information. This argument captures precisely the *systematicity* feature of compositional communication: if the receiver knows the meaning of 'pick up x' and the meaning of 'the book', but not the meaning of 'put down x', then she might understand the command 'pick up the book', though she does not understand the meaning of 'put down the book'. Even so, she may still understand that the latter expression has *something* to do with the book.

Though this argument specifically concerned the syntactic signalling model given in Barrett (2006, 2007, 2009), the same considerations apply to the model for combinatorial systems of communication proposed by Scott-Phillips and Blythe (2013). Since they explicitly focus on composition *qua* syntactic composition, this system cannot give rise to a genuine notion of compositional signalling. The same is true for spill-over reinforcement (Franke, 2016). If we add a novel signal to a pre-established signalling system that has evolved via spill-over RL, the receiver loses information in any string containing the novel signal; therefore, Brochhagen (2015) is correct in pointing out that the agents are not sensitive to a generalisation condition for compositionality—namely, the relations between constituent parts are not generalisable.

4. Conclusion

The insights of Section 2.3 suggest that focusing exclusively on syntax in discussing the evolution of compositionality under the signalling-game framework is misguided. Barrett et al. (2020); LaCroix (2022) show how compositional signalling might evolve in a hierarchical signalling game with two basic senders, one executive sender, one basic receiver, and one executive receiver. Here, the executive sender and receiver—called *hierarchical agents*—can learn to influence the behaviour of the basic senders and receiver—called *basic agents*. However, it is not the compositionality of the signals that drives compositionality in this signalling system. Instead, it is the *reflexivity* and *modularity* of the executive sender and receiver that drives compositionality in this context, insofar as the ball that the executive sender chooses *refers* to a component of the base game (LaCroix, 2020a). This can be seen by the fact that the base game (constituted by the base senders and base receiver) is functionally equivalent to the 4×4 syntactic signalling game, which does *not* evolve compositional signalling, as we have seen. These formal insights help to validate the claim from LaCroix (2021) that explanations of the evolution of compositionality are too focused on syntax and will not furnish an evolutionary explanation of the origins of *language*.

References

Barrett, Jeffrey (2006). Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and Evolution of Grammar. *Technical Report, Institute for Mathematical Behavioral Science*.

Barrett, Jeffrey (2007). Dynamic Partitioning and the Conventionality of Kinds. *Philosophy of Science*, 74: 527–546.

Barrett, Jeffrey (2009). The Evolution of Coding in Signaling Games. *Theory and Decision*, 67: 223–237.

Barrett, Jeffrey A., Calvin Cochran, and Brian Skyrms (2020). On the Evolution of Compositional Language. *Philosophy of Science*, 87(5): 910–920.

Barrett, Jeffrey A. and Brian Skyrms (2017). Self-Assembling Games. *The British Journal for the Philosophy of Science*, 68(2): 329–353.

Birch, Jonathan (2014). Propositional Content in Signalling Systems. *Philosophical Studies*, 171(3): 493–512.

Brochhagen, Thomas (2015). Minimal Requirements for Productive Compositional Signaling. In *CogSci*, pages 285–290.

Cover, Thomas M. and Joy A. Thomas (2006). *Elements of Information Theory*. John Wiley & Sons, Hoboken, 2 edition.

Crane, Tim (2003). *The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation*. Routledge, London, 2 edition.

Dretske, Fred (1981). *Knowledge and the Flow of Information*. The MIT Press.

Fodor, Jerry (1984). Semantics, Wisconsin Style. *Synthese*, 59: 231–250.

Franke, Michael (2016). The Evolution of Compositionality in Signaling Games. *Journal of Logic, Language and Information*, 25(3): 355–377.

Godfrey-Smith, Peter (1989). Misinformation. *Canadian Journal of Philosophy*, 19: 533–550.

Godfrey-Smith, Peter (2011). Signals: Evolution, Learning, and Information by Brian Skyrms (Review). *Mind*, 120(480): 1288–1297.

Huttegger, Simon M. (2007). Evolution and the Explanation of Meaning. *Philosophy of Science*, 74: 1–27.

LaCroix, Travis (2020a). *Complex Signals: Reflexivity, Hierarchical Structure, and Modular Composition*. PhD thesis, University of California, Irvine.

LaCroix, Travis (2020b). Evolutionary Explanations of Simple Communication: Signalling Games and Their Models. *Journal for General Philosophy of Science*, 51: 19–43.

LaCroix, Travis (2021). Reflexivity, functional reference, and modularity: Alternative targets for language origins. *Philosophy of Science*, 88(5): 1234–1245.

LaCroix, Travis (2022). Using logic to evolve more logic: Composing logical operators via self-assembly. *British Journal for the Philosophy of Science*, 73(2): 407–437.

Lewis, David (2002/1969). *Convention: A Philosophical Study*. Blackwell, Oxford.

Scott-Phillips, Thomas C. and Richard A. Blythe (2013). Why is Combinatorial Communication Rare in the Natural World, and Why is Language an Exception to this Trend? *Journal of the Royal Society Interface*, 10(88): 1–7.

Shea, Nicholas, Peter Gofrey-Smith, and Rosa Cao (2018). Content in Simple Signalling Systems. *The British Journal for the Philosophy of Science*, 69(4): 1009–1035.

Skyrms, Brian (2010). *Signals: Evolution, Learning, & Information*. Oxford University Press, Oxford.

Skyrms, Brian and Jeffrey A. Barrett (2019). Propositional Content in Signals. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 74: 34–39.

Steinert-Threlkeld, Shane (2016). Compositional Signaling in a Complex World. *Journal of Logic, Language, and Information*, 25(3–4): 379–397.

Steinert-Threlkeld, Shane (2020). Towards the Emergence of Non-trivial Compositionality. *Philosophy of Science*, 87(5): 897–909.