Agency as Planning for Known Unknowns

Abbreviated title: Agency as Planning

Mimi Liljeholm[1,2*]

[1]Department of Cognitive Sciences, UC Irvine

[2]Center for the Neurobiology of Learning and Memory, UC Irvine

* Correspondence should be addressed to:

Mimi Liljeholm

Department of Cognitive Sciences

2312 Social and Behavioral Sciences Gateway

University of California, Irvine

Irvine, CA, 92697-5100

E-mail: m.liljeholm@uci.edu

Conflicts of interest: There are no conflicts of interests.

Acknowledgements: This work was supported by a CAREER grant from the National Science Foundation (1654187) awarded to Mimi Liljeholm. The author thanks Michelle Liang and Kara Schwartz for assistance with cloud-based data management.

Abstract

The utility of a given experience, like interacting with a particular friend or tasting a particular food, fluctuates continually according to homeostatic and hedonic principles. Consequently, to maximize reward, an individual must be able to escape or attain outcomes as preferences change, by switching between actions. Recent work on human and artificial intelligence has defined such control in information theoretic terms. However, information theoretic measures do not necessarily correspond to an agent's ability to escape or obtain sensory-specific states. In the current study, choice behavior was best described by a forward consideration of action values based on possible changes in outcome utilities. Consistent with previous work on agency and personality, individual differences in dimensional schizotypy were correlated with choice preferences in conditions with the greatest and lowest levels of instrumental control. The results contribute to a growing literature on the role of agency in goal-directed choice.

Introduction

What does it mean to control your environment and why is it important? A substantial body of research has demonstrated a preference for voluntary choice, across a wide range of animal species, from pigeons and rats to monkeys and humans (1; 2; 3; 4; 5). Specifically, even when the values and probabilities of outcomes are equated, subjects prefer environments in which they are free to choose between action alternatives over ones in which they are forced to select a particular action. Such findings pose a problem for normative theories of economic choice, which postulate that preferences are determined solely by the probabilities and subjective utilities of decision outcomes. A rational explanation for the free-choice preference is that subjective outcome utilities often change from one moment to the next, and that free choice allows an agent to maximize long run rewards by switching between actions to flexibly produce whichever outcome is most preferred at any given time. Recently, it has been noted that free choice between actions that produce highly similar or identical outcomes affords no such flexibility and that, consequently, *instrumental divergence* – the degree to which action alternatives yield distinct outcomes – is an essential aspect of dynamic reward maximization (6; 7; 8). Here, the computational basis of a preference for instrumental divergence is probed, by contrasting the absolute and relative distances between outcome probability distributions.

As an illustration, imagine that you are a gambler choosing which of two "rooms" to gamble in, knowing that you will only have access to the slot machines in the selected room for several rounds of gambling. Imagine, further, that each slot machine yields three differently colored tokens with various probabilities, and that the values of the different colored tokens, which can be either positive or negative, are currently unknown. Now, consider the two room options illustrated in the top panel of Figure 1: In the room on the left, the two available slot machines, depicted as pie charts, have identical distributions across the three token colors, yielding zero instrumental divergence. In the second room, on the right, one of the machines uniquely delivers red tokens with a non-trivial probability, yielding a slightly greater, but still quite low, information theoretic distance between token probability distributions. Critically, while the distance between token distributions associated with available gambling options is low in both rooms, the 2nd room uniquely affords complete avoidance or possible access to one of the token colors – not knowing the monetary values of different colors, this is likely a desirable feature.

Conversely, in the bottom panel of Figure 1, both room options have relatively large distances between outcome distributions associated with alternative gambling options, and both afford the opportunity to completely avoid or gain access to a particular token color (i.e., red). Consequently, while the difference in distance between token distributions across room options is slightly greater than that in the top panel, the difference in potentially obtained outcomes is lesser – if the preference for instrumental divergence is indeed rooted in an affordance of flexible control as subjective utilities change, the preference for the left over the right room in the bottom panel should be significantly less than that in the top panel. In contrast, if the preference scales with the information theoretic distance between token distributions, used in previous work to quantify agency in biological and artificial systems (9; 10), the preference for the left room should, if anything, be slightly greater in the bottom panel of Figure 1.

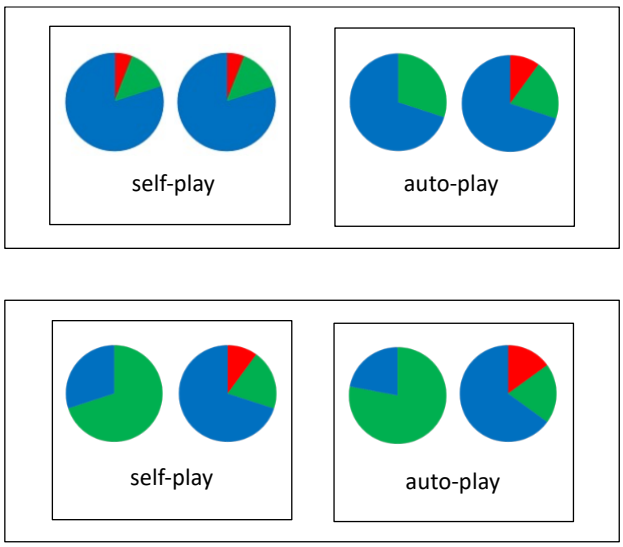

Figure 1: Task illustration. Choice screen shown at the beginning of a gambling round, where room options differ in terms of the outcome divergence of available "slot machines" and in terms of whether gambling in the room involves free (self) or forced (auto) choice. Slot machines are illustrated by pie charts with slices indicating the probabilities of different color outcomes for each machine. The difference across rooms in the information theoretic outcome divergence of available slot machines is 0.04 in the top panel, and 0.05 in the bottom panel.

Importantly, while free choice and outcome divergence are affordances of the environment, such information is filtered through the mind of an agent whose beliefs, discernments, and biases play a critical role in decision making. Substantial evidence suggests that positive symptoms of schizophrenia (e.g., hallucinations and delusions) are associated with an exaggerated *sense of agency* (SoA), whether assessed using explicit statements of self- vs. external attribution (11; 12; 13), or implicitly, as altered time perception (14) or motor interference (15). Here, extending such findings to the domain of motivated behavior, individual differences in positive schizotypy were used to predict a preference for flexible instrumental control in neurotypical adults.

Results

Participants (MTurk; n=120) were instructed to assume the role of a gambler playing various slot machines in a casino, with the goal of maximizing the amount of monetary gain. They were told that, in each of several gambling rounds, they would be required to first select a room in which only two slot machines were available and would be restricted to gamble only on the two machines available in the selected room on several trials within that round. Furthermore, they were informed that each slot machine produced three different token colors (red, green, and blue) with different probabilities, and that each token color was associated with a particular monetary gain or loss that would change from one gambling round to the next, and that would only be revealed once a room had been selected. Finally, participants were instructed that each gambling room was either *self-play* (participants freely choose between available slot machines on each trial) or auto-play (participants were forced to alternate between slot-machines across trials). The measure of primary interest was the choice at the beginning of each gambling round, between rooms that differed in terms of *instrumental divergence* – the information theoretic distance between outcome distributions associated with available slot-machines – and in terms of self- vs. auto-play (see Figure 1).

Figure 2: Task illustration. Choice & feedback screens on a trial inside a gambling room. Note that, once inside the room, the current values of the outcome colors are displayed at the top of the choice screen. These values stay the same throughout the gambling round but change once the participant selects a new room to gamble in (see choice screens in Figure 1).

Once a room was chosen, the two slot machines available in that room were displayed on the screen (see Figure 2), as were the current monetary values of the tokens (recall that these monetary values changed across gambling rounds, so that they were unknown prior to each room choice, but known, and static, once "inside" a room). Participants then gambled in the chosen room on four discrete trials, either choosing between the two slot machines available in the room themselves on each trial (self-play) or alternating between slot-machines across trials as indicated by a computer selection (auto-play).

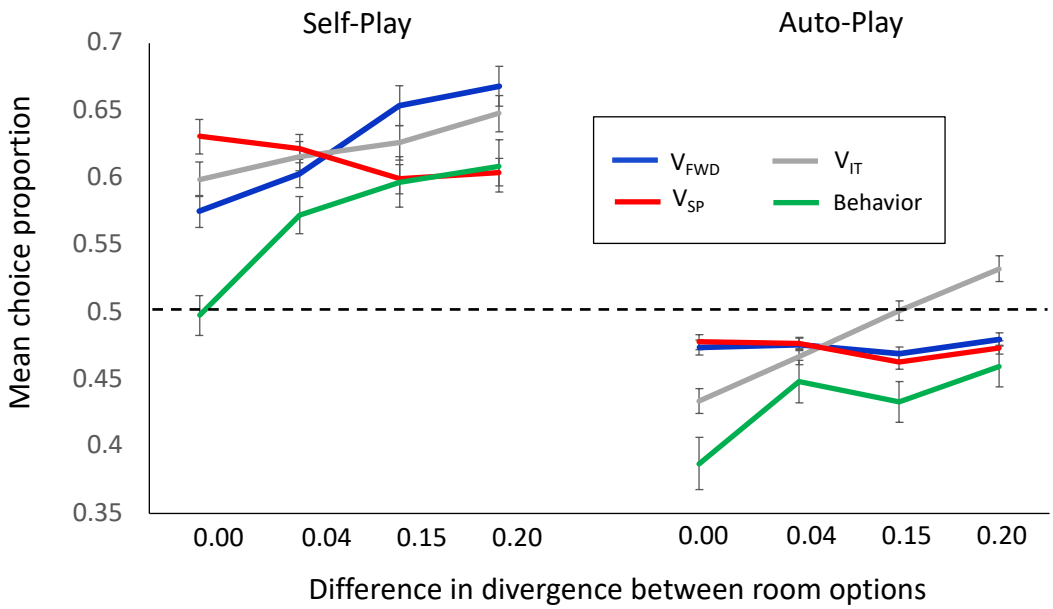

Figure 3: Results. Mean proportions of choosing a gambling room as a function of the information theoretic divergence of outcome distributions associated with available gambling options, for three expected utility models as well as observed behavioral performance. Dashed lines indicate chance performance and error bars = SEM.

Three expected utility models were used to characterize the computations supporting an integration of conventional, monetary, reward with the utility of flexible control (see Methods). In all three models, the acquired value of a particular gambling room (i.e., a particular pair of pie charts) was incrementally updated across gambling rounds using a weighted average of the total payoff in the current round and the previously acquired value of the room (Eq 1). In other words, despite explicit instructions that the monetary values of different token colors would change

from round to round, all models assumed that expectations about monetary outcomes depend, to some extent, on experienced payoffs. Likewise, all three models included a free parameter quantifying the subjective utility of free over forced choice (i.e., self-play). The models differed, however, on their treatment of instrumental divergence. In the first model, V_{SP}, the value of a particular room, informing the decision of which room to gamble in at the onset of each trial, depended only on the incrementally updated expected monetary payoff and whether the room was self- or auto-play. In a second model, V_{IT}, the decision value of a room depended additionally on the information theoretic distance, specifically the *Jensen-Shannon divergence*, between token probability distributions associated with the two slot machines available in the room, weighted by a free parameter to quantify subjective utility: Since the outcome distributions of slot machines available in a room were graphically represented by pie charts on the room-choice screen, participants could compute the IT distance prior to making a room selection.

Note that, in V_{IT}, utility increases linearly with the information theoretic distance between outcome distributions. Finally, in a third model, V_{FWD}, the value of instrumental divergence was defined, not in terms of an information theoretic distance, but by simulating a range of possible token values and computing a forward estimate of the anticipated payoff in each room option given the *possible* values of the different token colors. Critically, when a room was self-play, the slot machine with the greatest payoff given the combination of possible token values was used to compute the anticipated payoff; conversely, when a room was auto-play, computations over possible token values were averaged across machines, reflecting the alternating selection of machines across trials in the room (Eq 6). For all models, room values were transformed into room choice probabilities using a softmax rule (Eq 7).

Table 1: Mean Shannon entropy and expected monetary payoffs, given self- and auto-play respectively, at each divergence level.

Room Divergence	0.00	0.04	0.15	0.20
Mean Entropy	0.74	0.71	0.71	0.71
Self-play pay-off	$0.39	$0.36	$0.34	$0.38
Auto-play pay-off	$0.39	$0.37	$0.34	$0.37

Participants were more likely to select the room with greater divergence when both rooms were self-play than when both rooms were auto-play ($t(119)=2.17$, $p=0.03$), suggesting a specific preference for instrumental control, rather than just outcome divergence. The model-derived probability of choosing a room is plotted, in Figure 3, as a function of the information theoretic divergence between outcome distributions associated with available slot machines, for self- vs. auto-play rooms, and for each expected utility model, together with observed choice preferences. Note, first, that all models categorically discount the auto-play condition. Likewise, all models predict an influence of a room's anticipated monetary pay-off, which is negatively correlated with, thus countering the influence of, outcome divergence (see Methods), and which is the sole basis for V_{SP} predictions within self- and auto-play conditions. Finally, whereas the V_{IT} model scales with divergence in *both* self- and auto-play rooms, the V_{FWD} model only scales with divergence in self-play rooms, reflecting reward maximization based specifically on instrumental control.

As can be seen in Figure 3, consistent with the V_{IT} model, participants categorically avoided zero-divergence rooms, whether self- or auto-play. On the other hand, consistent with the V_{FWD} model, choice probabilities selectively increased with subsequent increases in divergence for self-play rooms, but not for auto-play rooms. Overall, the V_{FWD} model provided the best fit to behavioral choice proportions, reflecting both the strong influence of divergence for self-play rooms and the lack of that influence (except for the zero-divergence case) for auto-play rooms. Specifically, the corrected Akaike Criterion (AIC) score was significantly lower for the V_{FWD} model than for both the V_{SP} ($t(119)=4.82$, $p<0.001$) and the V_{IT} ($t(119)=10.93$, $p<0.001$) model. Model simulations across a large parameter space yielded excellent recovery of decision noise ($r>=0.92$) and self- vs. auto-play parameters ($r>=0.89$), robust recovery of the divergence parameter ($r=0.63$), but only moderate recovery of learning rates for monetary payoffs ($r>=0.32$), as might be expected, since participants were told that the monetary value of each token outcome would change from round to round.

Given the substantial literature implicating positive symptoms of schizophrenia in an exaggerated sense of agency (11; 15; 13; 14; 12), we used the Short Oxford-Liverpool Inventory of Feelings and Experiences (O-LIFE; (16)) to assess whether positive schizotypy predicted preferences for flexible instrumental control. Across participants, positive schizotypy was

positively correlated with the probability of choosing a zero-divergence auto-play room ($r=0.28$, $p=0.0019$) and *negatively* correlated with the probability of choosing a high-divergence self-play room ($r=-0.28$, $p=0.0017$); in contrast, schizotypy did *not* significantly predict the probability of choosing a zero-divergence self-play room ($p=0.95$), nor a high-divergence auto-play room ($p=0.12$). Finally, there was a strong negative correlation between positive schizotypy and the overall preference for self-play, collapsed across divergence conditions ($r=-0.34$, $p<0.001$).

Discussion

This study used a novel gambling task in which the decision of interest was a choice between different gambling "rooms", given information about the divergence with which a pair of slot machines available in respective rooms produced a set of distinctly colored token outcomes. Critically, in each round of gambling, the monetary values of the tokens changed, with new values being revealed only *after* a room had been selected, and with gambling restricted to the slot machines in the chosen room for several trials. In addition to token probability distributions associated with available slot machines, gambling rooms differed in terms of whether participants were free to choose between machines during the round (self-play) or forced to alternate between slot machines across trials (auto-play). Three expected utility models were contrasted – each model updated room values based on experienced monetary payoffs, and each treated self-play as a reward surrogate. One model also treated the information theoretic divergence of token distributions across machines as intrinsically rewarding while a third, forward, model maximized expected utility over *possible* future monetary payoffs, given the degree of flexible instrumental control. Participants' choice behavior was best explained by the forward model, suggesting that, rather than being intrinsically rewarding, instrumental divergence is used as a planning variable. The generality of the findings was confirmed with a replication of the results in an independent sample.

Recent work in human (8) and artificial (10) intelligence suggests that an explicit representation of instrumental control may serve as a reward surrogate, reinforcing and motivating decisions and representations that yield high-agency states. Formally, the information theoretic distance between transition probability distributions associated with action alternatives has been proposed to capture an agents control over its environment (6). However, divergence measures do not necessarily reflect an agent's subjective estimate of instrumental control. In particular, as

illustrated in Figure 1, there are instances in which differences in divergence do not correspond to an agent's ability to escape or obtain sensory-specific states. Moreover, the computation of a divergence measure adds an entire dimension to reward-based decision-making, thus decreasing parsimony. On the other hand, a forward consideration of action values based on changes in outcome utilities, the approach that best described choice behavior in the current study, is a constitutive element of model-based value estimation.

In addition to differing in magnitude, choice preferences at different divergence levels, in self- vs. auto-play conditions, were differentially predicted by individual differences in dimensional schizotypy. Specifically, the greater an individual's positive schizotypy score, the *more* likely that individual was to select the zero-divergence auto-play room, and the *less* likely that individual was to choose the greatest-divergence self-play room. A possible interpretation of this pattern of results is that high levels of positive schizotypy produce a bias towards an experience of agency where none is afforded by environmental contingencies; consequently, the zero-divergence auto-play room is perceived as providing more instrumental control than it actually does, making it more appealing; conversely, the high-divergence self-play room is rendered less discriminable, in terms of instrumental control, from lower-divergence or auto-play rooms, reducing its relative appeal. This interpretation is consistent with work on individuals with clinical schizophrenia, which has shown that positive symptoms of schizophrenia, in particular hallucinations and delusions, are associated with an enhanced sense of agency (11; 12; 15; 13; 14), whether assessed explicitly, as when asking an individual to attribute an event to themselves or to an external source, or implicitly, for example in terms of intentional binding – a perceived compression of the time interval between a voluntary action and its outcome .

In summary, the findings reported here suggest that a recently demonstrated preference for high-agency environment may reflect action values computed over potential changes in outcome utilities. Whereas the current study employed instructed changes in monetary payoffs, dynamic fluctuations in hedonic valence associated with sensory-specific features is a fundamental aspect of natural reward. Notably, computing the value of an action based on the current valence of its consequences is a resource intense process. Of course, so are estimates of information theoretic distances, and of action values based on *possible*, rather than current, outcome utilities. However, cognitive capital spent on gauging the degree of agency afforded by an environment

may be well invested, not only because it indicates an opportunity for reward maximization, but also because it informs whether deliberate, goal-directed, decision-making is likely to yield an advantage over less resource intense action selection strategies. An important future direction is the assessment of how instrumental control constraints the deployment of goal-directed decision making.

Methods

Participants One hundred and twenty individuals (45 females; mean age = 23.7 ± 3.6) participated on MTurk for a $5 baseline compensation and up to $7 in additional earnings based on experimental contingencies. All participants gave informed consent, and the study was approved by the Institutional Review Board of the University of California, Irvine.

Task & Procedure The task is illustrated in Figure 1. At the start of the experiment, participants were instructed that they would assume the role of a gambler playing various slot machines in a casino, with the goal of maximizing the amount of monetary gain. Participants were further told that, in each of several gambling rounds, they would be required to first select a room in which only two slot machines were available, that they would be restricted to gamble only on the machines available in the selected room on several trials within that round, and that they would get to keep the monetary earnings, up to $7, from three gambling rounds, randomly drawn from all rounds at the end of the study.

Each slot machine yielded three different colors (red, green and blue) with different probabilities, and each color was worth a particular monetary amount, which differed from gambling round to gambling round, and which was only revealed once a room had been selected. The probabilities with which a given slot machine produced the red, green and blue outcome were graphically illustrated using pie chart slices (Figure 1). The primary measure was the decision at the beginning of each round, between gambling rooms that differed in terms of the information theoretic divergence of outcome distributions associated with available slot machines.

Importantly, information theoretic divergence is defined with respect to the distributions associated with *available* action alternatives: if the relevant outcome distributions were instead associated with different cues, or with any other events not subject to an agent's volition, their divergence, although relevant to the predictability and diversity of outcomes, would not be

instrumental and, consequently, would have no implications for flexible instrumental control. Here, to experimentally isolate the role of agency, rather than outcome diversity, or any other aspect of divergence, we use a self- vs. auto-play manipulation, such that all room options were either self-play – participants choose freely between slot machines available in the selected room – or auto-play – a computer algorithm alternated between machines across trials in the selected room – as indicated by labels printed below options on the room-choice screen (see Figure 1).

There were 4 divergence levels (0.00, 0.04, 0.15 and 0.20 nats) yielding 6 unique "divergence-difference" conditions (e.g., 0.00 vs 0.04, 0.00 vs 0.15 etc.), combined with 4 self- vs auto-play combinations (both rooms self-play, both rooms auto-play, greater divergence room self-play, lesser divergence room self-play), repeated across two color schemes, for a total of 48 room choice trials. The Once a room had been selected, there were 4 discrete slot machine choices performed in that room, yielding a total of 256 gambling trials. The monetary values of color outcomes differed across rounds of gambling, and participants were explicitly instructed that these values would "change from round to round". To ensure that actual monetary earnings were not greater in high-divergence self-play rooms, thus explaining any observed preference for high divergene, reward distributions were constructed such that monetary pay-offs were largely balanced, and, if anything, biased *against* divergence. The expected monetary pay-offs, with the expectation taken over the products of the probability and $ amount associated with each outcome color for each slot-machine, (using the max for self-play rooms and mean for auto-play rooms), and the mean Shannon entropy of slot machines available in a room (i.e., the "flatness" of their outcome probability distributions) are listed for each divergence and self- v. auto-play condition in Table 1. Note that the mean Shannon entropy in a room was slightly greater for zero-divergence rooms, but constant across all other divergence levels.

Assessment of Positive Schizotypy The Oxford-Liverpool Inventory of Feelings and Experiences (O-LIFE) is a four-scale questionnaire intended to assess dimensions of schizotypy, with scales corresponding, respectively to "unusual experiences", "cognitive disorganization", "introvertive anhedonia", and "impulsive non-conformity". In particular, the "unusual experiences" dimension has been phenomenologically related to positive symptoms of schizotypy (17; 18) and, thus, was of primary interest here. We hypothesized that differences along this dimension would predict individual differences in the preference for instrumental

divergence (i.e., for a combination of high-divergence and self-play). The O-LIFE questionnaire was administered immediately after the gambling task for all participants. The Pearson correlation coefficient was computed between positive schizotypy and lowest (0)- vs. greatest-(0.2) divergence conditions, for self-play vs. auto-play rooms.

Computational Variables Instrumental divergence is defined as the Jensen-Shannon (JS) divergence (19) of the outcome probability distributions of available actions. Let P_1 and P_2 be the respective color probability distributions of the two slot machines available in a given room, let O be the set of possible color outcomes, and $P(o)$ the probability of a particular color outcome. The instrumental (Jensen-Shannon) divergence is:

$$ID = \frac{1}{2}\sum_{o \in O} \log\left(\frac{P_1(o)}{P_*(o)}\right) P_1(o) + \frac{1}{2}\sum_{o \in O} \log\left(\frac{P_2(o)}{P_*(o)}\right) P_2(o),$$

where log is the natural logarithm, yielding nats, and

$$P_* = \frac{1}{2}\left(P_1 + P_2\right)$$

Note that instrumental divergence is defined with respect to the sensory, rather than motivational, features of outcome states – since subjective utilities are constantly changing, motivational features are intrinsically unstable as a basis for agency – and with respect to distributions associated with *available* action alternatives, rather than observed actions or cues – since only freely chosen actions confer agency and flexible instrumental control.

We implemented three expected utility models to characterize the computations supporting an integration of conventional, monetary, reward with the utility of flexible control. In all three models, the acquired value of a particular gambling room (i.e., a particular pair of pie charts) was incrementally updated across gambling rounds, such that

$$E\$_{ij+1} \leftarrow \alpha * \Sigma\$_{ij} + (1 - \alpha) * E\$_{ij} \qquad \text{Eq. 1}$$

where $E\$_{ij+1}$ is the expected monetary pay-off in room i on round $j+1$, $\Sigma\$_{ij}$ is the sum of monetary outcomes in room i on round j, $E\$_{ij}$ is the expected monetary payoff in room i on round j, and a is a free learning rate parameter. In other words, despite explicit instructions that the monetary values of different token colors will change from round to round, all models assume that expectations about monetary outcomes depend, to some extent, on experienced

payoffs. Likewise, all three models included a free parameter, γ, quantifying the subjective utility of free over forced choice (i.e., self-play). The models differed, however, on their treatment of instrumental divergence. In the first model, the value of a particular room, V_{SP}, informing the decision of which room to gamble in at the onset of each trial, depended only on the incrementally updated expected monetary payoff, and whether the room was self- or auto-play:

$$V_{SP} = E\$ + \gamma P \qquad \text{Eq. 2}$$

where P is an indicator set to 1 for self-play and 0 for auto-play. In a second model, V_{IT}, the decision value of a room depended additionally on the subjective utility of instrumental divergence, defined as the information theoretic distance, D, between token probability distributions associated with the two slot machines (i.e., pie charts) available in the room, such that

$$V_{IT} = E\$ + \gamma P + \lambda D \qquad \text{Eq. 3}$$

where λ is a free parameter. Note that, in V_{IT}, the utility of instrumental divergence scales linearly with the information theoretic distance, in both self- and auto-play rooms. Finally, in a third model, V_{FWD}, the value of instrumental divergence is captured, not in terms of an information theoretic distance, but by simulating a range of possible token values and computing a forward estimate of the room payoff:

$$V_{FWD} = E\$ + \gamma P + f(\varepsilon\$) \qquad \text{Eq. 4}$$

where $\varepsilon\$$ averages across the products of the probabilities indicated by the pie charts and the *possible* values of the different token colors, such that

$$\varepsilon\$ = \sum_{o \in O} \sum_{z \in \mathbb{Z}} \big(p(o|m) * z(o)\big) \qquad \text{Eq. 5}$$

where $p(o|m)$ is the probability of a particular token outcome, o, given selection of a particular slot machine, m, and $z(o)$ is a possible value of token o from the set $\mathbb{Z} = \{-\$3 : \$1 : \$3\}$. Critically, when a room is self-play (i.e., P=1), the slot machine with the greatest payoff given the combination of values from $\mathbb{Z}$ is used to compute $\varepsilon\$$; conversely, when a room is auto-play,

computations over possible token values are averaged across machines, thus always summing to zero, reflecting the alternating selection of machines across trials in the room:

$$f(\varepsilon\$) = \begin{cases} \max\limits_{m \in M} \varepsilon\$, \ P = 1 \\ \sum_{m \in M} \varepsilon\$, \ P = 0 \end{cases} \qquad \text{Eq. 6}$$

The model-derived value of a room is plotted, for each expected utility model, as a function of the information theoretic distance between outcome probability distributions associated with available action alternatives, for self- and auto-play rooms, together with observed choice preferences, in Figure 3.

Finally, slot machines differed with respect to the unpredictability with which they produced the various colors – that is, with respect to the Shannon entropy of their outcome distributions, defined as:

$$H = -\sum_{o \in O} \log\left(P(o)\right) P(o)$$

Bibliography

1. *Choice as a value.* **Voss, S. C. and Homzie, M. J.** 3, 1970, Psychological Reports, Vol. 26, pp. 912-914.
2. *Preference for free choice over forced choice in pigeons.* **Catania, A. C. and Sagvolden, T.** 1, 1980, Journal of the experimental analysis of behavior, Vol. 34, pp. 77-86.
3. *The lure of choice.* **Bown, N. J., Read, D. and Summers, B.** 4, 2003, Journal of Behavioral Decision Making, Vol. 16, pp. 297-308.
4. *Selection of forced-and free-choice by monkeys (Macaca fascicularis).* **Suzuki, S.** 1, 1999, Perceptual and motor skills, Vol. 88, pp. 242-250.
5. *The inherent reward of choice.* **Leotti, L. A. and Delgado, M. R.** 10, 2011 , Psychological science, Vol. 22, pp. 1310-1318.
6. *Neural correlates of the divergence of instrumental probability distributions. .* **Liljeholm, M., Wang, S., Zhang, J., & O'Doherty, J. P.** 30, 2013, Journal of Neuroscience, Vol. 33, pp. 12519-12527.
7. *Agency and goal-directed choice.* **Liljeholm, M.** 2021, Current Opinion in Behavioral Sciences, Vol. 41, pp. 78-84.
8. *The Rostrolateral Prefrontal Cortex Mediates a Preference for High-Agency Environments. .* **Norton, K. G. and Liljeholm, M.** 22, 2020, Journal of Neuroscience, Vol. 40, pp. 4401-4409.
9. *Agency and goal-directed choice.* **Liljeholm, M.** 2021, Current Opinion in Behavioral Sciences, Vol. 41, pp. 78-84.
10. *Empowerment: A universal agent-centric measure of control.* **Klyubin, A. S., Polani, D. and Nehaniv, C.L.** s.l. : IEEE Press., 2005. Proceedings of the 2005 IEEE Congress on Evolutionary Computation. Vol. 1, pp. 128-135.
11. *Looking for the agent: an investigation into consciousness of action and self-consciousness in schizophrenic patients. .* **Daprati, E., et al.** 1, 1997, Cognition, Vol. 65, pp. 71-86.
12. *Defective recognition of one's own actions in patients with schizophrenia.* **Franck, N., et al.** 2001, Am. J. Psychiatry, Vol. 158, pp. 454–459.
13. *Aberrant sense of agency in patients with schizophrenia: forward and backward over-attribution of temporal causality during intentional action. .* **Maeda, T., et al.** 2012, Psychiatry Res. , Vol. 198, pp. 1–6.
14. *Altered awareness of action in schizophrenia: a specific deficit in predicting action consequences. .* **Voss, M., et al.** 2010, Brain , Vol. 133, pp. 3104–3112.
15. *(2016). Abnormal sense of agency in patients with schizophrenia: Evidence from bimanual coupling paradigm. .* **Garbarini, F., et al.** 43, Frontiers in Behavioral Neuroscience, , Vol. 10.
16. *Short scales for measuring schizotypy. .* **Mason, O., Linney, Y. and Claridge, G.** 2, 2005, Schizophrenia research, Vol. 78, pp. 293-296.
17. *Selectively reduced contrast sensitivity in high schizotypy.* **Harper, L., et al.** 1, 2020 , Experimental Brain Research, Vol. 238, pp. 51-62.
18. *Do measures of schizotypal personality provide non-clinical analogues of schizophrenic symptomatology?* **Cochrane, M., Petch, I. and Pickering, A. D.** 2, 2010, Psychiatry Research, Vol. 176, pp. 150-154.
19. *Divergence measures based on the Shannon entropy. .* **Lin, J.** 1, 1991, IEEE Transactions on Information theory, Vol. 37, pp. 145-151.

20. *When dyads act in parallel, a sense of agency for the auditory consequences depends on the order of the actions.* **Dewey, J. and Carr, T.** 2013, Consciousness and cognition, pp. 155-166.
21. *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.* **Sato, A. and Yasuda, A.** 2005, Cognition, pp. 241-255.
22. *Apparent mental causation: Sources of the experience of will.* **Wegner, D. and Wheatley, T.** 1999, American psychologist,, p. 480.