Journal of Experimental Psychology:
Human Perception and Performance
2003, Vol. 29, No. 2, 280–289

DOI: 10.1037/0096-1523.29.2.280

Vanishing Dual-Task Interference After Practice: Has the Bottleneck Been Eliminated or Is It Merely Latent?

Eric Ruthruff and James C. Johnston
National Aeronautics and Space Administration
Ames Research Center

Mark Van Selst
San Jose State University

Shelly Whitsell and Roger Remington
National Aeronautics and Space Administration Ames Research Center

Practice can, in some cases, largely eliminate measured dual-task interference. Does this absence of interference indicate the absence of a processing bottleneck (defined as an inability to carry out certain stages in parallel)? The authors show that a bottleneck need not produce any observable interference, provided that there is no temporal overlap in the demand for bottleneck stages on the 2 tasks. Such a "latent" bottleneck is especially likely after practice, when central stages are short. The authors provide new evidence that a latent bottleneck occurred for a participant who produced no interference in M. Van Selst, E. Ruthruff, and J. C. Johnston (1999). These findings demonstrate that the absence of dual-task interference does not necessarily indicate the absence of a processing bottleneck.

When people attempt to perform two speeded-response tasks at the same time, substantial slowing generally occurs. For novel tasks at modest practice levels, dual-task slowing tends to be severe (several hundreds of milliseconds; Pashler & Johnston, 1998). It is also very robust, having been found even for pairs of very simple tasks with no obvious input or output conflicts. There is strong evidence that dual-task slowing is primarily due to a bottleneck in central "thought-like" stages—while central stages on one task are underway, central stages on the other task must wait (Davis, 1957; Welford, 1952; for reviews see Lien & Proctor, 2002, and Pashler & Johnston, 1998). A generalized version of this central bottleneck model is shown in Figure 1A. Each task is decomposed into three processing stages: Stage A, a prebottleneck stage; Stage B, a bottleneck stage; and Stage C, a postbottleneck stage. By hypothesis, Stages A and C can proceed in parallel with any stage on another task. However, Stage B (the bottleneck stage) proceeds on only one task at a time. Therefore, the Task 2 bottleneck stage must wait for the bottleneck stage on Task 1 to finish, resulting in a dual-task cost.

Although dual-task costs are almost always found to be very large at low levels of practice, they are sometimes very small after extensive practice (Levy & Pashler, 2001; Ruthruff, Johnston, & Van Selst, 2001; Schumacher, Seymour, Glass, Kieras, & Meyer, 2001; Van Selst et al., 1999; see also Allport, Antonis, & Reynolds, 1972; Hirst, Spelke, Reaves, Caharack, & Neisser, 1980; Shaffer, 1975). Given small or nonexistent dual-task costs, it is tempting to conclude that participants have learned to carry out both tasks in parallel with no bottleneck, as shown in Figure 1B. There are at least three attractive explanations for why this might have happened. First, practice might reduce the resource demands of tasks. For instance, tasks might need a central supervisor that can work on only one task at a time (Norman & Shallice, 1986) at low practice levels but not at high practice levels. Second, participants might learn how to carry out one or both tasks using a completely new processing path that makes no demands on limited resources. For instance, with high practice one might develop "jumper cable" paths directly between stimulus brain areas and response brain areas (Johnston & Delgado, 1993). Third, it has been argued that at low practice levels, people choose to schedule tasks conservatively in a single-channel fashion, whereas with practice they learn a riskier parallel scheduling strategy (Schumacher et al., 2001).

Any of these theoretical possibilities, if true, would be very important for understanding how cognitive architecture changes with practice. We argue, however, that it would be premature to conclude that any of them are true, because the key assumption they have in common—that the limitation underlying the processing bottleneck has been eliminated—has not been established. The problem is that a very simple logical alternative, which we call the *latent bottleneck hypothesis* (Van Selst et al., 1999), has not yet been ruled out.

The Latent Bottleneck Hypothesis

The key idea of the latent bottleneck hypothesis is that an inability to perform certain mental operations in parallel need not

Eric Ruthruff, James C. Johnston, Shelly Whitsell, and Roger Remington, National Aeronautics and Space Administration (NASA) Ames Research Center, Moffett Field, California; Mark Van Selst, Department of Psychology, San Jose State University.

This research was supported by grants from the National Research Council and the Airspace Operations Systems Project of NASA's Airspace Systems Program. We thank Richard Ivry, Eliot Hazeltine, Mei-Ching Lien, Hal Pashler, Richard Schweickert, Werner Sommer, and an anonymous reviewer for helpful comments on a draft of the manuscript.

Correspondence concerning this article should be addressed to Eric Ruthruff, Mailstop 262-4, NASA Ames Research Center, Moffett Field, California 94035. E-mail: eruthruff@mail.arc.nasa.gov

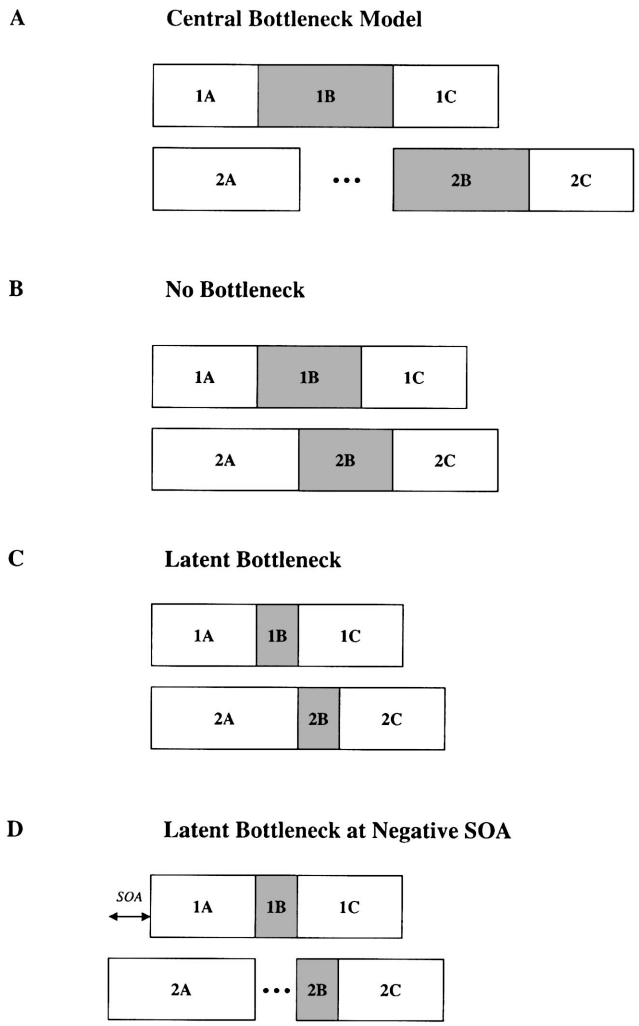

Figure 1. Stage-time diagrams where processing is decomposed into three stages, arbitrarily labeled *A, B,* and *C.* (A) Generalized bottleneck model. Stages A and C on one task can proceed in parallel with any stage of the other task. However, Stage B (the bottleneck stage) proceeds on only one task at a time. (B) Model with no bottleneck. (C) Latent bottleneck hypothesis. Stage B can proceed on only one task at a time. However, Stage 1B generally finishes before Stage 2B is set to begin. (D) Latent bottleneck revealed by using a sufficiently negative stimulus onset asynchrony (SOA).

necessarily result in substantial interference. The state of affairs envisioned by this hypothesis is shown in Figure 1C. Here, the bottleneck stages (1B and 2B) do not come into conflict, because Stage 1B is completed before Stage 2B is ready to begin (i.e., before Stage 2A has finished).

The latent bottleneck hypothesis is not only logically possible, but also is highly plausible under some circumstances. There are two key conditions that promote a latent bottleneck. The first condition is a substantial difference in the finishing times of the two prebottleneck stages (1A and 2A), causing a substantial asynchrony in the demand for the two bottleneck stages (1B and 2B). The second condition is that the bottleneck stages be short, which reduces the likelihood that Stage 1B will still be operating when Stage 2B is ready to begin. Previous single-task studies have found that central stages such as response selection are the ones shortened the most by practice (Fletcher & Rabbitt, 1978; Mowbray & Rhoades, 1959; Pashler & Baylis, 1991; Van Selst et al., 1999; Welford, 1976). It is therefore plausible that this second condition would be met after extensive practice.

Given the plausibility of the latent bottleneck hypothesis, it needs to be ruled out before concluding that the processing limitation underlying the bottleneck has been eliminated. Unfortunately, the latent bottleneck hypothesis has generally not been considered (but cf. Hazeltine, Teague, & Ivry, 2002). This oversight might arise partly from a subtle ambiguity in the term *bottleneck.* This term has been used to refer to the cognitive limitation, preventing certain stages from operating in parallel (e.g., a single-channel mechanism), and also to the processing delay that occurs as a consequence of the cognitive limitation (what is usually meant by a "bottleneck delay"). The absence of dual-task costs directly indicates the absence of a bottleneck delay, but it does not necessarily indicate the absence of the cognitive limitation.

We argue that researchers should focus not just on the practical issue of whether the observed bottleneck delay is large or small, but on the deeper theoretical issue of whether the underlying processing limitations have been eliminated. In this article, we take this suggestion to heart and look for evidence that a latent bottleneck was present in Van Selst et al. (1999), where one of six highly practiced participants (initials SW) showed no dual-task slowing.

Van Selst et al. (1999)

Noting that many previous failures to find reductions in dual-task costs with practice might be due to input or output conflicts, Van Selst et al. (1999) used a vocal response to an auditory stimulus on Task 1 and a manual response to a visual stimulus on Task 2. Participants performed 36 sessions (400 trials per session) with these two tasks in a Psychological Refractory Period (PRP) design with a variable stimulus onset asynchrony (SOA).

Figure 2 shows Task 2 response time (RT2) after practice

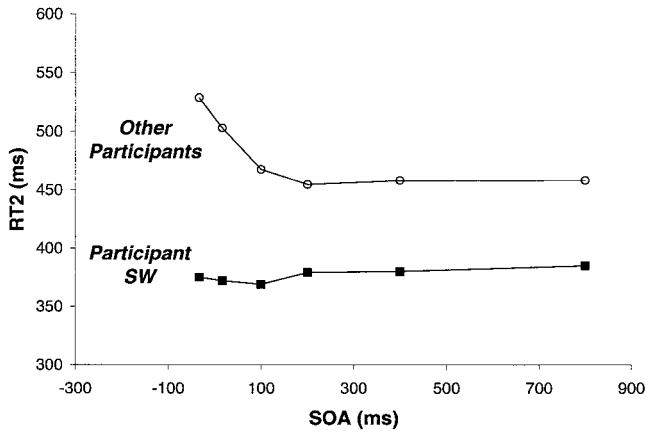

Figure 2. Data from Van Selst, Ruthruff, and Johnston (1999): Task 2 response time (RT2) as a function of stimulus onset asynchrony (SOA). Data are shown separately for SW (filled squares), the only participant who did not show a Psychological Refractory Period effect, and the average of the remaining 5 participants (open circles).

(Sessions 27–36) as a function of SOA. The data are shown separately for participant SW (solid squares) and the other 5 participants (open circles). Although practice dramatically reduced the amount of slowing at short SOAs (i.e., the PRP effect), significant slowing was still apparent for 5 of the 6 participants. This residual PRP effect and the pattern of factor interactions with SOA suggested that a processing bottleneck was still present for these participants (for supporting evidence see Ruthruff et al., 2001, who reported several transfer experiments on these participants). In contrast, participant SW showed no PRP effect. If anything, SW's RT2 was actually slightly slower at the long SOAs, which might reflect reduced alertness relative to the short SOAs.

Follow-Up Experiments

Did SW eliminate the processing bottleneck (Figure 1B) or did SW merely have a latent bottleneck (Figure 1C)? Given that SW performed RT1 relatively quickly after practice (240 ms in Sessions 27–36, compared with about 310 ms for the other 5 participants), a latent bottleneck is certainly plausible. Although Van Selst et al. (1999) considered the latent bottleneck hypothesis, it could not be tested on the basis of the available data. In the present study, we brought SW back into the lab (after a period of about 14 months) for a series of follow-up experiments using variations of the Van Selst et al. design. The key manipulation was to change the temporal alignment of the two tasks, so that the putative bottleneck stages would come into conflict. Given that RT1 was much faster than RT2 for SW, what was needed was to present Task 2 before Task 1 (i.e., at a negative SOA).[1] According to the latent bottleneck hypothesis, this manipulation could lead to the reemergence of a PRP effect, as shown in Figure 1D.[2] However, if the bottleneck has been eliminated (e.g., Figure 1B), then there is no clear reason to expect a substantial PRP effect at the negative SOAs.

In PRP research, one usually compares performance at short SOAs, where the tasks overlap in time, with performance at the baseline long SOAs. In the present context, however, the short versus long distinction is not the appropriate one. Instead, we focus on the comparison between the most negative SOAs (where Stages 1B and 2B might come into conflict) and positive SOAs (where Stages 1B and 2B generally should not conflict).

Procedure

Sessions consisted of 8 blocks of 50 trials (including 2 warm-up trials); the first block was considered practice. The Task 1 stimulus was a 150-ms tone; SW said "high" to 625- and 3125-Hz tones and said "low" to 80- and 400-Hz tones. For Task 2, SW responded by using the *H*, *J*, *K*, and *L* keys on the computer. The numbers 1–4 were mapped compatibly onto these four keys (i.e., numerically increasing from left to right), whereas letters were mapped incompatibly onto the same four keys in the order *D*, *C*, *A*, and *B*. The character stimuli (white on a black background) were either bright or dim. For more details of the methodology, see Van Selst et al. (1999).

SW was instructed to "Perform both tasks quickly and accurately, but try especially hard to make your tone-task responses as fast as possible." The instructions also emphasized that SW could respond to the tasks in any order (this instruction was especially important, in case SW could perform the tasks in parallel, because in later phases the Task 2 stimulus often came first). SW was not informed of the specific purpose of these follow-up experiments until after they had been concluded.

Analyses

Trials were eliminated from the analyses if either RT fell outside the cutoff values or if either response was an error. Trials following an error were also eliminated. Of primary interest was the effect of SOA. However, we also included Task 2 stimulus–response compatibility (compatible vs. incompatible) and Task 2 luminance (dim vs. bright) in the analyses of variance (ANOVAs). Session was used as a random factor, and the alpha level was .05.

Phase 1

The initial set of SOAs was the same as that used after practice in Van Selst et al. (1999): −33, 17, 100, 200, 400, and 800 ms, mixed within blocks. Phase 1 consisted of two refresher sessions (not analyzed), followed by three test sessions. RT1 fell outside the cutoff values (100–350 ms) on 1.1% of all trials, and RT2 fell outside the cutoff values (200–600 ms) on 1.8% of all trials. The data are represented by open squares in Figure 3. Table 1 shows the mean RT, standard deviation, and error rate for each task, along with the correlation between RT1 and RT2 as a function of SOA. The most important aspect of the Phase 1 data is that they show little or no PRP effect; in fact, RT2 is actually faster at the −33-ms SOA than at the long, positive SOAs. Thus, these data replicate the findings of Van Selst et al. (1999) for SW.

One ancillary but intriguing aspect of the data, also observed in Phases 2 and 3, was the large percentage of Task 2 errors at the 100-ms SOA (12.9%; see Table 1). Because mean RT2 was slightly faster at these SOAs than at the other SOAs, the effect could be due partly to a speed–accuracy trade-off. Alternatively, it might reflect some type of cross talk between certain stages of Task 1 and Task 2 that occurs only when those stages operate simultaneously.

Phase 2

As noted earlier, one way to reveal a latent bottleneck is to change the temporal alignment of the tasks so that the bottleneck stages come into conflict. In our case, it was necessary to present the Task 2 stimulus earlier in time, before the onset of the Task 1 stimulus (i.e., at a negative SOA; see Figure 1D). We found it convenient to simply subtract 83 ms from each of the six SOAs, so that the SOA range (833 ms) would stay the same and a 17-ms SOA would still be present. The 17-ms SOA provides a useful baseline because it did not show any interference (see Phase 1 and Van Selst et al., 1999), yet it is not so long that RT2 is artificially slowed because of reduced alertness (as appears to have happened

[1] Here, we define Task 1 as the task that usually comes first and Task 2 as the task that usually comes second (even if presented in the opposite order on a given trial).

[2] The nature of the PRP effect depends on whether the Task 1 or Task 2 bottleneck stage is performed first. As discussed later, however, the data suggest that the Task 1 bottleneck stage was almost always performed first.

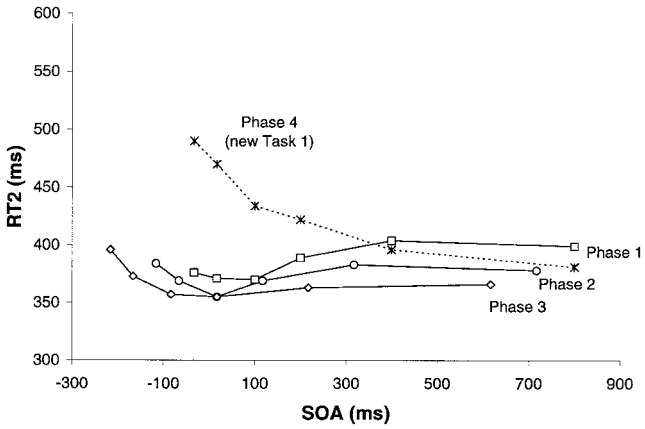

Figure 3. Task 2 response time (RT2) as a function of the stimulus onset asynchrony (SOA) for Phases 1–4. Positive SOAs mean that Task 1 came first; negative SOAs mean that Task 2 came first. Phase 1: open squares. Phase 2: open circles. Phase 3: open diamonds. Phase 4: asterisks.

at the 400- and 800-ms SOAs of Phase 1). Thus, the SOAs used in the six sessions of Phase 2 were −116, −66, 17, 117, 317, and 717 ms.

RT1 fell outside the cutoff values (100–350 ms) on 0.8% of all trials, and RT2 fell outside the cutoff values (200–550 ms) on 3.6% of all trials. The Phase 2 data are represented by open circles in Figure 3 (see also Table 1). Mean RT2 at the most negative SOA (−116 ms) was significantly slower (M = 384 ms) than it was at the three positive SOAs (M = 377 ms), $F(1, 5) = 14.5$, $MSE = 45.6$, $p < .05$. This small PRP effect (7 ms when measured using the three long, positive SOAs as the baseline) might indicate that a bottleneck was rarely encountered. On the other hand, as noted above, mean RT2 at the long positive SOAs might have been artificially slow because of reduced preparation or alertness. Using only the 17-ms SOA as a baseline, we found that the −116-ms SOA condition shows a mean PRP effect of 29 ms.

With regard to the size of the PRP effect, it is also important to note that mean RT1 was actually smaller at the most negative SOA (M = 209 ms) than it was at the three positive SOAs (M = 221 ms), $F(1, 5) = 31.3$, $MSE = 54.4$, $p < .01$. A similar pattern was observed in Phase 3 as well. If some of this RT1 speed-up were due to a reduction in the duration of Stage 1A and/or Stage 1B, then the reduction would carry over onto RT2, reducing the size of the observed PRP effect. As for the cause of the RT1 speed-up, note that it was accompanied by an increase in Task 1 error rates (10% vs. 3.7%; see Table 1); therefore, it might be due to a speed–accuracy trade-off. The bottleneck model, according to which Stage 2B often must wait until Stage 1B has finished, suggests a more specific explanation. Perhaps while Task 2 is waiting for some occupied resource, pressure is placed on Task 1

Table 1
Mean Response Time (RT; in Milliseconds), Standard Deviation, and Error Rate (ER) on Task 1 and Task 2, Along With the RT Correlation Between Task 1 and Task 2, as a Function of Phase and SOA

Phase and SOA	RT1	*SD*1	RT2	*SD*2	*r*	ER1	ER2
1 (3 sessions)							
−33	227	38.9	376	46.2	.10	8.1	0.7
17	225	32.0	371	62.1	−.02	3.9	5.9
100	232	34.4	370	65.5	.09	6.1	12.9
200	231	37.6	389	65.0	.02	4.9	7.6
400	234	37.9	404	61.8	−.06	2.9	4.4
800	226	34.7	399	61.6	.04	6.8	4.8
2 (6 sessions)							
−116	209	36.1	384	35.3	.19	10.0	1.0
−66	210	31.8	369	46.3	.10	8.0	3.8
17	221	29.3	355	56.1	.04	3.4	9.5
117	221	33.4	369	60.8	−.11	2.8	10.1
317	222	31.0	383	54.6	.03	4.6	4.6
717	220	32.4	378	52.2	.02	3.8	4.8
3 (9 sessions)							
−216	200	29.5	396	32.8	.40	7.0	0.2
−166	210	26.1	373	31.7	.30	7.5	1.4
−83	216	23.5	357	49.3	.11	4.4	3.5
17	219	23.0	355	64.7	−.03	2.0	12.1
217	223	23.6	363	59.2	.05	2.0	5.4
617	223	24.6	366	49.9	.04	2.7	3.6
4 (3 sessions)							
−33	400	101.1	490	111.4	.75	11.5	0.7
17	399	97.6	470	124.7	.67	12.4	0.0
100	401	100.9	434	93.2	.68	15.7	0.7
200	401	92.4	422	87.9	.45	11.5	0.0
400	406	106.7	396	69.4	.31	12.0	2.3
800	398	98.6	381	64.1	.05	11.5	0.7

Note. SOA = stimulus onset asynchrony.

to free up that resource, causing a faster and more error-prone Task 1 response.

Phase 3

In Phase 3 (nine sessions) we subtracted another 100 ms from each of the SOAs (note that this reduction again provides a 17-ms SOA, for continuity). Thus, the SOAs were −216, −166, −83, 17, 217, and 617 ms.

RT1 fell outside the cutoff values (100–350 ms) on 0.6% of all trials, and RT2 fell outside the cutoff values (200–550 ms) on 2.5% of all trials. The Phase 3 data are represented by open diamonds in Figure 3 (see also Table 1). As predicted by the latent bottleneck hypothesis, the use of more negative SOAs led to the further reemergence of a PRP effect: Mean RT2 was 35 ms longer at the −216-ms SOA (M = 396 ms) than it was at the three positive SOAs (M = 361 ms), $F(1, 8) = 86.5$, $MSE = 247$, $p < .001$. Using only the 17-ms SOA as a baseline, we found that the −216-ms SOA condition shows a PRP effect of 41 ms.

Processing order. Before evaluating a bottleneck model of these data, it is necessary to first assess in what order the bottleneck stages of Task 1 and Task 2 were performed. The usual assumption in classic PRP studies is that Stage 1B will always be processed before Stage 2B. However, those PRP studies rarely use negative SOAs, where the Task 2 stimulus precedes the Task 1 stimulus. At the −216-ms SOA of the present study, it is plausible (e.g., if access to bottleneck resources were "first come, first served") that Stage 2B would be performed before Stage 1B some of the time. If so, the bottleneck would have delayed Task 1 responses and not delayed Task 2 responses. Other things being equal, use of this processing order should cause an increase in mean RT1 at the −216-ms SOA relative to the baseline SOAs (i.e., the long, positive SOAs; viz., there should be a PRP effect on RT1). Contrary to this expectation, RT1 was actually faster at the −216-ms SOA (M = 200 ms) than at the three longest positive SOAs (M = 222 ms; see Table 1). Given this finding, it seems unlikely that Stage 2B was generally performed before Stage 1B at the negative SOAs.

It is nevertheless conceivable that this processing order (Stage 2B before Stage 1B) was used on a small proportion of trials, producing a modest increase in mean RT1 that was offset by some other factor that decreased RT1. This possibility, however, is not supported by the observed scatterplot of RT1 and RT2, which is shown in Figure 4 for the −216-ms SOA (Figure 4A) and for the longest positive SOA (Figure 4B). Note that the most likely scenario in which Stage 2B would be performed before Stage 1B is when the durations of Stages 1A and 1B are relatively long (which would lead to a long RT1) and Stage 2A is relatively short (which would lead to a short RT1). The ensuing bottleneck delay on Task 1 would only further lengthen RT1, without affecting RT2. Therefore, this processing order should produce a cluster of points centered in the lower right portion of the scatterplot for the −216-ms SOA, corresponding to a relatively slow RT1 combined with a relatively fast RT2. There are essentially no such points. On the contrary, at the −216-ms SOA there is a remarkable absence of fast RT2s; RT2 fell below 325 ms on 27% of trials at the baseline (positive) SOAs, but almost never (<1%) at the −216-ms SOA. Likewise, there was no apparent increase in the frequency of trials with a slow RT1 at the −216-ms SOA. This analysis therefore suggests that when a bottleneck occurred, Stage 2B was almost never processed before Stage 1B.

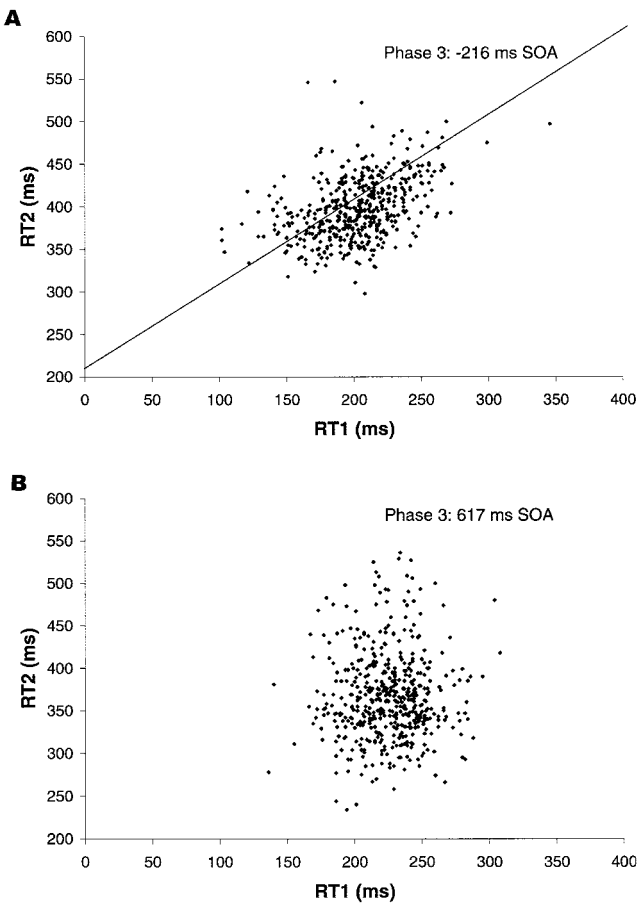

Figure 4. Task 2 response time (RT2) versus Task 1 response time (RT1) at the most negative stimulus onset asynchrony (SOA; A: −216 ms) and at the most positive SOA (B: 617 ms) in Phase 3. The diagonal line in A corresponds to an interresponse interval of 0 ms (i.e., simultaneous responses).

Although Stage 2B was apparently not processed before 1B, the Task 2 response was often emitted before the Task 1 response at the −216-ms SOA. These "response reversals" can be seen in Figure 4. The diagonal line in Figure 4A corresponds to simultaneous responses (the line does not appear in Figure 4B because the SOA was so long); points below and to the right of this line indicate that the Task 2 response was emitted before the Task 1 response. The bottleneck model shown in Figure 1A can account for these trials provided that the Task 2 stages after the bottleneck (Stages 2B and 2C) typically take less time than the Task 1 stage after the bottleneck (Stage 1C). It is highly plausible that Stage 2C (depressing a finger) would take less time than Stage 1C, which involves a movement of the tongue to begin speaking one of two words ("high" or "low"). However, in order for Stages 2B and 2C to still take less time than Stage 1C, it would appear to be necessary for Stage 2B to be very short in duration.

Why would SW perform Stage 1B before Stage 2B, rather than using a first-come, first-served strategy? One possible explanation is that SW was trying to prevent overt response reversals. This

hypothesis seems unlikely because our instructions emphasized that response reversals were permitted. Furthermore, SW actually did reverse the responses on more than half of all trials at the −216-ms SOA (see Figure 4A). It also seems unlikely that SW was attempting to synchronize the two responses by temporarily withholding whichever response was ready first. Although this response grouping strategy correctly predicts RT2 slowing, it also predicts RT1 slowing—in fact, RT1 slowing is the classic signature of response grouping. Yet RT1 slowing was not found even at SOAs where it would have been expected to be substantial. For instance, at the −83-ms SOA, Task 1 should have finished about 60 ms ahead of Task 2, on average (assuming from the positive SOA data that Task 1 and Task 2 can be completed in about 220 and 360 ms, respectively). Therefore, more than half of the Task 1 responses should have been slowed by grouping. Yet, no RT1 slowing at all was found at the −83-ms SOA (or at any other SOA) compared to the most positive SOAs.

A more attractive explanation for the fixed order of processing—Stage 1B before Stage 2B—is that experienced participants in multitask environments prepare in advance to perform bottleneck stages in a particular order (see De Jong, 1995). Throughout much of SW's extensive previous training, Task 1 almost always came first, and even in the present study Task 1 came first more often than Task 2. Given this history, it should perhaps be expected that SW would prepare to perform bottleneck operations on Task 1 before those on Task 2. Even if Stage 2A occasionally finished before Stage 1A, SW may well have been unprepared to carry out the bottleneck operations in the unexpected order. Rather than taking the time to establish a new preparatory state on these trials, it appears that SW stayed with the "Task 1 first" long-term strategy, at the expense of some modest delays in RT2.

Evidence for a bottleneck. Having shown that a bottleneck delay, if present, should primarily affect Task 2, we can now assess the evidence for such a bottleneck. Two converging lines of evidence support a bottleneck explanation for the observed PRP effect. First, bottleneck models correctly predict that RT1 and RT2 should be positively correlated at the most negative SOAs (see Pashler & Johnston, 1989, 1998). The rationale behind this prediction is that both RT1 and RT2 depend on the duration of Stages 1A and 1B. Assuming a bottleneck is encountered, RT1 = 1A + 1B + 1C and RT2 = 1A + 1B + 2B + 2C − SOA (Pashler & Johnston, 1989). Thus, if Stages 1A and 1B take an especially long time on a particular trial, both RT1 and RT2 are likely to be long. But if Stages 1A and 1B take only a short time, then both RT1 and RT2 are likely to be short. Other things being equal, variation in Stages 1A and 1B should produce a positive correlation between RT1 and RT2. As predicted by the latent bottleneck hypothesis, the correlation between RT1 and RT2 was positive and statistically significant at the most negative SOA ($r = .40$, $p < .001$; see Figure 4A and Table 1). In contrast, at positive SOAs (where the bottleneck should rarely occur) the correlations were close to zero (see Figure 4B and Table 1) and nonsignificant ($p > .2$ in each case). A similar pattern was seen in Phase 2. Note that if the PRP effect was simply caused by reduced alertness or preparation for Task 2 at the −216-ms SOA, there would be no obvious grounds for expecting a substantial positive correlation.

The second line of evidence in favor of a bottleneck comes from an analysis of RT2 variance across SOAs. Variance in RT generally increases as the mean RT increases (cf. Luce, 1986), therefore one might expect RT2 variance to be largest at the most negative SOAs (where mean RT2 was the longest). The latent bottleneck hypothesis, however, makes the counterintuitive prediction that, under the unusual circumstances of this experiment, RT2 variance should actually decrease at the most negative SOAs. Consider the stages that contribute to RT2. At the most negative SOAs (assuming there is always a bottleneck delay), RT2 = 1A + 1B − SOA + 2B + 2C; at the most positive SOAs (assuming there never is a bottleneck delay), RT2 = 2A + 2B + 2C (see Pashler & Johnston, 1989). Note that for the bottleneck to be latent at a zero or positive SOA, Stages 1A + 1B must take less time than Stage 2A. Combining these facts together, it is clear that the set of processing stages involved at negative SOAs (1A, 1B, 2B, 2C)—ignoring the SOA term, which has no variance within a condition—should actually take less time than those involved at the most positive SOAs (2A, 2B, 2C). Given that the processing stages at negative SOAs take less time, it is reasonable to expect that they would also have less variance (Luce, 1986). Below, we also demonstrate this prediction using a quantitative simulation of the bottleneck model.

In Phase 3, RT2 variance was indeed much smaller at the most negative SOA (1,074 ms^2; $SD = 32.8$ ms) than at the three positive SOAs (3,394 ms^2; $SD = 58.3$ ms), $F(741, 256) = 3.16$, $p < .01$. A similar pattern was observed in Phase 2 (see Table 1). This dramatic and counterintuitive finding—that RT2 variance decreases even as the mean RT2 increases—is consistent with the latent bottleneck hypothesis, but it is puzzling from the viewpoint of many alternative models. For instance, if Task 2 were performed slowly at negative SOAs because of reduced alertness, then the obvious expectation would be an increase in variance, not a decrease.

One piece of evidence that might appear to contradict bottleneck models is that the slope of the PRP function across the first two SOAs was much flatter than the −1 value sometimes predicted by bottleneck models. However, the slope need not be −1 if a bottleneck does not occur on every trial (see Schwarz & Ischebeck, 2001). Indeed, we show below that a quantitative bottleneck model can in fact simulate the observed slope of the PRP function (as well as other important aspects of the data).

Phase 4

According to the latent bottleneck hypothesis, SW showed no PRP effect at positive SOAs (or small negative SOAs, as in Phase 1) because Stage 1B finished very early, typically before Stage 2B was ready to begin. If so, then another way of revealing a PRP effect would be to substitute a new Task 1, causing Stage 1B to finish much later in time. The effect of inserting a new Task 1 should be similar to that of using negative SOAs (as in Phases 2 and 3): Both manipulations delay completion of Stage 1B, bringing the bottleneck stages into overt conflict. One advantage of inserting a new Task 1 is that it permits use of the original SOA range from Phase 1 (−33, 17, 100, 200, 400, and 800 ms), which is closer to the range commonly used in PRP experiments. The latent bottleneck hypothesis predicts a substantial PRP effect in Phase 4 even at this traditional SOA range. The alternative hypothesis that Task 2 had become automated—so that it no longer required any bottleneck operations—provides no reason to expect a substantial PRP effect.

A similar experimental logic was used by Ruthruff et al. (2001), who introduced a new Task 1 judgment to the 5 other highly practiced participants (all except for SW) from the Van Selst et al. (1999) study. Ruthruff et al. found that the new Task 1 caused a sharp increase in the PRP effect (194 ms in Session 1) and therefore concluded that Task 2 had not been automated to the point that it could bypass the processing bottleneck. The present experiment was intended to determine whether the same conclusion applies to participant SW as well.

SW performed three sessions with a new Task 1 involving a rapid sequence of three 150-ms tones separated by 50 ms. SW was to indicate whether the third tone was higher or lower in pitch than the first (i.e., ignoring the middle tone) by saying "gate" or "pike," respectively. RT1 fell outside the cutoff values (200–850 ms) on 1.4% of all trials, and RT2 fell outside the cutoff values (200–950 ms) on 1.0% of all trials. Figure 3 shows mean RT2 as a function of SOA with the new Task 1 (asterisk symbols); see also Table 1. As predicted by the latent bottleneck hypothesis, introduction of a new Task 1 resulted in a substantial PRP effect of 167 ms in Session 1. Averaged over three sessions of practice, the PRP effect was 109 ms, $F(5, 10) = 8.3$, $MSE = 2{,}102$, $p < .01$. In contrast, we found essentially no PRP effect in Phase 1 using the same SOA range with the old Task 1. The size of the initial PRP effect in Phase 4 (167 ms) is roughly what one might expect given a bottleneck model and the observed mean RT1 of 439 ms (see Van Selst et al., 1999, for a discussion of the quantitative relation between the PRP effect and RT1). Thus, the findings of Experiment 1 of Ruthruff et al. (2001) appear to also extend to participant SW.

The evidence from the Phase 4 transfer test refutes the hypothesis that Task 2 had been automated to the point that it could bypass the processing bottleneck. This evidence does not refute the logical possibility that Task 1 and Task 2 were automatized, not individually, but as a unit. However, this hypothesis seems unlikely given that very different SOAs were mixed within a block (see also Ruthruff et al., 2001).

Bottleneck Model Simulations

We discussed above how the latent bottleneck model can explain the key qualitative effects found at the most negative SOA, relative to the positive SOAs, including (a) the increase in RT2, (b) the stronger correlation between RT1 and RT2, and (c) the decrease in RT2 variance. However, one might question whether the bottleneck model is capable of accounting for the approximate size of these effects and whether it can account for all three effects with a single set of parameter values. To answer these questions, we simulated a bottleneck model.

We modeled RT1 as the sum of two independent, normally distributed stages: one prior to the release of the bottleneck (a combination of Stage 1A plus Stage 1B) and one after the release of the bottleneck (1C). We modeled RT2 as the sum of two independent normally distributed stages: one prior to the release of the bottleneck (2A) and one after the release of the bottleneck (a combination of Stage 2B plus Stage 2C). The Task 2 stage durations were assumed to be independent of the Task 1 stage durations except that, due to the hypothesized processing bottleneck, the late Task 2 stages (2B + 2C) could not begin until the early Task 1 stages (1A + 1B) had finished. Consistent with our findings above, we assumed that Stage 1B was always completed before Stage 2B.

The resulting model has eight parameters (1 *M* and 1 *SD* for each of the four stages). The means and standard deviations for the two stages of each task were stipulated to produce the means and standard deviations observed in our baseline condition (i.e., at the three positive SOAs, where it is assumed that the bottleneck did not occur).[3] Given these stipulations, only four free parameters remained. We set out to determine whether this model can, with some reasonable set of parameters, provide a close fit to the observed data from the −216-ms SOA (where a bottleneck is hypothesized to have occurred). Essentially, we question whether it is possible to decompose the observed baseline RTs on each task into two independent stages, such that if Task 2 begins 216 ms before Task 1 and there is a bottleneck, the three qualitative findings (a, b, and c above) are produced.

Basic bottleneck model. One complication concerns how to account for the speed-up in RT1 observed at the negative SOAs. In the first set of simulations, rather than adding additional parameters to the basic bottleneck model, we simply stipulated that the Task 1 parameters match the mean and the standard deviation of RT1 observed at the −216-ms SOA. The data left to be fit, therefore, are the mean and standard deviation of RT2 and the RT1–RT2 correlation at the −216-ms SOA.

We found a wide range of parameter values that produced satisfactory fits. We believe that what is critical is not finding the absolute closest fit, because the data include measurement error, but rather finding very good fits with plausible parameter values. The following is one set of parameters that fits the data well: For Stages 1A + 1B, $\mu = 83$ ms, $\sigma = 25.3$ ms; Stage 1C, $\mu = 117$ ms, $\sigma = 15.2$ ms; Stage 2A, $\mu = 279$ ms, $\sigma = 57.8$ ms; Stages 2B + 2C, $\mu = 81$ ms, $\sigma = 13.8$ ms.

The fit of this basic bottleneck model to the data (based on 1 million simulated trials) is shown in Table 2. As can be seen, the fitted values are quite close to the observed values. Thus, the simulated bottleneck model with a single set of parameter values successfully captured the approximate magnitude of the three key qualitative findings at the −216-ms SOA: (a) the modest increase in mean RT2, (b) the moderate correlation between RT1 and RT2, and (c) the strong decrease in RT2 variance.

Bottleneck model with RT1 speed-up parameter. The first bottleneck model simulation did not account for the observed speed-up in RT1 at negative SOAs. To account for these effects, we conducted a second simulation with one additional parameter (total of five free parameters). We assumed that when Stage 2A finishes before Stages 1A + 1B, the remaining time to complete Stage 1A + 1B is reduced by a fixed percentage (as might happen if the speed–accuracy criterion shifted). The following is one representative set of parameters that fits the data well: speed-up percentage = 37%; Stages 1A + 1B, $\mu = 103$ ms, $\sigma = 7.4$ ms;

[3] To compare the simulated data to the real data, it was necessary to apply the same RT cutoffs in both cases. Also note that the simulations were designed to predict within-cell variance, after removing variance due to the factors used in our ANOVAs (Task 2 S-R compatibility, Task 2 stimulus intensity, and SOA). In addition, if a normally distributed stage produced a negative value, we set that value to zero.

Table 2
Actual Data From Phase 3 and Values Predicted by Two Different Models: The Basic Bottleneck Model and a Bottleneck Model With an Extra Parameter to Account for the Speed-Up in RT1 at Negative SOAs

	−216-ms SOA					Positive SOAs				
Phase 3	RT1	*SD*1	RT2	*SD*2	*r*	RT1	*SD*1	RT2	*SD*2	*r*
Actual data	200	29.5	396	32.8	.40	222	23.8	361	58.3	.02
Basic model	200	29.5	*396*	*35.0*	*.39*	222	23.8	361	58.3	.00
Model with RT1 speed-up	*200*	*28.8*	*396*	*36.1*	*.40*	222	23.8	361	58.3	.00

Note. Data are shown separately for the −216-ms stimulus onset asynchrony (SOA) and for the average of the three positive SOAs (baseline). Shown are the mean response time (RT; in milliseconds) and standard deviation for Task 1 and Task 2, along with the RT correlation between Task 1 and Task 2. The model parameters were constrained to fit the nonitalicized values exactly. Of interest is whether the values in italics fit the actual data.

Stage 1C, $\mu = 119$ ms, $\sigma = 22.6$ ms; Stage 2A, $\mu = 266$ ms, $\sigma = 54.1$ ms; Stages 2B + 2C, $\mu = 94$ ms, $\sigma = 24.4$ ms.

The fit of the revised model to the data (based on 1 million simulated trials) is shown in Table 2. As can be seen, the fitted values again are close to the observed values. This revised model successfully captured the key qualitative findings mentioned above. In addition, it also accounted for the speed-up in mean RT1 at the most negative SOA. Although we did not set out to account for the concomitant increase in RT1 variance, the model did a reasonable job of that as well. Note that this model, and the more basic model presented above, would produce virtually no PRP effect at zero or positive SOAs, and in fact no PRP effect was found at those SOAs in Phases 1–3.

It is, of course, conceivable that some other model might provide an even better fit to these data. Nevertheless these simulations show clearly that SW's data, which might appear at first glance to reject the entire class of bottleneck models (especially in Phase 1, where there was no PRP effect), can in fact be well accounted for by quantitative models incorporating a processing bottleneck. Thus, these simulations support the conclusion that SW's bottleneck had not been eliminated by extensive practice but merely had become latent at particular SOAs.

Did Our SOA Manipulation Reveal a Bottleneck or Did It Cause One?

We have concluded that our experimental manipulations revealed the presence of a bottleneck that previously was only latent. It is worth considering, however, the logical possibility that our manipulations actually reintroduced a bottleneck that previously had been eliminated by extensive practice. That is, perhaps the introduction of substantial negative SOAs produced a bottleneck only because it disrupted performance. For instance, SW might have learned a strategy for overlapping bottleneck operations that was specific to a particular temporal arrangement (i.e., the near zero SOAs used in Van Selst et al., 1999). When negative SOAs were introduced, perhaps this strategy became ineffective and could no longer be used.

One argument against the disruption hypothesis is that the use of negative SOAs in Phase 3 did not disrupt performance over the range of SOAs (e.g., 17–617 ms) that were used in Phase 1. Indeed, performance at those SOAs actually improved in Phase 3 relative to Phase 1 (see Figure 3). In addition, note that SW had thousands of trials of practice in which to adapt to the new range of SOAs (six sessions in Phase 2 and nine sessions in Phase 3). Yet, even after all this practice, SW's PRP effect was still 46 ms (measured at the −216-ms SOA relative to the 17-ms SOA) in her final session. For comparison, note that in the Van Selst et al. (1999) study, it took SW only about 10 sessions to entirely eliminate the PRP effect, despite being initially unfamiliar with the tasks.

Locus of the Bottleneck

Given a bottleneck model like the one in Figure 1A, the standard approach to determining the bottleneck locus is to manipulate the durations of various Task 2 stages and see how the effects on RT2 interact with SOA. Underadditive interactions indicate that the manipulated stage comes before the bottleneck, whereas additive interactions indicate that the stage comes at or after the bottleneck (see Schweickert, 1978, and Pashler & Johnston, 1989, 1998, for more details regarding "locus-of-slack" logic). The present experiments included a stimulus–response (S-R) compatibility factor, intended to selectively influence the duration of response selection, which showed a significant underadditive interaction with SOA in Phase 3: The compatibility effect on RT2 was 19 ms averaged across the three positive SOAs, but it shrank to 8 ms at the most negative SOA, $F(1, 8) = 19.2$, $MSE = 54.9$, $p < .01$.[4] This result suggests that at least part of the Task 2 response selection stage took place prior to the bottleneck, which might be located at a later substage of response selection or even later at response initiation (see De Jong, 1993; Keele, 1973). Ruthruff et al. (2001), who used the same S-R compatibility manipulation, reached the same conclusion for the other 5 highly practiced participants of the Van Selst et al. (1999) study.

Converging evidence for a relatively late bottleneck comes from the fact that at the −216-ms SOA of Phase 3, SW responded to

[4] The experiment also included a Task 2 contrast manipulation, whose mean effect size was 17 ms in Phase 3. This effect was slightly larger at the positive SOAs (19 ms) than at the most negative SOAs (14 ms). Although not significant ($p > .2$), the direction of the effect is consistent with the hypothesis that contrast primarily affected a stage prior to the bottleneck.

Task 2 before Task 1 by an average of 20 ms. We noted above that given the bottleneck model shown in Figure 1A, this finding indicates that Stages 2B + 2C must take less time, on average, than Stage 1C. Although it is highly plausible that Stage 1C (moving the tongue to begin speaking one of two words) would take more time than Stage 2C (depressing a finger), it still seems unlikely that Stage 2B could consume much time. This conclusion is consistent with our bottleneck model simulations, which arrived at an estimate of only 94 ms for the combined duration of Stages 2B and 2C. At the same time, Stage 2A was estimated to last 266 ms. This estimate seems much too long for the identification of the alphanumeric characters, suggesting that after practice Stage 2A includes some later processes (such as response selection) that were part of Stage 2B before practice. Thus, there is converging evidence that although a bottleneck still remained in Phase 3, it included fewer component processes than it did prior to practice.

There is tentative evidence that in Phase 4, where we inserted a new Task 1, the bottleneck was not as late as it was in Phase 3. First, the interaction between SOA and S-R compatibility was no longer underadditive; instead, there was a nonsignificant trend toward overadditivity (11 ms at the three long, positive SOAs; 30 ms at the −33-ms SOA). Although this finding is very tentative because of the small sample (three sessions, compared with nine sessions in Phase 3), it suggests that the entire Task 2 response selection stage was once again subject to the processing bottleneck. A second piece of evidence for an earlier bottleneck in Phase 4 is that Task 2 responses occurred much later relative to Task 1 responses. For instance, at the −33-ms SOA, the Task 2 response occurred an average of 57 ms after the Task 1 response (compared to 20 ms before the Task 1 response in Phase 3). Bottleneck stage arithmetic (assuming a bottleneck occurred on every trial at the most negative SOA) indicates that the quantity 2B + 2C − 1C in Phase 4 increased by about 77 ms relative to Phase 3. It seems unlikely that Stages 1C or 2C changed much (because the responses were very similar in the two phases), leading to the conclusion that Stage 2B increased by about 77 ms. Given that the total RT2 was similar across phases, Stage 2A would have had to decrease by about the amount that Stage 2B increased. Thus, use of a less practiced Task 1 might have shifted the bottleneck to an earlier locus than in Phase 3. This tentative conclusion is reasonable, because one would expect the central stage of a new, relatively unpracticed Task 1 to have higher resource demands.

Relation to Hazeltine et al. (2002)

One of the few studies that explicitly considered the latent bottleneck hypothesis was a non-PRP dual-task study by Hazeltine et al. (2002). Following Schumacher et al. (2001), participants performed several sessions of practice with two types of block types: pure single-task blocks and mixed blocks containing some single-task trials and some dual-task trials (with an SOA of 0 ms). The auditory–vocal task required participants to say "one," "two," or "three" in response to low-, medium-, and high-pitched tones, respectively. The visual–manual task required a spatially compatible keypress to the location of a white circle. Hazeltine et al. found very little dual-task interference after practice, even when the visual–manual task was made more difficult by reducing discriminability (Experiment 2) or by switching to an incompatible S-R mapping (Experiments 3–4). A small within-block SOA manipulation (50, 0, −50 ms) also had little effect on the size of dual-task costs.

Although dual-tasks costs were small, mean RTs were generally very short (about 250–300 ms for the visual task, depending on the condition, and about 300 ms for the auditory–vocal task). We noted above that very short RTs after practice make it possible for the bottleneck to be latent, producing little or no observed interference. Hazeltine et al. (2002) addressed this possibility by conducting a set of simulations of their SOA data (Experiment 4), focusing on the easy (compatible) version of the visual–manual task. For simplicity, Hazeltine et al. assumed that stage durations were deterministic. They also assumed that the bottleneck stages were handled on a first come, first served basis. For many different combinations of plausible prebottleneck stage durations (1A and 2A), Hazeltine et al. determined what bottleneck stage durations were consistent with the observed dual-task costs. On the basis of these simulations, they concluded that only bottleneck stages lasting 30 ms or less could fit the data. Although the authors did not favor a latent bottleneck account of their data, they acknowledged that they could not rule it out.

Concluding Remarks

The existence of small dual-task costs after practice has been used as evidence that processing bottlenecks can be eliminated. In this article, we have argued that this inference is suspect when RTs are short, because a processing bottleneck can exist yet be latent (i.e., produce little or no observable interference). We presented new analytical techniques designed to uncover latent bottlenecks. Using these techniques, we found strong evidence that a latent bottleneck was in fact present in a previously published experiment (Van Selst et al., 1999). To our knowledge, this is the first solid evidence for the occurrence of a latent bottleneck. These findings support our contention that the absence of dual-task costs, by itself, does not indicate that all processing bottlenecks have been abolished. This cautionary note applies not only to studies of practice effects, but to any study in which a small dual-task cost is found (e.g., studies of ideomotor compatibility; Greenwald & Shulman, 1973; Lien, Proctor, & Allen, 2002).

The present data, along with Van Selst et al. (1999) and Ruthruff et al. (2001), suggest that a bottleneck (albeit a relatively late one) persisted even after extensive practice in a PRP design. However, we remain open to the possibility that the true elimination of processing bottlenecks is possible under other conditions. An important goal for future research is to determine what these conditions might be. Progress toward this goal will require determining whether the cognitive limitation underlying the bottleneck has been eliminated, rather than simply determining whether the observed dual-task costs are small or large at particular SOAs.

References

Allport, D. A., Antonis, B., & Reynolds, P. (1972). On the division of attention: A disproof of the single channel hypothesis. *Quarterly Journal of Experimental Psychology, 24,* 225–235.

Davis, R. (1957). The human operator as a single channel information system. *Quarterly Journal of Experimental Psychology, 9,* 119–129.

De Jong, R. (1993). Multiple bottlenecks in overlapping task performance. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 965–980.

De Jong, R. (1995). The role of preparation in overlapping-task performance. *Quarterly Journal of Experimental Psychology, 48A,* 2–25.

Fletcher, B. C., & Rabbitt, P. M. A. (1978). The changing pattern of perceptual analytic strategies and response selection with practice in a two-choice reaction time task. *Quarterly Journal of Experimental Psychology, 30,* 417–427.

Greenwald, A. G., & Shulman, H. G. (1973). On doing two things at once: II. Elimination of the psychological refractory period effect. *Journal of Experimental Psychology, 101,* 70–76.

Hazeltine, E., Teague, D., & Ivry, R. B. (2002). Simultaneous dual-task performance reveals parallel response selection. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 527–545.

Hirst, W., Spelke, E. S., Reaves, C. C., Caharack, G., & Neisser, U. (1980). Dividing attention without alternation or automaticity. *Journal of Experimental Psychology: General, 109,* 98–117.

Johnston, J. C., & Delgado, D. F. (1993, November). *Bypassing the single-channel bottleneck in dual-task performance.* Paper presented at the 34th Annual Meeting of the Psychonomic Society, Washington, DC.

Keele, S. (1973). *Attention and human performance.* Pacific Palisades, CA: Goodyear.

Levy, J., & Pashler, H. (2001). Is dual-task slowing instruction dependent? *Journal of Experimental Psychology: Human Perception and Performance, 27,* 862–869.

Lien, M.-C., & Proctor, R. W. (2002). Stimulus-response compatibility and psychological refractory period effects: Implications for response selection. *Psychonomic Bulletin and Review, 9,* 212–238.

Lien, M.-C., Proctor, R. W., & Allen, P. A. (2002). Ideomotor compatibility in the psychological refractory period effect: 29 years of oversimplification. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 396–409.

Luce, R. D. (1986). *Response times.* New York: Oxford University Press.

Mowbray, G. H., & Rhoades, M. V. (1959). On the reduction of choice reaction times with practice. *Quarterly Journal of Experimental Psychology, 14,* 1–36.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum Press.

Pashler, H., & Baylis, G. C. (1991). Procedural learning: I. Locus of practice effects in speeded choice tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 20–32.

Pashler, H., & Johnston, J. C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology, 41A,* 19–45.

Pashler, H., & Johnston, J. C. (1998). Attentional limitations in dual-task performance. In H. Pashler (Ed.), *Attention* (pp. 155–189). Hove, United Kingdom: Psychology Press.

Ruthruff, E., Johnston, J. C., & Van Selst, M. A., (2001). Why practice reduces dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 3–21.

Schumacher, E. H., Seymour, T. L., Glass, J. M., Kieras, D. E., & Meyer, D. E. (2001). Virtually perfect time sharing in dual-task performance: Uncorking the central attentional bottleneck. *Psychological Science, 121,* 101–108.

Schwarz, W., & Ischebeck, A. (2001). On the interpretation of response time vs. onset asynchrony functions: Application to dual-task and precue-utilization paradigms. *Journal of Mathematical Psychology, 45,* 452–479.

Schweickert, R. (1978). A critical path generalization of the additive factor method: Analysis of a Stroop task. *Journal of Mathematical Psychology, 18,* 105–139.

Shaffer, L. H. (1975). Multiple attention in continuous verbal tasks. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V* (pp. 157–167). San Diego, CA: Academic Press.

Van Selst, M. A., Ruthruff, E., & Johnston, J. C. (1999). Can practice eliminate the Psychological Refractory Period effect? *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1268–1283.

Welford, A. T. (1952). The "psychological refractory period" and the timing of high-speed performance—A review and a theory. *British Journal of Psychology, 43,* 2–19.

Welford, A. T. (1976). *Skilled performance: Perceptual and motor skills.* Glenview, IL: Scott, Foresman.

Received February 26, 2001
Revision received July 2, 2002
Accepted July 18, 2002 ■