

Cowie Loves a Curry: Negative evidence and nativism

Edward Windle

December 2024

1 Introduction

Why is it ungrammatical to say ‘hurdles horse gate the’? Moreover, why is it that children never make mistakes like these? After all, there are infinitely many ways to construct unnatural and ungrammatical strings, but only a (relatively) limited number of ways to construct grammatical sentences. Children, who are unfamiliar with the rules of grammar in the same way a mature speaker is, do not hypothesise grammars that allow for sentences like the one outlined. This implies a negative aspect of language acquisition: i.e., information on what not to say. The problem is that there is a dearth of this negative evidence in the primary linguistic data (*pld*) of children: evidence as to what is not in the language they are acquiring. Why, then, do children not make mistakes that require negative evidence? Briefly, this is the *logical problem of language acquisition* (hereafter, logical problem). Nativists maintain that domain-specific innate restrictions stop the learner from entertaining hypotheses that require negative evidence to disprove them.

Within her *What’s Within* (1999), Fiona Cowie proposes multiple arguments that aim to dispel the logical problem. In this essay, I show the failure of Cowie’s anti-nativist arguments. She maintains that the nativist, by making idealizations, obfuscates viable sources of negative evidence in a learner’s *pld*. This results in the logical problem making non-nativist alternatives seem untenable, when in reality cognition-general mechanisms can infer negative evidence from structure and statistical data, such as frequency of occurrence compared to frequency of non-occurrence.

In §2, I unpack the characteristics of the theses that constitute enlightened empiricism, weak nativism, and Chomskyan nativism regarding language acquisition. I also detail the logical problem and Cowie’s *reductio ad absurdum*: both of which premise much of the following discussion in §3, where I reject Cowie’s criticisms of the nativist’s idealizations that language acquisition ‘is as if it were’ instantaneous, and the convergence of all learners on the same theory. I will argue that these are forms of minimalist idealisations that help identify and represent otherwise unidentifiable systems that underlie the phenomenon of language acquisition. Finally, in §4 I show Cowie’s arguments for relevant

indirect negative evidence to be insufficient. I reinterpret the positive data that constitutes Cowie’s conception of indirect negative evidence, and show it to be best interpreted on an nativist parameter setting model.

2 The Logical Problem of Language Acquisition

Cowie characterises *weak nativism* concerning language acquisition as the conjunction of the following two theses (1999, p.176):

(DS) DOMAIN SPECIFICITY

Learning a language requires that the learner’s thoughts about language be constrained by principles specific to the linguistic domain.

(I) INNATENESS

The constraints on learners’ thoughts during language-learning are innately encoded.

So, (DS) maintains that a learner must have domain-specific foreknowledge concerning the properties of a language to prevent ‘wild and incorrigible over-generalisations’ (1999, p.209). (I) maintains that this foreknowledge be known innately—endowed in the mind at birth. In concomitance, *Chomskyan Nativists* argue that they constitute the further claim, characterised here by Cowie (1999, p.176):

(U) UNIVERSAL GRAMMAR

The constraints and principles specified in (DS) as being required for language learning are to be identified with the principles characterised in the Universal Grammar.

(U) is a ‘system of categories and mechanisms and constraints shared by all human languages’ (O’Grady et al. 1996, p.734) and is a form of *local nativism*—it concerns the innateness of the particular capacity of language acquisition. (I) is a necessary condition for the existence of (U) – i.e., universal grammar is only universal if it is innate.

The logical problem is a version of the poverty of stimulus argument inspired by the ‘subset sum problem’ in formal learning theory. Here, it relates to the ‘logical structure of the language-learning task’ (Cowie 1999, p.205) and is characterised by the unavailability of negative evidence to language learners. Negative evidence relates to evidence that a string is not in the language. The argument maintains that if a learner’s current grammatical hypothesis (H) is generally consistent with correct grammar (L), but (H) generates a string that is not a string in (L), then the dearth of negative evidence means that the child would have no way of correcting their mistake. Because all children’s grammar ultimately converges on the right theory, their grammar must have two restrictions to prevent such mistakes: it must neither (G) ‘hypothesise a grammar that is “too large”’, or (P) ‘never project *any* hypothesis implying statements of the form “String *s* is a sentence of L” where *s* is not in fact a sentence of L’

(*ibid.* p.209). As such, a theory of language acquisition must explain how this restriction is made.

Both weak and Chomskyan nativists maintain that the problem is solved by (DS) and (I). Firstly, domain-specific foreknowledge prevents overgeneralisations (G). Clearly, by constraining a learner's thoughts about a language to principles specific to the linguistic domains prevents the learner from hypothesising (G). Secondly, it restricts the potential for thoughts about language, thus preventing the learner from ever projecting any hypothesis that would require negative data (P). This point is further ratified if we adopt (U). To unpack how, consider the Principles and Parameters (P&P) model. Principles can be thought of as innate restrictions, premising language specific linguistic parameters that can be set like a switchboard: '[w]hen switches are set one way, we have Swahili; when they are set another way, we have Japanese' (Chomsky 2000, p.3). When learning, a Chomskyan could argue that when parameters are set in ways that allow for possible languages, the learner cannot project strings that require negative evidence to correct. Furthermore, for Chomskyans, adopting (I) guarantees that both (DS) and (U) are present in all humans, in the same way that there is a universal potential for humans to gain teeth.

Cowie's position on the nativist's response is to grant (DS), but reject both (I) and (U). This is *enlightened empiricism*. She accepts the former, maintaining that 'empiricist's requirement that the learner prefer simpler, more general, more elegant (etc.) hypotheses provides little in the way of guidance: too little, surely, to prevent a learner's falling irretrievably into the myriad possible pitfalls revealed by the Logical Problem' (1999, p.209). Essentially, the complexity of language means there are many ways to make mistakes; without the learner having foreknowledge about language, it seems highly likely that she would go seriously wrong. This is fine; the nativist, after all, accept (DS) also. Cowie's arguments against (I) (and thus (U)) are the main course here.

2.1 Cowie's Culinary Case

In assessing the nativist's support for (I), Cowie tries to show the impossibility of this commitment with a *reductio ad absurdum*:

Given some minimal gastronomic experience, virtually all normal humans are guaranteed to acquire 'culinary competence,' the ability to recognise and distinguish a variety of foods from each other and from non-foods. Yet, when learning about curries, say, no one ever systematically informs us that Irish stews, tacos, and quiches—let alone boats and babies and bison—are not curries. There are infinitely many things that might be curries, but that are not curries, and that we have no information about one way or the other. Yet, and despite the dearth of negative evidence, we all manage to converge on the view that a curry is itself and not another thing. We all, that is to say, manage to more or less converge in our intuitive judgements as to what is and what is not an instance of curry.

(1999, p.215)

What she aims to show by this case is that if the restrictions of (DS) necessary to answer the logical problem are innate, then we would require innate, domain-specific faculties of acquisition ‘for all learning involving projection beyond our experience’ (1999, p.215). Sensibly, one must concede that we do not have these mechanisms for learning about curry, just as we do not have them for learning about jungle music. In turn, there must be more *a posteriori* evidence than is let on, rendering (I)–and, in turn, (U)–superfluous conditions for the process of language acquisition.

Let us assume that *reductio* is successful¹. The upshot, then, is that the logical problem misconstrues the process of language-learning. Cowie claims that it does this in four ways: (1) by not accounting for indirect forms of negative evidence in the *pld*; (2) by failing to acknowledge the ‘piecemeal nature of theory construction’; (3) by ignoring the ‘holistic and probabilistic relations of evidence to theories’; and, (4) by assuming a convergence of opinion on the theory when disagreements regarding the theory show that we do not converge fully (1999, p.217). By acknowledging these misconceptions, much more evidence is considered significant to the language learning task; therefore, the problem is tractable without (I) and (U). The logical problem would thus dissolve and not support nativism.

3 Idealization

According to Cowie, the first ramification (2) is a problem because nativists idealise ‘to instantaneous learning [which is] a very unrealistic and misleading [idealization]’ (1999, p.218). By this, she is referring to the nativist claim that the ‘order of presentation of data [to a learner] is irrelevant, so that learning is “as if it were instantaneous,” so as if S_O maps the data onto S_S ’ (Chomsky 1986, p.52; S_O is the ‘initial state’ of the learner; S_S is the ‘steady state’ of knowledge required for the target language). The curry analogy apparently makes clear that language acquisition is not instantaneous, but requires intermediate hypotheses and gradual, piecemeal development. Though Cowie stops her point here, I think the further implication is that, if the above argument is true, the quixotism of nativists makes superfluous the requirement of (U). Think back to the switchboard example: if language acquisition is not ‘as if it were instantaneous’, then explanation of (U) is significantly more difficult, as it would have to account for the processes too (by, for instance, microparameters).

The problem with Cowie’s contention, however, is that the idealization of instantaneity serves a specific explanatory role. Specifically, instantaneity can be understood as a *minimalist idealisation*, which are the products of *minimal models*: models constructed to exclusively identify and represent the minimal and

¹See Laurence & Margolis (2001, pp.262-264) and Fodor (2001) for rejections of the *reductio*. In sum, they amount to the claim that curry is not the domain of interest here, but instead food. For it seems perfectly sensible to postulate innate information about food – primitive markers of what to eat and what not to eat.

‘core causal factors which give rise to a phenomenon’ (Weisberg 2007, p.642). The purpose of this in language acquisition is to idealise away from, to use Cowie’s own words, the ‘strange’ and ‘abstruse’ workings of language to assess the core causal factors facilitating it. Explicitly, idealisation of instantaneity can be taken as a minimal model – a component of the minimalist idealisation P&P – that helps *inter alia* account for language acquisition because the *pld* appears insufficient. If, as Cowie argues, the *pld* is in fact sufficient, then a case could be made against this minimalist idealisation. This, however, relies upon the later discussion of (1) and (3) which, I argue, concludes in favour of an insufficient *pld*.

Cowie makes the further argument that the concept of convergence in the ‘final states’ is an idealisation to *homogeneity* that leads us ‘badly astray in our thinking about language acquisition’ (Cowie 1999, p.222). The main bone Cowie has to pick is drawn from Soames (1984) conceptual bifurcation between linguistics and psycholinguistics – the former concerns natural languages and facts about languages; the latter concerns the psychology of the speaker and facts about speakers. Cowie thinks the idealization of convergence misleads the study of language acquisition by conflating an idealization that serves linguistics (i.e., convergence) with actual facts about the learners (i.e., psycholinguistics). For instance, consider the sentence, **The boys is laughing*—people may maintain that this is grammatical, others may consider it ungrammatical. These differences are explained by Chomskyans with idiolects: the brain’s innate language faculty realises an inter-individual way of speaking by ‘representing’ the faculty. To uncomplicate the explanatory task, the Chomskyan assumes an ‘ideal speakerhearer’ that is representative of a ‘completely homogenous speech-community’ (Chomsky 1965, p.3) that lacks these idiolects. Cowie’s argument is that, though the idealization may be perfectly useful in theorising about language, it is misleading to assume the idealisation veridical, and hence make the assumption that language acquisition must occur the same way. Seeing as the nativist uses it to answer the logical problem, this seems like question begging. To make such a claim ‘requires further empirical support’ (Cowie 1997, p.28), especially with the ‘amount of divergence-within-convergence in speakers’ “final states”’ (1999, p.221)

But does the nativist *really* conflate this conceptual distinction to the extent that the idealization ignores important facts about the way language is acquired? I see no reason to believe that they do. Firstly, Cowie goes wrong by assuming that convergence is being used to show that the learning process is not idiosyncratic. This is incorrect: convergence in S_S is a consequence of the conjunction of the minimalist idealisation of P&P and the assumption that the language acquisition module is universal. These in themselves are, as Cowie notes rightly, on the assumption of the ‘ideal speaker-hearer’ model. But this does not result in ignorance to differences. For the salient feature emphasised in a minimalist idealisation is the causal structure: convergence (i.e., that everyone converges on the same grammar) is a consequence of P&P *qua* (U)². This is

²(U) being the base from which language acquisition occurs through *pld* and the language

not to say that differences regarding the process of language acquisition do not result in different final states, but that the lack of difference in final states is explained by the starting point being the same for everyone (i.e., (I) and (U)). Convergence essentially says little about the final state and nothing about the process leading to it. So, in making the claim that ‘all children’s grammar ultimately converges on the right theory’ is to also bring in all the other descriptive apparatus of (DS), (I) and (U), not to mislead about the language acquisition process. It is, essentially, to say a lot about (G) and (P), both of which are claims argued against by (I). Recall that nativists appeal to convergence to establish the restriction on hypothesising a grammar that is too large and projecting hypotheses that require negative data. This is not, as Cowie claims, to ‘lead us badly astray’ or ‘blind us to the myriad sources of negative evidence in the *pld*’, but to claim that there is no negative evidence and then provide an explanation of language acquisition in face of that difficulty. The only way, then, Cowie’s criticisms can be entertained is if there truly are significant sources of negative evidence in the *pld*. *Contra* Cowie, I argue that there are not.

4 Cowie Loves a Curry

Seeing as Cowie’s criticisms of the nativist position thus far all rely upon the possibility of another solution, the crux of the problem is proving that there is significant and conceivable negative evidence in the *pld* of a learner. Recall that a successful form of negative evidence is evidence that a string *S* is not a sentence of *L*. Cowie, however, claims that the nativist is referring only to *negative data*, which are explicit utterances that *S* is not a sentence of *L*. Apparently, the nativist doesn’t even pay ‘lip service’ to negative evidence concerning the means to disconfirm that *S* is not a sentence of *L*. Cowie argues that non-occurrence of syntax can be taken as negative evidence

Cowie’s first argument for this starts by speculating that a child’s grammar may predict that all intransitive verbs can be causative. Her grammar would predict sentences like ‘I falled the cup off the table’. However, in witnessing her father knock over his coffee and hearing him exclaim, ‘I caused the cup to fall from the table’, the prediction she had generated as appropriate is disconfirmed by her father’s utterance. Cowie’s point is that the ‘datum is positive’ (i.e., a statement), but the evidence is negative, insofar as the non-occurrence of her prediction disproves it. I think this is well explained by the following: the child expected (C), heard (M) from her father, and thus concluded ($\neg C$). However, this conclusion is tentative (what if she misheard, or her father misspoke) and ultimately inconsequential for the nativist’s position (Cowie even cites the nativist Steven Pinker for this example 2004, pp.275-6!). So, building from this, Cowie then makes the further, more significant claim: that the non-occurrence of syntactic structure can also be considered a form of negative evidence, serving to prove that ‘whole classes of objects are not instances of the kind whose

acquisition process.

extension is being learned' (1999, p.224). Assume that the following sentences have never been uttered and are thus not part of the *pld*:

(5) Steve enjoyed the curry.

(6) Enjoyed curry Steve the.

Cowie wants to argue that while the non-occurrence of (5) is not evidence that it is not a sentence, the non-occurrence of (6) is evidence that it is not a sentence. This is because sentences of the same syntactic structure as (5) occur, but sentences of the same syntactic structure as (6) do not:

(5a) $[S[_{NP}[_{VP}[_{NP}[det][_{NP}]]]]]$

(6a) $[V][_N][_N][det]$

Sentences of the form (5a) do occur elsewhere, such as (5a') 'Cowie loves a curry', but sentences of the type (6a), such as 'loves Cowie curry a', never occur elsewhere in a child's linguistic environment. The disparity between occurrence of grammatical structures and nonoccurrence of (5a) constitutes negative evidence, such that the evidence shows a learner that ($\neg$ 5a) in (L). This would be *indirect negative evidence*: evidence *inferred* by the learner and about what is not in (L) from its non-occurrence in the *pld* compared to the occurrence of positive data the learner knows is in (L). Essentially, there are 'many strings that would serve as evidence that a particular structure is allowed by the language' (1999, p.224). If, as Cowie argues, non-occurrence of syntax structure qualifies as indirect negative evidence and it influences the learner, then the nativist's idealizations establishing (I) and (U) are not needed to answer the logical problem.

Cowie's argument thus relies upon the above conjunction, insofar as it amounting to a conclusions such as ($\neg$ 5a). The negative evidence comes from positive data that is the frequency of specific structural occurrences in the *pld*. But while Cowie acknowledges that a structure can have many strings, Laurence and Margolis note that she fails to acknowledge that 'for any given string there are [also] many structures it could have' (2001, p.267) – a many-many relation between structures and strings. This means that the significance of structure as negative evidence is entirely dependent upon the current hypothesis the learner is employing. Consider the following (*ibid.*):

- (7) a. *Does she could go?
b. Does he seem happy?

A learner at (H) where they do not yet individuate between lexical and auxiliary verbs could say the former freely, despite its non-occurrence in her environment. But learners do not do such for both this and the staggering 'number and variety of ungrammatical forms that fit the model of any string' (Laurence & Margolis 2001, p.267). Hence, non-occurrence is not a relevant form of evidence because if it were to serve as indirect negative evidence, then the additional relation

allows for huge numbers of ungrammatical sentences to also contribute to the hypothesis formed. They cannot, then, explain why a learner does not go wrong concerning (7a). This point is *prima facie* correct, but also, I think, unsuccessful at dispelling Cowie’s argument. For she is still ultimately showing that a child can conceivably learn that some strings of a specific (necessarily ungrammatical) structure are not in (L) without (I): that there is (be it indirect and potentially useless) negative evidence in the *pld*. This, I argue, is wrong. But let me return to this after unpacking Cowie’s final point. For my argument dispels both anti-nativist threats at once

Cowie’s argument clearly requires a way for a learner to actually theorise about implicit negative evidence whilst remaining detached from anything that amounts to the claims (I) and (U). To achieve this, Cowie suggests that learners utilise the holistic and probabilistic relations between theories and their evidence. This means that, through a sensitivity to, and facility for exploiting, statistical information in her environment, a learner can infer negative evidence—i.e., what not to say. The statistical information can consist of, for example, the likelihood that a parent repeats a toddler’s utterance if it was grammatical Hirsch-Pasek et al. (1984), or mothers’ contracted and expanded repetitions that break the flow of conversation (e.g., ‘What?’) typically following ungrammatical utterances Demetras et al. (1986). So, as implicit negative evidence, the latter is a statistical frequency that can be noticed by the learner that their current theory is wrong, even though the parent never exclaimed, ‘No, that is wrong.’ Equally, the statistical frequency of sentence structures in the *pld* amount to the significant hypothesis that sentence forms of (6a) are not part of (L). This is fundamental to Cowie’s argument for negative evidence, being the explanation of the other aforementioned conjunct: that the learner is able to infer from the negative evidence to a negative conclusion.

The argument above seems viable, so the question for the nativist is: what, if not implicit negative evidence, is facilitating language acquisition in this case? Answering this is twofold. Firstly, the nativist can argue that the holistic and probabilistic reasoning aspect is actually best explained by nativism. If not constrained by innate principles, then the non-nativist is in a sticky situation. For, in being provided with data, learning consists of identifying and interpreting the causal system underlying the observation. But, left unrestricted (and this is possible because (DS) does not necessitate that everyone has the restriction), probabilistic interpretation of the data can produce infinitely many potential hypotheses, which may make the learner either (G) or (P). An obvious way of restricting the overgeneration of potential hypotheses is (I) and (U)³. Recall P&P: one can think of the principles as corresponding to invariant properties universal in human languages. The parameters, however, are ‘set’ by positive data about what is in (L), and this can include statistical data too. However, consider Baker’s Paradox Baker (1979):

³Of course, almost all theories of language acquisition presume some degree of innateness: ‘A mind cannot be a blank slate, because blank slates don’t do anything’ (Pinker 2004, p.34) In turn, the focus is on whether statistical reasoning is better explained by (I) or just (DS).

- (8)a. John told Bill the story.
John told the story to Bill.
- (8)b. John donated the painting to the museum.
*John donated the museum the painting.

These sentences have the same structure and construction, but only the dative alternation of (8b) is ungrammatical. The problem is that, unlike (7), children do overgeneralise the restriction in the way of (8b), but eventually converge to not overgeneralise it. This was initially taken as a strong case for the innateness of morphophonological and semantic knowledge Gropen et al. (1989). However, Theakston (2004) found that restrictions on the dative alternation of verbs were learned at later ages for low-frequency verbs than for high-frequency verbs. Now, if we understand indirect negative evidence in Cowie’s sense, this looks brilliant for her and tragic for the nativist: a convergence on a grammatical restriction related to the frequency of verbs in the *pld*! But if we go back to Laurence and Margolis’s point concerning the inapplicability of negative evidence we see a very different picture for Baker’s paradox and Theakston’s evidence.

Take their argument that implicit negative evidence is unavailable in any relevant sense to learners. This argument, which I maintain is successful, does not exclude the positive evidence that empowers Cowie’s conceptualisation of implicit negative evidence. However, I argue that it is not even *conceivable* that non-occurrence constitutes implicit negative evidence. Instead of taking the relations between occurrence and non-occurrence as evidence, why can we not consider the positive data as having informative statistical relations to other positive datum? Cowie’s induction of negative evidence from positive data has been shown only the potential of very limited and problematic indirect negative evidence. In turn, instead of arguing for an inductive learning approach as Cowie does, it is possible to understand the process language acquisition as a parameter setting (i.e., based on P&P) approach, an example of this being the *Structural Triggers Learn* model Sakas et al. (2017). What is important to take from these strategies for present purposes is that the statistical inference of *positive* evidence can still be applied to language acquisition without incurring all the issues from inferring relations between, say, occurrence and non-occurrence. This is because positive data is still useful here, but one can disregard indirect negative evidence. For example, in (8) the convergence from frequency is a consequence of the positive data influencing the positive hypotheses that, say, a certain syntactic feature (*F*) triggers dative alternation. Moreover, parameter setting models usually imply at least (I) concerning principles; otherwise, we end up with the same issues of overgeneration, etc. highlighted by myself, and Laurence and Margolis earlier. Let us say, however, that we only accept (DS) concerning (8): we have to thus assume that every learner will eventually encounter sufficient foreknowledge to restrict dative alternation. Fundamentally, the proponent of just (DS) would have a much larger task ahead of her than the nativist. For the nativist can appeal to the inborn principles restricting mistakes and helping her narrow down her statistical hypotheses. Unlike, Cowie’s unsuccessful attempts to glean negative data from inductive relations, param-

eter setting models seems to successfully synthesises the positive evidence into a sufficiently explanatory role within the nativist system, without sacrificing or obfuscating the statistical patterns and role of word frequency. Equally, by retaining (I), it fails to succumb to any of the issues that Cowie raises. Strictly speaking, nativists provide a better explanation of how a child manipulates positive data than Cowie's remedy, and does so without the *myriad possible pitfalls* shown throughout this essay.

5 Conclusion

Cowie's case against the logical problem has ultimately failed. The success of her first two accusations concerning idealization as a means to obviate the existence of other sources of negative evidence failed. They ultimately hinged upon her arguments for indirect negative evidence, which I suggest point further to the inconceivability of indirect negative evidence as derived from unattested sentence form or structure. Moreover, I have suggested that she misinterprets what learners even do with positive datum: the approach to statistical inference by holistic and probabilistic means seems wholly unreasonable on Cowie's account.

However, Cowie's account has made clear certain weaknesses of the nativist's initial position. The assessment of indirect negative evidence, for instance, has been shown to be irretrievable; however, though I suggested a non-inductive method, the inductive method is not exclusively nativist. It is certainly possible that the richness of the positive datum can likely produce indirect negative evidence. However, Cowie would have to step away from her staunch anti-nativist path and concede more than what she lets on. For the logical problem continues to prove itself a strong argument for both weak and Chomskyan nativism.

References

- Baker, C. L. (1979), 'Syntactic theory and the projection problem', *Linguistic Inquiry* pp. 533–581.
- Chomsky, N. (1965), *Aspects of the Theory of Syntax*, MIT Press, Cambridge, MA.
- Chomsky, N. (1986), *Knowledge of Language: Its Nature, Origins, and Use*, Praeger, Westport.
- Chomsky, N. (2000), *New Horizons in the Study of Language and Mind*, Cambridge University Press, Cambridge.
- Cowie, F. (1997), 'The logical problem of language acquisition', *Synthese* pp. 17–51.
- Cowie, F. (1999), *What's Within: Nativism Reconsidered*, Oxford University Press, Oxford.

- Demetras, M., Post, K. & Snow, C. (1986), 'Feedback to first language learners: The role of repetitions and clarification questions', *Journal of Child Language* pp. 275–292.
- Fodor, J. A. (2001), 'Doing without what's within: Fiona cowie's critique of nativism', *Mind* **110**(437), 99–148.
- Gropen, J., Pinker, S., Hollander, M., Goldberg, R. & Wilson, R. (1989), 'The learnability and acquisition of the dative alternation in english', *Language* pp. 203–207.
- Hirsch-Pasek, K., Treiman, R. & Schneidermann, M. (1984), 'Brown and hanson revisited: Mothers' sensitivity to ungrammatical forms', *Journal of Child Language* pp. 81–88.
- Laurence, S. & Margolis, E. (2001), 'The poverty of stimulus argument', *The British Journal for the Philosophy of Science* pp. 217–276.
- O'Grady, W., Dobrovolsky, M. & Katamba, F. (1996), *Contemporary Linguistics: An Introduction*, Longman, London.
- Pinker, S. (2004), *The Blank Slate: The Modern Denial of Human Nature*, Viking, New York.
- Sakas, W. G., Yang, C. & Berwick, R. (2017), 'Parameter setting is feasible', *Linguistic Analysis* pp. 391–408.
- Soames, S. (1984), 'What is a theory of truth?', *Journal of Philosophy* **81**(8), 411–429.
- Theakston, A. L. (2004), 'The role of entrenchment in children's and adults' performance on grammaticality judgment tasks', *Cognitive Development* pp. 15–34.
- Weisberg, M. (2007), 'Three kinds of idealization', *Journal of Philosophy* pp. 639–659.