

Evidence against perfect integration of sensory information during perceptual decision making

Matthew A. Carland,¹ Encarni Marcos,^{2,3} David Thura,¹ and Paul Cisek¹

¹Department of Neuroscience, University of Montréal, Montréal, Québec, Canada; ²SPECS, Universitat Pompeu Fabra, Barcelona, Spain; and ³Department of Physiology and Pharmacology, Sapienza University of Rome, Rome, Italy

Submitted 16 March 2015; accepted in final form 23 November 2015

Carland MA, Marcos E, Thura D, Cisek P. Evidence against perfect integration of sensory information during perceptual decision making. *J Neurophysiol* 115: 915–930, 2016. First published November 25, 2015; doi:10.1152/jn.00264.2015.—Perceptual decision making is often modeled as perfect integration of sequential sensory samples until the accumulated total reaches a fixed decision bound. In that view, the buildup of neural activity during perceptual decision making is attributed to temporal integration. However, an alternative explanation is that sensory estimates are computed quickly with a low-pass filter and combined with a growing signal reflecting the urgency to respond and it is the latter that is primarily responsible for neural activity buildup. These models are difficult to distinguish empirically because they make similar predictions for tasks in which sensory information is constant within a trial, as in most previous studies. Here we presented subjects with a variant of the classic constant-coherence motion discrimination (CMD) task in which we inserted brief motion pulses. We examined the effect of these pulses on reaction times (RTs) in two conditions: 1) when the CMD trials were blocked and subjects responded quickly and 2) when the same CMD trials were interleaved among trials of a variable-motion coherence task that motivated slower decisions. In the blocked condition, early pulses had a strong effect on RTs but late pulses did not, consistent with both models. However, when subjects slowed their decision policy in the interleaved condition, later pulses now became effective while early pulses lost their efficacy. This last result contradicts models based on perfect integration of sensory evidence and implies that motion signals are processed with a strong leak, equivalent to a low-pass filter with a short time constant.

decision making; drift-diffusion model; urgency

MANY MODELS OF DECISION MAKING posit that the brain arrives at a decision by accumulating evidence in favor of competing choices until the accumulated total for one of the choices reaches a fixed decision bound (Bogacz et al. 2006; Gold and Shadlen 2007; Laming 1968; Ratcliff 1978). In this view, the rate of accumulation corresponds to the quality of sensory information, while the bound controls trade-offs between decision speed and accuracy (Reddi and Carpenter 2000). The most common of these models, the drift-diffusion model (DDM; Ratcliff 1978), assumes perfect integration of sensory samples to a fixed bound. This simple and elegant model replicates reaction time (RT) distributions and error rates across a variety of tasks (Ratcliff 2002; Ratcliff et al. 2004) and has been used to explain the buildup of neural activity to a threshold in many brain regions (Gold and Shadlen 2000,

2003; Kim and Shadlen 1999; Palmer et al. 2005; Ratcliff et al. 2007; Roitman and Shadlen 2002).

While such models capture the intuition that one should accumulate evidence until the accumulated total satisfies some criterion of accuracy, an important question concerns what constitutes evidence in a given situation. In particular, during the kinds of perceptual discrimination tasks usually studied, sequential samples of a static stimulus are increasingly redundant; taking a second look may be useful, but taking a tenth or eleventh is less so. This motivates one to only integrate information that is novel. Furthermore, in many natural scenarios the world can suddenly change, and the decision-making system must be quick to respond to such changes. For these reasons, several alternatives to the DDM have been proposed. In particular, it has been suggested that integration is leaky, with old information gradually being discounted as new information arrives (Busemeyer and Townsend 1993; Roe et al. 2001; Usher and McClelland 2001). This is consistent with the observation that decisions are dominated by information from a limited time window (Chittka et al. 2009; Cook and Maunsell 2002; Luna et al. 2005; Uchida et al. 2006; Yang et al. 2008), emphasizing novel information. It has also been suggested that the decision bound is not fixed but decreases over time, possibly because of a growing “urgency” signal that pushes neural activity toward commitment as time passes (Churchland et al. 2008; Cisek et al. 2009; Ditterich 2006; Hanks et al. 2014). This proposal is motivated by analyses showing that for tasks in which subjects are allowed to respond at any time, a decreasing decision bound yields higher reward rates than any setting of a fixed bound (Drugowitsch et al. 2012; Thura et al. 2012).

Importantly, if an urgency signal is present, then most existing data can be explained even if the integration is highly leaky (Cisek et al. 2009; Ditterich 2006), equivalent to a low-pass filter with a relatively short time constant (e.g., 200 ms). This is called the “urgency-gating model” (UGM). Crucially, in nearly all previous experiments the sensory information provided to subjects was constant within each trial, and under such conditions a perfect integrator with fixed bound behaves very similarly to a leaky integrator with an urgency signal (Cisek et al. 2009; Hawkins et al. 2015b; Thura et al. 2012). Thus data from experiments featuring constant-evidence tasks can support either model, and even sophisticated comparisons of model fits (Hawkins et al. 2015a) do not produce consistent answers across tasks, subjects, or modeling assumptions. Part of the reason is that these models are relative extremes on a wide continuum of parameters, whereby the DDM assumes zero leak and a fixed accuracy bound while the

Address for reprint requests and other correspondence: P. Cisek, Dept. of Neuroscience, 2960 chemin de la tour, Montréal, QC, H3T 1J4 Canada (e-mail: paul.cisek@umontreal.ca).

UGM proposes a strong leak and a growing urgency signal that implements a decreasing accuracy bound. Deciding on the appropriate settings of these parameters, even for a specific subject, is difficult using data that can be fit with a wide range of parameter combinations. The purpose of the present experiment is to place bounds on the values that these parameters may take, and therefore to constrain the set of possible models that may be used to account for behavior during perceptual decision-making tasks.

We recently conducted two behavioral studies (Cisek et al. 2009; Thura et al. 2012) and one neural recording study (Thura and Cisek 2014) featuring tasks in which sensory information changed over the course of each trial. In such conditions, different settings of leak and urgency parameters produce clearly divergent predictions that are easy to distinguish empirically. Indeed, the results of those studies were incompatible with perfect integration and instead supported a UGM with a filter time constant of 200 ms or less. However, it could be argued that those results were task dependent and do not generalize to the kinds of constant-evidence perceptual discrimination tasks usually studied. Consequently, here we aim to distinguish among perfect vs. leaky integration by observing the effect of brief motion pulses on subjects' RTs during the classic random-dot motion discrimination task. Previous studies using motion pulses have shown that they can influence RTs (Huk and Shadlen 2005)—but the critical test examined here is how the effects of these pulses change as subjects modify their decision policy.

The logic of our experimental design is illustrated in Fig. 1. As in a typical random-dot motion discrimination task, subjects are presented with a noisy coherent motion stimulus that remains constant within a given trial and are asked to respond as soon as they detect the direction of motion. Unknown to the subjects, some trials contain a brief 100-ms increase of motion coherence at various times after motion onset. If decisions are made using a perfect integrator such as the DDM (Fig. 1, *A* and *B*), then motion pulses presented prior to the time of decision will briefly increase the rate of integration and result in decisions that are made earlier, on average, than those made in

no-pulse trials. Pulses that appear later in the trial will have no effect if the decision bound is low (Fig. 1*A*), because they occur too late to affect the decision process. If the decision bound is high (Fig. 1*B*), then both early and late pulses will affect RTs. In other words, as subjects slow down their decision policy, the time window during which pulses can have an effect on RTs is predicted to expand. Critically, this model predicts that early pulses will always be at least as effective as late pulses in reducing RTs.

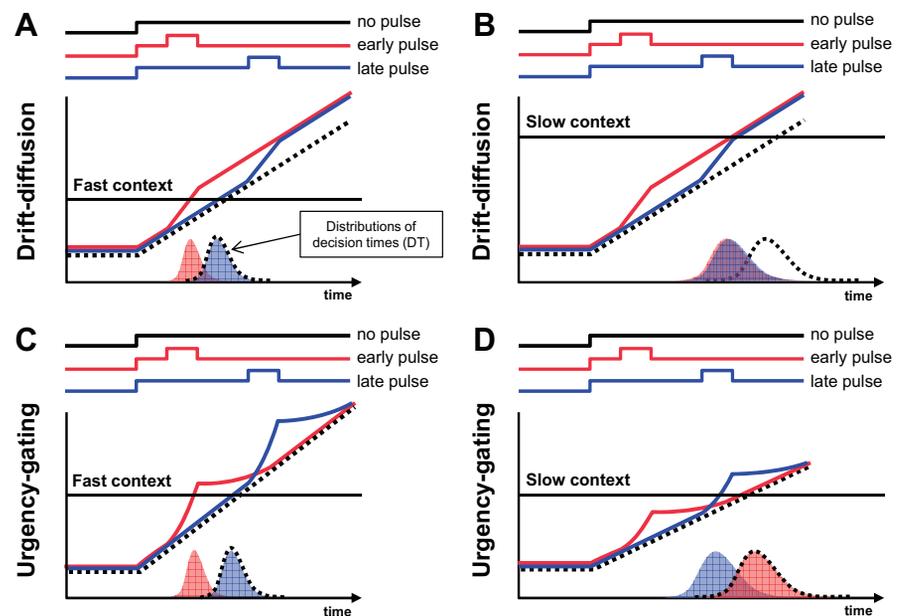
If the perfect integrator is replaced with a low-pass filter (a highly leaky integrator) and combined with an urgency signal, then early pulses will still have a stronger effect than late pulses under fast decision policies, i.e., when the urgency signal grows quickly (see Fig. 1*C*). However, the reverse will be true under slower decision policies, i.e., when urgency grows more slowly (Fig. 1*D*), because the information provided by early pulses will have leaked away by the time the decision bound is crossed. Thus as subjects slow down, the time window in which pulses are effective is predicted to shift, and later pulses will become effective while early pulses will lose efficacy.

Our approach for testing these predictions is to present subjects with an identical set of pulse trials in two different contexts: one in which they are motivated to respond quickly and another in which they are motivated to slow down. Some of these data have previously appeared in abstract form (Carland et al. 2013).

METHODS

Subjects and apparatus. Forty-four right-handed participants (24 women, 20 men; ages 19–37 yr) with normal or corrected-to-normal vision provided written consent and were naive to the purpose of our experiment. Participants were seated in front of a large digitizing tablet placed at arm level for recording movements of a handheld cordless stylus embedded within a vertical plastic handle (125-Hz sample rate with 0.013-cm accuracy). Stimuli, targets, and cursor position feedback were projected by an LCD monitor onto a half-silvered mirror positioned 16 cm above and parallel to the digitizer surface, and thus appeared to float on the plane of the digitizing tablet.

Fig. 1. Schematic of predicted effects of motion pulses assuming the drift-diffusion model (DDM) vs. the urgency-gating model (UGM). *A*: here, following the DDM, we assume a perfect integrator of motion signals with a fixed bound that is set to a low value, to emphasize speed in a “fast” task context. As a result, early motion pulses (red line) will result in significantly shorter reaction time (RT) distributions than in no-pulse trials (black dotted line), but late motion pulses (blue) will have no effect because they occur after the decision bound has already been reached. Schematic RT distributions are shown on *x*-axis. *B*: in a “slow” task context, the bound is set to a higher value, and as a result both early and late pulses cause a reduction of RTs compared with no-pulse trials. *C*: here, following the UGM, we assume that the motion signal is low-pass filtered and combined with a growing urgency that is steep, to emphasize speed in the “fast” task context. As in *A*, early pulses have an effect but late pulses occur too late to reduce the RT. *D*: in the “slow” task context, the urgency is shallower, and so late pulses now significantly reduce the RT. However, in contrast to *B*, early pulses no longer reduce RTs because their effect has leaked away by the time the threshold is crossed.



The subjects' task was to report the direction of motion of the stimulus by completing reaching movements toward one of two targets whose locations corresponded to the potential motion directions. The task and data collection was programmed in LabVIEW (National Instruments, Austin, TX), stored in a database (Microsoft SQL Server 2005, Redmond, WA), and analyzed with custom MATLAB scripts (MathWorks, Natick, MA). The experimental protocol was approved by the university ethics committee. The data presented in this report are available upon request from the corresponding author.

Behavioral task. Each trial began when subjects moved the cursor into a small circular target (1 cm in diameter) near the center of a white display. After 500 ms, two circular targets (3 cm in diameter) appeared 6 cm to each side of the stimulus display area, separated by 180°. The targets were projected to appear oriented along the natural direction of single-joint elbow movements, thereby minimizing any potential difference in biomechanical cost between the two peripheral targets (Fig. 2A). Three hundred milliseconds after the start target was entered, 200 black dots appeared in a borderless circular area (3-cm diameter) in the center of the display between the two targets. Each of the dots was redrawn in a new location 2 pixels away from its previous location on each frame (60 Hz). While most dot movements were random, a subset of the dots was redrawn along a vector corresponding to the location of one of the two targets. While the individual dots assigned to the coherently moving subset changed from frame to frame, the resulting percept was of a persistent motion signal whose direction subjects could reliably and accurately report with a degree of difficulty inversely related to the percentage of coherently moving dots (Kim and Shadlen 1999; Newsome et al. 1989).

Subjects were given up to 3,000 ms to report the direction of the coherent motion by moving from the initial start target to one of the two choice targets and were free to respond at any time. Movements had to be completed in <1,000 ms and had to land within the chosen target circle. The motion stimulus continued until the cursor crossed a target circle's border. The cursor had to remain within the chosen target for 500 ms, at which point the outline of the target turned green or red to indicate a correct or incorrect choice, respectively. After a brief intertrial interval of 500 ms, all on-screen objects disappeared except for the starting target, and a new trial began.

In the analyses reported here, response times (RTs) for each trial were obtained post hoc by determining the precise moment at which the cursor's velocity began to increase from a point of rest within the start target. However, an ad hoc estimate of RT based on the time at which the cursor exited the boundary of the start target was used to

obtain session-specific estimates of subjects' mean RTs; these were then used to provide online feedback during the experimental sessions, as described below.

Before each session began, we presented the subject with 40 very easy motion-discrimination trials in which the motion coherence was 100% and instructed them to respond as rapidly as possible. The average RT from these trials was then stored as a session-specific estimate of "non-decision time" comprising both sensory and motor delays (mean = 475 ms, SD = 103 ms). This measurement was used to classify correct and incorrect responses for trials in which the motion signal could change directions over the course of a single decision (see below). Importantly, these initial 40 trials were the only ones for which subjects were ever provided with explicit instructions about how quickly to respond: for the main experimental task, subjects were informed of the 3-s time limit but were told that they could make their decision whenever they liked, although most of our subjects very rarely took >1,800 ms to make their decision.

For any experimental session, subjects completed one of two session types: "blocked" or "interleaved." "Blocked" sessions consisted entirely of trials with a single, common baseline motion coherence value of 3%: we refer to these as "constant-motion discrimination" (CMD) trials. In 40% of such trials there were no additional changes to the stimulus, and we refer to these as "no-pulse" trials. The remaining 60% of the CMD trials contained brief motion "pulses" during which the coherence of the motion stimulus was doubled (to 6%) for 100 ms. Such brief coherence manipulations have previously been shown to affect response timing in motion discrimination tasks (Huk and Shadlen 2005; Wong et al. 2007), even though they are not consciously detectable by our subjects (confirmed by postexperiment interview and consistent with similar studies; see Kiani et al. 2008). Thus, as far as the subjects were aware, the motion coherence for all CMD trials appeared to remain constant throughout each trial, regardless of whether a pulse was or was not actually shown. These pulses could occur 100, 200, or 400 ms after stimulus onset, and we refer to such trials collectively as "pulse" trials. Each of these pulse timing conditions thus comprised 20% (60% total) of the total number of trials in the "blocked" condition.

"Interleaved" sessions consisted of a mix of trial types. Twenty percent of the trials in these sessions were CMD trials—including both pulse and no-pulse trials—identical in every respect to those presented in the blocked sessions. Each of the four CMD trial types comprised 5% of the total number of trials encountered during these sessions. These CMD trials were randomly interleaved among "variable-motion discrimination" (VMD) trials, which comprised the re-

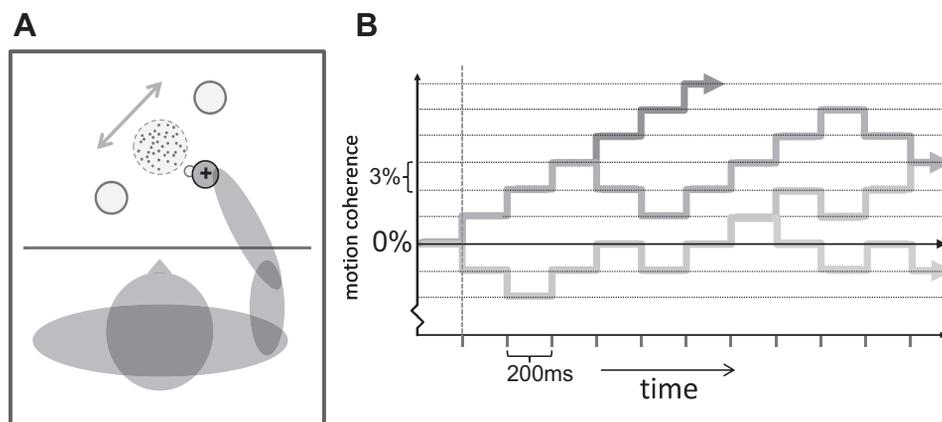


Fig. 2. Experimental setup and the variable motion discrimination (VMD) trials. *A*: overhead view of the experimental display. Subjects discriminated the direction of movement of a random-dot motion display contained within a central circle (the border of which was not drawn on screen) by moving a handheld stylus (solid black circle) from a start target (small circle) to 1 of 2 equidistant peripheral targets, each separated by 180° and oriented with respect to the elbow's natural direction of motion. *B*: example trajectories of the time course of motion evidence in VMD trials. After being initialized at 0% at the beginning of each trial, the motion signal is adjusted in increments of 3% motion coherence toward 1 of the 2 targets (with equal probability) every 200 ms. Note that when the motion signal already favors a target, a step "toward" the opposing target corresponds to a weakening of the strength of the motion signal.

maintaining 80% of the interleaved session. These trials began with a net motion coherence of $\pm 3\%$, and this motion signal was adjusted either up or down in 3% steps every 200 ms (sometimes reversing the direction of motion; see Fig. 2B). Of the VMD trials, 60% were random, such that each motion coherence change was given an independent and equal probability of favoring either of the two possible targets. The remaining 40% of trials were divided among a number of pregenerated trial types similar to those featured in previous studies (Cisek et al. 2009; Thura et al. 2012)—these included “easy” and “ambiguous” trial types, as well as a variety of “bias” trials. These trials were included to test whether effects previously reported (Thura et al. 2012) were still observed in our subjects. In brief, “easy” trials were those in which the motion signal reached $\geq 9\%$ within the first 800 ms and remained at or above 9% for the remainder of the trial; “ambiguous” trials were those in which the absolute motion signal remained within 6% of zero throughout the trial. “Bias” trials included brief biases (800 ms or 1,000 ms) either for or against the final motion direction and then resembled easy trials (see Thura et al. 2012 for details). Since the previously reported effects were indeed confirmed, we do not analyze these further and refer the reader to Thura et al. (2012).

Each session type consisted of a single pseudorandom, predefined sequence of trials that was the same for all subjects. Subjects had to achieve a total of 560 correct trials to complete one “blocked” session or 500 to complete an “interleaved” session. Correct trials were always defined with respect to whether the net direction of the motion signal indicated the chosen target at the time of the subject’s decision. Decision accuracy classification was straightforward for CMD trials, in which the motion signal always favors one of the two targets; in contrast, the motion signal in VMD trials could sometimes indicate different targets over the course of a single trial. We therefore determined decision accuracy for VMD trials by subtracting each subject’s estimated non-decision time from the approximate time of the start of the movement with which he/she reported his/her decision. The trial was counted as correct if the motion signal at this time indicated the chosen target, even if the signal had changed directions between the effective decision time and the offset of the motion stimulus. Both session types took ~ 50 min on average to complete, depending on an individual subject’s speed and accuracy on that day. Importantly, however, subjects were paid the same amount per session (\$20 CAD) regardless of how long it took for them to reach the required quota of correct trials. Thus, while we otherwise provided no explicit penalty for wrong answers, the structure of the task nonetheless implicitly motivated subjects to minimize the total session duration by finding a decision policy that maximized success rate for each experimental session.

Crucially, our two session types differed with respect to the value of stimulus observation time. In blocked sessions, the average success rate was not appreciably improved with longer observation times because the motion in CMD trials was essentially constant; in other words, all relevant decision information was fully present from the start of each trial and therefore motivated relatively rapid decisions. In contrast, the interleaved sessions predominantly featured VMD trials. In these trials not only does the motion signal itself change over time, but—crucially—the range over which it may vary increases in direct proportion to elapsed time. Thus the more one prolongs one’s decision, the greater the chance that the motion signal will reach a value of larger magnitude. Because this tendency is unique to the VMD trials, which only appear in interleaved sessions, we expected that subjects would exploit this by generally adopting a slower decision policy during the interleaved sessions relative to the blocked sessions. With respect to the two models being tested, this difference in decision policy can be achieved either by increasing the decision bound, as per Fig. 1, A and B, or by decreasing the urgency signal’s slope, as per Fig. 1, C and D.

Consequently, we expected that this difference in decision policy between task conditions would result in systematically different RTs

during the otherwise identical CMD trials common to both session types. Note that while subjects were generally able to discern that some sessions were different from others, postexperiment interviews revealed that they could not specify the precise nature of the differences in the stimuli across sessions; nor could they detect the presence of the pulses in CMD trials or specific types of VMD trials. Furthermore, even had they been able to tell the difference between session types, the fact that trial sequences were randomized would preclude them from being able to know in advance what type a given trial would be. Thus we can be confident that any differences in our subjects’ decision policies in CMD trials across session types indicated adaptations to the implicit reward structure inherent to each session type as a whole, rather than specific strategies adopted on a trial-by-trial basis.

Most importantly, this contextual manipulation of decision policy, if successful, would allow for an empirical discrimination of the divergent predictions of the DDM and UGM by comparing trials in which evidence is identical and only the effective decision policy differs (as per Fig. 1). Specifically, if early pulses have an effect on RTs that is always at least as strong as late pulses, then this would support a pure integration model such as the DDM. If, instead, early pulses lose their efficacy as decisions are slowed while late pulses become more effective, then this would support models in which evidence is not integrated over time but instead low-pass filtered with a highly leaky integrator.

To test these predictions, we initially ran 39 subjects for three sessions each (1 blocked, 2 interleaved), thereby allowing across-subject analyses (total trials = 90,302). Then, to obtain enough data to perform within-subject analyses, we ran four of these subjects plus an additional five new subjects for 10–24 sessions each (total trials = 71,736).

Modeling. To implement the DDM, we use the following equation:

$$\frac{dx}{dt} = aE + N \quad (1)$$

where evidence E is set to 1 to simulate 3% coherent motion and increased to 2 for 100 ms to simulate the motion pulse. The variable N denotes 1,000-Hz intratrial Gaussian noise with mean zero and standard deviation Θ . The variable a denotes an “attentional gain” that varied from trial to trial with mean 1 and standard deviation 1.5, effectively implementing endogenous variations in the signal-to-noise ratio that differed across trials but was constant within each trial (negative assignments to parameter a were redrawn from the source distribution until it was positive, and thus this parameter was never permitted to be negative). The decision was made when the variable $x(t)$ reached a threshold $\pm T$, and a non-decision delay of $t_0 = 300$ ms was added to yield the total RT. The T and Θ parameters were adjusted to fit the data separately in the blocked and interleaved conditions, using an exhaustive grid search to find the pair that minimized the mean squared error between the model’s estimate of the median RT in no-pulse trials and its estimate of the effect of the 100, 200, and 400 ms pulses on median RT (the latter 3 error terms were multiplied by 10 to emphasize the importance of these effects). After finding the best parameters with a grid search, we fine-tuned them by hand to further improve the fit. The purpose of all of our fitting procedures was to capture the qualitative differences among the various pulse conditions (e.g., relative effect of early vs. late pulses), with less emphasis placed on precise quantitative fits. For reasons explained below, in some simulations we also allowed the model extra parameters to implement a delay between the onset of motion and the start of integration.

To implement the UGM, we first low-pass filter the sensory information with a first-order linear differential equation:

$$\tau \frac{dx}{dt} = -x + (aE + N) \quad (2)$$

where the time constant is set to $\tau = 167$ ms. Note that the precise value of the time constant is difficult to establish with confidence from behavioral data, because any changes of the time constant can be “traded off” with changes to the intertrial variability of the urgency signal or other potential variance parameters. Thus not knowing the precise value of these parameters ahead of time either necessitates the introduction of an intractably large number of free parameters or else calls for assumptions about what these parameters may be. Consequently, we assumed a time constant of 167 ms on the basis of a number of previous behavioral and physiological studies that have suggested that it must be at least 100 ms (Cisek et al. 2009; Thura et al. 2012) and at most 200 ms (Thura and Cisek 2014). The evidence (E) and attentional gain (a) parameters were exactly the same as in the DDM, and intratrial noise was 60 Hz with mean zero and standard deviation set to $\Theta = 5$.

The resulting variable $x(t)$ is then combined with an urgency signal as

$$y(t) = x(y) \cdot U(t) \quad (3)$$

where $U(t)$ is the urgency signal that rises from zero with a slope that varies from trial to trial according to a log-normal distribution with parameters μ and σ . While our analyses of monkey behavior suggest that the urgency signal has a nonzero baseline value that can vary with speed-accuracy trade-offs (Thura et al. 2014), here we simply set its baseline to zero so as to avoid introducing another free parameter. The decision was made when the variable $y(t)$ reached a threshold of $T = \pm 300$, and a non-decision delay $t_0 = 300$ ms was added to yield a RT. To simulate each of the conditions (blocked and interleaved), we picked values of μ and σ that produced the best fit to the mean and standard deviation of the RT distribution from 3% coherence no-pulse CMD trials in each condition. We then used the same parameter settings when simulating pulse trials, trusting the effects of our different pulse timings to “fall out” out of the parameters used to fit the no-pulse trials.

Note that for each model two parameters were adjusted to fit the data. For the DDM these were the threshold T and the noise Θ , while for the UGM they were the μ and σ parameters used to determine the urgency slope. These pairs are functionally related: T and μ influence the means of RT distributions, while Θ and σ influence their variability. While the settings of the DDM parameters were determined through an exhaustive search for the least mean squared error fit to data from all trials, the UGM parameters were only adjusted to fit the no-pulse trials, and the effects of pulses expected to follow simply from the assumption of a short time constant. The models were used to simulate 5,000 trials for each trial type in each task condition, and the results were analyzed in the same way as the behavioral data.

RESULTS

Effects of sessions. The first step of our analyses was to determine whether our manipulation of decision policy succeeded in slowing subjects down in the interleaved condition. We did this by comparing RTs for identical no-pulse CMD trials across the two conditions. Mean (\pm SE) RTs of individual subjects are shown in Fig. 3A for constant-evidence, 3% motion coherence trials without pulses in both the “blocked” and “interleaved” conditions. All individual data points lay above the unity slope line, indicating that the mean RTs for identical CMD trials were slower when these were interleaved among VMD trials than when blocked together ($P = 3.3 \times 10^{-41}$, K-S test), thereby corroborating the efficacy of our decision policy manipulation. Similarly, cumulative RT distributions

for no-pulse CMD trials for the nine subjects who completed the greatest number of experimental sessions (Fig. 3, *B* and *C*) both show a clear rightward displacement along the x -axis, indicating later RTs in the interleaved condition.

The main effect on overall mean RT was found for nearly all of our subjects (42/44 subjects, $P < 0.05$ for each, K-S test) despite the otherwise large intersubject variability in overall speed. Thus while some subjects tended to be considerably faster than others, all of them individually slowed down during the interleaved sessions. Crucially, this behavior emerged despite the fact that no explicit instructions were ever provided to the subjects regarding the timing of their decisions. The strength and consistency of this result across all of our subjects thus strongly supports the effectiveness of our contextual manipulation on subjects’ decision policies. Moreover, because these differences obtain in no-pulse CMD trials that were otherwise identical, the most parsimonious interpretation of this effect implicates a slowed decision policy for the interleaved sessions relative to the blocked sessions. This corresponds within the framework of the DDM to increasing the decision bound during interleaved sessions or in the framework of the UGM to decreasing the slope of the urgency signal (Fig. 1).

Effects of pulses on reaction times. Next, we analyzed the effects of pulses in each task context to distinguish between the specific predictions of each model under conditions of changing evidence (Fig. 1). Similar to the above, this analysis also focused exclusively on CMD trials (both with and without pulses), which were identical in both session types.

As can be seen in Fig. 3A, collectively our subjects varied greatly in terms of average RTs within both task conditions. Such intersubject variability prohibited a general comparison of across subject means, as pooling together data from subjects with such different RT distributions could weaken any latent significant effects within each subject. Additionally, normalizing by RT would obscure the time course of the impact of evidence on the developing decision. Thus we instead pooled subjects into subgroups on the basis of the similarity of their mean RTs. Two such subgroups are indicated in Fig. 3A. The “fast subgroup” was defined so as to capture fast responders while including most of the subjects for whom we had 10+ sessions’ worth of data, while the “slow subgroup” aimed to capture a similar range of slower responders.

The pooled RT distributions for all CMD trials appear in Fig. 4. For the “fast” subgroup (Fig. 4A), the 100 ms and 200 ms pulses significantly sped up RT in the “blocked” condition (K-S test $P = 6.7 \times 10^{-16}$ and $P = 1.8 \times 10^{-9}$, respectively), while pulses at 400 ms had no significant effect ($P = 0.21$, K-S test). For the interleaved sessions, however, the 100 ms pulses lost their efficacy ($P = 0.052$, K-S test), while the 400 ms pulses became effective ($P = 2.1 \times 10^{-5}$, K-S test). The 200 ms pulses remained effective in the interleaved condition ($P = 0.002$, K-S test). This is consistent with the predictions shown in Fig. 1, *C* and *D*.

Results for a second, “slow” subgroup appear in Fig. 4B. For these subjects, the mean RTs were long even during the “blocked” condition, and none of the pulses had a significant effect on decision timing, although there is an apparent trend for faster decisions with later pulses (Fig. 4B, *top*). RTs for the “interleaved” condition are even later, and also not affected by any of the pulses (Fig. 4B, *bottom*). Note, however, that this

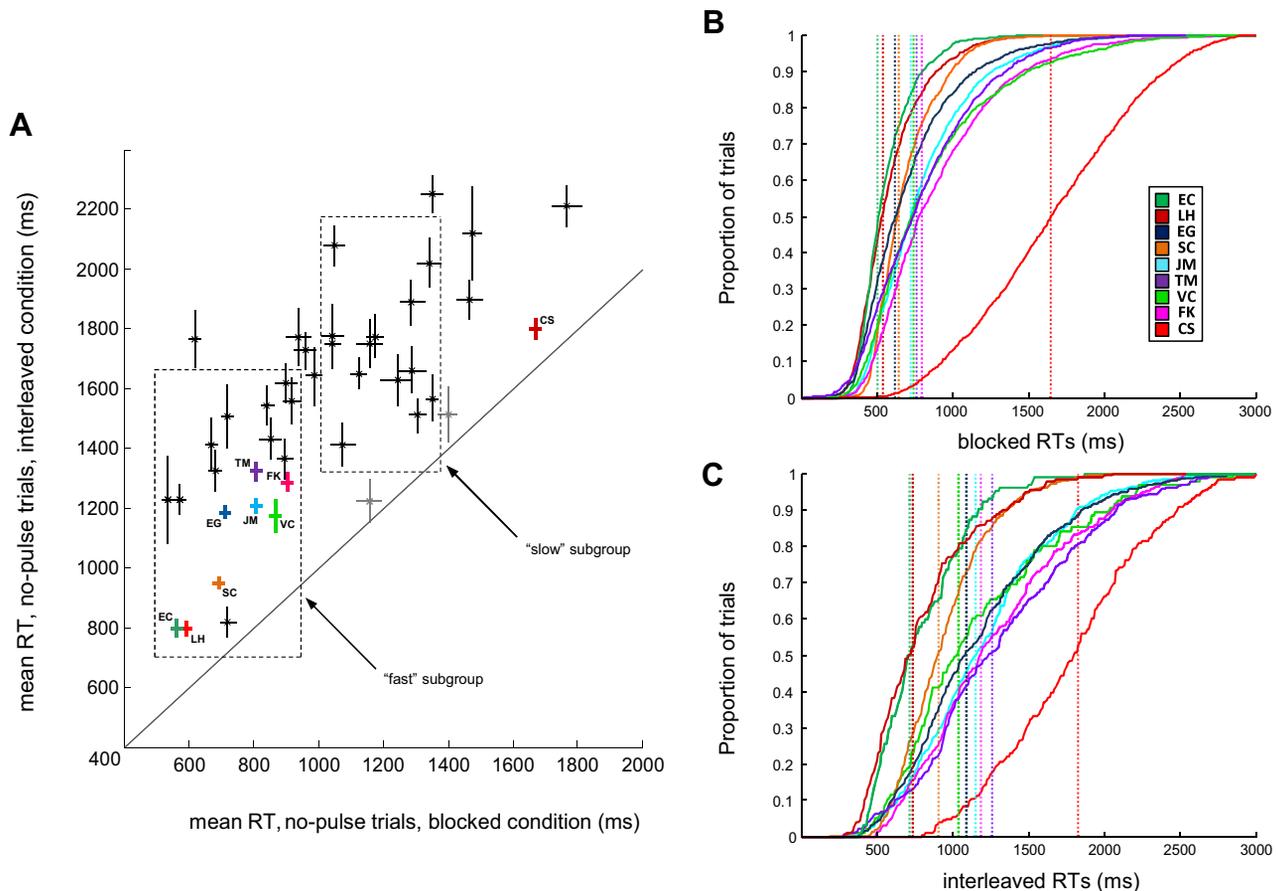


Fig. 3. RTs in no-pulse trials in the blocked and interleaved conditions. *A*: mean RTs of individual subjects in no-pulse trials during the blocked (*x*-axis) vs. the interleaved (*y*-axis) conditions. Crosses show SE, and colored points indicate subjects who performed 10 or more experimental sessions each. All but 2 subjects (shown in light gray) showed a significant difference between conditions ($P < 0.05$, K-S test). Dashed rectangles show subjects who were grouped together for the pooled analyses shown in Fig. 4. *B*: cumulative RT distributions in no-pulse trials during the blocked condition for the 9 subjects who performed >10 sessions (same colors as in *A*); vertical lines indicate medians. *C*: cumulative RT distributions and medians of the same subjects in no-pulse trials during the interleaved condition. In all cases, RTs for such trials are significantly shorter in the blocked condition than in the interleaved condition (all no-pulse trials pooled by session type, $P = 8.84 \times 10^{-43}$, K-S test).

subgroup consists only of subjects who completed three sessions each (see Fig. 3*A*) and therefore comprises substantially fewer trials than the "fast" subgroup, which may therefore account for the failure of this trend to reach statistical significance.

While here we have shown only two subject groups, other arbitrary groupings of subjects along similar lines (data not shown) yield effects that are qualitatively similar to those discussed above. Similar results were also obtained on the level of individual subjects. Figure 5*A* shows the cumulative RT distributions from blocked and interleaved sessions for one representative subject (*subject JM*; $n_{\text{sessions}} = 20$, $n_{\text{trials}} = 12,900$). Like 42/44 subjects, *JM*'s median RTs for all CMD trial types were significantly shorter during the blocked condition than during the interleaved condition (727 ms vs. 1,057 ms, $P = 1.53 \times 10^{-80}$, K-S test). Furthermore, in the blocked condition, the 100 ms and 200 ms pulses significantly sped up *JM*'s RTs ($P = 6.98 \times 10^{-12}$ and 2.42×10^{-6} , respectively, K-S test), while pulses at 400 ms were not significantly effective ($P = 0.21$, K-S test). When *JM* was completing an interleaved session, however, pulses at 400 ms now significantly sped up RTs (K-S test $P = 0.03$), while the RT distributions for trials with early pulses (100 ms and 200 ms)

were no longer statistically distinguishable from no-pulse trials (K-S test $P = 0.75$ and 0.98 , respectively).

A similar—though not identical—pattern obtained for other subjects as well. For example, for *subject VC* ($n_{\text{sessions}} = 16$, $n_{\text{trials}} = 9,957$; data summary shown in Fig. 5*B*) only the 200 ms pulse had a significant effect in the blocked condition (K-S test $P = 0.016$ in the blocked condition, $P = 0.27$ in the interleaved condition), whereas only the 400 ms pulse was effective in the interleaved condition (K-S test $P = 0.38$ in the blocked condition, $P = 0.001$ in the interleaved condition). For *subject SC* ($n_{\text{sessions}} = 16$), the first two pulses were effective in the blocked condition, whereas only the last two were effective in interleaved sessions. Overall, of the nine subjects who performed >10 sessions, six (*JM*, *EC*, *VC*, *SC*, *EG*, *TM*) showed patterns in their data that are qualitatively consistent with the results described above, insofar as the most effective pulse was earlier in the blocked than in the interleaved condition (although this only reached significance in 4/6 cases). Of the remaining subjects, two (*FK* and *CS*) showed no effects of pulses at all (these were our two slowest subjects), while one (*LH*) showed similar effects in both blocks (this subject was fast in both blocks) (Fig. 3*A*). This too is consistent with the UGM, which predicts that the relative difference between a

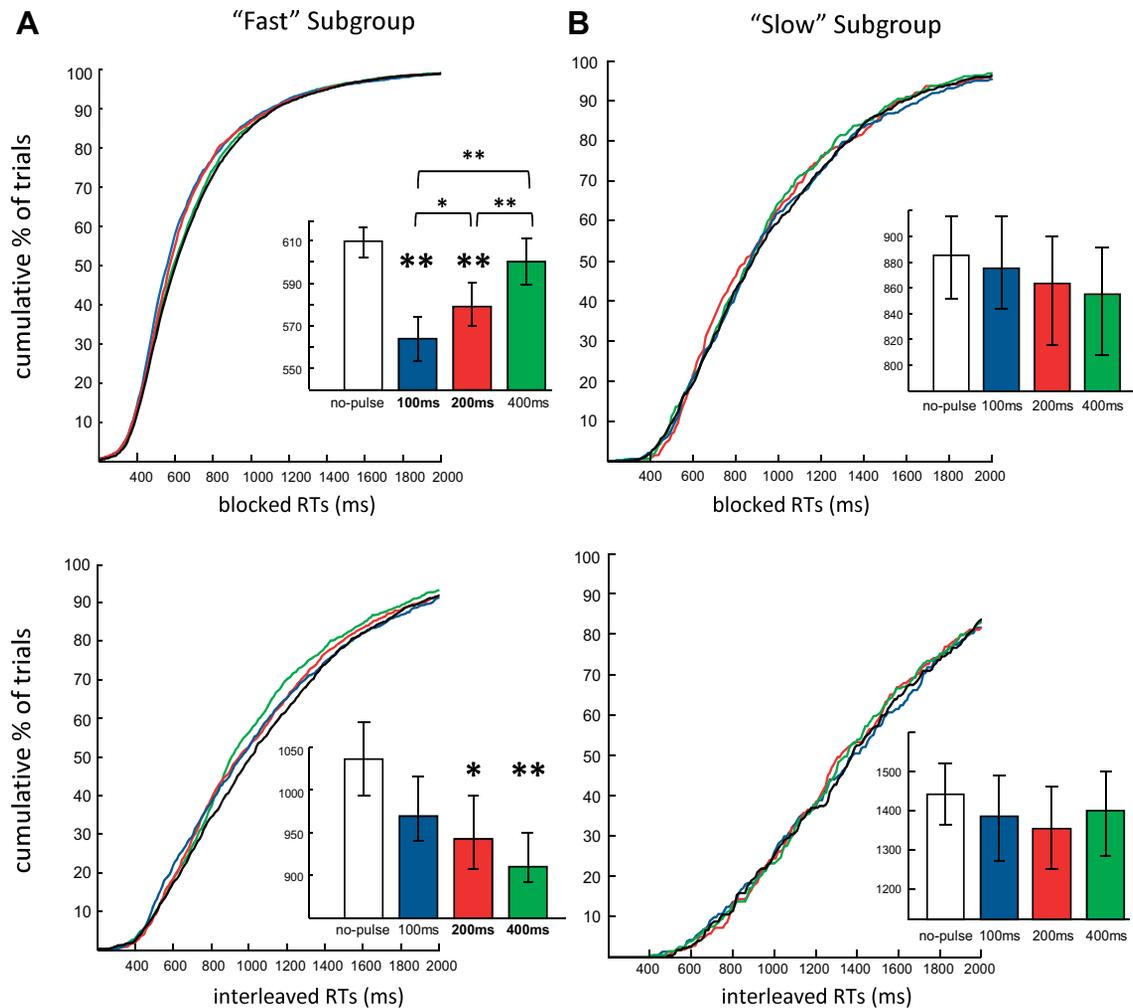


Fig. 4. Comparison of RT effects of pulses in constant-motion discrimination (CMD) trials. *A*: cumulative RT distributions for no pulse (black), 100 ms pulse (blue), 200 ms pulse (red), and 400 ms pulse (green) for the “fast” subgroup of subjects indicated in Fig. 3. *Insets* show the corresponding median RTs (with 95% confidence intervals); asterisks indicate significant differences (large asterisks indicate statistical comparisons against no-pulse trials, small asterisks show comparisons between pulse trials). * $P < 0.01$; ** $P < 0.001$. *B*: cumulative RT distributions for the “slow” subgroup of subjects. Same format as *A*.

subject’s overall mean RT and the timing of relevant changes in sensory evidence determines the extent to which these changes influence the timing of the decision. Because the pulse timings we selected occurred relatively early in a trial, the differences between the effects of these pulses on RT between the two session types are most pronounced for subjects whose mean RTs are relatively short. A clear demonstration of this comes from the fact that one subject’s RTs were extremely slow (Fig. 3*B*, outlying red line) and did not appear to be influenced by any pulses in either condition—as if they had leaked away.

The relationship between the median RT of an individual subject and the efficacy of pulses at different times is summarized in Fig. 6, where the fastest eight of our subjects with >10 experimental sessions are plotted together to illustrate how the effect of a given pulse on RT depends on the RT itself. In the blocked condition—where RTs are faster for any given subject—the 100 ms pulse has the strongest effect, followed by the 200 ms pulse, with the 400 ms pulse having little or no effect. However, for the slowest subjects (*TM*, *FK*) even the effects of 100 ms and 200 ms pulses are reduced. In contrast, the patterns of effects are inverted in the interleaved condition, where

median RTs are longer. Here, most of the subjects are most strongly influenced by the 400 ms pulse, while some relatively faster subjects are also influenced by the 200 ms pulse. While there is variability in the data, and many points do not reach significance, there is a clear tendency for early pulses to be stronger than late pulses in the blocked condition, while the opposite is true in the interleaved condition. Furthermore, there is a “window” of RTs in which the pulses influence behavior, and that window shifts as subjects slow down their decision policy. The 100 ms pulse is most effective on RTs between 650 and 750 ms (blocked sessions), while the 400 ms pulse is most effective (in interleaved sessions) ~ 300 ms later, as expected.

Analyses of performance accuracy. For the fast subgroup of subjects, accuracy in the no-pulse CMD trials was 83.0% (9,939/11,969) during the blocked condition and 86.4% (1,703/1,972) during the interleaved condition, and this difference was significant (binomial test, $P < 0.01$). Within the blocked condition, pulses at 100 ms produced a slight and significant increase in accuracy (2,633/2,997, 87.9%), as did pulses at 200 ms (2,587/3,002, 86.2%), but the increase in accuracy in 400 ms pulse trials (2,548/3,007, 84.7%) was not significant (binomial test, $P > 0.05$). Within the interleaved condition, none of

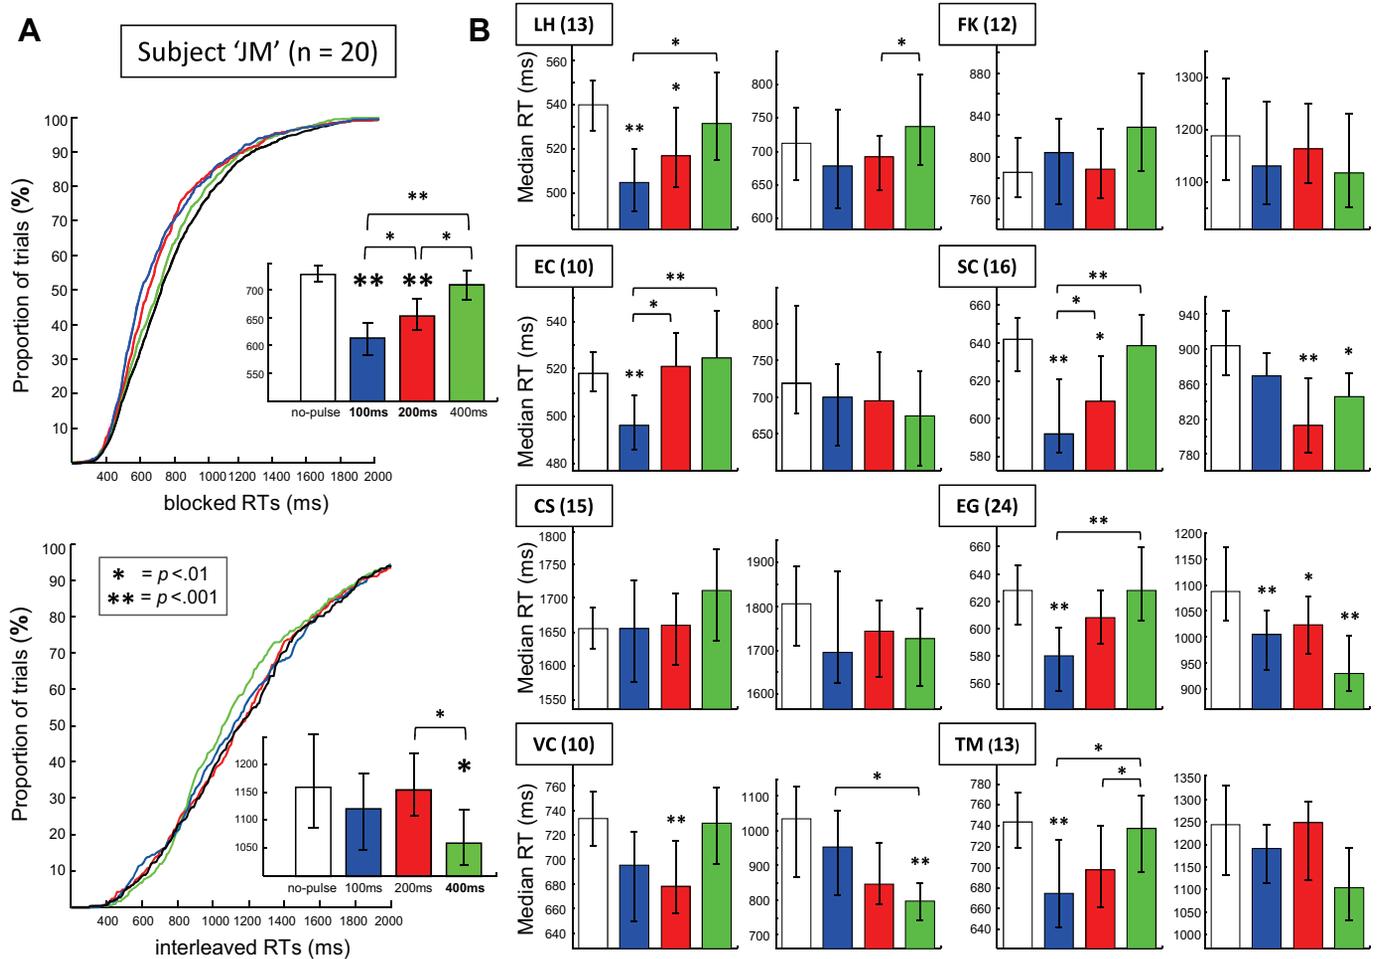


Fig. 5. Comparison of RT effects of pulses in CMD trials; individual subjects. *A*: cumulative RT distributions of *subject JM* during the blocked condition (*top*) and during the interleaved condition (*bottom*). Same format as in Fig. 4. *Subject JM*'s mean estimated non-decision time across all sessions was 451 ms. *B*: median RT data for all subjects presented in Fig. 3, *B* and *C*. Individual subjects are identified by initials and correspond to the data shown in Fig. 3, *B* (*left*) and *C* (*right*). Numbers in parentheses next to subject initials indicate the total number of experimental sessions completed by each subject.

the pulses had a significant effect on accuracy (100 ms: 926/1,120, 82.7%; 200 ms: 908/1,034, 87.8%; 400 ms: 996/1,132, 88.0%). These trends can be captured by both the DDM and UGM (data not shown), and therefore do not help to distinguish between the models.

Figure 7 shows that the accuracy of decisions decreases as RTs increase, for all conditions, both for individual subjects (*subject JM* shown) as well as for both the “fast” and “slow” subject subgroups. This may at first appear paradoxical, since longer viewing of a motion stimulus should allow one to better filter out noise and thus better estimate the underlying signal. An increasing accuracy with time is often reported in tasks in which observation time is externally controlled (Ratcliff and Rouder 1998, 2000; Ratcliff and Smith 2004). However, in our task the subjects are allowed to respond at any time, which means that the distribution of trials within each RT bin can be distorted by the distribution of trials in earlier bins. In such conditions, a decreasing performance for longer RTs can be easily explained if we suppose that a subject’s attention can vary between trials. If a subject happens to be more attentive on a given trial, then he/she will both respond more quickly and be more accurate than if he/she is less attentive. Thus the short RT trials will be biased toward those in which attention was high, while the long RT trials will be biased toward those

in which attention was lower. As a result, accuracy will tend to decrease for longer RT bins, and in theory this should hold for both models.

Modeling results. As shown in Fig. 8, both models correctly simulated the tendency for accuracy rates to decrease over time. This tendency simply results from variations in the “attentional gain” parameter *a*, included in both models, which effectively varies the signal-to-noise ratio from trial to trial. This causes differences in the distribution of otherwise identical trials in different RT bins, as described above. We added this parameter because the results shown in Fig. 7 may otherwise be taken as direct evidence for a dropping bound, and thus favor the UGM (which effectively implements a decreasing bound by including a growing evidence-independent urgency signal). However, that conclusion would not be accurate; both the UGM and the DDM can explain this phenomenon, so it cannot be used to distinguish between the models. The true distinction between the model predictions lies in analyses of the effects of pulses on RT distributions (see Fig. 1), to which we turn next. Importantly, all of the simulations described below have been performed both with and without the attentional gain parameter, yielding qualitatively identical results.

Figure 9 shows the RT distributions produced by a UGM with a filter time constant of 167 ms (see METHODS) and urgency

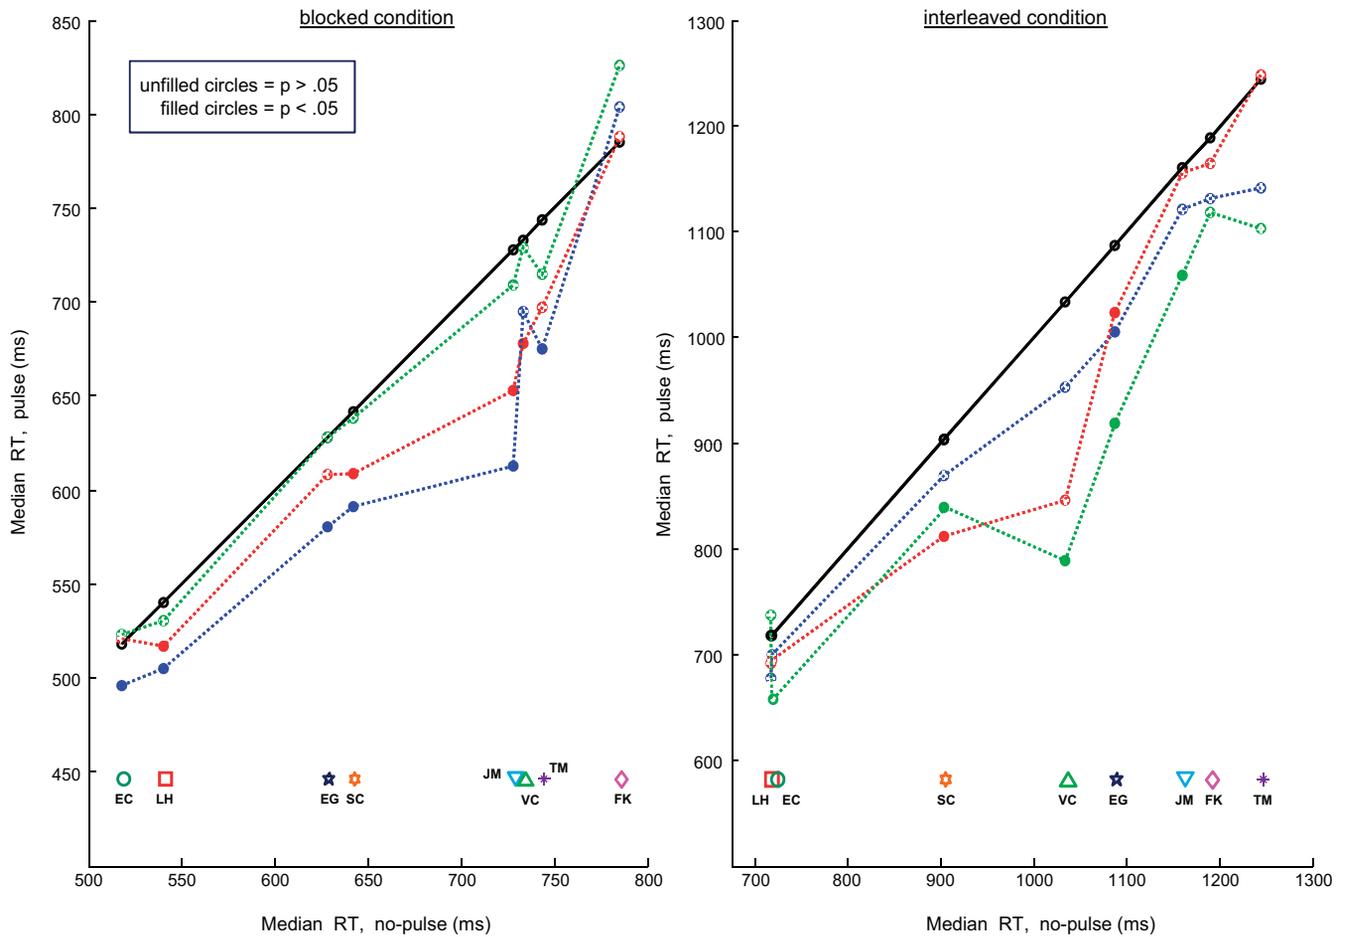


Fig. 6. Summary of effects of pulses on median RTs of 6 individual subjects for whom we collected >10 sessions of data. *Left*: data from the blocked condition. Each point compares an individual subject’s median RT in no-pulse trials (*x*-axis) against the median RT of that same subject during 100 ms (blue), 200 ms (red), and 400 ms (green) pulse trials (*y*-axis). The no-pulse trials (black) define the unity slope line, and colored points below that line indicate an effect of the given pulse. Filled circles indicate significance (K-S test, $P < 0.05$). *Right*: data from the interleaved condition. Same format as for blocked condition.

signal parameters that were condition dependent: for the “fast” subgroup, the slopes of the urgency signals were drawn from a log-normal distribution with $\mu = -0.50$ and $\sigma = 0.65$ for modeling the blocked condition and $\mu = -1.65$ and $\sigma = 0.7$ for modeling the interleaved sessions. These parameters were chosen so that the RT distributions produced by the model for no-pulse trials fit those of the fast subgroup (Fig. 4A). Next, the various pulses were added to the simulated input signal and their effects on RT determined by the same analyses performed on the real data.

As expected, the output of a model parameterized according to the data obtained from the “fast” subgroup resulted in the 100 and 200 ms pulses being effective during the blocked condition (Fig. 9A, *top*; K-S test $P = 7.2 \times 10^{-22}$, $P = 3.7 \times 10^{-4}$, respectively), while the 400 ms pulse was not ($P = 0.31$, K-S test). In contrast, during the interleaved condition, the 200 ms and 400 ms pulses were effective (Fig. 9A, *bottom*; K-S test $P = 3.7 \times 10^{-6}$, $P = 3.2 \times 10^{-6}$, respectively) while the 100 ms pulse was not ($P = 0.07$, K-S test), as expected (Fig. 1D).

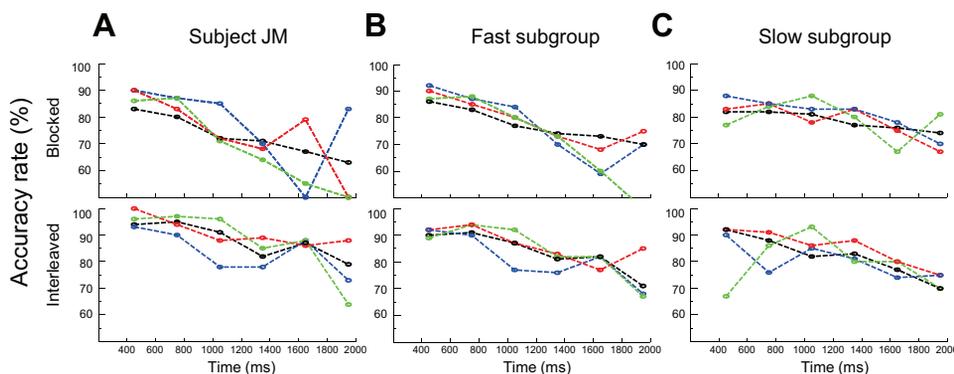


Fig. 7. Accuracy as a function of time. Trials were grouped into 300-ms bins on the basis of RT, and for each bin the accuracy was plotted separately for no-pulse trials (black) as well as trials with pulses at 100 ms (blue), 200 ms (red), and 400 ms (green). *Top*: results from blocked sessions. *Bottom*: results from interleaved sessions. *A*: data for subject JM. *B*: data for the fast subgroup. *C*: data for the slow subgroup.

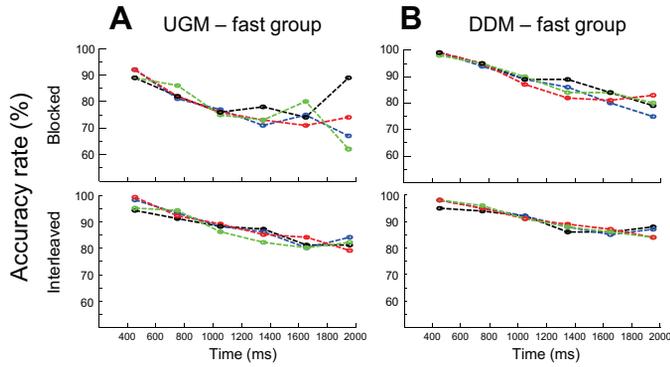


Fig. 8. Simulations of accuracy as a function of time, in the same format as Fig. 7. *A*: Simulations of the UGM fitted for the fast subgroup in the blocked (*top*) and interleaved (*bottom*) sessions. *B*: simulations of the DDM fitted for the fast subgroup in the blocked (*top*) and interleaved (*bottom*) sessions.

Data from the “slow” subgroup were modeled in a similar way (Fig. 9*B*), with urgency parameters $\mu = -0.41$, $\sigma = 0.30$ for the blocked condition and $\mu = -0.50$, $\sigma = 0.17$ for the interleaved condition. With these parameters, the UGM predicted that 400 ms pulses will have an effect in both conditions (blocked $P = 0.04$, interleaved $P = 0.01$, K-S test). While such a trend can be seen in the “slow” subgroup data from the blocked condition (Fig. 4*B*, *top*), it did not reach significance, perhaps because this subgroup contained only a few sessions of data for each subject.

Figure 10*A* shows simulations of the DDM with parameters adjusted to fit the “fast” subgroup. For the blocked condition, the best fit was provided by $T = 550$ and $\Theta = 16$, and these parameters simulated the major features of the data quite well—correctly producing an effect of 100 ms and 200 ms pulses (K-S test $P = 9.6 \times 10^{-39}$, $P = 2.2 \times 10^{-12}$, respectively) but not late pulses ($P = 0.07$, K-S test). For the interleaved condition, the best-fitting parameter settings were $T = 1,550$ and $\Theta = 35$, and they correctly predicted a significant effect of 200 ms and 400 ms pulses (K-S test $P = 1.9 \times 10^{-8}$, $P = 5.5 \times 10^{-4}$, respectively). However, the DDM also predicted that the 100 ms pulses had a significant effect ($P = 4.4 \times 10^{-11}$, K-S test), as is clear by examining the cumulative RT distributions. This prediction follows directly from the assumption of perfect integration, which is inherent in the DDM and which does not depend on parameter settings. For the same reason, the DDM predicts that for the “slow” subgroup of subjects all pulses always have significant effects in both blocked and interleaved conditions (Fig. 10*B*; $P < 10^{-4}$ for all pulses, K-S test). The best-fitting parameters for the slow subgroup were $T = 1,100$ and $\Theta = 25$ for the blocked condition and $T = 2,500$ and $\Theta = 40$ for the interleaved condition.

Delays in onset of integration. Because the motion signal during VMD trials can often reach conspicuously high values, but nonetheless cannot grow much in strength until several steps into the trial, it could be argued that during the inter-

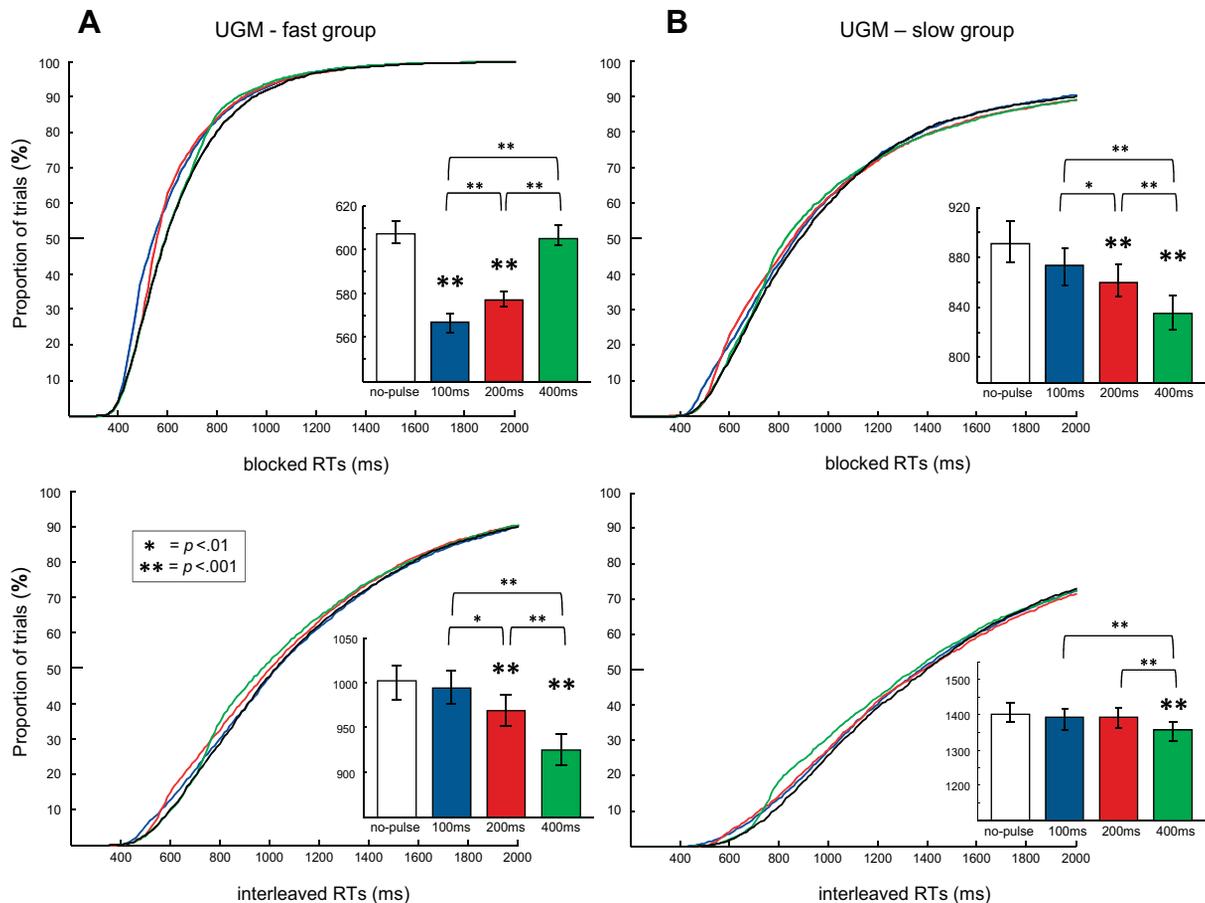


Fig. 9. Simulated cumulative RTs generated by the UGM, same format as Fig. 4. *A*: simulations of the fast subgroup in the blocked (*top*) and interleaved (*bottom*) sessions. *B*: simulations of the slow subgroup.

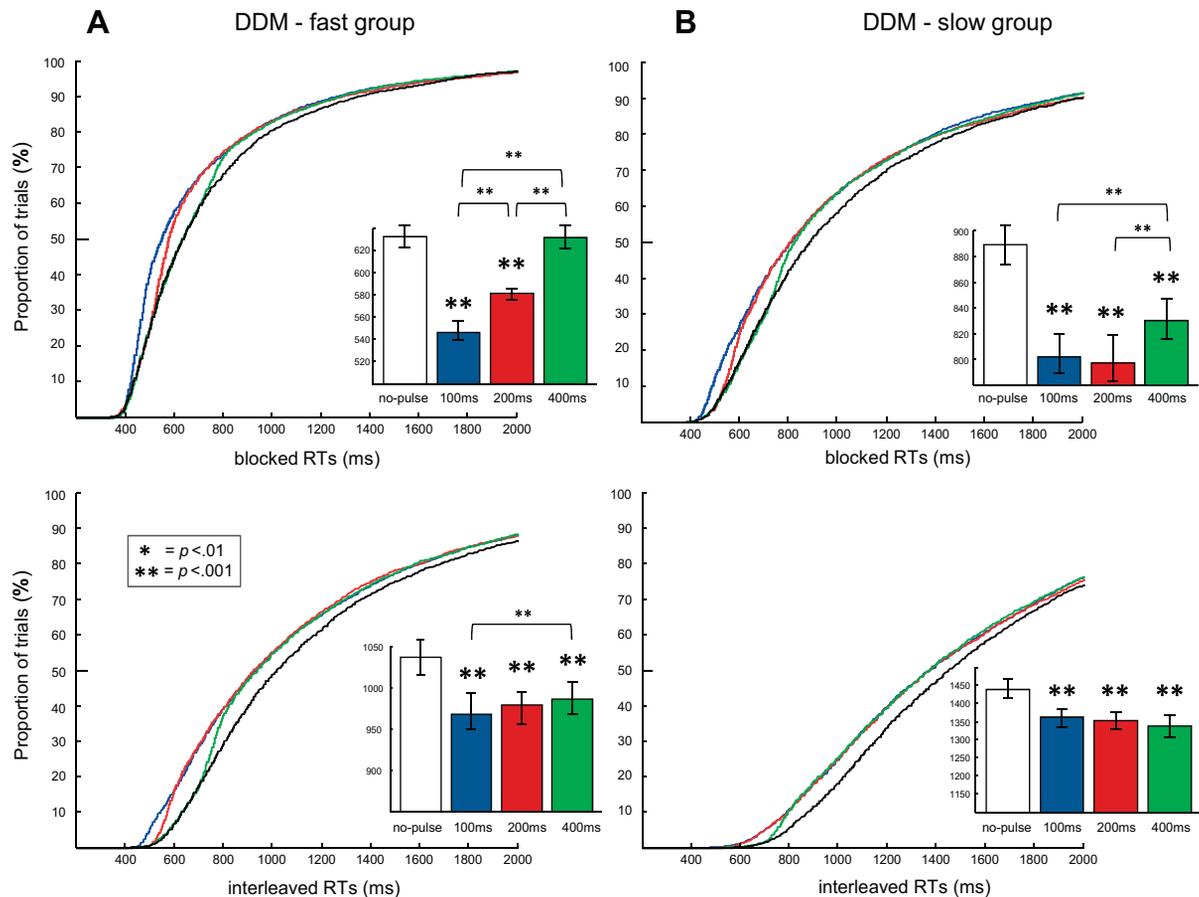


Fig. 10. Simulated cumulative RTs generated by the DDM, same format as Fig. 4. *A*: simulations of the fast subgroup in the blocked (*top*) and interleaved (*bottom*) sessions. *B*: simulations of the slow subgroup.

leaved sessions our subjects had simply learned to delay the onset of evidence integration for the first few hundred milliseconds so as to avoid having to discern weak early motion in favor of what could be assumed to be an easier discrimination later on. If this was true, it would cause them to show no effects of the earliest pulse timings (e.g., the 100 and 200 ms pulses) in the interleaved CMD trials, because these pulses would have already ended before the subjects began integrating evidence. Correspondingly, if the onset of integration was delayed by 200–300 ms, the 400 ms pulse would effectively become a 100 ms pulse, which would explain why such pulses suddenly became effective in the VMD sessions. This would explain not only the shifts in pulse efficacies obtained in the VMD condition but also the slower RT distributions for all VMD trials in general, providing an explanation for the data that would still be consistent with the DDM.

To test this conjecture, we added two more parameters to the DDM to allow a preintegration delay with a mean and standard deviation. For modeling the blocked sessions, we set both of these to zero because the DDM could fit that data without any preintegration delay. For modeling the interleaved session, we searched for the best-fitting setting of threshold and preintegration mean (with standard deviation set to 100 ms). With $T = 1,250$ and preintegration delay of 150 ± 100 ms, the DDM correctly simulated our finding that 100 ms pulses lost efficacy in the interleaved session, simply because most of them were ignored by the model.

However, for obvious reasons such a model could not produce decision times that are shorter than the preintegration delay. Only 0.7% of its decision times were made before 200 ms and only 9.1% before 400 ms. This contrasts with the data, in which we observed that 11.6% of decision times in no-pulse trials were shorter than 200 ms and 27.6% were shorter than 400 ms. Importantly, these early decisions were not merely random; they were correct 78.9% and 78.2% of the time, respectively, indicating that the early information was not in fact being ignored by our subjects.

In summary, the DDM could be parametrized either to capture the patterns of pulse efficacy in the interleaved trials or to produce early decisions comparable to those observed in the data, but we found no setting of parameters that could allow it to reproduce both of these findings. Nevertheless, it is possible that advocates of the DDM may find other ways to modify that model to explain our results. To that end, our data are available upon request.

DISCUSSION

The main result of our study is that when subjects slow down their decision policy the effect of early pulses becomes weaker while the effect of later pulses grows (Figs. 4–6). Although the specific pattern of pulse efficacy varied across subjects, the earliest pulse timings consistently lost their efficacy as subjects slowed down. This result is important because it cannot be reproduced by any model involving perfect integration of the

motion signal (Fig. 10). In particular, the DDM predicts that early pulses will always be at least as effective as late pulses in reducing RTs, because a perfect integrator retains all input until decision time regardless of the threshold setting. To explain our results, a strong leak would have to be added to the DDM, effectively turning it into a low-pass filter with a relatively short time constant. Our previous analyses of behavioral and neural data in changing-evidence tasks suggest that the time constant is between 100 and 250 ms (Cisek et al. 2009; Thura et al. 2012; Thura and Cisek 2014).

There may be many reasons why subjects slow down during the interleaved sessions. Above, we proposed that they do so because in the VMD trials higher motion strengths are reached later in time, and thus many individual trials tend to get easier with longer viewing. It is also possible that subjects slow down because sometimes the net direction of motion might change just after they make their decision, reducing their overall confidence. While such reversals occurred in <3% of trials, the very possibility of reversals might motivate subjects to be more conservative. There may still be other reasons. Regardless of why subjects slow down, different models suggest different mechanisms for how they do so. The DDM assumes that there is a change in the threshold, while the UGM assumes that there is a change in urgency. Other mechanisms may also be proposed, such as a reduction in the strength of recurrent feedback in attractor-type models (e.g., Wang 2002). What is important to emphasize, however, is that our present results and their interpretations are not dependent upon any of these issues. Regardless of why subjects slow down and regardless even of how they do so, it is the fact that they slow down that is most pertinent to the question of the time constant. If slowing down causes early pulses to lose their efficacy, then the time constant must be short.

Possibility of a “two-model” solution. One possible objection to the above interpretation, however, is that our subjects may have employed radically different strategies during each of our two session types. For example, perhaps they used a DDM in our blocked sessions (and in previous studies) and a UGM in our interleaved sessions. Because none of the trials presented in the blocked sessions ever exceeded 3% baseline motion coherence, all trials in the blocked condition can be considered fairly difficult, which may therefore have motivated our subjects to integrate motion evidence with a long time constant. By contrast, VMD trials presented during the interleaved sessions could often reach much larger values of coherence, and may therefore have motivated subjects not to integrate motion evidence across time but rather to simply wait for the motion signal to become highly conspicuous before deciding.

However, there are a number of reasons to doubt this account. First, in both sessions subjects were always given the same instruction: to respond as soon as they were able to detect any motion signal at all. This is the same instruction typically given to subjects in a large variety of studies using the random-dot task in conditions similar to our blocked sessions. Thus the only feature of the task that may have motivated subjects to change their strategy in interleaved sessions is the reinforcement provided regarding correct or incorrect choices. If our subjects’ default strategy was to use a perfect integration, the only reason they would switch to an alternative, short-time-constant strategy would be if they were penalized often enough

for choosing the direction indicated by the total net motion but contraindicated by the current motion at the time of their decision. Such cases—if they occurred often enough to be noticeable—might motivate subjects to begin to discard past evidence, and to instead weight recent evidence more heavily.

To examine this possibility, we compared overall accuracy rates that would ensue from a post hoc reclassification of “correct” and “incorrect” choices according to the total cumulative motion evidence, as opposed to our original classification based on the sign of the motion evidence at the estimated time of decision. This revealed that out of all VMD trials completed by our subjects, only 9% would have been incorrectly reinforced according to a cumulative-motion-evidence criterion, and that this discrepancy falls to only 3% when considering VMD trials in which the motion signal ever reached or exceeded 15% coherence, i.e., the trials that would have conspicuously stood out to the subjects. Consequently, it seems unlikely that the reinforcement provided to our subjects was significant enough to be responsible for the changes in decision strategy we observed across the session types.

Furthermore, the logic of this proposal can be reversed. In natural behavior, sensory evidence does change, and changes often. To react quickly to such changes, leaky integration is much more effective than perfect integration, which first needs to “undo” previously accumulated evidence before moving toward the new decision bound. Thus reinforcement in the real world would seem to favor mechanisms with short time constants (as long as the time constant is long enough to filter out noise) as the default strategy. This then raises the question of why humans would ever change their strategy and use perfect integration, even when performing a standard constant-motion discrimination task. In such tasks the motion does not change, so the reinforcement obtained by a perfect integrator would be no different than that obtained by a leaky model.

Finally, while a two-model solution is plausible, it lacks parsimony. It proposes two different models for explaining two different kinds of data and necessitates additional mechanisms for arbitrating between them. From a modeling perspective, it is certainly possible that such a hybrid, “switching” model could produce better quantitative fits to our data; however, by the same token, such a model would necessarily include a larger total number of parameters than either model by itself. Any measure of goodness of fit to any given data set would be penalized by the additional parameters of the switching mechanism as well as those of the “unused” model, and therefore rank low according to traditional model-comparison measures (e.g., AIC/BIC criteria, etc.). In contrast, the UGM can explain all of the data with just a single parameter change—a modification of urgency—that is motivated by optimization of reward rates. In a recent paper (Thura and Cisek 2014), we directly demonstrated that neural activity in dorsal premotor and primary motor cortex combines both urgency and evidence-related components, and that the latter is processed with a short time constant. While this used a different species and a different task, we propose that the conclusions generalize to a broad range of tasks and species. Indeed, although we acknowledge that perfect integration is commonly assumed when interpreting data, we are not aware of data that conclusively prove that assumption to be correct.

Narrowing down the value of the time constant. The present results could, in principle, be explained by the leaky competing

accumulator (LCA) model (Ossmy et al. 2013; Tsetsos et al. 2012; Usher and McClelland 2001), as long as the leak parameter is set high enough to produce a short time constant. With such a high leak, however, the evidence signal in the LCA model will equilibrate long before it reaches any appreciable bound, and until that point the temporal shape of the buildup will be a saturating exponential, not the linear buildup usually observed in neural data (cf. Roitman and Shadlen 2002). Consequently, something else is needed to explain the prolonged linear time course of neural activity growth and long decision times. We and others have proposed that this growth is at least partially caused by an urgency signal that pushes neural activity toward the threshold even in the absence of new evidence, effectively implementing a decreasing accuracy bound (Churchland et al. 2008; Cisek et al. 2009; Ditterich 2006; Thura et al. 2012). In fact, the UGM is equivalent to the LCA model with a high leak parameter and an urgency signal. Such a model can explain not only the present data but also the lack of biasing effects observed by Thura et al. (2012), suggesting that the effects observed in that study are not task dependent but rather indicative of a general strategy of low-pass filtering a noisy stimulus signal.

A recent analysis using the LCA (Ossmy et al. 2013) suggested that time constants can change between conditions with different distributions of signal durations. Across subjects, the best-fitting time constant values varied from 29 ms to 1,995 ms (mean 78 ms and 493 ms) in the two conditions (see their Table S2). However, because the LCA assumes a fixed bound, it forces the use of a long time constant for explaining any late decisions. We believe that if that assumption was relaxed, and the bound was allowed to decrease through the use of a context-dependent urgency signal, then long decisions could be explained even if the time constant was always short.

While a precise estimate of the time constant is difficult to make without direct neural data, the present results are nonetheless useful for providing bounds on the possible values that it may take. In particular, while results similar to Fig. 9 could be generated with a 200-ms time constant, a 250-ms time constant consistently predicts a significant effect of the 100 ms pulse even in the interleaved sessions, in contrast to our data. In principle, one could use a model fitting optimization procedure to try to find the best-fitting time constant, but in our view that would not be fruitful in practice. The reason is that once the time constant is set below 250 ms small variations of it can be traded off against changes in intertrial variability parameters (e.g., the σ parameter of the UGM), producing similar goodness of fit. Our approach is not to try to estimate these parameters precisely but rather to eliminate sections of the parameter space. The results of the present study as well as previous studies (Cisek et al. 2009; Thura et al. 2012; Thura and Cisek 2014) allow us to eliminate models with long time constants (>250 ms), including perfect integrators such as the DDM.

Alternative (non-integration-based) models. It is possible that during the interleaved sessions subjects neither integrate the motion signal nor combine it with a growing urgency but simply wait until a large coherent motion appears randomly during a trial. However, we consider it unlikely that this strategy—akin to a threshold detection process—was used in our task. Because our subjects could not predict ahead of time what the range of motion coherence would be on a given trial,

there was no way for them to set any particular decision threshold. For example, simply waiting until the motion coherence exceeded 15% would have failed to produce any decision at all on ~40% of VMD trials, whereas the number of “time out” trials we actually observed was <1%. Indeed, any threshold much higher than 3% would have failed to produce any decisions on the no-pulse CMD trials. Conversely, setting the threshold lower would have produced many more errors than what we observed. Thus, while a detection strategy is plausible and may be used in some situations, it seems unlikely to explain our data.

The UGM suggests that low-pass-filtered evidence is brought to threshold through combination with an independent urgency signal that controls the decision timing policy. After a pulse, activity related to evidence increases briefly but then quickly returns to the baseline “no-pulse” level. Thus a motion pulse reduces the RT on a given trial only if the decision is made before the effect of the pulse has leaked away (see Fig. 1, *C* and *D*). The UGM thus predicts that the efficacy of a given motion pulse will depend on its timing with respect to the decision time. By extension, any systematic change in average decision times will change which pulse timings reduce RTs. This is precisely what we observed. Moreover, this explains not only the general trends we found but also many of the differences obtained across subjects, with different individuals susceptible to a different set of motion pulses based on the average timing of their decisions in a given context (Fig. 6).

An alternative way to implement a buildup of neural activity is through recurrent feedback between the populations representing the different options, as in the model of Wang (2002). While such models differ from the DDM in many important ways, their recurrent feedback effectively implements a long time constant of integration. Indeed, simulations by Wang (2002) show that reduction in the strength of the recurrent connections causes the system to lose its slow buildup of neural activity and equilibrate within 200–300 ms after stimulus presentation. If the recurrent connections are strong enough to produce continued buildup, then they effectively implement temporal integration with almost no leak. Consequently, we conjecture that such models, like the DDM, would also predict a persistent effect of the 100 ms pulses even in the interleaved condition, in contrast to our data. However, variations of attractor models in which the interaction function governing the recurrent connections becomes steeper over time (see Standage et al. 2011) can produce dynamics that approximate a leaky integrator with urgency, and consequently are in principle able to reproduce our findings (Marcos et al. 2012). Indeed, such models were the original inspiration for our experiment, though their detailed implementation is beyond the scope of the present report.

Noise. Integrator models deal with noise by adding together successive samples on the assumption that the noise components will cancel each other out over time, leaving an estimate of the underlying signal. On the surface, the UGM may appear susceptible to noise because it privileges recent information. However, because the UGM includes a low-pass filter, it is just as effective as an integrator at dealing with noise. Indeed, for stimulus components above the filter cutoff frequency, a low-pass filter and an integrator are approximately equivalent. The time constant of the filter jointly determines both the frequency of input fluctuations that will be screened out as noise as well

as the amount of time required for the decision variable to respond to a genuine change in the underlying signal. This issue is relevant for two recent papers that attempted to argue against the UGM. Winkel et al. (2014) showed that early motion signals influence RTs and took this as evidence against the UGM. However, in their implementation of the UGM they did not include the low-pass filter. If a low-pass filter is added to the model, it can easily reproduce those data (Carland et al. 2015). Indeed, as shown here (Fig. 9), a low-pass filter with a short time constant is capable of simulating the effects of early motion signals, even on RT distributions whose means are $\sim 1,000$ ms. Churchland et al. (2011) suggested that the UGM cannot explain correlations in neural activity, but those authors also did not include the low-pass filter. If the UGM is implemented fully, including the low-pass filter, then analyses of high-frequency correlations cannot distinguish between it and an integrator model, because a low-pass filter and an integrator are equivalent with respect to reducing the gain of high-frequency components in their input.

Definition of “evidence” and its ramifications for modeling the decision process. Important to the present discussion is the question of what constitutes evidence in a given task. The DDM is often seen as equivalent to the sequential probability ratio test (SPRT), a statistical test that optimizes sampling time to attain a given level of accuracy (Wald 1945). However, the equivalence between the DDM and the SPRT holds only under the assumption that each sensory sample is statistically independent from preceding ones (Bogacz et al. 2006; Huang and Rao 2013; Thura et al. 2012). In any constant-evidence task, repeatedly sampling the stimulus means that each additional sample is increasingly redundant, providing less and less new information. Ideally, the decision process should be governed primarily by novel information. Thura et al. (2012) suggested that for simple tasks this may be approximated by a low-pass filter, which quickly adjusts the current estimate of evidence while ignoring fluctuations whose frequency is above the range at which the signal of interest is likely to change. This mechanism not only ensures that redundant information is ignored but also enables faster transitions between options under conditions in which evidence can change. A perfect integrator would be slow to reflect such changes, as it would have to first “undo” the previously integrated sum for the initial choice before it could begin to accumulate evidence in favor of the new choice. In contrast, a UGM could respond to the new choice after only a brief delay determined by its time constant.

In light of the above, we can consider some of the specific tasks that have been described in the literature. In some tasks subjects are being given new sensory information over time—this includes the “weather prediction” task (Kira et al. 2015; Yang and Shadlen 2007), the “tokens” task (Cisek et al. 2009; Thura et al. 2014; Thura and Cisek 2014), the VMD task (Thura et al. 2012), and the “click-counting” task of Brody and colleagues (Brunton et al. 2014; Hanks et al. 2015). In such conditions, information from each new piece of information is novel and should indeed be accumulated, as supported by both behavioral and neural data. In contrast, in some tasks the stimulus is completely static—this includes “brightness discrimination” (Ratcliff 2002; Ratcliff et al. 2007), “dot separation” (Ratcliff et al. 2003), and “color discrimination” (Coallier and Kalaska 2014) tasks. In such conditions, each sample of the visual stimulus is identical to previous samples and pro-

vides no new information, so it should not be integrated. The random-dot motion discrimination task (Britten et al. 1992) as well as noisy image categorization tasks (e.g., Tremel and Wheeler 2015) are a variant of the latter category—they are tasks in which the underlying signal is static and all that changes between sequential samples is noise. In such conditions, each sequential sample provides some novel information but with diminishing returns, because once the noise has been filtered out there is no more new information to be obtained. Thus, in such tasks, an accurate estimate of evidence should stop growing quickly and therefore could not be responsible for the prolonged buildup of neural activity.

Conclusions. Perfect integration models such as the DDM have a long history and have become well accepted because they capture a great deal of data within a simple and intuitively appealing framework. However, the intuition of accumulating sensory samples must be supplemented by considering the evidence actually provided by those samples, which are increasingly redundant in most tasks. Strictly speaking, only novel evidence should be accumulated. Furthermore, in nearly all previous experiments the sensory information provided to subjects was constant over the course of each trial. In such conditions it is difficult to distinguish between different parameter settings, and even sophisticated comparisons of model fits (Hawkins et al. 2015a) do not yield consistent conclusions across tasks, subjects, or modeling assumptions. In contrast, tasks in which evidence changes during each trial cause different models to make clearly divergent predictions that are empirically distinguishable. Our previous analyses of behavioral and neural data in such tasks strongly favored the UGM (Cisek et al. 2009; Thura et al. 2012; Thura and Cisek 2014), but those results could have been task dependent.

Here we used the same random-dot motion discrimination task that has been used many times to support the DDM. However, we designed our experiment to directly test situations in which the two models make empirically distinguishable predictions. We did this both by blocking trials into two distinct contexts that implicitly motivate either faster or slower responses and by presenting in each context some identical trials with brief motion pulses. The finding that early pulses become less effective as subjects slow down their decisions (Figs. 4–6) cannot be explained with any version of a perfect integrator (Fig. 10), regardless of parameter settings, and points to the necessity of including a strong leak term. The consequence of the leak term, however, is that with constant evidence the system will equilibrate quickly and therefore cannot explain long RTs or the prolonged linear buildup of neural activity. This suggests the necessity of including an urgency signal. Neural data supporting the presence of an urgency signal at the individual cell level have been found during constant-evidence tasks (Churchland et al. 2008; Hanks et al. 2014) as well as during changing-evidence tasks (Gluth et al. 2012; Kira et al. 2015; Thura and Cisek 2014). We conjecture that an urgency signal is a general mechanism for controlling speed-accuracy trade-offs in decision making and movement control (Thura et al. 2014) and that it is responsible for much of the neural activity buildup observed during decision making.

GRANTS

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada (Discovery Grant 327654), the Canadian Institutes of Health Research (MOP-102662), and the European Commission through the project Goal-Leaders (FP7-ICT-270108) and an infrastructure grant from the Fonds de la Recherche en Santé du Québec.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

E.M. conceived the experimental paradigm; M.A.C. collected and analyzed data and performed model simulations; all authors interpreted the results and wrote the manuscript.

REFERENCES

- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD.** The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113: 700–765, 2006.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA.** The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12: 4745–4765, 1992.
- Brunton BW, Botvinick MM, Brody CD.** Rats and humans can optimally accumulate evidence for decision-making. *Science* 340: 95–98, 2013.
- Busemeyer JR, Townsend JT.** Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 100: 432–459, 1993.
- Carland M, Marcos E, Thura D, Verschure P, Cisek P.** Decision-making is influenced by a context-dependent urgency signal (Abstract). *Neuroscience Meeting Planner* 2013: Program No. 286.11.2013, 2013.
- Carland MA, Thura D, Cisek P.** The urgency-gating model can explain the effects of early evidence. *Psychon Bull Rev* 22: 1830–1838, 2015.
- Chittka L, Skorupski P, Raine NE.** Speed-accuracy tradeoffs in animal decision making. *Trends Ecol Evol* 24: 400–407, 2009.
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN.** Variance as a signature of neural computations during decision making. *Neuron* 69: 818–831, 2011.
- Churchland AK, Kiani R, Shadlen MN.** Decision-making with multiple alternatives. *Nat Neurosci* 11: 693–702, 2008.
- Cisek P, Puskas GA, El-Murr S.** Decisions in changing conditions: the urgency-gating model. *J Neurosci* 29: 11560–11571, 2009.
- Coallier E, Kalaska JF.** Reach target selection in humans using ambiguous decision cues containing variable amounts of conflicting sensory evidence supporting each target choice. *J Neurophysiol* 112: 2916–2938, 2014.
- Cook EP, Maunsell JH.** Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci* 5: 985–994, 2002.
- Ditterich J.** Evidence for time-variant decision making. *Eur J Neurosci* 24: 3628–3641, 2006.
- Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A.** The cost of accumulating evidence in perceptual decision making. *J Neurosci* 32: 3612–3628, 2012.
- Gluth S, Reiskamp J, Büchel C.** Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *J Neurosci* 32: 10686–10698, 2012.
- Gold JI, Shadlen MN.** Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.
- Gold JI, Shadlen MN.** The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23: 632–651, 2003.
- Gold JI, Shadlen MN.** The neural basis of decision making. *Annu Rev Neurosci* 30: 535–574, 2007.
- Hanks T, Kiani R, Shadlen MN.** A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife* 3: e02260, 2014.
- Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD.** Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520: 220–223, 2015.
- Hawkins GE, Forstmann BU, Wagenmakers EJ, Ratcliff R, Brown SD.** Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *J Neurosci* 35: 2476–2484, 2015a.
- Hawkins GE, Wagenmakers EJ, Ratcliff R, Brown SD.** Discriminating evidence accumulation from urgency signals in speeded decision making. *J Neurophysiol* 114: 40–47, 2015b.
- Huang Y, Rao RP.** Reward optimization in the primate brain: a probabilistic model of decision making under uncertainty. *PLoS One* 8: e53344, 2013.
- Huk AC, Shadlen MN.** Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25: 10420–10436, 2005.
- Kiani R, Hanks TD, Shadlen MN.** Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28: 3017–3029, 2008.
- Kim JN, Shadlen MN.** Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1999.
- Kira S, Yang T, Shadlen MN.** A neural implementation of Wald's sequential probability ratio test. *Neuron* 85: 861–873, 2015.
- Laming DR.** *Information Theory of Choice-Reaction Times*. New York: Wiley, 1968.
- Luna R, Hernandez A, Brody CD, Romo R.** Neural codes for perceptual discrimination in primary somatosensory cortex. *Nat Neurosci* 8: 1210–1219, 2005.
- Marcos E, Carland M, Thura D, Cisek P, Verschure PF.** Decision-making depends on an urgency signal modulated by context (Abstract). *Neuroscience Meeting Planner* 2012: Program No. 494.11.2012, 2012.
- Newsome WT, Britten KH, Movshon JA.** Neuronal correlates of a perceptual decision. *Nature* 341: 52–54, 1989.
- Ossmy O, Moran R, Pfeffer T, Tsetsos K, Usher M, Donner TH.** The timescale of perceptual evidence integration can be adapted to the environment. *Curr Biol* 23: 981–986, 2013.
- Palmer J, Huk AC, Shadlen MN.** The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5: 376–404, 2005.
- Ratcliff R.** A theory of memory retrieval. *Psychol Rev* 83: 59–108, 1978.
- Ratcliff R.** A diffusion model account of response time and accuracy in a brightness discrimination task: fitting real data and failing to fit fake but plausible data. *Psychon Bull Rev* 9: 278–291, 2002.
- Ratcliff R, Cherian A, Segraves M.** A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol* 90: 1392–1407, 2003.
- Ratcliff R, Hasegawa YT, Hasegawa YP, Smith PL, Segraves MA.** Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J Neurophysiol* 97: 1756–1774, 2007.
- Ratcliff R, McKoon G.** The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* 20: 873–922, 2008.
- Ratcliff R, Rouder JN.** Modeling response times for two-choice decisions. *Psychol Sci* 9: 347–356, 1998.
- Ratcliff R, Rouder JN.** A diffusion model account of masking in two-choice letter identification. *J Exp Psychol Hum Percept Perform* 26: 127–140, 2000.
- Ratcliff R, Smith PL.** A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111: 333–367, 2004.
- Reddi BA, Carpenter RH.** The influence of urgency on decision time. *Nat Neurosci* 3: 827–830, 2000.
- Roe RM, Busemeyer JR, Townsend JT.** Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol Rev* 108: 370–392, 2001.
- Roitman JD, Shadlen MN.** Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22: 9475–9489, 2002.
- Standage D, You H, Wang DH, Dorris MC.** Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Front Comput Neurosci* 5: 7, 2011.
- Thura D, Beauregard-Racine J, Fradet CW, Cisek P.** Decision making by urgency gating: theory and experimental support. *J Neurophysiol* 108: 2912–2930, 2012.
- Thura D, Cisek P.** Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81: 1401–1416, 2014.
- Thura D, Cos I, Trung J, Cisek P.** Context-dependent urgency influences speed-accuracy trade-offs in decision-making and movement execution. *J Neurosci* 34: 16442–16454, 2014.
- Tremel JJ, Wheeler ME.** Content-specific evidence accumulation in inferior temporal cortex during perceptual decision-making. *Neuroimage* 109: 35–49, 2015.

- Tsetsos K, Gao J, McClelland JL, Usher M.** Using time-varying evidence to test models of decision dynamics: bounded diffusion vs. the leaky competing accumulator model. *Front Neurosci* 6: 79, 2012.
- Uchida N, Kepecs A, Mainen ZF.** Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat Rev Neurosci* 7: 485–491, 2006.
- Usher M, McClelland JL.** The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* 108: 550–592, 2001.
- Wald A.** Sequential tests of statistical hypotheses. *Ann Math Stat* 16: 117–186, 1945.
- Wang XJ.** Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36: 955–968, 2002.
- Winkel J, Keuken MC, van Maanen L, Wagenmakers EJ, Forstmann BU.** Early evidence affects later decisions: why evidence accumulation is required to explain response time data. *Psychon Bull Rev* 21: 777–784, 2014.
- Wong KF, Huk AC, Shadlen MN, Wang XJ.** Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci* 1: 6, 2007.
- Yang T, Shadlen MN.** Probabilistic reasoning by neurons. *Nature* 447: 1075–1080, 2007.
- Yang Y, DeWeese MR, Otazu GH, Zador AM.** Millisecond-scale differences in neural activity in auditory cortex can drive decisions. *Nat Neurosci* 11: 1262–1263, 2008.

