

This is an excellent book. I cannot think of another on this topic that matches its clarity, concision, accessibility, comprehensiveness, and argumentative rigor. I'm quite amazed that Scott-Phillips has managed to combine such seemingly antithetical virtues in one work. The discussion is also admirably honest: Scott-Phillips owns up to the obvious weaknesses with the view and offers strong responses.

I am a little embarrassed and anxious, therefore, because I disagree with most of the main theses of the book. Not all of them. Scott-Phillips persuades me that pure code theories of language origins are hopeless. I am also persuaded that some kind of inference is necessary to explain linguistic communication. There are also persuasive discussions regarding the dearth of combinatorial communication systems in nature, and the role of cultural attractors in the evolution of languages. However, I totally reject the main thesis of the book: that linguistic communication is entirely parasitic on ostensive-inferential communication, where this is understood in terms of metapsychological competence, in particular, the capacity to attribute recursive mental states, via the kinds of inferences that scientists use to infer causal hypotheses from observable data (chapter 1.4).

I detail my reasons for skepticism regarding the metapsychological roots of linguistic communication in the next section. I think that, given Scott-Phillips's background assumption about the only viable alternative theories of language evolution, the case he makes is plausible. However, another problem with the book is the assumption that the code model and Scott-Phillips's version of the ostensive-inferential model exhaust the possibilities. There is a third alternative that has been explored by some philosophers and psychologists: conceiving of language as a shared, normative practice. Scott-Phillips also conflates two questions: what is required for successful linguistic communication, and how human beings meet these requirements. He assumes that the requirements on successful communication – the production of relevant signals and their interpretation as such – can be accomplished only via metapsychological inferences to attributions of recursive mental states. But this is not the only possible mechanism for implementing relevance. It appears to be the only possible mechanism if one focuses entirely on mechanisms endogenous to individual communicators, as Scott-Phillips does. However, when one considers linguistic communication as a shared normative practice, it is possible to specify social mechanisms that obviate the need for metapsychological inferences to attributions of recursive mental states in order to implement relevance. In the last section, I explain this alternative and how it can evade many of the problems I raise for Scott-Phillips's view.

Despite its overall clarity, there is one fundamental topic about which the book could be clearer: what the central thesis is. As near as I can tell, this is the book's central thesis: "the common assumption that the linguistic code makes linguistic communication possible is simply false. Instead, linguistic communication is a type of ostensive-inferential communication, made possible by metapsychology" (page 21). But what is meant by "makes possible"? I think most would agree, certainly in the wake of Scott-Phillips's arguments against the code model, that linguistic codes are not sufficient for linguistic communication. However, even Scott-Phillips must grant that ostensive-inferential communication made possible by metapsychology is also insufficient for linguistic communication. Furthermore, if Scott-Phillips's arguments are correct, then surely both are necessary. So the thesis should be that linguistic communication is made possible both by linguistic codes and by ostensive-inferential communication. This is the view that I will urge below, identifying the different yet equally important roles each plays. I will also argue that ostensive-inferential communication is possible without metapsychology.

### How Is Metapsychology Supposed to Help?

Scott-Phillips marshals a series of compelling examples illustrating to what extent literal meaning

underdetermines speaker meaning. Words, as he repeatedly points out, can be used to mean anything in specific contexts, no matter what their literal meanings are. This is a huge problem for the code model of linguistic communication. But, argues Scott-Phillips, an ostensive-inferential model that puts pragmatics first can avoid this problem. The reason is that interlocutors can infer the specific mental states driving communicative acts in particular contexts, thereby determining precisely what the intended speaker meanings are.

But there are a number of problems with this proposal. First, the book is entirely silent on how a finite, computational system, like the human mind, can successfully infer mental states from observed behaviors and contexts. It is true that literal meaning underdetermines speaker meaning. But it is just as true that observed contexts and behaviors underdetermine the mental responses to and causes of them. This is the well-known problem of holism: any finite observed behavior or circumstance is compatible with an infinite number of distinct sets of mental states, and any finite set of mental states is compatible with an infinite variety of future behaviors, if we make appropriate adjustments to the other mental states that constitute an agent's whole set. Thus, it is unclear how attempting to infer a speaker's mental states from the circumstances and behavior that accompany their utterances can help establish what they mean. Or, at least, Scott-Phillips has not sketched a plausible computational mechanism for accomplishing this feat.

The only attempt to formally model propositional attitude attribution of which I am aware concludes that computationally bounded interpreters are not guaranteed to make accurate attributions of propositional attitudes to a target unless they also model the target's reasoning and belief revision strategies (Alechina & Logan 2010). And perhaps Scott-Phillips is gesturing at this when he likens the process of interpretation to scientific inference. The only problem is that we have no idea how the brain implements such inferences. They are what Fodor calls "isotropic" (1983): any information might be relevant to any inference. And it is hard to model isotropic inferences in a computationally tractable way. This is the so-called "frame problem" of artificial intelligence. I am not saying that this problem is insurmountable, just that Scott-Phillips owes us at least a sketch of how the brain might solve the seemingly intractable problem of interpreting behavior, if he thinks behavioral interpretation can help resolve under-determination of speaker meaning by literal meaning.

But there are more serious problems lurking in the background here. Even if we identify computationally tractable algorithms for scientific inference, it is not clear that these are at work in most cases of communication. Scientific inferences are laborious, conscious, time consuming, and often unreliable, or at least epistemically fraught. But, typically, quotidian linguistic interpretation is fast, automatic, highly reliable, and unconscious. And it appears equally effortless with people we know well as with complete strangers who speak the same language. [1] Surely, we have access to more relevant background information about familiar than about unfamiliar people. So inferring communicative intentions should be easier in the former than in the latter case. But, except in special circumstances where we are not using a conventional language, communication seems equally effortless with familiar as with unfamiliar interlocutors. How can this be the case if inferring communicative intentions is like scientific inference, and hence sensitive to the presence/absence of potentially relevant background information? So even if we identify computational mechanisms capable of implementing slow, laborious scientific reasoning, it seems unlikely that routine quotidian communication makes use of these same mechanisms.

One of the main advantages of code models over ostensive-inferential models of communication is that it is relatively straightforward to implement fast, automatic, efficient, and reliable decoding algorithms. As Scott-Phillips points out in an analogy to mathematics, decoding does not seem like a process that requires inference. One might add that it does not seem isotropic: in a code there is only a limited amount of information that could possibly be relevant. This is one way that some computations avoid the frame problem. And it suggests a way in which codes might help overcome

the prima facie computational intractability of interpretation. If not just linguistic items, but also the non-linguistic behaviors that accompany them, constitute biological or conventional codes for information relevant to interpretation then this might make the task of interpreting communicative acts more tractable. And this sets the stage for another problem with Scott-Phillips's metapsychological model of ostensive-inferential communication: many of the behaviors to which he refers when explaining how metapsychological inferences help determine speaker meaning are, arguably, themselves codes for emotions, reactive attitudes, and other relevant information.

Raised eyebrows, puffed cheeks, smiles, winks, flared nostrils, raised voices, forceful intonations, etc., can all carry an indefinite variety of meanings on their own. It is only in the context of certain culturally variable expectations that they take on determinate significance. Enculturation is the lifelong process of learning such significances from their repeated demonstration in witnessed conversations and other interactions. Interlocutors from the same culture arrive at complementary interpretations of non-linguistic contexts and behaviors only because they have been shaped by similar cultures to situate them in stereotyped scripts or frames that limit the set of viable interpretations. So such non-linguistic contexts and behaviors come to function as codes that simplify the task of interpretation, obviating the need for metapsychology. We need not engage in metapsychological inferences about what other subway riders expect from us, for example, because we have all been socialized to have complementary expectations in such contexts.

As Olivier Morin has pointed out to me, and as I concede below with respect to eye contact, there is much evidence that some such nonlinguistic communicative behaviors have universal significances across cultures. But it does not follow from this that they require science-like inference to be interpreted: there may also be biological codes linking such behaviors to relevant information. For example, on Csibra & Gergely's "natural pedagogy" hypothesis (2006; 2009; 2011), eye contact, "motherese", contingent interaction, and other such low-level behavioral cues carry a very unambiguous meaning for human infants: they signal the imminent demonstration of generic information regarding an object to which the performer is about to refer. I am not sure why infants require science-like inference to interpret such behaviors in context-sensitive ways; their meaning appears remarkably context-insensitive, as with any code.

The lesson here is the following. If interpreting non-linguistic behaviors, like puffed cheeks, etc., in specific contexts is supposed to help resolve under-determination of speaker meaning by literal meaning of utterances, then it must do so in a computationally tractable way; otherwise it cannot explain the speed, automaticity, and reliability of typical communication. But a speaker's mental states seem just as under-determined by such behaviors if all interpreters have to go on in metapsychology: after all, any finite set of behaviors is strictly compatible with any finite set of mental causes, due to holism. This problem can be avoided if non-linguistic behaviors themselves constitute a culturally or biologically determined code for communication-relevant information. But if this is the case, then there is no need for metapsychological inference to help in routine cases of interpretation. Instead, what is needed is mastery of the non-linguistic, embodied, communication codes of particular cultures, and of the species as a whole. [2]

The final problem with Scott-Phillips' appeal to metapsychology is one he addresses forthrightly, but in my view, unsuccessfully. He acknowledges that there are apparent counter-examples to the claim that successful ostensive-inferential communication requires metapsychological inference to recursive mental states: preschoolers and autistic individuals. His response is to cite evidence that metapsychological inference might be much less cognitively demanding than people assume. Unfortunately, the evidence he cites does not establish this. First, the evidence for automatic understanding of indefinitely higher orders of nested propositional attitudes comes from adults, with years of experience interpreting narratives that make explicit such states, using the recursive structures of language. Second, the evidence he cites from implicit false belief tasks with pre-verbal

infants does not support the hypothesis that preverbal infants understand recursive mental states. At best, they show some sensitivity to first-order mental states, like false beliefs, though even this has been questioned, or at least interpreted as a minimal version of the mindreading available to adults (Butterfill & Apperly 2013). In fact, Scott-Phillips claims affinity between his view and Apperly & Butterfill's (2009) minimal mindreading hypothesis, but they explicitly deny that minimal mindreading makes the attribution of recursive mental states possible. So, Scott-Phillips has given us absolutely no reason to think that pre-verbal infants meet the requirements he claims for ostensive-inferential communication. But clearly they are capable of it.

None of this is an indictment of Scott-Phillips' claim that codes are insufficient for explaining communication, or of the importance of inference to successful communication, or of the relevance-theoretic analysis of communication. It is a critique of his assumption that the inferences that make communication possible must be metapsychological, and that the only way the relevance-theoretic analysis of communication can be implemented is through the attribution of recursive mental states. I now turn to an alternative proposal that eschews metapsychology without neglecting the importance of inference and relevance.

### Interpretation without Metapsychology

According to the relevance-theoretic analysis of communication, here is what is necessary for linguistic communication on the interpreter's side: "In order to interpret B's utterance, A searches for an interpretation that optimizes relevance i.e. one that maximizes the positive cognitive effects, and minimizes the processing effort required" (page 59). Here is what is required on the signaler's side: "the Communicative Principle. It states that every ostensive stimulus carries a presumption of its own optimal relevance. What this means is that when signallers produce signals, they produce those signals that maximize the relevance of the stimulus to the audience" (page 60). Neither of these says anything about metapsychology or the attribution of recursive mental states. Those belong to hypotheses about how signalers and interpreters implement relevance, not to a specification of what is required for successful communication. And Scott-Phillips elides this distinction. His full description of the requirement that signalers produce optimally relevant signals includes the following: "... they produce those signals that maximize the relevance of the stimulus to the audience, given both the signaller's goals and preferences, and what the signaller knows about the receiver's goals and preferences" (page 60). But why does the signaler need to know about the receiver's goals and preferences? If there were some other way to produce signals that maximize the relevance of the stimulus to the audience, wouldn't communication be successful, whether signalers knew why it worked or not?

I think Scott-Phillips inherits from both the Gricean and the relevance-theoretic approaches a fixation on atypical and overly intellectualized forms of communication. It is true that there are circumstances in which signalers can maximize relevance only by thinking about what their audience believes about their beliefs and intentions, and audiences can interpret communicative acts only by thinking about what signalers believe about and intend for their beliefs. These usually involve non-standard, jury-rigged communicative signals, especially those targeted at specific audiences in specific circumstances, as when spies need to communicate in ways that others cannot understand. But there is no reason to suppose that this is a good model for most conventional, linguistic communication.

Suppose interlocutors typically think in this way. Conversations are treated as joint activities the goal of which is to share information, and interlocutors are, as a default, simply expected to play the appropriate roles in this joint goal. We have certain evolved behaviors that are highly reliable signals that ensuing behaviors have as their goal the sharing of information that is not public, e.g., eye contact. So, when a signaler wants to share such information, she makes eye contact with her

audience. The audience comes to expect that what will follow aims to make manifest information that she does not have and that is optimally relevant to her goals. This is not because she has inferred the mental states of the signaler; rather, eye contact is a highly reliable indicator that this will happen. The signaler then engages in a performance that, given culturally determined assumptions she shares with her audience, ought to make manifest such information to the audience. Again, she does not speculate about what the interlocutor thinks or knows; she simply makes use of a conventional communicative act which, in that context, it would be rational for anyone sharing her cultural background to interpret as making manifest the relevant information. If all goes well, there is no reason to entertain any hypotheses about what the interlocutor was actually thinking. We need only make assumptions about how people ought to respond to certain behaviors in certain circumstances.

It is important here not to conflate relations like being-informed-by and having-a-goal, with mental states like beliefs and intentions. Mental states have traditionally been conceptualized as theoretical posits aimed at causal, quasi-scientific explanation of observable behavior (Sellars 1997). To conceive of bodily behavior as caused by mental states, one must conceive of the body as animated by an enduring, unobservable object of which they are states: the mind. These mental states must be conceived of as interacting in complex, unobservable ways to yield behavior; otherwise there would be no point in positing them – tracking behavioral patterns would be sufficient. If positing mental states is relevantly like theorizing to unobservable causes in science, then it must support a robust behavioral appearance / mental reality distinction: the hypothesizer of mental states must be capable of conceiving the possibility that qualitatively indistinguishable, counterfactually robust behavioral patterns are caused by different mental states, as when medical diagnosticians conceive the possibility that the same patterns of symptoms are caused by different underlying conditions. But there is absolutely no evidence that infant interpreters, or adults in the heat of seamless, dynamic, communicative interaction, are attributing such explanatory, theoretical constructs. All that matters for quotidian interpretation is reliable behavioral anticipation, not limning the true mental causes of behavior. For these purposes, interpreters need only track what bouts of behavior are likely informed by, and what they aim at. These are relations between targets of interpretation and non-psychological facts that interpreters can represent independently of any interpretive project. Applying Gergely & Csibra's "teleological stance" (2003), interpreters need only think of behaviors as aiming at some observable alteration in the environment, and as guided by information relative to which the behaviors constitute rational means to that goal, whether the information is actual or not. Interpreters can represent such facts without any concept of unobservable, mental causes of behavior. [3]

Many human interactions are like this. Consider games like chess. Here, metapsychology is not necessary unless one's opponent starts making really irrational moves. Otherwise, one can simply use the norms that define chess to anticipate one's opponent's moves. Solving crossword puzzles is similar, and it involves linguistic interpretation, like communication. I know nothing about the psychological profiles of the persons who construct the crossword puzzles I solve every morning. What I do know are certain linguistic conventions, like word spelling, and certain culturally specific facts that underlie allusions, puns, and other kinds of cryptic clues. These are enough to infer the correct solution to the puzzle; no metapsychology is necessary. My suggestion is that everyday conversational interpretation is similar to this. We think about what people ought to know or infer, and because, due to similar socialization, we largely agree on this, we can communicate without thinking about each other's psychologies.

Of course background knowledge about specific individuals helps (although it is not necessary, as we communicate successfully with complete strangers). But there is no reason to construe even such background knowledge metapsychologically. We can look to recent behavior, line of sight, manner,

personal history, appearance, etc., to infer what a potential interlocutor is or is not likely informed of. But this requires no speculation about mental representations. Such informedness can be conceived of in terms of observable relations, like line of sight, to non-psychological facts of which we are aware, or stereotypes regarding certain types of people and their sensitivity to certain types of facts.

If I am picking berries with someone, it goes without saying (or thinking) that they like berries. If they have not seen the patch from which I am currently picking, or if they do not join me, they are clearly uninformed about the edibility of the berries. I automatically make eye contact, given that we are partners in a joint endeavor that includes the joint goal of sharing information about berries. They automatically expect a behavior that will make manifest information they do not have that is maximally relevant to their current goals - picking and eating berries, it just so happens. I slowly and exaggeratedly eat the berries I am picking off the bush. They wonder, what could this mean, given that it is maximally relevant to my goals, including picking and eating berries? Ah, they think, those berries are edible! Where is the attribution of recursive mental states?

Of course this episode can be interpreted as involving the attribution of nested beliefs and intentions, as Scott-Phillips does (chapter 3.4), but a far simpler explanation of how relevance is implemented is possible. Suppose the following claims are true of human language users:

1. They conceive of themselves as obligated to share information relevant to goals their partners in joint endeavors are expected to have.
2. They can tell through various low-level, behavioral cues whether or not their partners have some relevant bit of information.
3. They have at their disposal low-level, stereotyped, behavioral signals that, as a matter of fact, whether they think about it or not, indicate to their partners that sharing of relevant information is imminent (like eye contact).
4. They can follow such cues with performances which, as a matter of fact, whether they think about it or not, are interpreted as and, typically, succeed at making manifest such information.

My claim is that the relevance-theoretic requirements on successful information can be met under such circumstances without sophisticated metapsychology. Both biological and cultural evolution can insure that such mechanisms implement relevance without metapsychology. For example, biological evolution yields stereotyped behavioral cues of imminent sharing of relevant information. Cultural evolution yields capacities for context-sensitive performances, including both linguistic and nonlinguistic components that, in similarly enculturated individuals, are interpreted as, and typically succeed at making manifest relevant information.

## Conclusion

I have much more to say about this highly stimulating and insightful book, but very little room to say it in. For example, I do not think that Scott-Phillips is entirely fair to handicap theories of signaling. There are clearly linguistic phenomena that succumb to this analysis. For example, accents are excellent, costly signals of group membership. They are more costly to produce for people who have not been socialized in a particular linguistic group. In prehistory, analogous forms of communicative "filters" may have been a great way of discriminating between people based on likely trustworthiness or complementary interests. For example, Sosis (2003) proposes that rituals constitute costly signals that can filter reliable cooperation partners from unreliable mimics: mimics will see ritualistic preludes to cooperative endeavors as opportunity costs, while those socialized in a community will see them as routine and hence uncostly. If prehistoric demographics were relevantly analogous to those of contemporary hunter-gatherer societies (Powell et al. 2009; Mellars 2005; Hill et al. 2011), then individuals likely belonged to nested hierarchies of groups composed of other individuals with whom they had varying degrees of affiliation and familiarity (Caporael 2001). Besides immediate family members with whom they interacted daily, they also had to cooperate with

band-mates, members of hunting teams, and members of larger groups like tribes, with whom they interacted rarely. Despite their rarity, such interactions likely constituted some of the most biologically significant ones: e.g., mustering war parties or exogamous pair bonding. Complex communicative rituals would have been an especially important form of costly signaling in such contexts.

It is possible that the apparent excess expressive capacity of human language, made possible by recursive grammar, descends from costly rituals used to filter reliable from unreliable group members. This would make these structural aspects of language analogous to birdsong, the structural complexity of which derives from sexual selection for signals of mate quality (Fitch 2010; Miyagawa et al. 2013). Such content poor yet structurally complex communication systems can avoid the chicken-and-egg problem identified by Scott-Phillips for code-based models of language evolution (chapter 2.3), since capacities for producing structurally complex calls co-evolve with preferences for them, both in populations that use them to advertise sexual quality, like songbirds, and in populations that use them to advertise for cooperative commitment and competence, like, plausibly, prehistoric human populations. This could explain how humans came to have a communicative code that was structurally complex yet semantically impoverished, a “prosodic protolanguage” as Fitch (2010) calls it. Such a code could have then been employed to make ostensive-inferential (yet, if I am right, not metapsychological) communication properly linguistic.

I also share [Olivier Morin's concerns about gossip and reputation](#) as means of stabilizing honest communication: this explanation seems circular. And I do not share Scott-Phillips' skepticism about the significance of understanding shared goals to communication (chapter 3.6), as suggested by what I say above. All of these worries can be traced to what in my view is Scott-Phillips' excessively individualistic orientation. Linguistic communication is seen as a tool that one individual uses to manipulate another, who attempts to insure that the manipulation does not go against her interests. But I think if we conceive of communication as a norm governed practice, evolved through cultural group selection (Henrich 2004) to improve group coordination via practices of information sharing, then we can avoid some of the problems with Scott-Phillips' focus on sophisticated metapsychology. If people typically have complementary goals, similar assumptions about what is relevant and rational, transparent relations to information, and access to low-level signals of imminent information sharing (like eye contact), then successful communication in the relevance-theoretic sense does not require sophisticated metapsychology. Indeed, it is hard to see how increasing group size and complexity in human evolution can have led to increasingly sophisticated metapsychology, as Scott-Phillips assumes (chapter 6.3), given that it would make the problem of attributing mental states increasingly intractable, as people encountered increasing numbers of completely unfamiliar individuals. It is more likely that our ancestors coped with such demographic changes by instituting normative practices that made group-mates more easily interpretable to each other.

## Notes

[1] This is not obvious, and it is not obvious how to test this empirically. Also, there are clearly cases where it is false: inside jokes, etc., that only people who know each other intimately understand. But my point is that, typically, linguistic communication among speakers of the same language seems qualitatively similar whether or not it involves interlocutors that know each other well. Think of asking strangers who are native speakers of one's language for directions, or the time of day, or ordering in a restaurant, or countless other quotidian, communicative interactions we take for granted everyday. In my experience, these do not seem different from analogous interactions with people we know intimately. It is no more difficult to ask a complete stranger for the time of day than it is to ask members of one's family.

[2] I thank Olivier Morin for pointing out to me the evidence that many such embodied

communicative behaviors have universal, culturally invariable significances.

[3] Indeed, at the recent “Modeling Self on Others” workshop (May 2015), at CEU’s Department of Cognitive Science in Budapest, Gyorgy Gergely presented evidence that, as I understand it, infant interpreters do not attribute goals or information access to particular agents: if one agent is taken to have a goal or information access in some context, a different individual taking that agent’s place is automatically assumed to have the same goal or information access. This is hard to reconcile with the idea that infant interpreters are attributing states to an enduring, unobservable object located within a particular agent, i.e., the agent’s mind.

## References

Alechina, N., & Logan, B. (2010). Belief ascription under bounded resources. *Synthese*, 173, 179–197.

Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4), 953–970.

Butterfill, S. A., & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind and Language*, 28(5), 606–637.

Caporael, L. R. (2001). Evolutionary psychology: Toward a unifying theory and a hybrid science. *Annual Review of Psychology*, 52, 607–628.

Csibra, G., & Gergely, G. (2006). Social learning and social cognition: The case for pedagogy. In Y. Munakata & M. H. Johnson (Eds.), *Processes of change in brain and cognitive development*. London: Oxford University Press.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153.

Csibra, G., & Gergely, G. (2011). Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society of London: Series B*, 366, 1149–1157.

Fitch, W. T. (2010). *The evolution of language*. Cambridge: Cambridge University Press.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences*, 7(7), 287–292.

Henrich, J. (2004). Cultural group selection, coevolutionary processes, and largescale cooperation. *Journal of Economic Behavior and Organization*, 53, 3–35.

Hill, K. R., Walker, R. S., et al. (2011). Co-residence patterns in hunter-gatherer societies show unique human social structure. *Science*, 331, 1286–1289.

Mellars, P. (2005). The impossible coincidence. A single-species model for the origins of modern human behavior in Europe. *Evolutionary Anthropology*, 14, 12–27.

Miyagawa S., Berwick R. C., Okanoya K. (2013). The emergence of hierarchical structure in human language. *Front. Psychol.* 4:71 10.3389/fpsyg.2013.00071

Powell, A., Shennan, S., & Thomas, M. G. (2009). Late Pleistocene demography and the appearance

of modern human behavior. *Science*, 324, 1298-1301.

Sellars, W. (1997). *Empiricism and the philosophy of mind*. Cambridge, MA: Harvard University Press.

Sosis, R. (2003). Why aren't we all Hutterites? Costly signaling theory and religious behavior. *Human Nature*, 14, 91-127.