



Agency and goal-directed choice

Mimi Liljeholm^{1,2}

In philosophy, agency is construed in terms of desires and means-end beliefs as reasons for actions. Psychological theories of goal-directed behavior provide a formal bridge between objective contingency knowledge and subjective beliefs, missing from such accounts. In this review, I argue that, because they conflate contingency and reward, theories of goal-directed behavior are nonetheless themselves unsuitable as accounts of agency. I then review behavioral and neuroscientific data suggesting that the recently proposed construct of instrumental divergence might serve as a normative and descriptive psychological index of agency that constrains and motivates goal-directed choice.

Addresses

¹ Department of Cognitive Sciences, UC Irvine, United States

² Center for the Neurobiology of Learning and Memory, UC Irvine, United States

Corresponding author: Liljeholm, Mimi (m.liljeholm@uci.edu)

Current Opinion in Behavioral Sciences 2021, 41:78–84

This review comes from a themed issue on **Value-based decision-making**

Edited by **Laura Bradfield** and **Bernard Balleine**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 12th May 2021

<https://doi.org/10.1016/j.cobeha.2021.04.004>

2352-1546/© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Introduction

The topic of *agency* has received considerable attention across a wide range of disciplines, from philosophy and computer science to psychology and neuroscience. In philosophy, agency is commonly explained in terms of desires and beliefs as causes of, and reasons for, actions. Similarly, in psychological theory, goal-directed actions are motivated by their consequences [1–3].

Specifically, formal psychological models of goal-directed choice specify how the subjective utilities of possible outcomes are combined with knowledge about the probabilities of those outcomes given a particular action, in order to estimate the value of that action [4,1]. In this review, I argue that psychological theories of goal-directed choice provide a formal bridge between objective contingency knowledge and subjective beliefs, missing from prominent metaphysical accounts of agency, but that, due to their conflation of contingency and reward,

theories of goal-directed choice are themselves unsuitable as accounts of agency. I then introduce the recently proposed psychological construct of *instrumental divergence* [5•] and suggest that it provides a formal index of agency – specifically in terms of an organism's ability to differentially transform the world – as well as a generic boundary condition on goal-directedness.

The standard story

Contemporary philosophical accounts of the metaphysics of agency have largely focused on the merits of what is known as the *standard story of action* [6]. According to this view, a bodily movement is an *action*, and thus implemented by an *agent*, if it is caused by a desire to experience a particular state of the world and a belief that the bodily movement will, in fact, bring about that state. For example, imagine that you move your arm towards an item sitting on your desk. Whether or not the movement of your arm constitutes agency according to the standard story depends on the causal antecedents of that action: If moving your arm in a manner that places your hand next to the item is caused by a *desire* to have your hand next to the item, and by your *belief* that the movement is a means to the end of your hand being next to the item, then the movement of your arm constitutes agency.

Several objections to the standard theory have been put forth [7]: For example, what about situations where desires and beliefs cause the *absence* of a bodily movement, as when, for example, one declines to answer a phone call because one wishes not to be disturbed [8]; or cases in which a desire-belief pair would have caused a bodily movement, had the movement not been accidentally performed (a so-called *deviant causal chain*, commonly illustrated by the example of an individual whose desires and beliefs so unnerve him that he unintentionally, spastically, performs the intended bodily movement [9]).

Indeed, some detractors have argued that the standard story essentially takes the agent out of agency [10,8], in that the mental states that cause and rationalize actions are events that simply occur within a person but in which the person does not actually participate [11•].

Of particular interest here is the suggestion, by Hornsby [12,8], that what is missing from the standard story is an account of human knowledge about the causal structure of the world — that is, knowledge of how to bring about the consequences of actions. Hornsby [12] notes that, while a subjective belief might constitute a reason for

acting even when that belief is false or unwarranted, a belief-desire account of agency must also include knowledge of objective causal facts, and a connection between such knowledge and means-end beliefs. As detailed below, psychological theories of goal-directed choice formalize this connection by deriving means-end beliefs from objective contingencies (Box 1).

Psychological theories of goal-directed choice

In spite of the many criticisms levelled at the standard story, its basic notion – that agency is, at least partly, characterized by some form of interplay between actions, desires and means-end beliefs – is rarely disputed. This interplay is likewise at the core of psychological theories of goal-directed choice. There are, of course, several salient differences between the two classes of explanations. Most notably, they do not aim to account for the same thing — psychological theories of goal-directed behavior rarely make any explicit statements regarding the nature of agency. Nonetheless, the contents of psychological theories of goal-directed behavior map quite clearly onto the components of agency identified by the standard story. Moreover, theories of goal-directed behavior provide descriptions of how beliefs about the relative efficacy with which an action brings about a desired state may be computed from conditional outcome probabilities, thus identifying causal knowledge as a ‘reason for acting’.

Box 1 Glossary

Contingency: A causal or predictive relationship between two events, such as, for example, between an action and its outcome, commonly formalized as DP — a difference in the probability of the outcome across the absence and presence of the action (i.e. Eq. (1)).

Instrumental behavior: Behaviors that are acquired as a function of their consequences.

Instrumental divergence: The difference between outcome probability distributions associated with available action alternatives.

Action value: The affective or motivational properties of an action, usually a function of its association with a pleasant or unpleasant stimulus.

Action probability: The probability that an action will be performed, usually a function of its value.

Relative efficacy: The degree to which one action is more effective than an alternative action in bringing about an outcome state.

Causal induction: Inductive reasoning about cause and effects.

Sense of agency (SoA): The subjective experience of generating an action or event.

Means-end belief: The belief that a particular action is effective in bringing about an outcome state.

Goal-directed decision-making: Decisions that are based on the current utility and probability of outcomes. Usually contrasted with a more reflexive, ‘habitual’, elicitation of instrumental responses by stimuli based on reinforcement history.

Probabilities and rewards

Early accounts of goal-directed behavior formalized the strength of the action-reward relationship as the difference between two conditional probabilities: the probability of gaining a target reward, r , given that a specific action, a , is performed and the probability of gaining the reward in the absence of that action, $\sim a$:

$$\Delta P = p(r|a) - p(r|\sim a) \quad (1)$$

This ‘instrumental contingency’ [1,3] captures the expected *increase* in reward given an action, and has been reliably demonstrated to govern goal-directed choice in humans [13,14] as well as rodents [1,2].

A more recent formalization of goal-directed behavior [4], adopted from computer science [15], uses an internal model of the world such that, for each action, a , available in the current state, s , and for all possible outcome states, the probability of transitioning into a particular outcome state, s' , given a particular action, $T(s,a,s')$, is dynamically combined, at each choice point, with the reward associated with that outcome state, $R(s')$, to yield the value of the action:

$$Q(s,a) = \sum_{st} T(s,a,st) * \left[r(st) + \gamma \max_{a'} Q(st,a') \right] \quad (2)$$

where $Q(s',a')$ is the recursively defined value of an action performed in the outcome state and γ is a free parameter, discounting the value of delayed rewards. To obtain the advantage of performing a particular action (i.e. the *increase* in reward given that the action is performed), a policy is specified that translates $Q(s,a)$ into an action probability, $p(s,a)$, as a function of relative action values and decision noise. This approach, known as model-based reinforcement learning (RL), has gained immense popularity in decision neuroscience over the past couple of decades.

Mapping goal-directed actions to the metaphysics of agency

It is reasonable to map the reward term, r , in accounts of goal-directed behavior to the desire component of the standard story. Although reward is operationalized in most psychological experiments as a stimulus that has some apparent currency, such as food, money or points, it must nonetheless be presumed, if such models are to explain the full breadth of goal-directed behavior, that any outcome state that the organism desires to obtain, for example, having waved at a neighbor, typed a sentence or mailed a letter, is rewarding. It also seems fair to suggest that ΔP or $p(s,a)$ may correspond to means-end beliefs. It should be noted that neither model is appropriate as an account of causal induction (for example, neither

addresses the possibility of confounding factors); nonetheless, in so far as these postulated mental variables are derived from objective event probabilities, they provide a formal bridge between contingency knowledge and means-end-beliefs. Moreover, they emphasize the *relative* means-end potential of an action — that is, not only whether the action brings about a desired state, but whether it does so more effectively than alternative actions. It might seem, then, that although not intended as such, psychological theories of goal-directed behavior are themselves suitable as accounts of agency. I contend that they are not. Specifically, in the next section, I will argue that, because they conflate the contingency structure of the environment with the constantly changing subjective utilities of outcomes, the action values yielded by such theories do not reliably indicate whether the environment affords flexible instrumental control — a pre-condition of agency and, arguably, a normative boundary condition on goal-directedness [5*].

Agency as instrumental divergence

By defining means-end beliefs as the relative efficacy with which actions bring about desired outcomes, derived from objective conditional probabilities, psychological accounts of goal-directed choice highlight the critical role of causal knowledge in agency. However, as noted, they also conflate instrumental contingencies with dynamic outcomes utilities, so that actions are differentiated solely in terms of their *transient* values. In contrast, the recently proposed construct of *instrumental divergence* [16–18,5*,19*] is a measure of the differential effects of actions that is independent of changes in outcome utilities — as such, it provides a stable index of flexible instrumental control.

Illustration and formalization of instrumental divergence

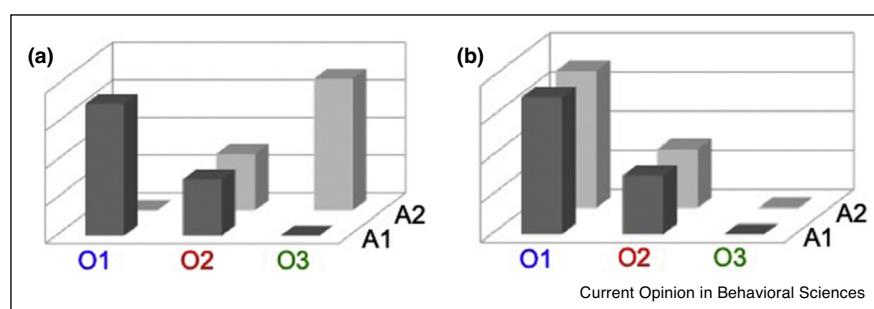
Conceptually, instrumental divergence is simply the difference between outcome distributions associated with alternative actions. As an illustration, consider the scenario in Figure 1a, which shows two available actions, A1

and A2, with bars representing the transition probabilities of each action into three perceptually distinct outcome states, O1, O2 and O3. Here, instrumental divergence is relatively high, since A1 is likely to yield O1 but never results in O3 and the opposite is true for A2. Now consider the scenario in Figure 1b, in which the probability distribution of A2 has been reversed across outcomes, yielding the same outcome variance as in 1a, but with zero instrumental divergence. Note that if the subjective utilities of O1 and O3 are the same, then according to conventional accounts of economic choice, including model-based RL and ΔP , all actions depicted in Figure 1 have the same *expected value*. Consequently, there should be no preference for the scenario depicted in Figure 1a over that in Figure 1b. And yet, if one considers the dynamic nature of subjective outcome utilities, the two scenarios clearly differ.

To appreciate the significance of this difference, imagine that you are required to spend an extended period of time in one of the two environments depicted in Figure 1. Further imagine that O1 and O3 represent food and water, respectively, and that, at the time of choosing between the two environments, you are as hungry as you are thirsty, rendering both outcomes equally desirable. However, having committed, for example, to Figure 1a, you might find that after a large meal without a drop to drink, your desire for O3 is suddenly greater than that for O1. A few hours later, having thoroughly quenched your thirst, you may again prefer O1. Unlike those in Figure 1b, the instrumental contingencies in Figure 1a allow you to produce the currently desired outcome as preferences change, by switching between actions [5*]. This ability, to differentially impact the world by selecting one action over another, is a defining feature of agency that cannot be reliably identified by goal-directed action values, since the dynamic outcome utilities on which such values are based may or may not differ across perceptually distinct outcome states at any given time.

The above concept of instrumental divergence can be formalized as the Jensen-Shannon (JS) divergence of

Figure 1



Probability distributions over three distinct potential outcome states (O1, O2, and O3) for two available actions (A1 and A2) across which instrumental divergence is high (1a) and zero (1b), respectively.

instrumental transition probability distributions [16]. Let T_1 and T_2 be the respective transition probability distributions for two available actions, O the set of possible outcome states, and $T(o)$ the probability of transitioning into a particular outcome state, o , for a given action. The instrumental (Jensen-Shannon) divergence is:

$$ID = \frac{1}{2} \sum_{o \in O} \log \left(\frac{T_1(o)}{T_*(o)} \right) T_1(o) + \frac{1}{2} \sum_{o \in O} \log \left(\frac{T_2(o)}{T_*(o)} \right) T_2(o) \quad (3)$$

where

$$T_* = \frac{1}{2}(T_1 + T_2)$$

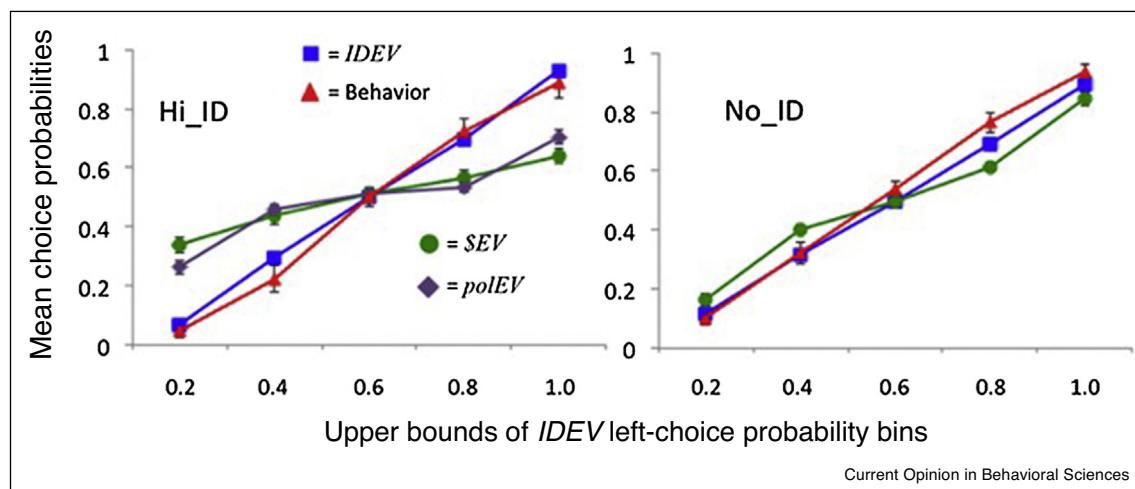
Critically, instrumental divergence is defined with respect to sensory, rather than motivational, features of outcome states. Again, since subjective utilities may change from one moment to the next (e.g. due to sensory satiety), a measure of divergence based on outcome utilities would fail to identify potential instances of flexible instrumental control. Moreover, instrumental divergence is defined over available action alternatives: If T_1 and T_2 were probability distributions associated with different cues, their divergence, while highly relevant for predictability and discriminability, would have no implications for instrumental control [5*].

Behavioral and neural effects of instrumental divergence

Given the dynamic nature of subjective utilities, the high instrumental divergence in Figure 1a is essential for long-run reward maximization and, as such, may have intrinsic value, serving to motivate and reinforce decisions that guide an organism towards high-agency environments. Conversely, the absence of instrumental divergence in Figure 1b renders any attempt at goal-directedness futile and, given the computational expense of such deliberations [20], plainly disadvantageous.

To assess the neural computations mediating the potential utility of instrumental divergence, Norton and Liljeholm [19*] scanned participants with functional MRI as they choose between gambling environments that differed in terms of the expected monetary payoffs and instrumental divergence of available gambling options, as well as in terms of whether participants were allowed to choose freely, or forced to alternate between options (the forced choice, 'auto-play', environments served to control for effects of the discriminability of options, and diversity of outcomes), in a particular environment. They found that a model of expected value that treats instrumental divergence as a reward surrogate provided a better account of participants' choice preferences than did conventional models, sensitive only to monetary reward (Figure 2). Moreover, activity in the rostralateral and ventromedial prefrontal cortex, regions implicated in

Figure 2



Behavioral results from Norton and Liljeholm [19*]: Mean probability of choosing the left over the right of option given a choice between two gambling environments that differed in terms instrumental divergence and monetary pay-offs, derived using an expected value model that treated instrumental divergence as a reward surrogate (IDEV; blue), an expected value model of the mean monetary pay-off available in a given environment (SEV; green), and a model of the maximum monetary pay-off available in a given environment (polEV; purple), respectively, together with actual behavioral left-choice proportions (red), plotted as a function of binned IDEV left-choice probabilities, for choice scenarios in which at least one room was both high-divergence and self-play, yielding high instrumental divergence (Hi_ID), and those in which the high-divergence room was auto-play or both rooms had zero divergence (No_ID). Error bars = SEM.

directed exploration and subjective value computations, respectively, scaled with the divergence-based account of expected value (Figure 3a).

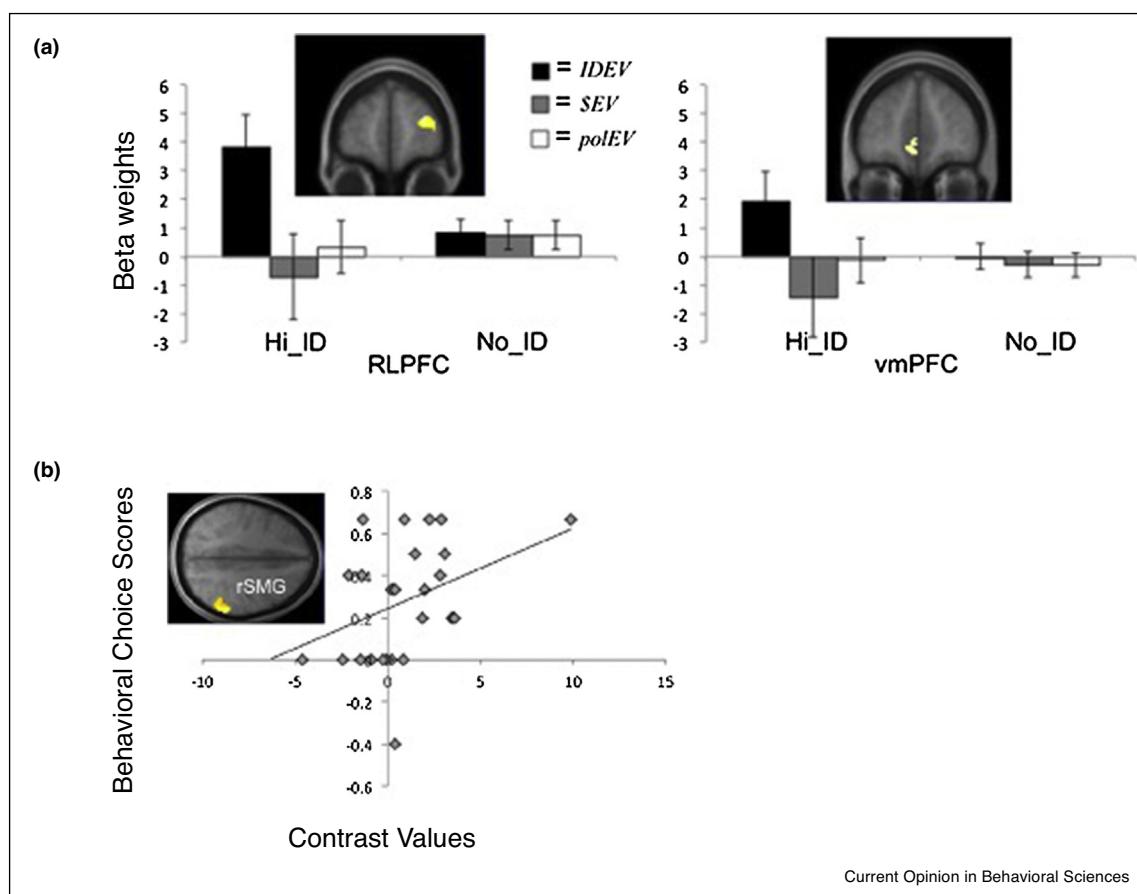
Norton and Liljeholm [19•] also found that, across subjects, the influence of instrumental divergence on economic choice preferences was predicted by the degree to which activity in the right supramarginal gyrus (rSMG) scaled with the divergence-based account of expected value (Figure 3b). Intriguingly, in a different study [17], training-dependent increases in rSMG activity were observed across blocks of instrumental acquisition in a high-divergence, but not a zero-divergence, condition; moreover, the degree to which rSMG activity discriminated between high-divergence and zero-divergence conditions predicted the degree to which those conditions

generated different levels of outcome devaluation sensitivity, a defining feature of goal-directed behavior. These results are important, because the rSMG has also been implicated in the evaluation of self-versus external event attributions [21] – a standard measure of agency in experimental psychology – providing a neural link between instrumental divergence, goal-directed decision making, and the subjective sense of agency.

Conclusions and open questions

In this review, I have proposed that, while it does not provide a comprehensive theory of agency, instrumental divergence – a measure derived from formal models of goal-directed choice – might serve as a normative and descriptive psychological index. What is the added explanatory value of such an index, beyond what is

Figure 3



Current Opinion in Behavioral Sciences

Neuroimaging results from Norton and Liljeholm [19•]: (a) Greater parametric modulation of neural activity by the difference across gambling environments in expected values generated by a model that treated instrumental divergence as a reward surrogate (IDEV) than a model of average expected monetary pay-off (\$EV), for choice scenarios in which at least one room was both high-divergence and self-play, yielding high instrumental divergence (Hi_ID), and those in which the high-divergence room was auto-play or both rooms had zero divergence (No_ID), in the right RLPFC (left) and vmPFC. Bar plots represent unbiased beta weights extracted from RLPFC (left) and vmPFC (right), for neural modulation by differences across gambling environments in IDEV, \$EV and in expected maximum monetary pay-off (polEV). Error bars indicate SEM. (b) Correlation, across individuals, between behavioral preference for high instrumental divergence and parametric modulation of neural activity in the rSMG by an expected value model that treats instrumental divergence as a surrogate reward.

provided by theories of goal-directed choice? As noted, given the dynamic nature of subjective outcome utilities, a relatively high level of instrumental divergence is essential for long-run reward maximization and, as such, may have intrinsic utility — a hypothesis supported by recent behavioral and neuroscientific evidence [19[•]]. Moreover, in the absence of instrumental divergence, the computational expense of goal-directed deliberation does not yield the return of flexible instrumental control, suggesting that a reflexive, or random, decision strategy might be more adaptive [17]. It seems plausible, therefore, that an explicit representation of agency based on instrumental divergence might promote autonomy and optimality in both animal and artificial intelligence.

In experimental psychology and cognitive neuroscience, the topic of agency has been dominated by research on the *phenomenology* of causing or generating an action or event, the so-called Sense of Agency (SoA) [22]. SoA can be assessed explicitly, using declarative statements of self- versus external attribution [23–25], as well as implicitly, as *intentional binding* — a perceived compression of the time interval between voluntary actions and their outcomes [26,27,28[•]]. Notably, while both explicit and implicit measures of SoA are strongly modulated by factors known to impact goal-directed choice, such as action-outcome contingency [29,30] and contiguity [31,32], the two literatures are rarely integrated. It is unclear, therefore, whether variations in SoA reflect variations in the representation of goal-directed actions or variations in agency *per se*. Indeed, exactly how the subjective experience of agency relates to a nonsubjective index, such as instrumental divergence, is an open question.

Finally, and crucially, none of the accounts of agency or goal-directedness discussed in this paper address how actions are acquired in the first place. The conditional probabilities used to compute ΔP , model-based action values, and instrumental divergence are all assumed to be known a priori, as are the representations of actions and outcomes over which the probabilities are defined. Likewise, the Standard Story identifies means-end beliefs as critical features of agency but is silent on the question of their origin. And yet, the innovation of more efficient courses of action in the pursuit of goals is likely a core aspect of agency. Note that, just as environments with low instrumental divergence may not warrant the expense of goal-directed computations, actions (or movements) with low divergence may not warrant the effort of representational discrimination. Future work will be aimed at exploring the role of instrumental divergence in the discovery and shaping of actions.

Conflict of interest statement

Nothing declared.

Acknowledgements

This work was supported by a CAREER grant from the National Science Foundation (1654187) awarded to Mimi Liljeholm. I am grateful to Sean Ostlund and Bernard Balleine for helpful comments on a previous version of this paper.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest

1. Hammond L: **The effect of contingency upon the appetitive conditioning of operant behavior.** *J Exp Anal Behav* 1980:297-304.
2. Balleine Bernad, Dickinson Antony: **Goal-directed instrumental action: contingency and incentive learning and their cortical substrates.** *Neuropharmacology* 1990:407-419.
3. Allan LG, Jenkins HM: **The judgment of contingency and the nature of the response alternatives.** *Can J Psychol* 1980:1.
4. Daw Nathaniel, Niv Yale, Dayan Peter: **Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control.** *Nat Neurosci* 2005:1704-1711.
5. Liljeholm Mimi: **Instrumental divergence and goal-directed choice.** In *Goal-directed Decision Making*. Edited by Morris, Bornstein, Shenhav. Academic Press; 2018:27-48.
6. Davidson Donald: *Actions, Reasons and Causes. Essays on Actions and Events.* Oxford: Clarendon Press; 1980, 685-700.
7. Smith M: **Four objections to the standard story of action (and four replies).** *Philos Issues* 2012:387-401.
8. Hornsby Jennifer: **Agency and actions.** *R Inst Philos Suppl* 2004:1-23.
9. Davidson Donald: *Freedom to act [book auth.]. Essays on Actions and Events.* Oxford: Clarendon Press; 1980.
10. Velleman David: **What happens when someone acts?** *Mind* 1992:461-481.
11. Dalsotto LM: **The agent's role in the causation of action: is Michael Smith's causal theory of action in trouble?** *Kriterion Revista de Filosofia* 2019:143-164.
12. Hornsby Jennifer: **Knowledge, belief, and reasons for acting. Explaining the Mental.** Newcastle: Cambridge Scholars Publishing; 2007, 88-105.
13. Shanks David R, Dickinson Antony: **Instrumental judgment and performance under variations in action-outcome contingency and contiguity.** *Mem Cogn* 1991:353-360.
14. Liljeholm Mimi et al.: **Neural correlates of instrumental contingency learning: differential effects of action-reward conjunction and disjunction.** *J Neurosci* 2011:2474-2480.
15. Sutton RS, Barto AG: **An adaptive network that constructs and uses an internal model of its world.** *Cogn Brain Theory* 1981:217-246.
16. Liljeholm Mimi et al.: **Neural correlates of the divergence of instrumental probability distributions.** *J Neurosci* 2013:12519-12527.
17. Liljeholm M, Dunne S, O'Doherty JP: **Differentiating neural systems mediating the acquisition vs. expression of goal-directed and habitual behavioral control.** *Eur J Neurosci* 2015:358-1371.

18. Mistry Prachi, Liljeholm Mimi: **Instrumental divergence and the value of control.** *Sci Rep* 2016;36295.

19. Norton Kaitlyn G, Liljeholm Mimi: **The rostral-lateral prefrontal cortex mediates a preference for high-agency environments.** *J Neurosci* 2020;4401-4409.
This study demonstrates the utility of instrumental divergence, showing behavioral choice preferences for environments with high instrumental divergence, predicted by activity in brain regions previously implicated in value-based decision making and self-versus external attribution

20. Otto AR et al.: **The curse of planning: dissecting multiple reinforcement-learning systems by taxing the central executive.** *Psychol Sci* 2013;751-761.

21. Koreki A et al.: **Dysconnectivity of the agency network in schizophrenia: a functional magnetic resonance imaging study.** *Front Psychiatry* 2019;171.

22. Haggard P, Chambon V: **Sense of agency.** *Curr Biol* 2012;R390-R392.

23. Dewey J, Carr T: **When dyads act in parallel, a sense of agency for the auditory consequences depends on the order of the actions.** *Conscious Cogn* 2013;155-166.

24. Sato A, Yasuda A: **Illusion of sense of self-agency: discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership.** *Cognition* 2005;241-255.

25. Wegner D, Wheatley T: **Apparent mental causation: sources of the experience of will.** *Am Psychol* 1999;480.

26. Haggard P, Clark S, Kalogeras J: **Voluntary action and conscious awareness.** *Nat Neurosci* 2002;382-385.

27. Moore JW, Obhi SS: **Intentional binding and the sense of agency: a review.** *Conscious Cogn* 2012;546-561.

28. Suzuki K et al.: **Intentional binding without intentional action.** *Psychol Sci* 2019;842-853.
A series of important studies demonstrating that, contrary to conventional interpretation, intentional binding can emerge in the absence of intentional action, suggesting that the phenomenon is more strongly related to multisensory causal binding than to intention or agency.

29. Moore J, Haggard P: **Awareness of action: inference and prediction.** *Conscious Cogn* 2008;136-144.

30. Moore JW et al.: **Feelings of control: contingency determines experience of action.** *Cognition* 2009;279-283.

31. Haering C, Kiesel A: **Was it me when it happened too early? Experience of delayed effects shapes sense of agency.** *Cognition* 2015;38-42.

32. Wen W, Yamashita A, Asama H: **The influence of action-outcome delay and arousal on sense of agency and the intentional binding effect.** *Conscious Cogn* 2015;87-95.