SPECIAL ISSUE

CHARACTERISING COMPENSATION

(Commentary on Ullman and Pierpont, "Specific Language Impairment is not Specific to Language: The Procedural Deficit Hypothesis")

Michael S.C. Thomas

(School of Psychology, Birkbeck College, University of London, UK)

ABSTRACT

This article considers Ullman and Pierpont's Procedural Deficit theory of Specific Language Impairment (SLI). The theory represents an innovative attempt to fill the gap between brain and cognition in SLI, and has the potential to explain the non-linguistic as well as linguistic deficits seen in this disorder. The theory is reviewed with regard to: (1) the claims it makes on the domain-specificity of language structures; (2) the falsifiability conditions of the theory; (3) the level of detail at which compensatory processes are specified; and (4) from a computational perspective, whether the inferences that the theory draws from uneven behavioural impairments to underlying structural deficits are necessary ones.

Key words: developmental disorders, compensation, Specific Language Impairment (SLI), past tense, computational modelling, connectionism

Cognitive level explanations of SLI have yet to reach a consensus on whether the disorder truly represents a language-specific deficit, or whether language abilities merely exhibit a particularly salient impairment arising from a more domain-general deficit. There is more agreement that the disorder – usually identified purely on behavioural grounds – is inherently heterogeneous. However, developmental accounts of SLI have rarely sought to relate proposals of cognitive deficits to the brain structures that may be involved. This would open up the possibility of a unified account of SLI from genotype through brain development and cognitive development to the subsequent behavioural impairments, something of a holy grail for researchers in language development.

The article presented by Ullman and Pierpont (this issue) offers two innovations in the study of SLI. The first innovation is an attempt to fill the gap between brain and cognition in SLI. The authors offer detailed proposals as to the particular brain structures that may be developing atypically in SLI and how these would lead to the observed behavioural deficits. Within this framework, Ullman and Pierpont (henceforth U and P) offer a possible means to explain the heterogeneity of SLI and its patterns of co-occurrence with non-linguistic deficits.

At the heart of U and P's "Procedural Deficit" (henceforth PD) theory is the claim that the uneven profile of deficits observed within the language abilities of children with SLI stems from the differential reliance of the normal language system on two separate, more domain-general memory systems. They propose that on the one hand, grammar development normally relies on the *procedural* memory system, whose characteristics are implicit learning, slow acquisition, fast automatic execution and sequence processing. In adults, the procedural system is mediated by brain structures related to the "dorsal" stream, a network involving the basal ganglia, structures predominantly within left cerebral hemisphere [including pre-motor frontal cortex, Broca's area, portions of parietal cortex, and portions of the superior temporal lobe] and the cerebellum. In addition, the procedural system has domain-general involvement in dynamic mental imagery, working memory, complex sequential/-hierarchical processing, and rapid temporal processing. On the other hand, vocabulary development is taken to rely on the *declarative* memory system. This is characterised by parallel processing, slow recall, explicit/conscious learning, and in adults is mediated by brain structures associated with the "ventral stream", a network involving medial temporal lobe regions [in particular the hippocampus], temporal and parietal neocortical regions, portions of ventral-lateral pre-frontal cortex and anterior fronto-polar cortex, and parts of the right cerebellum. This system is also involved in acquisition and representation of semantic knowledge, episodic knowledge, and learning arbitrary relations.

In SLI, behavioural impairments are particularly observed in grammar. According to U and P, SLI is associated with atypical development of the procedural memory system. Explanations of heterogeneity and associated non-linguistic deficits are then derived as follows. Given that the procedural system employs a network of brain structures, heterogeneity within SLI might be explained if the system-wide procedural memory problem can arise through developmental impairments to different component structures

within the network, each causing a subtly different constellation of effects. Given that the procedural memory system is involved in processing other cognitive domains, patterns of co-occurring non-linguistic deficits might be explained by shared reliance on certain processing resources.

The second innovation offered by U and P's article is that compensation is placed centre stage in explaining observed behavioural impairments in a developmental disorder. The profile of language skills in SLI is a consequence *not only* of the procedural system's sub-optimal attempts to acquire the structural aspects of language *but also* of the attempts of the declarative memory system to compensate for this shortcoming. Compensation – effectively a part of the developmental process – has often been left out of explanations of developmental deficits (see Karmiloff-Smith, 1998; Karmiloff-Smith and Thomas, 2003; Thomas, 2003; Thomas and Karmiloff-Smith, 2005, for discussion and examples). It is exciting to see its explicit inclusion in U and P's account of SLI. The exact nature of the compensation process within U and P's theory is the focus of this article. I will argue that specifying the compensation process is both a complex task and key to advancing our understanding of SLI.

Given the influential role of SLI in theories of language development, I begin by clarifying U and P's precise claims on what SLI tells us about the domain-specificity of language structures. Next, I set out the falsification conditions for the PD theory. I then seek to make explicit U and P's account of the compensatory process via the concrete example of inflectional morphology. Lastly, I briefly discuss some data from a recent computational model of atypical language acquisition, which indicate that without a clear understanding of the compensatory process, direct inferences from behavioural deficits to (selective) impairments in underlying structures can be problematic.

How Domain-Specific is SLI?

This issue may appear clear-cut given the title of U and P's article: 'SLI is not specific to language'. Given the broad definition of SLI, U and P do not claim that all subgroups and individuals identified with SLI have a dysfunction of the procedural system. But for the variety of SLI they do wish to explain, U and P's position on the specificity of SLI is a subtle one. In their view, SLI is *typically* not specific to language. But they leave it open as to whether the same or distinct circuits within procedural memory (basal ganglia-thalamocortical channels) underlie grammatical and non-linguistic processing. This neutrality extends to whether sub-domains of grammar such as phonology, morphology, and syntax, themselves rely on common or distinct circuits.

In terms of the atypical phenotype, U and P view it as unlikely that a developmental anomaly of procedural memory would be restricted only to the portions of frontal or basal-ganglia circuitry that subserve grammar. One can put this another way. It is possible that associated behavioural deficits outside language in SLI may occur because neurogenetic disorders tend to impair all the basal-ganglionic loops irrespective of domain. Or it could be that such loops are not domain-specific in the first place. The authors accept that evidence of an entirely "pure" behavioural sub-type of SLI associated with frontal/basal-ganglia abnormalities might point to the existence of domain-specific circuitry within procedural memory. However, they do not take the position that such data currently exist. Moreover, in common with other authors (e.g., Karmiloff-Smith, 1998), U and P recognise the methodological difficulty of demonstrating pure deficits given that these are contingent on the range and sensitivity of tests employed outside the domain of interest.

Falsification Conditions for the Procedural-Deficit Theory of SLI

Two aspects of the PD theory make it powerful in explaining various patterns of data. First, the PD theory can explain the heterogeneity of SLI as arising from differential types of developmental deficit to the procedural memory system. As indicated above, deficits may be more or less restricted to language. Second, the PD theory incorporates compensation from the (initially) normal declarative memory system, which may alleviate some (potential) behavioural deficits across development. As a consequence, PD theory is able to capture variations in apparent severity of behavioural deficits as well as variations in the breadth of the disorder – from apparently pure linguistic deficits to those with clear accompanying motor and other non-linguistic deficits, including the pattern exhibited by the KE family.

In order to place this explanatory potency in context, it is useful to set out the scope of the theory and the conditions under which it would be falsified. In terms of scope, there are realms of language deficit that the theory does not seek to explain, such as socio-pragmatic deficits. In terms of falsification conditions, several are easily identifiable:

1. If the heterogeneity of SLI is explained by multiple loci of deficits to the procedural system, each sub-type of SLI should be associated with a consistent pattern of brain deficits.

2. A type of behaviour inevitably tied by structure-function correspondence to the procedural system (that is, only computable by the procedural system) should not look normal when early damage to the procedural system is diagnosed. Such behaviour would be complex, hierarchically structured, sequential, productive, have elements

that are low frequency, be learned slowly but executed quickly on-line.

3. If compensation from the declarative system is invoked to explain more normal-looking behaviour in a domain usually (assumed to be) acquired by the procedural system, then there should be some independent method available to demonstrate the atypical involvement of the declarative system. Two are on offer: (1) declarative memory involvement should generate behavioural hallmarks (such as frequency and similarity effects); (2) brain imaging should indicate differential activation of declarative brain structures.

4. Compensation should only occur for those parts of 'procedural' language also amenable to acquisition by the declarative system. Where compensation is invoked, integrity of the declarative system in the individual should be independently demonstrable, for instance by testing vocabulary ability or paired-associate learning.

These conditions demonstrate that although the PD theory is flexible enough to account for many patterns of data in SLI, it is nevertheless open to falsification through future empirical work.

Detailing the Compensation Process – A Worked Example within the Procedural-Deficit Theory

U and P's appeal to developmental compensation to account for the patterns of behavioural deficits in SLI is impressive and brave. It is impressive because other explanations of developmental deficits frequently omit the possibility of compensation for simplicity's sake (Thomas and Karmiloff-Smith, 2002, 2003, 2005, Karmiloff-Smith and Thomas, 2003, for discussion and examples). Apparently selective behavioural deficits in developmental disorders are sometimes conveniently explained with reference to the atypical developmental of a single functional module (with the module concerned usually inferred from the selective breakdown of the normal *adult* cognitive function). Yet developmental disorders are essentially stories about why compensation does or does not take place. Development is characterised by plasticity and interactivity (Bishop, 1997). Selective behavioural deficits in healthy adults are often linked to focal brain damage (e.g., in aphasia). Selective behavioural deficits in developmental disorders seem to imply focal brain damage (e.g., as in SLI). Yet focal brain damage in young healthy children is usually followed by recovery across development, not selective deficit (e.g., for language following left hemisphere damage; Bates and Roe, 2001). So why is there no compensation-to-recovery in developmental disorders?

The appeal to compensation is also brave because once made, it demands specification. What are the constraints that predict the extent to which compensation will occur? How far can declarative memory compensate for grammatical deficits? Should the compensating system demonstrate deficits in acquiring its normal abilities through capacity limitations? If systems can compensate for each other, doesn't this imply a relatively greater degree of equipotentiality in the component systems? What, then, is the nature of the competitive process that drives the specialisation of these systems under normal circumstances? Why doesn't such equipotentiality lead to duplication and redundancy, so that both procedural and declarative memory systems learn both grammatical and vocabulary knowledge? (Thomas and Richardson, in press, for discussion).

These theoretical concerns can be illustrated by a more concrete example. U and P illustrate the application of their dual-memory model to SLI in the domain of inflectional morphology, and particularly acquisition of the English past tense (Pinker, 1999; van der Lely and Ullman, 2001; Ullman and Gopnik, 1999). Their proposal is as follows. The empirical data indicate that children with SLI are very poor at producing inflected past tense forms. (Illustrative data can be seen later as part of Figure 1). These children produce low levels of inflections on both regular and irregular verbs. They show very limited extension of the regular rule to novel forms (e.g., wug-wugged) and a much-reduced level of over-generalisation errors compared to normal children (errors where the regular rule is mistakenly over-applied to an irregular verb, e.g., think-thinked). U and P propose that this pattern is a consequence of (1) the failure of procedural memory to learn regular inflections and the general rule of adding –ed to form a past tense, and (2) developmental compensation from declarative memory. The behavioural evidence offered for compensation from declarative memory includes (a) the prevalence of high frequency past tense forms in the language of such children; (b) the absence of the normal advantage for regular forms so that regular and irregular verbs appear to be treated equivalently; and (c) the presence of frequency effects (hallmarks of associative memory) in producing regular past tenses, where such effects are normally absent or weak in typically developing children. U and P propose that the overall low level of inflections 'could also (in addition or instead) be accounted for by lexical retrieval deficits, which would depress performance at the production of irregulars and also of regulars if these are memorized as a compensatory strategy').

Let us now interrogate the PD theory, in an attempt to find out exactly how the compensation process works. We begin with the rationale for incorporating it in the first place.

Why Postulate Compensation?

In the PD theory, the procedural mechanism acquires regular suffixation and the declarative

mechanism acquires whole lexical forms. If there were no compensation, then initial impairment of the procedural mechanism should lead to impaired regular suffixation but normal development of irregular verb production. However, the data do not demonstrate this pattern (e.g., Ullman and Gopnik, 1999; van der Lely and Ullman, 2001). There are low levels of all inflection types. Some additional process must therefore be added.

In the PD-with-compensation theory, the declarative mechanism learns both inflected regular and irregular verbs at low levels. On its own, this makes the explanation under-constrained. What is the predicted capacity limitation of the declarative memory that prevents it from learning ALL inflected verb forms? If the declarative system could learn all verb forms as a long list, it would demonstrate 100% correct on regular and irregular past tenses, but zero performance on unknown novel verbs. PD-with-compensation must invoke a *limited capacity* for the declarative memory to explain why compensation isn't absolute – the system must be able to store only a limited number of whole inflected regular and irregular verbs. However, the declarative system is the home of the entire mental lexicon, potentially hundreds of thousands of word forms. Why should it have a problem learning, say, the thousand most frequent regular and irregular past tense forms so that the individual can get by undetected in every day conversation? Further detail from the PD account is clearly necessary here.

Evidence from Regular vs. Irregular Verbs

Within the PD-with-compensation theory declarative system compensation is justified on the basis that the normal advantage of regular verbs over irregular verbs is absent. In SLI, regular and irregular verbs are treated the same.

This is potentially important evidence, but it depends on data that are limited in two respects. Any comparison between regular and irregular verb inflections will depend on the verbs that are used. Irregular verbs include a subset of very high frequency items (e.g., go, be, come) that are among the earliest acquired words in a child's vocabulary. Whether or not these are included in the set of irregulars will affect whether children do or do not demonstrate greater proficiency with regular or irregular verbs. Second, assuming regular and irregular verb sets have been appropriately balanced, when a null effect of verb type is found in studies of inflection in SLI, it is typically in the context of overall inflection rates of around 20%. These rates are low enough to mean that large sample sizes would be needed to demonstrate a significant main effect. Several studies that fail to find the normal advantage of regular past tense formation over irregular in SLI nevertheless demonstrate a non-significant advantage for regular past tense formation (Norbury et al., 2001; Oetting and Horohov, 1997; Ullman and Gopnik, 1999; van der Lely and Ullman, 2001; Watkins et al., 2000). It is possible that with larger samples, this advantage for regular verbs would be shown to be reliable.

Nevertheless, the lack of a regular-irregular difference is not the only line of evidence used to support the PD-with-compensation theory. Involvement of the declarative system is further supported by evidence that regular past tense formation in SLI exhibits exaggerated frequency effects (illustrative data can be seen later in Figure 3). Thus, whereas in typical development the speed or accuracy of inflecting a regular form is largely independent of the frequency of the verb stem in the language, this is not the case in SLI, where the likelihood of generating the correct inflection is greater for high-frequency verbs.

However, the presence of frequency effects in regulars has another ready explanation. Researchers who argue that both verb types are learned in an associative (connectionist) system also argue that regulars experience an attenuated influence of frequency through ceiling effects, since regular verbs in English have high type frequency and this tends to squeeze out performance differences along this dimension. Any sub-optimal learning system that did not hit the ceiling would therefore exhibit an exaggerated regular verb frequency effect. Ullman, Hartshorne, Estabrooke, Brovetto and Walenski (under revision) have countered this idea by showing that some regular past tenses are produced more slowly than frequency-matched irregulars and with more response-time variability, arguing against ceiling effects in the normal case.

Over-Regularisation Errors

Most SLI data show a small residual productivity of regular inflection in terms of over-regularisation errors (e.g., thinked) and in terms of regularly inflected novel verbs. Where does this come from if rule-use is procedural and the procedural system does not work? Or is the procedural system working *just enough to* explain the low level of productivity?

The PD-with-compensation theory can explain the residual productivity if declarative memory is viewed as an associative connectionist network. Such networks demonstrate similarity-based generalisation that could produce extension of regular inflection to irregular or novel stems. However, if declarative memory has this property, why is compensation so poor? For example, children normally exhibit productivity of their irregular verb knowledge to certain novel verbs that rhyme with existing irregulars (e.g., *frink*, rhyming with drink, is sometimes inflected as *frank* instead of *frinked*). This supports the productivity of the declarative system in U and P's theory. Existing connectionist models confirm that associative networks can produce

similarity-based generalisation, such as the extension of irregular patterns. However, these models are much more powerful than the PD proposal implies. Along with learning *all* the regular and irregular inflections, existing associative networks generalise the regular rule with around 80% accuracy (e.g., Joanisse and Seidenberg, 1999; Plunkett and Juola, 1999). A challenge for the PD theory is to explain why, if residual function in SLI is traced to performance of the declarative system, performance should be so low.

The Role of Lexical Retrieval Deficits

One might imagine that a pure procedural deficit should not impair lexical function, which is handled by the intact declarative system. Nevertheless, production of both regular and irregular inflected past tense forms is very low in SLI. The PD-with-compensation theory explains this by arguing that the procedural system is normally involved in lexical retrieval (if not storage). A problem for this account is that it appears inconsistent with the fact that around 70% of the children's responses in SLI past tense elicitation tasks are context-appropriate, uninflected lexical forms. To explain these uninflected forms, the PD-with-compensation theory postulates a *separate* compensation process emerging from a non-grammatical part of the declarative system, which inserts conceptually appropriate word forms for the sentence context [Ullman and Gopnik's (1999) notion of "conceptual selection"]. However, the deployment of one compensatory process to explain the level of correct performance and a second compensatory mechanism to explain the nature of the errors makes the PD theory somewhat less parsimonious. It also leaves open the question of how exactly this new compensatory process works, and leads us to ask why inflected whole-forms cannot be stored in the non-grammatical part of the declarative system and accessed via a conceptual representation stipulated to refer to a past event.

The aim of this interrogation is not to criticise the PD theory *per se* – no current theory has answers to these sorts of questions. It is to demonstrate that when the theory is pushed hard, it becomes clear that the proposed compensation process is under-constrained. But the exact nature of the compensation process is key in explaining the pattern of behavioural data an atypically developing system will exhibit. In the PD account of SLI, it is not evident how the compensatory contribution of the declarative memory system can be made powerful enough to explain the residual abilities we see in SLI without making it so powerful so that it would predict compensation to recovery. This is not just a matter of *post-hoc* calibration – we need empirical evidence to constrain the proposed level of power.

Inferences from Behavioural Deficits to Structural Deficits – Results from Computational Modelling of Atypical Development

Simply estimating what compensation will 'probably' do in an atypically developing system is very tricky. One way to place these intuitions on a more solid foundation is to implement a computational model of the compensatory processing within an atypically developing language system. Staying within the concrete domain of past tense acquisition, I discuss two examples of developmental deficits and models that have simulated them. The first is a set of data drawn from empirical work on Williams syndrome (Thomas et al., 2001) while the second is drawn from SLI (e.g., van der Lely and Ullman, 2001). Despite the real empirical data, the intention here is to focus on the types of inferences that should be drawn from developmental deficits. Consider these scenarios and resulting inferences:

Scenario A: A group of individuals with delayed past tense acquisition exhibits an exaggerated influence of the semantic dimension of imageability on their irregular past tense formation. Inference: Semantic dimensions are hallmarks of the mental lexicon, and therefore of declarative system involvement. Perhaps in this disorder, the declarative memory system is playing an unusually strong role in past tense acquisition?

Scenario B: A group of children with delayed past tense acquisition exhibits a greater delay in learning regular verbs than irregular verbs, against a background of low overall levels of inflection, reduced over-generalisation errors, and poor generalisation of the regular –ed rule to novel strings. Regular verbs show increased frequency effects, similar to those exhibited by irregular verbs. Inference: there has been an initial deficit to a domain-specific mechanism for regular verb acquisition. Residual behaviour is through compensation from an alternative system.

The advantage of implemented computational models is that they allow such inferences to be evaluated. Are there other ways these patterns of behaviour can be generated, or do the observed data necessarily imply the derived inference? Recently, Thomas and Karmiloff-Smith (2003) explored a developmental model of past tense acquisition under a wide range of atypical learning conditions to evaluate the patterns of developmental deficits it might exhibit. The model combines information about phonological regularities in the mappings between verb stem and past tense form with information from the lexicon about word identity. In broad terms, this model may be seen as respecting the procedural-declarative divide

proposed by U and P (Lavric et al., 2001). Joanisse and Seidenberg (1999) demonstrated that such a model learns to differentially drive irregular past tense formation from lexical information (cf. declarative) and regular past tense formation from phonological regularities (cf. procedural). If the model is trained with atypical initial computational constraints, the success of the model in acquiring the past tense can be affected and the balance of the information sources it uses to drive performance can be altered. Results from these simulations indicated that the patterns of data in both scenario A and B could be reproduced. Interestingly, in neither case was the nature of the underlying deficit the one inferred in the above examples.

In scenario A, the model demonstrated that a pattern of exaggerated imageability effects in irregular past tense formation could be produced by a *reduced* rather than increased contribution of lexical-semantic information (Thomas and Karmiloff-Smith, 2003). In normal development, word-specific information from both abstract and concrete verbs is used to disambiguate irregular and regular mappings, facilitating acquisition of the former. Under the assumption that representations of abstract words are less robust, with fewer context invariant semantic features (Plaut and Shallice, 1993), a reduction in the overall strength of lexical-semantic input to the inflectional process means that the word-specific information from abstract irregular verbs is now too weak to facilitate acquisition. Abstract irregulars consequently experience interference from the overriding phonological mappings of the regular verbs. However, concrete irregulars still experience a strong enough lexical signal to disambiguate them from regulars. The result is a greater difference in irregular past tense formation along a semantic variable, even though the cause is a *reduction* in the contribution of lexical-semantic (declarative) information.

I will consider Scenario B in a little more detail, since it corresponds to the empirical pattern identified for SLI. Figure 1 depicts illustrative empirical data from a past tense elicitation task reported by van der Lely and Ullman (2001). Data for an approximate chronologically age matched group has been added from another study using the same task, since van der Lely and Ullman's control groups were matched on language ability. Figure 1 also includes simulation data from a past tense model trained under normal conditions, and data from the model when trained under atypical conditions of reduced processing unit discriminability (see Thomas and Karmiloff-Smith, 2003, for details). In each condition, the model experienced the same amount of training, so the two conditions are equivalent to chronological age matches.

Discriminability is a property of the layers of simple processing units that learn to compute the transformations between input (stem phonology and lexical information) and output (past tense phonology) in a connectionist network. Each processing unit has an activation function (usually a smoothed threshold) that determines how its output will be affected by changes in the input it receives. Figure 2 illustrates the normal activation function and the alteration applied to the atypical network. A sharper-than-normal function would mean that the processing unit was more able to distinguish

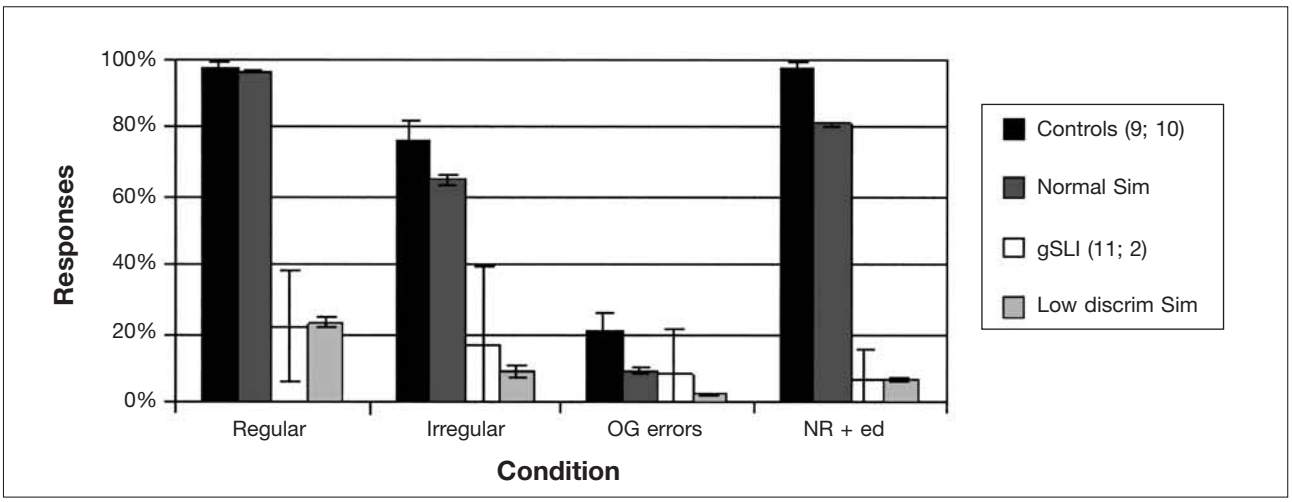

Fig. 1 – *Empirical and computational simulation data for correct use of English past tense inflections in SLI. Empirical data are for children with 'grammatical' SLI (gSLI) taken from van der Lely and Ullman (2001). An approximate age matched control group is included for the same task, taken from Thomas et al. (2001) (gSLI group mean age: 11 years 2 months [n = 12]; control group mean age 9 years 10 months [n = 10]). Simulation data are based on a model described in Thomas and Karmiloff-Smith (2003) and are presented for typical and atypical (low discriminability) developmental conditions, matched for level of training. Regular = regularly inflected verbs. Irregular = verbs irregularly inflected via an internal vowel change (e.g., sing-sang). OG errors = Over-generalisation errors on irregular verbs (e.g., sing-singed). NR + ed = application of -ed past tense rule to novel stems that do not rhyme with existing irregular verbs. For Regular and Irregular, responses = proportion correct. For OG errors, responses = proportion of all responses showing this error type. For NR + ed, responses = proportion of novel stems regularised.*

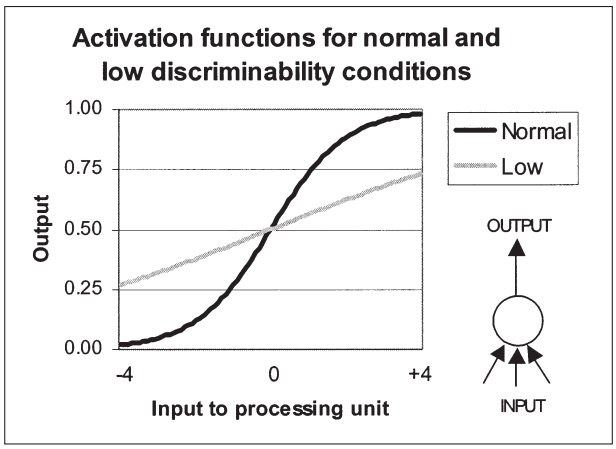

Fig. 2 – *The startstate computational manipulation in the atypical simulation condition. All processing units in the hidden layer and output layer of a 3-layer feedforward network had the discriminability of their activation functions reduced.*

between small changes in the signal it receives. The flatter function depicted in Figure 2 has the opposite effect, so that larger changes in the input signal are required to produce a given change in output. The consequence is a reduction in discriminability. In computational terms, the manipulation means that the network struggles to learn sharp category boundaries in representational space, and takes much additional training in order to do so.

Figure 1 demonstrates that this initial change to the computational properties of the learning system successfully simulates the qualitative pattern of the SLI data. In the atypical network, inflection levels are low, the delay is greater for regular verbs, and there is a small residual regularity effect, with small residual amounts of over-regularisation errors and rule generalisation. This simulation was not intended to be a detailed model of SLI. In its current form it has at least two shortcomings. First, the predominant error pattern in children with SLI is to produce uninflected stems. This is not what the simulation does.[1] Second, the model's deficits can be overcome by additional training, so that the system eventually exhibits ceiling performance on regulars and irregulars (though a deficit persists in generalisation). Its profile therefore demonstrates features in common with developmental delay rather than modelling the more persistent impairment typically seen in SLI. However, there is one further notable similarity between the model and SLI data. Figure 3 plots the frequency effects found in past tense production in the four experimental groups reported by van der Lely and Ullman (2001). These data exemplify the claim that regular verbs show greater frequency effects in children with SLI than in (younger) healthy children. Figure 4 illustrates the frequency effects in the network under normal and atypical training conditions. The model also exhibits an increased frequency effect in regular past tense production in the atypical condition. This effect arises because the network is less able to generate inflected stems based on phonological regularities, since this requires sharp category boundaries is representational space. As a consequence, the model tends to exploit lexical information to drive the output for all verb types. Such mappings do not reinforce each other by overlapping similarity at input, and therefore they are more influenced by token frequency of individual items. Figure 5 demonstrates that under normal conditions, a lesion to lexical input only impacts irregular verbs. In the low discriminability condition, however, it impacts both regular and irregular verbs, confirming the atypical role of lexical input in driving regular inflection.

The simulation of scenario B is important for two reasons. First, the compensatory pattern of development caused by atypical starting conditions comes fairly close to *implementing* some of U and P's claims. The system is particularly impaired in learning regularities and tends to rely on declarative memory to drive both regular and irregular inflections. The network has moved towards treating regulars and irregulars equivalently, so that both exhibit frequency effects. Second, and crucially, to simulate this pattern, the deficit applied to the startstate was *not to a domain-specific mechanism* for learning regulars. It was to a *domain-general computational resource* shared by both regulars and irregulars. It so happens that the altered property of the domain-general resource was one upon which regulars particularly relied. One might say this property was *domain-relevant* to regular inflection (Karmiloff-Smith, 1998). Regularity requires sharp category boundaries so that all items within a category are treated equivalently according to the 'rule'.[2] Reduced discriminability means that a great deal more training is required to generate the requisite sharp boundaries, and although regulars and irregulars were eventually acquired with high levels of training, the boundaries were never sharp enough to support good generalisation. Note that the "domain-generality" of the deficit here is only general with reference to the regular/irregular divide. But it is possible to imagine that, were this computational property to be impaired across a range of systems, it might also explain other features of SLI if they rely on sharp category boundaries. For instance, poor representational

[1] A more general model of inflection would be required to capture the prevalence of uninflected stem errors. A model that produces all inflectional paradigms (e.g., talk ⇒ talked, talks, talking, talker, talkest) can demonstrate stem errors under non-optimal conditions, since the verb stem emerges via a prototype effect as the most robust output – it is common to all regularly inflected forms (see Hoeffner and McClelland, 1993, for a demonstration of this effect in an early connectionist model of SLI).

[2] See Joanisse (2000, 2004) for a connectionist model of SLI in which a similar effect was achieved by adding noise to phonological representations, as part of a theory that a deficit in phonology lies behind inflectional impairments.

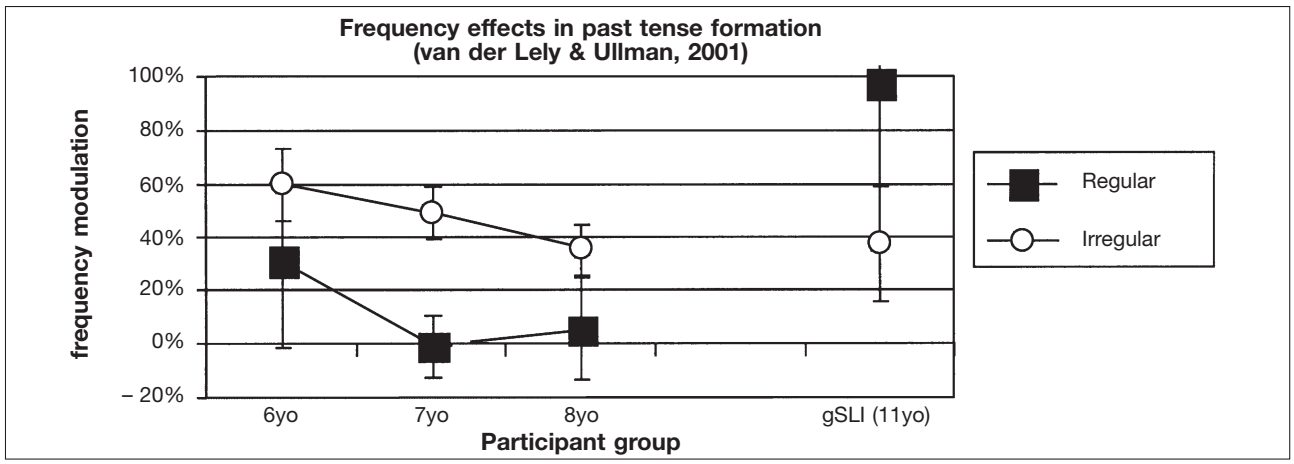

Fig. 3 – *Frequency effects in the accuracy of elicited past tense production in SLI and three younger control groups (van der Lely and Ullman, 2001) [yo = years old]. Scores show the modulation caused by frequency (derived by dividing the difference between performance on high and low frequency verbs by their mean level of performance). Error bars represent standard errors.*

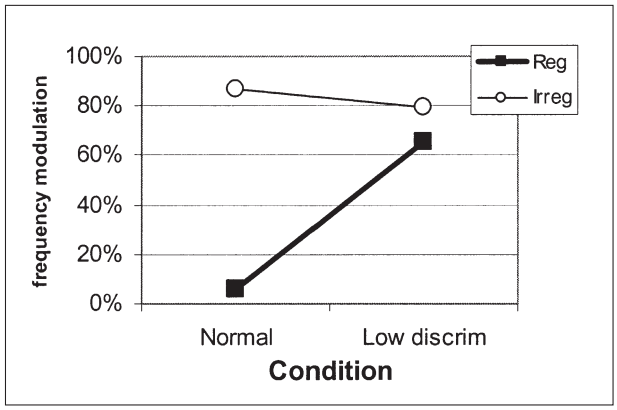

Fig. 4 – *The modulation caused by frequency in the model, for normal and low discriminability learning conditions.*

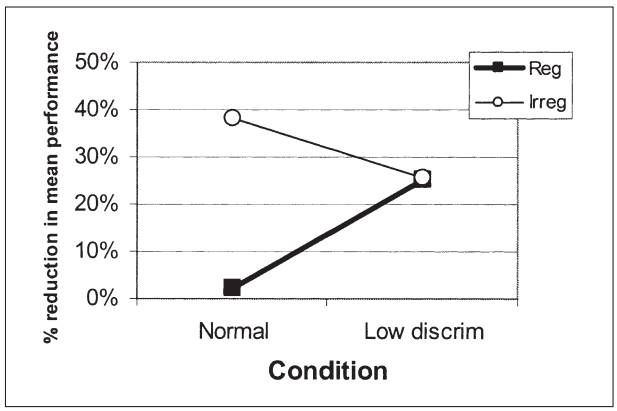

Fig. 5 – *The modulation of performance caused by removing lexical-semantic input to the model for regular and irregular verbs.*

acuity might explain a fast temporal processing deficit; or the more subtle deficits to the lexicon that are hard to incorporate into the PD theory, such as poorly defined semantic representations.

U and P would undoubtedly have reservations about the particular computational model employed in these simulations. However, the main thrust in presenting these results is that it is hard to guess the outcome of a compensatory developmental process without clear specification – and preferably implementation – of that process. The nature of the process affects the inferences that should be drawn from behavioural impairment to underlying structural deficit. From the behavioural pattern in SLI, U&P draw the inference that there is a deficit in a domain-specific mechanism for regular inflection. The modelling results do not disprove this theory, but they do indicate that, in a model with an explicitly implemented compensation process, the SLI pattern of data can be reproduced via a domain-general computational deficit that does not respect the regular-irregular / procedural-declarative distinction.

Conclusion

Ullman and Pierpont's Procedural Deficit theory, with its duality between procedural and declarative memory and potential for compensation between them, raises testable predictions about the pattern of impairments in SLI, and the underlying brain bases of this condition. It also places a welcome emphasis on processes of compensation, which must lie at the heart of an explanation of a developmental deficit, due to the plasticity and interactivity of the developing brain. The difficulty with compensation is that, unless it is clearly specified as a developmental process, we cannot properly evaluate any developmental theory that (quite rightly) seeks to include it. Computational models demonstrate that our intuitions about how compensation may operate can be seriously misleading. However, the very fact that U&P's model prompts us to address these issues is indicative of the progress it may stimulate. Moreover, the advantage of the PD theory is that data from brain imaging may offer independent justification of the involvement of different memory systems in language function.

Acknowledgements. This work was supported by Medical Research Council Career Establishment Grant G0300188 to Michael Thomas

References

BATES E and ROE K. Language development in children with unilateral brain injury. In CA Nelson and M Luciana (Eds), *Handbook of Developmental Cognitive Neuroscience*, 2001. Cambridge, MA: MIT Press, pp. 281-307.

BISHOP DVM. Cognitive neuropsychology and developmental disorders: Uncomfortable bedfellows. *Quarterly Journal of Experimental Psychology, 50A:* 899-923, 1997.

FOWLER A. Language in mental retardation: Associations with and dissociations from general cognition. In JA Burack, RM Hodapp and E Zigler (Eds), *Handbook of Mental Retardation and Development*, 1998. Cambridge: Cambridge University Press, pp. 290-33.

HOEFFNER JH and MCCLELLAND JL. Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A computational investigation. In EV Clark (Ed), *Proceedings of the 25th Child language research forum,* 1993. Stanford: Center for the Study of Language and Information, pp. 38-49.

JOANISSE MF. Connectionist Phonology. Unpublished Ph.D. Dissertation, University of Southern California, 2000.

JOANISSE MF. Specific language impairments in children: Phonology, semantics and the English past tense. *Current Directions in Psychological Science, 13:* 156-160, 2004.

JOANISSE MF and SEIDENBERG MS. Impairments in verb morphology following brain injury: A connectionist model. *Proceedings of the National Academy of Science USA, 96:* 7592-7597, 1999.

KARMILOFF-SMITH A. Development itself is the key to understanding developmental disorders. *Trends in Cognitive Sciences, 2:* 389-398, 1998.

KARMILOFF-SMITH A and THOMAS MSC. What can developmental disorders tell us about the neurocomputational constraints that shape development? The case of Williams syndrome. *Development and Psychopathology, 15:* 969-990, 2003.

LAVRIC A, PIZZAGALLI D, FORSTMEIER S and RIPPON G. Mapping dissociations in verb morphology. *Trends in Cognitive Sciences, 5:* 301-308, 2001.

MCDONALD JL. Language acquisition: The acquisition of linguistic structure in normal and special populations. *Annual Review of Psychology, 48:* 215-241, 1997.

NORBURY CF, BISHOP DVM and BRISCOE J. Production of English finite verb morphology: A comparison of SLI and mild-moderate hearing impairment. *Journal of Speech, Language and Hearing Research, 44:* 165-178, 2001.

OETTING J and HOROHOV J. Past tense marking by children with and without specific language impairment. *Journal of Speech and Hearing Research, 40:* 62-74, 1997.

PINKER S. Words and rules. London: Weidenfeld and Nicolson, 1999.

PLAUT DC and SHALLICE T. Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology, 10:* 377-500, 1993.

PLUNKETT K and JUOLA P. A connectionist model of English past tense and plural morphology. *Cognitive Science, 23:* 463-490, 1999.

PLUNKETT K and MARCHMAN V. U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition, 38:* 1-60, 1991.

TAGER-FLUSBERG H and SULLIVAN K. Early language development in children with mental retardation. In JA Burack, RM Hodapp and E Zigler (Eds), *Handbook of Mental Retardation and Development*, 1998. Cambridge University Press, pp. 115-131.

THOMAS MSC. Limits on plasticity. *Journal of Cognition and Development, 4:* 95-121, 2003.

THOMAS MSC and KARMILOFF-SMITH A. Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioural and Brain Sciences, 25:* 727-750, 2002.

THOMAS MSC and KARMILOFF-SMITH A. Modelling language acquisition in atypical phenotypes. *Psychological Review, 110:* 647-682, 2003.

THOMAS MSC and KARMILOFF-SMITH A. Can developmental disorders reveal the component parts of the human language faculty? *Language Learning and Development, 1:* 65-92, 2005.

THOMAS MSC, GRANT J, GSÖDL M, LAING E, BARHAM Z, LAKUSTA L, TYLER LK, GRICE S, PATERSON S and KARMILOFF-SMITH A. Past tense formation in Williams syndrome. *Language and Cognitive Processes, 16:* 143-176, 2001.

THOMAS MSC and RICHARDSON F. Atypical representational change: Conditions fot the emergence of atypical modularity. In M. Johnson and Y Munakata (Eds), *Attention and Performance XXI.* Oxford: Oxford University Press.

ULLMAN MT and GOPNIK M. Inflectional morphology in a family with inherited specific language impairment. *Applied Psycholinguistics, 20:* 51-117, 1999.

ULLMAN MT and PIERPONT E. Specific Language Impairment is not specific to language: The Procedural Deficit hypothesis. *Cortex,* (this issue).

ULLMAN MT, HARTSHORNE JK, ESTABROOKE IV, BROVETTO C and WALENSKI M.. Sex, regularity, frequency and consistency: A study of factors predicting the storage of inflected forms. (Under Revision)

VAN DER LELY HKJ. Language and cognitive development in a grammatical SLI boy: Modularity and innateness. *Journal of Neurolinguistics, 10:* 75-107, 1997.

VAN DER LELY HKJ and ULLMAN MT. Past tense morphology in specially language impaired and normally developing children. *Language and Cognitive Processes, 16:* 177-217, 2001.

WATKINS KE, DRONKERS NF and VARGHA-KHADEM F. Behavioural analysis of an inherited speech and language disorder: Comparison with acquired aphasia. *Brain, 125:* 452-464, 2002.

Michael S.C. Thomas, School of Psychology, Birkbeck College, University of London, Malet Street, Bloomsbury, London WC1E 7HX, UK. e-mail: m.thomas@bbk.ac.uk