CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Exploring early syntactic generalisation: evidence from a growth curve analysis of Spanish "se" constructions

Nick RICHES [ORCID]

Education, Communication and Language Services, Newcastle University, UK
Email: nick.riches@newcastle.ac.uk

Abstract

Children's early grammatical constructions, e.g., SVO, exhibit a learning curve with cumulative verb types (CVT) increasing exponentially. According to Ninio (2006), the fact that learning curves, though nonlinear, can be modelled by a continuous regression suggests instant generalisation. Moreover, differences in initial verbs across children indicate minimal involvement of semantics. This study tested these claims on the Spanish "se" constructions (SSCs) in two children, Juan and Lucía (Aguado-Orea & Pine, 2015). Ninio's findings were replicated. Nonetheless, exploratory analyses indicated that curves are driven by the temporal distribution of tokens (instances of the SSC irrespective of verb type) and therefore may reflect non-productivity-related mechanisms, e.g., retrieval-based learning. Furthermore, hapax verbs were relatively late to emerge in the children's data, suggesting emergent generalisation. Analyses of raw lexical frequencies indicated relative semantic homogeneity across the two children's verb types, suggesting a semantic prototype. Nonetheless, ecological factors may also explain these lexical similarities.

Keywords: syntax; semantics; usage-based approaches

Literature Review

Growth curve analysis and its theoretical implications

The syntactic productivity of children's early utterances is a central issue in language acquisition research. While a sentence such as *break it* employs a VERB-plus-pronoun schema, it is only genuinely productive if the child can insert numerous verbs to produce combinations unattested in the input (Childers & Tomasello, 2001). Productivity is a defining characteristic of adult-like competence, and therefore constitutes an end-state which children must attain. While many researchers argue for the rapid emergence of productivity (Yang et al., 2017), others propose an early stage of limited productivity during which children gradually acquire abstract syntactic rules from the input alone (Tomasello, 2003). Debates are informed by theoretical models. Researchers influenced by Universal Grammar presuppose innate representations which support early

productivity (Yang et al., 2017), while usage-based theorists who reject pre-existing knowledge assume a gradual emergence of productivity (Dąbrowska & Lieven, 2005; Tomasello, 2003). However, some accounts do not adhere to this early-late dichotomy. For example, maturational accounts argue that innate linguistic knowledge is activated at particular stages of development (Rice et al., 1995). Likewise, usage-based accounts are consistent with early productivity, if we assume that domain-general learning mechanisms are extremely powerful.

Both corpus studies and behavioural studies have investigated productivity. Corpus analyses by usage-based theorists argue that children's early language is uncreative, for example, depending on rote-learned Determiner-plus-Noun combinations (Pine & Lieven, 1997), or generating complex sentences such as questions from previously-acquired grammatically simple phrases (Dąbrowska & Lieven, 2005). However, techniques for determining productivity are problematic. For example, skewed lexical distributions can create an appearance of conservatism, and adult data also demonstrate limited productivity when similar algorithms are applied (Yang et al., 2017).

Behavioural data are likewise contested. Some production data suggest a gradual emergence of linguistic structure. For example, Akhtar (1999) found that, unlike older children, three-year-olds do not impose standard subject-verb-object word order when recalling sentences such as *Elmo the tree meeked.* However, production studies are confounded by numerous factors, including whether children understand the experimental task, are fully engaged or have phonological or lexical retrieval difficulties. Comprehension paradigms, by contrast, may be less confounded. Indeed, preferential looking studies find above chance comprehension of the transitive in children aged 1;9 (Gertner et al., 2006). However, as Abbot-Smith and Tomasello (2006) observe, children may exploit low-level heuristics rather than fully abstract representations, e.g., identifying the post-verbal argument as the patient.

Given these methodological challenges, we should explore alternative paradigms. The current study explores a relatively underutilised approach: growth-curve modelling. This involves plotting productions of the target construction over time in a longitudinal linguistic corpus sampled for a particular child. The shape of the developmental trend may indicate whether the construction is productive. For example, some researchers have argued that a curve characterised by a gradually-increasing slope suggests productivity from the outset. The technique is promising because, with a sufficiently dense sampling regime and appropriate developmental time window, we can track the emergence of a construction in production. This can help determine whether innate knowledge is present. By contrast, because experimental studies generally require children old enough to participate in an artificial interaction with a stranger, they often fail to capture early development.

The paradigm also offers several other advantages. While, in experimental studies, it can be difficult to create a context which elicits the target construction, in naturalistic data such contexts arise naturally. Because the context is relatively uncontrolled we can also explore lexical variation within the construction. Furthermore, because growth curves are plotted for individual children, data are not influenced by between-subject variability in processes unrelated to linguistic competence such as lexical retrieval and output phonology. Though growth curves themselves may conceivably reflect such processes, Ninio (2006) argues that learning trajectories do not reflect these confounds (see below).

Regarding developmental trends, most studies identify a non-linear trajectory characterised by decreasing intervals between targets. When cumulative targets are plotted against time, a relatively smooth curve is observed which transitions from approximately
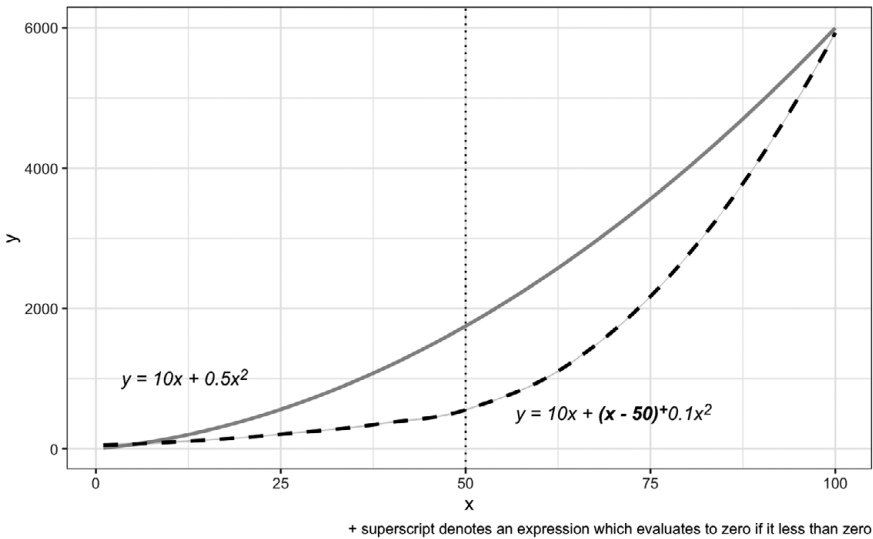
Figure 1. Plot demonstrating continuous and discontinuous growth curves

flat to approximately forty-five degrees. Two lines demonstrating such a curve are shown in Figure 1. Both exhibit a "concave" trajectory, bending downwards from a straight line. This shape is modelled by a regression with an exponentiated term (x^2) whereby a variable is raised to a power, in this case two. Regression models that use a constant exponent (e.g., 2), or series of exponents, are sometimes called polynomial models. Alternatively, a power law function may be used whereby the exponent varies, e.g., $y = \beta^x$. In practice, both techniques can describe very similar curves.

A key theoretical question is whether the trend is discontinuous. In Figure 1 the bottom curve is discontinuous because the exponentiated term is only included once x increases above 50, generating sudden curve. By contrast, the model describing the top curve is continuous. Though the rate of change varies, importantly, it changes at a constant rate. Whether a process is continuous or not has implications for theories of language development. A discontinuity may indicate the sudden onset of linguistic generalisation, suggesting that linguistic knowledge is not present from the outset.

The curve-fitting approach was pioneered by Ruhland and colleagues (Ruhland et al., 1995; Ruhland & Geert, 1998) who modelled the frequency of function words (pronouns, modal verbs, and articles) in 6 Dutch-acquiring children aged approximately 1;6 to 3;3. They were also the first to make theoretical claims based on the presence or absence of a continuous process. Though plots, based on fortnightly recording sessions, indicate a sudden acceleration during the third year, the developmental curve was best fitted by a nonlinear model (cubic logistic). Because this model is continuous, the authors argued that there is no sudden emergence of a linguistic rule.

More recently, Abbot-Smith and Behrens (2006) employed growth-curve analyses to investigate the acquisition of the German passive by a single German-acquiring child, Leo. The child's language was sampled during five one-hour sessions per week between 2;0 and 3;0 thereby achieving a good sampling density. They compared growth curves for the *sein* passive (*Der Reis ist gekocht* = 'the rice is cooked'), and constructions exhibiting partial

formal and functional overlap with the *sein* passive, e.g., copula plus past participle constructions, (*Der Ball ist gefallen* = 'the ball is fallen'). This analysis investigated whether passive-like constructions would, via analogy, "conspire" to support the development of the *sein* passive. Rather than model frequency, they modelled cumulative verb types, i.e. the number of different verbs used in the construction at a given timepoint. This, they argue, provides a good measure of productivity.

While the supporting constructions were best modelled by concave growth curves using a power law function, the *sein* passive exhibited a linear trajectory. According to Abbot-Smith and Behrens, the curves of the supporting constructions indicate emergent productivity, while the *sein* passive is productive from the outset, because it benefits from the supporting constructions. By contrast, for Leo's mother, all constructions exhibited a linear trajectory, thus indicating that learning curves are a characteristic of child language. The authors also explored Leo's use of the *werden* passive (*werden* = 'become'). This did not benefit from prior acquisition of potentially-supporting constructions and therefore exhibited a nonlinear trajectory, which characterises emergent generalisation.

Probably the most detailed analysis of learning curves was undertaken in a series of articles by Ninio (1999, 2005a, 2005b), culminating in a monograph entitled *Language and the Learning Curve* (2006). She focused on simple syntactic schemas such as Verb-Object, Subject-Verb-Object, and Verb-Indirect Object in Hebrew-acquiring children (Ninio, 1999, 2005a, 2005b). The samples of Hebrew-acquiring children are relatively large, ranging from fourteen (Ninio, 2005a) to twenty (Ninio, 2005b), and all, except one child, were sampled at weekly intervals. She also investigated Verb-Object and Subject-Verb-Object structures in a child acquiring English (Ninio, 1999) whose data were collected by Tomasello in a diary study (1992). Like Ruhland and colleagues, she argued that initial trajectories, characterised by a progressively steeper gradient, are best modelled by continuous regressions, thereby indicating early abstract knowledge. The learning curve itself, she claims, reflects a power-law-of-practice effect (Fitts & Posner, 1967) whereby the time taken to perform a complex task approximately halves with each new trial. This is best modelled using a power law. This arises from the domain-general process of analogy (Brown & Kane, 1988), which, she argues, underpins generalisation across abstract syntactic structures.

One key piece of evidence for the role of domain-general mechanisms is that the learning curve is similar for both early and late-acquired structures (Ninio, 2006). This excludes language-specific, developmentally-constrained processes, such as early difficulties with lexical retrieval or output phonology. As mentioned above, the hypothesised ability to minimise the impact of such confounds constitutes an important benefit of this approach.

Ninio's account is eclectic, adopting a usage-based emphasis on domain-general learning mechanisms, while advocating innate syntactic knowledge. Another key claim is that generalisation is guided by form, not meaning. She demonstrates this in numerous ways. Firstly, early verbs appearing in a particular frame/schema do not necessarily exhibit semantic properties consistent with the assumed semantic prototype underlying that schema. For example, prototypically transitive verbs involve a dynamic event with a volitional agent, and a highly affected patient (Hopper & Thompson, 1980). However, Ninio (1999) observes that many early transitive verbs are not consistent with this prototype. Secondly, she argues that there is little semantic uniformity among the early verbs and also among the thematic roles of their arguments (Ninio, 2005a, 2005b, 2006). Finally, she argues that semantic homogeneity among early verbs is not linked to rates of generalisation. For example, the deletion of items with many semantic antecedents does

not affect the shape of the learning curve (Ninio, 2005a). Importantly, she argues that semantically-driven generalisation may be present in older children (> 3;5) (Ninio, 2006:92), but, crucially, plays little role in early development. Ninio's argument for form-based generalisation runs counter to usage-based theories which argue that the semantic fit between a verb and a construction influences productivity, e.g., Goldberg's dynamic categorisation model (2016).

In summary, the syntactic growth curve literature is diverse in its methods, and claims. Sampling regimes range from a single child sampled on a daily basis (Abbot-Smith & Behrens, 2006) to twenty children sampled every week (Ninio, 2005b). A key difference across studies is the interpretation of nonlinear growth. While Ruhland and colleagues and Ninio emphasise that such curves reflect a continuous process, thereby suggesting early or instant generalisation, Abbot-Smith and Behrens assume that they reflect emergent productivity. Nonetheless, there is a consensus that non-linear growth characterises early syntactic development. There is clearly a need for further investigation of this phenomenon.

The current study conducts a growth curve analysis of Spanish *Se* Constructions (SSCs) in two children from Aguado-Orea and Pine corpus (Aguado-Orea & Pine, 2015). This enables an exploration of semantic generalisation because certain subclasses of SSCs impose relatively strong constraints on verb meanings (Maldonado, 2008). In addition, their adoption of a regular clitic form makes them amenable to corpus analysis.

Spanish *Se* Constructions (SSCs)

Though Spanish reflexive clitics, realised as *se* in the third person, frequently describe reflexive events (as in 1), they frequently lack a reflexive interpretation. In fact, the wide variety of meanings they signal has been described as "functional overload" (Moreno, 2021). The term SSC is borrowed from Mendikoetxea (1999) and Moreno (2015) to describe constructions employing reflexive clitics, whether or not they can be interpreted as true reflexives. The use of a form-oriented label is motivated by both the variety of SSC meanings, and the inherent difficulties of grouping these into discrete categories. For a detailed description of non-reflexive SSCs the reader is referred to Mendikoetxea (1999).

The reflexive (1) is used when an action is directed at the individual who performs it. The reflexive particle corresponds to the English pronoun *himself/herself* when placed in direct object position.

(1) Juan   se   vio   en   el   espejo
     John   REFL-3SG   see-3SG.PST in   the   mirror
     'John saw himself in the mirror'

The reflexive pronoun exhibits clitic properties, being phonetically insubstantial, unstressed, and with a syntactically-motivated distribution. It occurs before a tensed verb (as above), or after a non-tensed verb. The reflexive has a dual, or "Janus-like" nature (Kemmer, 1994; citing Faltz, 1977), whereby subject and clitic (*Juan* and *se*) refer to the same entity but have distinguishable roles. In this case, *Juan* is an EXPERIENCER, while *se*, describing his reflected image, is a THEME. Some authors use different terminology, e.g., Kemmer (1993) refers to INITIATOR and ENDPOINT. Distinct roles are possible because reflexive sentences use verbs which are conventionally transitive. For example, because *ver* (see) when used transitively takes a THEME direct object, this role is

expressed by the clitic when the verb is used reflexively. Kemmer (1993; as cited by Calude, 2017) describes this as "coincidental" co-reference because there is a potential for disjoint reference which is overridden in reflexives.

Other SSCs lose the duality of the reflexive. (2) employs the reflexive clitic with a change-of-state (inchoative) verb. Unlike the reflexive, roles are not distinct. The subject (*ella*) is the EXPERIENCER, and there is no other identifiable thematic role. Though coreference is formally marked via person and number agreement between subject and clitic, it is not semantically expressed.

(2) (Ella)   se   enfadó
     (She)   REFL-3SG   anger-3SG.PST
     'She became angry'

The reflexive clitic may also be used in passive, impersonal and reciprocal constructions:

(3) Los   pasteles   se   vendieron
     The   cakes   REFL-3PL   sell-3PL.PST
     'The cakes were sold'

(4) Se   vive   bien   en   España
     REFL-3SG   live-3SG.PRS   well   in   Spain
     'One lives well in Spain'

(5) Los   gamberros   se   pelearon
     The   hooligans   REFL-3PL   fight-3PL.PST
     'The hooligans fought each other'

Like (2) these all employ the reflexive clitic but do not exhibit the two key characteristics of reflexives – namely, distinct roles and coincidental co-reference. The term MIDDLE VOICE is often used by linguists to describe such non-reflexive usages of reflexive clitics, with the term SPANISH MIDDLE VOICE used to to describe the middle voice system in Spanish.

Other languages with productive reflexive and middle voice systems also use the reflexive morpheme across a similar range of constructions (Kemmer, 1993). In some of these languages, e.g., German and French, the same form is used for reflexive and middle voice interpretations. In others, e.g., Russian, Old Norse, Jola, a reduced form of the reflexive clitic is used in the middle voice. Though some languages employ different forms for reflexives and middle voice interpretations, e.g., Latin and Turkish, there is a strong cross-linguistic tendency for forms to overlap.

Theoretical linguists have attempted to explain the reflexive-versus-middle relationship across languages. A common claim is that the reflexive clitic reduces transitivity (Grimshaw, 1982; Kemmer, 1993; Rosen, 1988). This can be conceptualised as a continuum between prototypical two-participant events (*hit*) and prototypically one-participant events (*go*). When a verb frequently used to describe a transitive event is rendered reflexive (example 1), a participant is removed, thereby reducing transitivity. However, there are still two identifiable roles. Middle voice verbs such as *enfadarse* (become_angry-REFL) have only one role, and are consequently less transitive. *Se* passives reduce transitivity by syntactically profiling a single argument. Finally,

reciprocals collapse the distinction between roles as each participant becomes both an initiator and endpoint.

While transitivity reduction characterises many SSCs, as Maldonado notes (1992, 2008), the reflexive clitic can also INCREASE valence. For example, the verb *quitar* ('remove') requires at least two participants; a remover and the thing being removed. However, it is frequently used with an additional reflexive clitic:

(6) (Ella) se   quitó   el   sombrero
     She   REFL-3PS   remove-3PS.PST   the   hat
     'She took off her hat'

This clitic makes the action self-directed. For example, if you were removing someone else's hat, you would not use it (*quitó el sombrero a Juan* = 'she/he/they$_{SING}$ took off Juan's hat'). According to Maldonado the clitic identifies a region of conceptual space called the DOMINION to which the individual has mental or physical access (Maldonado, 1992).

Other examples of extravalent *se* have no clear referent. In *(Él) se fue* = 'he left', and *(Ella) se cayó* = 'she fell', the clitic results in a focus on the pivotal moment of change (Maldonado, 2008), e.g., the moment when someone leaves the room, or falls over. If there is no obvious moment of change, *se* is not used, e.g., *el coche iba lento* = 'the car went slowly', *las hojas caían* = 'the leaves were falling'. For other verbs, extravalent *se* yields a relatively idiosyncratic interpretation. With *bailar* ('dance'), it describes a situation whereby the dancers have a heightened emotional investment in the dance (Maldonado, 2008). With consumption verbs (*comer* = 'eat', *beber* = 'drink') it describes a situation where all of the food or drink is consumed. With *morir* ('die') it denotes unexpectedness.

Verbs taking extravalent *se* contrast markedly with SSCs which reduce transitivity (2 – 5). However, there are semantic commonalities, e.g., *enfadarse* (become_angry-REFL) and *caerse* (fall-REFL) both describe changes of state. A unified semantic account of SSCs encompassing both valence-reducing and valence-increasing verbs was proposed by Kemmer (1993). She argued that the reflexive clitic signals a relative reduction in granularity, or elaboration of events. Transitive scenes form the most granular end of this continuum. These involve referentially distinct and thematically-differentiated agents and patients, and a clear causal pathway, whereby the agent imparts its energy to the patient. Reflexives and middles reduce granularity by removing one or more of these characteristics. In reflexives the referential distinction between arguments is collapsed. In inchoatives the thematic distinction is collapsed. In *se* passives and impersonal *se* sentences, the agent is deprofiled. For verbs of movement (*irse* = go-REFL, *caerse* = fall-REFL), a sub-part of the movement path is profiled, thereby accounting for the aspectual properties of *se*.

Acquisition of SSCs

SSCs evidently constitute a complex system. Some relatively abstract generalisations, e.g., reduction in transitivity or granularity, capture a wide range of data. Other generalisations have more limited coverage, and are more semantically specific, e.g., verbs describing an emotional change-of-state (*enfadarse* = become_angry-REFL), or a focus on a pivotal moment of change (*caerse* = fall-REFL) frequently occur in SSCs. For even smaller verb groups *se* conveys a relatively idiosyncratic meaning, e.g., emotional involvement (*bailarse* = dance-REFL).

Frameworks differ in their view of generalisation. Generative theories presuppose that high-level generalisations, for example related to loss of transitivity (Grimshaw, 1982; Rosen, 1988) or the suppression of an external argument (Alexiadou & Doron, 2012) are encoded in Universal Grammar, and require sufficient input to be activated. Usage-based accounts, by contrast, view generalisations as an emergent phenomenon, and therefore, they merely need to be sufficiently "general" to allow productive language use (Langacker, 2009; as cited in Dąbrowska, 2008).

Though the type of linguistic generalisations made by competent adult speakers is much debated, there is a relative consensus regarding triggers for generalisation in children. Ninio (2006), adopting a minimalist framework, proposes an extended period of form-based generalisation. For SSCs, the key formal marker is the presence of the reflexive particle. Usage-based theories likewise propose initial attention to form, e.g., distributional patterns involving open and closed slots. For example, Childers and Tomasello (2001) propose that children initially categorise direct object *it* as a transitivity marker. Likewise, children may use closed-class *se* as a marker of SSCs. Such patterns constitute islands of "reliability" (Ibbotson, 2013). Eventually, via domain general mechanisms, they will become productive schemas.

One area where accounts differ is the rate of generalisation. According to generative accounts, children possess innate grammatical knowledge which allows them to make rapid generalisations. By contrast, many usage-based theories propose that generalisation is emergent, rather than instant, because it depends on analysis of the input (Dąbrowska & Lieven, 2005; Pine & Lieven, 1997). This broad dichotomy, of course, simplifies a complex theoretical landscape.

Another contested issue is the role of semantics, e.g., whether generalisation is guided by constructional or lexical meanings. Usage-based accounts emphasise semantic generalisation for two reasons. Firstly usage-based grammatical theories, e.g., Construction Grammar, propose that grammatical constructions are inherently meaningful (Fillmore, 1988), and therefore, semantic generalisation is likely to be an important and early process. Secondly, if children do not have innate knowledge, they must exploit all available generalisations even low-level ones, e.g., ones based on prototypical meanings of constructions.

Regarding SSCs, there is evidence for relative semantic conformity. Jackson-Maldonado et al. (1998) found that, in children aged 2;4 to 3;0, certain semantic classes of verbs predominated, with 32% involving motion, and 20% involving unexpected changes. The latter might indicate a semantic category based on unexpectedness, a concept closely related to Maldonado's (2008) "pivotal moment of change". Interestingly, reflexives constituted only 9% of tokens. This could reflect a relative absence of inherently reflexive verbs. Grooming verbs, e.g., *bañarse*, are one such group, but they constitute a relatively small set. It may therefore be more difficult to exploit semantic generalisation to acquire the reflexive.

Another issue related to acquisition of SSCs is whether children start with a unitary SSC category. Jackson-Maldonado et al. (1998, p. 408) argue that "differentiating between true reflexive forms and middles is… crucial to the proper understanding of the development of the *se* form". This implies that children rapidly acquire, or begin with, a system involving a reflexive-middle distinction. Alternatively, children may start with a unitary category which is progressively subdivided. For example, as proposed by both generative and usage-based theorists, children may begin with a formal category, e.g., one based on the presence of the clitic *se.* Given the distributional properties of the input (Jackson-Maldonado et al., 1998), most verbs in this category are likely to denote sudden agentless

change, e.g., *caerse* (fall-REFL) and *romperse* (break-REFL). Subsequent members of this formal category may be semantic outliers, e.g., *bañarse* (bathe-REFL), which assumes an agent, and *verse* (see-REFL), which is often used impersonally, e.g., 'one can see that…'. Over time, as these outlying groups acquire more members, they will form their own categories, e.g., grooming SSCs or SSCs denoting impersonal situations. Because grooming verbs, and the verb *ver* ('see') are frequently used in transitive situations, they can act as a means of acquiring the reflexive-transitive alternation.

This account is motivated by the limitations of high-level generalisations. These are either too specific to describe all SSCs (e.g., *se* as a marker of detransitivisation), or too abstract to be genuinely productive given the limited data available to young children (e.g., the notion of loss of "granularity"). Moreover, given the sheer multiplicity of SSCs children would need to acquire substantial data to derive these high-level generalisations in the first place. Starting with a small number of semantically-motivated subgroups provides a stepping stone to these higher-level generalisations, enabling children to acquire and practise the formal properties of this construction en route to developing a more abstract system.

Research aims and hypotheses

Planned analyses replicated Ninio's. Growth curves for cumulative verb types were modelled. Consistent with Ruhland and colleagues (Ruhland et al., 1995; Ruhland & Geert, 1998) and Ninio (Ninio, 1999, 2005a, 2005b, 2006), a curve reflecting a continuous process was assumed to indicate early or instant generalisation. This is based on the argument that, if there is no discontinuity in the underlying statistical model, we should not assume a discontinuity in terms of psycholinguistic mechanisms.

Further exploratory analyses, i.e. not planned in advance, were conducted to explore psycholinguistic mechanisms underlying growth curves. The first analysis investigated the cumulative frequency of SSC tokens and whether these could also be described via a growth curve. This is important as it is theoretically possible that patterns observed for cumulative verb types are a consequence of nonlinear trends in token frequencies of the target construction. These have not been modelled by existing studies.

Secondly, hapax legomena (items occurring only once in a particular frame) were investigated within the verb slot. This provides an alternative means of exploring productivity which has been utilised by numerous studies (e.g., Baayen & Lieber, 1991). This method originated in the analysis of morphological productivity, and has since been extended to investigate productivity in large-scale constructions, e.g., argument structure constructions (Barðdal, 2008; Zeldes, 2012), and the German comparative correlative construction (Zeldes, 2011). The underlying logic is that there is a correlation between frequency and the likelihood that an item-frame combination is productive. A high frequency combination is likely to result from a rote-learned chunk. Lower-frequency items are less likely to reflect chunk-learning. N = 1, as the lowest possible frequency, provides the best indicator of productivity. The potential of hapax legomena as an indicator of productivity was recently demonstrated by Pierrehumbert and Granell (2018). Using data from Wikipedia, they found a strong relationship between morphological processes identified from hapax legomena in artificially truncated corpora, and their frequency profiles in much larger corpora.

Finally, while lexical analyses conducted by Ninio investigated the order of occurrence of verb types, the current study also investigated raw frequencies for different verb types.

Though order of occurrence is more informative regarding generalisation, it is more impacted by the vagaries of sampling, e.g., whether regimes are sufficiently dense to accurately capture ordering effects. Because raw frequencies are based on multiple samples, they are less affected by this issue.

The study will investigate growth curves for SSCs in child language corpora replicating the methods developed by Ninio, in order to test her claims. In particular it will address the following questions:

(1) Are growth curves evident in cumulative verb types occurring in SSCs? Can they be described by a regression modelling as a continuous process?
(2) Do growth curves for cumulative SSC tokens mirror those for verb type frequency?
(3) Do hapax legomena exhibit (a) a growth curve (b) early onset?
(4) Does the order of occurrence for verb types within SSCs reveal semantic homogeneity?
(5) Does the frequency of occurrence for verb types within SSCs reveal semantic homogeneity?

Methodology

Ethical statement

Because the data are in the public domain, no ethical approval was sought.

Properties of the corpus

Data were taken from the Aguado-Orea and Pine Corpus (Aguado-Orea & Pine, 2015: https://childes.talkbank.org/access/Spanish/OreaPine.html) hosted by the CHILDES data exchange system (MacWhinney, 2000). The two children in the corpus, Juan and Lucía, were sampled from an early age; Juan from 1;10 to 2:05, and Lucía from 2;02 to 2;07. The corpus is relatively dense, with short intervals between recording sessions: approximately 4 day intervals until 26 months, and approximately 3-day intervals thereafter (see graphs in Appendix 1). Having short intervals maximises the temporal resolution of growth curves. The corpus has a denser sampling regime than those employed by Ninio (weekly) and Ruhland et al. (fortnightly). A further advantage is that individual sessions were relatively long, providing substantial data. There were on average 274 child utterances per session with a mean standard deviation of 162. Having large samples increases the coverage of the data, i.e. the proportion of Juan's total expressive language, thereby improving reliability.

A further consideration is whether the developmental time window captures the emergence of SSCs. Though research on early SSCs is sparse, data from the López Ornat corpus (Ornat, 1994), analysed by Domínguez (2003), suggest appearance late in the second year. María (pseudonym) was the single subject in this corpus, and her utterances were gathered during 30 sessions from 1;7 to 4;0. She produced a single SSC token aged 1;07. However, it was not until 1;10 that multiple SSCs were observed (15 tokens altogether in the sample). According to the longitudinal data plotted by Domínguez, María reaches around 50 SSC tokens per transcript by the age of 2;3, constituting approximately 7% of utterances in that transcript. These data indicate that SSCs develop late in the second year. Consequently, Juan's data may better reflect the emergence of SSCs than Lucía's.

Coding

For each script, the CLAN program COMBO was run to identify MOR lines containing a reflexive particle (combo +t%mor +t*CHI +s*refl* +fS). The results were transformed into an Excel spreadsheet using a bespoke R script. A further R script extracted verbs from the MOR tier, along with the English translation. All lines were manually checked by the author to ensure that the mor codes in CLAN had been correctly applied. The author is a proficient speaker of Spanish as a second language (Level C2 of the Common European Framework assessed via the *Diploma de Español como Lengua Extranjera*). A few problematic cases arose. Firstly, for some sentences containing multiple verbs, there was more than one possible attachment site for the reflexive particle. An example is *se va a caer* (REFL go-3ps.PRS to fall-inf). The most likely interpretation is that an entity will accidentally fall in the near future. In this case, the reflexive particle modifies the more distant verb, *caer* ('fall'), yielding an accidental/unintentional interpretation. Another frequently occurring ambiguous sentence employed a control verb, e.g., *se puede caer* (REFL can-3ps.PRS fall-inf). Again, *se* is best interpreted as modifying *caer*.

A further source of coding difficulty was the homophony between the reflexive particle (third person singular) and indirect object pronoun, e.g., *se lo quité el sombrero* (IND.OBJ. pronoun DIR.OBJ.pronoun take_off.1ps.past the hat = 'I took his hat off him'). These were identified by searching for the pronoun combination *se lo(s)* and *se la(s)*. Though the *se lo(s)/la(s)* combination can occur with genuinely reflexive verbs, in most cases, e.g., *se lo llevó el sombrero* = 'he took his hat with him', agreement cues can be used to disambiguate. For example, lack of agreement between *se* (3ps) and the main verb above, would exclude the reflexive interpretation.

The three potentially ambiguous cases above (going to future, control verb, and potential ditransitives) were identified using Excel searches for relevant morphemes (e.g., lemmas *ir* and *poder*, and *se lo(s) / se la(s)*), and manually coded by the author. Such utterances constituted 22% of all utterances identified as potentially containing SSCs (Juan, Lucía and respective caregivers). Where decisions were difficult, transcripts were inspected for context. Given the complexity of this coding process, potentially ambiguous cases were also checked by a native Spanish speaker with linguistics training, who had access to the transcripts. Agreement rates were high (85%). Where disagreements arose the author consulted the codings of the interrater, and in all cases agreed and recoded the data accordingly.

The reflexive-middle distinction was not coded. This is consistent with the theoretical claim, outlined above, that children initially treat SSCs as a single category. Moreover, practically, this distinction is difficult to code without rich contextual data. For example, *No se ve* (No REFL see-3SG.PRS), which Juan produces at 28 months 24 days, could describe an impersonal situation ('one cannot see it'), or a reflexive situation ('she cannot see herself'). Only a detailed analysis of context, e.g., video data, can distinguish these two meanings. However, video data was not archived with transcriptions. Given these theoretical and practical arguments, SSCs were coded as exemplars of a single category.

Statistical analysis

Curve-fitting was conducted using polynomials consisting of a linear term and a squared term ($y = \beta_0 + \beta_1 x + \beta_2 x^2$). This is a second order polynomial as the exponent is two. This approach is different from the models used by other researchers. Ruhland et al. (1995) employed a cubic logistic model, and Abbott-Smith and Behrens, and Ninio used a power

law function ($y = x^\beta$). Nonetheless, polynomials and power law curves can be used to describe similar learning curves. Moreover, polynomials offer an additional statistical advantage as they contain a nested linear model ($y = \beta_0 + \beta_1 x$). This enables one to statistically test the difference in goodness of fit between a straight line and a curve by comparing the full and nested model using an ANOVA. Though, to my knowledge, polynomials have not previously been used in the growth curve modelling of syntactic constructions, they have nonetheless been used in the exploration of lexical development (Tribushinina et al., 2014; van Veen et al., 2009).

All data and statistical analysis files may be found at: https://osf.io/e35jc/

Results

Testing the learning curve

The first analyses tested Ninio's claim that early construction use is characterised by a nonlinear growth curve. Age-in-months was plotted against cumulative number of verb types occurring in SSCs, focusing on the first two months after the construction appears, as investigated by Ninio (Figure 2). A second order polynomial model was fitted, shown with the solid line. For Juan, the polynomial model fitted the data significantly better than the linear model. For Lucía's data, the curve, though visually apparent, is weaker, with the p-value marginally exceeding the 0.05 cut-off.

An additional exploratory analysis was conducted for SSC token frequencies (instances of SSCs irrespective of verb type). This determined whether cumulative SSC tokens were also nonlinear, and could therefore impact on cumulative verb types. Indeed, when the data were plotted (Figure 2, second row), similar curves were observed, with a significantly better polynomial fit for Juan's data. Consequently, initial concavity for cumulative verb types may be an artefact of the temporal distribution of SSC tokens. With such a distribution, growth of verb types will appear nonlinear even when productivity,
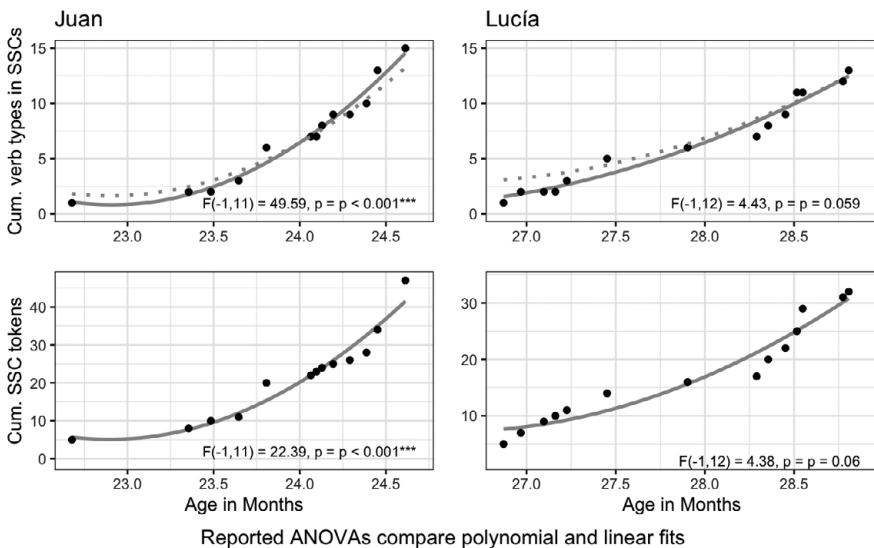

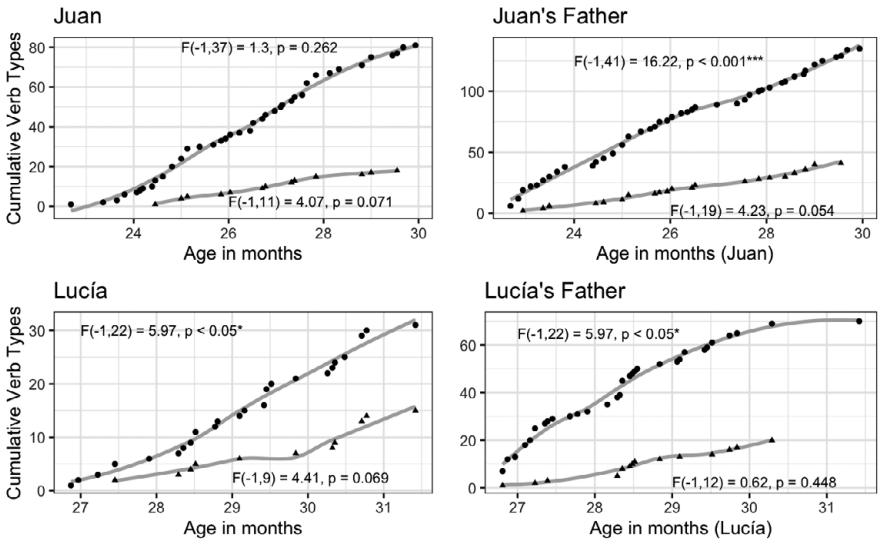**Figure 2.** Growth curves for SSCs during initial two months

Figure 3. Growth curves for all verbs versus hapax only, whole corpus

operationalised as the likelihood that each new SSC token will contain a new verb type, is held constant. To demonstrate this, cumulative verb tokens were multiplied by a constant such that the maximum value of this new variable equalled the maximum number of types within the two month period. A second order polynomial was conducted to model the relationship between age and this new measure. The trend lines are shown by the dotted line (first row). The trend lines for this new metric closely mirror the trend lines for cumulative verb types. This demonstrates that non-linear growth for verb types can be explained via the distribution of SSC tokens.

To further investigate productivity, a growth curve was plotted for hapax verbs. Assuming that productivity is emergent this was hypothesized to be strongly concave. The results are shown in Figure 3 with cumulative types for hapax verbs plotted against all verbs combined. The graphs show the data for the entire sampled period, in order to provide enough hapax verbs for analysis. Data for both children's fathers, the primary adult interlocutors, are included for comparison. To aid visual interpretation the trend line is drawn with the R loess function as it provides a close fit to the data points. However, statistical models, as previously, compared a linear and second order polynomial fit.

A statistically significant curve was detected only for the fathers' data (all verbs). With a negative coefficient for the squared term, these indicated a convex trend, i.e. curving upwards from a straight line. The curvature evident for the first two months of Juan's data is now undetected by the models, presumably because the initial curved phase is short in relation to the period sampled.

No significant curvature was identified for the hapax verbs in either children or adults. However, in contrast with the adult data, hapax verbs appeared relatively late in the child data. For Juan, the lag between the first SSC token and the first hapax verb (*juntar* = 'join') was 1.8 months. This contrasts with a mean lag between subsequent hapax productions of 0.5 months. For Lucía, the lag between the first SSC token and the first hapax verb was 0.6 months compared with a mean subsequent lag of 0.4 months.

Analysis of lexical distributions

Figure 4 shows the log-transformed rank (most frequent verb, second most frequent verb etc.) plotted against log-transformed frequency. The straight line indicates a Zipfian distribution. Where space allows, individual verbs have been listed for each rank. Translations, coded by CLAN, are broad, and not intended to convey the contextually-dependent meaning. Numerals in grey indicate the number of verbs observed for each rank.

Distributions are broadly similar across children. For both, *caer* ('fall') is the most frequent verb, and *romper* ('break') comes in second place for Juan and fifth place for Lucía. *Ir* ('go') is the third most frequent verb for Juan, and the sixth most frequent verb for Lucía. Interestingly, for Lucía, *caer* is super-Zipfian, falling above trend line.

To investigate the degree of lexical overlap, correlations between orders of occurrence were conducted using Kendall's method, as employed by Ninio (2005b, 2006). Results are displayed in Tables 1 and 2.

There was relatively little lexical uniformity across the corpora, with the exception of *caer* (fall) which came in the top three across all four corpora, and within the top two for Juan and Lucía's data. There were relatively few significant correlations between the rank orders across corpora (Table 2). In particular, the key correlation between the two child corpora was not significant. This is consistent with Ninio's claims.

A further exploratory procedure, which ranked verbs by overall frequency, was then performed. In addition, it was decided to investigate whether frequencies in child-context corpora, i.e. utterances from children or adults interacting with children, were consistent with frequencies in adult corpora based on online written materials. This explored whether the lexical characteristics of early verbs in SSCs reflect factors specific to child contexts. For example, children may be exploiting a semantic prototype, or alternatively, certain types of events may be more likely in child contexts. Consequently, a search was conducted in the Esteten18 corpus in the SketchEngine search engine. This is an approximately 18 billion word corpus of Spanish obtained from internet sources including Spanish Wikipedia. The corpus was searched using a Corpus Query Language (CQL).
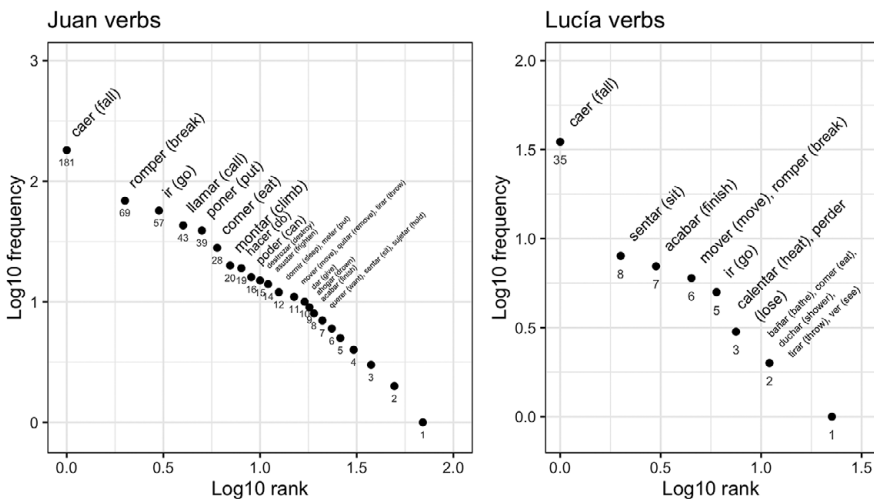

Figure 4. Lexical distributions of verb types

Table 1. Top ten 10 SSC verbs by order of occurrence

Juan	Lucía	Juan Adult	Lucía Adult
*caer (fall)	mover (move)	***caer (fall)**	**comer (eat)**
poner (put)	***caer (fall)**	*poner (put)*	olvidar (forget)
chocar (collide)	sentar (sit)	*tirar (throw)*	***caer (fall)**
ir (go)	abrir (open)	acordar (remember)	*acabar (finish)*
tirar (throw)	parar (stop)	llamar (call)	subir (rise)
ver (see)	*ir (go)*	*romper (break)*	tocar (touch)
romper (break)	**comer (eat)**	bañar (bathe)	**ver (see)**
comer (eat)	**ver (see)**	coger (take)	mirar (look)
dormir (sleep)	despertar (wake up)	hundir (sink)	tomar (take)
acabar (finish)	meter (put)	abrir (open)	estropear (spoil)

Formatting reflects number of corpora where verb appeared in the top 10: ***4 corpora**, **3 corpora**, *2 corpora*, 1 corpus

Table 2. Correlation matrix comparing SSC verb orders of occurrence across corpora

	Juan	Lucía	Juan Adult	Lucía Adult
Juan		−0.07	***0.45	**0.29
Lucía	−0.07		−0.1	0
Juan adult	***0.44	−0.1		0.07
Lucía adult	**0.28	0	0.07	

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Search strings are provided in the appendix. The top 10 verbs are shown in Table 3, and the correlation matrix is shown in Table 4.

In contrast to the sequential analysis, there was greater evidence for lexical-semantic homogeneity. Across all child-context corpora *caer* ('fall') was the most frequent verb. *Romper* ('break') and *ir* ('go') came in the top ten for all four child-context corpora. By contrast, distributions in the adult corpus (Estenten18) were very different from those in the child-context corpora. In the adult corpus, *poder* ('be able to') was the most frequent verb, but its highest position in the child-context corpora was seven. This may reflect the prevalence of impersonal sentences, e.g., *se puede comer bien en España* ('one can eat well in Spain'), in formal writing, which constitutes the main source for this corpus. It should be noted that many utterances with *se + poder* in the child data were reclassified, with *poder* interpreted as a control verb. This process was not possible for the Estenten corpus given its sheer size. Nonetheless, it can be seen that few of the top ten verbs for child corpora also appeared in the Estenten corpus.

This pattern of overlap is reflected in the correlation matrix. This indicates that speakers from the Aguado-Orea and Pine corpus, both children and adults, employed a similar range of verbs in SSCs, as demonstrated by rank order correlations. By contrast the lexical frequency data from the Estenten18 corpus correlated weakly with the individual speakers from the Aguado-Orea and Pine corpus.

Table 3. Top 10 SSC verbs by corpus

Juan	Lucía	Juan Adult	Lucía Adult	Estenten
*caer (fall)	***caer = fall**	***caer (fall)**	***caer (fall)**	poder (can)
*romper (break)	sentar (sit)	**llamar (call)**	**llamar (call)**	encontrar (find)
*ir (go)	*acabar (finish)*	***ir (go)**	pedir = ask for	**hacer (do)**
llamar (call)	mover (move)	***romper (break)**	***ir (go)**	tratar (treat)
poner (put)	***romper (break)**	*comer (eat)*	*comer (eat)*	***ir (go)**
comer (eat)	***ir (go)**	**poner (put)**	**poner (put)**	*ver (see)*
montar (climb)	calentar (heat)	**hacer (do)**	decir	dar (give)
hacer (do)	perder (lose)	**poder (be able to)**	***romper (break)**	realizar (realise)
poder (be able to)	bañar (bathe)	*ver (see)*	*acabar*	deber (must)
destrozar (destroy)	comer (eat)	quedar (stay)	hacer	convertir (change)

Formatting reflects number of corpora where verb appeared in the top 10: ***4 corpora**, **3 corpora**, *2 corpora*, 1 corpus (none identified across 5 corpora)

Table 4. Correlation matrix comparing SSC verb frequency distributions across corpora

	Juan	Lucía	Juan Adult	Lucía Adult	Estenten
Juan		*0.35	**0.60	***0.46	**0.27
Lucía	*0.35		*0.29	*0.40	0.11
Juan adult	***0.60	*0.29		***0.43	***0.24
Lucía adult	***0.46	*0.40	***0.43		**0.23
Estenten	**0.27	0.11	***0.24	**0.23	

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Discussion

The study investigated two claims of Ninio's Learning Curve Hypothesis; that generalisation is instant, and not semantically constrained. Though her empirical claims were supported, a number of exploratory analyses challenge her interpretation.

Are SSCs productive from the outset?

Firstly, the study replicated the finding that early syntactic development exhibits a learning curve based on a continuous process. In Juan's data the polynomial model fitted the data significantly better than a linear model. For Lucía, the curve was more moderate resulting in a p-value approaching significance. Nonetheless, according to research by Domínguez (2003) Lucía may be too old for us to capture the emergence of SSCs. Therefore, her weaker curve may indicate that she has moved, or is moving beyond, the initial stages of SSC-learning, characterised by a learning curve, and into a more linear

stage. In contrast, no concavity was present in the adult data, with the parents in Figure 2 exhibiting significant convexity, characterised by a steep initial slope, which then weakens.

While an early learning curve was observed, exploratory further analyses queried the claim that this relates to productivity. Cumulative SSC tokens, irrespective of verb types, exhibited a similar learning curve characterised by decreasing lags between tokens. Given this temporal distribution we would expect a curve for verb types, even if the probability of a new verb appearing in the verb slot is held constant. In fact, when SSC tokens were multiplied by a constant, the trend line closely matched the curve based on verb type frequencies, as shown in Figure 2.

This analysis indicates that the main driver of the learning curve may be the accelerating token frequencies. This is unrelated to productivity, as it is not influenced by variation in the verb slot. One potential explanatory mechanism is retrieval-based practice, sometimes called the "testing effect" (see Roediger & Karpicke, 2006 for a literature review). According to this framework, the robustness of a representation in long-term memory increases with each retrieval attempt, whether successful or not (Kornell & Vaughn, 2016). Initial retrievals will be difficult, with a low chance of success, but will become more successful over time, as each retrieval attempt boosts the probability of subsequent retrievals. We can explain the observed data via retrieval-based practice if we assume that SSCs, and indeed all grammatical constructions, have a baseline frequency, i.e. the rate at which they naturally occur in a discourse. This is demonstrated by the relatively straight trend line for cumulative verb types in both children and adults when the entire corpus is analysed (Figure 3). The learning curve constitutes a gradual approximation of this baseline frequency as, via retrieval-based practice, children become better able to retrieve the target construction.

This account contrasts with Ninio's, which assumes that the growth curve reflects a cognitive process with nonlinear dynamics. She refers to the power law of practice, whereby the time taken to perform a complex task decreases in a nonlinear fashion. However, there is little evidence that this effect governs language learning. Power law phenomena are typically observed in numerosity judgement tasks whereby participants see pairs of diagrams containing a haphazard array of dots, and must say whether the second diagram contains more or fewer dots (e.g., Rickard, 1999). Such tasks arguably bear little relationship to the process of retrieving and productively using grammatical constructions. By contrast, retrieval-based practice effects have been found in lexical learning tasks, e.g., Leonard and Deevy's (2020) study of word learning by children with language disorders.

Given the confounding nature of SSC token frequencies, an exploratory analysis of hapax verbs was conducted. This is based on arguments that hapax legomena are indicators of productivity (e.g., Baayen & Lieber, 1991). An analysis of cumulative types did not find a nonlinear trend, with the line appearing straight on a plot. Given that linearity, whether of types or tokens, is a feature of later language development, this might indicate full productivity from the outset. However, for both children, hapax verbs were relatively late to emerge, with larger lags between the first non-hapax verb and the first hapax verb than between subsequent hapax verbs. This indicates a delayed onset of productivity. Lucía's lag was smaller than Juan's which, again, suggests she may be moving beyond the initial stage of SSC-learning. In contrast with the children, the adults produced hapax verbs early, suggesting that this lag is a characteristic of child language development (Figure 3).

Are generalisations semantically constrained?

The study also investigated Ninio's claim that initial generalisations are not semantically constrained. This was based on a relative lack of semantic homogeneity and also semantic fit between the meaning of the verb and the construction. There was indeed limited semantic homogeneity among early occurring verbs. Nonetheless, *Caer* ('fall') was prominent, appearing first in Juan's data and second in Lucía's. Moreover, raw frequency data exhibited substantial homogeneity across both children and adult caregivers. Across all four speakers, *caer* was the most frequently used verb in the SSC. For Lucía the frequency of *caer* was 'super-zipfian', falling above the line of best fit. *Caer* was also the most frequent unexpected event verb in Jackson-Maldonado et al.'s (1998) corpus. Two other verbs were frequently produced by both children: *romper* ('break': 2[nd] most frequent SSC verb for Juan and 5[th] most frequent for Lucía) and *ir* ('go': 3[rd] and 6[th] most frequent).

Caer is consistent with Maldonado's (1998) semantic prototype focusing on a pivotal moment of change. Two other highly frequent verbs are also consistent with this prototype; *romper* ('break') and *ir* ('go'). This may indicate the semantic organisation of early SSCs. Moreover, only one of these verbs (*ir* = 'go') appears in the top ten for the Estenten corpus, in position five. Consequently verbs focusing on a moment of change may be a characteristic of child contexts, and may therefore reflect cognitive processes underpinning acquisition.

Nonetheless, some data conflict with early semantic generalisation. Many early or frequent SSC verbs were not consistent with the above semantic prototype. For example, *sentar* (sit, Lucía's second most frequent verb), describes a change-of-state which is not sudden or irreversible. *Llamar* ('call'), which was Juan's fourth most frequent verb and Lucía's second, is idiosyncratic, and does not conform to any subtype of verbs typically appearing in SSCs. When combined with the SSC it is used as a naming verb, e.g., *Me llamo Juan* = 'my name is John'. Moreover, the preponderance of verbs such as *caer* ('fall'), *romper* ('break') and *ir* ('go') does not, by itself, indicate semantic generalisation. Such verbs may be common for ecological reasons as the events they denote may be common in childhood play. If certain types of SSCs, e.g., those describing a change of state, are more frequent, change-of-state verbs describing events which frequently occur in child contexts will also occur more frequently. This does not necessarily imply that such verbs play a causal role in the development of a semantic category.

Study limitations

For both theoretical and practical reasons, the study overlooked the reflexive-middle distinction. Nonetheless, the SSCs exhibited two statistical properties identified by studies of construction learning; an initial learning curve, and a zipfian distribution of verb types. The latter mirrors distributions found in grammatical constructions such as ditransitives (Goldberg et al., 2004). If Spanish reflexives and middles were genuinely separate constructions in child learners, one might not expect these previously observed patterns to have been replicated.

A further issue is the reliability of growth curve analyses. Polynomials are useful for statistically testing curvature, but are too flexible in their capacity to model various nonlinear trends. For example, significant convex profiles were observed for both fathers, a pattern which is difficult to theoretically motivate. For this reason, Abbot-Smith and Behrens (2006) chose a power law function, which is more constrained in its ability to model different curves. The chosen time window likewise determines whether curvature

is identified. For example, Juan's initial growth curve in the first two months was rendered non-significant when a larger time window was chosen. Due to the complexities of curve fitting, arbitrary decisions, e.g., the length of the time window, will often need to be made.

Given that a key theoretical claim is the presence or absence of continuity, it may be theoretically informative to fit discontinuous models. Regression Discontinuity Designs (RDDs) can do this, by incorporating a term which is included only once the y variable exceeds a particular value. However, these are likewise influenced by arbitrary decisions: in particular, where to place the cut-off.

Many analyses were exploratory, and therefore a potential source of bias. Nonetheless, all exploratory analyses were theoretically motivated. The analysis of cumulative SSC tokens was conducted to explore this factor as a potential confound. The analysis of hapax legomena was motivated by substantial data suggesting their sensitivity to productivity. Finally, the analysis of raw lexical frequencies provided a distributional measure which was relatively unconfounded by sampling issues.

Finally, it should be acknowledged that many of the hypothesised developmental trends were weaker in Lucía's data. This may reflect her relatively advanced age which, according to Domínguez (2003), places her outside the initial learning stage. If learning curves are only evident in the initial stages of learning, a lag of even a few months may greatly impact research findings. Nonetheless, with only one study providing normative data on SSCs (Domínguez, 2003) we cannot be absolutely certain that Lucía's sampling period occurred too late.

Conclusions and future directions

The study provided further evidence that early construction-learning exhibits a learning curve. As observed by Abbot-Smith and Behrens (2006), this contrasts with linear trajectories among adults. However, unlike previous studies it found that growth curves do not necessarily relate to productivity, operationalised as the number of verbs in the verb slot. Analyses of SSC tokens found that the learning curve, assessed using verb types, may merely reflect an accelerating incidence of tokens of the target construction. This suggests an alternative mechanism, retrieval-based practice. Given this, the analysis of hapax legomena may provide better evidence regarding productivity. The relative lateness of hapax verbs in the child data, particularly Juan's, suggests that generalisation is emergent, rather than instantaneous, though clearly this finding is in need of further replication. Verb frequency data indicate the possible existence of a semantic prototype organised around verbs which signal a pivotal moment of change. However, it is acknowledged that verb frequency data, by themselves, provide limited evidence for causal mechanisms.

To test the above claims, future analyses of growth curves should adopt similar methods to the current study – namely, a comparison of type and token frequencies, an exploration of hapax legomena, and an analysis of both sequential and raw lexical frequencies. If further studies find that token frequencies mirror type frequencies, this would further undermine the latter as a measure of productivity and provide evidence for alternative accounts, e.g., retrieval-based practice. Regarding hapax legomena, the onset of the first hapax verb may be the key variable. Lexical analyses may place more emphasis on sequential data if sampling regimes are sufficiently dense.

Given the complexities of analysing growth curve data, distributional analyses of corpora could be complemented by other methods. Experimental studies might

investigate children's ability to generalise SSCs to novel verbs. By varying the meaning and aspectual properties of these novel verbs one could evaluate semantic generalisation using a cross-sectional approach to investigate developmental trends. Similarities across corpus and behavioural data would provide converging evidence for key findings. Finally, the analyses in the current study abstracted away from sequential patterns within corpora. For example, did productive use of SSCs reflect models provided by the adult within recent conversational turns? Controlling for this would enable a more fine-grained analysis of generalisation.

References

Abbot-Smith, K., & **Behrens, H.** (2006). How Known Constructions Influence the Acquisition of Other Constructions: The German Passive and Future Constructions. *Cognitive Science*, **30**(6), 32. https://doi.org/10.1207/s15516709cog0000_61

Abbot-Smith, K., & **Tomasello, M.** (2006). Exemplar-learning and schematization in a usage-based account of syntactic acquisition. *The Linguistic Review*, **23**(3), 275–290. https://doi.org/10.1515/TLR.2006.011

Aguado-Orea, J., & **Pine, J. M.** (2015). Comparing Different Models of the Development of Verb Inflection in Early Child Spanish. *PLOS ONE*, **10**(3), e0119613. https://doi.org/10.1371/journal.pone.0119613

Akhtar, N. (1999). Acquiring basic word order: Evidence for data-driven learning of syntactic structure. *Journal of Child Language*, **26**, 339–356.

Alexiadou, A., & **Doron, E.** (2012). The syntactic construction of two non-active Voices: Passive and middle. *Journal of Linguistics*, **48**(1), 1–34. https://doi.org/10.1017/S0022226711000338

Baayen, H., & **Lieber, R.** (1991). Productivity and English derivation: A corpus-based study. *Linguistics*, **29**(5). https://doi.org/10.1515/ling.1991.29.5.801

Barðdal, J. (2008). *Productivity: Evidence from case and argument structure in Icelandic: Vol. Constructional Approaches to Language 8.* John Benjamins Publishing.

Brown, A. L., & **Kane, M. J.** (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology*, **20**(4), 493–523.

Calude, A. S. (2017). Testing the boundaries of the middle voice: Observations from English and Romanian. *Cognitive Linguistics*, **28**(4). https://doi.org/10.1515/cog-2016-0046

Childers, J., & **Tomasello, M.** (2001). The role of pronouns in young children's acquisition of the English transitive construction. *Developmental Psychology*, **37**(6), 739–748. https://doi.org/10.1037/0012-1649.37.6.739

Dąbrowska, E. (2008). The effects of frequency and neighbourhood density on adult speakers' productivity with Polish case inflections: An empirical test of usage-based approaches to morphology. *Journal of Memory and Language*, **58**(4), 931–951. https://doi.org/10.1016/j.jml.2007.11.005

Dąbrowska, E., & **Lieven, E.** (2005). Towards a lexically specific grammar of children's question constructions. *Cognitive Linguistics*, **16**(3), 437–474. https://doi.org/10.1515/cogl.2005.16.3.437

Domínguez, L. (2003). Interpreting Reference in the Early Acquisition of Spanish Clitics. In S. Montrul & F. Ordonez (Eds.), *Linguistic Theory and Language Development in Hispanic Languages: Papers from the 5th Hispanic Linguistics Symposium and the 4th Conference on the Acquisition of Spanish and Portuguese* (pp. 212–228). Cascadilla Press, MA.

Faltz, L. (1977). *Reflexivization: A study in universal syntax.* Doctoral dissertation, University of California, Berkley, 1977.

Fillmore, C. (1988). *The Mechanisms of 'Construction Grammar'.*

Fitts, P. M., & **Posner, M. I.** (1967). *Human Performance.* Brooks Cole.

Gertner, Y., **Fisher, C.**, & **Eisengart, J.** (2006). Learning words and rules: Abstract knowledge of word order in early sentence comprehension. *Psychological Science*, **17**(8), 684–691. https://doi.org/10.1111/j.1467-9280.2006.01767.x

Goldberg, A. E. (2016). Partial productivity of linguistic constructions: Dynamic categorization and statistical preemption. *Language and Cognition*, **8**(3), 369–390. https://doi.org/10.1017/langcog.2016.17

Goldberg, A. E., **Casenhiser, D. M.**, & **Sethuraman, N.** (2004). Learning argument structure generalizations. *Cognitive Linguistics*, **15**(3), 289–316.

Grimshaw, J. (1982). On the lexical representation of Romance reflexive clitics. In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations* (pp. 87–146). MIT Press.

Hopper, P. J., & **Thompson, S. A.** (1980). Transitivity in grammar and discourse. *Language*, **56**(2), 251–299.

Ibbotson, P. (2013). The Scope of Usage-Based Theory. *Frontiers in Psychology*, **4**. https://doi.org/10.3389/fpsyg.2013.00255

Jackson-Maldonado, D., **Maldonado, R.**, & **Thal, D. J.** (1998). Reflexive and middle markers in early child language acquisition: Evidence from Mexican Spanish. *First Language*, **18**(54), 403–429.

Kemmer, S. (1993). *The middle* voice (Vol. 23). John Benjamins Publishing.

Kemmer, S. (1994). Middle Voice, Transitivity and the Elaboration of Events. In B. A. Fox & P. J. Hopper (Eds.), *Voice: Form and Function* (pp. 179–229). John Benjamins Publishing.

Kornell, N., & **Vaughn, K. E.** (2016). How Retrieval Attempts Affect Learning. In *Psychology of Learning and Motivation* (Vol. **65**, pp. 183–215). Elsevier. https://doi.org/10.1016/bs.plm.2016.03.003

Langacker, R. W. (2009). A dynamic view of usage and language acquisition. *Cognitive Linguistics*, **20**(3), 627–640. https://doi.org/10.1515/COGL.2009.027

Leonard, L. B., & **Deevy, P.** (2020). Retrieval Practice and Word Learning in Children With Specific Language Impairment and Their Typically Developing Peers. *Journal of Speech, Language, and Hearing Research*, **63**(10), 3252–3262. https://doi.org/10.1044/2020_JSLHR-20-00006

MacWhinney, B. (2000). *The Child Language Data Exchange System (CHILDES).* Laurence Erlbaum Associates.

Maldonado, R. (1992). *Middle voice: The case of Spanish se.* [Dissertation]. University of California, San Diego.

Maldonado, R. (2008). Spanish middle syntax: A usage-based proposal for grammar teaching. In S. De Knop & T. De Rycker (Eds.), *Cognitive Approaches to Pedagogical Grammar* (Vol. **9**, pp. 155–196). de Gruyter Mouton.

Mendikoetxea, A. (1999). Construcciones con" se": Medias, pasivas e impersonales. In I. Bosque & V. Demonte (Eds.), *Gramática descriptiva de la lengua española* (pp. 1575–1630). Espasa Calpe.

Moreno, C. (2021). Is there really an aspectual *se* in Spanish? *Folia Linguistica*, **55**(1), 195–230. https://doi.org/10.1515/flin-2020-2074

Moreno, C. de B. (2015). *Las Construcciones con 'se' desde una Perspectiva Variacionist y Dialectal / Constructions with 'se': A Variationist and Dialectal Approach.*

Ninio, A. (1999). Pathbreaking verbs in syntactic development and the question of prototypical transitivity. *Journal of Child Language*, **26**(3), 619–653. https://doi.org/10.1017/S0305000999003931

Ninio, A. (2005a). Accelerated learning without semantic similarity: *Indirect objects*. **16**(3), 531–556. https://doi.org/10.1515/cogl.2005.16.3.531

Ninio, A. (2005b). Testing the role of semantic similarity in syntactic development. *Journal of Child Language*, **32**, 35–61. https://doi.org/10.1017/S0305000904006713

Ninio, A. (2006). *Language and the learning curve: A new theory of syntactic development.* Oxford University Press.

Ornat, S. L. (1994). *La Adquisición de la Lengua Española.* Siglo Veintiuno.

Pierrehumbert, J., & **Granell, R.** (2018). On Hapax Legomena and Morphological Productivity. *Proceedings of the Fifteenth Workshop on Computational Research in Phonetics, Phonology, and Morphology*, 125–130. https://doi.org/10.18653/v1/W18-5814

Pine, J. M., & **Lieven, E. V. M.** (1997). Slot and frame patterns and the development of the determiner category. *Applied Psycholinguistics*, **18**(2), 123–138. https://doi.org/10.1017/S0142716400009930

Rice, M. L., Wexler, K., & Cleave, P. (1995). Specific Language Impairment as a period of extended optional infinitive. *Journal of Speech and Hearing Research*, **38**, 850–863.

Rickard, T. C. (1999). A CMPL alternative account of practice effects in numerosity judgment tasks. *Journal of Experimental Psychology: Learning Memory and Cognition*, **25**(2), 532–542. https://doi.org/10.1037/0278-7393.25.2.532

Roediger, H. L., & Karpicke, J. D. (2006). *Test-Enhanced Learning. Taking Memory Tests Improves Long-Term Retention.* 249–255. https://doi.org/10.1111/j.1467-9280.2006.01693.x

Rosen, C. (1988). *The relational structure of reflexive clauses: Evidence from Italian.* (D. M. Perlmutter & C. Rosen, Eds.; Vol. **2**). Harvard University.

Ruhland, R., & Geert, P. (1998). Jumping into syntax: Transitions in the development of closed class words. *British Journal of Developmental Psychology*, **16**(1), 65–95. https://doi.org/10.1111/j.2044-835X.1998.tb00750.x

Ruhland, R., Wijnen, F., & Van Geert, P. (1995). An exploration into the application of dynamic systems modelling to language acquisition. *Amsterdam Series in Child Language Development*, **4**, 107–134.

Tomasello, M. (1992). *First verbs, a case study in grammatical development.* Cambridge UP.

Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition.* Harvard UP.

Tribushinina, E., van den Bergh, H., Ravid, D., Aksu-Koç, A., Kilani-Schoch, M., Korecky-Kröll, K., Leibovitch-Cohen, I., Laaha, S., Nir, B., Dressler, W. U., & Gillis, S. (2014). Development of adjective frequencies across semantic classes: A growth curve analysis of child speech and child-directed speech. *Language, Interaction and Acquisition*, **5**(2), 185–226. https://doi.org/10.1075/lia.5.2.02tri

van Veen, R., Evers-Vermeul, J., Sanders, T., & van den Bergh, H. (2009). Parental input and connective acquisition: A growth curve analysis. *First Language*, **29**(3), 266–288. https://doi.org/10.1177/0142723708101679

Yang, C., Crain, S., Berwick, R. C., Chomsky, N., & Bolhuis, J. J. (2017). The growth of language: Universal Grammar, experience, and principles of computation. *Neuroscience & Biobehavioral Reviews*. https://doi.org/10.1016/j.neubiorev.2016.12.023

Zeldes, A. (2011). On the productivity and variability of the slots in German comparative correlative constructions. *Grammar & Corpora, Third International Conference,* Mannheim, 22, 2009.

Zeldes, A. (2012). *Productivity in argument selection: From morphology to syntax* (Vol. **260**). Walter de Gruyter.