Cost, Benefit, Tonic, Phasic: What do response rates tell us about dopamine and motivation?

Yael Niv[1,2*]

[1] Gatsby Computational Neuroscience Unit, UCL
17 Queen Square, WC1N 3AR, London, UK
[2] Interdisciplinary Center for Neural Computation
The Hebrew University of Jerusalem, Jerusalem 91904, Israel

Corresponding author and mailing address:

Yael Niv

Center for the Study of Brain, Mind and Behavior

3-N-12 Green Hall

Princeton University

Princeton, NJ 08544, USA

Email: yael@princeton.edu

Acknowledgments

---

[*]Currently at the Center for the Study of Brain, Mind and Behavior, Princeton University.

Abstract

The role of dopamine in decision making has received much attention from both the experimental and computational communities. However, because reinforcement learning models concentrate on discrete action selection and on phasic dopamine signals, they are silent as to how animals decide upon the rate of their actions, and they fail to account for the prominent effects of dopamine on response rates. We suggest an extension to reinforcement learning models, in which response rates are optimally determined by balancing the tradeoff between the cost of fast responding and the benefit of rapid reward acquisition. The resulting behavior conforms well with numerous characteristics of free-operant responding. More importantly, this framework highlights a role for a tonic signal corresponding to the net rate of rewards in determining the optimal rate of responding. We hypothesize that this critical quantity is conveyed by tonic levels of dopamine, explaining why dopaminergic manipulations exert a global affect on response rates. We further suggest that the effects of motivation on instrumental rates of responding is mediated through its influence on the net reward rate, implying a tight coupling between motivational states and tonic dopamine. The relationships between phasic and tonic dopamine signalling, and between directing and energizing effects of motivation, as well as the implications for motivational control of habitual and goal-directed instrumental action selection, are discussed.

Browsing through any random selection of experimental psychology papers will reveal that the dependent variable most commonly used to study animal behavior is response rate.[1] The effects of experimental manipulations as diverse as changes in the amount of reward that an animal can earn, alterations of the requirements or conditions under which rewards or punishments are delivered, lesions of neural structures, or the administration of drugs, are commonly discerned through changes in response rates. In terms of decision making and action selection, response rates are, in fact, inseparable from responding itself: accompanying any choice of which action to perform is a choice of how fast (or at what instantaneous rate) to perform this action. It may come as a surprise, then, that *normative* models of responding, such as reinforcement learning, which have done much to explain *why* it is appropriate for animals choose actions the way they do, have completely ignored the choice of response rates.

Response rates have played a more prominent role in descriptive models. These aim to quantify the relationships between experimental variables and response rates (eg. the Matching Law[2]), but not why, or in what sense these relationships are appropriate in different scenarios. In the absence of normative models (which deal exactly with these latter aspects) questions such as why does motivation influences response rates, and how should dopamine affect rate selection, are left unanswered. In previous work,[3–6] on which we focus this review, we proposed to remedy this by extending the framework of reinforcement learning to the optimal selection of response rates.

In our model,[3,6] animals choose with what latency (ie, how fast, or with what instantaneous rate) to perform actions by optimally balancing the costs of fast performance and the benefits of rapid reward acquisition. Focusing on this tradeoff, the model highlights the *net expected rate of rewards* as the important determinant

of the cost of delaying future rewards and the optimal rate of responding. We marshal evidence suggesting that this quantity is signaled by tonic levels of dopamine, and argue that this, in fact, explains why higher levels of dopamine are associated with faster performance, while low levels of dopamine induce lethargy. We further leverage the normative framework to argue that motivation and dopamine are, in fact, tightly linked in controlling response vigor, as the effect of motivation on response rates is mediated by a change in the expected net rate of rewards.

In the following we first detail the basic characteristics of response rates which we expect our model to reproduce. We then describe the new model in Section 2, emphasizing the tradeoffs that must be negotiated optimally in order to maximize reward intake. In particular, we focus on the role of the expected rate of reward in determining the opportunity cost of time and the optimal rate of responding. Section 3 relates this signal to tonic levels of dopamine, and discusses the implications for understanding the role of dopamine in action selection. In Section 4 we utilize this normative model of response rates to analyze the effects of motivation on responding. We first discuss how both the directing and energizing effects of motivation are manifest in the model. The results suggest a parcellation of motivational effects into outcome-dependent and outcome-independent, leading to a new understanding of the susceptibility of goal-directed behavior on the one hand, and habitual behavior on the other hand, to motivational manipulations. Finally, we argue that the outcome-independent energizing effects of motivation on response rates are mediated through changes in the expected net rate of rewards, implying a strong link between tonic dopamine and motivation. Last, in Section 5 we discuss some open questions such as the extension of the model to Pavlovian behavior, the relationship between phasic dopaminergic signals and motivation, and the neural locus of cost/benefit tradeoff computations.

1   What do we know about response rates?

Action selection has most frequently been studied in instrumental conditioning paradigms, on which we will focus here. In the commonly used *free operant* form of these,[7] animals (typically rats, mice or pigeons) perform an action (eg. pressing a lever, pecking a key) in order to obtain some coveted reinforcement (such as food for a hungry animal). Importantly, rather than performing actions at discrete, predefined time-points (as is typically modeled in reinforcement learning[8]), free operant responding is *self paced*, and animals are free to choose their rate of responding.

Numerous experiments have shown that the schedule of reinforcement (eg. ratio or interval), the nature or amount of the rewards used, and the motivational state of the animal profoundly affect the rate of instrumental responding. In general, responding is slower the longer the interval duration or ratio requirement,[9–12] and faster for higher magnitude rewards or more desirable rewards.[13,14] More refined characteristics of free operant behavior include the observation of higher response rates on ratio schedules compared to yoked

interval schedules[15–17] and response allocation that matches payoff rates when two interval schedules are concurrently available.[2,18,19]

The fact that response rates are affected by manipulations of the schedule of reinforcement suggests that animals choose with which rate to perform different actions as an adaptation to the specifics of the task they are solving. Furthermore, in most cases behavior in such schedules is well below ceiling rates, evidence that response latencies are not constrained by decision times, motor or perceptual requirements, but rather the particular response rate was selected as appropriate for the task at hand. In the following we will assume that the choice of response rate is the result of an optimization process that is influenced by two opposing goals: the desire to acquire rewards rapidly on the one hand, and to minimize effort costs on the other hand.

2   Optimal responding: Cost/benefit tradeoffs

Consider a situation in which a rat can choose between several actions: it can poke its nose into a (possibly empty) food well, it can press a lever that may cause food to fall into the food well, it can pull a chain that may cause water to pour into the food well and so forth (Figure 1a). The choice of which sequence of actions to take, and at what rate (or with what latency) to take each action, can be seen as an optimization problem, if we assume that the goal of the rat is to harvest rewards at as high a rate as possible, while incurring minimal effort costs. Because for free-operant tasks the problem can be defined computationally as a (semi-)Markov decision process, the optimal solution can be derived as a series of optimal decisions: the rat should first choose the currently optimal action and execute it with the optimal latency, and then, based on the consequences of this action (the resulting 'state' of the world, eg. whether the action resulted in food falling into the food well or not), choose the next optimal action and latency, and so forth. The optimal policy of which actions to choose in the different states[†], and with what latency to perform the chosen actions, can be found using reinforcement learning methods such as 'value iteration',[8] or online 'temporal difference' learning[8,20] (for a full computational exposition of the model equations and solution, see refs 3,5). In order to gain insight into the optimal policy, we will now analyze the factors which affect a single decision within the series of actions. In our model this consists of two parts: the rat must choose *which* action to perform, and *how fast* (or with what latency) to perform it. It turns out that these two sub-decisions depend on different characteristics of the task.

The choice of which action to perform depends on the utility of the rewards potentially available for each of the actions, the probability that the action will indeed be rewarded, and the effort cost of performing the action. For instance, if pressing the lever is rewarded with food with a probability of 20%, this would be

---

[†]The states we refer to here are states of the environment, such as whether there is food in the magazine, whether the lever is extended and available for pressing, etc. These should not be confused with the motivational state of the animal, which will be discussed in Section 4. For modeling simplicity, we assume that the animal's motivational state is constant during the experimental session.

preferable to an action that leads to the same outcome but with only 10% chance. What about a choice between actions that lead to different rewards? When comparing the worth of qualitatively different outcomes such as food and water, the motivational state of the animal must come into consideration, as it determines the utility of each outcome to the animal.[4] A hungry rat may prefer to press the lever for food, while a thirsty one might choose to pull the chain to obtain water. The choice of which action to perform also depends on how costly the action itself is, in terms of effort: for instance, if pulling the chain necessitates much effort to jump and reach it, the benefit of a small amount of water may not be worth this effort. To summarize, the optimal choice of which action to perform can be detemined by comparing the available actions in terms of the worth and probability of the potential reward for each action, and the effort cost of performing the action.

The optimal choice of how fast to perform the chosen action is determined by an altogether different cost/benefit tradeoff. First, we must assume that it is more costly for the rat to perform an action quickly rather than slowly (otherwise rats would always perform actions at the fastest possible rate, which is clearly not the case). Against what should this cost of fast performance be weighed, when deciding on the optimal speed of an action? Completing the chosen action slowly, for instance moving toward the lever and pressing it without haste, will of course delay the availability of the possible reward for this action. But, more important, all future actions and rewards will be delayed. So the effort cost of behaving quickly should be weighed against the cost of delaying *all* future rewards. Note that while the choice of which action to perform is affected by parameters local to the different actions and their (potentially long-term) outcomes, the choice of response rate influences the timing of all future actions, and is thus affected by global considerations.

How can the rat estimate the cost of delaying all future rewards? Average reward reinforcement learning techniques reveal a simple solution.[21,22] A specific policy of action choices and response latencies will lead to an average rate of rewards obtained per unit time, at an average effort cost per unit time.[23] The rate of rewards minus costs — the influx of net benefit per unit time, which we will refer to as the net reward rate — is exactly the worth of time under this policy,[24–26] or the *opportunity cost* of wasted time.[3,5] That is, in every second in which the current policy of responding will *not* be performed, on average this amount of net benefit will be lost. This means that when selecting a rate of performance, or a speed of execution for each individual action, the lower cost of performing the action more slowly should be weighed against the opportunity cost of the extra execution time, ie, the net reward rate which could have been obtained during this time. The formal average-reward reinforcement learning solution ensures that such a choice of actions and latencies will indeed lead to the highest possible net influx of benefit, and so will be the truly optimal solution.

Simulating a wide variety of free operant experiments using this model of optimal behavioral choice showed that the well known characteristics of free operant behavior indeed qualitatively match the optimal solution: simulated rats showed a higher response rate when the magnitude of reward was larger or the schedule was more rewarding (lower interval or ratio requirement), response rates were lower on interval schedules

compared to yoked ratio schedules, and when tested on two concurrent interval schedules the simulated rats matched their response rates on each lever to their payoff rates.[5]

3   Opportunity costs and tonic dopamine

Optimal action selection based on online learning of the values of different actions has previously been suggested as a model of action selection in the basal ganglia.[27–30] In one version of these, called Actor/Critic models, it has been suggested[29] that ventral striatal areas (the so called 'Critic') learn to evaluate situations or states of the world, by using a reward prediction error signal provided by dopaminergic neurons in the ventral tegmental area. The dorsal striatum (the 'Actor'), in turn, learns the values of different actions in these states, based on a similar dopaminergic prediction error signal originating in the substantia nigra pars compacta (for a review of the underlying neural data see ref[31]). These models have emphasized the role of phasic dopaminergic firing patterns, which signal temporally local errors in the prediction of future outcomes, in providing the basis of optimal learning of long-term values of actions and states.[28,32]

In addition to requiring this phasic prediction error signal to determine the optimal selection of actions and rates, our model highlights the importance of a new signal which should indicate the expected net rate of rewards, that is, the opportunity cost of time. In a certain class of problems to which this model is applicable, the net reward rate is a global, slowly changing term, common to all the states and to all actions and rates evaluated.[23] That is, whether deciding how fast to perform the next lever-press, or the next nose-poke, and regardless of whether a reward is currently available in the food well or not, the opportunity cost of time is the same — the long term average reward rate forfeited in that time.

What could be the neural bearer of such a global, slowly changing signal? We hypothesize this to be the *tonic level of dopamine* in basal ganglia and prefrontal areas.[5] The tonic level of dopamine is suitable to indicate the net rate of rewards on computational, neural and psychological grounds. Computationally, resulting from the very definition of temporal difference reward prediction errors, a slow averaging over time of the phasic dopaminergic prediction errors will exactly result in the correct average rate of reward. Neurally, dopamine concentrations in target areas such as the striatum are relatively homogeneous,[33] and a recent investigation using fast scan cyclic voltammetry indeed showed that time averaging of phasic dopaminergic activity in target areas results in a stable tonic level,[34] well within the range expected from microdialysis measurements.[35]

Finally, psychological theories of dopamine function have long focused on a putative role for dopamine in modulating the vigor of behavior.[36–43] The identification of tonic dopamine levels with the opportunity cost of time explains, for the first time, *why* dopaminergic manipulations affect response rates as they do. According to our theory, artificially elevating the tonic level of dopamine increases the opportunity cost of time,

with the effect of making the optimal response rates for all actions higher. Suppressing dopamine levels will lead to a reduced cost of time, and slothful behavior. Indeed the most prominent effect of dopaminergic interventions is an enhancement or reduction of overall response rates as a result of increased or decreased dopaminergic transmission, respectively.[42,44–52] Modeling dopamine manipulations as changes in the effective net reward rate, we can simulate and replicate many of these results.[5]

4   The effects of motivation on responding

Using this model of response rates, we can now analyze the effects of motivation on response selection. Understanding how motivation influences behavior is complicated by the fact that animals utilize a number of different action-selection systems, which are differentially sensitive to motivation, and with which we will separately deal below. But first, let us consider in general how motivation can affect the optimal cost/benefit tradeoff we have discussed above. One way to define motivational states, is as a *mapping* between outcomes (or significant events in the world) and the utility they confer to the animal.[4] For instance, food holds high utility for a hungry rat, but a low utility for a sated or thirsty rat (Figure 1b). Using this simple definition, a straightforward means by which motivation can affect action selection, is through the determination of the utility of the outcomes of the different available actions. This corresponds to the traditional 'directing' role ascribed to motivation, because by determining which are the most valuable outcomes motivation can direct action selection toward those actions that will lead to these outcomes.

But this is not the only way that the motivational mapping can affect responding: the outcome utilities will also affect the net rate of rewards (which is measured in units of utility per time). Because the net reward rate serves as the opportunity cost of time, motivation will affect the optimal response rates of all chosen actions. For instance, consider a rat pressing a lever for food pellets on a random interval 30 second schedule. On average, the net rate of reward is equal to the utility of two pellets per minute, minus the costs per minute of the actions emitted in order to obtain and harvest these pellets. If the rat is now made hungrier, the utility of each of the pellets increases and with it the net reward rate, thus increasing the opportunity cost of time and favoring faster responding. In this way, higher motivational states cause higher response rates, while lower motivational states such as satiety decrease the rate of responding. This corresponds to the 'energizing' role of motivation, and the much debated notion of 'Generalized Drive'.[53–55]

In sum, motivation can exert a two-fold influence on responding in our model: a 'directing' effect on the choice of which action to perform, and an 'energizing' effect on the rates with which all actions are performed.[3,5] Figure 1c,d illustrates these two effects, and their qualitative differences. The choice of action depends on a comparison of the local utilities of the outcomes of different actions, and so the 'directing' effect of motivation is *outcome specific* (ie, motivation differentially affects different actions, based on their consequent outcomes; Figure 1c). In contrast, the choice of response rate depends on the global opportunity

cost of time, thus motivation exerts a similar 'energizing' affect on *all* prepotent actions, regardless of their specific outcome (Figure 1d). This explains some hitherto paradoxical observations of 'Generalized Drive', such as the fact that hungrier rats will also work harder for water rewards.

Multiple action selection mechanisms

Although motivation can potentially influence action selection in two ways, different action selection mechanisms may be differentially sensitive to the 'directing' or 'energizing' effects of motivation.[4] In addition to the traditional distinction between Pavlovian and instrumental mechanisms for action control, a recent series of sophisticated studies has teased apart two different types of instrumental control, namely, goal-directed and habitual behavior, based exactly on their susceptibility to motivational influences.[56] The evidence points to two neurally distinct[57] behavioral controllers which employ different computational strategies to estimate what is the currently optimal behavior.[58] The goal-directed system uses a forward model (or action→outcome knowledge) to iterate forward to the expected consequences of a series of actions. As such, its decision-making process is directly sensitive to the utilities of the outcomes consequent on the different actions.[59–62] Conversely, habitual decision making eschews the online simulation of potential consequences of actions, relying instead on estimates of the long-term values of actions, which have been previously learned and stored. These value estimates summarize previous experience about the consequences of actions, but do not represent the outcomes themselves. As a result, habitual responding is not immediately sensitive to changes in action-outcome contingencies,[60,63–65] and similarly can not react to a change in outcome utilities without the relatively slow relearning of new values of actions.[58]

How do the two effects of motivation interact with the constraints of these instrumental action selection systems? We can expect goal directed action selection, which chooses actions based on the utility of their consequent outcomes, to express the 'directing' influence of motivation naturally, selecting those actions that lead to desired outcomes based on the current motivational state of the animal. This is, in fact, the characteristic hallmark of goal-directed behavior.[56] Moreover, the effects of motivational shifts on goal-directed responding have been shown to depend on a process of 'incentive learning' (in which animals experience the utilities of different outcomes in different motivational states[66–71]), testifying that motivational states indeed affect action selection through outcome utilities.

The habitual controller, however, can choose actions that are optimal for the current motivational state only if the animal has learned and stored the long term values of different actions in this motivational state. This means that a rat that has been extensively trained (to the point of habitization), in a state of hunger, to press one lever for food and another for water, will not be able to adjust its behavior flexibly and will continue to predominantly press the food lever, even when shifted to a motivational state of thirst. Only through subsequent learning of the new values of the lever-press actions in terms of the utility of their consequent outcomes in the new motivational state, will habitual behavior be sensitive to the 'directing' effects of

motivation.[58]

Does this mean that habitual behavior is initially totally insensitive to motivational manipulations? We argue to the contrary.[4] Because motivation also exerts a global effect on response rates, which is independent of the specific outcomes of the different actions, motivational states can 'energize' both habitual and goal-directed behavior. Assuming that the animal can estimate at least the *direction* in which the net reward rate will change in the new motivational state (which depends on whether the current motivational state is lower or higher than previously, and whether the animal has reason to expect the availability of outcomes that are relevant to this state), the rate of responding, whether habitual or goal-directed, can be adjusted appropriately so as to approximate the optimal solution. Our model thus predicts that habitual behavior should be sensitive to the 'energizing' aspects of motivation, while goal-directed behavior should be affected by both the 'energizing' and the 'directing' aspects.[4]

5 Discussion

Building on and extending previous normative models of action selection, we have suggested a model of optimal selection of response rates in free operant tasks. Our analysis focused on the critical tradeoffs that need to be negotiated in order to reap rewards at the highest possible rate and the lowest possible cost. This revealed that, different from the decision of which action to perform that is determined by outcome-specific considerations, decisions regarding response rates are determined by global considerations as the consequence of slow performance is to delay all future outcomes. This insight provided the basis for a novel outlook on the effects of motivation on the one hand, and of dopamine on the other hand, on instrumental responding.

In our model the global quantity used to evaluate the cost of delaying all future rewards, that is, the opportunity cost of time, is the net rate of rewards. We suggest that this quantity is reported by the tonic level of dopamine, which explains why high levels of dopamine are associated with generally high response rates, and lower levels of dopamine induce lethargy. Consequently, dopamine has a dual effect on behavior: an effect on action choice through learning, based on phasic aspects of dopaminergic signalling, and an effect on rate selection, mediated by tonic levels. Different from other roles that have been suggested for tonic dopamine,[72–74] our analysis is the first to suggest a normative role, and to imply that the tonic level of dopamine is a quantity which represents specific aspects of the task and of the animal's performance in it. From this follow computationally specific predictions: our model predicts that tonic levels of dopamine will be higher when performing a more rewarding or a less costly task, and lower when working harder or for fewer rewards.

We have further argued that motivation also exerts a two-fold effect on responding. By determining the

mapping between outcomes and their utility, motivation 'directs' action selection to those actions which are expected to yield the most valued outcomes, and 'energizes' all ongoing behavior through affecting the overall reward rate. However, due to the computational limitations of the habitual system, only the goal-directed system is susceptible to the 'directing' effect of motivation. The 'energizing' effect, in contrast, can influence both habitual and goal-directed behavior. It is this latter effect that we hypothesize to be mediated by tonic levels of dopamine, suggesting a strong link between motivation and dopaminergic control.[75] The direct prediction, which has yet to be tested, is that higher motivational states will be associated with higher tonic levels of dopamine (providing the animal has reason to believe that motivation-relevant outcomes are forthcoming).

Incentive motivation and dopamine

In our model, response rates are determined based on the vigor cost of the action and the overall net reward rate, but importantly, without regard for the outcome contingent on the specific action. However, behavioral results from discrete trial experiments show that specific outcome expectancies do affect response latencies, with responding to cues predictive of higher reward being typically faster than responding to less valuable cues.[76–79] Furthermore, although in our model the speed of responding is generally associated with the tonic level of dopamine, dopaminergic recordings have shown a linear relationship between reaction times and phasic dopaminergic responding.[80,81]

If the tonic average reward signal is indeed computed by slow averaging of the phasic prediction error signals, then this result is perhaps not surprising. Cues associated with higher reward expectancies induce larger phasic reward prediction signals,[82,83] which would transiently elevate dopamine tone,[81,84,85] influencing vigor selection and resulting in faster responding. This explanation is a slightly different outlook on ideas about 'incentive motivation', according to which different outcomes exert a motivational effect on responding by virtue of their incentive value.[56,86,87]

Pavlovian responding

We have accounted for the role of dopamine, and that of motivation, in controlling habitual and goal-directed instrumental responding. What about the third class of behavior, namely, Pavlovian responding? The answer to this is not straightforward. On the one hand, phasic dopamine reward prediction errors have been implicated in optimal learning of Pavlovian predictive values, as well as instrumental values. On the other hand, Pavlovian responding itself is not necessarily normative — rather than a flexible, optimal, adaptation to a task, it seems as if Pavlovian responding is adaptive only on an evolutionary timescale. Within an animal's behavioral repertoire, Pavlovian responses are characterized by their inflexibility, and tasks can be

constructed in which they are strictly suboptimal. For instance, Pavlovian behavior persists even in circumstances (such as omission schedules) in which the occurrence of the Pavlovian response *prevents* the delivery of a reward. It therefore seems that Pavlovian responses are an inevitable consequence of the predictive value of cues.[88] A normative model is thus limited in its applicability to Pavlovian responding.

There is another sense in which our model is ill-suited for Pavlovian behavior: a critical simplification of our model is that once a decision is made regarding the next optimal action and the latency with which to perform it, the validity of this decision does not change while the action is executed. That is, we have assumed that the state of the world (eg, whether a reward is available in the food well or not) does not change while an animal is executing an action. Though this is true in free-operant schedules, our framework can not be used, without modification, to model tasks in which this assumption is invalid, such as in instrumental avoidance conditioning (in which an aversive outcome occurs if a response is not performed fast enough). More generally, the model can not incorporate Pavlovian state changes, eg, stimuli appearing and disappearing, and rewards that are given regardless of the animal's actions.

Having said this, we can still derive some insight from the model as to the effect Pavlovian cues or rewards *should* have on instrumental behavior in a simplified setting. Consider the case of a rat performing an appetitive free operant task, to which we now add a 'free' reward that is delivered independent of the animal's actions, and does not require any harvesting actions (for instance, brain stimulation reward delivered with some fixed probability at every second). Extending our framework to this case is straightforward, and we can analyze the effect of this free reward on ongoing instrumental behavior. According to the optimal solution, and consistent with common sense, such a reward should *have no effect* on any ongoing instrumental behavior: any action and rate of responding that were optimal in the original task, are still optimal in the modified setting. This implies that the effective net reward rate used to determine the optimal rate of instrumental responding should be the same in both tasks, that is, that the net rate of rewards controlling instrumental behavior should comprise of only those rewards that are instrumentally earned.

However, to infer which rewards are earned instrumentally and which would have been delivered regardless of one's actions is not at all a trivial problem, especially when behavior is habitual. Indeed, although animals show sensitivity to the contingencies between actions and rewards and reduce responding on a lever if rewards are offered at the same rate whether the lever is or is not pressed (a 'contingency degradation' treatment[59,61,89,90]), responding in such cases is not completely eliminated, evidence for some confusion on the part of the animal. As a result of such overestimation of agency in obtaining Pavlovian rewards, the net instrumental reward rate would be overestimated, leading to instrumental response rates that are higher than is optimal.

A more obvious example is the phenomenon of Pavlovian to instrumental transfer (PIT) in which the onset of a cue that has been associated previously with Pavlovian rewards, enhances the rate of ongoing instrumental behavior. This is clearly not optimal: the Pavlovian cue does nothing to change the tradeoff determining the

optimal response rate. Nonetheless, PIT has been demonstrated in a wide host of settings.[91–95] It seems, then, that similar to the suboptimality of Pavlovian responding in general, Pavlovian effects on instrumental responding are suboptimal. Our model suggests that this is the result of erroneous inclusion of Pavlovian rewards in the expected net rate of instrumental rewards. Interestingly, there is an outcome-specific and an outcome-nonspecific component to PIT.[95] Based on our model and some suggestive experimental results,[40] it is tempting to propose that, like effects of motivation on behavior, the outcome-nonspecific effect of Pavlovian cues is indeed mediated by the tonic level of dopamine.

Where is the tradeoff resolved?

Finally, where in the brain is the tradeoff controlling response rate resolved is currently an open question. As this computation can be shared by both habitual and goal-directed controllers of instrumental behavior, it might not reside in either of these two neural systems. One potential candidate is the anterior cingulate cortex (ACC), and its projections to the nucleus accumbens, and to midbrain dopaminergic neurons.[96,97] The ACC has been implicated in monitoring conflict in cognitive tasks, specifically at the level of response selection, possibly as an index of task difficulty as part of a cost/benefit analysis underlying action selection.[98] Recent investigations using tasks specifically designed to probe cost/benefit tradeoffs,[48,99] confirmed that animals do indeed weigh up the amount of effort required for obtaining a reward on each of the available options, in order to decide which course of action to take.[96] In these same tasks, lesions to the ACC (but not to other medial frontal areas) affected animals' cost/benefit tradeoff, and caused them to prefer a low-effort/low-reward option to the high-effort/high-reward option preferred by non-lesioned rats.[96,97,100,101] Although a similar effect is seen with 6-hydroxydopamine lesions of the nucleus accumbens,[48,99] there are differences between the effects of ACC and accumbal dopaminergic lesions,[96] suggesting that the ACC and nucleus accumbens dopamine may fulfill different roles in the decision making process, with nucleus accumbens dopamine computing and signaling the opportunity cost of time, and the ACC integrating this with expected immediate costs and benefits to determine the tradeoff for or against each possible action. Results to the opposite direction, showing excessive nose-poke responding in a go/no-go task after ACC lesions,[102] indeed suggest that ACC lesions do not merely tilt the balance toward less effortful options (as is suggested for accumbal dopamine depletions), but rather disrupt the instrumental cost/benefit analysis such that a less sophisticated Pavlovian default response pattern is chosen. That is, in a lever pressing task in which the action is not the Pavlovian default, ACC lesions cause the animal to cease pressing, while in an appetitive approach task the Pavlovian default of approaching the nose port dominates as a result of the lesion.

Conclusions

To conclude, from a detailed analysis of the factors affecting response rates we have gained not only a normative understanding of free operant behavior, but also a new outlook on the effects of dopamine and motivation on responding. The tight coupling we suggest between motivation and dopamine is perhaps surprising: dopamine had been related to motivation in early theories, only to be dissociated from signaling reward motivation per se in contemporary normative models. However, we are not advocating to abandon ideas about reward prediction errors, and relapse to the 'anhedonia hypothesis' of dopamine. Rather, we suggest to take normative models of dopamine one step forward, to account for tonic as well as phasic signalling, two distinct modes of transmission which can carry separate computational roles.

References

[1] B. A. Williams. Reinforcement and choice. In *Animal Learning and Cognition*, chapter 4, pages 81–108. Academic Press, 1994.

[2] R. J. Herrnstein. *The Matching Law: Papers in psychology and economics*. Harvard University Press, 1997.

[3] Y. Niv, N. D. Daw, and P. Dayan. How fast to work: Response vigor, motivation and tonic dopamine. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *NIPS 18*, pages 1019–1026. MIT Press, 2005.

[4] Y. Niv, D. Joel, and P. Dayan. A normative perspective on motivation. *Trends Cogn Sci*, 10(8):375–381, 2006.

[5] Y. Niv, N. D. Daw, D. Joel, and P. Dayan. Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology (Berl)*, 2006.

[6] Y. Niv. *The effects of motivation on habitual instrumental behavior*. PhD thesis, The Hebrew University of Jerusalem, 2007.

[7] M. Domjan. *The principles of learning and behavior*. Thomson/Wadsworth, Belmont, California, 5th edition, 2003.

[8] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, 1998.

[9] J. A. Mazur. Steady-state performance on fixed-, mixed-, and random-ratio schedules. *J Exp Anal Behav*, 39(2):293–307, 1983.

[10] W. M. Baum. Performances on ratio and interval schedules of reinforcement: Data and theory. *Journal of the Experimental Analysis of Behavior*, 59:245–264, 1993.

[11] P. R. Killeen. Economics, ecologies and mechanics: The dynamics of responding under conditions of varying motivation. *J Exp Anal Behav*, 64:405–431, 1995.

[12] T. M. Foster, K. A. Blackman, and W. Temple. Open versus closed economies: performance of domestic hens under fixed-ratio schedules. *J Exp Anal Behav*, 67:67–89, 1997.

[13] C. M. Bradshaw, E. Szabadi, and P. Bevan. Relationship between response rate and reinforcement frequency in variable-interval schedules: The effect of concentration of sucrose reinforcement. *Journal of the Experimental Analysis of Behavior*, 29:447–452, 1978.

[14] C. M. Bradshaw, H. V. Ruddle, and E. Szabadi. Relationship between response rate and reinforcement frequency in variable interval schedules: II. Effect of the volume of sucrose reinforcement. *Journal of the Experimental Analysis of Behavior*, 35:263–270, 1981.

[15] G. E. Zuriff. A comparison of variable-ratio and variable-interval schedules of reinforcement. *J Exp Anal Behav*, 13:369–374, 1970.

[16] A. C. Catania, T. J. Matthews, P. J. Silverman, and R. Yohalem. Yoked variable-ratio and variable-interval responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 28:155–161, 1977.

[17] G. R. Dawson and A. Dickinson. Performance on ratio and interval schedules with matched reinforcement rates. *Q J Exp Psych B*, 42:225–239, 1990.

[18] L. P. Sugrue, G. S. Corrado, and W. T. Newsome. Matching behavior and the representation of value in the parietal cortex. *Science*, 304(5678):1782–1787, Jun 2004.

[19] B. Lau and P. W. Glimcher. Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the Experimental Analysis of Behavior*, 84(3):555–579, 2005.

[20] C. J. C. H. Watkins. *Learning with Delayed Rewards*. PhD thesis, Cambridge University, Cambridge, UK, 1989.

[21] A. Schwartz. A reinforcement learning method for maximizing undiscounted rewards. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 298–305. Morgan Kaufmann, 1993.

[22] S. Mahadevan. Average reward reinforcement learning: Foundations, algorithms and empirical results. *Machine Learning*, 22:1–38, 1996.

[23] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*. Athena Sc., 1996.

[24] N. D. Daw and D. S. Touretzky. Long-term reward prediction in TD models of the dopamine system. *Neural Computation*, 14(11):2567–2583, 2002.

[25] N. D. Daw. *Reinforcement learning models of the dopamine system and their behavioral implications*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA., 2003.

[26] N. D. Daw, A. C. Courville, and D. S. Touretzky. Representation and timing in theories of the dopamine system. *Neural Computation*, 18:1637–1677, 2006.

[27] A. G. Barto. Adaptive critic and the basal ganglia. In J. C. Houk, J. L. Davis, and D. G. Beiser, editors, *Models of information processing in the basal ganglia*, chapter 11, pages 215–232. Cambridge: MIT Press, 1995.

[28] P. R. Montague, P. Dayan, and T. J. Sejnowski. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *J Neurosci*, 16(5):1936–1947, 1996.

[29] J. P. O'Doherty, P. Dayan, J. Schultz, R. Deichmann, K. Friston, and R. J. Dolan. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, 304(5669):452–454, 2004.

[30] N. D. Daw, Y. Niv, and P. Dayan. Actions, policies, values, and the basal ganglia. In E. Bezard, editor, *Recent Breakthroughs in Basal Ganglia Research*, pages 111–130. Nova Science Publishers Inc., 2006.

[31] D. Joel, Y. Niv, and E. Ruppin. Actor-critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Networks*, 15:535–547, 2002.

[32] W. Schultz, P. Dayan, and P. R. Montague. A neural substrate of prediction and reward. *Science*, 275:1593–1599, 1997.

[33] G. W. Arbuthnott and J. R. Wickens. Space, time, and dopamine. *Trends in Neuroscience*, In Press.

[34] M. F. Roitman, A. Seipel, J. J. Day, B. J. Aragona, R. M. Carelli, and R. M. Wightman. Rapid onset, short-duration fluctuations in dopamine contribute to the tonic, steady-state level of dopamine concentration in the nucleus accumbens. In *Society for Neuroscience Abstracts*, volume 32: 254.15, 2006.

[35] C. J. Watson, B. J. Venton, and R. T. Kennedy. In vivo measurements of neurotransmitters by microdialysis sampling. *Anal Chem*, 78(5):1391–1399, Mar 2006.

[36] R. J. Beninger. The role of dopamine in locomotor activity and learning. *Brain Res Rev*, 6:173–196, 1983.

[37] K. C. Berridge and T. E. Robinson. What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Review*, 28:309–369, 1998.

[38] W. Schultz. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80:1–27, 1998.

[39] S. Ikemoto and J. Panksepp. The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. *Brain Res Rev*, 31:6–41, 1999.

[40] A. Dickinson, J. Smith, and J. Mirenowicz. Dissociation of Pavlovian and instrumental incentive learning under dopamine agonists. *Behavioral Neuroscience*, 114(3):468–483, 2000.

[41] I. Weiner and D. Joel. Dopamine in schizophrenia: Dysfunctional information processing in basal ganglia-thalamocortical split circuits. In G. Di Chiara, editor, *Handbook of Experimental Pharmacology Vol. 154/II, Dopamine in the CNS II*, pages 417–472. Springer-Verlag, Berlin, 2002.

[42] J. D. Salamone and M. Correa. Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behavioural Brain Research*, 137:3–25, 2002.

[43] A. Murschall and W. Hauber. Inactivation of the ventral tegmental area abolished the general excitatory influence of Pavlovian cues on instrumental performance. *Learn Mem*, 13:123–126, 2006.

[44] M. Lyon and T. W. Robbins. The action of central nervous system stimulant drugs: a general theory concerning amphetamine effects. In *Curr Dev Psychopharmacol*, pages 80–163. Spectrum, New York, 1975.

[45] D. M. Jackson, N. Anden, and A. Dahlstrom. A functional effect of dopamine in the nucleus accumbens and in some other dopamine-rich parts of the rat brain. *Psychopharmacologia*, 45:139–149, 1975.

[46] J. R. Taylor and T. W. Robbins. 6-Hydroxydopamine lesions of the nucleus accumbens, but not of the caudate nucleus, attenuate enhanced responding with reward-related stimuli produced by intra-accumbens *d*-amphetamine. *Psychopharmacology*, 90(390-397), 1986.

[47] G. D. Carr and N. M. White. Effects of systemic and intracranial amphetamine injections on behavior in the open field: a detailed analysis. *Pharmacol Biochem Behav*, 27:113–122, 1987.

[48] J. D. Sokolowski and J. D. Salamone. The role of accumbens dopamine in lever pressing and response allocation: Effects of 6-OHDA injected into core and dorsomedial shell. *Pharmacol Biochem Behav*, 59(3):557–566, 1998.

[49] J. E. Aberman and J. D. Salamone. Nucleus accumbens dopamine depletions make rats more sensitive to high ratio requirements but do not impair primary food reinforcement. *Neuroscience*, 92(2):545–552, 1999.

[50] J. D. Salamone, A. Wisniecki, B. B. Carlson, and M. Correa. Nucleus accumbens dopamine depletions make animals highly sensitive to high fixed ratio requirements but do not impair primary food reinforcement. *Neuroscience*, 5(4):863–870, 2001.

[51] M. Correa, B. B. Carlson, A. Wisniecki, and J. D. Salamone. Nucleus accumbens dopamine and work requirements on interval schedules. *Behav Brain Res*, 137:179–187, 2002.

[52] S. Mingote, S. M. Weber, K. Ishiwari, M. Correa, and J. D. Salamone. Ratio and time requirements on operant schedules: effort-related effects of nucleus accumbens dopamine depletions. *Eur J Neurosci*, 21:1749–1757, 2005.

[53] C. L. Hull. *Principles of behavior: An introduction to behavior theory*. Appleton-Century-Crofts, New York, 1943.

[54] J. S. Brown. *The motivation of behavior*. McGraw-Hill, New York, 1961.

[55] R. C. Bolles. *Theory of Motivation*. Harper & Row, 1967.

17

[56] A. Dickinson and B. W. Balleine. The role of learning in the operation of motivational systems. In Gallistel C.R., editor, *Learning, Motivation and Emotion*, volume 3 of *Steven's Handbook of Experimental Psychology*, chapter 12, pages 497–533. John Wiley & Sons, New York, 2002.

[57] B. W. Balleine. Neural bases of food-seeking: affect, arousal and reward in corticostriatolimbic circuits. *Physiol Behav*, 86(5):717–730, 2005.

[58] N. D. Daw, Y. Niv, and P. Dayan. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci*, 8(12):1704–1711, 2005.

[59] B. W. Balleine and A. Dickinson. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4-5):407–419, 1998.

[60] S. Killcross and E. Coutureau. Coordination of actions and habits in the medial prefrontal conrtex of rats. *Cerebral Cortex*, 13:400–408, 2003.

[61] H. H. Yin, S. B. Ostlund, B. J. Knowlton, and B. W. Balleine. The role of the dorsomedial striatum in instrumental conditioning. *Eur J Neurosci*, 22(2):513–523, 2005.

[62] H. H. Yin, B. J. Knowlton, and B. W. Balleine. Blockade of NMDA receptors in the dorsomedial striatum prevents action-outcome learning in instrumental conditioning. *Eur J Neurosci*, 22(2):505–512, 2005.

[63] C.D. Adams. Variations in the sensitivity of instrumental responding to reinforcer devaluation. *Quarterly Journal of Experimental Psychology*, 34B:77–98, 1982.

[64] A. Dickinson, B. Balleine, A. Watt, F. Gonzalez, and R. A. Boakes. Motivational control after extended instrumental training. *Animal Learning and Behavior*, 23(2):197–206, 1995.

[65] H. H. Yin, B. J. Knowlton, and B. W. Balleine. Lesions of the dosolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience*, 19:181–189, 2004.

[66] R. M. Colwill and R. A. Rescorla. Associative structures in instrumental learning. *The Psychology of Learning and Motivation*, 20:55–104, 1986.

[67] A. Dickinson and G. R. Dawson. Motivational control of instrumental performance: The role of prior experience with the reinforcer. *The Quarterly Journal of Experimental Psychology*, 40B(2):113–134, 1988.

[68] A. Dickinson and G. R. Dawson. Incentive learning and the motivational control of instrumental performance. *The Quarterly Journal of Experimental Psychology*, 41B(1):99–112, 1989.

[69] B. W. Balleine. Instrumental performance following a shift in primary motivation depends on incentive learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 18(3):236–250, 1992.

[70] A. Dickinson and B. W. Balleine. Motivational control of goal-directed action. *Animal Learning & Behavior*, 22:1–18, 1994.

[71] B. W. Balleine. Incentive processes in instrumental conditioning. In R.R. Mowrer and S.B. Klein, editors, *Handbook of contemporary learning theories*, pages 307–366. Lawrence Erlbaum Associates, Mahwah, NJ, 2000.

[72] N. D. Daw, S. Kakade, and P. Dayan. Opponent interactions between serotonin and dopamine. *Neural Networks*, 15(4-6):603–616, 2002.

[73] J. D. Cohen, T. S. Braver, and J. W. Brown. Computational perspectives on dopamine function in prefrontal cortex. *Curr Opin Neurobiol*, 12(2):223–229, 2002.

[74] A. D. Redish. Implications of the temporal difference reinforcement learning model for addiction and relapse. *Neuropsychopharmacology*, 30(Suppl 1):S27–28, 2005.

[75] P. Willner, K. Chawla, D. Sampson, S. Sophokleous, and R. Muscat. Tests of functional equivalence between pimozide pretreatment, extinction and free feeding. *Psychopharmacology*, 95:423–426, 1988.

[76] M. Watanabe, H. Cromwell, L. Tremblay, J. Hollerman, K. Hikosaka, and W. Schultz. Behavioral reactions reflecting differential reward expectations in monkeys. *Exp Brain Res*, 140(4):511–518, 2001.

[77] Y.K. Takikawa, R.K. Kawagoe, H.K. Itoh, H.K. Nakahara, and O.K. Hikosaka. Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res*, 142(2):284–291, 2002.

[78] J. Lauwereyns, K. Watanabe, B. Coe, and O. Hikosaka. A neural correlate of response bias in monkey caudate nucleus. *Nature*, 418(6896):413–417, 2002.

[79] G. Schoenbaum, B. Setlow, S.L. Nugent, M.P. Saddoris, and M. Gallagher. Lesions of orbitofrontal cortex and basolateral amygdala complex disrupt acquisition of odor-guided discriminations and reversals. *Learning & Memory*, 10:129–140, 2003.

[80] T. Satoh, S. Nakai, T. Sato, and M. Kimura. Correlated coding of motivation and outcome of decision by dopamine neurons. *J Neurosci*, 23(30):9913–9923, 2003.

[81] M. F. Roitman, G. D. Stuber, P. E. M. Phillips, R. M. Wightman, and R. M. Carelli. Dopamine operates as a subsecond modulator of food seeking. *J Neurosci*, 24(6):1265–1271, 2004.

[82] C. D. Fiorillo, P. N. Tobler, and W. Schultz. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299(5614):1898–1902, 2003.

[83] P. N. Tobler, C. D. Fiorillo, and W. Schultz. Adaptive coding of reward value by dopamine neurons. *Science*, 307(5715):1642–1645, 2005.

[84] P. E. M. Phillips and R. M. Wightman. Extrasynaptic dopamine and phasic neuronal activity. *Nat Neurosci*, 7:199, 2004.

[85] R. A. Wise. Dopamine, learning and motivation. *Nat Rev Neurosci*, 5:483–495, 2004.

[86] S. M. McClure, N. D. Daw, and P. R. Montague. A computational substrate for incentive salience. *Trends in Neurosci*, 26(8):423–428, 2003.

[87] K. C. Berridge. Motivation concepts in behavioral neuroscience. *Physiol Behav*, 81:179–209, 2004.

[88] P. Dayan, Y. Niv, B. Seymour, and N. D. Daw. The misbehavior of value and the discipline of the will. *Neural Networks*, 19(8):1153–1160, 2006.

[89] L. H. Corbit and B. W. Balleine. The role of the hippocampus in instrumental conditioning. *J Neurosci*, 20(11):4233–4239, Jun 2000.

[90] L. H. Corbit, S. B. Ostlund, and B. W. Balleine. Sensitivity to instrumental contingency degradation is mediated by the entorhinal cortex and its efferents via the dorsal hippocampus. *J Neurosci*, 22(24):10976–10984, Dec 2002.

[91] A. Dickinson and B. Balleine. Motivational control of instrumental performance following a shift from thirst to hunger. *The Quarterly Journal of Experimental Psychology*, 24B(4):413–431, 1990.

[92] R. M. Colwill and S. M. Triola. Instrumental responding remains under the control of the consequent outcome after extended training. *Behavioural Processes*, 57:51–64, 2002.

[93] L. H. Corbit and B. W. Balleine. Instrumental and Pavlovian incentive processes have dissociable effects on components of a heterogeneous instrumental chain. *Journal of Experimental Psychology: Animal Behavior Processes*, 29(2):99–106, 2003.

[94] L. H. Corbit and B. W. Balleine. Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. *J Neurosci*, 25(4):962–970, 2005.

[95] P. C. Holland. Relations between Pavlovian-instruemntal transfer and reinforcer devaluation. *Journal of Expperimental Psychology: Animal Behavior Processes*, 30(2):104–117, 2004.

[96] M. E. Walton, S. W. Kennerley, D. M. Bannerman, P. E. M. Phillips, and M. F. S. Rushworth. Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19(8):1302–1314, 2006.

[97] M. E. Walton, P. H. Rudebeck, D. M. Bannerman, and M. F. S. Rushworth. Calculating the cost of acting in prefrontal cortex. *Annals of the New York Academy of Science*, this issue.

[98] M. M. Botvinick, J. D. Cohen, and C. S. Carter. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci*, 8(12):539–546, 2004.

[99] M. S. Cousins, A. Atherton, L. Turner, and J. D. Salamone. Nucleus accumbens dopamine depletions alter relative response allocation in a T-maze cost/benefit task. *Behav Brain Res*, 74:189–197, 1996.

[100] M. E. Walton, D. M. Bannerman, K. Alterescu, and M. F. S. Rushworth. Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *J Neurosci*, 23(16):6475–6479, 2003.

[101] M. F. S. Rushworth, M. E. Walton, S. W. Kennerley, and D. M. Bannerman. Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci*, 8(9):410–417, 2004.

[102] T. C. Jhou, M. P. Saddoris, J. N. Maddux, T. Furuyashiki, and M. Gallagher. Lesions of the anterior cingulate cortex impair aversive response latencies in a go, no-go task. In *Society for Neuroscience Abstracts*, volume 32: 463.2, 2006.
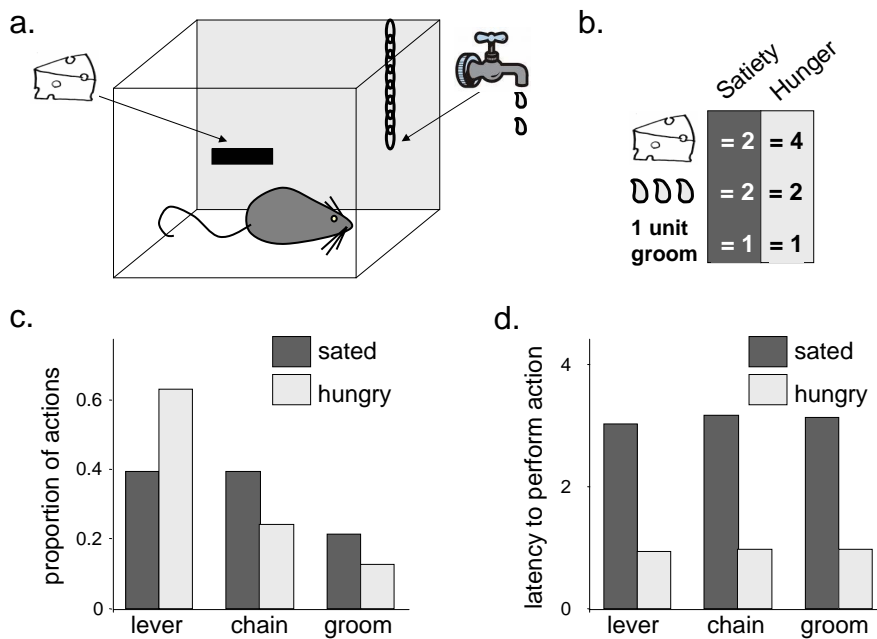
Figure 1: Two behavioral consequences of a motivational shift **a.** A simulated rat, trained in an operant chamber, can perform several actions: it can press a lever for a 20% chance of obtaining cheese, it can pull a chain for a 20% chance of obtaining water, or it can groom itself (with possibly an internal reward). **b.** Even when relatively sated, the cheese and water have slightly higher utilities than grooming. A shift to hunger, however, markedly enhances the utility of cheese, compared to the other utilities that are left unchanged. **c.** One effect of the shift from satiety to hunger is to 'direct' the rat to choose to press the lever (in order to obtain cheese) more often, at the expense of either grooming or chain pulling (which are still performed, albeit less often). **d.** A second orthogonal consequence of the motivational shift is that all actions are now performed faster. Measurements of the latencies to perform individual actions in the simulation reveal that not only is the rate of lever-pressing enhanced, but, when performed, grooming and chain-pulling are also executed faster. This 'energizing' effect of the motivational shift is thus not specific to the action leading to the favored outcome, and can be regarded an outcome-independent effect. Figure modified from ref 4.