

Response times from ensembles of accumulators

Bram Zandbelt¹, Braden A. Purcell¹, Thomas J. Palmeri¹, Gordon D. Logan¹, and Jeffrey D. Schall¹

Center for Integrative and Cognitive Neuroscience, Vanderbilt Vision Research Center, Department of Psychology, Vanderbilt University, Nashville, TN 37240

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved January 8, 2014 (received for review June 5, 2013)

Decision-making is explained by psychologists through stochastic accumulator models and by neurophysiologists through the activity of neurons believed to instantiate these models. We investigated an overlooked scaling problem: How does a response time (RT) that can be explained by a single model accumulator arise from numerous, redundant accumulator neurons, each of which individually appears to explain the variability of RT? We explored this scaling problem by developing a unique ensemble model of RT, called *e pluribus unum*, which embodies the well-known dictum “out of many, one.” We used the *e pluribus unum* model to analyze the RTs produced by ensembles of redundant, idiosyncratic stochastic accumulators under various termination mechanisms and accumulation rate correlations in computer simulations of ensembles of varying size. We found that predicted RT distributions are largely invariant to ensemble size if the accumulators share at least modestly correlated accumulation rates and RT is not governed by the most extreme accumulators. Under these regimes the termination times of individual accumulators was predictive of ensemble RT. We also found that the threshold measured on individual accumulators, corresponding to the firing rate of neurons measured at RT, can be invariant with RT but is equivalent to the specified model threshold only when the rate correlation is very high.

computational model | mathematical psychology | diffusion model | reaction time | neurophysiology

Response time (RT) is a core measure of human decision-making in experimental psychology (1). The random variation of RT across otherwise identical trials has been a puzzle since the mid-19th century. Since the 1960s, this variation of RT—measured in a wide range of perceptual, cognitive, and economic tasks (1–5)—has been explained through stochastic accumulator models. These models assume that a response is generated when evidence accumulates at a certain rate (v) over time to a threshold (θ) and that the stochastic variation of RTs arises primarily from random fluctuations in accumulation rates (Fig. 1A). Historically, these models were formulated and tested before data on the underlying neural processes were available.

Subsequently, neurons exhibiting accumulating discharge rates in various RT tasks have been found in sensory, sensorimotor, and motor brain structures; in premotor circuits for limb and eye movements it is known that the neurons with accumulating activity are necessary and sufficient for initiating movements (6, 7). Movements are initiated when the trial-averaged accumulating spike rate of these neurons reaches a fixed activation level (6) (A_{RT}) like a threshold, and the distribution of RTs is accounted for by the stochastic variability in the rate of growth of neural activity toward A_{RT} (Fig. 1B). This discovery inspired the conjecture that individual neurons instantiate the evidence accumulation process described by stochastic accumulator models (6). This conjecture has stimulated extensive research replicating the original observation and equating accumulator model parameters with measures of neural dynamics assessed by spike rates (8–17), EEG (18, 19), magnetoencephalography (MEG) (20), and functional MRI (21–23) and simulated with neural network models (24–28).

However, this productive line of research has overlooked a fundamental scaling problem. On the one hand, the behavior of specific single neurons seems sufficient to account for the RT

of the whole brain. On the other hand, we know that ensembles of tens of thousands of neurons are necessary to produce any action (*SI Text, How Many Neurons Produce a Movement?*). Hence, how can each individual accumulator neuron, recorded in isolation, seem sufficient to initiate a movement by crossing a unique threshold when no single accumulator neuron is necessary for a movement to occur? In other words, how is the accumulating activity of numerous redundant and idiosyncratic neurons in a large ensemble coordinated and combined to produce variable RTs that can be predicted by a model consisting of just a single stochastic accumulator? This question has not been addressed previously (*SI Text, Extension of Previous Work*).

This question is challenging to investigate empirically because the limited number of spikes emitted by individual neurons precludes reliable assessment of single-trial dynamics, and simultaneous measurement of numerous functionally homogeneous neurons is not possible with current technology. Therefore, we performed computer simulations of ensembles of stochastic accumulators.

We address four major issues. First, we investigate how RT distributions can be explained both by a single accumulator model and by the ensemble activity of many accumulators. Second, we explore how RT distributions scale with the accumulator ensemble size. Third, we investigate how the A_{RT} measured across trials from an individual accumulator can be invariant with RT even though RT is produced by a large ensemble of accumulators with different growth rates. Fourth, we explore how the measured A_{RT} from an individual accumulator relates to the actual threshold of that accumulator (θ).

To address these issues, we developed a unique ensemble model of RT, called *e pluribus unum* (EPU), which embodies the well-known dictum “out of many, one.” Stochastic accumulator models are typically designed to explain both RT and accuracy obtained in choice tasks. However, our questions are specifically centered on the basic variability of RT that is observed in responses in any task. Thus, this model does not address accuracy, although we envision natural extensions of this approach to racing or competing ensembles of accumulators embodied by simple differential equations or in more complex spiking network models.

Significance

The delay of responding to stimuli, known as response time (RT), is randomly variable. Psychologists explain this variability through models in which RT is dictated by the termination of a single random accumulation process. Neurophysiologists explain this variability through the dynamics of neurons sampled from very large networks. This paper explains how these radically different scales of explanation can both be correct.

Author contributions: B.Z., B.A.P., T.J.P., G.D.L., and J.D.S. designed research; B.Z. performed research; B.Z. and T.J.P. contributed new reagents/analytic tools; B.Z. analyzed data; and B.Z., B.A.P., T.J.P., G.D.L., and J.D.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence may be addressed. E-mail: thomas.j.palmeri@vanderbilt.edu, bram.zandbelt@vanderbilt.edu, braden@nyu.edu, gordon.logan@vanderbilt.edu, or jeffrey.d.schall@vanderbilt.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1310577111/-DCSupplemental.




Fig. 1. Response times predicted by ensembles of redundant stochastic accumulators. (A) Stochastic accumulator models describe RT in terms of an accumulation process (one trajectory per trial) that proceeds at a certain rate (v) to reach a fixed threshold (θ). Stochastic variation of RT arises from fluctuations of v between (η) and within trials (ξ). It is common to consider one accumulator associated with each of multiple responses; we considered instead the case of multiple accumulators associated with the same response (*Inset*). (B) RT can also be described by the time at which the evolving spike rates of certain neurons, averaged across bins of trials with common RTs (one trajectory per RT bin, replotted from ref. 49), reach an activation level that is invariant with RT (A_{RT}). These neurons have been argued to instantiate the process described by stochastic accumulator models. (C) Unless accumulators are perfectly correlated (*Inset*), it is unclear (i) how an ensemble of accumulators makes the transition from evidence accumulation to response execution, (ii) under what termination rules (p_N) and accumulation rate correlations (r_v) the dynamics of one accumulator (highlighted red) predicts RT distributions and the invariant relationship between A_{RT} and RT, as observed empirically, and (iii) how A_{RT} relates to the unobserved threshold of an accumulator (θ).

In single-accumulator models, RT critically depends on two key parameters: the accumulation rate (v) and the threshold (θ). Extrapolating these parameters to the ensemble case is not trivial (Fig. 1C).

First, how are accumulation rates coordinated across the ensemble? At one extreme, if all accumulators share identical dynamics, then the ensemble reduces to one accumulator (Fig. 1C, *Inset*), yet perfect correlation is implausible (29). At the other extreme, if all accumulators have uncorrelated dynamics, then unrealistic RT variability would occur. Moreover, uncorrelated dynamics would also be implausible from a biological perspective, given that ensembles receive common inputs, have recurrent connections, and are modulated by common neurotransmitter systems. We investigated this question by sampling correlated accumulation rates, with rate correlation (r_v) varying between 0.0 and 1.0. Though the range of rate correlations we simulated exceeds the noise correlation found among neighboring neurons (30, 31), they can arise naturally from redundancy in common inputs, recurrent connectivity, and modulation by a common source (32–34).

Second, how is ensemble activity combined to produce one RT? At one extreme, if RT is specified by the time when the fastest accumulator reaches threshold, the RT distribution will shrink with ensemble size. At the other extreme, if RT is specified by the time when the slowest accumulator reaches threshold, the RT distribution will expand with ensemble size. How large is the region between these two extremes where the RT distribution remains stable with ensemble size? We investigated these questions by assuming that each accumulator projects to a unit that tallies the proportion of accumulators having crossed a threshold activation (a “polling” mechanism akin to quorum sensing) (35) or monitors the average firing rate of the ensemble (a “pooling” mechanism akin to the vector averaging that guides movement dynamics in final common neural circuits that initiate movements) (36, 37). When this unit tallies a critical proportion of units hitting threshold (p_N , polling) or reaches a threshold of average activity ($\theta_{trigger}$, pooling), an overt response is triggered that is measured as RT.

We determined how RT distributions and the dynamics of individual accumulators were influenced by three ensemble properties: the number of accumulators ($1 \leq n \leq 1,000$), the correlation of accumulation rates across accumulators ($0.0 \leq r_v \leq 1.0$), and the termination rule of the accumulation process (polling: $0\% < p_N \leq 100\%$; and pooling: $\sum A_i(t)/N \geq \theta_{trigger}$). We explored two influential types of stochastic accumulator models, one assuming within-trial as well as between-trial variability in accumulation (diffusion model) (38) and one assuming only between-trial variability (linear ballistic accumulator model) (39), as well as four variants making additional assumptions. Conclusions based on simulation of these models agreed, so we present the simple linear ballistic accumulator model here and the diffusion model and other more complex models in *SI Text, Robustness of Findings*.

Results

RT Distributions from One and Many Accumulators. We began by identifying the conditions under which an individual accumulator model ($n = 1$) and a large-ensemble accumulator model ($n = 1,000$) predict RT distributions with similar shapes, defined as overlapping 95% confidence intervals over all five RT quintiles (0.1, 0.3, 0.5, 0.7, 0.9). We observed that an individual accumulator model and a large ensemble accumulator model predict RT distributions with virtually indistinguishable shapes if accumulation rates are at least moderately correlated ($r_v \geq 0.6$) with intermediate termination rules. Much higher rate correlations ($r_v \geq 0.9$) are necessary under extreme termination rules (Fig. 2). Similar results were obtained under a pooling mechanism (Fig. 2, rightmost column). Thus, RT distributions can be explained both by an individual model accumulator and by accumulating activity of large neuronal ensembles only if their activation dynamics are moderately correlated and RT is not governed by extremely fast or slow accumulators.

RT Distributions Over a Range of Accumulator Ensemble Sizes. We also investigated the invariance of RT distributions over a range of ensemble sizes to determine whether RTs may be invariant once some critical ensemble size is reached. Knowing that the same RT distributions are predicted whether an ensemble has 10




Fig. 2. Predicted RT distributions as a function of ensemble size (N), termination rule (p_N), and accumulation rate correlation (r_v). Each panel shows the 0.1, 0.3, 0.5, 0.7, and 0.9 RT quantiles on a log-log scale (the x axis ranges from 10^0 to 10^3 ; the y axis ranges from 10^2 to 10^3) as a function of N , p_N , and r_v vary across columns and rows, respectively. We identified conditions (p_N and r_v) under which RT distributions were (i) invariant over the entire interval of N (i.e., 1,1,000; white panels), (ii) invariant with N over the interval (10,1,000; light gray panels), (iii) invariant with N over the interval (100,1,000; medium gray panels), and (iv) not invariant with N (dark gray panels).

accumulators or 1,000 accumulators or more provides important insights into the properties of ensemble dynamics. It may be that ensembles need to be “large enough” but that the precise size of the ensemble has little effect on the RT that the ensemble generates.

Extending the analysis above, we investigated how RT distributions scale with accumulator ensemble size. We identified conditions under which a small-ensemble model ($n = 10$) and an intermediate-ensemble model ($n = 100$) predict RT distributions with similar shapes as a large-ensemble model ($n = 1,000$). RT distributions were invariant across ensembles with at least 10 accumulators if accumulation rates were at least modestly correlated ($r_v \geq 0.3$) and termination rules avoided the extremes ($10\% \leq p_N \leq 90\%$). RT distributions were invariant across larger ensembles ($n \geq 100$) with even lower rate correlations ($r_v \geq 0.1$). Only if accumulation rates were uncorrelated ($r_v = 0.0$) or termination rules were extreme ($p_N = \text{first}$ and $p_N = \text{last}$) did RT distributions vary dramatically in scale and shape with ensemble size (Fig. 2). Similar findings were observed when RT was determined by the pooling termination mechanism (Fig. 2, right-most column) and with other accumulator model variants we investigated (*SI Text, Robustness of Findings*).

Variability in RT remains remarkably constant across different ensemble sizes over a large proportion of parameter space. Only for uncorrelated accumulators and extreme termination rules ($p_N = \text{first}$ or $p_N = \text{last}$) does ensemble size affect RT variability, a lack of invariance anticipated by extreme value statistics. By analogy to the central limit theorem, we can perhaps anticipate why median RT remains invariant with ensemble size. However, there is no single mathematical property that might allow us to anticipate why variability in RT is invariant with ensemble size across correlated samples and intermediate termination rules, so

we need to rely on simulation. To begin with, we know that for $p_N = \text{first}$, variability decreases with ensemble size, and for $p_N = \text{last}$, variability increases with ensemble size. So at some point in the range of termination rules we might expect an invariance of variability with ensemble size. What is striking is that this invariance is observed across all of the intermediate termination rules we investigated, not just a single value of termination rule. Also, for small ensemble sizes, variability is largely dominated by sampling variability across those few accumulators, and low correlations between accumulator rates may have only a small influence on the predicted variability from trial to trial. By contrast, for large ensemble sizes, variability is largely dominated by the between-trial variability introduced by the correlation between accumulator rates. These counteracting effects of ensemble size and correlation largely cancel each other out, producing invariance in RT distributions over a range of model parameters and model architectures (*SI Text, Robustness of Findings*) that we did not anticipate.

Invariance of A_{RT} with RT. We then investigated how the trial-averaged A_{RT} from an individual accumulator can be invariant with RT even though RT is produced by a large ensemble of accumulators. Most accumulator models are based on thresholds that are invariant across RT (40–42), and multiple laboratories have observed invariant thresholds of neural discharge rate (6–17). However, the A_{RT} of an individual accumulator participating in the ensemble is not guaranteed to reach the same value on each trial because of the stochastic nature of its accumulation process—on some trials it has reached θ and contributes to the measured RT, but on other trials it has not yet reached θ and so does not contribute (Fig. 1C). Though it is trivially true for a single accumulator that A_{RT} will be invariant with RT, it is unknown whether large ensembles of accumulators with intermediate termination rules and accumulation rate correlations reproduce the invariance of A_{RT} with RT that is regularly measured in neurophysiology.

Just like a neurophysiology experiment would randomly sample one neuron in some brain region, we randomly selected one accumulator in the ensemble and measured A_{RT} for that accumulator on each simulated trial. We then quantified how the slope of the linear regression of A_{RT} over RT varied for ensembles of 10, 100, and 1,000 accumulators (Fig. 3), mimicking the approach used in neurophysiological analyses. For small ensembles ($n = 10$), A_{RT} was invariant over RT under intermediate termination rules ($10\% \leq p_N \leq 90\%$) and moderate rate correlations ($r_v \geq 0.4$). With many accumulators ($n = 1,000$), the invariance of A_{RT} with RT was only violated for the earliest termination rule ($p_N = \text{first}$) and low accumulation rate correlations ($r_v \leq 0.3$). Under a pooling mechanism, the invariance of A_{RT} with RT was never violated. Thus, the invariance of A_{RT} with RT emerges from the dynamics of individual accumulators operating in large ensembles, even though the dynamics of no single accumulator uniquely determine RT.

Relationship Between A_{RT} and θ . Finally, we explored how the A_{RT} measured from an individual accumulator relates to the actual threshold of that accumulator (θ). In the neurophysiology literature, it is commonly assumed that the A_{RT} of an individual neuron represents a threshold like that in stochastic accumulator models. However, because A_{RT} is a trial-averaged measure and the true threshold of a neuron (θ) is unknown, we do not know how closely the value of A_{RT} approximates the value of θ .

As expected, with $n = 1$, A_{RT} was constant with RT and identical across trials, and the measured A_{RT} equaled the model parameter θ . However, in ensembles operating under intermediate termination rules ($10\% < p_N < 90\%$) A_{RT} varied significantly between trials (Fig. 4). Thus, individual accumulators acting in ensembles do not reach the same activation level at RT




Fig. 3. Relationship between A_{RT} and RT as a function of ensemble size (N), termination rule (p_N), and accumulation rate correlation (r_v). Each panel shows the linear regression slope of A_{RT} on RT, expressed as colored pixels, for three ensemble sizes (Left, $n = 10$; Center, $n = 100$; Right, $n = 1,000$) and all combinations of termination rules and accumulation rate correlations. Hatched pixels indicate parameter combinations for which A_{RT} varied systematically with RT. Thus, beige, nonhatched pixels represent parameter combinations for which the slope of the linear relationship between A_{RT} and RT was zero and nonsignificant.

on each trial, meaning that measured A_{RT} is not necessarily equivalent to the threshold specified by the model (θ) for any given accumulator. Analogous nonequivalence was observed for pooling mechanisms. We further observed that the termination rule determined how closely A_{RT} approximated θ . Under early termination rules ($p_N < 50\%$), average A_{RT} was less than θ . Under late termination rules ($p_N > 50\%$), average A_{RT} was greater than θ . Under the median termination rule ($p_N = 50\%$), average A_{RT} equaled θ ; this entails that the relationship between A_{RT} and θ cannot be determined without knowledge of the termination rule. The accumulation rate correlation determined the magnitude of variability in A_{RT} . The more homogeneous the accumulators, the smaller the variability in A_{RT} and the closer the agreement with θ , which implies that the degree of stochastic variation in A_{RT} is indicative of the homogeneity of the accumulation process in the ensemble. Together, though these findings

demonstrate unanticipated complexity in the relationship between A_{RT} measured in an individual accumulator and the true θ that defines its dynamics, in conditions under which one accumulator resembles many, the average A_{RT} measured from neurons is a fair proxy of the relation of θ to RT.

Discussion

Before carrying out these simulations, we thought that different combinations of ensemble size, accumulation rate correlation, and termination rule might produce markedly different qualitative behavior. Instead, we observed that the RT distributions predicted by large ensembles of redundant accumulators were invariant with ensemble size, except in the conditions of extreme termination rules and low accumulation rate correlations. These results did not depend on the particular form of the accumulator, variation in parameters such as leakage or within-trial noise magnitude, and consistency of v and θ across accumulators (*SI Text, Robustness of Findings*).

These findings complement previous models of decision-making by incorporating stochastic variability across multiple redundant accumulators and specifying constraints on the degree of consensus necessary for robust performance across variation in ensemble size. The rate correlations we found exceed the noise correlation found among neighboring neurons (30, 31) but can arise naturally from redundancy in common inputs, recurrent connectivity, and modulation by a common source (32–34, 43). Consensus through correlation of accumulation rates also prevents extreme neural activity from governing behavior.

These findings also provide clarification and caution about the conjecture that the activation level reached by particular neurons before RT (A_{RT}) corresponds conceptually and quantitatively to the threshold parameter of accumulator models ($\theta_{trigger}$). This linking proposition cannot be taken for granted (44), and the current demonstration that mapping model parameters onto measures of individual accumulators depends on unobserved statistical properties of the ensemble in which these accumulators operate. However, the EPU model demonstrates the necessity of obtaining multielectrode recordings to assess correlations in neural accumulation rates. These recordings should be made from homogeneous ensembles of neurons at different sensorimotor levels, but most importantly in neurons projecting to brainstem and spinal circuits that innervate motor neurons; this can provide key insights into termination rules and variability of A_{RT} across trials. These observations are crucial to validating the mapping of model parameters onto neural measures. However, the robustness of the relationships between RT distributions and ensemble size may reveal how measurements at different scales (single neurons, multiunit activity, local field potentials, EEG, MEG, fMRI) can appear to relate so well to the parameters of accumulator models.




Fig. 4. Distribution of measured activation level around RT (A_{RT}) between trials in a randomly selected accumulator as a function ensemble size (N), termination rule (p_N), and accumulation rate correlation (r_v). The x axis ranges from 10^0 to 10^3 , and the y axis ranges from 10^2 to 10^3 . Other conventions as in Fig. 2. Individual threshold (θ , red line) was identical across accumulators. Thus, correspondence between A_{RT} and θ is indicated by overlap of distributions (black lines) and threshold (red line).

To summarize, the random variation of RT has been explained through models of a stochastic accumulation process and through measures from individual neurons that appear to correspond to that process. This juxtaposition entails a previously unaddressed scaling problem—how can the activity of a multitude of redundant neurons map onto a single model accumulator? We now show how coordinated stochastic accumulation among many redundant accumulators can produce realistic RT distributions and accumulator dynamics regardless of the number of accumulators over a wide range of ensemble parameters, requiring only modest accumulation rate correlations and prohibiting the fastest or slowest accumulators from governing performance. Under this design principle, the dynamics of individual accumulators predict the behavior of the ensemble. In future work, these scaling properties can be explored in two complementary approaches. First, the EPU model can be extended to more complex models explaining performance in choice and stopping tasks (*SI Text, Future Model Extensions*). Second, as technologies develop to map the activity of large ensembles of neurons in the brain (45), it will become more tractable to monitor the activity of ensembles of neurons in circuits instantiating accumulation and threshold mechanisms (*SI Text, Neural Threshold Mechanisms*), providing an opportunity to verify the predictions of our simulations.

Materials and Methods

EPU Model. Embodying the well-known dictum “out of many, one,” we simulated ensembles of N stochastic accumulators (2) to understand how one RT is produced from that ensemble. With bold letters used to represent N -dimensional vectors, the ensemble of accumulators is governed by the following stochastic differential equation

$$d\mathbf{A}(t) = (\mathbf{v} - k \cdot \mathbf{A}(t)) \frac{dt}{\tau} + \sqrt{\frac{dt}{\tau}} \boldsymbol{\xi}. \quad [1]$$

Eq. 1 implies that the change in activation $d\mathbf{A}$ at every time step dt depends on the accumulation rate \mathbf{v} driving the accumulators toward threshold θ , the leakage constant k pushing activation back to baseline as it becomes larger, and Gaussian random noise $\boldsymbol{\xi}$. The linear ballistic accumulator model presented in *Results* assumed no leakage and no within-trial Gaussian random noise, but we do consider models with these characteristics in *SI Text, Robustness of Findings*.

We sampled \mathbf{v} from an N -dimensional multivariate lognormal distribution, $\mathbf{v} \sim \ln \mathcal{N}(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)$, where $\boldsymbol{\mu}_v$ is a vector of identical location parameters (μ_v) and $\boldsymbol{\Sigma}_v$ is the covariance matrix. This covariance matrix was computed as

$$\boldsymbol{\Sigma}_v = \mathbf{r}_v \cdot \sigma_v, \quad [2]$$

where \mathbf{r}_v is the accumulation rate correlation matrix with off-diagonal elements equal to r_v , and σ_v is the scale parameter of the lognormal distribution.

We sampled \mathbf{v} from a multivariate log-normal distribution for three reasons. First, the lognormal distribution takes positive values only and is therefore a natural choice for modeling accumulation rates of movement neurons that increase firing rate before a movement (46). Second, log-normal race models with similar parameter values (see below) can account for the shape of RT distributions (47). Third, the multivariate log-normal and multivariate normal are the only distributions for sampling correlated random variables with simple analytic solutions. In additional simulations, we demonstrate that the type of sampling distribution does not change findings qualitatively (*SI Text, Robustness of Findings*).

We assumed identical accumulation distributions (mean = 1 and SD = 1) and thresholds ($\theta = 100$ units) across accumulators; this seems to contrast with the idiosyncrasy of neurons, but much of this idiosyncrasy is eliminated in the analysis of neurophysiological data through normalization of spike density functions. Moreover, additional simulations demonstrate that varying

accumulation rate distributions and thresholds across accumulators did not alter our findings qualitatively (*SI Text, Robustness of Findings*).

Following the stochastic accumulator literature (1, 3), RT was modeled as the sum of the duration of three processing stages: (i) a stimulus-encoding stage with fixed duration ($T_E = 100$ ms), during which activation level of all N accumulators was equal to zero, $\mathbf{A}(0) = 0$; (ii) an accumulation stage with variable duration (T_A), during which \mathbf{A} increased with rate \mathbf{v} toward θ until the termination rule was met (see Introduction); and (iii) a response-execution stage with fixed duration ($T_R = 15$ ms), during which \mathbf{A} continued to increase. T_E and T_R were set in accordance with values measured empirically (48) and used in previous neurally constrained stochastic accumulator models (49). If the accumulation process had not met the termination rule within 100 s, it was aborted and no RT was logged. This cutoff time was chosen so that an RT was obtained in >90% of all simulated trials under all combinations of ensemble size, accumulation rate correlation, and termination rule.

Monte Carlo Simulations. Simulations were performed in MATLAB (MathWorks Inc., version 7.13), running in parallel on the high-performance computer cluster at the Vanderbilt Advanced Center for Computing for Research and Education.

We manipulated three key parameters: ensemble size (N), accumulation rate correlation (r_v), and termination rule (ρ_N). We varied the ensemble size across 20 levels ($1 \leq n \leq 1,000$, in increments of 10 between 10 and 100 accumulators and in increments of 100 between 100 and 1,000 accumulators), the accumulation rate correlation across 11 levels ($0.0 \leq r_v \leq 1.0$ in increments of 0.1), and the termination rule across 12 levels (polling mechanism, $0\% < \rho_N \leq 100\%$, in 10% increments; pooling mechanism, $\sum A_i(t)/N \geq \theta_{trigger}$), yielding a total of 2,640 RT models. Although neuronal ensembles constitute many more accumulators, we did not go beyond 1,000 due to limitations of computational time and resources. In some versions of the models we investigated ensemble sizes of 5,000–10,000 and found the same results. Moreover, previous work has demonstrated that intrinsic noise correlations among neurons entail upper limits on pool size (50).

For each combination of those three key parameters, a simulation consisted of 1,000 Monte Carlo repetitions of 500 trials. On each trial, we simulated N correlated, redundant accumulation processes. When a critical proportion of these accumulators reached threshold (ρ_N , polling mechanism) or when the average activity across all accumulators reached threshold ($\sum A_i(t)/N \geq \theta_{trigger}$, pooling mechanism), a response was made that was measured as RT. Analogous to a neurophysiology experiment, we measured A_{RT} as the mean activation level 10–20 ms before RT in a single accumulator that was randomly selected from the ensemble once per session.

For each set of 500 trials, we computed a number of descriptive statistics. To characterize the distributions of RT and A_{RT} , we computed five quantiles (0.1, 0.3, 0.5, 0.7, and 0.9). To describe the relationship between A_{RT} and RT, we sorted trials by RT, binned them into groups of 10, and computed the linear regression slope of the relationship between A_{RT} and RT. The 1,000 Monte Carlo repetitions enabled us to compute 95% confidence intervals on the descriptive statistics by estimating the 2.5th and 97.5th percentile of the distribution across the 1,000 repetitions.

To determine the conditions under which RT distributions can be explained by 1 and 1,000 accumulators, we identified accumulation rate correlations and termination rules producing overlapping confidence intervals for each RT quintile. To determine how RT distributions scale with ensemble size, we repeated the same analysis for the $n = 10$ vs. $n = 1,000$ and $n = 100$ vs. $n = 1,000$ comparisons. To determine conditions under which A_{RT} was invariant with RT, we identified accumulation rate correlations and termination rules that produce regression slope confidence intervals including zero, separately for $n = 10$, $n = 100$, and $n = 1,000$.

ACKNOWLEDGMENTS. We thank J. Brown, R. Desimone, S. Everling, D. Godlove, and J. Kalaska for helpful comments on early versions of the manuscript. This work was supported by National Institutes of Health Grant R01EY021833, National Science Foundation Grant SMA1041755, the Vanderbilt Advanced Computing Center for Research and Education, and Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience.

- Luce RD (1986) *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford Univ Press, New York).
- Usher M, McClelland JL (2001) The time course of perceptual choice: The leaky, competing accumulator model. *Psychol Rev* 108(3):550–592.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111(2):333–367.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113(4):700–765.
- Ratcliff R, Van Dongen HPA (2011) Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proc Natl Acad Sci USA* 108(27):11285–11290.

6. Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. *Science* 274(5286):427–430.
7. Lecas JC, Requin J, Anger C, Vitton N (1986) Changes in neuronal activity of the monkey precentral cortex during preparation for movement. *J Neurophysiol* 56(6):1680–1702.
8. Ratcliff R, Cheria A, Segraves M (2003) A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol* 90(3):1392–1407.
9. Boucher L, Palmeri TJ, Logan GD, Schall JD (2007) Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychol Rev* 114(2):376–397.
10. Purcell BA, et al. (2010) Neurally constrained modeling of perceptual decision making. *Psychol Rev* 117(4):1113–1143.
11. Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27(3):161–168.
12. Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
13. Maimon G, Assad JA (2006) A cognitive signal for the proactive timing of action in macaque LIP. *Nat Neurosci* 9(7):948–955.
14. Fecteau JH, Munoz DP (2007) Warning signals influence motor processing. *J Neurophysiol* 97(2):1600–1609.
15. Bollimunta A, Totten D, Ditterich J (2012) Neural dynamics of choice: Single-trial analysis of decision-related activity in parietal cortex. *J Neurosci* 32(37):12684–12701.
16. Tanaka M (2007) Cognitive signals in the primate motor thalamus predict saccade timing. *J Neurosci* 27(44):12109–12118.
17. Ding L, Gold JI (2012) Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb Cortex* 22(5):1052–1067.
18. Schurger A, Sitt JD, Dehaene S (2012) An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci USA* 109(42):E2904–E2913.
19. O'Connell RG, Dockree PM, Kelly SP (2012) A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* 15(12):1729–1735.
20. Smyrnis N, et al. (2012) Single-trial magnetoencephalography signals encoded as an unfolding decision process. *Neuroimage* 59(4):3604–3610.
21. Forstmann BU, et al. (2008) Striatum and pre-SMA facilitate decision-making under time pressure. *Proc Natl Acad Sci USA* 105(45):17538–17542.
22. Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci USA* 107(50):21767–21772.
23. Bogacz R, Wagenmakers E-J, Forstmann BU, Nieuwenhuis S (2010) The neural basis of the speed-accuracy tradeoff. *Trends Neurosci* 33(1):10–16.
24. Wong K-F, Wang X-J (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26(4):1314–1328.
25. Beck JM, et al. (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60(6):1142–1152.
26. Ganguli S, et al. (2008) One-dimensional dynamics of attention and decision making in LIP. *Neuron* 58(1):15–25.
27. Lo CC, Boucher L, Paré M, Schall JD, Wang XJ (2009) Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *J Neurosci* 29(28):9059–9071.
28. Wang X-J (2008) Decision making in recurrent neuronal circuits. *Neuron* 60(2):215–234.
29. Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14(7):811–819.
30. Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370(6485):140–143.
31. Cohen JY, et al. (2010) Cooperation and competition among frontal eye field neurons during visual target selection. *J Neurosci* 30(9):3227–3238.
32. De Luca CJ, Erim Z (1994) Common drive of motor units in regulation of muscle force. *Trends Neurosci* 17(7):299–305.
33. Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci* 18(10):3870–3896.
34. Narayanan NS, Kimchi EY, Laubach M (2005) Redundancy and synergy of neuronal ensembles in motor cortex. *J Neurosci* 25(17):4207–4216.
35. Seeley TD, et al. (2012) Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* 335(6064):108–111.
36. Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233(4771):1416–1419.
37. Lee C, Rohrer WH, Sparks DL (1988) Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332(6162):357–360.
38. Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. *Psychol Sci* 9:347.
39. Brown SD, Heathcote A (2008) The simplest complete model of choice response time: Linear ballistic accumulation. *Cognit Psychol* 57(3):153–178.
40. Grice GR (1968) Stimulus intensity and response evocation. *Psychol Rev* 75(5):359–373.
41. Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: The urgency-gating model. *J Neurosci* 29(37):11560–11571.
42. Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11(6):693–702.
43. Haefner RM, Gerwinn S, Macke JH, Bethge M (2013) Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat Neurosci* 16(2):235–242.
44. Heitz RP, Schall JD (2012) Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76(3):616–628.
45. Alivisatos AP, et al. (2013) Neuroscience. The brain activity map. *Science* 339(6125):1284–1285.
46. Bruce CJ, Goldberg ME, Bushnell MC, Stanton GB (1985) Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol* 54(3):714–734.
47. Heathcote A, Love J (2012) Linear deterministic accumulator models of simple choice. *Front Psychol* 3:292.
48. Pouget P, Emeric EE, Stuphorn V, Reis K, Schall JD (2005) Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *J Neurophysiol* 94(3):2086–2092.
49. Purcell BA, Schall JD, Logan GD, Palmeri TJ (2012) From salience to saccades: Multiple-alternative gated stochastic accumulator model of visual search. *J Neurosci* 32(10):3433–3446.
50. Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16(4):1486–1510.

Supporting Information

Zandbelt et al. 10.1073/pnas.1310577111

SI Text

How Many Neurons Produce a Movement? Consider a saccadic eye movement. In asking how many neurons prepare a saccade, one discovers a specific lack of information about density and number of neurons in various structures. However, an approximate calculation (1) based on the results of inactivation, lesion, and anatomical investigations indicates that at least 10^5 neurons are necessary to produce a saccade. Note that the systematic simulations we report in the paper only show ensemble sizes up to 10^3 accumulators. For some models, we did explore ensembles up to 10^4 accumulators, but as we note in the article, predicted dynamics of ensembles of size 10^2 or greater were either invariant or not. Simulations of ensemble of size 10^5 would be computationally prohibitive and would not have provided further insight given their computational cost.

To determine our 10^5 estimate, we summed the number of neurons in the structures in which presaccadic activity related to the timing of the initiation of the movement; this includes the frontal eye field, superior colliculus, thalamus, basal ganglia, and brainstem. Estimates of the density of neurons and glia in the cerebral cortex and other structures have been determined (2–6). Our calculation is based on an estimate of 73,000 neurons per square millimeter in cerebral cortex, based on measurements ranging from 20,000 to 92,000 neurons per square millimeter of 146,000 cells of all types per square millimeter. We take the cortical area of frontal eye field (FEF) to be 50 mm^2 , so assuming a uniform 2-mm cortical depth, the total cell number in FEF is 7.3 million (low, 2 million; high, 9.2 million). However, if only the pyramidal cells in layer 5 are responsible for saccade generation, then this count must be reduced proportionally by estimating the thickness of layer 5 at 0.05 mm: 182,500 (low, 50,000; high, 230,000). Counts of cells in the superior colliculus (SC) arrive at a value of ~ 7 million total cells in the SC with $\sim 25\%$ of those being neurons (4). Assuming the intermediate layers constitute 40% of the thickness of the SC and that 50% of the neurons in the intermediate layers contribute to saccade generation, the number of neurons is 350,000. Restricting the thalamus contribution to the lateral sector of the medial dorsal nucleus and assuming again that 50% of these neurons contribute to saccade generation, the count is 100,000. Assuming that the number of neurons in the caudate nucleus and the substantia nigra pars reticulata that contribute to saccade generation are equivalent to that in the superior colliculus, then the basal ganglia contribution is 700,000. Finally, we assume that there are 10,000 long-lead burst neurons in the brainstem. Based on all these assumptions and estimates, the total number of presaccadic movement-related neurons amounts to 9.9×10^5 . Now, because presaccadic movement neurons have movement fields, not every neuron contributes to each saccade of a particular direction and amplitude. If we assume that one-third of the neurons are active before any saccade, then the total is 3.3×10^5 , and if the fraction is as low as 1/10, then the total is 9.9×10^4 ; this is the basis for our claim that 10^5 neurons are necessary for initiation of a saccade.

Extension of Previous Work. Previous studies have investigated population coding in sensory systems (7). These studies have focused on the effects of shared noise, population size, and noise on choice probabilities of perceptual judgments by assuming that the activity of sensory neurons is pooled or temporally integrated to represent evidence for particular responses (8). This work has provided crucial insights into the readout of sensory information,

but it does not address the fundamental scaling problems that are the focus of our work. Unlike our framework, these models make the simplifying assumption that evidence for each response is represented by a single activation. Thus, how and when the brain reaches consensus is unambiguous. Moreover, though this work was seminal in identifying the potential impact of correlated noise in sensory neurons on the probabilities of perceptual judgments, it does not address the role of rate correlations among downstream neurons responsible for generating a motor response at particular time.

Other studies have investigated population coding in motor systems. For example, some groups have described how pools of motor neurons with distinct response fields can encode something like movement direction through vector averaging (9, 10). The current investigation significantly advances this line of research because it is concerned with how stochastically variable accumulating activation is combined through time to produce an instantaneous change of state corresponding to measured response times (RTs). Our work also draws attention to the consequences of redundant coding of individual responses by multiple neurons.

Another line of work has analyzed response preparation as the migration of population activity in multidimensional state space (11). Though the current work can be extended in this direction, the compatibility of the rise-to-threshold mechanism with the measure of variance that inspired the state space hypothesis has been demonstrated (12). The advantage of our approach is that it explains why population activity exhibits specific dynamics and outlines the conditions that must be met by specific subpopulations of neurons to initiate motor responses.

However, other researchers have investigated how biophysical models of single neurons (13) can be mapped to macroscopic descriptions of neural activity through mathematical mean-field approximations (14–16) and how large networks of spiking neurons can reproduce both the neurophysiological and behavioral data (17–19). These neural mass or neural field models have provided intuitive and analytic insights into how activity patterns emerge in large networks, but these models do not explain how the dynamics of individual neurons in the network can account for RT variability. Moreover, we tested alternative methods by which the consensus to act among neurons evolves over time, which is often overlooked in neural-network modeling. The current investigation advances this line of research significantly by linking well-established abstract models of the dynamics of individual accumulators, the statistical properties of ensembles, and the full distribution of RT.

Robustness of Findings. In additional simulations, we tested the robustness and generality of findings obtained using the linear ballistic accumulator model by examining models assuming (i) within-trial variability in accumulation rate, (ii) leakage, (iii) variability in thresholds across accumulators, (iv) variability in mean accumulation rates across accumulators, and (v) alternative accumulation rate sampling distributions. As in the main text, we analyzed RT distributions (Fig. S1), neural activation level at RT (A_{RT}) as a function of RT (Fig. S2), and A_{RT} distributions (Fig. S3). The parameters that were kept fixed in the simulations are shown in Table S1.

Robustness of findings to within-trial variability in accumulation rate. Compared with the linear ballistic accumulator (LBA) model, the diffusion model not only assumes between-trial variability but also within-trial variability in accumulation rate (Table S1). Gaussian random noise (Eq. 1) was sampled from a multivariate normal

distribution, $\Sigma_{\xi} = \mathbf{r}_{\xi} \cdot \sigma_{\xi}$, where μ_{ξ} is a vector of zeros and Σ_{ξ} is the covariance matrix. This covariance matrix was computed as

$$\Sigma_{\xi} = \mathbf{r}_{\xi} \cdot \sigma_{\xi}, \quad [S1]$$

where \mathbf{r}_{ξ} is the noise correlation matrix with off-diagonal elements equaling r_{ξ} , and σ_{ξ} is the SD of the normal distribution, reflecting the magnitude of the noise. Based on previous findings (20, 21), we set $r_{\xi} = 0.10$. We set the magnitude of the noise to $\sigma_{\xi} = 0.5$. Simulation results for the diffusion-type model are shown in Figs. S14, S24, and S34. The results from the LBA model and diffusion model were largely in agreement, but we noticed a few small differences. We observed that large ensembles operating under late termination rules ($p_N \geq 70\%$) produced RT distributions with a longer upper tail, especially when the accumulation rate correlation was high (Fig. S14). Additionally, whereas in the LBA model the slope of the A_{RT} by RT relationship remained stable around zero when going from early to late termination rules, in the diffusion model we observed a shift from negative to positive slopes (Fig. S24). In fact, large ensembles operating under early termination rules and high accumulation rate correlations violated the invariance of A_{RT} over RT. Further analysis explained this finding (Fig. S4): because within-trial variability has a stronger impact on slow compared with fast RTs, under early termination rules, A_{RT} is often smaller than θ , resulting in a negative A_{RT} by RT relationship, whereas under late termination rules A_{RT} is often greater than θ , leading to a positive A_{RT} by RT relationship. Finally, in line with the RT findings, we found higher A_{RT} values in large ensembles under late termination rules ($p_N \geq 70\%$) than in the LBA model (Fig. S34). Despite these small differences, the diffusion model reproduced the robustness of RT distributions to variation in ensemble size and the invariance of A_{RT} over RT over a wide range of accumulation rate correlations and termination rules like we observed with the LBA model.

Robustness of findings to inclusion of leakage. Simulation results for a ballistic accumulator model with leakage (Eq. 1) are shown in Figs. S1B, S2B, and S3B. We chose the leakage parameter so that the maximum activation level for each accumulator reached an average asymptote of 200 units, twice as large as its threshold (Table S1). There were a few differences between models with and without leakage. We observed slightly longer RTs over the entire range of ensemble sizes, accumulation rate correlations, and termination rules tested (Fig. S1B). Furthermore, we found flatter slopes of the A_{RT} –RT relationship for late termination rules (Fig. S2B). Also, under perfectly correlated accumulation rates the invariance of A_{RT} with RT was violated, but examination revealed that this effect was very small (i.e., extremely narrow confidence intervals just above zero). Finally, we saw a reduction in stochastic variability of A_{RT} across trials that was most pronounced under low accumulation rate correlations and later termination rules (Fig. S3B), reflecting that leakage imposed a maximum on the activation level an accumulator could reach. Taken together, the general pattern of RT, A_{RT} by RT, and A_{RT} findings were in line with the results obtained with the LBA model assuming nonleaky accumulation.

Robustness of findings to variability in thresholds across accumulators. Simulation results for an LBA model with variability in thresholds across accumulators are shown in Figs. S1C, S2C, and S3C. The threshold activation level for each accumulator (θ) was sampled from a normal distribution truncated at zero, $\theta \sim \mathcal{N}(\mu_{\theta}, \sigma_{\theta}^2)$, $\mu_{\theta} \in (0, \infty)$; by contrast, the LBA model presented in the main text assumed that θ was identical across accumulators. Neurons are idiosyncratic and have different A_{RT} levels, and they may have different thresholds as well. We therefore analyzed a model in which θ varied randomly across accumulators according to a Gaussian distribution (Table S1). RT distributions were slightly more variable, especially under

high accumulation rate correlations ($r_v \geq 0.7$; Fig. S1C). There were hardly any differences in the slopes of the A_{RT} –RT relationship (Fig. S2C), and we noticed only small changes in stochastic variability of A_{RT} values when accumulation rates were high in combination with extreme termination values (Fig. S3C). Obviously, no differences were observed under a pooling mechanism, because RT production depends on the ensemble threshold, which remained the same. Taken together, these findings demonstrate that variability of thresholds between accumulators does not influence the general pattern of findings.

Robustness of findings to variability in mean accumulation rate across accumulators. The LBA model presented in the main text assumed identical accumulation rate distributions across accumulators. In this analysis we instead sampled the mean accumulation rate for each accumulator from a normal distribution truncated at zero, $\mu_v \sim \mathcal{N}(\mu_{v,a}, \sigma_{v,a}^2)$, $\mu_v \in (0, \infty)$, where $\mu_{v,a}$ is the mean accumulation rate across accumulators ($\mu_{v,a} = 1$), and $\sigma_{v,a}$ is the SD of the mean accumulation rate across accumulators ($\sigma_{v,a} = 0.05$). These means were then converted into a vector of location parameters that was used for sampling ν from an N -dimensional multivariate lognormal distribution (Eq. 1). In this way, some accumulators have relatively high rates, and others have relatively low rates, mirroring some of the heterogeneity seen in neurons. Simulation results for the LBA model with variability in mean accumulation rate across accumulators are shown in Figs. S1D, S2D, and S3D. It is important to note that the variability in mean rate across accumulators ($\sigma_{v,a}$) had to be somewhat small to obtain highly correlated accumulation rates from distinct accumulation rate distributions. Indeed, even at this small level of variability it turned out to be impossible to obtain highly correlated accumulation rates ($r_v \geq 0.7$) for very large ensembles ($n = 1,000$), as indicated by the cross-hatched panels in Fig. S1D and gray pixels in Fig. S2D. Nevertheless, there were hardly any differences in the RT distributions in other conditions (Fig. S1D). Also, the conditions producing invariant A_{RT} –RT relationships were very similar (Fig. S2D), even though the invariance was violated under slightly fewer conditions. There were no noticeable differences in the stochastic variability of A_{RT} values (Fig. S3D). Thus, these results show that that variation of mean accumulation rate between accumulators does not influence the variability pattern of findings.

Robustness of findings to alternative accumulation rate sampling distribution. In all our simulations so far, we used a multivariate log-normal distribution to sample accumulation rates with correlation. To test whether our findings depended on this particular distribution, we used a mixture approach (22) to instead sample accumulation rates from a multivariate gamma distribution, $\nu \sim \text{Gamma}(r, \lambda)$, where r is the shape parameter ($r = 2.5$) and λ is the rate parameter ($\lambda = 2$). Like the multivariate log-normal distribution, this distribution yields positive values only. The values of these parameters were chosen so that a model with one accumulator would produce RT distributions with location, spread, and shape similar to the linear ballistic accumulator model with one accumulator assuming a log-normal presented in the main text. Simulation results for the LBA model with accumulation rates sampled from a multivariate gamma distribution are shown in Figs. S1E, S2E, and S3E. As ensemble size increased, RT distributions produced under early termination rules showed a little less variability, whereas those produced under late termination rule showed a slightly later onset for low accumulation rate correlations and a slightly longer tail for high accumulation rate correlations (Fig. S1E). However, the overall pattern of results, as indicated by the various shadings, is strikingly similar to that obtained with the model in which accumulation rates were sampled from a multivariate log-normal distribution. We found a somewhat narrower area of parameter space under which invariant A_{RT} –RT relationships were produced (Fig. S2E), excluding the earliest and latest termination rules as viable mechanisms. Also, the direction of slopes showed a closer correspondence

to the diffusion model (Fig. S2A) than to the linear ballistic accumulator model presented in the main text. Finally, there was slightly less stochastic variability of A_{RT} values under low accumulation rate correlations (Fig. S3E). In sum, a model with accumulation rates sampled from a multivariate gamma distribution reproduced the robustness of RT distributions to variation in ensemble size and the invariance of A_{RT} over RT over a wide range of accumulation rate correlations and termination rules that we observed with the LBA models, even though this model appears a bit more sensitive to violations of invariance of A_{RT} over RT produced by extreme termination rules.

Neural Threshold Mechanisms. Neurobiologically, thresholds for ballistic movements are implemented by final common circuits in brainstem and spinal cord receiving inputs from many neurons in premotor structures with stochastically accumulating activity. For example, saccadic eye movements are initiated when omnipause neurons in the nucleus raphe interpositus release inhibition on burst neurons that produce the pulse of force producing the rapid eye movement (23). The descending, accumulating influence of cortical and subcortical accumulating activity ultimately inhibits the omnipause neurons, thereby releasing the saccade. However, the means by which the descending influence of ensembles of premovement neurons in different structures is combined or coordinated is entirely unknown. Other investigators have described evidence accumulation in brain regions that are removed anatomically and functionally from the response preparation process (24, 25). Though the neurobiology of the transition from preparation to execution is vaguely understood, outside recent modeling efforts (26, 27), a threshold mechanism for evidence accumulation triggering response preparation is even less understood.

Future Model Extensions. We addressed how ensemble size influences RT distributions in the simple framework of one-choice RT, which most directly corresponds to a simple go task or a very easy choice task with effectively no alternatives to consider; this enabled us to focus on the effects of ensemble size uncontaminated by the complexities of various interactions between competing processes common to models of more complex behaviors, such as choice and stopping. To further clarify the mapping between psychological models and neural measures of decision-making, future studies can extend this approach to account for choosing and stopping. Extending our framework to choosing and stopping is important because most research with cognitive paradigms involves choosing between alternatives (28), and response control requires explaining how responses are initiated and stopped (29, 30); however, this may be challenging because the framework needs to account for a much wider range of behavioral and neural data. To account for choice probabilities, the full distributions of error RTs, and the dynamics of individual neurons across all levels of choice difficulty, our framework will need to include multiple ensembles of stochastic accumulators, one for each response alternative, that race independently or compete interactively. To account for the effect of stop-signal onset asynchrony on stopping probability and the distribution of error RTs, as well as the latency of stopping, and the dynamics of individual neurons when stopping succeeds and fails, the choice framework needs to be further extended to include an ensemble of stochastic accumulators that can interrupt (29) or block* the accumulation process that leads to the production of a response (31).

*Yamaguchi M, Logan GD, Palmeri TJ, Schall JD, Poster Presented at Computational and Systems Neuroscience 2012, February 23–26, 2012, Salt Lake City, UT.

- Brown JW, Hanes DP, Schall JD, Stuphorn V (2008) Relation of frontal eye field activity to saccade initiation during a countermanding task. *Exp Brain Res* 190(2):135–151.
- Braitenberg V, Schüz A (1991) *Anatomy of the Cortex: Statistics and Geometry* (Springer, New York).
- Rockel AJ, Hiorns RW, Powell TP (1980) The basic uniformity in structure of the neocortex. *Brain* 103(2):221–244.
- Herculano-Houzel S, Collins CE, Wong P, Kaas JH (2007) Cellular scaling rules for primate brains. *Proc Natl Acad Sci USA* 104(9):3562–3567.
- Collins CE, Leitch DB, Wong P, Kaas JH, Herculano-Houzel S (2013) Faster scaling of visual neurons in cortical areas relative to subcortical structures in non-human primate brains. *Brain Struct Funct* 218(3):805–816.
- Carlo CN, Stevens CF (2013) Structural uniformity of neocortex, revisited. *Proc Natl Acad Sci USA* 110(4):1488–1493.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16(4):1486–1510.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13(11):1257–1269.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233(4771):1416–1419.
- Lee C, Rohrer WH, Sparks DL (1988) Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332(6162):357–360.
- Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006) Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci* 26(14):3697–3712.
- Purcell BA, Heitz RP, Cohen JY, Schall JD (2012) Response variability of frontal eye field neurons modulates with sensory input and saccade preparation but not visual search salience. *J Neurophysiol* 108(10):2737–2750.
- Herz AV, Gollisch T, Machens CK, Jaeger D (2006) Modeling single-neuron dynamics and computations: A balance of detail and abstraction. *Science* 314(5796):80–85.
- Fusi S, Mattia M (1999) Collective behavior of networks with linear (VLSI) integrate-and-fire neurons. *Neural Comput* 11(3):633–652.
- Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci* 8(3):183–208.
- Gerstner W (2000) Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. *Neural Comput* 12(1):43–89.
- Wang X-J (2002) Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36(5):955–968.
- Lo CC, Boucher L, Paré M, Schall JD, Wang XJ (2009) Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *J Neurosci* 29(28):9059–9071.
- Deco G, Rolls ET, Romo R (2009) Stochastic dynamics as a principle of brain function. *Prog Neurobiol* 88(1):1–16.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370(6485):140–143.
- Cohen JY, et al. (2010) Cooperation and competition among frontal eye field neurons during visual target selection. *J Neurosci* 30(9):3227–3238.
- Minhajuddin A, Harris I, Schucany W (2004) Simulating multivariate distributions with specific correlations. *J Statist Comput Simulation* 74(8):599–607.
- Scudder CA, Kaneko CS, Fuchs AF (2002) The brainstem burst generator for saccadic eye movements: A modern synthesis. *Exp Brain Res* 142(4):439–462.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22(21):9475–9489.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11(6):693–702.
- Lo C-C, Wang X-J (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat Neurosci* 9(7):956–963.
- Purcell BA, Schall JD, Logan GD, Palmeri TJ (2012) From salience to saccades: Multiple-alternative gated stochastic accumulator model of visual search. *J Neurosci* 32(10):3433–3446.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111(2):333–367.
- Boucher L, Palmeri TJ, Logan GD, Schall JD (2007) Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychol Rev* 114(2):376–397.
- Logan GD, Cowan WB, Davis KA (1984) On the ability to inhibit simple and choice reaction time responses: A model and a method. *J Exp Psychol Hum Percept Perform* 10(2):276–291.
- Palmeri TJ, Schall JD, Logan GD (2013) Neurocognitive modeling of perceptual decision making. *Oxford Handbook of Computational and Mathematical Psychology*, eds Busemeyer JR, Townsend J, Wang ZJ, Eidsels A (Oxford University Press, Oxford, UK), in press.




Fig. S1. RT distributions as a function of N , p_N , and r_v for (A) a diffusion model, (B) a ballistic accumulator model with leakage, (C) a linear ballistic accumulator model with variability in threshold (θ) across accumulators, (D) a linear ballistic accumulator model with variability in the mean of the accumulation rate distribution across accumulators ($\mu_{v,a}$), and (E) a linear ballistic accumulator model with accumulation rates sampled from a multivariate gamma distribution. Each subpanel shows the 0.1, 0.3, 0.5, 0.7, and 0.9 RT quantiles as a function of N on a log-log scale, with RT ranging from 10^2 to 10^3 ms and N ranging from 10^0 to 10^3 . Format and abbreviations as in Fig. 2.




Fig. S2. Relationship between A_{RT} and RT as a function of N , p_N , and r_v for (A) a diffusion model, (B) a ballistic accumulator model with leakage, (C) a linear ballistic accumulator model with variability in threshold (θ) across accumulators, (D) a linear ballistic accumulator model with variability in the mean of the accumulation rate distribution across accumulators ($\mu_{v,a}$), and (E) a linear ballistic accumulator model with accumulation rates sampled from a multivariate gamma distribution. Format and abbreviations as in Fig. 3.




Fig. S4. Single-trial accumulation trajectories for a fast trial (*Left*) and slow trial (*Center*) and A_{RT} -RT relationship (*Right*) for a diffusion model with perfect accumulation rate correlation and early termination rule ($p_N = \text{'first'}$; *Upper*) or late termination rule ($p_N = \text{'last'}$; *Lower*). Note the greater impact of within-trial noise on slow compared with fast trials; as a result, A_{RT} is often smaller than θ under early termination rules, but greater than θ under late termination rules. Red trajectory represents accumulation process of selected accumulator. Dots (*Right*) represent average activation level in selected accumulator per RT bin. Format and conventions as in Figs. 1 and 3.

Table S1. Values of fixed parameters for all models

Parameter	Diffusion (Figs. S1A, S2A, and S3A)	LBA with leakage (Figs. S1B, S2B, and S3B)	LBA with variation in threshold across accumulators (Figs. S1C, S2C, and S3C)	LBA with variation in mean rate across accumulators (Figs. S1D, S2D, and S3D)	LBA with rates from a multivariate gamma distribution (Figs. S1E, S2E, and S3E)
μ_V	-0.3466	-0.3466	-0.3466	-0.3466	—
σ_V	0.8326	0.8326	0.8326	0.8326	—
$\sigma_{V,a}$	0	0	0	0.05	—
μ_{ξ}	0	0	0	0	0
σ_{ξ}	0.5	0	0	0	0
r_{ξ}	0.1	0	0	0	0
μ_{θ}	100	100	100	100	100
σ_{θ}	0	0	20	0	0
k	0	0.005*	0	0	0
τ	1	1	1	1	1
dt	10	10	10	10	10
T_E , ms	100	100	100	100	100
T_R , ms	15	15	15	15	15

*The leakage constant was proportional to the accumulation rate and varied across trials and accumulators, so that the maximum activation level that each accumulator could attain on each trial was 200 units. μ_V , location parameter of log-normal distribution of accumulation rates; σ_V , scale parameter of log-normal distribution of accumulation rates across trials; $\sigma_{V,a}$, SD of the mean accumulation rate across accumulators; μ_{ξ} , mean of the Gaussian distribution of noise; σ_{ξ} , SD of the Gaussian distribution of noise; r_{ξ} , noise correlation; μ_{θ} , mean of the Gaussian distribution of thresholds; σ_{θ} , SD of the Gaussian distribution of thresholds; k , leakage constant; τ , time scale; dt , time step; T_E , stimulus encoding time; T_R , response execution time.