Research Article

THE MOSES, MEGA-MOSES, AND ARMSTRONG ILLUSIONS: Integrating Language Comprehension and Semantic Memory

Meredith Shafto and Donald G. MacKay

University of California, Los Angeles

Abstract—*This study develops a new theory of the* Moses illusion, *observed in responses to general knowledge questions such as, "How many animals of each kind did Moses take on the Ark?" People often respond "two" rather than "zero" despite knowing that Noah, not Moses, launched the Ark. Our theory predicted two additional types of conceptual error demonstrated here: the Armstrong and mega-Moses illusions. The Armstrong illusion involved questions resembling, "What was the famous line uttered by Louis Armstrong when he first set foot on the moon?" People usually comprehend such questions as valid, despite knowing that Louis Armstrong was a jazz musician who never visited the moon. This Armstrong illusion was not due to misperceiving the critical words (Louis Armstrong), and occurred as frequently as the Moses illusion (with critical words embedded in identical sentential contexts), but less frequently than the mega-Moses illusion caused when Moses and Armstrong factors were combined.*

How do people know when they understand a sentence, and why do they sometimes think they understand, but do not? These fundamental issues for understanding human communication and learning are mirrored in the *Moses illusion*, the fact that people often miscomprehend questions such as, "How many animals of each kind did Moses take on the Ark?" responding "two" rather than "zero" with high confidence, even though they know that Noah took animals on the Ark, but Moses did not (Erickson & Mattson, 1981). In this article, we first review current facts and theories about the Moses illusion and then develop a new theory that explains these facts, and predicts new types of Moses mistakes that no other theory predicts.

THE MOSES ILLUSION: CURRENT FACTS AND THEORIES

Moses mistakes are not due to hasty responding or unfamiliarity with "trick questions": When instructions illustrate Moses-like questions in advance and advise against hasty response, participants with unlimited time to respond still make Moses mistakes (e.g., Reder & Kusbit, 1991). Nor do Moses mistakes reflect hypercorrection, with participants detecting the Moses-for-Noah substitution and "correcting" the question before responding (see, e.g., Reder & Cleeremans, 1990). However, task, syntactic structure, sentential focus, instructions, and recent experience all influence the frequency of Moses mistakes (e.g., Bredart & Modolo, 1988; Buyer & Radvansky, 1995; Erickson & Mattson, 1981). For example, if participants first study a relevant fact about Noah (e.g., "Noah took two animals of each kind on the Ark"), they make fewer Moses mistakes (Reder & Kusbit, 1991), and if they learn a misleading association (e.g., "Moses-Ark"), they make more Moses mistakes (Kamas, Reder, & Ayers, 1996; Reder & Cleeremans, 1990; van Oostendorp & Kok, 1990).

Current theoretical accounts of Moses mistakes (Barton & Sanford, 1993; Erickson & Mattson, 1981; Reder & Kusbit, 1991; van Oostendorp & Kok, 1990) fall into two classes. According to one of these classes, everyday semantic processing tends to be overly general and therefore error-prone; according to the other, people check incoming information against only a subset of their stored semantic information because a complete check requires too much time and effort. Consistent with both of these semantic-level explanations[1] is the finding that Moses mistakes are not due to sensory-level misperceptions of "Noah" for "Moses": Participants still make Moses mistakes after correctly reading the questions aloud (Reder & Kusbit, 1991).

A NEW THEORY OF MOSES MISTAKES, THE ARMSTRONG PREDICTION, AND THE PRESENT TASKS

Our theory of the Moses illusion is an interactive activation model of memory and language known as node structure theory (NST; MacKay, 1987), which postulates a vast network of interconnected representational units called nodes. Nodes are organized into a *semantic system* (representing the meanings of words, phrases, and propositions) and a *phonological system* (representing hierarchically organized syllables, phonological clusters, and speech sounds). Unique to NST, the same nodes that perceive and comprehend a word within these systems also retrieve and produce the word. NST also makes a fundamental distinction between unconscious versus conscious processes: Node priming and node activation constitute two distinct but interrelated unconscious processes for perceiving and producing familiar words, but a third process (prolonged activation) is necessary to integrate word meanings into consciously comprehended phrases and propositions (see MacKay, 1990; MacKay & Burke, 1990; MacKay, Stewart, & Burke, 1998). Subsequent sections show how these processes explain Moses mistakes within the context of the two tasks used in the experiments we report here.

NST and the Partial Shadowing Procedure

Our first task was a new control procedure designed to exclude momentary attentional lapses as contributors to Moses mistakes: This

Address correspondence to Donald G. MacKay, Psychology Department, University of California, Los Angeles, 405 Hilgard Ave., Los Angeles, CA 90095-1563; e-mail: mackay@psych.ucla.edu.

1. Current theories have excluded phonology because of a single nonsignificant result in Erickson and Mattson (1981, Experiment 2). However, this result was problematic, involving only a single semantic context (biblical), very few sentence frames (two), a primitive manipulation of phonology, and confounds between phonological and semantic factors due to "difficulty in varying phonological features while holding semantic similarity constant" (Erickson & Mattson, 1981, p. 547). The field apparently ignored Erickson and Mattson's caveat that "an effect due to phonological similarity cannot be entirely discounted" (also on p. 547).

partial shadowing task demanded special attention to *Moses* and other critical words in our auditorily presented questions because participants silently read a written version of each question containing one or more blank slots, and had to shadow, or repeat aloud with minimal lag, whatever auditory word occupied each blank slot. For example, participants hearing the question "What was the famous line uttered by Louis Armstrong when he first set foot on the moon?" shadowed the words *Louis*, *Armstrong*, and *moon* (see Table 1).

Correct shadowing of *Louis Armstrong* in this *Armstrong question* illustrates how priming, activation, and prolonged activation function in comprehension-perception versus retrieval-production in NST. Figure 1 depicts some of the relevant nodes and their parallel bottom-up and top-down connections. Bottom-up connections are central to comprehension-perception of these units, whereas top-down connections are central to retrieval-production. Following acoustic presentation of *Louis Armstrong*, bottom-up connections automatically prime the hierarchy of phonological nodes for *Louis* and *Armstrong*, the lexical nodes for *Louis* and *Armstrong*, and the name-phrase node for *Louis Armstrong* in Figure 1. Priming automatically summates across all simultaneously active connections converging on a node, but directly causes neither activation nor prolonged activation (awareness) of a concept such as "Louis Armstrong." Node activation requires application of a domain-specific activation mechanism at a discrete point in time to a domain of nodes, here, the domain of name-phrase nodes. The activation mechanism then activates whatever node in the domain has most priming at that time, a principle known as "most-primed-wins."

This most-primed-wins activation principle provides the basis for *shadowing repairs*, the unconscious corrections that occur when shadowing a sentence that contains experimentally introduced mispronunciations such as "untegration" instead of "integration" (e.g., Cole & Scott, 1974). Word production in skilled shadowing begins with activation of whatever lexical node has accumulated most bottom-up priming, in this case, *integration*, because *untegration* has no lexical node. Activating *integration* then delivers top-down priming to its syllables *in* + *te* + *gra* + *tion*, so that *in*(syllable) will accrue more priming than *un*(syllable) and become activated under the most-primed-wins principle, thereby determining the shadowing repair (see MacKay, 1987, p. 132). Such corrections often lag the input by less than 300 ms (Marslen-Wilson, 1975), making it unlikely that higher-level phrases and propositions (which often last 1,000–2,000 ms) are guiding skilled shadowing. Shadowing tasks therefore illustrate how lower-level activation for shadowing the words *Louis* and *Armstrong* can in principle proceed independently of higher-level comprehension and awareness, here, prolonged activation of phrases and propositions in this Armstrong question.

NST and the Comprehension of Moses and Armstrong Questions

Our participants' second task was the main one: to comprehend and answer Moses and Armstrong questions. NST explains Moses mistakes as due to *convergent priming*. As can be seen in Figure 2, *Moses* receives one source of bottom-up priming in "How many animals of each kind did Moses take on the Ark?" whereas *Noah* receives two convergent sources of semantic-level priming. Preformed links in

Table 1. *Example auditory stimuli, corresponding text for partial shadowing, and results from Experiment 1*

Condition	Example auditory stimulus	Text for silent reading with blanks for partial shadowing	Mean number of "can't say" responses per participant	Mean proportion of "can't say" responses
Semantically related experimental questions	What was the famous line uttered by Alan Shepard when he first set foot on the moon?	What was the famous line uttered by _____ _____ when he first set foot on the ________?	3.00 (1.69)	.67 (.29)
Phonologically related experimental questions	What was the famous line uttered by Louis Armstrong when he first set foot on the moon?	What was the famous line uttered by _____ _____ when he first set foot on the ________?	3.82 (1.25)	.77 (.19)
Unrelated experimental questions	What was the famous line uttered by Dizzy Gillespie when he first set foot on the moon?	What was the famous line uttered by _____ _____ when he first set foot on the ________?	4.89 (1.13)	.91 (.17)
Valid experimental questions	What was the famous line uttered by Neil Armstrong when he first set foot on the moon?	What was the famous line uttered by _____ _____ when he first set foot on the ________?	0.50 (0.75)	.114 (.17)
Valid filler questions	How many letters are there in the alphabet?	How many letters are there in the ________?	0.21 (0.50)	.02 (.05)
Invalid filler questions	What is the name of the ferocious striped airplane found in India?	What is the name of the ferocious striped ________ found in India?	9.36 (0.73)	.95 (.07)

Note. Standard deviations are given in parentheses.

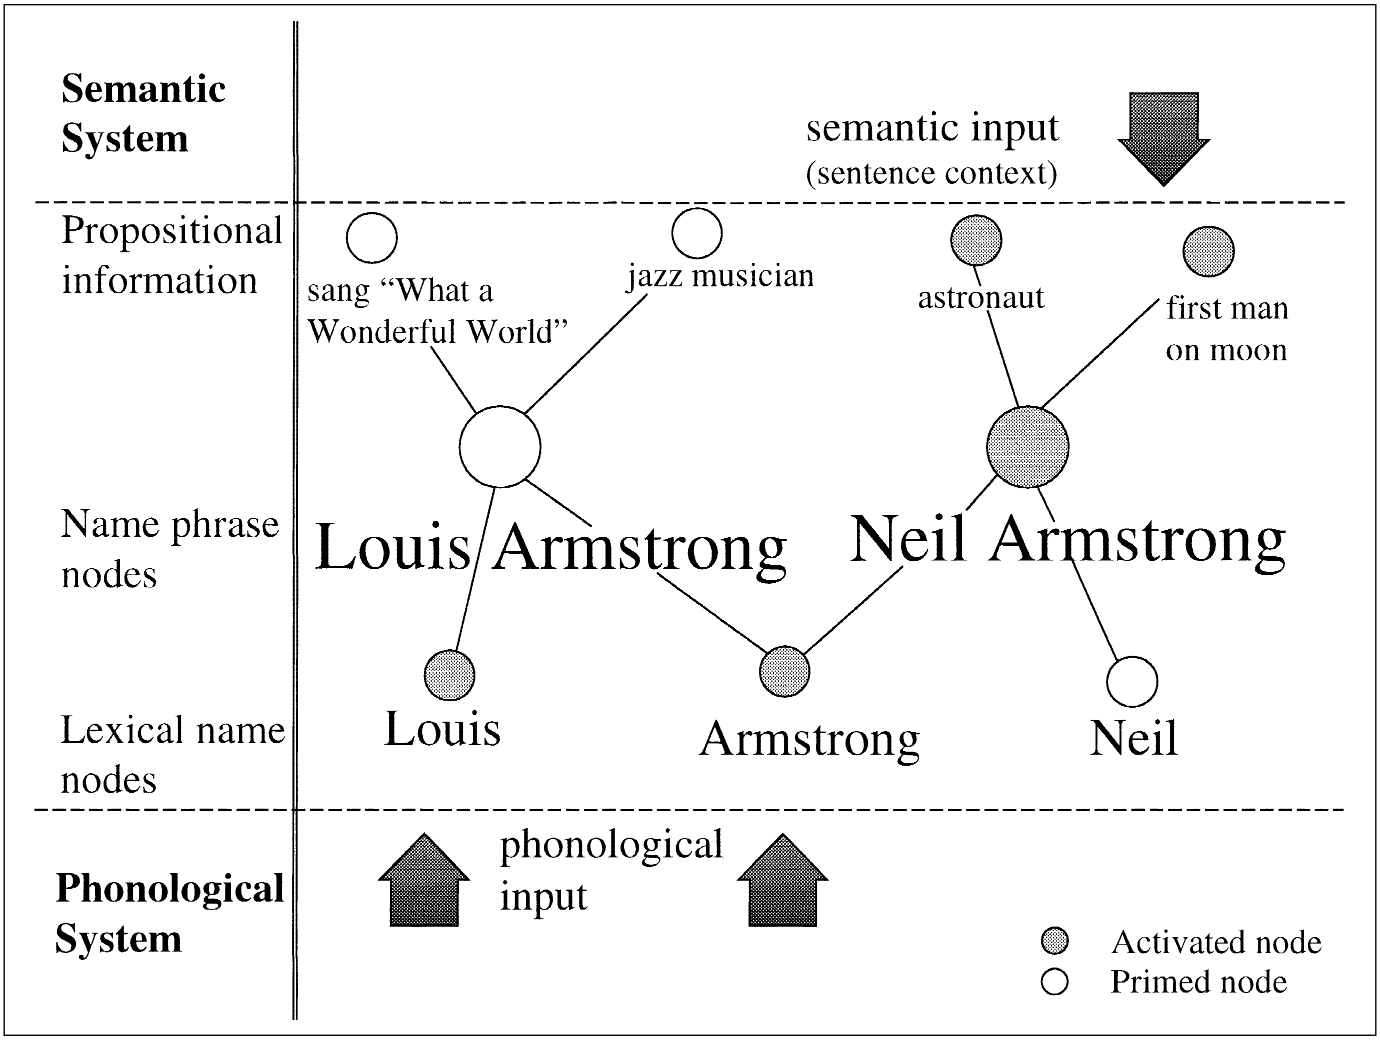

Fig. 1. A subset of semantic system nodes and their two-way connections in node structure theory for comprehending and producing selected words and concepts in the sentence, "What was the famous line uttered by Louis Armstrong when he first set foot on the moon?" Unfilled nodes receive priming, whereas filled nodes also undergo activation or prolonged activation during comprehension (see the text for explanation).

semantic memory between *Noah* and the concepts "built the Ark" and "animals of each kind" in the question provide one source of priming, and semantic similarities linking the critical names provide the other: Because two-way connections link *Moses* to *Noah* via identical predicates (e.g., *is male, is a biblical figure, spoke with God,* and *saved his people from adversity*), hearing *Moses* indirectly primes *Noah* via these shared predicate nodes. Receiving more priming than *Moses,*[2] *Noah* therefore becomes activated under the most-primed-wins principle, causing miscomprehension of *Moses* as *Noah*.

For Armstrong questions, NST predicts similar convergent priming due in part to shared phonology. As can be seen in Figure 1, *Louis Armstrong*(name phrase) receives bottom-up phonological priming from *Louis* and *Armstrong*, whereas *Neil Armstrong* receives convergent priming from a bottom-up phonological source (*Armstrong*) and from sentence-level semantic sources (i.e., preformed links in semantic memory all feed priming to *Neil Armstrong* from the concepts "moon," "first set foot on," and "famous line" in the question). Thus, even when participants correctly shadow the words *Louis* and *Armstrong*, *Neil Armstrong* will accumulate more convergent priming than *Louis Armstrong*, and become activated as the most-primed node in the name-phrase domain, causing miscomprehension of *Louis Armstrong* as *Neil Armstrong*. *Neil Armstrong* will therefore become integrated with the remainder of the sentence, thereby determining the conscious response, here, "The line was . . . " rather than "can't say." Although we have replicated this Armstrong illusion several times using different stimuli and procedures (Shafto & MacKay, 1998a, 1998b), the present report is the first to compare Armstrong and Moses effects in the same task with identical focus and instructional demand characteristics and using partial shadowing procedures to rule out attentional-lapse accounts.

2. Note that prior learning of the misleading association "Moses-Ark" will introduce a third source of convergent priming (bottom-up from *Moses* to *Ark* to *Noah*) that will further increase Moses mistakes under NST, consistent with data discussed earlier.

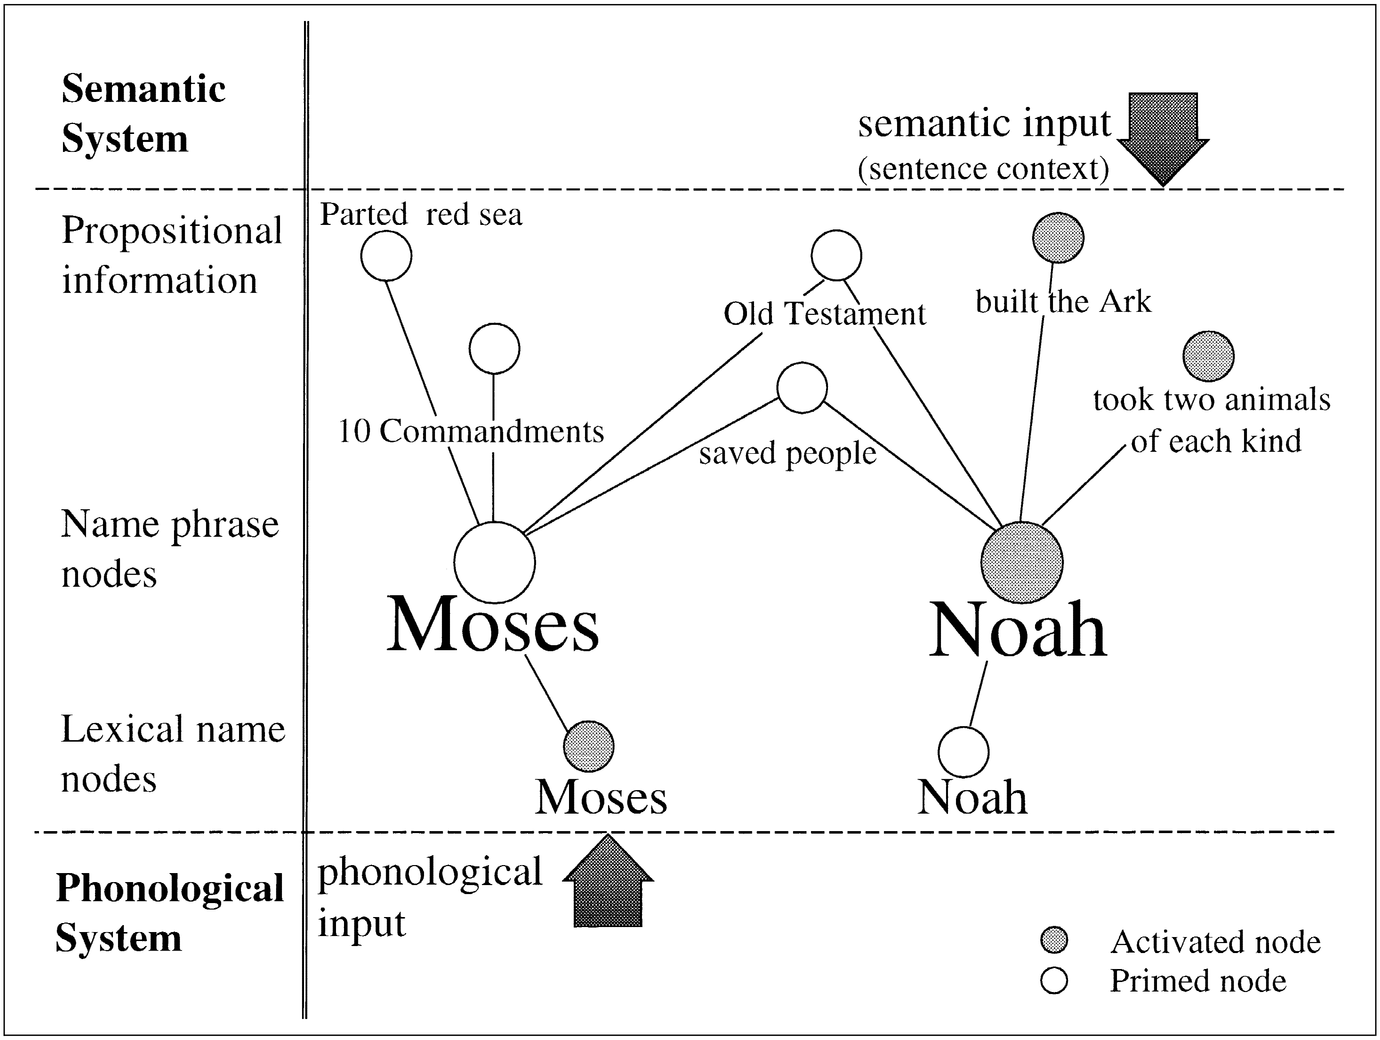

Fig. 2. A subset of semantic system nodes and their two-way connections in node structure theory for comprehending and producing selected words and concepts in the sentence, "How many animals of each kind did Moses take on the Ark?" Unfilled nodes receive priming, whereas filled nodes also undergo activation or prolonged activation during comprehension (see the text for explanation).

EXPERIMENT 1

Stimuli

The participants (28 undergraduates at the University of California, Los Angeles; mean age of about 20) heard three types of stimuli: practice questions ($n = 10$), experimental questions ($n = 24$), and filler questions ($n = 20$). Each experimental question contained a person's first and last names, and came in four different versions that differed only in proper name: a *valid* version, which contained the appropriate name, and three *invalid* versions, which contained names that were *semantically related*, *phonologically related*, or *unrelated* to the valid name (see Table 1 for examples). These four versions defined our conditions, and first names differed across all four conditions (e.g., *Neil*, *Alan*, *Louis*, and *Dizzy* in Table 1), whereas last names were identical in the valid and phonologically related conditions (e.g., *Armstrong*). The unrelated and phonologically related conditions involved semantically similar names (e.g., *Dizzy Gillespie* and *Louis Armstrong* were both jazz musicians) that differed semantically from the valid and semantically related names (e.g., *Neil Armstrong* and *Alan Shepard* were astronauts). Half the fillers were valid questions, and half were invalid or contained conflicting information (see Table 1).

In constructing our final stimuli, we ensured that participants were likely to know the correct answers to our valid and filler questions, and we equated familiarity of the famous names across the four conditions. To do this, we "filtered" more than 275 stimulus alternatives through four sets of pilot studies. In Filter Study 1, 10 participants rated their familiarity with the famous names on a scale from 0 to 4, and indicated the individuals' "reasons for fame." Names with reasons for fame that 3 or more participants identified incorrectly were filtered out, and we incorporated the remaining names into pairs of valid phonologically related questions (e.g., valid *Louis Armstrong* and *Neil Armstrong* questions). Then 10 new participants in Filter Study 2 answered these questions and rated confidence in their answers on a 5-point scale. Whichever member of a pair received more correct answers (or if the members were tied, higher confidence ratings) was then used to create a general knowledge question, which passed through a "knowledge filter" (Study 3), a five-choice recognition memory test involving names from Filter Study 1 plus a "don't know" response. For example, for the knowledge question "What astronaut uttered the famous line, 'That's one small step for [a] man, one giant leap for mankind,' " the response alternatives were James Brown, Neil Armstrong, Charles Lindbergh, and Louis Armstrong. Questions answered correctly most often in Filter Study 3 were then used to construct our final experimental questions. Filter Study 4 ensured that filler questions were easily answered, enabling a check on whether participants were using response alternatives appropriately in the experiment proper.

Procedure

Fillers were recorded on four audiotapes interspersed among experimental questions, with a 5-s gap between questions. Each tape contained the practice questions, plus one version of each experimental question, with versions counterbalanced across tapes so that conditions occurred in different orders and equally often across tapes.

Participants heard each question over headphones, and shadowed whichever words fit into blank slots in a written version that they read concurrently (see Table 1). Participants then selected one of four multiple-choice answers beside each written question: "don't know," "can't say," the correct answer to the valid version (e.g., "26" for the filler "How many letters are there in the alphabet?"), and an incorrect "ballpark" answer (e.g., "31" for the alphabet filler question). To obviate special strategies associated with proper names, fillers (unlike experimental questions) never contained shadowing slots for proper names.

Instructions illustrated correct and incorrect responses to valid and invalid fillers, and explicitly warned participants about invalid questions containing conflicting information. Following the practice trials, the experimenter again answered participants' questions, repeated the partial shadowing instructions, and redescribed the response alternatives with new examples.

Subsidiary results in Experiments 1 and 2 indicated that participants were motivated and capable of choosing the critical response ("can't say") appropriately. Experiment 1 was typical, in that participants responded "can't say" appropriately for 95% of the invalid fillers, and rarely responded "can't say" inappropriately for valid filler and experimental questions (2% and 11%, respectively; see Table 1). Analyses excluded trials involving skipped responses ($n = 2$) and incorrectly shadowed critical words ($n = 4$).

Results and Discussion

Table 1 shows mean numbers of "can't say" responses per participant by condition, with corresponding proportions that excluded "don't know" responses. A repeated measures analysis of variance on proportions of correct ("can't say") responses in the semantically related, phonologically related, and unrelated conditions indicated a main effect of condition, $F(2, 54) = 9.83$, $MSE = 0.04$, $p < .001$, that reflected a higher proportion of correct responses in the unrelated than in the phonologically related condition (the Armstrong illusion), $t(27) = -3.22$, $p < .01$, and in the unrelated than in the semantically related condition (the Moses illusion), $t(27) = -4.12$, $p < .001$. Identical tests with stimuli as the unit of analysis were also reliable ($p = .001$, .023, and .001, respectively). However, the phonologically and semantically related conditions did not differ reliably using either participants or stimuli as analytic units.

These results establish the Moses and Armstrong illusions within identical syntactic and semantic contexts in the same task, and rule out inattention to the critical names (e.g., *Louis Armstrong*) as a possible account of these illusions. Indeed, participants must have focused special attention on these critical names when accurately and selectively shadowing them within our sentences. This partial shadowing control therefore establishes a remarkable dissociation between perceiving words as phonological versus semantic entities. To accurately shadow *Louis Armstrong* in the phonologically related condition, participants must have perceived *Louis Armstrong* at phonological levels, but comprehended *Neil Armstrong* at phrase and proposition levels, so as to (erroneously) choose the response, "One small step for a man, one giant leap for mankind." Explaining this dissociation requires a theory resembling NST, in which the processes determining the final products of comprehension (prolonged activation) occur at semantic but not phonological levels (see the introduction).

The present results also rule out purely semantic theories (e.g., Erickson & Mattson, 1981; Reder & Kusbit, 1991; van Oostendorp & Kok, 1990) as viable accounts of the Armstrong effect (i.e., more "can't say" responses for unrelated than phonologically related questions). Unrelated and phonologically related names were themselves semantically similar, and therefore differed equally in semantics from target names (e.g., jazz musicians *Louis Armstrong* and *Dizzy Gillespie* are equally unlike astronaut *Neil Armstrong*).

However, ruling out purely semantic theories does not rule in purely phonological theories of the Armstrong effect. What is required is a theory resembling NST, in which phonology can influence semantics via one process (priming) even though other processes (prolonged activation and semantic priming from sentence context) determine final comprehension at semantic levels. More specifically, three summated factors (one phonological and two semantic) caused the unpresented name (*Neil Armstrong*) to accumulate more priming than the presented name (*Louis Armstrong*) and become activated: shared surname phonology, a shared lexical-surname node, and semantic-level sentence context (see Fig. 1).

However, not all three factors are necessary under NST. For example, even without surname identity, NST predicts weak but reliable effects of shared phonology on Armstrong and Moses mistakes. An interesting example concerns the original Moses question. If the many phonological nodes shared by *Moses* and *Noah* (representing bisyllabicity, stress pattern, initial vowels, onset nasality, and onset voicing) magnified the original Moses illusion, this would explain why replacing *Noah* with the semantically similar but phonologically dissimilar *Abraham* diminished without entirely eliminating the Moses effect in the study by Erickson and Mattson (1981): Whereas two factors determined their Moses effect (phonological and semantic similarities), only one (semantic similarity) determined their Abraham effect.

EXPERIMENT 2: THE MEGA-MOSES AND MISCOMPREHENSION PREDICTIONS

If, as the preceding discussion suggests, phonological similarity without surname identity enhanced the original Moses illusion, then phonological similarity *with* surname identity should further enhance the Moses illusion, causing an even greater, “mega-Moses” illusion under NST. Consider the effects of substituting three alternative names for *Andrew Johnson* in the question, “The 1868 impeachment trial involving former vice president Andrew Johnson followed what major American war?” NST predicts a standard Armstrong effect from substituting *Samuel Johnson* for *Andrew Johnson* due to shared phonological and lexical-surname nodes for *Johnson*. NST predicts a standard Moses effect from substituting *Theodore Roosevelt* for *Andrew Johnson* because these phonologically dissimilar names share many predicates within semantic memory (e.g., “was vice president” and “became president because of presidential assassination”). However, NST predicts an enhanced, or mega-Moses, effect from substituting *Lyndon Johnson* for *Andrew Johnson*. In this case, standard Moses factors (semantic similarities linking *Andrew* and *Lyndon Johnson*, both former vice presidents who became president following presidential assassinations) will combine with standard Armstrong factors (shared phonological and lexical-surname nodes for *Johnson*) to cause a supranormal effect.

Testing this mega-Moses prediction was the main goal of Experiment 2. A secondary goal was to test whether miscomprehending invalid names as valid ones is a primary determinant of Moses and Armstrong effects (see the introduction). Experiment 2 tested for miscomprehension indirectly and post hoc using a procedure developed by MacKay (1973): After the experiment, participants received a surprise recognition memory test for the gist or meaning of previously presented auditory questions. Participants knew that each visually presented recognition target resembled one of the original questions, but differed in word order; their task was to respond “yes” if the recognition target and prior question were synonymous, and “no” otherwise. In fact, recognition targets were always synonymous with valid versions of corresponding auditory questions, even for originally invalid questions. For example, an invalid auditory question resembling “What was the famous line uttered by Louis Armstrong when he first set foot on the moon?” had the recognition target “When first setting foot on the moon, Neil Armstrong uttered what famous line?” NST predicted reliably more “yes” (synonymous) responses to memory questions with auditory counterparts in related than unrelated (control) conditions for trials on which participants experienced the Moses and Armstrong illusions in the main task, but no reliable difference when participants responded correctly in the main task.

Method

Methods were identical to those of Experiment 1 except that 33 participants heard 32 fillers, 16 Armstrong questions, and 16 nonoverlapping Moses questions in the shadow-and-answer task. Valid and unrelated versions formed the standard control conditions, and the Armstrong comparison condition resembled the phonologically related condition of Experiment 1 (see Table 1). Moses comparison conditions were the standard (semantic similarity) and mega-Moses (combined semantic and phonological similarity) conditions. So we could assess comprehension of invalid names, participants performed an irrelevant 10- to 15-min (distractor) task following the shadow-and-answer task and then received the surprise recognition memory test involving rephrased versions of valid questions presented in the same order as the corresponding (valid or invalid) questions in the shadow-and-answer task. Finally, participants answered five-choice knowledge questions (resembling those in Filter Study 3) to ensure that they could correctly answer valid but rephrased versions of corresponding experimental questions. Shadow-and-answer trials involving inaccurately answered knowledge questions were excluded from all analyses.

Results and Discussion

The Experiment 2 results firmly established the mega-Moses effect. A repeated measures analysis of variance on proportions of correct (“can’t say”) answers to invalid Moses questions indicated a main effect of condition, $F(2, 60) = 20.13$, $MSE = 0.06$, $p < .001$, with a higher proportion of correct responses for unrelated ($M = .91$) than mega-Moses questions ($M = .51$), $t(32) = -8.04$, $p < .001$, and for semantically related ($M = .70$) than mega-Moses questions, $t(30) = -2.47$, $p < .05$ (the mega-Moses effect). Identical tests with stimuli as the unit of analysis were also reliable ($p = .001$, .001, and .004, respectively). Because Moses and Armstrong factors combine in mega-Moses effects, the fact that standard Moses and standard Armstrong effects in Experiment 1 did not differ statistically clearly reflects stimulus design rather than theoretical necessity.

The results of Experiment 2 also replicated standard Moses and standard Armstrong effects (smallest *p*s < .01 and .001, respectively), and supported NST miscomprehension predictions: For trials on which participants experienced the Moses and Armstrong illusions in the main task, the proportion of “yes” (synonymous) responses was higher for memory questions with auditory counterparts in related ($M = .50$) than unrelated ($M = .20$) conditions, $t(13) = 2.68$, $p = .013$; however, the proportion of “yes” responses did not differ reliably for related versus unrelated conditions (*M*s = .11 and .06) when participants responded correctly in the main task. This pattern comports with NST claims that Moses and Armstrong effects tend to occur when participants miscomprehend (and therefore misrecall) an invalid name (e.g., *Louis Armstrong*) as the valid one (*Neil Armstrong*).

GENERAL DISCUSSION

Are Armstrong effects basically parlor curiosities limited to a single experimental paradigm involving misleading questions? Parallel phenomena in studies of speech production suggest not (e.g., experimentally induced “Freudian slips,” Motley & Baars, 1976, and the “yolk phenomenon,” see Baars, 1988, p. 307; also Bohannon & Bonvillian, 1997): After participants repeatedly produce a word such as *poke*, they usually answer the question, “What is the white part of an egg called?” with the phonologically similar “yolk” (despite knowing the correct answer, *albumen* or *egg white*). These and other independently established language production phenomena (Cutting & Ferreira, 1999; Dell, 1988) involve the same perception-production nodes and the same summated phonological and semantic priming as Armstrong effects in NST.

Does NST fully account for all three illusions? As developed so far, NST has established a single unified answer to the question, when these illusions occur, what is their origin? These illusions occur when a specious (unpresented) name (e.g., *Neil Armstrong*) receives more priming than the presented (invalid) name (*Louis Armstrong*). Only the source of this specious priming distinguishes the three illusions under NST: phonological and lexical-surname nodes shared by valid and invalid surnames for standard Armstrong effects, predicates shared by valid and invalid name phrases for standard Moses effects, and both sources of specious priming for mega-Moses effects.

However, this specious-priming framework requires a supplementary viewpoint to account fully for Armstrong, Moses, and mega-Moses effects. This supplementary viewpoint focuses on correct responses (i.e., nonoccurrence of the illusions) and addresses the question: What enables participants to detect anomalies when they do? The process of novelty detection (see MacKay, 1990) is central to this question under NST: Participants detect semantic anomalies when they become aware of novel information that conflicts with simultaneously activated information in semantic memory. From this novelty-detection viewpoint, two events are necessary for detecting semantic anomalies, say, in the unrelated Armstrong condition: The name-phrase node for *Dizzy Gillespie* rather than *Neil Armstrong* must become activated, and *Dizzy Gillespie* information in semantic memory (e.g., "was a jazz musician") must become activated concurrently with conflicting, question-based information (e.g., "an astronaut first set foot on the moon").

The results for the unrelated Armstrong condition serve to contrast these two viewpoints. The novelty-detection viewpoint focuses on the high proportion of correct ("can't say") responses to these control questions (.91 in Experiment 1 and .86 in Experiment 2), whereas the specious-priming viewpoint focuses on the sizable proportion of errors (.09 and .14 in Experiments 1 and 2, respectively).[3] Although the novelty-detection viewpoint cannot readily explain these errors, they again reflect specious priming within the specious-priming viewpoint; that is, concepts such as "moon" and "first set foot on" prime the unpresented (valid) name in unrelated-name questions, causing errors in the same manner as Armstrong questions except that Armstrong questions contribute additional sources of specious priming, and therefore cause more errors.

Sharper contrasts between novelty-detection and specious-priming viewpoints arise from manipulating fame and familiarity, two higher-level (nonphonological) dimensions of similarity between valid and invalid names in Armstrong questions. For example, *Neil* and *Louis Armstrong* were both famous and familiar names for our participants, and under the specious-priming viewpoint, shared fame and familiarity may influence Armstrong effects in the same manner as shared surname and predicate nodes—by transmitting priming to the unpresented (valid) name. This being the case, replacing *Neil Armstrong* with an unknown *Armstrong* in questions resembling, "What famous line did Rick Armstrong utter when first setting foot on the moon?" should remove this hypothesized contribution of fame and familiarity, and so reduce without entirely eliminating the standard Armstrong effect. However, the novelty-detection viewpoint predicts that replacing *Neil Armstrong* with *Rick Armstrong* should increase the standard Armstrong effect because an unknown name lacks connected propositions that conflict with propositions in semantic memory and thereby facilitate anomaly detection. Recent tests of these and other predictions (MacKay & Shafto, 2000) strongly support novelty detection as a fruitful addition to the specious-priming viewpoint.

3. Note that the proportions of errors and "can't say" responses for Armstrong questions were not previously reported for Experiment 2.

Acknowledgments—Support from National Institute for Aging Grant R01 AG 097705 to Donald G. MacKay and National Science Foundation Graduate Fellowship to Meredith Shafto is gratefully acknowledged. The authors thank Bryan Chae for research assistance, and Deborah Burke, Lori James, Kay Livesay, and Helena Wall for helpful comments on an earlier draft of this article.

REFERENCES

Baars, B. (1988). *A cognitive theory of consciousness*. Cambridge, England: Cambridge University Press.

Barton, S.B., & Sanford, A.J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, *21*, 477–487.

Bohannon, J.N., & Bonvillian, J.D. (1997). Theoretical approaches to language acquisition. In J.B. Gleason (Ed.), *The development of language* (4th ed., pp. 259–316). Boston: Allyn & Bacon.

Bredart, S., & Modolo, K. (1988). Moses strikes again: Focalization effect on a semantic illusion. *Acta Psychologica*, *67*, 135–144.

Buyer, L.S., & Radvansky, G.A. (1995, November). *Decomposing the "Moses Illusion" II: Evidence for the incomplete matching hypothesis.* Poster presented at the annual meeting of the Psychonomic Society, Los Angeles.

Cole, R.A., & Scott, B. (1974). Toward a theory of speech perception. *Psychological Review*, *81*, 348–374.

Cutting, J.C., & Ferreira, V.S. (1999). Semantic and phonological information flow in the production lexicon. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 318–344.

Dell, G. (1988). The retrieval of phonological forms in language production: Tests of predictions from a connectionist model. *Journal of Memory and Language*, *27*, 124–142.

Erickson, T.D., & Mattson, M.E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, *20*, 540–551.

Kamas, E.N., Reder, L.M., & Ayers, M.S. (1996). Partial matching in the Moses illusion: Response bias not sensitivity. *Memory & Cognition*, *24*, 687–699.

MacKay, D.G. (1973). Aspects of the theory of comprehension, memory and attention. *Quarterly Journal of Experimental Psychology*, *25*, 22–40.

MacKay, D.G. (1987). *The organization of perception and action: A theory for language and other cognitive skills.* New York: Springer-Verlag.

MacKay, D.G. (1990). Perception, action, and awareness: A three-body problem. In O. Neumann & W. Prinz (Eds.), *Relationships between perception and action: Current approaches* (pp. 269–303). Berlin: Springer-Verlag.

MacKay, D.G., & Burke, D.M. (1990). Cognition and aging: A theory of new learning and the use of old connections. In T.M. Hess (Ed.), *Aging and cognition: Knowledge organization and utilization* (pp. 213–263). Amsterdam: North Holland.

MacKay, D.G., & Shafto, M. (2000). *Comparing two theoretical frameworks for explaining the Armstrong illusion: Specious-priming versus novelty-detection*. Unpublished manuscript, University of California, Los Angeles.

MacKay, D.G., Stewart, R., & Burke, D.M. (1998). H.M.'s language production deficits: Implications for relations between memory, semantic binding, and the hippocampal system. *Journal of Memory and Language*, *38*, 28–69.

Marslen-Wilson, W.D. (1975). Sentence perception as an interactive parallel process. *Science*, *189*, 226–228.

Motley, M., & Baars, B. (1976). Semantic bias effects on the outcome of verbal slips. *Cognition*, *4*, 177–187.

Reder, L.M., & Cleeremans, A. (1990). The role of partial matches in comprehension: The Moses Illusion revisited. In A. Graesser & G. Bower (Eds.), *The psychology of learning and motivation* (Vol. 25, pp. 233–258). New York: Academic Press.

Reder, L.M., & Kusbit, G.W. (1991). Locus of the Moses Illusion: Imperfect encoding, retrieval, or match? *Journal of Memory and Language*, *30*, 385–406.

Shafto, M.A., & MacKay, D.M. (1998a, April). *Effects of aging on the Moses illusion and related conceptual errors*. Poster presented at the biennial meeting of the Cognitive Aging Conference, Atlanta, GA.

Shafto, M.A., & MacKay, D.M. (1998b, July). *The Moses and Armstrong illusions: Implications for language comprehension and the structure of semantic memory*. Poster presented at the annual meeting of the Cognitive Science Association for Interdisciplinary Learning, Hood River, OR.

van Oostendorp, H., & Kok, I. (1990). Failing to notice errors in sentences. *Language and Cognitive Processes*, *5*, 105–113.

(RECEIVED 1/20/99; REVISION ACCEPTED 2/7/00)