What's better for me? Fundamental role for lateral habenula in promoting subjective decision biases

Colin M. Stopper and **Stan B. Floresco**[*]
Department of Psychology and Brain Research Centre, University of British Columbia, Vancouver, British Columbia, Canada, 2136 West Mall, Vancouver, BC Canada V6T 1Z4, (604) 827-5313

Abstract

The lateral habenula (LHb) is believed to convey an aversive or "anti-reward" signal, but its contribution to reward-related action selection is unknown. We found that LHb inactivation abolished choice biases, making rats indifferent when choosing between rewards associated with different subjective costs and magnitudes, but not larger/smaller rewards of equal cost. Thus, instead of serving as an aversion center, the evolutionarily-conserved LHb acts as preference center integral for expressing subjective decision biases.

When choosing between rewards that differ in terms of their relative value, subjective impressions of which option may be "better" can be colored by certain costs (e.g., effort, delays, uncertainty) that diminish the subjective value of objectively larger rewards. Decisions of this kind are facilitated by different nodes within the mesocorticolimbic dopamine system[1]. Recent studies have highlighted the LHb as a critical nucleus within this circuitry that acts as a "brake" on dopamine activity, via disynaptic pathways through the rostromedial tegmental nucleus (RMTg)[2–4]. LHb neurons encode negative reward prediction errors opposite of dopamine neurons, exhibiting increased phasic firing in expectation of, or after, aversive events, (e.g., punishments, omission of expected rewards) and reduced firing after positive outcomes. LHb stimulation promotes conditioned avoidance and reduces reward-related responding, suggesting that this nucleus conveys an anti-reward/aversive signal[5–9]. Yet, LHb neurons also encode rewards of dissimilar magnitude, displaying phasic increases/decreases in firing in anticipation, or after receipt, of smaller/larger rewards[9]. This differential reward encoding may aid in biasing decisions towards/away from subjectively superior/inferior rewards. Yet, how LHb signals may influence decision biases and volitional choice behavior is unknown.

We investigated the contribution of the LHb to different forms of cost/benefit decision making mediated by dopamine circuitry[10,11]. We initially employed a probabilistic discounting task (a measure of risk-based decision making), requiring rats to choose

between a small/certain reward (1 pellet) and a large/risky reward. During daily training sessions, the probability of obtaining a larger, 4-pellet reward changed in a systematic manner over blocks of discrete-choice trials (100–12.5% or 12.5–100%). After ~25 training days, rats (n=16) displayed appropriate shifts in their decision biases, playing riskier during the higher probability blocks (100–50%), and safer when odds were poorer (25–12.5%). This was apparent after control treatments within the LHb (Fig. 1a, circles).

On a separate day, rats received infusions of $GABA_{A/B}$ agonists to inactivate the LHb. Given that i) phasic LHb firing encodes an anti-reward or disappointment signal after reward omissions [6,7], ii) LHb stimulation reduces responding for reward[8] and iii) LHb inactivation increases mesolimbic DA release[12], a parsimonious expectation would be that this manipulation should increase responding for the larger reward. In fact, we observed an effect much more profound. LHb inactivation completely abolished any discernible choice bias, inducing random patterns of responding that, when averaged across subjects, yielded a choice profile reflective of rats selecting both options with equal frequency, with choice behavior not differing from chance (50%) ($F_{3,45}$=6.69, P=0.001; Fig. 1a and Supplementary Fig. 1). This effect was apparent irrespective of whether reward probabilities decreased (n=9) or increased (n=7) over time (Fig. 1b,c). LHb inactivation also increased hesitation to make a choice (control latency=0.7 +/–0.1 sec, inactivation=1.3+/–0.2 sec; t_{15}=2.53, P=0.02) and the number of trials where no choices were made (control=0.9+/–0.5, inactivation=6.45+/–1.5; t_{15}=3.22, P=0.006). Moreover, this shift to indifference was apparent during periods where subjects showed a prominent bias for either the large/risky or small/certain option (Supplementary Fig. 2). We also observed an identical effect in a separate group (n=7) trained on a simpler task wherein the odds of obtaining a larger reward remained constant over trials (40%; Fig. 1d), indicating that promotion of choice biases by the LHb is not restricted to situations where reward probabilities are volatile, unlike other nodes within dopamine decision circuitry[13].

Importantly, the effects of LHb inactivation were neuroanatomically specific as infusion 1 mm dorsal (hippocampus, n=8) or ventral (thalamus, n=5) to the LHb or near the ventricle (n=8) had no effect on decision making compared to control conditions (Fig. 1e and Supplementary Fig. 3). Thus, disruption of decision biases induced by inactivation treatments was attributable to suppression of neural activity circumscribed to the LHb, but not adjacent regions.

In addition to sending direct projections to midbrain DA neurons that promote aversive behaviors[5], the LHb projects to the RMTg (which in turn regulates DA neuron activity), and also to dorsal raphe serotoninergic neurons[3]. To clarify which of these projection targets may interact with the LHb to promote probabilistic choice biases, separate groups of rats were trained on the fixed probabilistic choice task. Inactivation of the RMTg induced a choice profile resembling indifference, similar to LHb inactivation (t_5=3.17, P=0.02; Fig. 1f and Supplementary Fig. 4). In contrast, dorsal raphe inactivation had no effect on choice (t_3=0.01, P=1.0.; Fig. 1g). Thus, modification of probabilistic choice biases by the LHb is mediated primarily via projections to the RMTg that in turn controls DA neural activity, but not via serotonergic pathways.

We next investigated whether the LHb is specifically involved in cost/benefit judgments entailing reward uncertainty or if it plays a broader role in promoting biases during other decisions about rewards of different subjective value. To this end, we used a delay-discounting task, requiring rats (n=6) to choose either a small reward delivered immediately or a larger, delayed reward. Here, the subject is always guaranteed the larger reward, yet, delaying reward delivery after choice (0–45 s) diminishes its subjective value and shifts bias towards the small/immediate reward (Fig. 1h, circles). In parallel to probabilistic discounting, LHb inactivation abolished delay discounting ($F_{3,15}$=3.99, P=0.03; Fig. 1h), as choice shifted to a point of indifference (t_5 vs. 50%=0.37, P=0.72). Therefore, the LHb appears to play a fundamental role in promoting biases in situations requiring choice between rewards that differ in their subjective value.

A separate group (n=5) were trained on a reward magnitude discrimination, choosing between 4 vs. 1 reward pellet, both delivered with 100% certainty. LHb inactivation did not alter preference for the larger reward which, in this instance, clearly had greater objective value ($F_{1,4}$=2.98, P=0.16; Fig. 1i). Choice latencies (t_4=0.35, P=0.74) and trial omissions (t_4=0.76, P=0.49) were also unaffected (Supplementary Table 1). Thus, it is unlikely that the profound disruptions induced by LHb inactivation on cost/benefit decision making can be attributed to motivational or discrimination deficits. Instead, the LHb contributes selectively to evaluation of rewards that differ in terms of their relative costs and subjective values, but not to simpler preferences for larger vs smaller rewards of equal cost, similar to other nodes within DA decision circuitry[11].

The fact that LHb inactivation reduced preference for the larger reward during the 100%/0 sec delay blocks of the discounting tasks but not on the reward magnitude discrimination likely reflects differences in the relative value representation of the larger vs. smaller rewards that emerge after experience with these two types of tasks. This notion was supported by an analysis of forced-choice response latencies on the larger and smaller reward levers (Supplementary Fig. 5). Rats tested on the magnitude task showed greater response latencies when forced to select the smaller vs. larger reward after both treatments. In contrast, for the discounting tasks, this difference during the forced-choice trials in the 100% or 0 sec blocks was significantly muted or non-existent. Thus, rats trained on the magnitude discrimination viewed the 1-pellet option as substantially inferior to the 4-pellet option, whereas for those in the discounting tasks, this discrepancy was not as apparent, similar to previous findings[11]. This may explain why preference for the larger reward during the no-cost blocks of the discounting task may have been more susceptible to disruption following LHb inactivation.

It could be argued that the lack of effect of LHb inactivation on reward magnitude discrimination was due to rats responding in a habitual manner on this task relative to the discounting tasks, which were clearly goal-directed. To address this, we conducted a subsequent behavioral experiment to show that choice behavior under these conditions was sensitive to reinforcer or response devaluation. A separate group of rats well-trained on the reward magnitude discrimination was given a reinforcer devaluation test (see Methods) which caused a marked reduction in preference for the larger reward (Supplementary Fig. 6a–e). An additional response devaluation test was conducted two days later, during which

responses on the large reward lever did not provide reward. This manipulation also decreased preference for the lever formerly associated with the larger reward (Supplementary Fig. 6f). Thus, because choice on the reward magnitude task was altered following devaluation of either the reinforcer or the response contingency, this suggests that animals maintained a representation of the relative value of the two options and were responding in a goal-directed, as opposed to a habitual manner. As such, the lack of effect of LHb inactivation on this task, combined with the profound disruption in decision making on the discounting tasks, renders it unlikely that these differential effects are attributable to differences in the contribution of the LHb to goal-directed vs. habitual behavior. Instead, these data add further support to the notion that the LHb plays a selective role in promoting choice biases in situations when larger rewards are tainted by some form of cost, but not in expressing a more general preference for larger/smaller rewards.

These findings reveal a previously uncharacterized role for the LHb in reward-related processing, in that it is critical for promoting choice biases during evaluation of the subjective costs and relative benefits associated with different actions. Disruption of LHb signal outflow rendered animals unable to display any sort of preference towards larger, costly rewards or smaller, cheaper ones. Rather, they behaved as if they had no idea which option might be better for them, defaulting to an inherently unbiased and random pattern of choice, but only when the relative value of the larger reward was tainted by some sort of cost (uncertainty or delays).

LHb stimulation induces avoidance behaviors and suppresses reward-related responding, and phasic increases in LHb neural firing encode aversive or disappointing events[6,7]. As such, an emerging consensus is that the LHb conveys some form of aversive or "anti-reward" signal[5,8]. Our findings call for a refinement of this view. Indeed, suppression of LHb activity did not enhance responding for larger rewards, but instead, disrupted expression of a subjective preference for rewards of different value. In this regard, it is important to note that LHb neurons encode both aversive *and* rewarding situations via dynamic and opposing changes in activity. Thus, while phasic increases in firing encode aversive/non-rewarded expectations or events, or smaller rewards, these fast-firing LHb neurons also show reduced activity in response to rewarding stimuli[6,9,14]. Our findings indicate that suppressing these differential signals, encoding expectation or occurrence of negative/positive events, renders a decision-maker incapable of determining which option may be "better". As such, it is apparent that the LHb does not merely serve as a disappointment or "anti-reward" center, but, more properly, this nucleus may be viewed as a "preference" center, whereby integration of differential LHb reward/aversion signals sets a tone that is crucial for expression of preferences for one course of action over another. Expression of these subjective preferences is likely achieved through subsequent integration of these dynamic signals by regions downstream of LHb, including the RMTg and midbrain dopamine neurons[15,16]. Indeed, the LHb exerts robust control over the firing of dopamine neurons[17], and like the LHb, mesolimbic dopamine circuitry plays a preferential role in biasing choice towards larger, costly rewards but not larger/smaller rewards of equal cost[11,18]. Collectively, these findings suggest that the LHb, working in collaboration with other nodes of dopamine decision circuitry, plays a fundamental role in helping an organism make up its mind when faced with ambiguous decisions regarding the cost and benefits of different actions. Activity

within this evolutionarily-conserved nucleus aids in biasing behavior from a point of indifference toward committing to choices that may yield outcomes perceived as more beneficial. Further exploration of how the LHb facilitates these functions may provide insight to the pathophysiology underlying psychiatric disorders associated with aberrant reward processing and LHb dysfunction, such as depression[19,20].

Methods

Experimental subjects and apparatus

Experimentally-naïve, male Long Evans rats (Charles River Laboratories, Montreal, Canada) weighing 250–300 g (60–70 days old) at the start of the experiment were single-housed and given access to food and water *ad libitum*. The colony was maintained on a 12 h light/dark cycle, with lights turned on at 7:00 AM. Rats were food restricted to no more than 85–90 % of free-feeding weight beginning 1 week before training. Feeding occurred in the rats' home cages at the end of the experimental day and body weights were monitored daily. Animals were trained and tested between 9:00 AM and 5:00 PM. Individual rats were trained and tested at a consistent time each day. All testing was in accordance with the Canadian Council on Animal Care. Testing occurred in operant chambers (Med Associates, St Albans, VT, USA) were fitted with two retractable levers on either side of central food receptacles where reinforcement (45 mg pellets; Bioserv, Frenchtown, NJ, USA) was delivered by a dispenser, as described previously[13]. No statistical methods were used to predetermine sample sizes.

Behavioral tasks

Rats were initially trained to press retractable levers within 10 s of their insertion into the chamber over a period of 5–7 days[10,11], after which they were trained on one of four decision making tasks.

Probabilistic discounting—Risk-based decision making was assessed with a probabilistic discounting task described previously[10,11]. Rats received daily training sessions 6–7 days/week, consisting of 72 trials, separated into 4 blocks of 18 trials. Each 48-min session began in darkness with both levers retracted (the intertrial state). Trials began every 40 s with houselight illumination and, 3 s later, insertion of one or both levers. One lever was designated the large/risky lever, the other the small/certain lever, which remained consistent throughout training (counterbalanced left/right). No response within 10 s of lever insertion reset the chamber to the intertrial state until the next trial (omission). Any choice retracted both levers. Choice of the small/certain lever always delivered one pellet with 100% probability; choice of the large/risky lever delivered 4 pellets but with a probability that changed across the four trial blocks. Blocks were comprised of 8 forced-choice trials (4 trials for each lever, randomized in pairs), followed by 10 free-choice trials, where both levers were presented. The probability of obtaining 4 pellets after selecting the large/risky option varied across blocks. Separate group of rats were trained on variants where reward probabilities systematically decreased (100%, 50%, 25%, 12.5%) or increased (12.5%, 25%, 50%, 100%) across blocks. For each forced and free-choice trial within a particular block, the probability of receiving the large reward was drawn from a random number generating function (Med-PC) with a set probability distribution (i.e.; 100, 50, 25 or 12.5%). Therefore,

on any given session, the probabilities in each block may have varied, but on average across training days, the actual probability experienced by the rat approximated the set value within a block. Latencies to choose were also recorded. Rats were trained until as a group, they (1) chose the large/risky lever during the 100% probability block on at ~90% of trials, and (2) demonstrated stable baseline levels of choice, assessed using an ANOVA analysis described previously[10,11]. Data from three consecutive sessions were analyzed with a two-way repeated-measures ANOVA with Day and Trial Block as factors. If there was no main effect of Day or Day X Trial Bock interaction (at $p>0.1$ level), performance of the group was deemed stable.

Probabilistic choice with fixed reward probabilities—Training on this task was very similar to the probabilistic discounting task, except that the probability of obtaining the larger 4-pellet reward was set at 40%, and remained constant over one block of 20 free-choice trials that were preceded by 20 forced-choice trials. Data from rats that displayed a preference for the large/risky reward were used in the analysis.

Delay discounting—This task shared similarities to the probabilistic discounting tasks in a number of respects, but with some key differences. Daily sessions consisted of 48 trials, separated into 4 blocks of 12 trials (2 forced- followed by 10 free-choice trials/block; 56 min session). Trials began every 70 s with houselight illumination and insertion of one or both levers. One lever was designated the small/immediate lever, that, when pressed, always delivered one pellet immediately. Selection of the other, large/delayed lever delivered four pellets after a delay that increased systematically over the four blocks: it was initially 0 s, then 15, 30 and 45 s. No explicit cues were presented during the delay period; the houselight was extinguished, and then re-illuminated upon reward delivery.

Reward magnitude discrimination—This task was used to confirm if the reduced preference for larger, costly rewards was due to a general reduction in preference for larger rewards or some other form of non-specific motivation or discrimination deficits. Rats were trained and tested on a task consisting of 48 trials divided into 4 blocks, each consisting of 2 forced- and 10 free-choice trials. As with the discounting tasks, choices were between a large 4-pellet and smaller, 1-pellet reward, both of which were delivered immediately with 100% certainty after a choice.

Devaluation tests—A separate behavioral experiment was conducted in intact animals to assess whether performance during the reward magnitude discrimination was under habitual or goal-directed control. A separate group of rats was trained for 9 days on a reward magnitude discrimination in an identical manner to those that received LHb inactivation. On day 10 of training, rats received a reinforcer devaluation test. One hour prior to the test session, rats received *ad libitum* access to the sweetened reward pellets in their home cages. If responding on this task had become habitual, the prediction would be that reinforcer devaluation by pre-feeding should not influence performance during the test. Conversely, if choice was goal-directed, the bias toward the large reward should be diminished during this test.

Following the sucrose devaluation test, rats were retrained for two additional days on the task under standard food restriction, after which they again were selecting the large reward on nearly every free-choice trial. On the following day, rats received a response devaluation test during which responding on the large reward lever no longer delivered reward (although selecting the other lever still yielded 1 reward pellet).

Surgery and microinfusion protocol

Rats were trained on the discounting tasks until they displayed stable levels of choice (20–25 days) after which they were fed *ad libitum* for 1–3 days and subjected to surgery. Those trained on the other tasks were implanted prior to training. Rats were anaesthetized with 100 mg/kg ketamine and 7 mg/kg xylazine and implanted with bilateral 23 gauge stainless-steel cannulae aimed at the LHb (flat skull: anteroposterior = −3.8 mm; mediolateral = +/−0.8 mm; dorsoventral = −4.5 mm from dura). Separate anatomical control groups were implanted with cannulae at sites, either 0.5–1.0 mm dorsal or 1 mm ventral to the LHb site. Separate groups of rats to be trained on the fixed probabilistic choice task were implanted with bilateral cannulae in the RMTg (ARN, flat skull at 10° laterally: anteroposterior = −6.8 mm; mediolateral = +/−0.7 mm; dorsoventral = −7.4 mm) or a unilateral cannula in the DR (flat skull with cannula at 20° laterally: anteroposterior = −7.6mm; mediolateral = 0.0; dorsoventral = −5.2 mm). Cannulae were held in place with stainless steel screws and dental acrylic and plugged with obdurators that remained in place until the infusions were made. Rats were given ~7 days to recover from surgery before testing, during which they were again food restricted.

Training was re-initiated on the respective task for at least 5 days until the group displayed stable levels of choice behavior for 3 consecutive days. One to two days before the first microinfusion test day, obdurators were removed, and a mock infusion procedure was conducted. The day after displaying stable discounting, the group received its first microinfusion test day. A within-subjects design was used for all experiments. Reversible inactivation of the LHb was achieved by infusion of a combination of GABA agonists baclofen and muscimol using procedures described previously[11] (50 ng each in 0.2 μl, delivered over 45 s). Injection cannulae were left in place for 1 min for diffusion. Rats remained in their cage for an additional 10 min period before behavioral testing.

On the first infusion test day, half of the rats in each group received control treatments (saline); the remaining received baclofen/muscimol. The next day, rats received a baseline training day (no infusion). If, for any individual rat, choice of the large/risky lever deviated by >15% from its preinfusion baseline it received an additional day of training before the next test. On the following day, second counterbalanced infusion was given.

Histology

Rats were euthanized, brains were removed and fixed in 4% formalin for 24 hrs, frozen, sliced in 50 μm sections and stained with Cresyl Violet. Placements were verified with reference to Paxinos and Watson[21]. Based on previous autoradiographical, metabolic, neurophysiological and behavioral measures[22–25], the effective functional spread of inactivation induced by 0.2 μl infusions of 50 ng of GABA agonists would be expected to be

between 0.5 and 1 mm in radius from the center of the infusion. Placements were deemed to be within the LHb only if the majority of the gliosis from the infusions resided within the clearly defined anatomical boundaries of this nucleus. Alternatively, rats whose placements resided outside this region, either because of direct targeting or from missed placements were allocated to separate dorsal (hippocampus), medial (third ventricle) or ventral (thalamic) neuroanatomical control groups, which resided beyond the estimated effective functional spread of our inactivation treatments. Data from these groups were analyzed separately.

Data analysis

The primary dependent measure of interest was the proportion of choices directed towards the large reward lever (ie: large/risky or large/delayed) for each block of free-choice trials, factoring in trial omissions. For each block, this was calculated by dividing the number of choices of the large reward lever by the total number of successful trials. For the probabilistic discounting experiment, choice data were analyzed using three-way, between/within-subjects ANOVAs, with treatment and probability block as two within-subjects factors and task variant (ie: reward probabilities decreasing or increasing over blocks) as a between subjects factor. Thus, in this analysis, the proportion of choices of the large/risky option across the four levels of trial block was analyzed irrespective of the order in which they were presented. For the delay discounting and reward magnitude experiment, choice data were analyzed with a two-way repeated measures ANOVA, with treatment and trial block as factors. Choice data from fixed probability experiments were analyzed with paired-sample two-tailed t-tests Response latencies (the time elapsed between lever insertion and subsequent choice) and the number of trial omissions (i.e., trials where rats did not respond within 10 s) were likewise analyzed with paired-sample two-tailed t-tests. Data distribution was assumed to be normal, but this was not formally tested. The use of automated operant procedures eliminated the need for experimenters to be blind to treatment.

Additional analyses were conducted on the latencies to make a response during forced choice trials of the different tasks to explore why LHb inactivation affected choice during the no-cost blocks of the discounting task but not the reward magnitude task. The rationale was that animals trained on the reward magnitude discrimination learn that the relative value of the larger reward is *always higher* than the smaller reward, while those trained on discounting tasks consistently experienced *changes in relative value* of the large reward option over a session and learn that the large reward lever is not always the best option available. To provide support for this hypothesis, we analyzed response latencies to select the large and small reward during all of the forced-choice trials for rats trained on the reward magnitude discrimination, and compared them to large and small reward forced-choice latencies displayed by rats performing the discounting tasks during the 100%/0-sec delay (i.e.; no-cost) blocks. If well-trained animals perceived the larger reward as considerably "better" than the smaller one, they should display faster response latencies when forced to choose the larger vs smaller reward. On the other hand, if the relative value of the two rewards is perceived as more comparable (even in the 100% or 0 sec delay blocks), the difference in response latencies when forced to select one option or the other should be diminished.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

1. Floresco SB, St Onge JR, Ghods-Sharifi S, Floresco SB. Cogn Affect Behav Neurosci. 2008; 8:375–389. [PubMed: 19033236]

2. Barrot M, et al. J Neurosci. 2012; 32:14094–14101. [PubMed: 23055478]

3. Hikosaka O, Sesack SR, Lecourtier L, Shepard PD. J Neurosci. 2008; 28:11825–11829. [PubMed: 19005047]

4. Jhou TC, Geisler S, Marinelli M, Degarmo BA, Zahm DS. J Comp Neurol. 2009; 513:566–596. [PubMed: 19235216]

5. Lammel S, et al. Nature. 2012; 491:212–217. [PubMed: 23064228]

6. Matsumoto M, Hikosaka O. Nature. 2007; 447:1111–1115. [PubMed: 17522629]

7. Matsumoto M, Hikosaka O. Nat Neurosci. 2009; 12:77–84. [PubMed: 19043410]

8. Stamatakis AM, Stuber GD. Nat Neurosci. 2012; 15:1105–1107. [PubMed: 22729176]

9. Bromberg-Martin ES, Hikosaka O. Nat Neurosci. 2011; 14:1209–1216. [PubMed: 21857659]

10. StOnge JR, Floresco SB. Neuropsychopharmacology. 2009; 34:681–697. [PubMed: 18668030]

11. StOnge JR, Stopper CM, Zahm DS, Floresco SB. J Neurosci. 2012; 32:2886–2899. [PubMed: 22357871]

12. Lecourtier L, Defrancesco A, Moghaddam B. Eur J Neurosci. 2008; 27:1755–1762. [PubMed: 18380670]

13. StOnge JR, Floresco SB. Cereb Cortex. 2010; 20:1816–1828. [PubMed: 19892787]

14. Bromberg-Martin ES, Matsumoto S, Hong S, Hikosak O. J Neurophysiol. 2010; 104:1068–1076. [PubMed: 20538770]

15. Bromberg-Martin ES, Matsumoto M, Hikosaka O. Neuron. 2010; 68:815–834. [PubMed: 21144997]

16. Jhou TC, Fields HL, Baxter MG, Saper CB, Holland PC. Neuron. 2009; 61:786–800. [PubMed: 19285474]

17. Ji H, Shepard PD. J Neurosci. 2007; 27:6923–6930. [PubMed: 17596440]

18. Salamone JD, Cousins MS, Bucher S. Behav Brain Res. 1994; 65:221–229. [PubMed: 7718155]

19. Meng H, et al. Brain Res. 2011; 1422:32–38. [PubMed: 21978548]

20. Sartorius A, Henn FA. Med Hypotheses. 2007; 69:1305–1308. Additional Methods References. [PubMed: 17498883]

21. Paxinos, G., Watson, C. The rat brain in stereotaxic coordinates. 4. San Diego: Academic Press; 1998.

22. Allen TA, et al. J Neurosci Methods. 2008; 171:30–38. [PubMed: 18377997]

23. Martin JH, Ghez C. J Neurosci Methods. 1999; 86:145–159. [PubMed: 10065983]

24. Arikan R, et al. J Neurosci Methods. 2002; 51–57:118.

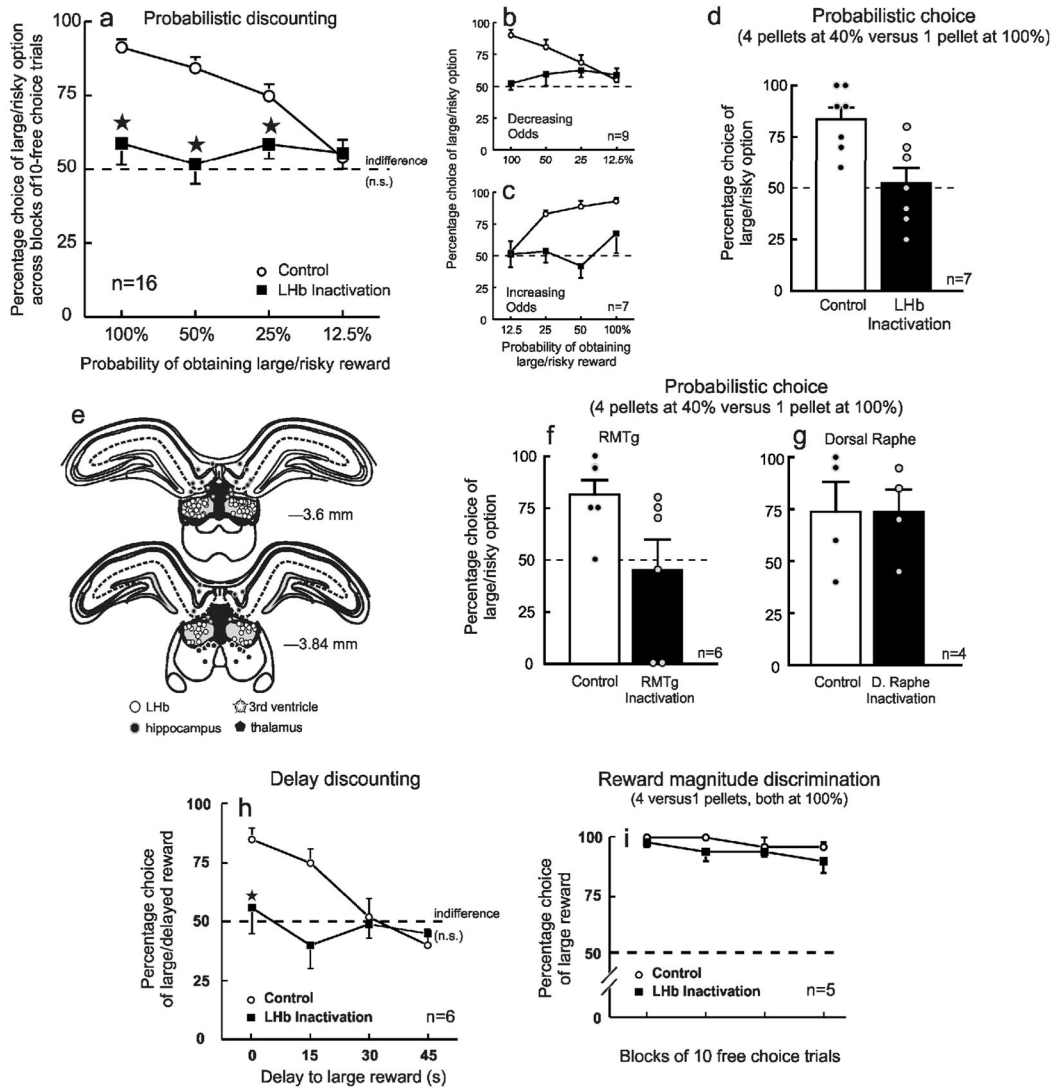25. Floresco SB, McLaughlin RJ, Haluk DM. Neuroscience. 2008; 154:877–884. [PubMed: 18479836]

Figure 1.

Inactivation of LHb circuitry abolishes choice biases during cost/benefit decision making. For all graphs, error bars represent S.E.M. (**a**) Percentage choice of the large/risky option across the 4 probability blocks for all rats trained on two variants of the probabilistic discounting task. LHb inactivation (*n*=16) abolished probabilistic discounting (treatment x block interaction, $F_{3,45} = 6.69$, $P = 0.008$), causing rats to randomly select both options with equal frequency (t_{15} vs 50%=1.44, P=0.17., dashed line). Choice after LHb inactivation did not vary across blocks ($F_{3,45}$=0.43, $P = 0.73$), resulting in a profile indicative of indifference. ★, $P<0.05$ vs control during a particular probability block. Data from subsets of rats trained on variants where reward probabilities decreased (*n*=9) or increased (*n*=7) over a session are presented in (**b**) and (**c**). LHb inactivation induced a comparable disruption in decision making in both groups (all effects of task variant, Fs<2.1, n.s.). (**d**) Data from a separate experiment where the probability of obtaining the large, uncertain reward remained constant (40%) across a session. Under these conditions, rats (*n*=7) chose the risky option on ~80% of trials following control treatments, but again, choice dropped to chance levels (50%) after

LHb inactivation ($F_{1,6}$=25.36, P = 0.002). (**e**) Location of infusions residing within the LHb for all experiments, and control placements in the adjacent hippocampus, ventricle or thalamus. Numbers correspond to mm from bregma. (**f**) RMTg inactivation (*n*=6) reduced preference for a large/risky reward to chance levels on a task where the probability of obtaining the large, uncertain reward remained constant (40%) across a session, whereas dorsal raphe inactivation (**g**) had no effect on choice (*n*=4). (**h**) Percentage choice of the large/delayed reward across the 4 blocks of the delay discounting task, wherein rats chose between a small, immediate reward and a larger, delayed reward (*n*=6). LHb inactivation abolished choice preference ($F_{3,15}$=3.99, p = 0.02), resulting in indifference (t_5=vs 50%=0.31, P = 0.76. dashed line). ★, P< 0.05 vs control during a particular block. (**i**) Rats trained on a reward magnitude discrimination task chose between a large and small reward, both delivered with 100% certainty (*n*=5). In contrast to the other experiments, LHb inactivation did not alter preference for larger, cost-free rewards.