A Dense Corpus Study of Past Tense and Plural Overregularization in English

Robert J. C. Maslen
Anna L. Theakston
University of Manchester,
United Kingdom

Elena V. M. Lieven
Michael Tomasello
Max Planck Institute for
Evolutionary Anthropology,
Leipzig, Germany

In the "blocking-and-retrieval-failure" account of overregularization (OR; G. F. Marcus, 1995; G. F. Marcus et al., 1992), the claim that a symbolic rule generates regular inflection is founded on pervasively low past tense OR rates and the lack of a substantive difference between past tense and plural OR rates. Evidence of extended periods of OR in the face of substantial correct input (M. Maratsos, 2000) and of an initial period in which nouns are more likely to be overregularized than verbs (V. A. Marchman, K. Plunkett, & J. Goodman, 1997) casts doubt on the blocking account and suggests instead an interplay between type and token frequency effects that is more consistent with usage-based approaches (e.g., J. Bybee, 1995; K. Köpcke, 1998; K. Plunkett & V. Marchman, 1993). However, previous naturalistic studies have been limited by data that account for only 1–2% of child speech. The current study reports analyses of verb and noun ORs in a dense naturalistic corpus (1 child, 2;00.12–3;11.06 [years;months.days]) that captures 8–10% of child speech and input. The data show (a) a marked difference in verb and noun OR rates; (b) evidence of a relationship between relative regular/irregular type frequencies and the onset and rate of past tense and plural ORs; (c) substantial OR periods for some verbs and nouns despite hundreds of correct tokens in child speech and input; and (d) a strong negative correlation between input token frequencies and OR rates for verbs and nouns. The implications of these findings for blocking and other accounts of OR are discussed.

KEY WORDS: language acquisition, morphology, overregularization, past tense, plural

Children's acquisition of English past tense and plural inflection has provided evidence for widely differing theories of morphological representation. The debate has tended to focus on the question of whether regular and irregular forms are processed by separate systems, as in rule-symbol models (Marcus et al., 1992; Pinker & Prince, 1988) or by a single associative mechanism, as asserted in connectionist (e.g., Plunkett & Marchman, 1991, 1993), competition (e.g., Kuczaj, 1977; MacWhinney, 1987) and schema-based (Bybee, 1995; Köpcke, 1998) accounts. Clearly, the outcome of this debate has considerable significance for the nature of language representation in both typically and nontypically developing children.

Symbolic accounts (e.g., Pinker & Prince, 1988) propose a dual mechanism in which regular forms are produced by a generative grammatical rule (plus -*ed*), while irregular forms are learned by rote and stored as individual memory traces. Marcus et al. (1992) refine this distinction with the "blocking-and-retrieval-failure" account. They claim

that following an initial period in which all past forms are rote learned and produced correctly, a child learns the regular affixation rule. From this point, the regular past tense is generated by the symbolic concatenation of stem forms and the tense marker *-ed*. Irregular past forms continue to be stored individually. An innate "blocking" process (seen as a link between the lexicon and the rule; Marcus et al., 1992; Pinker, 1999) inhibits the application of the rule to irregular stems with established past tense traces, but regular inflection is overextended to irregular stem forms without past tense traces, or with traces too weak to block the rule. As a memory trace becomes strong enough to be retrieved consistently, the blocking process is able to operate, and overregularizations (ORs) for that type die out. (This process of correct production followed by the onset of ORs and then by recovery to correct usage thus accounts for the characteristic U-shaped developmental curve.)

Marcus et al. (1992) found pervasively low OR rates in an extensive study of naturalistic data from the Child Language Data Exchange System (CHILDES) database (MacWhinney & Snow, 1985). Individual transcripts for 25 children yielded an average overall OR rate of 4.2%, and though 1 participant showed a far higher error rate (Abe, 24%), his data, at least for individual verbs, were considered "chaotic" and "haphazard" (Marcus et al., 1992, p. 52). These low rates were taken as evidence of "the occasional breakdown of a system built to prevent the error" (Pinker, 1999, p. 221). In other words, the system is biased in favor of correctly inflected forms. For the blocking account to be meaningful, in addition to low OR rates there must also be a reasonable limit on the number of occasions on which a child has to hear or use a correct irregular past tense before that form is sufficiently entrenched to facilitate the blocking process (Maratsos, 2000). This would seem to rule out extended periods of OR for irregular types with higher token frequencies.

Extending the blocking account to English noun plural morphology, Marcus (1995) found that past tense and noun plural (e.g., *foots*, *mans*) regularizations were not substantively different from each other, with each category showing a U-shaped curve and similar low OR rates. He noted that the contrasting relative proportions of regular to irregular types in the verb and noun systems of English should lead single-mechanism accounts to predict a difference in the two patterns, since in these accounts OR occurs when irregulars are attracted by the shared connection weight of regular forms (e.g., Rumelhart & McClelland, 1986) or by a strongly represented regular schema (e.g., Bybee, 1995), and high type frequency is critical in creating that attraction. The fact that he found no such differentiation seemed to support the conclusion that "regular inflection is independent of frequency" (Marcus, 1995, p. 457).

There are, however, some grounds for questioning the principal empirical foundations of the blocking account. Claims of pervasively low OR rates are contested by Maratsos (2000). He reanalyzes past tense regularizations in a subset of the Marcus et al. (1992) transcripts and identifies elements of the previous study that might have produced misleading results. The first of these is the calculation of the OR rate as an average across all forms. Echoing Rubino and Pine (1998) in their study of Portuguese person marking, he observes that a gross measure like this has the tendency to obscure even quite significant variations in the behavior of individual types (see also Behrens, 2001, and Köpcke, 1998, on the acquisition of the German plural, and Dabrowska, 2001, on the acquisition of the Polish genitive). In particular, types with low token frequencies are likely to be swamped by types with high frequencies in overall calculations. Breaking the data down into frequency groups overcomes this effect and reveals markedly higher OR rates, especially for the lowest frequency types, than the overall figure. For example, Adam's (Brown, 1973) overall error rate of 3.6% masks a rate of 54% for verbs produced 1–9 times in his data (Maratsos, 2000). This is despite the fact that the verbs in this group account for 22 of his 53 past tense irregular types (or 41.5% of them). Marcus et al. in fact exclude verbs with token frequencies lower than 10 from their analyses on the grounds that the change in OR rate induced by a single OR token is disproportionately large, making scores for these forms unreliable. Maratsos argues that since data for these verbs are grouped, they have greater statistical reliability than the data for any of the individual types alone, and that, in any case, all accounts of OR would predict low OR rates for all but lower frequency forms (Maratsos, 2000, p. 200). More important, forms sampled infrequently in a very thin corpus almost certainly represent much higher frequencies in actual speech.

A second innovation in the Maratsos (2000) paper is his attempt to estimate the real levels of input that children are hearing while still making ORs, a calculation with obvious implications for the blocking account, in which substantial input is expected to drive the child's recovery to correct usage. Using the heuristic that input and output token frequencies are roughly equivalent, Maratsos extrapolates from the 1–2% density of his data to a notional 100% sample, and he finds that there is robust OR for some verbs in periods during which his participants would have produced (and thus, he not unreasonably presumes, heard) hundreds, if not thousands, of correct tokens. Clearly, this needs to be substantiated with better naturalistic data. In principle, however, the higher OR rates for grouped low-frequency forms and robust OR production revealed by Maratsos pose a serious challenge to the adequacy of the blocking account. These findings are less problematic for single-mechanism

accounts, in which token frequency plays the role of entrenching correct irregular forms and thus overcoming ORs driven by the superiority of regular type frequency. As these accounts do not envisage an innately specified process operating to preempt ORs, more prolonged interference from the regular majority is not ruled out, as it would seem to be by blocking.

A further issue arises from evidence that there may be a significant difference in the patterns of past tense and noun plural ORs that children produce. As we have seen, the blocking account does not ascribe a role to frequency and predicts that there will be no discrepancy in OR rates or patterns over time. Marchman, Plunkett, and Goodman (1997) cast some doubt on the veracity of this claim, however, with their findings that children are in fact more likely to overregularize nouns than verbs—as a proportion of types—between the ages of 1;5 (years;months) and 2;6, and that the onset of ORs may be earlier for nouns than for verbs. As these results are based on parental report data there can be no comparison of token-frequency OR rates with those shown by Marcus (1995). However, this further highlights the tendency of OR rate scores aggregated across forms and over extended time periods to mask certain aspects of morphological behavior.

To sum up, symbolic dual-mechanism accounts of OR, such as Marcus et al. (1992) and Marcus (1995), cite in their support evidence of low OR rates and a similarity between OR patterns for verbs and nouns. Maratsos (2000) and Marchman et al. (1997) challenge the blocking account with evidence of higher OR rates, more enduring OR, and a differentiation in past tense and plural OR patterns. The principal aim of the present study is to test this contrasting set of empirical findings using a uniquely rich data source, which accounts for some 8–10% of the participant's speech and input between 2;0 and 3;2, augmented by a similarly dense sample from 1 week of each month from 3;2 to 4;0. These much denser data should provide a more authoritative contribution to a field in which researchers have previously noted the limitations of thin naturalistic corpora. As the blocking account explicitly rules out an effect of type frequency on the onset and rate of OR (Marcus, 1995), and as contrasting accounts ascribe a pivotal role to the interplay of type and token frequencies (e.g., Bybee, 1995), a further aim of the study is to examine whether there is any evidence in our data for such an effect.

Method
Participants

The data consist of 330 hour-long recordings, made over 2 years (child's age: 2;00.12 [years;months.days] to 3;11.06), of one child, "Brian," interacting at home with

his mother. The interaction between speakers was entirely spontaneous, and the contexts (almost exclusively mealtimes and play sessions with toys) and timings of the recordings were dictated by the mother.

Child and Input Data

From 2;0 to 3;2, Brian was recorded as close as possible to a regimen of "5 days on, 2 days off." For every five sessions, one video- and four audio-recordings were made. Maratsos (2000) estimates the "talk week" of a child of this age to be around 40 hours. Allowing for the few inevitable gaps in the recordings, this gives an estimated sample density of around 8–10% of Brian's total production. From 3;3 to 3;11, recordings were made on 4 or 5 consecutive days out of each month, thus retaining the earlier density of data for each week's sample.

Research assistants transcribed all of the tapes using standard CHAT procedures (MacWhinney & Snow, 1990). Transcription was subsequently checked twice by trained assistants. First, each transcript was "morphemized" to allow the subsequent application of the MOR automatic parser of the CHILDES system, and second, each utterance was linked to the sound file. Transcripts were checked for accuracy at each stage of the process. The 287 one-hr-long recordings from 2;0 to 3;2 were grouped into samples of 7, giving a total of 41 samples. The 4 or 5 recordings made during 1 week each month from 3;3 to 3;11 were also grouped together, these additional 9 samples making an overall total of 50. Where reference is made below to these grouped samples, the first recording of the sample is cited. Reference is also made to individual recordings (identified by age—e.g., 2;00.12), as well as to data grouped into months (2;0, 2;1, 2;2, etc.).

Procedure

Computerized language analysis (CLAN; MacWhinney, 2000) was used to extract all child utterances containing correct and regularized past tense forms of irregular main verbs, as well as correct and regularized plural forms of irregular nouns. Correct regular past tense and plural forms were also extracted for type-frequency comparisons. Zero-marked forms were not included in the analysis. Regularizations for verbs included both stem plus -*ed* and past plus -*ed* forms, and for nouns both singular plus -*s* and plural plus -*s* forms. Certain verb types were excluded from analysis, broadly on the basis of criteria set out by Maratsos (2000). Thus, *do*, *have*, *get*, copula *be* and no-stem-change verbs (e.g., *hit*) and irregulars with a permissible regular past (e.g., *spoil, spoilt, spoiled*) were not included in the present study. (Marcus et al., 1992, excluded *be, do,* and *have* from analysis of Abe's data but included these

[as main verbs] and no-stem-change verbs for Adam, Eve and Sarah.) In the case of the nouns, as with Marcus (1995), plurals with a permissible regular form (e.g., *fish*) and plurals of mass nouns were excluded, while compounds with irregular noun heads (e.g., *policemen*) were not. Unlike Marcus, we excluded compounds with irregular pronoun heads (e.g., *yourselves*), though there were few of these. The plurals of *roof* and *woman* were excluded, since only overregularized forms were recorded in the mother's data. Similarly, the plural of *leaf* was excluded, as a large number of regularized forms in the mother's data made it impossible tell whether the child's forms were actually ORs.

All the mother's utterances containing past tense and plural forms included in the child data were extracted. In contrast with Marcus, who included forms such as *a foots* and *a feet* (Marcus, 1995, p. 451), the ambiguity of the child's intention to convey a plural meaning was used as a criterion for excluding individual utterances, in line with the methodology of the past tense studies. Where it was not clear from an utterance itself that the child intended to convey a past tense/plural meaning, this was verified from the context. Thus, ambiguous utterances (e.g., verb phrases missing an overt subject) were retained if they clearly occurred in an appropriate context, but they were otherwise excluded. For example, the child utterance *bought some daffodils* (2;11.06) was judged to be an acceptable past tense utterance when it was given in response to the mother's question *what did we buy the other day?* Single word utterances, excluded by Marcus et al. (1992), were retained or excluded on the same grounds as other ambiguous utterances, although there were very few of these. Types with frequencies of fewer than 10 tokens were retained in the analysis; this allows comparison with the Maratsos (2000) findings, while not significantly affecting overall OR rate calculations. Past tense coding decisions (made by the first author) were checked by a second coder (the second author) for three samples taken from the corpus at 2;4, 3;0 and 3;9 (giving a total of some 15% of all past tense tokens). Past tense forms returned in initial searches were coded as correct irregulars, overregularizations, and correct regulars. Items were excluded from analysis if they were found to be immediate repetitions, imitations of an adult utterance, or perfective/adjectival uses, or were deemed to be ambiguous in some other regard. For the combined sample, coders agreed 91.19% of the time (Cohen's κ = .86, *p* < .0005).

Brian's OR rates were calculated for the whole corpus, for each of the 50 aggregate samples, for individual verb and noun types, and for types grouped according to child token frequency. Correlations between the child's and mother's token frequencies, and between these token frequencies and OR rates, were also calculated. Further calculations were made of regular and irregular verb and noun type frequencies over time.

Results

Table 1 gives a breakdown by month of Brian's mean length of utterance (MLU; morphemes), together with verb and noun type and token frequencies. MLU scores from 2;1 to 3;11 place Brian within the normal range for Brown's stages I to V, though, particularly early on, he is toward the lower end of the scale (Brown, 1973). Input frequencies for irregular pasts and plurals correlate highly with those reported by Francis and Kučera (1982). For irregular pasts, $r(24) = .67$, $p = .01$; for irregular plurals, $r(10) = .60$, $p = .05$.

OR Rates

Overall Rates: Verbs

Brian produced 1,263 past tense tokens of 52 irregular verbs in the 2-year period, with the mother's token frequencies for the same verbs being 10,081. At least one overregularized token was recorded for 24 of these verbs, leaving 28 types for which only correct past tense tokens were sampled. For all irregular verbs, Brian's correct past tense token frequencies varied from 1 to 218 ($M = 24.29$). Input token frequency varied from 4 to 1,796 ($M = 193.87$). The corpus contains 107 OR tokens; 91 of these are stem plus *-ed* (e.g., *goed*) and 16 past plus *-ed* (e.g., *droved*).

The rate of OR, following Marcus et al. (1992), was calculated as the percentage of all child irregular past tense tokens (correct and OR forms, not including bare stems) constituted by OR tokens. For the whole sample period (2;00.12–3;11.06) and for all irregular verbs, the overall OR rate was 7.81%. However, some verbs may have recovered from OR during the recording period. Clearly, all verbs recover eventually, and to continue to include tokens recorded after the OR production period has ended underestimates the true OR rate; all verbs could be shown to have negligible OR rates if recorded for long enough. Verbs were therefore assumed to have recovered once they had reached the second of two consecutive samples (a sample being seven recordings up to 3;02.11 and four or five thereafter) in which all recorded tokens were correct, and after which no further OR tokens were recorded. Under this definition of recovery, a second rate was calculated using revised values for *bring*, *find*, *go*, and *say*, giving an increase in the overall OR rate to 9.56%. Both the simple and revised rates are slightly higher than the overall figure reported by Marcus et al. (Marcus et al.'s median OR rate for participants with individual transcripts was 2.5%, with an average of 4.2%.) The revised rate would place Brian in third place out of the 25 individual children analyzed in that study, roughly equal to April (Higginson, 1985), though some way behind the dramatically overregularizing Abe (Kuczaj, 1977), at 24%.

Table 1. Brian's irregular past tense and plural types and tokens (monthly aggregates).

Age	MLU	Verbs (past tense)				Nouns (plural)			
		Reg Types	Irr Types	Reg Toks	Irr Toks	Reg Types	Irr Types	Reg Toks	Irr Toks
2;1	1.62	2	1	2	1	36	1	355	1
2;2	1.64	0	0	0	0	35	1	626	31
2;3	2.02	0	0	0	0	78	2	733	24
2;4	1.94	0	4	0	6	123	4	632	23
2;5	2.01	4	9	9	12	135	5	614	20
2;6	1.94	4	4	5	15	139	6	491	20
2;7	2.01	12	17	28	73	139	8	478	31
2;8	2.19	17	20	34	79	187	7	782	18
2;9	2.39	17	18	29	119	219	8	736	14
2;10	2.53	25	20	47	92	204	7	750	36
2;11	2.74	23	20	34	83	206	8	688	59
3;0	2.77	33	20	87	96	183	8	571	60
3;1	2.77	46	22	101	126	196	6	504	42
3;2	2.74	53	19	125	161	219	8	664	38
3;3	3.09	20	13	37	45	96	3	181	5
3;4	3.18	14	13	21	39	88	6	159	8
3;5	3.47	19	14	30	49	74	6	145	10
3;6	3.63	16	16	21	80	77	6	184	8
3;7	3.93	18	14	38	39	76	4	140	15
3;8	3.65	17	13	24	34	110	6	228	10
3;9	3.59	28	20	44	62	103	4	185	9
3;10	3.75	19	16	23	30	91	5	165	14
3;11	3.51	9	8	11	22	67	1	100	2
Total		[171]	[52]	750	1,263	[899]	[19]	10,111	498

Note. As a single type can contribute to more than one monthly figure, type totals (bracketed) are not sums of their respective columns. MLU = mean length of utterance.

The data were then broken down over time (groups of seven recordings for the 1st year, and of four or five for the 2nd). Figure 1 shows the rates of OR for each sample. It is clear that Brian's OR rates vary considerably. The first past tense ORs occur at 2;5.18 (Brian's age at the first of the seven recordings in the sample). The OR rate at this point is 20.0%. Of the 34 subsequent samples, 11 show zero rates. However, 9 of these lie before 2;10.13, where continuous OR begins. Eleven samples have rates between 1.0% and 10.0%, 9 between 11.0% and 20.0%, and 2 between 21.0% and 30.0%. ORs rise to a peak of 43.5% at 2;11.14. These findings lend some weight to Maratsos's (2000) doubts concerning the claim that "overregularization is a relatively rare phenomenon" (Marcus et al., 1992, p. 35). Viewed from a longitudinal perspective and at this sample density, regularizations in Brian's speech are in fact more prevalent than overall calculations would suggest.

Overall Rates: Nouns

Of Brian's 19 irregular plural nouns, 10 were overregularized at least once. Four hundred and seventy correct child tokens were recorded (*M* = 24.74), and

365 of these were of the 10 overregularized types. There were 1,757 tokens for these plural forms in the input (*M* = 92.47). Eight plural noun types were compounds of *men* or *women* (e.g., *dustbin men*, *police women*). These were treated as distinct types, since they exhibited clearly distinct error patterns. For example, non-compound *men* was produced in OR form some months before any of its compounds, and, while ORs of all compounds were either stem plus -*s* or plural plus -*s*, ORs of *men* were of both types. In total, the corpus contained 28 irregular plural OR tokens, giving an overall OR rate of 5.6%. Of these, 16 were of the form stem plus -*s*, and 12 of plural plus -*s*. Again, it is reasonable to apply the revised rates criteria, since several overregularized nouns could confidently be said to have recovered during the recording period. Revised figures for *dustbin men*, *feet*, *men*, and *mice* raise the overall OR rate slightly to 7.2%. This result is in line with the mean rate of 8.5% reported for the 10 children of the Marcus (1995) study.

It is again in the OR rate over time that we find the most striking characteristic of the ORs in Brian's speech (Figure 2). The first recorded noun plural OR occurs at 2;02.00, and the OR rate at this point is 5.9%. Eleven

Figure 1. Past tense overregularization rate over time (weekly samples).
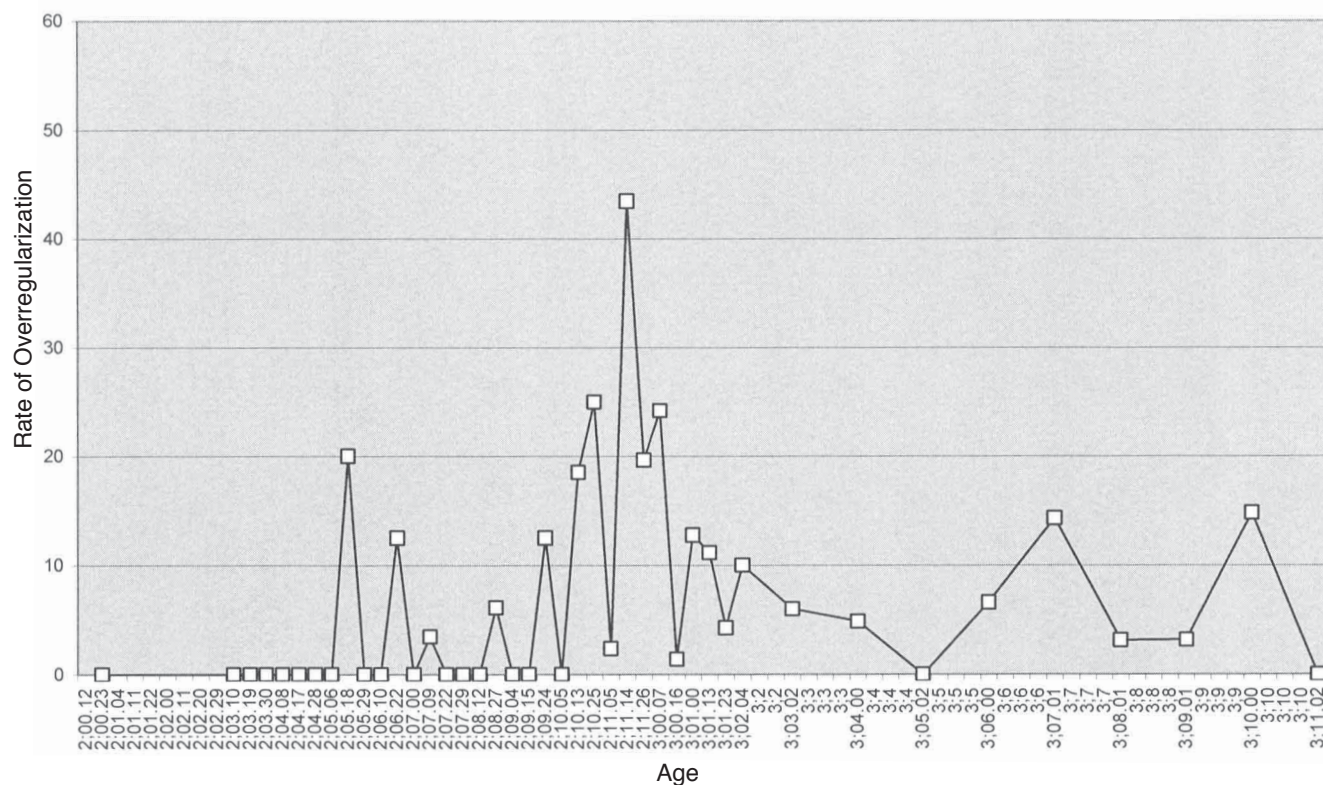
samples have rates of 10% or higher, 3 of these lying in the range 20-30%. The peak OR rate, at 2;08.27, is 50.0%. Examination of the relative distributions of noun and verb OR over time reveals two markedly different patterns. As with his verb ORs, Brian's noun plural ORs begin after an initial period of correct production. However, if we take the first recorded OR as the point of onset of OR, noun ORs begin 3 months sooner than verb ORs. Seven of the subsequent samples contain noun ORs, and 5 of these show rates between 10.0% and 22.0%. In the initial period in which noun and verb ORs are both produced (2;05.18–2;09.15) noun OR rates tend to be markedly higher. However, when the verb OR rate begins to climb (from 2;09.15), the noun rate settles for the most part between 0% and 8%, with 13 of the samples points between 2;09.15 and 3;02.04 returning zero noun OR rates. In that same period, only 3 samples show zero rates for verbs.

This comparison reveals the real value of denser data. The distinct patterns for verb and noun ORs in Brian's transcripts are not apparent in the data reported for Adam, Eve, Sarah, and Abe in Marcus et al. (1992) and Marcus (1995). With samples of 1–2% of a child's speech it is simply impossible to discern any meaningful difference between the use of irregular past tenses and plurals, given that both are overregularized during broadly similar periods; in the dense data, the greater

number of sample points and the increased likelihood of finding low-frequency items combine to reveal a significant period in which nouns are far more likely to be overregularized than verbs, followed by a period in which that pattern is reversed.

Taken together, these findings seem to provide evidence against the blocking account prediction that OR rates are consistently low throughout development (e.g., Marcus, 1995, p. 452). Perhaps more important, the clearly different OR patterns over time support Marchman et al.'s (1997) finding that there is indeed a substantive difference, at least initially, between verb and noun OR. It is argued below that this is better accounted for by very different input type-frequency patterns for nouns and verbs.

Frequency Groups and Correlations

Types Grouped According to Token Frequency: Verbs

Following Maratsos (2000), Brian's irregular verbs were grouped according to their past tense token frequencies. Table 2 shows the number of verbs in each frequency band, together with the mean number of child and input tokens for the group, and the mean OR rate for that verb group (nonoverregularized verbs are

Figure 2. Plural overregularization rate over time (weekly samples).
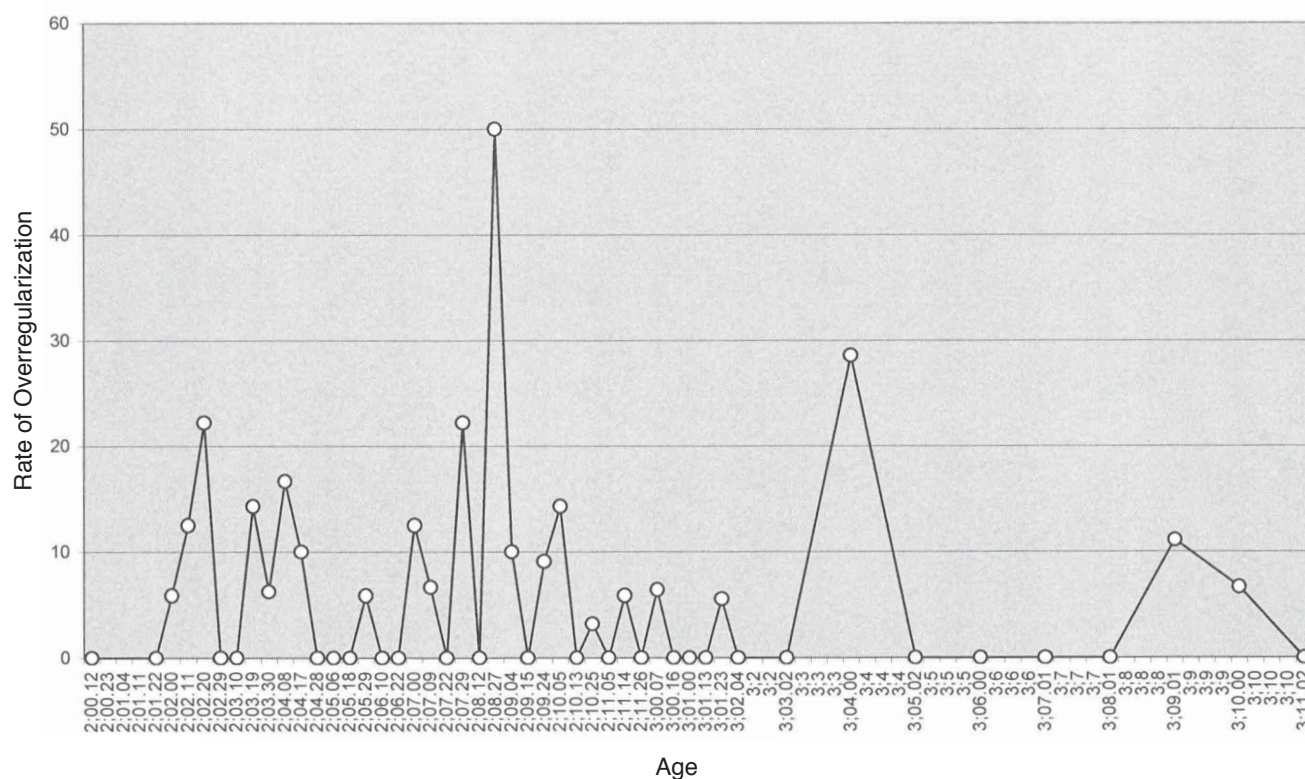
included in the analysis). Correlations were carried out to determine whether there was a relation between token frequency in the child's speech and the input, and between input frequency and the rate of OR in the child's speech. Results show that there is a clear correlation between input and child token frequencies, $r(52) = .73$, $p = .01$, although input tends to be an order of greater magnitude. They further show a strong negative relation between input and the rate of OR for overregularized verbs: The higher its input token frequency, the less likely a verb is to be overregularized, $r(24) = -.62$, $p = .01$. Token frequency in Brian's speech and his OR rate for particular verbs also correlate significantly, $r(24) = -.80$, $p = .01$.

The second important point to draw from the Table 2 data concerns the actual token frequencies that these

Table 2. Irregular verbs divided into groups according to past tense token frequency.

Frequency	No. of verbs	M tokens (child)	M tokens (input)	M OR rate
1–9	29	2.97	52.24	23.1%
10–49	16	25.06	157.19	13.75%
50–99	3	69	736.3	14.27%
100+	4	146.25	960.5	1.85%

figures represent. Maratsos (2000) argues that when sample density is taken into account, a single token might represent actual production of many times that number. Thus, with a sample of 2.5% of a child's speech, two tokens might be extrapolated to a "real" figure of 80 child tokens, and a similar number in adult input. In the case of the present study, where roughly 8–10% of the child's production is captured, two tokens might still represent "real" production of 20 or more. Adult token frequency for a given verb tends, in fact, to be somewhat greater than a child's, and so, in the light of these extrapolated figures, it would seem entirely valid to retain verbs in the low-frequency group for inclusion in group and overall calculations.

Taking the figures from Table 2, an average OR rate of 13.75% for verbs with a "real" frequency in Brian's speech of between 100 and 490 seems a remarkably high value. The OR rate is particularly striking when we consider that the mean figure for correct input of these same verbs, when similarly extrapolated, is well in excess of 1,000 tokens. An OR rate of 14.27% for verbs in the "real" 500–990 group, with extrapolated mean input well in the thousands, represents a similar dogged adherence to erroneous forms. Only the very highest frequency verbs, with "real" child token counts of a thousand or more, are relatively free from ORs.

Types Grouped According to Token Frequency: Nouns

Although the lower number of types and tokens in the noun data somewhat limits this kind of analysis, it is still possible to place Brian's nouns in frequency groups and discern a clear pattern. As was the case with the past tense data, input and child tokens for irregular plurals (excluding types with no recorded input or correct output) correlate very highly, $r(16) = .92$, $p = .01$. Input token frequency and the rate of OR show a stronger negative correlation than in the verb data, $r(10) = -.76$, $p = .05$, and this pattern holds for the relationship between Brian's plural token frequencies and OR rate, $r(10) = -.91$, $p = .05$. Brian's irregular nouns also show the same propensity as verbs to be overregularized despite significant levels of correct child production and input. For example, the group of nouns produced between 100 and 490 times and heard on average over a thousand times has a mean OR rate of 10.67% (see Table 3).

The pattern that emerges from these correlations and grouped calculations serves to underline two points. First, token frequency has a pervasive effect on OR: the more evidence Brian has of a correct form, the less he overregularizes it. Both dual- and single-mechanism accounts would expect higher frequency forms to be overregularized less than low-frequency ones. The second point, however, differentiates the accounts. If OR ends when an irregular memory trace is sufficiently reinforced to engage an innate blocking process, OR rates such as those reported here for medium frequency forms should not be possible; on the other hand, one might reasonably expect the type-frequency effect that drives ORs in single-mechanism accounts to be more gradually overcome. In a connectionist net, for example (e.g., Plunkett & Marchman, 1991), the weighting bias in favor of regular inflection might persist through a large number of training cycles, since, although each instance of a particular correct irregular form adjusts weighting in favor of that form, each instance of any regular form tends to resist that adjustment.

Variation Between Individual Verbs

Brian's overregularized irregular verbs can be grouped according to their pattern of ORs over time (see Table 4). (These categories correspond to those set out in Maratsos, 2000, p. 194, and are similar to three of the four "rough patterns" that Marcus et al., 1992, p. 52, use to break down Abe's data.)

1. *Infrequently sampled verbs.* Some verbs were sampled too infrequently for their behavior over time to be meaningfully categorized. Eleven verbs were thus discounted from detailed individual analysis, since they were sampled fewer than 10 times and so fell short of an extrapolated figure of 100 tokens in the 2 years of recording, which we take as sufficient evidence for recovery from OR to take place. Of course, there is no flawless basis for choosing this figure as a cutoff; Maratsos (2000) selects 50 child tokens (representing 50 inputs, as he assumes input to be roughly equivalent to output) as a reasonable number to facilitate recovery.

2. *Verbs showing low rates of OR.* The six verbs in this group provide either one or two OR tokens across the whole recording period, and that would appear to be a fair basis for categorizing them as "not substantially overregularized." However, if we bear in mind the assumption that a token in the sample might represent 10 or more times that number in terms of "real" production, and if we consider the overall recorded OR rates for these verbs (several of which are around 5%), then these data may in fact represent "competition" or "learning" periods, albeit brief ones, rather than momentary aberrations in the development of Brian's irregular past tensing.

3. *Verbs showing recovery after high rates of OR.* Only one verb in the corpus, *go*, clearly follows the pattern of being overregularized substantially before recovering to correct use. Across the full corpus, the correct past tense of *go* (the suppletive form *went*) was sampled 203 times in Brian's speech, first occurring at 2;03.19 and being continuously present from 2;10.13. Error tokens, 11 in total, were recorded in each of the 6 samples from 2;10.13 to 3;00.07, in proportions ranging from 3 out of 4 tokens at 2;11.14 (75% OR) to 1 out of 15 tokens at 3;00.07 (6.67%). Input for the whole 2-year period is 1,796 tokens.

Of particular interest is the sheer volume of input that occurs before the first OR of *go*, and the number of input tokens and correct child use of *went* between this point and the end of the OR period. Total child tokens before 2;10.13 number 21, with 1,040 tokens in the input. This would seem to be a huge weight of evidence for the correct form, bearing in mind the "real" figures that these totals represent.

However, the *go–went* pairing represents the greatest phonological distance between a stem and past form in English (Bybee & Slobin, 1982), and this presents children with a particular problem when

Table 3. Irregular nouns divided into groups according to plural token frequency.

Frequency	No. of verbs	M tokens (child)	M tokens (input)	M OR rate
1–9	11	1.64	9.90	23.12%
10–49	4	16.5	116.75	10.67%
50–99	3	93.33	263.33	5.21%
100+	1	106	349.00	0%

Table 4. Brian's overregularized verbs (past tense tokens).

Group	Verb	Correct tokens	Stem plus -*ed*	Past plus -*ed*	Input tokens	OR rate	Duration of OR
Infrequent in sample	*draw*	1	3		9	75.00%	0;03.19
	drink		1		15	100.00%	
	hold	2	1		15	33.33%	
	know		1		132	100.00%	
	make	5	1		159	16.67%	
	read	3	2		16	40.00%	0;06.00
	run	2	5		123	71.43%	0;11.07
	sell	1	1		13	50.00%	
	stick	4		2	19	33.33%	0;03.05
	sweep		1		8	100.00%	
	take	2	2		229	50.00%	
Low OR rates	*blow*	22	1		66	4.55%	
	bring	33	2		211	5.71%	0;00.11
	catch	24	1		76	4.00%	
	find	129	1	1	166	1.53%	0;02.29
	say	218	1		1,707	0.46%	
	see	39	1	1	730	4.88%	0;03.21
High OR rates/recovery	*go*	192	11		1,796	5.42%	0;01.24
Extended overregularization	*come*	60	20		898	24.39%	0;11.17
	drive	19	1	10	38	36.67%	0;08.02
	eat	6	5		104	45.45%	0;11.06
	fall	62	12	2	364	18.42%	1;01.13
	fly	9	7		16	43.75%	0;06.24
	throw	3	9		92	75.00%	0;07.06

Note. Duration of OR in years;months.days.

attempting to access the appropriate past tense form for *go*. Children may use different past forms of *go* in different constructions and may, on establishing that *went* expresses the past of *go*, initially be unaware that it does so exclusively (Kuczaj, 1977; Theakston, Lieven, Pine, & Rowland, 2002). These factors predict a certain proportion of ORs, regardless of theoretical standpoint, but there must be some reasonable limit on the number of correct forms that need to be processed before a blocking mechanism is expected to operate successfully. In fact, if we measure the tokens produced from 2;10.13 to 3;00.07 (inclusive), Brian produces a further 35 correct tokens and hears a further 262. Again, the "real" figure extrapolation would give child tokens in the hundreds and input tokens in the thousands for the same period. Thus, although the OR period is relatively short (roughly 2 months), it represents the kind of robust resistance to the correct form, in the face of massive evidence, that Maratsos (2000) finds in Abe's data. It is hard to see how simple reinforcement of a rote form to the point where it is strong enough to consistently block a rule can

account for Brian's dogged attachment to the regularized form of *go*.

4. *Verbs showing extended periods of OR without recovery.* Six verbs (*come, drive, eat, fall, fly,* and *throw*) show substantial OR rates for extended periods during recording, and they show little evidence of recovering to correct use before the end of Brian's 4th year. (The data for *fall* are slightly ambiguous in this regard, as no error tokens are recorded between 3;08.01 and 3;10.00, and no tokens are recorded at all at 3;11.02, which is the final sample.) *Come* is a clear outlier, having a higher OR rate (24.39%) and longer OR period than its frequency (child = 60, adult = 898) alone would predict.

The phonological characteristics of this group of verbs would lead us to expect higher rates of OR than we saw in group 2. Three of them, *come, drive,* and *fall,* undergo vowel change only in forming their past tenses. All end in final consonants—/m/, /v/ and /l/, respectively—that do not carry strong past tense salience (Bybee & Slobin, 1982). *Fly* and *throw* undergo vowel change in forming the past tense where

the vowel is word-final, again giving little information to the developing speaker of the form's "pastness." Even assuming some effect of their phonological makeup, however, the low frequency of *fly* and *throw* would in any case predict their relatively high OR rates, as it does for *eat* (see Table 4).

Past use of *come* appears early in Brian's speech. This would appear to be a function of its frequency in the input, as there are 110 recorded adult tokens before the first child utterance at 2;03. Brian uses the correct past of *come* 18 times across the next 20 samples, producing his first OR at 2;10.13. From this point to the end of the recordings, 12 samples contain error tokens of *come*, at error rates varying from 10% to 66.67%, with the last appearing at 3;10.00. As we saw with *go*, Brian overregularizes *come* despite both a huge volume of input and a considerable number of correct uses of his own. So we once again have a situation, if we multiply out these numbers to the "real" figures that they represent, where Brian continues to overregularize a verb despite using its correct form several hundred times and hearing it several thousand, in this case for a period of just under a year.

Drive is sampled a total of 30 times in Brian's speech and 38 times in his mother's. Of Brian's past tense tokens, 11 are ORs, giving an error rate of 36.67%. After the first error, at 2;10.13, *drive* appears in a further seven samples and is overregularized in five of these. Although this data is fairly sparse, there is some evidence that Brian continues to produce erroneous forms, since the penultimate occurrence of the verb—at 3;08.01—is an OR. By this stage, *drive* achieves an extrapolated figure of some 300 to 400 tokens of input and output. However, *drive* is particularly interesting for the nature of the errors that Brian produces. Ten of his 11 ORs are blending errors (past plus -*ed*, giving *droved*), a phenomenon not predicted by blocking (Patterson, Lambon-Ralph, Hodges, & McClelland, 2001), which states that a retrieved irregular form blocks the regular affixation rule. To create a blending error, the past form must by definition be retrieved. Why, then, is the rule not blocked? One could claim that the stem form has not been acquired and that the child is thus inflecting the past form as if it were the stem. Brian, however, has acquired the stem form by the time these errors are made. Occasional blends could be dismissed as online phonological mix-ups with no implications for the normal running of the system. Brian's persistent production of blending errors, however, casts doubt on such an assertion.

Variation Between Individual Nouns

Token frequencies for the nouns are considerably lower than those for the verbs. As we see in Table 5, 11 of the irregular nouns have frequencies below 10, with 9 occurring in Brian's speech only once or twice in either the correct or overregularized plural form. Of those that are produced more frequently, 2—*teeth*

Table 5. Brian's irregular nouns (plural tokens).

Noun (plural form)	Correct tokens	Stem plus -s	Plural plus -s	Input tokens	OR rate	Duration of OR
binmen	15		1	64	6.25%	
children	25			251		
dustbin men	87		3	120	3.00%	0;01.16
firemen	15		4	93	21.05%	0;11.07
feet	97	1	1	356	2.02%	0;00.17
geese	1			13		
halves	1			9		
hooves	1			2		
knives	2			18		
lives	1					
men	96	8	3	14	10.28%	0;08.05
mice	11	2		59	15.38%	
police men	6	1		23	14.00%	
police women		1			100.00%	
postmen	3	2		7	40.00%	0;00.09
rubbish men	2			7		
shelves		1		17	100.00%	
teeth	106			349		
workmen	1			13		

Note. Duration of OR in years;months.days.

and *children*—are not overregularized. In the case of *teeth*, this can probably be put down to its being the highest frequency irregular plural, with 106 occurrences. Combined with input of 349 for the 2-year recording period—and the extrapolated figures in the thousands that these numbers imply—this would appear to be a sufficiently overwhelming accumulation of experience to guarantee entrenchment of the correct form. Although token frequency in Brian's speech alone (25 tokens over the 2 years) would seem to predict that the plural form of *child* would occasionally be overregularized as *childs*, the zero error rate for *children* is perhaps not so surprising when we consider two factors: its phonological structure and input frequency. The addition of an extra syllable in *children* when the word refers to a plural entity may be a particularly salient feature and might thus facilitate the entrenchment of *children* as the appropriate form for plural contexts. *Children* also occurs very frequently in the input (251 times in the whole corpus).

The noun plural *men* presents an example of substantial OR over an extended period despite relatively high input and output frequencies. The input token count before Brian's first error is 27, representing an extrapolated "real" figure of 270, and he produces the plural correctly 14 (i.e., perhaps 140) times in the same period. Clearly, this mirrors the robust OR seen in some of the past tense data, and that fact is underlined by the volumes of output and, particularly, input recorded before Brian's errors cease. The total output during these 8 months is 61, perhaps representing some 600 actual correct uses; input for the same period is 202, which puts the mother's full production somewhere around 2,000. The reinforcement of the irregular form required to facilitate blocking does not seem ample explanation for this pattern, since, with these frequencies, that would surely predict a far shorter period of error production. Much the same argument applies to Brian's OR rate with *firemen*. It seems more likely that there is some pervasive force acting on the plural system, driving regularization and proving far harder to overcome than a simple weakness in memory trace. A strong candidate for this force is a heavy weighting bias in favor of regular inflection brought about by the vast superiority of regular type frequency. Evidence for the feasibility of this process will be examined below.

The last plural we will mention in detail—*feet*—seems to provide an example of the effect of intensive exposure to high token frequencies. After consistent input totalling 101 correct tokens, Brian produces his first correct form at 2;03.19. He overregularizes in the very next sample, 2;03.30, and once more at 2;04.17. The latter is his final error, and his total correct production at this point has reached 19, with input far higher, at 160 tokens. The total number of input tokens sampled in

the 3 weeks or so that Brian produces errors is 42. When extrapolated to a 100% sample density, the 420 tokens of input in such a short period, or the 1,600 in total directed at Brian by the time that errors cease, perhaps give an idea of the overwhelming volume of input that may be required to overcome the weighting bias of regular type frequency.

Summary of OR Rates

To summarize briefly, Brian overregularizes irregular past tenses and plurals in the face of considerable volumes of correct forms in the input and in his own speech. Analyses of frequency group data, individual types, and the rate of OR over time reveal far more persistent and frequent production of OR forms than overall rates would suggest. In addition, a considerable number of his ORs are blend forms such as *droved*. There is a striking contrast between the patterns of overregularization for verbs and nouns, both in terms of time of onset and initial rates of production of OR forms.

Type Frequency

Verbs

Dual- and single-mechanism accounts make very different predictions about the role of type frequency in OR. Marcus et al. (1992, p. 133) state that "children do not appear . . . to be influenced by either the relative or the absolute number of types" when deciding whether to internalize a pattern in the form of a regular rule or a list. They also assume that a child does not require a large volume of input to acquire a rule (and thus to begin to overregularize). They claim instead that the rate of regular past marking—that is, the proportion of regular pasts marked in obligatory contexts (Brown, 1973)—is central to this process. They show a correlation between increasing regular marking and OR rates during the initial stages of OR production (Marcus et al., 1992, p. 107). The "window" for OR production to begin is the period when a child goes from leaving most regulars unmarked in past contexts to marking them more than half the time, and regular suffixation is the "proximal trigger" (Marcus et al., 1992, p. 114). Bybee (1995) concurs on the importance of the child's recognizing obligatory marking. However, in her network account, no rule is required for the child to apply regular inflection. When a child fails to retrieve an irregular past form, the strongest alternative lexical pattern is accessed; for the English past tense, in terms of variety and openness, this will be the regular form. Other single-mechanism accounts also see regular type frequency as being of central importance in driving OR (e.g., Plunkett & Marchman, 1991, 1993), and the principal is similar. As regular past types overtake irregulars in a child's

speech, they provide a sufficient weight of analogical evidence—or shared connection strength—to make the immature morphological system entertain an erroneous form for an irregular verb because it matches the evidence from an increasing majority of existing representations (recall that regular pasts are individually represented in the single-mechanism account). Type frequency drives schema formation in single-mechanism accounts, while high token frequency acts to entrench irregular forms and protect them from OR. The OR rate correlations reported above lend weight to this argument.

We might then expect to see a relationship between the increase in regular past types in a child's speech—as the child's language comes to reflect the proportions of types in the language that he or she hears—and the onset of OR errors, although such productivity could be taken as the child's recognising obligatory marking.

Brian overregularized the past tense sporadically from around 2;6. At 2;7 and 2;8, as we can see from Figures 3 and 4, both within-month and cumulative irregular past type frequencies exceed their regular counterparts. This is in line with either dual- or single-mechanism predictions, since sporadic OR might be evidence of Brian's reaching a threshold point of regular weighting, but it could also be seen as the first sign that a rule has been learned, even if performance is at first

very limited. For the single-mechanism account to be fully plausible, input must be seen to be driving acquisition. As we would expect from all large corpus studies of adult language (e.g., Francis & Kučera, 1982), regular types considerably outweigh irregulars in the speech of Brian's mother at all stages of recording. We can hypothesize that input type frequencies have already had some effect on the strength of Brian's representation of regular past when he is only just beginning to show signs of productive regular marking.

Brian's within-month and cumulative type frequencies begin to show something like the adult pattern around 2;9, when both regular plots (within-month and cumulative) first exceed their irregular equivalents. It is shortly after this, at 2;10, that OR begins in earnest. During the following 6 months, when Brian's tendency to regularize his irregular pasts is at its peak, the rate of increase in cumulative regular past type frequency is strikingly rapid. This strongly suggests the kind of gradual strengthening of regular representations, or increase in their shared connection weight, that drives error production in the single-mechanism account.
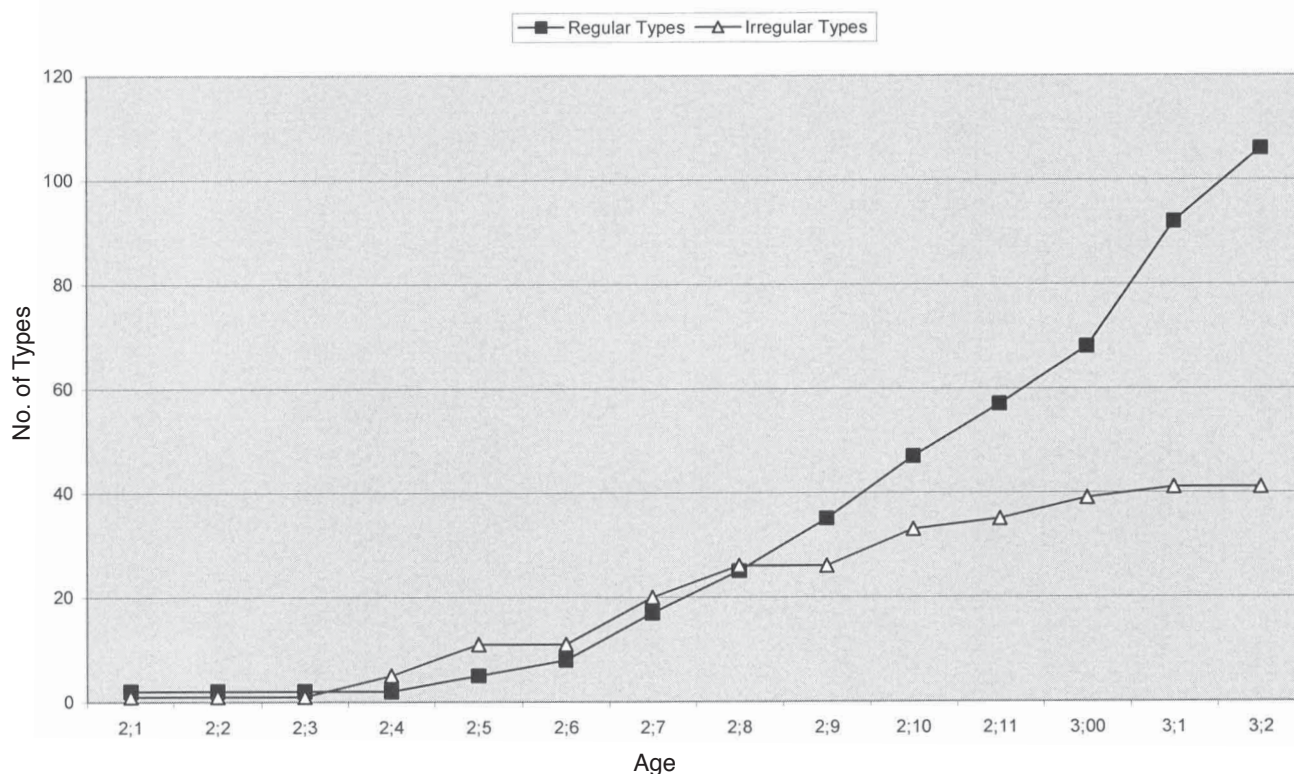
Nouns

The distribution of type frequencies for the plural forms of irregular nouns is very different from the pattern we have seen for the irregular past tense. Regular

Figure 3. Past tense type frequencies: Brian (figures are monthly aggregates).
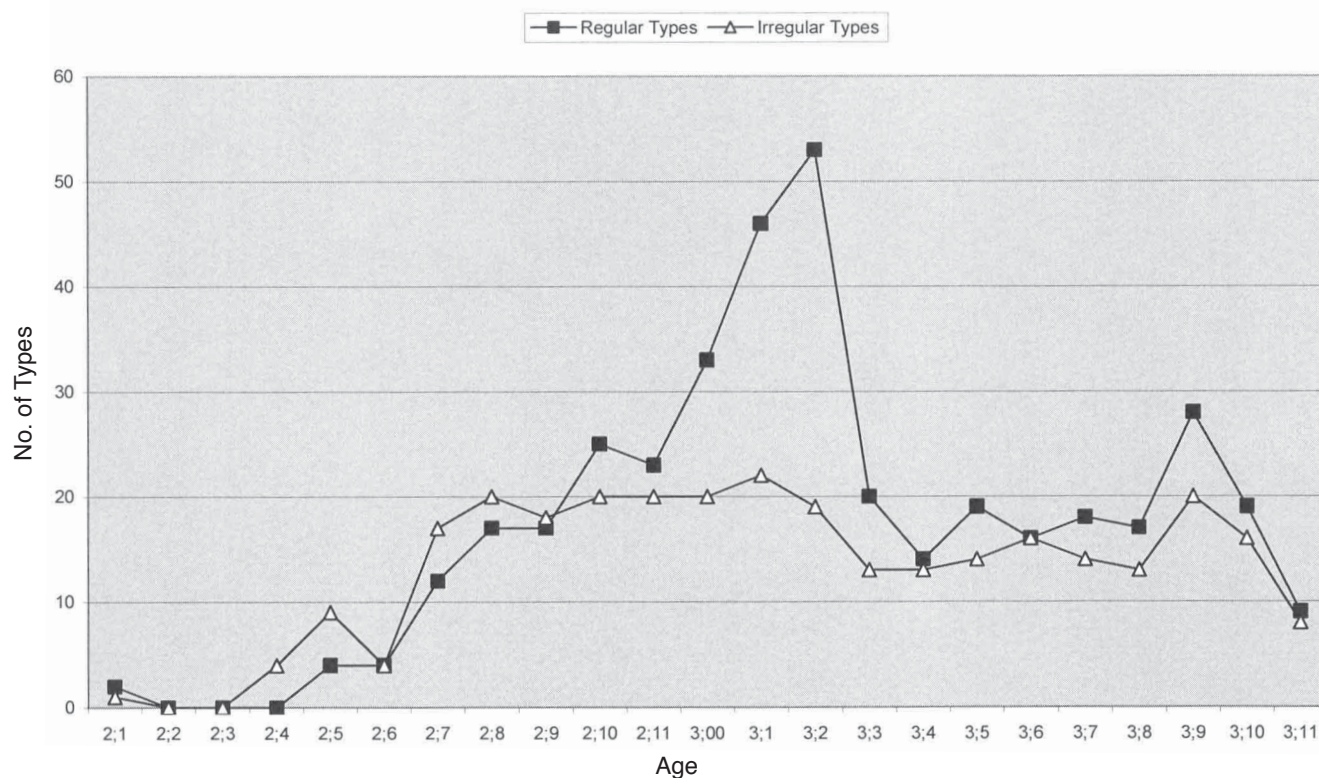
Figure 4. Cumulative past tense type frequencies: Brian (figures are monthly aggregates from 2;1 to 3;2).

types massively outweigh irregulars, and this difference is represented in Brian's speech from the earliest recording. There is no comparable gradual buildup to a verblike crossover point between irregular and regular types; the within-month and cumulative plots begin with a significant separation and continue to diverge. The driving force behind this—the input—shows the fully developed, extreme adult pattern, in which regular types outstrip irregulars by a factor of 20 or more from the first sample onwards.

The simple prediction of a single-mechanism account, when faced with such a discrepancy between the plural and past tense patterns, is that irregular noun plurals should be overregularized earlier and, ceteris paribus, at higher rates than irregular past tenses. The higher frequency of types with the regular inflection in the noun system should produce an earlier weighting bias and facilitate earlier productivity (and over-productivity) than is seen with verbs. Thus OR with nouns should be fully fledged while the same bias in the verb system is just beginning to develop. This is exactly what we see in Brian's data. Contrary to the blocking prediction (Marcus, 1995) that the two types of OR should show no substantive difference, Brian over-regularizes his noun plurals for several months before his first past tense OR, and during the first few months

where the two coexist he tends to overregularize nouns at higher rates (see Figures 1 and 2).

The question remains as to why irregular plurals are not entirely swamped by the regular inflection. As we have seen, only a handful of Brian's irregular plurals are very-high-frequency forms, so token frequency alone would not seem to fully account for their survival. Indeed, unlike the irregular verbs, irregular nouns are not the most frequent forms; *things* (1,808), *hands* (902), *flowers* (510), and *shoes* (507) are all considerably more frequent in the input than the 350 or so tokens of the most frequent irregulars. However, the type-token ratios of regular and irregular plurals in Brian's recordings reveal that although low token frequencies make the data for irregular nouns somewhat chaotic, type-token ratios are lower for irregulars than regulars (that is, there are on average more tokens per type) in 14 of the 22 monthly aggregates after the first month of recording. From this we can posit that there might be a general effect of greater entrenchment of irregular forms—with a concomitant protection from OR—due to the probability that the average irregular plural type has a higher token frequency than its regular counterpart. Marchman et al. (1997) propose in fact that noun plural OR is mitigated by a child's "familiarity" with irregular forms. Phonological group effects are also likely

to play an entrenching role. All but one of Brian's irregular plurals (the exception being *children*) fall into one of two phonological paradigms. Thus, experience of one may have a reinforcing effect on the usage of another. We should also note the high semantic salience of the referents of irregular nouns. All but two of Brian's irregular plurals refer to other people (*postmen*), parts of his body (*teeth*), animals (*mice*), parts of animals' bodies (*hooves*), or common household items (*knives*). This cannot but help him in the task of storing and retrieving their slightly unusual, but paradigmatic, forms.

The concept of cue reliability in Bates and MacWhinney's (1987) competition model may also suggest an explanation for the resistance of irregular plurals to OR by contrast to the pervasive OR of past tense verbs. Cue reliability reflects how often the cue indicates the required meaning relative to how often it does not. The cue reliability of the *-ed* ending is relatively high (it rarely means anything other than "past"). However this is not the case for the *-s* ending, which can also be the possessive marker on nouns and the third person inflection on verbs, for instance. This might "protect" the irregular plurals. There are thus several possible factors that may be defending irregular plural nouns from being engulfed by the effects of hugely superior regular type frequency.

Brian's verb and noun type-frequency data are both broadly consistent with a single-mechanism account of OR. Both show that the expected dominance of regular types. Crucially, the difference between the ratios and numbers of regular and irregular forms in the two systems, when combined with the effect of token frequency discussed above, seems to go some way to accounting for the very different patterns of OR in Brian's irregular past tense and plural production.

Discussion

OR is an important and well documented issue in the larger debate concerning the nature of linguistic representation. That debate, broadly defined, centers on the question of whether linguistic structure rests on innately specified foundations, or emerges from general cognitive mechanisms. Children's ORs of English past tense and plural forms provide us with a delimited subset of language data on which we can test empirically specific claims arising from these contrasting characterizations of language.

In this study, we have used a dense naturalistic database to test the predictions of the blocking account and examine the evidence for a single-mechanism alternative based on the interplay of type and token frequencies. Existing criticisms of blocking as an explanation of OR (Maratsos, 2000; Marchman et al., 1997) have

been largely borne out. Evidence for this conclusion comes principally from three sources. Firstly, group OR rates and enduring periods of OR production for individual verbs in the face of high-frequency correct input and output are not consistent with recovery from OR being solely reliant on the reinforcement of a weak memory trace. Although it might be argued that reinforcement sufficient to trigger blocking requires a volume of input/output of many hundreds, or thousands, of tokens, this would profoundly weaken the account, since it would predict, on the one hand, the failure of some forms ever to reach the blocking threshold, and on the other, extremely high overregularization rates that are not attested in any study of which we are aware. Secondly, the large number of blending errors with verbs and nouns whose stem forms had been acquired is not consistent with the operation of a blocking mechanism and the rule-driven inflection of stem forms. Finally, the differentiation in past tense and plural OR patterns runs entirely against the explicit prediction of Marcus (1995) that there be no "substantive" difference between the two.

A successful account of OR must, then, address this contrast between the way verbs and nouns behave. We suggest that the difference in the way Brian incorrectly inflects his irregular past and plural forms can be linked to the very different ratios of regulars to irregulars in the two systems. Past tense errors begin as Brian moves toward an adultlike distribution of past tense forms and regular types overtake irregular types; whereas, from the beginning of recording, his distribution of noun plurals resembles that of an adult, and ORs are recorded much sooner. This difference is precisely what one would expect to see if ORs occurred as the result of interference from a dominant schematic paradigm. Whereas blocking predicts no effect of type frequency, Brian's data exhibit the pattern that would result from, for example, a bias in weighting in a connectionist net in favor of an inflection with a high type frequency (Rumelhart & McClelland, 1986). This is underlined by the fact that Brian's data also show a differentiation between the effects of two dominant paradigms—regular plural and past tense—of different magnitudes.

Clearly, the main source of evidence that Brian has about these relative distributions comes from the language he hears. Input is thus of central importance in the process, a fact that is supported by the close correlations between input and output on the one hand and between input and OR rate for overregularized verbs and nouns on the other.

To summarize, there is evidence in Brian's data of regular type frequency driving ORs, and of token frequencies for individual irregulars mitigating that effect. This fits models of OR which hold that type frequency leads to the formation of a particularly strong schema for regular

inflection and that individual irregular forms are entrenched due to their high token frequency and are thus protected from ORs, or pushed to recover from them.

We do not, of course, claim to make the case here for single-mechanism, connectionist or network models—what are broadly termed usage-based models (Langacker, 1987). However, the evidence from our uniquely dense data highlights the substantial failure of the blocking model to account for OR. In contrast, explanations that explicitly invoke the interplay of type and token frequencies seem to us to provide the most promising avenues for future research.

Acknowledgment

References

Bates, E. & MacWhinney, B. (1987). Competition, variation and language learning. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 157–193). Hillsdale, NJ: Erlbaum.

Behrens, H. (2001). Learning multiple regularities: Evidence from overgeneralization errors in the German plural. In A. H.-J. Do, L. Domínguez, & A. Johansen (Eds.), *Proceedings of the 26th Annual Boston University Conference on Language Development* (Vol. 1, pp. 72–83). Somerville, MA: Cascadilla Press.

Brown, R. (1973). *A first language*. Cambridge, MA: Harvard University Press.

Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, 10, 425–455.

Bybee, J. L. & Slobin, D. I. (1982). Rules and schemas in the development and use of the English past tense. *Language, 58,* 265–289.

Dabrowska, W. (2001). Learning a morphological system without a default: The Polish genitive. *Journal of Child Language, 28,* 545–574.

Francis, N., & Kučera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar*. Boston: Houghton-Mifflin.

Higginson, R. (1985). *Fixing-assimilation in language acquisition*. Unpublished doctoral dissertation, Washington State University, Pullman.

Köpcke, K. (1998). The acquisition of plural marking in English and German revisited: schemata versus rules. *Journal of Child Language, 25,* 293–319.

Kuczaj, S. A. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behaviour, 16,* 589–600.

Langacker, R. W. (1987). *Foundations of cognitive grammar: Vol. 1. Theoretical prerequisites*. Stanford, CA: Stanford University Press.

MacWhinney, B. (Ed.). (1987). *Mechanisms of language acquisition*. Mahwah, NJ: Erlbaum.

MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk*. Mahwah, NJ: Erlbaum.

MacWhinney, B., & Snow, C. E. (1985). The Child Language Data Exchange System. *Journal of Child Language, 12,* 271–296.

MacWhinney, B., & Snow, C. (1990). The child language data exchange system: An update. *Journal of Child Language, 17,* 457–472.

Maratsos, M. (2000). More overregularisations after all: New data and discussion on Marcus, Pinker, Ullman, Hollander, Rosen & Xu. *Journal of Child Language, 27,* 183–212.

Marchman, V. A., Plunkett, K., & Goodman, J. (1997). Overregularization in English plural and past tense inflectional morphology: A response to Marcus (1995). *Journal of Child Language, 24,* 767–779.

Marcus, G. F. (1995). Children's overregularization of English plurals: A quantitative analysis. *Journal of Child Language, 22,* 447–460.

Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., & Xu, F. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, 57(4, Serial No. 228).

Patterson, K., Lambon-Ralph, M. A., Hodges, J. R., & McClelland, J. L. (2001). Deficits in irregular past-tense verb morphology associated with degraded semantic knowledge. *Neuropsychologia, 39,* 709–724.

Pinker, S. (1999). *Words and rules: The ingredients of language*. New York: Science Masters.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28,* 73–193.

Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition, 38,* 43–102.

Plunkett, K., & Marchman, V. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition, 48,* 21–69.

Rubino, R. B., & Pine, J. M. (1998). Subject–verb agreement in Brazilian Portuguese: What low error rates hide. *Journal of Child Language, 25,* 35–59.

Rumelhart, D., & McClelland, J. (1986). On learning the past tense of English verbs. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models* (pp. 272–326). Cambridge, MA: MIT Press.

Theakston, A., Lieven, E., Pine, J., & Rowland, C. (2002). Going, going, gone: The acquisition of the verb "go." *Journal of Child Language, 29,* 783–811.

Contact author: Robert Maslen, Max Planck Child Study Centre, Department of Psychology, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom. E-mail: robert.maslen@stud.man.ac.uk