TITLE

Dynamic dopaminergic activity controls the timing of self-timed movement

AUTHORS

Allison E. Hamilos[1,4*], Giulia Spedicato[1†], Ye Hong[1†], Fangmiao Sun[2‡], Yulong Li[2‡] & John A. Assad[1,3*]

AUTHOR LIST FOOTNOTES

[1]Department of Neurobiology, Harvard Medical School, Boston, Massachusetts, 02115, USA.

[2]State Key Laboratory of Membrane Biology, Peking University School of Life Sciences, Beijing, 100871, P.R. China.

[3]Istituto Italiano di Tecnologia, Genova, Italy.

[4]Lead Contact

*Correspondence: A.H. (ahamilos@mit.edu) or J.A.A. (jassad@hms.harvard.edu).

SUMMARY

Deciding when to initiate action is essential to survival. Insights from movement disorders and pharmacological studies implicate the neurotransmitter dopamine as a regulator of movement timing, but the underlying neural mechanisms are not understood. Here we show dynamic dopaminergic signaling over seconds-long timescales controls movement timing in mice. Animals were trained to initiate licking after a self-timed interval following a start-timing cue. Surprisingly, dopaminergic signals ramped-up slowly between the start-timing cue and the self-timed movement, with the slope predicting the movement time on single

trials. Steeply rising signals preceded early lick-initiation, whereas slowly rising signals preceded later initiation, reminiscent of a ramp-to-threshold process. Higher baseline activity also predicted earlier self-timed movements. Optogenetic activation of dopamine neurons during self-timing caused systematic early-shifting of movement initiation, whereas inhibition caused late-shifting. These results reveal a causal role for dynamic dopaminergic signaling unfolding over seconds in controlling the moment-by-moment decision of when to move.

KEYWORDS

Timing, Dopamine, Self-Initiated Movement, Nigrostriatal, Basal Ganglia, Decision-making, Volition, Parkinson's Disease, Cognition, Learning

INTRODUCTION

What makes us move? Empirically, a few hundred milliseconds before movement, thousands of neurons in the motor system suddenly become active in concert, and this neural activity is relayed via spinal and brainstem neurons to recruit muscle fibers that power movement (Shenoy et al, 2013). Yet, mysteriously, just before this period of intense neuronal activity, the motor system is largely quiescent. How and why does the brain suddenly and profoundly rouse motor neurons into the coordinated action needed to trigger movement?

In the case of movements made in reaction to external stimuli, it is possible that activity evoked first in sensory brain areas could be passed along to appropriate motor centers to trigger this coordinated neural activity, thereby leading to movement. But humans and animals can also self-

initiate movement without overt, external input (Deecke, 1996; Hallett, 2007; Lee & Assad, 2003; Romo et al., 1992). For example, while reading this page, you may decide, without prompting, to reach for your coffee. In that case, the movement cannot be clearly related to an abrupt, conspicuous sensory cue. What went off in your brain that made you reach for the coffee at this *particular* moment, as opposed to a moment earlier or later?

Human movement disorders may provide clues to this mystery. Patients and animal models of Parkinson's Disease (PD) experience difficulty self-initiating movements, exemplified by perseveration (Hughes et al., 2013), trouble initiating steps when walking (Bloxham et al., 1984), and problems timing movements (Malapani et al., 1998; Meck, 2006; Meck 1986; Mikhael & Gershman, 2019). In contrast to these self-generated actions, externally cued reactions are often less severely affected in PD, a phenomenon sometimes referred to as "paradoxical kinesis" (Bloxham et al., 1984, Barthel et al., 2018). For example, patients' gait can be normalized by walking aids that prompt steps in reaction to visual cues displayed on the ground (Barthel et al., 2018).

Because the underlying neuropathophysiology of PD includes the loss of dopaminergic neurons (DANs), the symptomatology of PD suggests DAN activity plays an important role in deciding when to self-initiate movement. Indeed, pharmacological manipulations of the neurotransmitter dopamine causally and bidirectionally influence movement timing (Dews & Morse, 1958; Schuster & Zimmerman, 1961; Meck 1986; Lustig & Meck, 2005; Mikhael & Gershman, 2019). This can be demonstrated in the context of *self-timed* movement tasks, in which subjects reproduce a target-timing interval by making a movement following a self-timed delay that is referenced to

a start-timing cue (Malapani et al., 1998). Species across the animal kingdom, from rodents and birds to primates, can learn these tasks and produce self-timed movements that occur, on average, at about the criterion time, although the exact timing exhibits considerable variability from trial-to-trial (Rakitin et al., 1998; Gallistel & Gibbon, 2000; Merchant et al., 2013; Meck, 2006; Mello et al., 2015). In such self-timed movement tasks, decreasing dopamine availability/efficacy (e.g., Parkinson's disease, neuroleptic drugs) generally produces late-shifted movements (Merchant et al., 2013; Meck, 2006; Malapani et al., 1998; Meck, 1986), whereas high dopamine conditions (e.g., amphetamines) produces early-shifting (Dews & Morse, 1958; Schuster & Zimmerman, 1961).

Although tonic manipulations of the dopamine neurotransmitter influence self-timing, it remains unknown whether the activity of DANs is involved in determining the timing of movement. DANs of the nigrostriatal pathway richly innervate spiny projection neurons of the dorsal striatum, which send outputs that modulate motor centers in the thalamus, cortex and brainstem (Grillner & Robertson, 2016). We hypothesized that the activity of nigrostriatal dopamine neurons (DANs) could control the moment-by-moment decision of when to move, potentially by modulating the balance of the direct and indirect pathways of the basal ganglia, which are poised to coordinate movement via their push-pull influence on motor centers (Albin et al., 1989; DeLong, 1990; Freeze et al., 2013; Grillner & Robertson, 2016).

To test this hypothesis, we trained mice to make a movement (lick) after a self-timed interval following a start-timing cue, and we asked how DAN signals during the timing interval related to the precise timing of movement initiation. The mice learned the timed interval, but, as observed in

other species, the exact trial-by-trial timing of movement was highly variable, spanning seconds (Rakitin et al., 1998; Gallistel & Gibbon, 2000; Merchant et al., 2013; Meck, 2006; Mello et al., 2015). We exploited this inherent variability in self-timing by examining how moment-to-moment nigrostriatal DAN signals differed when animals decided to move relatively early versus late. We found that DAN signals were highly predictive of trial-by-trial movement timing, even *seconds* before the movement occurred. In particular, DAN signals showed a pronounced seconds-long build-up ("ramping") from the time of the cue leading up to the time of the movement. Steeper ramping was associated with relatively early movements, whereas shallower ramping with later movements. Ramping was evident both in DAN intracellular Ca^{++} signals as well as in dopamine release within the dorsolateral striatum (DLS), where the neurotransmitter is positioned to tune the balance of direct and indirect pathways of the basal ganglia. Ramping signals tended to reach a similar amplitude around the time of movement, reminiscent of a ramp-to-threshold process. The amplitude of the baseline signal before the start-timing cue was also predictive of the upcoming movement time, with higher baseline activity portending earlier movements. Optogenetic DAN manipulation during the timing interval produced bidirectional changes in movement timing on single trials, with activation causing earlier self-timed movements and suppression causing later movements. Altogether, we propose a new role for slowly-evolving dopaminergic signaling in controlling the moment-to-moment decision of when to move.

RESULTS

We trained head-fixed mice to make self-timed movements to receive juice rewards (Figure 1A). Animals received an audio/visual start-timing cue and then had to decide when to first-lick in the absence of further cues. Animals only received juice if they waited a proscribed interval following

116 the cue before making their first-lick (3.3s in most experiments). First-lick timing exhibited a broad
117 distribution spanning several seconds (Figures 1B and S1A,B), as expected from previous studies
118 (Gallistel & Gibbon, 2000; Merchant et al., 2013; Meck, 2006; Mello et al., 2015). As mice
119 executed the task, we employed fiber photometry to record the activity of genetically-defined
120 DANs expressing the calcium-sensitive fluorophore GCaMP6f (12 mice, substantia nigra pars
121 compacta (SNc); Figures 1C-E, 2A, S2, and S3). We controlled for mechanical/optical artifacts by
122 simultaneously recording fluorescence modulation of a co-expressed, calcium-insensitive
123 fluorophore, tdTomato (tdt, Figure 1E), as well as body movements detected by neck EMG, high-
124 speed video and a back-mounted accelerometer (Figure 3A).
125

126 **DAN signals ramp up slowly between the start-timing cue and self-timed movement**

127 DAN GCaMP6f fluorescence typically exhibited brief transients following cue onset and
128 immediately before first-lick onset (Figure 1C), as observed in previous studies (Schultz et al.,
129 1997; Dodson et al., 2016; Coddington & Dudman, 2018; da Silva et al., 2018; Howe & Dombeck,
130 2016). However, during the timed interval, we observed slow "ramping up" of fluorescence, with
131 a minimum after the cue-locked transient and maximum just before the lick-related transient. We
132 asked whether this ramping differed between trials in which the animal moved relatively early or
133 late. Strikingly, when we averaged signals by movement time, we observed systematic differences
134 in the steepness of ramping that were highly predictive of movement timing (Figure 1D,E). Trials
135 with early first-licks exhibited steep ramping, whereas trials with later first-licks started from lower
136 fluorescence levels and rose more slowly toward the time of movement. The fluorescence ramps
137 terminated at nearly the same amplitude regardless of movement time. Similar ramping dynamics
138 and baseline differences were found in a dorsal lateral striatal "lick area" (DLS, Sippy et al., 2015),

139 in both the fluorescence of GCaMP6f in DAN axon terminals (Figure 2B), as well as $GRAB_{DA2m}$
140 (DA_{2m}) expressed in striatal cells (Figures 1F and 2C). DA_{2m} is a new, improved extracellular
141 dopamine indicator derived from the dopamine-2-receptor (Sun et al., 2020). Thus, ramping SNc
142 GCaMP6f dynamics are played out at the axon terminal and are reflected in striatal dopamine
143 accumulation, suggesting that DAN ramping dynamics may causally influence movement timing
144 via the interaction of released dopamine with downstream striatal neurons. Similar ramping
145 dynamics were also observed in GCaMP6f-expressing DAN cell bodies in the ventral tegmental
146 area (VTA, Figure 2D), reminiscent of ramping dynamics observed in VTA spiking and
147 mesolimbic dopamine release during goal-oriented navigation tasks as animals approached a
148 rewarded target (Howe et al., 2013; Kim et al., 2019).
149

150 **Higher pre-cue, baseline DAN signals are correlated with earlier self-timed movements**

151 In addition to ramping dynamics, slowly-modulating DAN signals were correlated with first-lick
152 timing even before cue-onset, with higher baseline fluorescence predicting earlier first-licks
153 (Figures 1E,F, 2A-D and 4A-B). Because dF/F correction methods can potentially distort baseline
154 measurements, we rigorously tested and validated three different dF/F methods, and we also
155 repeated analyses with raw fluorescence values compared between pairs of sequential trials with
156 different movement times (Figure S3, see Methods). All reported results, including the systematic
157 baseline differences, were robust to dF/F correction. The systematic correlation of baseline signals
158 with first-lick timing was not fully explained by prior trial outcome (reward/no reward) nor licking
159 during the intertrial interval (ITI) (Figure 4C). In fact, although a reward on the prior trial tended
160 to elevate signals during the ITI, there was an abrupt "resetting" of baseline signals during the
161 random delay (after lamp-off but before the start-timing cue), such that baseline amplitude became

162 abruptly and progressively better explained by the upcoming trial outcome (reward/no reward)
163 compared to the prior trial (Figure 4B,C). Mice trained on a variant of the self-timing task without
164 lamp-off/on events showed no systematic differences in their timing distributions, suggesting that
165 although DAN resetting occurred at lamp-off, the mice still referenced their timing to the start-
166 timing cue (Figure S1C).
167

168 **Timing-predictive DAN signals are not explained by movement or optical artifacts**

169 The systematic ramping dynamics and baseline differences observed with GCaMP6f and DA_{2m}
170 were not observed in the tdt optical control channel nor in any of the other movement-control
171 channels (Figures 1E, 2E, and 3B), making it unlikely that these ramping dynamics could have
172 arisen from optical artifacts. Nevertheless, because DANs show transient responses to salient cues
173 and movements (Schultz et al., 1997; Dodson et al., 2016; Coddington & Dudman, 2018; da Silva
174 et al., 2018; Howe & Dombeck, 2016), it is possible GCaMP6f and DA_{2m} fluorescence could reflect
175 the superposition of dopaminergic responses to multiple task events, including the cue, lick,
176 ongoing spurious body movements, and hidden cognitive processes like timing. For example,
177 accelerating spurious movements could, in principle, produce motor-related neural activity that
178 "ramps up" during the timed interval, perhaps even at different rates on different trials. We thus
179 derived a nested linear encoding model of single-trial GCaMP6f signals (Engelhard et al, 2019;
180 Park et al., 2014; Runyan et al, 2017), a data-driven, statistical approach designed to isolate and
181 quantify the contributions of task events (timing-independent predictors) from processes predictive
182 of movement timing (timing-dependent predictors) (Figures 5A,B and S4A-D). The model
183 robustly detected task-event GCaMP6f kernels locked to cue, lick and EMG/accelerometer events
184 (Figures 5C and S4E), but these timing-independent predictors alone were insufficient to capture

185 the rich variability of GCaMP6f signals for trials with different first-lick times, especially the
186 timing-dependent ramp-slope and baseline offset (68 sessions, Figures 5C and S4F,G). In contrast,
187 two timing-dependent predictors robustly improved the model: 1) a baseline offset with amplitude
188 linearly proportional to first-lick time; and 2) a "stretch" feature representing percentages of the
189 timed interval, which predicted a ramp from cue-to-first-lick with slope inversely proportional to
190 first-lick time (68 sessions, Figures 5B,C and S4E). Similar results were obtained for SNc DAN
191 axon terminals in the DLS, DLS neurons expressing DA_{2m}, and VTA DAN cell bodies (Figure
192 S4H).
193
194 In contrast to the GCaMP6f model, when the same procedure was applied to control photometry
195 signals (tdt), the timing-independent predictors (which could potentially cause optical or
196 mechanical artifacts—cue, first-lick, EMG/accelerometer) improved the model, but timing-
197 dependent predictors did not improve the model (Figures 5C and S4F-H).
198
199 Principle component (PC) analysis revealed ramp-like and baseline-offset-like components that
200 explained as much as 93% of the variance in GCaMP6f signals during the timing interval (mean:
201 66%, range: 16-93%), but similar PCs were not present in tdt signals (mean: 4%, range: 1.6-15%)
202 (Figure S5A,B).
203
204 **DAN ramping and baseline signals predict single-trial movement timing**
205 Given that ramping and baseline-offset signals were not explained by nuisance movements or
206 optical artifacts, we asked whether DAN GCaMP6f fluorescence could predict first-lick timing on
207 single trials. Using a simple threshold-crossing model (Maimon & Assad, 2006), we found that the

208 GCaMP6f signal was predictive of first-lick time even for low thresholds intersecting the "base"
209 of the ramp, with the predictive value of the model progressively improving for higher thresholds
210 (R^2 low: 0.34; mid: 0.64; high: 0.94, Figure 6A). To more thoroughly determine the independent,
211 additional predictive power of DAN baseline and ramping signals over other task variables (e.g.,
212 previous trial first-lick time and outcome (reward/no reward), etc.), we derived a nested decoding
213 model for first-lick time (Figure 6A and S5C). All predictors contributed to the predictive power
214 of the model. However, even when we accounted for the contributions of prior trial history, tdt
215 artifacts and baseline GCaMP6f signals, GCaMP6f threshold-crossing time robustly dominated
216 the model, alone explaining 10% of the variance in first-lick time on average (range: 1-27%, Figure
217 6B-D). Alternate versions of the decoding model showed similar results (Figure S5C).
218

219 **Moment-to-moment DAN activity causally controls movement timing**

220 Because the DAN ramping signal robustly predicted first-lick timing and was apparently
221 transmitted via dopamine release to downstream striatal neurons, ramping DAN activity may
222 causally determine movement timing. If so, causally increasing the activity of DANs during timing
223 should result in earlier self-timed movements, and vice-versa. We thus optogenetically activated
224 or inhibited DANs (in separate experiments) on 30% of trials (Figures 7A and S6A,B). Activation
225 significantly early-shifted the distribution of self-timed movements on stimulated trials compared
226 to unstimulated trials (12 mice), whereas inhibition produced significant late-shifting compared to
227 unstimulated trials (4 mice). Stimulation of mice expressing no opsin produced no consistent
228 effect on timing (5 mice, Figures 7B-D and S6C-E). The direction of these effects was consistent
229 across all animals tested in each category. Whereas bilateral activation of SNc DAN cell bodies
230 caused early-shifting, the effects were generally larger and more consistent when activating SNc

231 DAN terminals in DLS (2 mice, Figure S6C,D). Outside the context of the timing task, DAN
232 activation did not elicit immediate licking, nor did inhibition prevent licking, suggesting
233 optogenetic effects on timing did not result from direct triggering or suppression of movement
234 (Coddington & Dudman, 2018), but rather were expressed through a "higher-level" cognitive
235 process related to self-timing (Figure S7A,B).
236
237
238 **DISCUSSION**
239 Here, we show that both baseline and slowly ramping DAN signals predict the timing of self-
240 initiated movements. Trial-by-trial differences in these signals were finely tuned to movement
241 onset, whether these signals were recorded from SNc cell bodies, SNc terminals in the DLS, or
242 VTA cell bodies. Moreover, slow DAN dynamics were reflected in dopamine release in DLS,
243 demonstrating availability of this information to downstream striatal effectors positioned to
244 influence when movement occurs. Consistent with the direction of these effects, optogenetic
245 suppression and augmentation of DAN activity during the timing interval causally altered
246 movement timing. Thus, DAN activity is poised to control the moment-to-moment decision of
247 when to move.
248
249 **A causal role for slow ramping DAN activity in the timing of self-initiated movements**
250 A number of studies have reported short-latency (≤ ~500ms) increases in DAN activity in response
251 to sensory cues and immediately preceding self-initiated movements (Schultz et al., 1997; Dodson
252 et al., 2016; Coddington & Dudman, 2018; da Silva et al., 2018; Howe & Dombeck, 2016), similar
253 to the sensory- and motor-related transients we observed just after the cue and just before the first-

lick. However, the slow-timescale DAN signals we observed during self-timing were markedly different. First, the ramping signal unfolded over *seconds*, preceding the first-lick by as long as 10 s. Second, variations in both the ramp-slope and pre-cue baseline amplitude predicted trial-by-trial first-lick timing. To effectively model DAN signals on single trials, we had to incorporate two timing-dependent features: a baseline offset and a "stretch" parameter that scaled DAN signals along the time axis (Figure 5). Moreover, these features predicted movement timing independent of recent trial history (Figure 6). Combined with the optogenetic results, these findings suggest that variations in slow DAN dynamics affect trial-by-trial movement timing.

These slow-timescale DAN signals could be unique to the timing requirement of our task. However, when we averaged DAN signals aligned to "spontaneous" licks during the ITI, we also observed slow ramping similar to that observed during the timing interval, with signal building over seconds from the offset of the previous lick up to the time of the next lick (Figure S7C). Thus, slowly evolving DAN signals may be integral to self-initiated movements more generally. It is possible that slow ramping dynamics predictive of movement timing would emerge in previous datasets if DAN signals were similarly averaged according to the interval between self-initiated movement bouts.

Previous studies have reported slow ramping signals in the mesolimbic system in certain behavioral contexts, including goal-directed navigation (Howe et al., 2013); multi-step tasks culminating in reward (Hamid et al., 2016; Mohebi et al., 2019); and passive observation of dynamic visual cues indicating proximity to reward (Kim et al., 2019). It has been proposed that slowly ramping mesolimbic DAN signals could encode increasing value as animals approach

reward (Hamid et al., 2016; Mohebi et al., 2019) or alternatively could reflect moment-by-moment reward-prediction errors (RPE, Kim et al., 2019; Mikhael et al., 2019). The ramping signals we observed in the nigrostriatal system are consistent with either value or RPE interpretations. However, it has been unclear how the brain *employs* slowly ramping DAN signals in behavior. Our study moves beyond previous studies by finding that trial-by-trial variability in ramping dynamics explains the precise timing of a behavioral output—the self-timed lick—and that optogenetically manipulating SNc DAN activity causally alters the timing of that output. Thus, SNc ramping may not merely encode progress toward a goal, but could also play a *causal* role in the timing of movement initiation. This interpretation could be related to classic findings from Parkinson's disease, in which loss of nigrostriatal DANs results in difficulty initiating movements (Grillner & Robertson, 2016; Freeze et al., 2013).

Relationship to Pacemaker-Accumulator models of neural timing

Lesion and pharmacological studies have long suggested roles for the SNc and dopamine in timing (Merchant et al., 2013; Meck, 2006). Broadly speaking, conditions that increase dopamine availability affect timing as if speeding an internal "pacemaker" (Dews & Morse, 1958; Schuster & Zimmerman, 1961; Mikhael & Gershman, 2019), whereas conditions that decrease dopamine availability/efficacy generally have the opposite effect (Merchant et al., 2013; Meck, 2006; Malapani et al., 1998; Meck, 1986). Our results—in both DAN recordings and optogenetic manipulations—are consistent with this view. Moreover, the ramping signals and the anti-correlation of ramping slope with movement time bear striking resemblance to Pacemaker-Accumulator models of neural timing (Meck, 2006; Lustig & Meck, 2005), a longstanding conceptual framework that captures canonical features of timing behavior.

300

301 **A unified temporal difference learning model of DAN activity in neural timing**

302 Soares *et al*. recently reported findings complicating the standard view of dopamine in timing
303 (Soares et al., 2016). In mice performing a temporal bisection task, RPE-like transients in SNc
304 DAN GCaMP6f signals were observed immediately after the stop-timing cue. These transients
305 were smaller when animals overestimated the timed interval, which was interpreted as evidence
306 that *lower* DAN activity reflects a *faster* pacemaker, the *opposite* of our findings and most prior
307 work (Mikhael & Gershman, 2019). This finding may be unique to the temporal bisection task,
308 which has aspects of categorization as well as timing and thus is more complex than tasks that only
309 rely on movement to produce timed intervals (as in our task).

310

311 However, a recently proposed temporal difference learning framework for explaining dynamic
312 DAN activity could provide a unified explanation for these findings (see Hamilos & Assad, 2020
313 for detailed discussion, also see Kim et al., 2019; Mikhael et al., 2019; Mikhael & Gershman,
314 2019). The model assumes that DAN activity provides a continuous readout of RPE, which under
315 conditions of state uncertainty (as in timing) is shown to reflect the moment-to-moment derivative
316 of the value landscape (Mikhael et al., 2019). In this framework, variation in interval timing arises
317 from differences in the rate of traversing the internal model of the value landscape compared to
318 veridical time, which can be modeled as stretching/compression of the subjective value function
319 (Mikhael & Gershman, 2019). Critically, the amount of compression is taken to be controlled by
320 a pacemaker whose speed is proportional to the tonic level of dopamine from trial-to-trial (Mikhael
321 & Gershman, 2019). In both the self-timed movement and bisection tasks, we observed higher
322 baseline DAN signals associated with relatively fast timekeeping, consistent with relatively high

tonic DAN activity reflecting fast pacemaking in both tasks. Additionally, temporally discounted value should increase as the time of reward approaches in both tasks, and the model predicts faster pacemaking on a given trial would *compress* this function. During the self-timed movement task, compression would cause faster increases in value and thus steeper DAN signal ramping, as we observed. In the bisection task, compression of the value function would produce higher estimated value just before the stop-cue, and thus a *smaller* change in value (RPE) following the stop-cue, resulting in a blunted DAN transient. Soares *et al*. indeed observed smaller stop-cue-related transients when animals overestimated the timed interval—consistent with compression of the subjective value function due to *faster* pacemaking.

Altogether, we argue that relatively high DAN activity reflects faster pacemaking across timing tasks, with the specific timing from trial-to-trial influenced by the moment-to-moment dynamics of DAN signaling.

ACKNOWLEDGEMENTS

We thank D. Chicharro, J.G. Mikhael, S.J. Gershman, K. Reinhold, S. Panzeri and E.N. Brown for discussions on analytical methods; V. Berezovskii, J. LeBlanc, T. LaFratta, O. Mazor, P. Gorelik, J. Markowitz, A. Lutas, K.W. Huang, L. Hou and N.E. Ochandarena for technical assistance; and B. Sabatini, C. Harvey, S.R. Datta, M. Andermann and W. Regehr for reagents. The work was supported by NIH grants UF-NS109177 and U19-NS113201, and NIH core grant EY-12196. A.E.H. was supported by a Harvard Lefler Predoctoral Fellowship and a Harvard Stuart H.Q. & Victoria Quan Predoctoral Fellowship.

346

347 ## AUTHOR CONTRIBUTIONS

348 A.E.H. and J.A.A conceived the project. A.E.H. performed all experiments. G.S. assisted with
349 experiments using tapered fiber optics; Y.H. assisted with optogenetics control experiments. F.S.
350 and Y.L. developed the dopamine sensor, DA_{2m}. A.E.H. and J.A.A. analyzed the data and wrote
351 the paper.

352

353 ## DECLARATION OF INTERESTS

354 J.A.A. is a co-founder of OptogeniX, which produces the tapered optical fibers used in some
355 experiments.

356
357
358 ## METHODS

359 **Animals**. Adult male and female hemizygous DAT-cre mice (B6.SJL-Slc6a3$^{tm1.1(cre)Bkmm}$/J; The
360 Jackson Laboratory; Backman et al., 2006) or *wt* C57/b6 mice were used in all experiments. Mice
361 were housed on a reversed night/day cycle (12h dark/12h light) and behavioral sessions occurred
362 during the dark cycle. All experiments and protocols were approved by the Harvard Institutional
363 Animal Care and Use Committee and were conducted in accordance with the National Institutes
364 of Health Guide for the Care and Use of Laboratory Animals.

365

366 **Surgery.** Surgeries were conducted under aseptic conditions. Mice were anesthetized with
367 isoflurane (0.5-2% at 0.8L/min). Analgesia was provided by *s.c.* 5mg/kg ketoprofen injection
368 during surgery and once daily for 3d postoperatively. Virus was injected (50nL/min) and the pipet
369 remained in place for 10min before removal. 200μm, 0.53NA blunt fiber optic cannulae (Doric)

370 or tapered fiber optic cannulae (200μm, 0.60NA, 2mm tapered shank, OptogeniX) were positioned
371 at SNc, VTA or DLS and secured to the skull with dental cement (Metabond). Neck EMG
372 electrodes were constructed from two 32G pacemaker wires attached to a custom socket mounted
373 in the dental cement. Sub-occipital neck muscles were exposed by blunt dissection and electrode
374 tips embedded bilaterally.
375
376 **Stereotaxic coordinates (from bregma and brain surface)**.
377 Virus:
378 SNc: 3.16mm posterior, +/- 1.4mm lateral, 4.2mm ventral
379 VTA: 3.1mm posterior, +/-0.6mm lateral, 4.2mm ventral
380 DLS: 0mm anterior, +/- 2.6mm lateral, 2.5mm ventral.
381 Fiber Optic Tips:
382 SNc/VTA: 4.0mm ventral (photometry) or 3.9mm ventral (optogenetics).
383 DLS: 2.311mm ventral (blunt fiber) or 4.0mm ventral (tapered fiber)
384
385 **Virus.**
386 Photometry:
387 tdTomato ("tdt"): AAV1-CAG-FLEX-tdT (UNC Vector Core), 100nL used alone or in
388 mixture with other fluorophores (below), working concentration $5.3*10^{12}$gc/mL
389 gCaMP6f (at SNc or VTA): 100nL AAV1.Syn.Flex.GCaMP6f.WPRE.SV40
390 ($2.5*10^{13}$gc/mL, Penn Vector Core). Virus was mixed in a 1:3 ratio with tdt (200nL
391 total).

392 DA_{2m} (at DLS): 200-300nL AAV9-hSyn-DA4.4(DA2m) (working concentration: *ca.*
393 $3*10^{12}$gc/mL, Vigene) + 100nL tdt
394 Optogenetics (all bilateral at SNc):
395 ChR2: 1000nL AAV5-EF1a-DIO-hChR2(H134R)-EYFP-WPRE-pA ($3.2*10^{13}$gc/mL,
396 UNC Vector Core)
397 ChrimsonR: 700nL AAV1-hSyn-FLEX-ChrimsonR-tdT ($4.1*10^{12}$gc/mL, UNC Vector
398 Core)
399 stGtACR2: 300nL 1:10 AAV2/8-hSyn1-SIO-stGtACR2-FusionRed (working
400 concentration $4.7*10^{11}$gc/mL, Addgene/Janelia Viral Core)
401
402 **Water-deprivation and acclimation.** Animals recovered 1wk postoperatively before water
403 deprivation. Mice received daily water supplementation to maintain ≥80% initial body weight and
404 fed *ad libitum*. Training commenced when mice reached the target weight (~8-9d post-surgery).
405
406 **Histology.** Mice were anesthetized with >400mg/kg pentobarbital (Somnasol) and perfused with
407 10mL 0.9% sodium chloride followed by 50mL ice-cold 4% paraformaldehyde in 0.1M phosphate
408 buffer. Brains were fixed in 4% paraformaldehyde at 4°C for >24hr before being transferred to 30%
409 sucrose in 0.1M phosphate buffer for >48hr. Brains were sliced in 50μm coronal sections by
410 freezing microtome, and fluorophore expression was assessed by light microscopy. The sites of
411 viral injections and fiber optic placement were mapped with an Allen Mouse Brain Atlas.
412
413 **Behavioral rig, data acquisition and analysis.** A custom rig provided cues, recorded events and
414 delivered juice rewards under the control of a Teensy 3.2 microprocessor running a custom

Arduino state-system behavioral program with MATLAB serial interface. Digital and analog signals were acquired with a CED Power 1400 data acquisition system/Spike2 software (Cambridge Electronic Devices). Photometry and behavioral events were acquired at 1000Hz; movement channels were acquired at 2000Hz. Video was acquired with FlyCap2 or Spinnaker at 30fps (FLIR). Data were analyzed with custom MATLAB statistics packages.

Self-timed movement task. Mice were head-fixed with a juice tube positioned in front of the tongue. During periods when rewards were not available, a houselamp was illuminated. At trial start, the houselamp turned off, and a random delay ensued (0.4-1.5s) before a cue (simultaneous LED flash and 3300Hz tone, 100ms) indicated start of the timing interval. The timing interval was divided into two windows, early (0-3.333s in most experiments; 0-4.95s in others) and reward (3.333-7 s; 4.95-10s), followed by the intertrial interval (ITI, 7-17 s; 10-20s). The window in which the mouse first licked determined the trial outcome (early, reward, or no-lick). Importantly, nothing the mouse did could make the trial time-course go any faster, so there was no advantage to the mouse to lick early except its concern that it might miss the end of the 7 s trial period. An early first-lick caused an error tone (440Hz, 200ms) and houselamp illumination, and the mouse had to wait until the full timing interval had elapsed before beginning the ITI. A first-lick during the reward window caused a reward tone (5050Hz, 200ms) and juice delivery, and the houselamp remained off until the end of the trial interval. If the timing interval elapsed with no lick, a time-out error tone played (131Hz, 2s), the houselamp turned on, and ITI commenced. During the ITI and pre-cue delay, there was no penalty for licking.

437 Mice learned the task in 3 stages (Extended Data Fig. 1a). On the first 1-4 days of training, mice
438 learned a beginner-level task, which was modified in two ways: 1. To encourage participation, if
439 mice did not lick before 5s post-cue, they received a juice reward at 5s, 2. Mice were not penalized
440 for licking in reaction to the cue (within 500ms). When the mouse began self-triggering ≥50% of
441 rewards (day 2-6 of training), the mouse advanced to the intermediate-level task, in which the
442 training reward at 5s was omitted, and the mouse had to self-trigger all rewards. After
443 completing >250 trials/day on the intermediate task, mice advanced to the mature task (no reaction
444 licks permitted, day 4-7 of training). All animals learned the mature task and worked for ~400-
445 1500 trials/session.
446
447 **Online movement monitoring.** Movements were recorded simultaneously during behavior with
448 four movement-control measurements: neck EMG (band-pass filtered 50-2000Hz, 60Hz notch,
449 amplified 100-1000x), back-mounted accelerometer (SparkFun), high-speed camera (30Hz, FLIR),
450 and tdt photometry. All control signals contained similar information, and thus only a subset of
451 controls was used in some sessions.
452
453 **Photometry.** Fiber optics were illuminated with 475nm blue LED light (Plexon) (SNc/VTA:
454 50μW, DLS: 35μW) measured at patch cable tip with a light-power meter (Thor Labs). Green
455 fluorescence was collected via a custom dichroic mirror (Doric) and detected with a Newport 1401
456 Photodiode. Fluorescence was allowed to recover ≥1d between recording sessions. To avoid
457 crosstalk in animals with tdt expression, tdt was recorded at one of the 3 sites (SNc, VTA, or DLS,
458 550 lime LED, Plexon) while GCaMP6f or DA_{2m} was recorded simultaneously only at the other
459 implanted sites.

460

461 **dF/F.** Raw fluorescence for each session was pre-processed by removing rare singularities (single
462 points >15 STD from the mean) by interpolation to obtain F(t). To correct photometry signals for
463 bleaching, dF/F was calculated as:

464

465 $$\frac{dF}{F}(t) = \frac{F(t) - F_0(t)}{F_0(t)}$$

466

467 where $F_0(t)$ is the 200 s moving average of F(t). We tested several other complementary methods
468 for calculating dF/F and all reported results were robust to dF/F method (Supplementary Methods
469 A). To ensure dF/F signal processing did not introduce artifactual scaling or baseline shifts, we
470 also tested several complementary techniques to isolate undistorted F(t) signals where possible
471 and quantified the amount of signal distortion when perfect isolation was not possible (See
472 Methods: dF/F method characterization and validation, below, and Figure S3C).

473

474 **dF/F method characterization and validation.** dF/F calculations are intended to reduce the
475 contribution of slow fluorescence bleaching to fiber photometry signals, and many such methods
476 have been described (Soares et al., 2016; Kim et al., 2019; Mohebi et al., 2019). However, dF/F
477 methods have the potential to introduce artifactual distortion when the wrong method is applied in
478 the wrong setting. Thus, to derive an appropriate dF/F method for use in the context of the self-
479 timed movement task, we characterized and quantified artifacts produced by 4 candidate dF/F
480 techniques.

481

482 *Detailed description of complementary dF/F methods.*

1. Normalized baseline: a commonly used dF/F technique in which each trial's fluorescence is normalized to the mean fluorescence during the 5 s preceding the trial.
2. Low-pass digital filter: F_0 is the low-pass, digital infinite impulse response (IIR)-filtered raw fluorescence for the whole session (implemented in MATLAB with the built-in function *lowpass* with $f_c=5\cdot10^{-5}$ Hz, steepness=0.95).
3. Multiple baseline: a variation of Method 1, in which each trial's fluorescence is normalized by the mean fluorescence during the 5 s preceding the current trial, as well as 5 trials before the current trial and 5 trials after the current trial.
4. Moving average: F_0 is the 200 s moving average of the raw fluorescence at each point (100 s on either side of the measured timepoint).

Although *normalized baseline* (Method 1) is commonly used to correct raw fluorescence signals (F) for bleaching, this technique assumes that baseline activity has no bearing on the type of trial; however, because the mouse decides when to move in the self-timed movement task, it is possible that baseline activity may differ systematically with the mouse's choice on a given trial. Thus, normalizing F to the baseline period would obscure potentially physiologically-relevant signals. More insidiously, if baseline activity *does* vary systematically with the mouse's timing, normalization can also introduce substantial amplitude scaling and y-axis shifting artifacts when correcting F with this method (Figure S3C, middle panels). Thus, Methods 2-4 were designed and optimized to isolate photometry signals minimally distorted by bleaching signals and systematic baseline differences during the self-timed movement task. Methods 2-4 produced the same results in all statistical analyses, and the moving average method is shown in all figures.

506 *Isolating minimally-distorted photometry signals with paired trial analyses of raw fluorescence*.
507 Although slow bleaching prevents comparison of raw photometry signals (F) at one time in a
508 behavioral session with those at another time, the time-course of appreciable bleaching was slow
509 enough in the reported behavioral sessions that minimal bleaching occurred over the course of 3
510 trials (~1 min, Figure S3A). Thus, F was comparable on sets of *paired* trials. To observe the most
511 minimally-distorted photometry signals possible, average F on paired trials was compared (Figure
512 S3B,C). Because dF/F baseline DAN signals were systematically related to lick timing, we
513 compared F baseline signals between all paired trials in which an Early (first-lick between 0.7-
514 2.9s, unrewarded) trial was followed by a Late (first-lick between 3.4-7s, rewarded) trial ("ER"
515 comparison). To ensure systematic differences did not result from subtle bleaching in the paired-
516 trial interval, we reversed the ordering contingency and also compared all Late trials preceding
517 Early trials ("RE comparison"). The same systematic relationship between baseline signals and
518 first-lick time was found for paired trials analyzed by raw F (Figure S3C, left panels).
519
520 *Quantification of artifactual amplitude scaling/baseline shifts introduced by dF/F processing*.
521 Each Candidate dF/F Method was applied to the same Paired Trial datasets described above. The
522 resulting paired-fluorescence datasets were normalized after processing (minimum dF/F=0,
523 maximum=1). The amount of distortion introduced by dF/F was quantified with a Distortion Index
524 (DI), which was calculated as:
525 $$\text{Distortion Index, } DI(t) = abs(F(t)-dF/F(t))$$
526 where F(t) and dF/F(t) are the normalized, paired-trial raw fluorescence signal or dF/F signal at
527 time *t*, respectively. *t* spanned from the beginning of the n-1th trial (-20 s) to the end of the nth trial

528 (20s), aligned to the cue of the nth trial (Figure S3C, bottom panels). The DI shown in plots has
529 been smoothed with a 200 ms moving average kernel for clarity.
530
531 As expected, normalizing fluorescence to the baseline period (*normalized baseline*) erased the
532 correlation of baseline dF/F signals with first-lick time (Figure S3C, middle panels). More
533 insidiously, this also resulted in distortion of GCaMP6f dynamics *during* the timing interval,
534 evident in the diminished difference between E-signals compared to R-signals relative to the
535 shapes observed in the raw fluorescence paired-trial comparison (Figure S3C, middle-bottom
536 panel). However, dF/F Methods 2-4 visually and quantitatively recapitulated the dynamics
537 observed in the raw fluorescence comparison (Figure S3C, right panels).
538
539 These results were corroborated by time-in-session permutation tests in which datasets for single
540 sessions were divided into thirds (beginning of session, middle of session, and end of session). The
541 differences between baseline and ramping dynamics observed in whole-session averages were
542 present even within these shorter blocks of time within the session (i.e., faster ramping and elevated
543 baseline signals on trials with earlier self-timed licks). Furthermore, permutation tests in which the
544 block identity (begin, middle, end) was shuffled showed that this pattern held when trials with
545 earlier first-licks from the end of the session were compared with trials with later first-licks from
546 the beginning of the session (and vice versa).
547
548 **Baseline DAN signal encoding models.** To determine whether baseline DAN signals were best
549 explained by the prior trial outcome ($n-1^{th}$), current trial outcome (n^{th}), or some interaction between
550 the two, we employed two "paired-trial" strategies, one using raw F(t) (to check for robustness to

551 dF/F scaling artifacts, see Methods: dF/F method characterization and validation) and the other
552 using dF/F signals. Mean baseline activity was measured in sliding windows (100ms divisions, 2s
553 windows) from the n-1th trial ITI-start to the nth trial cue-onset. We abbreviate trial outcomes as
554 E="early" (first-lick between 0.7-3.333s, unrewarded) and R="rewarded" (first-lick at 3.334-7s,
555 rewarded), with the position of the letters indicating the order of consecutive trials (e.g., "ER" or
556 "RE").
557
558 Baseline differences between all paired, consecutive trials (EE, ER, RE, RR) were compared for
559 dF/F signals, controlling for the prior trial outcome by 4-factor ANOVA (factor 1: previous trial
560 outcome (E or R, 1 degree of freedom (df)); factor 2: subsequent trial outcome (E or R, 1 df);
561 factor 3: presence of licks in window (df: 1); factor 4: session ID (df: 112), comparison: baseline
562 dF/F activity within 2s sliding windows, Figure 4C, top). Relative influence on baseline dF/F by
563 the n-1th and nth trial outcome was estimated as the relative size of the 4-way ANOVA F-statistic
564 as a function of time (Figure 4C, bottom). Similar results were obtained on single sessions by 2-
565 factor and 3-factor ANOVA. Similar results were also obtained comparing triplets of consecutive
566 trials (EER, RRE, EEE, RRR). Both trial 2 and trial 3 in the triplet were preceded by a trial of the
567 same outcome type, and thus we could compare baseline raw F(t) signals with minimal bleaching
568 distortion.
569
570 **DAN signal encoding model.** To test the independent contribution of each task-related input to
571 the photometry signal and select the best model, we employed a nested fitting approach, in which
572 each dataset was fit multiple times (in "nests"), with models becoming progressively more

complex in subsequent nests. The nests fit to the GCaMP6f photometry data employed the inputs $X^{(j)}$ at each jth nest:

Null Model: $X^{(0)} = x_0$

Nest 1: $X^{(1)} = X^{(0)} + \text{tdt}$

Nest 2: $X^{(2)} = X^{(1)} + \text{cue} + \text{first-lick}$

Nest 3: $X^{(3)} = X^{(2)} + \text{EMG/accelerometer}$

Nest 4: $X^{(4)} = X^{(3)} + \text{time-dependent baseline offset}$

Nest 5: $X^{(5)} = X^{(4)} + \text{stretch representing percentages of interval}$

Overfitting was penalized by ridge regression, and the optimal regularization parameter for each nest was obtained by 5-fold cross-validation to derive the final model fit for each session. Model improvement by each input was assessed by the percentage loss improvement at the nest where the input first appeared compared to the prior nest. The loss improvement of Nest 1 was compared to the Null Model (the average of the photometry timeseries). The nested model of tdt control photometry signals was the same, except Nest 1 was omitted.

The GLM for each nest takes the form:

$$Y = \Theta X^{(j)}$$

Where Y is the $1 \times n$ vector of the photometry signal across an entire behavioral session (n is the total number of sampled timepoints); $X^{(j)}$ is the $d \times n$ design matrix for nest j, where the rows correspond to the d_j predictors for nest j and the columns correspond to each of the n sampled timepoints of Y; and Θ is the $d \times 1$ vector of fit weights.

Y is the concatenated photometry timeseries taken from trial start (lights off) to the time of first lick. Because of day-to-day/mouse-to-mouse variation (ascribable to many possible sources, *e.g.*, different neural subpopulations, expression levels, behavioral states, *etc.*), each session was fit separately.

The d_j design matrix predictors were each scaled (maximum amplitude 1) and grouped by input to the model. The timing-independent inputs were: 1. Null offset (x_0, 1 predictor), 2. tdt (1 predictor), 3. cue (24 predictors), 4. first-lick (28 predictors), and 5. EMG/accelerometer (44 predictors). The timing-dependent inputs were: 6. timing-dependent baseline offset (1 predictor), 7. stretch (500 predictors).

To reduce the number of predictors, cue, first-lick and EMG/accelerometer predictors (Figure S4C) were composed from sets of basis kernels as described previously (Park et al., 2014; Runyan et al., 2017). The cue basis kernels were spaced 0-500 ms post-cue and first-lick basis kernels were spaced -500ms-0ms relative to first-lick, the typically-observed windows of stereotypical sensory and motor-related neural responses. For nuisance movements (EMG/accelerometer), events were first discretized by thresholding (Figure S4B) and then convolved with basis kernels spanning -500 to 500 ms around the event. This window was consistent with the mean movement-aligned optical artifact observed in the tdt channel. The timing-dependent baseline offset was encoded as a constant offset spanning from lamp-off until first-lick, with amplitude taken as linearly proportional to the timed interval on the current trial. The timing-dependent stretch input was composed of 500 predictors, with each predictor containing 1's tiling 0.05% of the cue-to-lick

interval, and 0's otherwise (Figure S4D). Importantly, the stretch was not constrained in any way to form ramps.

Basis sets were optimized to minimize Training Loss, as calculated by mean squared error of the unregularized model:

$$\text{argmin}_{X^{(j)}}(\text{Training Loss}(\Theta) = 1/n \text{ sum}(Y - \Theta X^{(j)})^2)$$

Superfluous basis set elements that did not improve Training Loss compared to the Null Model were not included in the final model. Goodness of the training fit was assessed by Akaiki Information Criterion (AIC), Bayesian Information Criterion (BIC), R^2, and Training Loss. The optimal, regularized model for each nest/session was selected by 5-fold cross-validation in which the regularization parameter, λ_j, was optimized for minimal average Test Loss:

$$\text{argmin}_{\lambda_j} (\text{Test Loss}(\Theta,\lambda_j) = 1/n \text{ sum}(Y - \Theta X^{(j)})^2 + \lambda_j|\Theta|^2)$$

Test Loss for each optimal model was compared across nests to select the best model for each session. Models were refit with the optimal λ_j to obtain the final fit.

Model error was simulated 1000 times by redrawing Θ coefficients consistent with the data following the method described by Gelman and Hill, 2006, and standard errors were propagated across sessions. The absolute value of each predictor was summed and divided by the total number of predictors for that input to show the contribution of the input to the model (Extended Data Fig. 7g). To simulate the modeled session's photometry signal for each nest *j*, Yfit was calculated as

$\Theta X^{(j)}$ and binned by the time of first lick relative to the cue. The error in the simulation was shown by calculating $Yfit_{sim} = \Theta_{sim} X^{(j)}$ for 300 simulated sets of Θ_{sim}.

Principle component analysis (PCA)

Unsmoothed ramping intervals for GCaMP6f photometry timeseries were fit with PCA and reconstructed with the first three principle components (PCs). To derive a PCA fit matrix with ramping intervals of the same number of samples, the length of each trial was scaled up by interpolation to the maximum ramping interval duration:

7s–0.7s cue buffer–0.6s first-lick buffer=5.7s: 5700 sample ramping interval

Following PC-fitting, datasets were down-sampled to produce a fit of the correct time duration. Trials where the ramping interval was <0.1s were excluded to exclude noise from down-sampling.

First-lick time decoding model

A nested, generalized linear model was derived to predict the first-lick time on each trial in a session and quantify the contribution of previous reward history and photometry signals to the prediction. The model was of the form:

$$\log(y) = bx$$

where b is a vector of fit coefficients and x is a vector of predictors. The nested model was constructed such that predictors occurring further back in time (such as reward history) and confounding variables (such as tdt photometry signals) were added first to determine the additional variance explained by predictors occurring closer to the time of first-lick, which might otherwise obscure the impact of these other variables. The predictors, in order of nesting, were:

Nest 0: b0 (Null model, average log-first-lick time)

662 Nest 1: b1 = b0 + first-lick time on previous trial

663 Nest 2-5: b2 = b1 + previous trial outcome (1,0)*

664 Nest 6: b3 = b2 + median photometry signal in 10s window before lamp-off

665 Nest 7: b4 = b3 + median photometry signal from lamp-off to cue

666 Nest 9: b5 = b4 + tdt threshold crossing time**

667 Nest 10: b6 = b5 + GCaMP6f threshold crossing time**

668

669 where all predictors were normalized to be in the interval (0,1).

670

671 * Outcomes included (in order of nest): Reaction (first-lick before 0.5s), Early (0.5-3.333s),
672 Reward (3.333-7s), ITI (7s-17s). No-lick was implied by all four outcomes encoded as zeros.
673 ** Details on threshold-crossing time and alternative models included in Supplementary Methods
674 B.

675

676 To exclude the sensory- and motor-related transients locked to the cue and the first-lick events in
677 the threshold-crossing nests, the ramping interval was conservatively defined as 0.7s post-cue up
678 until 0.6s before first-lick, and the minimum ramping interval for fitting was 0.1s. Thus, for a trial
679 to be included in the model, the first lick occurred between 1.4s to 17s (end of trial).

680

681 Initial model goodness of fit was assessed by R^2, mean-squared loss and BIC. Models were 5-fold
682 cross-validated with ridge regression at each nest to derive the final models, as described above.
683 95% confidence intervals on model coefficients were calculated by 2-sided t-test with standard
684 errors propagated across sessions.

685

686 **B. <u>Derivation of threshold and alternative decoding models</u>**

687 ***Derivation of threshold models***

688 As a metric of the predictive power of ramping DAN signals on first-lick time, we derived a
689 threshold-crossing model. A threshold-crossing event was defined as the first time after the cue
690 when the photometry signal exceeded and remained above a threshold level up until the time of
691 first-lick on each trial. Importantly, while the analysis approach is reminiscent of pacemaker-
692 accumulator models for timing, we make no claims that the analysis is evidence for pacemaker-
693 accumulator models. Rather threshold-crossing times provided a convenient metric to compare the
694 rate of increase in signals between trials.

695

696 Photometry timeseries for GCaMP6f and tdt were de-noised by smoothing with a 100ms Gaussian
697 kernel (kernel was optimized by grid screen of kernels ranging from 0, 30, 50, 80, 100, 150, 200
698 ms to minimize noise without signal distortion). To completely exclude the sensory- and motor-
699 related transients locked to the cue and the first-lick events, the ramping interval was
700 conservatively defined as 0.7 s post-cue up until 0.6 s before the first-lick. To eliminate chance
701 crossings due to noise, we imposed a stiff, debounced threshold condition: to be considered a
702 threshold crossing event, the photometry signal had to cross the threshold from low-to-high and
703 remain above this level until the end of the ramping interval.

704

705 To derive an unbiased threshold for each session, we tested 100 evenly-spaced candidate threshold
706 levels spanning the minimum-to-maximum photometry signal during the ramping interval for each
707 session. Depending on threshold level, some trials never crossed, i.e., signal always remained

below threshold or started and ended above threshold. Thus, the lowest candidate threshold for which there was a maximum number of trials crossing during the timing interval was selected as the "mid-level" threshold-crossing point. This threshold was specific to each photometry signal tested on each session. Threshold-crossing time was included in the decoding model as the normalized time on the ramping interval (0,1). If a trial never crossed threshold, it was encoded as a zero. If no trials ever crossed threshold, the threshold predictor was encoded as a vector of ones, thus penalizing the model for an additional predictor but providing no new information.

Multi-threshold Model

An alternative model employed 3 unbiased thresholds: 1) the lowest threshold with ≥50 trials crossing ("min"); 2) the lowest threshold with the most crossings ("mid," described above); and 3) the highest threshold with ≥50 trials crossing ("max"). For tdt datasets, trials rarely met the monotonic threshold constraint (usually the signals oscillated above and below the threshold throughout the ramping interval, failing to meet the debouncing constraint). Thus, to include tdt signals as conservatively as possible, we relaxed the 50-trial minimum constraint, taking the threshold with the most trials crossing, which was usually around 10 or fewer. The addition of more thresholds did not substantially improve the cross-validated model compared to the single-threshold model (Figure S5C).

Principle component analysis (PCA) threshold-crossing models

In another version of the decoding model, the threshold-crossing procedures were applied to ramping intervals fit with the first three PCs (as described in Methods) to derive a PCA version of the single-threshold and multi-threshold models. PCA analysis on tdt datasets showed no

731 consistent PCs, and thus these PCs were not included in the decoding model. Instead, the actual
732 tdt data was employed in the threshold model as in the other models described.
733
734
735 **Optogenetics—naïve/expert control sessions.** To determine whether optogenetic stimulation
736 directly elicited or prevented licking, licking behavior was first tested outside the context of the
737 self-timed movement task on separate sessions in the same head-fixed arena but with no cues or
738 behavioral task. Opsin-expressing mice were tested before any exposure to the self-timed
739 movement task ("Naïve") as well as after the last day of behavioral recording ("Expert"). In ChR2
740 control sessions, stimulation (5mW 425nm light, 3s duration, 10Hz, 20% duty cycle) was applied
741 randomly at the same pace as in the self-timed movement task. stGtACR2 control sessions were
742 conducted similarly (12mW 425mW light, 3s duration, constant illumination); but to examine if
743 inhibition could block ongoing licking, we increased the baseline lick-rate by delivering juice
744 rewards randomly (5% probability checked once every 5s).
745
746 **Optogenetics—self-timed movement task.** SNc DANs were optogenetically manipulated in the
747 context of the 3.3s self-timed movement task. To avoid over-stimulation, light levels were adjusted
748 to be subthreshold for eliciting overt movements (Coddington & Dudman, 2018), and mice were
749 not stimulated on consecutive days.
750
751 Activation: SNc cell bodies were illuminated bilaterally (ChR2: 0.5-5mW 425nm blue LED light;
752 ChrimsonR 550nm lime or 660nm crimson) on 30% of trials (10Hz, 20ms up-time starting at

753 cue onset and terminating at first-lick). DAN terminals in DLS were stimulated bilaterally via
754 tapered fiber optics on separate sessions.
755 Inactivation: SNc cell bodies were illuminated bilaterally (stGtACR2: 12 mW 425 nm blue light)
756 on 30% of trials (constant illumination starting at cue onset and terminating at first lick).
757
758 **Quantification of optogenetic effects.** The difference in the distribution of trial outcomes between
759 stimulated and unstimulated trials on *each session* was quantified in four ways.
760 1. 2-Sample Unsigned Kolmogorov-Smirnov Test.
761 2. Difference in empirical continuous probability distribution function (cdf). The difference
762 in the integral of the stimulated and unstimulated cdf (dAUC) was calculated for each
763 session from 0.7-7s. Effect size was quantified by permutation test, wherein the identity of
764 each trial (stimulated or unstimulated) was shuffled, and the distribution of dAUCs for the
765 permuted cdfs was calculated 10,000x. Results were reported for all sessions.
766 3. Difference in mean movement time. Movement times on stimulated and unstimulated trials
767 were pooled and the distribution of movement time differences was determined by non-
768 parametric bootstrap, in which a random stimulated and unstimulated trial were drawn from
769 their respective pools 1,000,000x and the difference taken. The mean of each session's
770 bootstrapped distribution was compared across sessions by the 1,000,000x bootstrapped
771 difference of the mean between sessions of different categories.
772 4. Difference in median movement time. Same as above but with median.
773
774 **CODE AVAILABILITY**

All custom behavioral software and analysis tools are available at https://github.com/harvardschoolofmouse.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

FIGURE LEGENDS

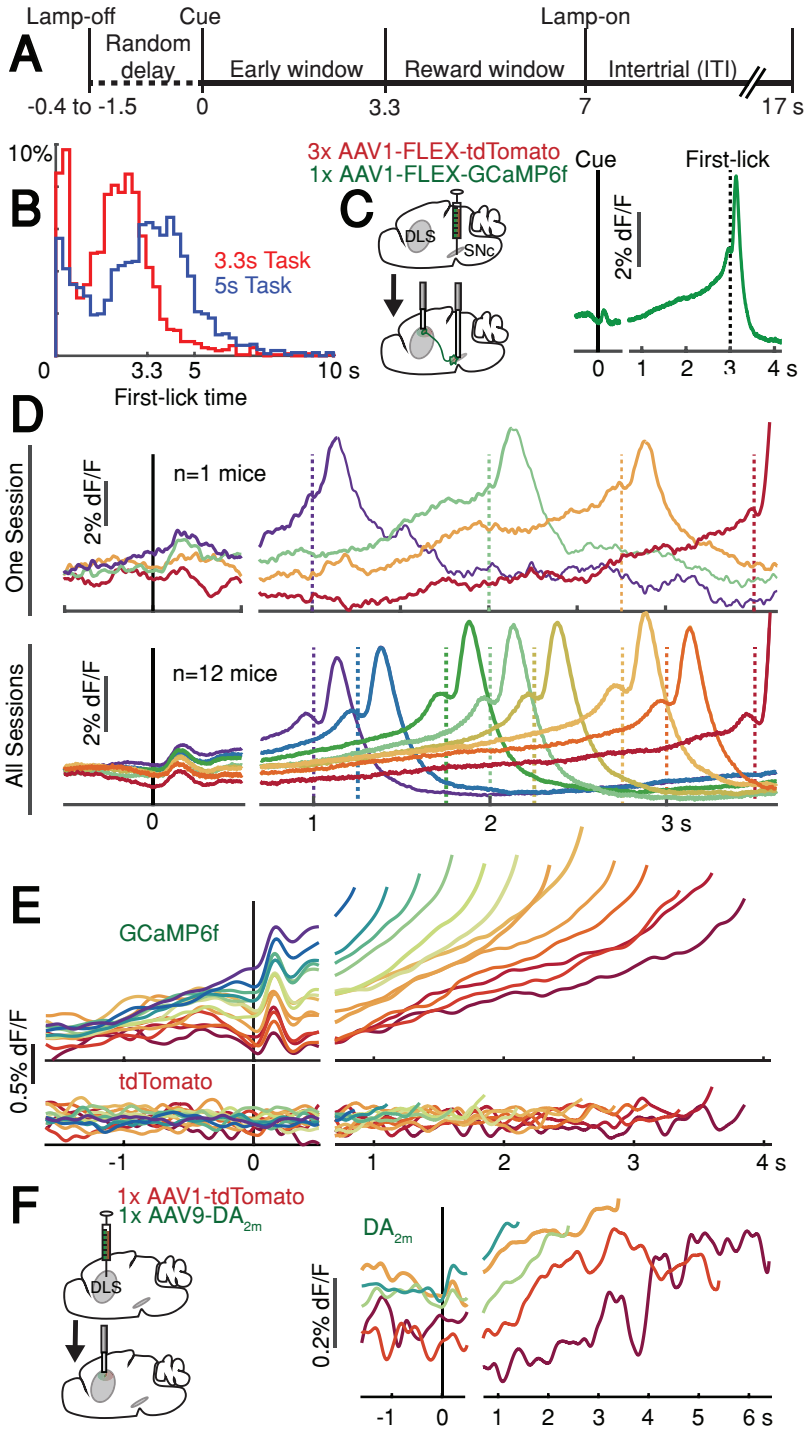

Figure 1. SNc DAN signals preceding self-timed movement

(A) Self-timed movement task. **(B)** First-lick distribution for two task variants performed by the same mouse. Red: 3.3 s-reward boundary (4 sessions); Blue: 5 s-reward boundary (4 sessions). For all mice, see Figure S1.

(C) Left: surgical strategy for GCaMP6f/tdTomato fiber photometry. Right: average SNc DAN GCaMP6f response for first-licks between 3-3.25 s (12 mice). Left of plot: cue-aligned average; right of plot: first-lick-aligned average. Vertical dashed line: first-lick time. Break in axis indicates change in plot alignment.

(D) Average SNc DAN GCaMP6f responses for different first-lick times (12 mice).

(E) Comparison of average DAN GCaMP6f and tdTomato responses (12 mice). Traces plotted up to 150 ms before first-lick.

(F) Average DLS DA_{2m} signals (4 mice).

See also Figures S1-S3.

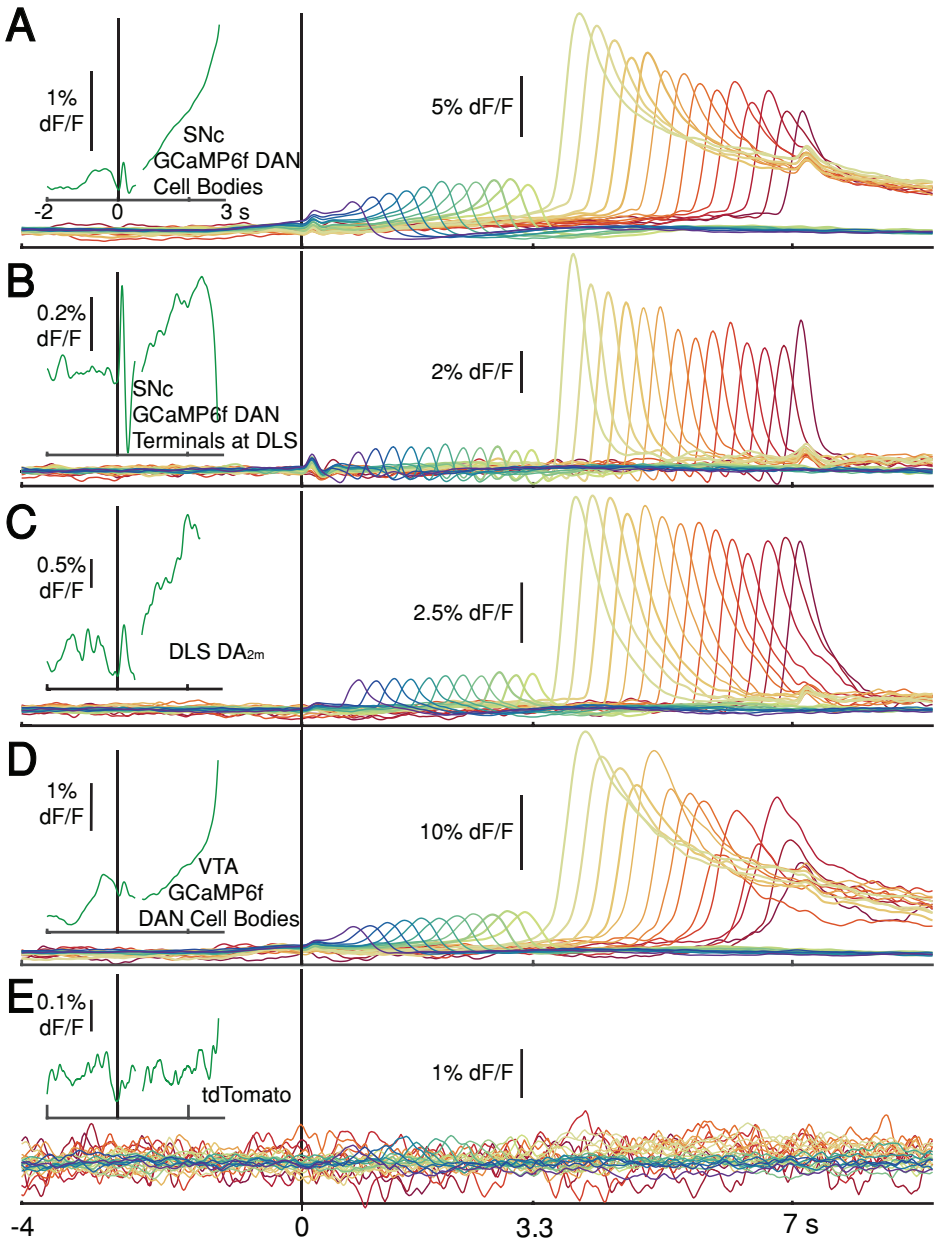

Figure 2. Cue-aligned average photometry signals showing reward-related responses: all mice, all fluorophores

(A) DAN GCaMP6f signals at SNc cell bodies (12 mice).

(B) DAN GCaMP6f signals at axon terminals in DLS. The sharp, downward deflection immediately prior to movement onset was observed in every mouse (12/12) on every session and was not explained by movement artifacts. There appears to be a rapid "off" response.

(C) Striatal DA_{2m} signals at DLS (4 mice).

(D) DAN GCaMP6f signals at VTA cell bodies (4 mice).

(E) tdTomato signals (all sites, all sessions).

Insets: Average signals for first-licks occurring between 3-3.25, aligned to cue (left of axis break) and aligned to first-lick (right of axis break). Traces plotted up till approximate movement onset (150 ms before first-lick).

See also Figures S2-S3.

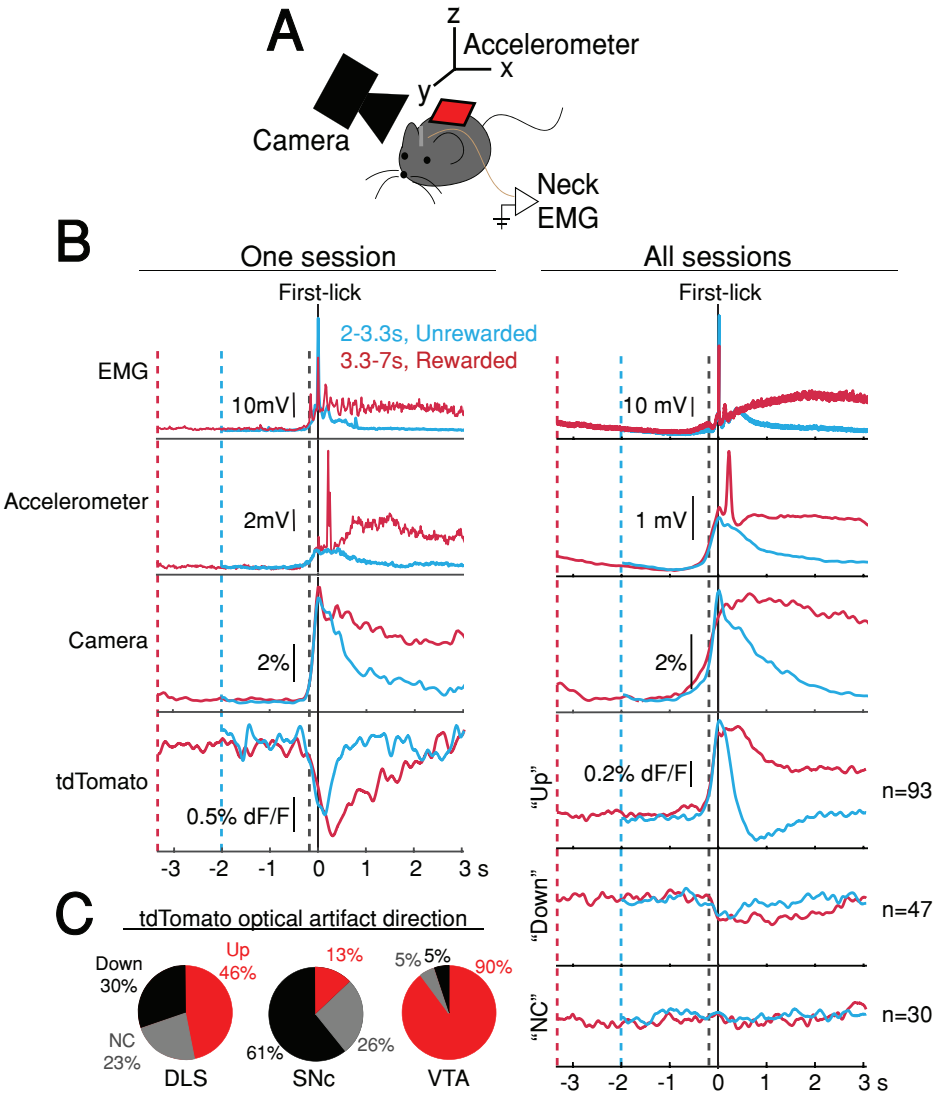

Figure 3. Movement controls reliably detect movements, but there is no systematic difference in movement before first-lick during the timing interval

(A) Schematic of movement controls.

(B) First-lick-aligned average movement signals on rewarded (red) and unrewarded (blue) trials. Pre-lick traces are truncated at the nearest cue-time for the averaged traces (dashed red, dashed blue). Left: one session; Right: all sessions. Dashed grey line: time of earliest-detected movement on most sessions (150 ms before first-lick). Average first-lick-locked tdt optical artifacts showed inconsistent directions even within the same session. Averages for all three types of artifact (consistently up, "Up"; consistently down, "Down"; and not consistent "NC") shown for all sessions.

(C) Breakdown of average tdt artifact direction by session at each recording site.

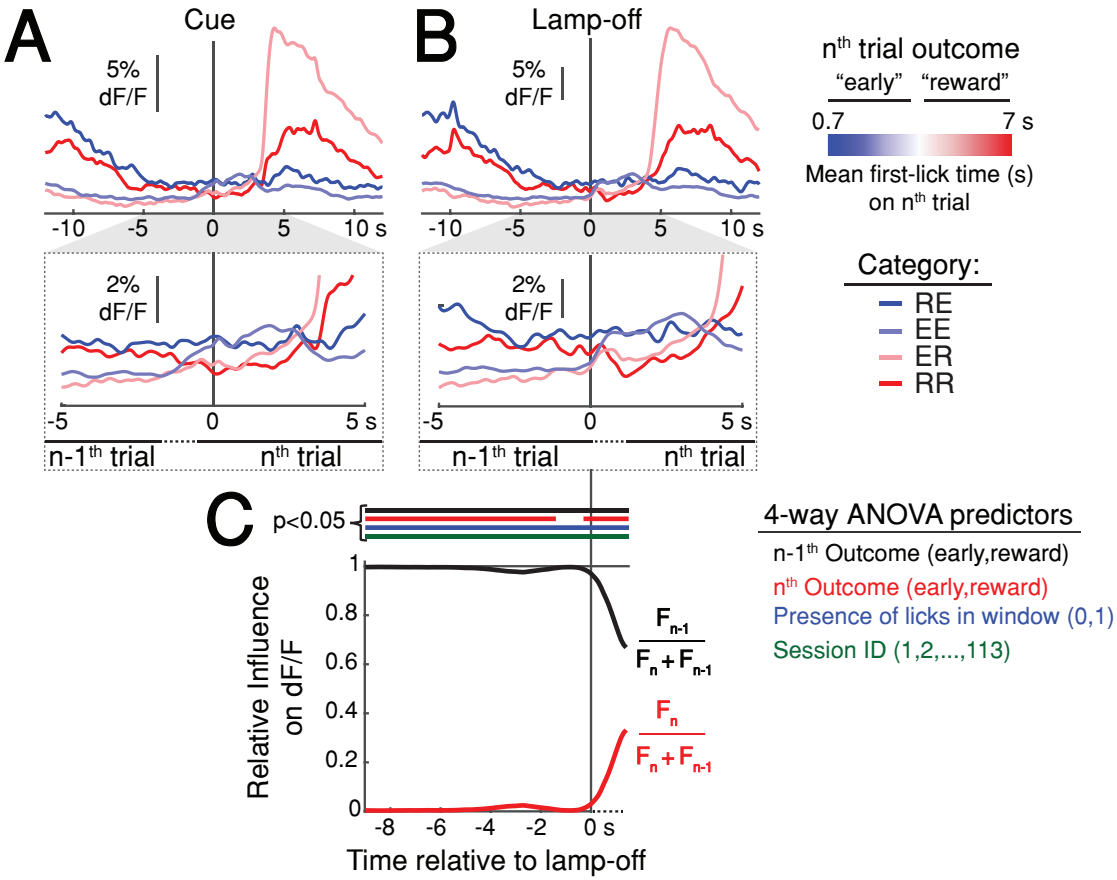

Figure 4. Baseline SNc DAN signals predict trial outcome, even when controlling for prior trial outcome and ongoing movement

Paired, consecutive trials were pooled into 4 categories based on the n-1th and nth trial outcomes (4 mice, the 17 sessions with highest signal to noise and number of trials). Categories: Early-Early (EE, where n-1th outcome is early, nth outcome is early: 2254 trials), Early-Reward (ER: 730 trials), Reward-Early (RE: 190 trials) and Reward-Reward (RR: 174 trials). To control for the contribution of movement to baseline signals, plots are shown only for trials with no licking during the last 5 seconds of the ITI before Lamp-off.

(A) Cue-aligned average DAN signals become more predictive of nth trial outcome as the cue time approaches.

(B) Lamp-off-aligned average DAN signals show "resetting" effect after the houselamp turns off. Before lamp-off, average DAN signals reflect the n-1th trial outcome; subsequently they reflect the nth trial outcome.

(C) Selectivity index taken on single trials quantifies the relative contribution of n-1th and nth trial outcomes to the prediction of the baseline signal. (Index calculated to exclude timepoints after the nth trial first-lick).

See also Figure S3.

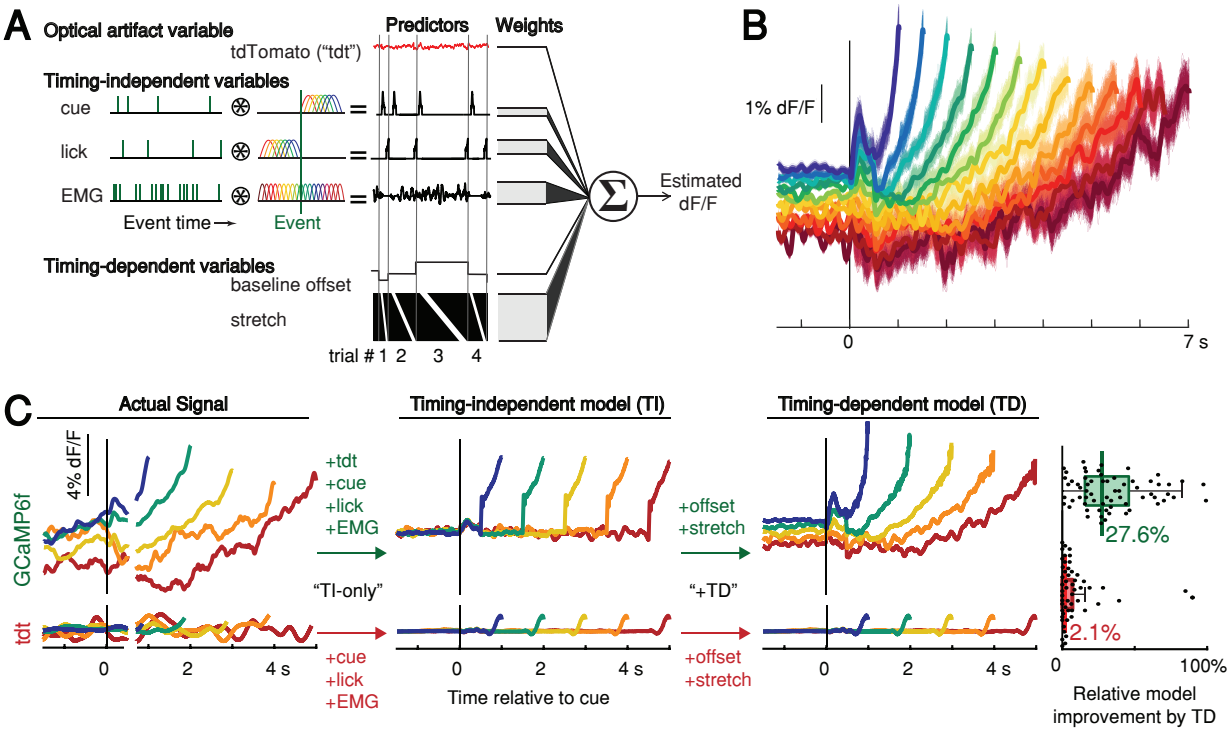

Figure 5. Contribution of optical artifacts, task variables and ongoing movements to SNc GCaMP6f signals

(A) Nested encoding model comparing the contribution of timing-independent predictors (TI) to the contribution of timing-dependent predictors (TD).

(B) Predicted dF/F signal for one session plotted up to time of first-lick. Model error simulated 300x (shading).

(C) Nested encoding model for 1 session showing the recorded signal (1st panel), the timing-independent model (2nd panel), and the full, timing-dependent model with all predictors (3rd panel). Top: GCaMP6f; Bottom: tdTomato (tdt). Right: relative loss improvement by timing-dependent predictors (68 sessions, line: median, box: lower/upper quartiles, whiskers: 1.5x IQR, dots: single sessions).

See also Figure S4.

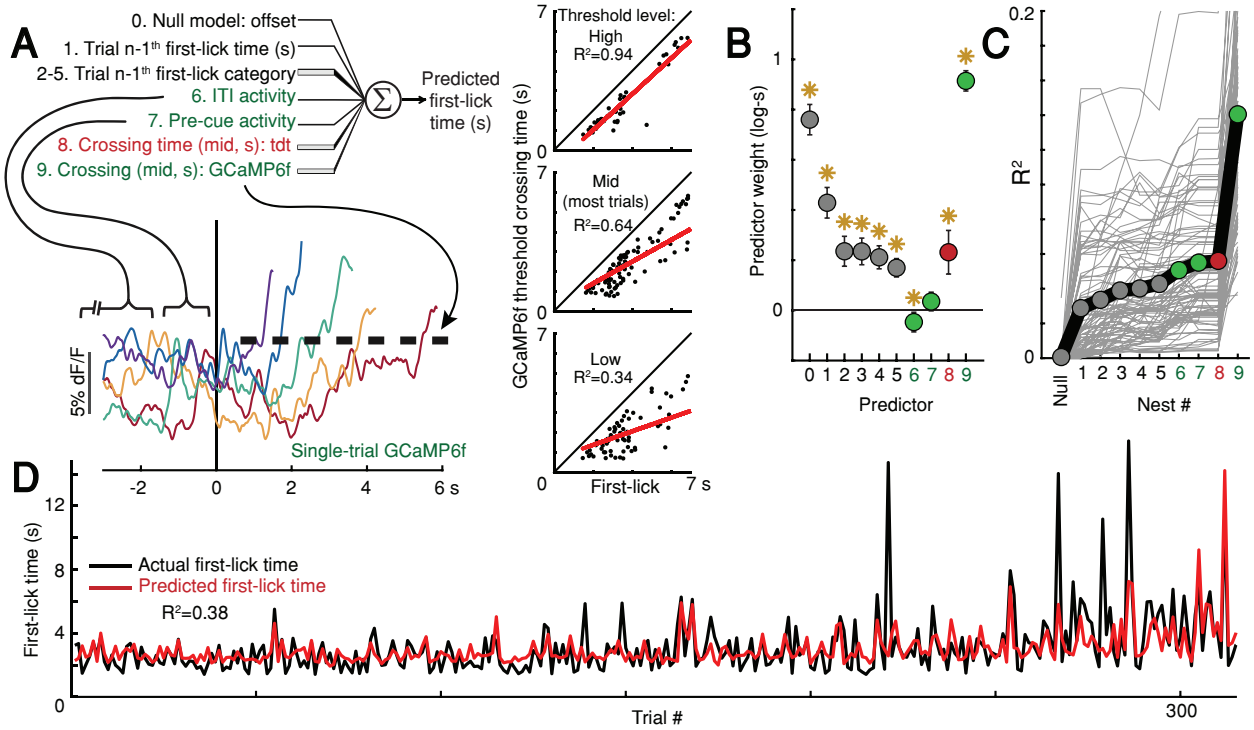

Figure 6. Single-trial baseline and ramping DAN signals predict first-lick timing

(A) Nested decoding model. Top-left: schematic. Bottom-left: single-trial cue-aligned SNc DAN GCaMP6f signals from one session (6 trials shown for clarity). Traces plotted up to first-lick. Right: threshold crossing model. Low/Mid/High: threshold amplitude. Grey dots: single trials.
(B) Model weights, 94 sessions. Error bars: 95% CI, *: p<0.05. Numbers indicate nesting-order.
(C) Variance explained by each model nest. Grey lines: single sessions; thick black line: average. For model selection, see **Figure S5C.**
(D) Predicted vs. actual first-lick time, same session as **6A**.
See also **Figure S5.**

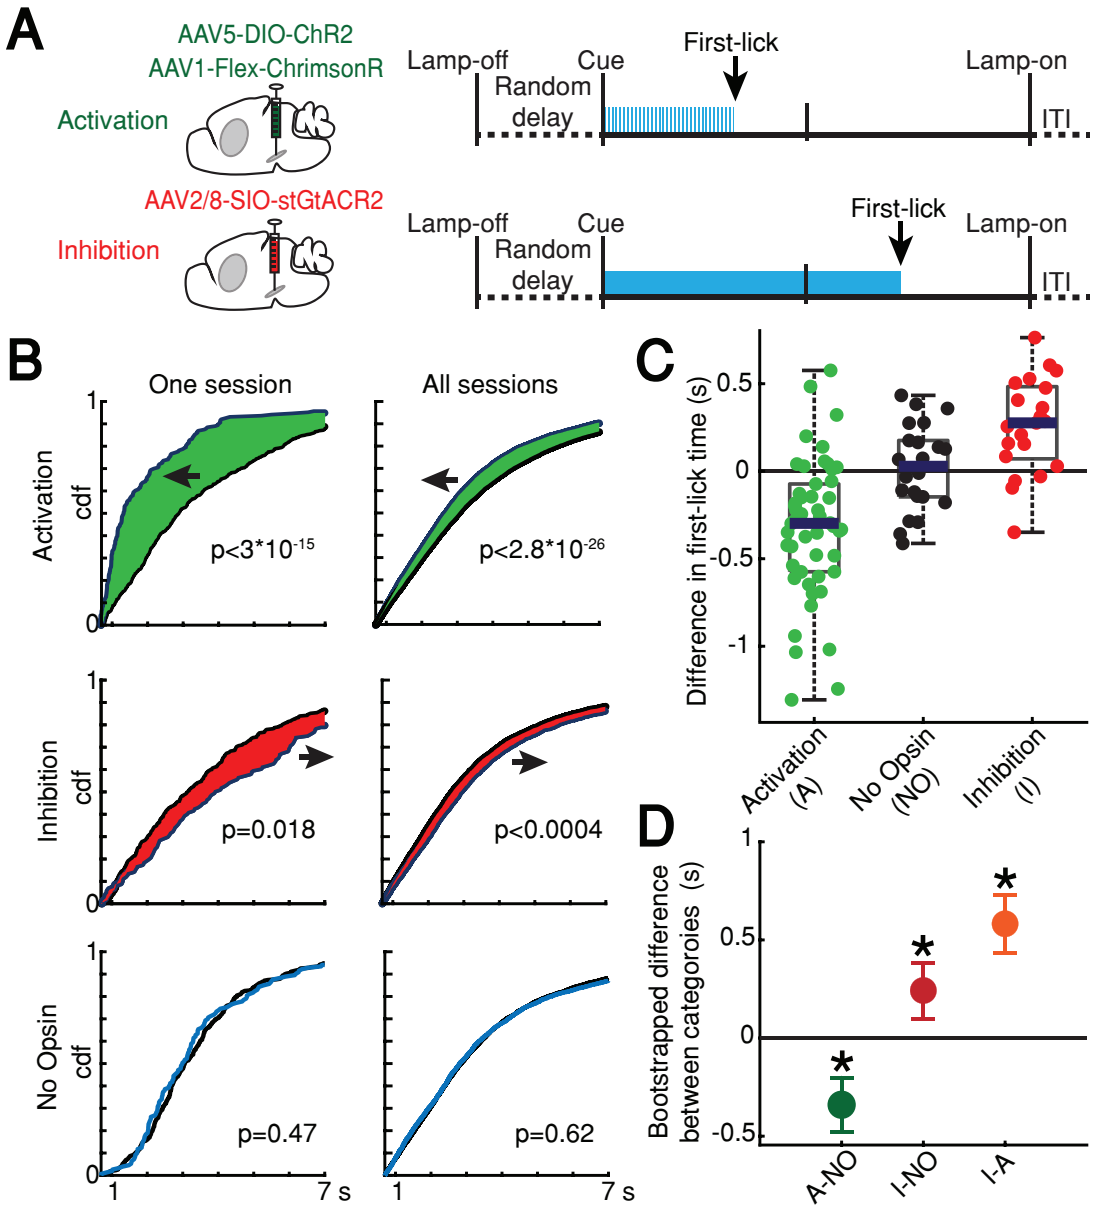

Figure 7. Optogenetic DAN manipulation systematically and bidirectionally shifts the timing of self-timed movements

(A) Strategy for optogenetic DAN activation or inhibition. Mice were stimulated from cue-onset until first-lick or 7 s.

(B) Empirical continuous probability distribution functions (cdf) of first-lick times for stimulated (blue line) versus unstimulated (grey line) trials**.** Arrow and shading show direction of effect. P-values calculated by Kolmogorov-Smirnov test (for other metrics, see Extended Data Fig. 9b-e).

(C) Mean bootstrapped difference in first-lick time, stimulated-minus-unstimulated trials. Box: upper/lower quartile; line: median; whiskers: 1.5x IQR; dots: single sessions.

(D) Comparison of mean first-lick time difference across all sessions. Error bars: 95% confidence interval (*: $p < 0.05$).

See also Figures S6 and S7.

SUPPLEMENTAL INFORMATION

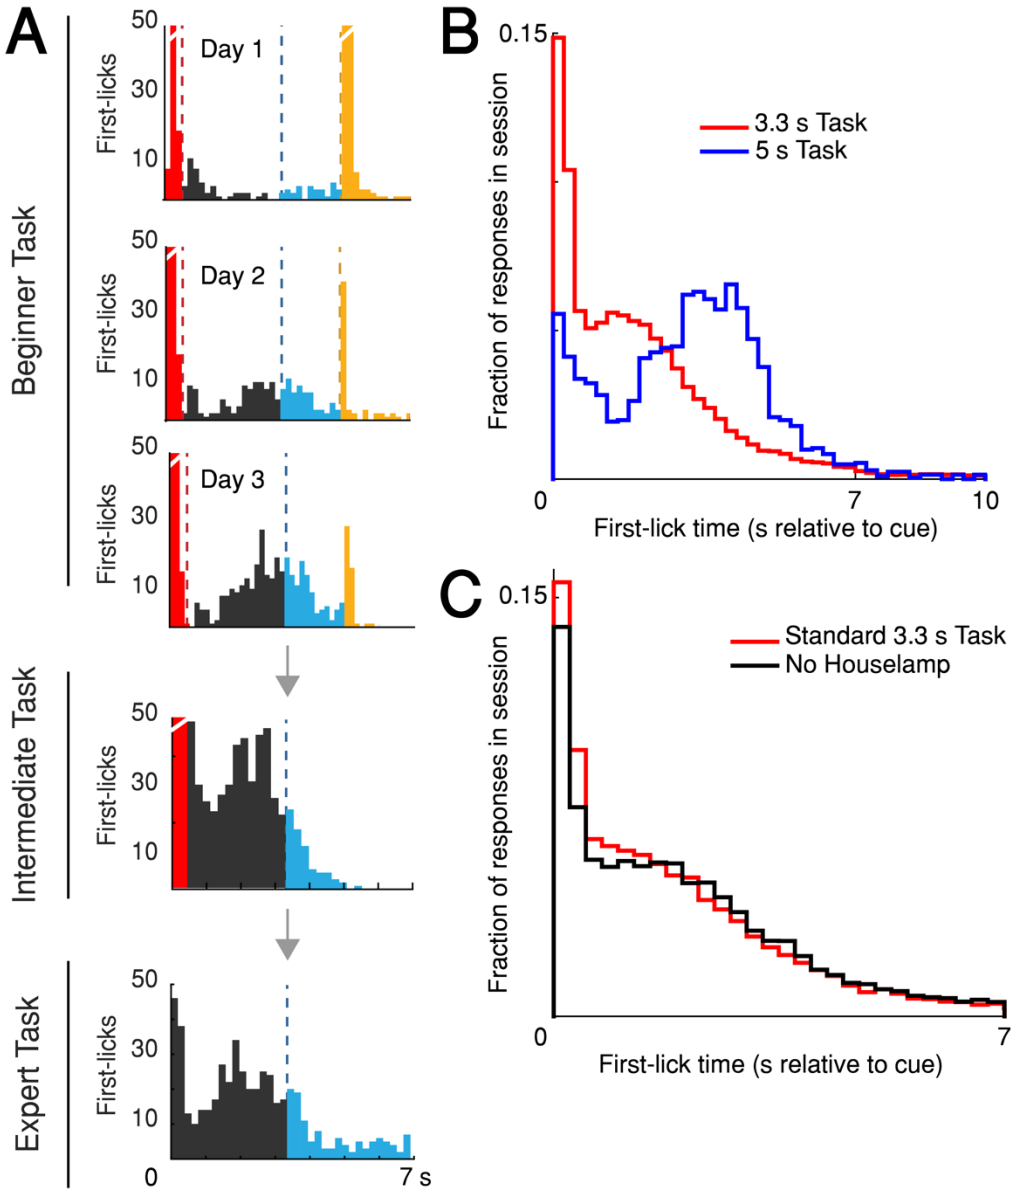

Figure S1. Self-timed movement task learning and variations, Related to Figure 1

(A) Task learning. Histogram of first-licks from single sessions at different stages of training. Bar color indicates first lick category (red: reaction, grey: early, blue: operant-rewarded, gold: Pavlovian-rewarded). Diagonal slash across bar top indicates bar height truncated for clarity.

(B) First-lick distributions from tasks with different target timing intervals (16 mice, 152 sessions). Red: 3.3 s reward-boundary. Blue: 5 s reward-boundary. Mice adjust behavior to the timing-contingencies of the task.

(C) First-lick distributions during behavior with and without houselamp cues. Red: standard 3.3 s task; Black: 3.3 s task omitting houselamp cues (4 mice, 4-5 sessions/mouse on each version of the task). Mice time their first-licks relative to the start cue, not the houselamp.

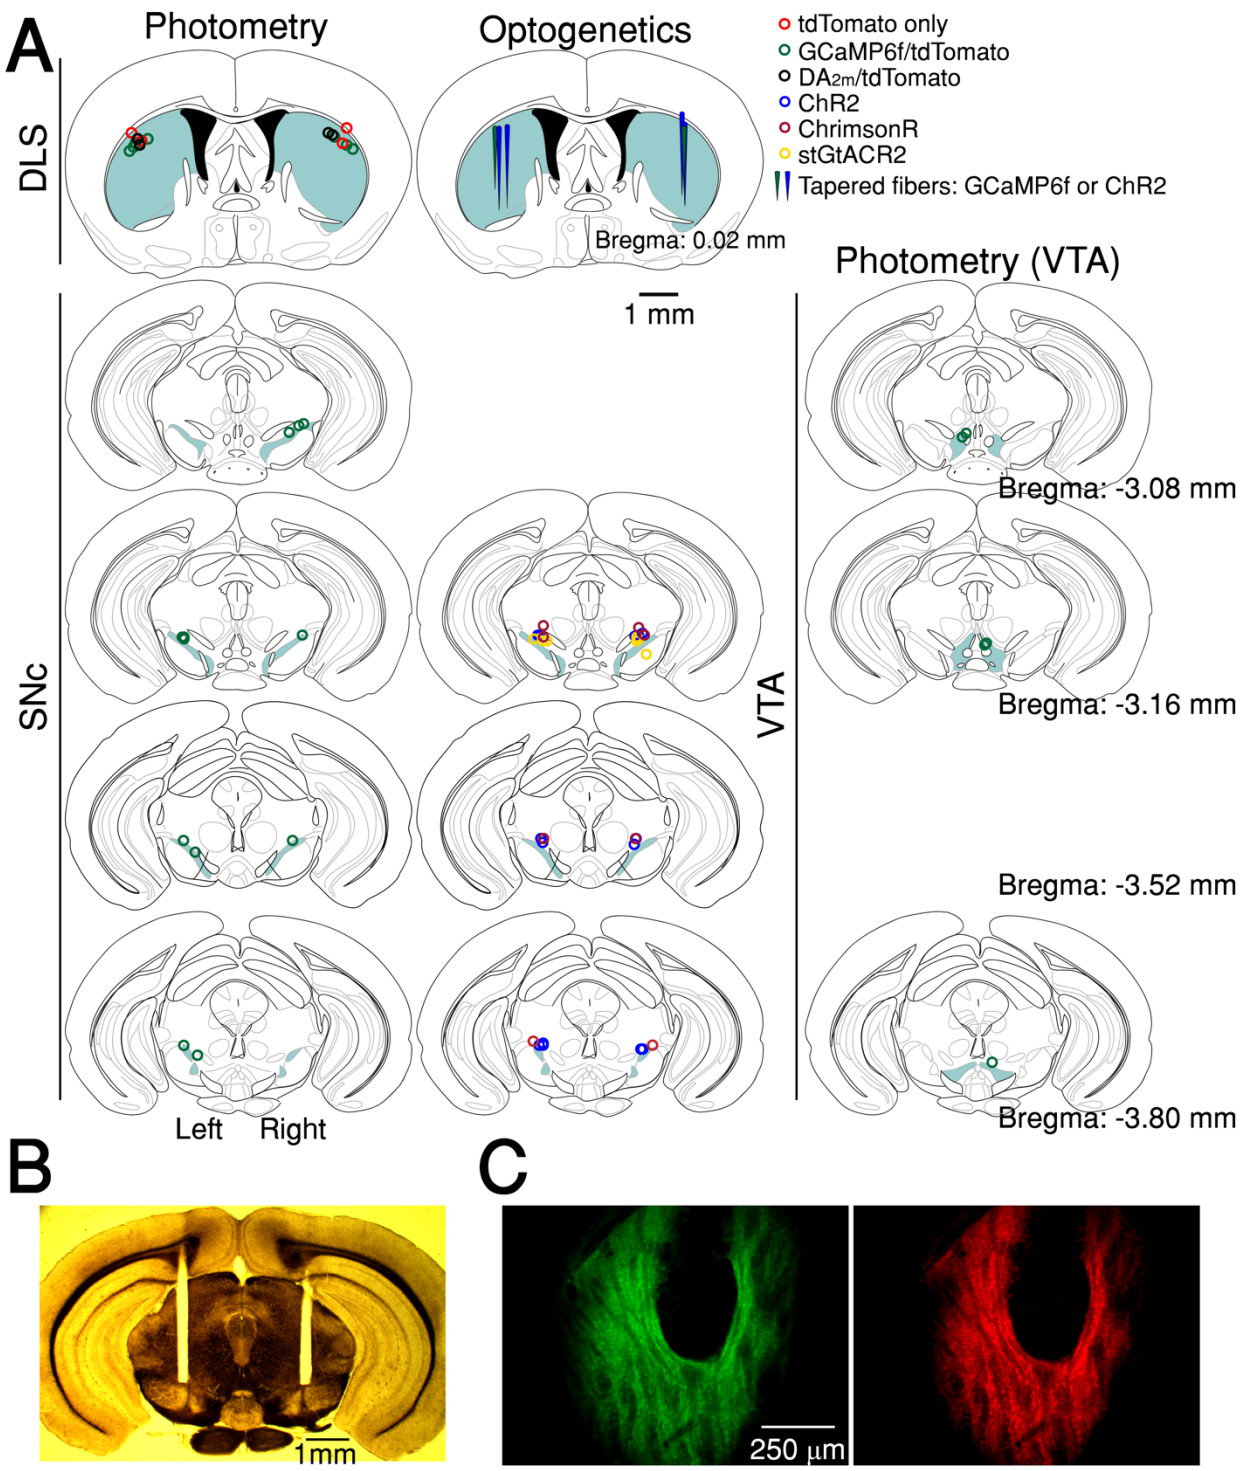

Figure S2. Fiber optic placement and histology, Related to Figures 1-2 and 7.
(A) Approximate fiber positions, all mice.
(B) Brightfield microscopy with polarized filter on a freshly-cut brain slice showing bilateral fiber placement at SNc (stGtACR2).
(C) Example of co-expression of green and red fluorophores relative to fiber optic tip (Left: DA_{2m}, Right: tdTomato).

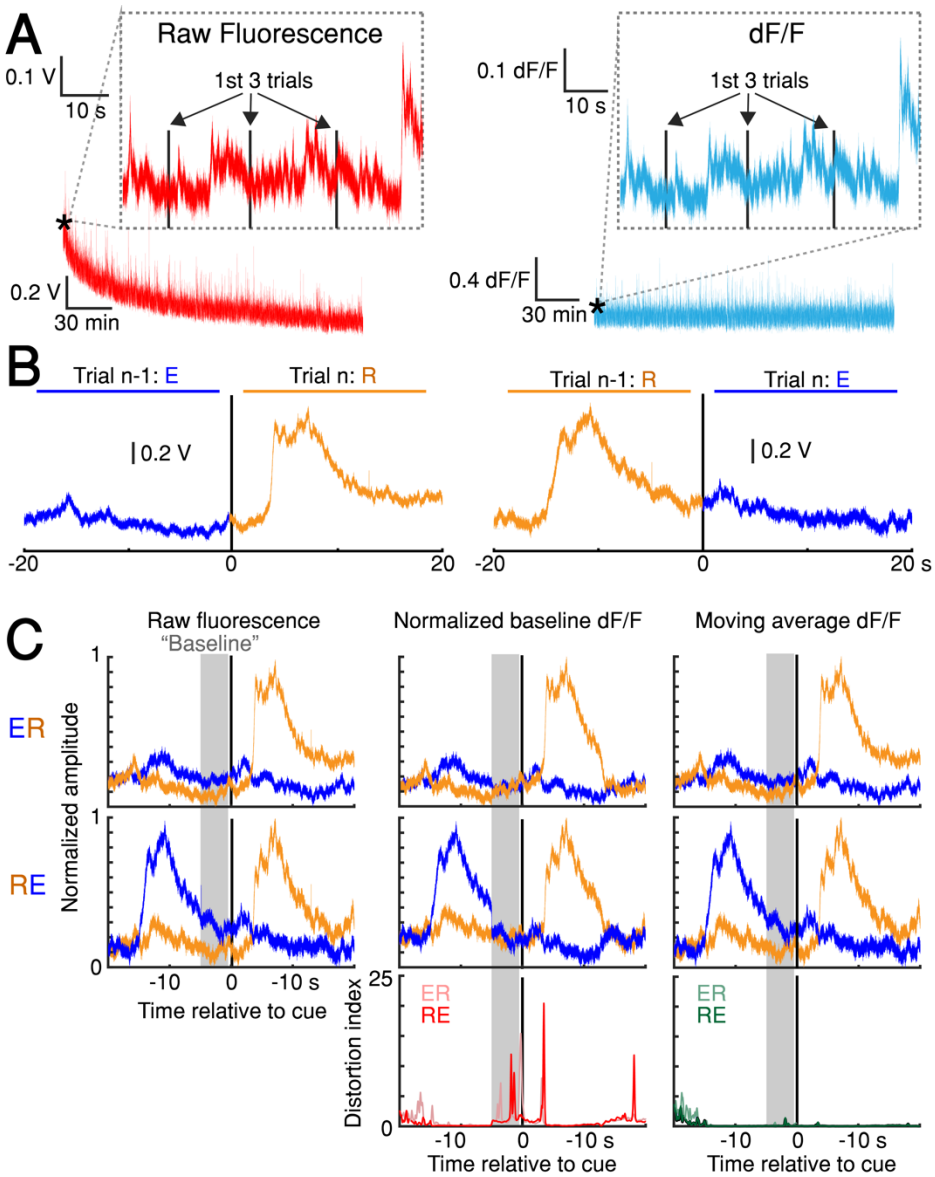

Figure S3. dF/F method validation, Related to Figures 1-2 and 4

(A) Minimal bleaching occurs over a 3-trial window. Left: slow, raw fluorescence bleaching across one session. Left inset: Comparison to bleaching across the first 3 trials (~1 min). Right: dF/F removes bleaching dynamics. Right inset: the same 3-trial window shown for dF/F-treated signal. **(B)** Average raw fluorescence on paired, consecutive trials from one session aligned to cue on the nth trial. Left: $n\text{-}1^{th}$ trial was early, n^{th} trial was rewarded ("ER" condition). Right: "RE" condition (See Methods: dF/F method characterization and validation).

(C) Comparison of baseline GCaMP6f signals on paired, consecutive trials aligned to cue. Columns: three different versions of the signal (Raw fluorescence, Normalized baseline dF/F method, Moving average dF/F method. Top row: ER condition; middle row: RE condition; bottom row: distortion index. Red distortion index plot shows only Normalized baseline method. Green distortion index plot shows overlay of Moving Average, Low-Pass Filter, and Multiple Baseline dF/F Methods because the difference in signal distortion between these methods was indistinguishable (See Methods: dF/F method characterization and validation).

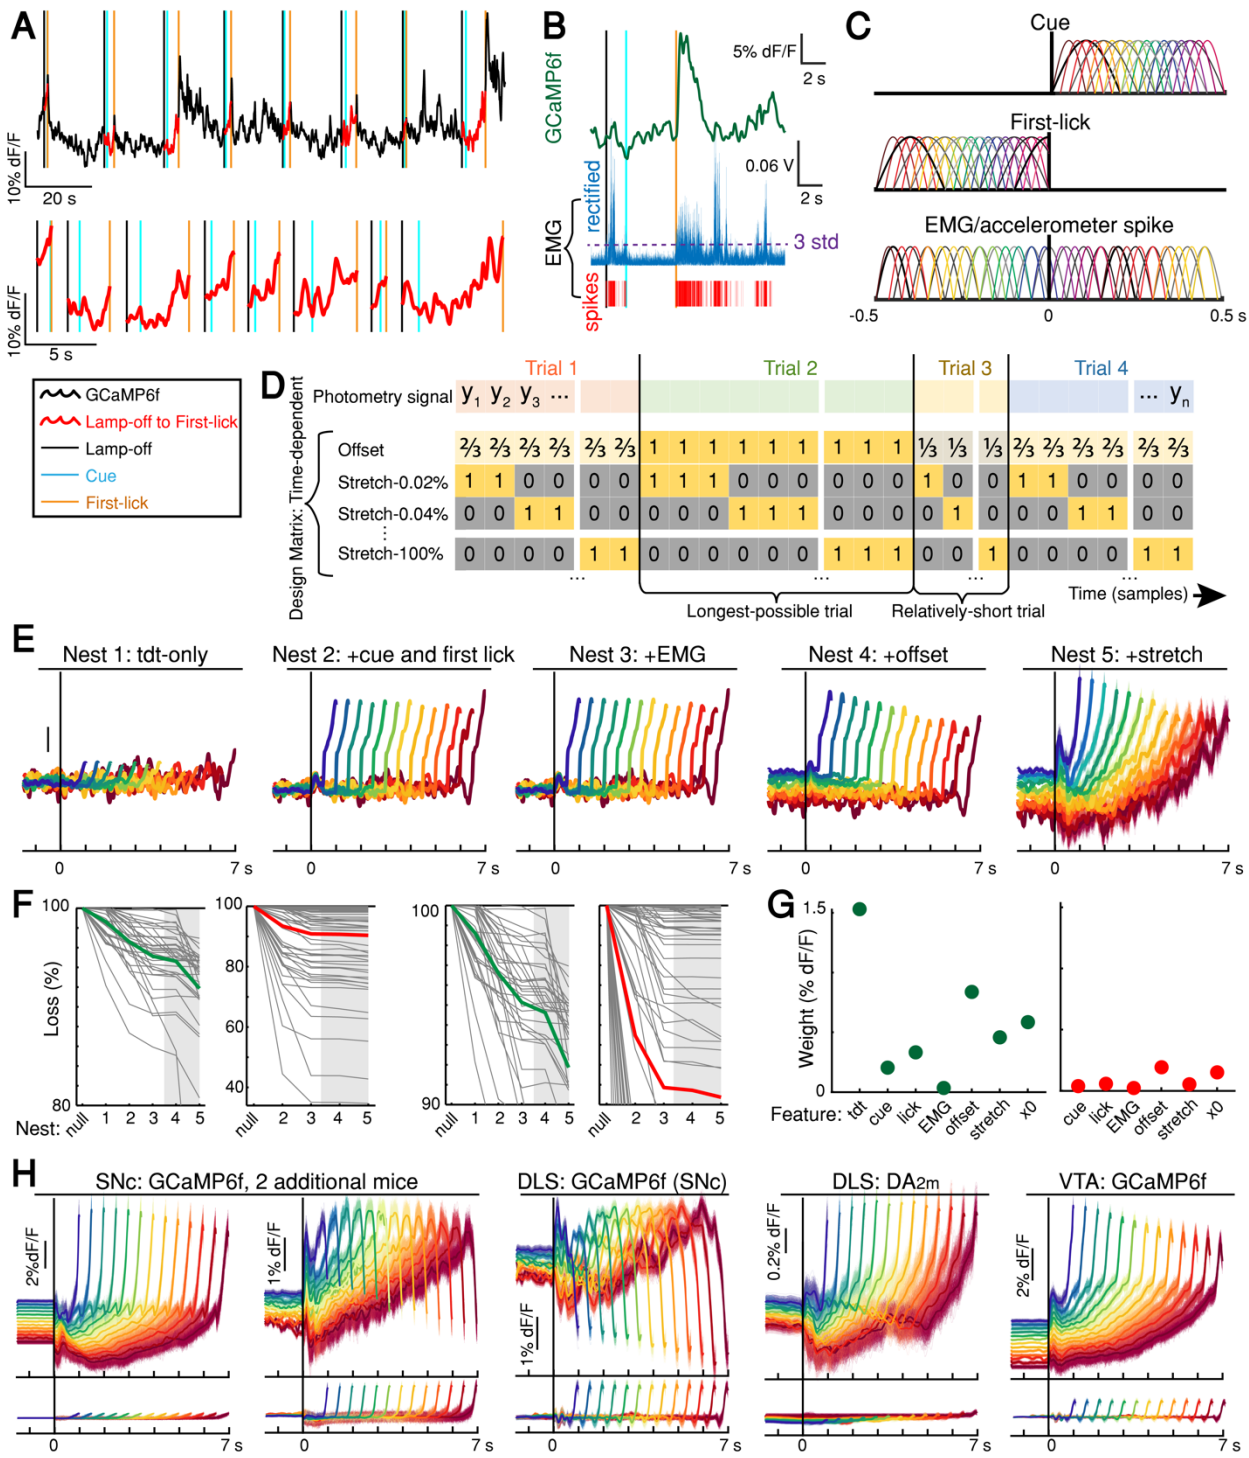

Figure S4. DAN signal encoding model parameterization and model selection, Related to Figure 5

(A) Schematic of photometry timeseries fit by encoding model. The lamp-off to first-lick interval was isolated from each trial in a session (top) and concatenated to produce the timeseries fit by the model (bottom).

(B) Derivation of EMG spikes from raw signals. Thresholding of rectified EMG at 3 standard deviations (std) during an example trial.
(C) Optimized basis kernels for cue, first-lick, and EMG/accelerometer spikes to produce timing-independent features.
(D) Schematic of d *x* n Design Matrix for timing-dependent features. Note: timing-independent features not shown for clarity in schematic.
(E) GCaMP6f model fits by nest for example session. Model error simulated 300x (shading).
(F) Model loss by nest. Green: mean loss for SNc GCaMP6f; red: mean loss for tdTomato (tdt); grey lines: individual sessions; grey shading: timing-dependent nests. Left: full-scale view of all datasets. Right: mean GCaMP6f and tdt loss compared on same scale.
(G) Summary of feature weights across SNc GCaMP6f (left) and tdt (right) models (68 sessions each). Coefficient weights were rectified, summed, and divided by the number of predictors per feature. Error bars: 2*standard error (too small to see). All features were significant in both GCaMP6f and tdt models.
(H) Top: examples of the full timing-dependent model (nest 5) from additional mice for all recording conditions. Bottom: tdt control channel fit. Model errors simulated 300x. Some mice show downward-going movement-related spikes at SNc cell bodies (second panel). All mice showed downward-going movement-related spikes from SNc terminals in DLS (middle panel).

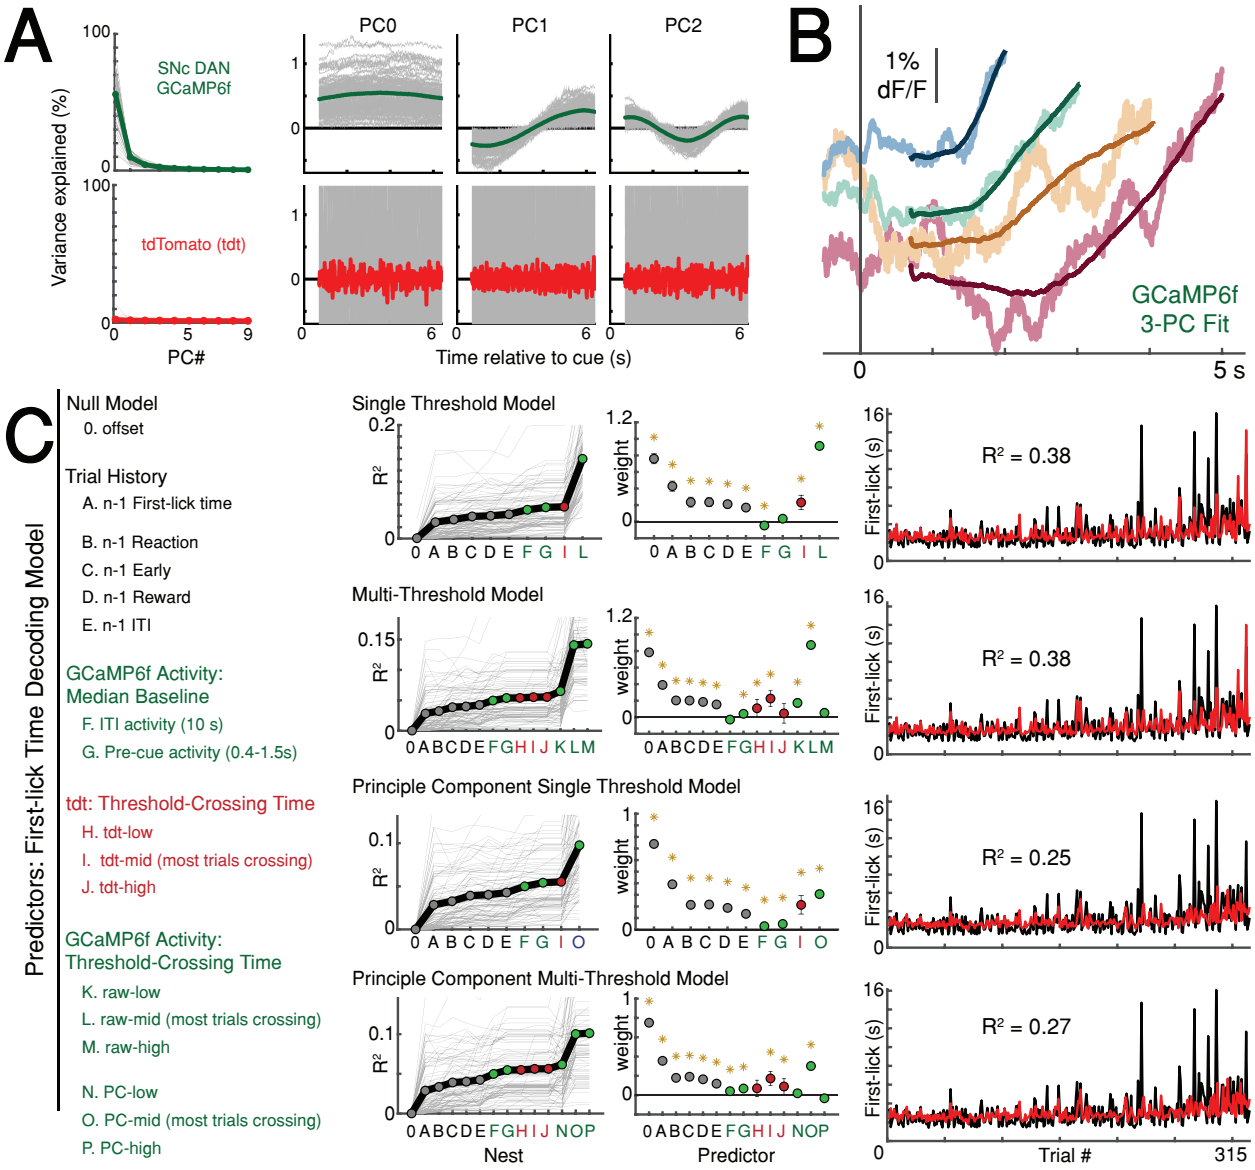

Figure S5. Variations of the first-lick time decoding model, Related to Figure 6

(A) Principle component analysis (PCA) of the ramping interval (0.7 s up to first-lick relative to cue). Green: mean GCaMP6f recorded at SNc; Red: mean tdTomato (tdt) recorded at SNc and VTA; Grey lines: single-session data. Left: Variance explained by first 10 principle components (PC). Right: Scores of first three principle components. X-axis shown for longest-possible interpolated trial duration; trials of shorter duration were interpolated to have the same number of samples for PCA.

(B) Example session data simulated with first 3 PCs. Light traces: actual averaged GCaMP6f activity truncated at first-lick onset; Dark traces: PC fits of the same trials.

(C) Decoding model variations. *: $p<0.05$, error bars: 95% confidence intervals. GCaMP6f threshold crossing time dominated every version of the model; n-1th trial first-lick time was consistently the second-best predictor.

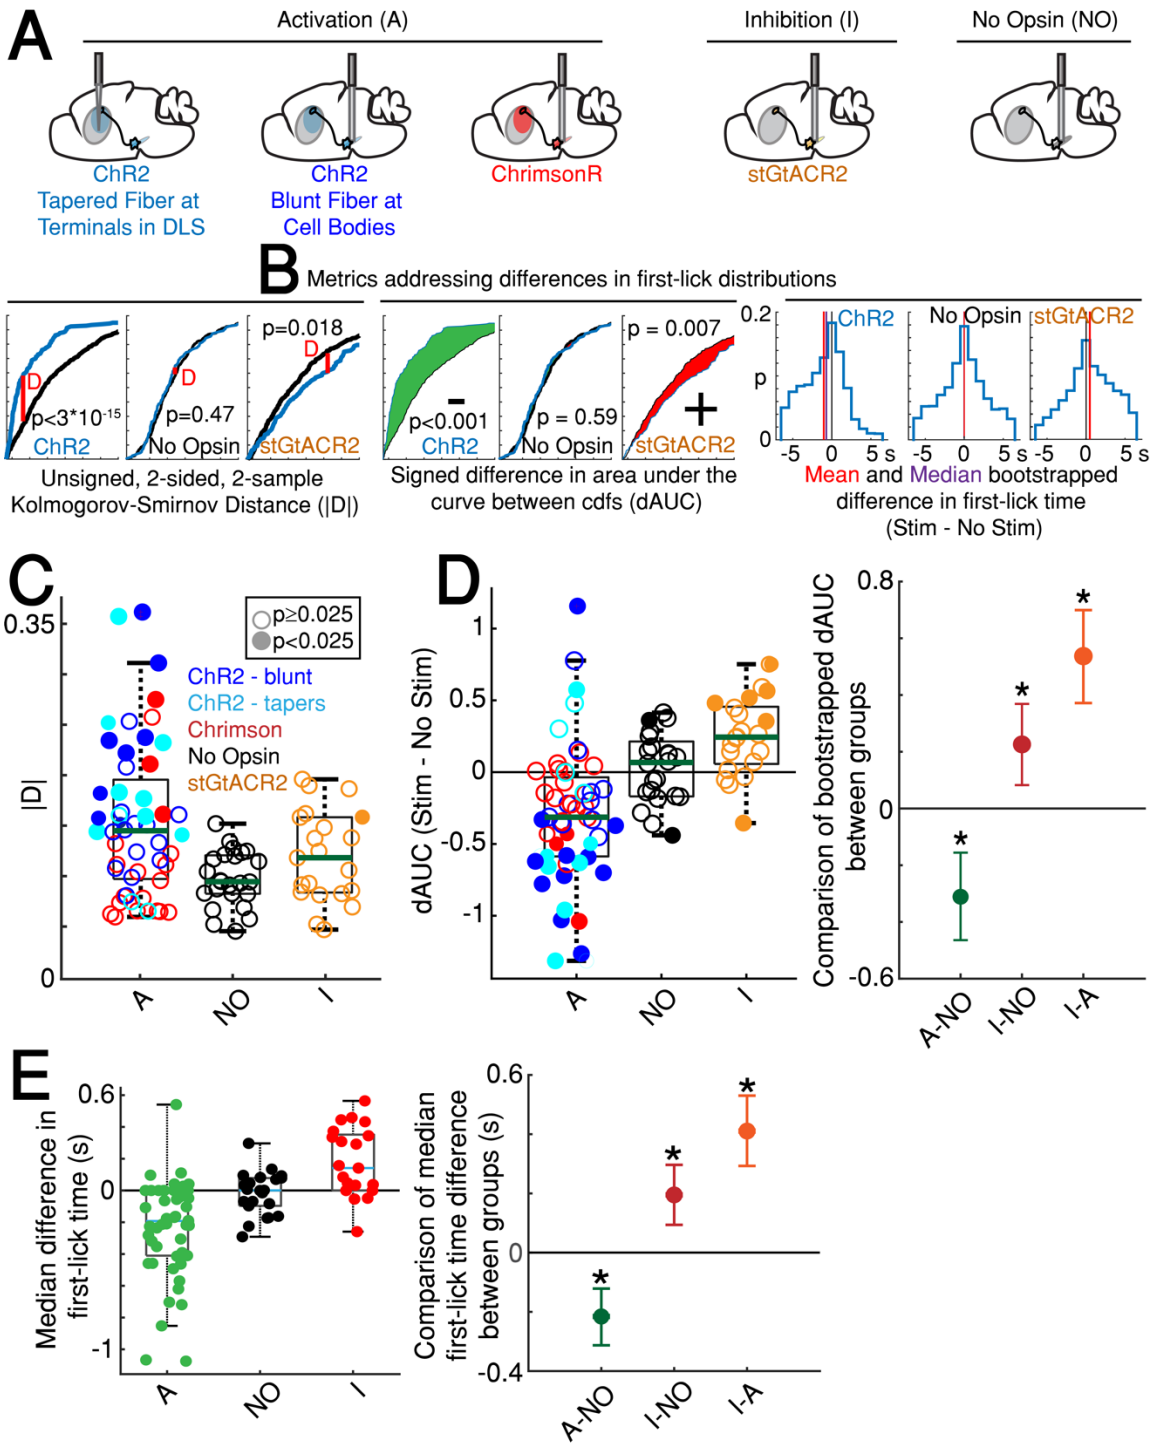

Figure S6. Variations on measurements of optogenetic effects, Related to Figure 7

(A) Strategy for optogenetic targeting of DANs.

(B) Comparison of four complementary metrics for addressing optogenetic effects. Left: unsigned Kolmogorov-Smirnov Distance (KS-D) assay of differences in first-lick time distribution. Center: signed, bootstrapped comparison of difference in area under the cdf curves (dAUC). Right: mean and median bootstrapped difference in first-lick time.

(C) KS-D Assay: all sessions. A: activation sessions; NO: no opsin sessions; I: inhibition sessions. Filled circles indicate significant difference between stimulated/unstimulated trials on single session ($p<0.025$, 2-sided, 2-sample KS test). Standard box plot, line: median, box: upper/lower quartiles; whiskers: 1.5x IQR.

(D) Left: bootstrapped dAUC Assay: all sessions, standard box plot as in **S6C**. Filled circles: significant difference on single session ($p<0.025$, 2-sided bootstrapped dAUC test, see Methods). Right: comparison of dAUC in first-lick distributions across all sessions between groups. Error bars denote bootstrapped 95% confidence interval (*: $p<0.05$).

(E) Left: median bootstrapped difference in first-lick time, stimulated-minus-unstimulated trials, standard box plot as in **S6C**. Dots: single sessions. Comparison of median difference in first-lick time across all sessions. Error bars denote bootstrapped 95% confidence interval (*: $p < 0.05$).

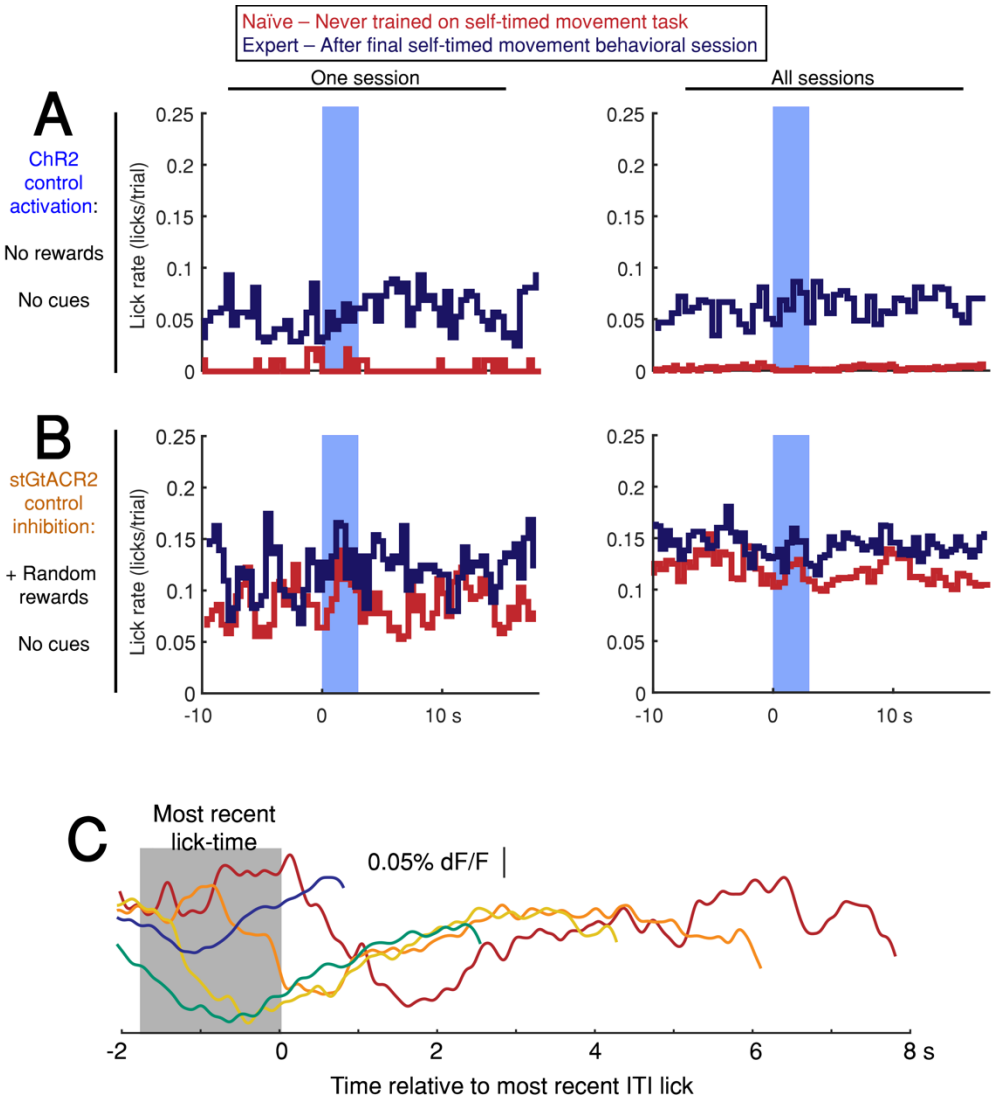

Figure S7. Optogenetic DAN stimulation does not cause or prevent licking, Related to Figure 7.

(A,B) Stimulation-aligned lick-rate during stimulation-control sessions. Animals expressing ChR2 or stGtACR2 were tested in 1-3 control sessions both before exposure to the self-timed movement task (red) and in 1-2 control sessions after the end of behavioral training (red). Blue bar indicates stimulation period (3 s starting at time 0 s). Left: one control session, Right: all sessions.

(A) Activation control sessions (no cues or rewards). Animals were head-fixed on the behavioral platform in the absence of any cues or rewards and were stimulated randomly at the same pace as the standard 3.3 s self-timed movement task. Activation did not elicit immediate licking in any single session.

(B) Inhibition-control sessions (no cues, + random rewards). Animals were head-fixed on the behavioral platform in the absence of cues while receiving juice rewards at random times. Inhibition did not prevent licking in any single session.

(C) Average ITI GCaMP6f signals aligned to most recent previous lick-time plotted up to onset of next spontaneous self-initiated lick during the ITI. (1 mouse, 5 sessions, truncated 150 ms before lick detection).

1054 **REFERENCES**

1055 Albin, R. L., Young, A. B., & Penney, J. B. The functional anatomy of basal ganglia disorders.
1056 *Trends in Neurosciences*, *12*(10), 366–375 (1989).

1057 Backman, C. M *et al*. Characterization of a Mouse Strain Expressing Cre Recombinase From the
1058 3' Untranslated Region of the Dopamine Transporter Locus. *Genesis* **44**, 383–390 (2006).

1059 Barthel, C., Nonnekes, J., Van Helvert, M., Haan, R., Janssen, A., Delval, A., … Ferraye, M. U.
1060 (2018). The laser shoes. *Neurology*, *90*(2), e164–e171.
1061 https://doi.org/10.1212/WNL.0000000000004795

1062 Bloxham, C. A., Mindel, T. A., & Frith, C. D. (1984). Initiation and execution of predictable and
1063 unpredictable movements in parkinson's disease. *Brain*, *107*(2), 371–384.
1064 https://doi.org/10.1093/brain/107.2.371

1065 Coddington, L. T. & Dudman, J. T. The timing of action determines reward prediction signals in
1066 identified midbrain dopamine neurons. *Nat. Neurosci*. **21**, 1563–1573 (2018).

1067 da Silva, J. A., Tecuapetla, F., Paixão, V. & Costa, R. M. Dopamine neuron activity before
1068 action initiation gates and invigorates future movements. *Nature* **554**, 244–248 (2018).

1069 Deecke, L. (1996). Planning, preparation, execution, and imagery of volitional action. *Cognitive*
1070 *Brain Research*, *3*(2), 59–64. https://doi.org/10.1016/0926-6410(95)00046-1

1071 DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in*
1072 *Neurosciences*, *13*(7), 281–285. https://doi.org/10.1016/0166-2236(90)90110-V

1073 Dews, P. B. & Morse, W. H. Some observations on an operant in human subjects and its
1074 modification by dextro amphetamine. *J. Exp. Anal. Behav.* **1,** 359–364 (1958).

1075 Dodson, P. D. et al. Representation of spontaneous movement by dopaminergic neurons is cell-
1076 type selective and disrupted in parkinsonism. *Proc. Natl Acad. Sci. USA* **113**, E2180–E2188
1077 (2016).

1078 Engelhard, B et al. Specialized coding of sensory, motor and cognitive variables in VTA
1079 dopamine neurons. *Nature* **570**, 509–513 (2019).

1080 Freeze, B. S., Kravitz, A. V, Hammack, N., Berke, J. D. & Kreitzer, A. C. Control of Basal Ganglia
1081 Output by Direct and Indirect Pathway Projection Neurons. *J. Neurosci.* **33,** 18531–18539 (2013).

1082 Gallistel, C. R. & Gibbon, J. Time, Rate and Conditioning. *Psychol. Rev.* **107,** 289–344 (2000).

1083 Gelman, A., & Hill, J. *Data Analysis Using Regression and Multilevel/Hierarchical Models.*
1084 Cambridge University Press, Cambridge, 43-45 (2006).

1085 Grillner, S. & Robertson, B. The Basal Ganglia Over 500 Million Years. *Curr. Biol.* R1088–R1100
1086 (2016).

1087 Hamilos, A.E. & Assad, J.A. Application of a unifying reward-prediction error (RPE)-based
1088 framework to explain underlying dynamic dopaminergic activity in timing tasks. Preprint as
1089 supplemental discussion section of
1090 https://www.biorxiv.org/content/10.1101/2020.05.13.094904v2 (2020).

Hallett, M. Volitional control of movement: The physiology of free will. *Clinical Neurophysiology*, *118*(6), 1179–1192 (2007).

Hamid, A. A. et al. Mesolimbic dopamine signals the value of work. *Nat. Neurosci.* **19**, 117–126 (2016).

Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. M. & Graybiel, A. M. Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature* **500**, 575–579 (2013).

Howe, M. W. & Dombeck, D. A. Rapid signalling in distinct dopaminergic axons during locomotion and reward. *Nature* **535**, 505–510 (2016).

Hughes, L. E., Altena, E., Barker, R. A. & Rowe, J. B. (2013). Perseveration and choice in parkinson's disease: The impact of progressive frontostriatal dysfunction on action decisions. Cereb. Cortex *23***,** 1572–1581

Kim, H. R. et al. A unified framework for dopamine signals across timescales. Preprint at https://www.biorxiv.org/content/10.1101/803437v1 (2019).

Lee, I. H. & Assad, J. A. Putaminal Activity for Simple Reactions or Self-Timed Movements. *J. Neurophysiol.* **89,** 2528–2537 (2003).

Lustig, C. & Meck, W. H. Chronic treatment with haloperidol induces deficits in working memory and feedback effects of interval timing. *Brain Cogn.* **58,** 9–16 (2005).

Maimon, G. & Assad, J. A. A cognitive signal for the proactive timing of action in macaque LIP. *Nat. Neuro.* ***9***, 948–955 (2006).

Malapani, C., Rakitin, B. C., Levy, R., Meck, W. H., Deweer, B., Dubois, B., & Gibbon, J. (1998). Coupled Temporal Memories in Parkinson's Disease: A Dopamine-Related Dysfunction. J. Cog. Neuro., *3*(May), 316–331.

Meck, W. H. Affinity for the dopamine D2 receptor predicts neuroleptic potency in decreasing the speed of an internal clock. *Pharmacol. Biochem. Behav.* **25,** 1185–1189 (1986).

Meck, W. H. Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Res.* **1109,** 93–107 (2006).

Mello, G. B. M., Soares, S. & Paton, J. J. A scalable population code for time in the striatum. *Curr. Biol.* **25,** 1113–1122 (2015).

Merchant, H., Harrington, D. L. & Meck, W. H. Neural Basis of the Perception and Estimation of Time. *Annu. Rev. Neurosci* 36, 313–36 (2013).

Mikhael, J. G. & Gershman, S. J. Adapting the flow of time with dopamine. *J. Neurophysiol.*, **121**, 1748–1760 (2019).

Mikhael, J. G., Kim, H. R., Uchida, N., & Gershman, S. J. Ramping and State Uncertainty in the Dopamine Signal. Preprint at https://www.biorxiv.org/content/10.1101/805366v1 (2019).

Mohebi, A et al. Dissociable dopamine dynamics for learning and motivation. *Nature* **570**, 65–70 (2019).

Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat. Neuro.* **17**, 1395–1403 (2014).

Rakitin, B. C., Penney, T. B., Gibbon, J., Hinton, S. C., & Meek, W. H. Scalar Expectancy Theory and Peak-Interval Timing in Humans. *Journal of Experimental Psychology: Animal Behavior Processes* **24**, 15–33 (1998).

Romo, R., Scarnati, E., Schultz, W. & Schultz, W. Role of primate basal ganglia and frontal cortex in the internal generation of movements II. Movement-related activity in the anterior striatum. *Exp. Brain. Res.* **91,** 385–395 (1992).

Runyan, C. A., Piasini, E., Panzeri, S., & Harvey, C. D. Distinct timescales of population coding across cortex. *Nature* **548**, 92–96 (2017).

Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).

Schuster, C. & Zimmerman, J. Timing behavior during prolonged treatment with dl-amphetamine. *J. Exp. Anal. Behav.* **4,** 327–330 (1961).

Shenoy, K. V., Sahani, M., & Churchland, M. M. Cortical Control of Arm Movements: A Dynamical Systems Perspective. Annual Review of Neuroscience, 36(1), 337–359 (2013).

1145 Sippy, T., Lapray, D., Crochet, S., & Petersen, C. C. H. Cell-Type-Specific Sensorimotor
1146 Processing in Striatal Projection Neurons during Goal-Directed Behavior. *Neuron* **88**, 298–305
1147 (2015).

1148 Soares, S., Atallah, B. & Paton, J. Midbrain dopamine neurons control judgment of time. *Science*
1149 **354,** 1273–1277 (2016).

1150 Sun, F. et al. New and improved GRAB fluorescent sensors for monitoring dopaminergic activity
1151 in vivo. Preprint at https://www.biorxiv.org/content/10.1101/2020.03.28.013722v2 (2020).