

A neural basis of choking under pressure

Adam L. Smoulder^{1,2}, Patrick J. Marino^{2,3}, Emily R. Oby^{2,3}, Sam E. Snyder^{2,4}, Hiroo Miyata¹, Nick P. Pavlovsky³, William E. Bishop⁵, Byron M. Yu^{1,2,6}, Steven M. Chase^{1,2,7,8,9*}, Aaron P. Batista^{2,3,8*}

¹Department of Biomedical Engineering, Carnegie Mellon University; Pittsburgh, USA

²Center for the Neural Basis of Cognition; Pittsburgh, USA

³Department of Bioengineering, University of Pittsburgh; Pittsburgh, USA

⁴Center for Neuroscience, University of Pittsburgh; Pittsburgh, USA

⁵Janelia Research Campus, Howard Hughes Medical Institute; Ashburn, USA

⁶Department of Electrical and Computer Engineering, Carnegie Mellon University; Pittsburgh, USA

⁷Neuroscience Institute, Carnegie Mellon University; Pittsburgh, USA

⁸These authors contributed equally

⁹Lead contact

*Correspondence: schase@cmu.edu, aaron.batista@pitt.edu

Summary

Incentives tend to drive improvements in performance. But when incentives get too high, we can “choke under pressure” and underperform right when it matters most. What neural processes might lead to choking under pressure? We studied rhesus monkeys performing a challenging reaching task in which they underperform when an unusually large “jackpot” reward is at stake and sought a neural mechanism that might result in that underperformance. We found that increases in reward drive neural activity during movement preparation into, and then past, a zone of optimal performance. We conclude that neural signals of reward and motor preparation interact in the motor cortex in a manner that can explain why we choke under pressure.

Keywords: Motor cortex; reward; movement preparation; reaching; paradoxical performance decrement; cognitive neuroscience; motivation; reward-mediated performance

Main Text

Introduction

Failing to perform to one's highest standard when the potential payoff is particularly great is known as "choking under pressure"¹. While failures in professional athletics often provide memorable examples of this phenomenon, people choke under pressure in diverse settings, including test-taking², video games³, puzzle-solving⁴, and more⁴⁻⁶. Neuroimaging studies have implicated the involvement of reward and motor structures in choking under pressure^{7,8}, but the neural mechanisms whereby the possibility of increased rewards leads to performance failure remain unclear.

We recently reported that animals also choke under pressure⁹. Rhesus monkeys performed a challenging task in which they had to perform a goal-directed reach that was both fast and accurate (**Figure 1A**). We cued the volume of the liquid reward they would receive for a successful reach. Success was more frequent for Medium and Large cued rewards than for Small. This behavioral improvement presumably reflects an increase in motivation to perform this challenging task¹⁰⁻¹³. When "Jackpot" (i.e., rare and exceptionally large) rewards were proffered, monkeys underperformed. This effect is "choking under pressure", and it is characterized by an "inverted-U" relationship between performance and reward (**Figure 1B**)^{4,7,8}. Here we leverage the fact that monkeys choke under pressure⁹, like humans do, to explore the phenomenon's neural basis at the resolution of individual neurons' activity and at the sub-second timescale at which neural activity controls behavior.

We report a novel neural explanation of choking under pressure: a deficit in motor preparation. Motor preparation benefits the execution of rapid, voluntary movements, like the reaches the animals performed in our task, by allowing time for neural computations that support the movement to be completed before execution¹⁴⁻¹⁶. Motor preparation is known to engage populations of neurons in the primary motor cortex and the dorsal aspect of the premotor cortex (referred to collectively here as MC)¹⁷⁻¹⁹. To study how motor preparation relates to choking under pressure, we recorded the spiking activity of hundreds of MC neurons and examined how cued rewards modulated neural population activity during movement preparation. If choking under pressure involves a failure in motor preparation, we might expect there to be some aspect of MC activity that exhibits an inverted-U relationship with reward size, like behavior does.

Previous work has demonstrated that the state of MC activity being closer to an "optimal zone" before a movement cue influences reaction time and reach execution^{17,20}. We may then expect that MC activity for Large reward trials is closer to this optimal zone, whereas Small and Jackpot trials are less so. Leveraging this optimal zone framework, we hypothesized three neural mechanisms that could link MC preparatory activity to the choking under pressure behavior that we observed (**Figure 1C**). First, we considered whether MC might paradoxically receive a weakened reward signal for Jackpot rewards. This lack of motivational drive seems feasible given the observation in our prior work that monkeys exhibited an excess of hypometric reaches on Jackpot trials⁹, and might indicate that Jackpot scenarios are, paradoxically, not as intrinsically rewarding as reason might suggest. We refer to this as the "Insufficient Drive" hypothesis (**Figure 1C**, left). While prior work has generally demonstrated monotonic trends between incentives and neural activity in areas upstream of MC^{7,8,21}, one could conceivably imagine that neural stress signals from the prospect of a loss might overwhelm those of positive incentives⁸ in a manner that manifests as an insufficient drive of MC activity.

A second possibility is that the anticipated reward interacts with aspects of motor preparation important to the upcoming reach. It could be that when Large rewards are offered, the patterns of neural activity affiliated with motor preparation are made more suitable for the behavioral goal, but that Jackpot rewards somehow make neural activity less so. We refer to this as the "Neural Bias" hypothesis (**Figure 1C**, middle), where the inverted-U we observe in success rates can be explained by rewards of increasing magnitude biasing average neural activity patterns to be closer to, and then farther from, the optimal zone for movement preparation.

A third possibility is that reward modulates the variability of MC activity such that more or fewer trials' preparatory states are nearer to the optimal zone. In monkeys, much of the variability in movements comes from variability in upstream cortical signals²². Some animals have been shown to alter behavior by modulating neural variability based on context, as happens when juvenile songbirds practice courtship singing in solo bouts versus performing in front of a potential mate^{23–25}. A similar, albeit counterproductive, mechanism may be at play here: Increasing reward from Small to Large might suppress neural variability, but Jackpots might paradoxically cause neural variability to increase, disrupting performance. We refer to this as the “Neural Noise” hypothesis (Figure 1C, right).

We note that these hypotheses are not mutually exclusive, and it could be that some variant of each is at play in our task. Further, behavior cannot disambiguate between these, as the inverted-U trends in performance could emerge from any mechanism that pushes neural activity out of the optimal preparatory zones. Of course, it could also be that choking under pressure in this task is explained by neural activity changes occurring during movement, instead of preparation, meaning none of these hypotheses need be observed.

We tested each hypothesis individually in search of a motor preparatory explanation for choking under pressure. Our data support the Neural Bias hypothesis. We find that rewards interact with target preparation signals to drive neural activity towards a region associated with improved reach execution, and then, at the highest rewards, away from this region. This finding provides an example of how cortical circuits can be pressed beyond ideal functionality by internal drives (such as reward expectation) with consequences for behavior.




Figure 1. Potential neural mechanisms of choking under pressure. (A) Monkeys were trained to prepare a brisk reach to a small target. The color (Monkeys E, P) or shape (Monkey R) instructed the reward size. Parameters bolded in green were selected individually for each animal to keep the task challenging and motivating (Table S1 shows all parameter values). A separate choice task indicated that animals understood reward cues (Table S2). Simultaneously, we recorded from primary motor (M1) and/or dorsal premotor cortex (PMd), together termed MC, using 96-channel microelectrode “Utah” arrays (Blackrock Microsystems, Inc., see Table S1 for array location details). (B) Success rates improved from Small to Large rewards (binomial proportion test, $***p < 0.001$), indicating that performance in this difficult task is influenced by motivation. All three animals choked under pressure, indicated by the significant decrease in success rates from Large to Jackpot. Error bars represent S.E. of overall success rate shown in dark gray. Light gray traces show individual sessions. See Table S3 for full statistics. (C) We considered three mechanisms that could relate motor cortical activity to the inverted-U in success rates that indicates choking under pressure. In each, Large reward neural activity lies best in the optimal preparatory zone for the upcoming movement (purple shaded area). (Left) The “Insufficient Drive”

hypothesis: Jackpot rewards induce paradoxically low reward drive to motor cortex. (*Middle*) The “Neural Bias” hypothesis: increasing offered rewards pushes neural activity toward, but then beyond, the optimal preparatory state for performance. (*Right*) The “Neural Noise” hypothesis: higher reward reduces neural variability to contain more of the distribution in the optimal zone, but the pressure of Jackpot rewards causes dysregulation by increasing variability.

Results

Reward tuning in motor cortex is monotonic, inconsistent with the Insufficient Drive hypothesis

We examined how the magnitude of the cued reward affected firing rates of individual neurons in MC. The Insufficient Drive hypothesis predicts that the main effect of reward on firing rates follows an inverted-U, where the reward signal for Jackpot trials would be low. We examined MC neuron firing rates at the end of the delay period, a time when the animal had information regarding both the target location and potential reward size to be received for a successful trial. Neural signals of anticipated reward have been previously reported throughout the cerebral cortex, with neurons in many brain areas exhibiting changes in firing rates when more valuable rewards are cued^{26–32}, including in MC^{33–36}. However, previous studies have not presented monkeys with rare and exceptionally large potential rewards that induce performance decrements, and thus the nature of the cortical response to these Jackpot rewards is not known. If Jackpot rewards induce individual neurons in MC to fire closer to their Small reward levels than Large, this would support the Insufficient Drive explanation for choking under pressure.

The reward tuning in MC was predominantly monotonic. Most MC neurons exhibited tuning to cued reward ($n = 300/459$ neurons, 65.4%; single neuron metrics are provided in **Table S4**), where most exhibited either monotonically increasing (179/459, 39.0%) or decreasing (95/459, 20.7%) changes in firing rate through the entire range of cued reward size (**Figure 2A-B**). We observed little “inverted-U” (18/459, 3.9%) or “U-shaped” (8/459, 1.7%) reward tuning in individual neurons’ firing rates.

We next considered whether neural signatures of Insufficient Drive might be visible at the population level. We analyzed patterns of covariance in the activity of simultaneously recorded neurons. We define a “neural population state space” by treating the activity of each individual neural unit as an axis in a high dimensional space. In this population space we can identify specific dimensions (i.e., linear combinations of neurons’ firing rates) that capture reward-related signals (**Figure 2C**). Because there were very few Jackpots given per session, we combined neural activity across days using a recently-developed “stitching” algorithm³⁷ (STAR Methods). We then used principal components analysis (PCA) on the neural activity in the analysis bin to identify the linear projection maximizing the amount of reward signal variance captured. We found that a single dimension, which we call the “reward axis,” captured most of the reward-related variance (Monkey E: 92.6%, P: 89.8%, R: 84.7%). Neural population activity projections onto the reward axis were monotonic with reward size (**Figure 2D**). As with the responses of individual neurons, this is inconsistent with the Insufficient Drive hypothesis.

Having identified a “reward axis” in neural population activity in MC, we considered whether neural activity along it might be simply explained by movements. First, we asked whether these reward-monotonic effects might reflect muscular co-contraction. However, electromyography (EMG) from arm and shoulder muscles showed only minimal activity during the reach preparation period, and this activity did not reliably predict reward axis projections (**Figure S1A-B**). The kinematics of the ensuing movement also did not consistently match the patterns of reward axis projections (**Figure S1C-G**). Finally, we identified a “target plane” that maximized variance captured in the average neural activity corresponding to reach preparation to different targets (**Figure 2E**). This plane captured most of the variance due to target direction (Monkey E: 92.7%, P: 99.7%, R: 90.8%). We found that the reward axis and this target plane were near-orthogonal for each animal (**Figure 2F**). This is consistent with the interpretation that effects observed along the reward axis are unlikely to be an artifact due to directional-encoding signals. In sum, we find that the encoding of reward information in MC is not on its own able to explain the performance drop observed for Jackpot rewards.




Figure 2. Reward tuning in motor cortex remains monotonic for Jackpot rewards, ruling out the Insufficient Drive hypothesis. (A) Individual neurons exhibited monotonic tuning to reward size. Firing rates from three example neurons from Monkey E are shown, averaged within each reward condition ($\pm$ S.E.). **Table S4** shows single unit reward tuning statistics. For further analyses, we calculated the firing rate for each neuron and each trial in an Analysis Bin at the end of the delay period from 150 ms before to 50 ms after the Go Cue. (B) MC neurons showed tuning to the target direction, which was largely separable from the reward tuning. (C) Simultaneous neural firing rates can be visualized in a state space in which the firing rate of each neuron corresponds to one dimension (axis) within the space. Three neurons were used here for illustration; in actuality, hundreds of neurons recorded over 6-12 days were used, and neural data was combined across sessions using a neural stitching algorithm (STAR Methods). Using the average neural activity for each reward condition, we can identify a “reward axis” capturing the majority of reward-related variance. (D) Projections along the reward axis are monotonic with cued reward, even though behavior is not. Small dots show single trial values. Large dots show the mean of the reward condition. Horizontal jitter is to aid visualization. Because the primary effect of reward is shifting activity along the reward axis monotonically (not an inverted-U), this rules out the Insufficient Drive hypothesis. (E) Using the average activity for each target direction (dots, color indicates target location), we identify two “target axes” forming a plane that captures the majority of direction-related variance (STAR Methods). For visualization, adjacent reach directions are connected to form a ring. (F) The reward axis is nearly orthogonal to the target plane (STAR Methods, Monkey E: 82°, 93rd percentile of random distribution, Monkey P: 74°, 95th percentile, Monkey R: 71°, 86th percentile).

Jackpot rewards position neural activity in a poor state for upcoming movements, corroborating the Neural Bias hypothesis

We next sought an explanation for choking under pressure by examining how reward signals might interact with reach preparation. Our Neural Bias hypothesis posits that as potential rewards grow from Small to Jackpot, average neural activity is pushed toward, and then away from, the optimal preparatory zone for the upcoming reach. Prior theoretical work has demonstrated that the quality of neural encoding in the preparatory period is important for the quality of the upcoming reach³⁸, a hypothesis supported by some experimental work³⁹. By this account, the optimal preparatory state for the upcoming reach may be one with greatest possible separability from other reach conditions. Hence, to test the Neural Bias hypothesis, we will examine if reward interacts with neural directional encoding signals, then see if these patterns of interaction correlate with behavior.

To look for an interaction between reward and reach preparation signals, we projected neural activity into the 3D space spanned by the reward axis and target plane. Intriguingly, we observed a non-monotonic interaction between reward size and target information. Comparing activity for Small and Large rewards, the average response for different upcoming movement directions grew farther apart from one another with increasing reward (**Figure 3A**). But then, for Jackpot rewards, the neural states for different target directions collapsed back toward each other. To conceptualize the structure we observed, one can picture neural activity on the surface of a prolate spheroid (a rugby ball shape): Neural activity moves along the long axis according to the reward size, while the activity’s location along the orthogonal plane is governed by the intended target.

To quantify this expansion-then-collapse of target-related neural states with increasing reward magnitude, we examined activity on individual trials. Within each reward condition, we identified “target preparation axes” by finding the average response for each target direction (large purple dot in **Figure 3B**) and then calculating the average across the targets (large white dot). We then constructed unit vectors that pointed to each target's average from the average response across targets. We projected the neural activity for each trial onto its corresponding target preparation axis. Like success rates, the average projection along the target preparation axis follows an inverted-U as a function of reward (**Figure 3C**).

To demonstrate that this non-monotonic trend of neural activity along the target preparation axis is consistent with the Neural Bias hypothesis, we must show that preparatory states further along the target preparation axis correspond to better reach preparation (**Figure 3D**). To assess this, we first separated the reaches into successes and failures, and then we categorized failed trials by their specific failure mode. The animals failed by executing reaches that either overshot or undershot the target (STAR Methods). The decrease in success rate between Large and Jackpot reward trials was dominated by undershoot failures (**Figure 3E**)⁹. These undershoot failures can be caused by any combination of slow reaction, slow reach speed, and planning a hypometric reach, and the three animals showed idiosyncratic mixtures of these for both Small and Jackpot trials (**Figure S1C-G**).

If the target preparation axis reflects quality of reach preparation, we would expect successful trials to have a greater projection along it than failed trials. Within each reward condition, we projected neural activity onto the target preparation axis and labeled it according to whether the trial was a success or failure. Within nearly every reward condition, neural preparatory activity prior to a failed trial had a smaller average projection along the target preparation axis than activity prior to a success (**Figure 3F, left**). Interestingly, even successful trials alone exhibited an inverted-U relationship between these projections and reward (**Figure S2C**). This implies that reward is biasing the average neural activity to better (for Medium and Large) or worse (Small and Jackpot) preparatory states for the upcoming reach.

When we split trials by failure mode, a stronger trend emerged: undershoots exhibited an even greater decrease compared to successes in target preparation axis projections (**Figure 3F, middle**). In contrast, there was little difference between projections for overshoot failures and successes (**Figure 3F, right**). This means that when the projection of neural activity onto the target preparation axis was smaller for a trial of a given reward type, the animal was more likely to fail by undershooting the target. Quantifying this across all conditions for each animal revealed that undershoot trials had significantly smaller target preparation axis projections than successes on average (**Figure 3G**). We also found that within direction-reward conditions, smaller target preparation axis projections were correlated with slower reaction time and, to a lesser extent, peak speed (**Figure S4**). It makes sense that neural preparatory activity that is better for the upcoming reach should correlate with faster reactions and reach speeds since this task requires brisk movements. We conclude that neural states further along the target preparation axis provide for better reach preparation. This directly links Jackpot-driven changes in neural activity to sub-optimal reach performance. As such, our data support the Neural Bias hypothesis.




Figure 3. An inverted-U interaction between reward and reach direction preparatory activity supports the Neural Bias hypothesis. (A) Neural population activity corresponding to motor preparation for different reach directions is pushed apart with increasing cued reward from Small through Large. However, for Jackpot rewards, the activity for different reach directions collapses back towards each other, diminishing their discriminability. We projected neural activity averaged separately for each reward and direction condition into a 3D space reflecting reward (Reward Axis) and target information (Target Axis 1 and 2). For visualization, adjacent reach directions (dot color) are connected to form a ring for each reward (line color), and insets (bottom) highlight one target in the target plane. (B) To quantify single trial separability of preparatory states, we found a “target preparation axis” for each reward and target direction (STAR Methods). (C) When neural activity for individual trials is projected onto these target preparation axes, it exhibits an inverted-U as a function of cued reward that parallels how behavior is influenced by reward. Dots represent single trials, and large filled circles show the mean within each reward condition (Welch’s t-test, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, n.s. not significant). We tested several methods for quantification and observed similar results (Figure S2). We note that this same effect of expansion and then collapse in neural encoding with reward can also be observed in single-neuron direction tuning curves (Figure S3). (D) We considered whether a relationship exists between target preparation axis projections (small purple dots) and the quality of reach preparation for individual trials. (E) Success rates (green) and failure rates broken down by failure type (light gray: undershoots; dark gray: overshoots). Compared to Large rewards, Jackpots had more undershoots, whereas Small rewards evoked both more undershoots and overshoots (binomial proportion test). (F) Failed trials showed a consistent decrease in the average target preparation axis projection across animals (shape) and rewards (color) compared to successes. When divided by failure mode, undershoots strongly exhibited the decrease, whereas overshoots showed little difference. Thick lines and accompanying stars indicate a significant difference within that reward condition (Welch’s t-test). (G) To summarize the relationship between target preparation axis projections and failure modes, we pooled across rewards after z-scoring within each (STAR Methods, mean $\pm$ S.E.). Undershoot trials (left, light gray) show a significant decrease in projections along the target preparation axis (Welch’s t-test), while overshoot failures (right, dark gray) show a much smaller effect. Because reward caused a shift in the average neural activity that corresponds with the inverted-U observed in behavior, our data support the Neural Bias hypothesis.

Neural Noise does not explain choking under pressure across animals

Our analyses so far support the view that a neural explanation for choking under pressure is that preparatory activity is poorly positioned with respect to an optimal region in MC’s neural population activity space. We also considered if neural variability could explain choking under pressure, as posited by our Neural Noise hypothesis. Note that our finding of a shift in *mean* neural activity due to reward size does not preclude there also being changes in *variability* with reward; anticipated reward could affect both.

To look for an explanation of choking under pressure stemming from reward-induced effects on variability, we calculated trial-to-trial variability at the population level (STAR Methods). The Neural Noise hypothesis predicts greater noise corresponds to lower success rates. We found inconsistent relationships between neural variability and reward across subjects (Figure 4A-B). For two animals (E and R), we observed a decrease in neural variability with reward, up to and including for Jackpots. In the third (P), we observed a monotonic increase with reward. These findings held whether variability was considered in the full neural population activity space or in the target plane (Figure 4C). As such, while the Neural Noise hypothesis may offer some degree of explanation for Monkey P’s choking under pressure, we do not find evidence of it explaining Monkey P’s Small reward failures, nor the choking behavior observed in the other two animals.




Figure 4. Trial-to-trial variability does not consistently explain choking under pressure, providing evidence against the Neural Noise hypothesis. (A) Example of “noise variance”, or trial-to-trial variability about a condition average, in the target plane. Points represent the average neural activity for each reach direction as seen in Figure 3A. Ellipses indicate the covariance of the single trials about the averages. (B) Total noise variance (STAR Methods). Error bars are S.E. calculated using bootstrapping. (C) Noise variance from the target axes’ projection. This refutes the Neural Noise hypothesis as a consistent explanation of the animals’ choking, as there are not consistent U-shaped trends of noise variance as a function of reward across subjects.

Discussion

We can now describe a potential neural basis for choking under pressure. Reward information interacts with preparation-related activity in the motor cortex. Moderate amounts of reward boost motor performance by driving neural activity closer to an optimal zone of reach preparation. However, when a Jackpot is proffered, neural activity pushes too far along the reward axis, and concomitantly, is biased away from the optimal state in the target preparation subspace for the upcoming reach. This phenomenon can be thought of as a collapse in neural information about the upcoming reach, as the preparatory activity for different movements becomes less differentiated. The ensuing behavior is adversely impacted by this poor preparation, with the result that animals underperform when they should be incentivized to do their best – they choke under pressure.

It was this “Neural Bias” hypothesis which found the strongest and most consistent support in our data. Two other candidate neural mechanisms for choking under pressure were not well supported by our data. “Insufficient Drive” was ruled out by the presence of strong reward signals evident in motor cortex (MC) through the entire range of rewards, continuing monotonically for Jackpots. The “Neural Noise” hypothesis received only weak and inconsistent empirical support, since trial-by-trial variability increased on Jackpot trials in only one animal. Of course, we cannot rule out that such explanations, or others, might pertain in other brain areas or tasks.

Here we offer a mechanistic explanation for choking under pressure in terms of patterns of spiking activity in the cerebral cortex. It was possible for us to do this because it turns out rhesus monkeys are also susceptible to this paradox of behavior⁹. It is well worth attempting to square our findings with the extensive literature on choking under pressure in humans. A few studies in humans have examined the neural mechanisms of choking under pressure using functional imaging. These studies generally show monotonic trends in signal with reward magnitude in various brain regions^{7,8,40}. This aligns well with our finding that the strongest effect of reward on MC spiking activity was a monotonic trend along the reward axis (Figure 2). Also, human studies of choking under pressure have demonstrated changes in functional connectivity between the motor cortex and the prefrontal cortex⁷ or striatum⁸. It is very likely that the changes we observe

in MC are driven by signals from other brain areas like these. Thus an interesting topic for future animal work would be to investigate if the inverted-U we observe in the encoding of reach direction in MC (**Figure 3**) might be caused by changes in how other areas, such as the prefrontal cortex or striatum, are communicating with MC, and/or with how excessive midbrain neuromodulatory drive (e.g., dopamine or norepinephrine) might push cortical activity outside of its optimal operating range⁴¹. It seems reasonable to conclude from the neuroimaging studies that choking under pressure is a whole-brain process, impacting many neural circuits and the communication between them. The neural correlate of choking under pressure that we see in MC is no doubt just one manifestation of how the brain contends with unusually high potential rewards, albeit one that influences behavioral performance in a particularly intriguing way.

Most studies of choking in humans focus on behavior without measuring neural correlates, and they provide explanations of choking in terms of psychological states and drives. Psychological explanations for choking under pressure include excessive self-monitoring^{42,43}, over-arousal (perhaps causing “jitters”)^{4,44–46}, distraction by the high stakes at hand^{2,47,48}, and stress induced by loss aversion^{8,21}. It is important but challenging work to build links between psychological explanations and their underlying neural mechanisms. With the recognition that such links may be very indirect, we wish to propose some speculative connections. Our “Neural Noise” model could be consistent with a “jitters” account of choking, where over-arousal could increase neural variability in MC, yielding detrimental behavioral variability as the result. Our “Insufficient Drive” hypothesis could connect with psychological descriptions such as distraction or loss aversion, whereby attentional resources are drawn away from the task by Jackpot rewards, and thus performance mimics that when Small rewards are proffered, and motivation and attentional engagement are low. Our “Neural Bias” model might link with excessive self-monitoring, where increased focus on the go cue could draw resources away from reach direction encoding or make it harder for the animal to release from its hold state to initiate movement. Our findings support the “Neural Bias” model, and thus perhaps are better aligned with the psychological explanation of over-monitoring, but it would be far too much to claim our study can resolve the issue of which psychological account is the “correct” one. Furthermore, the fact that a psychological explanation for choking under pressure has not already been uniquely identified reinforces the findings from MRI studies suggesting that choking under pressure is a brain-wide phenomenon, and that myriad neural processes, and a wide range of cognitive functions, are adversely impacted by outsized expected rewards.

By providing an animal model of choking under pressure and identifying a neural correlate at the resolution of neural population activity, animal studies can now proceed alongside human studies toward understanding the complex interplay between motivational drive and the fine structure of behavioral performance.

Acknowledgments

We thank Simon Borgognon for feedback on the manuscript, Hongwei Mao for assistance in experimental setup, and Alan Degenhart for contributing to task design. This work was supported by: Achievement Rewards for College Scientists' Pittsburgh Chapter Award (ALS) and Abraham–Martin–Ragni Award (NPP), the Bradford and Diane Smith Graduate Fellowship (ALS), National Science Foundation grants DGE1745016 (ALS), DGE2139321 (PJM), BCS1533672 (SMC,BMY,APB) and DRL2124066 (BMY,SMC) / DRL2123911 (APB), National Institutes of Health grants CRCNS R01NS105318 (BMY,APB), R01NS129584 (APB,SMC,BMY) and R01NS129098 (APB,SMC), and Simons Foundation grant 543065 (BMY).

Author Contributions:

	ALS	PJM	ERO	SES	HM	NPP	WEB	BMY	SMC	APB
Conceptualization	x	x				x		x	x	x
Data curation	x	x	x	x						
Formal analysis	x				x					
Funding acquisition	x	x				x		x	x	x
Investigation	x	x	x	x						
Methodology	x	x			x	x	x	x	x	x
Project Administration	x	x	x	x					x	x
Resources									x	x
Software	x	x					x	x		
Supervision									x	x
Validation	x				x					
Writing (original draft)	x	x							x	x
Writing (review/edit)	x	x	x	x	x	x	x	x	x	x

Declaration of Interest: None to declare

Main figure titles and legends

Figure 1. Potential neural mechanisms of choking under pressure. (A) Monkeys were trained to prepare a brisk reach to a small target. The color (Monkeys E, P) or shape (Monkey R) instructed the reward size. Parameters bolded in green were selected individually for each animal to keep the task challenging and motivating (**Table S1** shows all parameter values). A separate choice task indicated that animals understood reward cues (**Table S2**). Simultaneously, we recorded from primary motor (M1) and/or dorsal premotor cortex (PMd), together termed MC, using 96-channel microelectrode “Utah” arrays (Blackrock Microsystems, Inc., see **Table S1** for array location details). (B) Success rates improved from Small to Large rewards (binomial proportion test, $***p < 0.001$), indicating that performance in this difficult task is influenced by motivation. All three animals choked under pressure, indicated by the significant decrease in success rates from Large to Jackpot. Error bars represent S.E. of overall success rate shown in dark gray. Light gray traces show individual sessions. See **Table S3** for full statistics. (C) We considered three mechanisms that could relate motor cortical activity to the inverted-U in success rates that indicates choking under pressure. In each, Large reward neural activity lies best in the optimal preparatory zone for the upcoming movement (purple shaded area). (Left) The “Insufficient Drive” hypothesis: Jackpot rewards induce paradoxically low reward drive to motor cortex. (Middle) The “Neural Bias” hypothesis: increasing offered rewards pushes neural activity toward, but then beyond, the optimal preparatory state for performance. (Right) The “Neural Noise” hypothesis: higher reward reduces neural variability to contain more of the distribution in the optimal zone, but the pressure of Jackpot rewards causes dysregulation by increasing variability.

Figure 2. Reward tuning in motor cortex remains monotonic for Jackpot rewards, ruling out the Insufficient Drive hypothesis. (A) Individual neurons exhibited monotonic tuning to reward size. Firing rates from three example neurons from Monkey E are shown, averaged within each reward condition ($\pm$ S.E.). **Table S4** shows single unit reward tuning statistics. For further analyses, we calculated the firing rate for each neuron and each trial in an Analysis Bin at the end of the delay period from 150 ms before to 50 ms after the Go Cue. (B) MC neurons showed tuning to the target direction, which was largely separable from the reward tuning. (C) Simultaneous neural firing rates can be visualized in a state space in which the firing rate of each neuron corresponds to one dimension (axis) within the space. Three neurons were used here for illustration; in actuality, hundreds of neurons recorded over 6-12 days were used, and neural data was

combined across sessions using a neural stitching algorithm (STAR Methods). Using the average neural activity for each reward condition, we can identify a “reward axis” capturing the majority of reward-related variance. **(D)** Projections along the reward axis are monotonic with cued reward, even though behavior is not. Small dots show single trial values. Large dots show the mean of the reward condition. Horizontal jitter is to aid visualization. Because the primary effect of reward is shifting activity along the reward axis monotonically (not an inverted-U), this rules out the Insufficient Drive hypothesis. **(E)** Using the average activity for each target direction (dots, color indicates target location), we identify two “target axes” forming a plane that captures the majority of direction-related variance (STAR Methods). For visualization, adjacent reach directions are connected to form a ring. **(F)** The reward axis is nearly orthogonal to the target plane (STAR Methods, Monkey E: 82°, 93rd percentile of random distribution, Monkey P: 74°, 95th percentile, Monkey R: 71°, 86th percentile).

Figure 3. An inverted-U interaction between reward and reach direction preparatory activity supports the Neural Bias hypothesis. **(A)** Neural population activity corresponding to motor preparation for different reach directions is pushed apart with increasing cued reward from Small through Large. However, for Jackpot rewards, the activity for different reach directions collapses back towards each other, diminishing their discriminability. We projected neural activity averaged separately for each reward and direction condition into a 3D space reflecting reward (Reward Axis) and target information (Target Axis 1 and 2). For visualization, adjacent reach directions (dot color) are connected to form a ring for each reward (line color), and insets (bottom) highlight one target in the target plane. **(B)** To quantify single trial separability of preparatory states, we found a “target preparation axis” for each reward and target direction (STAR Methods). **(C)** When neural activity for individual trials is projected onto these target preparation axes, it exhibits an inverted-U as a function of cued reward that parallels how behavior is influenced by reward. Dots represent single trials, and large filled circles show the mean within each reward condition (Welch’s t-test, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, n.s. not significant). We tested several methods for quantification and observed similar results (**Figure S2**). We note that this same effect of expansion and then collapse in neural encoding with reward can also be observed in single-neuron direction tuning curves (**Figure S3**). **(D)** We considered whether a relationship exists between target preparation axis projections (small purple dots) and the quality of reach preparation for individual trials. **(E)** Success rates (green) and failure rates broken down by failure type (light gray: undershoots; dark gray: overshoots). Compared to Large rewards, Jackpots had more undershoots, whereas Small rewards evoked both more undershoots and overshoots (binomial proportion test). **(F)** Failed trials showed a consistent decrease in the average target preparation axis projection across animals (shape) and rewards (color) compared to successes. When divided by failure mode, undershoots strongly exhibited the decrease, whereas overshoots showed little difference. Thick lines and accompanying stars indicate a significant difference within that reward condition (Welch’s t-test). **(G)** To summarize the relationship between target preparation axis projections and failure modes, we pooled across rewards after z-scoring within each (STAR Methods, mean $\pm$ S.E.). Undershoot trials (left, light gray) show a significant decrease in projections along the target preparation axis (Welch’s t-test), while overshoot failures (right, dark gray) show a much smaller effect. Because reward caused a shift in the average neural activity that corresponds with the inverted-U observed in behavior, our data support the Neural Bias hypothesis.

Figure 4. Trial-to-trial variability does not consistently explain choking under pressure, providing evidence against the Neural Noise hypothesis. **(A)** Example of “noise variance”, or trial-to-trial variability about a condition average, in the target plane. Points represent the average neural activity for each reach direction as seen in **Figure 3A**. Ellipses indicate the covariance of the single trials about the averages. **(B)** Total noise variance (STAR Methods). Error bars are S.E. calculated using bootstrapping. **(C)** Noise variance from the target axes’ projection. This refutes the Neural Noise hypothesis as a consistent explanation of the animals’ choking, as there are not consistent U-shaped trends of noise variance as a function of reward across subjects.

STAR Methods

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Steven Chase (schase@andrew.cmu.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The data that support the findings of this study are available from the lead contact upon reasonable request. All original code is publicly available as of the date of publication (see Key Resources Table). Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Three adult male rhesus macaques, Monkeys E (9.0 kg, 10 years old), P (9.5 kg, 6 years old), and R (19.0 kg, 10 years old) were implanted with either 96 or 64 electrode arrays (Blackrock Microsystems) in the proximal arm region of the primary motor cortex and/or dorsal premotor cortex. All experimental and animal procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee in accordance with the guidelines of the US Department of Agriculture, the International Association for the Assessment and Accreditation of Laboratory Animal Care, and the National Institutes of Health.

METHOD DETAILS

Experiments and behavioral recordings.

Three monkeys were trained on delayed reaching tasks. We used standard water regulation procedures to maintain motivation and the valuation of reward. During experiments, each monkey sat in a primate chair facing a mirror ~8 cm in front of his eyes that reflected a computer monitor displaying task events. The monkeys performed all tasks by making hand movements (right arm for Monkeys E and R, left arm for Monkey P) in an open space in front of them. While the working arm and hand were unrestrained, the hand was not visible to the animal, as it moved in the space behind and below the mirror. 3D hand position was tracked with an infrared LED marker attached either to the monkey's index finger (Monkeys E and P, 120 Hz sampling rate, nominal resolution < 1 mm, PhaseSpace, Inc) or the back of the hand (Monkey R, 60 Hz sampling rate, nominal resolution < 1 mm, Optotrak 3020, Northern Digital Instruments). The monkey's hand movements corresponded to a cursor position displayed on the monitor. The software environment was calibrated such that 1 cm of hand displacement in the coronal (frontoparallel) plane corresponded to 1 cm of cursor movement. Any trials where tracking of the hand failed at any point in time were removed (fewer than 1% of trials).

For post-hoc analysis, we smoothed the hand position signals with a zero-phase low pass Butterworth filter (4th order, cutoff frequency 15 Hz). We then also calculated velocity by taking the first difference of the position data, dividing by the time difference between samples, and assigning each velocity sample a timestamp that was the midpoint of the timestamps for the two position samples used. We then used spline interpolation for both the position and velocity signals to upsample to 1000 Hz at identical timepoints. For Monkey R, we also simultaneously

recorded surface electromyography (EMG) from shoulder and arm muscles (see Electromyography section below).

Main task.

The main task performed by the animals in this study was a challenging delayed center out reaching task (**Figure 1A**) which has been described in detail in Smoulder et al. 2021⁹ (referred to in that work as the “speed + accuracy task”). Specific task parameters for each animal are listed in **Table S1**. We note that nearly all of the data in this study are new behavioral recordings, with only Monkey E’s behavioral data here being a subset of the data reported in our previous work⁹.

All trials began with a circular target appearing at the center of the display. The animal had to move its hand to position the cursor in the center target to initiate a trial. After a short period of time (ranging from 200-600 ms, different for each animal), a reach target appeared at one of several possible locations (Monkeys E and R: 8 possible target locations, Monkey P: 4) and remained visible through the rest of the trial. The appearance of the reach target began a delay period during which the animal could see their hand position cursor, the reach target, and the center target, but was not allowed to let the cursor exit the center target. The duration of this delay period was drawn randomly from a uniform distribution each trial (ranging from 200-950 ms, different distributions for each animal). After this duration, the center target disappeared from the screen, which cued the animal to move the cursor to the reach target (i.e., the go cue). The animal had only a brief amount of time after the center target disappeared to react and acquire the reach target (range across animals of 667-825 ms). If the animal moved the cursor into the reach target in time, he had to hold the cursor within the target for 400 ms, after which he received a liquid reward. The allowed reach duration and target size were titrated during training for each animal to make the task challenging, and then those values were maintained throughout experiments.

The magnitude of the liquid reward that would be dispensed upon successful completion of the trial was indicated by the color of the target cue (Monkeys E, P) or by an image within the target (Monkey R). Four reward sizes were used for each animal and were drawn randomly every trial: Small, Medium, Large, and Jackpot (magnitudes and cues differed by animal; see **Table S1**). Jackpot rewards were presented on only 5% of trials (which resulted in a range of 1-12 Jackpots proffered per reach direction per day), while the other rewards were presented with equal frequency for the remaining trials. For purposes of visualization in this work, we use the color scheme indicated in **Figure 1A**: Small is red, Medium is orange, Large is blue, and Jackpots are black (see **Table S1** for actual cue colors and shapes).

For Monkeys P and R, a small proportion of trials (5%) were “catch” trials, where the trial abruptly ended at the time of go cue and the animal was rewarded without making a reach. These were used to discourage “false start” reaches before the go cue. Catch trials were randomly interspersed throughout the session and could be of any cued reward and direction. The reward provided on catch trials corresponded to the cued reward.

Two target choice task.

To test whether the animals understood the reward cues, we instructed them to perform a two-target choice. This task proceeded in the same way as the main task with two differences: (1) two reach targets were available for the animal to select, diametrically opposed, and (2) the animal had a slightly longer amount of time to reach to their selected target. Target direction was selected randomly from the working animal’s target set used in the main task. Reward values were typically selected randomly each trial with equal probability, with the exception that Jackpot rewards were available for selection only on 5% of trials and the Jackpot reward was never presented at both available targets. Upon trial success, the animal would receive the reward associated with the target he selected.

We then evaluated the animals' selections in each combination of reward presentations to confirm the animals selected the greater reward when presented. Collectively, animals selected the larger reward on 98.9% of trials (**Table S2**, see Behavioral training and recording sessions below). Trials where both targets had the same reward cue were excluded from this analysis.

Trial failure

In both the main and two-target task, the animal failed the trial if he moved the hand position cursor outside of the center target during the delay period, did not acquire the reach target in time, or exited the reach target before 400 ms had elapsed after acquiring it. Upon failure, the cursor and reach target turned purple and the screen froze for a duration exceeding the maximum length remaining for a successful trial (i.e., if the animal failed during the delay epoch, the freeze time was the remaining time for the longest delay period, plus the maximum reach time, plus the 400 ms target hold). After this, a purple reach target appeared at a location other than that of the failed target. The animal had to reach to the purple target, for which no reward was given, before the next trial could begin. These "punishment" reaches were included to discourage the animal from aborting Small reward trials, as it makes task failure objectively worse than attempting the trial, even for the least rewarding (e.g., 0 mL) cases. After either a punishment reach was completed (for failed trials), or a reward was dispersed (for successful trials), an intertrial interval elapsed before the next trial began (Monkeys E and P: 500 ms, Monkey R: 600 ms). For Monkey R, intertrial intervals after successful Jackpot trials were extended to 1200 ms to give additional time for the animal to consume the larger reward.

Behavioral training and recording sessions.

All animals followed the same general training procedures prior to performing the task for the recordings used in this study. Animals were first trained to proficiently perform a delayed reaching task (> 90% success rate). Then, reward cues for Small, Medium, and Large were introduced into this task. We ran sessions of the two-target choice task after the main task each day to assess understanding of the relative reward cue values. After at least one week of performing the same task with these reward cues, we began making the task more challenging. We titrated the amount of time the animal had to acquire the reach target and the reach target size to reduce average success rate to about 70%. Making the task challenging enabled us to better measure if cued reward either improved or hurt success rate. If the difference in success rate between Small and Large rewards was small (i.e., less than 10%), we expanded the reward range by shrinking the Small reward magnitude and increasing the Large reward. We expanded the Small-Large reward range in this manner for two animals; Monkey R before introduction of Jackpot rewards, and Monkeys E after Jackpots had been introduced (see ⁹).

Once the animals indicated understanding of the relative value of the reward cues in the choice task and parameters were adjusted to the appropriate task difficulty, we introduced the Jackpot reward cue. For each animal, we ran the main task for one session, then the next day split the session between the main task and the two-target choice task to assess understanding of the Jackpot cue. Monkey E indicated full Jackpot cue understanding, after which we performed experiments. After several sessions of this, we then expanded Monkey E's Small-Large reward range and ran the choice task each day after the main task. These sessions with the expanded Small-Large reward range are what is reported in this study and are shown in **Table S2** (see ⁹ for information on the initial main task sessions). Monkey P indicated full Jackpot cue understanding during this session, after which we halted Jackpot behavioral experiments until after array implantation. After recovery, we ran a session of choice task including Jackpots to ensure Monkey P still understood cue values (the data of which are shown in **Table S2**), after which we collected data for the main task. Monkey R did not immediately select the Jackpot cue when presented, potentially because of a more challenging cue to learn and distinguish (the cue was a specific image, as opposed to a color). Additionally, his performance improved in the first 2-3 sessions after Jackpot reward introduction. We adjusted the parameters of the task to account for these behaviors (we shortened the reach period maximum time but slightly increased reach target diameter and increased cue contrast). After this, we repeated the procedure of one session with all four reward cues and one session of split main and choice task, during which the animal

indicated understanding of all reward cues. We began neurophysiological recordings the following day. The data from this session of the choice task and additionally run choice task sessions after a few days of the main task are reported in **Table S2**.

Electromyography.

For one animal (Monkey R), we recorded surface electromyography (EMG) on various muscles of the working arm and its shoulder.

Surface EMG recordings.

We used pairs of disposable adhesive electrodes for bipolar recordings (3M, model 2560), placed with a dab of conductive gel (Spectra 360) above muscles on shaved regions of the working arm (lateral biceps, triceps) and shoulder (anterior deltoid, posterior deltoid, trapezius). Anecdotally, shoulder muscle signals were more salient than those from arm muscles during experimental recordings. We also placed a pair of electrodes across the chest to record electrocardiogram (ECG) signals, with one electrode on the center of the chest and the other on the left upper abdomen. A ground electrode was placed on the skin of the lower abdomen, under which the animal had a fair deal of adipose tissue such that muscular signals were minimal during the task. Signals were sent through a differential amplifier (RA16LI-D, Tucker Davis Technologies (TDT), Inc.), then a pre-amplifier (PZ2 Pre-Amp, TDT), then to a signal processor that digitized signals (nominal 48KHz sampling frequency, RZ2 BioAmp Processor, TDT). Signals were then fed through a 2nd order 10 Hz biquad high pass filter, downsampled to 10KHz and stored. Two minutes of data were collected immediately preceding the start of each experiment while the animal rested to identify baseline EMG and ECG.

Surface EMG preprocessing.

For analysis, EMG and ECG data were zero-phase filtered using a comb of notch filters to remove line noise and a bandpass Butterworth filter to attenuate movement artifacts and high-frequency noise. The comb filter was 2nd order, with cutoff frequencies of 60 Hz and its harmonics up to 480 ± 1 Hz. The bandpass filter had cutoff frequencies of 20 and 500 Hz with stopband attenuation of -60 dB, implemented using MATLAB's "bandpass" function. ECG artifacts were visible in the EMG signals during recording, particularly during the two-minute baseline period. To remove these, we first fit a static, linear transfer function from the ECG data to each of the EMG recordings using the two-minute baseline period. We used a moving 50 ms window and used ordinary least squares to fit a transfer function from ECG signals to the EMG from each channel. We then applied this transfer function to the entire ECG recording, producing a predicted ECG artifact signal for each EMG channel. This predicted ECG signal was subtracted out for each channel. While we performed this as a data cleaning step, we note that results did not qualitatively change if ECG artifact removal was not performed. After this operation, signals were rectified and downsampled to 1000 Hz.

We wanted to combine EMG recordings across days. Because these monkeys are overtrained on the task, we expect that the ground-truth average EMG activation of a given muscle during reaching will be the same across days⁴⁹. However, due to differences in contact quality and exact electrode placement, the actual recording voltage range varies day-to-day. To combine signals across days, we first took the EMG power in a 200 ms window surrounding peak cursor speed (see Behavioral analysis) for each trial and each muscle, then averaged within each reach direction condition. This effectively made an EMG directional "tuning curve" for each muscle. We also added an additional point to this tuning curve to include EMG in the 200 ms period before target onset while Monkey R held the cursor at the center target. We thus had 9 data points per muscle per day. For each muscle each day, we calculated the mean and the standard deviation of the tuning curve across these 9 points. We then normalized all EMG data for that muscle/day by subtracting this mean and dividing by the standard deviation, effectively z-scoring so that the units for EMG for each day were comparable.

Neural recordings and preprocessing.

We recorded neural activity from the primary motor cortex (M1) and dorsal aspect of the premotor cortex (PMd) using multielectrode “Utah” arrays (Blackrock Microsystems). Monkey E had one 96-electrode array straddling the shoulder regions of M1 (~64 channels) and PMd (~32 channels), implanted approximately 1.5 years prior to these experiments. Monkey P had two 64 electrode arrays, one in the arm area of M1 and one in the arm area of PMd, implanted two weeks prior to these experiments. Due to equipment constraints, only 96 electrodes could be recorded simultaneously at the time of the experiments, so only half of the M1 electrodes were used. Monkey R had two 96 electrode arrays in the anterior and posterior aspects of M1, in the shoulder region, implanted approximately 4.5 years prior to these experiments. Recordings on most electrodes from both regions showed modulation during the late delay period analyzed here, and thus in this study we make no distinction between recordings from M1 versus PMd, referring to the combined site of recordings as motor cortex (MC). We include all recordings for all analyses.

Before each experiment began, we set voltage thresholds on each electrode at a multiplier of the root-mean-square voltage (Monkey E: -3.5x, P: -3x, R: -3x). We stored 1 ms waveform snippets (sampling rate 30 KHz; 30 samples per snippet) surrounding each threshold crossing during the experiment. Individual units were manually sorted on each electrode using offline spike sorting (Plexon). Units were identified visually using a combination of features, including waveform principal components and peak-to-trough amplitude. We included both well-isolated units and multi-unit waveforms for this study, yielding each day 145 ± 4 units for Monkey E (mean $\pm$ std across sessions), 183 ± 86 units for Monkey P, and 273 ± 25 units for Monkey R.

QUANTIFICATION AND STATISTICAL ANALYSIS

Behavioral analysis.

Success rate and failure mode analysis.

We analyzed success rates for the main task as a function of reward cue (**Figure 1B**). Average success rates were pooled across all sessions (black lines), and also calculated for individual sessions (gray lines). For error bars, we performed a bootstrap analysis in which we randomly sampled (with replacement) the trials within each reward size and calculated the success rate 10,000 times. Error bars are the standard deviation of these 10,000 bootstrap samples.

We analyzed the way the animals failed trials in the same manner as previous work⁹. For this study, we only considered trials where the go cue occurred (i.e., no failure during the delay period). Beyond this, all successful and failure trials were considered for behavioral and neural analyses. Once the go cue occurred, two types of reach failures could occur: overshoots and undershoots (**Figure 3E**, right). First, the animal could reach past the target, miss it, and run out of time prior to being able to make a correction, an “overshoot.” We also considered trials where the cursor blew through the target as overshoot failures. Second, the animal could “undershoot” by running out of time before acquiring the target. Undershoots could occur either due to a reach landing short of the target and the animal having insufficient time for corrective movements (approximately 25% of undershoots), or due to reacting and/or reaching slowly and running out of time mid-reach (approximately 75% of undershoots). Although these two types of undershoots are distinguishable based on kinematics, we reasoned that both could correspond to a failure to plan well, and thus we combined them for analyses. It could also be that reaches with high angular error could feasibly not fit easily into either category. However, when we calculated error from a straight line between the center and reach target, only one trial out of the total 15,721 exceeded 25 mm (this trial was excluded). The rest were generally straight towards the target and could fall into the classification of going past the target (overshoot) or coming up short (undershoot).

If the cursor stopped successfully in the target, it was considered a reach epoch success (a success for the purposes of **Figure 3E**); however, the animal could still fail if the cursor drifted out of the reach target during the target hold period (a failure for the purposes of overall success rate in **Figure 1B**). The specific parameters used to decide which label each failed trial was given are described in ⁹. We note that the failure mode trends in this study for comparing Small and Jackpot reward trials to Large reward trials were consistent with the results from ⁹.

Single trial behavioral metrics.

We report five kinematic metrics as a function of reward in **Figure S1C-G**. To find peak speed for each trial, we first calculated the cursor speed as the square root of the sum of the squared horizontal and vertical velocity signals. We calculated the peak speed of each trial as the max cursor speed occurring after the go cue. We calculated the reaction time as the time from the go cue that it took for the cursor to reach 20% of the animal's overall average peak speed. In this way, reaction time was calculated at the crossing of a static speed threshold for each animal.

We calculated the homing time and ballistic endpoint predictions in the same way as previous work⁹. Homing time was calculated starting at the time the cursor was within the last 1/3 of the distance to the target and ending when the cursor was within 1 mm of the target edge. Homing time trends were not sensitive to this exact start or end point. Ballistic endpoint predictions were calculated by first identifying the locations of the cursor 50 ms before reaction time and at the time of peak speed. The displacement between these points was calculated and doubled. This is mathematically equivalent to mirroring the velocity profile in time about peak speed and integrating, producing a prediction of where the reach would have landed with a symmetric, ballistic velocity profile. We report this quantity specifically projected along the vector connecting the center target to the reach target.

The off-axis error was the cursor location of at the time of peak speed along the axis orthogonal to the vector connecting the center and reach target for each trial (e.g., for a rightward reach, the peak speed position along the vertical axis). We defined positive values as being counterclockwise of the center to reach target vector. We calculated the standard deviation of this quantity (which we call the off-axis deviation) as a measure of the angular accuracy of the reach. The off-axis deviation was calculated for each reward-direction condition, and we report values averaged across directions. We calculated error bars via bootstrapping: within each reward-direction condition, we resampled with replacement and calculated the off-axis deviation, then averaged across reach conditions. We repeated this procedure 10,000 times and used the standard deviation of this distribution for error bars.

Surface EMG analysis.

We performed two analyses on Monkey R's surface EMG signals. First, we assessed the overall quality of the EMG signals by visualizing activity for each reach direction (**Figure S1A**). For all successful trials, we smoothed rectified EMG signals with a 50 ms boxcar filter and extracted three windows aligned to target onset, go cue, and target acquisition. We then averaged across all trials within each reach direction and plotted activity as a function of reach direction for each muscle.

Second, we assessed EMG during the delay period time bin coinciding with the neural analyses in the main text ([-150 50] ms around the time of go cue) to determine if muscle activation could explain changes in neural activity with increased cued rewards. For any trial that had a go cue (i.e., no delay failures or catch trials) we calculated the average EMG signal in this time bin to get one value per muscle per trial, then grouped trials by their reward and direction labels. We then performed a two-way ANOVA to evaluate statistical significance of reward and directional tuning in the EMG signals (**Figure S1B**).

Neural data analysis.

In this work, we focused on steady-state neural activity during reach preparation, a 200 ms period starting 150 ms before each trial's go cue and ending 50 ms after. As 50 ms is less than the time it takes for visual information to reach MC⁵⁰, we consider all of this activity to fall within the time the animal was prepared to move but had yet to begin reacting to the go cue. We calculated the averaged firing rate within this bin. For the visualizations of **Figure 2A**, we made peri-stimulus time histograms for the delay period using spike times convolved with a Gaussian kernel (25 ms standard deviation). All neural analyses performed were additionally tested after subsampling trial counts of Small, Medium, and Large rewards to match those of Jackpots. These are not reported, as beyond the decreased statistical power associated with lower sample counts, we observed no differences from the results presented.

Monkeys P and R experienced a subset of trials with very short delay periods (< 400 ms). To assess if neural activity sufficiently reached a steady state for these short-delay trials, we calculated the average distance of the neural state from a steady state endpoint (~450 ms after target onset) as a function of time after target onset for non-short delay trials¹⁹. We found that Monkey E's trajectories reached 90% of the average distance by 220 ms, P by 197 ms, and R by 317 ms. We excluded trials with delays shorter than this for all analyses. We note that this procedure only removes trials for Monkey R (shortest delay 200 ms), as Monkey E and P's shortest delay period is longer than this. Using a static threshold of trials with delay length greater than 400 ms did not meaningfully affect results.

Single unit analysis.

For single unit analyses (**Table S4**, **Figure S3**), we only included units that were present for at least 10 presentations of Jackpot rewards for each reach direction (leaving 42 units for Monkey E, 113 units for Monkey P, and 304 units for Monkey R). This typically required the unit to be present across at least 2-3 sessions of data collection (see *Combining neural data across sessions*). To analyze tuning as a function of reach direction and reward, we used a 2-way analysis of variance (ANOVA) on the binned firing rates for each identified unit. If the ANOVA indicated significant effects of reward on firing rate, we used Tukey's test with multiple comparisons correction (the "multcompare" function in MATLAB) to assess significant differences between reward conditions, from which we could evaluate the shape of the reward tuning curve. For this analysis only, we combined Medium and Large rewards due to their general similarity in effects on firing rate and to simplify the operation of evaluating reward tuning curve shape. Based on the pattern of the differences between the Small, Medium/Large, and Jackpot reward conditions, we identified 9 possible tuning curve shapes, as described in **Table S4**. We categorized each unit into one of these 9 shapes and tallied their results into **Table S4**.

We also assessed how directional encoding in single unit activity was affected by reward. For each unit, we fit directional tuning curves to average firing rates for each reach direction using linear regression with a sinusoidal basis set with period of 360 degrees as done in previous work⁵¹. Modulation depth is the amplitude of the tuning curve (**Figure S3A**). We plotted modulation depth for the units with the strongest directional tuning for each animal (**Figure S3B**). To combine values across units, we first focused on only units in the top quartile of average modulation depth value, as there are many motor cortical units with little-to-no delay tuning. Our results were robust to other reasonable fractions of units used (i.e., 1/5, 1/3). We then calculated the average modulation depth for each reward across this fraction of units (**Figure S3C**). To calculate error bars, we first subtracted the average modulation depth across rewards from each unit, then calculated the standard error of the mean within each reward.

We also evaluated how trial-to-trial variability was affected by reward. For each unit x reward x direction condition, we calculated the fano factor as the variance of firing rates divided by the mean (**Figure S3E**). This effectively creates a variability metric that controls for the signal-dependent noise inherent to Poisson-like neural firing rates⁵². Because our hypotheses related to neural noise were related to the entire neural population, we included all units for this analysis, though we note that our interpretations of results (i.e., not seeing a "U-shape" in noise as expected by the Neural Noise hypothesis) did not change when the same subset of units used in

the modulation depth analysis were considered. Otherwise, we averaged fano factor across reach directions for each unit and direction, and we calculated the average fano factor (and S.E.M.) across units the same way as described for modulation depth.

Combining neural activity across sessions.

Jackpot rewards occurred rarely (5% of trials), meaning that single sessions have a low quantity of Jackpot trials for a given reach direction. Combining across days is necessary to gain the statistical power sufficient to analyze the Jackpot trials. However, due to electrode recording instabilities, not all units are conserved session-to-session, and few neural units are present across all recording sessions (21 units for Monkey E, 3 units for Monkey P, and 6 units for Monkey R; see below for details about how this was assessed). Statistical algorithms can be used to “stitch” neural recordings across sessions^{53–57}. In this work, we combined neural activity across days using a modified version of the stabilization algorithm from Degenhart, Bishop et al. 2020³⁷, described below. It requires identification of common units across recording sessions but does not require consistent task or trial structure across sessions. Unless stated otherwise, all further methods and analyses were performed using the latent factors identified by this neural stitching method.

First, to identify units that were present in recordings across different sessions, we used the method described in⁵⁸. In brief, we calculated four quantities for every sorted unit using the first 100 trials of each session: 1) average waveform shape, 2) average firing rate, 3) spiking autocorrelation, and 4) spiking cross-correlations with all other sorted units from that session. We assumed that units recorded from different electrodes were not the same across days due to the physical distance between electrodes. Accordingly, to identify units putatively conserved across days, we calculated these four similarity metrics for units recorded on the same electrodes on different days as described in⁵⁸. To evaluate the stability of the units across days, we compared the distribution of values for same-channel-different-day to a “different” distribution that we created by computing these four values across units recorded on different electrodes on different days. If similarity scores computed for two units from the same electrode was below a selected false-positive threshold percentile in the “different” distribution (99th percentile for Monkeys E and R, and 95th percentile for Monkey P), the units were considered the same.

Using this procedure, we found a high number of units tracked over different pairs of recording sessions, primarily across consecutive sessions. For example, the minimum number of units present across two consecutive sessions was 71 for Monkey E, 28 for Monkey P, and 67 for Monkey R. Units could be present across non-consecutive sessions (i.e., present sessions 1 and 3 but not session 2). However this was rare and nearly exclusively occurred for Monkey P, whose array was implanted shortly preceding the study and, as such, exhibited greater amounts of instability in recordings.

Having identified and tracked units across sessions, we organized the neural activity into a $n \times T$ matrix of binned spike counts, where n is the number of unique units recorded (i.e., the union of units across all sessions), and T is the total number of trials across all sessions. If unit i was present during trial t , that element of the matrix was filled in with the spike count. Otherwise, that element was left empty.

To identify a common latent space across sessions (the T trials), we used factor analysis (FA). We chose FA because it is the most basic dimensionality reduction method that seeks to preserve variance shared across units. FA relates neural activity for trial t , $\mathbf{x}_t \in \mathbb{R}^n$, to latent factors, $\mathbf{z}_t \in \mathbb{R}^m$ ($m < n$) according to:

$$\mathbf{x}_t | \mathbf{z}_t \sim N(\Lambda \mathbf{z}_t + \boldsymbol{\mu}, \Psi) \quad (1)$$

$$\mathbf{z}_t \sim N(\mathbf{0}, I) \quad (2)$$

where $\Lambda \in \mathbb{R}^{n \times m}$ is the loading matrix that defines the relationship between latent factors and spike counts, $\boldsymbol{\mu} \in \mathbb{R}^n$ is a vector of mean spike counts for each unit, and $\Psi \in \mathbb{R}^{n \times n}$ is a diagonal matrix

that describes the variability independent to each unit. The latent factors $\mathbf{z}_t$ capture the variance that is shared across units.

To fit the FA model parameters Λ , Ψ , and μ , we used the expectation maximization (EM) algorithm. Notably, there are empty entries in the spike count vectors, $\mathbf{x}_t$, as described above. We treat this as a missing data problem, where units that are not present for a given session are treated as missing observations. The EM algorithm can handle missing data seamlessly by maximizing the probability of only the data entries that were observed. By fitting a single FA model across the T trials, we effectively “stitch” the neural activity across experimental sessions to identify a single common latent space. This method is described in ⁵⁶, with theory that can be used to derive how many units need to be in common across recordings for successful stitching in ⁵³. Other methods are also available for neural stitching, for example those that leverage the time course of activity^{55,57} or trial labels such as reach direction⁴⁹. Note that the stitching method we use here does not require trial labels; instead, it leverages the identification of stable units across (a subset of) sessions.

To determine the number of factors m , we used 5-fold cross validation to fit FA models ranging from 1 to 40 factors. We selected the number of factors m^* that maximized the cross-validated data likelihood (Monkey E: 16 factors, P: 5, R: 20). We then fit a single FA model for each animal across all observed activity $\mathbf{x}_t$ using m^* factors. This procedure was performed separately for each animal, yielding a single “stitched” FA model per animal. To facilitate visualizations and analyses involving the latent factors, we orthonormalized the columns of the loading matrix Λ and ordered the factors (i.e., the elements of $\mathbf{z}_t$) by the amount of shared variance explained⁵⁹.

To validate the stitching procedure, we assessed the stability of neural direction and reward representations across days. We note empirically that projections along the reward and target axes (see below) were highly consistent across days; more rigorous validation of the algorithm itself is available in ⁵⁶. Whenever possible, we reproduced results found in the stitched latent space with analyses on single units (**Table S4, Figure S3**).

We refer to the latent factors simply as “neural activity” for the remainder of the methods and in the main text from Figure 2D onward.

Reward axis calculation.

To find a population-level signature of reward encoding in the motor cortex, we identified a signal that captured reward-related variance across the population. We first calculated average neural activity from a single time bin in each trial (“Analysis Bin” in Figure 2A, from 150 ms before to 50 ms after the Go Cue). We then grouped trials by their direction-reward condition (e.g., leftward-Large) and averaged across trials within each condition. We then averaged across reach directions, yielding a number-of-reward-conditions by number-of-factors matrix (e.g., for Monkey E, a [4 x 16] matrix). We performed principal components analysis (PCA) on this matrix to identify the dimensions (of length 4) explaining the most reward-related variance. A single PC explained the overwhelming majority of the reward-related variance for each animal (E: 92.6%, P: 89.8%, R: 84.7%); because of this, we dubbed this PC the “reward axis”. Since PCA output is only unique to a sign-flip, we selected the sign of these PCs such that Large reward trials had greater average projections than Small reward trials.

We projected all trials along this single dimension for each animal (**Figure 2C-D**). We note that there was no part of the PCA objective that encouraged monotonic trends for Jackpot rewards; that is, if the main reward-related variance in the data exhibited an inverted-U, it would have come out of this procedure. We note that we found near-identical reward axis results when we used other algorithms to identify it, such as linear discriminant analysis (LDA) with reward labels or linear regression of neural activity to categorical reward size (i.e., $S = 1$, $M = 2$, $L = 3$, $J = 4$). We also saw similar results if we only fit the reward axis using Small/Medium/Large trials and projected all trials down, effectively extrapolating Jackpot responses.

For **Figure S4**, we also considered the second PC from the method described that explained nearly all of the remaining reward signal variance for each animal (E: 6.4%, P: 8.5%, R: 15.8%). We again aligned this projection such that the average projection for Large reward trials was higher than that of Small reward trials.

Target axes calculation.

We wanted to study how reward interacted with the representation of target direction during movement preparation. To determine the dimensions in neural space with the most target-related variance, we once again calculated the trial-averaged activity within each direction and reward condition, then marginalized across rewards to produce a number-of-directions by number-of-factors matrix for each animal. We then performed PCA on this to find the linear combinations of factors that capture the greatest variance about target response. We found that the top two PCs explained the overwhelming majority of this variance (E: 92.7%, P: 99.7%, R: 90.8%). (Note that this large amount of variance explained by two dimensions for a static time bin of neural data is not particularly surprising: the number of points used is equal to the number of directions, and accordingly the number of dimensions needed to fully account for these points' variance is the number of targets minus one. Further, the target positions were embedded in a 2D space in the coronal plane, meaning only two dimensions are needed to fully define the target locations.) We called the axes of this 2D space "target axes". For ease of visualization only (**Figure 2E, 3A, 4A, S3D**), after projecting to this 2D plane, we aligned all animal's spaces such that the average neural activity for the up-rightward reach target was pointed along the first quadrant diagonal unit vector (meaning that target axis 1 is positive for rightward targets, while target axis 2 is positive for upward targets).

Projecting to 3D space of target/reward axes.

The target axes are orthogonal in the latent space by definition when using PCA. However, as the reward axis and the target axes were found in separate steps, there were no constraints to the alignment of the reward axis with respect to the target axes. We assessed how aligned the reward and target axes were for each animal (**Figure 2F**). We first calculated the angle between the reward axis and the target plane (black vertical line). We compared this value against two null distributions: 1) random vectors, and (2) aligned vectors. For the first of these, we drew random vectors in the full dimensional space of the factor analysis model and measured their angle to the target plane. We did this 10,000 times to construct a distribution of angles between random vectors and the target plane (dark gray, labeled "Random"). For an aligned vectors distribution, we split the data into two halves 1,000 times, randomly apportioning equal numbers of trials for each reward and direction condition. We then found the 2D plane spanned by the target axes in one of the halves, and in the other identified the 1st target axis. We calculated the angle between these (light gray, labelled "Same"), reasoning that the target axis in one half should be aligned with the target plane calculated from the other half of the data. Thus, this distribution reflects what the variation in angle measurements might be between a vector and an aligned 2D plane that is due to sampling error. The angles found between the reward axis and target axes' plane were on the upper end of the "random" distribution, far above the "same" distribution, and close to orthogonal. This permits orthogonalization of the reward axis to the target plane with minimal distortion of reward axis projection values.

Given this near-orthogonality, for visualization (**Figure 3A, Figure S3D**), we orthogonalize the reward axis with respect to the 2D target plane using QR decomposition. We projected all trials into this 3D space, then calculated the average activity within each direction and reward condition (**Figure 3A**). We reproduced the results shown in **Figure 3A** using other algorithms to identify the target axes, including LDA with target direction labels and linear regression of neural activity to spatial target location (not shown due to space constraints but is available upon request).

Target preparation axes.

We saw qualitatively that the separation of average activity for different upcoming reach directions (i.e., the radius of the target ring) was modulated as a function of cued reward in **Figure 3A**. To quantify this effect, we performed an analysis on single trials calculating how close

or far the neural preparatory state was from the origin of the target plane. To do this, we calculated “target preparation axes” within each reward and direction condition as follows (**Figure 3B**): First, we calculated the average projections for each reward and direction condition in the target axes’ plane. Then, for a given reward condition, we calculated the average across directions. Intuitively, this is the center of the ring of targets for the given reward shown in **Figure 3A**. Then, for a given reward-direction condition, we calculated the vector connecting this ring center for the given reward to the average activity for the current reward-direction. We then projected all trials’ activity for that condition along the vector. Higher values indicate greater distance from the ring center along this vector. Lower values indicate activity closer to (or beyond) the ring center. We performed this procedure for each direction-reward condition. To combine values across direction conditions while maintaining the relationship between reward condition values, we pooled values across rewards within a given reach direction, then z-scored, such that the target preparation axis values for each reach direction overall had zero mean and unit standard deviation (**Figure 3C**). Results from **Figure 3C** do not meaningfully change when we use other algorithms to identify the target axes before calculating the target preparation axis (**Figure S2A**). We also note that similar results can be observed using offline decoders of reach direction on the delay period neural data from each reward condition (**Figure S2B**), and that these results are also observed even if only successful trials are included for analysis (**Figure S2C**).

We hypothesized that neural activity with smaller projections onto the target preparation axes indicated more poorly-prepared reaches, and hence, would more likely be failures (**Figure 3D**). To test this, we calculated the average target preparation axis projection (after combining across directions as in **Figure 3C**) for each reward separately for successes and failures (**Figure 3F**, left), then separately depending on whether the trial was a success, an undershoot, or an overshoot (**Figure 3F**, right). To combine across rewards and compare the target preparation axis for successes versus each failure mode, we z-scored values within each reward based on the mean and standard deviation from their successful trials, then pooled across them (**Figure 3G**). We tested this same procedure using data from the full stitched space (**Figure S2A**).

To further test the relationship between the target preparation axis projections and behavior, we correlated projections with the metrics quantified in **Figure S1**. We split data into each reward-direction condition ($n = 80$ total: 32 for Monkeys E and R, 16 for Monkey P), and we calculated the Spearman rank correlation between the target preparation axis projections and each metric. We used the *corr* function in MATLAB, which also calculates the p-value for each correlation based on a t-test. We report all correlation values as histograms in **Figure S4A** with median shown as a dashed line, and correlations that were significant ($p < 0.05$) were overlaid in a shaded histogram. We used a sign-rank test ($n = 80$) for each metric to assess if the overall correlation was significant. We also performed this test within each subject to assess consistency across animals. Note that we correlated projections with the absolute value of the off-axis error, which has a value per trial, as opposed to the off-axis deviation, which has a value per direction-reward condition.

Prior work has demonstrated correlations between neural preparatory states in PMd activity and reach accuracy¹⁹. We assessed if the projections in the target plane exhibited correlations with peak speed cursor locations (**Figure S4B-C**). Within each reward-direction condition, we first compared the “on-axis” quantities in both neural and cursor space. This means we compared the neural projection along the target preparation axis (TPA) with the location of peak speed along the vector pointing from the center to the reach target for that condition. We also compared “off-axis” values: we calculated the “off-TPA” projection as the projection along the vector orthogonal to the TPA in the target plane for a given direction-reward condition. Because the target plane was aligned to the reach target locations, we selected the orientation of the off-TPA to match the off-axis of the reach direction, such that positive values pointed counterclockwise of the TPA in the target plane. We correlated these projections with the off-axis error of the cursor position at the time of peak speed within each direction-reward condition and compiled the results in the same manner as the metrics in **Figure S4A**.

Trial-to-trial variability analysis.

Along with changes in average neural activity as a function of direction and reward, we wanted to assess how trial-to-trial variability was modulated as a function of reward (**Figure 4**). We calculated trial-to-trial variability within a given reward-direction condition ($\sigma_{d,r}^2$) for each factor of the latent space individually as:

$$\sigma_{d,r}^2 = \frac{1}{T_{d,r}-1} \sum_{t=1}^{T_{d,r}} (x_t - \mu_{d,r})^2 \quad (3)$$

where $T_{d,r}$ is the number of trials for this reward (r) direction (d) condition, x_t is the neural activity for trial t of this reward-direction condition (in this case, just one dimension of neural activity), and $\mu_{d,r}$ is the average neural activity for reach direction d and reward r . We then calculated the average trial-to-trial variability for each reward along each factor as the average across reach directions:

$$\sigma_r^2 = \frac{1}{D} \sum_{d=1}^D \sigma_{d,r}^2 \quad (4)$$

where D is the total number of reach directions (8 for Monkeys E and R, 4 for Monkey P). This produces a value of trial-to-trial variability for each individual factor of the latent space and each reward condition. To calculate the total trial-to-trial variability within each reward, we took the sum of the trial-to-trial variability across all factors of the latent space (**Figure 4B**). We then specifically calculated trial-to-trial variability for each reward along dimensions of the neural state space that we previously identified, specifically the target axes ($\sigma_{r,Targ.Ax.}^2$, summed across the two target axes (**Figure 4C**)). To calculate standard error bars for each variance calculation, we performed 1,000 bootstraps within each reward condition and used the standard deviation of the bootstraps.

Supplemental item titles

Figure S1. Neither muscle activation during cue presentation nor upcoming reach behavior explains reward-related modulations of neural activity along the reward axis.

Figure S2. Further validation of Neural Bias hypothesis.

Figure S3. Single unit tuning shows an inverted-U relationship between separability of directions with reward.

Figure S4. Correlations between target plane projections and upcoming reach behavior.

Table S1. Task conditions and experimental details for each animal.

Table S2. A two-target choice task indicates the animals understood the reward values.

Table S3. Full statistics for all main text figures.

Table S4. Single unit reward tuning statistics.

References

1. Baumeister, R.F. (1984). Choking under pressure: Self-consciousness and paradoxical effects of incentives on skillful performance. *Journal of Personality and Social Psychology*, 11.
2. Beilock, S.L., and Carr, T.H. (2005). When high-powered people fail. *Psychological Science* 16, 5.
3. Kimble, C.E., and Rezabek, J.S. (1992). Playing games before an audience: Social facilitation or choking. *Soc. Behav. Pers.* 20, 115–120. <https://doi.org/10.2224/sbp.1992.20.2.115>.
4. Ariely, D., Gneezy, U., Loewenstein, G., and Mazar, N. (2009). Large stakes and big mistakes. *Review of Economic Studies* 76, 451–469. <https://doi.org/10.1111/j.1467-937X.2009.00534.x>.
5. Lewis, B.P., and Linder, D.E. (1997). Thinking about choking? Attentional processes and paradoxical performance. *Personality and Social Psychology Bulletin* 23, 937–944. <https://doi.org/10.1177/0146167297239003>.
6. Gucciardi, D.F., and Dimmock, J.A. (2008). Choking under pressure in sensorimotor skills: Conscious processing or depleted attentional resources? *Psychology of Sport and Exercise* 9, 45–59. <https://doi.org/10.1016/j.psychsport.2006.10.007>.
7. Lee, T.G., and Grafton, S.T. (2015). Out of control: Diminished prefrontal activity coincides with impaired motor performance due to choking under pressure. *NeuroImage* 105, 145–155. <https://doi.org/10.1016/j.neuroimage.2014.10.058>.
8. Chib, V.S., De Martino, B., Shimojo, S., and O'Doherty, J.P. (2012). Neural mechanisms underlying paradoxical performance for monetary incentives are driven by loss aversion. *Neuron* 74, 582–594. <https://doi.org/10.1016/j.neuron.2012.02.038>.
9. Smoulder, A.L., Pavlovsky, N.P., Marino, P.J., Degenhart, A.D., McClain, N.T., Batista, A.P., and Chase, S.M. (2021). Monkeys exhibit a paradoxical decrease in performance in high-stakes scenarios. *Proc. Natl. Acad. Sci. U.S.A.* 118, e2109643118. <https://doi.org/10.1073/pnas.2109643118>.
10. Codol, O., Holland, P.J., Manohar, S.G., and Galea, J.M. (2020). Reward-based improvements in motor control are driven by multiple error-reducing mechanisms. *J. Neurosci.* 40, 3604–3620. <https://doi.org/10.1523/JNEUROSCI.2646-19.2020>.
11. Codol, O., Kashefi, M., Forgaard, C.J., Galea, J.M., Pruszynski, J.A., and Gribble, P.L. (2023). Sensorimotor feedback loops are selectively sensitive to reward. *eLife* 12, e81325. <https://doi.org/10.7554/eLife.81325>.
12. Summerside, E.M., Shadmehr, R., and Ahmed, A.A. (2018). Vigor of reaching movements: reward discounts the cost of effort. *Journal of Neurophysiology* 119, 2347–2357. <https://doi.org/10.1152/jn.00872.2017>.
13. Manohar, S.G., Chong, T.T.-J., Apps, M.A.J., Batla, A., Stamelou, M., Jarman, P.R., Bhatia, K.P., and Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology* 25, 1707–1716. <https://doi.org/10.1016/j.cub.2015.05.038>.
14. Haith, A.M., Pakpoor, J., and Krakauer, J.W. (2016). Independence of movement preparation and movement initiation. *Journal of Neuroscience* 36, 3007–3015. <https://doi.org/10.1523/JNEUROSCI.3245-15.2016>.

15. Ames, K.C., Ryu, S.I., and Shenoy, K.V. (2019). Simultaneous motor preparation and execution in a last-moment reach correction task. *Nat Commun* 10, 2718. <https://doi.org/10.1038/s41467-019-10772-2>.
16. Churchland, M.M., and Shenoy, K.V. (2007). Delay of movement caused by disruption of cortical preparatory activity. *Journal of Neurophysiology* 97, 348–359. <https://doi.org/10.1152/jn.00808.2006>.
17. Afshar, A., Santhanam, G., Yu, B.M., Ryu, S.I., Sahani, M., and Shenoy, K.V. (2011). Single-Trial Neural Correlates of Arm Movement Preparation. *Neuron* 71, 555–564. <https://doi.org/10.1016/j.neuron.2011.05.047>.
18. Tanji, J., and Evarts, E.V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology* 39, 1062–1068. <https://doi.org/10.1152/jn.1976.39.5.1062>.
19. Even-Chen, N., Sheffer, B., Vyas, S., Ryu, S.I., and Shenoy, K.V. (2019). Structure and variability of delay activity in premotor cortex. *PLoS Comput Biol* 15, e1006808. <https://doi.org/10.1371/journal.pcbi.1006808>.
20. Churchland, M.M. (2006). Neural Variability in Premotor Cortex Provides a Signature of Motor Preparation. *Journal of Neuroscience* 26, 3697–3712. <https://doi.org/10.1523/JNEUROSCI.3762-05.2006>.
21. Chib, V.S., Shimojo, S., and O'Doherty, J.P. (2014). The effects of incentive framing on performance decrements for large monetary outcomes: behavioral and neural mechanisms. *Journal of Neuroscience* 34, 14833–14844. <https://doi.org/10.1523/JNEUROSCI.1491-14.2014>.
22. Churchland, M.M., Afshar, A., and Shenoy, K.V. (2006). A central source of movement variability. *Neuron* 52, 1085–1096. <https://doi.org/10.1016/j.neuron.2006.10.034>.
23. Woolley, S.C., Rajan, R., Joshua, M., and Doupe, A.J. (2014). Emergence of context-dependent variability across a basal ganglia network. *Neuron* 82, 208–223. <https://doi.org/10.1016/j.neuron.2014.01.039>.
24. Kao, M.H., Doupe, A.J., and Brainard, M.S. (2005). Contributions of an avian basal ganglia–forebrain circuit to real-time modulation of song. *Nature* 433, 638–643. <https://doi.org/10.1038/nature03127>.
25. Ölveczky, B.P., Andalman, A.S., and Fee, M.S. (2005). Vocal Experimentation in the Juvenile Songbird Requires a Basal Ganglia Circuit. *PLoS Biol* 3, e153. <https://doi.org/10.1371/journal.pbio.0030153>.
26. Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature* 382.
27. Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238. <https://doi.org/10.1038/22268>.
28. Sugrue, L.P. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787. <https://doi.org/10.1126/science.1094765>.
29. Musallam, S., Corneil, B.D., Greger, B., Scherberger, H., and Andersen, R.A. (2004). Cognitive control signals for neural prosthetics. *Science* 305, 258–262. <https://doi.org/10.1126/science.1097938>.
30. Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2009). Fictive reward signals in the anterior cingulate cortex. *Science* 324, 948–950. <https://doi.org/10.1126/science.1168488>.

31. Kennerley, S.W., Behrens, T.E.J., and Wallis, J.D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* 14, 1581–1589. <https://doi.org/10.1038/nn.2961>.
32. Stănişor, L., van der Togt, C., Pennartz, C.M.A., and Roelfsema, P.R. (2013). A unified selection signal for attention and reward in primary visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 110, 9136–9141. <https://doi.org/10.1073/pnas.1300117110>.
33. Roesch, M.R., and Olson, C.R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 304, 307–310. <https://doi.org/10.1126/science.1093223>.
34. Marsh, B.T., Tarigoppula, V.S.A., Chen, C., and Francis, J.T. (2015). Toward an autonomous brain machine interface: Integrating sensorimotor reward modulation and reinforcement learning. *Journal of Neuroscience* 35, 7374–7387. <https://doi.org/10.1523/JNEUROSCI.1802-14.2015>.
35. Ramakrishnan, A., Byun, Y.W., Rand, K., Pedersen, C.E., Lebedev, M.A., and Nicolelis, M.A.L. (2017). Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proc Natl Acad Sci USA* 114, E4841–E4850. <https://doi.org/10.1073/pnas.1703668114>.
36. Pastor-Bernier, A., and Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *J. Neurosci.* 31, 7083–7088. <https://doi.org/10.1523/JNEUROSCI.5681-10.2011>.
37. Degenhart, A.D., Bishop, W.E., Oby, E.R., Tyler-Kabara, E.C., Chase, S.M., Batista, A.P., and Yu, B.M. (2020). Stabilization of a brain–computer interface via the alignment of low-dimensional spaces of neural activity. *Nat Biomed Eng* 4, 672–685. <https://doi.org/10.1038/s41551-020-0542-9>.
38. Kao, T.-C., Sadabadi, M.S., and Hennequin, G. (2021). Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. *Neuron* 109, 1567–1581.e12. <https://doi.org/10.1016/j.neuron.2021.03.009>.
39. Dekleva, B.M., Kording, K.P., and Miller, L.E. (2018). Single reach plans in dorsal premotor cortex during a two-target task. *Nat Commun* 9, 3556. <https://doi.org/10.1038/s41467-018-05959-y>.
40. Mobbs, D., Hassabis, D., Seymour, B., Marchant, J.L., Weiskopf, N., Dolan, R.J., and Frith, C.D. (2009). Choking on the Money: Reward-Based Performance Decrements Are Associated With Midbrain Activity. *Psychol Sci* 20, 955–962. <https://doi.org/10.1111/j.1467-9280.2009.02399.x>.
41. Sörensen, L.K.A., Bohté, S.M., Slagter, H.A., and Scholte, H.S. (2022). Arousal state affects perceptual decision-making by modulating hierarchical sensory processing in a large-scale visual system model. *PLoS Comput Biol* 18, e1009976. <https://doi.org/10.1371/journal.pcbi.1009976>.
42. Masters, R.S.W. (1992). Knowledge, knerves and know-how: The role of explicit versus implicit knowledge in the breakdown of a complex motor skill under pressure. *British J of Psychology* 83, 343–358. <https://doi.org/10.1111/j.2044-8295.1992.tb02446.x>.
43. Butler, J.L., and Baumeister, R.E. The Trouble With Friendly Faces: Skilled Performance With a Supportive Audience.
44. Yerkes, R.M., and Dodson, J.D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *J. Comp. Neurol. Psychol.* 18, 459–482. <https://doi.org/10.1002/cne.920180503>.
45. Easterbrook, J.A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review* 66, 183–201. <https://doi.org/10.1037/h0047707>.

46. Sosnowski, M.J., Benítez, M.E., and Brosnan, S.F. (2022). Endogenous cortisol correlates with performance under pressure on a working memory task in capuchin monkeys. *Sci Rep* 12, 953. <https://doi.org/10.1038/s41598-022-04986-6>.
47. Wine, J. (1971). Test anxiety and direction of attention. *Psychological Bulletin* 76, 92–104. <https://doi.org/10.1037/h0031332>.
48. Mesagno, C., Harvey, J.T., and Janelle, C.M. (2012). Choking under pressure: The role of fear of negative evaluation. *Psychology of Sport and Exercise* 13, 60–68. <https://doi.org/10.1016/j.psychsport.2011.07.007>.
49. Gallego, J.A., Perich, M.G., Chowdhury, R.H., Solla, S.A., and Miller, L.E. (2020). Long-term stability of cortical population dynamics underlying consistent behavior. *Nat Neurosci* 23, 260–270. <https://doi.org/10.1038/s41593-019-0555-4>.
50. Golub, M.D., Yu, B.M., and Chase, S.M. (2015). Internal models for interpreting neural population activity during sensorimotor control. *eLife* 4, e10015. <https://doi.org/10.7554/eLife.10015>.
51. Georgopoulos, A., Kalaska, J., Caminiti, R., and Massey, J. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2, 1527–1537. <https://doi.org/10.1523/JNEUROSCI.02-11-01527.1982>.
52. Harris, C.M., and Wolpert, D.M. (1998). Signal-dependent noise determines motor planning. *Nature* 394, 780–784. <https://doi.org/10.1038/29528>.
53. Bishop, W.E., and Yu, B.M. (2014). Deterministic symmetric positive semidefinite matrix completion. In *Advances in Neural Information Processing Systems*.
54. Pandarinath, C., O’Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat Methods* 15, 805–815. <https://doi.org/10.1038/s41592-018-0109-9>.
55. Jude, J., Perich, M.G., Miller, L.E., and Hennig, M.H. (2022). Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation. Preprint at arXiv.
56. Bishop, W.E. (2015). *Combining Neural Population Recordings: Theory and Application*.
57. Nonnenmacher, M., Turaga, S.C., and Macke, J.H. (2017). Extracting low-dimensional dynamics from multiple large-scale neural population recordings by learning to predict correlations. In *Advances in Neural Information Processing Systems*.
58. Fraser, G.W., and Schwartz, A.B. (2012). Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology* 107, 1970–1978. <https://doi.org/10.1152/jn.01012.2010>.
59. Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology* 102, 614–635. <https://doi.org/10.1152/jn.90941.2008>.